Bedford College,
Department of Psychology,
(University of London),
Regent's Park,
London NW1 4NS,
England.

COMPUTER SIMULATION OF A NEUROLOGICAL MODEL OF LEARNING

Thesis submitted for the degree of Ph.D.
in the University of London
1980

Samuel Donald Houtman, B.A., B.Sc. (Est.Man.)

10 Meadow Drive,
Hendon,
London NW4.

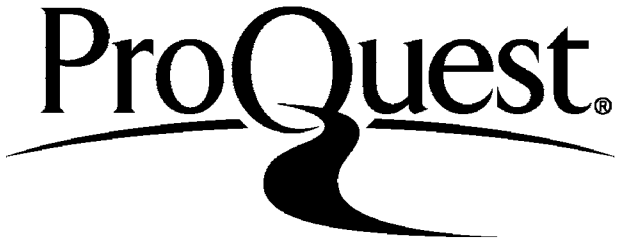

CONTENTS

ABSTRACT

COMPUTER SIMULATION OF A NEUROLOGICAL MODEL OF LEARNING

by

Samuel Donald Houtman

A number of problems in psychology and neurology are discussed to orient the reader to a theory of neural integration. The importance is stressed of the comprehensive temporal and spatial integration of sensory, motor and motivational aspects of brain function. It is argued that an extended neural template theory could provide such an integration.

Contemporary solutions to the problem of neural integration are discussed. The available knowledge concerning the structure of neural tissue leads to the description of a theory of neural integration which might provide such neural templates. Integrating Neurons are suggested to be organised in columns or pools. Subsets of Neurons are formed as a result of excitation and can preferentially exchange excitation. These sub-sets or Linked Constellations would act as spatial templates to be matched with subsequent states of excitation. Inhibition acts to restrict spike emission to the most highly activated sub-sets.

An initial computer simulation represented a simple learning or classical conditioning situation. In a variety of test computer runs the performance confirmed the main predictions of the theoretical model.

The model was then extended to include representation of instrumental, consummatory, motivational and other aspects of behaviour. The intention of these further simulations was not to demonstrate the predictions of prior formulations but rather to use the computer to develop simulations progressively able to represent behaviour.

Difficulties were encountered which were remedied by incorporating rhythmic mechanisms. A number of different versions of the model were explored. It was shown that the models could be trained to produced a different response to discriminative cues, when those cues had previously signalled different contingencies of obtaining the opportunity to perform consummatory behaviour.

A published experiment on the Spiral Illusion is reported, which confirmed predictions suggested by the model.

ACKNOWLEDGEMENTS

I am deeply indebted to my Supervisor, Dr. R. Meddis,formerly of Bedford College, and now at Loughborough. Ray's penetrating criticisms of my work have, I hope, helped me to avoid many pitfalls. He has given generously of his time and effort and has supported me when things were difficult.

I would also like to thank my Professor, Brian Foss, for his continuing support and wise counsel. My especial gratitude is also due to Dr Paul Pal, Computer Services Supervisor, and his staff at Bedford College. In particular I would like to thank Dr Tom Lake for his advice and discussions of the philosophy of programming. On a less formal level I am indebted to both Tom Lake and John Barker for my general discussions with them. I am also grateful to John Valentine, of Bedford College, for his advice on the statistical treatment of data, and to Dr Mary Pickersgill for her advice on the Spiral Illusion.

My thanks are offered to Sarah Chapman for typing and helping me to prepare this thesis.

I would also like to mention early encouragement given to me in the development of the ideas in this thesis by Ray Meddis, Dr A.R. Jonckheere, Dr T. Shallice, Mr E.A. Newman and Mr J. Cowan. My friend, Dr Harry Retkin of Hatfield, has also given me considerable support.

Finally I would like to express my appreciation of my wife and family for putting up with me whilst I have been engaged on this task.

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

BRAIN MODELS AND THEIR CONTRIBUTION TOWARDS UNDERSTANDING BRAIN FUNCTION

I Introduction

The idea that there is some special principle or discovery which would enable the working of the brain to be understood goes back to antiquity and continues to the present day in such models as those of Pribram (1971), Marr (1970), Hendrickson (1972), Mark (1974), Kilmer and Olinski (1974), Stanley and Kilmer (1975), McIlwain (1979) and many others. Over the years, however, the contribution which such a model can make towards the understanding of the brain has come to be seen in a truer perspective. Especially during the past twenty years, there has been an explosion of knowledge and technique in the disciplines surrounding the problem. These include advances as diverse as the development of the electron microscope, the placement of intra-cellular electrodes, computer analysis of the EEG, and, perhaps most remarkable of all, the elucidation of biochemical and hormonal processes. The result has been an enormous increase in the understanding of the physiology, anatomy, histology and biochemistry of the brain. Associated with these advances there has been an expansion of experiments correlating lesions, pharmacological procedures and electrical stimulation with behavioural changes. Psychological experiments designed to tease out the factors affecting behaviour have also grown in sophistication and refinement of technique.

In a more abstract context, the disciplines of communication and control theory, and of the logical aspects of computer programming known as Artificial Intelligence, have supplied models of the kind of processing a nervous system might perform in order to account for the ability of the animal to discriminate gestalt patterns and to seek out natural goals. In this latter respect, the concept of the feed-back from the environment as an essential part of a closed control loop (Wiener, 1949) has greatly modified earlier notions of stimulus-response relationships (Oatley, 1978).

These advances have led to a richer and deeper understanding of brain function than could have been foreseen half a century ago. The brain model has played a part in that progress. It is perhaps mostly through brain models that many principles have been elucidated which enable us to understand how intelligent functions might be mediated by systems of neurons. One may perhaps mention, as example, the work of Uttley (1956, 1959) in demonstrating how simple neural arrangements might represent conditional probability and temporal sequence, as for example in the notes of a melody.

Hebb resolved one of the apparent paradoxes: the apparent diffuse neural processes accompanying mental activity, of which increasing evidence was coming from Adrian (1941, 1947), and the precise patterning of neural events required at the periphery to execute behaviour.

Konorski (1967) has shown how intricate and apparently paradoxical features of behaviour can be brought under a unified concept by considering a system in terms of opposing motivational forces, rather than the presence or absence of a single force, as for example, his postulation of opposing Food and Hunger reflexes. His work followed Anokhin (1961) who broadened the concept of the conditioned reflex by introducing the notion of an acceptor of action which would cause discordance in the case of the return afferentation not matching its requirements.

In proposing another such model based upon a special principle, the author hopes that it may make a contribution of this kind to the general understanding of the working of the brain. The principle is that of a distributed diffuse system. It is, of course, not suggested that the brain can only be understood by reference to this principle, but rather that it is an explanatory principle of some power. The position at the present stage of neuroscience may perhaps be similar to that in physics when both a particle theory and a wave theory seemed to be required to explain light. The hierarchical network systems have received considerable representation in computer simulations, and it seems reasonable to see what can be done with a different kind of theory.

There is a second motive in undertaking a project of the present kind. With the rapid development of the digital computer, there is the possibility of constructing more complex and realistic working models than has previously been possible. Man has always been fascinated with making

models which exhibit life-like characteristics. This applies not only to the mechanical models of former centuries and to present day robots, but also to models like Grey Walter's Speculatrix (Walter, 1953) which exhibit some fundamental resemblances of principle with living things rather than a mere superficial imitation.

In pursuing this second aim of the project, the intention is to try to demonstrate that the model is capable of producing behaviour. As Oatley (1978) stresses, this is one of the cardinal requirements of theories of brain function. The more that the behaviour one is able to produce resembles that of animals, not superficially but in the principles which our new-found insights have given us, the more confidence we may have that the model is telling us something about the function of the brain. On the neurological side, the more resemblance we find between what is known of neural function, and the processes we find in the model, the greater our hope that, even if it does not in itself approximate a solution, at least by exposing the deficiencies of the model, it will lead us to see what is wrong with our ideas.

Accordingly, I would like to stress that I am not proposing that the solutions incorporated in the model, whether of a general kind or detailed mechanisms, represent what actually happens in the brain. Having a conviction that the general organising principle proposed possesses explanatory value, I have attempted to take the model stage by stage to generate behaviour of a more convincing kind. Each stage has presented a number of theoretical difficulties, many of them unforeseen. The solution chosen in each case has seemed to me the one most likely to take the model a stage further. Wherever possible, I have tried to select the solution most in accordance with the behavioural and neurological evidence. That was not always possible, for practical reasons of computing, and also because it would have meant diverting resources to simulate well known and irrelevant aspects of physiology. Nevertheless, it is hoped that the resulting model will be found not to be unlike the brain in its general structure and in the physiology represented.

From what has been said, it will be appreciated that I am aware that my model is wrong in many respects of general principle. It is certainly wrong in detail. However, the difficulties I have experienced in taking a model to the stage of producing behaviour in terms of neurons are possibly universal. The areas of difficulty present a list of

criteria which we might at least apply to other brain models. Have they at least solved these problems in a working simulation.

In many cases, the difficulties concern the control of physiological variables in the model; these problems suggest questions which neurophysiologists might wish to comment upon.

It was hoped that the present project would lead to a more realistic model of the brain in both neuro-anatomical and behavioural terms than has actually been achieved. The author found that this aim has been very much more demanding than envisaged, both in computing resources and in time. Furthermore, the finding of solutions to unforeseen problems in moving towards a more realistic model, e.g. brain (EEG) rhythms, has been far more difficult than was expected. The model as at present simulated does, however, represent a small step in that direction. Further simulations are planned, and it is hoped that the present effort does show the possibility of development in the direction of a more realistic model of a nervous system, given greater concentration of resources.

II The problem of neural integration

Richard Mark (1974) after summarising "all that is known in principle about the physiology of the brain", concludes that "somewhere in such a system is lying undetected a mechanism of modification of some aspect of brain function that can permanently alter the way it works. The result of one fleeting experience is instantly recorded and remains a potent influence for a lifetime of up to a hundred years".

However, it seems doubtful whether the discovery of a particular mechanism of neuronal alteration would in itself tell us much that we wish to know about the way the brain accomplishes its purposes. We need also to know the principle by which the brain is organised. The two questions may be closely related. For example, in Mark's own proposal, the modification is suggested to be a selective suppression of initially profuse connections. This is by way of being a principle rather than a mechanism. There are also different levels of principle; neither of the two examples given explain in themselves how the brain mediates behaviour. Moreover, principle is not always separable from mechanism as in McIlwain's

proposal that neuronal patterns of discharge may reproduce previously imposed patterns converging from several inputs (McIlwain, 1979). Neuronal mechanisms and their relationship with organising principles will be discussed further in Chapter 4.

Whatever this hidden mechanism or principle is taken to be, the problem which faces the model is to be able to demonstrate to some degree at least that it is able to explain the working of the brain. So in effect the problem to be solved by a brain model is to encompass the entire sum of neurological and psychological knowledge available at the time. For this reason the presentation of brain models has tended to fall under one of two approaches. It has often taken the form of a synthesis of existing knowledge, e.g. Pribram (1971); Konorski (1967). An opposite approach, especially in the field of computer simulations, has been to design the model as an expression of principle, and to leave its application to problems of behaviour to the reader, e.g. Kilmer and Olinski (1974); Rosenblatt (1958).

The present project differs from either of these approaches. The theory of neural integration, as initially conceived, is presented in the form of a computer simulation which is intended to demonstrate its basic characteristics. Progressive attempts are then made to extend the model so as to enable it to produce behaviour approximating closer to that of animals. In doing so, advantage is taken of the data available from contemporary psychology, neurology and other relevant disciplines. Unlike many other simulations of neurological models, the later versions of the model presented here do not comprise a computer simulation of a system which is theoretically designed beforehand. The computer simulations are themselves a tool which enables the basic theory of neural integration to be developed further towards a model which bears a closer resemblance to nervous systems. Working in this way, one cannot say that the model actually reflects the working of the brain either in general principle or in detail. One can, however, note the resemblances between the model and its subject in structure and function, and one can assume that where the model is able to reproduce some aspect of animal-like behaviour, it at least sheds light on that behaviour.

III Thesis outline

In the next two chapters some attempt will be made to orient the reader to the psychological approach of the present model. Chapter 4 deals with solutions provided by some other neurological models, and tries to classify them in terms of main principles of operation. This will give some background to the present attempt.

Chapter 5 outlines in broad neurological terms the main concepts of my model. Subsequent chapters describe the simulation of a simple version of the model and its performance, exhibiting the basic characteristics of the theory of neural integration. Mark (1974) in a brief but wide-ranging survey of contemporary knowledge of the brain, lists a number of features which a neurological model might be expected to clarify. These are Pavlovian Conditioning, Operant Conditioning and latent learning. On the neurological side he draws attention to space codes, time codes and overall organisation. In a simple presentation of a neurological model, these would appear to be the priorities, and the simulations presented are oriented towards the demonstration of behaviour in simple conditioning situations. This is not intended to convey a stimulus-response view of brain function, but rather the demonstration that a more holistic view of brain function can demonstrate the kind of behaviour seen in simple stimulus-response situations.

Chapter 10 serves the dual purpose of introducing a set of more advanced simulations of the model and of considering further some other models in the literature which were only briefly touched upon in the opening chapters. In subsequent chapters, the extended simulations are described and their performance reported. Chapter 15 relates some significant problems encountered in the development of the model. These problems were solved by the introduction of rhythmic mechanisms. The functions which such rhythms might serve in the theoretical model are discussed.

Chapter 16 reports a separate simulation intended to demonstrate that the basic proposal of the theory, the formation of neural templates termed Linked Constellations, could be mediated by orthodox synaptic connections obeying simple rules of connectivity. In Chapter 17 a short report of the computer programming is given. This is followed by an account of some further simulations which were in the course of

investigation at the cessation of the work. These programs were intended to extend the mechanisms used in earlier simulations to show how temporal sequences requiring a number of successive instrumental actions might be organised. Although definite demonstrations were not obtained from these programs, they may be of interest in pointing to the manner in which the model might eventually be able to represent more advanced behaviour.

An experiment was carried out to confirm predictions suggested by the theoretical model in relation to the Spiral Illusion. This is discussed in Chapter 18.

The concluding chapter briefly reviews the work carried out.

Simulations are described which are in the course of investigation. These programs extend the mechanisms used in earlier simulations to show how temporal sequences requiring a number of successive instrumental actions might be organised to exhibit purposive behaviour. Although at the time of writing definite demonstrations of these capacities have not been achieved, an account of these programs points to the manner in which the application of an appropriate theory of neural integration might be extended so that the more life-like attributes of brain function may emerge.

IV Terminology

In general the terms used in this thesis are not intended to carry any theoretical import. Where they are intended to do so, their theoretical significance will be specifically stated: otherwise they are merely used as convenient terms of reference to well known aspects of behaviour or neuro-physiology. For example, the terms Action and Response are used interchangeably to refer to the actual physical movements made by animals or notionally by the model. It is not intended to imply that they are released by some direct Stimulus-Response connection. Similarly, the term stimulus is often employed to refer to a particular permutation of the states of the receptor organs of the animal, or to the notional representation of a state by an Input Pattern of the model. Some of the aspects of the model have been given rather inappropriate terms. This has been principally to distinguish them from other parts of the model which might otherwise have been confused.

The aim of these labels has been mainly to avoid the necessity to include long specification phrases within a sentence. The author would apologise if they impose a burden on the reader to remember to which aspects of the model a term refers.

CHAPTER 2

SOME PROBLEMS OF NEURAL INTEGRATION: PATTERN RECOGNITION AND PERCEPTION

I Introduction

This chapter discusses a number of aspects of pattern recognition and perception. The aim is to make some points which will assist in orienting the reader towards the model to be presented. These points may be briefly summarised as:

(i) Pattern recognition as an aspect of animal perception is inseperable from the organisation of action. Accordingly a general model of neural integration should establish the main principle of motor organisation before tackling pattern recognition.

(ii) Temporal organisation of schemata is an important challenge to theories.

(iii) An enrichment model based upon a sophisticated template matching system has greater potentiality than has generally been credited to it.

II The problem of pattern recognition

It should perhaps be stated immediately that the simulations presented in this thesis do not provide a gestalt pattern recognition apparatus. It is not that the theory of neural integration upon which the simulations are based would not probably support such an apparatus. It is rather a question of priorities, as to what it is necessary to demonstrate in a model of neural integration before one tackles what seems to me a more challenging task, that of taking pattern recognition and perception a stage further than is offered by models already available. Possible ways of mediating the recognition of static gestalt patterns are no longer a mystery. Many artefactual devices are capable of doing so with a high degree of efficiency. What does not appear to have been clarified

to the same degree is the manner in which the animal integrates his pattern recognition capabilities to perceive a world of objects.

Many authors have assumed that the primary problem to be solved by a brain model is a demonstration of gestalt pattern recognition. This was probably justified before certain fundamental aspects of pattern recognition had been elucidated, as, for example, by Rosenblatt's Perceptron (1958) described in Chapter 4. Pattern recognition as an aspect of animal perception is, however, a function of basic capacities of neural integration, namely, feature analysis, sequential ordering, sensori-motor co-ordination and the ability to make appropriate selection from total input. By concentrating upon the static image, pattern recognition devices may miss clues to the general integrative properites of nervous systems. It is upon these general properties that the futher understanding of perception depends. Accordingly the strategy taken here is that the solution of basic sensori-motor functions should precede an attempt to further the understanding of pattern recognition.

An indiscriminate use of the term "pattern" has been criticised recently (Uttal, 1973). Generally in this thesis, the use will be that which implies least theoretical commitment. Pattern will be used simply to refer to a particular permutation of the states of a set of elements or fibres. A different use of "pattern" is implied in references to pattern recognition models, and will generally apply in the present chapter. Pattern in the latter usage implies a perceptual equivalence between certain different states of a set of elements, e.g. an equivalence enabling various forms of the letter 'A' to be recognised. The term gestalt pattern will be used as a convenient shorthand whenever there is likely to be ambiguity.

The strength of a demonstration of pattern recognition lies in the accomplishment of a task which seems to contain the essentials of human intelligence. Sutherland (1968) gives a list of twelve aspects of pattern recognition which illustrate this: (i) size variance; (ii) position variance; (iii) brightness variance; (iv) equivalence of outline and filled shapes; (v) lack of invariance under most rotations; (vi) confusions of pattern; (vii) disregard of jitter; (viii) segmentation; (ix) Man's recognition of complex scenes without recognition of detail; (x) perceptual learning; (xi) utilisation of redundancy; (xii) consistency with physiological evidence. The list is impressive.

One might perhaps add the capacity to recognise pattern under conditions of fragmentary and degenerate information. However, as models of the integrative activity of the brain, pattern recognition models have a number of weaknesses. A few words only will be given on some main points.

(1) Without any substantial exceptions, simulations of pattern recognition have theoretical gaps in the story of neural integration. In particular, certain models which have had a considerable influence on theories of neural integration specify a "reinforcement line" or "teaching input" (e.g. Selfridge's Pandemonium; Rosenblatt's Perceptron; see Chapter 4). Such inputs represent the feed-back from the environment. However, the source of the information in such a teaching input is external to the pattern recognition apparatus and its explanation in neural terms is missing. The assumption implicit in such models is that the theoretical gap is either irrelevant or trivial to the task of pattern recognition. The animal presumably takes action following its discrimination of pattern, and the reinforcement line represents the action of the environment on the animal, favourable or unfavourable. Such an assumption may be justified for the model as a pattern recognition device, but not as suggesting a general model of neural integration.

The question of how a model can assess the results of its behaviour as appropriate is by no means as simple as it seems. Even if it is seen simply as a matter of, say, obtaining food or not, the means of neural integration by which this assessment reaches the pattern recognition apparatus as the state of a reinforcement line is far from clear. More significantly, it seems fairly certain that neither animals nor Man learn to discriminate patterns by receiving direct physiological reinforcement such as food after each successful discrimination. Rather their recognition of pattern seems to be an essential part of their interaction with their environment, including activities which are remote from the receiving of physiological rewards and punishments.

In respect of reinforcement, a pattern recognition model faces a harder task than general models of neural integration. A model of say, classical conditioning, might incorporate a mechanism of physiological reinforcement, and that would be logically acceptible. Pattern recognition seems implausible in terms of such reinforcement. Even secondary reinforcement does not answer the difficulty. Many patterns in a sequence of behaviour deriving its organisation in secondary reinforcement

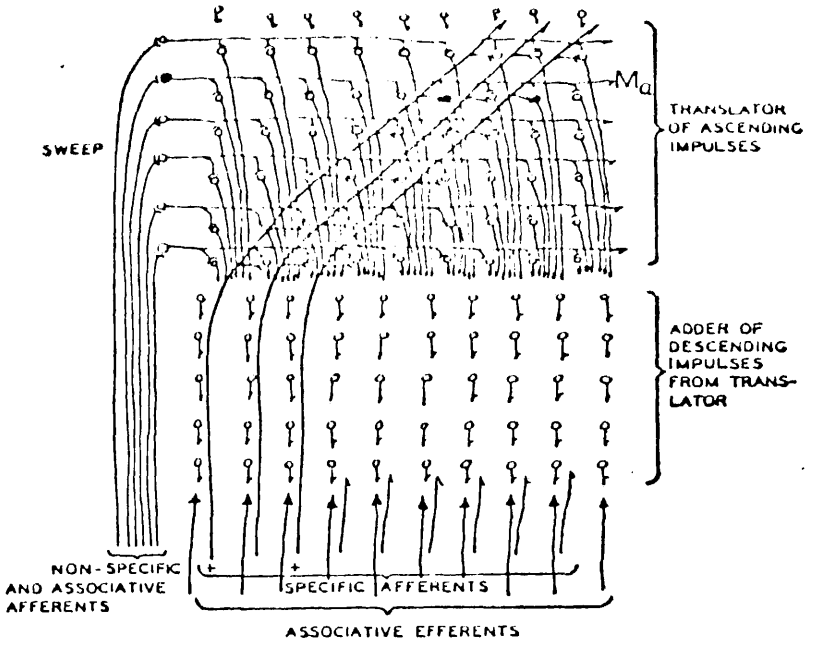

Figure 2.1 (reproduced from Pitts and McCulloch,1947) The impulses from a chord travelling up the specific afferents are facilitated when a non-specific afferent is active. They are then enabled to pass down to the depth of the tissue for computational comparison. Activity in the non-specific afferents moves up and down repetitively "so that the excitement delivered to the depths moves uniformly back and forth,as if the sounds moved up and down together in pitch,preserving intervals".

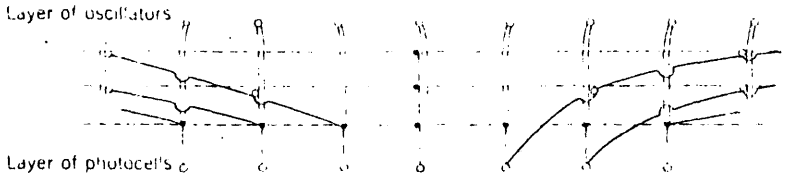

Figure 2.2 (reproduced from Wiener,1948). A device described by Wiener as designed by McCulloch. The single line represents the leads from the photocells, the double lines the leads to the oscillators, the circles on the dotted lines the points of connection between incoming and outgoing leads. The device was to assist the blind to read print by auditory signals. It would equivalence print of different magnifications. According to Wiener, when this design was shown to Dr.von Bonim,it suggested to him the fourth layer of the visual cortex.

terms from a single physiological event, say, eating, would not be equivalent in a gestalt sense. On the contrary, due to eye-movements, many of the retinal images which result even from the inspection of a single object, are clearly not equivalent, in a gestalt sense.

This difficulty clearly does not apply to pattern recognition models which are inherently capable of discriminating pattern. An early model of this kind is useful as illustrating the type of neural arrangement which might fit in with many theories of neural integration. Pitts and McCulloch (1947) suggested that the structure of the cortex, as a series of columns crossed by oblique fibres, would enable transformations of an image to be scanned (Figure 2.1). Undoubtedly the authors did not intend to propose a general explanation of neural integration, as it was presented as only part of their general argument about "How we know Universals". The model was simulated as a physical device which could read print of varying size and type-face (Figure 2.2). The principle of the model is a simple but powerful one, especially as its mode of neural implementation is so clear. The model fits in especially well with my model, because my model is based upon the columnar structure of the cortex also. However, such direct neural arrangements are not the only way in which geometrical transformations might be mediated, as will be discussed in the following section.

(2) "Computer models almost invariably neglect whole classes of evidence - usually behavioural - which should be taken into account ..." (Dodwell, 1970). Three examples are contextual information, sensory analysers and temporal sequential sensori-motor integration. Rosenblatt's Perceptron does not employ sensory analysers, whilst Selfridge's Pandemonium employs feature detection, analogical to sensory analysers, but does not employ sensori-motor integration. Neither uses contextual information. The omission of known features of animal performance and processing does not, of course, detract from these models as illustrations of a particular principle which may or may not operate in the nervous system. What it does suggest is that the models do not represent neural integration as it operates to mediate pattern recognition in the living animal, because animals appear to rely extensively upon other mechanisms. Eye-movements, for example, are an intimate part of the perceptual activity of mammals (see Section III (iv)).

Aside from its involvement in perception, sensori-motor integration is an important aspect of neural integration. The absence of a detailed

formulation of motor organisation in pattern recognition models accordingly leaves a considerable gap in such models. This absence has been noted by a number of workers. In particular, Kilmer and his associates produced several simulated models (Kilmer, McCulloch, Blum, Craighill and Peterson, 1968; Kilmer and Olinski, 1974) which applied principles of neural organisation similar to those proposed by Rosenblatt and Selfridge to control notional motor actions. Such motor responses are however mere names for the output. The structure and organisation of motor action in relation to the environment, e.g. appetitive and consummatory behaviour, is not differentiated. Nor is any neural structure or organisation represented which would discriminate the significance of the feed-back from such motor action. Consequently the relationship between motor behaviour and perception remains at the most abstract level. These models will be further described in Chapter 4.

Pitts and McCulloch (1947) supplemented the pattern recognition model mentioned earlier with a neurological model of the Superior Colliculi which would control eye muscles to converge upon a luminance centre of gravity. This was a step towards the neural representation of those motor actions which assist perception in bringing the retinal image to a standard position on the retina, what Neisser (1967) calls normalising processes. However, the perceptual implications of sensori-motor integration would appear to go considerably beyond normalising procedures (see Gibson, 1968).

(3) Many models which put forward an interesting principle do not qualify as general models of neural integration because they rely upon a single principle to the exclusion of others. Although the principle may enable pattern recognition to be effectively demonstrated, this only goes to show that pattern recognition is not a satisfactory test of a model of neural integration. For example, Rosenblatt's Perceptron gained tremendous support just because it was able to demonstrate a capacity to discriminate gestalt patterns by a single principle of organisation, the feed-back modification of its connections. The principle was of major importance in an abstract theoretical sense. However, the fact that it could acheive its object without sensory analysers, temporal sequential processing, or any distinction between the effects of its actions on the environment except that they were successful or

unsuccessful, suggests that the contribution of the basic principle to the actual pattern recognition of animals may well have been exaggerated.

Animal nervous systems enjoy a wealth of sensory analysis for contour, movement and colour. The feed-back from motor action can hardly be classified as simply successful or unsuccessful. Aside from the point made earlier, that such a classification would provide behavioural equivalence rather than gestalt equivalence, it does not take into account the qualitative learning which the animal gains from manipulating its environment. When we inspect an unfamiliar object, move our eyes or head, or reach out and examine it with our hands, what is it about the sensory feed-back which may be classed as success or failure?

(4) Feature analysers, sampling techniques and special codings play a large part in models of pattern recognition. Such arrangements are, however, equally possible under many different theoretical principles of neural integration. These particular aspects of neural organisation explain the chief items in Sutherland's list (above). Sutherland (1968) reports the ability of goldfish to respond to novel shapes projected upon different parts of the retina. He himself refers this capacity to special codings and feature analysers. It certainly seems possible that goldfish have an advanced visual analysing system. There are good reasons for believing that frogs have one (Lettvin, Maturana, McCulloch and Pitts, 1966). Octopus can recognise and discriminate horizontal and vertical shapes and co-ordinate its motor action accordingly (Young, 1964). In an earlier model related to Octopus, Sutherland (1957) described a neural arrangement which coded the vertical and horizontal extent of figures. Corcoran (1971) comments: There is a very good relationship between the similarity of the coded patterns and the tendency of the animal to confuse shapes . . . but the crucial question of how the similarity between the various codes is computed is not answered". The point made here is that although these models contribute to our understanding of the capacity of animals to recognise patterns, they do not provide models of neural integration.

(5) An important aspect of perception which is missing from pattern recognition models is the relationship of motor action to the geometrical transformation of the retinal image (Gibson, 1968). The retinal

projections of objects occupy successive positions on the retina as the object moves, and, perhaps more significantly, as the animal moves its head, eyes and body. The retinal images also change shape in a regular manner - circular objects present a series of ellipses, rectangular objects, rhombuses. Many of these regular transformations relate to key shapes which are common to many objects. Given the registration of such transformations in conjunction with the motor actions which accompany them and in conjunction with their temporal relationships, the equivalence of retinal images under such transformations follows as a fundamental aspect of neural integration. The concept of an object would appear to be intimately concerned with the temporal sequential aspects of such geometrical transformations (cf. Piaget, 1954). From the recognition of the equivalence of different views of an object to the recognition of the equivalence of different transforms of the retinal image is a step. It may be a step which animals such as rats cannot take. Their pattern recognition appears to be more closely bound to the discrimination of objects: they are able to discriminate two patterns when these are objects, i.e. cues for jumping. To discriminate thirty six alpha-numeric characters may be difficult for them because they lack the capacity to know the different uses to which these thirty six objects can be put (see below).

(6) The ability to identify alpha-numeric characters was undoubtedly a considerable step forward in explaining intelligence in terms of the concepts of the physiological sciences. Man can, of course, do far more, but it is significant that in many respects animals can do far less. One would have thought that if the identification of alpha-numeric characters could be accomplished by a simple arrangement of the nervous system containing simple feature analysers and feed-back adjustment of neuronal connections, then animals as highly developed as the rat, octopus or cat, could learn to do so quite easily. The inability of animals to do so must reduce the plausibility of a demonstration of pattern recognition as elucidating the basic principles of operation of the nervous system.

It is not suggested here that animals cannot recognise patterns of far greater complexity than alpha-numeric characters. Probably one of the difficulties is to get the animals to pay sufficient attention to the aspects of the pattern which are important for alpha-numeric identification. Reading books is not a frequently observed aspect of

cat behaviour. This points to attentional aspects of pattern recognition. Anyway, there are no experiments reporting the ability of animals to be able to identify all 36 alpha-numeric characters. This is an easy task for many pattern recognition models. The patterns used by Lashley (Hinde,1970, p.89) for testing pattern discrimination in rats, although complex were either very different from each other, or relied upon the discrimination of two simple shapes within the presentation, e.g. a cross or a triangle. Even so the animals were only required to discriminate between two of such patterns, not thirty six.

A related point is made by Neisser (1976). According to the kind of explanation given for pattern recognition "the image that he supposedly looks at is upside down, foreshortened, and the wrong size . . . such a homunculus would have a few problems of his own . . . Many sophisticated theories have been advanced to explain how he might solve these problems, but real perceivers do not have them. We do not see retinal images, we see the real environment of objects and events . . .". Moreover, if the object is upside down, we may still recognise it but we also see that it is upside down. We do far more than equivalence retinal patterns. We distinguish them at the same time.

In summary, the aim of this section has been to suggest that pattern recognition devices and principles enter into a description of the nervous system only as part of a general sensori-motor-temporal integration. It is the latter which comprises the main subject matter for a theory of neural integration.

In the context of attempts to simulate brain models, the successful representation of pattern recognition undoubtedly contributed enormously to our understanding of psycho-neurological function. Criticisms of the shortcomings of such models has not been intended to detract from that contribution. It is intended rather to indicate that the stage has been reached when it is desirable to see pattern recognition in the setting of a more general neural integration incorporating temporal and motor aspects.

In Chapter 4 the contribution made by some of the models mentioned briefly in this chapter will be examined in further detail.

One of the aims of this section is to defend in advance the relatively

simple performance of the simulations to be reported here against a possible expectancy that a simulation of a brain model must exhibit performance on a par with that of pattern recognition models.

III Templates and pattern recognition

Before leaving pattern recognition it may be helpful to the reader in following the later exposition of the present model, if we consider certain aspects of explanatory concepts used in pattern recognition theory, in particular the notion of a "simple template". The model to be presented here will probably be considered as a simple template theory. Most writers (e.g. Corcoran, 1971) are rather scathing about template models. Neisser (1967) says "one must choose between elaborating the template model considerably and abandoning it altogether . . . Hebb himself chose the second alternative". Although in other respects the present model may be seen as resembling Hebb's theory, this is one way in which it differs.

Pattern recognition theory may be somewhat loosely described in certain main ideas which appear in most models. Most of the models reject the template and substitute some kind of feature analysis as the first stage of the recognition process. Such feature analysers may perhaps themselves be characterised as (i) specific, (ii) random samples and (iii) sophisticated recognisers. Specific feature analysers are analogous to the receptive fields discovered by Hubel and Wiesel. Random sampling appears in many models where the effect is to distinguish features (see Chapter 4). The statistical implications of random sampling were developed by Rosenblatt into his theory of Separability (1958b). Selfridge's Pandemonium combines the notion of discriminability and feature analysis, by suggesting feature analysers which compete with each other and which are changed to achieve optimum discrimination. Sophisticated feature analysers are those which distinguish features which are closer to perceptual aspects of gestalt patterns, e.g. legs and arms of numbers and letters.

This point will not be pursued here. It is merely intended to show that a template process may be elaborated in quite simple ways to embrace the chief assumptions of pattern recognition.

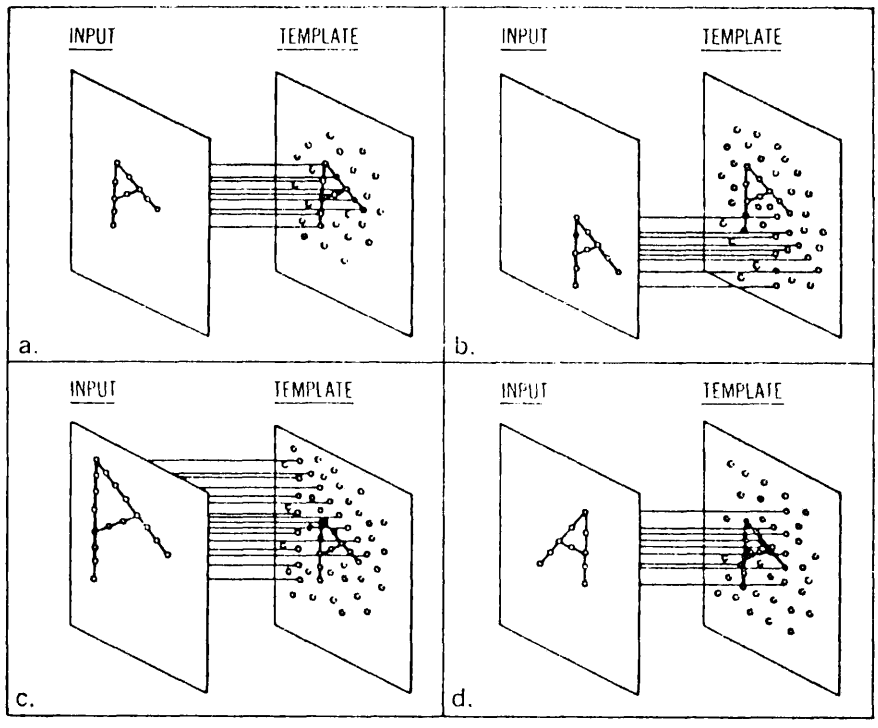

Figure 2.3 (Reproduced from Neisser, 1967) The customary criticism brought against template theories.

The simple template model is illustrated by Neisser (1967), reproduced here as Figure 2.3. It may be elaborated as follows:

(i) Coded templates

Templates need not represent complete isomorphic projections of the retinal patterns. The simplest feature analysers may be seen as simple templates reproducing fragments of gestalt patterns. These templates may be duplicated across the retinal field. For example, the line oriented receptive fields of Hubel and Wiesel (1959). A template with greater abstractive powers would be one taking a smaller number of inputs, say, one where the input to a particular location on the array represented the appearance of a line of a particular orientation anywhere on the retina, i.e. where a neuron fires whenever one or more of the line oriented receptive fields for one particular direction is fired anywhere in the system. A template with generalising capacity would be useful, e.g. where a neuron fires when any receptive field oriented within say, 15 degrees of the optic horizon is activated. These are illustrated in Figure 2.4.

The foregoing is of course merely a manner of expressing the well known principles of feature analyser function. By expressing these principles in simple template terms, it is hoped to remove the presumption that the presence of feature analysers implies some kind of hierarchical organisation of the nervous system. The present model can elaborate such templates without convergent-divergent neural circuits. There is a hierarchy of templates in the systems analysis sense but not in the neurological sense.

A simple template model is not an alternative to a feature analysis model. It is one way of mediating feature analysers.

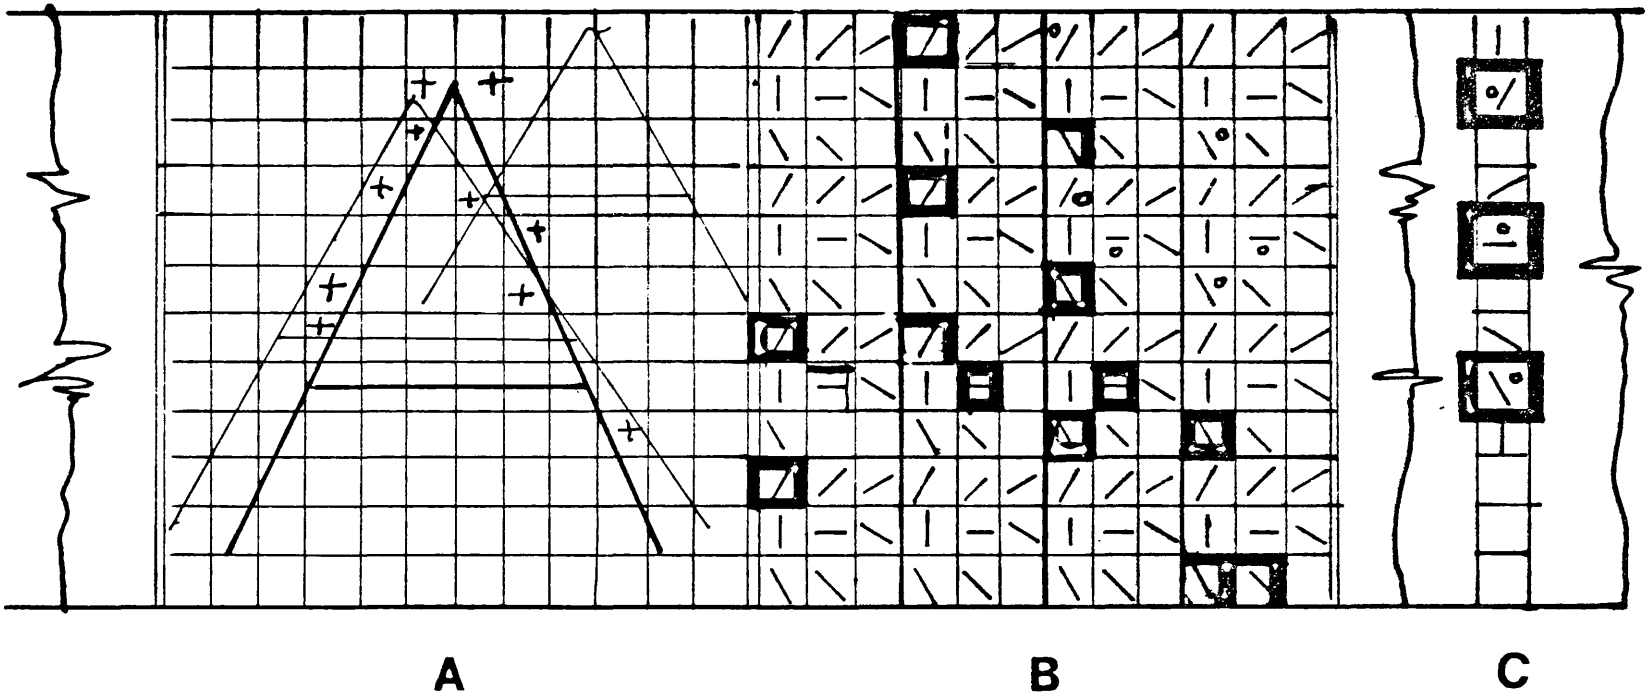

Figure 2.4 Diagram illustrating different analyses of sensory stimulation represented in an extended visual template. A. Simple template. Locations responding to both of the large A's marked +. B. Portion of template responding to line orientations in limited parts of the visual field. Locations responding to both large A's marked with full boxes. Those locations responding to one or other of the large A's marked with half box. Locations responding to small A, marked with small circles. C. Locations responding when a line of given orientation occurs anywhere in visual field.

A's of different size and angle produce little matching in template A, more in B, and complete in C.

(ii) Variety of codings: attention

Most models of pattern recognition explore the potentialities of one particular scheme of coding e.g. angles, lines, or some general sampling and weighting principle. Selfridge's Pandemonium suggests that there are arrays of feature analysers, and these can be altered. In template terms, there may be a multiplication of templates against which the retinal image or parts of it are matched. With any particular input pattern there will be matching in some of these templates and not others. We are not here referring to the matching of templates to the input pattern within a particular category of templates, but the fact that a matching may be found in one template category, say orientation of lines, and not in another, say colour patches. For example, two images of A in two different retinal positions will show a matching in the line orientation template system but not in the simple isomorphic template system. What gave Selfridge's system its power was that it could choose between different codings, according to which set of "demons" shouted the loudest. What is perhaps not so easily seen is that the template system offers a very simple way of doing the same thing; its simple neurological implementation may point to a fundamental mechanism of attention. If we imagine an array of templates with separate cortical locations for each category of encoding, it is a relatively simple matter to visualise how each matching may demand "attention" according to the degree of matching it finds within its own library of templates. These matchings within each category may conflict or confirm matchings in other categories. Such conflicts and confirmations may decide the particular matching which is selected in each category. A matching in the colour library of an input which may coincide with that of many grey, white or blue objects previously encountered, may be accompanied by a matching in the line orientation library of a circular image. To see how this combination of template matching may be used to discriminate against oranges, grey flags or circular dinner plates, or billiard balls, one must assume that the results of the two or more template matchings can be combined in some way. This combination is also often assumed to be necessarily hierarchically organised. To show that it can be seen in template form is the next step.

(iii) An extended template

Template systems are usually conceived as operating within self-contained

categories and self-contained sensory modalities. Generally in pattern recognition, the template analogy is only used of isomorphic images, anyway, but the idea may be extended, as suggested, to include templates which match particular sensory analyses. Each feature analysis category say, line orientation, colour, movement, would have its own separate match-mismatch process, but would be part of an attempt of find the best match for the whole of the extended template. In my proposed model, the principle will be extended to a single template containing the projections of all the sensory and motor modalities, including the projections of all innately organised feature analysers. We need not consider here how such a template is neurally organised: this is the function of the neurological theory to be presented. What is important is that the extended template should be able to find the best matching over its entire span.

With such a template system it may be seen how a tennis ball may be distinguished from an orange. If there is a template of a tennis ball with grey or white in the "colour" section, and "circular" in say a line orientation section, it will be selected in favour of templates of either oranges or grey-white rats, as the templates of such discriminated objects will not be matched in both sections.

It is perhaps not so easy to see how a template system would explain the interpretation of a particular wavelength of light as "grey" on some occasions and "white" on others. Such an explanation depends upon a schematic, contextual and temporal organisation of templates which it would be premature to call upon at this stage. It may be sufficient to stress the argument of the previous section. Perceptual problems require to be discussed in a system of neural integration in which the main problems or organisation have been solved, and in which accordingly, temporal sequential and contextual information may be utilised.

That the neural implementation of the visual system may conform to a template system is reported by some contemporary research. According to Bazier, Robinson and Dow (1977), colour, orientation, spot cells, light inhibited cells and border cells occupy separate populations in area 18 of monkey striate cortex. "The findings that different populations of cells in areas 17 and 18 are maximally sensitive to different stimulus parameters implies that, at least at this stage in visual processing, the neural representation of visual stimuli is accomplished through the simultaneous activity of cells in different groups rather

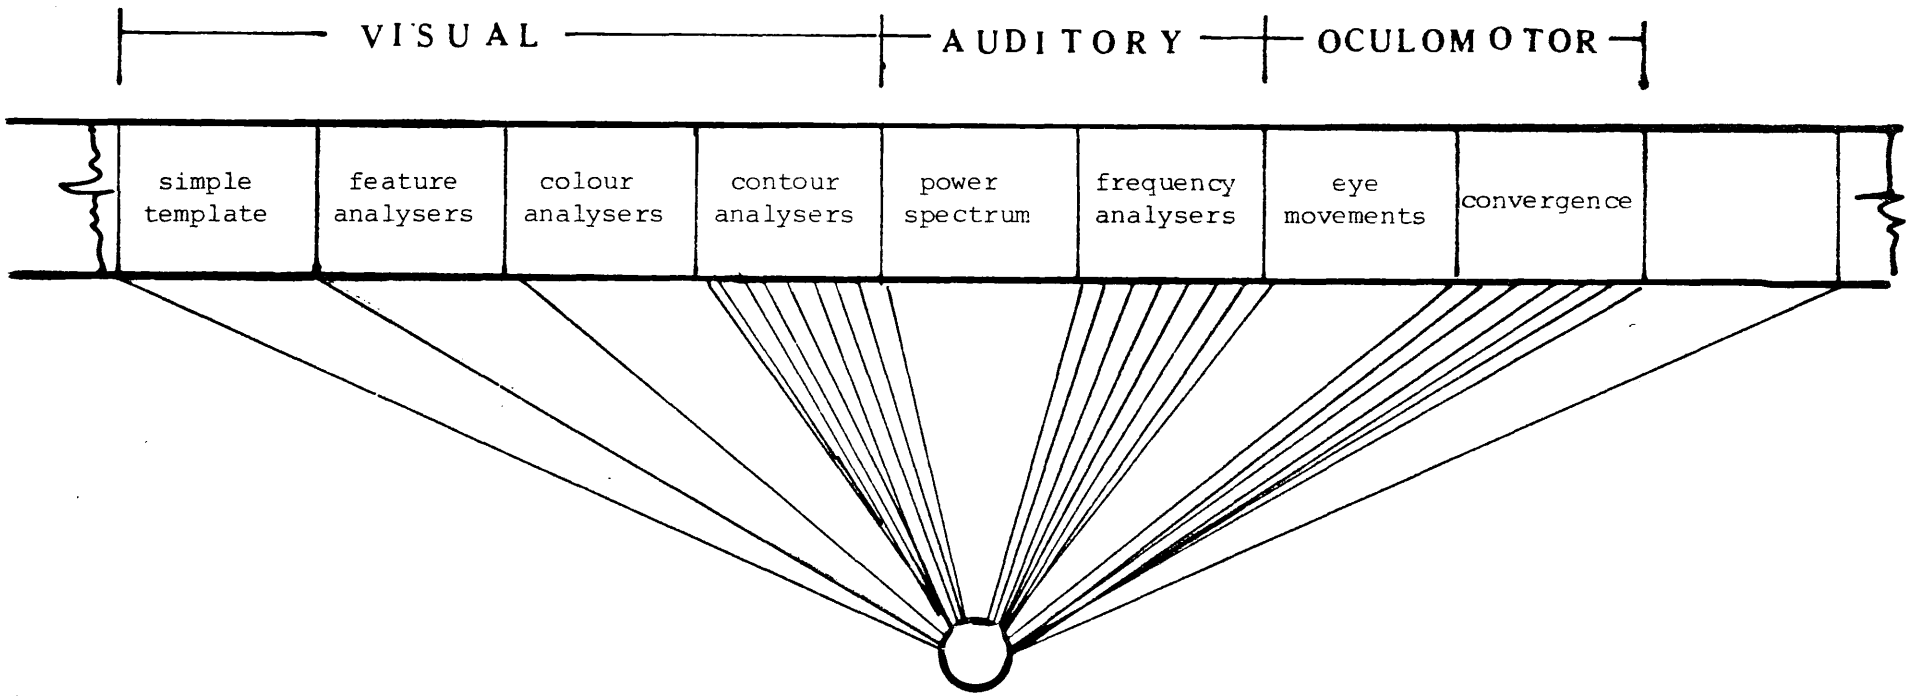

Figure 2.5 Attention in an extended template system. Excitation may be sent selectively to different portions of the template and so weight the information analysed by those portions. It may be noted that the template system includes motor representations.

than through the activities of single cells with compound specifities". The utilisation of simultaneous activity in spatially dispersed cells implies a template concept. Singer and Tretter (1976) interpret the results of their own experiments in a similar manner: "It rather appears as if the predetermined circuitry of striate cortex allows for interactive processes between fairly distant pattern elements". In the template concept, a particular "colour" receptive neuron may not fire simply because it is excited by an afferent fibre which is active when a particular wavelength falls on the retina. Its firing may be partly dependent upon the matching of other aspects of the retinal image occurring in other parts of the template. A template implies both a functional separation of elements and an interaction between them.

The variation of weightings which gives Selfridge's model its power may have its counterpart in neural terms with a simple variation of the effective excitation sent to different locations; or even to selected fibre systems within a location (see item (v)). One may envisage a display of template categories extended across the cortex. If these separate categorial locations were illuminated by separate lamps which might be extinguished individually or in combinations, it will be clear that even an on-off system offers a considerable degree of "attentional" flexibility. The whole display may be illuminated as a giant template, or certain parts only, effectively reducing the size of the template matching. (Figure 2.5).

(iv) Sensori-motor representation

The template may include the representation of motor actions. It will in fact be proposed that this is the manner in which sensory and motor aspects of events are integrated. Assume that a figure is first fixated in its centre. A left-hand saccade brings a new input to be matched to previous inputs following a left hand saccade. A right hand saccade then produces another input to be matched. Let us suppose that the system of establishing neural templates includes a representation of the actions taken just before and just after the sensory input. Then when we fixate the centre of say, a triangle, a left hand eye-movement will result in a selection of all those templates containing both a certain kind of angle and a left-hand saccade. Similarly saccades will serve to bind bundles of templates into schemata. The similarity will have been noted with Hebb's account of the perception of a tringle

and with his postulation of phase sequences (Hebb, 1949).

(v) Superimposition of templates

Templates within a single modality need not require separate cortical locations, although it seems probable from the researches of Bazier, Robinson and Dow, and of Singer and Tretter, cited above, that they are organised separately. Topological mapping of sense organs onto cortex would be uneconomically used if each sub-cortical encoding had to be separately mapped onto a different projection location. Fibres from a particular sensory location but carrying different information might map into neighbouring points on the cortex. Different analyses of the stimulus, for example, for line orientation and movement, may be seen as separate templates sharing approximate spatial co-ordinates but separately registering coincidences of either line orientation or movement. An analogy would be the three-colour screen of a television tube, in which each small point is actually made up of three light-sensitive phosphors. Although functionally separate, the juxtaposition of different "feature detectors" would enable several template systems to be packed into the same single topographical mapping. Nor need local hierarchical network arrangements be entirely excluded from the scheme. Templates matching conjunctions of simpler feature detectors, e.g. line orientation and movement might be mediated by convergences from two of these superimposed templates. One may then envisage a stack of templates, one above the other in the thickness of the cortex, so that at the upper layers the particular co-ordinate of the template would derive from the conjunction of the templates situated lower down in the stack. Such an arrangement would accord well with Hubel and Wiesel's findings of a change of function as one moves through the cortical column (Hubel and Wiesel, 1963).

(vi) Innate and learned templates

The evidence from ethology is that some templates are innate. The question of the learned or innate nature of templates may be applied separately to the proposed divisions of the template system and to the actual stock of templates present within the library held in any particular division. It would seem improbable that the receptive field organisation disclosed by Hubel and Weisel and others (see Chapter 5) would be other than innate. On the other hand, there does not appear to be a similar clearly defined organisation of neurons registering the reception of complex environmental patterns, such as cubes, or

perhaps legs and noses. It is not necessary for our purpose to decide the extent to which the template system is innate or acquired, provided some neural mechanism to implement either source is allowed for. What does seem to be of interest is the manner in which initially provided (genetically and developmentally determined) templates would influence the later acquisition of acquired templates.

(vii) Temporal integration

For the template matching principle to be applied not only to pattern recognition but to neural integration generally, it is necessary that there should be some way in which the matching of templates is influenced by the conditional probability of events (Uttley, 1959). The inspection of a triangle provides an example. The matching of a left-pointing angle followed by a right-handed saccade, increases the conditional probability that the next input will be a right-pointing angle. One method of introducing the conditional probability has been mentioned, the inclusion of the saccade in the template representation. The temporal sequential bringing forward of appropriate templates within a particular division of the template system would be assisted by the fact that over the whole system input does not generally change simultaneously. These common contents of templates which differ only in one division would tend to bind them into bundles. Another possibility is that there is some division of the template section which is specially representative of temporal sequence. These aspects of the model go beyond our present concerns. They are mentioned becuase it seems important that temporal sequential and conditional probability considerations are borne in mind when the neural mediation of a template system is examined.

(viii) Summary

Aspects of matching of patterns have been discussed in terms of an extended template system. It is argued that a sophisticated template system is acceptible. Of course, the description of the system as a template system is somewhat arbitrary. Use of the template concept has been applied to showing how information might be spatially organised and selected by reference to the total configuration of the sensory and motor systems.

IV The problem of perception

As one turns from pattern recognition to its theoretical explanation, one finds that one has moved into the field of perception. There have been a number of computer models of perception in terms of information processing, list processing and computer programming logic (see e.g. Minsky, 1958, 1968). In the field of Artificial Intelligence there have been models which can interpret two dimensional photographs as a three dimensional world (e.g. Roberts, 1965) and a model which can manipulate and even describe such a world in English (Winograd, 1972). In terms of neurons, however, there have been no simulations which might be described as models of perception rather than pattern recognition. As may have been gathered from the discussion of pattern recognition in the preceding section, an attempt to do so would be premature. Before a convincing model of animal perception can be simulated, the basic problems of neural integration must be overcome in respect of many aspects of sensorimotor integration and temporal sequential processing.

Nevertheless, at a verbal level, neurological models have attempted to explain perception. Some have had particular influence (e.g. Kohler, 1942; Hebb, 1949). Hebb's model was simulated but not at a level which would demonstrate its perceptual implication (Rochester, Holland, Haibt and Duda, 1956). Models of perception have also been implicit in many theoretical discussions. A few points will be discussed here only to orient the reader.

A. Enrichment versus differentiation

Gibson and Gibson (1955) distinguished between "enrichment models" and "differentiation models". My model is undoubtedly an enrichment model, but may be a differentiation theory also.

Gibson and Gibson also distinguished "differentiation" theory from those "discriminative" theories, following Pavlov, which explain perception merely in terms of attaching stimuli to motor responses. The essence of "differentiation theory" is that "instead of assuming that stimuli can give rise to only meager sensations, let us now consider the opposite assumption: that the environment is rich in varied and complex stimulus information capable of giving rise to diverse, meaningful, complex perceptions . . . In the normal environment there

is always more information than the organism is capable of registering . . . As perception develops the organism comes to detect properties of stimulation not previously detected even though they may have been present. With growth and continued exposure to the world of stimulation, perception becomes better differentiated and more precise . . ." (E.J. Gibson, 1969, pp75-76).

The strength of this approach lies in the complexity and difficulty of the experimental data concerning perception (reviewed comprehensively in Gibson (1969), where the main theories of perception are also reviewed). The data has been difficult to fit into an overall explanatory system. Differentiation theory avoided a growing sterility of former approaches by concentration upon the relationship of the stimulation to the perception. As Postman (1955) stresses, the theory is essentailly one of psycho-physical correspondence, the perception reflects the stimulation. But as he commented: "The fact that the organism has learned to discriminate more qualities is the very fact that we need to explain". Some purely discriminative models do in fact learn to do this. They do so however only by attaching sets of equivalent stimuli to responses. One might say that discrimination theories explain discrimination learning but not perception, unless one is prepared to equate the two. Apparently the Gibsons are not, and it is this which makes their position difficult.

Consider for example the perception of a cup. The differentiation theorist would say, quite reasonably, that having had experience of cups which are hollow, and other cylindrical objects which are not, the observer learns to detect the information which differentiates the two. This is apparently intended to be different from saying that the observer discriminates the two, for the latter idea only refers to what the observer can do with a cup, whereas I take it that the differentiation theorist considers that the observer learns to differentiate the objects in other ways, for example, to differentiate a hollow cup from a solid cylinder. However, the perception of a hollow vessel is something different from the information indicating its presence. It is the conceptual (perceptual) knowledge that one can, say, put an object or liquid into it. This is an implicit knowledge of the nature of objects which, whether innate or acquired, has in some way to be attached to the discriminated information. If all the information for perception were in the stimulus presentation, we could differentiate a cup from a solid cylinder, but we would not

perceive the cup as something we could fill with tea unless we could actually see into its receptacle. Whether one considers that perceptual knowledge of the world is innate, developmental or individually acquired, it is still necessary to explain how that knowledge becomes attached to the discriminative attributes of the stimulus. This is absent from differentiation theory. Moreover, the distinction against discrimination theory may be overstated. Implicitly, for example, our perception of hollowness may be that of attaching responses to a representation of a cup: certainly, what we can do with a hollow vessel is close to our perception of the nature of such a vessel.

Gibson's theory was attacked by Postman (1955). Postman's defence of enrichment theory is not one which would be taken here. He defended a certain kind of associationism - "physiological associationism". Gibson (1955) reports him in the following words: "he (Postman) objected to the distinction" because the "kind of enrichment hypothesis exemplified by the addition of images to sense data is no longer of interest to contemporary associationists . . ." However, "sense data" is a philosophical not a psychological or neurological term. Few enrichment theorists would wish to maintain that the brain stores "images". The brain encodes information in such a manner that when that encoding substrate is reactivated something similar to perception - imagery - occurs. As Postman pointed out, associationists have long since left behind the introspective pursuit of bare sensations and accruing imaginal contexts. Such a model was not really an enrichment model but one in which "ideas" were attached to a "sensation". In the kind of model we have in mind not only are there no sense data, but sensation is only a term for convenient reference. In the model to be presented there are only perceptions, that is, activities in the apparatus which result from the interpretation of sensory input by the current state of the apparatus. Perceptions change in the presence of identical stimuli. The differentiation theorists say that it becomes "better differentiated and more precise". In the enrichment model this is partly because greater selection is made from the environmental information, and in this respect the enrichment model does not differ greatly from the differentiation model. However, the enrichment model says that the selection is on the basis of what is already held in storage: it also says that something may be added to the information at present coming from the stimulus. The differentationist would one

assumes, disagree on just this pioint, and say that what one perceives still reflects only the input information.

Possibly, there is a different usage of the term "perception". Perception is not an all or none event. It is hard to see that in cases of fragmentary or degenerate information, the observer does not construct what he perceives. Oatley (1978) has a photograph of a student seated at a table in which only the smallest fragment of the chair frame is visible. Yet we have no difficulty in discerning a chair. Presumably the differentiation theorists would wish to maintain that we do not "construct the image" of a chair from the fragment, but that we differentiate the small fragment as part of a chair. Probably the differentiation theorist would not wish to maintain that we cannot describe the remainder of the chair, e.g. that it is formed of bent tubing and probably has a canvas seat. He would be forced to say that this knowledge is cognition, not perception. In doing so he would be drawing a hard line between the two. If we enter a room and perceive a telephone on the table, and then discover that it is not the telephone at all but a cardboard box placed there by a mischievous experimenter, we would then have to say that the original perception of the telephone was a false cognition!

One may agree with the differentiation theorists that given full inspection in adequate circumstances, perception reflects the stimulation. It is in circumstances where evidence is fragmentary or degenerate that enrichment becomes an important part of the process. It is just this ability to deal with fragmentary evidence which would seem to point to the intellectual and creative aspects of human psychology. Enrichment occurs "by association" because stored information gives us access to further information about the stimulus. Enrichment theory does not wish to deny, as Gibson suggests it does, that the "environment is rich in stimulation" or that stimuli "can give rise to only meager sensations".

Hierarchical implications of differentiation theory are summarised in Gibson's (1969) conclusions: "First, that some sort of memorial representation of the distinguishing and invariant features of a stimulus object or event is involved in differentiation of it from other objects or events in its absence. Second, that this representation is based on prior discovery of distinctive features of the

object, or abstraction of invariant features of a pattern. Discrimination is thus prior to recognition. And third, that both discrimination and representation are prior to production. In other words, perceptual learning . . . is a requisite for all three processes, discrimination, recognition, and production, and the processes form a kind of cognitive hierarchy".

As thus stated, the enrichment theorist would have little to quarrel with differentiation theory, except perhaps for the possible implication that the hierarchy defined is a neural hierarchy. In the kind of extended template model described earlier, all three of Gibson's postulated processes occur as a single process. Moreover, the template model, at least as embodied in my neurological hypothesis, is an enrichment model. The priority to which the differentiationists have addressed themselves is a logical not a neurological priority.

Considerable time has passed since the controversy between differentiation, discrimination and enrichment theories was a central issue. The matter has been discussed at some length because it may assist the reader in what to expect of the neurological hypothesis to be developed here.

B. Sensori-motor schemata

The work of the Gibsons served the purpose of drawing attention to many aspects of the sensory input which the brain is able to analyse. An example is the change of grain size with depth, as when we observe a beach of pebbles extending into the distance. Partly as a result of the work of the Gibsons, there are a number of excellent contemporary discussions of the mechanisms of perception, not necessarily sharing their viewpoint (e.g. Neisser, 1967, 1976; Zusne, 1970; Gibson, J.J., 1968). Accordingly it is unnecessary to consider such matters here.

It is perhaps worth stressing that sensori-motor integration and temporal integration have been seen as a key to perception by writers of many different viewpoints.

Corcoran attributes the first appreciation of these aspects to Hebb. "Pre-Hebbian theory" was dominated, on the one hand, by the Gestalt psychologists, whose implicit model of the nervous system was of a sensitive, but unstructured mass, and on the other hand, by the "learning theorists" who thought in terms of deterministic sensory-motor

connections in the brain'. Hebb's neurological theory will be mentioned again in Chapter 4. The psychological importance of the theory may be seen as at least partly due to its emphasis upon integration over both time and space. As is well known, his explanation of perception depends upon successive recognition of discriminated features. These were integrated into a perceptual unity because the neural process itself was conceived as a diffuse process in which "cell assemblies" were bound into "phase sequences".

Hebb's hypothesis is supported by the data concerning eye-movements, at least in normal perception. Scanning eye-movements recorded during inspection of visual displays are closely related to specific stimulus characteristics, e.g. eyes and nose in a face. "When looking at pictures the observer's eyes fixate more frequently on those features which are actual or potential carriers of information". (Zusne, 1970, referring to work of Buswell (1935) and Yarbus (1967)). He continues: "Information here is not information measured in bits or physically the most complex portions of the picture. It is rather the meaning and significance of a particular portion of the picture to the observer which attracts attention. Thus, in a painting of a hunter in a forest, it is not the (physically) very complex branches, bushes and grass that attract attention. Rather it is the very small, detail-less inconspicuous figure of the hunter that most fixations center upon". Such facts present a challenge to neurological theories of perception.

Gibson J.J. (1968) stresses the longer term possibilities of temporal integration: "The visual world as I once described it has the property of being stable and unbounded. By stability is meant the fact that it does not move when one turns his eyes or turns himself around. The phenomenal world seems to stay put, to remain upright and to surround one completely" (p253). "The man in the dark who gropes with his fingers, like the insect who gropes with his antennae, seems to get a succession of contact stimuli, not a simultaneous pattern of them (p251).

Piaget, in his various studies of the development of a concept of space and time in the infant, developed the concept of a sensori-motor schema. Neisser (1976) feels the necessity for such a concept. A neurological substrate for such a schema is one of the primary aspects of the present model. Piaget developed his concept into his "centration theory" of perception (Piaget, 1969). The relationship of this theory to the

present model is too distant to pursue here. Relevant to our present concerns are his earlier observations of children from which he developed a concept of the nature of the phenomenal object. To the infant the face which appears from no where in the game of peek-a-boo has not been in existence since it disappeared behind the screening article: it appears and disappears. The infant has to be older before a sweet hidden under a cushion is the same sweet as is found there later. Objects present successive transformations of shape, retinal size and position. These successive transformations are embedded in sensory motor relationships. The integration of appearances into objects is acheived by conservation. Piaget does not indicate how conservation is mediated neurologically. Such an answer would be extremely useful, as it would bridge the neurological gap between perception and intelligence.

Neisser in his later work (1976) on cognition elaborates Piaget's concept of sensori-motor schemata. He makes two proposals of especial relevance to the model to be presented here: (i) "Because schemata are anticipations, they are the medium by which the past affects the future . . ."; (ii) In addition, however, some schemata are temporal in their very nature. When an object moves, for example, continuous and complex changes take place in the optic array . . . We need not suppose that this expansion is picked up as a series of discrete and individually anticipated frames . . ." (op.cit. p22). In my model, it is in a way, picked up by such discrete frames: these are the neural templates. The frames are however capable of being bonded into temporally integrated sequences. Moreover, change itself can be embodied in a single discrete frame, as when the activity of a movement detector is registered. To record the continuous and complex changes which take place it is not necessary to propose that the succession of frames registers everything that occurs. Even within a single frame, the template may be only fragmentary. A template of an 'A' can be discriminated from that of a 'B' if only small fragments are present. In a temporal sequential record, one or two frames selected from a sequence may discriminate a particular sequence from another. (Compare Bledsoe and Browning's pattern recognition model, 1959).

C. Information process analysis distinguished from neuroscience

Neisser (1976) expressed growing conviction in an information processing view of perception. The success of Artificial Intelligence models of

perception supports that view. Nevertheless it is necessary to be aware of the inherent limitations of that approach. Except in so far as such models borrow from psychology and neurology, they are neither psychological models nor neurological models. Information processing and systems analysis are conceptual abstractions common to many different kinds of apparatus, biological and artefactual. As such, like mathematics and logic, they are one way of furthering our knowledge of brain function in terms of an analysis of the way the components of the brain and its assembly into a system, handles the information passing through it. Such analyses are powerful tools in elucidating the psychological and neurological functions of brains. They may be alternative languages in which psychological and neurological hypotheses may be expressed, but they are not in themselves such hypotheses.

There may be a level at which one can speak of an information processing or systems theory of a psychological or neurological process. That level is not the psychological or neurological level, but a more abstract one which the theory will share with quite other entities than the animals with which the psychological or physiological theories are concerned.

One or two examples may indicate what is intended. A television set may be described as the coding of a spatial set of variables into a temporal sequence with an inverse process at the receiver. Such a description gives no indication of cathode ray tubes, of the timing of circuits, aerials, and electro-magnetic radiation. Compare a modern television receiver with Baird's original set, connected by wires instead of radio. From an information processing viewpoint the two receivers are in principle identical, except perhaps in the quality of performance and the kind of repair man one would call in. The analogy may give a rather fair picture of both the strengths and weaknesses of the information processing approach!

There are two further weaknesses of the information-processing approach to psycho-neurological problems. Two pieces of apparatus may acheive the same informational results by completely different means, according to the particular strengths which the different hardware possesses. It is not that the hardware is the same but the software different. It is that one program is quite different from the other. Compare a pendulum analogue clock with a quartz digital. Both have as fundamental a

natural process with a fixed cycle time. From then on their information handling is entirely different. One counts discrete changes. The other controls an analogue motion. Of course, once can still see analogy at abstract level: if not they would not both be clocks. But knowledge of the pendulum clock does not help much in understanding the quartz digital clock. This is the trouble with comparing the behaviour of the brain with that of the computer. Even if the information handling is similar, each may acheive its end in very different manners.

The other weakness is that because information theory is at an abstract level, it is easy to mistake the knowledge, which we as human beings possess of the world, with the knowledge which the apparatus we are proposing would possess or be able to obtain on its own account. This is the controversial field known as heuristics. An example, is the computer model which can interpret a two-dimensional line drawing as a three-dimensional world (Guzman, 1969). The model interprets conjuntions of lines as evidence of three-dimensional aspects of the world. For example a conjunction of three reasonably equally spaced lines is a "corner". A T-junction implies that a nearer surface is partially obstructing the view of a more distant object, and so forth. The knowledge of such coincidences between the conjunctions of lines and objects of the three-dimensional world is not discovered by the program itself: it is supplied by the programmer who has applied his own intelligence in discovering such relationships. This can be defended upon the basis that evolution has undoubtedly furnished the animal or Man with a similar knowledge. Such an answer does not provide any hypothesis of how the knowledge of the relationship is mediated in the brain. Even in relation to the psychology of the process, it is doubtful whether the ability to recognise the three-dimensional meaning of conjunctions lies at the base of our ability to interpret line-drawings. It seems quite possibly the other way round. We immediately perceive two-dimensional line drawings as three-dimensional objects, and because of that we are able to infer the relationships between conjunctions of lines and three-dimensional objects.

The intention of this discussion is not to call into question the value of Guzman's model in interpreting perceptual processes. In fact Oatley (1978) is able to show some close analogies between a similar model (Roberts, 1965) and neurological processes. The point is to suggest that information processing models are an adjunct to but not a direct approach to understanding brain function and psychological attributes.

V Conclusion

A number of topics have been discussed with a view to orienting the reader to the psychological and neurological ideas which underly the neurological model to be presented.

Pattern recognition as it appears in neurological models has been contrasted with the perception of gestalt pattern as an inherent aspect of ongoing sensori-motor integration. It was argued that for pattern recognition models to become models of perception, it would first be necessary for models of effective action to be simulated. Accordingly in the present project, the simulations would be directed towards modelling motor action rather than pattern recognition.

The template analogy was discussed and it was argued that an extension of the template concept can be used to explain many aspects of perception and motor interaction. There are, however, limitations in the application of any simple analogy to the problems of brain function. Especially in relation to temporal sequential organisation and to the formation of schemata, the template matching system is necessarily superseded by considerations of neural mechanisms. Nevertheless, the template concept will, it is hoped, assist the reader in following the main neurological ideas of the model. Essentially the model proposes how extended neural templates may be constructed and manipulated in terms of simple neural connections.

Differentiation theory has been discussed, partly because that theory was especially antagonistic towards enrichment theories, and partly to give an opportunity to suggest certain implications of the model which will be proposed; namely, its mediation of the constructive aspects of the animals interaction with its environment. Sensori-motor schemata and information processing models served as further opportunities for orienting the reader to the approach taken here.

CHAPTER 3

SOME FURTHER PROBLEMS OF NEURAL INTEGRATION: MOTOR ASPECTS AND THE GENERAL CO-ORDINATION OF FUNCTION

I The problem of motor activity

In the previous chapter emphasis was placed upon the role of motor activity in establishing sensori-motor relationships which contribute to pattern recognition. This is of course only a limited aspect of motor activity. Motor activity is the means by which the animal satisfies its biological needs, physiological and psychological. The considerations discussed in relation to pattern recognition must be extended to these functions. Accordingly, the primary problem which a neurological model must deal with is that of sensori-motor integration, in the sense of showing how the action of the system on the environment produces sensory feed-back which in turn controls the action of the system on the environment. Closely associated with this problem is the problem of motor learning, and in turn this is concerned with concepts like reinforcement and motivation.

The literature of motor activity is perhaps more extensive than that of any other branch of psychology. In contrast, the simulation of neural motor mechanisms does not appear to have been attempted so far in anything more than a formal labelling of output as "motor". (Some simulations with stronger motor representation will be described in the next chapter). In attempting to take the simulation of motor organisation one step further, the present project should not be taken as putting forward a unified theory of motor learning, such as have been proposed by Hull, Skinner, Sheffield, Konorski, Deutsch and others. The simulation is used here as a tool for developing a theory, rather than the presentation of a theory which has been worked out beforehand. The simulation starts with the representation of the simplest learning situation in which the full cycle - sensori-motor-environment-sensory - can be seen, the classical conditioning situation. It then attempts to build upon this base. In doing so, the modeller naturally has in mind the more thorny problems, but the main aim remains the demonstration

that the theory of neural integration can cope with the control of motor action. Naturally, one should bear in mind the level of representation in the simulation.

The following discussion of one or two topics may assist in orienting the reader.

(i) TOTE

Miller and Pribram suggested that behaviour is organised in an hierarchy of units of a negative feed-back kind. These units compare the input to some criterion within the organism. A mismatch gives rise to a response, and the sequence continues until a matching occurs, i.e. Test-Operate-Test-Exit. This is a powerful concept in explaining goal directed behaviour without necessarily involving the notion of primary goals, such as food, dominating all behaviour. Hinde (1970, p.630) considers the hierarchical organisation of TOTE units proposed by Miller and Pribram as important "for this permits the possibility that behaviour can be goal directed (in some sense) at one level and not another". Whilst agreeing with Hinde, one might draw attention to the point earlier made with respect to hierarchical organisation in connection with perception. A system may be hierarchical in action without being reflected in an hierarchical neural organisation.

A difficulty with the TOTE hypothesis is to find a plausible neural implementation. Whilst Pribram (1971) gives some notional examples of neural structures which might be organised on a Tote basis, these only apply to simple feed-back adjustment loops, where a clear physiological variable, say a tension in a muscle spindle, is to be balanced. What is required to explain motor behaviour is a TOTE mechanism with psychological criteria of success. An attempt is made in the present model to show how an image of anticipation could act as an internal criterion which the system tries to match. Unfortunately, it must be confessed at this point that the simulation of the present model has had to stop short of actually demonstrating this capacity. There is, however, some demonstration of the capacities required to implement such mismatch-match tests. A simulation,

programmed but not yet run, will be described to show that the idea, although rather ambitious in terms of present resources, is not nebulous.

(ii) Classical Conditioning concepts, instrumental conditioning concepts, and more generalised concepts of motor action

Research on motor action has grown very much from Thorndike's early experiments and from Pavlov's demonstrations of Classical Conditioning. A considerable advance was the isolation of the concept of operant or instrumental conditioning by Skinner, Konorski and Miller, and Hull. Accordingly, many neurological models have aimed at demonstrating the neural implementation in simple form of those concepts. A number of models of the control of motor action are in effect pattern recognition models the outputs of which are categorised as responses (e.g. those of Kilmer and his associates, see Chapter 4 and Chapter 1).

A useful exposition of the problem of motor learning in neural terms is that of Burke (1966). He sets out the main forms of learning in terms of conditioning theory, i.e. classical, operant conditioning, drive and reinforcement implementation. His neurological "models" involve very few neurons. The idea would seem to be that these models could be duplicated and combined. They have the merit of making the requisite conceptual components of learning very clear. The models typify a certain approach to behavioural function: they assume that the mechanisms within the brain reproduce the same overall structural relationships which can be seen in overt function. Each neural network is really a model of the conditioning process itself. However, the components are without the complexities which surround the concepts they represent, e.g. reinforcement and motivation (see below). Moreover, the problems of neural integration in combining such models to deal with behaviour is not shown. Their value remains, therefore, as conceptualisations of the factors which enter into conditioning rather than a neurological model of nervous function.

Two process theory

One of the controversies which has occupied the theoretical literature is whether there are two kinds of motor learning or one kind. Put in its simplest terms, the controversy is whether instrumental learning (operant conditioning) represents a different kind of process from classical conditioning (Miller and Konorski, 1928; Skinner, 1938;

Mowrer, 1947). This issue is discussed in more recent perspective by Rescorla and Solomon (1967) and Mackintosh (1974). The distinction has been made in a number of ways, some of which imply no more than a reflection of the situations in which they are observed. Essentially the two-factor theorists maintain that in classical conditioning learning occurs between stimuli, and in operant conditioning between stimulus and resonse. Another way of putting the distinction is that the classical conditioning principle of stimulus substitution states that the occurrence of a reinforcing stimulus causes the responses elicited by the reinforcer to antedate its delivery, and those responses become attached to an accompanying or antedating stimulus. The principle of operant conditioning, following Thorndike's law of effect, is that the occurence of a reinforcing stimulus will strengthen any response which happens to precede its delivery (Mackintosh, 1974, pp. 125 ff).

The resolution of the controversy may depend upon further research clarifying learning and conditioning mechanisms generally. Where a neurological model may be useful is in demonstrating that such issues may not be quite as clear cut as they seem. Neither Classical Conditioning nor Instrumental Conditioning might represent fundamental modes of learning as such. They might merely be situations in which certain aspects of the system appear more prominently. One may bear in mind that over and above the basic capacity of nervous tissue to mediate learning, there may be superimposed mechanisms which have developed to deal with the more specific requirements of a more sophisticated nervous system. Thorpe (1963) suggests that mechanisms of learning may differ in different species. Tolman (1932, 1959) related performance in standard learning situations to six capacity laws, including for example, the ability to form an expectancy of a stimulus and to act in accordance with it. Certain conditions, e.g. the classical conditioning paradigm, may give prominence to one aspect of the process. The present model suggests that classical conditioning does not represent a behavioural mechanism at all, but rather allows a feature of nervous processing to become observable. The possibility is that the Classical CR is not intended by nature as an act upon the environment but is a side effect of the build up of excitation in a system directed towards learning to anticipate events.

Because of the importance of classical and operant conditioning in the

literature, and because well documented experimental situations provide a test of the performance of the model, these paradigms are the subject of the simulations reported here. It is not, however, intended that the simulations show "classical conditioning" or "operant conditioning". They are intended as demonstrations of how the model would act in those situations.

(iii) Broader approaches to motor organisation

Many workers have considered that the conditioning approach fails to illumine a wide area of brain function. Amongst views still related strongly to the laboratory, one may mention those of Tolman (1959). He showed that reinforcement may not influence learning so much as the performance of learning. He suggested that cognitive maps guide behaviour. The cognitive psychologists have developed concepts of behaviour which workers in Artificial Intelligence have expressed in more definite algorithmic terms, although such implementations are rarely in forms which provided models of neural integration.

The ethologists following Lorenz and Tinbergen (Schiller, 1957; Hinde, 1970) have developed concepts which offer some promise of being expressed in neural form, e.g. imprinting, innate releasing mechanism, search image, hierarchy of behavioural organisation. The present project does not reach a stage where the broader aspects of behaviour as described by the ethologists could be discussed except in general verbal terms. The simulations do, however, contain features by which an ethological description of behaviour might be implemented in some model of the future. For example, the basic formulation of the model contains something analogical to the Innate Releasing Mechanism, and which could provide the substrate for a search image.

Although the limitations of the simulations necessarily tie them to the representation of emission of single notional abstract "responses", it is possible to see the organisation proposed as the basis for a more complex explanation in terms of ethological descriptions. Behaviour may be described either in relation to its effects on the environment or in terms of its implementation by muscles or other effectors. A neurological model is naturally concerned with the latter aspect, but it must bear in mind the eventual application of its explanatory concepts to broader descriptions. Hinde (1970) distinguishes three sources for the control of movements (i) environmental stimuli (ii) proprioceptive

stimuli and (iii) co-ordinating mechanisms within the central nervous system. At the basis of the co-ordinating systems are fixed action patterns which are able to be oriented towards the environment: superimposed upon this system is that of skilled movements. Underlying these mechanisms may be seen a propensity for learning. In devising a neurological model of behaviour it seems necessary to provide for a combination of influences upon the emission of muscular and other effector movements. These influences must express the environment, the proprioceptive information which guides the development of the muscular actions, and the factors within the organism which cause behaviour to be generated towards the environment. Behaviours like walking are oriented fixed action patterns. They are influenced by learning in that they occur and are shaped accordingly when, for example, at a given time the animal moves towards a food bowl where it knows it will be fed at a particular time. Yet the act of walking must be guided by feed-back information from the paws (Gray and Lal, 1965, Shik and Orlovsky, 1976). Such problems of organisation cannot be dealt with in a simulation at the level attempted here. However, the mechanisms for response emission and control must be capable of being developed to cope with such behaviour.

At another level, the cognitive aspects of behaviour seem to be dependent upon anticipatory capacities of the brain. Expectancy is an aspect of the theory of neural integration presented here which arises from the possible build up in an analogue fashion of excitation within a distributed network.

II The problem of reinforcement

The incorporation of "reinforcement lines" into many pattern recognition models suggests that the concept of reinforcement is not restricted to motor learning. Moreover, reinforcement is involved in questions of motivation (see next section). The absence of a representation of actual neural mechanisms of reinforcement in simulations of neurological models makes it a matter of some importance in the present project. The aim is to incorporate in a model at the level of neurons, some means of showing how learning can be dependent upon what happens to the animal following its actions upon the environment. This capacity should

appear in the model without some external intelligence deciding for the model whether the response was correct or incorrect.

The concept of reinforcement is theoretically controversial. For the purpose of a neurological model one can, without committing oneself theoretically, make a start at representing the reinforcing effects of "reward" and "punishment". Certainly the satisfaction of physiological needs and the avoidance of physiological injury assists learning. The strategy of representing these mechanisms does not exclude the possibility that there are other ways of supporting learning. Certain aspects of the mode of neural integration in the model to be presented here, suggest that even at the stage of simulation so far acheived, the model has some relevance to these more general aspects of motor learning. In particular, the model shows a certain ability to anticipate events (see below).

A. Emission of responses: innate and learned control

There is considerable doubt whether physiological reinforcers, such as food and shock, act quite in the way assumed by theories of behaviour like those of Skinner and Hull. Premack (1971) has demonstrated that the role of a customary "reinforcer", candy, can be exchanged with a customary "instrumental response" - manipulation. "Running" is often used as a reinforcement for rats. The reinforcing aspect of a response may depend upon the likelihood that the response would occur at this time spontaneously, in relation to the likelihood of the response it "reinforces". To deal with such problems it would be desirable to have some clear neurological picture of how responses are emitted. This is an aspect of neurological modelling which appears to have been neglected.

Skinner distinguished operants from respondents. Respondents appeared when specific stimuli occurred: operants appeared in a more spontaneous manner to a much wider range of stimuli. This distinction is in the same direction as Premack's analysis of reinforcement in suggesting that all responses are not solely dependent upon environmental stimulation. Many actions appear to be innately programmed and under the control of stimulus feed-back. Walking is such an action. Yet the occasion and emission of such an action appear to be subject to learning, as when a dog walks to where it expects its master to appear. In many so called learned actions it does not seem to be the form of the actions

themselves which is modified, but the occasion of release, the orientation, the co-ordination and application to the environment.

Neurologically, the emission of responses has not been clarified. Theoretical behavioural approaches refer to probability or occasion of emission. One might represent such theories in a model by a random number generator. Reinforcement would alter the distribution from which the random number was drawn. Motivational states would also alter the distribution. One might label input patterns with an operator which indicated their probability of eliciting a response. Such a model would not, of course, be a neurological model at all.

An unexpected difficulty encountered in devising a neural mechanism of response emission was the conflict between the learned control and the spontaneous emission. This problem may be specific to the model to be presented here. That seems unlikely, as response emission seems to be necessarily under dual control. Even if it is assumed that all responses are initially emitted in response to specific stimuli, e.g. bar-pressing, those responses are brought under learned control. The question which must be answered neurologically is how the innately controlled emission of response is able to give way to learning. This question, which is not raised in most neurological models of action (e.g. Pringle, 1951), is far from simple. If learning takes precedence over innate emission, then one will learn not to employ a response in a given situation, and perhaps employ some other response. Suppose the contingencies are changed. The animal will then be unable to experiment by emitting the response required for success, because it has learned not to do so. This takes precedence over its innate power to emit the response in trial and error.

Something similar was found in the simulation of the model. In the model, if learned responses were given precedence, then the model would continue to emit that response even if changed contingencies made it unsuccessful. The continuing failure eventually extinguished that response. However, since other responses had already been extinguished by the earlier learning, the model was reduced to a standstill. The model could be allowed to escape from this learning trap by giving innate emission occasional precedence. However, then the opposite kind of error occurred. Even when the model had learned to produce a correct response, it would occasionally produce some other response. This was because the latter spontaneous response had been given greater precedence

at that particular time. Naturally, the difficulty would not have arisen if the model could have had unlimited time to learn its task. The original learning would have dissipated and released spontaneous responses. However, the latter remedy merely serves to underline the problem of control.

It is appreciated that there are other ways than those mentioned by which a nervous system could escape from its own learning. One way would be the operation of a reactive neural inhibition which suppressed a response when it had been repeatedly emitted. However, the time scales for behavioural inhibition would seem to make it unlikely that it could be mediated simply by the release of inhibitory transmitters in reaction to excitatory spikes.

Some progress was made towards a representation of motor emission which could deal with problems of this kind, and which could be related to Premack's analysis.

B. Key trigger patterns: preset patterns

Whatever may be the eventual resolution of the problem of reinforcement, one way the animal's nervous system can monitor the appropriateness of its actions is by reference to innate "wiring" which tells it so, e.g. food acting upon the chemical and metabolic system to signal that its actions have been favourable; electric shock informing it that the situation is unfavourable. Evidence, principally of specific releasers from ethologically oriented psychology, suggests that innate recognition of stimuli can be far more complex than can be explained by a simple stimulus-response connection, analogical to the spinal reflex. Neural integration must handle these innate or developmental recognitions. More than this, if such recognition processes enter into reinforcement in the sense that "food" and "sex" are reinforcers, it must be shown how they act in the required manner to support learning.

Konorski (1967) deals extensively with these problems in terms of conditioning theory. Eventually his neural solution must depend upon the hierarchical neural network system which constitutes his neurological model. His solutions, although of considerable psychological interest are not developed in terms of neurons, only the general principles of the neural model being given. A solution closer to that of the model proposed here is that of Pringle (1951). In a system of neural oscillators capable of forming pattern configurations, he

proposed that "key trigger patterns" are present. When a sensory pattern matches such innate patterns, motor action is triggered. A similar mechanism operates in my model. It is not restricted to the mediation of physiological events. It is capable of being interpreted as the Innate Releasing Mechanism proposed by Lorenz (1957). It is capable of providing the innate "sensory analysers" porposed by Deutsch (1960) which lead to primary links with motivational excitation. (See below).

C. Cognitive confirmation

A question related to reinforcement is whether an animal can judge whether it has succeeded in producing the effect which it has "intended" independent of predetermined satisfactions or unpleasantness. To human beings, this seems a natural capacity. When we try to draw a circle, we know to what extent we have succeeded. This may be a learned capacity obtained by comparing the result of our handiwork with a remembered circle. Moreover, we can say what we expect to happen in all sorts of circumstances. Our ability to do so does not seem to be related to whether we wish what we expect or not.

At the neural level, key experiments show changes of processing of "novel" in contrast with "habituated" stimuli, reflecting a control of orienting response (Sokolov, 1960; Anokhin, 1961). The implied capacity of the nervous system to lay down and label a "neuronal model" (Sokolov, 1960) might contribute to the organisation of behaviour itself. In Pribram's illustration of the TOTE mechanism, of hammering nails, an inner criterion is required by which we recognise that the nail is flush with the surface. Deutsch (1960) would say that the input from the flush nail affects an "analyser" which switches off the "link" providing excitation to the hammering action. The "link" (flush nail) was earlier connected to excitation deriving from a primary link (innate analyser) because it was excited just before an analyser one step nearer to the primary link. This theory suggests that we cease hammering to proceed to the next step in a sequence approaching some innately satisfying reward. According to Deutsch's account, the afferent pattern switches off drive excitation and connects that excitation to the next link and so on to the primary link. We may note in passing that this does not explain why the flush nail ever became the link for ceasing hammer action.

However, the main question for a model of neural integration is simply the neural implementation for the setting up of the inner criterion and its action upon motor movement. Does such a criterion have to be tied to a primary source of motivation? Deutsch's model says that it does. Yet a device for hammering pegs into a board is a well known toy, which presumably provides its own reinforcement.

One way in which an animal might judge the consequences of its actions is in comparison with some kind of expectation. Such an expectation is implied not only by Sokolov's experiments, but later work by Grey Walter on the CNV (Walter, 1969). Konorski (1967, p.290) says: " . . . of particular importance is the old concept put forward by James, stating that the chief, if not the only, source of voluntary movements is provided by their images (ideo-motor action)". Konorski does not give the neural implementation of such particulars beyond the general statement of his hierarchical network theory (see Chapter 4), but he clearly has in mind the concept of conditioning as an anticipation of events. Neisser (1976) considers that sensori-motor schemata act as anticipations.

As will be reported, the present model is able to anticipate events when a familiar sequence is repeated. This anticipation takes the form of a rise in excitation in the neural records of the event which is about to occur. The anticipation can act both as a prescription for the next response, as proposed in James' notions of ideo-motor action, and also as an inner criterion. According to this explanation, hammer knocking would stop because the anticipated result of further hammering is "no change" in the position of the nail.

Unfortunately, it was not possible to simulate the model to a stage when these concepts could be definitely demonstrated. This was mainly due to the delays in getting an efficiently operative model, attributed to difficulties solved by rhythmic mechanisms. It was possible to demonstrate that the anticipatory excitation could control instrumental actions. Further simulations were programmed which were intended to show how these anticipations could confirm the environmental success or failure of an action, independently of primary reinforcement. Unfortunately, these were unable to be run before programming ceased. Accordingly, the discussion of such mechanisms rely upon the ability to express the ideas in a computer program, supplemented by the considerations which make such a view plausible or otherwise.

III The problem of motivation

Like the problem of reinforcement, with which it is closely associated, motivation has appeared in simulations only in a minimal form. In contrast the subject has occupied a large part of the theoretical literature and has given rise to deep-seated controversy. The more difficult aspects of motivation are beyond present prospects of simulation. However, within a system of neural integration one would wish to see the general direction in which motivational problems would be handled.

A number of neurological models in the literature, although not simulated, give a fairly specific account of the neural implementation of motivational factors. These give some idea of the problems involved.

(a) Spinelli's neural implementation

Spinelli's (1970) neural hologram model is described in Chapter 10. He suggests that the neural networks mediating particular input patterns might be completed by adding networks representing reinforcement. In this way an input pattern leading to reinforcement can be discriminated from one which does not. Further, the networks might be biassed by adding inputs representing motivational states. Because of the special character of the hologram model as a diffuse distributed memory network, this is not quite the same as merely adding formal drive and reinforcement inputs to the sensory pattern. The drive inputs would make some of the memory networks more easily stimulated than others, a bias effect. My model is also a diffuse distributed network model, and drive inputs would have a similar effect. Although the neural hologram model was simulated, the reinforcement and motivational proposals were not. In any case, the simulation was not in behavioural terms. Accordingly, it is not possible to say what the actual effects of the motivational inputs would be. In the simulation of my model, the introduction of motivational sources of excitation had a most disturbing effect. In particular it raised two problems never quite resolved. However, some progress was made towards their solution (see below).

(b) Olds' neural implementation

Olds (1969) proposed a neurological model of motivation related to the neuro-anatomical structure and input output relations of the Hippocampus. Research had revealed hypothalamic centres, the stimulation of which can support eating and drinking. Olds discovered neurons in the Hippocampus which were active while the animal "expected" food, and inactive while

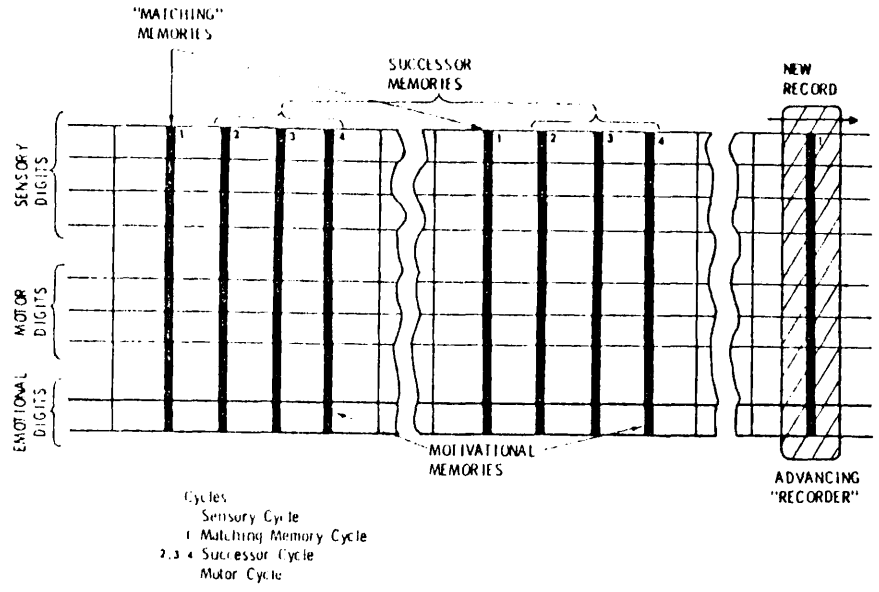

Figure 3.1 (Reproduced from Olds, 1969) The basic wiring required by Olds' motivational model of the Hippocampus, based on the cross-wire effect of the core store of a digital computer.

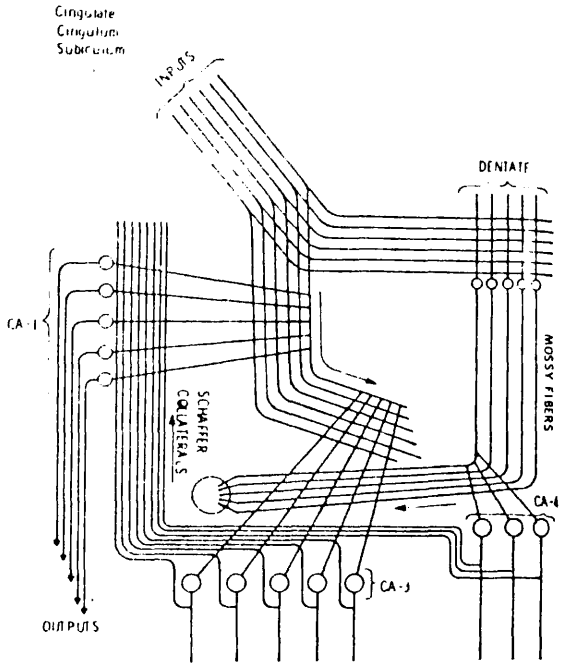

Figure 3.2 (Reproduced from Olds, 1969) Olds' interpretation of the neural structure of the Hippocampus as a neural mediation of his model. It may be noted how the inputs from the lines representing the perforant pathway (marked inputs) sometimes bifurcate to innervate the dentate on the one side and the CA1 and CA3 fields on the other.

the animal "expected" water. He found other neurons with opposite characteristics. Moreover, he found neurons in the reticular system which were correlated with the neurons in the Hippocampus. This relationship was not permanent, and Olds attributed it to a memory function.

The basic concept of Olds' neural model (Figure 3.1) is modelled on the wiring of a digital computer hardware memory, such that the magnetic cores at the crosswires are analogical with synapses. The "digit" lines represent emotional, motor and sensory information. Successive lines of the grid represent the patterns which occur in temporal sequence, i.e. "successive memories". To deal with the problem that computer cores retrieve information by an addressing system, whereas the brain recognises patterns, Olds suggests a scanning mechanism which identifies similar patterns. "The memory element itself has to be a single line, with all its sensory elements crossing its afferent side, or its dendrites, and all the motor elements crossing its efferent side or its axons. The structure required is that of the Hippocampus". (Figure 3.2).

Although superficially very different from my model, the two share many assumptions. However, Olds' model is not sufficiently developed neurologically to see exactly how it would work. For example, the scanning mechanism is not defined. Neuronal integration is restricted to an analogy with the crosswire magnetic cores. Each unit would therefore be isolated, holding one bit of information defined by the selective action of axons and dendrites. It is not clear how this would operate as a general neural integrative system. Olds' interpretation of the histological structure is tight and convincing. However, Kilmer and Oliński's model (1974), described in a later chapter, also utilises the detailed layout of the Hippocampus to support a very different hypothesis.

Olds proposes: "if there was a reward among the events that followed last time, the biological memory would yield an output which would instigate remembered behaviours . . .". A psychologically similar strategy was adopted in my model. However, the difficulties encountered were with just those aspects which Olds does not define: (i) how scanning looks forward through the learned records; (ii) how the system takes a record of sensory and motor patterns, and (iii) how the motor output is switched on and off.

Olds' model is perhaps best regarded as summarising the main advances in understanding of the neurology of motivation. It suggests that there is a centre where sensory, proprio-ceptive and motivational factors meet and are made the basis of motor decision, and that this centre is the Hippocampus (Adey, 1970; Vanderwolf, Bland and Whishaw, 1973; Kilmer and Olinski, 1974). The model co-ordinates much contemporary research on the pharmacology of the rewards system of the brain, including his own work on brain stimulation (see for example, Olds, 1976; Deutsch and Koopmans, 1976; Mogenson and Philips, 1976; Margules and Margules, 1973; Sigg, 1971; Isaacson, 1974).

A particular difficulty with Olds' model as one of neural integration is that it relies on the special histological arrangements of the Hippocampus, i.e. crossing fibres, for its function. It seems unlikely that the properties of integration are peculiar to a single location with somewhat unusual anatomical structure, the Hippocampus. More likely the neural organisation of brain tissue provides a general integrative process, and the Hippocampus is a special exploitation of that process.

Olds' model is essentially a computational decision model, which issues motor instructions on the basis of past experience. Many workers consider that motivational states drive behaviour dynamically, and modify perceptual as well as motor processes to a degree dependent upon the current state of the motivational system. There is little scope in Olds' model for such pervasive effects. In this respect, it may perhaps be constrasted with Spinelli's neural hologram model, in which motivational input biasses the operation of the system. Nevertheless, Olds' model firmly indicates the neural relationships which have to be provided between reinforcing, motivational and sensory influences. In particular it makes the point that a reinforcing event may not occur immediately after an action but may have to be searched for in the "successive memories".

(c) Response suppression as a learned response

Much motor learning seems to be learning not to do, rather than learning to do. Such learning is implicit in most behavioural experiments of instrumental learning or operant conditioning: the animal learns in the specific situation not to make incorrect movements. Margules and Margules (1973) summarise the behavioural data in relation to motivation: "The brief presentation of a motivationally significant stimulus to an organism results in the activation of a broad spectrum of exploratory

movements. These movements have consequences that exert a powerful effect upon the future movements of an organism. Some of these movements may be unsuccessful in producing the re-occurrence of the motivationally significant stimulus, whereas others may produce harmful consequences. When a successful movement occurs, an active organism will refine this movement by the suppression of unsuccessful variations . . ."

This aspect of motor learning has been neglected in most neurological models, and is not, as far as I know, represented in any simulation of a neurological model. Konorski (1967) is one of the few theorists who examines this aspect of behaviour in great detail. Using Pavlovian terms, adapted to his own concept of a Type II Conditioned reflex (i.e. operant conditioning), Konorski refers to the suppression of responses due to failure of reinforcement or an "aversive Unconditioned Stimulus" as Inhibition. He argues, on analogy with his researches into the similar suppression of classically conditioned responses, that there are two centres of conditioning, one promoting the response and one antagonistic to it. These centres are both excitatory in the neurological sense. "It is worth remembering that in all our experiments dealing with motor act inhibition, the antagonistic movement was always conspicuous . . . when the dog had been trained not to flex his leg in the presence of the metronome, because its flexion prevented receiving food, the metronome began to elicit a strong extension movement of that leg". (Konorski, 1967, p.453). Basically, what Konorski suggests is that the receptor signals which indicate absence of a UCS (e.g. no food in the mouth) form their own Inhibitory Conditioned reflex. This conditioned reflex is driven by an Anti-food Drive.

As already mentioned, there is a considerable theoretical gap between Konorski's outline of his hierarchical network model and his motivational and similar explanations of behaviour. Nevertheless, his basic mechanisms have the strength of being based upon an exceedingly thorough and sophisticated analysis of results of behavioural experiments.

In my model, concepts similar to those of Konorski are developed. There are some differences. These differences are probably due to the particular way the simulations developed, rather than to an intention to challenge Konorski's conclusions in the respects in which they differ. What both models have in common is a conviction that the suppression of

behaviour is important: that it is mediated by excitatory connections in special suppression centres; and that these centres are driven by motivational sources.

(d) Disturbing effects of motivational inputs

It was mentioned earlier that two problems arose in the model in connection with the incorporation of the motivational system. The first of these was predicted by Hebb (1955), who argued that motivation would be disorganising at high levels. Accordingly the appearance of such disorganisation in my model does not perhaps represent so much a deficiency of the model as a challenge which nervous systems have overcome more successfully. In my model, the effect was exaggerated because of control troubles to be recounted later.

Briefly, in my model there was demonstrated a conflict between motivational bias and stimulus recognition. If the motivational inputs were given a sufficiently low value, trouble was avoided. In that case the motivational excitation merely added slightly to the pattern discrimination of one learned record from another. The model remained one in which motor actions are elicited by stimuli, albeit with the motivational state as a discriminative factor: i.e. a Stimulus-Response model. There is then no explanation of how behaviour might be initiated by motivational factors. If, on the other hand, the motivational inputs were made powerful enough to initiate behaviour, then the excitation from such inputs overbalanced the discriminatory effects of the sensory inputs. In the model this had the effect that the model could be easily taught to repeat an action which had previously been reinforced: it found it very difficult to learn to make one response under one notional stimulus cue, and another response under another such cue.

The second problem concerned attempts to make the model learn a number of successive actions to reach a goal. Once it had succeeded in reaching the goal and being "reinforced", on subsequent trials the successful response was driven by the motivational input. Although the motivational excitation spread "backwards" temporally speaking to responses emitted earlier in the sequence, it acted most strongly on the response nearest to the goal. The model therefore produced this response at earlier inappropriate steps of the sequence.

Although these problems were not adequately solved in the simulations to be reported, some progress was made. A solution was developed towards

the end of the project, and was programmed, although not run sufficiently to demonstrate its effectiveness definitely. It is mentioned here because, although developed independently in the present model, it shows considerable similarity to solutions proposed by Konorski (1967) and Deutsch (1960). Essentially, the idea is that every representation of an emitted response is balanced by a representation of the suppression of that response. Motivational excitation is fed in through links with these representations. When any input pattern is in the input other than the particular input pattern which preceded a previous emission of the response concerned, then both representations are equally biassed by the motivational inputs. The particular input pattern which preceded the previous emission triggers the release on the current occasion because it energises one of the two representations more than the other, thus disturbing the equilibrium in the direction of release. A similar strategy of balanced motivational inputs was used to label responses as "successful" or "unsuccessful" in relation to a goal which occurred later in the sequence.

Further details of the proposed mechanism will be given in due course. As far as can be seen, the representations of suppression would act rather like the switches suggested by Deutsch, permitting the currently required response to be released, and afterwards switching off the excitation to it. One slight difference is that in my proposal motivational inputs are counterbalanced rather than actually switched off. This enables the whole of a sensori-motor pathway to a particular goal to be biassed above general levels, whilst the release of responses is held in check. The mechanism also resembles Konorski's formulation by providing an opposing force which brings a response to cessation.

(e) Motivation and arousal

In the early simulations of my model, an arousal system was introduced. The level of arousal was made dependent upon the notional biological significance of the input, i.e. its temporal relationship to innately recognised patterns representing satisfaction of physiological needs. One of a number of problems encountered was the appearance of error accompanying high levels of arousal, as forecast by Hebb. Another was its neurological implausibility. The behaviour of the simulated system bore little resemblance to the known behaviour of the nervous system in arousal (see e.g. Sokolov, 1960; Magoun, 1954). The main difficulty

now seems to have involved theoretical difficulties in distinguishing the operation of arousal and motivation. Hebb saw this difficulty also: "If we tentatively identify a general state of drive with degree of arousal, where does this leave hunger, pain and sex drives?" (Hebb, 1955). Since Hebb's statement, an enormous amount of research has been devoted to this aspect of neuroscience (see e.g. Deutsch and Koopmans, 1976; Olds, 1976).

Some progress was made towards incorporating the insights of such research into the later development of the model. In particular, certain problems of control encountered in the model were solved by a notional periodic gating of afferent inputs. These brought the behaviour of the simulated system more into conformity with the observed activity in the nervous system in arousal. It also introduced a change of perspective into the representation of both arousal and motivational inputs to the system. Arousal became an attentional rather than a directly motivational control. The concept of arousal came closer to that described in connection with an extended neural template system. Arousal can be directed to any part of the template, or to all or none of them. Arousal in specific pathways enables weight to be given to one or another aspect of the total configuration of sensory and other inputs to the integrating system. Motivational inputs, by comparison, form part of the neural template system. States of arousal can thus drive either motivational or sensory discriminative aspects of the system, with differing behavioural effects.

A diffuse distributed model in which motivational inputs may have biassing effects would appear to differ from a computational one in that, as in the simulations of my model, motivational inputs at too high a level led to the model "perceiving" incorrectly as well as behaving incorrectly.

(f) The nature of the reinforcing event

A question which a model of neural integration representing motivated action must answer is the nature of the reinforcing event. To find a solution concordant with the data is far from simple.

Perhaps the most widely accepted proposal was that of Hull (1937, 1943). He said that reduction of drive was the essential condition for reinforcement of learning. Hull distinguished drive (D) from Habit Strength (sHr). Reinforcing states are those involving a diminution of the

receptor discharge characteristic of a need, or, alternatively, the reduction in drive receptor response (Hull, 1943). Essentially this ties learning firmly to specific physiological needs. Later research has thrown considerable doubt upon the reduction of drive concept. Clearly many aspects of eating behaviour have little to do with the metabolic need states resulting from lack of food (Mogenson and Phillips, 1976). On the other hand, the discovery of the diencephalic reward systems makes the physiological aspects of the hypothesis more plausible. Considerations of this kind led to Konorski's (1967) formulation, in which there is a separate "food reflex" and "hunger reflex", each having its own reinforcing event.

In Hull's theory, D and sHr are multiplicative, so that an absence of drive means that no learning takes place. In contrast, Guthrie and his followers supported a pure contiguity theory of learning in which motivational conditions play the part of stimuli only. In this context, Tolman's classic work (1932) makes Hull's hypothesis extremely doubtful.

Drive reduction offers no particular difficulty of representation in my model. It was not actually adopted as the definition of reinforcement because of my predelictions in other directions. Contiguity, as a principle of learning, is the basis of the neurological model proposed here. The question, however, is whether such a principle is a sufficient explanation of learning. An exhaustive attempt to demonstrate effective learned instrumental responses using only the contiguity aspects of the model failed. This may of course have been due to inadequate solutions of quite different problems of neural integration at the time the attempt was made. The reader will have gathered from foregoing discussion that the present author considers that contiguity is a principle which underlies all learning; learning as it appears in animals is, however, assisted by special mechanisms. Amongst the most important of those mechanisms may be mechanisms of arousal.

Instead of Drive reduction, Sheffield (1966) proposed what in effect is a drive induction theory, if it is a drive theory at all. Sheffield and Campbell (1954) had suggested that the effect of hunger is to lower thresholds at which responses are elicited by external stimuli. It is not clear how the threshold reduction is to be accomplished. Selective addition of drive inputs lowers threshold. This necessarily applies to

all the motivational models discussed. Whether such a reduction of thresholds is to be regarded as driving behaviour or facilitating response to stimuli is a question of the weight given to the inputs, as already discussed. This question is closely associated with the so-called incentive view of motivation. These views of motivation are perhaps not so far apart as may at first seem. If motivational inputs can be applied selectively to records of sensory events which have preceded reinforcements, then such inputs would reduce the thresholds of the records concerned. Thus stimuli would acquire incentive properties.

A feature of Sheffield's model is that it is the consummatory response rather than drive reduction which is the reinforcing event. This is quite close to Deutsch's model (1960). In Deutsch's view the reinforcing event is the arrival of excitation to a stimulus analyser connected to a primary link (innately connected to motivational excitation). For him the reinforcing event is the arrival of a sensory pattern. The resemblance of this assumption to the "key trigger patterns" proposed by Pringle (1951) may be noted. Key trigger patterns also release responses. A key trigger pattern could mediate Deutsch's Sensory Analyser connected to a Primary Link. A similar assumption is made in my model. In so far as behaviour is guided by innate needs, it is reinforced by the occurrence of an environmental situation which matches an inbuilt configuration.

However, such an assumption does not in itself explain reinforcement. The mechanisms still have to be defined whereby responses which lead to a key trigger pattern gain precedence over responses which do not. Some indication has already been given of the general approach I have taken to this question. The shaping of behaviour is not suggested to be exclusively tied to direct physiological satisfactions, like food, or even psychological satisfactions, like sex. Nevertheless, in developing the simulations a number of mechanisms have been used. As already mentioned, the aim has been primarily to take the model a stage further towards the mediation of behaviour which can be compared to that of living things. The kind of picture reached towards the end of the project resembles most closely that of Deutsch (1960) and perhaps Sheffield (1966). In Deutsch's model, excitation becomes transferred from the primary link to the stimulus analyser excited just before the excitation of the Stimulus analyser operating the primary link.

This theoretical position has its shortcomings. It ties behaviour strongly to a primary link. Moreover, it implies that all behaviour is built up in atomic chains, whereas animals, especially Man, often act as though they possessed cognitive maps of the world in which they live. The adoption of mechanisms similar to this scenario does not mean that the view promoted here is that behaviour conforms generally to such narrow limits. On the contrary, the basic attributes of my theory of neural integration would appear to contain potential for a broader view of neural organisation, e.g. anticipation. Although the kind of step-by-step assembly of motivated sequences, as proposed by many theorists, provides a rather narrow concept of behaviour, it is not necessarily inimical to a broader view. In my model, such mechanisms are to be regarded as specialisations which have developed from the basic capacity of neural tissue to mediate integrations. Other mechanisms would have developed, less tied to physiological events. Even the mechanisms developed show promise that they could be applied to the broader kind of behavioural organisation with which we are more familiar in Man. The development of such behaviour is beyond the scope of the present project, but some further reference to the issue will be made in later chapters.

IV The problem of brain function: a problem of integration

A. Behavioural and psychological aspects

The author has raised a number of questions in the foregoing sections. One of the main problems which confronts the designer of a neurological model of the brain is not the solving of any one of the problems, but proposing a model which at least may seem to hold the promise of contributing to the solution of them all. A model must of course fail in that respect. It will fail in two ways. First there will be the theoretical shortcomings in the principles of the model which make it inherently unable to deal with some aspects of brain function in a plausible way. Second, the difficulties of simulation will make it progress only a short distance, in some directions at the expense of others, and incorporating deficiencies due to the inadequate design of the simulations.

What it is hoped has been established in the foregoing sampling of psychological and behavioural problems is that the essential task of a neurological model is to show how all the aspects of function can be

co-ordinated into a unified organismic action. Perception cannot be entirely separated from motor action, nor from reinforcement and motivation, nor can any of these be separated from each other.

B. Neuro-anatomical integration

Alongside the integrative activites of the brain in respect of behavioural function, is the co-ordinating function of many discrete interconnected ganglia, nuclei and centres of specialised function. What is the contribution of the archi-cortex and of the neo-cortex? What is the function of thalamus and basal ganglia? These are not questions for a model of neural integration. However, they are questions upon which the neurological model should be able to say how their activities are co-ordinated.

Few neurological models have pursued this question of the co-ordination of parts and their separable and distinguishable functions. Old's model describes the function of a notional Hippocampus, but does not fit that model into an overall neurological model of the brain. Kilmer, McCullich, Blum, Craighill and Peterson (1968) proposed a model of the reticular formation and Stapley and Kilmer (1974) one of the Hippocampus. Again these are not fitted into a model of the overall integrative properties of the brain. On the other hand, simulations which do show the capacities of overall integration (e.g. Spinelli, 1970) do not show the differentiation of the function of the parts and therefore do not demonstrate how the separate functions would be co-ordinated. What is aimed at here is a demonstration that separation of function by evolved specialist mechanisms is compatible with overall integration without ad hoc co-ordinating mechanisms having had to be evolved.

V Conclusion

A number of problems of psychological function have been discussed. These are topics which it is hoped will orient the reader to the implications of the model of neural integration to be presented. It is hoped that they will enable the reader to see something of the rationale behind the neural representation in the model, and the decisions taken in the simulations. In considering a model of neural integrations, one should bear in mind the tremendous burdens which the integrative activities of the brain places upon any model. It is hoped that the foregoing discussion has illustrated this burden.

CHAPTER 4

SOME PRINCIPLES OF ORGANISATION IN CONTEMPORARY MODELS OF NEURAL INTEGRATION

In the previous chapters some psychological aspects of brain function were discussed. In this chapter, some consideration will be given to theories of the way in which neurons might be integrated to mediate those functions.

I Mechanism, technological analogies

Mechanism and organising principle are often difficult to distinguish in brain models. Before reviewing some of the models which have been proposed, one may note that one way of considering a model is in relation to the technology of the era in which it arises. Griffith (1967) argues against the employment of such analogies. " . . . from an experimental viewpoint we are in complete ignorance of a great number of basic facts . . ., therefore many people have slid away from the real problem and have tried to impose upon the brain an alien structure - that of electronic circuitry, of a mechanical servo-mechanism, or of symbolic logic."

Such attempts have often led to valuable insights. Given essential ignorance of facts, analogies with complex artefacts may be the only way to proceed. At a time when hydraulic engineering had just come into its own, and fountains and waterworks were appearing all over Europe, Descartes proposed that "animal spirits" flowed through the nerves rather like water through a pipe. It was sent through the "pores of the brain", and once the pores had been used, the flow through them was facilitated. This may not be so different from contemporary views of transmitter function.

The telephone analogy of the brain was a powerful step forward. Its full implications have only become apparent with Shannon's development of the boolean algebra of telephone switching networks, and was taken even futher in the theorems of Pitts and McCulloch in the forties. (McCulloch and Pitts, 1943; Pitts and McCulloch, 1947).

More recent examples of importance are Pringle's model of neurons acting as coupled oscillators, and the "neural hologram" model of Pribram (1971) and Spinelli (1970). There is even a sense in which models which proposed some principle of organisation directly related to the brain derived those principles from recent technological developments. Rosenblatt's Perceptron (see below) can for example be referred to contemporary progress in cybernetic engineering.

Valuable as such insights may be, we now appear to be reaching a stage where the facts which are becoming available are best understood not by analogy with any other subject, but by direct reference to the brain itself. In particular, the data coming from electrical implantation, neuro-pharmacology, electron microscopy and neuro-behavioural studies point to an understanding of the brain in terms of its neurons, pharmacology and micro-anatomy. As Griffith urges: "we seek to understand the brain not as a computer or anything else, but as an entity in its own right". Accordingly in the following section, brain models will only be considered insofar as they seem to offer principles which aid the understanding of the brain as a biological entity.

II Three principles distinguished

With the elucidation of the neuronal functions in the nervous system two earlier types of model have fallen into the background. In one the brain was conceived as a connective system of a completely determined kind based upon analogy with a reflex arc. In the other, the nervous tissue was seen as a single substance in which field forces act, more or less independently of the neuronal individuality of the tissue. Like the particle and wave theories of light, both these ideas are now seen to be aspects of something more complex than either. Most, if not all, constructors of contemporary models, interpret the nervous system in a statistical sense. A connective theory, even one such as that of Griffith (1967, see below), in which connections are determined by "neuronal recognition", take into account that such connections are probabilistic. In view of the number of elements in the brain, the recognition of the probabilistic factors, tend to bring all models into the class of diffuse processing.

There does, however, seem to be a fundamental difference between models, such as were initiated by Hebb (1949), which interpose a diffuse central process between sensory and motor processing, and those where the

complexities are introduced only in the pathways between sensory and motor processing. These two kinds of model may again be distinguished from those in which there are successive levels of orderly processing culminating at an apex from which again there are successive levels down to the periphery.

For convenience these may be termed (i) hierarchical models, (ii) discriminative feed-back weighting models, and (iii) distributed memory models.

III Hierarchical models of the nervous system

A. The hierarchical organisation of the brain

Much of the support for a hierarchical neural model is the apparent hierarchical organisation of the brain seen in clinical neurology and in the general anatomical organisation of brain areas. At the neurophysiological level this argument has been developed most powerfully by Luria (e.g. 1973) and Konorski (e.g. 1967). Konorski has developed a well-defined neurological model in conjunction with his general hierarchical view. This will be discussed further in this thesis. Luria does not propose a precise neural model. However, his structural and functional analysis serves as a background to hierarchical models generally.

Luria distinguishes three functional units in all sensory modalities, the primary projection area, a secondary area "which convert(s) the somato-topic projection of the individual parts of the retina (in the visual area) into its functional organisation", and tertiary "overlapping zones". The latter, Luria thinks, are exclusively human developments. (Luria, 1973, p.73). They are concerned with the integration of excitation arriving from different analysers. He proposes three basic Laws: the Law of Hierarchical Structure of cortical zones; the Law of diminished Specificity of the hierarchically arranged cortical zones; and the Law of Progressive Lateralisation of function. The last law implies that there is a progressive transfer from primary cortical areas to secondary and finally to tertiary areas.

In considering Luria's work, it must be borne in mind that the aim is not the understanding of neural integration, but rather the relationship of structure to function. Luria may well be right in thinking that anatomically that the relationship is hierarchical. However, the point which was made in connection with pattern recognition models must be made again here: functional hierarchy does not necessarily imply the existence of an

hierarchical neural structure. Consider the extended template system described in Chapter 2. The section of the template which matches "strokes" is functionally hierarchical to the section matching "letters" (i.e. the simple template section). Non-hierarchical theories of neural integration, including that to be described in this thesis, can mediate an extended template of this kind.

There is now evidence throwing doubt upon the strictness of the hierarchy proposed in Luria's Laws. Singer and Tretter (1976) investigating the receptive field properties of contour deprived cats, found receptive fields in Area 18 very similar to those in Area 17. Their methods of investigation, which included electrical stimulation from the LGN, suggests that these projection cells in Area 18 were at least partially innervated by direct projection from the LGN. Baizer, Robinson and Dow (1977) investigated visual responses of Area 18 neurons in awake behaving monkey with the specific aim of determining whether such neurons were organised on an hierarchical basis from Area 17 neurons or by "independent chains". Neurons in Area 18 were found to be preferentially responsive to six kinds of stimulus parameter: orientation, colour, border, direction, spot, light inhibition. The receptive fields served by Area 17 neurons described by Hubel and Wiesel (1962) were simple, complex and hypercomplex. All preferred elongated stimuli and all were sensitive to stimulus orientation. Baizer et alia conclude: "The finding that different populations of cells in Areas 17 and 18 are maximally sensitive to different stimulus parameters implies that, at least in this stage in visual processing, the neural representation of visual stimuli is accomplished through the simultaneous activity of cells in different groups rather than through the activity of cells with compound specificities".

Luria also proposed Three principal Functional Units: (1) Tone or level of activity, associated with the reticular formation, hypothalamus and metabolic links with medulla and mesencephalon (2) Unit for processing and storing information, comprising principally the lateral regions of neo-cortex, convex surface of the hemispheres and posterior regions, (3) Programming, regulating and verifying activity, located in the anterior regions of the hemispheres, i.e. the frontal lobes. One of the difficulties with hierarchical neural models is how these functional links would be brought into the hierarchical neural network. Perhaps the most detailed effort to explain this was that of Konorski (1967).

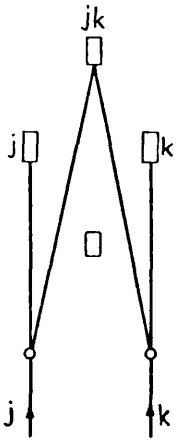

FIGURE 3a

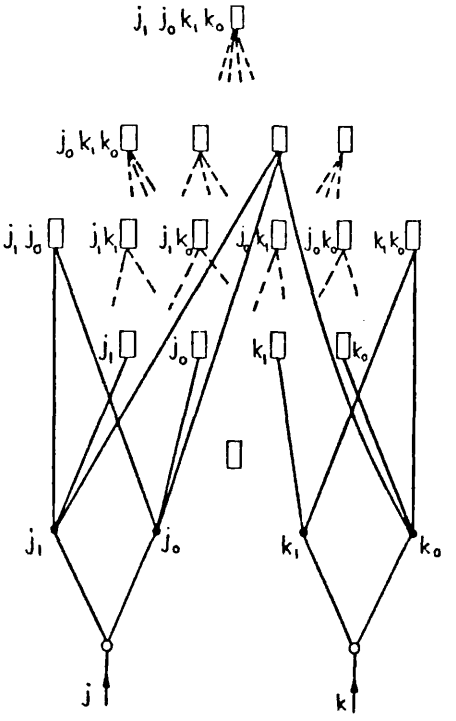

FIGURE 3b

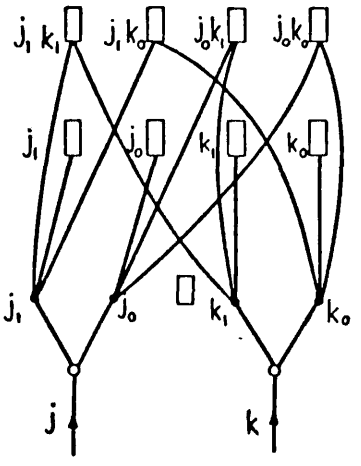

FIGURE 3c

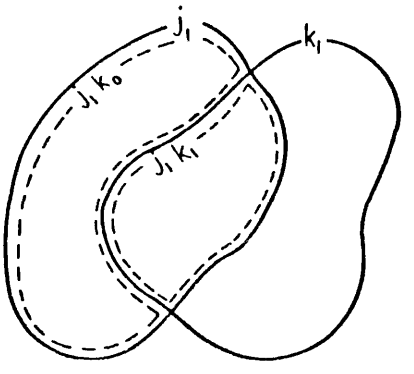

FIGURE 4

Figure 4.1 (reproduced from Uttley,1956a) Hierarchical organisation of neurons.

Luria himself draws attention to certain aspects of behaviour which one would have thought difficult to fit into an hierarchical scheme of neural organisation. For example, "This motor task, or model of the future need, is constant or invariant and it demands an equally constant or invariant result . . . It is a most important fact that the invariant motor task is fulfilled not by a constant, fixed set, but by a varying set of movements which, however, lead to the constant invariant effect". (Luria, 1973, p.248). Again, with respect to memory, "All, however, unanimously agree that system of connections into which traces of information reaching the subject are introduced are coded with respect to different signs, and consequently they form multi-dimensional matrices from which the subject must choose each time the system which, at that particular moment, will form the basis for coding". (Luria's italics). It would be out of place to pursue the difficulties in hierarchical systems here in depth. They are argued by Feeney, Pittman and Wagner (1974) and Pribram (1971). Some further aspects of the difficulties in such systems will appear in discussions of models below.

B. Conditional probability, temporal sequence: Uttley

Shannon's elucidation of the boolean algebra of networks led to the realisation of the possibilities of logical networks as an explanation of brain function. Pitts and McCulloch developed further the concept of a "formal neuron". This was a neuron which fired or failed to fire as a function of excitatory and inhibitory inputs from other formal neurons. The "Formal neuron" enabled a logical calculus to be applied. In such analyses Pitts and McCulloch challenged the apparent rigidity of logical networks by demonstrating that "reverbertory circuits" could remove the dependence of the networks from the particularities of a specific input event and so mediate Universals (McCulloch and Pitts, 1943; Pitts and McCulloch, 1947).

Uttley (1956, 1959, 1961) put forward an hierarchical model which differed from former network proposals in taking into account not merely the conjunctions and disjunctions of inputs, but the probability of those conjunctions. He had already outlined (Uttley, 1954) an analog-digital conversion system. The essential form of the nervous system, as proposed in Uttley's model, can be seen in Figure 4.1. This figure, however, only represents the hierarchical convergence of a network. The essential idea of Uttley's "conditional probability machines" is that units count not merely the occurrence or non-occurrence of inputs (or

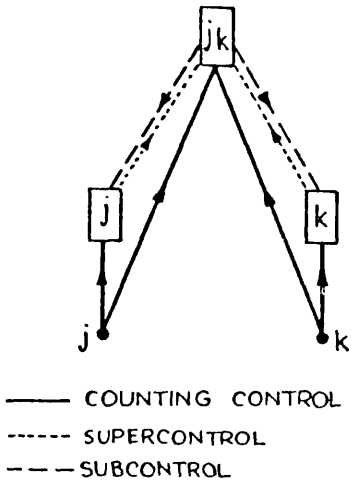

Figure 4.2 (reproduced from Uttley,1956b). Conditional probability in an hierarchical system implemented by Super contro and Sub control connections. These enable the units at the j and k levels to register the probability of jk given j or k.

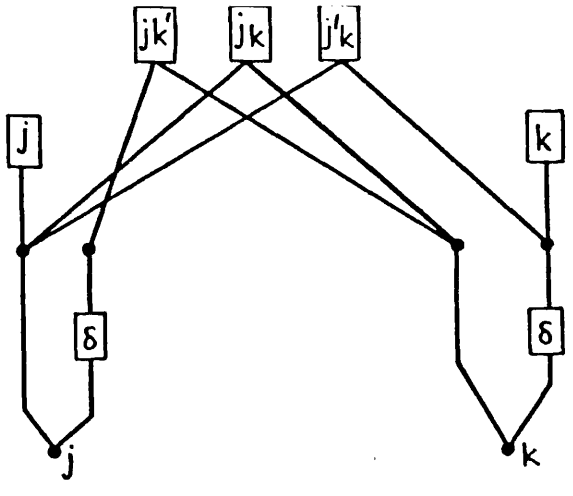

Figure 4.3 (reproduced from Uttley,1956b). The introduction of a delay into one of two convergent pathways enables units to register the sequence jk'.

conjunctions and/or disjunctions of inputs), but that they also count the probability that such an input should occur if another neural event has occured. Thus the units at the J and K level of Figure 4.1. would contain "a quantity representing unity probability" if "the set J" actually occurs. If neither J nor K occurs, the unit is to contain "not zero, as in the classification machine, but the unconditional probability (my italics) of J as determined by prior experience. If K occurs, the unit must contain the probability of J given that K occurs". Uttley compares such a machine to an Unconditional probability Machine. The unconditional probability machine computes only the unconditional probability of J when J does not appear, or unity when it does appear.

For calculation of the conditional probabilities, it is necessary for the lower (J,K) level of units to refer to the next highest level (JK). This necessitates what Uttley refers to as "Supercontrol" and "Subcontrol" (See Figure 4,2 from Uttley, 1956b). There are, of course, other units than K to be considered. This brings in the difficulty of multiple control. First it must be noted that "There must be a connection from every unit to each of its superunits" and "the system of subcontrol is a duplicate of the supercontrol system, but it has a different function, operating in a reverse direction". There are also necessary rules for sub-control and super-control which avoid false control, i.e. changes supporting false inferences.

Lack of simultaneity of events was expressly excluded from Uttley's analysis of the conditional probability machine. In a separate paper (1956b) he shows how the network can be extended to compute temporal sequence. For the unextended conditional probability machine, events (impulses) must be considered as occurring together or not at all. By introducing delays into the pathways, then events which occur in temporal sequence can be represented as conjunctions at points in the hierarchy (see Figure 4.3). This can, of course, be extended to include representations of conjunctions resulting from two, three or more delays.

Although it is frequently assumed that Uttley's demonstration of delay in pathways as a mediator of temporal sequence is peculiar to hierarchical network models, this is not in fact so. A considerable use of this principle to mediate temporal sequence is made in later versions of my model.

The main objection to Uttley's proposals lies in the number of units and pathways involved (Uhr, 1973). Sutherland (1959) points out that to make a discrimination of two points on a retina as to their relative position and brightness, "we would require more units than there are units in the central nervous system". Dodwell (1964) proposed that a shape coding system, similar to that described by Deutsch (1960) or Sutherland (1957) is interposed between sensory receptors and a "Recognising system" of the kind contained in Uttley's Unconditional Probability machine. The outputs from this can then be applied to the Conditional Probability machine. Dodwell points out that his "coupled system" would provide an economy of units, serving to remove the objection raised against the Uttley type of system.

Although connections would be reduced by Dodwell's coupled system other questions are raised. There is the nature of the precoding device. If he means feature analysers of the kind discovered by Hubel and Weisel, then, according to the hierarchical model, this only constitutes a lower hierarchical stage, and the overall numerical questions remain unchanged. If, however, as seems to be intended by Dodwell, the precoders are of the kind suggested by Sutherland, one must ask where the neural apparatus would be situated. If in the cortex, the non hierarchical computational network would be immersed within the recogniser and conditional probability apparatus, leading to further complications of structure and function. If sub-cortical, it is difficult to explain the topological mappings of sense organ onto cortex. These would be obviated by codings of the kind suggested by Sutherland. Moreover, there would have to be a precoder for every different form of analysis available to the animal. There would also be the difficulty noted by Luria (cited above) that the subject can choose coding according to his immediate purposes.

In general it seems to the present author that there is an implausibility in the interpretation of the neuro-anatomy and histology in terms of hierarchical systems. Nature appears to be capable, not only in nervous systems but elsewhere, of controlling the most precise ordering of fibres. Although it is admittedly possible to account for an hierarchical system in terms of random diffuse connections, if the main principle of neural organisation were really hierarchical, one would expect to see the kind of layout of fibres implied by such a system. Instead one finds a highly specialised, morphologically developed, histology of a different kind (see Chapter 5).

In summary, there are the numerical difficulties inherent in hierarchical systems; the evidence from Singer and Tretter (1976) and Baizer, Robinson and Dow (1977) against the proposed hierarchical relationship of areas 17 and 18, and there are the difficulties posed by the apparent lability and coding of attention.

Finally, many will find it implausible that the perception (recognition) of unitary concepts in the world should depend upon the response of a single neuron, to which impulses have converged. This will be more fully discussed in connection with Konorski's gnostic unit model, and in the last section of this chapter.

C. A complete hierarchical model of the brain: Konorski

Konorski (1967) has presented what is perhaps the most definitely specified verbal model of the brain in neural terms. My model is in some respects at the opposite end of the spectrum from hierarchical models like Konorski's. Yet there are several strong resemblances with his formulations. Accordingly, some space is devoted to his theories, both here and in later chapters.

Konorski uses the evidence of receptive fields to support his proposal of convergent and divergent networks. "As it was indicated by Hubel and Weisel, the units of higher levels of the visual afferent system are formed by the convergence of appropriate units of the lower levels. In other words, a unit of the highest level represents the top of a pyramid whose basement consists of receptive organs . . . However, neither humans nor animals notice lines, corners, or "tongues" or "rods" . . . It may be supposed that particular units of the so-called associative areas of cortex become interconnected in various ways, forming what Hebb has called cell assemblies corresponding to particular perceptions. . . But having at our disposal the recent data from Hubel and Weisel's experiments, we can extrapolate their findings and explain the origin of perceptions according to the same principles which were found to operate on the lower levels of the afferent system". (This was, of course, written before the data from Singer and Tretter (1976) and Baizer, Robinson and Dow (1977) called this assumption into question). "In other words, we can assume that perceptions experienced in humans' and animals' lives, are represented not by assemblies of units but by single units in the highest levels of particular analysers. We shall call these levels gnostic areas and the units responsible for particular perceptions, gnostic units". (Konorski, 1967, p.75).

In respect of the proposed convergence upon a single cell, Konorski qualifies his proposal by introducing the idea of redundancy: "each unitary perception may be represented in a given gnostic field not by a single gnostic unit but by a number of them, because, if in a state of arousal of that field a new stimulus pattern is presented, all unengaged units which potentially include its elements are capable of becoming actual gnostic units representing that pattern". (Konorski, 1967, p.90). He uses this explanation to deal with the fact that brain damage in older persons damages recent memory more than older memories. "In younger subjects the number of gnostic units largely prevails over the number of actual units . . .". Consequently in older persons, the sets of gnostic units representing recent memories are in general smaller than those of early memories and more vulnerable to damage. A prediction of this kind, which appears to be specific to a particular model, may in fact be general to all models incorporating certain assumption, e.g. a stock of memory elements. In my model, the provision of a finite stock of elements yields similar predictions, although in a very different kind of system. In my model the stock of elements is being constantly turned over: briefly, the allocation to the storage of a memory undergoes decay, upon which the memory element becomes available for contributing to the storage of another memory. In such a system, there is a change in the statistical distribution of the probable duration of memories. This will be referred to later. These aspects are not specific to either my model or Konorski's.

Another feature of Konorski's model which is in common with mine is the assumption that representation in the nervous system is generally in the form of complementary balanced systems. For every unit signalling an event there is a unit which would be active when that event did not occur, e.g. on-receptor and off-receptors. At the time this was introduced into my model, I had not read Konorski's (1967) book; nor had I appreciated the significance of Uttley's provision of similar elements. Konorski makes extremely powerful use of the principle. For example, letters of the alphabet which have parts in common, say, E and F, are not confused because the extra tail of one appears as a "not-tail" in the other. These considerations apply equally within my model. In my model, the balanced system was introduced because otherwise the model would have suffered from great inadequacies of processing.

Konorski's model has not itself been simulated. However, a model of

attention utilising lateral recurrent inhibition was simulated by Walley and Weiden (1973). The basic idea is that gnostic units do not inhibit each other within a set but do so between sets. This idea seems to be very close to the mutual inhibition between cell assemblies introduced into Hebb's model at Milner's suggestion (Rochester, Holland, Haibt and Duda, 1956). "The encoding of one stimulus will tend to interfere with the simultaneous encoding of other stimuli. This interference is termed cognitive masking, and is attributable to recurrent lateral inhibition between gnostic units". (Walley and Weiden, 1973).

Feeney, Pittman and Wagner (1974) have criticised the Walley and Weiden model and take the opportunity to mount some criticisms of hierarchical neural interpretations of perception in addition to those previously mentioned here. They also raise objections to the attribution of attention solely to inhibition.

Konorski's neurological model has been considered above only in its embodiment of an hierarchical structure of neurons. His other theoretical formulations, especially his discussions of behaviour in conditioning, are discussed elsewhere in this thesis.

IV Feed-back weighting adjustment: discriminative models

In Chapter 2 reference was made to a number of pattern recognition models, e.g. Rosenblatt, 1958, Selfridge, 1958 which exercised a profound influence on notions of neural processing. Essentially these are "discrimination models" in the sense distinguished in the discussions in that chapter. Their function is to connect responses to input patterns, patterns being distinguished, i.e. discriminated, only by the responses attached to them.

The characteristic feature of these models is that the connections defining the pathways through the system remain unchanged by experience. What changes is the weighting given to the messages which pass along the pathways. This variation of weighting may be as simple as an all-or-none passing on of the signal at a substation. Although the direction of the pathway of signals is afferent-efferent, there are interposed sub-stations at which conjunctions of afferent fibres determine the state of the fibre which emerges at that stage to proceed to the next. These connections are not changed with learning. The conjunctions are generally random selections of fibres, taken in bundles the size of which

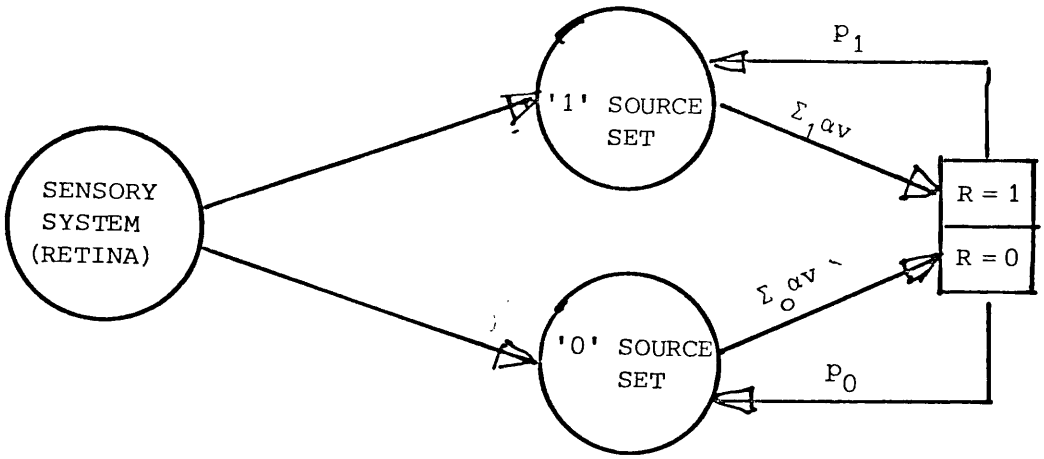

Figure 4.4 (Redrawn from Rosenblatt, 1958b) A simple perceptron.

is also random. The only efferent-afferent connections in the system are feed-back fibres which vary the weighting of the transmissions at the sub-stations. The feed-back system provides a discriminative mixture of correlations of inputs at each sub-stage.

The purest form of this principle is seen in Rosenblatt's Perceptron (Figure 4.4). (Rosenblatt, 1958(a), 1958(b), 1960). In this model one can see most clearly the implication of the principle as providing a direct through-put of signals between sense-organ and effector. In conjunction with the model, Rosenblatt proposes a principle of separability. In effect this argues the self-organising nature of the system. Any set of input patterns will be discriminatied into sub-sets if a randomly distributed feed-back function, acting in accordance with some criterion assessed at the output, varies the weightings given to the conjunctions of the transmission fibres at a substation. The perceptron has been critically analysed by Minsky and Papert (1969) and summarised in Scott (1977).

A similar principle may be seen in a number of models which however contribute other insights. In the well known Pandemonium of Selfridge (1958) weights are used to form what may be seen as feature analysers. There is a stage of "computational demons" which "shout" according to some feature of the input. "The computing demons are constructed from only a very few operational functions which are carefully non binary". These sub-demons shout at a set of "cognitive demons" whose "sole task is to add a weighted sum of the outputs of all the computational demons. The system offers the prospect of being able to vary the feature analysis so that those are found which optimise the operation.

Selfridge discusses the concept of "hill climbing". As this concept conveys the mathematical flavour of many models (e.g. Kilmer, McCulloch, Blum and Craighill, 1968; Marr, 1969, 1970), it will be briefly summarised. ". . . the complete set of weights for all the cognitive demons is a vector . . . Now for some (unknown) set of weights the behaviour of this whole Pandemonium is optimum, and the problem of feature weighting is to find that set". Considering the vector space, one climbs the hill to the optimum point by taking small random steps in all directions. When a direction is found which improves the "score", take that step and start again. The technique is similar to that employed in solving equations in the discipline of statistical mechanics. One problem is discerning false peaks. Such models work by small steps, computing a

score at each step, following false trails to false peaks, until finally it finds a peak.

This process would seem to throw light upon the general principles of evolutionary development. As an explanatory model of an highly evolved nervous system, there are difficulties. What is the principle by which each step in the hill climbing is assessed? According to the model, the only changes required to be brought about in the neural apparatus by evolution are the weightings of the connections. In that case the highly evolved nervous system should retain essentially the structure of the primitive one except for modified feed-back connections.

A model which attempts to apply the same kind of principles to discovering what those evolutionary changes have been is that of Marr (1968, 1970). The model is too complex to be summarised here, and only an indication will be given. In the (1968) model of the cerebellum, the Purkinje Cells are the key elements, each controlling an elemental movement and receiving a feed-back via the climbing fibres from the Olivary Nucleus. These feed-back signals act essentially as reinforcements. Alternatively they can be seen as mediating the feed-back modifications of a Rosenblatt Perceptron. The climbing fibres control the synaptic modifications of the input of the parallel fibres to the Purkinje cells. A key concept of the model, which differentiates it from other feed-back weighting modification systems, is the Codon. The idea is analogical with the codon by which the structure of DNA was analysed. An almost unlimited differentiated structure can be built up of different combinations of relatively small samples.

In Marr's model, Mossy Fibres are input fibres reporting sensory contexts. The mossy fibres act on granule cells which act upon the parallel fibres which act upon Purkinje cells. The only synapse which is modifiable is the feed-back control synapse of the climbing fibre upon the Purkinje cell. The other arrangements provide selective sampling. A Codon is a small sub-set, about 4-12, of the mossy fibres. There are about 200,000 parallel fibres capable of exciting any one Purkinje Cell. The reinforcement action on the parallel-fibre-Purkinje Cell synapses serves therefore to vary the selection of codons which will excite that Purkinje Cell. Marr (1970) suggests that the system can be generalised to provide a model for the neo-cortex: "The key idea behind the present theory is that the brain decomposes its afferent information into what are essentially its natural cluster classes". During life the

clusters are refined and decomposed until they fall inside classes "which in real life have to be discriminated".

Albus (1971) considered the anatomy and probable function of the cerebellar cortex. Whilst not working out a detailed model, he suggests that the cerebellum acts as a Perceptron. His contribution was to suggest that a "list processing" could be undertaken, the iterative aspects of the processing being implemented by the recursive loop consisting of motor efferents reporting to cortex, reporting back to cerebellum.

These models have been mentioned in some detail to give some indication of their theoretical flavour. That character can only really be appreciated by a study of the original presentations (see e.g. Marr, 1968). There seems to be little doubt that the principles offered in such models are valid. In particular they enable one to understand how information carried by neuron-like systems can be analysed merely by chunking and sampling information, and by feed-back variation of connectivity. To what extent those principles actually describe the structure and function of the brain is another matter. Of course, the answer need not be simply yes or no. Marr gives a convincing description of the cerebellum, and it may be that the operation of that organ is at least partly carried out by a feed-back codon adaptation as he suggests.

Kilmer and his associates (Kilmer, McCulloch, Blum, Craighill and Peterson, 1968; Kilmer and Olinski, 1974; Stanley and Kilmer, 1975) have used feed-back weighting adjustment to construct models of various anatomical structures in the brain, i.e. the Reticular Formation and the Hippocampus. These will be discussed further in a later chapter.

One of the purely mechanical difficulties in feed-back adjustment models is that of the retention of the values of the weightings. For example, in the Kilmer models, the feed-back weighting is measured by the number of synapses which have been facilitated from an input upon the integrating cell. This raises questions concerning the permanence of the connections. If they are permament, there would appear little difficulty in retaining the invariance of the feed-back. If they decay it would be necessary for the decay to operate equally upon all the synapses concerned: otherwise decay would lead to a change of the feed-back and accordingly a change of the pattern discriminated. Moreover to mediate

a large number of patterns, the system imposes a tremendous burden on the synaptic aspect. This is increased if no decay is permitted.

However, it is not really upon such points that the discriminative model will be judged. The question is rather whether the model is considered to be satisfying in relation to the general problems of neural integration discussed in earlier chapters. The strength of this kind of model is its efficiency in pattern recognition. Essentially it is a Stimulus-response model, and what the Gibsons have categorised as a Discriminative model (see Chapter 2). In the weighting feed-back model, memory takes the form of the retention of the values of weightings applied to particular random correlations of input fibres. It is difficult to see a place in such a system for the registration and recall of individual discrete memories, enrichment and imagination would be even more difficult to explain. Perhaps the salient point is what is missing from its explanation. How do we perceive essential differences between objects which are equivalent in the reinforcement they obtain from the environment: for example, how do we distinguish between a brown cup and a blue one?

In such models, the neurons performing the integrative functions are isolated from each other. Each neuron receives inputs from neurons or receptor cells nearer the periphery and passes them on to cells further along the line. There is little communication between neighbouring cells. In this respect the feed-back weighting model is a long way from the isomorphic view that there is a psycho-physical correspondence between physical structure and perceptual function. It is similarly at a distance from views of living things as co-operative relationships between cells.

V Collateral Interaction models

In a large number of models, it is the collateral communication between neurons which is the basic principle of operation. The characteristic feature is that afferent fibres stimulate cells which do not pass on their excitation directly to the next stage. Rather they enter into diffuse contacts with their neighbours. Pathways through the system are complex. The activity in the efferent fibres cannot be related in any simple way to the activity in the afferent fibres: it is determined by the lateral connectivity of the integrative tissue interposed between the afferent and efferent fibres.

The essential idea of Hebb's (1949) model is probably that such lateral connectivity supports a diffuse central process which is dynamic in its own right: in that respect it proposed a model of "thought processes". The detail of Hebb's model is too well known to require description here. He assumed that cells with diffuse random connections would tend to form reverberatory circuits if they were excited in sequence. Rochester, Holland, Haibt and Duda (1956) simulated this aspect of the theory. By incorporating an inhibitory modification of Hebb's proposals, suggested by Milner, cell assemblies did form. Hebb further argued that if cell assemblies were excited in sequence a larger scale relationship would form capable of mediating more complex processing. These "phase sequences" were the core of Hebb's discussion of psychological processes.

A number of models utilising lateral connective systems have analogies to wave functions in optical and electrical artefacts. As long ago as 1942, Lashley had suggested that a solution to the dynamic apsects of perception and memory (see Hebb, 1949) would be to regard the cortex as built up of "millions of transmitting loops of various lengths and refractory periods, actually consisting of an indefinite number of resonators". Excitation started at any point would give rise to waves. "With several or many points of excitation, interference patterns would be formed". Lashley acknowledged earlier sources of the idea.

Beurle (1956) presented a probabilistic mathematical analysis of connectivity and argued that, provided that the activity was unsaturated by inhibitional damping, waves would be propagated. A simple synaptic facilitation mechanism would lead to the learning of such waves. "The mass will always tend to reproduce as a trial response the one which has been found most satisfactory in the most nearly similar circumstances in the past". (Beurle, 1959). Learning would particularly concentrate at interference points where waves crossed each other.

With the development of the hologram, a number of authors proposed that the brain might operate upon similar principles. (Van Heerden, 1963; Willshaw, Buneman, and Longuet-Higgins, 1969; Pribram, 1971; Pribram, Nuwer and Baron, 1974; Landfield, 1976. Spinelli (1970) simulated a neural hologram. These models are examined in a later chapter when their implications can be better related to the model to be presented here. Essentially they are the extension of Lashley's proposal that what is registered is the interaction (interference patterns) between neural excitations.

A model which appears to be even closer to Lashley's suggestion was that of Pringle (1951). This has already been mentioned in relation to "key trigger patterns", and the relevance of that idea to innate patterns as represented in my model. Pringle stressed the "parallel between learning and evolution" and sought a mechanism which would parallel "replication": the presence of a unit in the system causes more similar units to appear. Oscillators would provide such units, because of their tendency to couple. " . . . the beats of the chambers of the vertebrate heart, (which) behave like relaxation oscillators. The coupling here is so strong that the whole heart normally beats as a single unit . . ." (Pringle, 1951, p.189). He argued that a reverberatory loop would act as a non-linear oscillator. The non-linearity imparts a unique configuration of response of the tissue to each afferent pattern, and this is registered in terms of coupling. Although there are many differences between the organisation proposed in the model to be presented here, there are many similarities. In both, the input induces a unique configuration of coupled units. In both there are the equivalent of "key trigger patterns". In both there is a similar assumption of a sensori-motor configuration, so that motor action follows from sensory reactivation of the sensori-motor configuration. On the other hand there is an essential difference. Pringle's philosophy centres upon the idea of replication of units. Consequently the configurations caused by a particular pattern, or set of patterns, would grow as a mass with experience. One of the particular attributes of my model is the notion that essentially the record laid down by a particular event is "isolated" and unmodifiable, except in certain fairly obvious ways.

Kilmer, McCulloch, Blum, Craighill and Peterson (1968) prelude their model of the Reticular Formation by confirming Pringle's view that "any chunk of nerve tissue . . . can be viewed as an assembly of coupled non-linear oscillators" . . . Nevertheless, "we (Kilmer et alia) abandoned all coupled oscillator and neural net approaches to RF theory construction as utterly hopeless". Their model accordingly takes a form dictated by an interpretation of the reticular formation in canonical information theory equations. In philosophy the resulting model is quite close to the feed-back models. There is a sense in which certain feed-back weighting models have both hierarchical and lateral interactionist characteristics. In Kilmer et alia's (1968) model, the outputs from stage to stage converge. At the same time the transmission lines are subject to weighting modification. In this model of the Reticular Formation, the idea is that a

number of "modal" computers represent anatomical divisions of the core. These incorporate weighting functions from neighbouring computers, so simulating lateral coupling in addition to feed-back. It would seem that the model applies iterative network feed-back weighting techniques to a more general interactionist concept of neural function. In implementation the model resembles the iterated network expressed in canonical equations.

A similar marriage of lateral interaction and hierarchical convergence was used in a pattern recognition model by Taylor (1959). Taylor extended his ideas to a neurological interpretation of the structure of the neocortex not unlike the later interpretation by Marr (1969, 1970). An unusual aspect of Taylor's model was that it was essentially analogue.

Taylor's model is used in some detail in J.Z. Young's "A model of the brain" to illustrate his argument for an "engineering" approach to brain function. In spite of its title, Young's (1964) work is not so much a neurological model as the interpretation of his work on Octopus to suggest some principle of brain function. He is concerned with coding, filter stages and general organisation. He suggests that the vertical and horizontal pattern recognition attributes of Octopus may be mediated by the shape of the dendritic trees of specific "detector cells". He proposes simple paradigm models of their use to control a dichotomous "attack or retreat" pathway.

In past times both hierarchical and feed-back models were thought to suffer from severe limitations. Collateral interaction theories, such as those of Hebb and Pringle, arose mainly to overcome those limitations. It might now be argued by proponents of the switching models that many of the limitations they were supposed to suffer from have been shown to be inapplicable. For example, it was thought that the ability to recognise a gestalt pattern regardless of its size and position on the retina could not be accomplished by a switching model. In fact many such models have since been produced. It might be argued further that many of the neurological considerations which argued against switching models have also been reinterpreted. Lashley's well known experiments in search of the engram had made it extremely unlikely that a model of fixed connections could mediate learning. It is now known that there is a measure of localisation of memory. Nevertheless, something of Lashley's concepts of mass-action and equipotentiality remain. For one thing, the connectivity

of the brain is so complex that complete localisation of function is out of the question. One can attempt to explain this connectivity within an hierarchical model. Konorski (1967) enters into much clinical detail to do so. To some workers, however, it seems simpler to assume some fundamental lateral interaction principle rather than an ad hoc interconnective convergence. This matter will be looked at again in a later chapter.

Other advantages of the distributed memory approach to neural integration have been argued by Hebb (1949, 1955), Pribram (1971) and Feeney, Pittman and Wagner (1974). At root they concern aspects of brain function which demonstrate an elasticity and unification which seem foreign to other kinds of model. This impression may of course be misleading. Nevertheless, models such as those of Hebb and Pribram seem to many more plausible explanations of the overall integrity of action of nervous systems than do those based upon either hierarchical convergence or feed-back discrimination.

If hierarchical and feed-back models suffer a certain "credibility gap" in some respects, then lateral interaction models suffer a "credibility gap" in others. They have an indefiniteness of structure which gives them a vagueness in application. Although some have been simulated (Hebb's and Spinelli's), the demonstration was only of the general characteristics of the system. How such systems would actually mediate the intricate peculiarities of behaviour is far from clear. Moreover, systems like those of Beurle, Pribram, Pringle and Hebb, do not seem to possess sufficient structure to mediate the highly specific innately programmed aspects of behaviour. It is true that these limitations apply generally to all simulations in the present stage of development. But, in contrast to the simulations of distributed memory models, those of hierarchical and iterative networks actually do something, e.g. pattern recognition or emission of responses. Also, it appears quite clear what form further development of such models would take: they would be extended simply by including more transmission lines, more points of convergence of more functions of feed-back weighting adjustment.

This definiteness of network models appears to proponents of the distributed memory models as a limitation in itself. What attracts them is the very lack of rigidity in diffuse models. In such lack of definition they may see the possibility of the whole being greater than the part, in unknown attributes arising from the complexity.

One of the aims of my model is to pursue the simulation of a diffuse lateral connectionist model to a point where it loses some of the vagueness generally attributed to such models.

VI Models related to neuronal mechanisms

To give a general indication of the present state of the art of explaining brain function by a neurological model, it is perhaps necessary to mention a number of formulations which do not fit easily into the classification suggested so far in this chapter. Many of these models concern the neuronal processes by which it is suggested that nervous tissue incorporates experience.

My model, like many of those already described, is based upon a principle of organisation. It is not more demanding of neuronal mechanism than of most of such models, e.g. Marr (1969, 1970); Hebb (1949); Kilmer and Olinski (1974). Nevertheless, the interpretation of my model might be affected by the mechanism of neuronal alteration by which it is considered to be implemented. Essentially this may be seen as a question of the intimacy of the connections between neurons. My model proposes a co-operative network between neurons. That network could be implemented by simple reciprocal orthodox synaptic connections or by the more intimate kinds of connection disclosed by contemporary research, e.g. gap junctions, tight junctions, dendro-dendritic synapses (Schmitt, Dev and Smith, 1976; Reese and Shepherd, 1972; Gilula, 1975; Palay, 1975). To consider the model in the light of different possible mechanisms of neuronal communication would be premature at this stage. However, some mention of models utilising different neuronal mechanisms helps to complete the background.

At a formal level, Brindley (1967) classified modifiable and unmodifiable synapses according to the logical functions they generate. He listed ten examples of synaptic arrangments appearing in models, e.g.: "(9) Excitatory, facilitated but not excited by axo-axonal synapse". He produced paradigm networks capable of mediating conditioning, extinction and temporal separation of stimuli. In this respect his work has affinities with that of Uttley (1954, 1956) and Burke (1966) cited earlier. Brindley developed a number of theorems relating to the properties of networks. By extending his classification of synapses, incorporating "repeater cells" and delays, he is able to develop networks which he proposes provide analogies in principle of verbal sequence and association learning. It

would be going too far from our present concerns to consider them here. Like those of Burke, the models of Brindley are essentially logical abstractions. Marr (1969) acknowledged his debt to Brindley's (1969) paper.

Brindley exhaustively analyses the logical possibilities provided by different kinds of neuronal connection without becoming involved in the micro-biological implementations and controversies. There is however, one form of connection which he does not consider, that which would enable two or more neurons to act so that their excitation is effectively averaged. This would be the effect of reciprocal axonal-somatic synapses, or of dendro-dendritic connections, or perhaps tight junctions. The logical effect of such "lateral" interconnections is not easy to see and therefore they do not feature in neural network models. It is this kind of relationship required for my model.

Griffith (1967) developed a matrix algebra of connections to analyse a very different type of model from that presented in logical network theory. The basic idea is a biochemical principle of neuronal recognition. The model is embryological in presentation, tracing the structural and connective form of the nervous system in ontological development. Utilising the notions of genetic suppression and release of complementary enzymes, Griffith explores the possibilities of cell division on connectivity. Synaptic links from the axonal tip of cell A to the dendrites of cell B are assumed to be maintained by a growth substance. There are a number of such substances, and cell A will only synapse with cell B if that cell produces the particular substance which it requires. Upon division, the daughter cells may differ in enzymatic specificity from the parent or be the same. The set of daughter cells resembling the parent is called a "clone". Griffith uses variants of this basic idea to show how an orderly system of cells linked to each other would be generated. He then uses the resulting connective sets and principles to build a dynamic model of the nervous system. Only the briefest indication of the model can be given here. What perhaps cannot be conveyed here is the sophistication of the concepts of neural processing by the resulting system.

In Griffith's model a set of cells linked in an excitatory manner is called a "mode". Two "modes" may be linked through two sets of inhibitory meurons. When both such "modes" are active, then, because of

habituation, first one and then the other will dominate, forming a "neural oscillator". In the mammalian brain there is a "centre" (subcortical) divided into a set of N mutually exclusive "modes". When a "mode" in the centre is active it excites "certain corresponding cells in the cortex" which after a "very complicated transformation" in turn excites cells back in the centre. There is consequently a shuttling of patterns of excitation between "centre" and "cortex". In addition to "centre" and "cortex" there are "moderators", each of which has "modes" of its own. These control sequences of "modes" at the centre. The intention of these moderators seems to be the implementation of innately programmed behaviour. Learning is effected by modification of "synaptic area" by the intervention of "facilitating molecules". Molecule X facilitates excitatory synapses, Molecule Y, inhibitory synapses.

Griffith uses the model to discuss conditioning and some more general aspects of neural integration, e.g. sleeping and breathing. It is in fact possible to follow his accounts of these functions in terms of the overall operation of the model. Its actual operation to mediate these functions is less clear. As I understand the model, the "centre" is fairly rigidly structurally defined. The "cortex" acts as a kind of random scrambler. The "moderators" modulate the process. How they do this is not clear. It is difficult to grasp the rationale of the model. Is it purely random, or is the "complicated transformation" effected by the cortex, a directed process?

The model has been partly simulated in a computer. However, as far as can be discerned the simulation was in the form of the solution of canonical equations concerned with the formation of connections and the oscillation of modes. Consequently it is difficult to judge how the system would act as a whole. Would it really effectively mediate behaviour?

The model has been described at some length because it is one of the few which attempt to derive a highly defined global neural system from principles derived from contemporary research. The biochemical and enzymatic mechanisms utilised in Griffith's model are well established aspects of present-day genetic, immunological and developmental research. At the same time it does not seem to have been developed with sufficient specificity of detail to say whether it could perform what its designer

claims and, if it did, whether the resulting system would really resemble the brain.

Landauer (1964) summarises suggestions of biochemical mechanisms mediating learning: "an unspecified mechanism involving direct coding of information on RNA molecules in terms of base sequences (Gaito, 1961); enzyme induction with respect to transmitter substances (Briggs and Kitto, 1962; Smith, 1962); changes in proteins in terminal vesicles (Hyden, 1961); and changes in ionic binding which influence firing rates. Landauer himself offers a more highly defined neurological model. He suggests that glial cells produce RNA which renders the membrane of a particular cell receptive to particular frequencies of the electrical signals. Neurons become "tuned" to the neural activities representing the CS.

Another model which utilises the work of Hyden and his associates is that of Hendrickson (1972). He proposes that the timing of nerve impulses codes the sequence of nucleotides in the RNA chain. The code consists of the permutation of four possible interval values between successive action potentials in spike trains.

McIlwain (1979) has applied the results of recent discoveries of biochemical mediators of synaptic events to suggest a complexity of function of the individual neuron. Biochemical mediators, e.g. cyclic AMP and GMP, may last for some seconds, and even have effects for up to 30 minutes. He suggests that such mediators are cytoplasmically transported and activate in sequence a "precisely ordered system of synaptic connections". The augmentation of activity at the synapses by the modulators expresses "an expectancy based upon terminal placement". The idea is that the travelling biochemical mediators would simulate previous actual activiy. He develops the idea to deal with actual neural structures, such as "complex" cells in visual cortex. McIlwain's proposal seems important in stressing the possibilities of neurons being able to carry out truly complex functions. This contrasts with the simple functions attributed to neurons in most models including my own. I do not doubt that neurons can carry out such complex functions. However, it seems to me possible that such complex functions would be superimposed upon a simpler basic system rather than representing a fundamental mode of operation of nervous systems.

The mechanisms proposed by the models mentioned are not perhaps

implausible in themselves. Intra-cellular transport, the codon basis of RNA and DNA and complex interactions of macromolecules undoubtedly play an important part in neuronal function. The difficulty is in deciding exactly what that part is. To utilise the knowledge which is rapidly appearing from biochemical research it seems necessary to show how it fits into a neural organisation. A model proposing such an organisation it seems should show more than a likely neuronal mechanism, or even a likely principle of neuronal communication. It should also show a working system which bears a close correspondence to what is known in other respects about the brain.

VII Concluding remarks

In this chapter some attempt has been made to give a general picture of models which have been proposed in recent years to explain neural integration in the brain. In order to bring a vast amount of material within a reasonable short account, models were classified under three types: hierarchical, discriminative feed-back, and distributed memory. In addition some models which centre around the mechanism of neuronal alteration were mentioned. Individual models do not, of course, conform closely to the principles of any one category, nor can their contribution be conveyed by the relegation to one of a number of simple categories. Nevertheless it is hoped that the presentation has enabled some of the main differences between models to be seen.

This account has been intended to serve two purposes. First it may have served to remind the reader of other solutions to the problems which I have attempted to solve in ways characteristic of my model. In that respect it provides the customary introduction to work presented. However, it also serves the purpose of summarising the many insights which brain models have made towards the understanding of brain function. It does not seem necessary that brain models be seen as exclusive and mutually antagonistic explanations. Of course in some respects explanatory principles are empirically opposed. An hierarchical network is not the same as a distributed memory system, such as the neural hologram. Even in the case of such oppositions, it is possible that one principle acts in one location or level of processing, and another in a different situation. They may even act hand in hand. Although my model proposes a neural template system very different from an hierarchical network, the latter seems too good an idea for nature not to have utilised

it where appropriate. Moreover some models may be seen as exemplifying principles which are essentially independent of the form of model in which they are embodied. For example, the use of temporal delays to convert temporal sequence into spatial conjunction (Uttley, 1956b). Although feed-back modification of connections dependent upon results has been presented generally in a certain kind of model, it may be taken as a principle which is merely brought into prominence in that kind of model. In every model where the results of action leads to a modification of connectivity, the same principle necessarily operates.

Models may be seen as aids to understanding a very difficult subject matter rather than as different descriptions of what the brain is supposed to be like. At the present stage of development we may not be ready to propose a single model which definitively describes the function of the brain. It is hoped that the model to be presented will provide a contribution in this sense to our understanding of neural integration.

CHAPTER 5

OUTLINE OF A THEORY OF NEURAL INTEGRATION

This chapter gives the main elements of the theory of neural integration underlying the simulations. The presentation of the model is preceded by the citation of some contemporary research concerning the anatomy and histology of brain tissue which is intended to show that the requirements of the model are concordant with what is known of the structure of nervous tissue.

I The anatomy and physiology of brain tissues

A. Columnar organisation

The neural tissues of Man and the higher animals present an abundance of highly diverse neuronal constituents (Szentagothai, 1976) interacting in models of great complexity which are little understood. Amidst this complexity, what has emerged as a most widely applicable and well established feature of the mammalian cortex, and of other centres, is the columnar organisation.

In simple terms the functional organisation of the cortex consists of replicated columns, the cells of which are "activated by stimulation of the same single class of peripheral receptors" (Mountcastle, 1957). Mountcastle recorded from single units in the cat's somatic cortex using penetrating microelectrodes. The stimuli were mechanical deformations of peripheral tissue, skin, hair and joints. "The neurons encountered in a perpendicular traverse of the cortex are activated from almost identical peripheral receptive fields at latencies which on the average are not a function of the position of the cell in depth within the cortex". Modality pure columns were intermingled in a mosaic like fashion. Within the limits of the experimental numbers, there was a topographical mapping. Cells belonging to the different modalities were found in all cell layers. Cells driven by joint movement "signal the steady position and phasic chnages in positions of joints and are suitably arranged to subserve position sense. Pairs of closely adjacent cells were seen to respond reciprocally to alternating joint

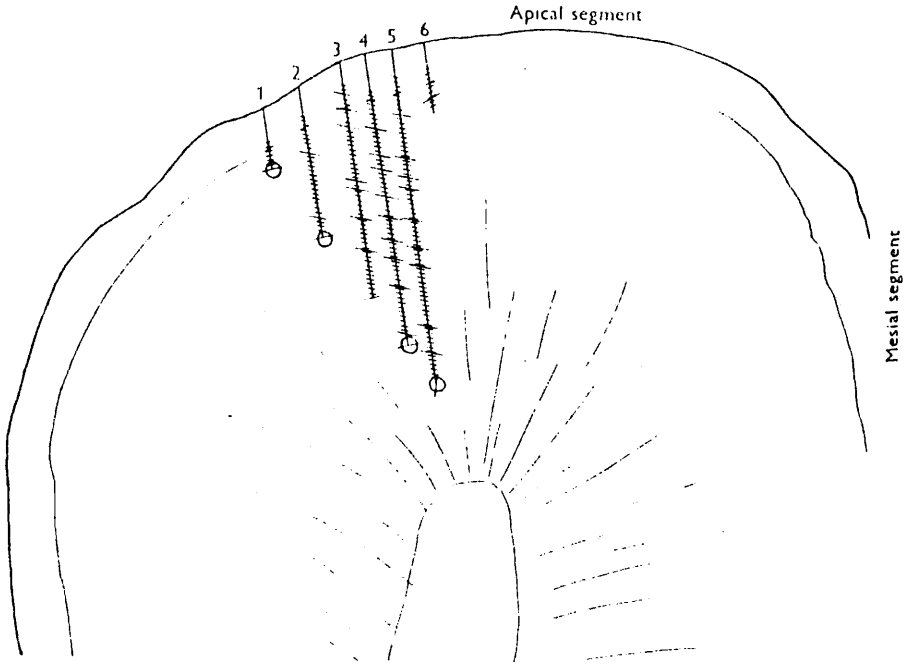

Figure 5.1 (Reproduced from Hubel and Wiesel, 1963a) Reconstruction of two parallel micro-electrode penetrations in the anterior part of the post lateral gyrus. Longer lines intersecting the electrode tracks represent cortical cells. Receptive-field axis orientations are shown by the direction of these lines. Shorter lines represent regions in which unresolved background activity was observed.

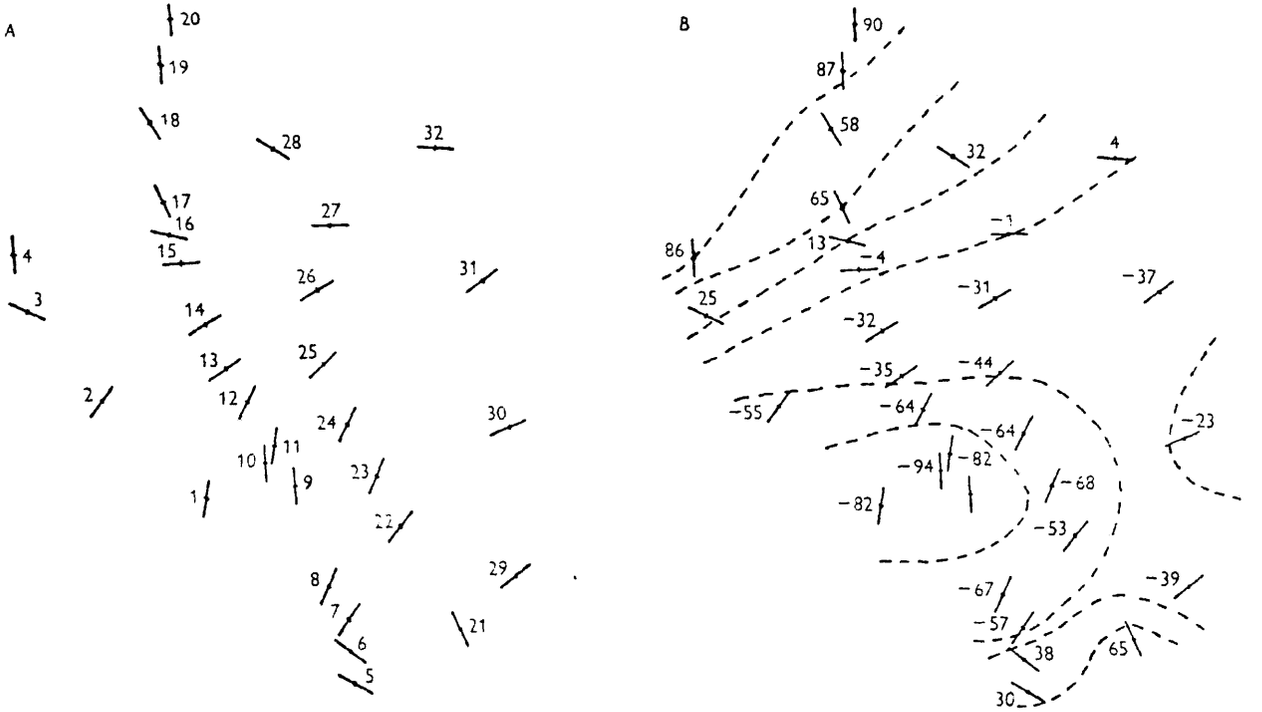

Figure 5.2 (Reproduced from Hubel and Wiesel, 1963a) Photograph of a region in post lateral gyrus mapped by Hubel and Weisel in their experiments on receptive field orientations. Orientations are shown for 32 superficial penetrations.

movements, a form of reciprocal afferent innervation". The cells of the vertical columns "are related to peripheral fields which are very much alike in size and shape and which occupy more or less the same peripheral area". ". . . when slanting penetrations are made . . . the receptive fields of neurons successively encountered shift gradually, in conformity with the surface pattern of representation directly above the successive positions of the electrode tip in the depths". Mountcastle suggests that the width of the columns has an upper limit of about 0.5 mm.

The work of Hubel and Wiesel (1959, 1962, 1963(a) & (b)1968; Wiesel, Hubel and Lamm, 1974; Drager and Hubel, 1976) is too well known to require elaboration here. The relevance for the present context is the finding for the visual cortex of the cat and monkey a columnar organisation similar to that discovered by Mountcastle for the cat's (Mountcastle, 1957) and Monkey's (Powell and Mountcastle, 1959) somatic cortex. As is well known, Hubel and Wiesel discovered neurons responding to specific receptive fields, especially oriented bar stimuli (see sub-section C below). Each column was found to be organised in relation to a single orientation of the stimulus. The receptive fields extended about 4 degrees. They suggested a range of 100 μ to 2 mm for the width of the columns (Hubel and Wiesel, 1963). Figure 5.1 (Tex-fig 1 of Hubel and Wiesel, 1963) illustrates two parallel penetrations of electrodes about 175 μ apart. Hubel and Wiesel report that the transition of orientation from one column to the next is not continuous but in regular shifts. Figure 5.2 (Hubel and Wiesel, 1963) shows the surface shapes of columns and the receptive fields of sample penetrations, and the general distribution of orientations in a small area of striate cortex. These figures illustrate the appropriateness of the representation for the mediation of the "letter stroke" template described in Chapter 2 (see Figure 2.4).

In the somato-sensory region of mouse cerebral cortex, Woolsey and Van der Loos (1970) found a columnar organisation taking the form of "barrels". The walls appeared densely populated with cells, with relatively sparse "hollows" within, and separated by sparse divisions. They are centred on Layer IV and are about 250 μ across and 300 μ high. Each of these barrels represents the sensory innervation of a single contralateral whisker follicle (Van der Loos and Woolsey, 1973). The barrels were further investigated in rats and mice by Killackey, Belford, R. and D. Ryugo (1976) who found that whisker removal in the newborn

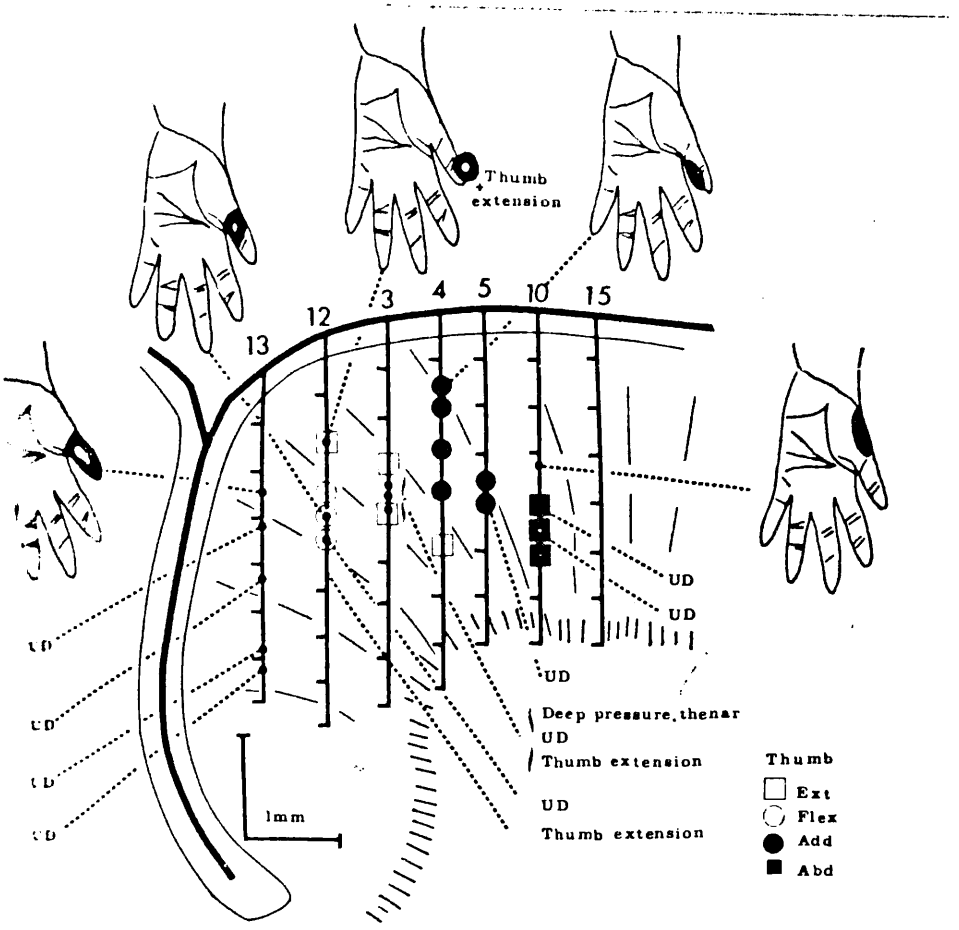

Figure 5.3 (Reproduced from Asanuma, 1975) Asanuma and Rosen's 1972 reconstruction of the electrode tracks and associated input and output relations in the thumb area of monkey motor cortex. The cells activated by tactile stimuli were shown to have receptive fields within cortical efferent zones projecting to finger muscles. Each efferent zone was found to receive tactile inputs exclusively from an area of finger or hand located in the direction of movement produced by stimulation of the zone.

animals disrupted barrel development.

The columnar organisation of the motor area was demonstrated by Asanuma and his associates (e.g. Asanuma and Sakata, 1967; Asanuma and Rosen, 1972) and by Welt, Aschoff, Kameda and Brooks (1967, cited in Woolsey and Van der Loos, 1969). Asanuma et alia mapped corticospinal projections which facilitate or inhibit monosynaptic reflexes of foreleg muscles of the cat. "The low threshold points were confined within a cylinder of cortex which extended along the radial fibre into the gray substance. The diameters of the sector ranged from 0.5 to a few millimetres and the fringes overlapped". Asanuma (1975) reviews work carried out by his group indicating that the sensory feed-back from movement is mapped into columns close to the efferent column. Afferent signals reached the efferent zone from deep structures such as periosteum, joint and connective tissue within the muscle. A given efferent zone in the hand of the monkey received proprio-ceptive inputs from joints or muscles which that zone projected. (See Figure 5.3 from Asanuma and Rosen, 1972). (This work is relevant to topographic mapping, and is complimentary to that of Werner et alia, cited at B, below).

There remains the cortical areas receiving cortico-cortical fibres. The columnar organisation of somato-sensory cortico-cortical fibres was demonstrated by Jones, Burton and Porter (1975) and Shanks, Rockel, A.J. and Powell, T.P.S. (1975). Because they felt that the columnar organisation might still be attributed to a representation of the receptor surface, Goldman and Nauta (1977) investigated by autoradiography cortico-cortical fibres in the frontal association, limbic and motor cortex of developing Rhesus monkey, including the U fibres joining gyri. "Labelled axons originating in these various regions of the frontal lobe have topographically diverse ipsilateral and contralateral destinations, but virtually all these projections share a common mode of distribution: they terminate in distinct vertically oriented columns, 200-500 μ wide that extend across all layers of cortex and alternate in regular sequence with columns of comparable width in which grains (of label) do not exceed background".

With regard to sub-cortical organisation, Szentagothai (1973a) reviews evidence of projection columns in the lateral geniculate nucleus (LGN) and concludes that they consist of overlapping cylinders containing Golgi type II cells (cells whose axons terminate locally) which in turn

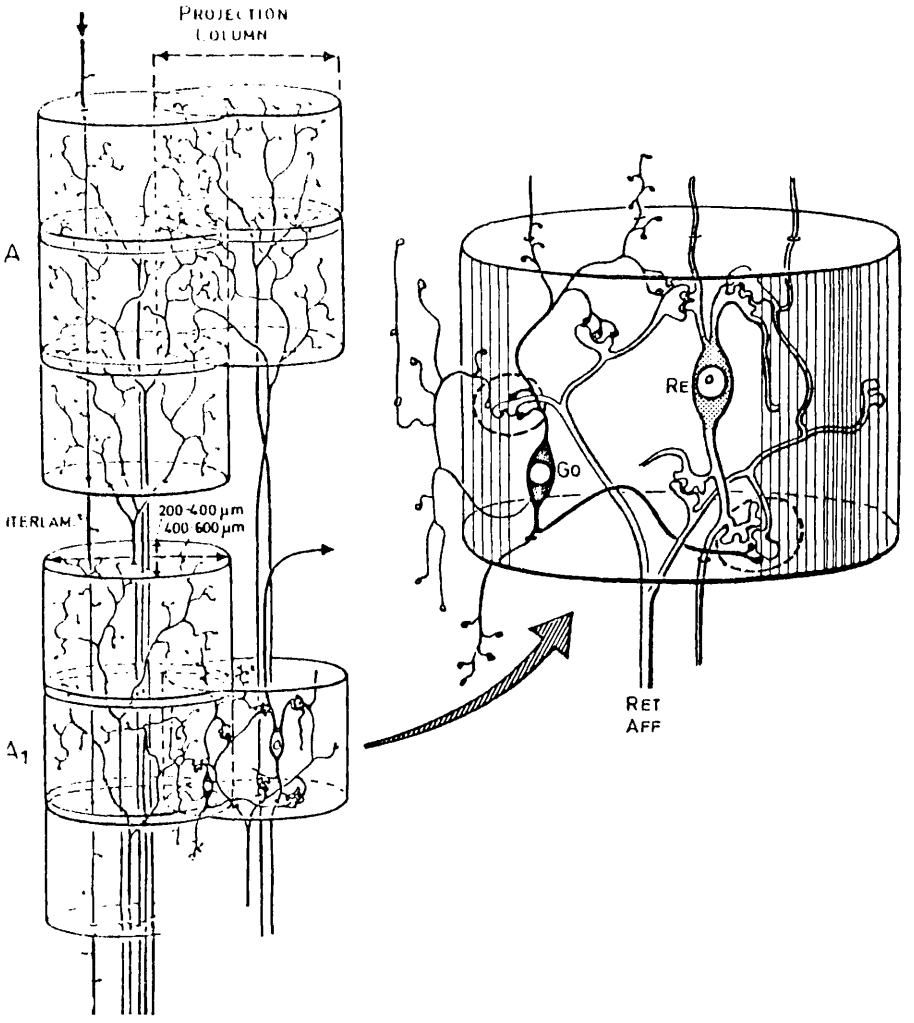

Figure 5.4 (Reproduced from Szentagothai, 1973a) Szentagothai's hypothetical reconstruction of the lateral geniculate nucleus, illustrating the columnar structure. Diagram at left illustrates over-all arrangement in the termination of bundles of retinal afferents and their relation to relay cells (Re) and Golgi 2nd type cells (Go).

synapse upon relay cells (see Figure 5.4 from Szentagothai, 1973a). A number of authors (Hubel and Wiesel; Bishop, Kozak and Levick; and Vakur; reviewed by Freund, 1973) consider that neurons of the LGN have a common direction and are arranged along bundles of optic tract axons originating from the retina. The topographic mapping of the LGN indicates that neurons are organised into local neuronal pools, if not columns, and this is one of the basic assumptions of my model.

An interesting columnar organisation is reported of the Superior Colliculus. Drager and Hubel (1976) showed that in the mouse, the three upper layers responded only to visual stimuli, whereas the lower layers are driven by somatosensory and auditory stimuli. Each of these representations is topographically mapped. The scale of the mapping differs with modality, so that in any column are represented visual and somatosensory stimuli which more or less correspond to the external environment. The environment is coded in terms of the visual co-ordinates of (that) eye and other sensory modalities merely endorse the visual plane. Stein, Megalhaes and Kruger (1976) report almost identical results for the cat Superior Colliculus. Tactile representation revealed a somatotopic plan which was in register with the overlying visuotopy.

In my model it will be proposed that although all the cells within a column are directly activated by afferent signals, either monosynaptically or via one or two synapses only, there is a lateral distribution of excitation across column boundaries. This lateral connectivity is either polysynaptic or by some non-synaptic mechanism of neuronal communication. Communication between columns is also by means of radial fibres of stellate cells which exit from the cortex (see later chapters).

Relevant to the proposed lateral communication may be Mountcastle's (1957) comment that the analysis into columns was based strictly upon the first response of cortical cells to a repetitive stimulus. According to my model the first response would reflect the afferent specificity, subsequent responses the integrative effects. Many studies show rather a change of threshold across columns (Albus, 1976; Lee, Heggelund, Holme and Creutzfeldt, 1976) rather than the all or none effect which would be expected from a totally radial (vertical) organisation. This gradation of threshold is to be distinguished from the discrete organisation of columns: the latter relates to the allocation of neurons to a particular

class as a consequence of their maximum response characteristics. Burns et alia (1962) in their investigation of the effects of an oscillatory dark boundary stimulus showed that the response of a particular single neuron responded in some degree to all positions of the stimulus, although maximally in a single position.

In an important paper, Mountcastle (1978) examined the parietal lobe Area 7 of primates. This area is associated with Body-Image, visuotopy and eye-movements. He found evidence of mini-columns of about 30 μ radius. Each hypercolumn of Hubel and Wiesel (about 500-1000 μ) contained about 225 minicolumns. Mountcastle estimates about 140,000,000 minicolumns in the neocortex of man. Other aspects of Mountcastle's paper relate to the response characteristics of individual neurons and to possible neuronal wiring implications. These are dealt with in the next two sections. It must be mentioned that although his data and even his general interpretations are cited in support of my model, his own conclusion differs. He proposes that each column is an independent module without lateral communication, an idea close to that proposed by Szentagothai (1973) and Scheibel and Scheibel (1970).

In summary, the cerebral cortex of mammals appears to be organised in discrete columns perpendicular to the surface. This applies to sensory projection areas, motor areas, afferent innervation to motor area, and "intracortical fibre distribution of all the conventionally distinguished categories of the cortex". (Goldman and Nauta, 1977). The size of columns vary according to report in the 0.3 to 1.5 mm range. These columns do not appear to be entirely unitary (e.g. Mountcastle, 1978; Hubel and Weisel, 1972; Asanuma, 1975). Asanuma suggests that these large columns represent several afferent columns measuring between 0.1 to 0.4 mm which together control an individual muscle in a column about 1 mm wide.

B. Topographic mapping

An aspect of neural organisation closely associated with the columnar structure is the topographic mapping of sense organs at all stages to the cortex. Summarising the results of a "prodigious burst of investigative work along traditional hodological lines", Palay (1972) comments: "These studies emphasised the precision with which the nervous system is organised. For example, the retention of the topographic order as the synaptic links were followed from level to level became more impressive".

This mapping is not directly represented by any particular computational feature in the simulations of the proposed model. It has, however, a considerable relevance to its rationale. It may be seen as mediating a relational equivalent to the isomorphism of earlier models of the Gestalt school (e.g. Kohler, 1942). Implicit in the proposed model is the idea that the topographical mapping enables spatial templates to be constructed: these templates may act as models of the external world within the brain (Craik, 1943).

Almost all the work cited in the previous section on columnar structure discloses also that both the sensory field and the effector array is mapped precisely into the nervous system. Earlier work on the general representation of the body surface and the "sensory homunculus" (Penfield and Rasmussen, 1952) needs little mention. This work has been extended recently especially by Werner, Whitsel, Mountcastle (e.g. Werner, Whitsel and Petrucelli, 1972; Werner, 1970). Werner et alia (1972) demonstrated the regular mapping of parts of the digits, palm and leg surfaces of the macaque somatosensory area I, disclosing the customary columnar arrangement. This work complements that of Asanuma and Rosen (1972). Werner et alia developed a scheme of topographical mapping. Iwamura and Tanaka (1978a) confirmed that receptive fields such as the pad and digit of the cat are topographically mapped onto the somatosensory cortex in a focal sense. They also demonstrated a second representation of the forepaw with a different receptive field organisation from the first (1978b).

The Odour perceptual system is topographically mapped (Moulton, 1976).

Retinotopic mappings of the retina on the cortex have been found by Talbot (1942), Cowey (1964) and Daniel and Whitteridge (1961) in the monkey; Choudbury (1978) in the rabbit, where it develops over the first 14-21 days of life. A mirror-like topographical arrangement of area 17 and 18 with adjoining location of the vertical meridian was found. It may be relevant to theories of pattern recognition and the confusion of mirror reversals, that Hubel and Weisel (1965) report a third retinotopically organised projection. Kinston, Vadas and Bishop (1969, cited in Freund, 1973) report a second mirror image of the visual field in the medial interlaminar nucleus of the cat.

A retinotopic organisation of Area 18 in the rat was not changed by lesions in Area 17 (Olavarria and Torrealba, 1978). If area 18 was

innervated in an hierarchical manner from Area 17, as has been suggested by proponents of hierarchical models (see Chapter 4), one would have expected degeneration of the retinotopic representation in Area 18. This supports the interpretations by Baizer, Robinson and Dow (1977) of their experiments.

The topographical mapping of the Superior Colliculus is very precise (Drager and Hubel, 1976; Stein et alia, 1976); Weiskrantz and his group (Brooks and Jung, 1973) demonstrated that cortically ablated monkeys can locate and grasp objects when only their brain stem visual connections including their tectum opticum remain intact. Brooks and Jung also cite work in the same volume (Sprague, Berlucchi and Rizzolati, 1973; Doty, 1973) which suggests that in cats and rabbits considerable form vision may be mediated in the tectum. They propose that the tectum acts in close co-ordination with the visual cortex.

A similar conclusion may be drawn from the organisation in the Superior Colliculus of eye-movements (Mohler and Wurtz, 1977). As mentioned in Sub-section A, there is a layered visual, tactile and somato-topic topographical mapping in the Superior Colliculus, which does not, however, arise from direct sensory projections (Drager and Hubel, 1976; Stein, Magalhaes-Castro and Kruger, 1976). In Chapter 2 it was suggested that templates might be superimposed so as to employ the topographical mapping of projection fibres with maximum economy. In the Superior Colliculus it would appear that Nature may have gone one better. By superimposing templates immediately upon each other with adjustments of scale to ensure coincidence with external events, the greatest measure of integration would be achieved of externally co-ordinated events with the shortest possible connections.

C. Feature analysis and receptive fields

The concept of feature analysis has been prominent in many approaches to perception and pattern recognition (see Chapter 2). The discoveries by Lettvin, Maturana, McCulloch and Pitts (1959) of selective signalling by the frog's retina, by Young (1964) and Sutherland (1968) of feature detection of octopus and goldfish, and the demonstration of structured receptive fields by Hubel and Wiesel, cited above, played a large part in this conceptualisation. If feature analysis is used to imply an explanation of how the nervous system recognises high level characteristics of objects, e.g. noses and legs, the idea should perhaps be

treated with reserve. The term "feature analyser" implies that there are inbuilt neural arrangements by which we recognise aspects of objects, whereas the organisation disclosed by neuro-anatomy and physiology is of elements, such as orientated lines, spots of colour, only.

Work on visual receptive fields in the Lateral Geniculate Nucleus is reviewed in Freund (1973) and in respect of the cerebral cortex, in Brooks and Jung (1973, same volume). In the lateral geniculate nucleus, the receptive fields are generally concentric with either an excitatory centre and inhibitory surround, or vice versa. In the cortex, Hubel and Wiesel encountered primarily orientation selective neurons (Hubel and Wiesel, 1962). They characterised as simple fields those having, for example, excitatory and inhibitory strips lying side by side, with fine orientation selectivity. They are binocularly activated, but not necessarily without retinal disparity. They prefer slow moving stimuli. They are recorded most often in Layer IV (see next section).

The edge of a moving bar elicits the response (Henry and Bishop, 1971). Moving stimuli in all non-preferred orientations were shown to inhibit the simple cells. Brooks and Jung point out that only a small proportion of the total neuronal population representing a restricted region of the visual field will be activated at a given moment by an appropriate stimulus.

Neurons responding to a wider range of angular displacement and not so easily mapped into excitatory and inhibitory regions, were classed as complex. Of relevance to the next section (D), Hubel and Wiesel found that simple cells were generally located in Layer IV, and complex cells in upper layer of the same column.

Hubel and Wiesel (1965) reported complex fields in Areas 18 and 19. Special higher order receptive fields were encountered in Areas 18 and 19 which responded selectively to a variety of complicated angular forms. They inferred that these were mediated by hierarchical neural convergence from Area 17, but as already noted, Singer and Tretter (1976) and Baizer et alia (1977) consider that this is not so. Bazier et alia divide area 18 cells into six groups, orientation, colour, border, direction, spot and light inhibited. Although they do not report upon neurons of the hypercomplex type, they do consider that most of the orientation type cells encountered were of Hubel and Wiesel's complex type.

Zeki (1978) has found that whilst cells responsive to various aspects of the stimulus are found in all parts of five prestriate areas of rhesus monkey cortex, there is very considerable specialisation from area to area in, say, orientation detection, colour or motion.

The analysis by the nervous system of information relating to modalities other than vision parallel the feature analysers found in that system. The work cited in connection with columnar organisation and topographical mapping also indicates that signals relating to touch, hair movement, movement of whiskers, deformations of tissue, movement and positions of joints report isolated aspects of overall sensory or motor configurations. Single neurons respond to one particular form of stimulation. Their receptive fields relate to small local areas of the sensory field.

Mountcastle (1978) found that in the parietal lobe of primates, individual neurons in a mini-column organisation reported specific behavioural correlates. He distinguished neurons responding during (i) visual fixation, (ii) visual tracking, (iii) saccade, (iv) light sensitive (v) special and (vi) hand manipulation and projection.

Figures 5.2 and 5.3 illustrate what may be termed the "atomic" character of sensory organisation in vision and somatic sense respectively. In a convergent hierarchical organisation such atomic elements would fall at the lower levels of the hierarchy. If the hierarchy were organised within a column then one would expect to find cells within each column which represented higher order convergences beyond the atomic level. Although some "complex" visual fields are found within projection columns, these are still essentially atomic and relate to the topographic locality. There is nothing reported corresponding to a gestalt or even a configurational convergence of inputs from the whole visual or somatosensory field. If the hierarchy were mediated from area to area of cortex, then it is surprising to find that Areas 17, 18 and 19 are organised in the same atomic manner.

Looking at the feed-back adjustment kind of model, feature analysers would be formed by converging a variety of random samples of inputs and varying the weightings given to the samples. In that case, one would expect to lose the topographical mapping and to see the "successful" feature analysers differing from element to element. In contrast, there is a regular variation of the character of fields from column to

column, the variation relating to only one parameter which changes in some precise characteristic, e.g. orientation. The atomic nature of the columnar organisation and the precise ordering of the representation is as difficult to reconcile with this kind of model as with the hierarchical.

In comparison, the preservation of the localised atomic nature of receptive fields would be ideal for a template system. In each set of fields, e.g. orientation bars or joint-movements, a sub-set of the reporting neurons chosen from selected columns could be combined to form the basis of a template representing some aspect of the whole stimulus configuration. The subsets of different analysers can be considered as providing separate template systems, each looking at the total stimulus event in a different way. These template systems may be set side by side, that is to say, in adjacent areas of cortex, as in the case of Areas 17 and 18, or may be overlaid as in the Superior Colliculus or parietal cortex. In area 18 colour, border and orientation elements are intermingled in a topographical mapping. The sub-sets of different analysers in such an area may be considered as templates stacked close together to take advantage of the overall topographic projections.

The atomic nature of columnar representation appear only during investigations in which it is being investigated, that is when the stimulus is appropriately controlled. It is the first response. The template interpretation is supported by the contrasting complex neural response of cortical neurons during normal complex sensory stimulation.

D. Tissue structure in relation to the model

"It would indeed be unrealistic to underestimate the difficulties that any attempt at constructing the interneuronal network is bound to meet in a structure of such lack in apparent regularity and such an abundance of highly diverse neuronal constituents" (Szentagothai, 1976). Lorente de No identified more than 40 basic cell types (Jones, 1975b). Recently there have been a number of exhaustive studies utilising contemporary techniques (Valverde, 1971; Hubel and Wiesel, 1972; Lund, 1973; Szentagothai, 1973, 1975, 1976; Jones, 1975a and b, to mention only the most relevant).

These studies have greatly clarified three aspects of cortical neural structure: (a) the morphology of the various cell types (b) the

distribution of these cells, and therefore the clarification of the Layers of the cortex, and (c) the termination of the afferent fibres. By comparison, there are two aspects which remain more or less a mystery: (i) the morphological basis of the columnar structure and (ii) the actual inter-relationships by which neural processing is mediated.

In briefly summarising this work, it is hoped to show that the evidence is consistent with the model to be proposed. Essentially this is that afferent fibres synapse upon a large number of neurons within a localised region. Some of these neurons are assumed to act simply as distributing neurons, transferring excitation to sub-sets of neurons within the column layer. A large population of integrating neurons are to be regarded as excited by afferent terminals. These integrating neurons are of different morphological types, the variation of which, in accordance with the model, is directed towards ensuring a wide distribution of inter-connection between them. Thus as will be seen from the following review, the majority of such neurons have locally organised dendritic and axonal arborisations. Other integrating neurons at or near the termination of the afferent fibres in Layer IV have ascending axons ending in diffuse arborisation, whilst others in upper layers (i.e. nearer the cortical surface) have descending axons. Some have axonal and dendritic fields extended to great distances horizontally. Many neurons have axons and dendrites extending into Layer I, which forms a dense neuropil (see Jones, 1975a and b). According to the model, these varied shapes enable the integrating neurons to interchange excitation with each other in a dense network which extends across column boundaries. This explains the failure of research to find a morphological equivalent of the column boundary other than at afferent level, or efferent level. According to the model, the column is purely the effect of afferent and efferent organisation.

For the purposes of considering the model, the integrating neurons may be loosely identified with stellate-like neurons. Many stellate neurons have long fine axons which leave the local region of cortex. According to the model, these enter the cortex elsewhere, and thus join up the network of integrating cells in one location of cortex with the network of another. It will also be tentatively suggested that such axons form part of the reticular network of the brain, the large

disc like neurons of the reticular formation acting as integrating neurons. The model suggests that the integrating neurons may be considered as an interconnective network extending throughout the entire brain.

A similar columnar structure may be seen to underly efferent organisation. The model suggests that corresponding to the columnar afferent termination, the integrating neurons drive the pyramids. The pyramids in upper layers may act partly as collecting cells feeding the larger pyramids of Layers V and IV, and partly as a fine differentiation of the columnar structure. In approximate terms, the columns are suggested to be functionally organised with a main pyramid in Layer V acting as an efferent neuron, and a number of other secondary efferent neurons from subdivisions of the column, the total efferent neurons constituting a bundle. In motor areas, the efferent organisation would be organised around a Betz cell.

In addition to the afferent terminals, integrating neurons and efferent neurons, the model requires a variety of inhibitory neurons. These are assumed to be reactively stimulated by the neurons already described.

Before considering the data to support this arrangement, it is perhaps necessary to anticipate by stating that my model distinguishes two kinds of connections between neurons. One kind, termed for convenience morphological connections, mediate the structural layout of the brain, e.g. afferent, efferent connections and columnar organisation. The other kind of connection, integrative connections, mediate learning and genetically determined connectivity analogical to learning. Some very important recent evidence relating to the termination of afferents upon spiny-dendrite stellate cells, may be speculatively fitted into this scheme. It is speculatively suggested that the spines which may be seen on the dendrites of the stellate-like cells which receive afferent terminals, and the spines on dendrites of pyramidal cells indicate morphological connections, the integrative connections being without spines.

E. Cytoarchitectural Data

Detailed light microscopic studies of the cortex disclosed a great variety of forms densely packed (Cajal, 1899; Lorente de No, 1949). There was little to suggest that the cortex would be found to be functionally organised in columns. The laminar structure seemed to contradict the

physiological findings of Mountcastle and Hubel and Wiesel.

One of the difficulties in clarifying the structure was the lack of information concerning the termination of the afferent fibres. Since electron microscopic studies showed degenerating afferent terminals mostly on dendritic spines, it was assumed that these were the dendrites of pyramids. Scheibel and Scheibel (1970) and Szentagothai (1967, 1971 and 1973) were amongst those who evolved a "modular concept" of columnar organisation on this supposition. The clarification of the distribution of the various cells in the Layers and the precise location of the afferent terminals made it more probable that the afferents terminated on a certain kind of spiny dendrite stellate cell (Valverde, 1971; Hubel and Wiesel, 1972; Lund, 1973; Jones, 1975 and b; Szentagothai, 1975).

In visual cortex of monkey, Valverde found stellate cells in Layer IV. "Short axon cells have been found in clusters or glomeruli upon which terminal branches of cortical afferents have been observed. It is suggested that periglomerular stellate cells with ascending axons connect clusters of clewed* cells with small, medium and large stellate and pyramidal cells of Layers II and IVa, which in turn activate the giant solitary cells of Layer V through synaptic contacts along vertical descending axons" (Valverde, 1971).

Valverde considers that Layer IVc is about 250-300 μ in thickness and the lower limit is 950-1,200 μ from the pial surface. "In Nissl preparations it appears filled with small round bodies with scanty cytoplasm and 6-10 μ in diameter". "A columnar arrangement in lower levels is most patent in Nissl sections. The cell bodies appear piled up forming vertical columns of eight to ten elements, or are grouped into elongated clusters (glomeruli) with major axis in a radial direction, and separated by clear vertical narrow passages . . ."

Hubel and Weisel (1972) using an original technique of limited lesions in the LGN, were able to show the termination of afferents from the parvocellular layers to Layer IVc with a second minor input to Layer I. Magnocellular layers of the LGN sent afferents only to IVb. They concluded ". . . Certainly the results of the present study, given the dense packing in Layer IVc of an almost pure stellate population (my italics) provide little encouragement to the idea that the afferents end primarily upon pyramidal cells".

* Valverde defines "clewed" cells as having tangled dendrites resembling a ball of thread, sometimes called spider cells or granular cells.

Incidentally, by destroying separately the afferents from either eye, they were able to demonstrate two orthogonally organised systems of columns or slabs, mediating respectively the orientation of receptive fields and ocular dominance. They also qualified earlier assumptions concerning columns, by demonstrating clear differences in responses to visual stimuli from different layers. "Not surprisingly, cells of Layer IV . . . have the simplest properties and show the least intermingling of inputs from the two eyes. The more complex cell types, for the most part binocularly driven, are found in the layers above and below".

Lund (1973) divides the neuron population of Area 17 of the monkey into pyramidal neurons, stellate neurons with spinous dendrites and stellate neurons with spine free or sparsely spined dendrites. The spinous dendrite stellate cells are restricted to lamina IV. Lund found three groups of thalamo-cortical axons, synapsing respectively to lamina 1, the upper half of IVc, and the deeper half of IVc and to IVa.

Jones published two exhaustive researches on the somatic sensory cortex of monkey, one (1975a) dealing with distribution of afferents, the other (1975b) with the morphology and distribution of non pyramidal neurons. A feature of the latter was its quantitative assessments. Most of his work confirms that reported for other areas of cortex, particularly the striate, in primates and cats. Two differences were (a) the relatively low numbers of spine bearing stellate cells of the type reported in layer IV of the striate cortex. The majority of cells found by Jones did not have a true star shape, but had a major dendrite from Layer III; and (b) the thalamic afferents terminate not only in Layer IV, but also in Layer III. His research also covered the motor area, and in that area terminals fill much of Layer II. In all areas a small proportion also end in layer I.

In support of the proposed tissue structure outline for my model, it may be worth summarising Jones' nine classes of cell: Type 1 are large multipolar cells with ascending axons found predominantly in layers III and V. The dendrites may extend over three or more layers and over 600 μ. Type 2 resemble type 1, but in miniature. They have thin sinuous dendrites in fusiform or ovoid fields. By contrast with the mainly horizontal disposition of type 1 cells, the terminal axon branches of type 2 are oriented primarily in the vertical direction. Type 3 cells are small

multipolar or bitufted cells with vertical axons as long as 1,000 μ. The somata are small and round and found exclusively in layer II and the upper part of III. Type 4 cells are small with radially arranged dendrites and local bush like axonal ramifications. They have true stalked dendritic spikes, although sparse. Those with descending axons are in Layer II, those with ascending in Layer IV and III. Type 5 cells have round somata in Layer II and pronounced dendritic tufts in Layer I. The axons first descend then ascend, giving rise to horizontal and recurrent branches at each stage.

Type 7 cells are spiny cells with recurrent axons found only in Layer IV. They include both the "star cells" of Cajal, and the "star pyramids" of Lorente De No. There is no basal dendritic display as in pyramids, the soma being surrounded by dendrites radiating in all directions. The axon arises from the deeper surface of the soma, then divides into two or three branches "which recurve very sharply and often divide again, so that up to six vertically oriented branches ascend over the cell somata and into Layer III".

Type 8 cells are modified pyramids, the dendrites bearing true spines. The axon descends into the white matter. In addition to these eight types, Jones found only three examples of a giant type of cell. The somata lay at the junction of layers II and III, and the longest dendrite measured 600 μ.

Jones considers that thalamic afferents basically terminate on type 7 cells in Layers IV and part of III. He points out that the cylinder of cortex which could be influenced by a single type 7 cell is too small to account for the column size of the order of 300 μ. To this must be added the density of packing of such cells. "The number of type 7 cells contacted by a single afferent, however, remains unknown". The tentative interpretation for my model, is that the afferent fibres directly contact a number of type 7 neurons, which act primarily as distributing neurons to the integrating cells, which are the class of stellate cells generally. The extent to which neurons receive direct rather than indirect thalamic afferentation is reflected by the number of dendritic spines.

The modular concept of Scheibel and Scheibel and Szentagothai was mentioned above. Scheibel and Scheibel had argued that "all the elements are exposed to the same input, although the degree of

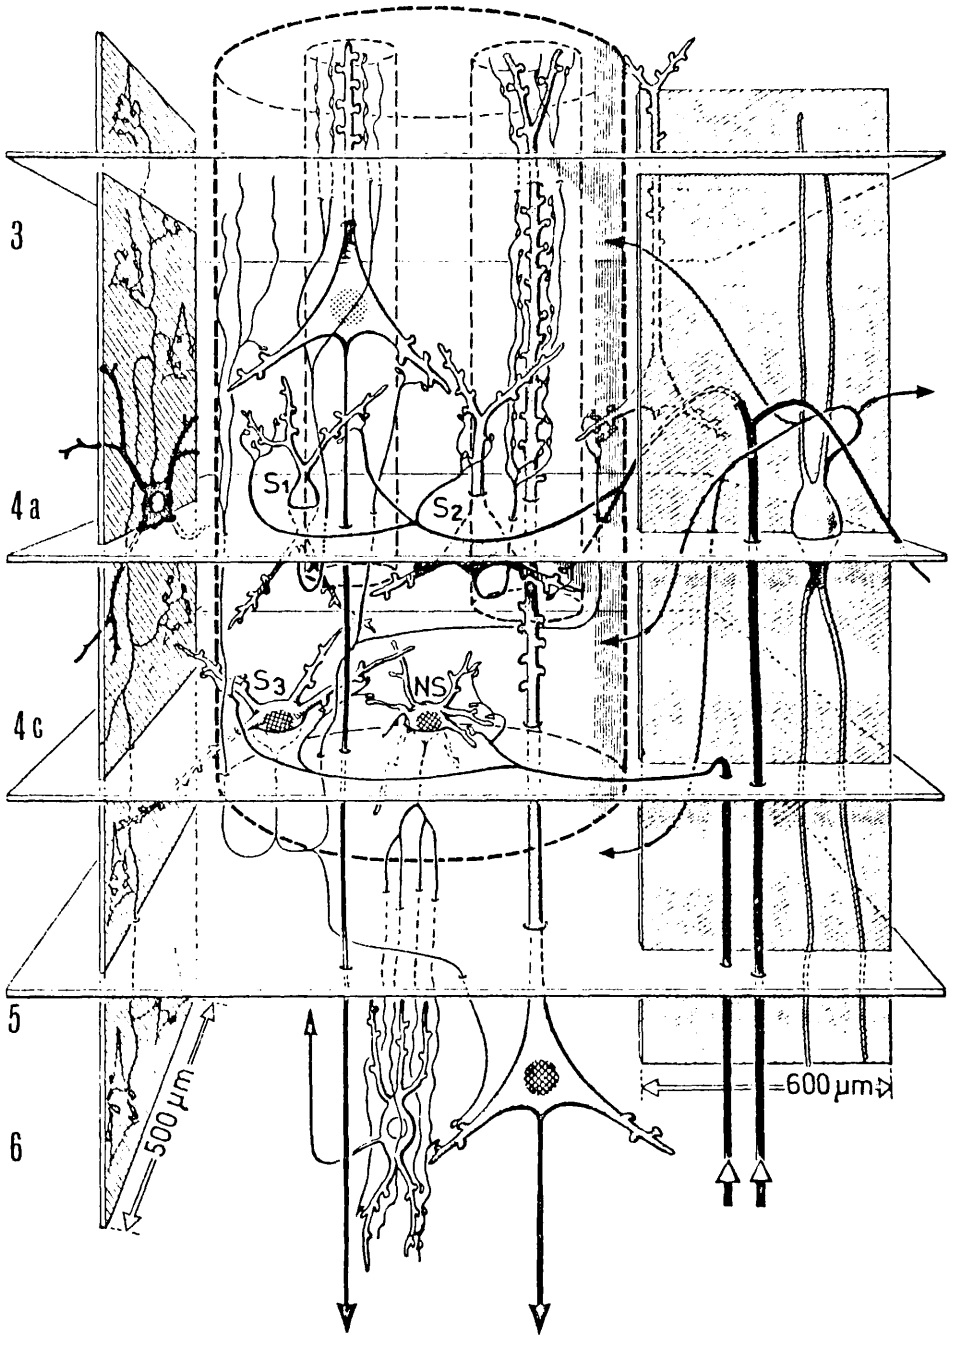

Figure 5.5 (Reproduced from Szentagothai, 1975) Szentagothai's modified hypothetical elementary neuron circuit incorporating evidence of the termination of afferents. Specific sensory afferents (heavy black vertical lines at right) terminate (separately) on spiny stellates S_1, so called star-pyramids S_2, and another type of spiny stellate S_3 as well as on neurogliform non-spiny stellates NS, in Layers IVa and IVc. The vertical plane at left shows the strictly oriented axonal arborisation of a large basket cell, which Szentagothai considers inhibitory.

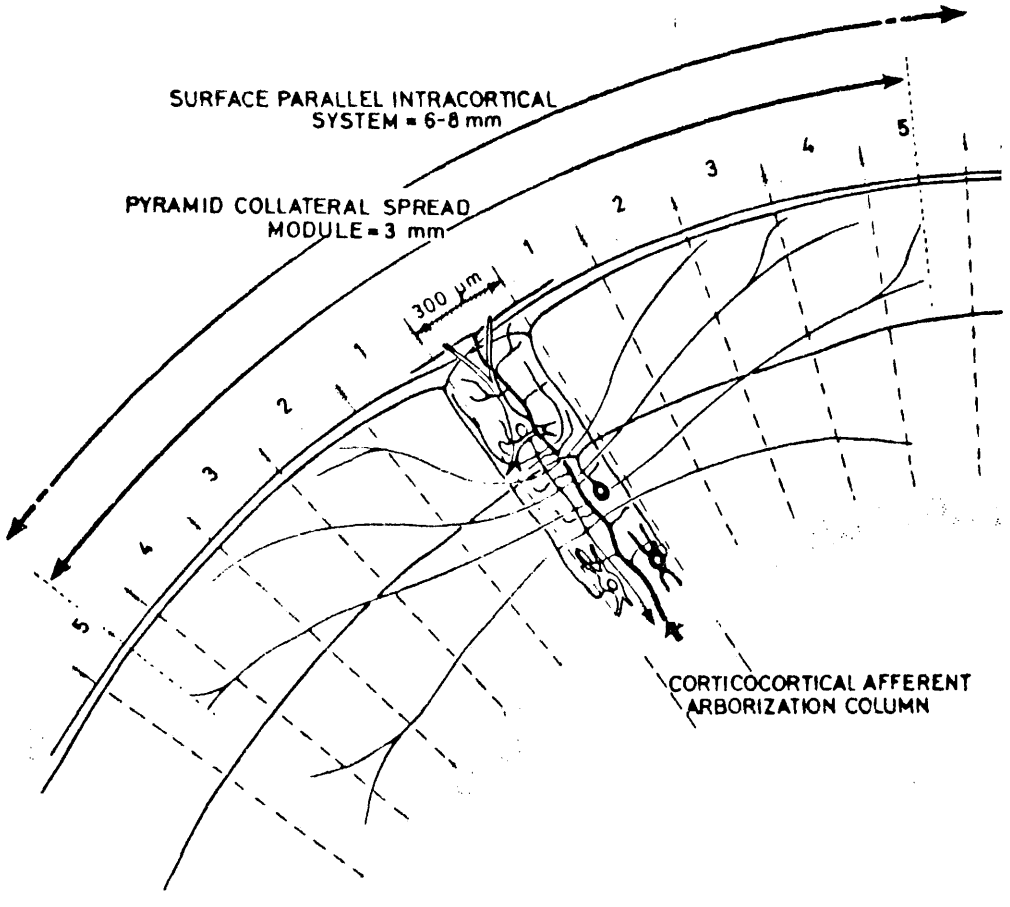

Figure 5.6 (Reproduced from Szentagothai, 1978) Illustrating the diffuse connections described by Szentagothai. Note that all cells shown in his diagram have extensive spread parallel to the surface. Many authors assume that pyramid collaterals are inhibitory.

contamination by inputs by immediately adjacent components varies . . ." This conclusion implying a modular concept would remain valid, even if one modifies the original target of the afferent fibres from pyramids to spiny-dendrite stellate cells.

Although not explicitly stated, both Szentagothai's and Scheibel and Scheibel's modular concept would appear to differ from that of the model presented here, and described in the opening of this sub-section. Like Mountcastle's (1978) interpretation based on parietal cortex, the concept appears to be of entirely self-contained cylinders within which some unknown computational process takes place. Although in general terms a modular concept remains valid after the modification of assumed afferent termination from pyramids to stellate cells, it would seem that the notion of independent cylinders is weakened by that change. As far as is known the processes of stellate and other morphological varieties of cell, other than pyramids, does not conform to any particular columnar boundary. Szentagothai produced a modified neuronal wiring diagram, with afferents terminating upon stellate neurons (Szentagothai, 1975) reproduced here as Figure 5.5. The only aspect of that diagram which really needs to be contradicted in my model is the representation of stellate cells as falling necessarily entirely within a cylinder. My model assumes additionally that stellate cells communicate with each other, and that these connections extend across column boundaries.

Szentagothai (1978) has further qualified his earlier assumptions by describing "more diffuse local connections" which extend beyond the boundaries of columns as defined by afferent terminals. Figure 5.6 from his paper illustrates these connections. "The remarkable feature in the connections illustrated . . . particularly in the pyramid collaterals is that this system is completely (and circularly) symmetric in the surface parallel plane . . . In such a system the conditions for co-operativity and the emergence of dynamic patterns are given also at the macroscopical level of the neuron network . . ."

Szentagothai stresses the pyramids, but his figure shows other neurons with an even greater spread. Moreover, the descriptions of neuron morphology cited earlier, (e.g. Jones, 1975) provide suitable candidates for the task. The fact that my model predicted both the columnar structure and the diffuse connections shows that whether those predictions

turn out to be right or wrong in the long run, the model has strong empirical predictive value.

A feature of Szentagothai's proposal is the function of inhibitory neurons. For this he utilised the inhibitory and excitatory synaptic distinctions argued by Colonnier (1968). He assumes that the large disc like cells are inhibitory and provide the substrate of the boundaries of columns. He assumes that Jones' type 4 cells are inhibitory, which he describes as "chandelier cells". Other candidates for inhibitory cells are "columnar baskets" and "large basket cells" (Szentagothai, 1975). Except that his inhibition assumptions are far more specific than need to be adopted for my model, there would probably be no conflict with his proposals.

Some comment may be necessary to relate the template idea to some complex neuronal evidence: (1) receptive fields are not entirely identical within a cylinder, complex fields being nearer the surface (Hubeland Wiesel, 1972); (2) binocular dominance columns; (3) the size of the columns compared with the size of any particular kind of cell. The only corresponding magnitudes are between degenerating afferent terminals and the functional column size. The nearest in size is the large pyramid shown by Szentagothai to have a spread of over 300 μ; (4) composite nature of cylinders controlling individual muscles; (5) Mountcastle's discovery of mini-columns in the parietal cortex.

The spatial template idea does not require any particular representation of detail in the projection. The finer the detail, the greater the discriminative power of the template system, but clearly the smaller the differences to be detected. As already pointed out, local conjunctions and convergences of afferent information need not be ruled out simply because my model is not an hierarchical network one. The idea of superimposition of templates has already been mentioned. The most obvious arrangement would be for the main template system to be defined roughly at the level of the afferent terminals, and to have a grain size equal to one of Hubel and Wiesel's or Mountcastle's columns. The binocular dominance columns provide at this level the necessary basis for a template matching taking into account what appears on both retinas. Nearer the surface the afferentation may be result of convergences from nearby afferent fibres. Thus the mini-columns may be seen either as separate finer grain template systems or, functionally, as fine tuning

apparatus for the main template system. In the case of Hubel and Wiesel's complex cells, such a neuron may be driven by a conjunction between two neighbouring afferent fibres, the one having a delay relative to the other, and so register temporal sequence.

It would be useful in my model to assess the number of excitatory neurons within a column. This must depend on how one defines the column. Sholl (1956) estimated 5,000 neuron somata within the compass of any single afferent fibre. Cragg (1975) counted by light microscopy on cresyl violet stained sections. His density was considerably lower than Sholl's and was around 15.6×10^6 per cu cm. Taking a column as measuring 0.5 across and 2.5 mm deep, gives about 9,500 neurons per column. This figure includes pyramids and inhibitory neurons. Nevertheless, the order of magnitude is close to that of Sholl. Mountcastle (1978) suggests about 110 neurons in each mini-column, which gives about 25,000 cells per hypercolumn. In considering these numbers, the possibility must be taken into account that the neuron may not be the element of integration. Each integrating neuron might contribute a large number of such elements. This possibility will be considered in greater detail when more speculative aspects of the model are discussed.

Summarising a vast amount of information, there is now strong evidence for the following: afferent terminals enter columns: in striate areas they synapse upon stellate neurons with spiny dendrites in Layer IV: in somatic areas they synapse upon both stellate and modified stellate neurons with spiny dendrites in Layers III and IV. Some afferents proceed in all areas to Layer I. In motor areas the main destination is Layer II. The neurons of the cerebral cortex can be classified into a number of morphological types, with characteristic soma and dendritic and axonal arborisation. All workers agree upon the main distributions. There are pyramids in various layers. Most of these have axons arborising in Layers V and VI which contain mostly large pyramids. Amongst the stellate cells some have ascending, some descending and some horizontal processes. These reach distant cells or terminate locally. Afferent terminals form profuse and extended arborisations capable of reaching many stellate neurons.

II The theoretical model

A. The general layout

In simple terms the anatomical organisation of the mammalian cortex may be regarded as consisting of arrays of replicated columns of neurons. Each column contains a variety of neurons, amongst them inhibitory neurons, efferent neurons and distributing neurons, the function of which will be described in due course. A proportion of the neurons to be found within any column are assumed to be excitatory neurons capable of forming connections. The columns may each be regarded as containing a neuronal pool of these integrating neurons. For the purpose of visualising the system, each column may be considered to contain a few thousand integrating neurons.

Afferent columns and efferent columns are defined functionally, the neurons of an afferent column being possibly also the neurons of an efferent column. Afferent columns are each served by one afferent fibre. Efferent columns are each organised around one efferent fibre. However, for the purpose of following the main idea, the columns may be regarded as segregated into separate spatial regions containing only either afferent or efferent columns respectively. Also for purposes of exposition, the columns may be regarded as discrete; that is, the integrating neurons within a particular column receive excitation only from the one afferent fibre which serves it; the integrating neurons of an efferent column drive one efferent fibre only. These restrictions are not considered to apply in the brain, but simplify the presentation. Initially it is convenient to consider the cerebral cortex as consisting of a sensory area comprising a set of afferent columns, and a motor area comprising a set of efferent columns.

The set of afferent columns may be considered as receiving a topographical mapping of the receptive fields of the sensory system. Similarly, the set of efferent columns maps onto the elements or effector fields of the motor system.

Each of the integrating neurons within a given neuronal pool serves an essentially identical role as able to represent that pool's participation in neural activity. All the integrating neurons reached by a particular afferent fibre "are exposed to approximately the same input", (Scheibel and Scheibel, 1970). For the purposes of exposition, the

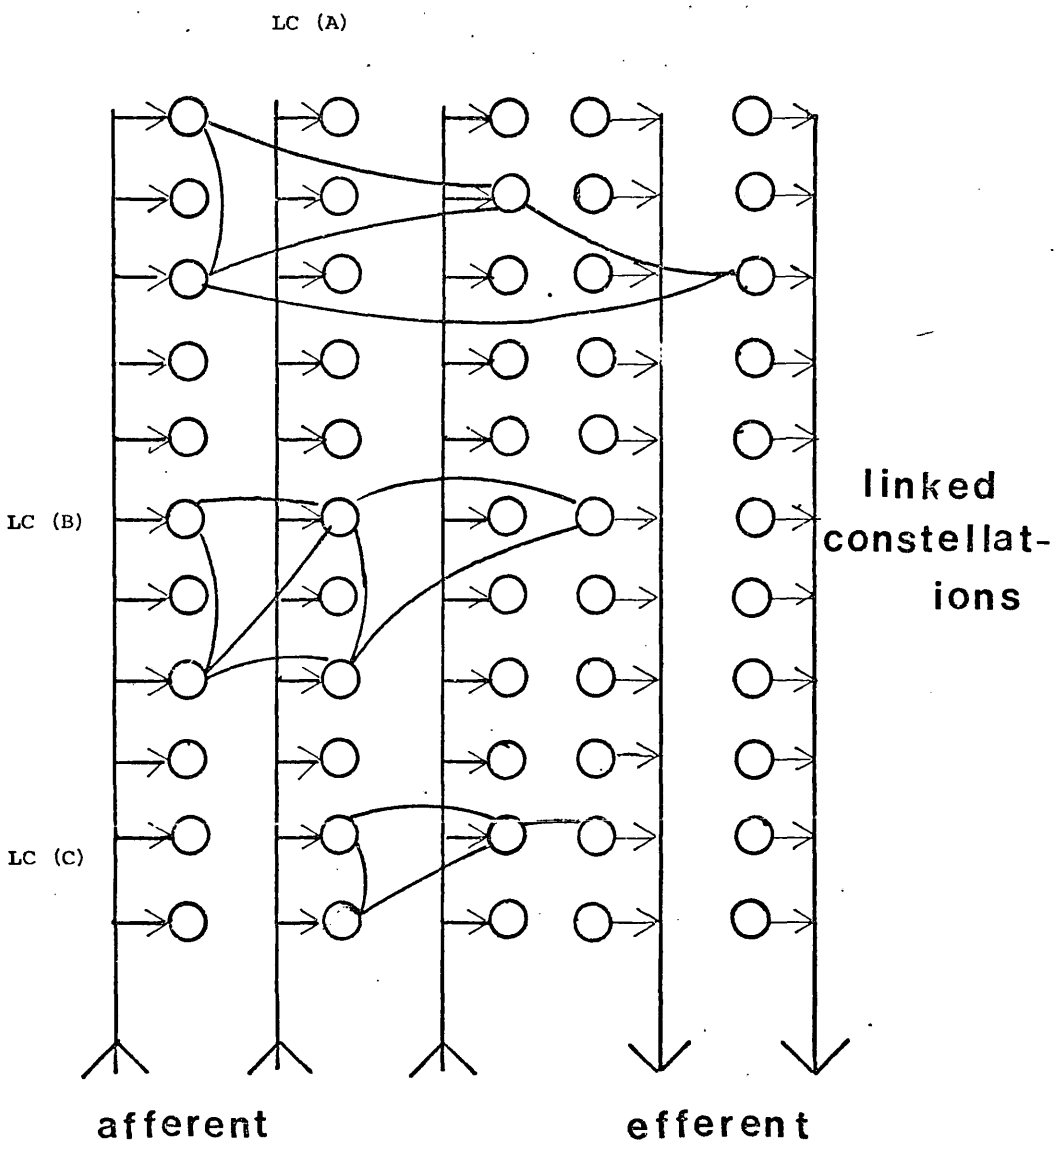

Figure 5.7 The concept of Linked Constellations. Each column of circles represents a neuronal pool of integrating neurons. Each afferent fibre is assumed to activate all the neurons of its pool uniformly. Two columns are shown in which the integrating neurons within a column send their impulses to a neuron with a fibre leading out of the tissue. The heavy curved lines represent selective diffuse connections by which the neurons concerned are able to distribute their excitation amongst members of the sub-set with which they share mutual connections. Different Linked Constellations will reach maximum states of excitation depending upon the pattern of activity in the afferent fibres. The differential excitation is transferred to neurons in the efferent pools.

activity of the afferent fibres may be taken to be all-or-none. At any instant a particular selection of afferent fibres will be active, the remainder passive. Within those neuronal pools served by active afferent fibres all the integrating neurons will be excited simultaneously. The model assumes that all integrating neurons which are excited simultaneously above an appropriate threshold and are within "connective range" of each other (see below) become linked to each other. The linked system formed is exclusive and unique: that is, neurons linked together in this manner cannot form connections elsewhere until they are released from their current association. "Connective" or "Synaptic range" simply means that the neurons concerned each possess a process (e.g. axon) which approaches closely enough to the process of another excited neuron (e.g. a dendrite) for a connection to become effective. The connections form both between integrating neurons within a neuronal pool and also across neuronal pool boundaries between neurons belonging to simultaneously excited pools.

The resulting linked structure will contain neurons from excited neuronal pools (columns) but not from non-excited neuronal pools. It will represent the pattern of excitation across the set of afferent fibres at the time it was formed.

Integrating neurons can only belong to one such linked structure at a time. The model as simulated assumes that linkage connections decay. In this way neurons are released and become available again for connection into a new structure. This assumption may however be replaced by others; the idea is not central.

The joining together of integrating neurons from excited neuronal pools to form linked structures is illustrated in Figure 5.7. Such structures or their residues would remain in the system as learned records, or more informally, memory records. The many structures would be intertwined within the feltwork of the cortex. A feature of the model is that they are both structurally and functionally separate. Subject to certain qualifications to be detailed, they are functionally isolated from each other. They remain in the system unaltered except for the possible decay of the structure due to decay of connections.

The linked structure is not to be considered as restricted to a local region of cortex, or even necessarily to the cortex. Linked structures

would be formed locally within many regions of the cortex: at the same time the local structures would also be joined together to form wide-ranging networks, linked constellations, extending throughout the cerebrum.

B. Integrative and morphological connections

As mentioned in an earlier chapter, it is convenient to refer to two classes of neuronal connection or communication. In the first class are those which are innately or developmentally determined and which define the morphological structure of the brain. For example, columns are defined by the synaptic connections which afferent fibres make upon the integrating neurons, or in the case of an efferent column, by the synapses which the integrating neurons make upon a neuron with an efferent fibre, e.g. a pyramid. A second class of connections are those which mediate linked structures. These are mainly across column boundaries. They are proposed as the substrate of learning and of the genetically determined adaptations analogical to learning (e.g. innately recognised patterns). The first class will be referred to as morphological connections, the second as integrative connections.

Morphological connections are almost certainly Gray type synapses. Current knowledge of neuronal alteration mediating learning (and analogical inbuilt information) is too uncertain to enable the form of integrative connections to be specified. Provisionally the model will be described as though both morphological and integrative connections were mediated by orthodox synapses. A non-synaptic mechanism would not appreciably modify the model as simulated, but might have considerable relevance to its further interpretation.

In terms of orthodox synapses, integrative synapses formed to mediate learning would require a rule assumed by many other neurological models (see Chapter 4); for a connection to become effective both the pre-synaptic and post-synaptic neuron must be simultaneously excited. The synapses may be considered to be structurally existing, but ineffective prior to the fulfilment of the condition. In my model such connections are not progressively facilitated by use. When the linking conditions are fulfilled, i.e. post-synaptic and pre-synaptic cell excited above threshold, the connection becomes fully operative.

C. Formation of linked constellations

Consider a number of neurons within synaptic range of each other, some

only of which are simultaneously excited. One of the excited cells will probably possess a potentially effective synapse upon each of several of the other excited cells. These synapses will become effective, whereas the potentially effective synapses it possesses upon unexcited cells will not. Some of the excited neurons upon which it possessed a potentially effective synapse may possess potentially effective synapses back upon itself. If not one or other of them will probably possess a potentially effective synapse upon another excited cell which possesses a potentially effective synapse upon the first cell, even though that first cell may not possess a potentially effective synapse upon it. Only the connections between the simultaneously excited neurons will have been made effective. Accordingly, after excitation, the subset of excited neurons would have a randomly distributed synaptic system in which every neuron would receive synaptic feet from some other neurons of the subset and possess synaptic feet upon others. By contrast none of the neurons of the subset would possess synaptic feet upon any neuron outside the subset, nor would any neuron outside the subset possess effective synaptic feet upon any neuron within it. The neurons of the subset form a Linked Constellation (LC).

If upon a subsequent occasion some of the neurons belonging to the LC are excited directly whilst others are not, the excitation will spread from those directly excited to those which are not. The effect will be to average the excitation received by the directly excited cells over all the neurons of the Linked Constellation.

In the simulation, the hypothesised exchanges are mediated by notional Spikes. Those neurons which are receiving excitation direct from afferent terminals impart spikes to those in neuronal pools which are not currently receiving direct excitation but which are members of the same LC because they were excited simultaneously on a former occasion. These latter neurons, being at a lower level of excitation, do not return so much excitation to the directly excited neurons; so, as just described, the level of excitation within the linked subset moves towards the mean of the excitation of all the neurons of the LC. In theory these exchanges could be by way of a non-synaptic mechanism: or perhaps, although synaptic, achieved by local micropotential changes within the membranes concerned, without actual spike emmission, as suggested by Pribram (1971).-

D. Operation of Linked Constellations

In brief, each afferent fibre excites only one neuronal pool. At any instant an afferent pattern would excite a corresponding selection of neuronal pools above an assumed threshold for discharge and formation of integrative connections. A Linked Constellation would be formed comprising cells from only the excited pools. Linked neurons can exchange excitation more readily than non-linked neurons.

Upon repetition of the sensory pattern the same selection of neuronal pools would become excited. The matching LC formed upon the former occasion would become more excited than other non-matching Linked Constellations within the system, formed during different sensory stimulation. The matching LC has all its member cells excited (on the sensory side). The non-matching has fewer member cells excited, since some of the neuronal pools in which its members reside are not excited by the current afferent pattern. Consequently the level of excitation rises faster in the matching Linked Constellation than in non-matching ones. The effect is one of degree. Linked Constellations containing a greater number of neurons situated in currently activated pools would rise to higher levels of excitation than those having a smaller number of directly excited neurons.

Effectively the afferent excitation may be said to be shared amongst the member neurons of each Linked Constellation. However, the idea of sharing implies the idea of the conservation of the amount of excitation. This is an interesting idea, which will not however be pursued here. It may perhaps be mentioned in anticipation, that a simulation of the Linked Constellation concept in terms of orthodox synapses showed that the system could be made either attentuating, amplificatory or conservative, depending principally upon the quantitative distribution of the potential connections. The amplificatory version showed certain advantages, the attentuating system, other advantages. In fact these issues proved to be rather different in operation from what was expected of them in advance, and their discussion will therefore be left for later. For the purposes of exposition, the system may be thought of as essentially conservative, with, however, the facility to reduce the excitation in a Linked Constellation after it has served its immediate function of being raised to dominance by the current input pattern.

E. Linked Constellations as extended templates

Each Linked Constellation provides a learned record of the state of

excitation of the set of neuronal pools (columns) occurring at the time it was formed. In a simple sense it is a template. If the sense organ were mapped in a strictly isomorphic fashion onto the array of columns, the Linked Constellation would provide a simple template (Chapter 2 and Neisser, 1967) of the stimulation falling upon the sense organ. For example, when a leter 'A' fell upon such a retina, the Linked Constellation formed in the isomorphically mapped visual array of columns would also represent an 'A'. In reality the mapping of sense organs is not strictly isomorphic. There are complications such as orientated receptive fields. There are some aspects of neural processing where it would not make sense to speak of a topographic mapping. Nevertheless, Linked Constellations are templates in the sense of providing a spatial configuration to be matched against the spatial configuration of the input pattern.

The templates provided by Linked Constellations are not suggested to be templates of the external world, nor even of the sense organ. Each Linked Constellation would be a template of one particular state of activity of the set of neuronal pools. Because of the distributed character of each Linked Constellation, extending between different projection areas, it would provide an extended template as illustrated in Figure 2.4

III Inhibition

A. Main role of inhibition in the model

A class of neurons are assumed to exert an inhibitory effect upon the integrating neurons. These will be termed inhibitory control neurons, or less formally, inhibitory neurons. Inhibition is used here purely in a neuronal sense, and does not refer to behavioural inhibition.

The general purpose of inhibition, according to the model, is to control the excitation of the integrating neurons. To provide a comprehensive control one may envisage the system of inhibitory control neurons as organised at various levels of range of action. There would be inhibitory neurons acting locally. Such an inhibitory neuron acting recurrently upon an individual integrating neuron would ensure that when that integrating neuron discharged it was subsequently inhibited. At the local level some inhibitory neurons would act laterally upon neighbouring

integrating neurons. At a broader spatial level, neuronal pools would contain inhibitory neurons which recurrently inhibit all the integrating neurons within that pool, and laterally all the integrating neurons of neighbouring pools. At a still broader spatial level, there would be a system of inhibitory neurons which controlled the general levels of excitation over a set of pools. It is this latter level with which we are particularly concerned at this point.

For purposes of exposition we are here concerned with the main function of inhibition in the model, which is to keep the average amount of activity in the system below threshold levels. In a dynamic system of distributed excitation, some excitation will be reaching most integration neurons most of the time. It is necessary to distinguish those which are receiving more excitation from those which are receiving less, rather than those receiving some excitation from those receiving none. In relation to the formation of integrative connections, large numbers of integrating neurons receiving little excitation must be stopped from being linked into a Linked Constellation. In relation to spike emission, it is necessary to prevent all the Linked Constellations becoming suprathreshold and discharging at once.

For the purposes under discussion, the inhibitory neurons are activated either directly or indirectly both by afferent fibres and by the integrating neurons themselves. Accordingly inhibition is reactive. The precise method of activation of the inhibitory neurons is not central to the model, but it is important that their activity can be finely adjusted to the spike activity of the integrating neurons.

In neurological models inhibition may be looked at in a number of ways. In models involving logical conjunctions, an inhibitory spike is considered as a discrete signal of opposite sign to an excitatory spike. In distributed memory models, there are massive excitatory effects and inhibition may be considered as a means of control. Magoun (1958) proposed a concept of this kind. The present model takes into account that a neuron may receive many thousands of synapses. Cragg (1975) gives 40,464 as the mean per neuron.

Integrating neurons may be considered to be excited by large numbers of synapses from afferent terminals (See Section I). The inhibitory neurons are assumed to possess abundant synapses (morphological connections) upon the integrating neurons. All the integrating neurons within a pool will

receive equal excitation from the afferent fibre. The level of excitation will however vary from neuron to neuron (see below) according to the integrative connections each possesses. If they are all subject to the same inhibitory influence, many of the less excited integrating neurons will have their spike activity totally inhibited. Only the most highly excited will discharge.

For the present purposes it is not necessary to consider the physiological mechanisms of inhibition within the cells. The treatment is essentially no different from that of other neurological models. In effect inhibition is algebraically subtracted from excitation. The treatment was in fact somewhat more complicated in the main simulations, in an attempt to guard against a too simplistic view of inhibitory neuronal processes. The representation of inhibition in these versions of the simulation was to provide a variable threshold for the discharge of integrating neurons and for the forming of integrative connections between them. This representation was also found to have its limitations. This aspect will be left until later in this chapter.

During development of the simulations a number of different strategies were adopted to try to control the application of the inhibitory system with maximum effectiveness. Eventually a rather simple conceptual approach was adopted. In this approach, the main inhibition system was assumed to apply its influence equally upon all the integrating neurons within an array, i.e. across a set of neuronal pools. This also helped to define an array as a set of columns having common inhibitory control. Lateral and recurrent inhibition, at various spatial ranges, was then considered as serving only the functions commonly attributed to them (see below) and were given subsidiary systems which modulated the main system.

B. Control of inhibitory activity

In the model an important aspect of differences of excitation of integrating neurons is in which columns the most excited neurons are situated. The aim was to adjust the output of the inhibitory system so that there would be some neuronal pools in which no integrating neurons were discharging, whilst in other pools discharges were taking place. This condition would enable the discharge pattern to be equated with a selection of neuronal pools, and accordingly with an afferent pattern. Referring to Figure 5.7, it will be seen that whatever pattern of afferent input

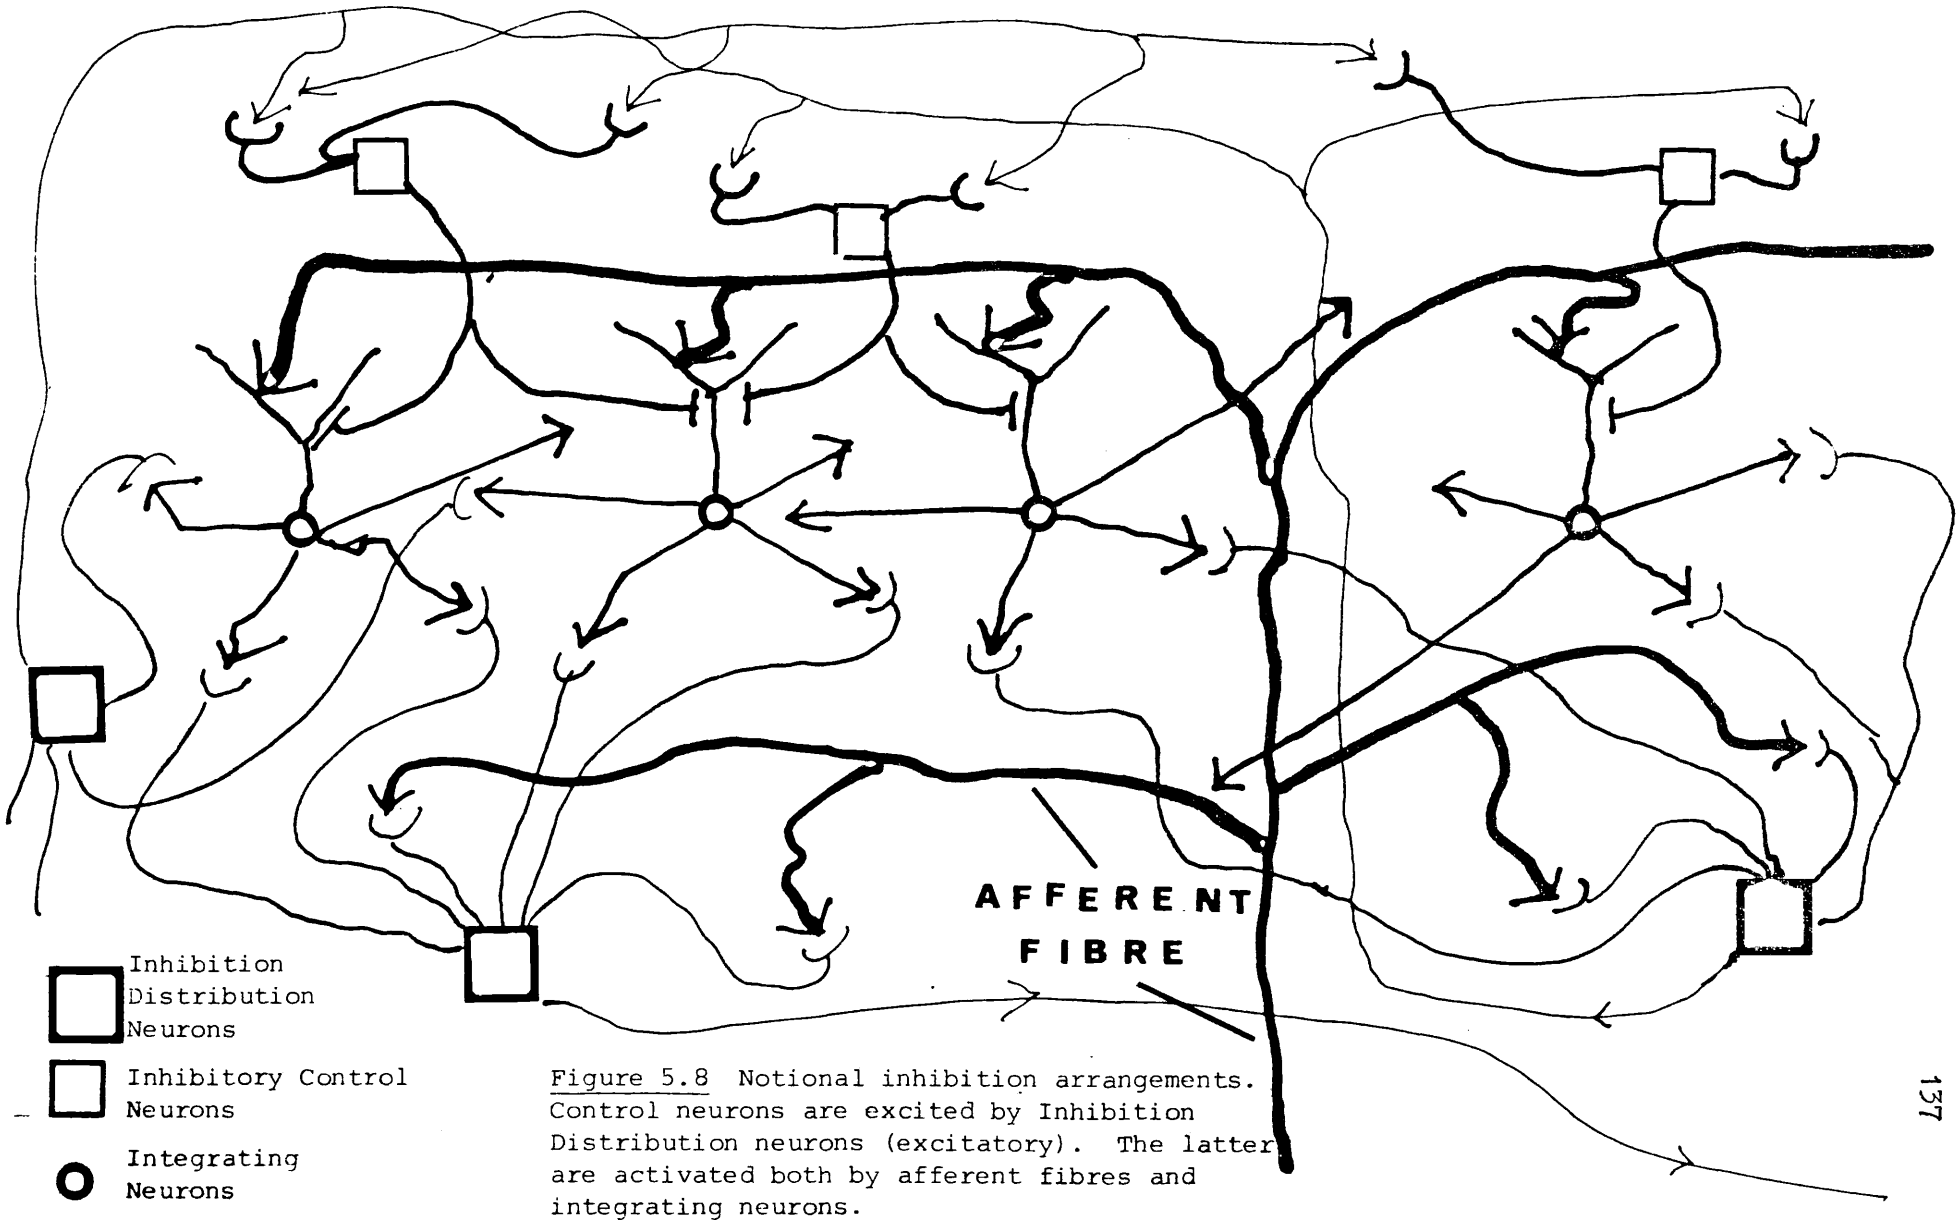

Figure 5.8 Notional inhibition arrangements. Control neurons are excited by Inhibition Distribution neurons (excitatory). The latter are activated both by afferent fibres and integrating neurons.

occurs, many Linked Constellations will be excited to some degree (numerical examples will be given in Section IV). Selecting the most hightly excited Linked Constellations is a matter of suppressing the partial excitation of the less excited ones. This could be acheived by applying inhibition across the whole array at an appropriate level.

The level of inhibitory activity of the main inhibitory system was controlled by feed-back from the excitatory output of the integrating neurons. Unexpected and poorly understood difficulties were encountered when it was attempted to develop the model beyond a simple stage. It was relatively easy to define a feed-back criterion which allowed those integrating neurons to discharge which were in columns whose afferent fibres were currently active. To take full advantage of the theoretical model it was required, however, to discern both the direct effects of afferent input and the effects of the most highly reactivated Linked Constellations, whilst excluding the effects of less highly reactivated LCs. It was impossible to find a criterion for the feed-back system which would do this. In fact, the solution of problems of inhibitory control, and of related excitatory controls, was not achieved until very late in the development of the simulations, when notional electrical rhythms were introduced.

C. Neuronal mediation

To achieve the equalising of inhibitory action across an array, it would seem probable that the integrating neurons and afferent fibres do not activate the inhibitory control neurons directly, but through intermediary neurons. These intermediary neurons would need to possess connections across column boundaries if they are to average inhibition across an array. Such connections could be orthodox excitatory synapses genetically or developmentally determined. The intermediary neurons would be excitatory, but may perhaps be termed inhibition distribution neurons, since they would control an inhibitory effect upon the integrating neurons (Figure 5.8).

The connection of afferent fibres direct to such neurons would provide a measure of "feed-forward inhibition" (cf. Scheibel and Scheibel, 1970). This would enable the integrating neurons to be prepared for the massive arrival of afferent excitation. It would protect the tissue and the system from sudden changes. It would adjust the tissue crudely to its task, whilst the contribution of the discharges of integrating neurons effected the fine adjustments. These neuro-histological speculations are, of

course, intended merely to indicate a possible neuronal implementation of the model concordant with contemporary knowledge.

D. Recurrent and collateral inhibition

In the theoretical model in the form just described, recurrent and lateral inhibition would be superimposed upon the main system of inhibitory control neurons. These would have the effects customarily ascribed to them in the literature (see e.g. Mach, 1959; von Bekesy, 1967; Ratliff and Hartline, 1959; Jung, 1972).

After the introduction of the revised inhibition approach just described, these aspects of the system played little part. They were however, included in many of the simulations, especially the earlier versions. Both lateral and recurrent inhibition was reactive to the spikes emitted by integrating neurons. Lateral inhibition acted immediately whereas recurrent inhibition was subject to delay. Lateral inhibition acting between neuronal pools, enhanced signal contrast. Recurrent inhibition has the effect of cutting short the discharges of neurons in excited neuronal pools and gave rise to a "negative after effect" in those pools which had received active afferent volleys.

These functions, although apparently not central to the main idea of the model, may be useful in rounding out the account of brain function. There is, moreover, one particular aspect in which lateral and recurrent inhibition might prove crucial in the interpretation of the model beyond the stage to which it is developed in this thesis. This relates to the possibility of the system being able to reproduce a sequence of events. Lateral and recurrent inhibition acting between individual integrating neurons would enable the neurons of the most highly reactivated Linked Constellation to rise to dominance: its activity might then be reduced by recurrent inhibition, so making way for the neurons of another Linked Constellation to become dominant.

IV Formal statement of the model

Some simple numerical examples may make the system clearer. The formal statements are only intended as an aid to exposition. They are intended to provide reasonable algorithms to express the theory.

A. Some rules of the model

1. Integrating neurons receive excitation from:
   (a) afferent volleys
   (b) other integrating neurons with which they have integrative connections: i.e. the averaging effect over Linked Constellations.
2. Integrating neurons discharge excitation upon:
   (a) efferent neurons
   (b) other integrating neurons with which they have integrative connections: i.e. the averaging effect over Linked Constellations.
3. Integrating neurons become linked together into Linked Constellations if:
   (a) the excitation in the cell is above the sum of the inhibition and the intrinsic threshold
   and
   (b) the cell is free to be linked (see next chapter).
4. At any given moment only one Linked Constellation involving all supra-threshold cells is formed. The time taken for the formation of a Linked Constellation is finite.
5. An integrating neuron is either linked or unlinked. If linked it has integrative connections with members of one Linked Constellation only.
6. At the end of a randomly determined duration, an integrating neuron ceases to be a member of a Linked Constellation and becomes available to form fresh integrative connections with a new Linked Constellation. This is known as decay (see next chapter).

B. Some further algorithms

(i) (a) Integrating neurons are either linked or unlinked
   (b) only unlinked cells are available for linking into new Linked Constellations.

(ii) Integrating neurons which are unlinked and are simultaneously excited above a threshold level (TH_{np}) become linked into the jth Linked Constellation.

(iii) TH_{np}, being the threshold of the nth integrating neuron in the pth neuronal pool, is partly attributable to Inhibition (I) and partly a natural threshold individual to each neuron, NTH_{np}. So:

$$TH_{np} = I + NTH_{np} \qquad \ldots 1$$

(iv) The jth Linked Constellation is that formed at a particular instant t_j. The instant t_j comprises the finite time interval between t_j and $t_j + 1$.

(v) Integrating neurons whether linked or unlinked emit spikes when excited above threshold. The frequency of the spikes emitted is a function of the level of excitation above threshold, i.e.

$$D = k(E - TH) \qquad \ldots 2$$

where D is the frequency of emitted spikes, E is the level of excitation and k is an operator. k is not necessarily linear, but may be taken for convenience as some simple proportion, say 1/3.

(vi) E, the excitation, is expressed as the number of spikes which would be emitted by an integrating neuron: (1) in the absence of further received excitation; (2) assuming the threshold remained constant; and (3) until such time that an instant occurs during which no spikes are emitted.

(vii) The values of excitation of all of the integrating neurons of a particular Linked Constellation tend to move to equality over the passage of time. If E_{mj} is the excitation of the mth neuron in the jth Linked Constellation, then ignoring additions or reduction of excitation by neurons within the LC or changes in threshold, all E_{mj} move to the same value.

$$E'_{mj} = \frac{1}{M} \sum_{m=1}^{M} E_{mj} \qquad \ldots 3$$

(viii) The application of equation 3 is qualified by the assumption that all neurons, whether linked or unlinked, lose excitation by emission of spikes, and possibly by dissipation over time. As will be noted from (vii) above, the assumption was adopted that only a proportion of excitation may be discharged over a single instant. This gives expression to the fact that a neuron may continue to emit spikes for some time after being stimulated.

Let AVV (the afferent volley value) be the number of spikes impinging upon the neuron N_{np} at some instant, t_i. Let R_{np} be the reduction of excitation occurring in this neuron at t_i. R_{np} will be dependent upon the discharge of N_{np} and upon other assumptions concerning dissipation of excitation. Then:

$$E'_{np} = E_{np} + (AVV - R_{np}) \qquad \ldots 4$$

(ix) For exposition, all the neurons of a neuronal pool may be considered to be receiving excitation from the afferent fibre of that pool or not.

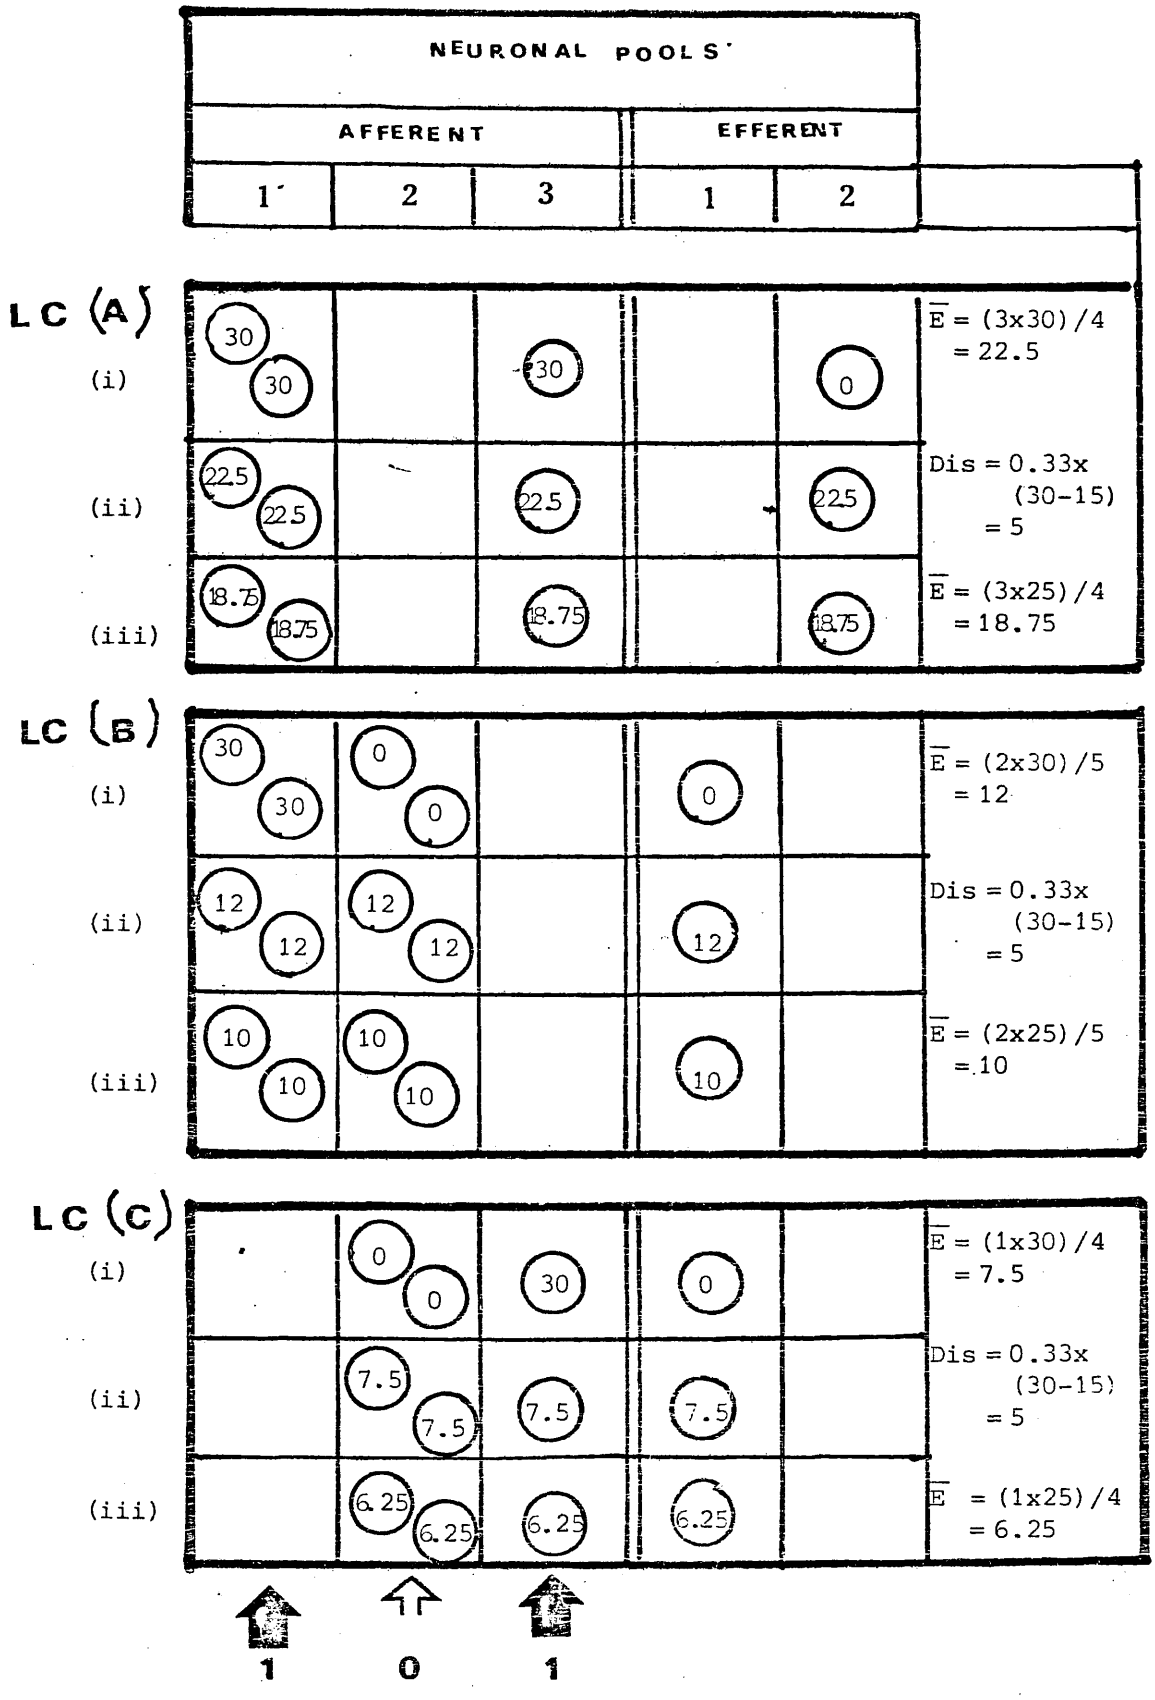

Figure 5.9 (a) Integration across Linked Constellations.
(i) Before integration; (ii) integration ignoring Discharge and Inhibition; (iii) integration assuming Discharge before transfers of Excitation and that Inhibition is maintained at 15. (See text and further calculations in Figure 5.9 (b) next page.
Dis = Spike Discharge. E = Excitation. $\bar{E}$ = mean Excitation after distribution. LC(A), LC(B), LC(C) are the Linked Constellations shown in Figure 5.7.

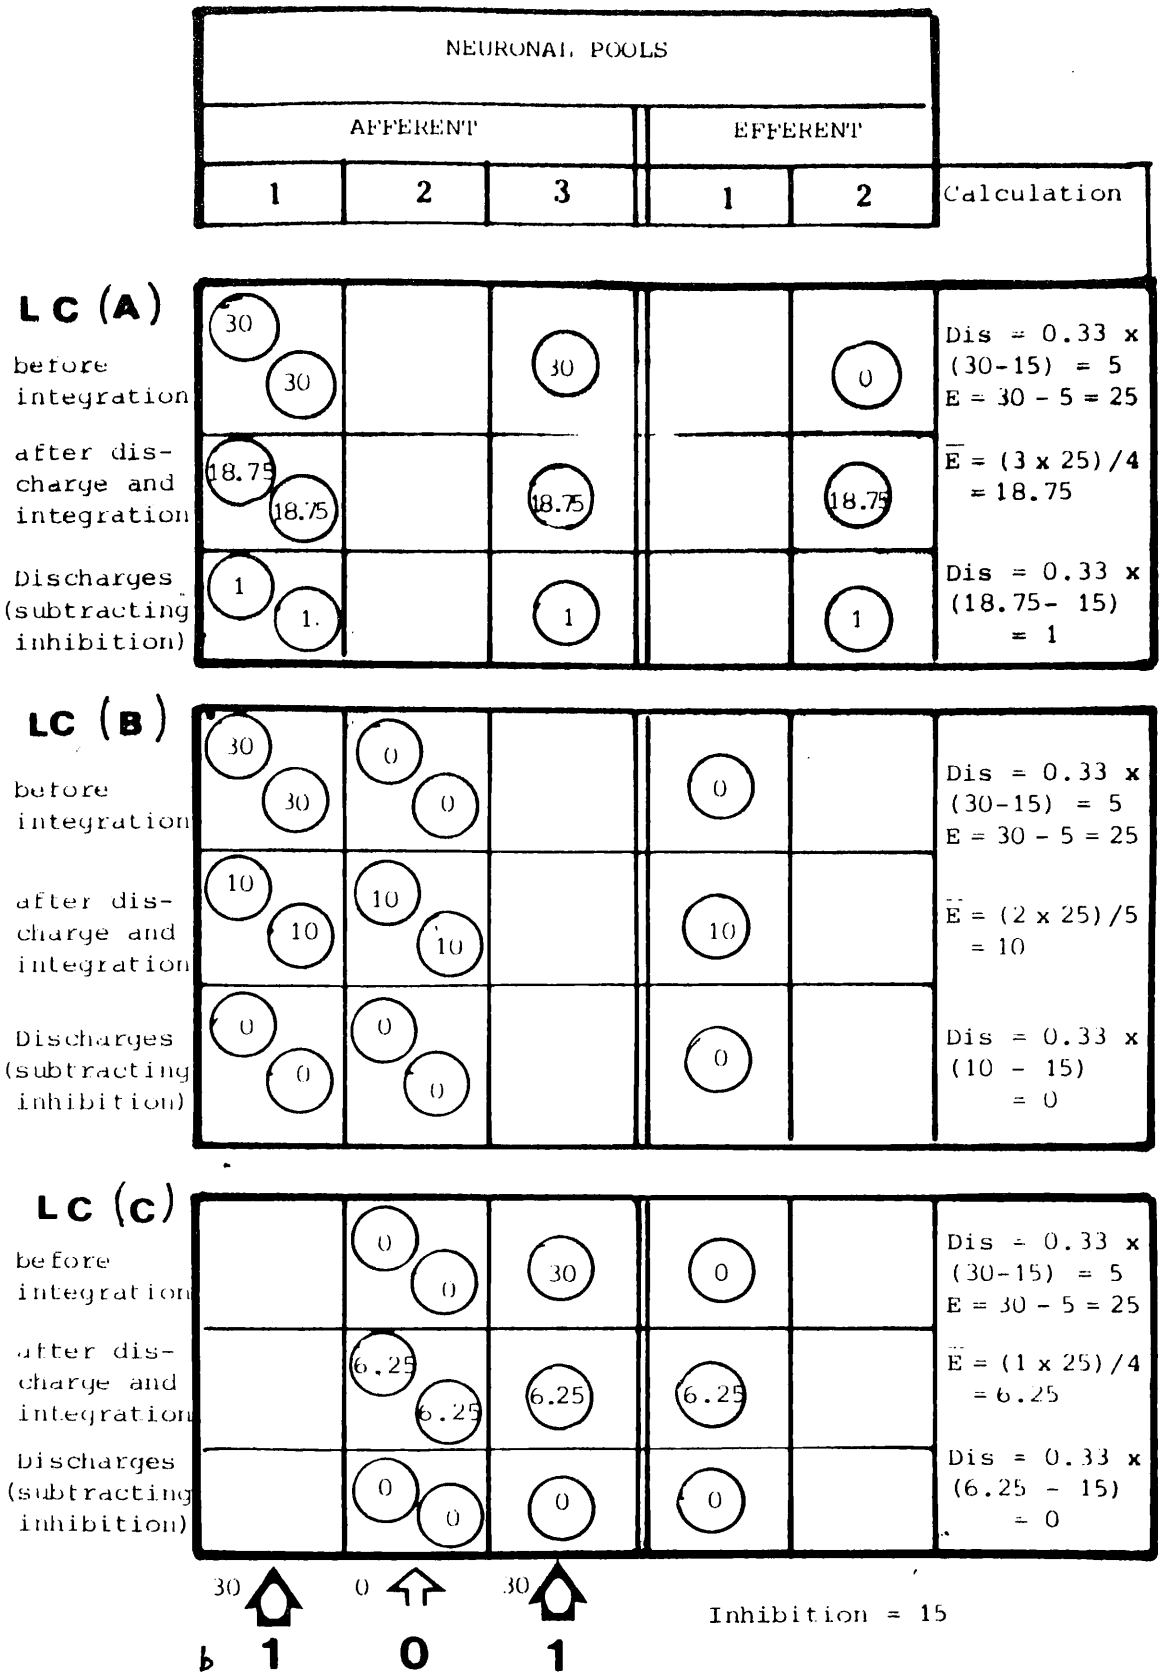

Figure 5.9 The effect of inhibition in assuring that only the most highly excited Linked Constellations produce discharge. The calculations naturally only indicate crudely by steps what would be a continuous analog process. The level of inhibition is assumed to be 15 and constant throughout. Dis = Spike Discharge; E = excitation; $\bar{E}$ = distributed excitation. LC(A), LC(B), LC(C) are the Linked Constellations shown in Figure 5.7.

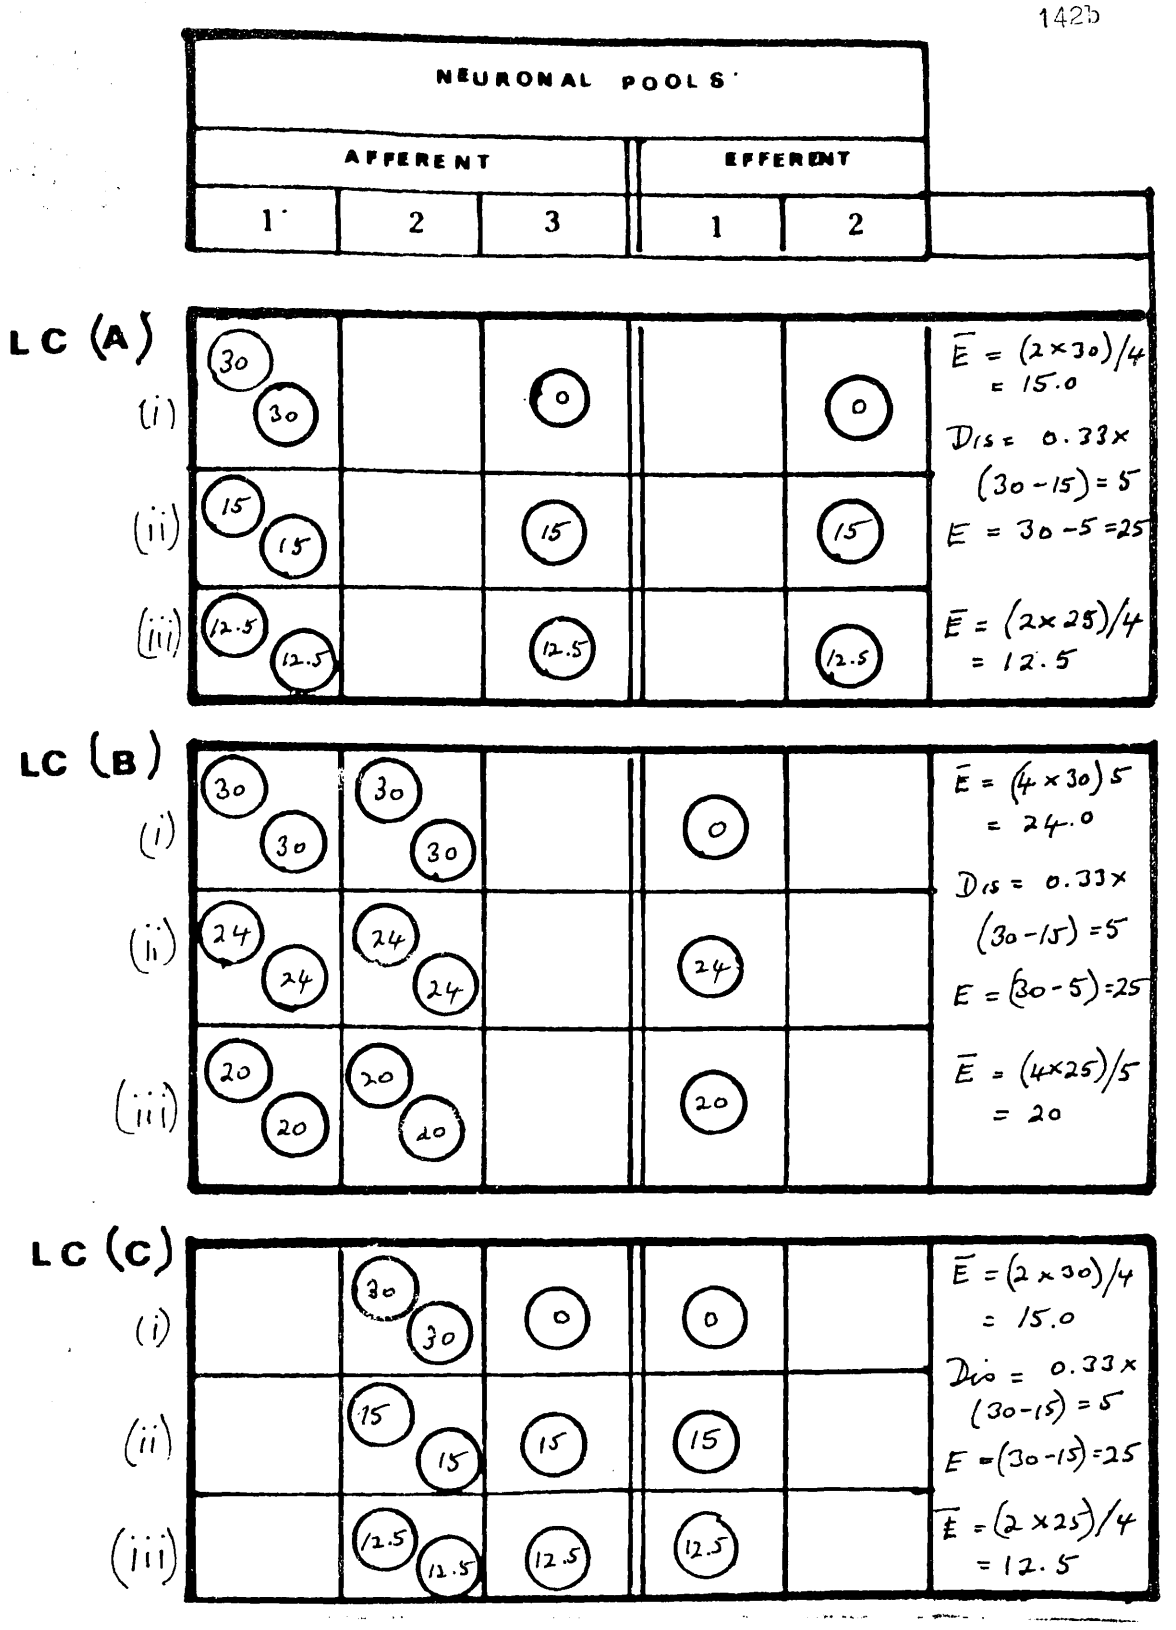

Figure 5.9c The above calculations are given for comparison with those shown in Figure 5.9a. The figures are now based upon the assumption that the first and second Afferent Fibres are Active instead of the first and third, as shown in Figure 5,9a; i.e. the Input is 110 instead of 101. The Linked Constellation which matches the revised Input becomes that with the greatest mean Excitation.

C. Some simple examples

Consider the simple arrangement shown in Figures 5.7 and 5.9a. Suppose, for illustration, that the first and third fibres are active at a frequency of 30 (spikes per unit time). Then ignoring all the above qualifications for the moment, each integrating neuron within an active column will have received 30 spikes by the end of the first time unit. Take each of the Linked Constellations in turn. LC (A) has two cells in the first column and one cell in the third column, all of which will receive 30 spikes, since the first and third columns were assumed to be active at the instant under consideration. The constellations LC (A) will therefore have received 90 spikes in all. The fourth neuron belonging to this Constellation is in efferent column 2, and we shall assume that this received no excitation from an afferent fibre. Then, assuming that the received excitation has been averaged over the neurons of this Linked Constellation, the excitation of each neuron will be 90/4 = 22.5. LC (B) also has two neurons in the first column, receiving 60 spikes between them. It has two in the second column. These receive no excitation from the afferent fibre of that column, since that fibre is assumed to be inactive. It has a neuron in efferent column 1. Accordingly the total excitation received by LC (B) is 60 and the mean excitation across this LC is 60/5 = 12. LC (C) has no neurons in the first column, two in the second which are unexcited, one in the third, receiving 30 spikes, and a neuron in efferent column 1. Its mean excitation is 30/4 = 7.5. LC (A), which may be said to match the input pattern, has a mean excitation of 22.5 as compared with 12 and 7.5.

If the input pattern is changed so that the first and second columns are now active (Figure 5.10 (b)), it will be readily seen that figures become: LC (A), $\bar{E}$ = 60/4 = 15.0; LC (B), $\bar{E}$ = 120/5 = 24.0; LC (C), $\bar{E}$ = 60/4 = 15.0; Linked Constellation B is maximally excited, as it is the matching LC.

This is the essential aspect of the model. However, as will be seen below, inhibition also has an essential part to play. Discharge may also be taken into account quite simply.

Assume that the inhibition level of all neurons is 15 at the end of the first instant. Spike emitted by a neuron in an active column will be $(30 - TH_{np})$, according to Expression 2. Ignoring the individual threshold NTH_{np}, $TH_{np} = I = 15$, for all neurons, so D = 5. For convenience it may be assumed that all reduction of excitation is attributable to spike discharge, i.e. R = D. Then at the end of the first instant, the excitation of a neuron in an active column has been incremented by 25, allowing for discharge.

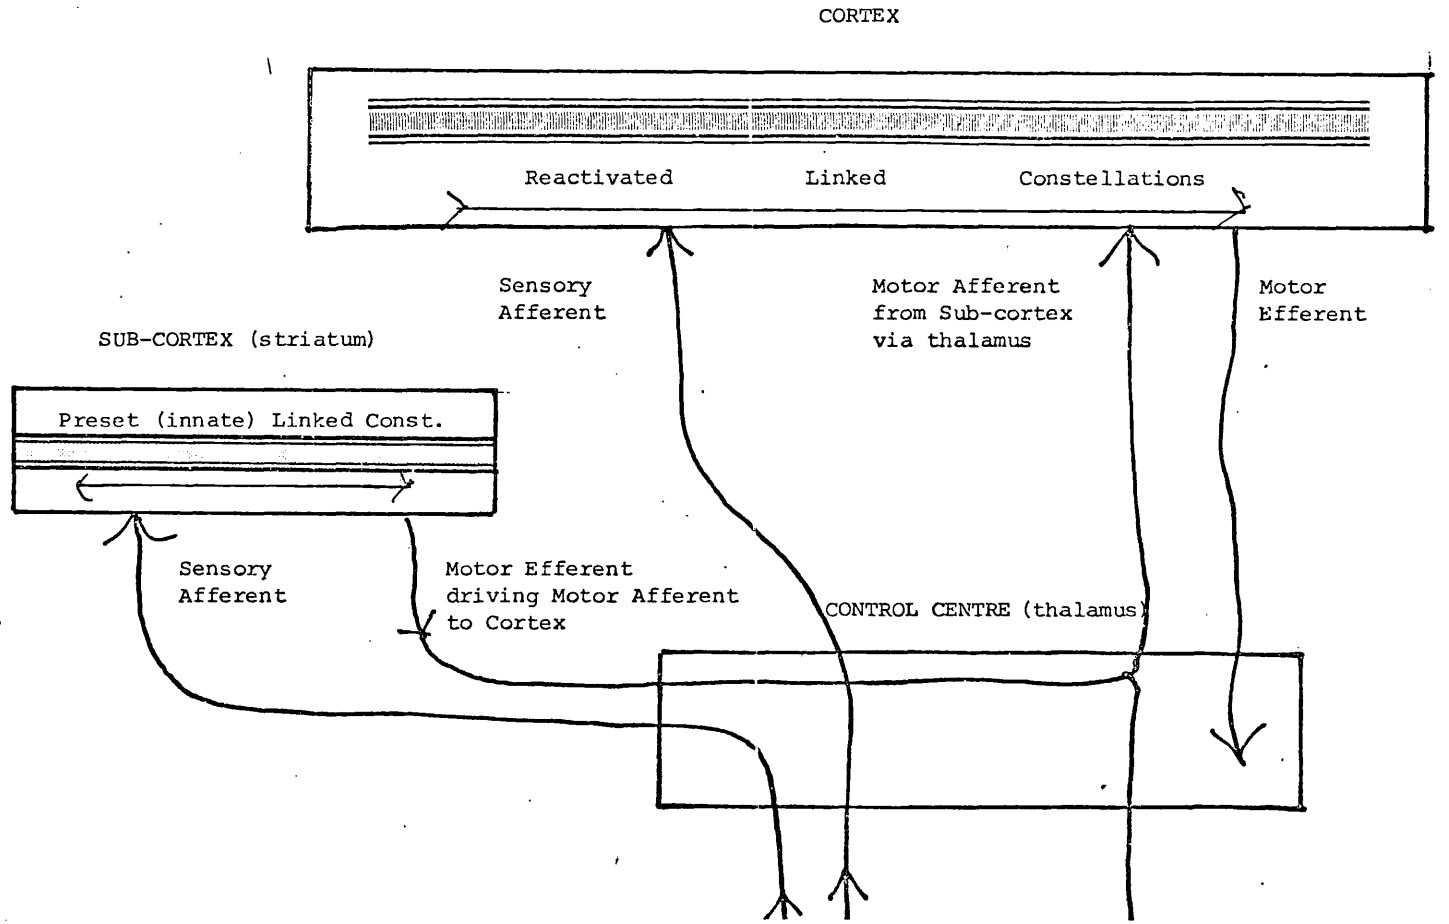

Figure 5.10 Relationships proposed in the model between a sub-cortical centre, termed for reference purposes striatum, and the Cortex. Preset (innately programmed) Linked Constellations in the Sub-cortex are activated by afferent excitation and produce a motor output when specific patterns occur across the afferent fibres. This motor output proceeds to the periphery but also causes a volley in a motor afferent fibre to the Cortex. Consequently a Linked Constellation is laid down in the Cortex which registers the conjunction of the emitted Sub-cortical motor response and the particular pattern of sensory input which occurs at that instant in the Cortex.

After distribution across Linked Constellations the resulting excitations for different LCs will be lower than those calculated above, but the relationships are preserved. For example, with input pattern "101", i.e. first and third columns active, the revised figures are:
LC (A), $\overline{E}$ = 75/4 = 18.75; LC (B), $\overline{E}$ = 50/5 = 10.0; LC (C), $\overline{E}$ = 25/4 = 6.25.

D. Physiological treatment of inhibition

It will be noted that in the last example, although inhibition was used to calculate discharge, it was not subtracted from excitation. The reason is that, in most of the simulations to be reported, inhibition and excitation are treated as though they are independent neuronal processes except for their effect upon spike initiation. They act algebraically to determine spike discharge in accordance with orthodox neuro-physiology. This may be contrasted with the treatment in most formal neurological models, in which inhibition immediately cancels out excitation, leaving excitation at a lower level.

In order to compare the two treatments, the simulations were also run in versions in which inhibition was immediately subtracted from excitation. In many ways this was easier to handle. The main postulates of the Linked Constellation concept was equally demonstrated in both forms. The effects upon the basic assumptions may be readily compared by recalculating the above illustration, subtracting inhibition (see Figure 5.9):
LC (A), $\overline{E}$ = (3 x (30 - 15 - 5))/4 = 7.5; LC (B), $\overline{E}$ = (2 x (30 - 15 - 5))/5 = 4.0; LC (C), $\overline{E}$ = (1 x (30 - 15 - 5))/4 = 2.5. It will be seen that the matching Linked Constellation still bears the same relationship to the others.

It was later appreciated that either form of simulation of inhibition is inadequate to represent inhibition more than in a very rough manner. Inhibition and excitation are useful terms at the overall level of description of a neural system. They are however only labels for extremely complex processes taking place at neuronal level.

Inhibitory transmitter does not, as far as is known, act directly upon excitatory transmitter. They have separate loci on the membrane (see for example, Bradley, 1968; Iverson, 1979; Stevens, 1979). Accordingly it is not necessarily the case that a given quantity of inhibitory transmitter reaching a neuron destroys a given value of excitation. Aside from what may happen within neurons, there may well be an accumulation of transmitters at synaptic terminals which would effectively make excitatory and

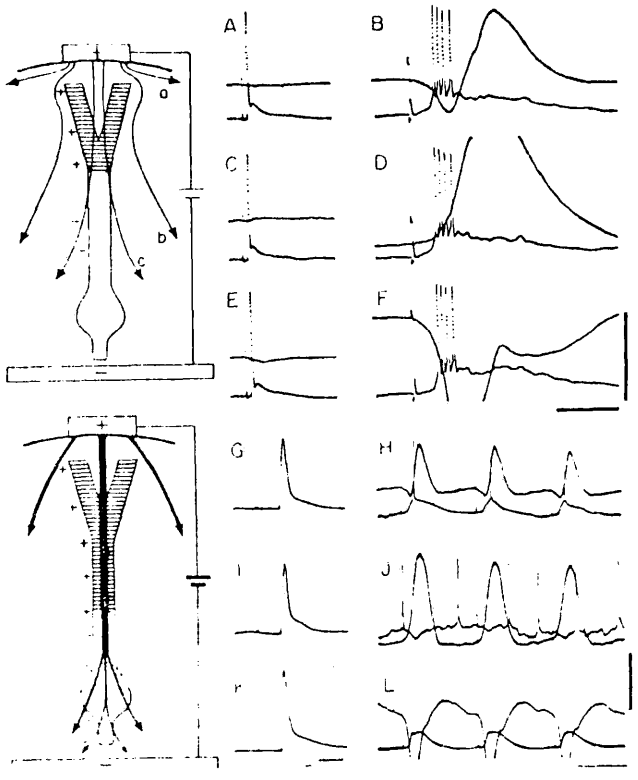

Figure 5.11 (Reproduced from Purpura, 1967) Purpura and McMurtry's diagrams showing probable distribution of currents and intracellular recordings of activities of a pyramidal tract neuron in the motor cortex during stimulation of ventrolateral thalamus and surface cortical polarisation. Upper channel records are surface responses. Weak anodal surface polarisation (C and D) and weak surface cathodal polarisation leave patterns of intracellularly recorded activity during dramatic changes in surface evoked responses. a, fraction of current flowing along surface; b, extraneuronal current flow; c, proportion of current inwards at terminals of apical dendrites, outwards across proximal dendrites.
Strong cortical polarisation (G-L) has different effects. Strong surface anodal polarisation (I and J) produces depolarisation of soma regions and increase in cell discharge. Cathodal polarisation (K and L) hyperpolarises soma and prevents EPSPs from attaining firing level. Cited to illustrate that neurons may be both inhibited and excited at the same time.

inhibitory action independent in the short term, i.e. within one instant of the simulation.

What is not yet clear is whether a neuron can remain excited although inhibited, so that after release from inhibition it emits its excitation as "post-inhibitory rebound". (Compare Andersen, 1974). Although largely independent of the operation of the model as simulated, the answer to this question would be relevant to the further interpretation of my model.

Purpura (1967) summarised work indicating that the dendrites of a neuron may well be in an excitatory condition whilst the axon hillock is inhibited. Moreover, the extent of the inhibitory influence affects the behaviour of the neuron (See Figure 5.11). Andersen and Lomo (1970) support the notion of excitation and inhibition concurrently at different location of the same neuron. They argue, however, for a dendritic spike.

The presentation of my model in this thesis is as a principle of neural organisation. The consideration of neuronal processes of excitation and inhibition accordingly lies outside the scope. However, at the time the main simulations were designed it was thought to be better to reserve one's options by treating the two aspects of cell function as independent except in relation to the action of spike discharge. Practically, the difference between the two modes of representation was probably not significant. Inhibition was still reduced at each iteration of the simulation, as was excitation. In relation to very short intervals, the state of a neuron is perhaps better expressed by the values of accumulated inhibitory and excitatory substances than by the opposing effects of inhibitory and excitatory postsynaptic potentials which occur within the membrane.

In summary, in some respects inhibitory influences on a neuron may be seen as cancelling out excitatory influences. In other respects inhibition may be seen as imposing a threshold upon the neuron which then requires to receive more excitation before it fires. In the model, the important role expressed by either formulation is that inhibitory influences control the capacity of neurons to emit spikes and to form integrative connections.

V Motor aspects

A. General principle

It will be seen from Figure 5.7 that the motor aspects of the model are essentially the reverse of the sensory aspects. Linked neurons belonging to the more highly reactivated Linked Constellations will be at higher levels of excitation than those belonging to less highly reactivated ones. The corresponding efferent columns within which they are situated will be more highly driven and the efferent fibres serving those columns will carry spikes at higher frequency.

B. Significance of inhibition

Consider the efferent columns in the examples given earlier. The figures can be followed from Figure 5.9. First, assuming no inhibition across the efferent columns and that the input pattern is 101 (first and third afferent fibres active), the integrating neuron in efferent column 1 belonging to LC (B) will have an excitation of 10.0 and that of LC (C) 6.25. The only linked integrating neuron in efferent column 2 is that belonging to LC (A) which has a notional excitation of 18.75. With a Discharge Rate of 0.33 (k in Expression 2, Section IV (vi)), the spikes emitted by the two linked neurons in efferent column 1 are 3 and 2 respectively, whilst 6 spikes are emitted by the linked neuron in efferent column 2. Efferent column 1 will receive 5 spikes whilst efferent column 2 will receive 6. It will therefore still be possible to distinguish the efferent output indicating which Linked Constellation was most highly reactivated. Suppose, however, that there were other Linked Constellations with linked neurons in efferent column 1, but no more with linked neurons in efferent column 2. The discharge driving efferent fibre 1 might then exceed those driving efferent fibre 2. This might occur even though none of the linked neurons in efferent column 1 had excitations as great as that of the single linked neuron in efferent column 2. It is a matter of a large number of integrating neurons emitting a small number of spikes in unit time overpowering a smaller number of integrating neurons emitting a greater number of spikes.

If the level of inhibitory action is set at the level of excitation of the nearest challenger to the most highly excited neurons, the competition is abolished. Setting inhibition at 15.0, the discharge of the integrating neuron in column 2, belonging to the matching LC, becomes (18.75 - 15.0) x 0.33 = 1 spikes. The integrating neurons in efferent

column 1 do not discharge at all, since their excitations of 10.0 and 6.25 do not exceed the inhibition level. At this level of inhibition it does not matter how numerous such less-excited neurons may be.

The principle can be applied to sensory arrays. However, as already mentioned, the application to sensory arrays discloses difficulties which at first were poorly understood. If one ignores the effects of continuous afferent input and assumes that the integrative connections have had their effect in equalising excitation within Linked Constellations, there is little difficulty in setting inhibition so that only the neurons belonging to the most highly reactivated Constellations may discharge. Whilst the afferent fibres are constantly active problems arise. In the presence of active afferent fibres, it is possible to set inhibition so as to exclude discharge from neurons within columns having silent afferent fibres. The pattern of columns which contain discharging neurons then directly reflects the afferent pattern. A problem arises when one wishes to consider the pattern of discharge resulting from reactivation of Linked Constellations rather than the afferent pattern. This made it impossible to distinguish the effects of secondary transfers of excitation through Linked Constellations from the direct effects of afferent input.

As it turned out the same problem affects efferent columns, if they are more realistically represented. Although no afferent fibres are shown to the efferent columns of Figure 5.7 this was an omission of convenience to make the main principle of the model clear. As will appear, it is necessary that efferent columns should possess afferent fibres.

These problems were eventually solved in the model by interrupting inputs periodically. During the gaps in afferent transmission, the effects of the integrative connections brought the system to the conditions indicated in the analyses of this and preceding sections. The inhibition levels could accordingly be adjusted at different times to detect either the afferent pattern or the discharges of the most highly reactivated Linked Constellations. This matter will be discussed at greater length after the presentation of the model.

C. Linking of integrating neurons in efferent columns: Afferent motor fibres

One may ask the question as to how the linked neurons shown in the efferent columns of Figure 5.7 came to be there. The linked neurons in the afferent columns are explained by the hypothesis of connectivity

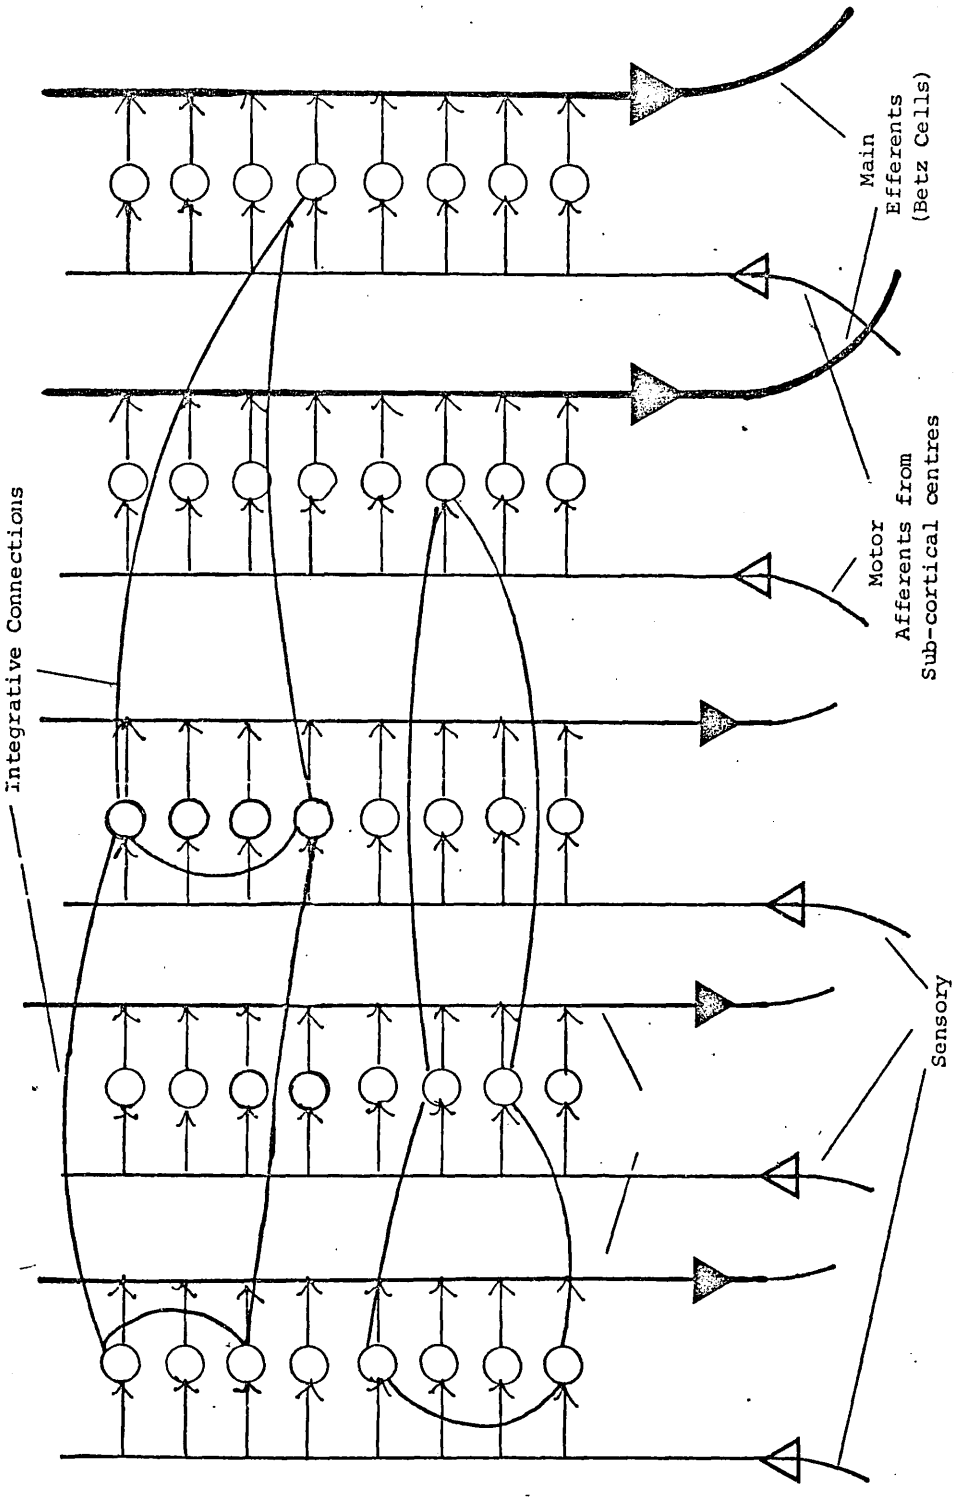

Figure 5.12 An elaboration of the arrangements shown in Figure 5.7. All Neuronal Pools (columns) possess both Afferent and Efferent Fibres. The afferent fibres to Motor Neuronal Pools, termed Motor Afferents, are dependent for their Activity on Sub-cortical Arrays.

when afferent fibres are active simultaneously (Section IV (ii)). The explanation of the linked neurons in the efferent columns is that there are assumed to be afferent fibres to motor columns. These may be regarded as afferents arising from sub-cortical motor centres, the operation of which will be described in due course. Such afferent fibres to motor cortex do occur as, for example, in the cerebellar-thalamo-cortical pathway through the ventro-lateral nucleus of the thalamus.

The presence of afferent fibres in both sensory and motor areas accounts for the integrative connections shown in Figure 5.8. The Linked Constellations shown may be regarded as having formed as the consequence of activity in afferent fibres from (a) sensory sub-cortical centres and (b) motor subcortical centres, arriving at the cortex simultaneously.

In a similar way afferent columns may be considered as possessing efferent fibres. These are able to transport the pattern of excitation across an array of sensory columns to some other centre where they may be required.

The most likely candidates for such efferent fibres would seem to be the pyramids which possess axons leading outside the cortex. These might be compared to the more prominent Betz neurons, which may correspondingly be regarded as the principal efferent output of a motor column. Within the theoretical model, such neurons possessing efferent fibres may be referred to as efferent neurons. Speculating further on neuronal circuitry, it seems probable that the secondary pyramids within a column act as "collectors" of efferent output, which they convey to the Layer V pyramids (Jones, 1975; Valverde, 1971; Szentagothai, 1975). Figure 5.12 illustrates the modifications discussed applied to the arrangements shown in Figure 5.8.

VI Innate Linked Constellations

Linked Constellations have been described as occurring as a consequence of neurons forming integrative connections when simultaneously excited. Such integrative connective systems are assumed also to occur by genetic determination. Recent research on the developmental specificity of the nervous system discloses that the morphology is genetically

defined in great detail (see e.g. Gaze, 1970; Barondes, 1976). Griffith's (1967) model, described in Chapter 4, illustrates how complex patterns of activity of the kind required might result from relatively simple labelling assumptions.

In the model, Linked Constellations present at "birth" are termed Preset Linked Constellations, Developmental or maturational Linked Constellations are those that form by genetic influence after "birth".

Although some plasticity is certainly attributable to sub-cortical centres, it is small in relation to cortical plasticity (Chow, 1970; Buchwald and Brown, 1973). On the other hand, innate patterns of activity remain substantially unaltered when cortex is removed (e.g. Shik and Orlovsky, 1976; Buchwald and Brown, 1973; Flynn, 1976), although their control may be fundamentally modified, e.g. sham rage. For simplicity it is assumed in the model that the cortex is the home of learning, whilst sub-cortical centres,especially the basal ganglia, contain Preset Linked Constellations.

Pringle (1951) proposed that innately determined "key trigger patterns" would occur in his oscillatory model of the cortex. These innate sensori-motor connections would pass excitation to motor neurons only when the key trigger pattern appeared in the environment. He proposed this as a bridge between the proposals of the ethologists, Lorenz and Tinbergen, and those of conditioning theory. A key trigger pattern could be simply an Unconditioned Stimulus with its Unconditioned Response, or it could be an Innate Releasing Mechanism with its corresponding Innate Behavioral response. The Preset Linked Constellations may be seen as fulfilling similar roles in the present model.

Although there are considerable similarities between Pringle's formulation and the present one, there are also significant differences. Pringle assumed, in company with most theorists, that integrating neurons could form further connections upon subsequent occasions. The innate wiring would provide cores around which learned aggregates accumulate. Upon subsequent occasions, in which similar but not identical input occurs, further neurons are added to the aggregate. The neurons of the associations (oscillatory couplings) formed earlier become further coupled to the neurons added on the subsequent occasions. This goes together with the assumption that connections are progressively made more effective with use. This is in contrast to the assumption of my model that connections are formed upon one, that is, the first, occasion. An essential

principle of my model is that Linked Constellations once formed, remain unaltered. They do not gather further similar connective structures into their orbit.

This aspect of my model may be both a strength and a weakness. The capacity of a neuron to make further connections on subsequent occasions would be much more economical of neurons. In my model, one does get strengthening effects on repeated presentation. They are however differently mediated from the Pringle model. The strengthening effects in my model arise because, although each Linked Constellation remains separate, the number of them representing similar input patterns increases. The reactivation of these LCs contributes activity to the current pattern of excitation across neuronal pools. Accordingly the Linked Constellation formed on each occasion is partly determined by the response of LCs already in the system. Consequently in my model, the Linked Constellation formed on any occasion represents the input as enriched by stored information. Of course these differences may be of little significance. The Pringle type model is also an enrichment model. It does, however, seem difficult to discern exactly how the aggregates implied in such a model would operate in practice.

There would appear to be no reason in principle in my model why Preset Constellations should not occur at both sub-cortical and cortical levels: nor why integrating neurons available to register events should not be provided at all levels. In fact the first simple simulations to test the theory were carried out in a system in which a single set of neuronal pools contained both integrating neurons initially linked into Preset Constellations, and others available for integrative connections. The system was effective but at a very low level of efficiency. This efficiency was greatly improved by allocating the Preset Constellations to their own set of Neuronal Pools, whilst another set was provided with initially unlinked Integrating Neurons. This led to the general design of the models simulated. In those models a notional sub-cortex and cortex are each represented by a set of neuronal pools. The basic organisation is the same in both, except that the integrating neurons of the sub-cortical arrays are assumed to have been already linked into Constellations ab initio. These Preset Constellations represent Innately Recognised Patterns on the sensory side and Innately Determined Responses on the motor side. (This description is for purposes of exposition only).

The division into sub-cortex and cortex is to be regarded as one of convenience, but one nevertheless representing a real division of emphasis in nervous systems. The capacity for learning is greatly reduced in decorticate animals and may be of a limited kind. Nevertheless, the terms sub-cortical and cortical describe the design of the simulations rather than any necessary implication of the theoretical model.

This question arises mainly from the difficulty, at this stage of development of the model, of allocating parts of the proposed apparatus to any specific anatomical location in the cerebrum. For example, for concreteness of exposition, the neural apparatus supposed to contain Preset Constellations may be loosely identified with the corpus striatum. It is known that the basal ganglia are heavily implicated in the organisation of innate behavioural systems, and that they have a sensory input (Barr, 1974; Vanderwolf, Bland and Wishaw, 1973). The question then arises as to differences of function of paleo-striatum and neo-striatum. It might be that the paleo-striatum acts in conjunction with archicortex to organise automatic behaviours, whilst instrumental behaviours are mediated by neo-striatum in conjunction with neo-cortex. These matters are regarded as beyond the scope of the present project. They are only mentioned to illustrate the general intention underlying the design of the simulations.

Figure 5.10 represents the basic relationships proposed. Sensory fibres activate neuronal pools at both sub-cortical and cortical levels. The sensory information need not be identical. If the sensory pattern at the Lower (sub-cortical) level matches a Preset Constellation, an efferent response appears on the motor side at that level. The efferent fibre acts as an afferent fibre to the Upper (cortical) level. In the simulation the efferent fibre passes to a general co-ordinating centre, which passes it on to the Upper (cortical) array as an afferent. The thalamus may perhaps be visualised as such a centre. At the Upper Level the sensory imput lays down a Linked Constellation record on the sensory side, whilst on the motor side the Lower Level-coordinating centre-Upper Level volley lays down a record of the released motor response, forming a Sensori-Motor Linked Constellation. Consequently a learned Linked Constellation is laid down at the Upper (cortical) level. This records the temporal conjunction of a certain sensory event with the release of a certain innate response at the Lower level. In one sense the learned sensori-motor Linked Constellation is a copy of the innate Linked

Constellation. In another sense it may differ considerably. The sensory information recorded in the cortex is likely to be far more precise and detailed than the sensory information required to excite the Preset Constellation. The information recorded in the motor cortex concerning the action taken is also likely to be far more detailed than its Lower Level equivalent, and will include sensory feed-back (Asanuma, 1975).

The information represented by afferent fibres entering the Lower Level will be required only to act upon innate structures mediating species specific behaviour, and behaviour which can be characterised as reflex and instinctive. Lorenz has suggested that such arrangements need only represent the skeleton of the actual adaptive behaviour. In the model, the Preset Constellations would mediate these skeletons, and the learned copies of behaviour in the cortex would expand its detail from actual events.

VII Temporal sequential effects

There is a qualification of the rule that a Linked Constellation is independent and functionally isolated from other integrating neurons of the population. This qualification concerns the temporal sequential overlapping of Constellations. The overlapping would appear to follow from the continuous nature of neuronal processes. Physiological processes are rarely discrete synchronised events (but see the effects of rhythmic interruption of afferent inputs, Chapter 15). It later appeared that the overlapping is not central to the model, as the same effects can be obtained by other means.

In the model integrating neurons form integrative connections when they are simultaneously activated; thereafter they are assumed to make no further connections until released from their former connections. Some admittedly ad hoc possibilities of how these rules might come about serve to illustrate the temporal implications. One mechanism might be that the making of the first connection might release an inhibitor of further connectivity, which however takes some time to become effective. During the interval a large number of connections might be made effective. Alternatively, the making of each connection might contribute towards an inhibition of further connectivity, so that only a given number of connections can be made. Another possibility is that the capacity to make

effective connections depends upon a slowly maturing biochemical process; cells become "ripe" at a critical time. Whatever mechanism is assumed, if the rule is to be implemented, the capacity to make effective connections is restricted to a finite time interval.

Suppose that the state of activity of the integrating neurons is not strictly synchronised, either because of the non-synchronous arrival of afferent volleys, or by differences of thresholds. One sub-set of integrating neurons might be sufficiently activated to form connections before another set, and so forth. The first sub-set will not make integrative connections with the second sub-set until the latter becomes active. Suppose that the first sub-set has already been making connections with other integrating neurons active earlier. Then by the time the second sub-set becomes active, it will have already made some of its permitted connections. Consequently it will cease making connections with the second sub-set whilst the second sub-set still has some unexpired time and/or connections. By that time the third sub-set will bocome active and the second sub-set will be able to make some connections with the third.

In such circumstances each Linked Constellation would have some integrative connections with neurons of the one formed just before and that formed just afterwards. A Linked Constellation is to some extent a concept of analysis pointed towards discrete simultaneous processing. It is modified by continuous non-synchronous processing.

Most theoretical models since Hebb possess temporal sequential properties. One characteristic of the present model is that when a Linked Constellation becomes reactivated it passes some of its excitation forward, temporally speaking, into the neurons of the Linked Constellation formed just after its own formation. It also passes excitation backwards temporally: such backward movement of excitation however, has little effect on the model. First, in a repeated sequence of events, the Linked Constellation formed just before that matching the current input will have just been reactivated by the preceding input, and will be subject to recurrent inhibition. Secondly, there would be a roll forward effect, each subsequent Linked Constellation in a reactivated temporal sequence adding a little more feed-forward excitation to prepare the next one. This provides a priming excitation in the model, similar to that hypothesised by many theorists, e.g. Grey Walter (1969), Konorski (1967), Neisser (1967).

It seems probable that any system which does not have simultaneity of neural events would possess similar temporal sequential properties. In the early formulation of the model, considerable stress was laid upon this aspect. This was because of the emphasis which I placed upon temporal integration in the nervous system (see Chapter 2). It now seems to me that this aspect of the system is merely a special example of the principle proposed by Uttley (1956,b). Uttley proposed that delays introduced into pathways would result in the representation of temporal sequential events by a spatial pattern. (See also Brindley, 1969). In terms of the model now being proposed, consider a delay of two sequential steps introduced into the sensory pathway between sense organ and cortex (Figure 5.10). A sensory pattern matching a Preset Constellation (an Innately Recognised pattern, or Innate Releasing Mechanism) would reach the Lower Level (sub-cortex) and release the efferent response in time for it to reach the cortex at the time the step just before the releaser arrived. A Linked Constellation would be formed in the cortex containing a record of the sensory stimulation preceding the IRM and the response released by the IRM. These Uttley style delays are utilised considerably in later developments of the model, but were not employed in most of the runs of the early versions. The latter relied upon the general sequential properties attributed to Linked Constellations themselves.

In later developments it was in fact found that temporal sequential properties can be an embarrassment. A strong sequence can lead to such a high priming excitation that the model does not adapt to changed inputs. This may be one of the reasons for the very slow rhythms of some phases of sleep. Assuming that processing only occurred at the peak of the waves, excitation would have time to dissipate between patterns. The temporal sequential process would be broken up. This would prevent too strong a "set" arising in the system from the constant repetition of similar sensory information.

VIII Summary

Neurological data is interpreted as indicating that neural tissue consists of neuronal pools of integrating neurons. Each pool is driven by an afferent fibre and drives an efferent fibre. Representative

neurons from excited pools are connected together across pool boundaries to form Linked Constellations. Sensori-motor Linked Constellations are formed by sending active afferent volleys to motor pools at the same time as sensory volleys arrive at sensory pools. Linked Constellations are isolated from each other and remain in the brain unaltered except for decay. There is some communication between a Linked Constellation and the one formed just before and that formed just after its own formation.

A Lower Centre is similarly organised except that its neurons are assumed to have formed integrative connections innately. This may correspond to the Striatum. Preset Linked Constellations mediate Innately Recognised Patterns. Preset Constellations may be sensori-motor so that IRP's act as releasing mechanisms or Unconditioned stimuli, triggering responses at the Lower Level. The efferents from the Lower Level act as afferent fibres to the motor pools of the Upper Level.

Upon repetition of a sensory pattern the Linked Constellations which registered the earlier occurrences are reactivated. This results in a learned motor output from the Upper Level (cortex). When a Linked Constellation is reactivated some of its excitation passes forward temporally speaking to the next Linked Constellation in a familiar sequence, priming it.

Inhibition acts as a variable threshold. Only the most highly excited neurons are permitted to discharge and to form integrative connections.

CHAPTER 6

SOME FURTHER ASPECTS OF THE MODEL

In this chapter some aspects of neural integration will be considered which are not exclusive to the model presented, but to which the model lends a distinctive character.

I. Balanced receptor and effector systems

A limitation of the model as so far described appears if one considers the effects of patterns of stimulation involving unequal numbers of active afferent fibres. Consider, for example, receptors in an hypothetical primitive nervous system signalling on some occasion a retinal image of a large filled circle, and on others, a small spot falling within the same retinal field. The Linked Constellation representing the large circle would possess neurons in many columns; that representing the spot, neurons in only a few. Moreover those few columns would be activated equally by both images. The Linked Constellation representing the spot would be maximally reactivated by many large retinal images which happened to cover the retinal area of the spot, in particular by the larger circle. A large and complex image would maximally reactivate many Linked Constellations representing fragments of that image. By comparison there would be few images which would maximally reactivate the Linked Constellation representing the large image.

The consequences of such inequalities would be far reaching. A retinal image covering a large number of receptors would maximally reactivate so many Linked Constellations in addition to the one representing the truly matching pattern, that discrimination or identification of that pattern would be impossible.

There are a number of ways of alleviating this weakness. One way is by the development of feature analysers, e.g. contrast sensitive receptors, which would not signal the large circle and small spot in an identical manner. Such analysers are of course known to be present in sophisticated nervous systems.

Another way of alleviating the problem in a system with temporal sequential effects is by a temporal sequential variation of patterns associated with an object. If perceptual events consist of integrated sequences, the number of active fibres would balance each other from step to step by randomisation. An individual pattern might not be recognised, but a non-matching sequence would be rejected.

A more reliable way of dealing with this difficulty is to provide a rule whereby every sensory pattern stimulates an equal number of fibres. Such a rule would result from a system of receptors which signals the absence of stimulation as well as its presence.

In sophisticated animals Nature seems to have provided both the mechanisms required to obviate the difficulty.

With regard to the randomising effects of sequence, eye movements randomise the relative proportions of bright parts to dark parts. In a system reporting relative intensities whis would equalise the number of active afferent fibres. The elements reported by Hubel and Weisel (1959) and Jung (1961) would also have the effect of transforming a perceptually unitary event into a sequence of events. Some units would signal the beginning of illumination, some the end. Eye movements and the shuttering of eye-blinks do not only randomise the image, they also convert a continuous pattern into a kind of tachistoscopic viewing.

With regard to neural arrangements whereby all patterns would activate an equal number of fibres, sophisticated nervous systems exhibit mechanisms directed to this end. Hubel and Weisel (1959) reported some elements which have their spontaneous firing inhibited when a spot of light falls upon the centre of the field; their firing indicates an absence of light. Others fire when light falls upon the centre of the field. More generally, the receptors reported by Hubel and Weisel show a considerable variety of coding. Any pattern is likely by random sampling to balance the stimulation of a large number of receptors of one kind by the stimulation of a smaller number of receptors of another kind, so that the total number activated over the whole system remains constant. Jung (1969, 1972) described the system found in his researches with Baumgarten: "Our concept regarding the coding of luminance is based upon the existence of two reciprocally organised and interacting sub-systems, B and D, which signal brightness and darkness respectively".

A similar effect may be achieved in nervous systems from naturally reciprocal states. Agonist and antagonist muscles move in concert. A joint in one position cannot be in another at the same time. Mountcastle (1957) shows the interlocking responses of two reciprocal elements which fire alternately as the limb is extended and retracted at the elbow.

Mason (1975) argues that "the three effector or motor systems which mediate the integrative functions of the brain" are all regulated by a balance of opposed forces. The endocrine system, the Autonomic nervous system, and the skeletal muscle system each have a positive and a negative component. Hunger and thirst are also mediated by reciprocal systems (Deutsch and Koopmans, 1976).

Some of the sensory systems of sophisticated animals are intrinsically balanced by their nature. Whether colour vision is mediated by reciprocal pairs of colour sensitive elements, as has sometimes been suggested, or not, colour vision implies an overall equality of the number of active fibres in most visual patterns. A greater proportion of an image stimulating receptors sensitive to one colour implies that fewer receptors sensitive to other colours are being stimulated. Similar considerations apply to the sense of taste.

A number of other model designers have incorporated reciprocal receptor systems in their formulations (Konorski, 1967; Uttley, 1956; Marr, 1970). These authors do not describe what would be the effect in their respective models of omitting the reciprocal aspects. Uttley demonstrated formally that a signal reporting the absence of stimulation is necessary to identify inputs in an hierarchical system. As earlier mentioned, Konorski also utilised the idea extensively in his model.

It is difficult to argue with authority concerning models other than one's own without going too far from relevant concerns. There does however appear to be one way in which the effect of these considerations is different in my model from those mentioned. In an hierarchical or other logical switching model of the nervous system, the overlapping of a pattern involving a small number of active fibres by one involving a large number produces calculation errors which in theory prevent identification. The same effect appears in my model, and in both kinds of model can be obviated by the means described above. However, the effects of inadequately compensating for the weakness would appear to be

different in the two kinds of model. My model, being an analogue distributed process, would mean that a pattern involving a large number of active fibres would cause a large inflow of excitation into the system whereas a pattern involving a small number of fibres would only energise the system at a much lower level. This consideration would appear to have two consequences: first, there is an additional powerful reason for equating all input patterns, either by equalising the number of active fibres in each possible pattern, or by temporal sequential randomisation: secondly, failure to equalise would produce an arousal effect. The latter aspect is dealt with further in the next section. Here it may be noted that a large bright retinal image seems to have greater attentional power than a small bright image.

II Binary coding of intensity: arousal

Binary coding of information is a generally accepted convention in neurological models. It represents a simplification enabling the main princple of à model to be seen without complication. In my model binary coding of information is regarded as an actual, although secondarily evolved, aspect of the organisation of sophisticated nervous sytems. The binary coding leaves a parameter of neural activity, the spike frequency, free to represent a second aspect of events, their biological importance to the animal.

A. The coding of intensity

It is well established tha the transduction of energy at the sense organ produces a frequency coding of intensity (Granit, 1955; Ochs, 1965; Horridge, 1968; Uttal, 1973; Kandel, 1976). In contrast the neural coding of intensity in the higher centres is by no means clear (Uttal, 1973). Although there is considerable evidence of a partial binary coding there is much evidence to the contrary.

The phenomena of "range fractionation" and "line labelling" (Horridge, 1968) are more or less general attributes of nervous systems. Range fractionation combined with a graded response of the receptors leads towards binary representation of intensity. For example, in the eye of the worker bee, light of varying intensities will stimulate either one, two, three or four of the four categories of colour receptor (Autrum and von Zwell, 1964, cited in Horridge, 1968). Many of the

studies of feature analysers and columnar organisation of the cortex cited earlier showed graded responses of more than one column depending upon the intensity of stimulation.

Kluver demonstrated in 1941 that in the absence of striate cortex macaques can discriminate the brighter of two panels but cannot distinguish a smaller from a large figure if the total luminous flux is the same (Doty, 1973). Doty also cites work by Schilder (1966) and Schilder, Pasik and Tauba (1967) which adds further circumstantial evidence that brightness is coded by numbers of units.

On the other hand there is abundant direct evidence that frequency of discharge of cortical receptor cells is related to intensity of stimulation at the sense organ. This is implicit even in the range fractionation results cited above, and in the response of feature analysers by Hubel and Weisel, Asanuma and his associates and Mountcastle cited earlier. For example, Jung (1969, 1972) found a frequency relationship between intensity at the retina and spikes at the cortex. "The brighter the light stimulus, the more rapidly the B neurons fire and the stronger is the inhibition of the D neurons." (Jung and Baumgartner, 1955).

Uttal (1973) reviews his own and other work on somato-sensory, visual and auditory generator potentials. Uttal considers as candidate codings for central processing (i) the number of responding elements, (ii) response amplitude as measured either with the amplitude of a graded potential or (iii) with the integral of a pooled response. He also considers favourable the possibility of "the interpulse irregularity of the spike train". Granit (1955) considers that the frequency code cannot in itself account for intensity discrimination and suggests a combination of frequency coding and the number of units activated.

In summary, what is known of neural coding, suggests that the frequency aspect of the representation of intensity remains but is overlaid by the complexity of stimulus analysis.

B. Stimulus intensity

For purposes of analysis one may assume that in a primitive nervous system organised according to my model, differences of intensity are reflected in the frequency of the afferent fibres driving columns (Figure 5.7). Fibres from elements stimulated intensely would contribute

more excitation to the distributed excitation matching process than those from less highly stimulated elements. This would provide an effective weighting of some elements over others.

It may be noted that, in the absence of other forms of coding than frequency or amplitude, the registration of intensity implies a statistical distribution of individual neuronal thresholds. If all integrating neurons possessed identical thresholds, all the unlinked neurons would either be linked or not by an afferent volley, depending upon whether its frequency was sufficient or not to bring the neurons above the threshold level. A statistical distribution of individual thresholds was in fact provided in my model. Accordingly the number of neurons linked in a column was related to the frequency of the afferent volley.

One of the effects of this system is that Linked Constellations recording high intensity events would contain a greater number of neurons than those registering low intensity events. Given a memory decay, these Linked Constellations remained effective longer in the system. Similar considerations would apply to portions of a Linked Constellation registering intense stimulation of selected receptors. Such inequalities would have considerable attentional effects. For example, compare the discriminative value of a pattern with or without a bright spot, as against with or without a spot of the average brightness of the presentation. Upon subsequent presentations of the pattern with the bright spot, the matching Linked Constellation will be reactivated relative to the non-matching one to a greater degree than would be accounted for by the number of columns activated by the spot. The greater number of linked integrating neurons in those columns contribute a relatively large proportion of the total excitation. Frequency representation of intensity provides a natural discriminative system in which intensity is identified with importance of discrimination. The function of intensity, insofar as there is a frequency coding, may be seen as an attentional mechanism.

In terms of the extended template system described in Chapter 2, high intensity stimulation reaching one of the template systems will draw attention to the information contained in that part of the system.

It is probably not possible to discover what a primitive nervous system would be like. Those that are with us, however low we might choose to

put them on the evolutionary scale, have all been subject to a similarly long period of evolutionary development. However, looking at what we consider primitive aspects of animals and their behaviour, the proposition does not seem unreasonable. Loud noises, bright light, strong smells and similar intense stimuli seem naturally to attract the attention.

C. Arousal

Although it may be true, especially in the environment of primitive animals, that intensity of stimulation is correlated with significance, this would be a dangerous generalisation for an animal to make. Upon occasion it might be urgent to pay attention to an aspect of the environment with low intensity stimulation which by an association of either learned or innate, signalled some important event. The hissing of a snake against a background of louder animal noises is an example. The development of feature analysers may have played its part in converting what would be a low intensity aspect of a mixed template presentation into a high intensity aspect of a specialised or "feature" template.

Sophisticated animals can learn to pay attention to cues of low intensity and small magnitude if they signal events of importance.

The effect in the model of increased frequency of volleys in afferent fibres resulting from intensity, suggests that the frequency of spikes might be utilised to control the level of arousal of the system or of parts of the system, by mimicking intensity. The development at the same time of neural mechanisms moving towards a binary reciprocal coding of both intensity and of the number of active elements, would render the animal more independent of actual intensity and image size in adjusting the system to its needs. Such an hypothesis suggests that as animals became able to judge the meaning of events rather than their natural stimulus qualities, the arousal system moved from one based upon intensity towards one which was more independent of intensity.

In the model as so far presented, the information from the environment is coded in binary form. The binary pattern can be fully specified by the pattern of activity and inactivity regardless of the actual frequency of spikes which characterises an active volley. Moreover, the binary pattern will remain fully specified even if the

nominally "inactive" fibres carry spikes, provided that the "active" fibres have a higher frequency. The difference in frequency of trains of spikes signalling activity or inactivity of fibres provides a variable which may be adjusted to weight the operation of the model in ways analogical to the effects of intensity.

In the model this variable is termed the Discriminative Volley Value (DVV). The frequency of the inactive fibres is termed the Diffuse Volley Value (Diffav). By reason of its definition Diffav will be applied to all fibres whether active or not, the active fibres having a frequency of the sum of DVV and Diffav. These variables are to be regarded as parameters of the system in the sense used by Ashby (1952), that is, they change only as step functions to set the level of operation of the neural apparatus. They would be independently adjustable in different parts of the system, parts corresponding with the separate template aspects of the system. The action of DVV and Diffav would be analogical to the contrast and brightness controls respectively of a television receiver.

It must be confessed that although the notion of DVV and Diffav as attentional and arousal mechanisms were considered important in the formulation of the model, their incorporation led to very great problems. As will be seen when the results of the initial simulations are reported, they worked well in the early versions of the models. Nevertheless, the way they worked seemed quite different from what is known of real physiological attentional and arousal mechanisms. Moreover, in more ambitious versions, the attentional and arousal mechanisms overwhelmed the integrative functions. The theoretical solution of these problems was found to be involved with similar problems relating to inhibition. A more convincing hypothesis of these functions could not be formulated until a more sophisticated model of the overall control functions of the brain had been worked out. Anticipating, it was found that control of excitation and inhibition in the model could be readily achieved by introducing notional Inhibitory Gates into the afferent pathways to the arrays. These Gates periodically blocked the volleys which would otherwise have entered the arrays. By varying the proportion of the cycle during which the Gate was open or closed, the amount of excitation entering an array could be controlled even though the frequency of spikes in the active fibres remained constant. In effect the basic

concept of arousal remained as before, being dependent upon the amount of excitation penetrating to a particular array.

In summary, this Section proposes that there is a natural connection between intensity and attention. The complexities of reciprocal representation, range fractionation, line labelling, and feature analysis frees the system to a large extent from this natural dependence upon intensity and magnitude of the sensory image. Attention may then be controlled by the brain itself according to the meaning attached to the stimulation. The brain accomplishes this by adjusting the amount of excitation reaching the various arrays so as to control the relative contribution of each array to the ongoing overall distribution of excitation. Essentially, the amount of excitation is achieved by frequency coding, although in sophisticated nervous systems this frequency aspect is complicated by hypothetical rhythmic mechanisms.

III Decay of connections: ripeness for linking

One or two matters which are not of central importance to the model must be mentioned before reporting the simulation.

The extent to which learning decays in the central nervous system is not clear. Recent studies suggest that what diminishes with time may be the ability to retrieve. Nevertheless, a decay factor has been incorporated into the model as the most likely substrate of logarithmic forgetting. There are a number of reason for this. One may assume that the neuronal mechanisms of alteration which underlies learning, whatever it may be, is a complex biochemical process involving molecular changes. It is now well established that such processes have characteristic half-lives. Secondly, although some forms of neuronal mechanism by which the model might be mediated would provide an ample number of "elements" without having to use them over and over again, this would not apply if the element is an entire neuron. One of the difficulties with any model utilising neurons as its elements is the apparently limitless capacity of the brain to store information. Even with the very large number of estimated neurons in the brain (10^{10}) it still seems difficult to reconcile the system of the present model with the ability to remember in great detail an incident which occurred many years earlier. This difficulty is relieved by a

decay system in which elements released by decay of connections become available to be used again. Thirdly, the decay of a Linked Constellation by the gradual dropping out of the elements of which it was originally formed, seems to coincide very well with our subjective fading of the vividness of memory.

To accord with the general rules of the model, it is necessary to assume that neurons which have made connections with one subset of the neuronal population can make no further connections until they have been completely released from their former connections.

It is possible that learning elements, such as neurons, may have characteristic "memory-life". The statistical distribution might vary from area to area of the brain, and would be a useful parameter. An implication of such an assumption would be that each Linked Constellation would contain a sample of the distribution of memory-life of the population at the time of linking. This distribution would change with the age of the animal. As Linked Constellations decayed, short-term memory-life neurons would be released and come back into the general population. Long-term neurons would remain locked into Linked Constellations. Consequently the proportion of short-term neurons in the general population would increase. This would be in accordance with the general opinion that memories laid down in youth are more permanent than those laid down in age. Other possibilities that suggest themselves are a correlation of duration of connections and threshold. This would mean that learning laid down during arousal would last longer. These examples are merely intended to indicate possibilities. They are probably not specific to the present theory.

One matter of mechanics must be mentioned. As described so far, the first afferent volley reaching a column would link all the available integrating neurons in that column. Virtually the whole of the information storage capacity would be exhausted. This difficulty can be overcome in a number of ways. First, it may be assumed that only a limited number of neurons (or other elements) mature at each instant. Secondly, a proportional rate of decay would mean that immediately after the formation of a Linked Constellation, a number of neurons would return to the general pool to be available for use in further learning. More probably there would be random factors which restrict linking to only a random proportion of those neurons fulfilling the

required conditions for linking. Such factors might not be random. Since integrative connections are likely to be controlled by complex biochemical processes, it would be reasonable to assume that they might be the subject of cyclic metabolic syntheses. The neuron might be only "ripe" for linkage at one particular phase. When released from a Linked Constellation, a certain time might need to elapse before the cell would be once again "ripe" to make integrative connections.

IV Perception and sensation

An impression may have been given that, according to the model, Linked Constellations register only the pattern of afferent excitation: that what is laid down in the brain is simply a record of sensory and motor events. This would reduce the model to an older kind of enrichment model which consisted of "bare sensations" enriched by additions of stored information (Gibson and Gibson, 1955a and b; Postman, 1955).

On the contrary, the concept of a bare sensation is foreign to the model. One may for simplicity provisionally accept that a particular sensory event results in a certain pattern of afferent excitation to the cortex. (Even this is probably untrue: there is considerable evidence of afferent neuronal inhibition at the sense organ and abundant evidence of modification of sensory signals at thalamic level, e.g. Sokolov, 1960). However, the pattern of afferent excitation is immediately modified when the afferent impulses interact with the integrating tissue.

It is true that the arrival of afferent excitation will have as its first effect the activation of all the integrating neurons within the respective neuronal pools activated. Immediately the recipient neurons will begin exchanging excitation with each other differentially according to their integrative connections. The pattern of excitation in the tissue will be partly dependent upon the afferent excitation and partly upon the Linked Constellations present in the tissue. The rule for formation of integrative connections relates to the state of excitation. Consequently, the Linked Constellations laid down will not register the afferent excitation itself, but rather the afferent

excitation as modified by the tissue response.

The implications are increased by temporal sequential processing. In a sequence of patterns, the state of excitation of the tissue will grow. In effect the afferent excitation will meet a tissue already excited by a dynamic central process (compare Hebb, 1949).

In the model, what the animal perceives at any instant is assumed to be correlated to the activity of the population of integrative neurons. This will not be a function only of the afferent signals but of the resultant of the interaction of the signals with the stored information. The Linked Constellations which are laid down of later events will be more complex than those laid down by earlier events. Linked Constellations are records of perceptions not sensations.

It may also be noted that the model is not perhaps atomic in quite the sense which has been conveyed. Emphasis has been laid upon the atomic character of the receptive fields, which it is suggested are the elements from which neural templates are composed. However, in contrast with such atomicism is the integrated unity of action of the templates, i.e. the Linked Constellations. The atomicism of the model may be compared with the atomicism of contemporary physics as contrasted with an earlier concept of atoms which were immutable elemental objects, unchanging even when combined with others to form compounds. In contemporary physics an atom is itself a complex object, and when combined with other atoms, loses its individual identity. The resulting molecule can only be considered a unique object in its own right.

V Brain organisation: a simplified model

The main characteristics of the model may be seen by ignoring many of the complications of motor organisation and considering a classical conditioning paradigm.

Volleys in afferent sensory fibres activate the neuronal pools in both a sub-cortical and a cortical array. The sensory representation is not necessarily identical in each. The sub-cortical (striatum) level has Innate Linked Constellations. These mediate Innately

Recognised Patterns, which may be thought of as UCS or as the Releaser part of Innate Releasing Mechanisms. The innate Linked Constellations are sensori-motor and so correspond to Unconditioned Reflexes or to IRM's.

The cortical level lays down a record of events in the form of Learned Linked Constellations. When the sensory input at the sub-cortical level matches the sensory side of one of the innate Linked Constellations, that Linked Constellation is maximally reactivated: a motor response appears across the efferent fibres of the sub-cortical level. This efferent volley is sent via the Thalamus to the cortex, where it acts as an afferent volley to its corresponding efferent neuronal pool. This afferent motor volley causes integrating neurons to be linked into the Linked Constellation forming at that time. The latter accordingly registers both the sensory input to the cortex and the response emitted at the sub-cortical level. It is a sensori-motor Linked Constellation.

When a familiar sequence of sensory patterns is repeated, the sequence of matching Linked Constellations formed earlier in the cortex is reactivated. These Linked Constellations have temporal sequential integrative connections. When the first pattern of the sequence arrives it brings the matching Linked Constellation to maximum excitation. Additionally, some of the excitation flows forward, temporally speaking, to the second Linked Constellation of the sequence. This sequence is primed before the second pattern arrives. After the second pattern has arrived, the third Linked Constellation becomes primed, and so forth. Because each subsequent Linked Constellation receives an increment from the preceding one, the level of excitation grows through the sequence.

Accordingly, the system should anticipate events when a familiar sequence is repeated. The anticipation takes the form of a rise in the excitation of those Linked Constellations about to be stimulated by the next pattern in the sequence.

A special case of this anticipation occurs in situations corresponding to the classical conditioning paradigm. The CS may be thought of as a sequence of sensory patterns, or cue sequence. The UCS is represented by an IRP (Innately Recognised Pattern). However often it has been repeated the cue sequence will have no especial effect upon the innate

Linked Constellations of the sub-cortical level. It will however reactivate the matching Learned Constellations in the cortex. Before the step in the sequence when the Learned Linked Constellations representing the IRP are reached, they will have been primed. In trials when the IRP is actually presented ("reinforced trials") the sensory input reactivates both the sub-cortical level Preset Linked Constellation and the cortical Learned Constellation. A motor response occurs at both levels, due to transfer of excitation from sensory to motor neurons of the Linked Constellations.

Suppose the IRP is not presented after the cue sequence (an "unreinforced trial"). The sensori-motor Linked Constellations registering previous presentations of the IRP are nevertheless primed by the cue sequence (CS). Consequently there is a discharge from the cortical motor pools. It is proposed that this cortical motor response corresponds to the CR of the classical conditioning paradigm.

In the model the Conditioned Response can be distinguished from the Unconditioned Response, because the latter occurs primarily at the sub-cortical level, the former primarily at the cortical level. This is necessarily a simplification. There is evidence of circular arcs, so that either response might activate the other level. Nevertheless, the model indicates that the CR is the effect of reactivation of learned records and is an independent process from the UCR which is a reactivation of innate arrangements.

It will be noted that CRs are only one special case of anticipations in the model. Such anticipations occur whenever a familiar sequence is repeated. In later developments of the model these anticipations at cortical level are used to initiate motor action. However, this does not imply that instrumental motor action is controlled by Classically Conditioned Responses, as discussed in the controversy between single process and dual process conditioning theories (see e.g. Rescorla and Solomon 1967). What is suggested is that anticipation in the form of a priming excitation underlies both the phenomena seen in classical and in instrumental situations.

The main points about the model at this stage concern: (a) the laying down of a comprehensive learned record of events; (b) the release of innate responses by innate connective systems at sub-cortical levels and (c) the reactivation of sensori-motor records at cortical level when familiar events occur.

CHAPTER 7

SIMULATION OF INITIAL VERSIONS OF THE MODEL IN THE ATLAS COMPUTER

I Introduction: general form of the model

The simulations consisted of a series of programs which underwent continuing modification and development. The aim of these modifications was to improve performance and to introduce additional notional neural mechanisms. A natural break in programming occurred when the University of London Atlas computer was closed down. Computing was then transferred to the CDC 6600 computer. The opportunity was taken at this point to revise the general form of the model. These later programs will be reported as a separate set later in the thesis.

In presenting the simulations, terms will be given initial capitals when it is desirable to distinguish physiological concepts, as they appear in their simulated or more speculative form in the theoretical model, from those concepts in their well established forms in contemporary neurophysiology. For example, Inhibition refers to aspects of the theoretical model and its simulations, whilst neuronal inhibition refers to aspects of neurophysiology as well established.

One aim of the simulations was to demonstrate some of the properties claimed for the system; for example that it would store inputs in the form of spatial templates and notionally recognise those inputs when they were presented on subsequent occasions.

A second aim was to find a design which would demonstrate the system as a basis of overall neural integration and at the same time would produce behaviour comparable in some sense with that of animals. The classical conditioning paradigm offers a number of advantages as a test of the basic characteristics of the model. It avoids many of the complications of motor organisation whilst providing a measurable response. The kind of response seen in classical conditioning (e.g. salivation, chewing, eye-blink, leg-withdrawal) corresponds to the kind of response attributable in the model to the action of Preset

Constellations (genetically determined Integrative Connections). Moreover one can avoid many of the theoretical problems concerned with the instrumental modification of behaviour and the relationship of that modification to the consequences of action upon the environment.

A simple model built around a classical conditioning paradigm offers an essential ingredient for a first step in pursuing the aim of simulating sensori-motor integration: it provides an observable response with a particular character. The initial model attempted to show that the particular behaviour observable in the classical conditioning situation would occur in an extended neural template system in the form of Linked Constellations, in which effector activity was included within the template scheme. This scheme was later developed to account for behaviours outside the range of those observable in a classical conditioning experiment.

II The basic layout

The basic form of the initial simulations followed the general arrangements shown in Figure 5.10. The essentials of that diagram are a Cortex, a Sub-cortex (possible Striatum) and a Thalamus, acting as a routing centre.

The Cortex and Sub-cortex each consisted of a set of Neuronal Pools organised as shown in Figure 5.7. The Cortex followed the further arrangements shown in Figure 5.12, i.e. it incorporated Sensory efferents and Motor afferents.

Motor action in the Cortex depends upon inputs from the Motor Afferents. It will be recalled that when those Fibres are active, the Integrating Neurons of the Motor Columns of the Cortex become excited and so are able to become linked to the Integrating Neurons excited in the Sensory columns. In this way Sensori-Motor Constellations are formed. In the simulation, activity was initiated in the Motor Afferent Fibres to the Cortex by the output of the Motor Efferent Fibres of the Sub-cortex. The Motor Efferent Fibres of the Sub-cortex passed to the Thalamus from where they were sent as Motor Afferents to the Cortex.

The functions of the Thalamus shown in Figure 5.10 comprised various processing procedures. Although most of these procedures seemed to be analogical of processes which might be mediated in the thalamus and associated reticular tissues, to avoid undue theoretical commitment, the Thalamus of the model will be termed a Control Centre. It was more than a notional relay station. Before the output of the Sub-cortical Array could be applied to the Cortical Array a certain amount of processing took place. Processing has to be applied to determine the notional frequency of the spikes in Afferent Fibres. This was partly concerned with the formal distinction between an Active and Inactive Fibre and partly to give effect to a primitive form of "arousal" assumptions discussed in Chapter 6. As the system developed it was also found necessary to control levels of Excitation and Inhibition in the Arrays. This control seemed to derive naturally from a variation of the Afferent Excitation passing through the Control Centre.

It will be appreciated that the suggestions of correspondence between parts of the model and regions of the brain is intended to be suggestive only. The aim is to assist in visualisation of the system and in the general interpretation of the brain organisation proposed.

On the physiological side, the initial simulations implemented (a) the concept of Linked Constellations (b) afferent and efferent excitation (c) neuronal inhibition (d) spike discharge of neurons (e) a primitive form of the arousal proposals discussed in Chapter 6 and (f) a representation of the decay of Integrative Connections.

These will now be discussed in detail.

III Simulation of Linked Constellations

In the main body of the simulations the actual connections by which Linked Constellations are mediated were not directly simulated. A separate simulation demonstrated the mediation of Linked Constellations by orthodox formal synaptic connections. Although the model is not restricted to such a connective medium, the separate simulation served to demonstrate the practicability of the Constellation concept. It will be reported in due course.

In the main simulations, a Linked Constellation was simulated by assigning a marker to the representation of an Integrating Neuron in the computer. This marker, the Constellation Label indicated which Constellation, if any, the Neuron belonged to. The Constellation Label was an integer which became attached to the Integrating Neuron when it satisfied the conditions for forming Integrative Connections. All Integrating Neurons which bore an identical Constellation Label were regarded as being interconnected in a manner necessary to fulfil the functions attributed to Linked Constellations. These functions were then simulated by processing the notional set of physiological variables by which Integrating Neurons were represented in the computer.

In particular the two main functions simulated were:
(i) the Excitation of the Integrating Neurons belonging to a particular Linked Constellation was pooled and redistributed so that all Neurons in the Constellation had the same Excitation, equal to the arithmetic mean of the individual Excitations of the Neurons concerned.
(ii) as discussed in Section VII, Chapter 5, some of the Excitation of a Linked Constellation was transferred to the Integrating Neurons of those Constellations which had been formed just before and just after the time of formation of the one being considered.

Each main cycle of the program indicated a notional passage of time and was termed an Instant. A new Linked Constellation was formed at each Instant. The Constellation Label used to indicate membership of a particular Linked Constellation was in fact the number of the Instant at which the Constellation was notionally formed. This number was used to identify all those Neurons to be treated as a Linked Constellation and served merely as a computing device to simulate Integrative Connections.

IV Simulation of Integrating Neurons

Each Integrating Neuron was represented by a set of six variables and Constants held in computer store. These comprised (1) the Constellation Label; (2) the Excitation Count; (3) the Inhibition Count;
(4) the Natural Threshold as discussed in Section IV, Chapter 5;
(5) Connection control constant characteristic of the individual Neuron

(6) Connection control representing passing of time.

These sets of variables were arranged in computing arrays, so as to represent the neural organisation proposed. Each Neuron belonged to a particular Neuronal Pool or Column. These Neuronal Pools were in turn organised in Neural Arrays (or Arrays) which represented the columnar arrangement of cerebral tissue discussed in Section II, Chapter 5. The variables representing Integrating Neurons were updated at each main loop, i.e. Instant, of the program.

A. Excitation: Spikes

The Excitation Count recorded the state of Excitation of each Integrating Neuron. For convenience an arbitrary unit of Excitation was termed a Spike. A Spike was the Excitation which an Integrating Neuron received from an action potential at one of the notional synapses upon it, or the Excitation which it imparted to other Integrating Neurons when it underwent an action potential.

The Excitation Count was initially zero (in early simulations) and subsequently recorded the net results of any additions or subtractions of Excitation. These transactions, which represented events taking place during the Instant, may be considered under three heads:

(a) Excitation received from Afferent Fibre (addition); Excitation lost in emitted Spikes (subtraction). The Spike emitted by a Neuron within a particular column was added to the Spike Frequency of the Efferent Fibre of that column.

(b) Transfers of Excitation within Linked Constellations (Linked Constellation sharing effect).

(c) Dissipation of Excitation (e.g. cell wall leakage, metabolic effects).

With regard to (a) at each Instant a Spike Frequency was calculated in the Control Centre (Thalamus) for each Afferent Fibre. This value was added into the Excitation Count of every Integrating Neuron notionally driven by that Fibre. After processing the various neuronal exchanges of Inhibition and Excitation (see below) in each Integraing Neuron, the Discharge of each Integrating Neuron was calculated (Section IV, item (v) of Chapter 5). The Discharges of all the Integrating Neurons within a column were supposed to excite the Efferent Neuron (e.g. pyramid) of that column and so contribute to the Spike Frequency of the Efferent Fibre of that column. In practice the Spikes emitted by the individual Integrating Neurons within each Pool were simply

aggregated to provide the Efferent Spike Frequency.

Within each Integrating Neuron the calculation was as detailed in Section IV, Chapter 5. The number of Spikes emitted by each Neuron was the excess of Excitation Count over Inhibition Count multiplied by a proportional (i.e. reducing) constant. Neurons would be unlikely to reduce their state of Excitation to the point of cessation of Spike emission within the short interval represented by an Instant.

The exchanges of Excitation through the Integrative Connections were simulated by adding the Excitation Counts of all the Neurons in a Constellation and sharing the total equally among the Neurons in the Constellation. A proportion of the total of Excitation Counts in Constellation 'T' was then transferred to the Neurons in Constellation 'T + 1' and Constellation 'T - 1'. At each cycle these operations were carried out for all Constellations.

The effects of the dissipation of energy in the action potential and its possible further dissipation over time was represented by reducing at each Instant the Excitation Count of each Integrating Neuron. The reduction consisted of the subtraction of a fixed arbitrary number of Spikes and a proportion of the Excitation Count of each particular Integrating Neuron in the system. This was intended to represent in rough form losses through a leaky channel. This aspect of the simulation was later seen to be related to the particular form which the operations in the simulation had taken. This became clearer in later versions of the model.

B. Inhibition

Inhibitory Control Neurons impose inhibitory synapses upon the Integrating Neurons. The number of Inhibitory Spikes received by an Integrating Neuron controlled a variable threshold at which the Neuron propagated an action potential (emitted a Spike). It also controlled a variable threshold at which the Neuron could form Integrative Connections.

It was unnecessary to simulate Inhibitory Control Neurons individually, since in these simulations their effect was always considered to be distributed evenly across all the Integrating Neurons of a Neuronal Pool.

The Inhibition Count was a variable which recorded in each Integrating Neuron the results of any additions or subtractions of Inhibition on that Neuron during the Instant. The mode of calculation of the quantities of Inhibition to be added or subtracted from Inhibition Count varied from version to version of the simulations. In principle an increment represented a release of inhibitory transmitter from the Inhibitory Control Neurons, and a subtraction the destruction of inhibitor transmitter or the dissipation of inhibitory effects within the Integrating Neurons. The increment to Inhibition Count was always made partly dependent upon some aspect of the Excitatory system at the preceding Instant, i.e. a representation of reactive inhibition. It was sometimes also made dependent upon the Spike Frequency of the Afferent Fibre at the Instant concerned, i.e. a feed-forward inhibition.

The simplest form was to add to the Inhibition Count of every Integrating neuron within a particular Neuronal Pool an amount intended to balance precisely the excess of Excitation Count over Inhibition Count, taken as an arithmetic mean over all the Neurons within that Neuronal Pool, as it had stood at the preceding Instant. To effect this as a reactive variable, Inhibition was related to the sum of the Spikes emitted by all the Integrating Neurons within the Pool, i.e the Efferent Spike Frequency of that Pool, i.e.

$$\text{Increment to Inhibition Count (to each Integrating Neuron in Pool)} = \frac{1}{k} \times \frac{\text{Efferent Spike Frequency}}{\text{Number of Integrating Neurons in Pool}}$$

Here k has the same meaning as in item (v), Section IV, Chapter 5, i.e. it represents the proportion of Excitation emitted as Spikes in one Instant. Accordingly multiplying by its reciprocal calculates the amount of Excitation within the Neurons from which the Efferent Spike Frequency derived.

C. Threshold

The assumption adopted in the initial simulations was that there is an excitation threshold at which an Integrating Neuron (a) discharges and (b) is able to form Integrative Connections. This threshold was dependent partly upon a natural threshold attributed to each cell, and partly to the influence of neuronal inhibition. As mentioned above, initially a single value was used as threshold for both functions (a)

and (b). It was, however, found later that thresholds applied to emission of Spikes introduced too great a variablility into the operation of the simulations consistent with the number of Neurons simulated. In later versions, Spike emission depended purely upon the excess of Excitation Count over Inhibition Count, and the Intrinsic Natural Threshold provided in each Integrating Neuron applied only to formation of Integrative Connections.

To simulate this idea an Intrinsic Natural Threshold was entered into each Integrating Neuron at the beginning of each run. This was a random number drawn from a rectangular or other statistical distribution. The range was chosen so that maximum Spike Frequencies would cause those Integrating Neurons with the highest Thresholds to form Integrative Connections, whilst minimum Spike Frequencies (i.e. low arousal) would cause at least some of the Integrating Neurons to form connections, i.e. those with Thresholds at the bottom of the range.

Since the random numbers entered in Neurons at the beginning of each run differed from run to run, no two runs of the model were entirely identical.

D. Formation and destruction of Linked Constellations

Reasons were given (Section III, Chapter 6) for including a notional decay principle. Since the method adopted is not central to the model, it will be briefly indicated without comment. Each Integrating Neuron contained a pair of numbers, one a constant and the other a counting variable indicating the passing of time. The constant, termed Memory Life Duration specified the number of Instants the Neuron could remain part of a Linked Constellation. It was an integer selected as a random number. This constant was used in conjunction with the counting variable, Memory Life Count, which indicated how long the Neuron had already been linked into its current Linked Constellation.

When an Integrating Neuron satisfied the conditions for formation of Integrative Connections, i.e. if it was (a) currently unlinked and (b) above threshold, then its Memory Life Count was set to zero. The Instant Number was written into the Constellation Label location of the Neuron to signify that it was now linked into the Constellation formed at that Instant. At each Instant all Memory Life Counts in the system were incremented by unity. When, in any Integrating Neuron,

Memory Life Count equalled Memory Life Duration, the Constellation Label was reset to zero, thus effectively releasing the Neuron from the Linked Constellation and returning it to the unlinked pool.

One of the rules which an Integrating Neuron had to satisfy before it was linked was that its Memory Life Count should exceed its Memory Life Duration. This had three consequences: (a) Each Neuron is a member of no more than one Constellation at any given Instant. No Integrating Neuron could be made part of the currently forming Linked Constellation whilst already a member of a Constellation. (b) Integrating Neurons were released from Constellations in a statistically determined manner, dependent upon the distribution of Memory Life Durations. The system gave an approximately logarithmic decay of the number of Neurons held in a Linked Constellation. (C) At any Instant only a proportion of the Integrating Neurons were available for linking, even if in other respects they satisfied the conditions for the formation of integrative Connections.

E. Summary

Integrating Neurons were represented in the computer by variables and constants which simulated in a simple way the neurophysiology of neurons as seen in the theoretical model. The aspects represented were Excitation, Inhibition, Integrative Connections in the form of a Constellation Label, and decay of Integrative Connections controlled by Memory Life Duration and Memory Life Count. There was also an Intrinsic Threshold.

V Some further particulars

As mentioned, the notional anatomy simulated was that shown in Figure 5.10 of Chapter 5. Both the Cortex and Sub-cortex were divided into a Sensory Division and a Motor Division. These were simulated in the form of Arrays. Each Array consisted of a set of Neuronal Pools (columns) and each Neuronal Pool consisted of a column of Integrating Neurons.

The Integrating Neurons of the Sub-cortical Arrays differed from those of the Cortical Arrays only in two respects: first, they were already members of Preset Constellations written in at the outset of the run;

second, they were given an indefinite Memory Life Duration, so that they would remain linked throughout the run, i.e. they represented permanent connections.

Each Neuronal Pool in the Cortical Arrays consisted of 100 Integrating Neurons. There were 16 Neuronal Pools in the Sensory Division and 4 Neuronal Pools in the Motor Division of the Cortical Arrays, i.e. a total of 2,000 Integrating Neurons.

Each Neuronal Pool of the Sub-cortical Arrays (Striatum) consisted of 16 Integrating Neurons. This was sufficient to accomodate the number of Preset Constellations written into the model, with some to spare. There were 16 Neuronal Pools in the Sensory Division and 4 Neuronal Pools in the Motor Division of the Sub-cortical Arrays. (See Table 7.1A).

The numbers just given represent the first layout of the simulation and enable the basic design to be seen. These were in fact increased for most of the runs reported as part of the initial set of simulations. This increase was to accommodate a system of reciprocal (balanced) representation, discussed in Section I, Chapter 6. The number of Integrating Neurons was also increased to 200 per Neuronal Pool (Table 7.1B).

The Thalamus (Control Centre) acted as a relay station which performed processing functions, such as the calculation of the Spike Frequencies of the Fibres leaving it. Theoretically its functions would be performed by Neuronal Pools, but in view of the simplicity of the operations, they were not formally simulated in terms of Neuronal Pools. The main function of the Thalamus was to send Afferent Fibres from the Cortex and transmit them through Afferent Fibres to the Sub-cortical Array. Accordingly the Thalamus may be regarded as having an equal number of Locations as there were Neuronal Pools in the Cortex (Table 7.1).

Effectively the Sensory Input of the model consisted of a set of binary digits, the Input Pattern. This consisted of 16 bits as first run, and 32 bits in later versions. This could arbitrarily and conveniently be regarded as a simple retina of 4 x 4 elements. The justification and significance of the isomorphic aspects of the model will not be pursued in this thesis.

TABLE 7.1

SUMMARY OF LAYOUT OF MODEL
IN INITIAL SET OF SIMULATIONS

	Sensory Division		Motor Division	
	No. of Neuronal Pools	No. of Int. Neurons in each N.P.	No. of Neuronal Pools	No. of Int. Neurons in each N.P.
A. As first run				
Cortex	16	100	4	100
Sub-Cortex	16	16	4	16
Thalamus	16 (Locations)	Not represented	4	Not represented
B. In later versions				
Cortex	32	200	8	200
Sub-cortex	32	16	8	16
Thalamus	32 (Locations)	Not represented	8	Not represented

VI Arousal: functions of Control Centre

The Spike Frequency of both Motor and Sensory Afferent Fibres was calculated as the sum of two parameters: (a) The Diffuse Volley Value, added to all Fibres of an Array; and (b) The Discriminative Volley Value, which had the value zero if the Fibre was notionally inactive and some other value, set by the system, if the Fibre was notionally active.

The events of biological significance to the model were the occurrence of Innately Recognised Patterns (IRPs) in the Input, that is, patterns which matched the Sensory part of Preset Linked Constellations. The appearance of such patterns in the Input could be identified by the occurrence of increased Motor output in the Subcortex. As a consequence Motor output Spike Frequency was used to vary the arousal level.

The values of the two parameters were set as functions of the Spike Frequencies in the Efferent Fibres of the model. One of the functions of the Control Centre (Thalamus) was to calculate these values and to apply them to the Sensory Afferent Fibres. In the case of these Fibres, the signal status, active or inactive, was determined simply by reference to the corresponding bit of the Input Pattern.

In the case of the Motor Afferent Fibres, i.e. the Fibres whose state derived from the Efferent Fibres coming to the Thalamus from the Subcortex (see Figure 5.10, Chapter 5), the Control Centre acted as an inhibitory filter. The Spike Frequencies of the Sub-cortical Efferents were subjected to simple arithmetic procedures which simulated a competitive inhibition system in a simple fashion. For example, in the simplest version, the arithmetic mean of the Spike Frequencies in all the Efferent Fibres was subtracted from each, any negative results being put equal to zero. The resulting Spike Frequencies were added to the Diffuse Volley Value and applied as the Spike Frequencies of the Motor Afferent Fibres to the Cortical Array.

VII Concluding remarks

In this chapter the mode of simulation of the essentials of the theoretical

model have been outlined. The structure of the model as simulated has been described and some indication given of the representation of physiological variables. The mode of calculation of Spike Frequencies in Afferent Fibres, including an arousal system, was also given.

CHAPTER 8

PERFORMANCE OF INITIAL SET OF SIMULATIONS

In this chapter will be reported the behaviour of the set of simulations which were in the form which has been described in the previous chapter. These were run in the University Atlas Computer and programmed in Atlas Basic Language.

I Demonstration of Linked Constellations as spatial templates

A. Registration of Input Patterns

In order to see whether the model functioned internally as predicted, samples of the Integrating Neurons taken from every Neuronal Pool in the model were monitored at intervals. The Constellation Label printed out for each Neuron enabled the identification of which Integrating Neurons belonged to a particular Linked Constellation. This indicated, for the benefit of the programmer, the Instant at which that Neuron had become notionally linked with other Neurons.

To form a spatial template of an Input Pattern, the Integrating Neurons belonging to the Linked Constellation concerned should be located in those Neuronal Pools which were excited by that Input Pattern. Clearly this was the intention of the design of the model, and in the case of a single Instant and a single Input Pattern, the demonstration would be trivial. It is however, a different matter when the process is set within a neural integrating system as a whole. The question arises as to whether it is practicable to suppose that adequate templates can be laid down of a large number of inputs occurring successively within a system containing a limited number of Neurons. It also seems necessary to demonstrate that this is possible in a system of variable physiological units like neurons, in which there are myriad exchanges of excitation and inhibition.

To confirm that the registration complied with the template notion, the Input Pattern presented at a particular Instant was considered bit by bit. The Neurons bearing the Constellation Label for that Instant

were then searched for in the printout and the Neuronal Pools to which they belonged ascertained. It was then checked that the pattern of Neuronal Pools containing Neurons with that Constellation Label matched the Input Pattern, i.e. that where a 1 appeared in the Input Pattern then in the corresponding Pool there were Neurons bearing the relevant Constellation Label. Conversely, where a 0 appeared in the Input Pattern there were to be no Neurons in the corresponding Pool which bore the relevant Label.

The system worked well. In the initial runs there were between 115 and 140 Instants. For purposes of reference a 'run' of the model means a run of the computer program in the computer for a given number of Instants during which its performance and operation could be investigated. A 'run' may be alternatively referred to as an investigation or experiment as appropriate in context. In each run some 15 to 20 different Input Patterns were presented. It was found that the location of the Linked Neurons corresponded to the Input Patterns presented at the Instant they were notionally linked.

The representation of the most recently presented pattern was naturally most profuse. It also had the greatest number of errors, that is, Neurons linked from Pools which had been notionally inactive at the relevant time, and an absence of Neurons bearing the relevant Constellation Labels in the Pools where they should have been found. The representation of Patterns presented earlier in the run was more limited in numbers, because of the timed destruction of Connections, but many of the errors had disappeared, that is, Neurons linked into the wrong Pools had disappeared. The spatial templates provided by these "older memories" were incomplete, but often contained a smaller proportion of errors. There was usually some representation of almost every Input presented. For example, in a printout made at Instant 73, there was adequate representation of the Linked Constellation formed at Instant 28.

B. Recognition by spatial template system

When a Pattern is presented we expect a matching Linked Constellation to be formed. When the same pattern is presented later we expect this Constellation to be energised to a degree greater than any other Constellation. This was indeed found to be the case. In that sense, the model could be said to be 'recognising the input' in terms of previously presented patterns.

Instant	agg. excit.	Input patt.		No. of Linked Neurons			Mean excit.	
66.	31329.	52	0.0	41.0		0.0	764.	
67.	20480.	2	0.0	IRP 25.0		0.0	819.	
68.	55810.	2	0.0	IRP 68.0		0.0	967.	
69.	17656.	2	0.0	IRP 31.0		0.0	569.	
70.	0.	2	0.0	IRP 0.0		0.0	0.	
71.	0.	54	0.0	0.0		0.0	0.	
72.	55862.	56	0.0	106.0		0.0	527.	
73.	33464.	37	0.0	94.0		0.0	356.	
74.	22056.	33	0.0	48.0		0.0	472.	
75.	33210.	47	0.0	CURRENT 54.0		0.0	615.	
76.	42172.	49	0.0	INPUT 52.0	→	0.0	809.	*
77.	80140.	51	0.0	60.0		0.0	1335.	
78.	51170.	53	0.0	60.0		0.0	851.	
79.	0.	53	0.0	0.0		0.0	0.	
80.	58432.	1	0.0	IRP 106.0		0.0	551.	
81.	9222.	1	0.0	IRP 29.0		0.0	318.	
82.	0.	1	0.0	IRP 0.0		0.0	0.	
83.	0.	54	0.0	0.0		0.0	0.0	
84.	43680.	56	0.0	96.0		0.0	455.0	
85.	31640.	32	0.0	70.0		0.0	452.0	
86.	32016.	38	0.0	69.0		0.0	464.0	
87.	31240.	23	0.0	55.0		0.0	568.	
88.	57614.	46	0.0	86.0	PATTERN	0.0	669.0	
89.	43762.	48	0.0	64.0	IN	0.0	683.	!
90.	58268.	50	0.0	58.0	INPUT AT INST	0.0	1004.	.
91.	48930.	52	0.0	60.0	122.	0.0	815.	
92.	22180.	52	0.0	28.0		0.0	792.	
93.	129660.	2	0.0	IRP 128.0		0.0	804.	
94.	15469.	2	0.0	IRP 54.0		0.0	454.	
95.	0.	2	0.0	IRP 0.0		0.0	0.	
96.	0.	54	0.0	0.0		0.0	0.	
97.	93639.	56	0.0	147.0		0.0	657.	
98.	37765.	36	0.0	91.0		0.0	415.	
99.	22230.	23	0.0	59.0		0.0	571.	
100.	70564.	37	0.0	106.0		0.0	663.	
101.	55 14.	47	0.0	CURRENT 61.0		0.0	574.	
102.	89838.	49	0.0	INPUT 125.0	→	0.0	718.	*
103.	51792.	51	0.0	83.0		0.0	624.	
104.	29784.	52	0.0	51.0		0.0	584.	
105.	21766.	53	0.0	27.0		0.0	769.	
106.	78818.	1	0.0	81.0		0.0	973.	
107.	37127.	1	0.0	37.0		0.0	1009.	
108.	18468.	1	0.0	25.0		0.0	738.	
109.	0.	54	0.0	0.0		0.0	0.	
110.	69762.	56	0.0	154.0		0.0	453.	
111.	29880.	33	0.0	83.0		0.0	360.	
112.	43920.	38	0.0	90.0		0.0	488.	
113.	78971.	46	0.0	CURRENT 108.0		0.0	731.	
114.	136294.	49	0.0	INPUT 191.0	→	0.0	818.	*
115.	78273.	51	0.0	112.0		0.0	698.	
116.	35100.	53	0.0	88.0		0.0	625.	
117.	36493.	53	0.0	63.0		0.0	611.	
118.	45552.	54	0.0	73.0		0.0	624.	
119.	105168.	54	0.0	132.0		0.0	795.	
120.	148383.	37	0.0	181.0		0.0	819.	
121.	135754.	33	0.0	148.0		0.0	917.	
122.	185146.	50	0.0	144.0	CURRENT INSTANT	0.0	1285.	
123.	111527.	49	0.0	133.0	⇐	0.0	838.	*

Figure 8.1 State of Excitation in Constellations at Instant 123. Note effects of temporal sequence and Arousal in Constellations registering IRPs.

It is necessary to anticipate somewhat and say that these predictions were fully confirmed in later versions of the model. They may be seen with clarity in the outputs of the separate simulation which demonstrated the mediation of Linked Constellations by orthodox synaptic connections, (see Figures 16.1 and 16.2, Chapter 16). In those Figures it may be seen that the Neurons belonging to the Constellation which matches the Input Pattern are at the highest level in the system, with the exception of those just formed by the current Input. Similar results were confirmed from the numerous runs reported. The Linked Constellations showing the maximum Mean Excitation Counts were in fact those formed at Instants when the same Input Pattern was currently in the Input.

In the monitor outputs of the initial simulations, this effect occurred but was seriously complicated by the temporal sequential and arousal effects. Figure 8.1 reproduces part of a table in a monitoring output. Each line in the table summarises a Constellation and reports (a) the sum of the Excitation Counts of all the Integrating Neurons identified as belonging to that Constellation; (b) the number of Integrating Neurons belonging to that Constellation; and (c) the arithmetic mean of the Excitation concerned.

Examination of many such monitor printouts disclosed that the matching effect was only clearly seen at certain phases of the experimental runs, notably when the model went step by step through the sequence of Input Patterns leading to the Instant of presentation of the IRP. Because of the build up of Excitation in the system at that point, the Input Patterns which arrived after the IRP, were unable to influence the main pattern of distribution of Excitation.

In Figure 8.1 the position is shown when a Neutral Pattern is in the Input following a presentation of the IRP. The Linked Constellation matching the Input Pattern responds by rising to an Excitation level slightly higher than its neighbours, but does not approach the high levels still remaining in the Constellation representing the consummatory climax of the sequence.

The full implication of these complications were not appreciated at the time. They became of considerable importance in the later development of the model where they will be discussed more fully. Here it may be noted, that where a familiar sequence is repeated, the Excitations of the Linked Constellations representing that sequence tend to dominate an individual input which does not match one of the patterns included in the sequence.

II Testing the initial simulations: design of investigatory experiments

The runs were designed to test the general account of brain organisation given in Section IV of Chapter 6. As mentioned earlier, a "run" in this context means a run of the program in the computer during which the performance of the model could be compared with prediction, i.e. an investigatory experiment. Each run simulated a simple classical conditioning or cue learning experiment. The model was expected to emit a notional response when presented with a specific stimulus pattern. The aim was the detection of a response when a cue stimulus or Conditioned Stimulus (which during training had preceded the Unconditioned Stimulus) was presented without being followed by the UCS (specific stimulus pattern).

For the purposes of the tests, the Motor Response of the model was indicated by the Efferent Fibre which carried the greatest Spike Frequency. For convenience the Efferent Fibres were numbered 1 to 4, so the Motor Response could be a number, 1 to 4, or a combination of numbers. The Motor Response to be conditioned was that obtainable by presenting to the model an Innately Recognised Pattern (IRP). It will be recalled that an IRP was an Input Pattern which matched the Sensory portion of a Preset Constellation written into the Sub-cortex. The Motor Response to be conditioned was therefore determined by which Efferent Columns of the Sub-cortex contained Neurons linked into the Preset Constellation concerned. In some runs the Preset Constellations were written so that the response of two Efferent Fibres was elicited when its corresponding IRP appeared in the Input. In many early runs, in order to produce a clear demonstration, the Responses consisted of either Efferent Fibres 1 and 2 or Efferent Fibres 3 and 4, designated respectively Action A and Action B.

In addition to the designation of Motor Response as a particular Fibre or Fibres carrying maximum Spike Discharge, the actual value of that Discharge (Efferent Spike Frequency) gave a measure of the strength of the Response.

A preprogrammed sequence of Input Patterns was prepared for each run. A Cue Sequence (or, less formally, a Cue) simulated a sensory event. It consisted of a set of three or four successively presented Input Patterns in a fixed order. These patterns were selected from an

A. TRAINING SESSION

N Trials

N = 7 in most experiments

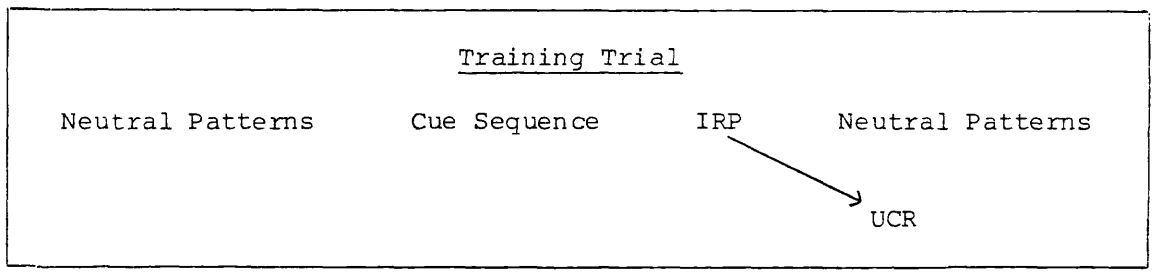

Followed by

B. TEST SESSION

M Trials

M = 3 in most experiments

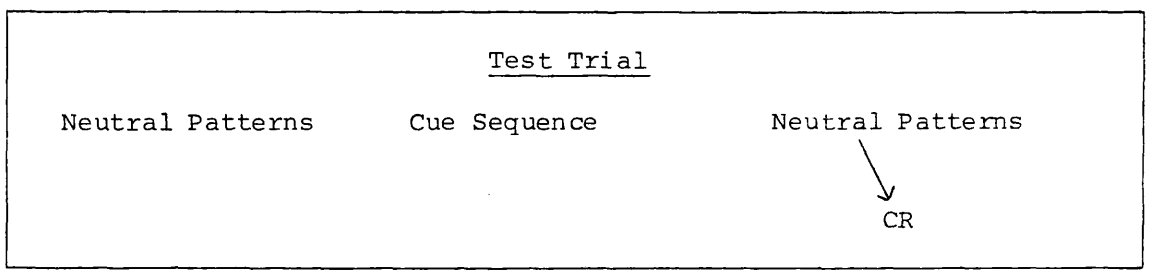

Figure 8.2 General design of computer runs to demonstrate the behaviour of the model in a situation analogous to a Classical Conditioning situation. IRP - Innately recognised Pattern which elicited a Motor Response from the model analogous to an unconditioned Response. Cue Sequence - a set of four indifferent Input Patterns in given order analogous to the Conditioned Stimulus. CR - output spikes analogous to a Conditioned response.

Input Pattern Library as being indifferent, that is they were shown to produce no appreciable Motor Response before training. A further number of indifferent patterns were selected to be used as Neutral or inter-trial patterns. These represented the experimental environment in the absence of a special stimulus.

A Training Trial consisted of a few Neutral Patterns followed by a Cue Sequence, followed in turn by an IRP (Innately Recognised Pattern or Releaser).

A Test Trial generally consisted of the Neutral Patterns followed by a Cue Sequence only, i.e. without the IRP. The preprogrammed sequence of patterns for the whole run consisted generally of a Training Session, consisting of a number of Training Trials, and a Test Session, consisting of Test Trials, i.e.

A. Training Session

TRAINING TRIALS each comprising:

Neutral Patterns - Cue Sequence - IRP - Neutral Patterns

B. Test Session

TEST TRIALS each comprising

Neutral Patterns - Cue Sequence - Neutral Patterns.

This design corresponded in many ways to a Classical Conditioning situation. Like the Conditioned Stimulus, the Cue Sequence does not eleicit a response prior to its pairing with a special kind of stimulus. Like the UCS, the IRP is a special kind of notional sensory pattern which elicits a response in the absence of prior experience. The choice of this design is not meant to imply that the model explains learning in terms of Classically Conditioned Responses. The explanation of learning in the model is the formation of Linked Constellations. What the model suggests is that in conditions analogous to the Classical Conditioning situation, the theory of neural integration would produce phenomena comparable to those seen in the Classical Conditioning situation.

III Demonstration of notional classical conditioning

The experimental run illustrated in Figure 8.3 was one of a number which approximated most closely to a straightforward classical conditioning experiment. In the particular run illustrated, a single

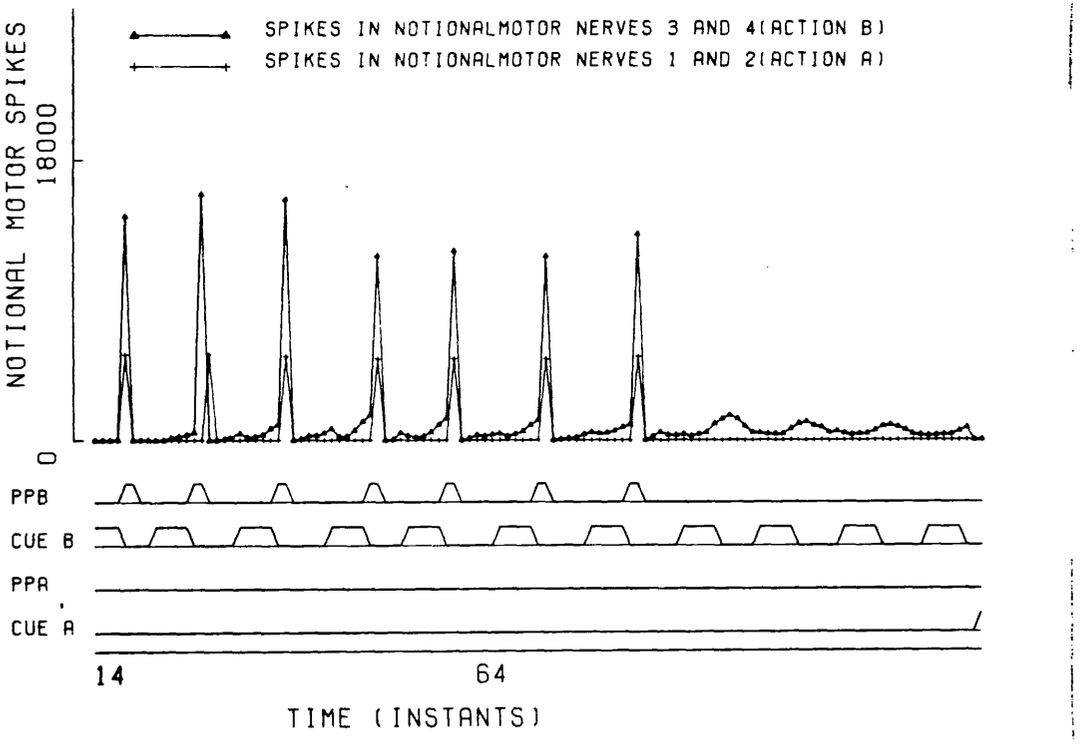

Figure 8.3 Example of an experimental run analogous to a classical conditioning situation. Lines marked Cue A and PPB show the times of presentation of the Cue Sequence (CS) and the Innately Recognised Pattern (UCS). In the Test Trials only the Cue Sequences was presented. The curve (Action B) shows the motor output Spikes produced by the model for each Instant of the entire run. The output referred to as Action B was relevant to the Innately Recognised Pattern presented (PPB). Action A was irrelevant (see Text).

Cue Sequence, of 4 Input Pattern, termed Cue B, was used throughout the run. In seven Training Trials it was followed by a particular Innately Recognised Pattern. In three subsequent Test Trials Cue B was presented without being followed by an IRP. The general experimental design is illustrated in Figure 8.2.

The behaviour of the model is indicated in the Figure by curves showing the number of Spikes in each of the two sets of Response Fibres. The curve labelled Action A shows the sum of the Spike Frequencies in Efferent Fibres 1 and 2. That labelled Action B shows the sum of the Spike Frequencies in Fibres 3 and 4. The abscissa represents successive Instant of the run. The time of presentations of the IRP is shown by deflections of the line marked PPB. The Line marked PPA was reserved for indicating the presentation of another IRP, but as only one IRP was used in this run, it was not needed. Generally PPA was used to refer to the IRP which elicited an unconditioned response from Fibres 1 and 2, i.e. Action A, and PPB for the IRP which elicited a response from Fibres 3 and 4, i.e. Action B. Cue A was used to refer to the Cue Sequence which was presented prior to the presentation of PPA, i.e. the CS for PPA; similarly Cue B was the CS for PPB. The presentation of Cue B is shown in Figures 8.3 and 8.4 by deflections in the appropriately labelled line. In this run and in related experimental runs sharing the design shown in Figure 8.2, only one IRP and one Cue were used.

It will be noted that during Training (Figure 8.3) each presentation of PPB (the particular IRP used in this run) produces a sharp peak in the relevant Motor Response output (Action B). This signifies the directly elicited Motor Response analogous to the Unconditioned Response.

During Testing, a clear peak in the curve for Action B follows each of the three presentations of Cue B, notwithstanding that during the Test Trials no IRP was presented. These peaks represent Spike Frequencies occurring in the relevant Fibres 3 and 4, and are analogous to the Conditioned Respone.

Latency of Response. In the first Training Trial, Spikes only appear in Action B upon actual presentation of PPB. In the second Training Trial, some Spikes occur during the presentation of the Cue Sequence, i.e. before the presentation of PPB. In subsequent Trials these 'anticipatory' Spikes increase in number and occur earlier from Trial to

Trial. The latency of the notional Conditioned Response may be said to diminish as the number of Training Trials increases. However, this statement really only applies to the relatively small 'anticipatory' rises in the curve and not the main peaks. The latter are clearly seen as occurring only when an IRP, in this case PPB, is actually in the Input. The response curve during Training indicates an overlaying of the 'conditioned' and 'unconditioned' responses. Action A, the irrelevant response. Although Action A was not intentionally utilised in the run concerned (Figure 8.3) small peaks occur in the curve for Action A whenever, during Training, a large peak for Action B occurs. This effect is attributable to the arousal system. Although Action A is not implicated in the experimental presentations of the run concerned, the Preset Constellation capable of emitting Action A was present in the system. The presentation of PPB elicited a massive output in Action B. It may be recalled that two arousal parameters, Diffuse Volley Value and Discriminative Volley Value, were in early versions of the model, both adjusted by reference to Spike Frequencies across the Efferent Fibres. It may be assumed that the small peaks seen in Action A were attributable to the increase in Diffuse Volley Value, which affected both Responses. It was effects of this kind which disclosed the problems associated with the use of Diffuse Volley Value as an arousal parameter. It may be noted that the disturbances in Action A occur only at the actual Instants when Action B has been elicited. Elsewhere, and in particular throughout the entire Test Session, Action A remains at the baseline.

In a number of different runs which followed the design indicated in Figure 8.2, PPA was substituted for PPB. A different set of Input Patterns was used as a Cue Sequence preceding the presentation of PPA, and was accordingly termed Cue A. PPA naturally elicited Action A during Training Trials and Cue A elicited a notional Conditioned Response of Action A during Test Trials. Further testing of the program was carried out by using different sets of Input Patterns as Cue A or Cue B, as the case may have been in the particular run. In all these variations the basic kind of performance achieved was similar to that illustrated in Figure 8.3.

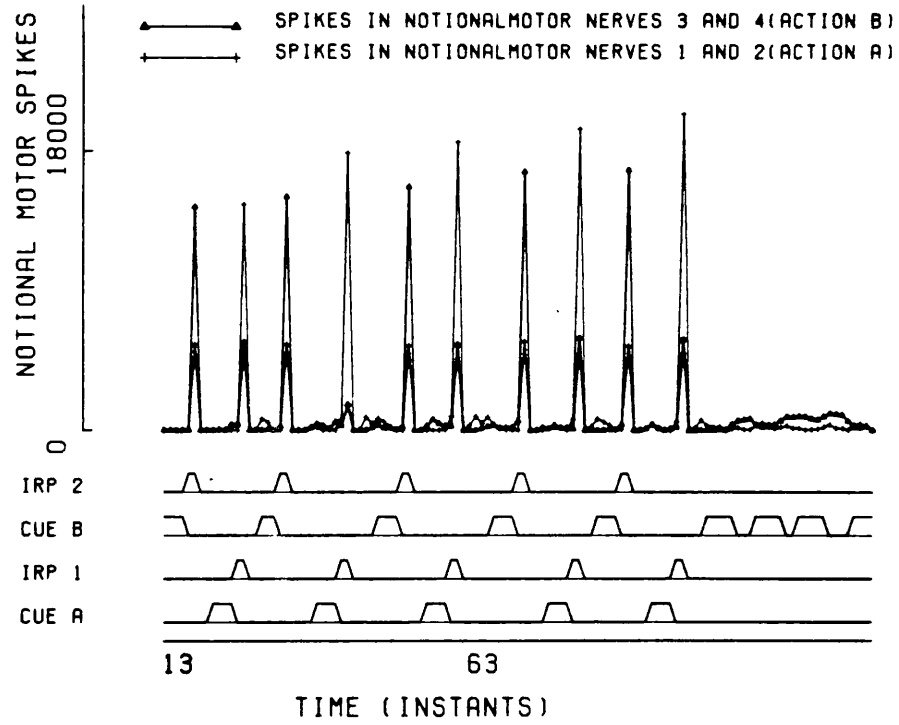

Figure 8.4 (a) An example of a run in which the model was required to discriminate between two Cue Sequences (Cue A and Cue B) each of which were followed during Training by a different Innately Recognised Pattern (UCS), IRP 1 and IRP 2. On alternate Trials Cue A was followed by IRP 1, and on the other Trials, Cue B by IRP 2. In all the Test Trials Cue B alone (i.e. without the UCS) was presented. A magnification of the curves for the Test Trials may be seen in Figure 8.4 (b).

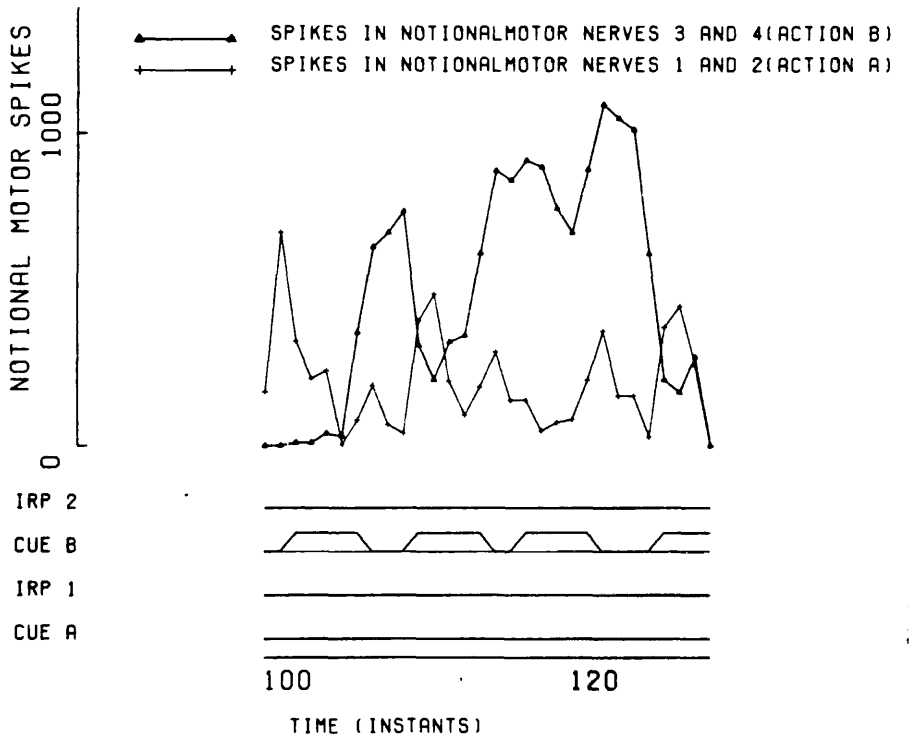

Figure 8.4 (b) An enlargement of the portion of the curve relating to the Test Trials of the experimental run shown in Figure 8.4 (a). During Training, on alternate Trials Cue A (CS) was followed by IRP 1 (UCS) and Cue B (CS) by IRP 2 (UCS). During the Test Trials, shown above, only Cue B was presented "unreinforced". The motor output Spikes in Action B, which was the output in which Spikes were elicited during Training by a presentation of IRP 2, rose above the output in Action A, which was the output in which during Training motor Spikes were elicited by a presentation of IRP 1.

IV Discrimination of Cues

Figure 8.4 illustrates one of many experimental runs in which two different Cue Sequences and two different IRPs were used in the same run. The general design is indicated in Figure 8.5. The design of this and related computer investigations does not correspond with arrangements generally seen in classical conditioning experiments, although some experiments have not been entirely dissimilar (e.g. Varga and Pressman, 1963). In most classical conditioning experiments only a single UCS is conditioned. The system of "double-conditioning" was adopted here to make sure that the model could not only discriminate two different Cue Sequences but could produce a discriminate "anticipatory" response to each of them when each had been consistently followed by a different IRP.

As may be seen from Figure 8.4, there were ten Training Trials. On odd-numbered Trials Cue B was followed by IRP 2; on even-numbered Trials Cue A was followed by IRP 1. The Instants of presentation of Cue Sequences and IRPs can be seen from the deflections in the appropriately labelled lines in the Figure. IRP 1 was the Innately Recognised Pattern which caused the model to emit Action A (i.e. an 'UCR'); IRP 2 was the IRP which caused the model to emit Action B (the other 'UCR'). Cue A and Cue B were two different Cue Sequences, each made up of four indifferent Patterns.

Figure 8.4 (b) is a magnification of that part of the curve shown in Figure 8.4 (a) which relates to the Test Trials only. In the particular run illustrated, Cue B was presented at each of the three Test Trials, a fourth Test Trial being interrupted by expiration of computer time. It will be seen that the curve for Action B, i.e. the CR to the relevant CS, Cue B, has three clear peaks, each coming at the end of the presentations of the Cue Sequence.

In this run, the curve for Action B does not return to base line between Presentations of the Cue, as it did in the run illustrated in Figure 8.3. This was because fewer Neutral Patterns could be interposed between presentations of the Cue if three Test Trials were to be accommodated within the available computer time. After the third Test Trial a larger number of Neutral Patterns were presented and it will be seen that the curve falls.

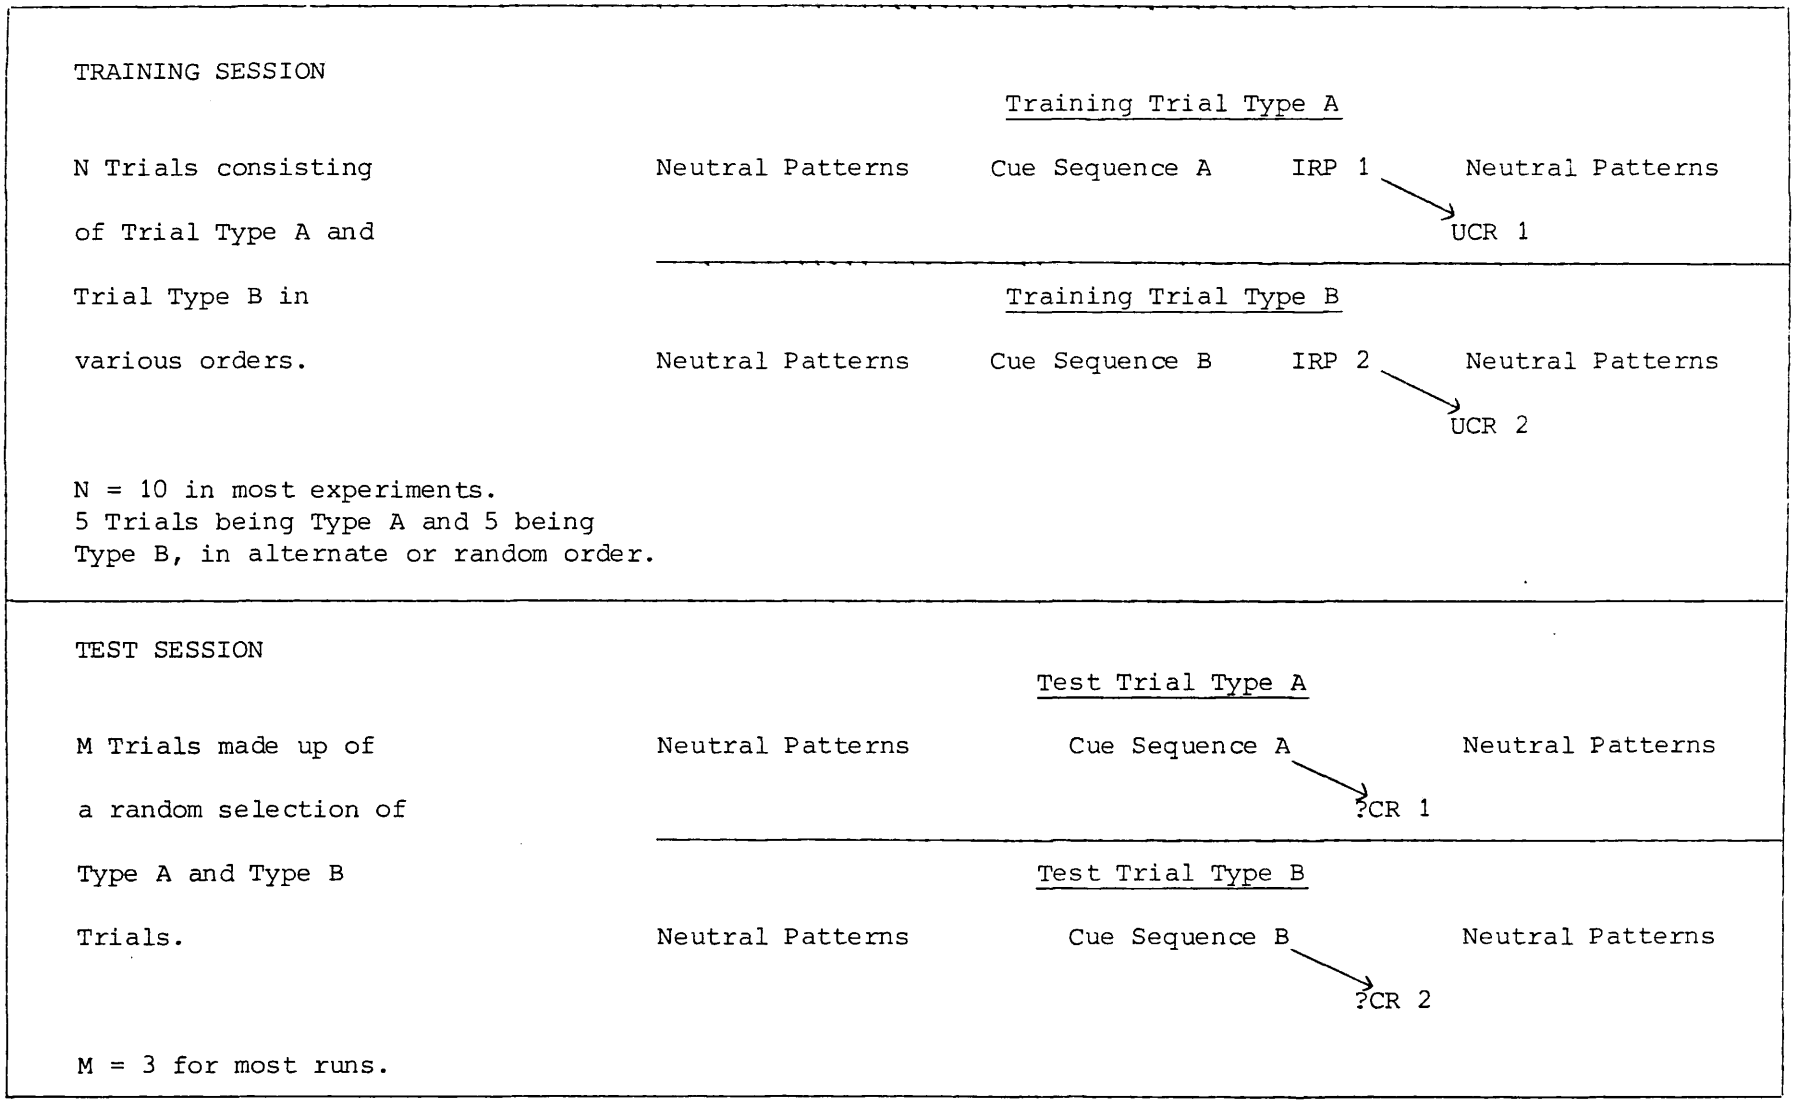

Figure 8.5 The general design of experimental runs in a double-conditioning or Cue discrimination Test. IRP 1 and IRP 2 were two different Innately Recognised Patterns which elicited different Motor Responses from the model (UCR 1 and UCR 2). Cue Sequence A and Cue Sequence B were each a set of four indifferent Input Patterns in a given order. All the Patterns were different.

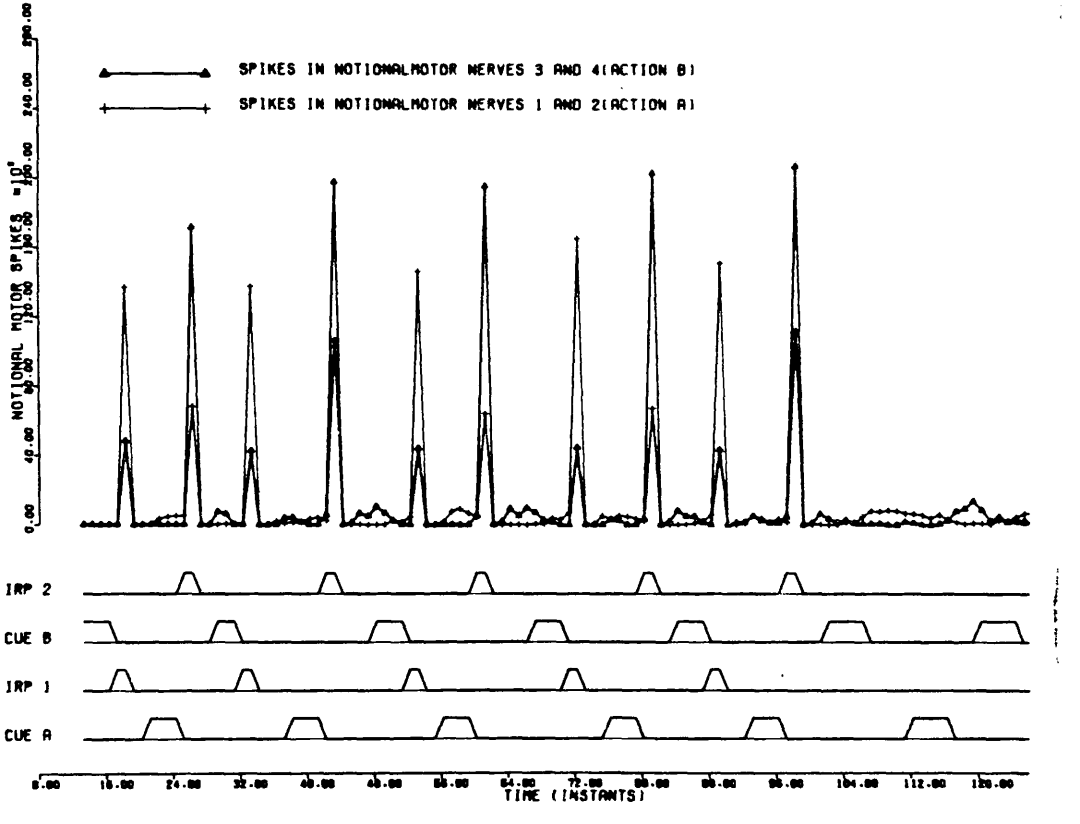

Figure 8.6 (a) A similar experimental run to that shown in Figure 8.4 except that the pairing during Training was reversed, i.e. Cue A was now paired with IRP 2, Cue B with IRP 1. At the Test Trials the Cues were presented in alternation without "reinforcement". Only one Test Trial is shown in the above Figure but in all three Test Trials (see Figure 8.6 (b)) the motor output varied appropriately as the different Cues were presented.

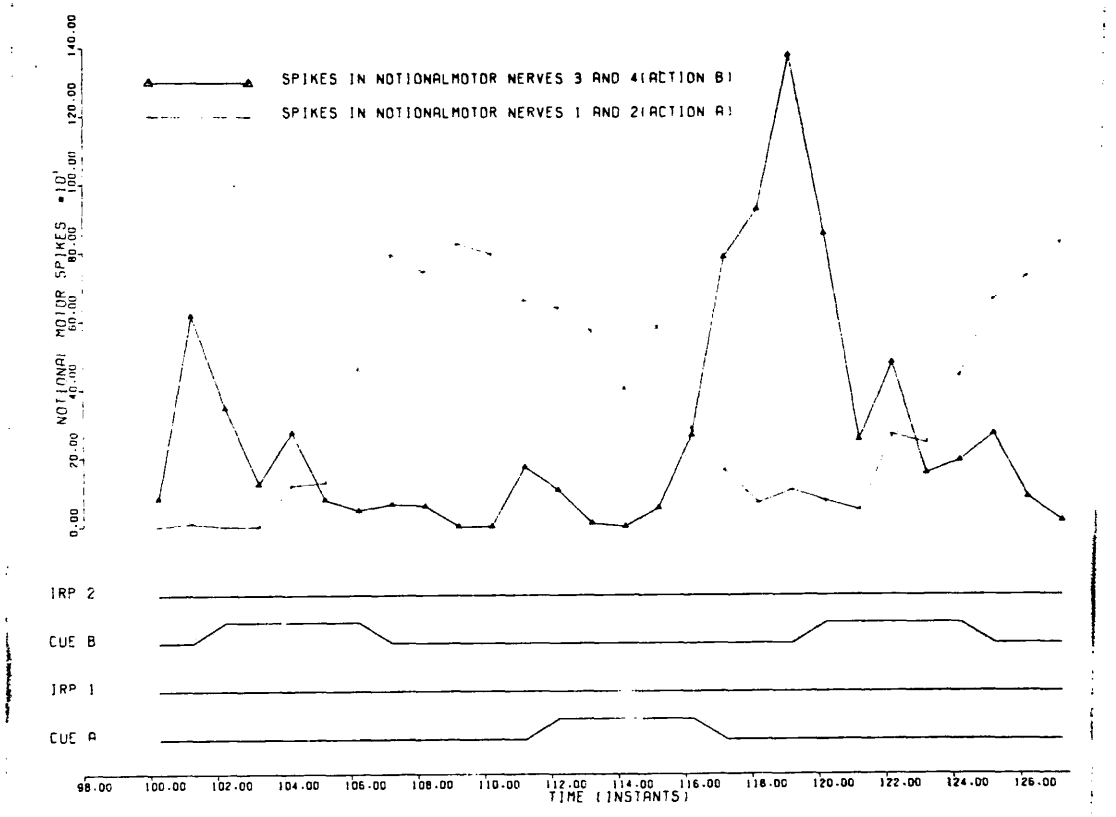

Figure 8.6 (b) An enlargement of the curve for the Test Trials of the experimental run illustrated in Figure 8.6 (a). The discriminated response to each of the two Cues may be seen in the variations of the curves representing each of the two motor outputs, one relevant to the CR to Cue A and one to the CR to Cue B.

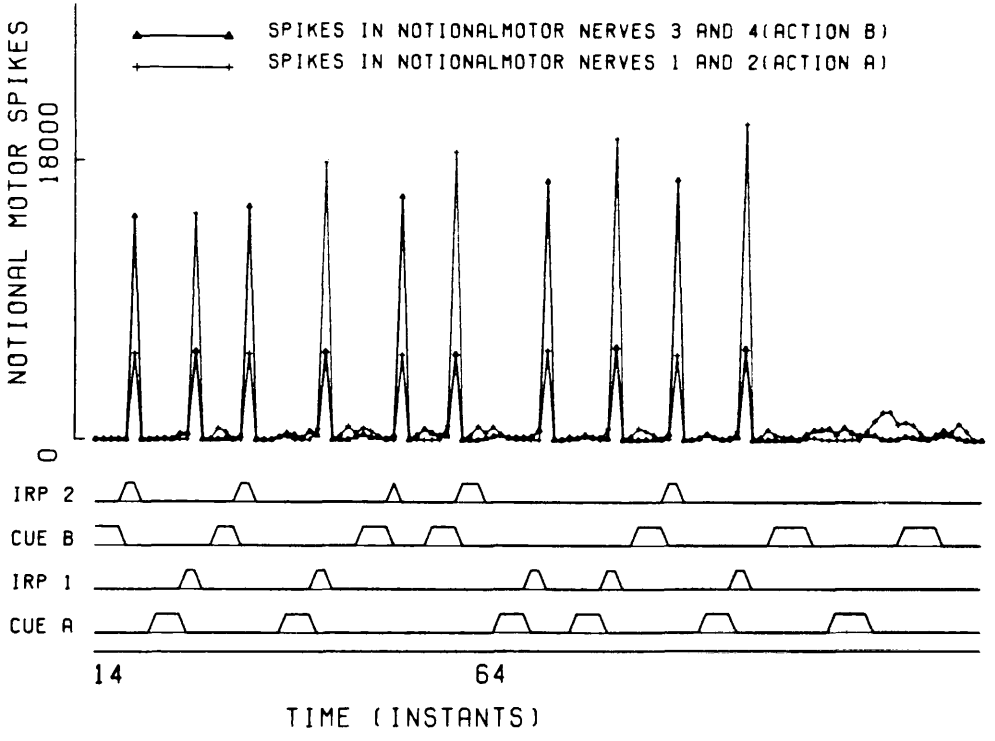

Figure 8.7 Another experimental run in which two Cue Sequences and two different Innately Recognised Patterns (PPA and PPB) were used to "reinforce" them.

After each presentation of the Cue, the curve rises and then falls. It then begins to rise again. It will be noted that it does this in the middle of the next presentation of the Cue Sequence, demonstrating an 'anticipatory' effect similar to that demonstrated in Figure 8.3. In this run, as compared with that shown in Figure 8.3, the curve for Action A does not remain at the baseline during Test Trials utilising Cue B. Nevertheless, the response of Action B is clearly distinguishable.

To test that the Responses shown in Figure 8.4 were not an artefact resulting from the particular patterns used for Cue A and Cue B, these patterns were varied in different experimental runs, different selections from the Pattern Library being used altogether. In some runs completely different patterns were used as "pseudo-cue sequences" in the Test Trials, to show that no analogical CR appeared.

In the run shown in Figure 8.6 (a), the experimental run shown in Figure 8.4 was repeated except that the pairing in Training Trials was reversed, i.e. Cue A was paired with IRP 2, Cue B with IRP 1. All three Test Trials produced an analogous CR appropriate to the new pairing. (See Figure 8.6 (b).

Many runs further tested the operation of the model and demonstrated that the notional Conditioned Responses were in fact elicited by the pairing of Cue Sequence and IRP and were not due to the variability of the system, chance or artefact. For example, Figure 8.7 shows an experimental run in which in the Test Trials a presentation of Cue A is interposed between two presentations of Cue B. The output of Action A and Action B varies appropriately. (A magnifiction of the curve for the Test Trials alone may be seen in Figure 9.1).

In the runs described Cue A followed by IRP 1 was presented alternately with Cue B followed by IRP 2 during Training Trials. The order of presentation was varied in many experimental runs. For example, that shown in Figure 8.7 has the order broken up somewhat. In Trials numbers 5 and 6 Cue B was presented followed by IRP 2, whilst in Trials 7 and 8 Cue A was followed by IRP 1. The appropriate analogous CR's were still obtainable. The order of presentation was, however, found to have some effect. In early versions of the model, the first Training Trial was found to have an undue effect. This was corrected by a modification of the model (see next chapter). The last Training Trial,

i.e. that immediately preceding the Test Trials was also found to have an undue effect, which may be seen, for example, in Figure 8.7. These aspects of the model, which were interpreted as resulting from the decay and linking control system, were considered to be analogical to possible attributes of actual nervous systems.

V Conclusions

(1) The Linked Constellations responded to input patterns appropriately.

(2) When a particular sequence of Input Patterns consistently preceded a Special Input Pattern (and IRP) which produced an inbuilt response from the model, the Linked Constellations recording the events could be trained to induce similar Motor Responses following the appropriate input patterns, i.e. a form of conditioned responding was demonstrated.

(3) The system was able to distinguish two sequences of input patterns which had been followed by different specific eliciting patterns (IRPs), by producing appropriate different responses to each .

These conclusions confirm the main predictions of the model as a possible mediation of a neural template hypothesis in the form of Linked Constellations. The performance of the initial version of the model was examined further, but before reporting those investigations it is necessary to report some modifications which were introduced in the course of developing the initial set of simulations. These are dealt with in the next chapter.

CHAPTER 9

SOME FURTHER ASPECTS OF THE FIRST SET OF SIMULATIONS

I Some modifications introduced in the course of simulation

As explained at the outset of the thesis, the simulations were used as a tool of development, not as the exposition of a fully developed model. In fact the first trial simulations were very simple indeed, and the results of computer experiments reported in the previous chapter were in relatively advanced versions in which the general structure described earlier had been developed. Most of the modifications introduced into the initial set of simulations have been discussed theoretically, and the following is mainly a matter of bringing the reader up to date before reporting further results.

A. Balanced systems: receptor system

The theoretical implications of sensory patterns of activity involving different numbers of active fibres were discussed at some length in Chapter 6, Section I. It will be recalled that it was there argued that patterns with a large number of active fibres would reactivate a large number of Linked Constellations which represented only fragments of the Input Pattern. Moreover, an Input Pattern with a large number of active Fibres would send more Excitation into the system and so have an arousing effect. These two effects would disturb the retreival of information from the system.

In the first versions of the simulations Input Patterns consisted of permutations of 16 bits. Virtually the model did not work in this form. Some improvement was made by selecting patterns at random with varying numbers of notionally active fibres (i.e. 1's). However, to obtain a high level of performance, it was necessary to restrict all the Input Patterns in the Library to those having the same number of active Fibres, i.e. 1's.

It was considered somewhat unsatisfactory to demonstrate the model in

a form subject to this restriction. Accordingly the simulation was extended to permit the Input Patterns to be increased from 16 to 32 bits. In this extended form, the second 16 bits were the 1 complements of the first 16 bits. The whole model had to be expanded, that is, the number of Neuronal Pools doubled (to 32), so that each bit of the Input Pattern still had its corresponding Neuronal Pool. The system then represented in principle a set of receptor elements with the property that when one receptor was excited by a certain stimulation, there was another which had formerly been in an excited state in the absence of such stimulation, which was now suppressed. The theoretical implications of this hypothesis have already been discussed in Chapter 6.

The remainder of the results illustrated were all in this extended version.

B. Balanced systems: motor system

It was suggested in Chapter 6 that motor systems might be organised in some way analogical with the balanced input aspect of the receptor system. The reciprocal aspects of effector systems may offer certain explanatory possibilities, as indeed they have to others, notably Konorski (1967), and it was wished to try to introduce such aspects at the earliest moment. The theoretical aspects of such systems will be discussed later. Here it is sufficient to say that animals learn not to do things just as they learn to do things. It seems plausible that there would be a representation of response suppression as well as a representation of response activity.

The representation of this idea in the initial set of simulations was rudimentary. It differed from version to version. These versions conformed however to a general principle. In the versions having a balanced effector system representation, the number of Motor Neuronal Pools was increased from 4 to 8 in the Cortical Array. As will be explained below, in some of the versions the Sub-cortical representation was also increased from 4 to 8 Neuronal Pools, and in some versions the Sub-cortical representation remained at 4 Neuronal Pools.

In the versions in which the Sub-cortical representation had been increased to eight Neuronal Pools, the Preset Constellations in the Sub-cortical Array were then rewritten so as to include the representation of 8 Motor bits, the second four bits being the 1 complements

of the first four (in analogy with the balanced receptor system). This is not of course the only possible interpretation of a balanced effector system, and indeed others were investigated (see below, and later chapters). In the system just described, if one of the Preset Constellations had Fibres 1 and 2 excited by its sensory pattern, it also had Fibres 7 and 8 excited. Interpreting the second four Motor Neuronal Pools as Response Suppression Pools, this meant that if the model notionally performed the action represented by a Discharge in Efferent Fibres 1 and 2, it also emitted Discharges indicating that it was suppressing the actions represented in Fibres 3 and 4.

There is an alternative manner of implementing a balanced effector system, which was not investigated in the model. This is the implementation of response suppression based upon the active neuronal inhibition of fibres representing the emission of the response. As mentioned in Chapter 6, Konorski (1967) discusses these two alternatives in considerable depth. The present author follows Konorski in believing that Response Suppression is the function of a separate excitatory system which actively signals the suppression of responses. These matters do not play a significant part in the initial set of simulations, but were of more importance in subsequent extensions of the model.

The arrangement just described worked, but not well. An alternative method which was tried was a reciprocal system in which whenever a fibre to one of the first four Neuronal Pools was active, its corresponding fibre to the second four Neuronal Pools was inactive, and vice versa. In this version it was only necessary to have four Neuronal Pools in the Motor Division of the Sub-cortical Array, because the eight Afferent Fibres could be controlled by the pattern of output across the four Sub-cortical Efferent Fibres.

A number of other similar arrangements were experimented with. However, it was appreciated that they were all equally unsatisfactory as a representation, however simple, of actual neural arrangements. One of the main troubles was that, since the model was not doing anything most of the time, the fibres to the suppression Neuronal Pools were active most of the time. This soon exhausted their stock of available unlinked Neurons. For another thing, the suppression representation was being applied to the kind of directly elicited response analogical

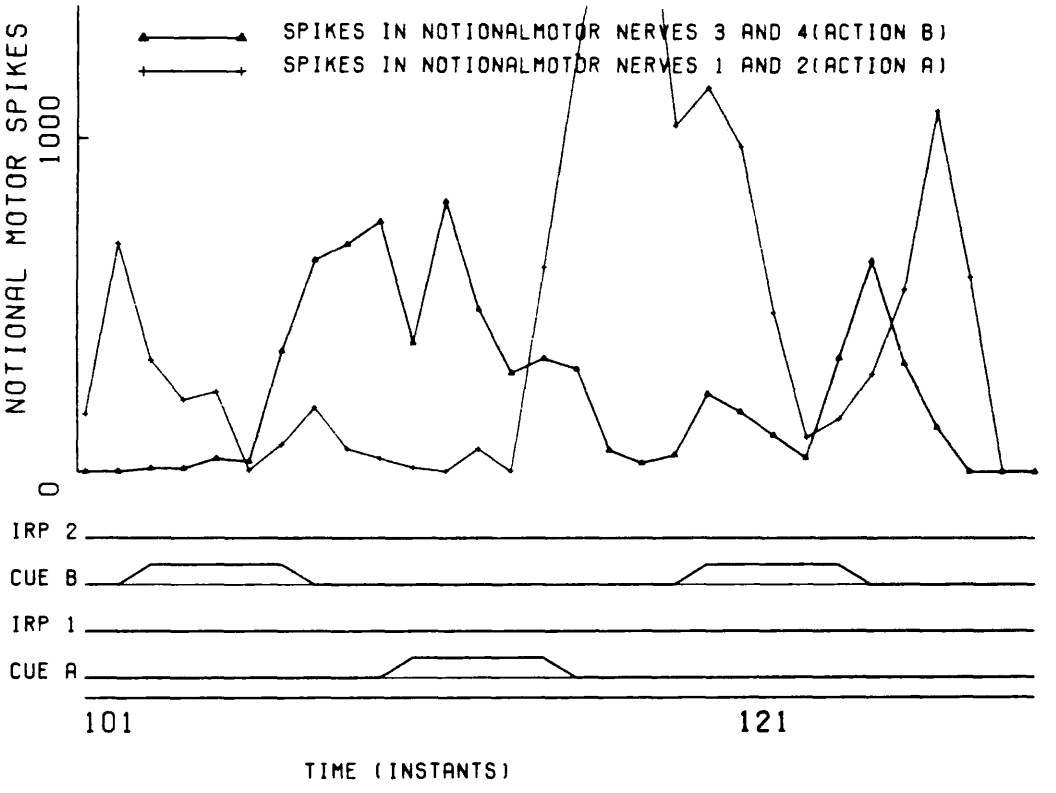

Figure 9.1 Conditioned "arousal". The Test Trials of an experimental run in which during Training two different Cue Sequences had been paired with different Innately Recognised Patterns. The output curves show an appropriate variation according to the particular Cue presented without "reinforcement". In this run, the exaggerated and somewhat erratic effect of the feed-back arousal system can be seen in the Test Trials. It will be noted that towards the end of the Test Trials, the Output in Action A is building up to an explosive inflationary loop. This loop began before the first Test Trial as an aftermath of the Training Session when Innately Recognised Patterns were presented and directly elicited responses (UCR's) were elicited.

to a UCS, and it is doubtful whether the animal learns not to do such things in quite the same way as it learns not to move in a certain situation. These initial attempts were however found useful when the later extended versions of the model were attempted. They are in any case reported for the sake of completeness.

C. Conditioned Arousal system

In the runs so far illustrated the effect of the arousal system can only be seen when an Innately Recognised Pattern (innate releasing pattern) is actually in the Input. It seemed to be a good idea to show that an increase of arousal could be made to occur during the presentation of the Cue Sequence and in anticipation of the IRP. Such an arousal might be termed a conditioned arousal. This was feasible because, as noted in connection with the computer experiments already reported, the output of the relevant Efferent Fibre began to rise before the actual presentation of the Innately Recognised Pattern. It will be recalled that the arousal parameters varied according to the output across the Efferent Fibres. Accordingly arousal should have resulted from the increase in Motor Nerve activity immediately preceding an anticipated innate releasing pattern.

In the runs so far reported, an arousal effect can be clearly seen when the IRP is actually in the Input. This results from the large Motor output initiated by the reactivation of the Preset Constellation at Sub-cortical level. The operation of arousal as an effect of the Conditioned Response was less definite. An attempt was made to increase the effect by magnifying the relationship between Motor Discharge and arousal. However this made the output during an Unconditioned Response too great for the model to accommodate reasonably. The system was then made non-linear. Small variations in Discharge of Motor Nerves were made to have a larger proportionate effect than large ones, i.e. arousal levels were a logarithmic function of Motor Nerve activity. The result of these modifications was to make performance of the model unstable. The trouble was a self-fuelling loop. A small variation of Motor Output caused an increase of arousal; this magnified the next variation of Motor Output, which in turn magnified the next. Nevertheless a number of runs gave a good analogous Conditioned Response and also gave evidence that the system was responding at a higher level of activity in anticipation of the Innately Recognised Pattern (Figure 9.1).

D. Notional maturation of Integrating Neurons

It was mentioned that in early versions of the model, the first Training Trials had an undue influence on later performance. This appeared to be because at the outset of the run the system had a large number of unlinked Integrating Neurons.

Such a state of affairs was considered not to be an adequate representation of a neural system, since it implied that a nervous system came into being fully operative. Accordingly it was thought justified to introduce a correcting assumption. At the outset of each run only a proportion of the Integrating Neurons provided was actually made operative. This number was increased from Instant to Instant during the first 15 to 20 Instants of the investigatory run. This expressed the idea that at birth many of the Integrating Neurons would not participate in neuronal integration, and would gradually come into the system as they mature.

The revised system worked well. It reduced the undue effect of early Training Trials (see later in this chapter), and it made later Trials more effective by providing economy of the use of Integrating Neurons to store information at successive Instants.

E. Summary of modifications

During the course of the simulations a number of modifications were introduced. These included a balanced receptor system. This was effective in freeing the model from restrictions upon the notional patterns of stimulation presented to it. Less successful was an extension of the principle to provide a balanced effector system. It worked but failed to add to the efficiency of the model in the manner expected. However, it disclosed a number of characteristics of balanced effector system which proved useful in later extensions of the model.

Experiments were made to increase the effects of the arousal system. This led to instability in performance. The difficulties in solving the consequent problems of control played a large part in the later emergence of a fresh concept of arousal and attention, notably to the introduction of rhythmic mechanisms.

The number of Integrating Neurons in the system was increased progressively over the first Instants of the computer runs. This reduced the undue effects of early Training Trials and assisted in economy of storage.

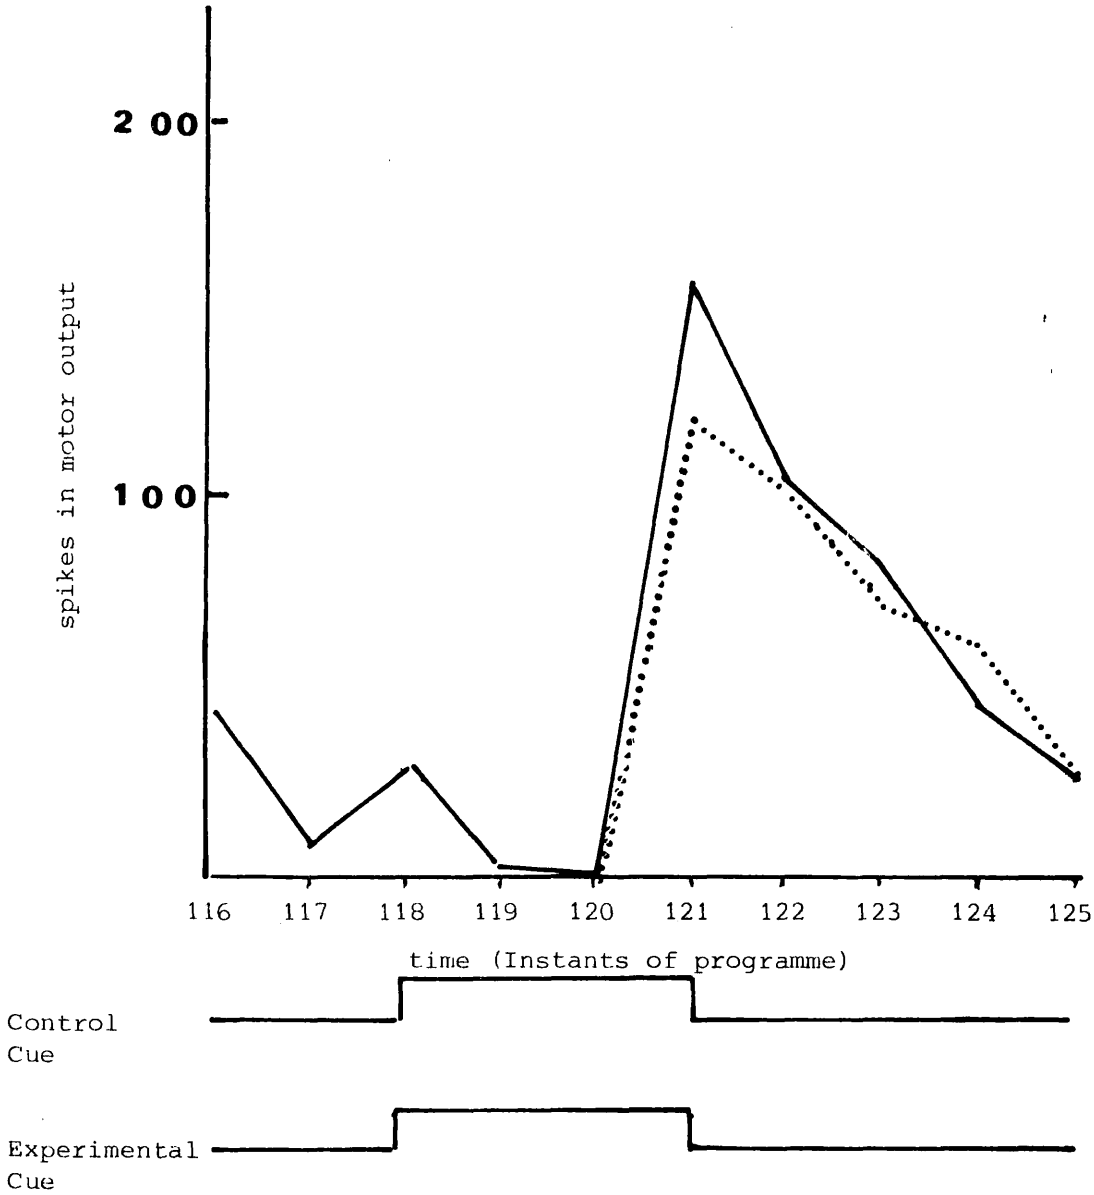

Figure 9.5 Stimulus generalisation and descrimination. The full line and the dotted line show respectively the responses in two different runs to a Cue Sequence presented at a Test Trial without "reinforcement". The runs were identical except for the composition of the Cue Sequences presented at this Test Trial. In one run (shown dotted) the Cue Sequence used in the other run was varied. The last pattern of the four used in the Sequence was varied by one pair of complementary bits from that used in the run shown by the full line.

II Some further results of the initial simulations: Stimulus generalisation

The effect of a stimulus which resembles the conditioned stimulus but differs from it along a dimension such as the pitch of a tone was somewhat loosely simulated by presenting patterns in the Test Trials which differed only slightly from the patterns used as Cue Sequence in the Training Trials.

Figure 9.5 compares the outputs during a Test Trial in two almost identical experimental runs, one being the control for the other. The only difference between the two runs was the Input Pattern presented at Instant 121. In both runs a sequence of Input Patterns was presented over Instants 118 to 121. In the run whose output is shown in the heavy line the sequence presented was identical with Cue Sequence B as used in both runs to precede the IRP which induced an emission of Action B. In the run whose output is shown in the dotted line the sequence presented in the Test Trial illustrated was identical with Cue Sequence B except for one pattern, the last of the sequence, presented at Instant 121. This pattern differed from the last pattern of Cue Sequence B by only one pair of complementary bits, e.g.

1001 0110	1000 0111
1100 0011	1100 0011
0101 1010	0101 1010
0101 1010	0101 1010
(a) Indifferent Input Pattern forming part of Cue Sequence during Training Tests	(b) Example of Input Pattern differing from (a) by one complementary pair of bits

It will be noted that the notional Conditioned Response is definite in both runs. However, the peak reached by the curve showing the output for the experimental run which substituted the slightly different pattern at the end of the Test Cue Sequence does not rise quite so high as that for the run using the true Cue Sequence B. The output in the former reached a Spike Frequency of 123 against a Spike Frequency of 160 induced by the true Cue Sequence B.

These runs demonstrated that the system could respond critically to differences of input as small as the minimum consistency change in the notional pattern of stimulation, i.e. a pair of complementary bits. It was also demonstrated that although the system discriminated between the two similar patterns, it nevertheless produced a response to the slightly different pattern which only differed by a small degree of strength from the response it produced to a complete match. It seems important that a system of neural integration should have this capacity so that useful generalisations can be developed to essentially identical (i.e. behaviourally equivalent) inputs. Nevertheless, if one such pattern should prove to be biologically desirable and the other harmful, it is important for the system to be able to distinguish them. This, the model has been shown to be able to do.

III Competing responses: recency and primacy effects

Some runs investigated the effects of inconsistent pairings between Cue Sequence and IRP. For example, in some Training Trials within the same experimental run, Cue A would be followed by IRP 1 and in others by IRP 2, whilst in other Training Trials Cue B would be followed by IRP 1 and in yet others Cue B would be followed by IRP 2.

This resulted at Test Trials in emissions of Responses in both Action A and Action B to any Cue presented, with one or other dominating. It will have been noted in the experimental runs already reported that when in Test Trials a Cue was presented and a notional Conditioned Response obtained, the curve for the output of the irrelevant Action rarely remained at base level as it did in those runs where only one Cue Sequence and one IRP was employed. In one sense this emission of Spikes in both Action A and Action B merely pointed to inadequacies in the Inhibition system, which should have reduced one output to zero. However, the actual emission of Spike Frequencies in both notional Motor Nerves gave the clue to what was happening within the model. In the case of the experimental runs now under discussion this aspect of the output became exaggerated. This was to be expected; the presentation of a Cue Sequence in a Test Trial would reactivate several matching Linked Constellations which would however contain a representation of Action B in some of them and Action A in others.

The point of these runs was to see to what extent the order of presentation of particular pairing affected the results at the Test Trials. The effects were not studied systematically, and are reported here for completeness to explain how it was known that the first Training Trials had an undue effect on the results.

It was in fact found that when the pairing was equally balanced over the run, the Response which dominated at the Test was that emitted in the last of the Training Trials, i.e. a recency effect. It was also found that a pairing in the first Training Trial of a run could outweigh two or even three Trials with a different pairing in the middle of the Training Session. It has already been mentioned that this effect was reduced by introducing a notional period of maturation for the Integrating Neurons.

IV Equipotentiality and Mass Action

Lashley's (1949) principles have been considerably modified by the more detailed knowledge which we now possess concerning the structure and function of the localised regions of the cortex. Some of his results may be explained in terms of redundancy arising through multiple representation (see for example, Zeki, 1978). Nevertheless, many contemporary workers remain convinced that there is a core of necessary truth in Lashley's demonstration that the engram must be distributed throughout the brain (e.g. Oatley, 1978). Some workers consider that specialisation in the cortex is less simple than it appears at first sight and is dependent upon inter-regional interaction (e.g. Iversen, 1973). Others that a theory of neural function which relies upon the integrity of individual circuits is implausible (Pribram, 1971).

The present model offers a reconciliation of a local specialisation of function with a distributed memory concept. In the extended template idea, the Linked Constellation provides an extended system of specialised parts situated in different regions. Partial destruction of such a structure would be expected to affect the function of the whole only insofar as the information contributed by that part was missing or destroyed.

An analogy of Lashley's experiments in destroying parts of the cortex was to remove parts of the Neuron Arrays of the Cortex of the model

Table 9.2 Notional ablation of Cortex after Training Session and before Test Session

Each figure displayed is a learning mark assigned to an experimental computer run. The learning mark was based on the numbers of Spikes appearing in the Motor Outputs relevant to the expected Conditioned Response, in excess of the numbers of Spikes in the Outputs relevant to the Conditioned Response to the discriminated stimulus. This was the aggregate of Spikes over the three Instants following the presentation of the last of the Input Patterns of the Cue Sequence.

No. of Neuronal Pools notionally ablated	Cues and their expected Responses			
	Cue A Response 1	Cue A Response 2	Cue B Response 1	Cue B Response 2
Control (none ablated)			957	1200
1 Neuronal Pool		1443	2237	
2 Neuronal Pools	871 0		2659 1225	1112
3 Neuronal Pools		540		
4 Neuronal Pools			2276	1287
8 Neuronal Pools		587 0	0	

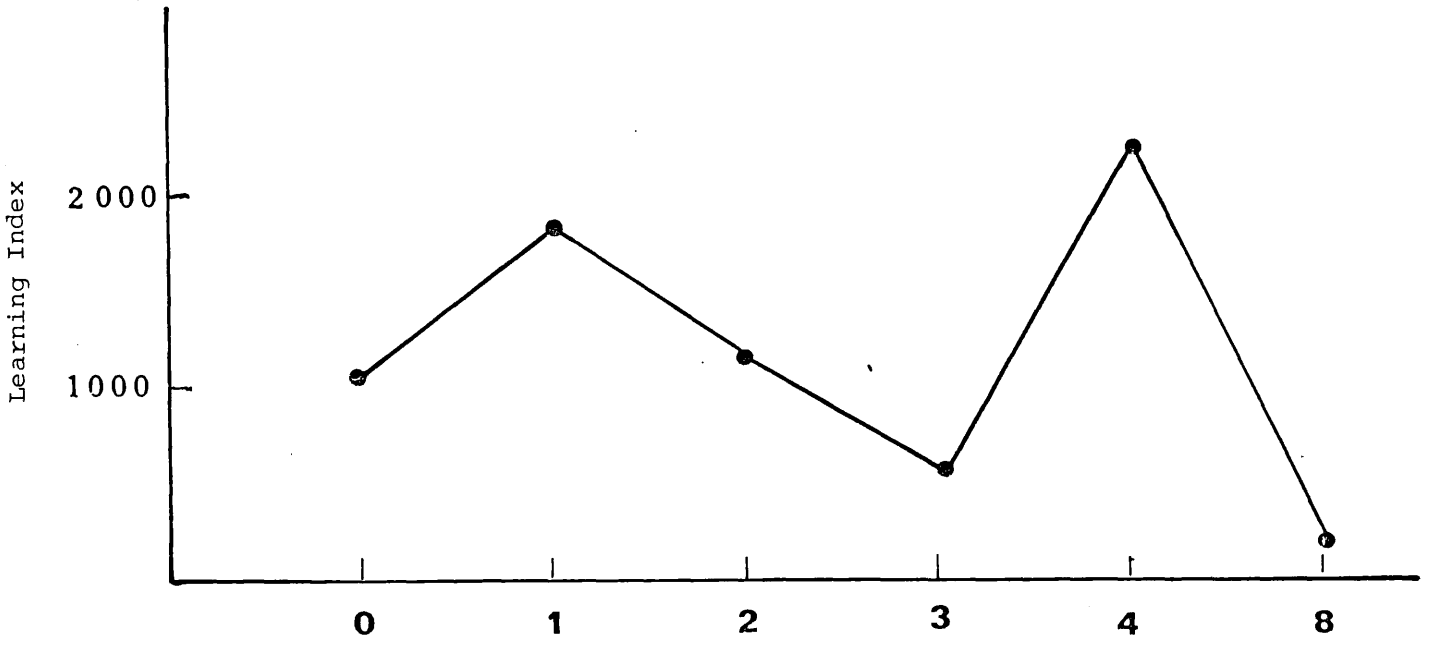

Figure 9.3 Notional ablation of Cortex. In a number of experimental computer runs of the model, after the Training Session selected Neuronal Pools were removed from the model. The results of the runs are shown in Figure 9.2. The above graph displays the results from the 13 runs which were most closely comparable. The Learning Index was calculated as described in Figure 9.2.

either before or after learning. A large number of runs were made in which various randomly selected Neuronal Pools were skipped by the program, simulating lesions. These varied from a single Pool to one quarter the total number of Pools. In some runs the notional lesions were made from the beginning of the run and in others they were made after the Training Trials, just before the Test Trials.

The strategy in which the notional lesions were made from the beginning of the run was not extensively investigated, since it soon became apparent that the procedure of removing Neuronal Pools from the beginning of a run merely reduced the effective size of the model. Provided the effective input pattern was taken as consisting of only those bits of the whole input pattern which corresponded with the Neuronal Pools allowed to remain, the system differed in no way from the whole model.

The effects of notional lesions after Training and before Test was investigated more extensively. Figure 9.3 incorporates the results of 13 such experimental runs.

The standard of performance in the versions of the model used for the runs concerned was less tidy than those reported so far. At this time efforts were being made to improve the Inhibition system. These efforts were unsuccessful, and this was one of the troubles which impeded the development of the model and which are reported late in the thesis. Consequently in these runs there were copious outputs in all eight Efferent Fibres. The information presented in Figure 9.2 is based upon a Learning Index calculated for each run. This Learning Index was produced by summing the effective outputs over the three Instants which followed the termination of presentation of the Cue Sequence. The effective output was calculated by simple arithmetic operation on the Spike Frequencies of the Efferent Fibres which established which response was dominant. Notwithstanding the necessity for these subsidiary procedures there is little doubt that the results represent a valid assessment of response of the models. In Table 9.2 the results in each of the runs concerned is displayed separately and it may be seen that notwithstanding permutations of design of the run pairing Cue Sequences and IRP's, the notional Conditioned Response of the models, as measured by the Learning Index, is always dominant in the appropriate Response, and considerable. As in the case of the

tidier performances of earlier versions, the permutations of presentation remove any possibility that the results could be artefacts or chance variation.

In summary it will be seen that the magnitude of the notional Conditioned Response varies considerably, and not necessarily consistently with the notional ablation of Neuronal Pools. What is quite apparent is that the conditioned response remains reliable with up to 4 Neuronal Pools removed, after Training. When 8 Neuronal Pools were notionally ablated, and these Neuronal Pools were removed completely at random, an appropriate notional Conditioned Response was still obtained. However, when the 8 Pools removed consisted of a set of four Pools and the four Pools representing the 1 complements of those Pools with reference to the Input Pattern, no relevant conditioned response was obtained. This suggests that at least part of the capacity of the model to continue to display learning after notional lesions is to be attributed to the redundancy implied in the balanced receptor system.

These results are hardly surprising. The removal of part of the model leaves a smaller version operating upon the same principles. The random variability of the system remaining depends upon the relationship of the missing pools to the information which would have been held in those pools. Although this relationship is necessarily more complex than the more obvious implications of removing Pools at the beginning of a run, the basic principles are the same.

It may be concluded that the system is of the kind indicated in Lashley's Laws. Removal of parts of the model after training does not prevent the successful performance of a notional CR but renders it more variable. Since only a few runs were made under each condition of ablation, it is not surprising that increasing unreliability of response should produce a complete failure with the removal of a large number of pools.

V Sequential dependence

In view of the emphasis placed upon temporal integration in the introductory chapters, it was felt desirable to demonstrate that the output of the simulations was not the function only of the last pattern of the Cue Sequence (of four patterns), and was the result of the whole

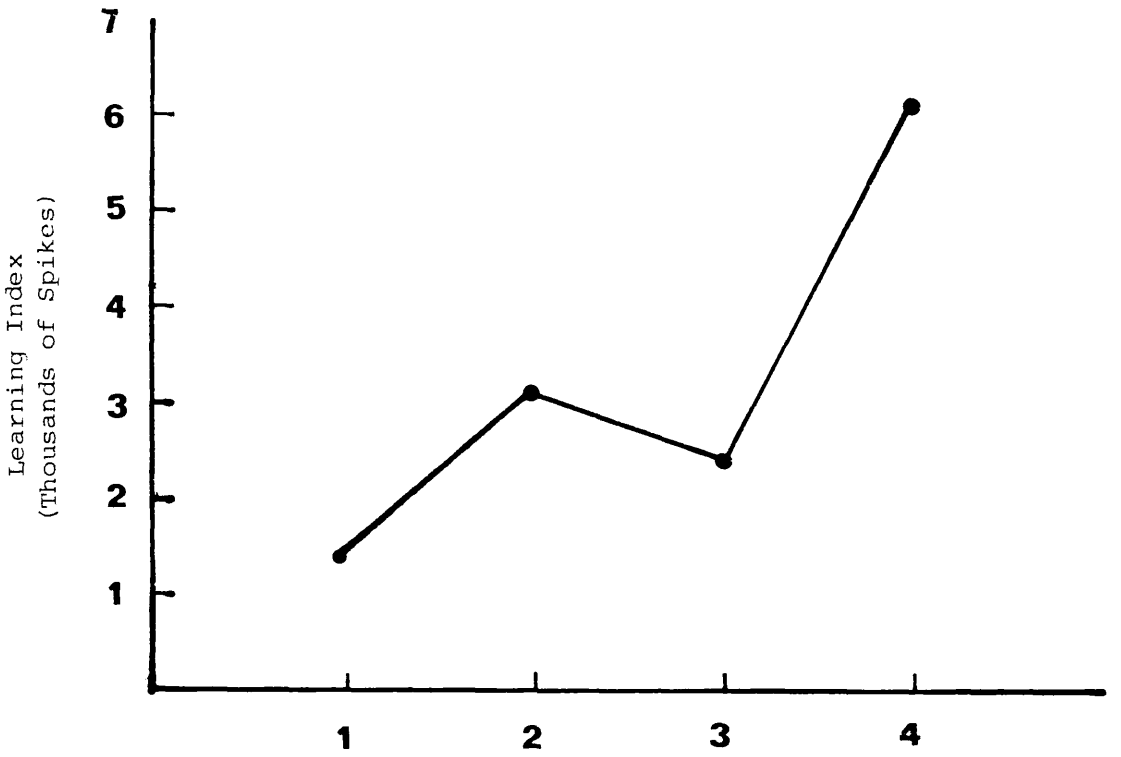

Figure 9.4 Temporal Sequential Effect. In a number of experimental computer runs of the model, during Training the Cue Sequences consisted, as previously reported, of four Indifferent Input Patterns presented in a given order. In the Test Session the Test Cue consisted only of 1,2,3, or 4 (i.e. complete) patterns of the Training Sequence. The above curve illustrates 4 of these runs. The Motor Response (CR) of the model was measured by a Learning Index, which was calculated by taking the number of Spikes in the Motor Output relevant to the expected CR over the three Instants following presentation of the Cue, and deducting the corresponding Spikes in the Motor Output relevant to the irrelevant CR. These results were not systemised, and are to be regarded as illustrative rather than significant.

sequence. As now seen, the temporal sequential effect in the model is not necessarily to be derived from an inherent sequential attribute of the proposed integrative connections. Differential delays in sensory pathways (i.e. Uttley-type delays) would produce a similar effect in my model. Nevertheless, the demonstration of temporal sequential effects was carried out in versions of the model without Uttley-type delays, that is, without specific delays introduced into selected pathways.

A large number of experimental runs of the model presented Cue Sequences at the Test Trials which differed in various ways from the Cue Sequences used in the Training Session. In general the same Input Patterns were employed but the Test Cue Sequence differed from the Training Cue Sequence either in the number of patterns in the sequence or in the order in which they were presented.

In one set of such experimental runs, reported in Figure 9.4, in different runs instead of the full Training Cue Sequence of four patterns, a sequence of either three, two or even a single pattern from the Cue Sequence was used. The versions of the model in which these results were obtained were similar to those reported in the preceding Section, and a similar Index of Learning was used.

The results may be summarised by saying that the magnitude of the Response and its reliability was greater the more closely the Training Cue Sequence was reproduced in the Test Cue Sequence. A consistent regression curve cannot be seen, and in any case, the derivation of such a curve would be of little validity in the circumstances of the investigation. It is however clear that the notional conditioned response is diminished greatly when in the Test Trial only the last pattern or last two patterns of the Training Sequence are presented as compared with the response when the complete sequence or three of the four patterns are presented.

VI Discussion of initial simulations

A. Demonstration of basic attributes of the model

The simulations which have been reported in this and the preceding chapters showed a number of deficiencies and disclosed difficulties.

The performance of which they were capable was very simple. This is the result of trying to simulate a system resembling in principle the physiological characteristics of neural tissue, even in a simplified manner. Many aspects of neural processing which do not arise in a formalised model become the occasion for difficult decisions and solutions. For example, the control of excitation and inhibition and their precise function in the nervous system are only poorly understood at the present time, and there is therefore no clear direction towards which the system can be guided.

In spite of these drawbacks, the simulations confirmed the essential properties attributed to the theory in its verbal form. These properties may be considered in relation to earlier discussions of the problems with which systems of neural integration must eventually deal.

(i) Neural templates. The simulation has demonstrated that a system of the kind proposed could register the pattern of excitation in both sensory and motor fibres in the form of extended spatial templates encompassing both aspects of nervous activity. A large number of such templates could be held in the system at one time in the form of integrative neuronal connections. It was shown that such templates would be differentially reactivated by subsequent inputs, the state of excitation of each of the templates being dependent upon the extent to which the current pattern of afferentation matches each template. The discriminative capacity of the system was effective at the level of one complementary-pair of bits difference in sequences of four consecutive patterns each having 16 pairs of complementary bits.

(ii) Temporal sequence. It was shown that a temporal sequence of patterns is more effective than a single pattern in reactivating the system of neuronal templates. This demonstrates that, aside from the appearance of special effects in relation to inputs of a special kind, i.e. innate releasing patterns, the system acts generally as a conditional probability machine, that is, its operation is modified at each step by what has happened at the preceding step. The special effects which appear when a Motor Response is elicited, i.e. the conditioning response, are not in themselves the mediators of learning in the system; they are events which enable the underlying process to be observed in motor behaviour. The temporal sequential aspects of the system are relevant to the general position urged in the introductory chapters that behaviour is implemented by an extended sensori-motor

template system which is capable also of temporal sequential integration.

(iii) Innate and Learned behaviour. The model provides a structural organisation which shows how innately determined neuronal connections could interact with those formed as the result of individual experience.

(iv) Classical conditioning paradigm. The model showed motor behaviour analogical to that observed in classical conditioning experiments. Notional conditioned responses appeared after training even when Cue Sequences, which were indifferent prior to training, were presented without the notional UCS (IRP), i.e. in unreinforced trials. In tests in which two notional reflexes were conditioned in the same run, the conditioned responses discriminated between two notional CS. (See also next Section). In other respects the behaviour of the model resembled the conditioned response which appears in animals, although many aspects of conditioning, e.g. partial reinforcement, were not tested for. It may perhaps be mentioned that one of the reasons for not testing the model further in such respects were the limitations on the length of computer runs with reasonable turn-round. The number of Trials necessary to investigate such effects would have been prohibitively large, even if a run of sufficient length had been appropriate to a model having the number of Integrating Neurons provided.

Insofar as the behaviour of the model may be judged to represent the characteristics of classical conditioning in animals, one or two points may be noted.

The first is that the model exhibits a very clear distinction between the UCR and the CR. The UCR occurs primarily as the direct response of the genetically determined network. The CR does not involve the genetically-determined network at all, but is the indirectly induced response in the Learned Records.

In the model, as compared with that of, say, Pringle (1951) that part of the neural apparatus which might be compared to the reflex concept in Conditioning Theory, remains unaffected by experience. The model does not constitute a substitution theory in the sense that the CS, substituting for the UCS, now acts through the reflex apparatus. The model proposes a separate apparatus which consists entirely of traces laid down in experience. In the model this apparatus is shown

as structurally separate from the apparatus mediating genetically-determined response. This complete structural separation is probably an oversimplification. More probably there would be evolutionary specialisation in which one structure specialised in Learned traces, the other in innate mediations. Still the essential idea remains that the two are functionally separate. It is in the separate apparatus mediating traces laid down in individual experience that the response, corresponding to the CR of classical conditioning, is initiated.

A second point concerns aberrations and lack of tidiness in the outputs of some of those versions representing something akin to a double classical conditioning, for example, that shown in Figure 9.1. The curve in that Figure demonstrates a very clear notional CR of Action B to Cue B in the first Test Trial, and a clear CR of Action A to Cue A in the second Test Trial. In the third Test Trial there is a clear CR in Action B to the presentation of Cue B, but there is also an unexplained rise in Action A. Outputs of this kind are, it is suggested, not to be considered as demonstrating deficiences in the model, but rather characteristics of the system which may, or may not, be analogical to characteristics of nervous systems.

In the case in point it will be recalled that the aberrations were attributed to the action of the feed-back "conditioned-arousal" system. However this only exaggerated tendencies in the system which must have been present. In this context it may be borne in mind that although the performance of the simulations is simple in the sense that it involves only one or two outputs, so is the behaviour of the animal as measured in a classical conditioning situation. In fact the introduction of two Unconditioned Responses and two different Conditioned Stimuli to operate on those responses is rare in the laboratory, and the resulting performance is difficult to analyse. That the operation of the system simulated here may not be unlike that of an actual nervous system does not seem implausible in view of the results of Varga and Pressman and Lelord (see Konorski, 1967, p.293 et seq.). These workers used two "relatively neutral" unconditioned responses and two CS. Each of these CS established its own CR. However, when presented separately each CS produced in addition to its own CR occasionally the CR associated with the other. These responses are analysed by Konorski in terms of pseudo-conditioning and backwards conditioning, in terms which would apply to the simulations.

(v) Independence of external assessment of events. "Reinforcement" may be used without theoretical commitment to refer to an event which increases the frequency or strengthens the kind of behaviour which is associated with it. The point was made earlier that many simulations of neurological models have relied upon an external reference or reinforcement line which signals to the model whether its responses have been judged as correct or incorrect. In my model the assessment of the character of the notional stimuli and the response to them is determined entirely by the internal structure and function of the system.

(vi) Behaviour and arousal: two modes of response to events. It may be noted that the presence of an internal criterion of that character of events is evidenced in the simulation in two ways: behaviour and arousal. With regard to arousal the point has been made that, in spite of its inadequacies, the system did enable the model to vary the number of Integrating Neurons incorporated into its Neural record of events according to the assessment made in relation to this internal criterion. This capacity could be seen in a straightforward sense in the effect of the presentation of an IRP on the Motor Response curves. (Although the system is referred to as an "arousal", this term is used for convenience of reference only).

It was also shown that the arousal system could vary the level of activity, both in retrieval and registration, according to its notional conditioned response, that is, in anticipation of actual events. However, the hope to show that it could do this without undermining its basic processes of neural integration was only partially successful. There were some indications that the response of the arousal system was not always entirely appropriate to the behavioural response and produced distortions, e.g. in the third Test Trial shown in Figure 9.1. Nevertheless, as may be seen in the first two Test Trials of this run, the basic operation of the concept was demonstrated.

There were a number of clues to the causes of the disturbances of function which were exaggerated when it was attempted to increase the effectiveness of the system. One was the inappropriateness of variations of a diffuse excitation parameter which had the effect of magnifiying irrelevant responses. Another was the explosive feed-back implications of the methods employed. Because these were embarassments to the obtaining of tidy performance does not necessarily mean that they are not

aspects of real nervous systems (see next sections).

It seems an important step in relating attentional mechanisms to theories of neural integration that it can be shown that levels of activity of retrieval and registration may be varied during neural processing. This applies not merely to concepts of general states of activation, as in sleep and wakefulness, but also in connection with directed attention. The present approach was outlined in Chapters 2 and 3, where selective attention was referred to as a combination of temporal and sensori-motor integration in conjunction with a variation of the weighting applied to parts of the proposed extended template system. Temporal integration and sensori-motor co-ordination co-operates with such a differential weighting. For example, if the signals reaching the auditory part of the template indicate (by the matching of templates already present in the nervous system) that important information may accompany the source of sound, then the head and eyes will turn in that direction: it is suggested that visual processing would then be highly activated not only in the optical mechanisms such as contraction of pupil, convergence and focussing, but also in the neural template dealing with visual signals.

Although it seems important to establish the principle of the possible variation of levels of activity in the system, it is not suggested that the control need take the form simulated, that is, of a simple variation of the parameters of the Spike Frequencies. On the contrary, difficulties in fact arising from the arousal aspects of excitation and inhibition led to a considerable modification of concepts of the mechanisms by which levels of activity are modified in the system. These will be related in due course.

(vii) Anticipatory discharges. In Chapter 3 the possible relationship of anticipation to two aspects of behaviour was discussed: one was the anticipatory motor image as the source of motor action; the other was the possible function of anticipatory images as test criteria in feed-back systems of the kind envisaged in Pribram's TOTE paradigm. It was therefore felt to be important to demonstrate that such anticipatory images would arise in the system of neural integration.

The occurrence of a possible neural substrate for such images was clearly demonstrated in the simulations in a number of ways.

The rise in the curves representing Motor Response is evidence that when a familiar sequence is repeated there is a rise in the Excitation of those Linked Constellations about to be stimulated by the next pattern in the sequence. This was also confirmed by printouts showing the state of the system. This rise demonstrated an actual temporal anticipation of events. However, for the purposes mentioned it is not necessary that the indication given of future events should temporally precede the customary times of occurrence of those events. It is sufficient for the system to provide a logical anticipation of events which have previously followed those which have occurred on the current occasion. The notional CR of the model, as does the actual CR in the animal, gives a clear indication that within the nervous system such a logical anticipation is present, since the CR occurs in the absence of the UCS. As already mentioned, in my model the CR is an indicator of this anticipatory process rather than the instrument of its mediation.

With regard to the use of the anticipatory aspects of the system to monitor its own behaviour (i.e. the relating to Pribram's TOTE concept) it seemed possible that changes in the level of the anticipatory Discharges might be used to inform the system as to the probable outcome of its actions. Briefly, if Excitation were rising in the output relevant to the CR then the model must be "on the right scent". If it were falling, then something was wrong.

These possibilities were pursued in attempts to extend the model to provide more ambitious explanations of behaviour. In the event they offered greater difficulty of implementation than expected. They will be discussed further in connections with later versions of the model.

B. Some difficulties and drawbacks

These will only be mentioned briefly here, since they will be referred to again in conjunction with the action taken later to remedy them.

Problems in the simulation revolved around the control of the levels of variables representing the physiological aspects of the system, i.e. inhibition and excitation. In the later versions of the model fluctuation of the variables tended to swamp the differential responses of the model. These difficulties extended into the versions of the model to be reported later, and were only understood late in the development of the model.

At this point it may be sufficient to indicate two sources of these difficulties:

(a) Inhibition and Excitation were represented as being independent neuronal processes: i.e. Inhibition did not cancel out Excitation. In the nervous systems of animals it may be assumed that although inhibition and excitation act at different neuronal sites, they interact in the membrane in the actual spike function. Yet it is inadequate to represent the interaction as a simple cancellation; that would ignore the facts of temporal and spatial integration of excitatory events in the initiation of action potentials. For this and other reasons mentioned earlier, rather arbitrary assumptions were incorporated in the model aimed at retaining one's options. The alternative to the choice made, that is, cancelling out Inhibition also seemed inadequate, since it did not allow for the accumulation of excitatory and inhibitory transmitters as a carry over from Instant to Instant of the simulation, nor for a possible post-inhibitory excitatory rebound, as proposed by a number of authors, notably Andersen (1974). The difficulties could accordingly be artefactual in the model and due to an inadequate representation of neuronal processes. On the other hand there is the possibility that the difficulties encoutered in the model represent actual difficulties which the nervous system has had to remedy. Experience in the later development of the model suggests that this latter alternative is at least possible.

(b) It may be thought that it would be possible to solve the problem by reducing the excitation level of a given cell by some fixed proportion at each Instant. This was in fact the method used, and it did work in the sense that given short periods of operation it gave sufficiently adequate control to produce the results which have been presented. What was not recognised in the early stages was that whatever proportion of Excitation is removed from the system at each Instant, so long as some remains as a carry over to the next Instant, that carry over must build up and eventually swamp the system, given continuous, that is, uninterrupted operation. On the other hand, if there were no carry over there would be no temporal sequential build up of Excitation and the temporal sequential aspects of the system would be negated.

(c) The explosive implications of feed-back arousal assumptions has already been mentioned. With continuous operation of the system a

feed-back arousal would inevitably lead to the eventual magnification of fluctuations of output and the corruption of processing.

In support of the work presented in this thesis, it may be considered that the difficulties and deficiencies in the simulation emphasise the desirability of simulating models of a physiological rather than a formal logical character, even if the simulation of physiology is rather crude. The simulation disclosed problems relating to levels of Excitation and Inhibition not contemplated at the outset. The attempt to devise a working system led to concepts of the control of these variables, which, even if they prove to be mistaken, demonstrate the capacity of simulation to highlight problems which may otherwise be generally overlooked.

C. Simulation of Linked Constellations mediated by orthodox synapses

As mentioned in Chapter 7, Section III, a separate simulation demonstrated the possibility of mediating Linked Constellations by orthodox formal synaptic connections. Since that simulation utilised methods of control of the physiological variables which had not yet been developed in the initial set of simulations run in the Atlas computer, it will be reported later.

D. Summary of discussion on the results of the initial simulations

Notwithstanding the difficulties referred to, the simulations were able to confirm that the theory of neural integration outlined earlier would in practice possess the general properties claimed for it.

The important aspects of the theoretical model demonstrated were the practicability of the neural template idea, together with the demonstration of a notional response analogous to the conditioned response in a classical conditioning experiment. The simulations demonstrated that the performance of the model was not due to artefact or chance.

It was shown that an anticipatory rise in the Excitation of Linked Constellations occurred when a familiar sequence of inputs was repeated. The model was able to identify inputs and to "recognise" them in the sense of relating them to earlier presentations of the same input. The model also demonstrated a possible functional relationship between genetic programming of behaviour and learning from experience.

CHAPTER 10

DISCUSSION PRELIMINARY TO PRESENTATION OF LATER VERSIONS OF THE MODEL

I Introduction

At this point some preliminary remarks may assist the reader to orient to the rather changed approach of the later simulations of the model. Moreover, having considered a simulation of the basic theory of neural integration it is possible to comment futher upon a number of other solutions in the literature.

II A notional brain

The computer simulations described so far may be considered to bear some resemblance to nervous systems. Nevertheless they fail to demonstrate many attributes which seem fundamental to animals with brains. In particular they cannot deal with their notional environment in any effective way. Putting this another way, the simulations represent classical conditioning rather than instrumental conditioning situations. The model as simulated is both structurally and functionally too simplistic to represent a brain: it has no apparatus by which it could alter the contingencies of its environment.

Another way in which the model fails to represent a brain is its neglect of functional and structural features which research has shown to be important, e.g. the limbic system and motivational factors.

Of course an attempt to model a brain must be taken in the context of an explanatory model and of a computer simulation. The aim is to model the overall capacity rather than its detail. In many cases a rather simple solution to a problem will be adopted, even though there is evidence that it is pretty far from the truth, simply because to model that function realistically would take one too far from the main purpose.

A difficulty in attempting a more realistic model of a brain may be expressed in the question: what brain? Many of the capacities of the brain which most interest us only appear in a definite form in the case of the higher animals. Some of them only in Man. An attempt to simplify the model by trying to represent the simplest of nervous systems directly would fail to deal with this problem. One can however to some extent resolve the difficulty by modelling a notional brain, that is a brain which could not really exist but which represents some of the capacities which real brains possess. It would be immeasurably more simple than any real brain. Nevertheless it might exhibit explanatory principles in a form in which the essentials rather than the complexities are reproduced. In the context of the present effort, the main aim would be to show how an overall sensori-motor integration might be achieved.

III Some continuing aims

One aim of endeavouring to extend the model so as to bring within its compass a wider concept of behaviour is as a demonstration of the explanatory potential of the basic theory of neural integration. Clearly, the greater the capacity of the system to mediate life-like behaviour then the greater the plausibility of the theory as an account of neural integration. This would apply whether the ways in which such behaviour is achieved bear much resemblance to the ways in which it is achieved in the animal or not.

One is also interested in developing a convincing model of the brain. For example, one wishes to see to what extent a sensori-motor temporal-sequential integration would explain brain function. At another level one may wish to explore hypotheses in general circulation and to follow up those suggested by the model itself.

In general one may consider the attempts at extended simulations as utilising a tool, the computer, to develop one's thoughts concerning brain function in a definite and directed fashion. A simulation provides a unique form of understanding. It is hoped that the reported attempts to develop simulations in a computer may take the effective representation of brain function a step forward.

IV Ad hoc solutions and general principles

To bring a model closer to the representation of behaviour, it is necessary to adopt mechanisms of action which are able to solve the notional problems of environmental choice. These mechanisms are not inherent in any obvious sense in the theory of neural integration. For example, there is the manner in which the model is motivated towards its goal, the manner in which responses arise, and the manner in which they are selected and witheld.

There is bound to be a certain ad hoc character in the means chosen to implement behaviour in a model. This is justified on the basis that in the nervous system mechanisms have evolved as the result of interaction of the organism with its enviornment, that is, that they reflect ad hoc solutions taken by evolution. Of course, simulations cannot hope to imitate evolution. They merely hope to approximate in a general way to the kind of mechanism which has developed. In that sense the simulations represent a general interpretation of behavioural mechanism rather than a set of hypotheses of actual neural function.

There is however a sense in which the solutions taken by evolution are not ad hoc. The mechanisms of living things reflect the potentiality of living matter. This is expressed in the model in two ways. One is by using a particular theory of neural integration to implement the mechanisms. The other is by attempting not merely to deal with isolated problems but to look for solutions which may have some generality.

The problem of modelling a notional brain may be considered from two aspects. First there are the problems of the nature of specific mechanisms mediating learning, perception, selective motor action and so forth. Second, there is the problem of the integration of these activities into a co-ordinated whole.

Both aspects are closely related and are dependent upon the way neurons are organised. Evolutionary development over millions of years has operated at all levels. At the level of the neuron, there are molecular mechanisms which impart to neural tissue its power to adapt and respond. In turn these molecular mechanisms are integrated into specific neural mechanisms of behaviour and perception. At the

level of the unity of the animal these mechanisms are themselves integrated to perform what Konorski (1967) calls the "Integrative Activity of the Brain" and Hebb (1949) "The Organisation of Behaviour".

The rationale of the neurological model presented here is that all three levels of integration express a single principle. This principle has been expounded at some length. It may be summarised as the capacity of the elementary particles of living matter to respond to the environment by forming a selective association. Each such association represents a particular instantaneous state of the environment as signalled by the energy configuration falling upon the particles at that instant. The development of multi-cellular organisms is seen as an expression of the principle. Within such organisms a certain modicum of the unity of the single-celled organism is preserved by such selective associations. In animals with complex nervous systems the principle is suggested to find expression in the formation of Linked Constellations. These intimate associations of selected neurons provide unified structures within the system which represent the unity of the organism at the moment they are formed.

The problem which confronts the designer of a brain model is two-fold: "The premise underlying the "integrative" or "synthetic" approach . . . is that ultimate understanding of the living organism lies not only in knowing its component parts, but that a unique and fundamental task in biology is to determine how the many separate parts or processes are integrated into the organism as a whole". (Mason, 1975).

V Some other models revisited

At this point it may be useful to make some further mention of other models in the literature, which were cursorily mentioned earlier.

A. Anatomically inspired models

Generally no model of neural integration has been simulated to represent an overall brain model in which the different functions of parts of a nervous system are represented and their overall integration explored. However, some models have utilised a principle of neural

integration to represent particular aspects of the brain in a separate model. For example, various models utilising different applications of the feed-back weighting principle have been used to explain: (i) pattern recognition (e.g. Rosenblatt, 1958a, 1958b, 1960; Selfridge, 1958); (ii) The structure and function of the cerebellum (Marr, 1969; Albus, 1971); (iii) selection of motor response (Kilmer and Olinski, 1974); (iv) temporal sequence (Stanley and Kilmer, 1975); (v) setting modes of behaviour (Kilmer, McCulloch, Blum, Craighill and Peterson, 1968).

It is no criticism of these models, in respect of the purpose for which they were framed, to suggest that they do not demonstrate how the overall activities of the brain might be co-ordinated. Of course, by implication they suggest that the problem of overall integration is merely a higher step of the integration shown in the separate models. That implication is not, however, a simulation. The point is made in order to clarify the main aim of the simulations which follow, which is to simulate the manner in which different aspects of neural process integrate different aspects of behaviour.

One of the aims of many of the models which have been mentioned is to interpret the neuro-anatomy of some particular part of the brain in terms of some particular integrative model. To illustrate what has been achieved in this direction, two examples will be considered.

Kilmer et alia (1968) examined the available anatomical and histological data relating to the Reticular Formation. They interpreted this data on the analogy drawn by Scheibel and Scheibel (1962, 1967) as a stack of poker chips. Each disc-like anatomical unit is then considered to be a modular logic unit. The model converges upon a particular disc, as representing the best selection from input.

This selection of discs acts, it is proposed, to determine the mode of behaviour of the animal. "Its primary job is to commit the organism to one or another of about 16 gross modes of behaviour - i.e. to run, fight, sleep, speak, as a function of the nerve impulses that have played on it during the last fraction of a second, and also to send out control directives to the other more specialised nerve centres so that in turn they can behave in an integrated fashion".

The psychological value of this principle is not to be doubted. The

model has been cited because it demonstrates, in relation to one particular anatomical region, one possible approach to the overall integrative activities of the brain. In effect, this approach is an hierarchical one. There is a controller which sets "the mode" within which separate activities can be integrated.

The computational sophistication of the model is extremely powerful. It shows a practical, if exceedingly complex, manner in which bundlings of a large number of inputs may be rearranged to select the sixteen (sic) best groupings, as a result of feed-back from results.

Although not explicitly spelled out, the model implies a solution to the problem of overall co-ordination. At the top of the hierarchy the RF* (the conceptual Reticular Formation) selects a mode. Presumably the control directives supposed to be issued by the selected module in turn select basic input conditions for a similar computational unit elswhere, whose function is to apply the mode to the particular circumstances of the moment. What is not demonstrated in the model is how this co-ordination would actually work. It is not clear how an effective strategy would develop, and how it would be organised. Does the "speak" mode merely switch to a particular computational unit handling speech? In which case, how is the mode of speech, e.g. conciliatory, aggressive, determined? Or are the control directives concerned with the latter aspect, that is, do they control parameters of the functions, so as to determine that speech is agressive or conciliatory?

At a purely practical level it is not clear what happens. It is conceivable that one of about sixteen, or even for that matter, sixteen hundred, modes might be selected in the manner proposed. At the next lower stage of the hierarchy, either a selection of motor responses is fixed in relation to each RF* disc, or a combination of lower level responses is subject to further computation. If the former, then the behavioural repertoire would be too inelastic to represent behaviour. If the latter, then it is not clear upon what inputs the weightings would operate upon at that level and what criterion of success the weightings would use. Would it be the same criterion as used by the RF*? In which case the computational conjunctions are enormously increased numerically.

As has been said, these criticisms do not fall upon the model for the

main purposes for which it was constructed. Such models aim at showing that the neuro-anatomy of particular regions could perform certain computational functions, and demonstrate that such a computation, in this case one of considerable psychological importance, could be achieved. Nevertheless, the function of the model as representing a particular brain function, must remain the analysis of a number of input lines and the derivation of an output. In order that a model may truly represent the function of a portion of the brain, it seems necessary to show how the inputs to that region and the output from that region fit into the overall pattern of inputs and outputs across the brain. Otherwise, we are left with a number of parts to fit together.

The principle of co-ordination suggested by Kilmer et alia's (1968) model may be compared with the approach in the ensuing simulations. In those simulations, the individual computing power of the parts is developed to a far less degree of sophistication than in Kilmer et alia's model. What is aimed at is an overall co-ordination of function in which the different aspects of behaviour - instrumental activity, consummatory activity, protective activity, the maintenance of metabolic and humoral conditions, - are themselves simulated in relation to a notional environment. The overall co-ordination is achieved by the principle of an extended neural template, the Linked Constellation. Thus the two sets of models (those of Kilmer and his group, and my models) demonstrate two approaches to neural integration. In particular they represent two conceptions of hierarchy. In the Kilmer model, the hierarchy is of a controller and controlled units. In my models, the hierarchy is rather that observed by Tinbergen (1951), a progressive targeting of behaviour, each step in the hierarchy being a comprehensive co-ordination. Of course, that is the aim of Kilmer et alia's model of the RF. The comparison is between the different ways of achieving the same end, co-ordinated behaviour.

Some of the answers to the questions raised with regard to the relationship between the RF* model and the subordinate computational systems controlled by it, may have been in mind in Kilmer and Olinski's model of the Hippocampus (1974), as supplemented by a model of temporal sequence by Stanley and Kilmer (1975). In a separate paper Kilmer (1975) reviews the anatomical and functional research concerning the Hippocampus. Kilmer and Olinski interpret the particular neural arrangements of the

Hippocampus as an iterative feed-back weighting adjustor, similar in principle to those described. Each Hippocampal pyramid is seen as an integrator, registering a particular permutation of input lines, and enabling their adjustment on principles similar to those laid down in Selfridge's Pandemonium and Rosenblatt's perceptron. Naturally this brief description does not do justice to either the computational ingenuity, nor to the correspondence drawn between the known inputs to the Hippocampus and the function performed by the model. The model computes an output based upon a sequence rather than an individual input pattern.

The implementation of the model is by means of canonical equations rather than neuronal in concept. For example, the temporal sequential aspects are simulated by the inclusion of a "Markovian factor". "We interpret this term ($p_i(q)$) to represent an intracellular or multineuronal process whereby sequential response habits are taken. The tendency of neural tissue to take such habits is not contested, though the mechanisms remain unknown. Our purpose is to obtain an estimate of a Markovian predictor term for our model . . ." For a criticism of the application of the Markovian model to behaviour the reader is referred to Mackintosh (1974, p.13). To some extent this defect is remedied by the model produced by Stanley and Kilmer (1975). The latter model is too complex to be described here. It also relates to the Hippocampal region. By a complex selection mechanism involving impulses travelling along parallel transmission lines, temporal sequential conjunctions are represented. The model has a rhythmic aspect which is apposite to any explanation of the Hippocampus, in view of the importance of that region as the origin of the Theta Rhythm (Landfield, 1976).

To return to Kilmer and Olinski's (1974) model, although the function of the model is motor, the output of the model remains undefined as to its precise effects of the environment. The action of the Hippocampus is controlled by a notional "reinforcement line". It is trained by rewarding some responses and not others by an external decision. In other ways as a model demonstrating the contribution of the Hippocampus to behaviour it lacks the representation of the differing character of different aspects of behaviour as well as the probable complexity of the contribution which the Hippocampus makes.

As Kilmer himself points out (Kilmer, 1975) a number of complex functions have been attributed to the Hippocampus. These include attention, habituation, inhibitory control, arousal and movement, initiation of voluntary movement, match and mismatch between memory store and current input. Kilmer cites Gloor as referring limbic structures to experienced emotion. As Mogenson and Phillips propose "the noradrenaline and dopamine pathways are at the interface between neural systems concerned with the "intention to respond" and those concerned with "motor control" . . . the ascending dorsal NA bundle projects profusely to the cerbral cortex and Hippocampus by a number of routes". It is of course, implicit in Kilmer and Olinski's model that the inputs themselves signal these functions. As has earlier been stated, the point made here is not intended as a critism of the models cited in respect of their aims, but rather to show that they are unable to tackle the overall integration of differentiated function except in the abstract sense in which functions do not differ from each other but are merely sets of inputs and outputs.

It is perhaps permissible to interpolate here a few remarks in support of the present simulations as representations of the neuro-anatomy of the brain. The models of Kilmer and his associates are closely related to the neuro-anatomy of particular brain regions. However, their interpretations are by no means the only ones possible. For example, Olds' motivational model, cited in Chapter 4, also closely follows the known anatomical configuration of the Hippocampus, yet it is a very different model from that of Kilmer and Olinski. In the present simulations the aim has been towards an anatomical correspondence of a broader kind. It would be possible to relate parts of the models to be presented here to different anatomical regions, i.e. Septum, Hippocampus, Amygdala and Dentate Gyrus, Putamen, Globus Pallidus and so forth. In fact I have succumbed frequently to the temptation to do so. However, it would appear that at the present state of knowledge to allocate functions and interpret working in this way would be unduly speculative. What has been attempted in the simulations to be reported is the representation of function, e.g. motivational drive, energisation of innate response, motor selection, the relationship between consummatory and appetitive mechanisms, and so forth, by neural implementations which are, it is hoped, similar to the general kinds of cytological and anatomical arrangements seen in terms, for example, of specific transmitters. It is not suggested that these arrangements

directly resemble the actual neural tissue, but rather that it approaches a realistic representation.

B. A distributed memory solution: the neural hologram

In giving some indication of the present state of the art considerable attention has been paid to those solutions which in general are at the other end of the spectrum to that proposed here, that is, to hierarchical and iterative feed-back models. Before proceeding to develop the present model further, one should perhaps give some further attention to a proposed solution which is a diffuse distributed model, the neural hologram.

The development of the hologram as a physical principle led to the formulation of a number of neurological models based on the principle (e.g. van Heerden, 1963; Willshaw, Buneman and Longuet Higgins, 1969). Pribram(1971) uses the hologram principle in the setting of a synthesis of existing knowledge on brain function. His work is closely associated with the simulation by Spinelli (1970).

Pribram's proposals, insofar as they specify a model of neural integration, may be summarised under three heads:

(1) He proposes that the function of neurons is expressed in two processes rather than a single one, that is the propagation of spikes usually ascribed to neurons. The second process consists of changes in the "slow potential micro-structure". It is through these that he considers that neurons communicate with each other in the felt work of the cortex. In this my model is close to Pribram's view of neural function. Although I have been at pains to specify my model in terms of orthodox synaptic connections, I have hinted that the intimate mode of communication between neurons of a Linked Constellation would probably be mediated by some non-synaptic or at least some more subtle mode of communication.

(2) Pribram carries forward the two process idea by distinguishing the "state" from the "operator on the state". This conceptualisation was not present in my model initially, but it may be seen to be implied in later developments. In my model it became necessary to introduce rhythmic gate mechanisms. These slice up the inputs into discrete temporal frames and give intermediate periods of silence in which the response of the integrative system can develop. In my model the input function is mediated by what were termed morphological connections,

whilst during the period on input silence it is the Integrative Connections which mediate the Linked Constellation function. Analogically, the input phase may be seen as creating the "state" and the silent phase, "the operator upon it". There are therefore considerable affinities between Pribram's formulation and that developed in this thesis.

Pribram considers the two-process mechanism in terms of his well known TOTE hypothesis. A number of points of interest arise in that connection which are not directly related to the neural hologram model. These will be discussed in a subsequent section of this chapter.

(3) Pribram proposes that the junctional microstructure gives rise to a neural hologram which stores the total configuration of sensory and other events, in the same way as the holographic photographic process stores the total information concerning the visual scene in a distributed form across the whole negative. It will be immediately apparent that this view coincides with that underlying the model developed here in the sense of the comprehensive nature of the record of information. It differs fundamentally in mechanism, because my model depends upon the precise spatial location of information, i.e. it is a spatial template, whereas the neural hologram is at the other extreme, information is distributed so that one can say that all the information is at one spatial point, and at the same time at every other point.

Pribram stresses this aspect. "The hypothesis is based on the premise that neural representations are not photographic but are composed not only by an initial set of feature filters but by a special class of transformations which have considerable formal resemblance to an optical image reconstruction process devised by mathematicians and engineers. This optical process, called holography, uses interference patterns". (Pribram, 1971, p.141).

Pribram's proposals can be understood by reference to a neural implementation simulated by Spinelli (1970). The design incorporates all three of Pribram's concepts; junctional microstructure, nerve-impulse input and output; match and mismatch. The neural hologram works as follows. "An input fibre enters the cortex and connects itself to a receiving cell; this cell in turn gives rise to many branches that connect themselves to a number of interneurons . . . the interneurons

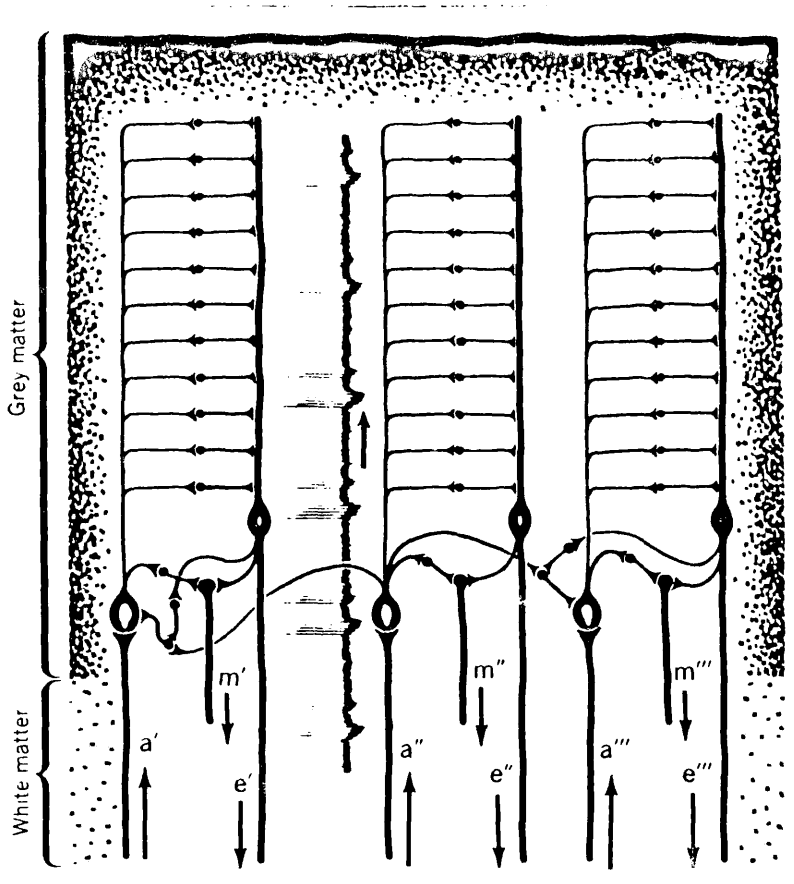

Figure 10.1 (reproduced from Pribram,1971) Spinelli's OCCAM. A neural implementation of the hologram concept, (Spinelli,1970)

in turn connect themselves to the dendritic ramifications of a further cell. This last cell generates an axon which leaves the cortex and is therefore part of the output system . . . Both the input cell and the output cell generate collaterals which connect themselves to a third cell called the Match cell. The Match cell in turn generates an axon which is also part of the efferent system". (Figure 10.1).

Spinelli himself appreciates that the arragements proposed could only learn a single pattern. His solution of this problem is to assume that different "words of addressible memory", that is, the ladder structures described, would respond differentially to different patterns of input. Therefore for each pattern one such structure responds more than the others and inhibits them (after Milner's modification of the Hebb model). Therefore only one or a few such structures learn each pattern.

Whilst one cannot but admire the ingenuity with which a hologram is reproduced in neural terms, one cannot help thinking that the assumptions just outlined detract considerably from advantages of the hologrammatic idea as an explanation of brain function. If each pattern is to be restricted to one or a few of the ladder structures, then the registration of the patterns is no longer distributed generally. Rather, the neural hologram is simply a means of holding a pattern at a localised region of the nervous system. The system then approximates in effect to a convergent model. Konorski's Gnostic Units are replaced by a ladder mechanism.

Spinelli's simulation successfully demonstrated that when a wave-form is input to a collection of the structures and is repeated some 50 times, then a further application of the wave-form produces a response in that model of the wave-form. Moreover a second wave-form can then be learned. "Upon presentation of parts of wave-form number 1 or wave-form number 2, Occam retrieves the remainder of the wave-form appropriately".

As mentioned in Chapter 3, Section III (a), Spinelli proposed that motivational and reinforcement could be represented. These aspects of the model were not, however, simulated.

Spinelli's simulation serves to demonstrate some of the inherent

limitations of the neural hologram concept. His model demonstrates that neurons can effectively act as mediators of neural holograms. It also demonstrates that the storage of more than one pattern involves departing from the hologram concept. Certainly one of the major conceptual difficulties in applying the hologram or any other interference pattern concept (e.g. Beurle, 1956; Lashley, 1942) is in solving the problem of multiple representation of patterns.

Essentially the hologram stores information concerning an entire configuration by recording in every location of the holographic record the interference pattern resulting from the unique configuration. The registration requires to be able to reproduce both amplitude and phase. The neural hologram must thereforehold considerable information at adjacent points. In Spinelli's implementation this is held by the strength of relative connections in a ladder formation. To store more than one pattern holographically requires that the points recording one set of interference patterns be distinguished from the points recording another. Spinelli solved this problem in the manner described, by making the registration apparatus differentially sensitive to different patterns. Put in this way, the difficulty of principle can be seen. An apparatus which is differentially sensitive to different patterns is already a registration system, and the holographic aspects are subsidiary.

Since the holographic explanation receives much support in contemporary literature (see for example, Landfield, 1976), and is a distributed memory model which is supported by many of the conceptual arguments used to support my model, some further points may be made.

There is the question of the topological mapping in the nervous system. As considered in some detail in Chapter 5, this topological mapping is everywhere evident. Palay (1975) has emphasised the precision with which the nervous system is organised, "the retention of topographic order as the synaptic fields were followed from level to level became more and more impressive".

Chow's (1970) experiments are amongst those which Pribram uses to reject the proposition that "some sort of "image" (is projected) from the receptor surface onto the cortical surface much as a photographic image is projected onto the plane film surface in a camera". He mentions that Lashley removed 80-90% of the striate cortex of rats

without impairing their ability to discriminate patterns. Galambos, Norton and Frommer (1967) cut up to 98% of the optic tracts, and Chow removed more than three fourths of the visual cortex without preventing the animals from discriminating patterns. In the case of Chow's experiment, visual discrimination became disturbed initially but the animals relearned the task in about the same number of trials required to learn prior to surgery. Roger Sperry and his group (1955) surgically cross-hatched a receiving area and even placed mica strips into the resulting brain troughs . . . Lashley, Chow and Semmes (1951) tried to short-circuit the electrical activity of the brain by placing strips of gold foil over the receiving areas. And I (Pribram) have produced multiple punctate foci of epileptiform discharge within a receiving area of the cortex by injecting minute amounts of aluminium hydroxide cream". . . "Such multiple foci" . . . "do not interfere with its execution (pattern discrimination) once it has been learned".

These are impressive experiments of the kind which made Lashley wonder how learning was possible. These results would not however seem to support the neural hologram more than they do the Linked Constellation concept. If anything they may be taken as favourable to the latter.

Taking Chow's experiments as an example, Chow points out that Galambos, Norton and Frommer (1967) cut about 90% of <u>both optic tracts</u>. This left, in Chow's estimation some 4,000 fibres intact out of 200,000. This is less reconcilable with the hologram hypothesis than with the "photographic" (i.e. Linked Constellation). Only this limited information was available to form the hologram, and if this were further distributed over the reception area, which is the whole point of the hologram hypothesis, then each point of the hologram could hold only the limited information. The further blow to the hologram hypothesis was dealt by Chow when he conducted an experiment in which he made combined lesions of optic tract and visual cortex. In this case the small amount of information entering the cortex is further diffused, and although in fact the registering holograms could only take effect in the intact cortex, the hologram model would imply that the pathways which would normally have distributed the hologram over the entire receiving area would be lost. In the neural hologram model the information would be reduced not by a factor of 2% but by a further factor of approximately 25% (approximately three-quarters of cortex having been destroyed). This leaves only 0.5% to support the learning.

A point made by Chow is that feature analysers must enter into the problem. "One model for explaining this redundancy is to have a small number of cell types, each responding to one specific feature of a patterned stimulus, and at the same time to have a large number of duplicates of them". It may be relevant in this respect that in describing the visual pathways, he says: "The LGD (lateral geniculate) cells send axons to the visual cortex, area 17 . . . and axons and collaterals to area 18". It may be that different feature analysers are projected to different mapping systems, an hypothesis which would accord well with Hubel and Wiesel's finding of adjacent columns which map different degrees of the orientation of receptive fields.

One may contrast the answer given by the neural hologram to those difficult facts, and the answer given by the Linked Constellation concept. First one may note that discriminable patterns are usually discriminable over the whole of their area. This is certainly true of the upright and inverted triangles used by Chow in one of his experiments. Suppose that a small area taking in part of the base of such a figure remains in the intact pathway and projection area. In Chow's figure, the area would contain in one case (the upright triangle) a fragment of a line, in the other (the inverted triangle) either nothing at all or part of the angle. Perhaps, the two comparable patterns might be a pair of corner angles as against a central inverted angle. Actually his animals could move their heads and he notes that they moved their heads and held them in a peculiar fashion. They were therefore able to focus the two base angles of the upright triangle, involving left hand and head movements, as against the centrally seen apex. Either way, it is clear that this task is equally as easy in information terms as the discrimination of the whole triangles, if not easier because of the removal of the complexity of irrelevancies.

Applying the Linked Constellation concept to the case where the tracts only were severed, by attentional inhibition of the remaining areas of projection cortex, the task became one of more or less straightforward mathcing of constellations of smaller size. That is, the input of only 4,000 fibres was compared with the portions of the already present Linked Constellations covering those 4,000 projections. This would be no harder than the original discrimination.

The case of the combined cortical and tract lesions demonstrated by Chow would be even easier to account for in the Linked Constellation model. Provided the lesions did not provide disturbing inputs, there would be no need to control attention, that is, to inhibit unwanted parts of the Linked Constellations. They would have been extirpated, leaving only the parts to be matched which were served by the remaining 4,000 fibres.

The case of the epileptic foci produced by Pribram is admittedly more difficult to explain. Epileptic foci might disturb the Linked Constellations. An argument could be made that they would disturb all Constellations approximately equally and so preserve discrimination. That is probably doubtful. However, they are just as damaging to the neural hologram hypothesis as to the Linked Constellation one. The preservation of function in face of insults points to the amazing stability of the nervous system. Probably the answer given by Chow is closest. The redundancy and specificity of analysis given by multiple representation and feature analysis protects the system from all but the greatest injuries.

An opportunity has been taken to consider the hologram hypothesis as an alternative to the spatial template hypothesis developed here. For present purposes the differences have been somewhat emphasised. In the next section some note is taken of similarities of interpretation.

VI Beyond the stimulus-response concept

In Pribram's (1971) formulation he considers his TOTE concept in relation to his two-process view of neuronal function. "Classically the function of the nervous system has been conceived as a direct input-output, stimulus-response device, expressed neurologically as a reflex arc. The ubiquitous presence of feed-back and feed-forward mechanisms (e.g. central nervous system control of receptor function) necessitates modification of this view. Feed-back and Feed-forward are best conceptualised as a Test-Operate-Test-Exit servo mechanism, an elementary logic structure of which more complex organisations are composed".

Perhaps the main difficulty in achieving a satisfactory model of the

brain is in providing an actual neural mechanism which will accomplish this. It should perhaps be said immediately that the goal is not reached in the work reported in this thesis. However, some progress may have been made, and it is perhaps worth considering the problem a little further.

Unfortunately, TOTE is a systems analysis solution, not a neural mechanism. Pribram's example, the muscle system servo mechanism, does not give any indication as to how the system would be implemented in the feltwork of the cortex. It is true that match and mismatch are represented in Spinelli's implementation of a neural hologram, as they are in my Linked Constellation model, but in itself it is not sufficient that there should be a match and mismatch mechanism. It is also necessary to show how a mismatch is detected and acted upon. In Spinelli's model as in mine, in the normal operation, it is the match which is operated upon.

The solution seems fairly clear in principle in my model but proved difficult to demonstrate. Perhaps this was because the key to its solution, the rhythmic mechanisms necessary to discern what Pribram calls the Test phase, were not discovered until late in the development of the model.

In the initial simulations an anticipatory discharge was demonstrated. This discharge, in the form of the notional Conditioned Response of the model, represented a logical expectation of the occurrence of an event. When the model goes through a familiar sensori-motor sequence, the anticipatory response should increase. Any sudden drop in this expectancy discharge should signal that things are not going according to expectation. Unfortunately, as already stated, the mechanism proved impracticable, at least in the state of the apparatus at the time. As will be reported, considerable conceptual changes were introduced late in the development. Some progress was made towards setting conditions in the model where it might be possible to detect a mismatch of the kind described. These will be related in due course.

On the conceptual level, an example suggested by Pribram may illustrate how a match-mismatch detection in an extended template system would explain how a model may transcend the Stimulus Response concept. As Oatley (1978) points out, most models which attempt to avoid that concept remain essentially limited by it, because in the event the only

discernible purpose of the model is to produce a response in relation to an input. In its standard mode of operation that is how my model works, and that mode is later termed the Stimulus-Response mode. The ability of a model to examine its own operations, and to declare a halt when a mismatch occurs introduces a new dimension.

Pribram asks what happens when we try to ascend a stationary escalator. How do we know when we move our eyeball and see the world move, that it is really stationary? Both of these examples demonstrate the results of match and mismatch. In terms of the Linked Constellation model, when we step onto the escalator, the sequence of Linked Constellations registering the familiar sequence of events which have occurred in the past, is reactivated. There is an anticipatory energisation of the next Constellations in the sequence. These register from a former occasion the feed-back to be expected from a moving escalator. The actual feed-back from the stationary escalator is quite different. Consequently there is a mismatch. It is this mismatch which a brain is capable of immediately detecting.

In the case of a movement of the eye-ball resulting from our own actions, the effect is the reverse. At first the unaccustomed conjunction of events, i.e. finger pressure and a moving world, may lead the neural apparatus to attempt a wrong match. The strong matching between Linked Constellations in the system registering retinal motion and, say a moving head, will lead to a mismatch, because the signals from the muscle and joint receptors will fail to show head movement. The first time we attempt this, the system may be somewhat upset. As soon however as we have established Linked Constellations registering the conjunction of finger on eyeball and a moving retina, a satisfactory matching occurs, and we "understand" the situation.

The factors limiting the stimulus-response concept are well stated by Horridge, 1968. "Aspects of the theory which have been available if not acceptable for the whole of this century are that simple reflexes and two neuron arcs have no place in the elementary nervous system. Even in quite advanced vertebrates we find that co-ordination mechanisms are dominated by local pacemakers forming centrally controlled programs which can be released or inhibited by particular stimuli . . .".

According to the philosophy of my model, the essential concept by which the Stimulus Response view may be transcended is by considering the

nervous system as having developed from the single celled creature as a three-neuron rather than a two neuron arc. Two of the neurons have taken over the functions of internal transport. In the single celled creature, the stimulus affects the organism directly. In the multicellular organism, the sensory neuron carries the signal from the now distally placed sensitive apparatus, the sense organ, up to the integrative apparatus, the brain. Similary, the motor neurons carry the signal from the central mechanism to the specialised motor organ which is now distally placed at the periphery of the multicellular organism. The central neuron must be seen as more than a simple connecting device between input and output, if the apparatus is not to remain a stimulus response device. Elaborated from one neuron into many, the central neuron remains between input and output. The Linked Constellation preserves the integrity of the central neurons, which in the form of an integrated structure within the organism, is the central core of the organism.

Although the simulations fell short of demonstrating attributes which go beyond the stimulus-response concept, much of the later development of the model was pointed in that direction.

VII Conclusion

Prior to describing attempts to extend the model, some rather general matters have been discussed. Some changed aims and a proposed use of the computer as a tool for theory development were discussed.

Some models which had received only cursory mention in the introductroy chapters were more fully described in relation to the present effort.

A distributed memory approach, the neural hologram, was compared with the extended template concept of the model presented here.

Finally some attempt was made to consider ways in which a model of the brain might conceptually go beyond a stimulus-response formulation.

CHAPTER 11

GENERAL DESIGN OF THE EXTENDED SIMULATIONS

I Introduction: appetitive-consummatory sequences

As discussed in the previous chapter, amongst the principal aims of extending the model was to represent motor behaviour more adequately. In this chapter it is proposed to describe the main features of the simulations which attempted to achieve this aim. The remarks in this chapter are not intended as a theoretical discussion of behaviour, but merely to enable the reader to follow the rationale of the arrangements in the simulations.

The model was directed primarily at representing appetitive-consummatory sequences. Activities such as seeking and obtaining food, shelter, and a mate, are prominent in the behaviour of animals. Such sequences may be described without theoretical commitment as related motor activities which culminate in the elicitation from the environment of specific patterns of stimulation which release inbuilt responses. They also have the characteristic of approach rather than withdrawal. (Maser, 1973; Konorski, 1967). The initial simulations already included a notional neural apparatus for producing specific responses, which might be considered analogical to consummatory actions. This made the extension of the model by a representation of appetitive behaviour a convenient subject. In relation to the general aims of the model appetitive-consummatory sequences give an opportunity for dealing with the integration of what may be somewhat loosely termed "learned" and "innate" behaviour.

The specifically elicited responses represented in the initial simulations could also represent a response caused by a harmful or other unwanted stimulus, e.g. shock. It was in fact intended later to extend the model further with an aversive-behaviour system paralleling that for appetitive behaviour. Unfortunately time did not permit this. However, there are aspects of appetitive behaviour which suggest the necessity of including a representation of stimuli which are innately

recognised as undesirable and which produce appropriate action, e.g. spitting out unwholesome food. For this reason provision was made in the simulation for Protective Responses (see below).

The integration of "learned" and "innate" behaviour has been extensively studied in Instrumental Conditioning (Skinner's operant conditioning or Konorski's Type II Conditioning) and in maze running and problem box experiments. Many such experiments represent appetitive-consummatory sequences, for example, pressing a bar to obtain food. The behaviour of the model may therefore be analogically compared to experiments of this kind. The model was however intended to have a wider and more informal interpretation. A kitten walking to its accustomed spot to find its saucer of milk is a case in point. The Responses in the model are, of course, mere indications of a principle of organisation. In the broadest sense much of behaviour can be understood in terms of the integration of "operants" and "fixed action patterns" released by specific stimuli (see Hinde, 1970, Chapter 25).

To model an appetitive consummatory sequence it was decided to provide four sub-divisions of Motor Response in place of the single kind of Response unit included in the initial simulations. These were termed: (i) Instrumental Responses; (ii) Preconsummatory Responses; (iii) Consummatory or Final Responses; (iv) Protective Responses.

These divisions were not intended as hard and fast theoretical categories. Such rationale as they possess will appear in due course. In part they were suggested by the form of the initial simulations. There are certain effector actions, such as those associated with eating, drinking and copulation which are strongly under the control of specific sensory input, e.g. food or water in the mouth, intromission, and which customarily form the climax of behavioural sequences. This character seemed to be reasonably well represented by the motor organisation provided in the initial simulations. In an early chapter it was suggested that neurological models should differentiate the character of motor response, say the difference between manipulation of a lever and swallowing food.

In part the divisions were suggested by functional distinctions uncovered by research. For example, there is evidence that consummatory and

CORTEX

Sensory		Motivational	Motor							
Projection 1	Projection 2		Protective		Preconsummatory		Consummatory		Instrumental	
			Do	Not	Do	Not	Do	Not	Do	Not

SUB CORTEX

P1	P2	M	Pr.	Prec.	Cons.	Instr.

CONTROL

P1	P2	M	Prot.	Prec.	Cons.	Instr.

Figure 11.1 Showing the main Arrays of Neuronal Pools provided in the Cortex of the extended models. Each Array in the Cortex corresponded to an Array in the Sub-cortex, and hypothetically the sub-cortical cortical relationships would be dealt with in a corresponding location in the Control Centre. An except were the motivational arrays, where the relationship was less easily defined.

appetitive (operant) responses are sensitive to different pharmacological agents and that the two can be dissociated both in activation and suppression. (Review in Margules and Margules, 1973). In part the divisions appear to reflect a more or less obvious difference in behavioural effect of phases of the appetitive consummatory sequence. The assumption was that these differences of functional relationship to the environment would have led to the evolution of corresponding differences of mechanism or differences of emphasis of aspects of a common mechanism.

It will be remembered that the extended simulation was not the embodiment of a pre-existing theoretical model of the brain. It was to be used as a tool to develop such a model. In effect the principles of action of all four divisions of the motor apparatus would be the same. Their separation into different Neuronal Arrays permitted variations and special modifications to be introduced into each division.

II Extension of the notional neuro-anatomy

The main aim of this section is to acquaint the reader with the general form of the extended model .

The basic layout followed that of the initial simulations, being based upon Figure 5.10 and 5.12. One of the main extensions was to provide for the differentiation of motor actions under the four subdivisions described in the previous section (Figure 11.1).

A. Instrumental Response Array

The Motor Response apparatus in the initial simulations consisted of a Sub-cortical Array and a Cortical Array. The Sub-cortical Array contained notionally inbuilt representations of the motor actions in the form of a Preset Linked Constellations. The activation of a Preset Linked Constellation led to activity in the notional Fibres which acted upon the Cortical Array and so a Learned Record of the occurrence of the Motor Response was laid down in the Cortical Array.

The overall organisation of all four kinds of notional motor actions followed this basic pattern. For each of the four sub-divisions a Sub-cortical and a Cortical Array was provided. When a particular Neuronal Pool in the Sub-cortical Array was active, its Efferent Fibre

became, after passing through the notional Thalamus, or Control Centre, the Afferent Fibre to the corresponding Pool of the corresponding Cortical Array.

The only kind of Response provided for in the initial simulation was that which occurred primarily following the occurrence in the input of an Innately Recognised Pattern (IRP). The latter activated the Preset Linked Constellations which included representation of Response. Responses elecited by this kind of Preset Constellation action will be termed specifically elicited responses. Essentially this mechanism continued to apply unaltered in the extended model to the Preconsummatory, Consummatory and Protective Responses. Although these three divisions had this in common there were of course intended to be other differences, many of which were not explored. The set of Instrumental Responses was to differ from them more fundamentally. Although basically the same Sub-cortical-Cortical principle would apply, it was aimed to make emission of Instrumental Responses less tightly bound to the appearance of Innately Recognised Patterns in the Input.

The Instrumental Responses may be regarded as representing operant actions. The reader may wish to object that the distinction between operants and respondents is not necessarily a function of either their neurological or psychological mediation. The categorisation of behaviour as operant or respondent may be regarded as a function of experimental contingency (see Mackintosh, 1974, p.4) or as referring to the most usual application of particular responses rather than their intrinsic operation (Moore, 1973; Jenkins, 1973). These objections are not, it is suggested, swept aside by the use which is made in the simulations of convenient divisions. To pursue such issues would introduce unnecessary complication into the presentation and would in any case go beyond the scope of the project.

The simple representation of the Instrumental Array is intened to stand in place of a great deal of complexity in the brain. It is perhaps worthwhile giving some indication of this complexity, and its manner of symbolisation in the model.

In the approach of the author all motor actions are morphologically determined; that is to say, they are innately programmed in conformity with the anatomical and physiological attributes of the particular

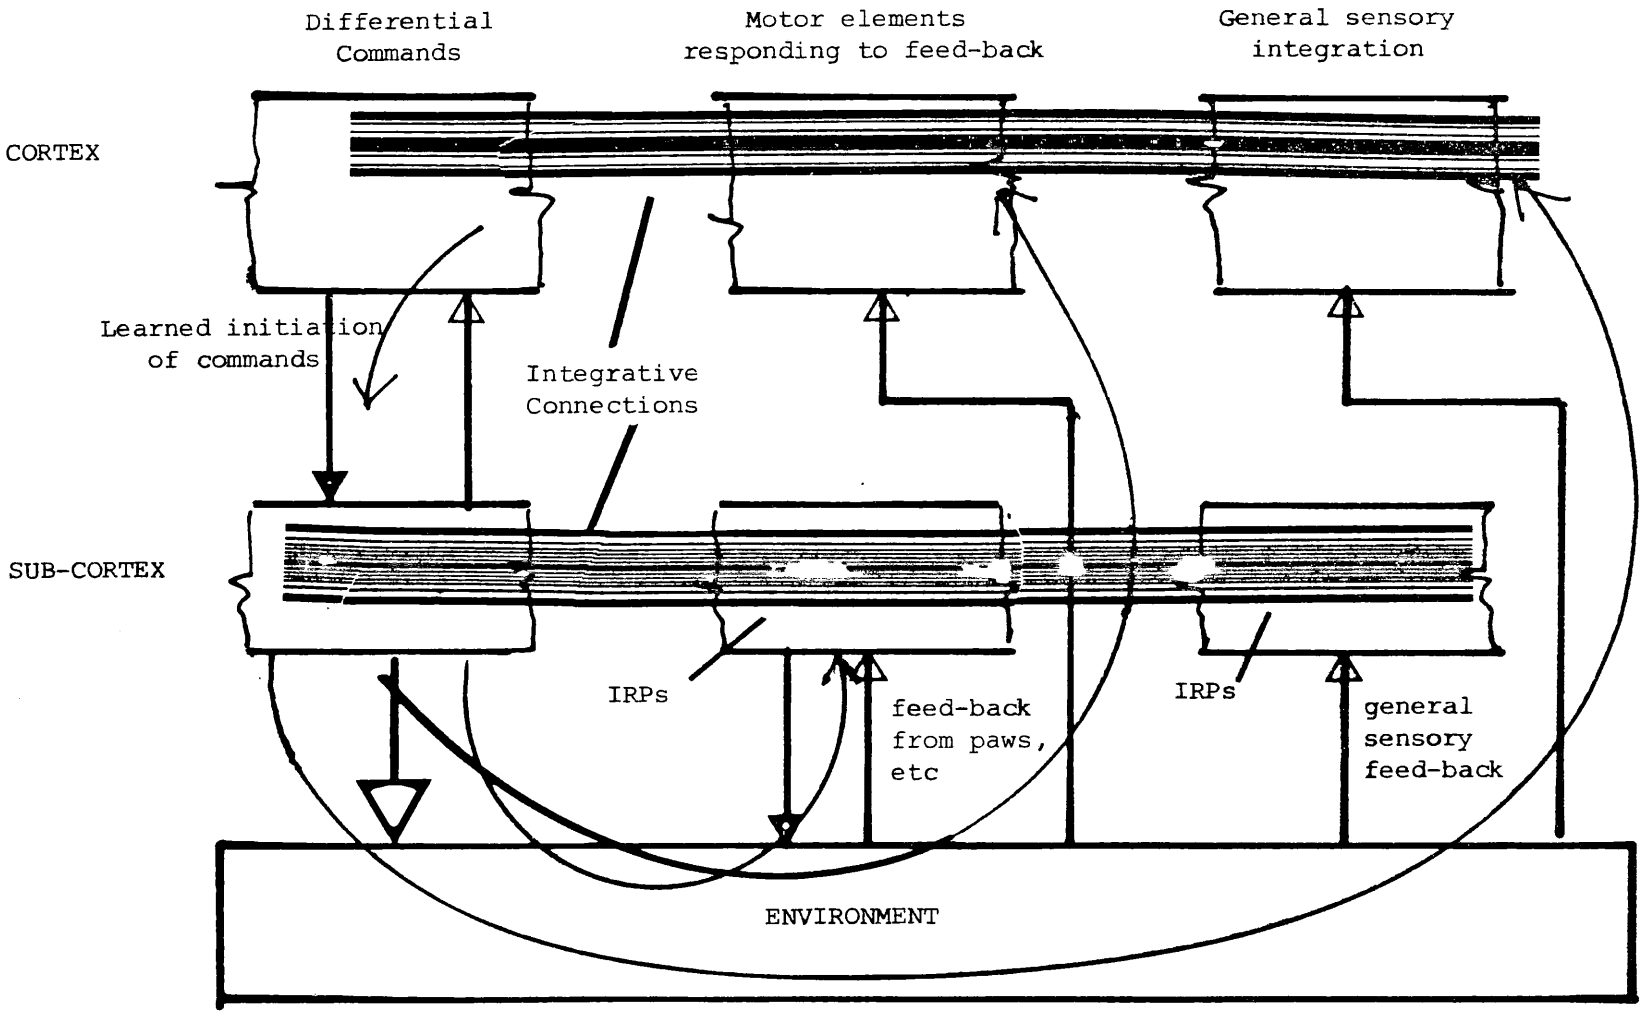

Figure 11.2 Illustrating co-ordination of various levels of feed-back from the environment in organising complex actions. The small loop represents immediate sensory control, as in feed-back from paws during walking. The intermediate loop represents feed-back from various motor elements, tactile, proprioceptive, muscle spindles. The large loop overall control of command behaviour resulting from the general "gestalt" feed-back from the environment.

animal. The kitten walking to its milk has been mentioned as the kind of motor action intended to be represented by a Response of the Instrumental Array. This suggestion may be slightly amplified.

Walking is not independent of species-specific stimulation. "Stepping" behaviour of cats is dependent upon sensory feed-back from the pads of the paws (Gray and Lal, 1965) and can occur in the absence of cortical control (Shik and Orlovsky, 1976). Walking to a particular spot is also dependent on control by species-specific neural arrangements. Rats are known to approach dark places rather than lighted ones. Upon the first approach or avoidance of an object, the emission of the response controlling the locomotion has probably considerable species specific influences; e.g. smell, colour and manipulative possibilities. Other examples of approach and avoidance may be influenced by experience. Moreover walking to an object is not so much simply an act of walking, but of walking integrated with the other movements, e.g. orienting the body.

The aim of the model is to represent appetitive behaviour, that is, essentially approach. In the animal such approach can be considered as a sequence of a number of co-ordinated morphologically controlled motor movements. These co-ordinations may be considered to fall under a number of different classes, each of which is probably also integrated by innate morphological neural arrangements. Thus walking to the left, walking to the right or walking forward are probably innately programmed variations of a basic organisation. Although this account clearly greatly simplifies the neurology, one may concede that walking to the left, forward or to the right, or turning tail are complex integrations subject to relatively simple commands (Isaacson, 1974). This assumption is supported by such clear relationships as the turning of the head in response to specific stimulation of certain pathways (Hassler, 1978). Nevertheless such actions are subject to learned adaptations and to control by sensory feed-back. (See Figure 11.2).

In the nervous system the various combinations of integrated movements used to approach an object from a given position would be mediated by activity in an enormous number of fibres. This activity would consist of sequential configurations of activity. Parts of these configurations would be controlled by feed-back at lower levels: for example the

stepping movement of the cat is controlled by sensory feed-back from the paws. These aspects of the act of approach, however complex, may be seen as species specific stimulus-response arrangements, which, in principle at least, could be mediated by Preset Constellations. What we are concerned with is the activation of that variation of the integrated action which leads the animal in one direction rather than another, i.e. to approach the object. It is this motor decision which is intended to be symbolised by activity in one or other of the Efferent Fibres of the Instrumental Response. This representation is only a pointer to the principles of operation of a large number of fibres in the nervous system: as in most models, the activity of a few fibres stands for the activities in a very large number of fibres.

Essentially the model, as so far described, suggests that particular co-ordinations of motor actions, e.g. moving to the right or left or forward, occur initially at a sub-cortical level at times which are less dependent upon "eliciting stimuli" than are consummatory responses, such as eating or drinking or copulation. When such Instrumental Responses occur then what follows is similar to what has been represented of elicited responses in the initial simulations. Their occurrence and the occurrence of accompanying and preceding sensory stimulation from the environmnet is registered in the cortex.

It may be noted that this function of the motor cortex is compatible with the fact that stimulation of the "electrically sensitive motor cortex" produces isolated movements rather than co-ordinated movements (Towe, 1973). The Response represented by activity in one or two Neuronal Pools of the simulation symbolises the activity in the brain of a large number of Neuronal Pools co-ordinated by Linked Constellations. According to the model, the cortex merely registers a co-ordination provided to it by sub-cortical emissions. That co-ordination is probably or primarily innately engineered at sub-cortical level (cf. Flynn, 1976). Stimulation of the motor cortex would therefore merely stimulate a few odd Neuronal Pools.

B. Preconsummatory Responses

It was felt desirable to distinguish between different function, and consequently probably differing mechanisms, of that class of responses elicited by species specific stimulation. In appetitive-consummatory sequences one can, in particular, distinguish between the specifically

elicited activity which brings the sequence to an effective end, e.g. swallowing food or drink, and those stimulus-released responses which are on the way to a final act, e.g. chewing, pecking in chicks.

The mechanisms of both Preconsummatory and Final (consummatory) Responses were essentially the same as that used in the initial simulations for the single class of Response. At Sub-cortical level the emission of the Response depended upon the occurrence of an IRP in the input. The distinction between the two sets of Responses did not relate to their status as fixed action patterns notionally elicited by specific stimulation, but rather their functional relationship to behaviour and to other aspects of the model. In this respect it may be mentioned that the reasons for introducing them were rather different from the character which they later took on.

Briefly, the original concept was related in a general way to the ethologists concept of a hierarchy of responses leading to a behavioural goal. It was felt that the introduction of a notionally inbuilt sign-post on the way to the final goal of a sequence would assist the model, as it was presumed to help the animal, to find the goal. Later in the development of the model, the motivational possibilities of specifically elicited fixed action patterns, acting as innately programmed instrumental acts on the way to final goals, was more fully appreciated.

"Eating, drinking and sexual behaviours are highly stereotyped activities, the components of which are recognised by ethologists to be fixed action patterns" (Maser, 1973). The presence of a pellet in the mouth may be taken to be what is aimed at in providing Preconsummatory Responses. The sensory feed-back of taste, smell and palpation by the mouth surfaces release mastication and salivation. The sensory feed-back may be regarded as IRPs releasing fixed action patterns mediated by Preset Constellations. The correct consistency of treated food-mass leads to a bolus being passed back to the throat, another Preconsummatory Response. The bolus releases swallowing: apparently the texture of the bolus is critical to the act of swallowing. The latter was originally termed the Consummatory Response, a term which was subsequently amended to the less committed one, Final Response. Of course, the term Final is to be taken with suitable reservations. Further events take place in the stomach and metabolism. To consider these would take us too far away from what was done.

At the beginning of extending the simulations, the Preconsummatory Responses were principally used to demonstrate the place of Protective Responses in the scheme of appetitive behaviour. If a Preconsummatory Response were emitted to some other input than the IRP intended to elicit that Response, the programme, simulating the environment, could present a special IRP representing some harmful effect. The protective Response (see below) would then be elicited. This might represent ejecting an object if chewing movements disclosed it to be other than food.

Later in the development of the model, and principally as the result of the influence of Konorski's discussion of the motivational effects of different stages of the eating process, the relationship between Preconsummatory and Consummatory Responses was used to control motivational mechanisms.

C. Consummatory Responses

Little more need be said of these than that they represented the specifically elicited Responses, which at one level of activity at least, closed the behavioural sequence. In simple terms they represented swallowing, drinking or copulating. As just mentioned, in later versions of the model they differed from Preconsummatory and Instrumental Responses in their relation to motivational mechanisms. Briefly, whereas Consummatory Responses induced a reduction of Excitation from Motivational sources, Preconsummatory Responses induced an increase (compare Konorski, 1967). Moreover, in terms of reinforcement terminology, Preconsummatory Responses, or rather the occurrence of the IRPs which induced them, were the primary reinforcing events. In early versions, Preconsummatory Responses released behaviour which increased the probability of a Consummatory IRP, e.g. in a maze situation a notional "standstill".

Like the Preconsummatory Respones they were implemented by mechanisms similar to those for the single class of Responses in the initial simulations: Cortical and Sub-cortical Array with Preset Constellations written into the sub-cortical Array.

D. Protective Responses

It has already been explained that the Protective Responses were intended to represent what occurs in an appetitive sequence when, instead of the IRP of a Preconsummatory or Final Response, some other

IRP representing, say, noxious food, is presented. As mentioned, it was intended eventually to incorporate an aversive behaviour system paralleling the appetitive one. Unfortunately time and resources made this impossible. The protective Responses represented aspects of aversive behaviour which it was thought necessary to include in the model to give a reasonably complete picture of appetitive behaviour.

The program was designed so that if the model emitted a Response which was wrong, in the sense that in the situation modelled it would not lead to the goal, then the model could be "punished" by being presented with an IRP which elicited an aversive response. The Protective Responses differed in their conceptual base. One of them was typical of the notional results of taking a wrong Instrumental Action. In terms of the Instrumental Responses signifying moving Left, Right and Forward, the IRP eliciting the Protective Response could be considered as representing a "bump" obtained by crashing into the wall of a maze and the Protective Response as moving backwards. Another Protective Response was related to the emission by the model of a Preconsummatory Response to some input other than the IRP which was supposed to elicit that Response. This represented what might happen if the animal attempted to chew a hard object e.g. regurgitation.

In fact, although these Responses worked well, they did not play a large part in the later development. It may be said here that this was because they were withdrawn at a time when the model was simplified in an attempt to cure the troubles with Excitation and Inhibition already mentioned. In the early versions of the model they were quite effective in producing a (superficially perhaps) life-like behaviour. If the model emitted a response considered by the programmer to be wrong, it was presented with an input pattern representing the "bump" and thereafter the offending Response failed to appear for some time.

E. Computer implementation of Response Arrays

In summary of the foregoing Sections, the Motor Division of the Sub-cortical Array of the extended simulations was divided into four sub-sets of Neuronal Pools. In turn each of these sub-sets consisted of 4 Neuronal Pools (total 16 pools). They were referred to respectively as the Sub-cortical Instrumental, Preconsummatory, Final (or Consummatory) and Protective Arrays (see Figure 11.1).

Corresponding to the 4 Neuronal Pools in each Subcortical Response Array, 4 Neuronal Pools were provided in a corresponding Array at Cortical Level, i.e. Cortical Instrumental, Preconsummatory, Final and Protective Arrays. Each of these Cortical Arrays possessed in fact 8 not 4 Neuronal Pools in toto. The first four of each Array, as just stated, corresponded to the Neuronal Pools in the Sub-cortical Array of the same name. They represented the neural emission giving an effector response. The second four Neuronal Pools in each Array represented the neural emission required to suppress an effector response, i.e. Not-do Pools. This was upon the principle already discussed in relation to the initial simulations.

As mentioned, the actual arrangements differed from version to version of the simulations. In some runs, for example, eight Neuronal Pools were provided in each of the Sub-cortical Arrays as well as in the Cortical Arrays. In these versions there was a representation of effector and suppressor action at both Sub-cortical and Cortical Level. The essential principles of operation did not, however, differ greatly.

The main Motor Connections were, as already mentioned, similar to those of the initial simulations. Essentially, the Efferent Fibres of the various Sub-cortical Arrays passed to the Control Centre (Notional Thalamus) from which they issued as Afferent Fibres (Motor Afferents) to the appropriate columns of the corresponding Cortical Array. As in the case of the initial simulations, return connections from the Cortical Arrays to the Control Centre and back to the Cortical Arrays were necessary to ensure registration of Responses emitted by the latter. These arrangements are omitted from the Figure and will be illustrated when the mechanisms of operation are more fully described.

F. Extensions to Sensory Arrays

The basic arrangements were as in the initial simulations. However, at the Cortical Level the sensory representation was increased by adding an additional Array of Neuronal Pools equal in size to that provided in the initial simulations.

This duplication of the Cortical Sensory representation allowed for the introduction of an Uttley-like delay, that is a notional difference of latency of pathway, to each of the two Arrays (see Chapter 4, Section III, B, and in Chapter 5). Multiple representations have been found in the cortex of projections of a number of modalities.

Celesia (1976) found differences of latency between the arrival of signals from acoustic stimuli in two separate projection areas, a primary and secondary auditory projection. The effect of the notional difference in latency of the pathways to the two Arrays in the simulation was that the same Input Pattern, notionally presented to the model at a given time, was processed in one Array several cycles before it was processed in the other. Consequently the Linked Constellation formed at any particular Instant represented both the Input Pattern in the input at one given moment and that appearing in the input somewhat later. This is the appropriate incorporation in the present model of Uttley's principle.

In respect of temporal delays, the present model differed from the initial simulations. In those, a single Uttley-like delay of 1 Instant, between sensory and motor pathways, was later removed to demonstrate the independent temporal sequential aspect of the basic Linked Constellation concept. This aspect was later seen to be of less importance, since the same effects could be obtained by Uttley-like delays. In fact the two principles were seen to be variations of degree of a single principle of lack of strict simultaneity of action of operation of a system. In the extended model a number of notional delays of this kind were introduced and will be described in due course.

It was intended to increase the Cortical representation of Sensory Input further by the introduction of further Sensory Arrays. The purpose was to demonstrate that a fuller and different sensory representation could be registered at the Cortex from that Innately represented at the Subcortical Level. Some experiments were carried out and showed that, at the least, the model still operated in a similar manner. However, the additional representation had to be withdrawn to make way for other developments of the model utilising the limited computer resources.

Figure 12.2 shows the revised sensory arrangements and temporal delay.

G. Motivational Arrays

The function of notional motivational mechanisms will be described in a later chapter. It is, however, necessary to know that a number of Neuronal Pools were provided to mediate motivational aspects of the system.

Two sets of Motivational Neuronal Pools were provided, one at Sub-cortical and one at Cortical Level. However, the relationship of these Arrays differed from those of the Sensory and Motor Arrays previously described. The Motor Arrays, for example, had "corresponding" Neuronal Pools at each Level, with characteristic connective arrangements. In comparison, the two Motivational Arrays (cortical and subcortical) were relatively independent in their action. In fact, the notional function of the Cortical and Sub-cortical Motivational Arrays differed from each other. It must be mentioned that it is not intended to imply that the function of each of these sets of neuronal pools, insofar as they might actually exist in the brain, would be devoted exclusive to Cortical or Subcortical locations. As with the allocation of Learning and Innate programming between Cortex and Subcortex, the location of the Motivational Arrays was a simplification of convenience. Nevertheless, in this case also it is thought that there is a closer relationship between certain motivational functions and cortex on the one hand, and between other motivational functions and sub-cortical centres on the other.

As in the case of the other Arrays, the Motivational Arrays consisted essentially each of 4 Neuronal Pools, in some versions this being increased for representation of Suppression Pools to 8.

There was however an important difference of principle in operation which will be described more fully when the Motivational mechanisms are dealt with. Essentially this concerned a distinction between two modes of operation of Afferent Excitation of the Motivational Pools at the Cortical Level. In one mode, the Integrating Neurons of the Pool were driven but did not form Integrative Connections, so in that mode the Afferent Fibres provided Motivational Excitation but did not provide Motivational registration. In a second mode Integrative Connections formed in the customary manner of the model. This may be regarded as representing a conjunction theory of synaptic or other integrative connectivity (see Eccles, 1972).

III Revisions of notional neuro-physiology

The neuro-physiological aspects of the model remained essentially the same as in the initial simulations. The Linked Constellation concept

was simulated in almost precisely the same manner as in the initial simulations. Excitation and Inhibition were treated in essentially the same way, although there were very big differences introduced into the notional flows of excitation to Excitatory and Inhibitory Neurons, concerned with the control of their levels of activity. Most of the revisions of neuro-physiology were not seen as matters of physical principle. They were attempts to deal with problems of control within the model, which have already been mentioned a number of times.

It may be mentioned here that these problems turned out to be the principle obstacle in obtaining satisfactory performance from the simulations. Many features of the model had to be removed before they had been effectively demonstrated so as to simplify the model in order to understand its weaknesses better. It was intended to reintroduce them at a later stage, but time did not permit. At one stage, the difficulties were considered to be the Achilles heel of the theory of neural integration. They were not solved until an adequate concept of the difficulties had been developed. The solution involved notionally interrupting the Afferent flows intermittently, thus producing rhythms in the processing within the Neuronal Pools. This solution was not seen until very late in the development of the model.

These problems and the steps taken to try to remedy them will be left until a later chapter. They really pose a different set of problems from the mechanisms used to implement sensory-motor integration. The latter will be dealt with in the next chapter. The presence of these problems does, however, serve to explain why more sophisticated implementations of the motor mechanisms were not achieved during the course of the development of the computer simulations.

One or two points of changes of principle introduced into the neuro-physiological aspects are mentioned here for completeness.
(1) The control of levels of Excitation and Inhibition in each Neuronal Array of the model was dealt with independently. This was to prepare the way for the introduction of attentional mechanisms arising from the notion of differential emphasis on different portions of an extended

template system (Chapter 2, Section III).

(2) A binary coding of afferent fibres was introduced at the Control Centre (i.e. notional Thalamus) so that they were either active or inactive with no intermediate states. Whereas in the initial simulations, the Spike Frequencies of the Efferent Fibres of the Sub-cortical Motor Arrays were applied as re-Afferent Spike Frequencies to the Cortical Arrays, in the extended model the set of Spike Frequencies of the Sub-cortical Efferent Fibres was first converted into a set of binary digits by the introduction of notional collateral inhibition and relay neurons. It may be recalled that Sensory Afferents were already coded in binary form. This meant that the activity of Afferent Fibres to any Neuronal Pool in the system, from whatever source it might stem, could be treated as a set of notionally active or inactive Fibres. The notional Thalamus (Control Centre) could then encode into Spike Frequencies the activity or inactivity of all fibres in terms solely of the Discriminative and Diffuse Volley Values as described for the earlier simulations.

(3) In the Neurons in the initial simulations, Excitation and Inhibition Counts were set initially at zero and later transactions were recorded from this zero base. This led to an imbalance in the effects of reduction compared with additions of Excitation and Inhibition. When some of the Integrating Neurons notionally lost more Inhibition or Excitation than they received, their Excitation and Inhibition Counts would have fallen below zero, had an arbitrary rule setting them to a minimum of zero not been introduced. In the extended simulations Excitation Count and Inhibition Count were initially set at an arbitrary Optimum Excitation and Optimum Inhibition Level. Flows of notional Diffuse Excitation and Diffuse Inhibition into the system worked in conjunction with notional destruction and dissipation of Excitatory and Inhibitory factors to attempt to maintain the system at the Optimum Levels. Reductions contained both a fractional and constant component. These systems of additions and deductions were intended to represent a combination of intracellular metabolic processes and extracellular activation of both Excitation and Inhibitory Neuron systems. Signal volleys then modulated the Opitmum Level, and when a particular pattern of input ceased, the system attempted to bring levels back to the optimum level.

(4) In addition to the overall Inhibition system, collateral and recurrent Inhibition was introduced between and within Neuronal Pools. When

a fresh Input Pattern reached an Array it first produced a collateral Inhibition of non-activated Pools. After a slight temporal delay, the recurrent Inhibition was effective, bringing the levels of Inhibition in all Pools back to one level. Thus when an Input arrived it first effected a heightened contrast between activated and non activated Pools, and if the pattern continued to be active, its effect was progressively diminished by the recurrent Inhibition.

IV Conclusion

In this chapter the general design of the extended simulations has been described. The aim was to demonstrate a sensori-motor integration in the form of the integration of notional learned Instrumental actions and notional species-specific innate behaviours. The original Motor Response Arrays were subdivided at both Cortical and Sub-cortical levels, to provide facilities for differences of mechanism between actions performing different functions in relation to the animals relations with its environment. Instrumental, Preconsummatory and Consummatory Responses, and Protective Responses were each provided with an Array of Neuronal Pools at both levels. Motivational Pools were also provided: the Motivational Pools at Sub-cortical level were however to be independent of those at Cortical level; in this respect the Motivational Pools broke away from the general Sub-cortical-Cortical relationship expressed in the model.

Changes in notional neuro-physiology were directed to solving problems of control. Certain major changes were described which were not considered to be matters of principle. These changes were chiefly the result of modifying a system of initially'empty' neurons, to one in which the neurons were maintained at an optimum level of Excitation and this level was modulated by Excitation and Inhibition. Other major changes which were matters of principle have only been briefly mentioned. Problems relating to the control, of this system will be discussed in a later chapter when the scheme by which these problems were solved will be described.

CHAPTER 12

GENERAL PRINCIPLES OF OPERATION OF THE EXTENDED SIMULATIONS

In this chapter it is proposed to describe further the main mechanisms which were introduced in order to operate the various extensions to the notional neuro-anatomy of the model described in the previous chapter. These were mainly concerned with the Instrumental Response Array and the Motivational Arrays.

I Response emission control: general mechanism

It will be recalled that in the initial simulation two kinds of response of the system could be detected. In the initial simulations these were suggested to be analogical respectively to the unconditioned Response and the Conditioned Response seen in animals in a Classical Conditioning situation. In future, for convenience in referring to them without theoretical implications, they will be referred to as the Directly Elicited Response and the Anticipatory Response.

The Response mechanisms of the extended model can perhaps best be followed in relation to the detection of these two kinds of response of the system. In this section only the Directly Elicited Response will be considered.

It will be recalled that in the initial simulations, when an IRP actually appeared in the input, a massive motor response occurred. The mechanism was briefly as follows. The effects of the IRP, reaching the Sub-cortical Arrays, activated the Preset Constellation to which that IRP related. Excitation was transferred from the sensory neuronal pools to the selected Motor Neuronal Pools represented in that Preset Linked Constellation. At the same time as this reactivation of the Preset Constellation in the Subcortical Array was taking place, a sensori-motor registration of the IRP and its Sub-cortical Response was being laid down in the Cortical Array. On subsequent occasions, when the effects of the IRP reached the Cortical Array, the Learned Constellations were reactivated at this level also. Consequently a

Directly Elicited Response was also eventually produced at the Cortical Level. This joined with the Sub-cortical Response and was sent notionally to the periphery. The effectiveness of the IRP is in this way enhanced by Learning.

This same principle was applied to the Consummatory, Preconsummatory, and Protective Arrays of the extended model. (It could not be applied without considerable modification to the Instrumental Array, since the Instrumental Responses were not elicited by IRPs (see next section)).

Differences between the mechanism of detection in the initial simulations and the extended simulations were chiefly matters of more sophisticated notional Inhibitory Control. When an IRP reactivated a Preset Constellation, the Excitation was transferred to member Neurons of the Constellation located in the Pools of the Motor division. A similar process occurred in the Learned Constellations of the Cortical Array. In the numerical examples given in Chapter 5, Section IV, it was demonstrated that, although all the Neurons on the Motor side might be activated in differing degrees, by adjustment of the level of Inhibition across the Neuronal Pools of the Motor division it could be arranged that only the most highly Excited Neurons would be permitted to Discharge. This ensured, in principle, that the only Efferent Fibres to show a Discharge would be those containing Linked Neurons from the relevant matching Linked Constellations.

In the event it will be remembered that the results reported of the initial simulations showed that this control of Inhibition was not adequately achieved in the initial simulations. This problem was part of overall problem of control of Excitation and Inhibition levels which will be discussed separately.

II Response emission control: Anticipatory response of system

A. Stimulus-Response chains

It was demonstrated in the initial simulations that after an appropriate training session, if a Cue Sequence of patterns were presented without being followed by the customary IRP, a response appeared in the relevant Efferent Fibres of the Cortical Motor Array. This response was smaller than the Directly Elicited Response, but reliably detectable.

Moreover, even during trials when the IRP was actually presented, an anticipatory Descharge could be detected earlier than the occurence of the Directly Elicited Response.

The anticipatory discharge, analogical to a Conditioned Response, was demonstrated to be result of an anticipatory effect of the system. Anticipation in the system took the form of a rise in the Excitation of the Linked Constellations about to be stimulated when a familiar sequence was repeated. The anticipatory rise could be observed by inspecting a monitor printout of Excitation Counts. It could also be detected as an observable effect in the Motor Efferent Fibres. In the initial simulations these Discharges were regarded as analogous CRs and arose from reactivation of Sensori-Motor Constellations which had been formed on rather special occasions, i.e. occasions when an IRP in the Input had reactivated a Preset Constellation and elicited an UCR. However, such an anticipatory Discharge should be able to be detected whenever a sequence which was undergoing reactivation contained a Sensori-Motor Constellation. This would apply if on an earlier occasion the emission of a Motor Response had been registered in the Cortex in conjunction with a Sensory Pattern, even if the initiation of the Response had been effected by some mechanism other than reactivation of a Preset Constellation by an IRP.

The Anticipatory Response of the system was regarded as the basis of motor decision underlying the integration of motor action. As argued in Chapters 2 and 3, such an anticipatory response might mediate a sophisticated control, of many aspects of behaviour: for example, the TOTE paradigm of Pribram. The primary aim, however, was to demonstrate its efficacy by using it to control the motor decisions of the model in respect of Instrumental Responses.

The most basic aspect of the production of operants may be described as Stimulus-Response chaining. Pierrel and Sherman (1963) trained food-deprived rats to successively climb spiral stairs, push down and cross a drawbridge, climb a ladder, pull in and pedal a car through a tunnel, climb stairs, and perform many further feats to reach food (Hinde, 1970, p611).

To see how this chaining might be mediated in the model, for the moment let the problem of production of Instrumental Responses at the Sub-cortical Level be ignored, and assume that an Efferent Discharge occurs

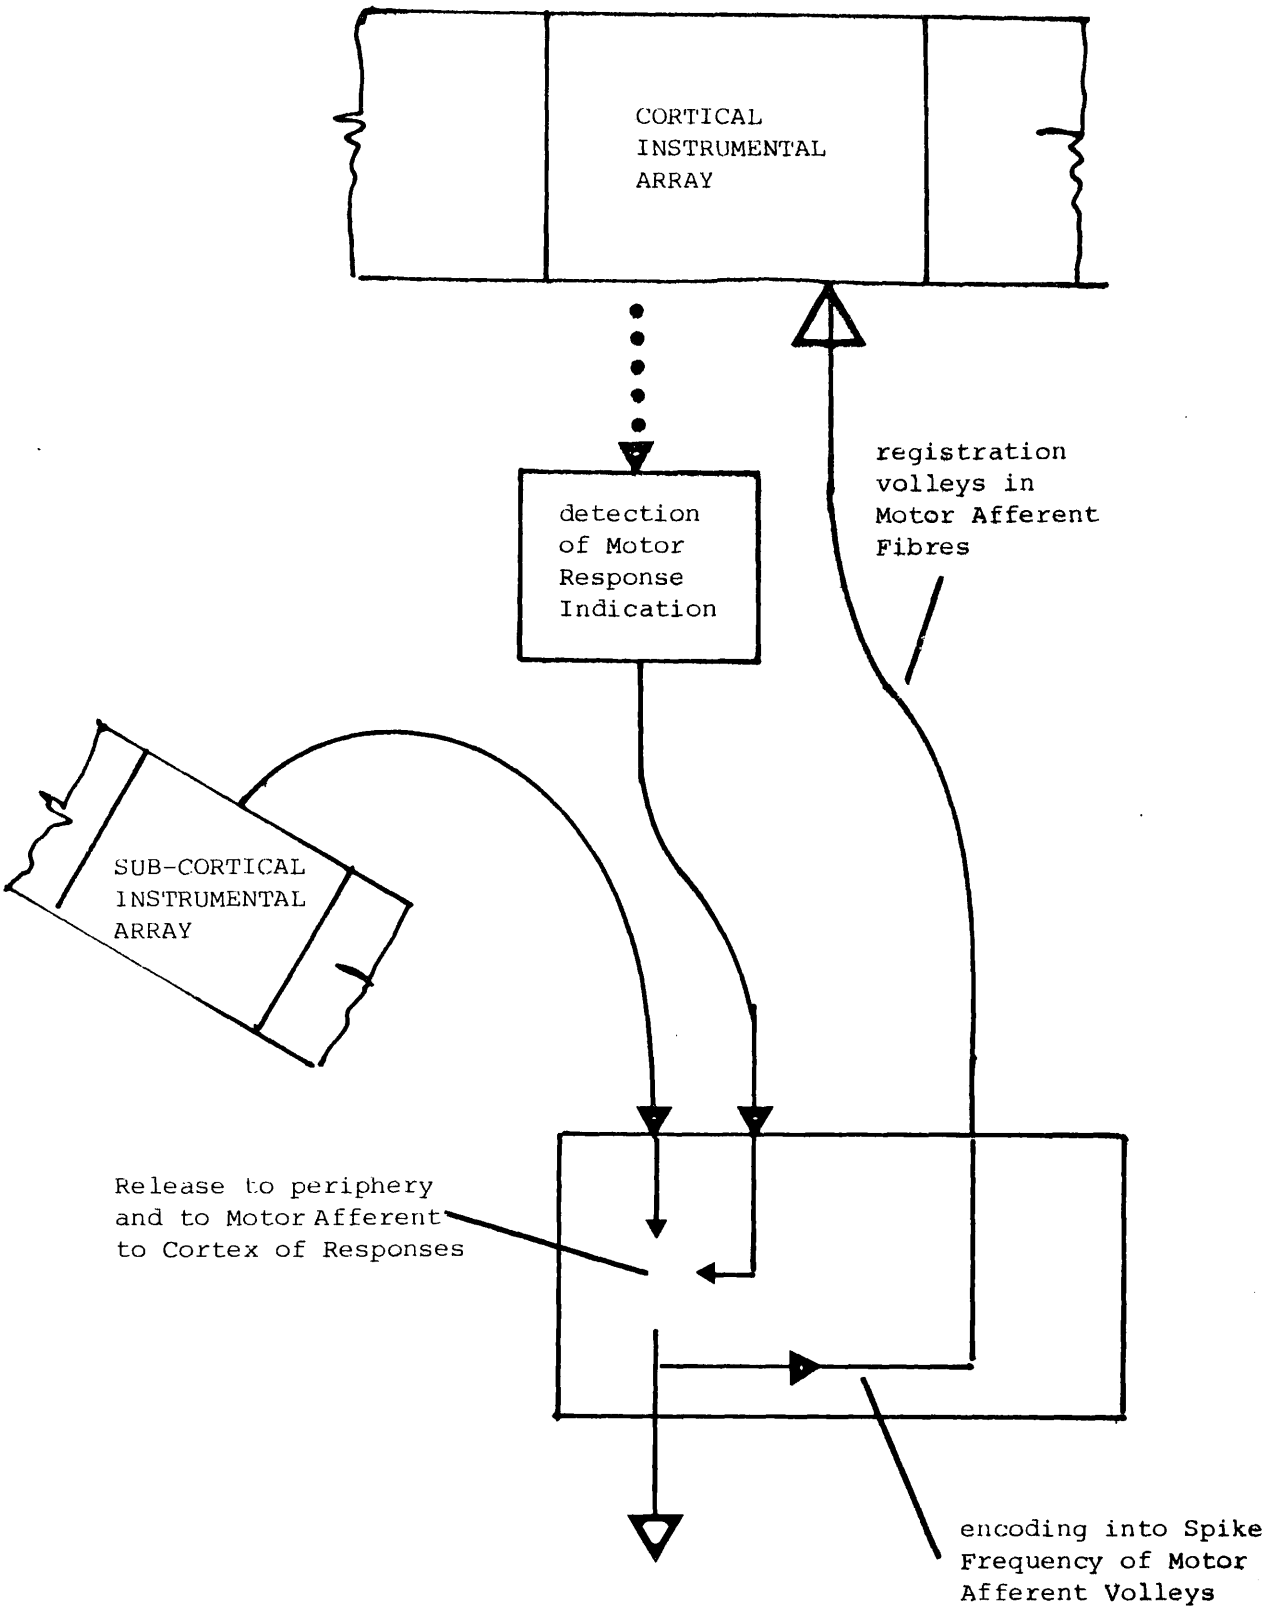

Figure 12.1 The Control Centre (Thalamus) receives the Instrumental Motor Response Indication and the Sub-cortical Response Indication, and must decide between them (see Figure 12.3). At each Instant it must encode the set of Motor Afferent Fibres to the Cortex into Active and Inactive Fibres ('0's and '1's) and calculate appropriate Spike Frequencies for them.

in one or other of the Efferent Fibres of that Array from time to time. These Sub-cortical Instrumental Responses pass to the periphery as notional Instrumental actions. The Efferent Discharges are also passed via the Control Array to the Cortical Array where they are registered with the ongoing Sensory input, so that a Sensori-motor Learned Linked Constellation is laid down. This provides for the registration of a sequence of sensory events integrated with a series of motor responses.

Upon repetition of the sensory event which forms the first step of a familiar sequence, the matching Linked Constellation is reactivated. If this Linked Constellation contains a representation of the motor action leading to the next sensory input in the sequence, then a Discharge will appear in the Efferent Fibres of the Cortical Array, indicating which action is next required. In principle the Discharge could actually constitute the initiation of the required Motor Response. For various reasons which will be seen in due course, it was preferable to have a detecting system aimed at finding in which Fibres the Motor Discharges appeared. These were then dealt with in the manner already described as typical of the action of the Control Centre.

The reader may wish to be reminded of what this action was. The Control Centre (i) determined which of the Efferent Fibres reaching it from some Array were to be classed as Active and which Inactive. It did this by adjusting a notional Inhibition across them, and (ii) it computed the Spike Frequency of the Fibres leaving it to become Afferent Fibres to some Array. It did this by applying the Diffuse Volley Value to all Fibres and the Discriminative Volley Value only to those Fibres which were Active (Figure 12.1).

In the case under consideration the Fibres which had to be coded as Active or Inactive were the Efferent Fibres of the Cortical Motor Array concerned. This determined which Fibre showed a sufficient Anticipatory Discharge to be justified in initiating a notional behavioural Response. Such an indication, based on the Anticipatory Discharges of a Cortical Array, was termed the Motor Indication. Accordingly the Motor Indication could be Consummatory, Preconsummatory, Protective or Instrumental, depending upon which Cortical Array it originated from. The processes by which the Motor Indication was determined was called the Motor Release Mechanism.

When a sufficient Motor Indication had been detected, an Efferent Volley was notionally sent to the periphery to effect the Motor Response of the model. In addition an Active volley was sent up the Afferent Fibre to the same Neuronal Pool of the Cortical Array from which the Motor Indication had originated. The purpose of sending back to the Cortical Array a controlled volley reproducing the Motor Indication was to ensure accurate and sufficient registration of the action released. It also enabled timing of the registration with particular sensory inputs accompanying the taking of the action.

The above arrangements implemented in the Instrumental Array provide for the mediation of Stimulus-Response chaining. When a familiar sequence of chained stimuli and responses has once been laid down, the recurrence of the first Input Pattern of the sequence leads to the emission of the first Instrumental Response. Because of the notional status of this Response as able to influence environmental events, its emission by the model justifies the environmental routines of the program in presenting the next Input Pattern of the sequence. This causes the emission of the next Instrumental Response, and so the Stimulus Response chain is reproduced.

It will have been noted that for this chaining to be effective, timing of events within the system is important. The Linked Constellation which registers the <u>current</u> Input Pattern must register the <u>next</u> Response. Actually, the model was more tolerant than this. Because of the temporal sequential effect, it was sufficient that the next Motor Response was represented in a Linked Constellation which had been formed after the Linked Constellation registering the Input Pattern currently in the Input, and that no Linked Constellation intervened which contained some other Motor Response. Nevertheless, the planning of Uttley-like delays in the pathways enabled the maximum effect to be obtained.

Stimulus-Response chaining of this kind was relatively easy to obtain in the model once a familiar sequence of Input Pattern and Instrumental Response had been established. The difficulty was to establish the sequence which led to the goal: in the case of the model, an IRP yielding a Preconsummatory or Consummatory Response.

Possible mechanisms which enable an animal to learn the required operants to reach a "reinforcement" raise all the controversies of learning and conditioning theory. It is not proposed to enter into those controversies in this thesis. The mechanisms which were incorporated were primarily intended to demonstrate the potential neural mediating power of the theory of neural integration. Secondarily they represented mechanisms which, in the opinion of the author, seem likely to be present in the brain: their essential relationship to the learning process is another matter which is not intended to be broached.

B. Contiguity theory: Frequency Principle

The Stimulus-Response chaining described above represents a Contiguity theory of learning (e.g. Guthrie, 1935; Estes, 1959).

Only the briefest mention of these theories will be made here. It is however necessary to report that before the introduction into the simulation of any further mechanisms of Instrumental action, an attempt was made to demonstrate the ability to reach a goal utilising the stimulus-response chaining principle only.

The general rationale of that attempt rested upon the following argument. Suppose that there are a number of Instrumental Responses which appear at random, and that one of these responses yields the goal. The attainment of the goal, according to the contiguity theorists, brings a sequence of behaviour to a natural close. In the simulation the trial was ended when the goal was attained. Now since the emission of the "correct" response is the event which signals the end of the trial, it must always appear no more than once in every trial. The other responses will only appear in those trials where their random selection causes them to be emitted before the correct response. Consequently the correct response should occur more frequently and establish itself by the single principle of frequency (i.e. without the introduction of a special reinforcement principle). This was one of the characteristic arguments put forward by contiguity theorists.

Considerable efforts were made to produce a convincing performance of the model utilising this principle. They were unsuccessful. It should however be emphasised that this result is not put forward to suggest that the argument is not valid, nor that contiguity theory is

not sound. Even if the test had been far more exhaustive, a computer model cannot in principle provide such a test. Moreover, for a number of reasons the attempt was not conducted in the best of conditions. The major problems with the model had not yet been solved. However, a brief test was made in the latest and most efficient version which demonstrated Instrumental Learning with a reinforcement mechanism. The test consisted of running the model without reinforcement mechanisms being active (i.e. suspending those mechanisms). No convincing learning resulted. The author is of the opinion that this was because the runs could not be long enough for the frequency principle to establish a superiority of the "correct" response. The size of the model restricted its learning capacity and hence the practical length of a run, so it was not possible to overcome this obstacle.

These attempts are reported for the sake of completeness. It is appreciated that many aspects of Contiguity theory were not satisfacorily represented, for example, Estes' sampling principle, and the importance of contextual association. In some respects the theory of neural integration upon which the present model is based is a contiguity theory. It suggests that, subject to attentional selection, the whole panoply of events, external and internal, can be registered. The author's private opinion is that the contiguity principle is the fundamental from which later evolutionary developments have arisen. That opinion is however an unsupported conjecture as far as the simulations are concerned.

C. A mechanism of reinforcement

The mechanism of reinforcement used in the simulations reported, is not put forward as the only or main mechanism by which instrumental learning is mediated in animals and Man. They may nevertheless point to principles which are utilised in neural systems.

The simplest form of reinforcement suggested by the model was to arrange for the registration of the Instrumental Response in the Cortical Array to take place only if the Response was successful. Success, in the appetitive-consummatory analogy of the model, is defined as the appearance in the Input of a Preconsummatory IRP. Registration means the incorporation in a Constellation of a representation of the emitted Response and could be effected or with-held by controlling the

Afferent Volleys to the Cortical Array. The model could not, of course, merely be told whether a Preconsummatory IRP had appeared in the Input. It would have to discover this by its own powers, and this it could only do by waiting before registration until the Input Patterns which followed either did or did not reactivate a Preset Constellation in the Sub-cortex.

The computer implementation of this scheme was effected by delaying the notional relay of the Motor Response Indication in its passage through the Control Centre for a number of Instants before passing the encoded Spike Frequency to the Motor Afferent Fibre (i.e. the Fibre from Control Centre back up to the Cortex). If during this delay a Preconsummatory Response was not forthcoming from the Sub-Cortical Array, the notional Volley in the Motor Afferent was blocked. On the other hand, if such a Response had been emitted, the volley in the Motor Afferent was sent to the Cortex where the Response was registered in the Instrumental Array (the same Neuronal Pool from which the Motor Indication had arisen).

The basic idea was later implemented without involving a protracted delay by incorporating a consolidation analogy. It will be specifically mentioned where employed.

In the initial simulations, in which all Uttley-type delays had been removed, the Sensori-Motor Constellation which produced the Anticipatory Response consisted of the Learned Record of the IRP on the Sensory side and the directly elicited Motor Response on the Motor side. The extended model necessitated the delay of the transmission of the Instrumental Response to the Cortex until a Preconsummatory Response (or absence of one) had been detected. If no further delays had been introduced the Instrumental Response would have been registered in the Linked Constellation which registered the occurrence of the IRP. Accordingly a notional delay was introduced into the Sensory Pathway to the Cortical Array. This approximately matched the delay of the registration of the Instrumental Response. Consequently, the signals registering the Instrumental Response arrived at the Cortical Array at the same time as those from the Input Pattern present just before the IRP.

It was mentioned that there were two notional Sensory projections to the Cortical Array, differing in notional time of arrival. Both of

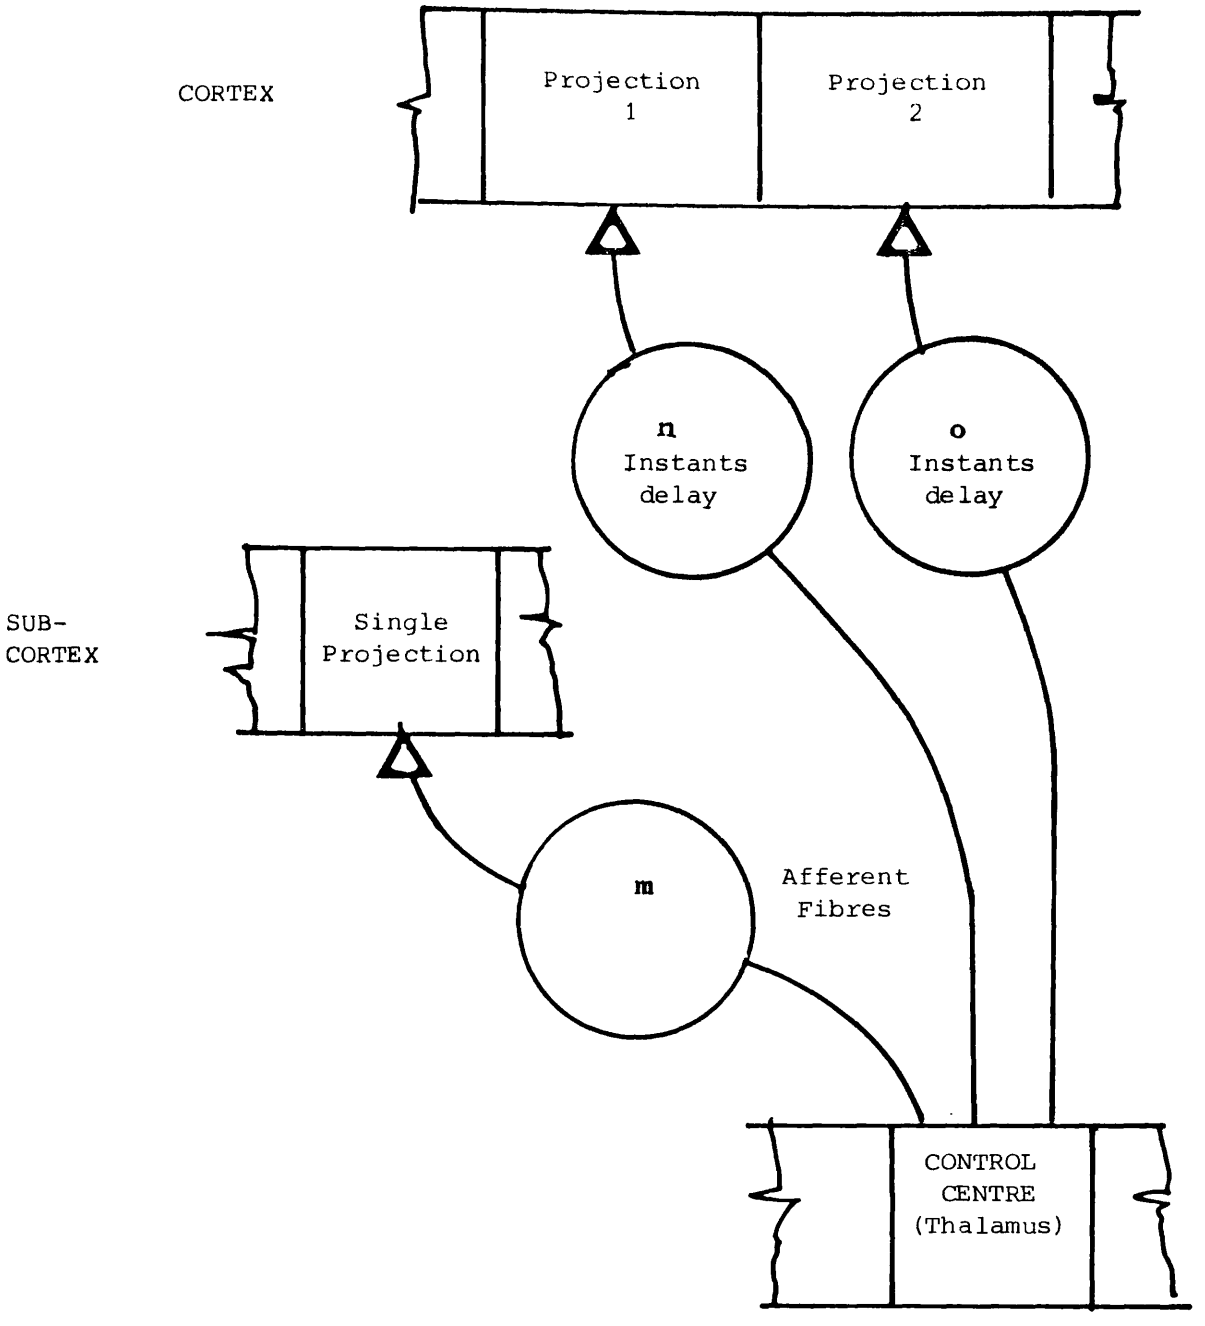

Figure 12.2 The Sensory Input Pattern was projected to three locations in the model; to the Subcortex and to two Neuronal Arrays in the Cortex. The pathway to projection 2 was delayed in relation to the pathway to projection 1. Accordingly two representations of Input Pattern were registered in the Linked Constellation formed at any one Instant. That in projection 1 was the Input Pattern currently in the input, whilst that in projections 2 was the Input Pattern which had been in the input some 13 Instants earlier.

these projections were now further delayed in relation to the notional time of arrival of the signals at the Sub-cortical Array. In effect, a pattern presented to the model at time T reached the Sub-cortical Array at time T + m, the first projection of the Cortical Array at T + n, the second projection at T + o, where o > n > m. The Instrumental Response reaching the periphery at T was registered at T + n. (Figure 12.2).

Aside from the tie-up between Innately Recognised stimulation and reinforcement implied in this scheme, the temporal implications appear somewhat implausible. In Chapter 3, it was pointed out that the dependence of reinforcement upon innate propensities was considered as only one of the methods of assessment of results available to an animal. Nevertheless, the close relationship between Need satisfaction and Instrumental performance must make such a tie-up plausible as one of the mechanisms available. The temporal implications are harder to deal with and will be left until the problem can be seen in relation to other possible mechanisms and the modification of the one just described. However, many instrumental actions of simple animals are of a continuous rather than a single event kind. Moreover the sensory imput of the animal, unlike that of the simulation, is only partly composed of configurations which change. Accordingly the Instrumental Response might continue right up to the time that the IRP appears, and even beyond its appearance, and so make the temporal delay required minimal or even unnecessary. Also, in the advanced nervous system there would presumably be neural mechanisms available to the animal which enable sequences of patterns to be viewed as units. These might be "labelled" by the Instrumental Response in a manner similar to that in the simulation. Finally, as mentioned above, the system could be supplemented with a consolidation hypothesis. A provisional registration could be laid down without utilising delay, and this could either be confirmed or not according to subsequent events.

D. Negative reinforcement by suppression

Another mechanism of reinforcement utilised the representation of suppression of effector action (i.e. Not-Do Neuronal Pools). The mechanism described in C above could be modified so that when a Preconsummatory Response followed an Instrumental Response (behavioural success), the Instrumental Response was registered in the Activation Neuronal Pool of the Cortical Array (Do Neuronal Pool). When a Preconsummatory

IRP failed to appear, the Instrumental Response was registered in the corresponding Suppression Pool of the Array. When on a subsequent occasion the Constellation was reactivated by the Sensory Input, then in the case of the Instrumental Response having been successful, on the former occasion, an Anticipatory Discharge was obtainable from the Activation Pool: in the case of its having been unsuccessful, the Anticipatory Discharge was obtained from the Suppression Pool. These could be weighed against each other to determine whether or not the Response should be emitted, i.e. to determine the Instrumental Motor Indication.

E. Summary of Learned Instrumental Response Arrangements

The output of the Cortical Instrumental Array was examined for the Anticipatory Response. The appearance of such an Indicated Motor Response led to the immediate dispatch of an Efferent volley to the notional periphery (Motor Output of the model). A Spike Frequency was prepared but not transmitted to the Cortical Array until an appropriate Response had been emitted by the sub-cortical Array. If no such Response were received within a given time (13 Instants), the equivalent volley was sent over the corresponding Afferent Motor Fibre to the suppression Pool.

It will be appreciated that this scheme avoided one of the difficulties encountered in the initial simulations in relation to the Suppression Pools. In the initial simulations the Fibres to the Suppresssion Pools were active whenever the model was notionally doing nothing, i.e. most of the time. The new arrangements involved activity in the Motor Pools, effector or suppressor, only when it was necessary to register the effects of an instrumental action. Neither solution seems entirely satisfactory. Konorski's (1967) suggestion that the state of arousal governs this activity seems entirely plausible. Such an assumption and its incorporation into the general ideas of the model will be discussed in a later chapter.

F. Combined Do and Not-do Registration

One scheme of furthering Instrumental Learning must be mentioned, as neurologically it seemed to the author to be the most plausible. Unfortunately, it was not possible to obtain a convincing performance utilising it. This may however have been due to the stage of development of the model in which it was investigated. It would be useful

to be able to try it out again in a more advanced version having both effective Motivational Respresentation and an adequate Excitation and Inhibition level control system.

This scheme was neurologically very similar to those which have just been described. The only difference was the (neural) set of rules permitting registration of the Effective Instrumental Response and the Suppressive Instrumental Response (i.e. Do and Not-Do). It will be recalled that in the schemes described above, either the Do was registered only when successful; or in another scheme, the Do was registered when successful and the Not-do when unsuccessful. In the main scheme intended for the model and now being described, both the Do and Not-do Pools were always activated after an Instrumental Response, regardless of success or failure. However, the registration of the Not-do came somwhat later than the registration of the Do.

The rationale of this mechanism was quite close to that of Contiguity theory. In a problem situation the animal emits several responses, which we categorise as wrong in the sense that they do not succeed in changing the situation, that is, in eliciting from the environment an input radically different from that which pertained before the emission of the Response. These wrong Responses are registered in one Linked Constellation as the activity of a Do Pool and in a slightly later Linked Constellation as a Not-do Pool. These balance each other. When the animal produces a "right" Response, the input is assumed to change radically. The registration of the Do is in a Linked Constellation the sensory part of which matches the sensory pattern appertaining to the situation before the emission of the Response. The delay in registration of the Not-do, leads to this being registered in the Linked Constellation containing the Sensory pattern appertaining to the situation after the emission of the Instrumental Response. Consequently on a subsequent occasion when the animal is faced with the sensory input preceding its emissions, i.e. the problem situation, all its previously emitted responses except one will have matched Learned Records in which Do and Not-do are balanced. The exception will be the Instrumental Response which successfully changed the input pattern. This will have only the Linked Constellation reactivated which contained the Do. The Not-do will be reactivated only after the situation has been changed by the emission of the Response. This scheme has the advantage that it also allows for an

Instrumental Response to be suppressed after it has acheived its purpose. It may also be noted that there is some evidence that, at least in the periphery, both effector and inhibitory excitation of motor fibres occurs, the balance being modified (See e.g. Groves and Thompson, 1973).

III The Sub-cortical and Cortical Control of Motor Actions

The introduction of Instrumental Responses into the model raised a problem which appears to have received little attention in other neurological models. This may be because the problem is specific to the present model. It is the problem of the balance of control between the emission of responses as the result of an innate propensity to do so, and the emission of responses under the influence of experience. To avoid a lengthy discussion at this point, the theoretical aspects will not be considered here, and only the arrangements necessary to deal with the problem will be described.

The difficulty may be seen in relation to the release of responses discussed in the previous section. The model shows how a learned indication of the required motor action can be obtained from the Cortical Array when the current Sensory input matches that in a Learned Constellation laid down on a previous occasion, provided that Constellation has a motor representation, i.e. it is a Sensori-motor Constellation. This means that the learned indication of what to do next comes from the cortical Array. Yet the Cortical Array cannot produce the Motor Indication unless there has at some time previously been an emission from the Sub-cortical Array. This is the only way in the model that a Sensori-motor Linked Constellation can arise in the Cortical Array. The Sub-cortical emmission is in this sense primary.

In the present context it does not matter whether the emission from the Sub-cortical Array is totally dependent, partially dependent, or totally independent of external stimulation for its primary emission. The point is that on its first appearance, whatever the mechanism, the emission of a response must be independent of experience. With subsequent experience the emission of the response comes under the control of that experience.

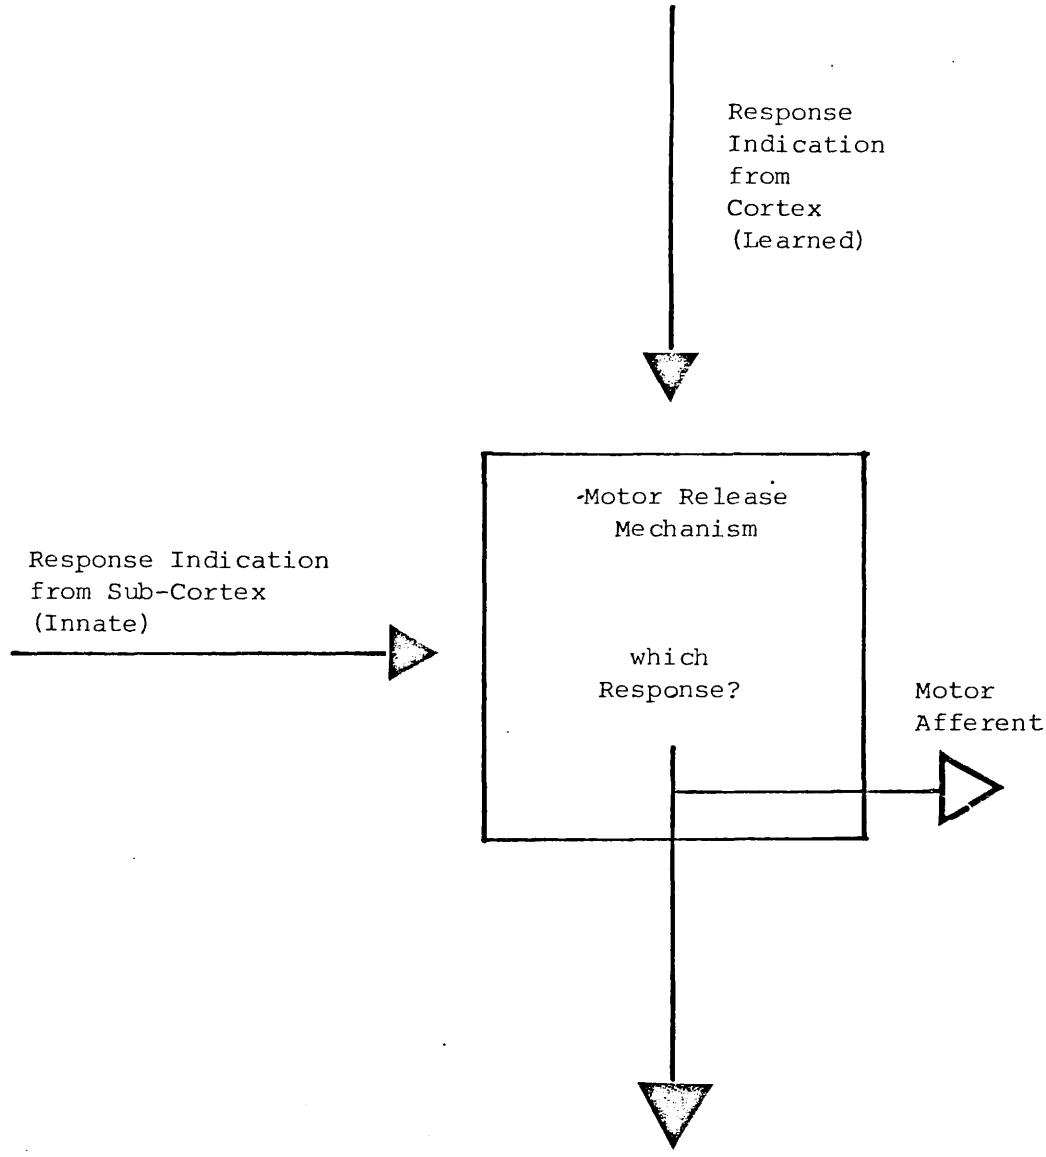

Figure 12.3 Illustrating the function of the Control Centre (notional Thalamus). The Motor Release Mechanism has to choose between a Response indicated by the Sub-cortical Instrumental Array or the Cortical Instrumental Array. At each Instant the Control Centre must encode the set of Motor Afferent Fibres as Active or Inactive and then calculate their Spike Frequencies accordingly.

In the model this control must derive from the Learned Records in the Cortical Array. As pointed out above, the first emission of the Response, at least, must come from the Sub-cortex. A possible solution which was considered was to have the first emission only coming from the Sub-cortical Array, and subsequent emissions by reference to the Cortical Array. However, not only would this hardly do from the point of view of plausibility, but in fact it did not work in the simulation. The organism will have to emit Responses in a variety of situations in its later life: unless the subsequent situation resembles the former situation sufficiently for the response to be derived from an indication in the Cortical Array, then it will be unable to emit the Response in that situation. Its power in novel situations would be seriously diminished. Additionally, if the situation were not novel, but was associated with some alternative Response, it could only make the latter. Moreover, it may have completely forgotten (lost its learning) relating to its earlier use of a Response before it needed to emit the Response again. In such a system, an animal could 'forget' to be able to emit its most characteristic instrumental actions. In the model, a particular difficulty occurred. If by any chance it learned a Response as "correct" which afterwards became "incorrect" due to a change in contingency, it could never remedy the position. As long as it "remembered" the occasion of the first emission of the Response, it would repeat it in the "wrong" situation, and it could never emit the Response in any other situation because there would be no Learned Indication from the Cortical Array. After it had forgotten the occasion of the first emission, it could never emit the Response at all.

An alternative is to have the mechanism which controls the emission of Instrumental Responses sensitive to both the Sub-cortical Array and the Cortical Array, and to initiate that Response which is represented by the most powerful Efferent volley. This was tried and produced a disruption of learning which prevented a convincing performance by the model. (Figure 12.3).

The reason for this disruption may be seen by reference to the reasonable suppositions controlling the emission of Responses from the Sub-cortical Array, that is, from factors other than those due to learning. Although opinions may differ as to what these factors may be, it will probably be agreed that they would include: the time interval since the last

emission of that response; motivational factors especially related to the response in question; the tendency of the animal to emit certain responses in the presence of certain kinds of stimulation, e.g. in darkness. These factors, being by definition those which do not represent learning, would operate in a probability sense to produce the response more frequently in some situations than others. Unless such responses can be suppressed, they are as likely as not to appear in the middle of a learned performance and to disrupt it. This, in fact, is what happened in the model. Just as the model was about to make the "right" Response it would be interrupted by a Sub-cortical emission.

The next method to be tried out imposed an alternating gate successively upon the Cortical and the Sub-cortical emission. This seemed plausible having regard to the general ubiquity of rhythmic inhibitory effects in the system. It worked considerably better than the method in which both sources of Response were consulted simultaneously, but the performance was still subject to disruption.

Finally the scheme was introduced which seemed both most reasonable and which provided a satisfactory performance. In this version, the mechanism which controlled the initiation of Instrumental Responses looked first to the Cortical Array for the Learned Motor Indication. Any such Indication suppressed the emission of a Response from the Sub-cortical Array. Consequently it was not possible for a Sub-cortical emission to disturb an emission indicated by the Cortical Array.

In the early versions of the extended model, this method still gave trouble. In these versions the active Suppression had not yet been introduced as a negative reinforcement to eradicate "wrong" Responses. Although the system should only have registered "correct" Responses (see previous Section) it often made mistakes. Once it had registered an "incorrect" Response, it was unable to remedy its error, because it went on indefinitely emitting this Response under Cortical influence. Theoretically it should have ceased to do so by decay of memory. Unfortunately the length of runs could not be made long enough to test this. In some respects the shortness of the runs magnified theoretical difficulties. A primitive learning system would probably only learn with great slowness: consequently, one or two errors would not prevent learning in this way. To produce a convincing Response in the short runs

available to the simulation, learning had to be made very strong on a single Instant.

To combat this difficulty, before Suppression was introduced, the rule that Response Indications from the Cortical Array had preference was qualified. If the model had not achieved a Consummatory Response within a certain number of Instants, then the Sub-cortical Array had preference. This effectively cured the model of its perseverance of an error. Actually although replaced by later alternatives, the method just mentioned does not seem implausible to the present author. The satisfaction of certain needs of the animal are so imperative that it seems feasible that if learned behaviour fails, the animal should revert to behaviour of a kind from which it gained its original learned skills: in the model, to actions emitted under the control of the Sub-cortical Arrays.

However, in the simulations the method just described became unimportant, because of the introduction of the principle of Suppression as an aspect of Reinforcement. A "wrong" Response, even if rewarded by an IRP, would on its subsequent "unreinforced" emissions be registered in a Linked Constellation which included representation in a Not-do Neuronal Pool. The Motor Indication from the Cortical Array after a number of ineffectual emissions was suppressed by the greater Efferent Discharge from the Not-do than from the Do Pool. Consequently, a Learned Indication disappeared. Either it gave way to another Cortical Motor Indication, or if no such Indication was available, to a Sub-cortical emission. In either case the model was released from the trap of its own error.

In effect, this behaviour in the model demonstrated the importance of extinction in the motor learning process: effectively the negative reinforcement of ineffectual Responses mediated a form of extinction.

This aspect of the model has been discussed at some length, because it absorbed much effort and time in the development of the model. The background of the investigation was the variable and uncertain performance of the model due to lack of control of Excitation and Inhibition. In this sense the latter difficulties may have contributed towards the realisation of the problem of competition between innate and learned sources of Instrumental Actions. A model which had been more reliable may have produced a learned Response more easily, and tolerated some

logical indefiniteness in this problem.

Although theoretical implications were to be avoided in this chapter, it may perhaps be suggested that the problem discussed above is not specific to the present model. At first sight it may appear that the problem arises from the individual characteristics of the model: the allocation of Learning to the Cortical Array and other aspects of Response emission to the Sub-cortical Array; or to the preservation of the key-trigger mechanism entirely separate from the Learned aggregates of experience (discussed in Chapters 4 and 5 in connection with Pringle's 1951 model). This does not seem to be so upon fuller consideration. The competition between the emission of responses under the pressure of factors independent of learning, and those urged by learning, would appear to arise regardless of the particular arrangements by which these factors are implemented. Provided one accepts that the animal would emit responses independent of its having learned the most favourable occasions upon which to do so, then it would appear that a neurological question arises as to the neural basis of the integration of the two influences.

IV Mechanisms of Motivational Arrays

It remains to explain the principles by which the Motivational Neuronal Pools were operated. As the principles of action of the Cortical and Sub-cortical Motivational Arrays differed, they will be dealt with separately.

A. Cortical Array

The Afferent Fibres of the Cortical Motivational Array were conceived as signalling metabolic and hormonal events, such as hunger, thirst, sexual urge and so forth. It was intended that each Pool signalling a certain condition would be balanced by a reciprocal Pool signalling the opposite condition, e.g. hunger and satiation (see e.g. Deutsch and Koopmans, 1976). Unfortunately there was no room in the structure of the simulation to fit these in. They were planned for a larger simulation which was programmed but not run.

When any Afferent Fibre of the model, e.g. a Sensory Fibre, is active, then in accordance with the general principles of the theory, Linked

Neurons within the Neuronal Pool served by that Fibre receive Excitation. This Excitation is then distributed across the Linked Constellations to which those Neurons belong. In this sense any active Afferent Fibre may be seen to drive or bias the system. Accordingly if Neurons within the Motivational Pools were to become members of Linked Constellations, the activity of the Motivational Afferents would bias those Linked Constellations. Had Neurons been linked, in the manner characteristic of other kinds of Neuronal Pools, whenever the Afferent Fibre was active, then when the Fibre representing, say, hunger, was active, Neurons in the Pools would be linked into the Linked Constellation forming at that time. The Learned Record of an event occurring when the animal was hungry would include such Linked Neurons. When the animal was hungry on a subsequent occasion the Learned Record of the earlier event would be primed by the current activity in the Motivational Afferent signalling Hunger. Consequently, the threshold of recognition of the stimuli present on the earlier occasion would be lowered. Such a lowering of threshold would represent a motivational attentional mechanism, by which the animal would be disposed to respond more readily to stimuli present upon earlier occasion of hunger.

Although such a mechanism appears plausible as a neural attribute of nervous systems, it was not considered to attack the root of motivational function. If the Cortical Motivational Pools operated on the same principles as other Pools in the system, they might be considered as merely an addition to the Sensory Arrays, the only difference being that their Afferent Fibres signalled internal metabolic events rather than external or proprio-ceptive events. It is true that such events would act as selective cues, biassing in the system in the same way as sensory stimuli would bias the system. However, in the context of Instrumental behaviour, they would act equally upon "unsuccessful" as upon "successful" actions.

So far described the metabolic or hormonal Afferent Fibres are not a mechanism of reinforcement. They do not assist the animal in discovering which Instrumental Response led to success in which particular environmental situation. For this purpose it is necessary for the Excitation from metabolic afferents to feed selectively only to the Learned Records of successful Instrumental Actions and not to unsuccessful ones. In terms of my model, the inputs from Motivational Afferents

must feed only to Linked Constellations containing Linked Neurons in the "right" Response Pools and not the "wrong" ones.

Ideally the way to accomplish this would be to have a double representation of each need Afferent, i.e. two Fibres for the particular condition of, say, Hunger, and two others for the condition of Satiation. One of these would act as the Contextual Need Input. Its Neuronal Pool operates on the same principles as other Pools in the system. The activity of the Pool acts merely to bias the system as described above. The other Fibre may be considered as a reinforcement of "successful results" Fibre. In fact these two Fibres were amalgamated into a kind of compromise in the simulation, since there was inadequate space in the structure to accommodate both.

In the simulation, the Afferent Motivational Fibres operated as follows. When a Motivational Fibre is active, the Neurons within its Pool are Excited. They are not, however, linked. For linking to occur in these pools, there is a special switch which had to be turned on only when success appears. In the implementation in the simulation, this mechanism was devised as the simplest which would be neurologially plausible and concordant with the principles of the model. It was not however considered to represent the actual mechanism which might perform the same function in the brain: in the brain it is suggested that the switch would be neuro-chemical in nature(see below). In the simulation the switch was modelled as follows. The Integrating Neurons of the Motivational Cortical Pools were given a second notional threshold, an Integrative Connection Threshold. The distribution of these thresholds had a mean which was approximately double the mean of the intrinsic Thresholds of the Neurons of non-Motivational Pools. These higher Thresholds did not affect Discharge. Their range was set such that the Excitation arriving from an Afferent Fibre carrying the Spike Frequency customary for an ordinary Active Fibre did not succeed in linking any Neurons. When an event signifying "success" occurred, the Spike Frequency of the Afferent Fibre was increased to double the Active level, thus causing Neurons in the Motivational Pool to be Linked.

The occasion for the switch to be on, that is the assessment of "success" was the occurrence of a Preconsummatory IRP in the input. The particular Motivational Pool in which the switch was turned on was that Pool which was (i) already active, e.g. the animal was notionally hungry when it

received the sensory signals of food in its mouth, and (ii) the Pre-consummatory Response was one which was neurally related to the Motivational Pool, i.e. the food-in-the-mouth IRP would be related to the Pool signalling Hunger.

This scheme is naturally intended to be interpreted fairly liberally in relating it to the animal nervous system. There is evidence that there are neural receptors which give rise to activity related to metabolic states, e.g. sugar and salt content of the blood: these would give rise to the normal level activity of the Motivational Fibres, exciting, say, the Hunger Pool. This is far too crude a description but expresses the general idea. However, there is no reason why an empty stomach should not activate an "empty stomach" Pool. Similarly, hormonal influences may act as long term clocks ensuring that motor actions are released with reasonable frequency: this, however, related more closely to the interpretation of the Sub-cortical Motivational Pools, dealts with below.

The reader is asked not to construe this theoretical structure as a narrow "needs" hypothesis of behaviour. As will be discussed in a later chapter, the approach is that all instrumental and consummatory responses are equipped with long-term clocks of similar character. If a rat does not run for a long time, then the 'run' Motivational Pool Afferent will be active. One must assume that the motoviational representations in the brain are not a few primary Need Pools, but are a vast system of monitoring the behaviour of the animal, with every important effector need system represented.

The Motivational Pools in the Cortical Array therefore worked somewhat as follows.

The model was assumed to be motivated (say, Hungry) and in a situation where the taking of a certain action would bring stimulation appropriate to that condition, say, food-in-the-moutn. The Input Pattern represented the sensory and proprio-ceptive aspects of the situation. The state of hunger (or perhaps one should say, food-deprivation) was represented by having the Afferent Fibre to one of the Cortical Motivational Pools active. It will be recalled that in the particular implementation just described, the Spike Frequency of an Active Afferent Fibre to a Motivational Pool could take one of two values. At one value the Fibre was only capable of biassing and at the other it linked Neurons. For

Patterns registered in

LINKED CONSTELLATION FORMED AT CRITICAL EVENTS

(ten to twelve Instants after Instrumental Response emissions)

Location in Cortex					
Sensory		Motivational	Motor		
Projection 1	Projection 2		Instrumental	Preconsummatory	Consummatory
Current Input Pattern i.e. Input Pattern which appeared as a result of emitted Instrumental Response	Delayed Input Pattern i.e. Input Pattern which preceded emission of Instrumental Response	Motivational Pool relevant to elicited Preconsummatory Response	Neuronal Pool relevant to emitted Instrumental Response	Neuronal Pool relevant to elicited Preconsummatory Response if one occurred in past 12 Instants	In some versions only, Neuronal Pool relevant to elicited Consummatory Response registered at this time. In other versions, registration delayed

Figure 12.4

convenience these may be referred to as bias-activity and linking-activity. In the starting situation the Motivational Afferent was only bias-active.

On the first Trial, the activity of the Motivational Fibre had little effect on the system, because there were no Linked Neurons in the Motivational Pool served by the bias-active Afferent. When the model emitted a right Instrumental Response, the Input Pattern was changed by the notional environmental part of the program, to the Preconsummatory IRP (notionally, say, food-in-the-mouth). The IRP elicited the Preconsummatory Response from the Sub-cortical Array. The Preconsummatory Response, in addition to releasing notional Preconsummatory behaviour, e.g. chewing and masticating, initiated an increase in the Spike Frequency of the Afferent Fibre to the relevant Motivational Pool bringing it to Linking-activity. Consequently the Constellation being formed at that time would contain a representation of the Motivational Pool concerned.

The timing of the change from bias to Linking in the Motivational Pool was arranged so that it coincided with the registration of the Instrumental Response (Section II,C). It will be recalled that the registration of <u>both</u> the Instrumental Response and the Sensory Pattern in the Input at the time the Instrumental Resposne was initiated, was deferred until a Preconsummatory IRP had had time to appear. If it did appear then the IRP, the Instrumental Response, the <u>pre-response</u> Input Pattern were all registered in one Constellation (or a series of temporally adjacent Constellations). Now, with the Introduction of Motivational representation, in the case of a successful Response, the Motivational Pool was also registered in the Constellation (Figure 12.4).

On a subsequent test trial, when the same Input Pattern was presented, the Sensory input reactivated all the set of Linked Constellations matching it. The Motivational Fibre relating to Hunger was active at the lower level. On this occasion the Afferent Fibre found Neurons in its Pool which were Linked. Consequently Excitation was passed by these Neurons selectively to the Linked Constellation which contained also a representation of the current input pattern and the Instrumental Response. The Input Pattern therefore merely acted as a releasing trigger bringing the Instrumental Response above threshold. In the absence of a Motivational input, the Instrumental Response would have had to be powered entirely by the Sensory input.

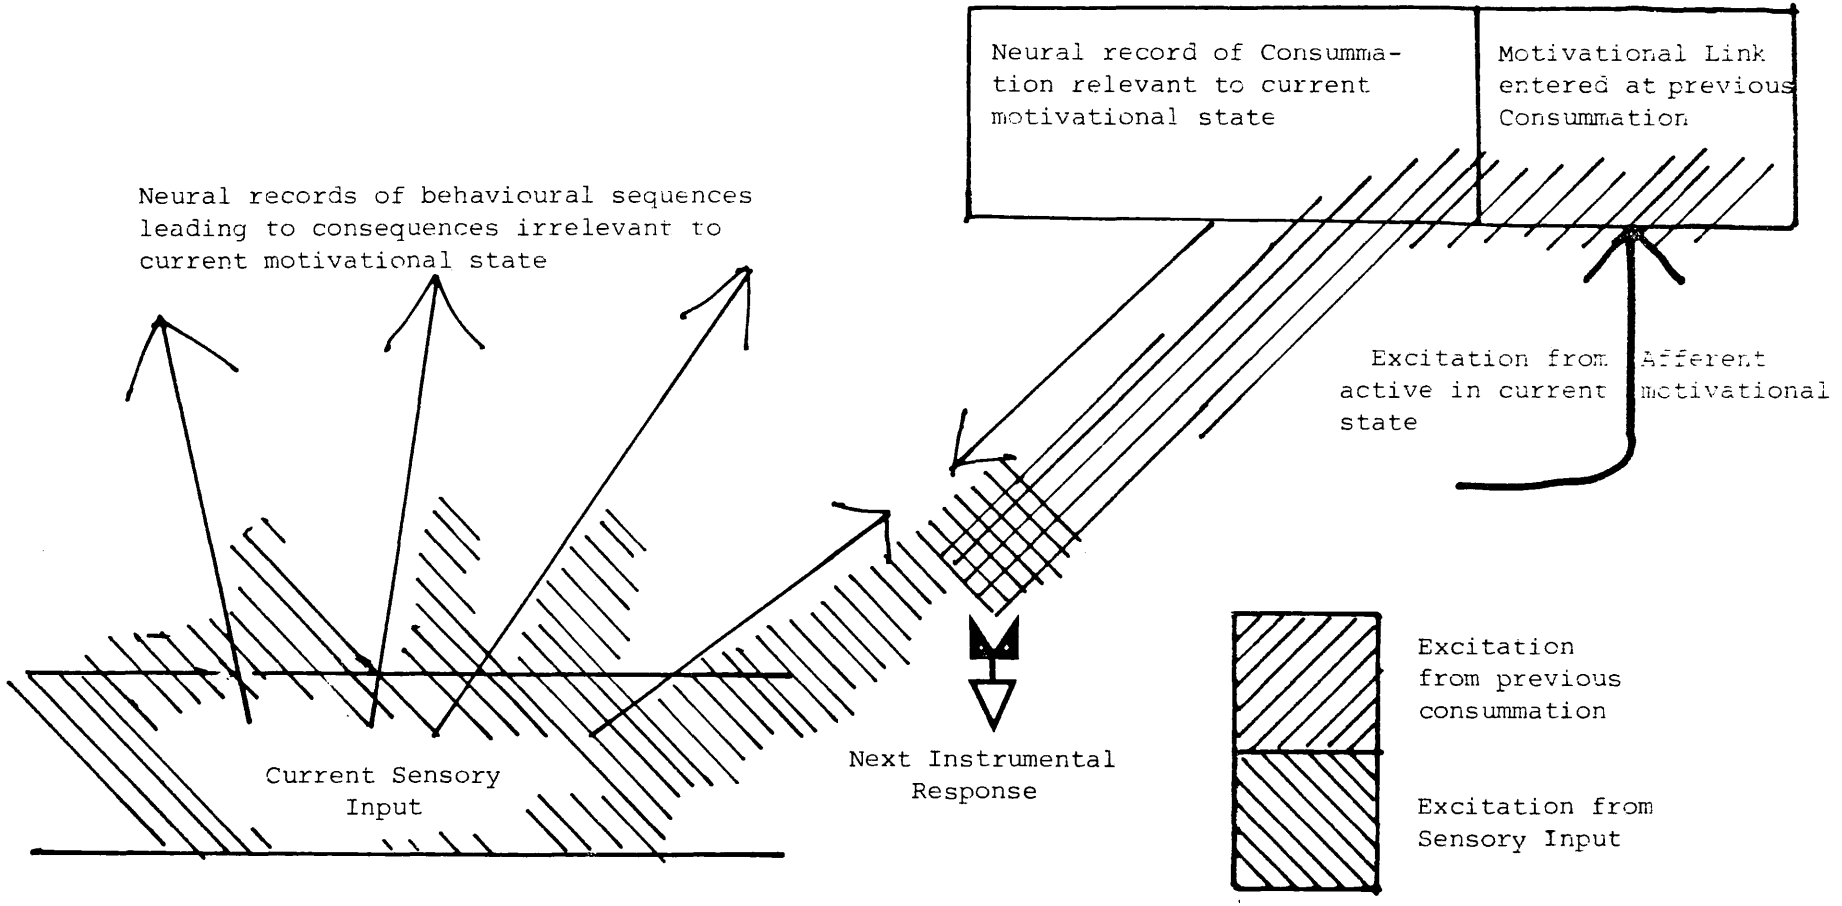

Figure 12.5 Need satisfaction learning (Freud's Primary and Secondary process). Excitation entering through motivational link in Constellation registering a previous consummation spreads through system and biasses sensori-motor pathway leading to consummatory situation (Primary process). Excitation entering through current sensory Input Pattern spreads through several pathways. In the sensori-motor pathway leading to consummation, this excitation is reinforced by biassing excitation and response results (Secondary process).

It may have been noted that the mechanism which has just been described is in itself a selective mechanism which makes the presence of the reinforcement mechanisms described earlier in this chapter somewhat redundant. For example, suppose that all Responses, right or wrong, were registered as was assumed by contiguity theory. In this case, an Input Pattern would reactivate all the Linked Constellations matching it, and all the Responses, right or wrong, included in the various Sensori-motor Constellations laid down, would be equally reactivated. However, only one of them would have the Motivational Input. It would be going beyond the terms of reference of this thesis to try to decide which mechanism finds a place in actual neural mechanisms. What seems probable is that in a sophisticated animal various mechanisms of selection of motor action are available. They represent modes in which the system can operate to its best advantage. For example, there will be occasions when the motivational guide to activity gives wrong answers. It would then be useful to be able to suppress the motivational mechanism and revert to mechanisms relying only upon the conjunction of sensory patterns and results: and pari passu with other mechanisms.

It may have been noted that the Motivational mechanism just described is not unlike Freud's theory of Primary and Secondary Processes. The Motivational Pool energises the "goal image" (Primary Process), from which Excitation flows into the neural apparatus for achieving that goal, the Instrumental Response (Secondary Process). This relationship is illustrated in Figure 12.5.

It may be mentioned at this point that considerable difficulty was experienced in practice with the effects of Motivational Representation, and consequently with the implementation of a concept of motivational bias proceeding from "goal" to "current action". At a late stage some possible solutions were programmed, although not actually run sufficiently to assess the validity as a possible mechanism of behaviour.

B. Response initiation biassing: Sub-cortical Motivation Array

Unlike the other Cortical and Sub-cortical parts of the model, the Motivation Arrays do not bear the customary structural and functional relationship to each other. In fact the term Sub-cortical Motivation Array is probably misleading, and should be perhaps be something more like the Unlearned Response Initiation Biassing Array, or some such. The author must apologise for introducing inappropriate terms into the

model. The aim is simply to be able to identify shortly which part of the model is being referred to. The difficulty is that parts of the model do not necessarily correspond to aspects of psychological function: according to the model psychological function is the product of the sum of the parts. On the other hand it would be dangerous, not to say presumptuous, to label them with anatomical regions of the brain.

The term Sub-cortical Motivation Array was chosen because the function of the Array resembles that of the Cortical Motivation Array in that it biasses Linked Constellations. It was termed Sub-cortical not only because its sites would probably be sub-cortical, but because it is concerned primarily with innate rather than learned proclivities of the organism: it biasses Preset Constellations.

The function of the Cortical Motivation Array was to bias the model towards the production of a particular "learned" Instrumental Response: that is a Response which may have been emitted originally entirely without reference to experience, but would be emitted again from the Cortical Array in circumstances where on former occasions its emission had been successful. The state of one or other of the Cortical Motivational Fibres defined the kind of environmental result which the system was aiming at. The activity of these Fibres was intended to represent hormonal and metabolic states such as Hunger, Thirst and Sex.

By contrast, the action of the Sub-cortical Motivation Array related to the emission of Instrumental Responses at the Sub-cortical Level, i.e. the initial emissions which are not the result of individual experience. Since the general mechanisms of Sub-cortical emissions of Instrumental Responses have not yet been described, their biassing by the Sub-cortical Motivation Array will be left until the next Section, where it will be included in the general account.

The general effect was to bias the Sub-cortical emission of Instrumental Actions according to the notional passage of time since the last emission of a particular response. This gave effect to the view that there is a disposition to utilise the motor effector system regardless of the consequences of doing so, i.e. an "exercise" urge. The manner in which this ties up with the general Instrumental strategy of the model will appear in due course.

C. Preconsummatory, Consummatory and Protective Responses: motivation

No mention has been made of motivational influences on types of response other than Instrumental. This is not because such influences were considered to be theoretically unimportant, but because at this stage of development of the model their representation was not immediately relevant to the purposes in hand. Theoretically it would appear that the mechanisms applied to the Instrumental Cortical Array could be applied to the other Arrays. In effect, the threshold of elicitation would be lowered by Neurons in Motivational Pools linked into Constellations which recorded Consummatory, Preconsummatory or Protective IRPs. It seems probable that the system might be developed. The appearance of a Consummatory IRP might have the effect of linking sources of Motivational Excitation into the Constellations recording the Preconsummatory event, rather than the Consummatory event itself.

As these possibilities were not investigated in the simulation, the matter will not be futher discussed.

V Conclusion

In the previous chapter the main notional anatomical parts of the extended model were described. In this chapter the main principles of operation were described insome detail. The manner in which a Response could be initiated by a Cortical Motor Indication was explained, and other details of Response mechanism given. The Cortical Motivation Array and its method of operation in two modes, a bias-only mode and a linking mode, enable a source of Motivational Excitation to be connected to the Linked Constellations representing particular events, such as the occurrence of IRPs. The question of selection of response emission from cortical or sub-cortical centres was discussed.

It remains to describe the principles of organisation of the sub-cortical mechanisms which determined which particular Response would be emitted when the Motor Release Mechanism referred to the Sub-cortical Array. This will be dealt with in the next chapter.

CHAPTER 13

RESPONSE EMISSION: SOME SUB-CORTICAL MECHANISMS

I Organisation of Sub-cortical Instrumental Responses

In Section III the question of the emission of Instrumental Responses was discussed in relation to the decision which the Motor Release Mechanism must make as to whether to emit the Response indicated by the Cortical ("learned") or Sub-cortical ("innate") apparatus. In the version most frequently adopted the Motor Release Mechanism looked first to the Cortical Array for a Motor Response Indication. Failing such an Indication, it permitted the release of the Sub-cortical Response. (See Figures 12.1 and 12.3).

It remains then to consider the arrangements whereby one Response rather than another would be urged by the Sub-cortical Array, that is, which Efferent Fibre of the Sub-cortical Instrumental Array would be dominant when the Motor Release Mechanism had occasion to refer to it.

A. Some comments on the status of differing versions of the model

As in other aspects of the model, the arrangements for the activation of Sub-cortical Instrumental Responses differed from version to version of the simulation. Since this variation of simulated mechanism applied to a number of aspects of the model, it is perhaps necessary to clarify the strategy adopted.

It is not the case that the most recent version of a mechanism necessarily represents the most plausible theoretically, or even the one which worked best. In this respect a distinction may be made between the means used to implement the basic theory of neural integration, and the various mechanisms employed to try to develop a more realistic brain model. The changes introduced into the simulation of the basic theory were aimed at overcoming certain problems which had presented themselves in connection with that theory. These were mostly concerned with control of Excitation and Inhibition and the formation of Integrative Connections. Not all modifications which

were introduced proved to be improvements, but the aim of such modifications was the remedy of difficulties, and it is felt that eventually a solution was obtained. In regard to such basic problems, the most recent version may be regarded as the most definitive form of the model of neural integration.

The position was rather different in respect of the aspects of the model simulating behavioural mechanisms. These mechanisms do not go to the root of the model. They are merely illustrations of how such mechanisms might be implemented by the theory of neural integration. Theoretically, their aim is suggestive only. A number of possible mechanisms seemed worthy of simulation. These are not necessarily alternative explanations; they may be complementary. ". . the brain consists of very large numbers of sub-divisions . . . each with a special architecture and circuit diagram; to describe one is certainly not to describe them all." (Hubel, 1979). For example, more than one of the mechanisms of "reinforcement" described in a previous section may find a place in a sophisticated nervous system, and as argued in Chapters 2 and 3, these may only represent a part of the reinforcement repertoire of an advanced animal.

It would have been satisfying to have been able to simulate a larger scale model in which different versions of a behavioural mechanism could have been incorporated into a single model. However, this would not only have involved a very large model to accommodate these variations, but would also have required the designing of mechanisms whereby the model could select between the use of one variation and another according to circumstances. This was beyond the resources available to the author. Accordingly the strategy was adopted of simulating one version of a mechanism in one set of runs, noting its characteristics and then passing on to another version. These experiments were not necessarily in runs which gave a convincing overall performance in the sense of a decisive demonstration of learning to solve a particular behavioural problem. Nevertheless the manner in which the mechanism worked could be seen, and implications noted which may have been difficult to see at a verbal level. No claim is therefore made that all the modifications reported were demonstrated as necessary to, or even contributory to, the improvement of performance. They are intended to illustrate how behavioural mechanisms might be neurally implemented.

Because a number of variations were simulated it is difficult to avoid confusing the reader as to what was done on particular occasions. On the other hand, the variations can best be understood when described together. Accordingly, it is proposed first to describe the mechansim of Sub-cortical Instrumental control as it appeared in the most recent version of the simulation. This was the version which will in due course be reported as demonstrating overall performance analogical with animal behaviour. In the subsequent Sub-sections of this Section, a number of other arrangements will be described, and their characteristics briefly noted. The reader will of course be reminded in due course of the main features present in any run when its performance is reported.

B. "Time since last emission" mechanism

This was the mechanism used to activate Sub-cortical Instrumental Responses in the most recent versions of the simulation.

It will be recalled that the Sub-cortical Instrumental Array consisted of 4 Neuronal Pools. These Pools did not differ in principle from other Neuronal Pools in the Sub-cortical Arrays of the model: they contained Neurons which belonged to Preset Constellations. However, in the version now being described, the Preset Constellations used to mediate Sub-cortical Instrumental Responses differed from those used to mediate other Sub-cortical Responses, e.g. Preconsummatory, Consummatory or Protective.

The Preset Constellations to which the Integrating Neurons in the Instrumental Array belonged were not Sensori-motor Constellations, as were the Preset Constellations mediating Responses from other Sub-cortical Arrays (at least in this version). Instead they were what may be termed Drive-Motor Linked Constellations. In the case of other Responses (e.g. Preconsummatory), the Sub-cortical Response was activated when an Innately Recognised Pattern (IRP) occurred which matched the Sensory portion of the Preset Constellation. This did not apply to the Neurons of the Sub-cortical Instrumental Array. The Preset Constellations to which they belonged lacked Sensory representation. They did, however, include linked Neurons located in one or other of the Neuronal Pools of the Sub-cortical Motivational Array. For example, a number of Integrating Neurons in the first Pool of the Sub-cortical Instrumental Array were linked to a number

of Neurons in the first Pool of the Sub-cortical Motivational Array: this formed one Preset Drive-Motor Constellation. A number of Neurons in the second Pool of the Sub-cortical Instrumental Array were linked to a number of Neurons in the second Pool of the Sub-cortical Motivational Array, and so forth. (Of course, as with the model generally, the idea was intended to extend to selections of Pools acting together, i.e. Motor Patterns, rather than single columns, but room within the model did not permit the simulation of this).

The action of the Sub-cortical Motivational Array in driving Sub-cortical Instrumental Responses may now be considered. This was extremely simple. Like all other Neuronal Pools, those of the Sub-cortical Motivational Array possessed Afferent Fibres. When one of these Fibres was active, the Neurons within its Pool became Excited. This Excitation was conveyed across the Preset Linked Constellation to the Linked Neurons of the corresponding Instrumental Neuronal Pool. If only one of the Sub-cortical Afferent Fibres was active, then the Efferent Fibre of the Corresponding Sub-cortical Istrumental Pool gave a differential Motor Discharge. If more than one Afferent Fibre was active, then the Discharges of the Sub-cortical Instrumental Pools were in competition with each other. The Motivational Pools drove the corresponding Instrumental Response.

The activity in the set of Afferent Fibres to the Sub-cortical Motivational Array determined which of the Efferent Fibres of the Sub-cortical Instrumental Array would be dominant. In the simulation the Afferents were controlled by a set of notional biochemical clocks. Each Afferent Fibre to the Sub-cortical Motivational Array was provided with a Clock, i.e. a counting variable which incremented at each Instant. The clock was set to zero whenever the Motor Response to which the Fibre related was performed by the model. When the clock registered a certain time interval since being set (the Sub-cortical Drive Threshold) its Afferent Fibre was alloted an increase in Spike Frequency. At each subsequent Instant the Spike Frequency was increased by a further increment. The Four Afferent Fibres were subjected to a notional collateral Inhibition so that only the dominant Fibre carried a Discharge. In the customary manner of the model, this dominant Fibre was defined as Active and allotted the appropriate Spike Frequency due to an active Fibre. Consequently the active Fibre to the Sub-cortical Motivational Array was that Fibre whose related Response had not been exercised for the

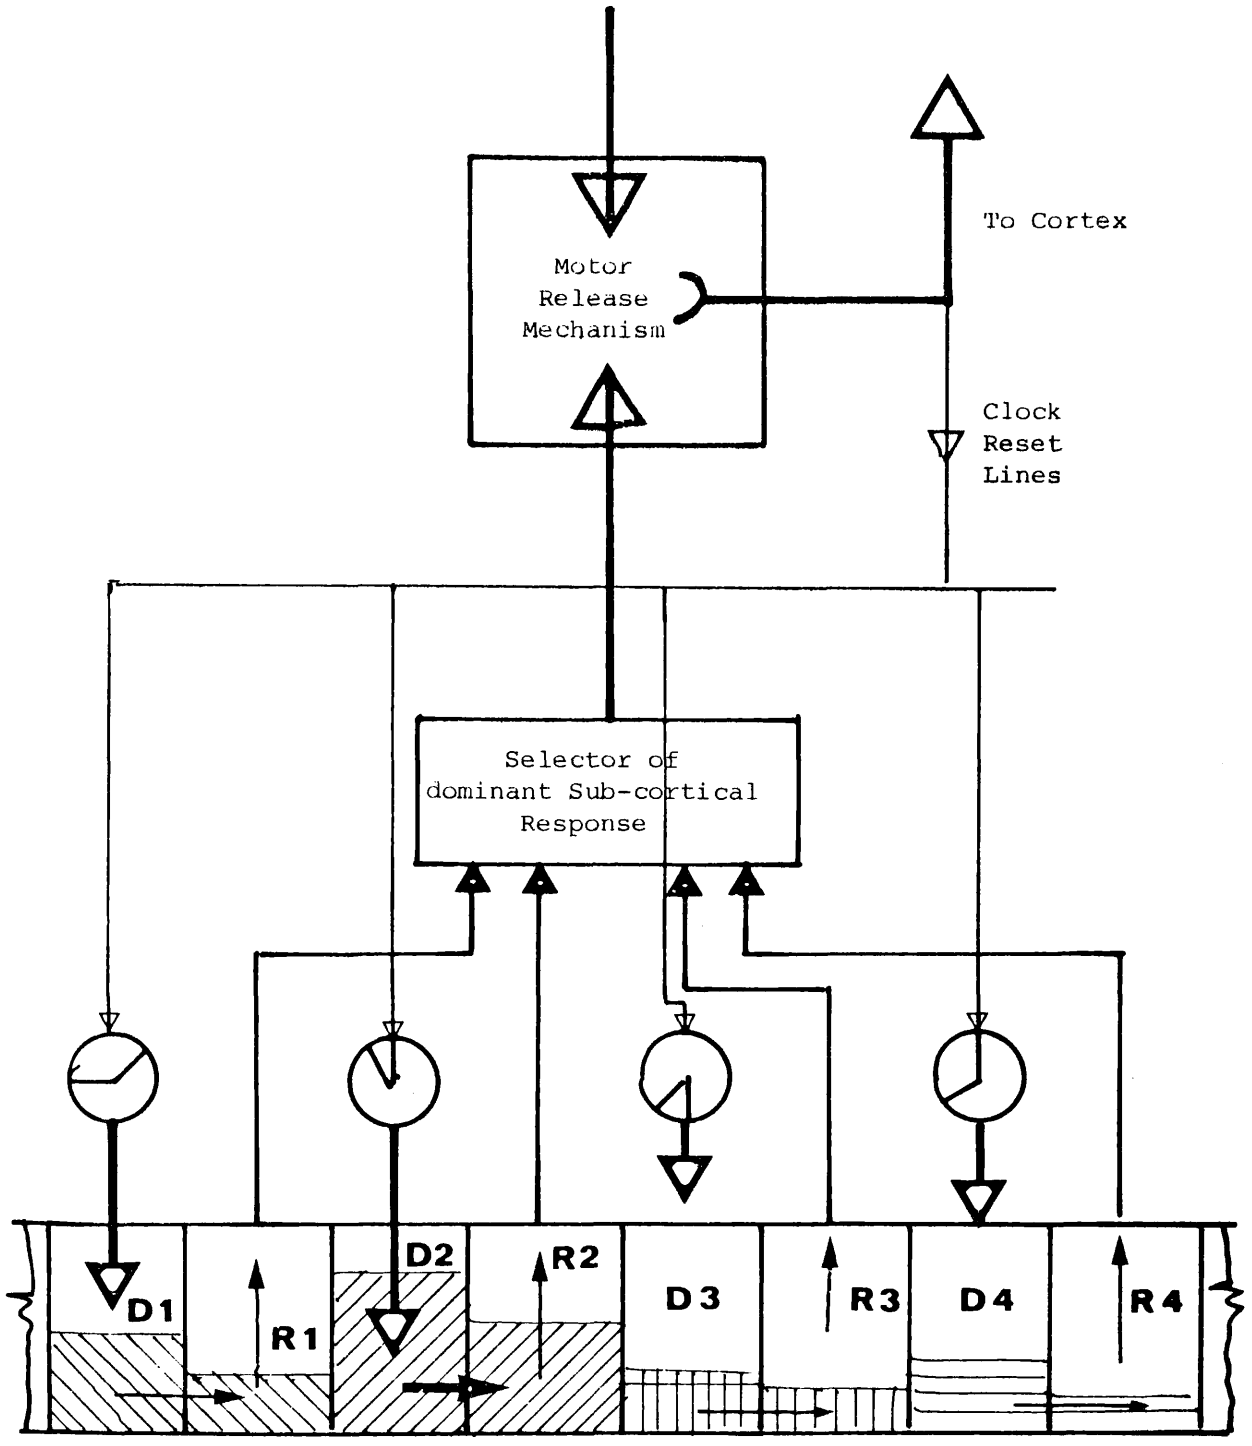

Figure 13.1 Sub-cortical mechanism of motor initiation. Each Response Neuronal Pool (R 1, etc) has an accompanying Sub-cortical Motivational Pool (pressor pools) (D 1, etc). Excitation from Afferent Fibres driven by metabolic or hormonal generators activate the pressor pools. Each metabolic generator acts as a biological clock, providing increasingly intense excitation according to the time since the clock was last reset. The most powerfully driven Sub-cortical Response is passed to the Motor Release Mechanism. The emission of a Response through that mechanism resets the relevant clock.

longest time interval. It will be appreciated that the urging of Sub-cortical Instrumental Array did not guarantee that it was emitted. It was only emitted when the Motor Release Mechanism had occasion to refer to that Array, i.e. upon a failure of Learned Response Indication from the Cortical Array. Nor did it mean that the Sub-cortical Array urged the emission of Instrumental Responses according to some kind of rota. The emission of a Response as the result of a Cortical indication still reset the Clock to zero. In fact, the emission of a Response according to a Cortical Indication was in effect a "release" of the Sub-cortical Response, triggered by the Cortical Indicator (see Figures 12.1 and 13.1).

This system was intended to summarise the notion that for every action the animal can take, there is a pool of Neurons which acts as a biochemical drive unit to press for the exercise of that action. Although expressed in the form of only two independent Motivational Arrays, each having only a few Fibres, the idea is of a principle which may be applied both atomically to the myriad motor elements of the animal nervous system, and more integratedly to organisation of those elements.

Each effector unit, i.e. muscle system, would have its own driving Neuronal Pool. Co-ordinated innately programmed organisations of effectors would have their own driving Neuronal Pools, e.g. pressing for the exercise of walking, head turning. These are the Sub-cortical "Instrumental Responses" of the model. Higher co-ordinations, effected by Cortical functions, relating effector integrations to innately recognised behavioural states, e.g. hunger, sex deprivation, would also have their own "pressor" Neuronal Pools. These latter co-ordinations are the Cortical Instrumental Responses of the model. Their driving Pools are the Pools of the Cortical Motivational Array. In this way the apparently different principles of operation of the Sub-cortical and Cortical Motivational Arrays may be seen as part of a more comprehensive system.

These pressor Neuronal Pools are envisaged as being biochemically controlled, and providing the clocks referred to above. In the case of the main metabolic drives, hunger and thirst, there can be little doubt that such drives exist (Deutsch and Koopmans, 1976). The notion of the biochemical clocks driving individual and co-ordinated effector actions may seem a little strange. However, it would seem plausible

that there are mechanisms which are designed to ensure that an animal's repertoire is utilised. The emergence of the concept that neuro-transmitters are within the class of hormonal messengers (Nathanson and Greengard, 1977; Iversen, 1979) makes the notion more plausible that the neuro-transmitters powering certain classes of neuron might act in a manner analogical to the metabolic and hormonal processes driving hunger and thirst. If such biochemical states can alter the probability of emission of food-seeking and sexual activity, it seems possible that transmitter functions could govern the probability of emission of other more atomic behaviours.

One may conjecture that the afferent fibres to the pressor neuronal pools driving a particular effector action would be driven at an increasing intensity with the passage of time. When the excitation in such a fibre exceeded threshold, there would be a resulting (sub-cortical) output, endeavouring to activate that effector action at the periphery. The biochemical process driving the pressor fibres would be subject to complementary substances directed to its suppression, i.e. neutralisers, the action of which would be to reduce their activity, i.e. to set the Clocks back to zero. In terms of the model, these neutralisers would be released when the action was performed. Roughly speaking, one may imagine that the sensory-feedback from the movement of a muscle, or set of muscles (e.g. walking) would release the neutralising substance (inhibitory transmitter). Although in the simulation the Sub-cortical Motivational Array is separately organised from the Instrumental Pools, this is a matter of convenience only. If the Motivational (driving) fibres are mediated by a distinctive transmitter, its fibres may interpenetrate the Instrumental Pools. This would make the anatomical arrangements for operating the scheme relatively simple.

It would be out of place in this thesis to endeavour to support this speculative system at the theoretical level. It will either seem plausible to the knowledgeable reader or it will not. In the simulation its representation provided a means whereby the Sub-cortical Instrumental Array would urge the emission of a particular Response upon the failure of the Cortical Array to provide a solution. The fact that the Response emitted was that which had not recently been exercised proved a useful strategy in getting the model to emit the required Response.

The reader may be reminded that in no case in the model was the emission of a Response, whether from Cortex or Sub-cortex entirely determined by the Motivational System. In every case the Excitation from the Motivational Neuronal Pools merely biassed the system. It was necessary for other influences to converge to produce the Response. This is intended to apply to the main metabolic drives. As Mogenson and Phillips (1976) argue, citing Fitzsimmons (1972), most feeding is anticipatory of hunger rather than in satisfaction of it. One advantage of the motivational ideas represented in the model is that they leave room for the non motivational influences on behaviour, the effect of input pattern, of time since last emission, and for the learned influences of the Cortical Array.

C. Generalised Preset Constellations

Prior to the introduction of the above scheme, the mechanism for the production of Sub-cortical Instrumental Response was merely a variation of the mechanism used for the other kinds of Response, say, Pre-consummatory, was the control by the sensory part of the Linked Constellation. Instead of Preset Constellations which would be uniquely matched by a single IRP (Innately Recognised Pattern) certain devices were employed to make the emission of a particular Response susceptible to a greater generality of input pattern.

One such method was to have a number of Preset Constellations with linked Neurons within a particular Instrumental Pool. Each of these Preset Constellations had a different Sensory Portion, so that several input patterns would power the Response. Another method was to have the sensory portion of the Preset Constellation deficient in its representation, i.e. only a few active Sensory Fibres were needed to maximally energise the Constellation. Such Constellations were activated by a large number of different Input Patterns.

The effect of these methods was to produce a fairly random Sub-cortical Instrumental emission. There were however difficulties and implausibilities. One of these was the repetition of an error Response, alternately from the Cortical Indication and the Sub-cortical Array, since the response urged by both sources was strongly related to the input pattern. In the simulation this was offset by the general Inhibitory characteristics of the system. It seems unlikely, both in the model and in the nervous system, that alternation of Responses could

result from recurrent neural Inhibition. The time scales for neuronal Inhibition and behavioural suppression do not seem to be the same for this to be effective. This disadvantage did not apply to the clock-like action of the notional pressor pools described in item C.

It would seem that there is some merit in a mechanism in which the probability of a Sub-cortical (i.e. Unlearned) emission is influenced by sensory stimulation. Rats will swim rather than walk when placed in water. Although different ideas of Sub-cortical emission were simulated separately in the model, this was chiefly for the clarity of observation of the results. For example, it would seem quite possible for a biological clock mechanism, such as that described in item C above, to be combined with Preset Constellations with fairly definite sensory representation. The advantage of the extended template concept is that no difficulty arises in combining such influences. One may envisage that the feed-back from the paws is involved in stepping movements, and other sensory feed-back in swimming. This would not interfere with the clock mechanisms which ensure that the muscular co-ordinations involved are exercised upon occasion.

D. Maturational Preset Linked Constellations

This scheme for the Sub-cortical emission of Instrumental Responses departed slightly from the design so far described in that a form of "learning" was introduced at Sub-cortical level. This may recommend the idea to some readers who may consider the Cortical-Sub-cortical division of the model too radical. As discussed further below, the scheme would provide an explanation of imprinting, but is perhaps better regarded as having affinities with developmental processes generally, with imprinting as a particular example.

Although not used in the final versions used to demonstrate overall instrumental performance, this scheme was used extensively throughout the development of the model.

The method was really an extension of that described in C above. In that method, Preset Constellations were designed so that a variety of Input Patterns would favour one rather than another of the Responses at Sub-Cortical Level. In the method now being described, the Preset Constellations were formed by a notional maturational development combined with a special kind of "learning". The Responses "matured" during the course of the run, and the resulting Maturational Constellation

incorporated the Input Pattern current at the time of Maturation. In other respects the Maturational Constellation resembled a Preset Constellation, i.e. it was located at the Sub-cortical level and was fixed for the duration of the run.

The Sub-cortical Neuronal Pools were provided with a number of Unlinked Neurons. The unlinked Neurons were only permitted to form Integrative Connections upon certain specified occasions. These were Instants of notional maturational growth of Connections. The formation of these connections may be considered to represent dependency upon a unique neuro-transmitter system, so that they are only formed when certain specific afferent fibres are active. Upon such occasions the Neurons of the Sub-cortical Array Sensory Pools were able to register whichever Input Pattern happened to have reached the Sensory Array at the Instant concerned. At this Instant, by a preprogrammed instruction, the Neurons of one or other of the Instrumental Pools were excited. Consequently a Sub-cortical Maturational Linked Constellation was formed. This served the same purpose as did the Preset Constellations in relation to the Preconsummatory and Consummatory (Final) Responses. The Maturational Constellation was a Sensorimotor Constellation. It contained on the Sensory side a representation of the Input present at the Instant when its Motor Response had been emitted. To ensure that thelatter Response actually reached the notional periphery at the Instant concerned, the Cortical Indication mechanism had to be suspended at that Instant, i.e. the Motor Release Mechanism gave preference to a Maturational Response, over and above a Cortical Response.

In the later versions of the Maturational Scheme, the scheme was combined with the Sub-cortical Motivational Array method dealt with in Sub-section B. This was effected by including linked Neurons within the Sub-cortical Motivational Array in the Maturational Linked Constellation. The actual implementation of this method was consistent with the ideas sketched out above. The entire set of Sub-cortical Neuronal Pools included Neurons which were Unlinked but which could only be linked by Afferent Excitation bearing the required label (actually a formal label indicating that the Excitation came from Maturationally inspired sources). To form a Maturational Constellation, labelled Afferent volleys were sent through the Afferent Fibres to the Sub-cortical Instrumental and Maturational Pools to be included.

The Input Pattern to the Sub-cortical Array was also labelled so that it would be registered. The result was the Maturational Linked Constellation.

The runs utilising these Maturational Constellations were programmed roughly as follows: a program of Maturation was worked out. Each of the Instrumental Responses and its associated Motivational (pressor) Pools was introduced in turn. Each Response was introduced not just once but several times at spaced intervals. This was so that there would be an increased chance that the Instrumental Response would be "successful" upon at least one of the occasions upon which its Maturation was furthered. The second Instrumental Response was then dealt with in the same way, and then the third and so on.

At each time that Maturation took place, the Response incorporated in the Maturational Constellation was emitted by the model. In the usual manner of operation of the system, a Sensori-motor Constellation was laid down in the Cortical Array. In the case where the Response was successful (elicited a Preconsummatory IRP from the notional environment), then at least in some versions of the model, the Activation (i.e. Do Pool) Instrumental Response was included in the Cortical Learned Constellation. In the case of its being successful, it was hoped that the model would have a good start in learning the implications of the relationship between a Response and the characteristic features of its environment.

To attempt to give the background of ideas on which this scheme was based would, as in the case of the Motivational Arrays, go far beyond the terms of reference of this thesis. Some of the ideas to which it gave expression are merely hinted at to assist the reader in orienting to the model.

One set of ideas which the author had in mind resulted from Piaget's well known descriptions and analysis of the development of behaviour in the human infant. The development of effective Instrumental behaviour in the human infant does not apparently develop in a manner very reminiscent of a problem box or maze in which the infant must choose from amongst a number of competing Responses that Response which will be successful on the specific occasion. Rather the infant develops a particular co-ordination at a particular maturational juncture. It then applies this "technique" to almost every situation it encounters.

In this way it acquires a library of situations when the Response is likely to be of use. As is well known, Piaget outlined a Sensori-motor period. "During the period the infant progresses from a few elementary activities to a wide range of well adapted behaviour" (Etienne, 1973). "This progress involves a constant interaction with the environment, in which the infant tries to apply his pre-existing action patterns to persons and objects, and changes these patterns if they do not lead to the desired results." . . . "The infant is born with a set of elementary action patterns which are relatively unconnected with each other. These include sucking and swallowing, crying, grasping, or visual tracking . . . - these are exercised through repetition and applied to a variety of different objects." . . . "A baby who sucks the nipple of a bottle or the corner of his blanket incorporates or assimilates these external elements into his sucking scheme; in other words, he applies a pre-existing general action pattern to these objects . . ." (Etienne, 1973). Piaget also describes how at a later stage co-ordinations of actions become part of the repertoire by being applied to almost everything within sight. For example, an infant which has learned to open drawers will open every drawer available. Although Piaget's observations were restricted to the human infant, Etienne (1973) has succeeded in applying his ideas, especially the development of object permanence, to other animals, notably kittens.

Another set of ideas leading to the introduction of Maturational Constellations was that of the ethologists concerning "sensitive periods" (see for example, review in Hinde, 1970, pp 460, 462, 516, 517, 565). The discovery of Imprinting by Lorenz and Tinbergen, and the importance of sensitive periods may have been exaggerated in the early days of ethological research. However, it seems to be well established, that sensitive periods exist not only in respect of Imprinting, but also in respect of the development of birdsong, recognition of eggs, and many other behaviours (Hinde, 1973). Research on both Imprinting and the sensitive periods associated with the kind of learning which seems to bear resemblance to Imprinting, has tended to concentrate on the sensory aspects of the process. However, considerable work has more recently been carried out on the motor aspects. Hogan (1973) has examined the sensitive period in relation to the recognition of food by the action of pecking; Baerands and Kruijit (1973) have examined preference of colour in relation to egg-rolling. Bateson trained day

old chicks to press a pedal which turned on a flashing light. He critically examined the notion of a sensitive period and came to the conclusion that it was part of a complex in which experience played a part. "So strong are the links between the imprinting situation and the one in which the flashing light was used as a reinforcer, that it seems valueless to attribute the reinforcing effects to non-specific stimulus change . . . I had thought that once a bird locks onto an appropriate stimulus, the role which that animal plays in the subsequent learning process was a reactive one. Recently, however, I have obtained some evidence suggesting that chicks continue to play an active part in determining what they learn even after the imprinting process is under way." (Bateson, 1973).

One of the advantages of the present model is that it explains such complex interactions without having to relinquish either the idea of a sensitive period, nor the influence on the animal of species specific bias in stimulus selection, nor the influence of "operant conditioning". Maturational Constellations express part of the idea; Preset Constellations express other parts. The Preset Constellation would contain a skeleton representation of a stimulus configuration linked to a Motor Response. When a stimulus occurs which is capable of activating the skeleton network at Sub-cortical level, a full account of that stimulus is registered in the Sensori-motor Constellation laid down at Cortical Level. This is sensori-motor learning guided by the influence of genetically determined stimulus-response probabilities. It would lead to the detailed "learning" of situations of the kind described by Piaget. If this principle is applied to the Maturational Constellation, it assumes in addition the special characteristics of Imprinting: (a) the sensitive period is represented by the period when Neurons at the Sub-cortical level are permitted to form Integrative Connections; and (b) the permanent effect of Imprinting is accounted for in the model by the assumed permanance of the Sub-cortical connections. Moreover, the Imprinted Response takes on the character of an innate behaviour.

It will be noted that the model does not suggest that the normal Cortical mode of learning may not be operative at the same time. It also suggests that the Imprinted Response, once formed, takes on a similar character to the Preset Constellations subserving "appetitive" (or at least, "approach") behaviour. It is not thus surprising that it can act as a reinforcer (see Bateson, 1973).

E. Summary of mechansims of Sub-cortical Response activation

When a "Learned" Motor Indication was not available from the Cortical Array, the Motor Release Mechanism looked to the Sub-cortical Array for the emission of an Instrumental Response. During the development of Aries a number of neural mechanisms to simulate the Sub-cortical Response apparatus, were tried out.

These were:

1. Neuronal Pools whose function was to drive an Instrumental Response at Sub-cortical level. These "pressor" Pools were located in the Sub-cortical Motivational Array. They urged that the Instrumental Response which would be released upon a reference to the Sub-cortical Instrumental Array would be that one which had not been exercised for the longest period. This method was used in the runs designed to demonstrate a definite Instrumental performance in Aries (next chapter). The method was combined with that of Maturational Constellations in many runs.

2. Generalised Preset Constellations were Preset Constellations designed to be activated by a large class of Input Patterns. When the Motor Release mechanism referred to the Sub-cortical Array, the Instrumental Response which would be released was the one most strongly activated by the current Input Pattern. That Response would then be Inhibited by a recurrent Inhibition mechanism, giving way to another Response. This Inhibition did not work satisfactorily, because of temporal delays between occasions of reference to the Sub-cortical Array. As a result the model failed to solve its problem because one or other of the Responses would fail to appear at Sub-cortical level. This defect was considered to be mainly due to limitations of the realisation of the model rather than inadequacies of the principle.

3. Maturational Linked Constellations were formed at Sub-cortical level at preprogrammed Instants of the program. Each Maturational Constellation registered in notional permanent form (i) the current Input Pattern and (ii) a preselected Instrumental Response, and (iii) a Neuronal Pool of the Sub-cortical Motivational Array. Several Maturational Constellations were formed for each Instrumental Response. The Response released upon a reference to the Sub-cortical Instrumental Array was influenced partly by the Input Pattern and partly by the time which had elapsed since the last exercise of the Response. For reasons which will be discussed, in some runs the representation in the Motivational Pool was omitted.

II Protective Responses

Since these did not play a large part in the development of the model, they will only cursorily be dealt with. The principle was similar to that of the Consummatory and Preconsummatory Responses. They were emitted when a Protective IRP (Innately Recognised Pattern) was in the Input. The presentation of such a pattern was arranged by that part of the program representing the environment. It occurred when the model emitted a Response which was notionally an "error" in the context of the test trial. One of the Protective Responses was named Bump, and was elicited when the model attempted to carry out an Instrumental Response which was notionally impossible. The other was termed Hurt, and was elicited when the model attempted to perform a Consummatory or Preconsummatory Response to some Input Pattern other than the appropriate IRP, e.g. when it attempted to notionally drink food.

The Protective Responses were initially involved with the simulation of a complex system of notional orienting reflexes. These had to be cut out when the model was streamlined to produced adequate performance. Unfortunately time did not permit their reintroduction. They will however be mentioned briefly in a later chapter.

One aspect of Protective Responses which was successfully simulated was a hypothetical variation between the neural mediation of "pain" mediated responses and appetitive responses. In the simulation this took the form of a difference between the sensory portion of the Protective Preset Constellations and other Preset Constellations. On the Sensory side, the Protective Preset Constellations did not consist of a complete representation of an Input Pattern, and the Protective IRP did not consist of a self-contained complete Input Pattern.

The Protective IRP consisted of an addition to the current Input Pattern. This addition had the distinctive quality of consisting of a number of "positive" bits together with the corresponding complementary "negative" bits. The Protective Preset Constellation detecting this notional pain was also distinctive in containing a representation of just these positive and negative bits. This arrangement represented the theoretical possibility that pain occurs when both the receptors signalling activity and inactivity are activated simultaneously, a condition which should only occur in abnormal circumstances.

III Conclusion

In this and the immediately preceding chapter, the structure and function of the extended model has been described. It would not be feasible to report in detail the actual performances of the many versions simulated. Moreover, general problems with the model meant that although the operation of a particular mechanism might be seen its overall effect on actual "successful learning" was less clear. Accordingly where mechanisms or modifications are not critical to the demonstration of the model as a whole, the account of performance will be restricted to the indications which have been given in this chapter. However, a number of test computer experiments, similar in intent to those reported for the initial set of simulations, were carried out on the extended model. These are aimed at demonstrating its essential rather than conjectural characteristics. These will be reported in the next chapter.

CHAPTER 14

SOME DEMONSTRATIONS OF PERFORMANCE IN ARIES

I Introductory remarks

In the previous chapter, the general characteristics of a number of mechanisms incorporated at various times into the simulations, were described. It has already been mentioned that whilst the effects of these mechanisms could be discerned by inspection of the internal operation of the model, the overall performance of the model in many of these runs was not such as could be cited to demonstrate to a sceptical observer that its performance resulted from more than chance. Some of the failures of the model to reach a high standard of performance could be attributed to the exaggeration in the model of certain limitations which might apply to animal intelligence. To justify such interpretations and generally to increase the observer's faith in the basic model it seemed desirable to demonstrate that it was capable of performance which could quite definitely not be attributed to chance factors or artefacts of the Input Patterns or run design.

It may be recalled that the computer simulations were referred to by the mnemonic Aries, denoting Arrays of Relatively Isolated Excitation systems.

II Design of tests for the model

The tests of the initial set of simulations approximated to a classical conditioning situation. In a rather general way, the tests of the extended simulations could be considered as analogical to an operant conditioning situation, or to a simple maze situation.

An Environmental Array was provided in the form of a simple two dimensional computing array, consisting of rows and columns. In the demonstration runs this was effectively restricted to a 4 x 4 unit.

Each location contained the identifying Input Pattern Number of the Input Pattern which was to be presented to the model when it was notionally positioned in that location of the Array. The notional organism (referred to for convenience as "the model" was assigned to a particular location in the Array; this was its notional spatial position in the maze. The Input Pattern which was presented to the model was that indicated by the Input Pattern Number written in to that location of the Environmental Array. A modification of the rule that the Input Pattern was that indicated in the location of the model, occurred when the model emitted a Preconsummatory, Consummatory (Final), or Protective Response. In that case the environmental aspect of the program changed the Input Pattern accordingly. For example, if the model were in a location containing an Input Pattern Number referring to an IRP 'representig, say, Food in the mouth, i.e. a Preconsummatory IRP, then when the model emitted the relevant Preconsummatory Response, representing say, chewing, the environmental routines replaced the Preconsummatory IRP by its relevant Consummatory (Final) IRP, which in turn would elicit the Consummatory Response from the model.

The model was moved about the Environmental Array, i.e. notional maze, by its own actions. Three of its Instrumental Responses were interpreted as Left, Right and Forward respectively. It was placed in a starting location in the Array, and then moved into other locations in accordance with its emitted Instrumental Responses. A notional maze could be laid out in the array by making certain locations unenterable. This was effected by entering a specific Number in that location of the array. If the model attempted to enter such a location, either nothing happened, i.e. the Input Pattern and the location of the model was unaltered. Alternatively, in runs where Protective Responses were being investigated, the Input Pattern was modified by the environmental routines to give the model a Bump IRP.

During the early development of the model, the mazes laid out in the Environmental Routines, involved several choice points. In the demonstrations now being reported these were simplified to a simple T-maze with a single choice point. In most of the runs, at this choice point the model could make one of three Instrumental Responses, termed Left, Right or Forward. It was then moved into the location indicated by that Response; unless, of course, that location contained an Unenterable Number. If it succeeded in entering one of the three

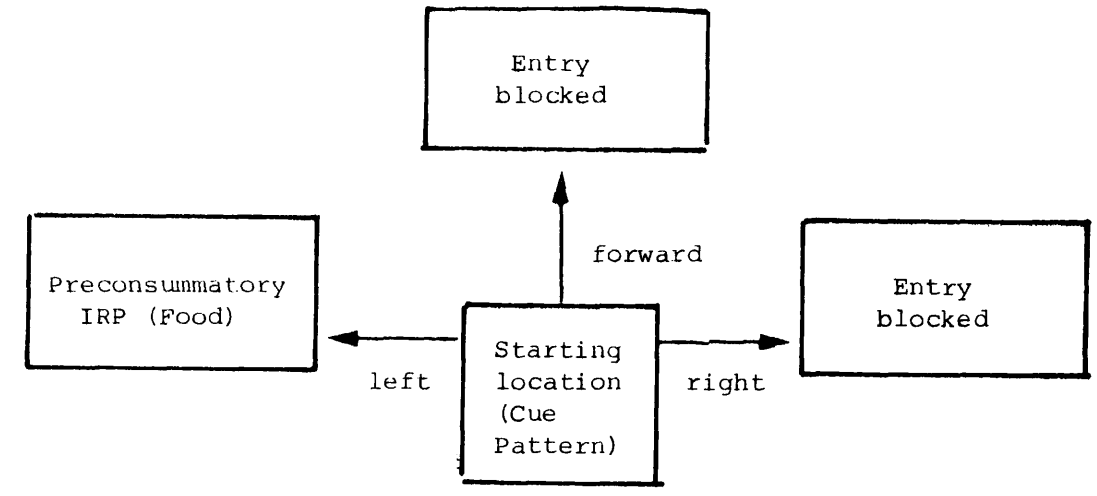

B

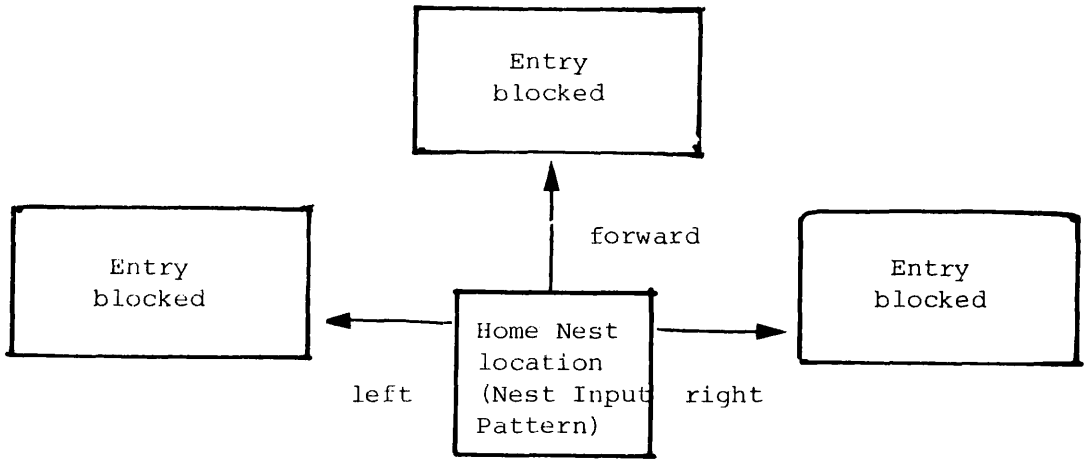

A

Figure 14.1 Test maze. Experimental runs consisted of a number of Trials, at the beginning of which the model was notionally placed in a situation such as that illustrated at B. Whilst in the starting location it was presented with a particular Indifferent Pattern. One of its three Instrumental Responses took it into a location in which it was presented with a Preconsummatory Innately Recognised Pattern (notional Food). Either of the other two Instrumental Responses did not cause the Cue Pattern to change. Between Trials the model was placed in the situation shown at A. In this location it was presented with a particular Input Pattern representing a notional Home Nest, in which case none of its Responses were successful.

possible locations, then its Input Pattern was changed to the one indicated by the Input Pattern Number in that location. That Input Pattern could either be an Indifferent Pattern, as defined in connection with the initial simulations; or it could be a Preconsummatory IRP. In the latter case, a Preconsummatory Response would be elicited from the model, and the Input Pattern would be changed to the relevant Consummatory (Final) IRP (Figure 14.1).

In the demonstration runs, the problem which the model had to solve was to emit the "correct" Instrumental Response to be able to enter the branch of the maze in which resided a Preconsummatory Input Pattern. It then ate the food. Incidentally, the Preconsummatory IRP, representing say, food, was altered to the Final IRP, representing say, "food ready for swallowing", whenever the model emitted the appropriate Preconsummatory Response: it was altered back to the Preconsummatory IRP when it ceased to emit that Response. Consequently the model could spend several Instants in the notional goal box, emitting Preconsummatory and Consummatory Responses, i.e. notionally eating the food. This behaviour was brought to an end by one of two processes: (i) internal motivational processing, not fully investigated and (ii) removing the model from that location, i.e. ending the Test Trial. In the demonstration runs, the latter method had to take preference in order that a sufficient number of Trials could be accommodation within a run of the model.

In many runs, the Test Trial was brought to an end by notionally removing the model from the goal location to a notional Home Nest. In this part of the environmental Array, the location in which the model was placed was surrounded on all sides by Unenterable Locations. The Input Pattern in the Home Nest location was a Neutral Pattern, i.e. an Indifferent Pattern which did not occur during Test Trials. The model was allowed to remain in the Goal location for a given number of Instants before the Trial was brought to an end. This number of Instants was independent of whether the Input Pattern in the Goal location had been an IRP or an Indifferent Pattern.

In the demonstration Runs, the program ran for about 734 Instants. The number of Trials varied, as for some purposes it was necessary to cut down the time spent in the Home Nest location, so as to produce a sufficient number of Test Trials. The number of Trials varied from about six to about seventeen.

III Recap of general mechanism of operation

In most runs, before the Learning Session proper began there was a preliminary portion of the run during which the model was notionally placed in the Home Nest. In runs where Maturational Constellations were incorporated, this gave an opportunity for the Maturationally induced Responses to appear in a Neutral Input situation so that the subsequent behaviour could not be attributed to a manipulation of the Input Patterns occurring at the time of Maturation. Although it was desirable to exclude such relationships between events at time of Maturation and later events, in runs designed to demonstrate the effectiveness of Instrumental Learning, it was of course supposed that in the animal there would be a useful relationship of that kind. In runs where Maturational Constellations were not employed, the preliminary Session in the Home Nest gave some opportunity for the scheme of Subcortical Motor initiation to be seen independently of the Learning Trials.

As earlier mentioned, the Motor Release Mechanism was provided with a gate which preferentially passed a Cortical Motor Indication and suppressed the Sub-cortical Indication. At the beginning of each experimental run there was no Learned Indication, so the emission of a Response for the First time was always from the Sub-cortical mechanism.

At the commencement of the Learning Session the model was notionally placed in a location of the Environmental Array which represented a starting box. After a variable amount of time the model emitted an Instrumental response. The delay was not under the control of the programmer, but resulted from the operation of the Motor Release Mechanisms already described. Suppose that the emission was dictated by the Sub-cortical mechanism: the event was nevertheless registered in the Cortical Array.

Suppose that the emitted Response was the one which the programmer had decided would take the model into the goal box: then, that location would contain the Input Pattern representing a Preconsummatory IRP. To this the model should respond with a Preconsummatory Response. As earlier described, if the model produced such a Response within a given number of Instants then an Active Afferent Volley was sent to the

Effector Neuronal Pool of the Cortical Motor Array. If no such Preconsummatory Response occurred within the given time, then an Active Afferent Volley was sent into the complementary Not-do Neuronal Pool of the Instrumental Cortical Array.

It will be recalled that the Linked Constellation in which this positive or negative aspect of the emitted Response was registered contained a representation not of the Sensory Input Pattern actually accompanying or following the emission, but the Pattern which preceded that emission. Because of the notionally longer Sensory pathway, the Sensory Input Pattern which preceded the emission of the Response reached the Cortical Array at the same Instant that the delayed registration Volley representing the Instrumental Response reached the Cortical Array. Notionally the Sub-cortical emission of a Response acted very rapidly at the periphery, i.e. on the notional environment. The notional pathway of the Sensory information to the Sub-cortical Array was also very short. Accordingly it was possible for the Input Pattern representing a Preconsummatory IRP to produce a Sub-cortical Preconsummatory Response well before it was required to decide whether to register the Do or Not-do aspect of the Instrumental Response. It is interesting in relation to the plausibility of this arrangement to cite some experiments of Libet (1978). In these experiments a shock pulse to the skin was accompanied by direct electrical stimulation of the cortex with a time interval. Libet tested the hypothesis that suitable neuronal activations at cerebral levels must proceed for a substantial minimum time period (about 500 ms) in order to give rise to a conscious sensory experience. "There was a subjective referral backwards in time, after neuronal adequacy was acheived, which antedates the experience to correspond to the time of early cortical responses to specific afferent projection signal". This would appear to be compatible with the idea that the registration of a sensory event can be deferred at the cortex by the requirement of a neuronal adequacy of about 500 ms.

When the model emitted a successful Instrumental Response the Do-Pool was activated; when unsuccessful the Not-do Pool was activated. The Learned Linked Constellation registering the Instrumental Response contained: (i) A representation of the Input Pattern just preceding the emission. (ii) the Instrumental Response of its corresponding Complementary Suppression Pool.

If the Response were successful, the model Responded to the Preconsummatory IRP. The Input Pattern was changed by the Environmental Routine to the Consummatory or Final Pattern. All these events were of course registered in the continuing sequence of Learned Constellations laid down in the Cortical Array. The operation of the Preconsummatory and Consummatory Responses was almost the same as in the initial set of simulations. With regard to the Instrumental Array, the only essential difference was that the Anticipatory rise in Excitation was in a Learned Sensori-Motor Constellation which did not mirror an equivalent Sub-cortical Preset Constellation. It was in a Learned Constellation which registered the conjunction of an Input Pattern and the emission of an Instrumental Response. The rise in Excitation was used to trigger the Cortical Release of that Response on the occasion of a similar Input Pattern being repeated which had on an earlier occasion been followed by reinforcement. If it had been followed on the earlier occasion by "failure" then it was the Not-do Pool which was reactivated by the rise in Excitation, and the Response was suppressed.

The model may have emitted a number of wrong responses before it emitted the right one. Then there would have been written into its Cortical library of Learned Linked Constellations a number of records. In some of these the representation of the critical Input Pattern would be linked to Not Do Responses, and in the case of the correct Response, to the Do Response Pool.

When the model had produced the correct Response which took it into the goal box, it was removed after a prescribed Interval into the Home Nest.

On the next Trial, the model was, say, faced with the same maze situation as at the first Trial. Upon presentation of the Input Pattern in the starting location, all the matching Linked Constellations in the Cortical Array were activated. Since the Input Pattern was not an IRP it did not match any of either the Preconsummatory, Consummatory or Protective Responses, so it did not attempt to emit any of these. (Had an IRP appeared in the input, the elicited Response would have overpowered the Instrumental Response: this was a feature of the model which has not yet been mentioned). The motor Release mechanism detected a differential positive Response from the Cortical Array, resulting from the Excitation of the matching Linked constellation

containing a representation of the "successful" Response i.e. the Instrumental Motor Indication. It emitted this Response, was moved into the goal position, and the events followed as at the end of the preceding Trial. On this second occasion, the Instrumental Response had been emitted as a result of a Cortical Indication, whereas on the first Trial it had occurred as the result of a Sub-cortical initiation.

Barring failure of the notional neural apparatus, the model would have then gone on emitting the correct Response and receiving "reinforcement". It did not always do so however. In many runs there were imperfections of the system. Moreover, the model had to register a sufficiently powerful Learning of events to overcome the differences in activation of Linked Constellations resulting from the chance interaction of non-matching Linked Constellations.

Whenever the model made an error, that is produced an unsuccessful Instrumental Response, it laid down an additional Record in the Cortical Array which tended to inhibit the production of that Response at Cortical Level. Unlike many former versions of the model, this one did not persevere indefinitely in an error. When the error had been sufficiently suppressed by the Cortical Array, then no Cortical Motor Indication was forthcoming at the next Trial. The Motor Mechanism then permitted the Subcortical Array to emit an Instrumental Response. This was the one which had not been emitted for the longest time. In other versions, utilising different Subcortical mechanisms, it might not be the one which had been longest since emission, but in any case, there was an alternation of Subcortical Response. Eventually, therefore, the successful Response was found at Sub-cortical Level.

This mechanism also operated when the notional contingencies of the situation were changed from Trial to Trial. An Instrumental Response which was Correct in the First Trial would be produced at the Second Trial. Unlike earlier versions of the model, this change was adapted to. The model Learned that the once successful Response was now unsuccessful.

As the Protective Responses and associated mechanisms were partially suspended for most of the demonstration runs, they need not be recapitulated here. The Cortical Motivation Mechanism has not been mentioned.

The Motivation mechanism was made ineffective during most learning runs, for reasons which will be mentioned in the next Section. They were however tested in runs designed for that purpose, and the reader will be reminded of their mechanism in context.

IV Demonstrations in extended model

A. Introductory

It has already been mentioned that the performance of the model during development was hampered by a number of problems concerned with the control of Excitation and Inhibition. It was also difficult in these circumstances to distinguish these theoretical problems from those arising from errors in computer programming.

The general progress in obtaining convincing results can be briefly summarised. Before the solution of the general problems, performance was unreliable and marred by a large number of inexplicable behavioural errors. The acceptance of the intelligent aspects of performance required the tolerant co-operation of the observer: it was necessary to penetrate into the intricacies of operation to be convinced that the apparently learned acts of the model really resulted from the characteristics of the system, and its stupid acts from deficiencies of the kind referred to above. After the solution of these general problems, although errors of programming and simulation remained, the performance of the model was sufficiently convincing that its characteristics could not be reasonably doubted.

Even in the final versions of the simulation, the methods by which these general problems of control were solved were grafted onto a simulation which had not been designed with them in mind. In apology for the deficiencies in the final versions, it may be claimed that the simulation served its true purpose. This may be seen not as the production of a technologically high performance simulation but as a tool in discovering the problems which have to be solved in proposing neurological models.

B. Frequency of Response emission

The first reasonable demonstration that the performance of the model obeyed animal-like laws concerned the relationship between the frequency of emission of a Response and its having been Reinforced upon

a previous emission. This was one of the fundamental relationships established by Skinner (e.g. Skinner, 1938; Ferster and Skinner, 1957). It appears implicitly in many demonstrations of Instrumental Learning ever since the pioneer work of Thorndike, Konorski and Miller, Hull and others, and is one of the most frequently used measures of instrumental learning (Hall, 1966, Chapter 3). An alternative related measure is the reduction of the emission of "error" responses, where the error response is defined as some response other than that being reinforced.

Although it has been mentioned several times that clear-cut demonstrations of learning were not obtainable in early versions of Aries in which general problems had not been solved, this statement perhaps overstates the position. Within the model evidence of learning was fairly robust in the face of considerable imperfections. This is illustrated in the run now reported, which was in a version in which control problems were still present, and in which a relatively ineffective mechanism of reinforcement was used.

Table 14.1 summarises the results of a typical run demonstrating the increased rate of emission of the "reinforced Response" and the reduction of emission of competing Responses. In relation to the simulation, the "operant" response of Skinner's experiments was to be equated with the Instrumental Response; the "reinforcement" was to be equated with the presentation by Environment Routines, of a Consummatory IRP, when the Instrumental Response which was emitted was the one arbitrarily designated as correct.

Table 14.1 shows the emission of an Instrumental Response at particular Instants of the run, and the Input Pattern present at that Instant. This may give an erroneous impression of the temporal character of the simulation. The simulation did not consist of a presentation of a different input pattern at each Instant, and recorded an Instrumental Response emitted at each Instant. The run reported ran for over 400 Instants. Instant 1-100 were occupied in permitting Responses to Mature (Maturational Constellations, see Section V (D) of Chapter 10). Input Pattern 39 was then presented at each successive Instant until the notional animal emitted a "successful" Instrumental Response at Instant 135. To give the reader some flavour of the pace of events and the temporally continuous character of the simulation, Table 14.2

TABLE 14.1 Response summary of a run of Aries demonstrating effects of reinforcement of Response emission

Response order	Input Pattern when R emmitted	Instrumental R emitted	Correct or Incorrect Response Score	Reinforced or not
1	39	"2" - Left	-	No
2	39	"4" - Right	+	Yes
3	39	"4" - Right	+	Yes
4	47 Home Nest	"4" - Right	+	No
5	39	"3" - Forward	-	No
6	39	"2" - Left	-	No
7	39	"2" - Left	-	No
8	39	"4" - Right	+	Yes
9	39	"4" - Right	+	Yes
10	47 Home Nest	"4" - Right	+	No
11	39	"2" - Left	-	No
12	39	"4" - Right	+	Yes
13	39	"4" - Right	+	Yes
14	47 Home Nest	"4" - Right	+	No
15	39	"4" - Right	+	Yes
16	47 Home Nest	"4" - Right	+	No

Summary: Whole Program

	Response No. 2	Response No. 3	Response No. 4
Number of responses emitted:	4	1	11

First Half compared with Last Half of run

	First 7 Responses		Responses 8-16	
Number emitted	"Correct" 3	"Errors" 4	"Correct" 8	"Errors" 1

Note: Reinforcement Rule: Pattern No. 47 in Input — No Response reinforced

Pattern No. 39 in Input — Instrumental R. "4" reinforced

TABLE 14.2 Details of Run shown in Table 14.1 illustrating continuous temporal character of simulation

Instants of Program	Input Pattern	Instrumental R. emitted	Preconsummatory R. emitted	Consummatory R. emitted
100 to 111	39	-	-	-
112	39	"2"	-	-
113 to 134	39	-	-	-
135 to 137	39	"4"	-	-
137 to 138	9 (PC. IRP)	-	-	-
139	9	-	"2" (PC. R.)	-
140 to 141	1 (C. IRP)	-	-	"2" (C. R.)
142 to 147	9	-	-	-
148 to 157	47 Home Nest	-	-	-

Notes: At Instant 112 an Instrumental Response "2" was emitted. Since this was not notionally the "correct" Response, the Input Pattern was not changed by the Environmental Routines. No Preconsummatory Response was emitted and the Response was not reinforced by the internal neural mechanisms. At Instant 135 Instrumental Response "4" was emitted. This was notionally "correct". The Environmental Routines changed the input pattern to 9, the IRP for the Preconsummatory Response (see text). The emission of the Preconsummatory Response at Instant 139 operated the notioanl internal neural mechansims of reinforcement.

PC. R. = Preconsummatory Response

C. R. = Consummatory Response

summarises events occurring during Instant 100-157. Although the program was not always consistent in the relationship of an Instant to notional real time, some idea may be obtained of what was in mind by taking an Instant equal to about 50 ms.

The mechanism of reinforcement in this run, although obeying the general principle of the reinforcement mechanisms described in Section III of Chapter 12 differed in a significant detail. In the run shown in Table 14.1, all Instrumental Responses were permitted to be registered in the Cortical Array: that is, the Afferent Volley which would ensure the Motor representation in the Linked Constellation, was not blocked. This registration took place whether the Instrumental Response was successful or unsuccessful. The internal criterion of success remained the same in this run as in the general principle already laid down, that is, the subsequent occurrence of a Preconsummatory Response. It was however only implemented in relation to the Suppression (Not-do) aspect of the neural processing.

Briefly, after an Instrumental Response had been emitted, an Active Afferent Volley was sent to the relevant Effector (DO) Neuronal Pool of the Cortical Instrumental Array. This registration of the effector aspect of the system ("Do Response") happened whether or not the Response was later judged to have been successful or not, that is, whether a Preconsummatory Response followed the emission of the Instrumental Response. Since the Instrumental Response was to be registered in the Effector Pool regardless of its notional environmental effect, there was no need to delay its registration, as was described in the mechanism detailed in Chapter 10. If, within a given number of Instants following the emission of the Instrumental Response no Preconsummatory Response had been emitted by the Subcortical Preconsummatory Array, an Active Afferent Volley was sent to the complementary Neuronal Pool (Not-Do), so that a Linked Constellation containing the Not-Do representation was laid down. If however within the number of Instants allowed a Preconsummatory Response had been elicited from the Preconsummatory Sub-cortical Array, the Afferent Volley to the Not-Do Pool was blocked. Consequently, if a Preconsummatory IRP occurred within a certain number of Instants (13) of emitting an Instrumental Response, that Response was reinforced by the absence of a Not-do represention.

This system of reinforcement combined two principles. On the one hand it provided a system whereby "successful" Responses were

represented by Linked Constellations having only "Do" representation, whereas "unsuccessful" Responses had balanced "Do" and "Not-Do" representation. On the other hand it possessed features in common with the combined "Do" and "Not-do" system described in Chapter 10 as a variety of Contiguity Learning. In the case of the unsuccessful Response, both the Do and Not-do representation was in Linked Constellations in which the Sensory part was the pre-instrumental Input Pattern. If the Instrumental Response succeeded in effecting a change in the Input Pattern, then even in the absence of a Preconsummatory Response, the Not-do aspect was registered in a Linked Constellation of which the Sensory part was the post-instrumental Input Pattern. The model would then be partially independent of having to emit a Preconsummatory Response with sufficient rapidity to block the Not-do registration. On a repetition of the pre-instrumental Input Pattern, in the case where a previous emission of the Response had notionally succeeded in changing the environment, only the effector aspect would be directly reactivated. The suppressor aspect would only be directly activated after the emission of the Response. Of course this would have dangers. The change effected by the Instrumental Response might not be beneficial. However, in that case a Protective Response might be elicited so that any such results would be clearly labelled in the Learned Records.

In the run reported in Tables 14.1 and 14.2, the action of the Motivational Cortical Array was suspended. The behaviour of the model resulted from the effects of notional sensory discrimination. There was no attachment of "drive excitation" to the Linked Constellations representing previous successes.

The reason for the suspense of the Cortical Motivational Mechanism in this run illustrates one of the characteristics of the model. Even in the early versions of Aries, it was relatively easy to obtain a demonstration that the frequency of emission of Responses was a function of the reinforcement of previous emissions. This resulted from the very powerful effect of the Motivational Array. This aspect of the model will be reported more fully below. Here it is necessary to state that in early versions of the model in which the Cortical Motivational mechanism was active (i.e. those described in Section III, Chapter 10) once a successful Instrumental Response had been emitted, it continued

to take preference in the Learned behaviour more or less indefinitely. The only way to obtain another Response was for the Motor' Release Mechanism to refer to the Sub-cortical Instrumental Array. The Instrumental Response was no longer under discriminative control. Probably some of this effect was due to the making of the Motivational effect too strong. Nevertheless, it was felt desirable to demonstrate the relation between reinforcement and probability of emission in the absence of this Motivational effect.

Even without the Motivational effect, Table 14.1 demonstrates some of the difficulties of demonstrating a clear-cut effect in the model. It will be noted that the Input Pattern was not restricted only to that of the notional trial situation, i.e. No. 39. Pattern 47 was also presented, representing the Home Nest (see Section II). This came about because the notional maze described in Section II was used for the runs. When the model made a Response which took it into one of the branches, after some Instants it was placed in the Home Nest for some Instants. During this period it might emit a Response. Had it responded with sufficient intelligence to this situation it would have learned not to Respond when 47 was in the Input. It failed to do this; what it learned to do was to Respond with the reinforced Response in both situations. However, it can be seen from the Table that the unreinforced Responses account for the errors. In fact, from internal monitoring, the progress in learning was found to be as follows. The model learned too well on the first two occasions of being reinforced that Response 4 was the one to emit. When it emitted that Response when 47 was in the Input, it was not reinforced. It then learned too well, not to make this Response. On the next Trial Response 4 was suppressed, even although Input Pattern 39 was presented. This confused it. Again at Response No. 10, it Responded to Pattern No. 47 with the previously reinforced Response No. 4. It failed to gain reinforcement. Its next Response, which should in fact have been Response No. 4, the Input Pattern having changed, was instead an error Response. However, on this occasion the performance was only disturbed for 1 Instance, as against 3 Instances on the earlier occasion.

What seemed clear from this and later runs to be reported was that the model could not learn to suppress its Responses or Release them entirely as the result of sensory discrimination. Although in this run the model failed to achieve a high level of discriminative Response

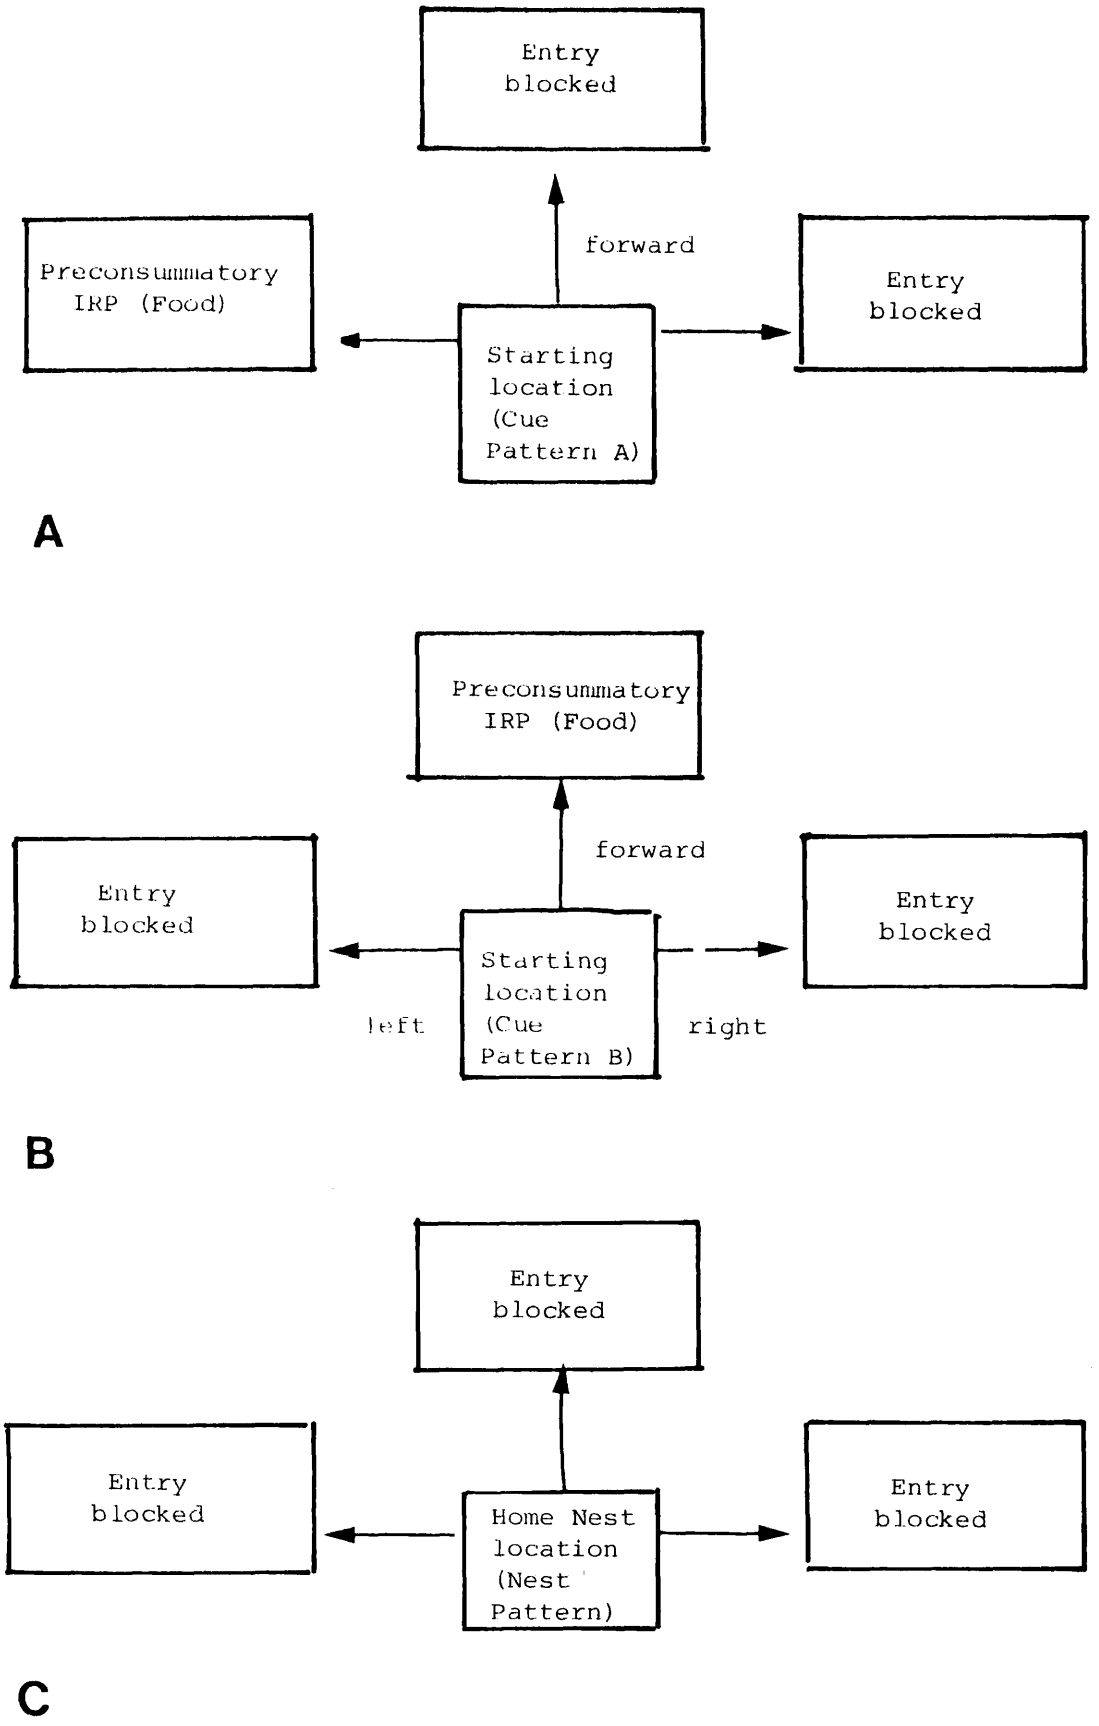

Figure 14.2 A variant of the Test situation shown in Figure 14.1. On some Trials one Input Pattern would be used as a Cue, i.e. presented to the model in the starting location, and on other Trials another Input Pattern was used. The Cue Pattern presented to the model indicated in which position it could find the Food. Between Trials it was, as before, placed in the Home Nest.

control, its errors were explicable in terms of a tendency to jump to conclusions, and in this respect did not seem unlike behaviour shown by animals and even the author.

C. Demonstration of Discriminative Instrumental Performance

In order to demonstrate in a clear way that the model was capable of producing a discriminative instrumental performance, the notional maze described earlier had to be adapted slightly, Figure 14.2. It will be recalled that each run was divided into a number of trials. At the beginning of each trial, the notional organism (the "model") was placed in the starting location of the maze Array. The three Instrumental Responses capable of being emitted by the model, may be thought of as moving it to the Left, Right or Forward of the starting location. The Input Pattern placed in the starting location was the Discriminative Cue for that Trial. It was presented to the model at each successive Instant of the Trial until the model emitted an Instrumental Response to take it into another location. Entry to a location could be notionally blocked by inserting a particular Input Pattern in that location. In only one of the three locations surrounding the starting location, the program placed an Input Pattern representing a Preconsummatory IRP (Innately Recognised Pattern). The selected Input Patterns and their location could be changed from Trial to Trial. The Input Pattern in the starting location, that is, the Discriminative Cue, was generally used to indicate the location of the Preconsummatory IRP. For example, in one run, when the Discriminative Pattern was No. 39, the Preconsummatory IRP was to be found by the model emitting Instrumental Response NO. 4; that is, it was in the location to the Right of the starting location. When the Discriminative Pattern was No. 43, the Preconsummatory IRP was situated on the Left of the starting location and could be reached by the model emitting Instrumental Response No. 2.

In many runs of the model, the model was permitted to enter any of the three locations by emitting the relevant Response. It was then presented with the particular Input Pattern in that location. This arrangement had been useful in investigating a number of aspects of the model already mentioned. However, it impeded obtaining a definite demonstration of Instrumental performance for a number of reasons. The principle reason was the restriction on the number of Trials obtainable in relation to the number of Trials necessary to establish an intelligent performance. If the model were allowed to enter locations not containing the reinforcing

IRP, i.e. locations which contained only Indifferent Patterns, many Trials were unreinforced, and the development of the intelligent performance was deferred. Although, by the end of the run the model could choose the right branch, there was insufficient time left of the run to show that the choice resulted from the intelligence of the model and not from pure chance. What was necessary was to show that the overall choice of the correct position of the reinforcement (Preconsummatory IRP) was at a greater frequency than could be explained by chance. For this purpose it was necessary to produce a relatively large number of Responses over the length of the run.

Accordingly, in the runs used to demonstrate Discriminative Instrumental Performance, in each Trial the only location which the model was permitted to enter was that containing the Preconsummatory IRP. In effect, this meant that the model could emit one of the three Reponses, but only one of these, the one that led it into the location containing the IRP, had any effect on the Input Pattern. The other two left the model in the starting location, and the Input Pattern unchanged.

The procedure for each Trial was then as follows. The model was presented with a particular Input Pattern, the Discriminative Cue, until it emitted the particular Response which had on that Trial been designated by the program as successful. Emission of the other two Responses had no effect on the Input Pattern. Emission of the Response designated as successful caused the program to change the Input Pattern to the Preconsummatory IRP. This elicits the Preconsummatory Response from the model. The Preconsummatory Response operates the internal reinforcing mechanisms. Its emission also causes the Environmental Array to change the Input Pattern once again, this time to the Consummatory (Final) IRP. The Consummatory and Preconsummatory Reponses were permitted to alternate for some Instants, representing the eating of food or drinking of liquid. After a given number of Instants in this situation, the Trial ended. The model was then presented with an Indifferent Pattern, one unrelated to any of those occurring within Trials. This represented placing the animal in a Home Nest, in which it spent its time between Trials. After a given number of Instants the next Trial began.

A run which has been selected as showing in the clearest form the basic

TABLE 14.3

Discriminative Instrumental Performance

Trial No.	Instants	Input Pattern	First Response of Trial	Subsequent Responses 2nd 3rd 4th	First Resp. Correct	Second Resp. Correct
1	85 - 107	43	4	3	No	Yes
2	129 - 150	43	3	-	Yes	
3	169 - 246	39	3	2, 2, 4	No	No
4	265 - 285	39	4	-	Yes	
5	304 - 327	43	3	-	Yes	
6	346 - 363	39	4	-	Yes	
7	382 - 403	39	4	-	Yes	
8	422 - 444	39	4	-	Yes	
9	463 - 484	43	3	-	Yes	
10	503 - 521	43	3	-	Yes	
11	540 - 559	39	4	-	Yes	
12	578 - 595	39	4	-	Yes	
13	614 - 674	43	4	2, 3	No	No
14	693 - 744	43	2	end program		

Total Number of Responses: 20

Total Number of Trials: 14

Correct First Time: 10

Total number of Correct Responses: 13

Reinforcement Rule: Input Pattern 39 - Response No. 4 Reinforced

Input Pattern 43 - Response No. 3 Reinforced

Probabilities

UACTOG6: 17.1.79

Correct first response of Trial: 10 out of 14. Z = 3.047
Probability < 0.002

Total number of correct responses: 13 out of 20. Z = 3.019
Probability < 0.002

performance of the model is illustrated in Table 14.3. In this run it will be seen that in the first three Trials the model performs little better than chance. If the performance of the model were entirely due to chance, that is, if any of its three Instrumental Responses are equally likely to be "correct" on any Trial, then the probability of a correct Response is 1/3. The probability, on a chance basis, that the First Response of any Trial will be correct is similarly 1/3. In the next nine Trials, the model produces nine correct Responses, that is, its First Response on each of these nine Trials was correct. The probability of this occurring by chance is less than 0.002.

Of course, it is appreciated that statistical procedures can only be applied to the results of computer simulations with great caution. They must also be properly interpreted. The probability is only mentioned to make it clear that the performance shown in this run can hardly have occurred by chance.

It is not to be thought that, in general, runs of ths model showed this level of performance.

In order to detect a progressive learning effect, the beginning of runs were compared with the latter portions of the runs. The division of Trials in this way was difficult, because the presentation of Cues was different from run to run. Some orders of presentation presented more difficulty to the model than others. Consequently the number of Trials to be allocated to the first portion to display the learning effect differed from run to run. Accordingly an arbitrary division was made between the first four Trials and the remainder. This gave a positive result in every run, but the statistical significance was much lower, around 12 per cent. This was not surprising, as in effect the model learned quite a lot from a single Trial. One of the troubles with the simulation was that to produce an effect within the duration of a run, learning had to be speeded up beyond what would be analogical in animal learning, in which many trials are necessary to modify behaviour.

The general level of performance achieved in the average run of the model is shown in Table 14.4 in which a number of runs are summarised. It will be seen that the number of correct responses achieved varies from run to run, and in some cases is not superficially impressive as

TABLE 14.4

Discriminative Instrumental Performance

(Results of several experimental runs)

A. Overall results

Run Identifier	No. of Correct First Responses	No. of Trials	Total No. Correct	Total No. of Responses
UACTO1H	6	13	13	24
UACTOG1	8	15	15	24
UACTOAU	7	13	13	22
UACTOG6	10	14	13	20
UACTOJ4	5	13	13	24
	36	68	67	114

B. Runs divided between results of first four Trials and remainder of run

Run Identifier	First Four Trials		Remainder	
	Correct	Total of Responses	Correct	Total of Responses
UACTO1H	4	7	9	17
UACTOG1	4	7	11	17
UACTOAU	4	9	9	13
UACTOG6	4	8	9	12
UACTOJ4	4	8	9	16
	20	39	47	75

Probabilities

The probability of a Response being correct by chance is 1/3.

(1) Correct Fist Responses: 36 out of 68 compared with 22 out of 68 by chance

$Z = 4.15$. Probability less than 0.002 per cent.

(2) Total Correct Responses: 67 out of 114 compared with 38 out of 114 by chance

$Z = 5.901$. Probability infinitesimal.

(3) Comparison of first four Trials with remainder (progressive learning) First part of run: 20 out of 39. Later part of run: 47 out of 75 $Z = 1.17$. Probability due to chance about 0.12.

an example of Instrumental Learning. Nevertheless in all these runs it could be seen from Internal monitoring that the learning process operated. Taken together the overall probability of the number of correct Responses occurring by chance (67 out of 113) is infintesimal ($Z > 5.0$), even less than that of the individual result shown in Table 14.3. There are reasons, which will be mentioned below, why the number of correct Responses within the set of First Responses in Trials are considerably poorer than would be expected from the overall correctness of Responses. Even in this category, the probability of the result occurring is only 0.002.

It is perhaps worth examining a run which indicates the problems of obtaining the kind of reliable performance sometimes obtained from computer simulations of neurological models, from a model which attempts to approach more realistically to the neural subject matter. Table 14.5 illustrates such a run. Actually this version learned a little too rapidly. It will be recalled that rapidity of Learning was a function of one of the parameters of an Integrating Neuron, the Memory Life. In conjunction with a counting variable, this controlled both the probability of a Neuron becoming a member of a Linked Constellation and its being removed from the Linked Constellation (decay). Consequently, if the model learned too rapidly it tended to forget rapidly also.

In the run illustrated, the model found the correct result at its third attempt. It then emitted in the next two Trials two correct Responses first time. In the next Trial, (No. 4) it was presented with an Input Pattern which it had not seen before. In the ideal simulation it would have responded with an orienting reaction. In the present simulation it emitted the Response which had proved successful on three previous Trials. However, on this Trial the contingencies were different. It had to emit two further Responses to find the one which was successful in the revised contingency indicated by the Discriminative Cue, Input Pattern No. 39.

The next Trial, No. 5, reverted to the earlier contingency, and the model was able to produce a successful Response first time. Although on the next Trial the contingency was changed once again to that which it had experienced only once, in Trial No. 4, it succeeded in producing a correct Response upon both this and the next Trial. Then it was

TABLE 14.5

Discriminative Instrumental Performance

A run illustrating some typical errors of performance

Trial No.	Instants	Input Pattern		First Resp. of Trial	Subsequent Responses 2nd 3rd 4th	First Resp. Correct	Second Resp. Correct
1	86 - 133	43		"4"	"4", "2"	No	No
2	152 - 184	43		"2"	-	Yes	
3	203 - 218	43		"2"	-	Yes	
4	237 - 289	39		"2"	"3", "4"	No	No
5	307 - 326	43	*	"2"	-	Yes	
6	344 - 363	39		"4"	-	Yes	
7	418 - 436	39		"4"	-	Yes	
8	455 - 473	43		"2"	-	Yes	
9	491 - 525	39	*	"2"	"4"	No	Yes
10	543 - 578	43		"4"	"2"	No	Yes
11	603 - 629	39	*	"3"	"4"	No	Yes
12	648 - 681	43	*	"3"	"2"	No	Yes
13	699 - 717	39		"4"	-	Yes	

* Responses emitted too early to be attributed to Response Indication from Cortical Array.

Total Number of Responses: 21

Total Number of correct: 13

Total number of Trials: 13

Total Correct First Time: 7

Total correct First and second time: 11

Reinforcement Rule: Input Pattern 39 - Response No. "4" Reinforced

UACTOA: 25:10 78: Input Pattern 43 - Response No. "2" Reinforced

able to react to a reversion to the earlier contingency when Pattern No. 43 was the Discriminative Cue. However, Trials 9 - 12, show a deterioration of performance. Nevertheless, its second Response is the correct one. In that respect the performance of the model differs clearly from what would be expected by chance. This behaviour is in fact clarified by a Printout monitor at the Instant of emission of the incorrect first Responses of each Trial. Three of the four Responses were emitted too early in the Trial to be the result of control by the Discriminative Cue.

To explain the last statement, it is necessary to remember that because of the Uttley-type delay, the Input Pattern reaches the Subcortical Array some Instants before it reaches the Cortical Array. The Learned Indication of Motor Response arises in the Cortical Array from reactivation of the Learned Linked Constellations, and takes several Instants to develop. The correct Anticipatory Response related to a Discriminative Cue cannot be expected to be forthcoming until the Excitation representing it has been reaching the Cortical Array for several Instants. The formal timing of the beginning of a Trial was the Instant at which the Discriminative Cue was first presented to the model, which was some 14 Instants before it reached the Cortical Array. If it emitted a Response during this interval, the Response was related to the Input Pattern representing the Home Nest, and not to the Discrimintive Cue presented in the current Trial. Upon inspection, it was in fact found that in three of the four Trials concerned, Nos 9, 11 and 12, the incorrect first Response was emitted only two Instants after the Discrminative Cue had reached the Cortical Array.

The analysis of this run illustrates the problems of representing behaviour on a time scale which has to be greatly contracted in comparison with that of the organisms. Some of this problem should disappear with the expansion of resources of computers.

This run also pointed to an inadequate suppression of Responses by the Indifferent Pattern representing the Home Nest location. This inadequacy was also partly attributable to the limitations on the time scale. The emission of a Response whilst the model was notionally in the Home Nest led to the inclusion in the Linked Constellation registering that Response of a Suppression (Not-Do) Pool. Responses emitted in the Home Nest occurred only infrequently, because the model was in the

Home Nest for only a short time. Moreover, since no Response was reinforced in the Home Nest, there was no single Response which became dominant in the Home Nest situation. The Motor Release mechanism tended to try to find a dominant Response, and would often find one when in fact there was no Response which had been sufficiently reinforced. There were a number of possible remedies for this situation. For example, it would seem probable that when an animal keeps still it is not always because there is no indication of a Learned Response. There may well be a Standstill or similar concept of Response. The Cortical Array might then indicate a Standstill Response and thus not seek a Sub-cortical Response Indication. Some attempts were made to simulate an arrangement of this kind, but unfortunately time did not permit its development of a successful conclusion.

Before leaving the discussion of runs used to demonstrate Instrumental Performance, the reader may wish to be reminded of the actual versions of mechanisms incorporated in the run illustrated in Table 14.3. In that run, the mechanism of reinforcement was that described in Section II (D) of Chapter 12; that is, the registration of an emitted Response was delayed to give an opportunity for a Preconsummatory Response to be elicited. If such a Response did appear, the Effector aspect was registered by an active Afferent to the relevant Neuronal Pool (the "Do" Pool). If it did not appear within the allotted time, the active Volley was sent to the Not-Do Pool. This was found to be the most efficient of all the mechanisms experimented with.

The method of selecting all Sub-cortical Responses was that described in Section V (B) of Chapter 10; that is, Sub-cortical Responses were driven by Sub-cortical Motivational Pools controlled by notional Clocks, giving a "time since last emission" dominance.

The Cortical Motivational mechanism was suspended in its main effect. It will be recalled that the main effect of this Array was to feed Excitation into Linked Constellations representing previous occasions when a goal had been successfully reached. Because of reasons already mentioned in relation to earlier runs, these were made inoperative. Runs investigating their effect are reported in the next Section.

D. Demonstration of Motivational Behaviour

The run illustrated in Table 14.6 demonstrates the influence of the notional motivational state of the model on its learned instrumental

TABLE 14.6

Demonstration of control of Instrumental Responses by notional Motivational States

Trial No.	Instants	Input Pattern	Notional Motivational State	Instrumental Responses First	Instrumental Responses Subsequent 2,3,4,5,6,7,8,9	First Response Correct	Notional character of Reinforcement	Consummatory Activity
1	109 - 122	43	Hungry and Thirsty	F		Yes	Drink	Drinking
2	141 - 239	43	Hungry and Thirsty	F	F,F,R,F,F,F,L	No	Food	Eating
3	258 - 276	43	Hungry and Thirsty	F		Yes	Drink	Drinking
4	295 - 342	43	Hungry and Thirsty	R	F,L	No	Food	Eating
5	361 - 374	43	Only Thirsty	F		Yes	Drink	Drinking
6	393 - 411	43	Only Thirsty	F		Yes	Drink	Drinking
7	430 - 454	43	Only Thirsty	F		Yes	Drink)	Failed to
8	473 - 493	43	Only Thirsty	F		Yes	Drink)	Drink (did not respond to IRP)
9	512 - 531	43	Only Hungry	L		Yes	Food	Eating
10	550 - 572	43	Only Thirsty	F		Yes	Drink	Drinking
11	591 - 609	43	Only Thirsty	F		Yes	Drink	Drinking
12	629 - 642	43	Only Hungry	L		Yes	Food	Eating

Key: F = Forward, Instrumental Response No. 3; L = Left, Instrumental Response No. 2; R = Right, Instrumental Response No. 4. The successful Response is underlined.
Hungry only: Motivational Afferent No. 2 Active; Thirsty only: Motivational Afferent Fibre No. 3 Active; Hungry and Thirsty: Motivational Afferent Fibres Nos. 2 and 3 both Active.
Eating = Preconsummatory Response No. 2 emitted. Drinking = Preconsummatory Response No. 3 emitted.

Run Identifier: UACTOQZ 81079

behaviour. For the sake of ease of following the performance, the IRP (Innately Recognised Pattern) eliciting Preconsummatory Response No. 2 has been labelled "Food"; that eliciting Preconsummatory Response No. 3 has been labelled "Drink". The conditions of the run were as follows. Except for the preliminary Session, the Input Pattern was always 43 during Trials. During inter-trial periods when the model was notionally in the Home Nest, the Input Pattern was 47. The Food was always in the location requiring the emission of Left (Instrumental Response No. 2) to bring the model into that location. The Drink was always in the location requiring the emission of Forward (Instrumental Response No. 3) to bring the model into that location. The location which would have been entered by the emission of Right (Instrumental Response No. 4) was always blocked.

At the beginning of each Trial the model was notionally placed in the starting location. From the beginning of the run until Trial No. 5 the model was both Hungry and Thirsty. This meant that the Afferent Fibres to Cortical Motivational Pools Nos. 2 and 3 were active at the Lower intensity i.e. at a level which biassed the system but did not induce linked Neurons in the Neuronal Pools concerned. It will be recalled that the occurrence of an IRP sent the appropriate Motivational Afferent to the Upper intensity of activity inducing linkage. Only one of the two maze locations containing the Preconsummatory IRPs could be entered. The other was blocked. This ensured that whilst the model was notionally both hungry and thirsty it experienced both forms of reinforcement and "discovered the whereabouts of the Food and Drink". It will be seen that on the First Trial, its first Response was to the location which was open, and contained Drink. On the second Trial this contingency was changed. It was the Left location, containing Food, which was now open. The model took eight Instrumental Responses to discover this.

On Trial No. 3, the contingencies reverted to those of the first Trial. The model preferred the correct Response and once again achieved reinforcement first time. On the next Trial, No. 4, the contingency was changed again. The model could not enter the Forward location. This time it took only three Responses to find the open Left compartment. This part of the run was directed to forcing the model to gain experience of the "position" of both Food and Drink.

In Trial No. 5 and all subsequent Trials both locations containing IRPs were notionally open. On Trials 5-8 inclusive, the model was made only thirsty, i.e. only the Afferent Fibre to Cortical Motivational Pool No. 3 was Active at the lower intensity level. If the model made a Response inappropriate to its motivational state, it could enter the compartment containing the reinforcement inappropriate to its motivational state, instead of going to the compartment containing appropriate reinforcement. However, as will be seen from Table 14.6, it did not do so. For the remainder of the run, the model emitted only those Responses which would take it into the compartment containing reinforcement appropriate to its Motivational State, and when that state was changed it changed its Response accordingly. Its behaviour was almost perfect in this respect.

It may be noted in passing that on two Trials, Nos. 7 and 8, although the model was notionally thirsty and entered the appropriate location to drink, it did not actually Drink. It failed to emit the Preconsummatory Response. Consequently the Input Pattern was not changed by the environmental routines so as to elicit a Consummatory Response either. An investigation into this failure made it probable that the defect in behaviour was partly due to some changes made in this version of the model quite unconnected with Motivational Behaviour.

It was also partly due to a satiation mechanism which had been built into the model at an early stage. In this early mechanism, when the model had eaten or drank several times (emitted Consummatory Responses) Excitation was withdrawn directly from Neurons mediating the Preconsummatory Response, effectively causing the latter Response to fail. This mechanism has not been mentioned previously, because it was considered of little theoretical interest. A later implementation was intended which conformed to the general motivational concepts developed in later versions. In this scheme, the activity in the Motivational Pools driving Preconsummatory Responses would be reduced when Consummatory Responses occurred (compare Konorski, 1967). However, the simulations had to be abandoned before this development was possible.

Table 14.7 illustrates an experimental run in which the model was required to discriminate the position of the food according to a Cue Pattern. This run differed from earlier runs reported in that one of the Motivational Afferent Fibres was active throughout the run, i.e. it

TABLE 14.7

Demonstration of discrimination of Input Patterns whilst in notional Motivational State

Trial No.	Instant	Input Pattern	Number of Responses in Trial to correct		Actions emitted Correct 1st	Actions emitted Correct 2nd
1	85 - 129	43	2	F,R		X
2	148 - 167	43	1	R	X	
3	185 - 277	39	4	R,R,L,F	-	-
4	295 - 314	39	1	F	X	
5	332 - 369	43	2	F,R		x
6	387 - 420	39	2	R,F		x
7	438 - 470	43	2	F,R		x
8	489 - 507	39	1	F	x	
9	528 - 550	39	1	F	x	
10	570 - 587	39	1	F	X	
11	607 - 640	43	2	F,R		x
12	661 - 677	43	1	R	X	
13	697 - 737	43	2	F,R		x

Key: R = Right, Instrumental Action 4
L = Left, Instrumental Action 2
F = Forward, Instrumental ACtion 3

The Model was notionally hungry, i.e. Drive 2 active, throughout the run. The reinforcement was notionally Food.

Summary: Total number of Responses 22

Total number correct 13

Total number of correct to be expected be chance 7

Run Identifier: UACTOTA 251179

was notionally Hungry. As has been reported, the presence of an Active Motivational Input generally disturbed learning. This disturbance could be reduced by moderating the strength of the Motivational Fibre Spike Frequency, and with such moderation the model could be persuaded to give quite good performance even in a situation where it was required to discriminate Pattern under Motivational pressure.

E. Summary of Computer demonstrations

It will be seen from the runs reported that the extended versions of the simulations demonstrated Instrumental behaviour analogical with that of animals in simple conditioning or learning experiments. In case the reader may have been given the impression from a discussion of the details of performance that the demonstration of learning is subject to doubt, it is stressed here that this is not the case: the demonstration of Instrumental Learning in the model was beyond any reasonable doubt. As will be seen from the summary of performances of many different versions in different notional environmental situations, the model performed in a manner so far removed from chance probability that its behaviour could only be reasonably attributable to the mechansims with which it had been provided. Moreover, the operation of these mechanisms was followed by monitor printouts, so that it can be unequivocally stated that this was so.

It is true, of course, that the model did not reach in all runs the level of performance which would be expected of even the most lowly animal. Its mechanisms often misled it, causing it to make stupid errors. On the other hand, in some runs and in the most recent versions in which programming errors had been corrected, its performance reached a very high level. Its performance must be considered in relation to the primitive character of the simulated model in comparison with the sophistication of any real nervous system, as well as the limitations on length of run in a computer compared to the real time of the animal. The point has been made elsewhere that a higher standard of performance would have carried with it its own implausibility.

Two main kinds of behaviour were demonstrated: (i) responding with an appropriate Instrumental Response to a Discriminative Cue which indiated where the model would be likely to find reinforcement on that Trial; and (ii) the ability to emit discriminatively that Instrumental

Response which would enable it to obtain whichever of two kinds of reinforcement was appropriate to its notional motivational state at that Trial. That is, in metaphorical terms, the model learned where to find food and where to find drink; and to go to the food when hungry and the drink when thirsty.

Reinforcement in this context meant a specific input pattern which elicited from the model a Preconsummatory Response. In view of the presence in the model of a representation of hunger, thirst and in general metabolic or hormonal state, and in view of the simulation of the following Consummatory Response, it is felt justified in claiming that reinforcement does represent to some extent those substances which have been termed reinforcers in animal experiments, i.e. food and drink.

CHAPTER 15

SOME PROBLEMS SOLVED BY RHYTHMIC MECHANISMS

I Introduction

A number of models possess a rhythmic quality which has been suggested to be analogical with the electrical rhythms of the brain. The occurrence of rhythms in my model is a different case. In the basic concepts of my model, i.e. Linked Constellations, there is no implication of rhythmic function. The model was developed at the verbal level to a considerable complexity without the necessity of rhythmic mechanisms appearing. It was only when the model was simulated that intractable complications arose which threatened to block the model from further development. What was surprising was that once the introduction of rhythmic mechanisms had been accepted, most of these problems were immediately solved.

Many of the models which have an intrinsic rhythmic quality may have been partly inspired by the known fact that there are electrical rhythms in the brain, and that there are artefactual devices with similar rhythmic qualities. For example, Pringle's (1951) oscillator model, Hebb's reverbertory circuit model, Beurle's (1956) propagated wave model, Lashley's interference oscillator model, Stanley and Kilmer's timed conjunction model (1975), Landauer's frequency induction model (1964) and Hendrickson's spike interval coding model (1972), all explain the EEG by putting forward a rhythmic model of neural integration.

The neural hologram models also explain the EEG by an essentially rhythmic mechanism. Spinelli's Occam (1970) registers and reproduces waveforms. Landfield (1976) proposes specifically that Hippocampal Theta is a "reference beam" for the hologram, an approach which contrasts with that of Pribram and his group, who favour the idea that a neural hologram would not require a reference beam (Pribram, Nuwer and Baron, 1974).

Although these and many other rhythmic models differ substantially from each other, they all have in common that rhythms are an essential feature of the actual coding of information. In contrast my model proposes a non rhythmic mode of neural integration - the Linked Constellation - and suggests that rhythmic mechanisms serve certain quite definite functions in enabling that model of integration to be developed to a stage which would be impossible without rhythmic mechanisms.

Rhythms appear in my model purely as a form of control. The model suggests that rhythms arise in nervous systems as a natural and highly effective way of controlling what would otherwise be an unstable and confused superposition of different patterns of excitation. The model represents neurons not as formal switching units but as analogue units handling streams of excitation. In the model, as would also seem to occur in the brain, there is at times a superabundance of excitation.

In brief, the control was exercised in the model by providing notional inhibitory gates across the afferent pathways to the arrays of neurons. These gates opened and shut rhythmically to allow only sufficient input of excitation into the array appropriate for its particular current function. Moreover, the rhythmic gates sliced up the input to the integrating tissue into discrete frames, giving periods of freedom from input during which the response of the system can form without being confused with the input pattern or being overwhelmed by it. The rhythmic mechanism also enabled the amount of excitation entering a particular array to be varied and so provided a method of implementing attentional weighting alternative to that used earlier, that of varying the Spike Frequency of Afferent volleys.

As in other aspects of the present project, it is not suggested that the following account describes what actually happens in the brain. To consider the suggested function of the rhythmic mechanisms which appear in the model in relation to actual nervous systems, would be beyond the scope of this thesis. Accordingly the account will be limited to the place of rhythms in the model, with perhaps some suggestions of supporting evidence that the approach adopted is plausible.

II Brief description of rhythmic mechanism in the simulations

The implementation was extremely simple. Each Neuronal Array was provided with an Inhibitory Gate across its set of Afferent Fibres. The Gate to each Array was under separate control. The Gates were notionally situated in the thalamus either as part of the relay or as acting upon the post-synaptic element of such relays. In the model they could be conveniently regarded as acting on the Afferent Fibres between the Control Centre (Thalamus) and the Cortex.

At any Instant the Gate could be in one of two States. When the Gate was Shut no volleys were permitted to enter the Array. (In some versions the Spike Frequency of the Volleys were reduced rather than abolished). When the Gate was Open, the Volleys were unimpeded, and entered the Array just as though the Gate did not exist.

Each Gate had its own characteristic Gate Period or Gate Frequency (to be distinguished from Spike Frequency of a Fibre). In terms of Instants the Gate Period used throughout the demonstrations reported was 14 Instants.

In the simulations only two Gate Conditions were used. The Gate Condition is to be distinguished from the Gate States. In one of these Conditions the Gate was Open for seven successive Instants and then Shut for seven Instants. In the other Condition, the Gate was always Open. These two Conditions were intended to represent two points on a graded range of control conditions of the Gate. The Period of the Gate did not vary. Control of the Gate consisted of variation of the proportion of the Gate Period during which the Gate was Shut.

Two points are to be stressed: (i) Volleys passing when the Gate was Open had the Spike Frequencies calculated for them by the Control Centre as before. When the Gate was Shut, the Spike Frequencies of all Fibres was zero (or in some versions, reduced by some simple formula). (ii) The terms Shut or Open apply to the State of the Gate at a given Instant. The Condition of the Gate, which determined its function, was a temporal condition, e.g. Shut for half of the Period.

For convenience of reference the Open-all-the-time Condition will be referred to as the Uninterrupted Condition, the other as the Interrupted Condition. In further development it would be possible to explore

different proportions of the Period during which the Gate might be Shut, in which case this dichotomous terminology would be inadequate.

The effects of this mechanism and the principles of its control will be related in due course.

III The nature of problems encountered

The development of the model was severely hampered by a number of poorly understood difficulties concerned with the levels of Excitation and Inhibition. Some indication of these difficulties has already been given in context. In a complex system in which many variables interact it is not possible to point with certainty to particular causes of the overall behaviour. What can be definitely stated is that prior to the introduction of rhythmic mechanisms the model performed so badly as to make further progress impracticable. After the introduction of those mechanisms, virtually no problems of control were encountered at all.

In the particular form in which they appeared these problems may be specific to my model. Some of them undoubtedly are. On the other hand it seems probable that similar problems would occur in any system with similar characteristics, that is, a system in which integration is mediated by differences of excitation in distributed sets of neurons.

The problems encountered may be summarised as the occurrence of variation resulting from the operation of the system which was larger than the variation representing the response of the system. To some extent these problems were caused by efforts to retain and demonstrate certain ideas which were felt to be important implications of the basic neurological theory. One of these was the idea that attentional weighting could be acheived by varying the amount of Excitation entering a particular portion of the neural template system. Another idea was that in nervous systems Excitation can build up sequentially so that even the smallest implications of the processing can acquire decisive power.

IV Arousal problems

A. Feed-back characteristics of original arousal system

In the initial simulations the Spike Frequency in each of the Afferent Fibres was the sum of the Diffuse Volley Value and the Discriminative Volley Value. The former applied to all Fibres regardless of their signal status. The latter distinguished Fibres designated as Active from those designated as Inactive. Arousal was mediated by adjusting the values of these variables according to the number of Spikes appearing across the Efferent Fibres. A slight variation of output in an Efferent Fibre caused an increase of arousal, which in turn caused a magnification of the next variation, and so on, so that an inflationary loop was entered. Fortunately the loop rarely overwhelmed the system within the limited duration of a run, so that results were still obtainable.

The inconvenience of these effects does not necessarily mean that they do not occur in nervous systems. It seems probable that they do (Chapter 14,). If they do occur then, as in the model, they emphasise the need for adequate control so that the system can recover from them.

As reported, because of the difficulties encountered, the arousal system had to be removed. It was, however, able to be reintroduced in a revised form when the Inhibitory Gates had been incorporated. In this version of Arousal, Spike Frequency in Afferent Fibres was no longer a function of the output of Efferent Fibres, so the inflationary feed-back loop was broken.

As may be remembered, in the new system the whole concept of arousal was revised. Two quite different forms of varying the level of activity were distinguished.

One way of increasing activity in the system was by sending Active Afferent Volleys to one or other of the Motivational Neuronal Pools. This selectively energised the information stored in specific Neural records. The other means of varying activity was to vary, by means of the adjustment of the Inhibitory Gates, the amount of Excitation able to reach the different Arrays. The Inhibitory Gates enabled the weighting given to selective parts or to the whole of the neural template to be reduced, or to be at maximum level. In this revised

concept of an attentional system, weighting could be varied to any aspect of the neural template, including the motivational aspect.

B. Gates as attentional mechanisms

When a Gate was Open, sufficient Excitation entered through the Afferent Fibres to cause the Excitation within the Array to rise. If one of the Arrays was to be given attentional weighting, i.e. arousal, its Gate would be put into the Uninterrupted Condition. Accordingly Excitation continued to rise in that Array, much as it had done in earlier versions, except perhaps more so. At a certain point a feed-back control system (described below) checked this rise effectively. This feed-back system was able to operate satisfactorily in the new version for two reasons. In the first place, Diffuse Volley Value and Discriminative Volley Value did not vary as a result of feed-back as it had in the earlier arousal system, and the major source of instability was thereby removed. Secondly, the Array could always be brought under control by putting the Gate into the Interrupted Condition.

If one of the Arrays was to receive reduced weighting (low attention) the Inhibitory Gate to that Array was put into the Interrupted Condition. In this Condition it was Open for seven Instants and then Shut for seven Instants. Whilst the Gate was Open, Excitation within the Array rose. When it was Shut, the Excitation began to fall. The Excitation within the Array fluctuated regularly when in this Condition. Its average level was, of course, half that of an Array to which the Gate was in an Uninterrupted Condition. At the time that this mechanism was introduced it was thought that the effects would only be as selective weighting and as an aid to control. Consequently it was thought that an Array, the Gate of which was in an Interrupted Condition, could only serve a useful purpose during the Open phase of the Gate. It later appeared that the other half of the Gate cycle, when it was Shut, was of equal importance.

C. Interpretation of EEG in the model

A few words may be necessary to relate the effects of the rhythmic mechanisms to the electrical rhythms seen in the brain.

It is not appropriate to derive a direct correlate of EEG from the simulations. A waveform was printed out which might seem to represent

in an approximate form the response of a gross electrode recording. To try to make this resemble the EEG in principle, the variable plotted included a summation over all the Integrating Neurons within a single Array of both the Excitation Counts and the Spikes emitted. These curves were superficially quite like EEG records. For example, when the Gate to the Array was in the Uninterrupted Condition the waveform showed rapid alternations (i.e. from Instant to Instant) of low amplitude. When the Gate was in the Interrupted Condition the waveform had high amplitude variations which followed the period of the Gate, i.e. 14 Instants. Nevertheless, it is recognised that there are many considerations which make an identification between any particular aspect of the model and the EEG of doubtful validity. For example, it is known that in addition to the envelope of emitted spikes, the EEG reflects both intra-cellular and extracellular potentials at various depths in the tissue. Moreover, it is exceedingly easy to obtain superficially similar waveforms to the EEG with many rhythmic mechanisms.

Accordingly the correspondence between the mechanisms of the model and possible causes of electrical rhythms in the brain must remain at this state purely a matter of the general plausibility of the model. As a general interpretation it seems reasonable that when Excitation within an Array is allowed to rise, the interneuronal exchanges between Integrating Neurons, mediating the Linked Constellation concept, would also increase. Assuming that Neurons differ in the rate at which their intracellular potentials grow and in the consequent latency of the Spike, the higher the level of interneurnal activity the greater would be the desynchronisation of the firing of Neurons. In the initial version of the arousal system, there was a considerable implausibility in this assumption. Since the increased input to the tissue was in the form of Volleys carrying higher frequencies of Spikes, the higher activity within the system was accompanied by a more rapid arrival of afferent excitation. Accordingly it was difficult to see how desynchronisation could arise from the enhanced Volleys. Moreover, it is known that a sudden or intense sensory stimulus or an artefactual stimulus to the afferent fibres produces a peak (evoked response) rather than a desynchronisation.

This latter implausibility disappears when Attention is mediated by the rhythmic gate mechanism. The increased Excitation in the integrating

tissue arises not because of enhanced afferent volleys but because they continue uninterrupted. Consequently as the level of Excitation rises in the Array, the balance of interneuronal exchange to the direct effects of afferent streams changes.

It seems therefore reasonable to equate an Array, the Gate to which has been for some time in the Interrupted Condition, as in a "synchronised state" in the following sense. During the Open phase an Input Pattern enters and during that phase the Excitation rises first in the recipient Neuronal Pools, the interneuronal exchanges being relatively small because of the recent increase in Excitation. When the Gate moves into the Shut phase, the Excitation falls more rapidly from the Unlinked Neurons and Neurons belonging to non-matching Constellations, because matching Constellations are shuttling their Excitation back and forth. Nevertheless, both interneuronal and Excitation within Neuronal Pools falls rapidly. The effect therefore is that of a synchronisation of the discharges within the Neuronal Pools receiving direct afferent Excitation.

Although we are not concerned with applying the model at this stage to an explanation of rhythms as they actually appear in nervous systems, it is desirable to avoid a possible misunderstanding.

When an inhibitory gate to a notional projection area is in the interrupted condition, it does not necessarily mean that during the Shut phase of the gate, no neural processing may proceed in that area. The activity would, it is suggested, fall because the levels of Excitation falls within the area. It is also possible that reactive inhibition, and in the case of profound slow rhythms, externally imposed inhibitory rhythms (see Jacobs, Mosko & Trulson, 1977 and Section VIII, B), play a part in subduing interneuronal activity during these phases. However, these inhibitory aspects are separate from the gating mechanism on the afferent fibres. It is quite conceivable that the afferent input could be blocked by the shut gate, but that activity continues in the area to which that gate applied because the integrating tissue is sufficiently activated by other sources of excitation. In the model these other sources would be the Motivational Pools. That sources of excitation other than sensory signals support fast activity has been suggested by Jouvet (1974) and Dement and Mitler (1974) who refer to "phasic generators".

In the simulations, during an episode when an Array is in the Interrupted Condition, it was not necessary to assume that neuronal interchange occurred during the nadir of the shut phase. The downward fall of the cycle was, as will be seen, used to derive the response of the system independent of the effects of the afferent input. This will be discussed further in a subsequent section.

V The control system: feed-back loops

The problem of inflationary feed-back loops did not result solely from the feed-back control of arousal. After removal of arousal, problems were still encountered. It later appeared that these difficulties resulted from a fundamental problem in the system which may have some theoretical interest. Before discussing that aspect of neural processing, a brief account will be given of the control system developed in the model and some of the difficulties encountered with it.

Excitation entered the system at each Instant through the Afferent Fibres. At each Instant both Inhibition and Excitation were destroyed through the emission of Spikes. They were also notionally lost as dissipation through a leaky system. Unless the total of Excitation received at each Instant was destroyed at each Instant then the amount of Excitation within the Integrating Neurons rose continually from Instant to Instant. On the other hand if at the end of each Instant all Excitation was destroyed, then the sequential and conditional probability aspects of the processing were also terminated at each Instant. The latter did not represent the kind of system I had in mind. An Instant was merely a convenient fiction for analysing the system in temporal terms. Accordingly the system was purposely designed so that some Excitation would remain from Instant to Instant to be carried over.

The resulting build up of Excitation was satisfactorily solved in early versions by periodically clearing all Neurons of Excitation and Inhibition at arbitrarily chosen times, e.g. immediately after a Preconsummatory or Consummatory Response. This worked well but was considered a mere device of convenience. As it later appeared that is exactly what the rhythmically operating Gates do: by blocking input periodically, they allow both Inhibition and Excitation to return to optimum levels.

One of the parameters by which the system was adjusted so that some Excitation was carried over from Instant to Instant was the value of Rate of Discharge. It will be recalled that this was a fractional multiplier the effect of which was that at each Instant only a proportion of the excess of Excitation over Inhibition was emitted as Spikes. Obviously the greater this parameter, the more Excitation was Discharged at each Instant and the less the problem from the build up of Excitation. As reported,this parameter was set at 0.3. This is now seen to be an unnecessarily low rate, and undoubtedly the use of this rate, which was quite purposely set low to increase the build-up, added to difficulties. I now think this was fortunate as it drew attention to characteristics of the system, which I should have otherwise probably ignored.

The problem of replacing the occasional clearing strategy of the early versions by a more acceptible physiological mechanism to control the build up was, as it now seems, erroneously seen as a problem of providing an adequate feed-back control of levels of Excitation and Inhibition. A reasonable control system was to have a basically leaky system which was replenished from Instant to Instant: the resulting level of Excitation could then be adjusted by the quantity of Excitation introduced at each Instant. The most obvious variable for this purpose was the Diffuse Volley Value. By feed-back adjustment from the average level of Excitation across the Integrating Neurons in an Array, a value of Diffuse Volley Value could be calculated which returned the system to its optimum level. This use of Diffuse Volley Value conflicted with its arousal functions, but since these were by that time considered as unsatisfactory in any case, Diffuse Volley Value became the control variable, rather than an arousal variable.

At this stage Diffuse Volley Value was interpreted as representing all those sources of Excitation within an Array which did not actually express the input pattern to the array. These would include intra-neuronal metabolic processes. Consequently there was no inherent implausibility in its use in the system. It seems probable that neural tissue possesses homeostatic mechanisms of this kind which serve the double purpose of maintaining the tissues at optimum conditions and protecting it from extremes induced in special circumstances, e.g. high Spike Frequencies. What was implausible was the particular prominence this control factor took in attempts to control the wild fluctuations of the system.

No difficulty was eventually encountered in arranging a simple feedback variation of the value of Diffuse Volley Value which corrected minor variations of Excitation level and which returned the system at each Instant to an optimum level. The difficulty arose from two sources. One of these is dealt with in the next Section. The one with which we are now concerned was the desirability of carrying Excitation forward from Instant to Instant in order to increase the strength of the Anticipatory Responses of the model.

In addition to the general increases of Excitation which resulted, there were considerable build-ups of Excitation in particular integrating Neurons. In the initial simulations Input Patterns changed from Instant to Instant, so these differential build-ups tended to average out. In the later versions a single Pattern, say the Cue Pattern, might be in the Input for many Instants. Of course, this was an unrealistic feature of the simulation which would have been obviated by replacing the single Cue Pattern by a sequence, as used in earlier versions.

Within limits this build up of Excitation was a necessary and desirable prediction of the system. What made it an embarrassment was its excessive magnitude, and its continued presence in the system after it had served its purpose.

VI Distinguishing the response of the system from the direct effects of the input

It is now possible to consider a problem the implications of which may have some theoretical interest. It was obviated by the rhythmic Gate mechanism. Essentially the problem concerned the difficulty of discerning the results of neural processing within a location of tissue which is subject to continual afferent input. In the simulations the problem arose mainly in the Motor Arrays when it was required to detect the relevant Motor Response Indications. However, its theoretical implications would appear to apply generally to the system.

The model suggests that one function of neural rhythms is to separate temporally within an Array the direct effects of the Afferent Input to that Array from the effects of interneuronal exchanges across the whole template. It does this in two ways. (i) It reduces the level

of Excitation entering a particular Array. Consequently the Excitation induced in the Neurons of that Array, by activity in some part of the most highly reactivated Linked Constellation, stands proud of the general levels of Excitation engendered within the Array itself. (ii) The rhythmic Gate can, when required, slice up the Input to that Array into discrete frames, giving the tissue in that Array a periodic freedom from the direct effects of Input to enable the response of the most highly activated Linked Constellation to be seen.

For convenience of reference one may refer to a primary and a secondary response of the tissue. The meaning of the signal input is derived in the model from the differences of Excitation induced in sub-sets of Integrating Neurons (Linked Constellations) by secondary exchanges. It was found that this secondary response could not be effectively detected against the effects of the direct afferent Excitation.

It will be noted that where a Constellation is present in the system which matches exactly the current Input in all locations, the primary and secondary responses of the system will match in one sense: the sub-set of Neuronal Pools directly Excited by the Input will be the same sub-set which contains the Neurons belonging to the most highly reactivated Constellation. In the case of the most closely matching Constellation differing from the Input Pattern, the two responses of the system will differ, i.e. the sub-set containing the Neurons of the most highly reactivated Constellation will not match the sub-set of Neuronal Pools directly activated by the Input.

The interesting case is where a mismatch occurs in a particular location of the extended neural template, and the model needs to detect the implications of the whole input as it relates to that location. For example, in the simulations the release of a Motor Response required the detection of a Motor Response Indication in the Motor Array. This Constellation response, induced by the Sensory Input could not be detected because of the direct effects of the Afferent Volleys sent up to the Motor Arrays when Responses were released.

A. Importance of detecting match and mismatch

Before considering the difficulties in the way of detecting the Constellation response in an Array, a short digression may be justified to establish the necessity of doing so in a more general theoretical sense.

The inability of the system to distinguish the secondary from the primary response will not matter if the aim of the system is simply that of establishing Stimulus-Response Connections; at least it will not matter on the Sensory side of the apparatus. The near match, even if not complete match, between Constellations present in the system and the current Input will ensure the rapid reactivation of the Constellations. In some circumstances it is important that the system should respond to the input as if it were the same stimulus as that on previous occasions and ignore minor discrepancies. This would assist a clear and prompt response. It would represent one mode of operation of the model, which may be termed the Stimulus-Response mode.

Ignoring for the moment the possible difficulties of detecting the Motor Response, one can envisage a well learned sequence of sensori-motor behaviour running off without trouble in a mode in which, on the sensory side at least, no attempt is made to distinguish between the primary and secondary response of the system. An equivalence is assumed between the most highly reactivated Constellation and the current Input. Each successive input would reactivate a relevant set of Linked Constellations. The resulting Motor Action interacts with the environment to modify input at the next Instant. The modified Input reactivates the next step of the sequence of Linked Constellations laid down on an earlier occasion. In this mode the system exemplifies the Stimulus-Response view of behaviour as a sensori-motor chain.

There are however circumstances where such a model of operation fails to take advantage of the information stored in the system. Such a model cannot "take care" and avoid some of the disasters which might result from faulty equivalences of slightly differing situations.

The interesting aspect of the model is where some part of the information represented in a Linked Constellation already present in the system is different from or absent from the current input, or conversely where information in the current input is different from that stored in Constellations. In the Stimulus Response mode, the most closely matching Constellation is simply identified with the currently arriving pattern of Excitation. In a purely discriminative model, as defined by the Gibsons (1955) (see Chapter 2) such faulty equivalences will be challenged only by the eventual environmental failure of the responses

released by the equivalenced neural records: the mismatch can only be detected by the subsequent (and possibly fatal) environmental consequences of an unjustified identification of stimuli. An alternative consequence of a faulty equivalence is that a valuable opportunity for the animal may be lost. A dog which goes out of a room to find food when a buzzer of the wrong frequency is sounded will not be in the room to hear the correct cue when it occurs.

Two rather different cases of mismatch may be considered. The first occurs when the system is already on its guard and the suspected input is to be examined. This may be because of the failure of the well known sensori-motor chain to be activated: it may be because some significant aspect of the environment has already differed from what is familiar, e.g. a new keeper arrives with food. In such circumstances it will be desirable that any sensory stimulation which differs even slightly from what has been previously encountered should be carefully examined for mismatch. When a mismatch is detected, it is desirable to proceed very cautiously with motor action, so as to be able to withdraw or proceed to the next step, and to label each Input Pattern as safe or otherwise. To detect the mismatch it is necessary in my model to discern the response of the reactivated Constellations in comparison with the direct effects of the Input Pattern somewhere in the system.

That the sophisticated nervous system is able to operate in this way, at least in critical situations, may be seen in the experiments of Sokolov (1960), Anokhin (1961) and Walter (1969), and their schools. These experiments show that a novel stimulus produces an orienting reaction (having distinct EEG parameters). The orienting reaction becomes habituated after familiarity with the stimulus. The reaction appears again if the stimulus is changed, even in a minor way.

A second case in which it would appear to be valuable to the animal to be able to detect the secondary response of the system is where some current stimulation signals that something else is likely to occur. A familiar example is the classical conditioning situation. Say that a buzzer signals food. There is no food present when the buzzer is sounded. Nevertheless the animal looks attentively at a place where food is customarily placed. According to the model, the <u>meaning</u> of the acoustic stimulus is seen in the reactivation of the matching Constellations. These Constellations, or those closely associated with

them by temporal sequence, will contain representations of the visual appearance of food, registered upon former occasions. It seems reasonable that the animal has a visual expectation of food. The representations of food laid down on former occasions will be in the visual location of the neural template (or some location receiving a conjunction of visual and accoustic inputs). The reactivation of the portion of the Constellation falling within the visual location, must be detected against the pattern induced in the same location by the current input, an input which signals that no food is present.

Such an interpretation is not of course essential to an explanation of Conditioned Responding. The interpretation may be distinguished from the substitution theory in which the CR, say, salivation, is simply attached to the CS, say, a buzz in substitution for the original stimulus capable of eliciting salivation, say, food.

Having argued for the necessity of distinguishing the response of the reactivated Constellations, we may now consider how the rhythmic Gate mechanism enabled it to be detected.

VII Difficulties in detecting Response of Constellations

A. Magnitude of Response of Constellations

The response of Constellations is induced in the system as the result of interneuronal exchanges between Neurons of a Linked Constellation. In the conservative Excitation system used in the main simulations, the differences of Excitation between Constellations were small as compared with the effects of Afferent Volleys on the Integrating Neurons situated in the Pools whose Fibres were Active.

The magnitude of the response of Constellations could be increased by departing from the conservation concept and assuming that Excitation is amplified by interneuronal exchange. In the separate simulation which investigated the mediation of Linked Constellations by orthodox synapses it was found that the system could be made amplificatory, conservative or attenuative simply by adjusting the threshold at which Integrating Neurons underwent a notional action potential in relation to the average number of synapses which Integrating Neurons made upon each other. In the amplificatory system, the problem was the converse of that experienced

in the main simulations. Instead of Excitation building up in the Neuronal Pools receiving direct afferent input, it built up in the Linked Constellations. Once a Constellation had exceeded a critical state it continued to rise in Excitation until it exploded at an impossible level. At subsequent Instants further Constellations reached the critical level and the whole model ran out of control. This control could only be regained by introducing a rhythmic Inhibitory Gate, and in addition a profound Inhibition, which by cancellation of Excitation, brought the Constellation back below the critical level.

To return to the simpler case of the conservation system in the main simulations, the response of the Linked Constellation was small in comparison with that of the Neuronal Pools. This seemed to be the required state of the apparatus for most purposes. The direct effect of inputs must be relatively large at least at the times when it is necessary to ensure accurate registration of sensory events.

B. Overlapping of directly excited and Constellation responses

In locations to which afferent excitation is continually arriving, the reactivated pattern is overlaid by the signal pattern. With rhythmically spaced Input Volleys, the two patterns are separated temporally. During the Open phase of the Gate, Excitation enters the selected Neuronal Pools and the afferent Excitation to that Array becomes dominant. At the peak of the response of the directly driven Neuronal Pools registration may take place. Interneuronal exchanges take place, but they cannot be distinguished from the direct affects of the afferent volleys. When the Gate Shuts, the Input Volleys are cut off. The Excitation in the system begins to fall. Where a Linked Constellation is at a high level because of Excitation reaching it across the whole Neural template, it will now stand proud.

The foregoing description is not the only way in which the rhythmic system might operate in the model. It was, however, the way in which they were used in the simulations, and which were accordingly shown to be effective. To consider further possibilities of the effects of the rhythmic Gates would be to go beyond the aims of the present discussion. However, some possibilities are mentioned in a later section.

It may be important to stress that the above account does not need to assume that interneuronal activity can take place when the Neurons

concerned are inhibited. Although that might be so (see Section IV, D of Chapter 5), the above account does not require it. This would remain true even if there is a profound externally induced Inhibition during slow waves, as was conveniently introduced to control the amplificatory aspects of the separate simulations utilising notional orthodox synapses. The Constellation response could be detected during the downward phase rather than at the nadir of activity within the Array.

It would seem probable that if rhythmic Gates perform the roles attributed to them in the model, there would be room for considerable variation of their mode of activity, both from location to location and also from moment to moment. For example, it may be that a profound externally imposed inhibition only occurs in cortical tissue under special circumstances (Jacobs, Mosko and Trulson, 1977).

Verzeano (1970) interprets his results as follows: " . . . different neurons behave in different ways in relation to the (gross) slow waves. . . . the largest majority of neurons, which produce spikes of low amplitudes, discharge in consistent relation to the negative slopes of the gross response. A small minority of neurons, which produce spikes of high amplitude, discharge in a more variable fashion: in some instances in relation to the negative, in other instances in relation to the positive slopes of the gross response".

C. Inhibition

Inhibition problems in the model could not be satisfactorily solved until the rhythmic Gates were provided.

Within an Array the average level of Inhibition Count rose so as to silence all but the most highly Excited Integrated Neurons. A feed-back criterion had to be found for this system. This raised the question of the distinction between the response of the system, that is, the permitted discharge, and the "noise" which was to be rejected.

The external intelligence of the programmer was obviously not to be utilised in recognising when a discharge "pattern" had been optimised. The model offered its own criterion of pattern. A principle of the model, already discussed, is that all valid signal patterns involve approximately the same proportion of Active to Inactive Fibres (both in Motor and Sensory Arrays). Accordingly, the system aimed at that level

of Inhibition at which the sub-set of Neuronal Pools containing Discharging Neurons conformed to a balanced pattern.

Inhibition was easily controlled to give this effect. The Discharge from the Pools representing the balanced pattern could be maximised by seeking the lowest level of Inhibition at which the required number of Neuronal Pools showed no Discharge.

During admission of signal volleys, the system naturally maximised the response of the Pools receiving direct Excitation and Inhibited the others, i.e. it isolated the directly excited response of the system. In the Uninterrupted Condition of the Gate (and, of course, in earlier versions of the model without Gates) this was the only response which could be isolated by the Inhibition system. In the Interrupted Condition of the Gate, during the Open period of the Gate the system homed in to the directly excited response, and during the early part of the Shut Condition it was able to home in to the Constellation response.

The system was able to distinguish the most highly Excited "relevant" reactivated Constellation from less Excited "irrelevant" Constellations because at the appropriate level of Inhibition, the most highly Excited could be isolated by the requirement that the sub-set of Pools in Discharge conformed to a balanced Pattern. If the Inhibiton were lowered further, the balanced pattern was broken up by the intrusion of Discharge in further Neuronal Pools, due to the inclusion of less highly Excited Constellations.

The precise setting of the system could be used to determine the degree of generalisation permitted. If the Criterion used for the feed-back control of Inhibition were satisfied when a greater number of Pools were allowed to contain Discharging Neurons, say, two or three more than those required for a balanced pattern, then the Constellations permitted to contribute to the Constellation response might include Linked Constellations which did not completely match the most highly reactivated Constellation. If the feed-back loop homed in to precisely the condition that just half the Neuronal Pools might contain Discharging Neurons, then only the Constellations reactivated to the highest levels which exactly matched each other, could contribute.

The important feature of the temporal separation of directly excited and Constellation response, by the action of the Gate in the Interrupted Condition, was that the Inhibition system could then optimise each of

the two responses of the system separately.

Before the foregoing features of the system had been understood, much time had been spent in searching for a criterion which would enable both the directly excited and Constellation response to be seen together. Even if it had been found, the difficulties associated with the simultaneous appearance of both the signal and the secondarily-reactivated pattern would have remained. Accordingly the Inhibition problem converged upon the other problems which were solved by the rhythmic mechanism.

D. Detection of Motor Response Indications

It might perhaps be thought possible to avoid the foregoing problems by looking at some location where Afferent Volleys are not present at that time. For example, as may be seen in Figure 5.7, in principle the meaning of the sensory input can be discovered by looking at the Discharge of the Efferent Fibres, i.e. the Motor Location. It was, however, pointed out that the freedom from Afferent Input of the Motor Neuronal Pools was a simplification used for exposition only. These pools were provided with Afferent Fibres and received Volleys which registered emitted Responses. Such Volleys occurred whenever a Response was initiated either by a Cortical Motor Response Indication or by a Sub-cortical initiation. There was also Diffuse Volley Value, the control variable which entered all Fibres at each Instant.

Although the present discussion concerns rhythmic mechanisms in the model rather than the brain, in reality motor cortex is subject to considerable direct excitation, some of it sensory, but much of it motor, e.g. from the ventral lateral nucleus of the thalamus. As already mentioned, Diffuse Volley Value is rather specific to my model and may be difficult to relate to specific aspects of neurophysiology. Nevertheless there are profuse sources of excitation to motor locations. It seems implausible to consider that any location of the cortex would be sufficiently free from afferent volleys for the purposes suggested by the model, unless special arrangements were made for that purpose, e.g. an inhibitory gate (see below).

However that may be, in the model it was found impossible to obtain the Motor Response Indication (i.e. the Reactivated Constellation response of the Motor Array) with accuracy. The main source of trouble was the effect of Volleys in the Motor Efferent Fibres sent from the Control

Centre to register Motor emissions. Although these did not occur all the time, the effect was that the Motor Response Indication was likely to reflect the Excitation left behind by a recently emitted Response registration, rather than reflecting the Constellation response to the current Input.

This difficulty was immediately cured when the rhythmic Gate to the Motor Array was put into the Interrupted Condition. The Motor Response could be taken at just the point when, after the Gate had Shut, the average Excitation in the Array had fallen to a level which maximised the response of Reactivated Constellation. The Motor Response was then found to be almost without exception a reliable product of the Sensori-Motor Constellation laid down at an earlier occasion. Even when the model made behavioural errors, which it frequently did, these were rarely found to be caused by an unreliable Motor Response Indication, but rather to the dominance of a "wrong" Indication.

VIII Attentional aspects

A. Economy of storage

In the initial simulations it was demonstrated that the amount of information stored at each Instant was adjusted according to the significance of events as assessed by the model. Linked Constellations which registered the occurrence of Innately Recognised Patterns (IRPs) contained a greater number of Integrating Neurons than those formed at other Instants. In the "conditioned arousal" versions, this differential storage could be applied to the registration of Patterns of the Cue Sequence, although these latter patterns were indifferent patterns before being repeatedly presented as a Cue.

As reported, the arousal system based upon variation of Spike Frequency had to be withdrawn for many reasons. In relation to economy of storage it had several inadequacies. These related chiefly to what happened when the arousal level was low. The reduced Spike frequencies produced smaller differences of Excitation between Active and Inactive Pools. Accordingly the signal to noise ratio was severely reduced and the number of errors increased. Secondly, the level of Excitation introduced being low at each Instant, only a few Neurons were linked into

each of the Constellations formed. This number was rapidly reduced by decay. The overall effect was that, if a substantial economy was made in information storage, no adequate record was maintained of the period during which the arousal level was low.

It would seem that in the initial arousal system, the only effective alternative to taking a full record was to take none at all. Theoretically this might be a considerable disadvantage if it were later found that events to which "no attention had been paid" were in fact important. This was one of the difficulties in implementing the "conditioned arousal" scheme. An indifferent pattern in a Cue Sequence could never give rise to "conditioned arousal" unless it had been adequately registered in the first place.

In the new system these drawbacks were largely remedied. To acheive Attention to an Array, the relevant Gate was put into the Uninterrupted Condition. As described, Excitation rose within the Array. In this Condition the new system behaved very much like the original system in the number of Neurons Linked, although of course without the control problems.

A difference regarding registration was seen when the Array was put into the Interrupted Condition (i.e. low arousal or Inattention). In this Condition, at the peak of the Open phase of the Gate cycle, the linked Constellations formed had approximately the same number of Neurons as did those formed when the Gate was in the Uninterrupted Condition. For the remainder of the Gate period, the Constellations consisted of relatively few or no Neurons, which decayed rapidly. Unlike the record of events taken when the original arousal system was at a low level, the Constellation record now consisted of a few populous Constellations separated temporally by gaps in which there was virtually no record. Whereas the original system had produced a continuous but inadequate record, the new system produced a periodic but adequate sample of the Input events.

B. Arousal and attention: monitoring sensory input

In the simulations as reported, the main purpose of the rhythmic mechanisms was the maintenance of control over the system and the derivation of the Motor Response Indication in spite of the effects of Afferent Excitation on the Motor Array. In the theoretical model there are some further implications which may be mentioned briefly.

It has been pointed out that the Reactivated Constellation response of the system could be derived in the absence of afferent input. Accordingly it should be possible to derive the sensory implications of signals to one cortical location, say, auditory, by blocking the input to another location, say, visual, as in the example mentioned earlier of a buzzer signalling food.

Assuming that the general proposals of the theoretical model are accepted, there is still no reason to consider that the required blocking of Afferent Input need take the form of a rhythmic inhibitory gate. The reduction of signals by habituation would achieve the required effect.

There would however appear to be a considerable advantage in interposing a rhythmic inhibitory gate between the source of signals and the central integrating tissue. Such a Gate might form an intermediate step in excluding unwanted signals from the cortex, with this difference, that the signals excluded by the Gate might be blocked by reason of their meaning rather than their sensory quality. This form of control might be especially useful in excluding those signals which do not in themselves lack variety or information content, and which would not easily be habituated at the peripheral stations, even when the animal had learned that they had little consequence.

The advantage of the exclusion by the Gate in the Interrupted Condition would be (a) an adequate although greatly economised record could be taken of them (b) when habituated at the periphery the signals passing through the Open phase of the Gate would have no effect, so nothing would be lost by interposing the Gate (c) the signals reaching the integrating tissue would be able to interact to some extent with the information stored there. At each cycle of the Gate some Excitation enters to be able to confirm the reactivation of Constellations indicating "safe" or "of no interest". In the event of the afferent signals changing character or becoming dishabituated, other Constellations become activated, and the Condition of the Gate could be changed.

In brief, rhythmic inhibitory Gates might be used to monitor input in locations to which Attention had been reduced, and the effects of monitoring used to control, the Condition of the Gate itself.

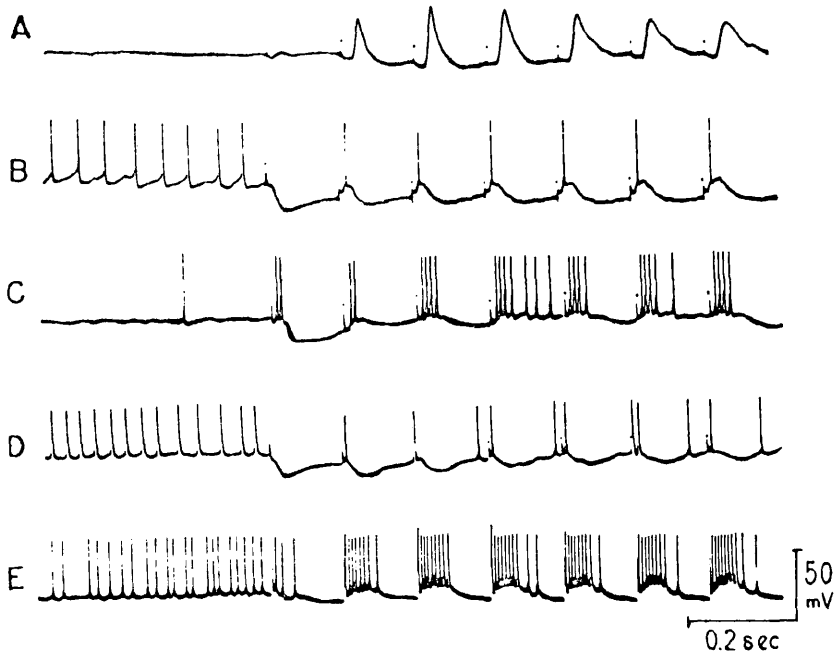

Figure 15.1 (Reproduced from Purpura and Shofer, 1963) Patterns of intracellularly recorded activities of thalamic neurons during cortical recruiting responses evoked by 7/sec. medial thalamic stimulation. A. surface negative recruiting response from Motor Cortex. B. Neuron in ventral anterior region of thalamus. C. Relatively quiescent ventrolateral neuron develops discharge. D. similar to B. E. Neuron in intralaminar region. Stimulus interrupts a spontaneous discharge. Successive stimuli evoked IPSP's and evoked EPSP's.

IX Some neurological considerations

As stated, the rhythmic mechanisms have been presented as part of the model, not as a theory of electrical rhythms in nervous systems. Nevertheless, it is reasonable for the reader to expect some mention of the neurological data relating to the mechanisms which have been incorporated into the model.

The essential mechanism used in the model to implement the various matters which have been discussed is the Inhibitory Gate. In the model this consists of notional neural mechanisms capable of blocking the transmission of afferent volleys from the Control Centre (Thalamus) to the Cortex.

In brief, there is ample neurological evidence that inhibitory arrangements, of just the kind required for the Inhibitory Gates of the model, exist in the mammalian nervous system. The subject, which has a vast literature, will not be fully reviewed, but a resume given.

Dempsey and Morison (1942), Bremer and Adrian carried out pioneer work which led to the recognition that certain rhythms originate in the thalamus. A number of neuro-anatomical hypotheses, backed by the research of certain well-known groups, involve inhibitory gating of afferent pathways in the thalamus. Three will be described.

(i) Recurrent axon collaterals, acting through Renshaw-like interneurons inhibit discharges flowing through specific thalamic nuclei (Andersen and Eccles, 1962; Andersen and Andersson, 1968; Andersen, 1974). Since rhythmic activity persists in the thalamus after complete decortication, such activity cannot be dependent completely upon cortical feed-back (Morison and Basset, 1945). The pacemaker effect proposed by Dempson and Morison (1942) is explained by Andersen as an intrinsic rhythmic excitability of neurons "responding in inhibitory rebound" (Andersen, 1974)

(ii) Purpura and his group demonstrated that stimulation of medial thalamic nuclei produce frequency-specific EPSP-IPSP sequences in neurons of both specific and non-specific thalamic nuclei. He favours a synaptic theory implicating "complex organisations of excitatory and inhibitory interneurons" (Purpura, 1970). A reciprocal relationship between specific and non specific nuclei was reported (Desiraju, Santini and Purpura). Purpura (1974) reviews the work of the group, disclosing rhythmic inhibition of specific, non-specific, motor and sensory nuclei,

both spontaneously and under medial thalamic artefactual stimulation. An example of the inhibitory gating demonstrated by Purpura and his associates is illustrated in Figure 15.1.

(iii) In an anatomical study, Scheibel and Scheibel suggest that the reticularis nucleus acts as a frequency specific gate. This work would appear to follow the early discoveries of Schlag-Chaillet, (1963) who observed that when the mesencephalic tegmentum was destroyed, i.e. caudal to the thalamus, cortical desynchronisation could no longer be obtained by high frequency stimulation of the medial thalamus. In Purpura's experiments, high frequency stimulation of the medial thalamus produced desynchronisation at the cortex, low frequency stimulation, synchronisation and thalamic inhibitory gating.

Scheibel and Scheibel describe a major projection of the thalamic non-specific system rostrally through the reticularis nucleus, and a caudal projection to the tegmentum. They summarise: "Low frequency stimulation of the medial thalamus is immediately effective in synchronising activity over huge areas of the thalamus and cortex because this roughly matches the frequency ranges established by the reticularis gate".

More recent reviews of the subject are in Landfield (1976), Brazier (1977) and Verzeano (1977). Essentially, however, the concepts of thalamic inhibitory gating in the field, as relevant to my model, remain as indicated above.

X Application of Inhibitory Gates in the simulations

To assist in clarifying the practical use of the Gates in the simulations, the control of the Gates at various points in the experimental runs reported in Chapter 14 will be given. Further, rather more informal, explorations of the possible implications of the use of inhibitory gates were carried out, but as they fall beyond the scope of this project, they will not be pursued here.

At the commencement of a Trial in a typical run, the model would be seeking an appropriate Instrumental Response. This "seeking" mode was the basic one in the simulations, and the model was set in this way unless some other indication occurred to indicate a resetting.

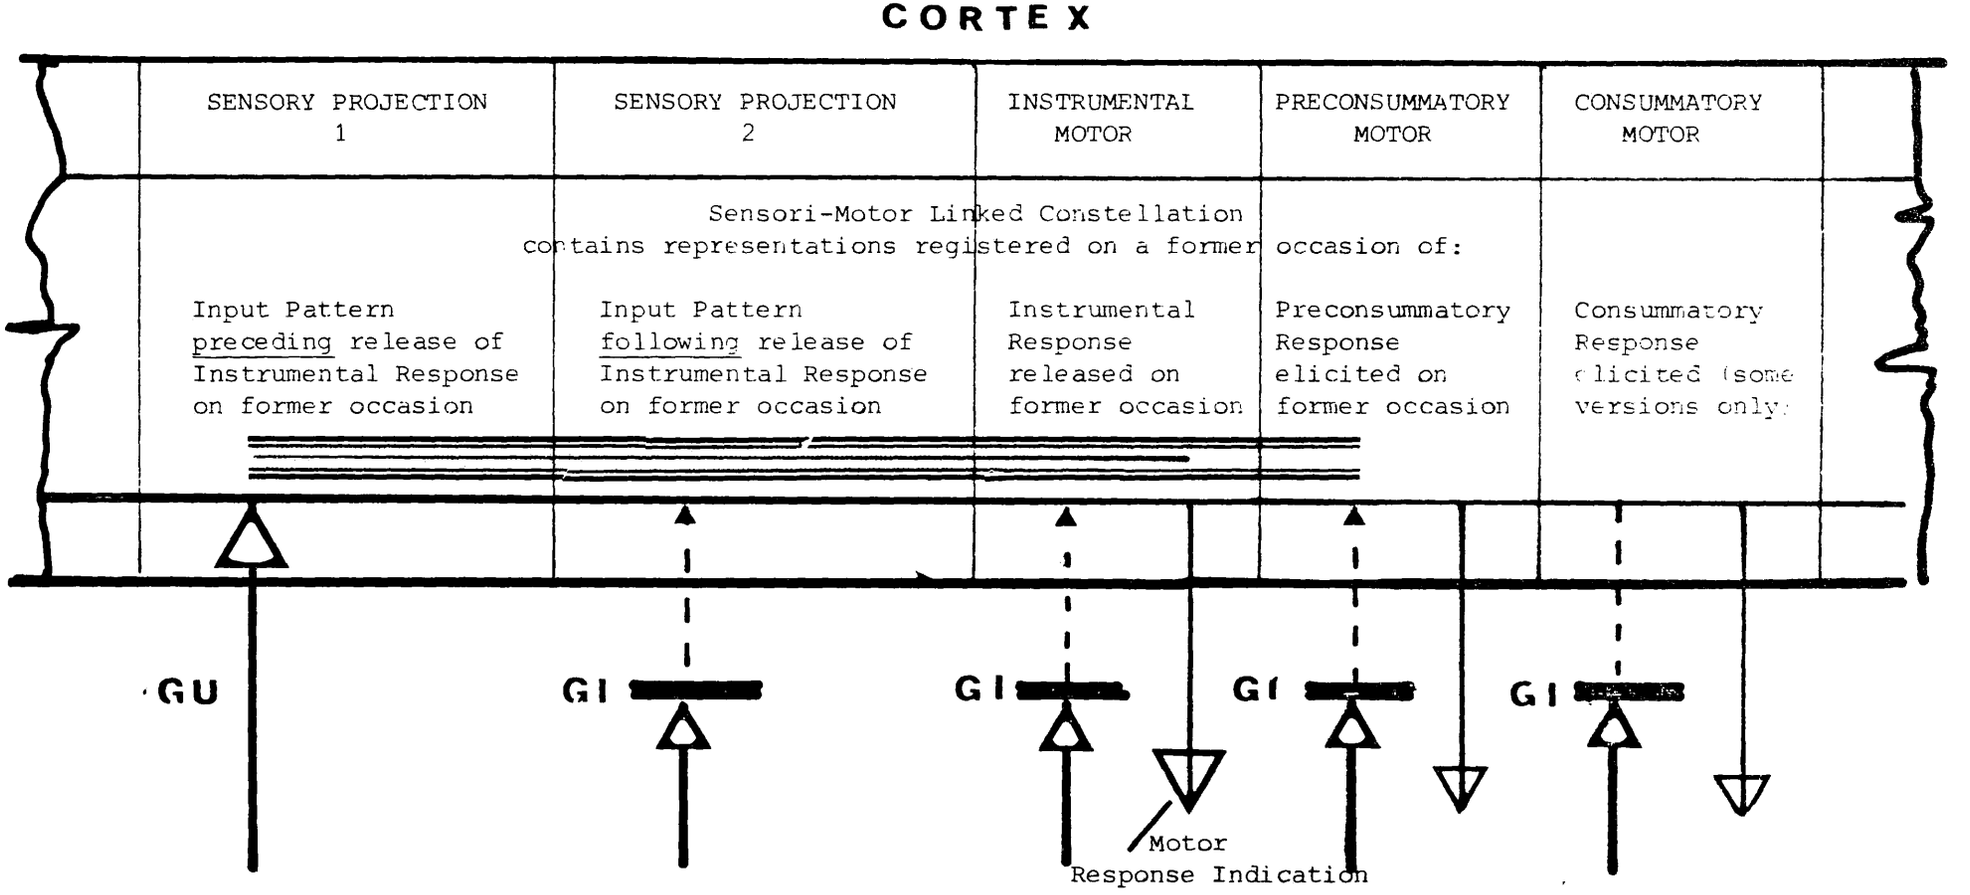

Figure 15.2 The Condition of Inhibitory Gates when the model seeks a Motor Response Indication. Before the next Instrumental Response is taken, the current Input Pattern matches the Input Pattern registered in Sensory Projection 1 of the Linked Constellation which registered a previous successful emission of the Response. In all the other Arrays, the events which were registered in the Constellation on the former occasion have not yet occurred on this occasion. If Excitation were admitted to those Arrays a mismatch would occur. Accordingly the only Gate which is in the Uninterrupted Condition is that to Sensory Projection 1. This effectively drives the relevant Linked Constellation and induces a Motor Response Indication. GU - Gate Uninterrupted Condition; GI - Gate Interrupted Condition.

In the seeking mode, the system was awaiting a Motor Response Indication from the Instrumental Cortical Array. It may be recalled that this Motor Response Indication was to be obtained from the reactivation of a Sensori-Motor Constellation laid down on a former occasion of emission of a successful Instrumental Response in that situation (i.e. location of the maze bearing the same Input Pattern). That Constellation had been formed some 13 Instants after the release of the Instrumental Response. It contained in different Cortical Sessory Arrays representations of two Input Patterns. One of these was the Input Pattern present before the release of the Instrumental Response (delayed pathway), the other was the Input Pattern which appeared as the result of the Instrumental Response (more immediate pathway). To obtain the Motor Response Indication when the model was once again receiving the pre-response Input Pattern, only the Cortical Sensory Array containing the representation of that projection was relevant. A mismatch would have occurred if the pathway to the other Cortical Sensory Array had been unimpeded. Accordingly the Inhibitory Gate to the latter Array was in the Interrupted Condition during the seeking mode; the Cortical Sensory Array containing the pre-response Input Pattern in the relevant Constellation, was in the Uninterrupted Condition.

In the seeking mode, the reactivation of the Constellations was being driven mainly by the Input to the Cortical Sensory Array receiving the delayed signals. Since these energised the part of a Constellation containing an event which had occurred later (i.e. the Response) this was the anticipatory pathway. It was required to detect that anticipatory Constellation response in the Cortical Instrumental Motor Array. This Constellation response was the Motor Response Indication. To do this effectively, as discussed earlier, the Inhibitory Gate to that Array was put into the Interrupted Condition. This illustrates a case where it might be said that "attention" was paid to an aspect of the neural process by reducing the Excitation to an Array (Figure 15.2).

In this mode, at each cycle of the Gate to the Cortical Instrumental Array at a particular phase, i.e. when Excitation within the Array had fallen below a criterion, the Motor Response Indication was taken. Notionally this meant opening another inhibitory gate (not formally simulated) across the Efferent Fibres of the Array, timed to coincide

with the main Inhibitory Gate to the Array.

As soon as an adequate Cortical Motor Response Indication had been obtained, or alternatively, a Response emission accepted for release by the Motor Release Mechanism (Thalamus), the seeking mode was brought to an end. In the typical runs referred to, the model was put for a short time into a "recuperative" mode. (In some exploratory runs still in progress, this time was used to obtain another secondary response of the system, see below). In this recuperative mode, all Gates were put into the Interrupted mode for about eight Instants. This was found to be important as ensuring that a sufficient number of Unlinked Neurons were available for the crucial registration of the Sensori-Motor Constellation some 13 Instants later.

At the end of the "recuperative mode", the Gates to all the Arrays whose activity was to be registered in the critical Sensori-Motor Constellation were put into the Uninterrupted Condition to ensure adequate registration, i.e. both of the Sensory Arrays, the Consummatory, Preconsummatory and Motivational Arrays.

In experiments still in progress, the Consummatory (Final) Response pathway was delayed. At the point under consideration, i.e. registration of the Instrumental Response, the Gate to the Consummatory Array was put into the Interrupted Condition and only put into the Uninterrupted Condition somewhat later. This was to give effect to the function of the Consummatory Response as being the event which caused the reduction of Motivational Excitation (see Chapter 19).

Whenever an Instrumental Response had been registered, and in the absence of a Preconsummatory or Consummatory Response, the model was put into the recuperative mode (all Gates in Interrupted Condition) for some time. In the case of the emission of a Preconsummatory Response, the Gates continued in the Uninterrupted Condition, (full arousal) for a further interval to allow for the reception and registration of the Consummatory (Final Response) and then the recuperative mode intervened, and was made to last for longer than after an unsuccessful Instrumental Response. At the end of the recuperative phase, the model was put once again into the seeking mode.

The modes represented are not intended to indicate the mode of action of the apparatus in a Stimulus-Response mode. The simulations aimed at

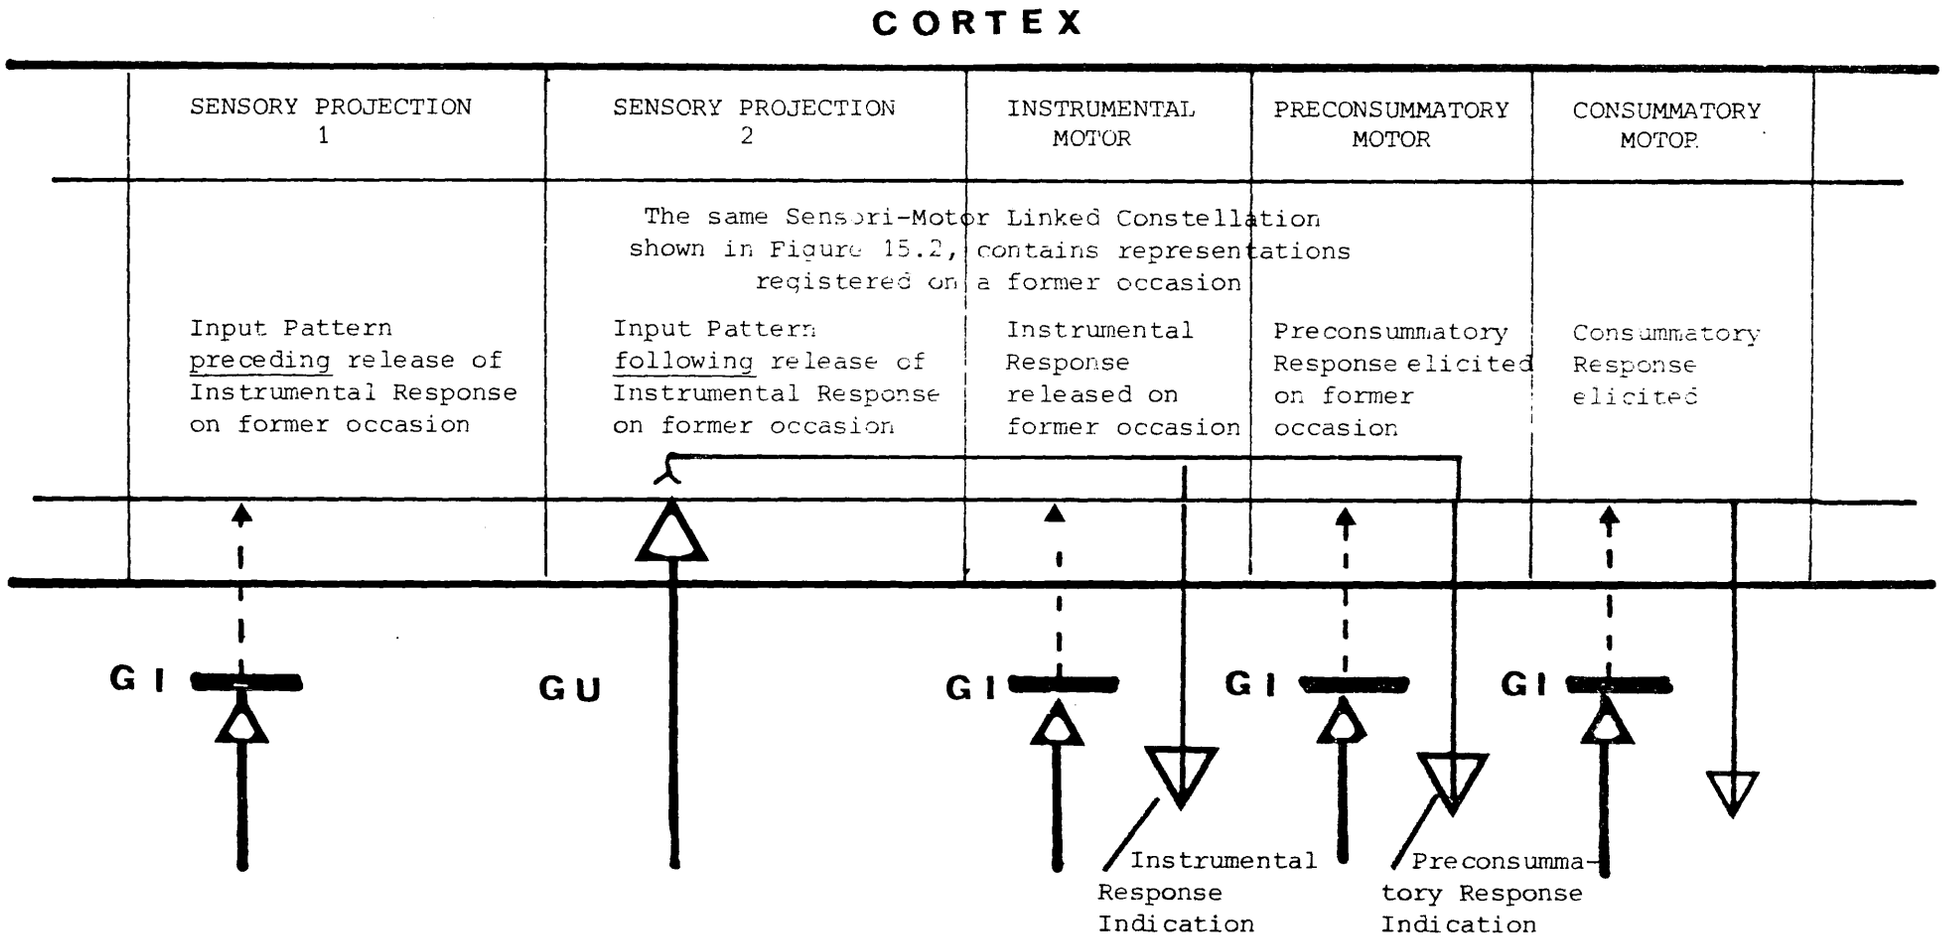

Figure 15.3 Illustrating the possible use of Inhibitory Gates to provide an internal test of the consequences of an emitted Response. Immediately after release and before registration of the Instrumental Response the Gate to Sensory Projection 2 is put into the Uninterrupted Condition and the other Gates into the Interrupted Condition. The same Linked Constellation which gave the Motor Response Indication is now driven by the current Sensory Input Pattern rather than the delayed one. The current input represents the sensory feed-back from the Instrumental Response, i.e. the post-emission Input Pattern. If it matches the post-emission pattern registered on the former occasion a Motor Response Indication will now occur which matches the one leading to the emission. Also an Anticipatory Response can be obtained from the Preconsummatory Array confirming that on the former occasion a "reinforcing event" followed the Response emission.

optimising learning, rather than representing what might happen when a well-learned behaviour was being performed. It seems possible that a recognition of a well learned stimulus could be devised which would put the model into the Stimulus Response Mode. The Gates would then be operated as described to obtain each Motor Response Indication, but immediately afterwards a further Motor Response Indication would be obtained without the more elaborate arrangements for registering and adjusting the system.

In some exploratory programs still under investigation, the Inhibitory Gates became of primary importance. In these runs, the critical Sensori-Motor Constellations registering the Instrumental Responses were labelled with the Preconsummatory Response which had become associated successfully with the emission of the Instrumental Response concerned. Opportunity was taken of the interval immediately after the release of an Instrumental Response, and before its registration, to obtain the Motor Response Indication of the Preconsummatory Array as driven by the post-instrumental-response Input Pattern. This provided a TOTE like test as towhether the Instrumental Response which had just been emitted was in fact a "correct" one. If it had been correct, then it would have had the effect on the environmental routines of producing from the Input the identical post-emission Pattern as had occurred on previous occasions. For this purpose, the Inhibitory Gate to the Cortical Sensory Array with the delayed pathway was put into the Interrupted Condition and the Gate to the Sensory Array with the more immediate pathway was put into the Uninterrupted Condition. The Gate to the Preconsummatory Motor Array was put into the Interrupted Condition in order to detect the secondary Response in that Array. The detection of an appropriate Response Indication in the Preconsummatory Array confirmed that the Instrumental Response had achieved its environmental effect. This Motor Indication was driven by the Input Pattern which had appeared in the Input after the release of the Instrumental Response. (Figure 15.3).

In summary, the Inhibitory Gates were used in the simulations not only to acheive a solution to the problem of control, but also as a selective attentional mechanism. They were also used to enable the Reactivated Constellation responses to be detected without disturbance by direct Inputs.

XI Rhythmic mechanisms and brain function

The literature relating electrical rhythms observed in the brain to behaviour, both animal and human, is considerable and complex. As already stated, it is considered that such matters go beyond the reasonable terms of reference of this project. The author has naturally given the matter some thought and is of the opinion that it is possible to reconcile much of the evidence with an hypothesis of rhythmic interruption of afferent pathways in combination with the intrinsic properties of neural tissue. However, to pursue such a theory would require considerable research. It would probably also require more knowledge than is at present available in respect of the structure and function of nervous systems. The approach in this project has been to find mechanisms which will enable the basic theoretical model to be taken a stage further in the representation of the brain. At the same time some of the implications of the mechanisms used are explored.

Accordingly, it is not proposed here to consider the behavioural correlates of electrical rhythms. Some idea of the relationship between behaviour and rhythms will be implicit in the aspects of the model which have been discussed.

In the most general sense one may say that the use of rhythmic Inhibitory Gates in the model implies some general relationships. Uninterrupted Conditions of Sensory Gates with Interrupted Conditions of Motor Gates suggest a sensory driving of Motor decisions. Interrupted Conditions of Sensory Gates imply either a state of relaxation or that some implication of meaning of the sensory input to some other sensory location, in which the Gate is in an Uninterrupted Condition, is being derived. A variety of Interrupted and Uninterrupted Conditions suggests selective attention.

To avoid misunderstanding, one or two points may be mentioned. The Interrupted Condition of the Gates, as discussed, need not be identified with the very slow waves which, for example, appear in Slow Wave sleep. Nor need the Uninterrupted Condition be identified with the fast episodes of Paradoxical or Fast Wave sleep. It is probable that special uses of inhibitory gates would intervene during sleep. As mentioned earlier, an Interrupted Condition of a Gate still permits the

relevant Array to be in a highly Excited ("desynchronised") state. Such a state might be driven by Excitation which reaches that Array through the interneuronal exchanges of the Linked Constellation system, or by means of special sources of Excitation, e.g. motivational or perhaps P.G.O. spikes.

In addition to the special attentional and detection functions used in the simulations, there is also the overall control function achieved by adjusting the Gates.

XII Concluding remarks

A number of difficulties encountered in the development of the model have been described in some detail. These difficulties threatened to bring the further development of the model to an end. The introduction of motional rhythmic Inhibitory Gates effectively remedied these problems.

In addition to their practical effect, the rhythmic mechanisms drew attention to a number of aspects of Neural function, as represented in the model. If the problems encountered in the present effort are taken to be indicative of the kind of problems which should be faced generally in endeavouring to provide a realistic simulation of nervous tissue, then these effects of rhythmic mechanisms might have further significance.

CHAPTER 16

SEPARATE SIMULATION INVESTIGATING NOTIONAL SYNAPTIC MEDIATION

The main simulations have been developed upon the basis that the concept of Linked Constellations is independent of the neuronal mechanisms by which they might be mediated. Nevertheless, it has also been proposed that the Linked Constellation concept did not necessarily require an unorthodox form of communication between neurons. Orthodox Gray-type synapses would suffice, provided that profuse interconnections are assumed, and subject to some simple rules for the formation of effective connections from potential connections.

It was considered necessary to test this supposition. Accordingly a separate series of simulations was run. This did not attempt behavioural implementation in the same way as the main set of simulations. They were designed to test that the main properties attributed to Linked Constellations in the main simulations, and represented in those by simple arithmetical procedures, would arise in a system of diffuse synaptic connections obeying simple rules.

I Design of model

The model followed the general layout of Figure 5.7 and 5.12. The number of Neuronal Pools was restricted to a minimum, in order that the number of synaptic connections per Neuron might be maximised within the computer resources available. Six Sensory Neuronal Pools were provided. The Input Patterns were designed so that each Active Fibre ('1') was always accompanied by its complementary ('0'), so that the set of Fibres represented a balanced receptor system.

Two Motor Neuronal Pools were provided, and initially these did not represent a balanced effector system. It was later discovered that some balancing was essential to maximise performance, and accordingly additional Neuronal Pools were added to the Motor representation (see below).

Separate Sub-cortical and Cortical Arrays were not provided in this model. The general idea was to present a succession of Input Patterns, then re-present them and see if Linked Constellations laid down on the first presentation were appropriately reactivated upon subsequent presentations.

It was desirable to include some Motor representation so as to be able to confirm the sensori-motor implications of the Linked Constellation idea. The equivalent of Preset Constellations were induced in the model by reserving the first ten Instants of an experimental run to the presentation of special Input Patterns. These Input Patterns were the same as the Input Patterns presented later in the program to test the model, except that they possessed extra digits corresponding to the Motor Neuronal Pools.

Consequently a number of Sensori-Motor Constellations were formed during the first ten Instants of the run. Following the initial ten Instants, Input Patterns were presented which matched Patterns presented during the first ten Instants on the Sensory side, but always had zeros in the locations corresponding to the Motor Neuronal Pools.

II Excitation arrangements

As in the main simulations, each Afferent Fibre which was Active delivered an equal amount of Excitation to each Integrating Neuron in its Pool. The other assumptions of the simulation now under discussion were relatively simple. An Afferent Fibre whose corresponding Input Pattern digit was a '0' delivered no Excitation to the Neurons within its Neuronal Pool.

In order to examine more closely the assumptions concerning inter-neuronal exchange of Excitation, in later versions of this simulation, the Excitation delivered at each cycle of the program was reduced. The number of cycles constituting an Instant was increased (usually to about ten cycles). Consequently there was a finer temporal analysis. All Neurons obeying the rules for forming Integrative Connections within an Instant, i.e. ten cycles, were still labelled as members of the same Linked Constellations.

These variations were introduced into this set of simulations to test that the preliminary results obtained with a somewhat cruder system, would still hold up under this more realistic representation. In the first versions of this simulation, an Active Fibre would deliver 1000 Spikes to each of the Neurons in its Pool. The thresholds of all Integrating Neurons in the system were set at 75 Spikes. With these figures, at Each Instant each Neuron in an Active Pool discharged a very large number of Spikes. These were exchanged through Synaptic Connections by a sub-cycle within the Instant. In the revised version, the amount of Excitation delivered to each Integrating Neuron at each sub-cycle was in the region of 100 Spikes, i.e. just enough to cause the recipient Neuron to fire once.

An attempt was also made to simplify the physiological assumptions concerning Excitation. All Excitation which did not result in Spikes within a sub-cycle was notionally destroyed (i.e. Excitation Count was set to zero). However, Excitation resulting in a Spike was preserved until the next sub-cycle upon the assumption that the Spike leaving the axon hillock of that Neuron would take some finite time to travel to the Neuron which was to receive it. Accordingly, at the next Sub-cycle, Neurons received not only the Excitation receivable from their Afferent Fibre but also 1 Spike of Excitation in respect of every Neuron which possessed a synaptic foot upon it and which had emitted a Spike at the previous Sub-cycle.

Inhibition assumptions were also simplified in a physiological sense. Inhibition was calculated by the system on a basis comparable to that in the main simulations. However, the action of the Inhibition was significantly different from that of the main simulations. In this version, Inhibition cancelled out Excitation. At the beginning of each Sub-cycle, the Excitation Count of each Integrating Neuron in the system was adjusted as follows:

(i) it was reset to zero

(ii) its value was put equal to the sum of the Afferent Excitation delivered by its relevant Afferent Fibre and the Interneuronal Excitation received from other Neurons in the system which possessed notional synaptic connections upon it

(iii) this sum was reduced by the Inhibition Count of the Neuron concerned

(iv) Inhibition Count was reset to zero

(v) the number of Spikes emitted was then calculated as the nett Excitation Count divided by the Threshold.

An important feature of this representation concerned the quantitative effect on an Integrating Neuron emitting an impulse. If the Neuron possessed N effective synaptic feet upon other Neurons in the system, then when it emitted a Spike, N Spikes were delivered into the system. On the other hand, a Neuron required to receive more than a single impulse to be able to pass on a Spike. It will be readily seen that if the Threshold is set equal to the average number of synaptic feet which each Neuron possesses upon other Neurons, then in principle the system is neither amplificatory nor attentuating. Actually this rule, although approximately indicating the properties of the system does not adequately reflect its operation. The destruction of Excitation which does not reach the required Threshold level without the creation of a Spike, introduces some distortion.

III Simulation of synaptic connections

Each Integrating Neuron in the system was assumed to possess profuse potential synaptic connections upon randomly distributed other Neurons in the system.

The actual method was briefly as follows:
At each Sub-cycle the Excitation Counts of all the Neurons were first adjusted as described above. Then each Neuron was considered in turn. If its Excitation was below Threshold it was rejected. If it was above Threshold, then notionally it could form effective synaptic feet upon other Neurons upon which it had potential connections provided that the postsynaptic Neuron was in a similar condition. Random numbers selected another Integrating Neuron in the system, and if, and only if, that Neuron also fulfilled the linkage condition, the synaptic foot upon that Neuron was notionally made effective. This was done by entering the address of the post-synaptic Neuron. The number of such synaptic feet that a Neuron could make effective was limited to a certain number. In most runs this was 175 for each Neuron. This number exhausted the core store of the computer.

Accordingly when the foregoing process had been completed, each Integrating Neuron which was Unlinked at the beginning of the Instant and which had exceeded threshold, now possessed a list of addresses of other

Neurons upon which it possessed synaptic feet. It only remained at the end of each Instant to check through all the Neurons, and select those Neurons which had been Unlinked at the commencement of the Instant, and which now possessed synaptic feet. The relevant Constellation Label was then entered into the Linkage Label location of the Neuron concerned. Thereafter it could not form Integrative Connections. Decay was not represented in this set of simulations.

IV Interneuronal exchanges

At the commencement of each Sub-cycle, the Excitation Counts of all the Neurons in the system were, as described above, set to zero. To perform the next step, that is to enter the Excitation received from both Afferent Excitation and Interneuronal exchanges, it was necessary to calculate the latter figure for each Neuron in the system. This was done in the following manner. Each Neuron was considered in turn. The number of Spikes which it had emitted during the previous Sub-cycle had been recorded in a special location within the Neuron. This number was then added into the Excitation Counts of each of the Integrating Neurons whose addresses appeared in the list of Neurons upon which the reference Neuron possessed a notional synaptic foot (or feet). In this way the Interneuronal exchanges mediated by synaptic connections of an orthodox one-way axonal-dendritic type were roughly simulated. Obviously the simulation was limited by failing to repeat this interchange sufficiently frequently to mimic a continuous analog system. However, within the limits of the simulation, it was considered to represent the Linked Constellation idea mediated by orthodox connections.

It will be noted that the Constellation Label played no part in the implementation of this system. It was purely the result of the exchanges recorded between the synaptic feet. The only minor purpose served by the Constellation Label was to indicate that the Neuron concerned had its full complement of effective synaptic feet. The actual purpose of writing in the Constellation Label was to enable the programmer to analyse the performance of the system. These Labels enabled the Instant at which particular Neurons had been linked to be identified, and so the theoretical attributes of the proposed Linked Constellations could be checked out.

There is one minor variation which should be mentioned. With the system as just described, and lacking a "ripeness" or decay system, all the Neurons of an Active Pool would have been Linked at the first occurrence of the Active Volley. To prevent this a probability control was introduced. Only if a random number chance system indicated accordingly, was the Neuron regarded as being in the condition to form Integrative Connections.

V Performance of the system

A Confirmation of formation of Linked Constellations and their attributes

As was predicted, at each Instant of the program the sub-set of Integrating Neuons excited by the Afferent Input Pattern formed an interconnected sub-set. In all the versions of the program, the basic attributes of Linked Constellations were confirmed. In particular, when an Input Pattern which had been presented earlier in the program, say at Instant I, was re-presented at a later Instant, the Integrating Neurons bearing the Constellation Label I, that is, those Integrating Neurons which had formed interconnections at the earlier presentation of the particular Input Pattern, were found to be the most highly Excited in the system.

Secondly, if at Instant I, the Input Pattern had possessed an Active Motor digit, then Integrating Neurons were naturally found Linked into Constellation I situated in Motor Pools. Similarly at Instant J, when another Input Pattern had appeared with an Active Motor digit, Neurons were linked in another of the Motor Pools. It was confirmed that when the Input Pattern presented originally at Instant I was subsequently repeated, but this time lacking the Active Motor digit, nevertheless, the Neurons in the relevant Motor Neuronal Pool rose to maximum Excitation. This could be compared with the Neurons belonging to Constellation J, and situated in the other of the two Motor Neuronal Pools. The latter Neurons were found to be at lower Excitation levels and to Discharge fewer Spikes than the Neurons situated in the Motor Neuronal Pool which was Active at the first presentation of the same Sensory Input Pattern and which was once again in the Input.

IBINPAT 1 0 1 0 1 0 0 0 ILIFE 12 INPUT PATTERN: 42

ICOUNT FOR RINHIB EQ 6

LINKED CONSTELLATIONS ; DISTRIBUTION OF EXCITATION THROUGH NOTIONAL ORTHODOX EXCITATORY SYNAPSES

TABLE SHOWING CONSTELLATION LABELS AND EXCITATION OF NOTIONAL NEURONS,DISPLAYED IN NEURONAL COLUMNS

SENSORY (AFFERENT)												MOTOR (EFFERENT)			
COLUMN 1		COLUMN 2		COLUMN 3		COLUMN 4		COLUMN 5		COLUMN 6		COLUMN 1		COLUMN 2	
C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.
3.0	33.1	11.0	11.1	5.0	63.1	0.0	0.0	0.0	0.0	0.0	0.0	4.0	67.1	0.0	0.0
7.0	22.1	3.0	24.1	6.0	14.1	10.0	0.0	0.0	0.0	11.0	20.1	0.0	0.0	0.0	0.0
3.0	30.1	0.0	0.0	5.0	61.1	6.0	15.1	0.0	0.0	7.0	26.1	4.0	92.1	0.0	0.0
12.0	81.1	0.0	0.0	4.0	73.1	0.0	0.0	5.0	59.1	0.0	0.0	0.0	0.0	0.0	0.0
3.0	37.1	3.0	30.1	4.0	73.1	0.0	0.0	7.0	22.1	11.0	17.1	0.0	0.0	0.0	0.0
2.0	25.1	6.0	18.1	5.0	57.1	6.0	12.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3.0	20.1	0.0	0.0	4.0	108.1	6.0	15.1	0.0	0.0	3.0	26.1	0.0	0.0	6.0	17.1
3.0	29.1	0.0	0.0	4.0	113.1	0.0	0.0	7.0	16.1	7.0	17.1	4.0	94.1	6.0	17.1
2.0	16.1	0.0	0.0	0.0	0.0	9.0	0.0	0.0	0.0	11.0	15.1	4.0	102.1	0.0	0.0
4.0	89.1	0.0	0.0	6.0	17.1	6.0	20.1	4.0	106.1	7.0	18.1	0.0	0.0	6.0	19.1
3.0	27.1	11.0	21.1	6.0	18.1	5.0	73.1	7.0	19.1	0.0	0.0	0.0	0.0	0.0	0.0
11.0	15.1	0.0	0.0	2.0	25.1	5.0	63.1	5.0	51.1	7.0	20.1	0.0	0.0	0.0	0.0
7.0	22.1	6.0	14.1	5.0	69.1	6.0	20.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
11.0	11.1	11.0	16.1	4.0	81.1	0.0	0.0	0.0	0.0	0.0	0.0	4.0	107.1	0.0	0.0
3.0	24.1	3.0	39.1	5.0	67.1	6.0	11.1	7.0	17.1	7.0	24.1	0.0	0.0	0.0	0.0
4.0	77.1	0.0	0.0	6.0	9.1	0.0	0.0	0.0	0.0	11.0	17.1	0.0	0.0	0.0	0.0
12.0	89.1	0.0	0.0	5.0	77.1	9.0	0.0	0.0	0.0	7.0	27.1	0.0	0.0	0.0	0.0
2.0	20.1	11.0	19.1	4.0	72.1	10.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4.0	86.1	6.0	11.1	0.0	0.0	10.0	0.0	0.0	0.0	3.0	24.1	0.0	0.0	0.0	0.0
11.0	14.1	11.0	13.1	6.0	16.1	0.0	0.0	0.0	0.0	11.0	11.1	0.0	0.0	0.0	0.0
11.0	10.1	2.0	30.1	5.0	77.1	5.0	67.1	7.0	20.1	0.0	0.0	0.0	0.0	0.0	0.0
4.0	66.1	2.0	25.1	4.0	95.1	5.0	55.1	0.0	0.0	7.0	22.1	4.0	105.1	0.0	0.0

ICOUNT FOR RINHIB EQ 6

Figure 16.1

These results are illustrated in Figures 16.1 and 16.2 which reproduce tables printed out by the computer during the runs. The numbers displayed under Column headings consist of (a) the Constellation Label, and (b) the Spikes emitted, in respect of sample Integrating Neruons located within the Neuronal Pool named at the head of the column.

Figure 16.1 represents the state of the model at the third sub-cycle of Instant No. 12. At this Instant the Input Pattern was Pattern No. 42. This Input Pattern had been previously presented at Instant No. 4, with an Active digit (a '1') in Motor Column 1. It may be noted that in general the Integrating Neurons bearing the Constellation Label '4' emitted more Spikes than others in the system. There were exceptions, some of the Neurons did not make adequate synaptic connections with other members of the Linked Constellation to which they were supposed to belong.

It may also be noted that the Neurons in Motor Column 1, bearing the Constellation Label '4', emitted considerably more Spikes than those in Motor Column No. 2, bearing Constellation Label '6'.

Figure 16.2 shows the state of the model in the same computer run, but at Instant No. 15. Input Pattern No. 28 is now in the Input. This is the Pattern which formed the Sensory part of the Input at Instant No. 6, there being at that time an Active digit in Motor Column 2. The Integrating Neurons which now generally show the highest Excitation are those bearing the Constellation Label '6', as predicted by the Linked Constellation concept. Moreover, it is now Motor Column 2 which contains Neurons at the highest Excitation Levels, the Neurons in that Column bearing Constellation Label '6'.

It may be noted that the system is not without irregularities. These support the observations already made earlier in respect of the main Constellations. For example, the Integrating Neurons in Column No. 6, bearing Constellation Label '10' are at the highest in the system. It will however be noted that all the Neurons bearing this Label are situated within Column 6. This can be explained on the hypothesis, argued in Chapter 6 , Section I, in connection with balanced receptor systems, that a Linked Constellation resulting from a small number of Active Fibres would be reactivated maximally by a large number of

IBINPAT 0 1 1 1 0 0 0 0 ILIFE 15 INPUT PATTERN: 28

ICOUNT FOR RINHIB EQ 6

LINKED CONSTELLATIONS ; DISTRIBUTION OF EXCITATION THROUGH NOTIONAL ORTHODOX EXCITATORY SYNAPSES

TABLE SHOWING CONSTELLATION LABELS AND EXCITATION OF NOTIONAL NEURONS,DISPLAYED IN NEURONAL COLUMNS

SENSORY (AFFERENT)												MOTOR (EFFERENT)			
COLUMN 1		COLUMN 2		COLUMN 3		COLUMN 4		COLUMN 5		COLUMN 6		COLUMN 1		COLUMN 2	
C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.
3.0	8.6	11.0	17.6	5.0	92.6	0.0	0.0	0.0	0.0	0.0	0.0	4.0	16.6	0.0	0.0
7.0	0.0	3.0	3.6	6.0	105.6	10.0	172.6	14.0	0.0	11.0	30.6	0.0	0.0	0.0	0.0
3.0	11.6	0.0	0.0	5.0	102.6	6.0	119.6	14.0	0.0	7.0	0.0	4.0	21.6	0.0	0.0
12.0	0.0	0.0	0.0	4.0	16.6	0.0	0.0	5.0	112.6	0.0	0.0	0.0	0.0	0.0	0.0
3.0	8.6	3.0	13.6	4.0	16.6	0.0	0.0	7.0	0.0	11.0	21.6	0.0	0.0	0.0	0.0
2.0	16.6	6.0	108.6	5.0	75.6	6.0	91.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3.0	15.6	0.0	0.0	4.0	21.6	6.0	89.6	0.0	0.0	3.0	3.6	0.0	0.0	6.0	97.6
3.0	7.6	0.0	0.0	4.0	24.6	0.0	0.0	7.0	0.0	7.0	0.0	4.0	23.6	6.0	142.6
2.0	22.6	0.0	0.0	13.0	0.0	9.0	166.6	0.0	0.0	11.0	17.6	4.0	21.6	0.0	0.0
4.0	22.6	0.0	0.0	6.0	114.6	6.0	112.6	4.0	22.6	7.0	0.0	0.0	0.0	6.0	128.6
3.0	12.6	11.0	14.6	6.0	123.6	5.0	90.6	7.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
11.0	25.6	0.0	0.0	2.0	7.6	5.0	112.6	5.0	77.6	7.0	0.0	0.0	0.0	0.0	0.0
7.0	0.0	6.0	99.6	5.0	105.6	6.0	120.6	14.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
11.0	20.6	11.0	16.6	4.0	22.6	0.0	0.0	0.0	0.0	0.0	0.0	4.0	21.6	0.0	0.0
3.0	3.6	3.0	7.6	5.0	82.6	6.0	103.6	7.0	0.0	7.0	0.0	0.0	0.0	0.0	0.0
4.0	19.6	0.0	0.0	6.0	93.6	0.0	0.0	0.0	0.0	11.0	15.6	0.0	0.0	0.0	0.0
12.0	0.0	0.0	0.0	5.0	89.6	9.0	166.6	14.0	0.0	7.0	0.0	0.0	0.0	0.0	0.0
2.0	21.6	11.0	9.6	4.0	14.6	10.0	157.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4.0	24.6	6.0	102.6	15.0	0.0	10.0	166.6	14.0	0.0	3.0	13.6	0.0	0.0	0.0	0.0
11.0	19.6	11.0	14.6	6.0	125.6	0.0	0.0	0.0	0.0	11.0	17.6	0.0	0.0	0.0	0.0
11.0	18.6	2.0	13.6	5.0	114.6	5.0	100.6	7.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4.0	14.6	2.0	12.6	4.0	19.6	5.0	81.6	0.0	0.0	7.0	0.0	4.0	21.6	0.0	0.0

Figure 16.2

This figure is a duplication of that which appears on page 386.

Patterns. Of course, in the simulation now being described the Linked Constellation concerned (LC (6)) was not formed from an Input Pattern activating a smaller number of Fibres. All Input Patterns activated three Sensory Fibres together with a Motor Fibre on special occasions (Instant 4 and 6). However, by the time the run had reached Instant No. 10, most of the stock of Unlinked Integrating Neurons had already been exhausted, so it is not surprising that the only Neurons linked at that Instant were drawn from that Pool which had been least fequently called upon up to that time. The efect was, however, just the same as if the Constellation had been formed by a Pattern having a single Active Fibre.

In summary, the simulations confirmed that subsets of Neurons formed by diffuse random connections obeying the rules laid down for the model behaved in the same way, as far as could be seen, as the Linked Constellations of the main simulations. The latter were simulated by computing the values of Excitation Counts according to the hypothesised attributes of the Linked Constellation hypothesis, as defined in Chapter 5. Accordingly the currently described simulation served as an effective demonstration that the proposed attributes of Linked Constellations would be possessed by sub-sets of Neurons which form mutual orthodox synaptic connections of the kind represented in the currently described simulation.

As has been mentioned, it is further suggested that subsets of Neurons categorised by interconnections of any kind which transmitted Excitation between Neurons, would serve the same purpose.

VI Inhibiton, Excitation, rhythmic mechanisms and other aspects of the separate simulation

A few points may be noted concerning the behaviour of the system in this separate set of models, since they are relevant to matters raised in relation to the main set of simulations.

(i) Inhibition In this set of models, Inhibition cancelled out Excitation in the manner customarily used in neurological models at a symbolic level. Nevertheless, as will appear below, similar problems arose in regard to Excitation as in the main set of simulations.

Inhibition could be used to control Excitation, and was in fact used. However, the mechanism by which Inhibition was able to control the build-up of Excitation was entirely analogical, except that it related to Inhibition rather than Excitation, to the rhythmic mechanisms used in the main simulations to control Excitation. That is, if the system were flooded with Inhibition at periodic intervals, then since the Inhibition cancelled out Excitation within Neurons, the system periodically cleared the Neurons of their built up Excitation. Necessarily however, this Inhibitory rhythm, if it were to be effective in controlling the build up of Excitation, would prevent the temporal sequential aspects of the system from acting in the sense of a priming Excitation.

(ii) The Excitation introduced into the system through the Afferent Fibres could be amplified, conserved or attenuated by adjusting the parameters of the system in a simple manner. This aspect depended upon the average number of synaptic connections. If an Integrating Neuron had a threshold of 50 Spikes and made 150 synaptic connections, then each time it emitted a Spike, the number of travelling Spikes in the system was increased threefold. Actually, as earlier mentioned, the system was considerably more complex than is indicated by such arithmetic, but the basic relationship was confirmed.

If the system were made conservative or attenuative, then it behaved approximately as the main set of simulations. The build up of Excitation could be controlled by introducing rhythmic Inhibitory Gates across the Afferent Fibres. In these cases it was unecessary to depart from the reactive Inhibition assumptions of the main simulations, i.e. it was not necessary for Inhibition to clear the system, the lack of Excitation did that effectively during the time the Gate was Shut.

If, however, the system were made amplificatory the Linked Constellation system only worked in a very limited sense. What happened was that the first Constellation formed in the system, since its Excitation was amplified as it was shuttled between member Neurons, grew continually in Excitation from Instant to Instant. The next Constelltion formed in the system did the same. Since it started one Instant behind the first, it never quite caught up. The third Constellation formed performed in a similar manner. Consequently, the Constellation whose Neurons stood at the highest Excitation level was always the first one

formed in the system, the next highest the second formed, and so forth. Changes in Input Pattern were unable to influence this situation significantly. Of course, eventually the system ran completely out of control.

This behaviour in the amplificatory version, could not be controlled by rhythmic Inhibitory Gate across the Afferent Fibres. It could however be controlled by periodically flooding the system with Inhibition.

Whilst one must, as already mentioned, be cautious about applying the behaviour of these simulations to the brain, it is an intriguing idea that both Inhibition and Excitation might be subject to rhythmic mechanisms. First, one may note that in nervous tissue, upon the assumption of the model, the system need not be amplificatory, conservative or attenuative in an absolute sense. This might be adjusted by a system of thresholds which changed the relationship of operationally effective synaptic feet to average Spike threshold. In a conservative or attentuative mode, inhibitory flooding might not be needed, but in an amplificatory mode an imposed external source of Inhibition might be essential to maintain stability. At the same time the sequential aspects of the system would disappear. The advantage of this would be that the maximum analysis of the single input pattern permitted to enter could be obtained through the continuing growth of Excitation in the matching Constellation. It would be curtailed by a profound Inhibition.

Whatever may be the relevance of these speculations to brain functions, it must be reported that these implications could be seen in the simulation under discussion.

VII Conclusion

A separate simulation investigated the feasibility of mediating Linked Constellations by orthodox formal synaptic connections. The simulations confirmed, within the limits attempted, that this was in fact the case.

The utilisation of rather different physiological assumptions concerning neuronal Excitation and Inhibition in this set of simulations, served to demonstrate that the characteristics of the main set of

simulations were not the product of the particular notional physiology incorporated.

The model also enabled the implications of amplificatory, conservative or attentuating systems of Excitation to be explored. The effects of rhythmic mechanisms applied to both Inhibition and Excitation were also investigated.

IBINPAT 0 1 1 1 0 0 0 0 ILIFE 15 INPUT PATTERN: 28

ICOUNT FOR RINHIB EQ 6

LINKED CONSTELLATIONS : DISTRIBUTION OF EXCITATION THROUGH NOTIONAL ORTHODOX EXCITATORY SYNAPSES

TABLE SHOWING CONSTELLATION LABELS AND EXCITATION OF NOTIONAL NEURONS,DISPLAYED IN NEURONAL COLUMNS

SENSORY (AFFERENT)												MOTOR (EFFERENT)			
COLUMN 1		COLUMN 2		COLUMN 3		COLUMN 4		COLUMN 5		COLUMN 6		COLUMN 1		COLUMN 2	
C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.	C.L.	EXCIT.
3.0	8.6	11.0	17.6	5.0	92.6	0.0	0.0	0.0	0.0	0.0	0.0	4.0	16.6	0.0	0.0
7.0	0.0	3.0	3.6	6.0	105.6	10.0	172.6	14.0	0.0	11.0	38.6	0.0	0.0	0.0	0.0
3.0	11.6	0.0	0.0	5.0	102.6	6.0	119.6	14.0	0.0	7.0	0.0	4.0	21.6	0.0	0.0
12.0	0.0	0.0	0.0	4.0	16.6	0.0	0.0	5.0	112.6	0.0	0.0	0.0	0.0	0.0	0.0
3.0	8.6	3.0	10.6	4.0	16.6	0.0	0.0	7.0	0.0	11.0	21.6	0.0	0.0	0.0	0.0
2.0	16.6	6.0	108.6	5.0	75.6	6.0	91.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3.0	15.6	0.0	0.0	4.0	21.6	6.0	89.6	0.0	0.0	3.0	3.6	0.0	0.0	6.0	97.6
3.0	7.6	0.0	0.0	4.0	24.6	0.0	0.0	7.0	0.0	7.0	0.0	4.0	23.6	6.0	142.6
2.0	22.6	0.0	0.0	13.0	0.0	9.0	166.6	0.0	0.0	11.0	17.6	4.0	21.6	0.0	0.0
4.0	22.6	0.0	0.0	6.0	114.6	6.0	112.6	4.0	22.6	7.0	0.0	0.0	0.0	6.0	128.6
3.0	12.6	11.0	14.6	6.0	123.6	5.0	98.6	7.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
11.0	25.6	0.0	0.0	2.0	7.6	5.0	112.6	5.0	77.6	7.0	0.0	0.0	0.0	0.0	0.0
7.0	0.0	6.0	99.6	5.0	105.6	6.0	120.6	14.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
11.0	20.6	11.0	16.6	4.0	22.6	0.0	0.0	0.0	0.0	0.0	0.0	4.0	21.6	0.0	0.0
3.0	3.6	3.0	7.6	5.0	82.6	6.0	103.6	7.0	0.0	7.0	0.0	0.0	0.0	0.0	0.0
4.0	19.6	0.0	0.0	6.0	93.6	0.0	0.0	0.0	0.0	11.0	15.6	0.0	0.0	0.0	0.0
12.0	0.0	0.0	0.0	5.0	89.6	9.0	166.6	14.0	0.0	7.0	0.0	0.0	0.0	0.0	0.0
2.0	21.6	11.0	9.6	4.0	14.6	10.0	157.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4.0	24.6	6.0	102.6	15.0	0.0	10.0	166.6	14.0	0.0	3.0	13.6	0.0	0.0	0.0	0.0
11.0	19.6	11.0	14.6	6.0	125.6	0.0	0.0	0.0	0.0	11.0	17.6	0.0	0.0	0.0	0.0
11.0	18.6	2.0	13.6	5.0	114.6	5.0	100.6	7.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4.0	14.6	2.0	12.6	4.0	19.6	5.0	81.6	0.0	0.0	7.0	0.0	4.0	21.6	0.0	0.0

Figure 16.2

CHAPTER 17

SOME MISCELLANEOUS TOPICS

I Introduction

This chapter considers some miscellaneous topics which require some mention.

II Some notes on the computer programming

A. Languages

The initial set of simulations were run on the ULCC Atlas computer and written in Atlas Basic Language. That language approximated to a machine instruction code with supplementary facilities. Iterative loops, return addresses, conditional jumps and other structural facilities had to be constructed by the programmer.

The extended version of the model (Aries) were written in Fortran, supplemented by machine code. They were run on the ULCC CDC 6,600. A version written in Algol was not actually used.

B. Special facilities and techniques

The main difficulty in programming stemmed from the limitations inherent in the computer resources available. These resources were considerably greater than those generally in existence at the time but were still not adequate for an effective demonstration of the model without resource to special techniques. These techniques in themselves necessarily complicated the programming. Only a brief indication of these matters will be given.

The limitations were closely connected with the theoretical basis of the model. Effective memory duration in the model is determined by the number of neurons represented in each neuronal pool of the model. Consequently to run the model long enough to demonstrate its principles demanded large numbers of represented neurons.

With the number of neurons which would fit into the computer core store, a further complication arose. Even when the number of neurons permitted a reasonable demonstration run of the model from the point of view of mediating retention of learning, the actual computer time necessary to process the model was greater than that permitted to the student, given a reasonable turn-around.

The solution was to run the program for as long as permitted at one time. Then the entire core store was copied onto magnetic tape. The next computer job began with copying the magnetic tape into core store, and so on. So the notional lifetime of the model organism took a number of computer runs to complete. However, in the thesis, an "experimental run" of the model means the succession of these individual runs.

To accommodate this strategy a number of special techniques were necessary. A single example will be given. During the relatively slow peripheral operations, the program would have monopolised the central processor. A pseudo-parallel processing was used. By means of special instruction "XRCL" the central processor could be held up. The program was split into two halves, AM and BM; whilst one was being processed the other was being read in and out of core.

Although in some respects, the change to the CDC computer made some of these difficulties obsolete, others appeared. With the extension of the model, the whole of the notional system could not be held in store at one time. Fortunately disc storage had by now been introduced. The model was split into 24 parts and these were buffered in and out onto disc, by much the same technique as had earlier been used to extend the duration of the run on the Atlas.

Most of the experimental exercises on the model required some 750 Instants. Each computer run could only handle 45 Instants within the time allotted to the student. Consequently an "experimental run" took over 15 runs on the computer. In the early days, the turn-round took several days. Later this improved to an overnight service. Even so an effective experimental run took a minimum of 15 days to complete. Even a minor error, a misplaced comma, took a day to discover and rectify.

These constraints brought many difficulties in their wake. Only one will be mentioned. Normally, in a program a variable can be relied upon during the program. However, the shuttling in and out of core memory did not permit the preservation of all the subsidiary variables. Consequently the programmer himself had to set up a system to preserve the values of the main variables.

Typical of the mysterious difficulties which arise when attempting to overcome the natural limitations of a machine, was the appearance of the number "32" in various key places within Integrating Neurons. They often appeared in the locations of Constellation Labels, so before the error was discovered, the model seemed to suffer from some strange disease, that of "proliferating 32's". This baffled all the computer advisory service personnel, until it was discovered that there was a conflict of coding during the tape copy operation which corrupted the information.

In leaving this merest hint of computer difficulties, it may be mentioned that recent developments in the architecture of computers may make the task of programming immeasurably simpler. The current development of disc storage which may be directly addressed without the involvement of the programmer, would mean that the whole of a very large model, could be designed and written in a simple direct manner. Such facilities are not however yet available.

C. Some brief statistics

The Aries program consisted of some 5000 Fortran Instructions. Each computer run took between 600 and 900 seconds on the computer, so the total time of an experimental "run" was about 10,000 seconds. Monitoring output was on Microfiche, and over 600 pages of notional printout, i.e. micro-photo output, was printed for each run. Accordingly there were approximately 9000 pages of printout monitoring a single experimental try-out of the model.

It was thus possible to say with some certainty exactly how the model operated in many respects. An example of such monitoring output may be seen in Appendix A, consisting of the microfiche of one computer run only. An extensive monitoring process of this kind is, in my opinion, an essential point in modelling a subject which is extremely complex and little understood, by means of a system which is also complex unless one has some means of understanding with certainty what is going on in the model.

D. Processing time and real time

This brief mention of computing matters should not ignore a special difficulty related to neurological models.

In the nervous system a large number of neurons do not necessarily take longer to perform their function than a smaller number. In the computer, the processing time is directly related to the number of neurons to be represented. This is so even where the system being modelled, as in the present case, is an analogue parallel processing system. This is because the computer is essentially a serial processor, each neuron involving an additional iteration of those operations, which, taken together, comprise its simulation.

This difficulty offers a very substantial difficulty to the adequate representation of neuronal systems. It was also responsible for characteristics in the models which were not really attributable to their theoretical basis. For example, to produce some kind of performance within a reasonable length of computing, learning had to be greatly "speeded up".

In summary, the computing aspects have in the past, and at the present time, present a considerable obstacle to simulating a brain model of the kind pursued in this thesis. There are however, in the rapid developments of computer technology, hopes that these will be greatly reduced in the near future. The present exercise has shown some progress is possible even with the limitations of current technology and resources.

III Some simulations in the course of development

A. Background remarks

At the time of writing, some further simulations were being developed but had not progressed sufficiently to say whether they would be successful. They are briefly reported here, not to suggest that they might or would work, but because the principles developed in their design may have some bearing upon the theoretical model as a whole, particularly in regard to the difficulties encountered.

The next stage in the development of the models would have been to try to enable the performance of a sequence of Instrumental Actions to

reach a goal, in place of the single Instrumental Action used in the simulations reported. This would be an important step in the effective simulation of brain function in neural terms. With such a model one could go on to simulate co-ordinated perceptual and motor activities. Moreover, the restriction of the previous models to a single Instrumental Response begs one of the major theoretical questions concerning the organisation of behaviour. It avoids the problem of how the effect of the reinforcing event is able to get from the neural records of a formerly reached goal to the initial steps which must be taken to reach that goal.

Although the latter question has been subject to much research and many theories, the mechanism is still not understood, nor its fundamental attributes agreed upon. The modeller may therefore feel fairly free to use what suggests itself.

In the models so far simulated the effect of the reinforcing event, the elicitation of the Preconsummatory Response, was applied to the representation of the Input Pattern which immediately preceded that event by introducing a delay into the sensory pathway. This introduced a basic implausibility, discussed earlier. The reinforcing event must follow sufficiently closely upon the Instrumental Response for the delayed pathway to be successful in integrating them. Of course, the unit of time represented in the models is a notionally short one, so one is really looking at a temporally continuous process in which the sensory pattern preceding an Instrumental Resonse overlaps the emergence of the Preconsummatory Response.

The implausibility of this time scale can be readily avoided by introducing a delayed consolidation concept. Such a concept is extraneous to the principles laid down for the theory of neural integration. It was accordingly desirable to demonstrate the main attributes of the theory without resort to such a mechanism. Having done so, it seems justifiable to utilise biochemical reactions to extend the power of the model.

This is especially the case because consolidation mechanisms seem to have been substantially confirmed by contemporary research. It is not proposed to review the field here, but two sources may be referred to in support of this statement. Landfield (1976) reviews those hypotheses closely connected with the electrical rhythms. Brazier's

(1979) anthology contains many papers devoted to this subject. Routtenberg & Haing-Ja have shown the interaction of two transmitters (1978).

In the model, the mechanism was simple to incorporate. The model remained essentially as before. A difference was introduced into the representation of Integrative Connections. Instead of the Constellation Label being inserted at the time that the Integrating Neuron fulfilled the linking conditions, a provisional marker was entered into the Neuron location. If within a certain number of Instants an event had occurred which it was considered enabled the consolidation to be confirmed, then the marker was replaced by the Constellation Label which would have been entered at the time of the Provisional Linking.

The mechanism worked satisfactorily from a practical point of view. Of course, all Integrative Connections in the model were not made subject to consolidation. In the short time that the mechanism has been experimented with, consolidation was restricted to the Connections in the Motor Neuronal Pools, in particular to those of the Instrumental Pools. In particular, the event which would confirm the consolidation of the provisional connections was the occurrence of a Preconsummatory Response within a given number of Instants.

The system was run sufficiently to demonstrate that in this way the temporal proximity between reinforcing event and Instrumental Response could be overcome. It was not run sufficiently to explore the behavioral implications of introducing long delays between provisional formation of Linked Constellations and their consolidation. Insofar as experiments were run, it seemed that if the delay were made longer than the average time taken to emit an Instrumental Response, complications of learning followed. The model would fail to learn the correct order of a sequence of responses. Although nothing definite can be claimed from the simulations run, it appeared that a notional biochemical consolidation process in itself would not enable the model to learn a sequence of successive Instrumental Responses.

There is however one theoretical aspect of the model which would appear to be broadened by the consolidation hypothesis.

A continuous sensori-motor chain cannot be the only mechanism of motivated behaviour. The type of learning where associations are made only if events occur within a few seconds of each other may be contrasted

with that which occurs between events which are separated by hours (Oatley, 1970) as when an animal suffers illness after eating a specific discriminable food and afterwards avoids it. Some theorists consider that purposive behaviour reflects the animal's concepts of causation derived from his general associations (Mackintosh, 1974).

A biochemical hypothesis of consolidation, especially one in which specific labelled biochemical processes interact selectively, enables the model to be extended to explain such time gaps. One may assume that the consolidation of actions concerned with the alimentary processes are bound together because their consolidations are mediated by a single sub-set of biochemical consolidators. In the hypothesised model, the Constellations indicating the acts of eating would be registered provisionally in both the Activator (Do) and Suppression Pools (Not-Do). In normal circumstances the Activator aspect would be consolidated, but in the case of some trouble in the alimentary system (mediated by a Preset Constellation which would recognise such a trouble) then the Suppression representation would be consolidated. In view of the biochemical specificity assumed to be specific to alimentary processes, this widely spaced temporal integration would be independent of intervening events provided they were not also alimentary events.

As to the biochemical plausibility of such mechanisms, Iversen (1979) remarks in reference to the neuro-peptides as neural transmitters, "It therefore seems that the neuro-peptides may be chemical messengers of a character different from that of previously identified transmitters: they appear to represent a global means of coding for patterns of brain activity associated with a particular function, such as body fluid balance, sexual behaviour, pain or pleasure".

The biochemical consolidation hypothesis as a corollary to the model would appear to broaden its scope considerably. Except for the theoretical implications just discussed, the inclusion of the mechanism does not seem to modify the model which is the subject of this thesis very greatly. However, it seemed of some importance because it greatly assisted another development of programming in hand. This was the mediation of extended sequences of instrumental actions leading to a goal. This is discussed further below.

B. Temporal sequences: motivational links

One means which has been suggested by many theorists (e.g. Hull, Sheffield) as the basis of organisation of successive instrumental actions, is that the effects of reinforcement become attached progressively to objects and situations which occur earlier in the sequence on each subsequent experience with the situation. It is not proposed to discuss the question here except to note that animals do learn sequences of instrumental behaviour with some ease (for review see Hinde, 1970, p.609). The approach here follows the general line already indicated, that special mechanisms exploit the basic attributes of the system.

Two forms of the hypothesis suggest themselves. In one, successive labelling of sensory input would enable the system to ascertain whether an Instrumental Action taken in the presence of that Input would be likely to lead to a goal situation (a Preconsummatory IRP). The second form of the hypothesis concerns the attachment of Motivational Links to progressively earlier steps in the sensori-motor chain, enabling the behaviour to be "driven" towards the goal, as in, for example, Deutsch's (1960) hypothesis. Both these forms are the subject of models in development, and both seemed to be subject to similar difficulties. It is in fact these difficulties, rather than the ideas themselves, which are the subject of the present section, as they are relevant to the models already reported.

In the simulations reported earlier, the emission of a Preconsummatory Response caused the Instrumental Action immediately preceding it to be registered in a Positive (Do) Neuronal Pool. The Preconsummatory Response in those models had to be a full one, that is, it had to arise from a direct reactivation of a Preset Constellation in the Sub-cortical Array, and the Sub-cortical Preconsummatory Response was then transmitted to the Cortex for registration. In the versions now being described, the reinforcement control of registration of Instrumental actions was not restricted only to the direct full elicitation of a Preconsummatory Response in this manner. A sufficiently strong Cortical Motor Response Indication from the Preconsummatory Array (i.e. the Anticipatory Response detected in that Array) served as a reinforcing event, just as had the direct elicitation of the Response served as a reinforcing event in earlier versions.

The question was: how was such an Anticipatory Response to be forthcoming? In essence, that Anticipatory Response was obtained only "one jump" ahead, by exactly the same principles as the Instrumental Response Indication had been obtained "one jump ahead" in the simulations already reported. However, in addition, at each repetition of the situation, the registration of the Anticipatory equivalent of the Preconsummatory Response was moved one step temporally backwards.

In brief, the idea was as follows: After each emission of an Instrumental Response, and at the time of registration of that Response, a provisional registration was also included in the Linked Constellation in the appropriate Preconsummatory Response Pool. The particular Preconsummatory Pool chosen for this provisional registration could not of course, be determined by an actual elicited Preconsummatory Response, because none would have yet occurred in the current trial. Accordingly, the Pool chosen was the one relating to the Cortical Motivational Pool which was active at the biassing level at that time. In simple terms, if the animal was hungry, then the Preconsummatory Pool which would be provisionally registered was that relating to the consumption of food. After each Instrumental Response, then, there was present in the Linked Constellation recording that Response, a provisional registration of a Preconsummatory Response.

The determination of whether that response was to be consolidated (in a positive or suppressive aspect) was taken at the emission of the next Instrumental Response. Suppose, in the simplest case, that this next Response was the one which had led directly to the direct elicitation of a Preconsummatory Response, as in the successful Instrumental Responses of the models already reported (call this Instrumental Response A). Then immediately after the Instrumental Response had been taken a Preconsummatory Response would be elicited. This was the primary reinforcing event which confirmed the consolidation of the Preconsummatory 'label' in the Constellation recording the preceding Instrumental Response (i.e. Instrumental Response B).

So far so good, as regards an Instrumental Response (B) which immediately preceded the one which elicited the Preconsummatory IRP. What about an earlier Instrumental Response, say the one which preceded that just discussed? It will be seen that upon emission of this earlier Instrumental Response (call it Instrumental Response C) it would be

given a provisional registration of the relevant Preconsummatory Label. When the model emitted its next Resonse (B), i.e. the one preceding the final Instrumental Response of the sequence, it would not elicit a direct Preconsummatory IRP from the notional environment. That IRP was only elicited later by Instrumental Response A. However, it will be recalled from the above,that the earlier registration of Instrumental Response B did include a representation of the Preconsummatory Response elicited by Instrumental Response A, termed a Preconsummatory marker or label. Accordingly, immediately after the emission of Instrumental Response B, on this second occasion, one could ascertain whether Instrumental Response C had been a successful one by referring to the Linked Constellation registering B on the earlier occasion.

The reference to the Linked Constellation registering B was made by obtaining the Anticipatory Response from that Constellation. Immediately after the emission of Instrumental Response B, the most highly reactivated Constellation in the system was that registering a former occurrence of Instrumental Response B. Accordingly, at this time the Inhibitory Gates of the model were put into the appropriate Conditions to maximise the Response Indication to be derived from the Preconsummatory Cortical Array; if Response B had led to Resonse A on a former occasion, a Motor Response Indication would be forthcoming from the relevant Pool of the Preconsummatory Array. This could then be used as the signal to consolidate the Preconsummatory Label provisionally attached to Instrumental Response C, and so on indefinitely.

The model reached the stage of being run, but failed to work for the following reason. As soon as a Preconsummatory Response had been registered by the occurrence of a primary reinforcing event, i.e. a directly elicited Preconsummatory Response, then a reference to the Preconsummatory Array to try to detect a Preconsummatory Motor Response Indication, as described above, was always successful. Consequently, all Instrumental Responses emitted were consolidated as though they were successful in leading to the goal, regardless of their actual effect.

The reason that the detection of the Anticipatory Response of the Preconsummatory Array was always "positive" once a Preconsummatory Response had been registered, seemed obvious once it had been seen in the computer simulation. Initially the only registration in that Array was

the result of an actual occurrence of an IRP. Consequently, at this stage of the development of behaviour, the only linked neurons located in a Preconsummatory Cortical Pool were those representing this actual reinforcing event. Accordingly a reference to this Array produced a maximum Response Indication from the Pool concerned simply because there was no competition.

The solution to this problem seemed clear, but unfortunately it was at this time that computing had to cease. Nevertheless, the solution seems to be of some theoretical interest, especially in relation to the motivational problem discussed below. Briefly, it was the registration on every occasion categorised as successful, of both the positive and negative (Activator and Suppression) aspects of the Preconsummatory Response and its Label equivalent. However, the Negative (Not-do) aspect was to be registered later than the Activator (Do) aspect. The rationale for this was as follows.

The Positive and Negative aspects of the labelling, being in Constellations close together in proximity, normally balanced each other out. Consequently, on reference to the Preconsummatory Array for an Anticipatory Response Indication, no positive Indication would be obtained. The only case where that balance would not be maintained was when the Input Patterns in both Cortical Projections matched the Linked Constellation recording a Positive Preconsummatory Label, and did not match the Constellation recording the negative Preconsummatory Label. Such a combined match and mismatch only occurred in respect of those Linked Constellations which had been registered upon former emissions of the Instrumental Response currently emitted in identical Sensory circumstances. Only a combined match and mismatch of the Constellations containing Positive and Negative representations of the Preconsummatory Label disturbed the balance of Positive and Negative and produced a Preconsummatory Response Indication (Anticipatory Response).

To make this clearer, the Linked Constellation registering a former Instrumental emission and containing the Positive Preconsummatory Label contained also a representation of both the pre-emission and post-emission Sensory Input Patterns, that containing a representation of the Negative Preconsummatory Response Label contained only the Post-emission Input Pattern.

Consider the emission of an Instrumental Response which on a former occasion had been given a Preconsummatory Response Label in the Linked Constellation by which it was registered: correspondingly, the Linked Constellation laid down just after the one containing the Positive Label, had been given the Negative Label. Upon the current occasion, immediately after emission of the Instrumental Response, the Preconsummatory Response Indication was sought. Since this Indication was looked for just after emission, the Input Patterns in Projections 1 and 2 consisted respectively of the Input Pattern just before emission and that arriving just after. This matched precisely the Linked Constellation laid down of the Instrumental Response emission on the former occasion, and containing the Positive Preconsummatory Label. Accordingly the Input at this time reactivated the Positive Label.

The position in relation to the Negative registration of the Preconsummatory Label was different. Since this Constellation had been laid down somewhat later, the temporal delay which preserved the pre-instrumental-emission Input Pattern to Projection 2, would have expired. Both Projections would have received the identical Input Pattern, i.e. the post-emission Input Pattern. At the time the Motor Response Indication was taken, on the current occasion, these Patterns were not in the Input in both Projections. At this time, (just after emission) the Input Pattern to Projection 1 contained the post emission Pattern, but that to Projection 2 contained the earlier, pre-emission Pattern due to the Uttley-type delay. Accordingly at this time the Constellation containing the Negative Label was mismatched.

Accordingly, provided the Anticipatory Response of the Preconsummatory Array (Preconsummatory Response Indication) were taken at the right moment, there was a combined match of the Constellation containing the Positive Label, and a mismatch of that containing the Negative Label.

The idea, then, was that after every emission of an Instrumental Response, Preconsummatory Labels, Positive and Negative were provisionally registered. These were consolidated if the <u>next</u> Instrumental Response were "confirmed" by a Positive Preconsummatory Response Indication. With each subsequent repetition of a familiar sequence, the Preconsummatory Label would be transferred to successively earlier steps in the sequence.

The advantage of this scheme was that, other than the use of the

consolidation delay, no new mechanism or concept had been introduced.

C. Motivational bias: Freud and Deutsch

(i) Some early abortive attempts

In the verbal, pre-simulation, formulation of the model, the organisation of extended behavioural sequences was visualised as resulting from motivational bias introduced selectively into the sensori-motor pathway leading to a goal. Essentially the idea was that shown in Figure 12.5 (Chapter 12) but extended over several successive instrumentally conditional responses. The idea was essentially the same as that used later in connection with the motivational mechanisms in relation to a single Instrumental Response. When an action was successful in reaching a primary goal, motivational links were entered into the Constellation registering the goal event. On subsequent occasions, Motivational Excitation entering through these links spread backwards (temporally speaking) through the sequence of Constellations and reinforced the Excitation feeding forwards from the current sensory input.

This was tried out in quite early versions of the simulations. It was however found to be impracticable. There were two difficulties. First, the Excitation had to build up massively in order to reach back as far as the Constellations registering the early steps of the sequence on Instrumental Responses. This difficulty may have been solved by rhythmic gating, had it been envisaged at the time. Before the introduction of rhythmic mechanisms, the build up continued from Instant to Instant and so overwhelmed the discriminative function of the system. With rhythms it would have been possible to allow the build up to occur just prior to the detection of the Instrumental Response Indication, its meaning ascertained, and then system cleared by the operation of the Gate, ready for the next Instrumental Response.

The second difficulty was more fundamental. The Motivational Excitation was in fact found to be effective in biassing the sensori-motor pathway, but not evenly. Since the Excitation entered the system through the goal Record, and spread temporally backwards, it naturally biassed the Responses nearer the goal more than those appearing earlier in the sequence. Consequently the Instrumental Responses for the early steps were not emitted. The model could only emit abortive Instrumental Responses which would have been effective had they been emitted later.

This uneven operation of the Motivational bias had, of course, been anticipated, but it had not been thought to be important. The required Instrumental Response in an early step of the sequence, would, it was thought, be triggered because it received double Excitation, that from the current Sensory Input, and that from the Motivational Excitation. The trouble actually experienced was that in order for the Motivational Excitation to penetrate far back enough from the goal to affect the initial steps, it had to be more powerful than the Excitation from the current Sensory Input.

At the time, the difficulty was not fully understood, and was abandoned because of troubles attributed to the general inefficiency of the system rather than to a difficulty of principle. With the introduction of rhythmic mechanisms and consequent control of the system, the difficulty is seen as more fundamental.

(ii) Solutions under investigation

One way out of the difficulty, whilst preserving the overall rationale of motivational bias spreading from former goal attainment, was to have the Motivational Excitation move backwards by steps. That is, on the first trial the Motivational Excitation only needs to influence the Instrumental Response taken just previously: on the next trial it moves backwards one step. The mechanism than resembles that described for the attachment of Preconsummatory Labels. The latter mechanism can readily be extended to include Motivational Links.

Whenever an Instrumental Response was emitted than, in addition to the provisional registration of the Preconsummatory Label, there were also included provisional Motivational Links. These Motivational Links were in the Motivational Neuronal Pool relevant to the current state of Motivation. Whenever the Preconsummatory Label was consolidated, then the Motivational Links were consolidated as well. Consequently, at successive repititions of the sequence, Motivational Links were attached to progressively earlier and earlier steps of the sequence. (The reader may wish to be reminded that to register Motivational Links all that was required was for an Afferent Volley at the Linking Value of Spike Frequency to be sent along the relevant Afferent Fibre).

The advantage of this scheme over earlier formulations was that Motivational Excitation would be distributed evenly throughout the steps of the sequence. Motivational bias could be kept within bounds

and need only enter for a few Instants at a time. The next step to be released would be that triggered by the current Sensory Input.

Actually, the two aspects discussed, i.e. Preconsummatory Labelling and Motivational bias, were developed together. Accordingly, the implementation was extended slightly to include both a Positive and a Negative (Drive and Anti-Drive) representation of the Motivational Links also. It would however be going to far into programs which did not produce reportable results to discuss this further. In general it may be taken that the use of Drive and Anti-Drive follows the use of Positive and Negative representations of Labels.

The rationale behind the scheme was roughly as follows. When a Motivational State occurred (e.g. hunger) both Drive (Activator Motivational Pool) and Anti-Drive (Suppression Motivational Pool) received Afferent Input at the bias-only Spike Frequency. Consequently Constellations containing representation of Drive were equally biassed as those containing Anti-Drive representation. The Constellations containing Anti-Drive Links also included registration of the Suppression of the Instrumental Response(i.e. links in the Not-Do Instrumental Pool) and of the Suppression of the Preconsummatory Label. Consequently both the Instrumental Response Indication and the Preconsummatory Label Indication were always balanced except at the critical times when one particular Constellation (i.e. the one containing the Activator) was more highly reactivated by the current Input Patterns than the other one containing the Suppression.

In this particular scheme, the Motivational Excitation did <u>not</u> contribute towards the selection of a particular Instrumental Response in the sequence leading to the goal. It was the whole of the sequence which was biassed above the general levels by the selective Motivational Inputs. The selection of particular successive Responses required to be emitted in sequential order was entirely due to the imbalances created by the Input Patterns.

The general resemblance of the scenarios described above to that of Deutsch (1960) may have been noted. There is however a significant difference. In Deutsch's scheme, the Motivational Link to each successive Instrumental Response is turned on by the operation of the Sensory Analyser which it elicits from the environment. This switching action may perhaps be seen as realised in the above scheme relating

to the use of Preconsummatory Labels to control "reinforcement". The Preconsummatory Motor Response Indication would, in the scheme described, be turned on by the Input Pattern occurring just before and turned off by the Input Pattern occurring just after, the Motor Response Indication is detected.

That switching action would not apply to the scheme sketched out for the Motivational bias. In fact, the individual Instrumental Action is not selected at all in the scheme. Motivation simply acts to pick out the whole of a sequence. A switching action could be seen in the version of Motivation in which Suppression (Anti-Drive) was not represented, but here again the switching was merely the effect of Input Pattern on reactivation rather than an actual switching in and out of Motivational Excitation.

These differences from Deutsch's formulation are not reported as being in any way superior. On the contrary, Deutsch's account was sufficiently persuasive to lead the author to compose a considerably more complicated scheme in which the Efferent Fibres of Suppression Pools actively inhibit the Afferent Fibres of Activatory Motivational Pools. To discuss that scheme would however take us too far from programs actually devised.

What seems of interest are the empirical questions raised. Insofar as Deutsch's formulation impresses as fitting the behavioural evidence, it begs a neuronal means of implementation. The difficulty in devising such a means may lead to a re-examination of whether in fact individual instrumental responses are facilitated by motivational conditions, or whether it is not rather the whole sequence of behaviour which is facilitated.

D. Summary of Programs not completed

The above programs cannot be cited in support of the theory of neural integration nor of the mechanisms of behaviour incorporated in them. Since they were not run sufficiently to determine their effects on the operation of the model, they remain merely possible versions of the model which might work. However, unlike verbal theories, they do have the character of being in the form of actual algorithms written into computer programs. Also they fit into the general scheme of the model, and can therefore be seen as implementations of definitely described neural structures.

They have been reported for several reasons. First, they give some indication of how the model would progress in the future. Second, they exemplify how the model may be used to build quite complex neural theories implementing ideas current in the literature, in this case, Deutsch's Motivational hypothesis. Thirdly, they indicate that difficulties encountered in first attempts to implement a model may in fact be overcome by perseverence, provided that it is logically acceptible to adopt the course of trying to adapt the system to the known facts (see discussion in final chapter). Finally, the mechanisms give some indiction of how mechanisms which at first seem limited to exceedingly simple acts may be broadened to explain behaviour of a more complex kind.

IV Cognitive Confirmation

The suggestion was made in Chapters 2, 3 & 10 that behaviour might be confirmed by a kind of TOTE test of the kind proposed by Miller, Galanter and Pribram (1960). The idea was that the understanding of the world need not be interpreted as deriving only from the satisfaction of Needs, whether physiological or psychological. There is also the capacity to understand that things are as they are expected to be.

Having burdened the reader with an account of programs in progress intended to label events with indications of their implication (i.e. Preconsummatory Labels, discussed in the preceding Section) it is now easy to suggest how such labelling could be extended to provide a confirmation of the fulfilment of expectancy (cognitive confirmation).

All that is required is to assume that some Preset Constellations have evolved, entirely analogous in the function to the Preconsummatory Preset Constellations discussed, except that they function internally by signalling to the remainder of the system that some state of affairs exists. Such a Preset Constellation might, for example, be activated by the absence of any internal event of significance. Any event of significance in this context would be the reactivation of any other Preconsummatory Preset Constellation or its Label, whether appettive or Protective. Such a Neutral Preset Constellation would act entirely in conformity with the labelling principle laid down for the usual Preconsummatory Responses. That is, after every emission of an

Instrumental Response, if no other Preconsummatory Response (primary reinforcement) or Preconsummatory Response Label (secondary reinforcement) were forthcoming, provisional registration of this General confirming Preset Constellation would follow, and this would be consolidated if nothing happened on the next Instrumental Response.

The effect of such a Neutral Preset Constellation in the system (what might term a Non-Preconsummatory Constellation) would be to label all Input which had no consequence for the system. Of course it is unnecessary to assume such a strict Neutrality. It seems plausible that, given the system developed so far, there would be Preset Constellations able to recognise all kinds of internal and external conditions, e.g. abdominal discomfort, unduly low temperature. Generally the idea is that all Input would be labelled in some way or other. At every Response emission some confirmation would be obtained from the Preconsummatory Response Indication mechanism described in the preceding Section.

Of course such an Indication would not be forthcoming in the case of novel stimuli. The complete absence of a confirming Preconsummatory Response Indication of any kind could therefore be the occasion of an Orienting Response.

That these speculations are not entirely divorced from the world of the laboratory may be seen by considering the well known experiments of Sokolov (1960), Anokhin (1961) and more recently Walter (1969). These experiments show that an Orienting reaction, defined neurologically and behaviourally, is obtained when a novel stimulus is encountered. Subsequent familiarity with the stimulus, especially in association with motor decision, reduces the neurological signs of orienting, and habituates the reaction. It reappears immediately novelty is introduced into the stimulus configuration.

V Conclusion

Opportunity has been taken to acquaint the reader with aspects of the work not reflected in actual computer demonstrations. The intention was to indicate some characteristics of the models which may assist in their general interpretation as early examples of applying the theory to the representation of brain functions.

CHAPTER 18

AN EXPERIMENT IN PERCEPTION: LEARNING AND THE SPIRAL ILLUSION

I Introduction

The generally accepted explanation of the after-effect of observed movement, as seen for example in the Spiral Illusion, ascribes the phenomenon to the fatigue or inhibition of receptors sensitive to specific directions of motion. However, the theory of neural integration which has been presented in this thesis, qualifies this explanation substantially and proposes a new contribution to the explanation of the illusion. The modified theory suggests a prediction upon which the accepted theory is silent, namely that the duration of the illusion can be extended by learning.

An experiment was carried out to test this prediction. The experiment is reported in a published paper (Houtman, 1974) which is bound into this thesis. The experiment is described in the paper, and accordingly only a limited account will be given here with a view to relating the experiment to the general concerns of the thesis.

II Background

The account of perception and cognition offered by the model presented in this thesis is an "enrichment theory", in the sense discussed in Chapter 2. According to the theory, what is perceived at any instant will depend not only upon the sensory signals sent to the brain from the periphery. Perception in the model is determined by the pattern of excitation induced in the integrative tissue. That pattern will depend upon the ongoing activity before the signals arrive, upon the immediate effects of the signals on the neuronal pools to which they deliver their excitation, and upon the patterns of activity induced by these two sources of excitation in the Linked Constellations already present in the tissue.

If this account is correct, then what is perceived is correlated with the Linked Constellations being formed at the current instant, rather than with current sensory event. Of course, as described in an earlier chapter, in a waking state the signal effects will dominate. Nevertheless, according to the model the meaning of events derives from the reactivation of the stored neural records. Sensation, in the pure sense of the term, is foreign to the ideas of the model. Perception is what occurs when afferent volleys interact with Integrating Neurons. "Bare" sensations could only occur if nothing were added to the afferent signals by the reactivated Constellations. Even at birth it may be assumed that sensory signals meet a tissue which already has ongoing activity and which already contains Linked Constellations, both those which are genetically determined and those laid down during intra-uterine development.

The question which presented itself was whether such an account of perception can be applied to the explanation of actually observed peculiarities of perception. A suitable candidate was the Spiral Illusion.

III The Spiral Illusion: a prediction from enrichment theory

The Spiral Illusion goes back to antiquity and undoubtedly contributed to the magic properties attributed to spirals. After a prolonged inspection of a rotating spiral, there is a disturbance of perception. The illusion is now seen as an instance of the after effects of observed motion. When a contoured pattern passes across the visual field for some time and one then observes a stationary figure, there is an illusion of movement in the opposite direction.

Wohlegemuth (1911) reviewed the early history of the illusion in considerable detail. According to him, Purkinje suggested the first sceintific explanation in 1825. "After having looked for more than an hour at a cavalry procession, the houses appeared to Purkinje to move in the opposite direction to that of the procession. In attempting to explain the phenomenon, Purkinje states that the eye, in endeavouring to fixate each individual soldier, moves unconsciously in the same direction. " This often repeated movement becomes for the time being habitual and continues even after the procession has

p. 406

Br. J. Psychol. (1974), **65**, 2, *pp.* 205–211
Printed in Great Britain

LEARNING AND THE SPIRAL ILLUSION

BY S. D. HOUTMAN
Department of Psychology, Bedford College, University of London

A generally accepted explanation of the spiral illusion is that the receptor cells sensitive to the specific motions being observed become 'suppressed' by inhibition consequent upon that stimulation. However, an original neurological model suggests that inhibition is a short-term process, whereas perception is also dependent upon longer-term 'learning'. The model predicts that the duration of the illusion will be increased by the accumulation of 'learned records' of the moving spiral.

In the experiment, 'learning sessions' followed daily tests of duration, and an increase was demonstrated on the next day, more than 22 hours after such a session. No increase was detected where the 'learning session' was omitted or replaced by a 'fatigue session' employing an oppositely drawn spiral. In this way the 'learning effect' was distinguished from 'suppression', or a possible improved 'discrimination' due to practice. The results confirm the predictions of the new model and are difficult to reconcile with the 'suppression' hypothesis.

The after-effect of observed movement (e.g. the spiral illusion) has attracted an extensive literature because of its challenge to theories of perception. However, following the demonstration of stimulation in the frog's eye (Lettvin *et al.*, 1966) and in the cat's eye (Hubel & Weisel, 1959, 1962), the explanation which ascribes the illusion to the 'suppression' by fatigue or inhibition of receptors sensitive to specific directions of motion was suggested by Sutherland (1961). Further work which led to the general acceptance of the 'suppression' hypothesis included the demonstration of the existence of receptors in the rabbit's retinal field sensitive to particular directions of motion and of their suppression following stimulation (Barlow & Hill, 1963). Sekuler & Ganz (1963) showed that in human beings an increased threshold to motion followed prolonged inspection of moving contours. In short, the hypothesis is that after continued inspection of any particular motion, the population of receptors will have a skew distribution and so the perceived motion is displaced in the opposite direction.

However, this explanation would appear to have weaknesses. For example, if darkness or a field lacking contours is interposed during the 'post-exposure' period, then the illusion is prolonged. That is, its duration extends beyond the time after which it would otherwise have ceased (Spigel, 1962; Pickersgill & Jeeves, 1964). This is difficult to understand in terms of the 'suppression' hypothesis. It would appear that 'suppressed' cells would have at least an equal opportunity to recover whether stationary contours, a dark field, or a field lacking contours intervened, just so long as the interposed field did not repeat the particular motion which had hypothetically caused the suppression. If anything, it would seem that an absence of contour would aid recovery and therefore shorten rather than prolong the illusion.

A neurological model upon which the author is engaged suggests that the role of neural inhibition in perception is principally in the control of attention to the stimulus. Levels of inhibition may be determined by longer-term learning processes, but inhibition itself rapidly disperses. The suppression due to inhibition is a short-term adaptive process which should be distinguished from more permanent 'learning'. In

the model this more permanent learning takes place by the formation of neural 'linkage structures' which record and represent the pattern of discharge in nervous tissue from moment to moment, a pattern which in turn is a function of both the current stimulation and the learned records. Perception which continues for any length of time or has been previously experienced is therefore partially a function of these memory records, whether of previous occasions or of the learning just laid down. The model suggests that the principal factor in the production of the after-effect is the interaction of activity in the 'learned records' of the moving spiral on the one hand, with the 'learned records' of the stationary contours on the other. This interaction is interpreted by the brain as a paradoxical reverse motion.

If the hypothesis is correct then it should be possible to detect longer-term effects than can be explained within the 'suppression' hypothesis. Inhibition or fatigue of receptors can only explain a 'shift' in the perceived motion for the very short durations during which inhibition can be considered effective.

METHOD

The purpose of the experiment was to ascertain whether the duration of the spiral illusion could be extended by a 'learning procedure' which would distinguish a longer-term learning effect from short-term adaptive changes due to inhibition. It was also desirable to distinguish a 'learned effect' in the sense of the model from a 'practice' effect due to an improved discrimination resulting from acquaintance with the after-effect.

The 21 subjects, college students and staff, were divided randomly into three groups. A 10 in. Archimedes spiral was in all cases rotated clockwise at 35 r.p.m. Every subject was given a set of 'test trials' at about the same time on each of three consecutive days, the purpose of which was to obtain a measure for each subject on each of the three days of the average duration of illusion experienced after a standard stimulation of 6 sec. rotation.

The control group were simply given the measurement trials on each of the three consecutive days and then released. The second group were required to undergo a 'learning session' after the test trials on each of the first two days, so that the test trials on the following day should reflect the effect, if any, of the previous 'learning session' after a delay of approximately one day.

The 'learning session' consisted of fixating the rotating spiral for 25 sec. with one eye, closing the eyes, resting for 25 sec., then fixating with the other eye for 25 sec., and so on until five fixations with each eye had been accomplished. Then a resting period of 7 min. was allowed, with eyes open and conversation permitted. Then the procedure was repeated until a total of 15 fixations had been finished on that day. This group were given the identical spiral used for all subjects in the test trials. (See Fig. 1.)

The 'learning session' was given this form for two reasons. First, the theoretical model suggested that the after-effect is a function of the records laid down not only of the moving contours but also of the stationary contours. In the normal experience of the illusion both of these would be equally learned and to produce the predicted prolongation a technique was necessary which would increase the strength of the records of the moving figure without giving the opportunity for strengthening the records of the stationary figure. Secondly, it was necessary to show that any increase of strength of the illusion following the learning procedure was not an improved discrimination due to greater experience of the after-effect. With the procedure followed all three groups had equal practice with the actual illusion, since the learning procedure involved no discrimination nor even an experience of the after-effect.

The third (or 'fatigue') group followed exactly the same procedure as the second, except only that the spiral used in the repeated fixations was drawn in the opposite direction to that used in the test trials, although rotated in the same direction.

The general design of the experiment is diagrammatically indicated in Fig. 2.

The predictions of the 'suppression' hypothesis and of the new model were clearly opposed. According to the suppression hypothesis one would expect no change in duration in any of the

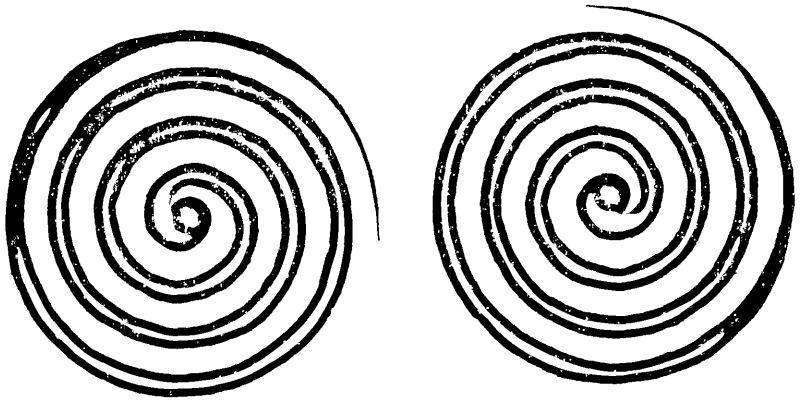

Fig. 1. (*a*) Reduced reproduction of the 10 in. diameter spiral used in the measurement of duration of the illusion for all groups. It was also used in the 'learning' sessions for experimental group 1. (*b*) Oppositely drawn spiral used in 'fatigue' sessions for experimental group 2. All spirals were rotated clockwise.

Daily routine, repeated for each subject on each of three consecutive days

Time (approx.)	Group 1 (control)	Group 2 (learning)	Group 3 (fatigue)
2.00 p.m.	Measurement trials: six readings of duration of illusion after 6 sec. rotation of Spiral A rotating clockwise	Measurement trials: exactly as in Group 1	Measurement trials: exactly as in Group 1
2.30 p.m. to 3.30 p.m.	No action: subjects released	'Learning session' with Spiral A rotating clockwise (i) Fixate rotating spiral with right eye for 25 sec., close both eyes, resting, for 25 sec., fixate spiral with left eye for 25 sec., close both eyes for 25 sec., fixate with right eye for 25 sec.... continue until five fixations with each eye are completed (ii) 7 min. rest, eyes open, conversation permitted. (iii) Repeat procedure (i) as above, i.e. with Spiral A (iv) 7 min. rest, eyes open, conversation permitted (v) Repeat procedure (i) as above.	'Fatigue session' with Spiral B rotating clockwise Procedure as in (i) to (v) of Group 2 but using Spiral B instead of Spiral A

Fig. 2. General design of the spiral illusion experiment.

groups over the three days, since any fatigue or suppression would have dissipated in the 22 hours which elapsed between the subjects' last experience of the spiral and the next test measurement. The neurological evidence is clear that the threshold returns to base level almost immediately after the illusion has ceased (Sekuler & Ganz, 1963) but in any case to claim that the inhibition

Table 1. *Group means (in sec.) of duration of spiral illusion as measured on each of three successive days*

(Each entry represents the rounded mean of 42 readings, i.e. six trials for each of seven subjects in each group.)

Group	Day 1	Day 2	Day 3
Control	7·5	6·4	6·4
Learning	5·0	7·9	11·0
Fatigue	9·4	10·8	7·5

Table 2. *Prolongation or curtailment of duration of after-effect after various experimental treatments*

(In seconds for each subject. Each entry is derived by subtracting subject mean for Day 1 from subject mean for Day 3.)

	Control group	Learning group	Fatigue group
	−6·0	5·9	0·9
	−0·8	2·2	−6·0
	0·5	4·1	−4·9
	0·4	5·2	−0·9
	−0·6	8·0	−0·0
	−0·0	14·8	−2·3
	−0·9	1·4	−0·0
Means of duration changes	−1·1	5·9	−1·9

lasted for 22 hours would make the function of neural inhibition so different from that normally accorded to it as to render the 'suppression' theory on this basis a different hypothesis from that implied in that theory. It is to be noted that an increase of duration does not occur upon mere acquaintance on subsequent days with the illusion (Pickersgill & Jeeves) and that it is only the present technique which predicts such an extension in accordance with the theoretical model proposed. It is difficult therefore under the suppression hypothesis for any prolongation to be accounted for in any of the groups.

However, if it should be argued that perhaps some effect of inhibition might remain even after the deferment, it should in any case be equalized, or nearly equalized, in the case of the groups which were exposed to repeated rotations of the two different spirals. The two spirals contained the same amount of contour and moved in the same direction. Even assuming that the 'suppression' effect was different in them as a result of the illusion of contraction in one and expansion in the other, they both share the same main motion, a motion which is clearly perceived, the circumferential rotation. If any inhibitory effects remained after the 22 hr. delay they would be expected to apply to both of these groups, if not equally then nearly so.

In contrast the new model predicts that the control group will show no change of duration, and the 'fatigue group' which watched the reverse-drawn spiral will be clearly differentiated from the 'learning group', since the reverse-drawn spiral may be expected to interfere with the 'learned records' of the rotating 'forward-drawn' spiral in the test trials.

RESULTS

Each subject provided a set of six scores each day. Table 1 shows the means of the three groups on each successive day. It will be seen that the 'learning group' shows an increase from 5 to 11 sec., whilst both the other groups show reductions. Table 2 gives the prolongment or curtailment of the illusion between the first and third day for each individual. No subject in the 'learning group' failed to show an increase of duration. In the 'fatigue group' three subjects showed substantial decrements, and the remainder registered changes of less than 1 sec., mostly negative. In

Table 3. *Summary of analysis of variance according to model by Winer* (1962)

(Differing treatments applied to separate groups of subjects.)

Source of variation	S.S.	D.F.	M.S.	F
Between subjects:	6023·3	20		
A (treatment groups)	391·2	2	195·6	0·63
Subjects within groups	5632·1	18	312·9	0·63
Within subjects:	3803·0	357		
B (days)	85·5	2	42·7	1·57
AB (groups × days)	923·1	4	230·8	8·47
B × subjects within groups; error B	980·8	36	27·2	—
C (trial scores)	4·1	5	0·8	0·13
AC (groups × scores)	52·6	10	5·3	0·84
C × subjects within groups; error C	560·7	90	6·2	—
BC (days × scores)	16·2	10	1·6	0·28
ABC (groups × days × scores)	155·9	20	7·8	1·37
BC × subjects within groups; groups; error BC	1024·1	180	5·7	—

Table 4. *Partition of sum of squares for AB (groups × days)*

	S.S.		D.F.		M.S.	F (using error B)
Contrast of linear trends	778		2		389	14·5*
Partitioned into:						
Contrast of A (control) against B (learning)		517		1	517	19·0*
Contrast of C (fatigue) against (A+B)		261		1	261	9·6**
Contrast of quadratic trends	145		2		72·5	2·7
Partitioned into:						
Contrast of A (control) against B (learning)		8		1	8	0·3
Contrast of C (fatigue) against A+B		137		1	137	5·0
Total S.S.	923	923				
Total D.F.			4	4		

* Significant: $P < 0{\cdot}001$. ** Significant: $P < 0{\cdot}01$.

the control group six subjects changed less than 1 sec., and the seventh showed a reduction of 6 sec.

An analysis of variance for differing treatments applied to separate groups disclosed that the only significant interaction was that for groups × days, as predicted by the new model ($F = 8{\cdot}47$; d.f. = 4, 36; $P < 0{\cdot}001$). A partition of trends showed that the contrast of linear trends was significant beyond the 0·001 level. The contrast of the learning group against the control group was significant beyond the 0·001 level, and the fatigue group against the other two, beyond the 0·01 level.

DISCUSSION

The results confirm the prediction of the 'learning' model, and support the hypothesis (Spigel, 1965) that the illusion is subject to complex determination, one factor of which is now shown to be learning in a sense not covered by short-term adaptive changes ascribable to inhibition or fatigue, nor to improved discrimination. The augmenting of the illusion remaining after 22 hr. cannot be explained within the ordinary concepts of neural inhibition, and furthermore the fact that exposure to the reverse-drawn spiral produces no increase, but if anything a reduction, cannot be explained by 'suppression', since at the very least the rotation of the two spiral must share a common 'suppressive' factor.

Although the new model does not, as presented, give a detailed explanation of the illusion, a fuller exposition suggests some possibilities.

One possibility is that the 'suppressive' hypothesis is right in its description of the actual mechanism giving rise to the perception of the 'reverse' movement, but that the level of inhibition is controlled by higher-order 'learning'. If inhibition is primarily an instrument of attention, as it is in the present neurological model, then it would be understandable for it to be subject to such control. Douglas & Pribram (1966) have suggested the central control of inhibition as an instrument of learning. However, this explanation would still not explain why the 'suppression' does not produce an apparent slowing down of perceived motion whenever motion is continuously observed.

Another possibility is that the illusion results from an after-discharge of the neurones involved in the perception of the stimulating motion (perseveration), giving a faint continuing perception of 'forward' motion which acts as a ground against which the now stationary 'figure' is seen as 'moving backwards'. Wallach (1959) showed that whichever really moves, it is the figure which is perceived as moving. The autokinetic illusion which occurs when a stationary light is observed in an otherwise dark room demonstrates the difficulty of discriminating motion without a frame of reference. The after-effect may not be due to a direct response of cells reporting 'reverse' motion, but rather a figure–ground effect resulting from a subliminal 'forward-moving' image.

It would seem desirable to ascertain experimentally whether there is in fact a 'forward-moving' image during the perception of the after-effect, or whether the perception is a simple perception of 'reverse' motion of both figure and ground. Various attempts have been made over the history of the illusion to discover the precise nature of the images accompanying the illusory movement: the results are contradictory and uncertain. The reason for this inconsistency is suggested by some pilot experiments by the present author. Any image is very difficult to obtain and the after-image of the rotation is confused with after-images of the stationary aspects of the situation, for example the retinal negative after-images of the boundary of the moving figure (the disc), and even of the pupil itself. A considerable methodological difficulty results from the necessity of a frame of reference to judge the movement of any image which may occur. Attempting to fix such an image as 'figure' in itself makes other images move in the opposite direction. Attempts to judge introspectively whether a 'forward' or 'reverse' moving image is present are therefore confused

by figure–ground effects inherent in the judgement itself. However, further attempts to investigate this question continue.

The now proposed model would appear to be effective in generating further questions for investigation and to make one think anew about an old problem.

The author is grateful to Dr R. Meddis for his valuable suggestions and criticisms.

References

Barlow, B. & Hill, R. M. (1963). Evidence for a physiological explanation of the waterfall phenomenon and figural after-effects. *Nature, Lond.* **200**, 1345–1347.

Douglas, R. J. & Pribram, K. H. (1966). Learning and limbic lesions. *Neuropsychologia* **4**, 197–220.

Hubel, D. H. & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.* **148**, 574–591.

Hubel, D. H. & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106–154.

Lettvin, J. Y., Maturana, H. R., McCulloch, W. S. & Pitts, W. H. (1966). What the frog's eye tells the frog's brain. In C. R. Evans & A. D. J. Robertson (eds.), *Brain Physiology and Psychology*. London: Butterworth.

Pickersgill, M. J. & Jeeves, M. A. (1964). The origin of the after-effect of movement. *Q. J. exp. Psychol.* **16**, 90–103.

Sekuler, R. W. & Ganz, L. (1963). After-effect of seen motion with a stabilized retinal image. *Science, N.Y.* **139**, 419–420.

Spigel, I. M. (1962). Contour absence in the inhibition of the decay of a movement after-effect. *J. Psychol.* **54**, 221–228.

Spigel, I. M. (1965). *Readings in the Study of Visually Perceived Movement*. New York: Harper & Row.

Sutherland, N. S. (1961). Figural after-effects and apparent size. *Q. J. exp. Psychol.* **13**, 222–228.

Wallach, H. (1959). The perception of motion. *Scient. Am.* **201**, no. 1, 56–60.

(*Manuscript received* 4 *August* 1972; *revised manuscript received* 8 *June* 1973)

passed. The eye then wants to fixate the stationary objects in a similar manner to which it has learned to fixate the moving one; it unconsciously slides in the accustomed direction, which makes the objects appear to slip away in the opposite direction".

This explanation sounds rather quaint today. Reference of the illusion to eye-movements was dismissed many years ago. For example, Sekuler and Ganz (1963) studied the after-effect with a stabilised retinal image in order to avoid contamination with eye movement effects however small.

Nevertheless there are two features which Purkinje's explanation has in common with that suggested by my model. It suggests that learning is implicated in the illusion, and it suggests that there is a perseveration of the physiological correlates of movement in the same direction as the observed movement.

One may consider the explanation of the illusion offered by my model at two levels. First there are the general implications of the model as a particular form of enrichment theory. Secondly, there is a more detailed, and naturally more problematical, suggestion of neural mechanisms which may contribute to the illusion.

At the general level, the model suggests a consequence of fixating a visual field which contains moving contours. Whatever the underlying explanation of the after-effect may be, according to enrichment theories in general, and mine in particular, subsequent perception should be added to by neural records of earlier experiences. Since the immediate effect of inspecting a moving field is the customarily perceived illusory after effect, it is reasonable to suppose that whatever it may be in the first experience which gives rise to the illusion would be augmented by subsequent experience.

It may be noted that, at first sight, this predicts that the duration of the illusion would be increased with familiarity. It would also appear at first sight that this prediction is the opposite of that implied by discrimination or differentiation theories. According to differentiation theory, the effect of experience is to facilitate the recognition of those features of the stimulus which differentiate it from other stimuli, and so enable a more veridical perception of the object, as signalled by the sensory information. Acquaintance with the illusion should enable the subject to identify the sensory

realities more promptly, i.e. that the spiral has ceased to rotate, and so shorten the duration of the illusion.

However, such a theoretical opposition oversimplifies the issues. Although the model presented here opposes the differentiation theorist approach, it does so because that approach excludes enrichment rather than because the model suggests that differentiation is not an important aspect of the perceptual process. Accordingly what enrichment theory would predict is that whatever produces the illusion is strengthened by repetition. On the other hand, familiarity with the after-effect does enable the subject to identify the realities more promptly.

With this more sophisticated analysis, one might predict that with repetition of the illusion, the factor which causes the illusion is strengthened but also the ability to discriminate the stationary contours when the spiral stops moving, is also strengthened. If one may be permitted to apply the term "learning" to the effects of laying down records in the system, then the actually observed movement is learned and the stationary contours which occur when the movement stops is also learned.

Accordingly, as may be noted from the attached paper, in designing the experiment it was necessary to separate these two aspects. A procedure was introduced which, on the basis of the above analysis, would increase the learning of the actually observed movement whilst impeding the learning of the stationary contours which bring the illusion to an end. This procedure was to introduce "learning sessions" during which the subject fixated the moving spiral but after which, instead of seeing the contours of the stopped spiral he simply closed his eyes.

IV Two hypotheses of the illusion

As sketched out in the paper, following the discovery of receptors sensitive to particular directions of motion (Lettvin, Maturana, McCulloch and Pitts, 1966; Hubel and Wiesel, 1959, 1962; and Barlow and Hill, 1963) an explanation proposed by Sutherland (1961) gained general acceptance, and remains the generally accepted hypothesis (Sekular, Pantle and Levison, 1978). The hypothesis assumes that

perceived motion depends upon the proportion which receptors sensitive to motion in a particular direction and stimulated by sensory stimulation bears to the total population of motion receptors. During observation of actual motion, receptors sensitive to that motion become fatigued, adapted or otherwise "suppressed". The effective removal of a proportion of receptor elements sensitive to motion in a particular direction leaves the distribution of the remainder of the population biassed towards the perception of movement in the opposite direction. The inspection of stationary contours in this condition accordingly gives rise to a perception of movement in the opposite direction to that viewed during the adaptation period.

My model offers an alternative, or possibley supplementary, explanation. During inspection of a moving field, Linked Constellations are laid down which record what is seen. According to the model, excitation builds up in learned records when matching sensory input continues for some time. After prolonged inspection of a moving field the learned records of the motion are in an excited condition. When the direct sensory signals from the moving contours cease, these records continue to discharge. This after-discharge gives a pattern of excitation similar to that caused when motion continues in the same direction as that occurring when the Linked Constellations were laid down. The new input, that is the signals arising from the stationary contours, now enter the neural tissue, and as emphasised in the exposition of the model, dominate the reaction of the tissue. Nevertheless the pattern of excitation includes the perseveration of the after-discharge representing a subliminal background which is moving in the same direction as the observed motion. This combined pattern of excitation is matched in the neural template system. The matching of a dominant excitation (Figure) against a fainter excitation (Ground) representing moving contours, is interpreted as a movement of the Figure in a direction opposite to that of the Ground, i.e. opposite to that of the observed movement.

The further discussion of these two hypotheses will be left until after the results of the experiment have been considered. It may be mentioned here that they are not necessarily as much in opposition as may at first appear. The explanation proposed by my model does not refute that elements sensitive to specific directions of motion may

be suppressed following observation of a moving field. It does not even need to argue that such elements are not involved in the after-effect. The important element in the explanation offered by my model is that learned records are laid down during inspection of motion and that they contribute to the after-effect.

V The experiment

As the experiment is fully described in the paper, only the general idea will be given here.

There were three groups of subjects, a Control Group, a so-called Learning Group and a so-called Fatigue Group. The Control Group were exposed to a standard "control" opportunity to experience the Spiral Illusion at approximately the same time on each of three successive days. The duration of after-effect reported on these Tests, formed a basis of comparison for the two other Groups. The Learning Group and the Fatigue Group were also exposed to exactly the same procedure as the Control Group, that is, on each of three successive days they were exposed to the standard "control" illusion and the duration of after-effects recorded. Their scores on these Tests provided the scores of the Learning Group and Fatigue Group respectively. Accordingly all three Groups underwent precisely the same test of duration of the illusion on three successive days.

The Learning Group and the Fatigue Group differed from the Control Group and from each other by the Treatment they received each day after the measurements of duration for that day had been taken.

The Control Group was dismissed. The Learning Group was asked to undertake repeated sessions of fixation of the rotating spiral. During these sessions they did not have the opportunity to experience the illusion, at least in its normal form, because after each period of fixation they were asked to close their eyes. The Spiral which was used in the sessions of the Learning Group was the same Spiral as was used for all three Groups for the standard Test Sessions.

The Fatigue Group followed exactly the same routine as the Learning Group, except that the Spiral which they fixated during their Treatment sessions was drawn in the opposite direction to that used in the

standard Test Sessions. Consequently the Spiral used in the Treatment sessions of the Fatigue Group was different from that used to obtain the measurement of durations of subjects in that Group.

The scores of the three groups were compared from day to day. The Control Group did not differ greatly in their scores on any day. The Fatigue Group showed a shortening of the duration of the after-effect from a mean of 9.4 seconds to 7.5 seconds. The learning Group showed an increase over the three days from a mean of 5.0 seconds to 11 seconds. The significance of the result for the partition of variance for the Contrast of the Control Group against the Learning Group was at the 0.001 level.

VI Discussion

The results of the experiment clearly confirm the main prediction of the model. When the opportunity for improving discrimination of the stationary stimulus (or of the circumstances in which it appears) is restricted, then familiarity with the moving stimulus extends the duration of the after effect. Rephrased in terms of the model, if the opportunity to lay down neural records of the stationary stimulus is restricted, then the neural records of the moving stimulus dominate more strongly and the after effect is augmented.

Although the results confirm the prediction of my model, they do not refute the positive contention of differentiation theory, that with experience perception approaches more closely to veridical perception. The Control Group did show a slight shortening of the duration of the after-effect over the three days. However, this result is also concordant with the predictions of the model. The model suggests that an opportunity was taken to lay down more records of the circumstances in which the stationary stimulus appear, i.e. that enrichment is the basis of discrimination.

The implications of the experiment concerning the causes of the after-effect are less clear. Upon fuller consideration, the arguments in the paper concerning the interpretation of the results of the Fatigue Group seem less strong than they did at the time.

The shortening of the duration of the after-effect in the Fatigue Group

is explicable in terms of the suppression hypothesis. Although the oppositely drawn spiral moved in the same direction as the standard spiral, effectively the contours moved in the opposite direction. Pantle (reported as in preparation in Sekuler, Pantle and Levison, 1978) devised a cancellation technique. After establishment of an after-effect, a super-imposed second stimulus moving in the opposite direction simultaneously with the first adapting stimulus, could be adjusted so as to cancel out the after-effect. The conditions of the Fatigue Group in my experiment bear some resemblance to these conditions. If the sessions with the oppositely drawn spiral is regarded as simply exposure to contours moving in the opposite direction, then the shortening of the after-effect resulting from the sessions with the standard spiral is comparable to the subtraction or cancellation effect seen in Pantle's experiment. Of course Pantle's experiment provides no evidence that the illusion is due to the suppression hypothesis. It is equally explicable by the perservera-tion hypothesis suggested here.

What does appear somewhat weaker now is the argument that exposure to an oppositely drawn spiral would provide an equivalent suppression effect to the standard spiral, since it contains the same amount of contour and moves in the same direction. It remains true that there would be a residual effect, as mentioned in my paper. Insofar as the two spirals are <u>perceived</u> to move in the same direction, there would be an equivalence of suppression effect. However, insofar as one spiral is <u>perceived</u> to expand and the other to contract, the suppression effect is not equivalenced. The argument that the results of the Fatigue Group bear against the suppression theory is thereby weakened.

It may be noted however, that whatever the interpretation of the cause of the after effect, the results from the Fatigue Group confirm the main prediction of the model. The effect is enhanced considerably by the third day.

In general the acceptance of the suppression theory has resulted from the overwhelming evidence that receptor units specific to particular directions of motion do exist in nervous systems, together with the plausibility in the light of the neurological evidence, that such units would be inhibited or otherwise adapted by prolonged use. However, these facts are equally acceptible to the perseveration theory.

The suppression theory requires in addition a theory of the perception of motion-direction as a function of the proportion of receptor units able to respond to stimulation of movement in a particular direction. Moreover, even if such an explanation of motion-direction perception were established, it would still not exclude the perseveration hypothesis (see below).

In the literature considerable emphasis is placed upon two pieces of research, Sekuler and Ganz (1963) and Barlow and Hill (1963). From subsequent reported work, there is certainly no need to doubt the validity of their findings. However, the interpretation as supporting exclusively a supression hypothesis is less certain.

Sekuler and Ganz showed that the threshold for the perception of motion after prolonged inspection of a moving field is raised for the direction of movement of the stimulus. Whether this evidence can be taken as critically rejecting the perseveration theory depends considerably upon the approach taken to the action of inhibition or excitation in the nervous system. If one is prepared to accept that neurons may be simultaneously excited and inhibited, as discussed earlier in the thesis, then the perseveration hypothesis is reconcilable with the observed raised thresholds. After prolonged inspection of a moving stimulus field, then excitation will be high in the neurons stimulated by and recording the movement. Inhibition due to recurrent processes would also be high. Upon cessation of the real movement, the excitation and inhibition would decline in step yielding both suppression and a perseveration of discharges.

Barlow and Hill (1963) recorded directly from units in the retinal ganglion of rabbit. During initial viewing of the stationary presentation, there was a maintained discharge. When motion of the stimulus started, the impulse frequency immediately rose to about 60/sec. but declined to about 25/sec. during the first 15-20 seconds. During the next 40 seconds of viewing the moving stimulus, it fell only slightly further. Motion was stopped 57 seconds after it had started and the impulse frequency dropped abruptly to zero - "below the maintained level preceding the stimulus". "During the following 0.5 seconds it climbed slowly back to the level of the maintained discharge".

This evidence argues strongly for the suppression of active units

following stimulation. However, it gives no evidence with regard to the theory of motion perception required in order to establish the suppression of units as the explanation of the after-effect of seen motion. On the contrary, it would seem to have some slight counter indication.

As pointed out by Barlow and Hill, the evidence relates to the rabbit's retina and it "seems unlikely that similar effects occur at a retinal level in man". Perhaps more significant is the lack of verbal corroboration from the rabbit that it perceived a motion after-effect. The point is not whether the rabbit's nervous system resembles that of man, nor whether perception in man resembles perception in rabbit, nor even whether the rabbit did or not perceive an after-effect. Most workers would be content to assume resemblance but not identity in these respects. The point is rather that the "suppression" theorists cannot have it both ways. If the reported results are to be used to support the suppression hypothesis, then the counter indications of these results must also be taken into account.

There is good evidence that although the after-effect requires an active retina, in order to provide contours against which to perceive it, the phenomenon itself is a central not a retinal phenomenon. (Barlow and Brindley, 1963; Pickersgill and Jeeves, 1964). Taking this into account, it would have supported the suppression hypothesis if Barlow and Hill had <u>not</u> found units which are suppressed in the retina! Barlow and Hill suggest that similar units to those which they found in the retina of the rabbit occur at a higher level in man. It seems more likely that units serving analogous functions to those found in the retina of rabbit would be situated in the retina of man. However, even this is not the critical point. The point is that to find units which are suppressed in the retina rather than at a higher level argues against the suppression as the explanation of the illusion, because the latter is a higher level function.

Barlow and Hill's evidence may be used against the perseveration theory upon the argument that no rebound occurred, such as is implied in a perseveration theory. Here again however, the fact that the units were found in the retina stops it being used to support that argument. The rebound or perseveration results in my model from the excitation built up in learned records. These would only occur at the higher levels responsible for the after-effect. There is no need to argue that

neurons at higher levels do not operate upon entirely similar principles to those found by Barlow and Hill, except that they do not form Integrative Connections.

Perhaps the main evidence throwing doubt upon the suppression theory is the demonstration that the after-effect is extended in duration by the interpolation of contour free periods (Spigel, 1962) or by blocking the signals from the eye by anoxia (Pickersgill and Jeeves, 1964). It is difficult to see how this would arise in the suppression model. It is concordant with the perseveration model. Since discrimination results from the learning of the conditions when stationary contours appear, and improved discrimination counter balancing the augmented perseveration effect is the function of reactivation of the stationary contours, absence of stationary contours impedes the discrimination of the veridical conditions.

The experiment reported here adds further evidence. The suppression theory is silent in respect of a long term enhancement. Although it could be suggested that this long term enhancement is due to long term suppression effects, as pointed out in the attached paper, this would make the suppression theory a different one from that customarily put forward. Moreover, long term suppression effects would be expected to affect perception generally. Perception appeared to be quite normal in all subjects when inspecting contours generally. It was only in relation to the particular stimulus that the effect was demonstrated. The author can confirm that during the time that the experiment was conducted his perception of motion was quite normal except when he happened to glance at the spiral. He then immediately experienced a strong disturbance of perception notwithstanding that the spiral was stationary.

When all is said and done, my experiment does not produce conclusive results as to the explanation of the illusion. It may serve to throw some doubts upon the unqualified acceptance of the suppression theory.

Probably the most powerful evidence that there is a perseveration of image moving the same direction as the observed motion would be a report to that effect in the absence of contour. Such images moving in the same direction as the motion are frequently reported in the anecdotal literature. With a view to supporting his, I conducted a number of pilot experiments. I found a methodological difficulty.

It was almost impossible to establish a positive after image of the movement without also establishing an after image of some stationary contour against which the positive image of the movement could be seen as moving.

This difficulty echoes the general difficulty in distinguishing the two explanations of the after-effect. Looking at the literature since the date of the experiment, there seems to be no change in this ambiguity. Almost every experiment which has been carried out can be explained equally in terms of either of the two explanations. For example, complex temporal and spatial frequency effects, cancellations and superimposed effects, reviewed by Sekuler, Pantle and Levison, (1978) can be explained equally by perseveration of an image moving in the same direction and forming a ground, or by a suppression of selective receptors. Other peculiar effects, e.g. the vectorial effects of head tilting (Carpenter, 1977, p.275) or some rather strange effects of apparent movement (Anstis, 1978) are equally explicable, or perhaps inexplicable, according to the two theories.

Of course it may be that both explanations describe processes which contribute to the illusion. This may not be as far-fetched as it at at first seems. Accepting for this purpose, the suppression hypothesis of the perception of motion direction, then immediately upon cessation of stimulation, one might expect a suppression of the stimulated units, giving rise to a suppression induced illusion of motion in the opposite direction. Upon the assumptions of the perseveration theory, this would be followed by an after discharge in the neurons concerned, giving a representation of motion in the direction of the adapting motion. These periodic suppressions and after-discharges might alternate, as has been shown to occur in the cortical neurons receptive to brightness and darkness following a flash of light (Jung, 1972). In the case of alternate suppression and after-discharge of direction sensitive neruons, the perceptual effect, on the combined template theory and suppression theory, would be additive, since both hypotheses predict a perceived movement in the direction opposite to that of the adapting stimulation. This explanation supports Spigel's conclusion that the illusion is subject to complex determination.

VII Conclusion

An experiment was carried out to test predictions suggested by a perseveration theory of the after-effect of observed motion, upon which the generally accepted suppression theory is silent. The perseveration theory predicted a prolongation of the illusion under specific conditions, and this prolongation was confirmed.

An opportunity was taken to discuss the theoretical issues involved. From this discussion it was concluded that the causation of the illusion has not been established with certainty. Although the experimental results may not constitute a challenge to suppression theory, they should cause a reassessment as to the general acceptance of the suppression theory without qualification.

The experimental results support in a general way an enrichment interpretation of perception.

CHAPTER 19

CONCLUSION

I Recapitulation

A theory of neural integration has been presented, and attempts to simulate progressive versions of the theory in a computer described.

In the early chapters of this thesis an attempt was made to isolate the main problems, both psychological and neurological, which an explanation of brain function must eventually deal with. Against this background and against the background of contemporary research on the structure of nervous tissue, the theory of neural integration was presented. It was then proposed, using the computer as a tool of theory building, to construct brain models which might progressively be able to deal with these problems. Naturally it was not to be expected that the eventual goal would be reached, in the sense that all problems would be overcome. On the contrary, the aim was to see how far and with what difficulty a particular theory of neural integration could be applied to that end.

The endeavour was to be regarded as a step towards constructing more comprehensive brain models, that is, to combine a theoretical assessment of contemporary data and opinion with computer simulation to provide effective neurological models. As argued by George (1965), the computer offers a unique opportunity to provide models which are effective in the sense that the implications are demonstrated unequivocally.

An initial set of simulations aimed to demonstrate the main implications of theneural theory. They were comparable in that respect with the simulations which have been presented in the literature in support of other models, (e.g. those of Kilmer and his associates, Spinelli, Walley and Weiden).

Unlike most other simulations of neurological models, the computer models reported here were extended progressively to try to encompass a

broader representation of behaviour and a more realistic brian model in terms of anatomical structure and function.

The models progressed to the stage where a simple sequence of appetitive consummatory behaviour was represented. A number of mechanisms were explored by which successful behaviour, as assessed by the model, would take precedence over unsuccessful behaviour. Within the limitations imposed by the capacities of the model at the stage to which it had been developed, an examination was made of arrangements by which motor behaviour might be initiated, and also the possible relationship between mechanisms which originate behaviour as the result of innate programming and those which adapt that behaviour as the result of experience.

In attempting to draw some tentative conclusions from the work which has been presented, one may consider it from three aspects:
(a) in relation to the particular theory of neural integration proposed;
(b) as a step towards developing a more realistic brain model;
(c) in relation to the use of the computer as a tool of theory building.

These aspects are closely involved with each other, and accordingly it is not possible to separate them entirely in the following discussions.

II General discussion: in relation to the particular theory

At a formal level the simulations have shown that the logical implications of the theory fulfil the predictions made. It was demonstrated that sub-sets of neurons obeying quite simple rules could form neural templates, Linked Constellations. These templates could represent events which occurred from instant to instant. A small number of written-in sub-sets (Preset Constellations) could effectively shape the further development of the system in conjunction with the inputs which occurred from instant to instant. It was also demonstrated that the inbuilt patterns could act as criteria of success and failure and enable the model to assessthe significance of the inputs in relation to those criteria. Storage space could be allocated and retreival activity adjusted accordingly. The system could operate effectively when, after learning, substantial portions of the system were removed.

It was demonstrated that the system of neural integration was capable of

controlling the emission of a notional instrumental response and learn to make the response which had on former occasions led to notional consummatory behaviour. In making a decision as to which response to emit, the model was influenced by its supposed motivational state.

In addition to these formal demonstrations, there were a number of characteristics of the simulations which appeared to be analogous to the behaviour and neurology of animals. It was shown , for example, that when a classical conditioning situation was simulated, a response was obtained from the model which was analogous to the Conditioned Response of animals. These similarities are naturally to be treated with caution as supporting the theory of neural integration. It seems probable that almost any reasonable system of neural integration can be applied to produce a notional conditioning. What seems to be of more importance than a demonstration of a conditioned response is the manner of its production and its particular characteristics. A number of features of the response of the model were found to be similar to that occurring in animals in classical conditioning situation. For example, the equivalent of stimulus generalisation and stimulus discrimination was expressed in the model as a modulation of the strength of the conditioned response much as the degree of salivation varies in the animal in analogous circumstances. Nevertheless, these similarities would no doubt extend to many models.

These less formal aspects of the demonstrations are perhaps best seen as particular hypotheses of the aspects of behaviour which they represent. Clearly when differing theoretical models of neural integration produce analogous CRs, the explanations which they offer will differ from model to model. The present model develops a rather special explanation of classical conditioning. The CR reflects a discharge from memory records about to be activated by a familiar sequence of events. It is clearly distinguishable analytically from the UCR, which acts directly through the inbuilt system. To ascertain whether this is a better theory of classical conditioning than other explanations is an empirical question which falls beyond the scope of this project. What the simulation is able to do is to show a way in which a CR may occur.

The extended set of simulations may be considered as taking the same kind of demonstration further. The aim is to explore the possibilities of the theory. The simulations obviously do not demonstrate that the mechanisms employed are those which occur in actual nervous systems. What they show is that the system is capable of handling behaviour at the level represented.

It may be objected that it is almost certain that a similar attempt to model the same behaviour by some other system would be equally if not more, successful. This is the trouble with theories of neural integration. There are so many contenders. This is true, but there are two answers to the objection.

The first is that although the same behaviour might be modelled in another theory, that result might only be accomplished, as it has in the present effort, after encountering many troubles and difficulties. Those troubles and difficulties exhibit the characteristics of the theory and so enable one to judge better the plausibility of the solution. They also enable one to find out some of the difficulties which are to be faced in constructing a model nervous system. The second answer follows closely upon the first. The important question is not the behavioural performance achieved, but to what extent the resulting models are life-like in important senses.

To answer the latter question is bound to be partly subjective. That is not because it is not a strictly empirical question, but rather because the answer involves such a large body of evidence. Moreover, the total evidence available may be inadequate. The observer is, in such circumstances, bound to be influenced by the weight he attaches to certain aspects of the evidence. In this lack of certainty there is scope for the philosophical predelictions of the individual. In the case of the present author, an endeavour was made to outline the aspects of brain function which were considered relevant to a theory of neural integration. Of course, the simulations did not solve these problems, but one may perhaps judge the life-like quality of the models in terms of steps on the way to attempting to solve them. These aspects are discussed further in the next section.

III Some aspects of the models

A. Overall co-ordination

One of the functions of neural integration which was considered to be important in the earlier discussions in this thesis was the overall co-ordinating and unifying characteristics of the process. It was this impressive characteristic of brain function which led to the concept of a Linked Constellation. It was argued that an extended neural template incorporating all aspects, sensory and motor, of the events would provide an inherent and automatic co-ordination of the separate systems.

Within the limits to which the simulations were developed, this unifying capacity of the system was demonstrated. A number of different aspects of behavioural organisation were represented in the models, including genetically-determined factors, sensory stimulation, temporal relationship, hormonal and metabolic states, different categories of motor action in relation to their interactions with the environment, the frequency of exercise of actions, and the relationship of action taken in the past to the consequences of those actions.

No special mechanisms had to be incorporated to obtain the overall co-ordination and unity of behaviour displayed in the models. Of course, it is true that care had to be taken to ensure the right temporal conjunctions within the model, and that inhibitory and other switches focussed attention on one part of the model rather than another at the right time. These adjustments and developments were however all simply variations of the general parameters and characteristics of the Linked Constellation model.

One must be careful not to overstate the case for the particular theory of neural integration. Konorski (1967) undoubtedly had a similar overall co-ordination in mind when he reviewed the neurological evidence for transcortical connections between different modalities. His system of Gnostic Units solves the problem of overall co-ordination upon the assumption that there is a comprehensive assessment of the entire configuration of sensory, motor and humoral events, similar to that mediated in my model by an extended template. In his system it is mediated by an overall convergence of all inputs to a single apex. A point made earlier may be repeated here. It does not seem clear how

an hierarchical system of this kind could operate unless complete matchings were obtained. A part matching would leave fragmented portions of the system in which matching occurred isolated at a lower level in the hierarchy, and the unity would be lost. That argument relates to his model rather than mine, and undoubtedly the last word has not been said in the debate which will undoubtedly continue between hierarchical and distributed models.

B. Differentiation of Response classes

Another aspect of neural function which characterises the behaviour of living systems in comparison with artefacts, is the differentiation of action in relation to functions which seem to be peculiar to organisms. Actions are mainly organised around acts of consummation, i.e. food gathering and mating, and protection, i.e. withdrawal. There are more than superficial differences between acts which lead to the situations in which these functions can be accomplished and the acts by which the functions are themselves executed. These differences were represented in the model by Consummatory, Preconsummatory, Instrumental and Protective Actions. Although, as was emphasised earlier, these were not to be regarded as theoretical categories, it is reasonable to maintain that their realisation in a computer model gives a life-like character.

Here again, one must be careful not to be misunderstood. The mere fitting out of a model with a variety of responses would not make it more life-like for present purposes. A clever toy might be equipped with teeth and jaws and swallow food. It might possess claws for fighting. Equipped with a sufficiently powerful digital computer and artificial intelligence programs, it might be able to emit appropriate instrumental and protective responses. In fact it is not difficult to imagine that such a robot might perform far more powerfully in life-like behaviour than the excessively limited simulations reported here.

The important point is the theoretical basis of the differentiation displayed. In the present project the initial simulations developed a basic paradigm of motor action and developed differentiated kinds of action from it. A Preset Constellation, when activated by specific input, released a Response. Such innately provided mechanisms are a central feature of nearly all behavioural and neural theories, e.g. key trigger patterns, Innate Releasing Mechanisms, Unconditioned Responses.

In the computer models, this idea was taken further and an attempt made to develop responses required for behaviour as a set of variations on the main theme.

In one respect the basic mechanism of my model, that is, the equivalent to the key trigger pattern of Pringle (1951) differs from other neural formulations of the innately provided mechanism. In my model, the innate mechanism remains unaffected by experience. The effect of experience is to lay down a separate neural record of events, independent of the innate mechanism. The occurrence of an activation of the innate mechanism is simply part of the overall record of events taken from instant to instant.

By a kind of speculative evolutionary process, this basic paradigm was varied so as to meet the requirements of the behaviour to be generated. Consummatory, Preconsummatory and Protective Responses had in common that they all symbolised actions which would be released by specific stimulation, so no variation was required of the basic relase mechanism. Their separation enabled their different relationship to behaviour to be represented. For example, Preconsummatory Responses could be used as an indicator that motivational inputs were to be stepped up, whereas Consummatory Responses indicated that the behavioural sequence had reached its climax, and motivational inputs could be reduced.

This distinction between the Preconsummatory Response and Consummatory Responses of the model was, of course, intended only to represent a wide range of organisational differentiation. Although in the models, each kind of behavioural sequence, say, eating or drinking, was represented by a single Preconsummatory Response and a single Consummatory Response, this was intended to be an abbreviation of sets of Preset Constellations. Each would have the relationship of a Preconsummatory Response to the Preset Constellation which followed in the sequence, and a Consummatory Response to the one which it followed.

In addition the extended template concept enables the Preset Constellation, as a releasing mechanism, to be applied interpretatively to the initiation of action by signals other than sensory. Hormonal, metabolic and other internal signals could contribute or provide the IRP (Innately Recognised Pattern) which induces action. The relationship of Preconsummatory to Consummatory behaviour symbolises in simple form

the kind of hierarchical organisation of specifically elicited behaviours described by Tinbergen (1951).

In the case of Instrumental Responses, the basic Preset Constellation mechanism required further modification. That class of behaviour which is exhibited in instrumental conditioning situations differs from behaviour which is reliably elicited by specific sensory stimulation. The validity of the distinction between "operants" and "respondents" was touched upon in an earlier chapter. It has been argued that the distinction is subject to considerable confusions (e.g. Moore, 1973; Jenkins, 1973). The approach adopted for the purposes of taking the present models further was to assume that all motor organisation is derived from fundamentally similar neural arrangements. There are however differences of emphasis. The idea of isolating a separate class of Instrumental Responses was to be able to concentrate on arrangements which might be notionally present in all responses but which come into prominence in instrumental learning situations.

Like Consummatory, Preconsummatory and Protective Responses, Instrumental action was produced in the model initially by the reactivation of innate templates, i.e. Preset Constellations. In all types of Response, initiation is primarily from the Preset Constellation, but their effect can be mimicked by reactivation of Learned Records. Instrumental Responses were made to differ from the more definitely stimulus-bound responses by variations of degree. A number of the arrangements explored have been described in detail. Essentially they were adaptations of the Preset Constellation designed to permit the primary (genetically-determined) initiation of the response to be less tied to specific sensory input and to be influenced by other more intrinsic factors. A prominent change of emphasis was towards the precedence of the Learned Indication over the Innate Indication in initiating responses. Although developed separately it may be seen that Instrumental Responses do not involve concepts absent from the other kinds of responses. Whereas as displayed in the model, Preconsummatory Responses were only effectively emitted as the result of a Sub-cortical initiation, elicited by specific stimulation, Instrumental Responses gave precedence to the Cortical Motor Response Indication. The differences between appetitive response types may therefore be seen as differences of threshold.

It does not matter in the present context whether these ideas are right or wrong. What is claimed for them is that they arise in the model as developments of basic concepts, and that they generate behaviour.

C. Reinforcement

In the simulations a number of specific mechanisms of reinforcement were explored. In some versions "reinforcement" consisted of selective registration of either the positive (activation) or negative (suppression) expression of the Response (Do and Not-Do Pools). In some versions both were registered. In some versions motivational links fed Excitation into the representations of successful events. Perhaps more than any other aspect of the simulations, these mechanisms of reinforcement are to be regarded merely as useful abbreviations enabling the model to achieve some kind of successful performance. They show possible ways in which the inherent meaning of events can be pointed up into specialised neural mechanisms by exploitation of the extended template idea. Although it is to be hoped that they may have something in common with the kind of specific neural mechanism developed in actual nervous systems, it is doubtful that they closely reflect such mechanisms.

The particular mechanisms of reinforcement represented in the models may be distinguished from the overall concept offered by the theoretical model. Reinforcement and motivation are important questions in the general interpretation of brain function. Pribram (1971) cites an old adage that all of the problems of psychology ultimately reduce to two: the nature of similarity and the nature of reinforcement. He then proposes that these two are in fact one, that of discerning meaning. At this general level, the model presented here points in the same direction. Essentially the concepts of motivation, reinforcement, and recognition of similarity are merely aspects of the overall organisation of the model. That organisation is directed towards matching the signals of the external world with the desirable state of affairs represented by innate templates, and towards avoiding matchings with undesirable states of affairs represented by other templates.

Accordingly, the model does offer a specific concept of motivation and reinforcement. Motivation refers to the influence of metabolic

and hormonal inputs to the system as compared with sensory stimulation. This does not mean that the two do not interact. A template which records a particular environmental object at a time when motivational inputs are high, may be reactivated by the sensory stimulation from the object as well as by the recurrence of activity in the motivational inputs. Reinforcement in the model is the reactivation of particular Linked Constellations. Such Linked Constellations are either themselves Preset Constellations, and so mediate an innately prescribed recognition of desirable or undesirable circumstances: or else they derive from Preset Constellations through the system of Learned Records.

Such a statement of reinforcement may appear to tie the model down to a rather narrow concept similar to that of earlier Needs satisfaction models. That is not actually the case. There is no reason why the desirable state of affairs towards which the model moves, and the undesirable state of affairs, away from which it moves, should be defined in Preset Constellations representing direct physiological requirements, such as deficits in food, drink, and copulation. With the evolutionary development of systems, Preset Constellations may have developed which monitor the state of the system itself rather than the external inputs to it. The desirable states of the system might be defined in terms of optimum states of variables of the system, for example, when Excitation is temporally randomised across the Pools of Arrays, thus ensuring variety of stimulation. Another manner in which the definition of desirable states may have broadened is by the change of emphasis of Preset Constellations which originally may have acted as signposts to major physiological goals, becoming themselves the goals of behavioural sequences, and becoming "reinforcing events".

Aside from the development of systems mediating the innate propensities of animals, there are possibilities of the broadening of goals through the Learned Records. Such secondary reinforcement comes about in the model in two ways. First, there are the general associative properties of the system, as for example, the transmission of excitation from a Cue Pattern to a Learned Record of a Preconsummatory event. Secondly, and perhaps more significantly, there are the possibilities of <u>labelling</u> Constellations which register events which were originally neutral. Such labelling is illustrated in the reported simulations by the registration of a Preconsummatory Response in a Constellation which

records an input pattern preceding rather than accompanying the reinforcing event.

One of the next steps which would be attempted if one were to attempt to develop the model further would be to see if the Preconsummatory and Motivational labels could be transferred to Linked Constellations registering progressively earlier steps in the behavioural sequence. It was also hoped by utilising a delayed consolidation hypothesis to show that labelling could be made relatively independent of temporal contiguity. These demonstrations remain in the future, and may, of course, be beyond the scope of the neural theory to mediate. Their mention serves to show the general rationale of the models.

D. Notional Neuro-anatomy

The life-like nature of a model does not depend only upon its principles of operation. It also depends upon its "flesh and blood". We may be reminded of Olympia, the doll in the Tales of Hoffman. Her behaviour has made her indistinguishable from a human being. When she falls to the floor and the trapdoor in her back is opened to expose the works, we all know that she is not a girl but a machine. This is not an idea which has lost its validity since the eighteenth century. We still know the difference between an "android" and a man, when we expose the works of the former. At the level with which we are now concerned, the choice between two models of the nervous system may well depend not on the overt behaviour of the systems, but the extent to which the notional neuro-anatomies resemble that of nervous tissue itself.

Here again resemblances between what is known of neural structure and the representation in a model must be treated with caution. Most models, including the present one, are designed with an eye on the results of contemporary research. The question is not whether there are superficial resemblances, but to what extent the resemblance of structure arises from the basic implications of the model itself.

Contemporary research on the histology, neuro-anatomical and general structural arrangements in the nervous system has been cited in context. There appear to be certain resemblances which arise in the model from its basic propositions. First, there is the Neuronal Pool or columnar structure. This is fundamental to the model. It also appears to be fundamental to the mammalian nervous system. Second, there is the

modality layout and the topographical representation as a series of elements, which seems ideally suited to an extended template layout.

These resemblances are those which arise from the basic propositions of the neural theory. In the case of the models developed in the simulations, these structural resemblances seem to progress further, not as an ad hoc imitation of what is known of actual neural arrangements, but as the necessary realisation of the model. For example, the presence of corresponding Neuronal Pools in sub-cortical, cortical and control arrays bears some resemblance to the topographical organisation of the regions of the central nervous system, thalamus, striatum, Superior Colliculus and Cortex. The simple system of afferent and efferent fibres between these Arrays of the model seems not unlike the tracts and pathways seen in the nervous system.

Once again, it must be remarked that such resemblances might arise in other models. What can be said of the present model is that the neural arrangements seem to be concordant with the data. It may be remarked in passing that such an arrangement is not at first sight what one might expect from the feed-back type of model nor from the hierarchical convergence kind of model. It may perhaps be claimed in this context that the structural and histological resemblances between the model and the nervous system goes beyond that achieved by the great majority of simulations at the neurological level, and equals that of systems which have been specifically designed to resemble neural structures, e.g. Kilmer and Olinski's (1974) model of the Hippocampus.

IV Some weaknesses and problems referable to the theory of neural integration

Certain difficulties in the theory were exposed by the simulations. Others were more or less obvious from the start, and will be mentioned only briefly.

The first point to be made was that the development of the model was so slow, and what was achieved was very modest in relation to the problems which must eventually be overcome by a neurological theory. This limited progress may be interpreted in a number of ways.

One may argue that if the Linked Constellation concept is as powerful

an explanatory concept as was claimed for it, then no difficulty should have been experienced in using it to develop systems of far greater behavioural sophistication than were in fact achieved.

Against this two arguments may be brought. The first of these in itself involves more than one consideration.

(1) The main cause of the delay was the trouble experienced with Excitation and Inhibition. These go to the root of the Linked Constellation concept, so trouble with those variables is admittedly trouble with the theoretical concept itself. There are however two possible contributions to the troubles experienced. One explanation which saves the Linked Constellation hypothesis is that Inhibition and Excitation may have been inadequately represented in the model. Two rather different implementations of these variables were adopted respectively in the main simulations and in the separate simulation investigating mediation of the theory by orthodox synapses. Similar troubles were encountered in both. Moreover, it was felt that neither representation of Inhibition and Excitation adequately represented the physiological facts. It may be that the concepts of inhibition and excitation as they must appear in the model, may require further clarification.

An alternative to this explanation is that the troubles experienced with Inhibition and Excitation in the model represent real problems which have had to be solved by the nervous system. This would mean that the Linked Constellation concept on its own is inadequate to explain neural function, but that it is effective when combined with some means of control of excitation and inhibition. The means found to be effective in the model was the introduction of Inhibitory Gates: an alternative, or perhaps complementary mechanism, was a similar rhythmic gating of externally imposed inhibitory influences.

Accordingly one may argue that the delays due to trouble with Inhibition and Excitation were due to an inadequate conceptualisation of the original theory. This inadequacy was disclosed by the subsequent attempts at simulation, and once the problems had been exposed and solved, the development of the model might proceed more rapidly. The questions raised by this argument are empirical. Do troubles such as those experienced with the model actually occur in nervous systems? Is it one of the functions of rhythmic mechanisms to control accumulations of excitatory and inhibitory effects? If research suggests that

such is the case then the Linked Constellation concept would be reinforced by those discoveries.

(2) There is additionally a second kind of argument against the contention that slow progress with the development of the models implies a weakness of the Linked Constellation hypothesis. This argument has been raised elsewhere in the thesis. It is that accomplishing too much too easily would be against a theory as an explanation of brain function. Brains have been in existence for hundreds of millions of years. As far as is known, those aspects of behaviour which typify the higher mammals and Man developed exceedingly slowly. The ethological concepts of Innate Releasing Mechanisms and fixed action patterns, as mediated by Preset Constellations, would account for much of the behaviour which one might attribute to primitive creatures (Maser, 1973). Although primitive creatures would be expected to show some signs of the further capacities which were to arise, those signs, like those to be seen in the models, would be in considerable contrast with the sophisticated behaviours observed in the more advanced animals.

Certain weaknesses in the performance of the models seemed to arise paradoxically directly from efforts to provide them with apparatus similar to that presumably possessed by animal systems. The often erratic and stupid performance which resulted in place of the expected "improvement" may be thought to argue against the overall design of the models, and of their rationale. That this may not be the case may be seen in the example of motivational representation.

It will be recalled that when motivational features were in operation in the model, the models tended to lose their power to discriminate Input Patterns and instead would repeat Responses which had previously yielded reinforcement. However some success was achieved. There is accordingly reason to suppose that the performance could be improved. The empirical question which is raised is whether there is evidence of trouble in animal systems similar to that encountered in the models. Most readers will agree that it is not implausible to suggest that animals "lose their reason" when acting under intense stimulation. Of course, the failure of the model cannot be cited in its support. But the fact that some success was obtained, and that the failure of performance bears resemblances to weaknesses seen in animals, must qualify one's judgement.

Three important steps in the implementation of behaviour, as proposed in the opening chapters, were not reached in the development of the simulations.

One was the representation of sequences in which more than one Instrumental Response was required to reach a behavioural goal. A second was the demonstration that events widely spaced in time and separated by events which have other meanings, can be integrated into a single pattern of action, as when a dog eats poisoned food and does not become ill for hours, yet will avoid the source of food (Oatley, 1970). A third is perhaps more speculative, but accords with the general approach taken here. It is the capacity to recognise that events are as expected, independently of the consequences of those events, i.e. what has been termed here, cognitive confirmation.

These absences from the demonstrations cannot be considered as weaknesses of the theory of neural integration in themsleves. In defence of the models, it is argued that they are merely transitional examples of realisations intended to develop the explanatory potential of the theory of neural integration progressively. That there is a long way to go before demonstrating such capacities is clear, but this does not in itself point to a weakness in the explanatory potential of the underlying theory. What raises a question mark against that explanatory potential are the particular difficulties encountered in the simulations which have been undertaken.

An example may serve to give some indication of questions raised. Initially the view of the motivation and direction of behaviour was related to the distribution of built up excitation in the system. It was visualised that excitation would spread backwards from the representations of previously achieved goals and would reinforce the forward spreading (in a temporal sequential sense) excitation from current sensory input. The possibility of this was seen in the initial simulations which showed a powerful temporal sequential and arousal effect centred around the Constellations representing former reinforcing events (see Figure 8.1). Eventually, however, these very aspects of the system threatened to overwhelm it. Solutions were found in the form of rhythmic inhibitory gates. A revised solution is possible in the sense that the flows of excitation from former achieved goals to meet current sensory excitation can only be permitted to occur

during short episodes, during which gating is suspended. Other solutions centre around the possibility of labelling the representation of one event with a marker which indicates its relationship to another temporally distant event.

Another difficulty surrounding the idea of a feed-forward and feed-backward of excitation concerns the difficulty of mediating a sequence of behaviour leading to a distant goal. What was encountered in the demonstrations was an inappropriate driving of the intermediate responses. A solution was devised, although not demonstrated, using the idea of a balanced application of motivational excitation to both the suppressive and activating representation of the motor activity. Whether these solutions would work, and what the implications of their working would be, it is not possible to say at this time. The questions raised are empirical, but their usefulness may be limited by what it is possible to discover in the laboratory at the present time. Meanwhile, it is difficult to say whether the difficulties encountered are weaknesses of the original formulation or strengths pointing to necessary mechanisms in the nervous system enabling its full potential to be exploited.

A more general weakness of the theory, not directly concerned with the work presented in this thesis, relates to the numerical implications. Its full consideration lies outside the scope of this thesis, but a few words may be necessary in acknowledgement of its existence. Arguments have been brought against many theories, in particular against those following the ideas of Uttley (see e.g. Sutherland, 1959) to the effect that too many elements would be required. A similar attack may be made upon the theory presented here.

It is difficult to estimate the information storage capacity of my model. In a single Linked Constellation, it appears that each linked Integrating Neuron holds one bit of information. However, the temporal sequential conditional probabilities of the system, the appreciation of the significance of which we owe to Uttley, would tremendously increase the information capacity of the system. The amount of this increase is impossible to assess unless we were prepared to say to what extent the system operates in temporal sequential mode, a question which has been consistently raised, but not answered, in the simulations.

Moreover, estimates of the information held in the nervous system are

themselves dependent upon equivalent arbitrary decisions. When, as phrased by Mark (1974) we say that "one fleeting experience is instantly recorded and remains as a potent influence on subsequent behaviour for a lifetime of up to one hundred years" we have no way of knowing to what extent it is in the information laid down in that particular experience which is recorded unaltered. The Gibsons' contribution was to show that such information could be represented as changes effected in the discrimination system itself. In the present model those changes take the form of the influence of earlier experiences on the Linked Constellations laid down to record later experiences. Even the ability to recreate earlier memories suggested by the experiences related by persons undergoing cortical stimulation (Penfield and Rasmussen, 1952; Penfield and Roberts, 1961) do not determine the extent to which recall of memory is subject to reconstruction and filling of gaps by more recent memories.

Nevertheless, when all is said and done, we remain with the impression that the actual storage capacity of the brain is unlimited.

If this is true, the principle of the Linked Constellation concept can be saved if the elements from which a linked structure is composed is taken to be a sub-unit of the neuron, say, a macromolecule. It is not immediately apparent how such a model could be implemented in a practical sense concordant with what is known of the structure and function of neurons. The idea may not be impossible, but lies beyond the present prospect of discussion. Accordingly at the present time it must be concluded that the numerical aspects of the theory, although not conclusively against it even assuming its mediation with individual neurons as its elements, is a weakness.

This short discussion of the weaknesses of the model is by no means comprehensive. However, to go further into the more abstruse implications would take us too far away from what was actually simulated.

V Some further aspects of the models

An aim of the project was to try to approach more closely to a realistic simulation of brain function. As will be apparent from the preceding sections, this aim is not clearly distinguishable from the attempt to demonstrate the explanatory potential of a particular

theory. Both eventually involve finding solutions to problems of generating plausible behaviour in a model.

Accordingly, it is only necessary to emphasise what has already been said. The main aim was to explore the possibility of producing life-like systems in the computer. The mechanisms developed are not put forward as definite hypotheses of neural mechanisms. Their aim is suggestive. In some cases the solutions adopted may be at variance with the evidence. In many cases the author would suggest that the existing evidence, although often extensive, is still inadequate to say definitely whether the mechanisms are plausible or not. What is hoped is that the models succeed in raising empirical questions. Where the mechanisms do not succeed in generating the required behaviour, or where the resulting neurology is at variance with what is known, or where the models seem particularly implausible, these very inadequacies may lead to the formulation of better solutions.

By way of completing the recapitulation of the features of the model, two or three further points will be mentioned.

The models relied fairly heavily upon a representation of the suppression of activity, in an active, that is to say, excitatory sense. In this respect, the model closely resembles that of Konorski (1967). There are differences. Konorski's formulation is an extension although an original one, of Pavlovian concepts. His views of organisation of behaviour tend to be expressed as interactions of specific "reflexes". The function of active suppression in my model is very similar to that proposed by Konorski, although my model was developed without reference to his formulation. In effect, response suppression is an instrument of reinforcement. The mechanisms simulated in my models are similar in function, although not neural implementation, to Konorski's description of the transformation of "positive" into "negative" Conditioned Reflexes. There is however a difference in respect to the organisation of Drive and Anti-Drive. He envisages these as operating as the result of the interaction of different reflexes, e.g. the Food Reflex operates the Anti-Drive to the Hunger Reflex. In my model, behaviour is conceived as organised into sequences by sets of Preset Constellations. These Preset Constellations bear an hierarchical relationship to each other, much as described by Tinbergen, and it is this hierarchical relationship which controls the balance of Drive and Anti-Drive.

These differences are probably not significant, and in any case, my formulation was developed not from a close consideration of the behavioural evidence, but from the requirements of the model.

An aspect of neural organisation raised in the models concerned the problem of deciding when an inbuilt mechanism for the initiation of action should take precedence over a learned indication of which action should be taken next. This problem in the form in which it occurs may be an artefact of the present model. However, Maser (1973) considers that it is one of the main questions to be answered. It certainly presented a considerable and unforeseen problem in generating behaviour in the models.

The fact that the solutions generated in the model seem to have a certain biological flavour in terms of current knowledge may lend support to the supposition that the problems as seen in the model reflect real problems of neural organisation. It may be remembered that a number of different possibilities were explored. These included a supposed maturation of inbuilt response to the environment, a maturation of motor response which incorporated the current sensory input at the time of maturation, processes which were in some ways akin to imprinting and other developmental characteristics of behaviour. Another avenue explored was a form of Preset Constellation which would respond to a class of inputs rather than to a specific input. The most effective method in the simulations was that of notional hormonal or metabolic clocks which acted as pressor drives to the exercise of motor actions. As time passed since the clocks were last reset, excitation rose and so increasingly urged the expression of that particular action driven by the clock concerned.

Whether these mechanisms have any reality in the brain, it is unnecessary to urge in the present context. In their defence it may be said that they are not complete phantasy. They were guided by a fairly extensive acquaintance with the literature (see e.g. Glickman and Schiff, 1967; Glickman, 1973; Milner, 1970; Valenstein, 1970; Margules and Margules, 1973; Isaacson, 1974; Mogenson and Phillips, 1976; and many others). Contemporary views of the biochemical aspects of neural processes, sub-cortical organisation of efferent functions, and the general behavioural characteristics of response may make them plausible to the reader. It is true that there is no sufficient evidence that the mechanisms occur, but then there is no sufficient evidence that they do not.

One may take a similar line with regard to the treatment of attention and activation (arousal) in the model. Here there is the advantage that the principle of attention arises directly out of the basic theory of neural integration. True it encountered considerable difficulties and complications. In fact, it may be said that the implications of the attentional aspects of the model were the aspects which most threatened its development. These difficulties were remedied by mechanisms which are admittedly extraneous to the original theory, that is, rhythmic inhibitory gates. Nevertheless, the attentional aspects of the model remain clearly tied to the extended template concept as mediated by the particular neural theory proposed. They are peculiar to a system in which quantity of excitation is a valid variable of the system. The quantitative values of the excitations of afferent fibres, as well as the pattern of excitation, determines the weight given to the information carried. This may be seen as a direct implication of the theory of neural integration. In this respect, the attentional model may be compared and contrasted with Walley and Weiden's (1973) attentional model based upon Konorski's Gnostic Unit concept.

As to whether the rhythmic mechanisms themselves reflect the functions of the electrical rhythms observed in the brain is, it is suggested, an open question for further investigation. In the models they solved many problems. It may be permissable to point out that in addition to the direct problems which they solved, they offer considerable help in visualising how the model might be developed to represent a broad view of behaviour. For example, in the absence of rhythmic mechanisms, the model cannot really be distinguished from a Stimulus-Response system, as categorised by Oatley (1978) in his review of conceptualisations of brain processes. Notwithstanding the distributed memory and analogue character of the model, in its operation before the introduction of rhythm, the net effect was the emission of a response. With the introduction of rhythms the possibility of obtaining definite Anticipatory Responses, exemplified in the simulations by the Motor Response Indication, became practicable. Thereby the eventual development of such Indications into cognitive attributes become possible.

VI Some concluding remarks concerning the theory of neural integration

Many of the weaknesses seen in the simulations have been argued away as failing to establish the necessary weakness of the fundamental concepts. Nevertheless the encounter with considerable difficulties discloses that the theory is not to be applied in a facile manner to explain psychological attributes. The fact that difficulties which did seem almost insuperable at the time were eventually able to be overcome to take the model to the next state, may in itself suggest that a basic theory of neural integration is insufficient in itself to explain brain function. Perhaps the full understanding of brain function will arise only when one can see, not only the underlying mode of communication between elements of the system, but also the manner in which the deficiences of the system have been overcome.

The theory presented here is by no means the only contender for the role of explaining the basic mode of communication between the elements of the system. In the present case, the general characteristics of the model, with its troubles with excitation levels, inhibition levels and disconcerting variablility, may make one perhaps be more drawn to the more computational approaches to a solution of brain function. It is natural to seek for the remedy of the deficiencies of one model in the strengths of another. Memory and thinking seem in some way too precise to be mediated by the kind of system depicted in the model.

In comparison the hierarchical and iterative network models may seem to offer a more powerful approach. It must, however, be remembered that much of the definiteness of the models of such systems is due to their representation in formal units from which the physiological properties of the living units have been extracted. Moreover, the formal network and hierarchical models have other, and perhaps more fundamental, weaknesses as explanations of brain function.

When confronted with these difficulties one is perhaps tempted to fall back upon the idea that there is some secret the discovery of which would make everything clear.

Whilst this idea is attractive, and undoubtedly has some truth, since every step forward in research probably brings a solution nearer, the experience of the present project supports a rather different approach.

The impression that we must look for some magic answer may be mistaken. As Oatley (1978) has argued, it may be the emergent properties rather than the immediately obvious properties of a system which display the exceptional capacities seen in the higher animals and Man. The present project may be seen as an attempt to discover some of the complications and some of the possible solutions, which display the properties of a system.

The theory, as developed in the simulations, showed aberrations and problems which make it difficult to envisage as the mediator of the instrument of higher mathematics, and highly skilled manipulation of the world. It remains possible that the evolutionary refinement of the system may have brought it to a level of efficiency far beyond that achieved in the simulations. Simple rhythmic mechanisms were able to achieve very great improvements of performance. With improved performance, the problem of applying the system to solve behavioural and psychological problems becomes simpler. The brain has had a very long period of evolutionary time to develop these solutions.

The approach has been taken that any theory of neural integration which is too efficient in computation fails as an explanation of brain function for that very reason. The resulting computational mechanism would be too clever to account for the slow progress and development of animal intelligence. Within the limitations of the simulations, and the stage reached, they have confirmed the predictions of the theory. However, those predictions have only been able to be confirmed by adding subsidiary hypotheses, e.g. rhythmic gates. The mechanisms simulated may be quite different from those actually occurring in nervous systems, but they accord in a general sense with the known data and possess a life-like character. In a general sense, therefore, the simulations confirm the explanatory power of the theory. It must be considered a contender amongst the theories which at the moment represent our insights into brain function.

VII Simulation as a tool of theory building

The models presented in this thesis may be considered to have raised more questions than they have answered. Does excitation build up in the nervous system? Is it necessary to hold up sensory inputs

before the results of neural processing can be discerned? Does the system possess the capacity to develop a "secondary tissue response" from which it can test its expectancies? Is behaviour organised with reference to sets of genetically determined innately recognised patterns of neural excitation?

These questions are not, of course, the original product of the model. The last question, for example, has been asked by the ethologists for a number of years. What the models do is to put these questions into a unique concrete form which is only possible in a simulation.

It is clear that a computer simulations can never confirm that the manner in which it achieves its performance is similar to, or even analogous to, the manner in which its subject matter, in this case, the nervous system, accomplishes similar ends. What a computer simulation can do is to prove that a model can operate in the manner claimed for it. Perhaps it can do slightly more than this. It can show that certain objections which may have been brought against the theory can be overcome.

It has been customary to use the computer in the neurological field for formal demonstration. The simulations in this thesis serve that purpose by exhibiting what the basic postulates of the theory can and cannot mediate. However, in the present project, the use of the computer was not restricted to formal demonstration. It was used as a tool of theory building. Ideas were tried out and adapted to the system until their performance was improved.

Of course, it is not suggested that the use of the computer as a tool of theory building originates in the present effort. Every simulation research necessarily involves some adaptation and struggle with the system. In most computer demonstrations this work is not presented, and it is the result alone which is regarded as important. That approach is partly determined by the formal logical structure of many simulations.

In the simulations reported here, especially in those developed later in the work, there is a change of emphasis. The work is no longer of a formal nature. The rationale of such work may need some clarification.

If a computer simulation is immediately successful in that the system performs entirely according to the verbal predictions made for it, the situation is clear. The position is more complex if the simulation is

only partly successful: if the system acts according to prediction in some senses, but shows difficulties. The failure does not prove that the theory underlying the simulation was not capable of fulfilling its predictions, only that there are complications. But complications are part of the nature of the world.

The failure of a simulation may be due to an number of causes. The one relevant to the present project is where the original formulation has failed to take into account some unforeseen complication. The computer simulation serves to expose this weakness. In a logical sense, one can then go back and prepare a new theory in which the complication is avoided. To some extent, this is indeed patching up a theory to make it fit the facts. On the other hand, that may be exactly what is required when one is trying to represent a subject matter which has been subject to such patching up by evolution. This kind of approach to simulation may therefore be especially apposite to an attempt to apply an explanatory principle to the simulation of a complex system like the brain, so many aspects of which are not well understood.

In such circumstances, when difficulties are encountered, the model may suggest a solution. The resulting simulation can no longer be pointed to as an elegant demonstration of the logical implications of a formal theory. However, it may now perform a rather different, and it is suggested, an equally useful function. It serves as a generator of a new theory about the complex system. The new theory has the advantage of posing rather pointed empirical questions.

In this light, the extension of the simulation to try to model a more comprehensive and realistic nervous system is not to be seen as attempting to provide anwers but rather to generate questions. The resulting models are not to be considered as proposing a unique answer, but rather as approximations which may be refined and modified as more information and better understanding becomes available. In this sense, the computer becomes a companion to thought processes, just as pencil and paper have been to former generations.

VIII Final remarks

Since the student is required to state his claims clearly, the following points are made:

(1) A theory of neural integration has been presented and its implications examined. In view of the vast historical background of the subject, it is difficult to claim of any neurological theory that it is original. Nevertheless in the form in which it has been presented and developed, the theory may be claimed to be original in a sense comparable to those in the literature.

Compatible with restricting the thesis to a reasonable length, the data available from contemporary research in psychology and neurology has been examined in relation to the plausibility of the theory, and to a limited extent, of rival theories.

(2) A computer simulation demonstrating the immediate implications of the theory as an explanation of neural integration has been undertaken. This simulation is comparable to other simulations of neurological models which have been presented in the literature.

(3) An attempt has been made to apply computer simulation to build models of mechanisms which may enable brain models to approach more closely to a realistic representation of brain function. In doing so a number of problems implicit in the original theory have been exposed. Solutions have been found, and questions raised as to what extent those problems reflect actual problems and solutions of neural integration.

(4) The theoretical model was applied to the explanation of the Spiral illusion. New facts were predicted and confirmed.

REFERENCES

ADEY, W.R. (1966). Neurophysiological correlates of information transaction and storage in brain tissue. Prog. Physiol. Psych., 1:1-43

ADEY, W.R. (1970). Spontaneous electrical brain rhythms accompanying learned responses. In The Neurosciences: Second Study Program. Editor-in-chief, F.O. Schmitt. Editors: G.C. Quarton, T. Melnechuk, G. Adelman. Rockefeller Univ. Press, New York, pp.224-243

ADRIAN, E.D. (1947). The physical background of Perception (Waynefleet Lectures), Oxford Univ, Press

ADRIAN, E.D. (1941). Afferent discharges to the cerebral cortex from peripheral sense organs. J. Physiol. (London), 100:159-191

ALBUS, J.S. (1971). The Cerebellum: a substrate for list processing. 23-2064 Spartan Books - Cybernetics - 11-13 CAL, 11 Artificial Intelligence and Robotics. Goddard Space Flight Centre, Greenbelt, Maryland, USA

ALBUS, K. (1976). On continuous representation in cat's Area 1. In Max-Planck Institut fur Geschicte, Afferent and Intrinsic organisation of laminated structures in the brain. Edited by O. Creutzfeldt. Spirnger-Verlag, Berlin, Heidelber, New York, pp.362-365

ANDERSEN, P. (1974). Physiological mechanism of barbiturate spindle activity. In Basic Sleep Mechanisms. Edited by O. Petre-Quadens and J.D. Schlag. Academic Press, pp.127-141

ANDERSEN, P. and ANDERSON, S.A. (1968). Physiological Basis of the Alpha Rhythm. Appleton-Century-Crofts, N.Y.

ANDERSEN, P. and ECCLES, J. (1962). Inhibitory phasing of neuronal discharge. Nature, 196:645-647

ANDERSEN, P. and LOMO, T. (1970). Model of control of hippocampal pyramidal cell discharges. In The Neural Control of Behaviour. Edited by R.E. Whalen and others (Conference, 1968). Academic Press, pp.5-26

ANOKHIN, P.K. (1961). A new conception of the physiological architecture of the conditioned reflex. In Brain Mechanisms and Learning. Symp. Council for Inter. Org. of Med. Sciences. Edited by J.F. Delafresnaye. Blackwell, Oxford

ANSTIS, S.M. (1978). Apparent movement. Chapter 21 in Perception, Vol. VIII Handbook of Sensory Physiology. Edited by R. Held, H.W. Leibowitz and H.L. Teuber. Springer Verlag, pp.655-673

ARBIG, M.A. and KILMER, L. (1976). Internal models and maps in laminated structures. In Afferent and Intrinsic organisation of laminated structures in the brain. Edited by O. Creutzfeldt, Max-Planck Institut fur Geschichte, Springer-Verlag, pp.477-483

ASANUMA, H. (1975). Recent developments inthe study of the columnar arrangement of neurons within the motor cortex. Physiol. Rev., 55, No.2, April 1975, 143-156

ASANUMA, H. and ROSEN, I. (1972). Topographical organisation of cortical efferent zones projection to distal forelimb muscles in the monkey. Exptl. Brain Res., 14:243-256

ASANUMA, H. and SAKATA, H. (1967). Functional organisation of a cortical efferent system examined with focal depth stimulation in cats. J. Neurophysiol., 30:35-54

ASHBY, W.R. (1952). Design for a Brain. Chapman and Hall, London

BAERENDS, G.P. and KRUIKT, J.P. (1973). Stimulus selection. In Constraints on Learning. Edited by R.A. Hinde and J. Stevenson-Hinde. Academic Press, pp.23-49

BAIZER, J.S., ROBINSON, D.L. and DOW, B.M. (1977). Visual responses of Area 18 neurons in awake behaving monkey. J. Neurophysiol., 40, No.5, Sept. 1977, 1024-1037

BARLOW, H.B. and BRINDLEY, G.S. (1963). Interocular transfer of movement after-effects during pressure blinding of the stimulated eye. Nature, No. 4913, Dec. 28, 1347

BARLOW, H.B. and HILL, R.M. (1963). Evidence for a physiological explanation of the waterfall phenomenon and figural after-effects Nature, No. 4913, Dec. 28, 1345-1347

BARONDES, S.H. (editor) (1976). Neuronal Recognition. Chapman and Hall

BARR, R.L. (1974). The Human Nervous System. Harper International, Hagerstown, USA

BATESON, P.P.G. (1973). Internal influences on early learning in birds. In Constraints on Learning. Edited by R.A. Hinde and J. Stevenson-Hinde. Academic Press, pp.100-116

von BEKESY, G. (1967). Sensory Inhibition. Princeton Univ. Press, N.Y., USA

BEURLE, R.L. (1956). Properties of a mass of cells capable of regenerating waves. Phil. Trans. Royal Soc. Vol. 240

BEURLE, R.L. (1959). Storage and manipulation of information in the brain. J. I. Elec. Eng. Feb, 75-82

BLEDSOE, W.W. and BROWNING, I. (1959). Pattern recognition and reading by machine. Proceedings of the Eastern Joint Computer Conference, 1959, 225-232. Reproduced in Pattern Recognition. Edited by L. Uhr. John Wiley, 1966, pp.301-316

BRADLEY, P.H. (1968). The pharmacology of synapses in the central nervous system. In Recent Advances in Pharmacology. Edited by J.M. Robson and R.S. Stacey. J. & A. Churchill Ltd., London

BRAZIER, M.A.B. (1977). Electrical Activity of the Nervous System. Pitman Medical, Tunbridge Wells, England

BRIGGS, M.H. and KITTO, G.B. (1962). The molecular basis of memory and learning. Psychol. Rev. 69:537-541

BRINDLEY, G.S. (1967). The classification of modifiable synapses and their use in models for conditioning. Proc. Royal Soc. B. 168:361-376

BRINDLY, G.S. (1969). Nerve nets of plausible size that perform many simple learning tasks. Proc. Roy. Soc. of London (Biol). 174:173-191

BROOKS, B. and JUNG, R. (1973). Neuronal physiology of the visual cortex. In Visual Centres in the Brain. Part B of Vol. VII/3 of Handbook of Sensory Physiology. Springer-Verlag, pp.325-440

BUCHWALD, J.S. and BROWN, K.A. (1973). Subcortical mechanisms of behavioural plasticity. In Efferent Organisation and the Integration of Behaviour. Edited by J. Maser. Academic Press, pp.100-136

BURKE, W. (1966). Neuronal Models for conditioned reflexes. Nature, 210, No. 5033, April 16, 1966, 269

BURNS, B.D., HERON, W. and PRITCHARD, R. (1962). Physiological excitation of visual cortex in cat's unanaesthetized isolated forebrain. J. Neurophysiol., 25:165-181

BUSWELL, G.T. (1935). How people look at pictures. Chicago Univ. Press, 1935

CAJAL, R. (1911). Histologie au systeme nerveux. Maloine, Paris

CARPENTER, M.B. (1973). Comparison of the efferent projections of the Globus Pallidus and the Substantia Nigra in the monkey. In Efferent Organisation and the Integration of Behaviour. Edited by J. Maser. Academic Press, pp.137-174

CARPENTER, R.H.S. (1977). Movements of the eyes. Pion, London NW2

CELESIA, G.G. (1976). Organisation of auditory cortical areas in man. Brain, 99:400-414

CHOUDBURY, B.P. (1978). Visual field representation in the newborn rabbit's cortex. Brain Res., 153:27-37

CHOW, K.L. (1970). Integrative functions of the thalamocortical visual system of cat. In The Biology of Memory. Edited by K.H. Pribram and D. Broadbent. Academic Press, N.Y., pp.273-293

COLONNIER, M. (1968). Synaptic patterns on different cell types in the different laminae of the cat visual cortex. An electron microscope study. Brain Res., 9:268-287

CORCORAN, D.W.J. (1971). Pattern Recognition. Penguin Science of Behaviour, general editor B.M. Foss. Penguin Books

COWEY, A. (1964). Projection of the retina on to striate and prestriate cortex in the squirrel monkey, Saimiri sciureus. J. Neurophysiol., 27:366-393

CRAGG, G.G. (1975). Density of synapses and neurons in normal, mentally defective and ageing human brains. Brain, 98:81-90

CRAIK, K.J.W. (1943). The Nature of Explanation. Cambridge Univ. Press

CREUTZFELDT, O.D., WATANABE, S. and LUX, H.D. (1966). Relations between EEG phenomena and potentials of single cortical cells II. Spontaneous and convulsoid activity. Electroencephalogr. clin. Neurophysiol., 20:19-37

CREUTZFELDT, O.D., GRUNEWALD, G., SIMONOVA, O. and SCHMITZ, H. (1969). Changes of the basic rhythms of the EEG during performance of mental and visuomotor tasks. In Attention and Neurophysiology. Edited by C.R. Evans and T.B. Mulholland. Butterworths, London, pp.148-168

DANIEL, P.M. and WHITTERIDGE, D. The representation of the visual field on the cerebral cortex in monkeys. J. Physiol. (London), 159:203-221

DEMENT, W.C. and MITLER, M. (1974). An introduction to sleep. In Basic Sleep Mechanisms. Edited by O. Petre-Quadens and J.D. Schlag. Academic Press, pp.272-296

DEMPSEY, E.W. and MORISON, R.S. (1942). The production of rhythmically recurrent potentials after localised thalamic stimulation. American J. Physiol., 135:293-300

DESCARTES, R. (1949). The passions of the soul.

DESIRAJU, T., BROGGI, G., PRELEVIC, S., and PURPURA, D.P. (1969). Inhibitory synaptic pathways linking specific and non-specific thalamic nuclei. Brain Res., 15:542-543

DEUTSCH, J.A. (1960). The structural basis of behaviour. Cambridge U.P.

DEUTSCH, J.A. (1960). The plexiform zone and shape recognition in the Ocotpus. Nature, Vol. 185:221-9

DEUTSCH, J.A. (1973). The cholinergic synapse and the site of memory. Chapter 3 in The Physiological Basis of Memory. Edited by J.A. Deutsch. Academic Press, pp.59-77

DEUTSCH, J.A. and KOOPMANS, H.S. (1976). Hunger and thirst. In Biological Foundations of Psychiatry. Edited by R.G. Grenell and S. Gabay. Raven Press, N.Y., pp.297-319

DODWELL, P.C. (1964). A coupled system for coding and learning in shape discrimination. Psychol. Rev., 71:148-159

DODWELL, P.C. (1970). Visual Pattern Recognition. Holt, Rinehart and Winston

DOTY, R.W. (1973). Ablation of visual areas in the central nervous system. In Central Processing of Visual Information, Vol. VII/3, Part B, of Handbook of Sensory Physiology. Edited by R. Jung. Springer-Verlag, pp.483-541

DOUGLAS, R.J. (1972). Pavlovian conditioning and the brain. In Inhibition and Learning. Edited by R.A. Boakes and M.S. Halliday. Academic Press, London, pp.529-554

DOUGLAS, R.J. and PRIBRAM, K.H. (1966). Learning and limbic lesions. Neuropsychologia, 4:197-220

DRAGER, U.C. and HUBEL, D.H. (1976). Topography of visual and somatosensory projections to mouse Superior Colliculus. J. Neurophysiol., 39:91-101

ECCLES, J.C. (1972). Possible synaptic mechanisms subserving learning. In Brain and Behaviour. Edited by A.G. Karczmar and J.C. Eccles. Springer-Verlag, pp.39-61

ECCLES, J.C. (1978). An instruction-selection hypothesis of cerebral learning. In Cerebral Correlates of Conscious Experience. Inserm. Symp. No. 6, Edited by Buser and Rougeul-Buser. Elsevier, North Holland Biomedical Press, pp.155-175

ESTES, W.K. (1959). The statistical approach to learning theory. In Psychology, a study of a science: Study 1, Vols. 1 and 2. Edited by S. Koch. McGraw-Hill

ETIENNE, A.S. (1973). Developmental stages and cognitive structures as determinants of what is learned. In Constraints on Learning. Edited by R.A. Hinde and J. Stevenson-Hinde. Academic Press, pp.371-395

FEENEY, D.M., PITTMAN, J.C. and WAGNER II H.R. (1974). Lateral inhibition and attention: comments on the neuropsychological theory of Walley and Weiden. Psychol. Rev., 81, No.6:536-539

FERSTER, C.B. and SKINNER, B.F. (1957). Schedules of reinforcement. Appleton-Century-Crofts, N.Y.

FITZSIMMONS, J.T. (1972). Thirst. Physiol. Reviews, 52:468-561

FLYNN, J.P. (1976). Neural basis of threat and attack. In Biological Foundations of Psychiatry. Edited by R.G. Grenell and S. Gabay. Raven Press, N.Y., pp.273-295

FREUND, H.J. (1973). Neuronal mechanisms of the lateral geniculate body. In Central Processing of Visual Information, Vol. VII/3, Part B, of Handbook of Sensory Physiology. Edited by R. Jung. Springer-Verlag, pp.177-246

GAITO, J. (1961). A biochemical approach to learning and memory. Psychol. Review, 68:288-292

GALAMBOS, R., NORTON, T.T. and FROMMER, G.P. (1967). Optic tract lesions sparing pattern vision in cats. Exper. Neurology, 18:8-25

GAZE, R.M. (1970). The Formation of Nerve Connections. Academic Press.

GIBSON, E.J. (1969). Principles of Perceptual Learning and development. Appleton-Century-Crofts, N.Y.

GIBSON, J.J. (1968). The senses considered as perceptual systems. George Allen and Unwin

GIBSON, J.J. and GIBSON, E.J. (1955a). Perceptual learning - differentiation or enrichment? Psychol. Review, 162:32-41

GIBSON, J.J. and GIBSON, E.J. (1955b). What is learned in perceptual learning? A reply to Professor Postman. Psychol. Review, 62, No.6, 446-450

GILULA, N.B. (1975). Junctional membrane structure. In The Nervous System. Editor-in-Chief Donald B. Tower. Vol. 1, The Basic Neurosciences, Raven Press, N.Y., pp.1-11

GLICKMAN, S.E. (1973). Responses and reinforcement. In Constraints on learning. Edited by R.A. Hinde and J. Stevenson-Hinde. Academinc Press, pp.207-242

GLICKMAN, S.E. and SCHIFF, B.B. (1967). A biological theory of reinforcement. Psychol. Review, 74:81-109

GOLDMAN, P.S. and NAUTA, W.J.H. (1977). Columnar distribution of cortico-cortical fibres in the frontal association, limbic and motor cortex of developing Rhesus Monkey. Brain Res., 122:393-413, Research Reports

GRANIT, R. (1955). Receptors and Sensory Perception. Yale Univ. Press, USA

GRAY, J.A.B. and LAL, S. (1965). Effects of mechanical and thermal stimulation of cats' pads on the excitability of dorsal horn neurones. J. Physiol., 179:154-162

GREY WALTER see WALTER, G.

GRIFFITH, J.S. (1967). A View of the Brain. Oxford Univ. Press

GRIFFITH, J.S. (1971). Mathematical Neurobiology. Academic Press

GROVES, P.M. and THOMPSON, R.F. (1973). A dual process theory of habituation: neural mechanisms. Chapter 6 in Habituation. Edited by H.V.S. Peake and M.J. Herz. Vol. II, Physiological substrates. Academic Press, N.Y. and Lond., pp.175-205

GUTHRIE, E.R. (1935). The Psychology of Learning. Harper, N.Y.

GUZMAN, A. (1969). Decomposition of a visual scene into three-dimensional bodies. In Automatic Interpretation and Classification of Images. Edited by A. Grasselli. Academic Press

HAITH, M.M. (1978). Visual competence in early infancy. In Perception, Vol. VIII of Handbook of Sensory Physiology. Edited by R. Held, H. Leibowitz and H.L. Teuber. Springer-Verlag, pp.319-356

HALL, J.F. (1966). The Psychology of Learning. Lippincott & Co.

HASSLER, R. (1978). Interaction of reticular activating system for vigilance and the trunco-thalamic and pallidal system for directing awareness and attention. In Cerebral Correlates of Conscious Experience. INSERM Symposium No. 6. Edited by Buser and Rougeul-Buser. Elsevier/North Holland Biomedical Press, pp.111-130

HEBB, D.O. (1949). The organisation of behaviour. John Wiley, N.Y.

HEBB, D.O. (1955). Drives and the CNS (Conceptual Nervous System). Psychol. Review, 62:243-54

van HEERDEN, P.J. (1963). A new method of storing and retrieving information. Applied Optics. 2:287-392

HENDRICKSON, A.E. (1972). An integrated molar/molecular model of the brain. Psychological Reports, 30:343-368

HENRY, G.H. and BISHOP, P.O. (1971). Simple cells of the striate cortex. In Contributions to sensory physiology. Edited by W.D. Neff. Academic Press, N.Y., pp.1-46

HENRY, G.H., LUND, J.S. and HARVEY, A.R. (1978). Cells of the striate cortex projecting to the Clare-Bishop area of the cat. Brain Res., 151:154-158

HINDE, R.A. (1960). Energy models of motivation. Symposium. Soc. for Exper. Biol. Vol. 14:199-213. Cambridge Univ. Press

HINDE, R.A. (1970). Animal Behaviour. McGraw-Hill

HINDE, R.A. (1973). Constraints on learning: an introduction to the problem. In Constraints on Learning. Edited by R.A. Hinde and J. Stevenson-Hinde, Academic Press, pp.1-20

HOGAN, J.A. (1973). How young chicks learn to recognise food. In Constraints on Learning. Edited by R.A. Hinde and J. Stevenson-Hinde. Academic Press, pp.119-139

HORRIDGE, G.A. (1968). The origins of the nervous system. In The Structure and Function of Nervous Tissue. Edited by Bourne. Methuen, pp.1-31

HORRIDGE, G.A. (1968). Interneurons. Freeman and Company, London and San Francisco

HOUTMAN, S.D. (1974). Learning and the spiral illusion. Br. J. Psychol., 65,2:105-211

HUBEL, D.H. (1963). The visual cortex of the brain. Scientific American, Nov. 173:226-228

HUBEL, D.H. and WIESEL, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. J. Physiol. (London), 148:574-591

HUBEL, D.H. and WIESEL, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J. Physiol. (London), 160:106-154

HUBEL, D.H. and WIESEL, T.N. (1963). Shape and arrangement of columns in cat's striate cortex. J. Physiol. (London), 165:559-568

HUBEL, D.H. and WIESEL, T.N. (1965). Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat. J. Neurophysiol., 28:229-289

HUBEL, D.H. and WIESEL, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. J. Physiol. (London), 195:215-243

HUBEL, D.H. and WIESEL, T.N. (1972). Laminar and columnar distribution of geniculo-cortical fibres in the Macaque Monkey. J. Comp. Neurol., 146:421-450

HULL, C.L. (1937). Mind, mechanism and adaptive behaviour. Psychol. Review, 42:491-516

HULL, C.L. (1951). Principles of Behaviour. Appleton-Century-Crofts. N.Y. (1943)

HYDEN, H. (1961). Satellite cells in the nervous system. Scientific American, 205:62-70

ISAACSON, R.L. (1972). Neural systems of the limbic brain and behavioural inhibition. In Inhibition and Learning. Edited by R.A. Boakes and M.S. Halliday, Academic Press, N.Y.

ISAACSON, R.L. (1974). The Limbic System. Plenum Press, New York

IVERSEN, S.D. (1973). Brain lesions and memory in animals. Chapter 9 in The Physiological Basis of Memory. Edited by J.A. Deutsch. Academic Press, pp.305-364

IVERSEN, L.I. (1979). The chemistry of the Brain. Scientific American, Sept. 118-129

IWAMURA, Y. and TANAKA, M. (1978a). Functional organisation of receptive fields in the cat somatosensory cortex: I: Integration within the coronal region. Brain Res., 151:49-60

IWAMURA, Y. and TANAKA, M. (1978b). Funtional organisation of receptive fields in the cat somatosensory cortex: II: Second representation of the forepaw in the ansate region. Brain Res., 151:61-72

JACOBS, B.L., MOSKO, S.S and TRULSON, M.E. (1977). The investigation of the role of serotonin in mammalian behaviour. In Neurobiology of Sleep and Memory. Edited by R.R. Drucker-Colin and J.L. McGaugh. Academic Press, pp.99-133

JACOBSON, M. (1976). Neuronal recognition in retinotectal system. In Neuronal Recognition. Edited by S.H. Barondes. Chapman and Hall, pp.3-23

JENKINS, HM. (1973). Effects of the stimulus-reinforcer relation on selected and unselected responses. In Constraints on Learning. Edited by R.A. Hinde and J. Stevenson-Hinde. Academic Press, pp.189-206

JONES, E.G. (1975a). Lamination and differential distribution of thalamic afferents within the sensory-motor cortex of the Squirrel Monkey. J. Comp. Neurol., 160:167-204

JONES, E.G. (1975b). Varieties and distribution of non-pyramidal cells in somatic cortex of the Squirrel Monkey. J. Comp. Neurol., 160:205-268

JONES, E.G., BURTON, H. and PORTER, R. (1975). Commissural and cortico-cortical 'columns' in the somatic sensory cortex of primates. Science, 190:572-574

JOUVET, M. (1972). Some monoaminergic mechanisms controlling sleep and waking. In Brain and Human Behaviour. Edited by A.G. Karczmar and J.C. Eccles. Springer-Verlag, pp.131-164

JOUVET, M. (1974). The role of monoaminergic neurons in the regulation and function of sleep. In Basic Sleep Mechanisms. Edited by O. Petre-Quadens and J.D. Schlag. Academic Press, pp.107-236

JUNG, R. (1961). Neuronal integration in the visual cortex and its significance for visual information. In Sensory Communication. Edited by W.A. Rosenblith. MIT Press, pp.627-674

JUNG, R. (1969). Neuronal integration in the visual cortex. In Perception and Action, Brain and Behaviour 2. Edited by K.H. Pribram. Penguin Books, pp.14-46

JUNG, R. (1972). Neurophysiological and Psychophysical Correlates in Vision Research. In Brain and Human Behaviour. Edited by A.G. Karczmar and J.C. Eccles. Springer-Verlag, pp.209-258

JUNG, R. and GAUMGARTNER, G. (1955). (cited in Jung, 1969). Hemmungsmechanismen und bremsende Stabilisierung an einzelnen Neuronem des Optischen Cortex. Pflugers Arch. Ges. Physiol. Vol. 261 pp.434-356

KANDEL, E.R. (1976). Cellular Basis of Behaviour, an Introduction to behavioural neurobiology. W.H. Freeman & Co. San Francisco

KILLACKEY, H.P., BELFORD, G., RYUGO, R. and RYUGO, D.K. (1976). Anomalous organisation of thalamocortical projections consequent to vibrissae removal in the newborn rat and mouse. Brain res., 104:309-315

KILMER, W. (1975). Biology of decisionary learning mechanisms in mammalian CA3-Hippocampos. A review. Int. J. Man-Machine Studies, 7:413-437

KILMER, W.L., McCULLOCH, W.S. and BLUM, J. (in collaboration with GRAIGHILL, E. and PETERSON, D.) (1968). Towards a theory of reticular formation. In The Mind. Edited by W.C. Corning and M. Balaban. John Wiley, pp.177-232

KILMER, W. and OLINSKI, M. (1974). Model of a plausible learning scheme for CA3 Hippocampus. Kybernetic, 16:133-144 (Springer-Verlag)

KINSTON, W.J., BADAS, M.A. and BISHOP, P.O. (1969). Multiple projection of the visual field to the medial portion of the dorsal lateral geniculate nucleus and the adjacent nuclei of the thalamus of the cat. J. Comp. Neurol., 136:295-316

KOHLER, W. (1942). Dynamics in Psychology. Kaber and Faber

KONORSKI, J. (1967). Integrative Activity of the Brain. Univ. Chicago Press

LANDAUER, R.K. (1964). Two hypotheses concerning the biochemical basis of memory. Psychol. Review, 71 No.3:167-179

LANDFIELD, P.W. (1976). Synchronous EEG rhythms: their nature and their possible functions in memory, information transmission and behaviour. In Molecular and Functional Neurobiology. Edited by W.H. Gispen. Elsevier Scientific Publishing Company, pp.390-424

LASHLEY, K.S. (1942). The problem of cerebral organisation in vision. In Biol. Symposia, Vol. VII, Visual Mechanisms. Edited by H. Kluever. Jacques Cattell Press, Lancaster, pp.301-322. Reproduced in Perception and Action, Brain and Behaviour 2. Edited by K.H. Pribram. Penguin Modern Psychology, pp.235-258

LASHLEY, K.S. (1949). In Search of the Engram. Exper. Biol. Symposia IV, pp.455-481. Reproduced in Brain Physiology and Psychology. Edited by C.R. Evans and A.D.J. Robertson, Butterworths, 1966, pp.1-32

LASHLEY, K.S., CHOW, K.L. and SEMMES, J. (1951). An examination of the electrical field theory of cerebral integration. Psychol. Review, 58:123-126

LEE, B.B., HEGGELUND, P., HULME, M. and CREUTZFELDT, O.D. (1976). Representation of orientation in vertical cortical penetrations in Area 17 of the Cat. In Max-Planck Institut fur Geschichte Afferent and Intrinsic Organisation of laminated structures in the brain. Edited by O. Cretuzfeldt, Springer-Verlag. (Author: Max-Planck Institut), pp.366-369

LELORD, G. (1966). Etude chez l'animal et chez l'homme d'un mode d'association specifique distinct du condionnement Classique, l'acquisition libre. Acta Biol. Exper. (Warsaw), Vol.20, No.4 pp.379-405

LETTVIN, J.Y., MATURANA, H.R., McCULLOCH, W.S. and PITTS, W.H. (1959, 1966). What the frog's eye tells the frog's brain. Proc. Inst. Radio Engineers, 47:1940-1951. Reproduced in Brain Physiology and Psychology. Edited by C.R. Evans and A.D.J. Robertson, Butterworths, 1966, pp.95-122

LIBET, B. (1978). Neuronal vs. subjective timing for a conscious sensory experience. In Cerebral Correlates of Conscious Experience. Edited by P.A. Buser and A. Rougeul-Buser. Inserm Symposium No.6. Elsevier/North Holland Biomedical Press, pp.69-82

LIVINGSTON, R.B. (1976). Sensory processing, perception and behaviour. In Biological Foundations of Psychiatry. Edited by R.G. Grenell and S. Gabay. Raven Press, N.Y., pp.47-143

LORENTE de NOR. (1949). Cerebral Cortex: architecture, intracortical connections, motor projections. In Chapter XV of Physiology of the Nervous System by Fulton, J.F., Oxford Univ. Press

LORENZ, K. (1957, 1952). The past twelve years in the comparative study of behaviour. Zoological Convention, Freiberg, 1952. In Instinctive Behaviour. Edited by C.H. Schiller, Methuen

LUND, J.S. (1973). Organisation of Neurons in the Visual Cortex, Area 17, of the monkey (Macacca mulatta). J. Comp. Neurol., 147:455-596

LURIA, A.R. (1966). Human Brain and Psychological Processes. Harper and Row

LURIA, A.R. (1973). The Working Brain. Penguin Education. Reprinted 1976

MACH, E. (1959). The analysis of sensation and the relation of the physical to the psychical. Dover Publications, N.Y.

MACKAY, D.M. (1956). The place of meaning in the theory of information. In Symposium on Information Theory. Edited by C. Cherry, Butterworths, London

MACKAY, D.M. (1965). A mind's eye view of the brain. In Progress in Brain Research, Vol. 17, Cybernetics of the Nervous System. Edited by N. Wiener and J.P. Schade. Elsevier Publishing Co. pp.321-332

MACKAY, D.M. (1978). The dynamics of perception. In Cerebral Correlates of Conscious Experience. Edited by P.A. Buser and A. Rougeul-Buser. Elsevier/North Holland Biomedical Press, pp.53-68

MACKINTOSH, J.J. (1974). The psychology of animal learning. Academic Press

McCULLOCH, W.S. and PITTS, W. (1943). A logical calculus of the ideas immanent in nervous activity. Bull. Math. Biophysics, 5:115-132

McILWAIN, H. (1979). Intracellular synaptic mediators and the endogenous simulation of neural input to the brain. In Brain Mechanisms in memory and learning: from the single neuron to Man. Edited by M.A.B. Brazier, Raben Press, N.Y., pp.71-78

MAGOUN, H.W. (1954). The ascending reticular system and wakefulness. In Brain Mechanisms and Consciousness. The Council for Inter. Org. Med. Sciences. Edited by Delafresnaye, Blackwell.

MAGOUN, H.W. (1958). The Waking Brain. Thomas, Illinois, USA

MARGULES, D.L. and MARGULES, A.S. (1973). The development of operant responses by noradrenergic activation and cholinergic suppression of movements. In Efferent Organisation and the Integration of Behaviour. Edited by J. Maser, Academic Press, pp.203-228

MARK, R. (1974). Memory and Nerve Cell Connections. Clarendon Press, Oxford

MARR, D. (1969). A theory of cerebellar cortex. J. Physiol., 202:437-470

MARR, D. (1970). A theory for the cerebral Neo-cortex. Proc. Royal Soc. B., 262:23-81

MASER, J.D. (1973). Efferent Response Processes: relationships amongst Stimuli, Movements and Reinforcement. In Efferent Organisation and the Integration of Behaviour. Edited by J. Maser, Academic Press, pp.1-19

MASON, J.W. (1975). Emotion as reflected in patterns of endocrine integration. In Emotions - Their Parameters and Measurement. Edited by L. Levi, Raven Press, N.Y.

MILLER, G.A., GALANTER, E. and PRIBRAM, K.H. (1960, 1969). Plans and the structure of behaviour. Holt, Rinehart and Winston, N.Y. Part reproduced in Perception and Action, Brain and Behaviour 2. Edited by K.H. Pribram. Penguin Modern Psychology, pp.541-556.

MILLER, S. and KONORSKI, J. (1928). Sur une forme particuliere des reflexes conditionels. C.R. Seance. Soc. Biol., 99:1155-7

MILNER, P.M. (1970). Physiological Psychology. Holt, Rinehart and Winston, N.Y.

MILNER, P.M. (1974). A model for visual shape recognition. Psychol. Review, 81, No.6:521-535

MINSKY, M.L. (1958). Some methods of artificial intelligence and heuristic programming. In Mechanisation of Thought Processes, National Physical Laboratory, H.M.S.O.

MINSKY, M.L. (1968). Semantic Information Processing, MIT Press

MINSKY, M. and PAPERT, S. (1969). Perceptrons, MIT Press

MOGENSON, G.J. and PHILLIPS, A.G. (1976). Physiological substrates of motivation. Progress in Psychobiology and Physiological Psychology. Edited by J.M. Sprague and A.N. Epstein. Vol. 6. Academic Press, pp.189-244

MOHLER, C.W. and WURTZ, R.H. (1977). Role of striate cortex and Superior Colliculus in visual guidance of saccadic eye movements in monkeys. J. Neurophysiol. Vol. 40, No.1:74-94

MOORE, B.R. (1973). The role of directed Pavlovian reactions in simple instrumental learning in the pigeon. In Constraints on Learning. Edited by R.A. Hinde and J. Stevenson-Hinde, Academic Press, pp.159-187

MORISON, R.S. and BASSETT, D.L. (1945). Electrical activity of the thalamus and basal ganglia in decorticate cats. J. Neurophysiol., 8:309-314

MORISON, R.S. and DEMPSEY, E.W. (1942). A study of thalamo-cortical relations. American J. Physiol., 135:281-292

MORRELL, F. (1967). Electrical signs of sensory coding. In The Neurosciences: A study program. Edited by G.C. Quarton, T. Melnechuk and F.O. Schmitt. Rockeffer Univ. Press, pp.577-602

MORRELL, F. (1972). Integrative properties of parastriate neurons. In Brain and Human Behaviour. Edited by A.G. Karczmar and J.C. Eccles. Springer-Verlag, pp.259-289

MOULTON, D.G. (1976). Spatial patterning of response to odours in the peripheral olfactory system. Physiological Reviews 56: Current topics in Physiology, p.578

MOUNTCASTLE, V.B. (1957). Modality and topographic properties of single neurons of cat's somatic sensory cortex. J. Neurophysiol., 20:408-434

MOUNTCASTLE, V.B. (1978). Some neural mechanisms for directed attention. In Cerebral Correlates of Conscious Experience. Edited by P.A. Buser and A. Rougeul-Buser. Inserm Symposium No.6. Elsevier/North Holland Biomedical Press, pp.37-51

MOWRER, O.H. (1947). On the dual nature of learning - a reinterpretation of "conditioning" and "problem-solving". Harvard Educ. Review, 17:pp.102-148

MULHOLLAND, T.B. (1969). The concept of attention and the electroencephalographic alpha rhythm. In Attention in Neurophysiology. Edited by C.R. Evans and T.B. Mulholland, Butterworths, pp.100-127

NATHANSON, J.A. and GREENGARD, P. (1977). Second Messengers in the brain. Scientific American. August, 108-119

NEISSER, U. (1967). Cognitive Psychology. Appleton-Century-Crofts, N.Y.

NEISSER, U. (1976). Cognition and Reality. W.H. Freeman and Company, San Francisco

OATLEY, K. (1970). Brain Mechanisms and Motivation. Nature, Vol.225, Feb.28, 797-801

OATLEY, K. (1978). Perceptions and representations. Methuen, London

OLAVARRIA, J. and TORREALBA, F. (1978). The effect of acute lesions of the striate cortex on the retinotopic organisation of the lateral peristriate cortex in the rat. Brain Res., 151:386-391

OLDS, J. (1969). The Central Nervous System and the Reinforcement of Behaviour. American Psychologist, 24:114-132

OLDS, J. (1976). Behavioural studies of the Hypothalamic Functions. In Biological Foundations of Psychiatry. Edited by R.G. Grenell and S. Gabay, Raven Press, N.Y.

PALAY, S.L. (1972). The Structural Basis for Neural Action. In Brain Function. Edited by

PALAY, S.L. (1975). An essay on Neurocytology. In The Nervous System. Editor-in-chief D.B. Tower. Vol.1 The Basic Neurosciences, Raven Press, N.Y.

PENFIELD, W. and RASMUSSEN, T. (1952). The cerebral cortex of Man. Macmillan, N.Y.

PENFIELD, W. and ROBERTS, L. (1959). Speech and Brain Mechanisms. Oxford Univ. Press

PIAGET, J. (1969). The Mechanisms of Perception. Routledge and Kegan Paul

PIAGET, J. (1954). The Construction of Reality: the Child. Basic Books, N.Y.

PICKERSGILL, M.J. and JEEVES, M.A. (1964). The origin of the after-effect of movement. Quart. J. Exp. Psychol., 4:90-103

PIERREL, R. and SHERMAN, R.J. (1963). Barnabus, the rat with college training. The Brown Alumni Monthly, Brown University

PITTS, W. and McCULLOCH, W.S. (1947). How we know Universals. Bull. Math. Biophysics, Vol.9:127-147

POSTMAN, L. (1955). Association theory and perceptual learning. Psychol. Review, Vol.62, No.6:438-445

POWELL, T.P.S. and MOUNTCASTLE, V.B.M. (1959). Some aspects of the functional organisation of the post-central gyrus of the monkey. Bull. John Hopkins Hosp., 105:133-162

PREMACK, D. (1971). Catching up with common sense or two sides of a generalisation: Reinforcement and Punishment. Chapter 5 in The Nature of Reinforcement. Edited by R. Glaser. Academic Press, pp.121-150

PRIBRAM, K.H. (1971). Languages of the Brain. Prentice-Hall, N.J. USA

PRIBRAM, K., NUWER, M. and BARON, R. (1974). The Holographic Hypothesis of Memory Structure in Brain Function and Perception. In Contemporary Developments in Mathematical Psychology. Vol. VII Measurement, Psychophysics and neural information processing. Edited by D.H. Krantz, R.C. Atkinson, R.D. Luce and P. Suppes. W.H. Freeman and Company, San Francisco, pp.416-457

PRINGLE, J.W.S. (1951). On the parallel between learning and evolution. Behaviour, 3:174-215

PURPURA, D. (1967). Comparative physiology of dendrites. In The Neurosciences: a study program. Edited by G.C. Quarton, T. Malnechuk and F.O. Schmitt. Rockefeller Univ. Press, N.Y., pp.372-393

PURPURA, D (1970). Operations and processes in thalamic and synaptically related neural systems. In The Neurosciences: Second Study Program. Edited by G.C. Quarton, T. Malnechuk and G. Adelman. Rockefeller Univ. Press, pp.458-470

PURPURA, D.P. (1974). Intracellular studies of thalamic synaptic mechanisms in evoked synchronisation and desynchronisation of electrocortical activity. In Basic Sleep Mechanisms. Edited by O. Petre-Quadens and J.D. Schlag. Academic Press, pp.99-126

PURPURA, D.P. and SHOFER, R.J. (1963). Intracellular recording from thalamic neurons during reticulo-cortical activation. J. Neurophysiol., 26:494-505

RATLIFFE, F. and HARTLINE, H.K. (1959). The response of Limusus optic nerve fibres to patterns of illumination on the receptor mosaic. J. Gen. Physiol., 42:1241-1255

RAUSCH, R. (1977). Cognitive strategies in patients with unilateral temporal lobe excisions. Neuropsychologia, 15:385-395

REESE, T.S. and SHEPHERD, G.M. (1972). Dendro-dendritic synapses in the central nervous system. In Structure and Function of Synapses. Edited by G.D. Pappas and D.P. Purpura. Raven Press, N.Y.

RESCORLA, R.A. and SOLOMON, R.L. (1967). Two process learning theory: relationship between Pavlovian Conditioning and Instrumental Learning. Psychol. Review, Vol. 74, No.3. May, 151-182

ROBERTS, L.G. (1965). Machine perception of three-dimensional solids. In Optical and Electro-optical information processing. Edited by I.J.T. Tippett. MIT Press, Cambridge, Mass. USA

ROCHESTER, N., HOLLAND, J.H., HAIBT, L.H. and DUDA, W.L. (1956). Tests on a cell assembly theory of the action of the brain, using a large digital computer. PGIT 2. No.3, 80

ROSE, S. and HAYWARDS, J. (1977). Experience, learning and brain metabolism. Chapter 8 in Biochemical Correlates of Brain Structure and Function

ROSENBLATT, F. (1958a). The Perceptron. Psychol. Review, Vol.65, No.6:386-408

ROSENBLATT, F. (1958b). Two theorems of statistical separability. Proc. Symp. Mechanisation of Thought Processes. National Physical Laboratory, HMSO

ROSENBLATT, F. (1960). Perceptron simulation experiments. Proc. Inst. Radio Engs, 301

ROUTTENBERG, A. and HAING-JA MING (1978). The substantia nigra and neostriatum: substrates for memory consolidation. In Cholinergic-Monoaminergic interactions in the brain. Academic Press, pp.305-331

SCHEIBEL, A.B. and SCHEIBEL, M.E. (1962). Interhemispheric relations and cerebral dominance. In Conference on Interhemispheric Relations and Cerebral Dominance. Edited by V.B. Mountcastle, Johns Hopkins, Baltimore

SCHEIBEL, M.E. and SCHEIBEL, A.B. (1967). Structural organisation of non-specific thalamic nuclei and their projection towards cortex. Brain Res., 6:60-94

SCHEIBEL, M.E. and SCHEIBEL, A.B. (1970). Elementary processes in selected thalamic and cortical sub-systems - the structural substrates. In The Neurosciences: Second Study Program. Edited by G.C. Quarton, T. Melnechuk and G. Adelman. Editor-in-chief, F.O. Schmitt. Rockefeller Univ. Press, 443-457

SCHILDER, P. (1966). Loss of brightness discrimination in the cat following removal of the striate area. J. Neurophysiol., 19:888-897

SCHILDER, P., PASIK, P. and PASIK, T. (1967). Total luminous flux: a possible response determinant for the normal monkey. Science, 158:806-809

SCHILLER, C.H. (Editor) (1957). Instinctive Behaviour. Methuen.

SCHMITT, O., DEV, P. and SMITH, B.H. (1976). Electronic processing of information by brain cells. Science, July, 114-120

SCHLAG, J.D. and CHAILLET, F. (1963). Thalamic mechanisms involved in cortical desynchronisation and recruiting responses. Electroencephalogr. clin. Neurophysiol., 15:39-62

SCOTT, A.C. (1977). Neurophysics. John Wiley.

SEKULER, R.W. and GANZ, L. (1963). After-effect of seen motion with a stabilised retinal image. Science, February, 419-420

SEKULER, R., PANTLE, A. and LEVINSON, E. (1978). Physiological basis of motion perception. Chapter 3 in Perception, Vol. VIII of Handbook of Sensory Physiology. Edited by R. Held, H.W. Leibowitz, and Hans-Lukas Teuber. Springer-Verlag, pp.75-96

SELFRIDGE, O.G. (1958). Pandemonium: a paradigm for learning. Proc. Symp. Mechanisation of Thought Processes. H.M.S.O. (reprinted in Pattern Recognition. Edited by L. Uhr, Wiley, N.Y., pp.339-348)

SHANKS, M.F., ROCKEL, A.J. and POWELL, T.P.S. (1975). The commissural fiber connections of the primary somatic sensory cotex. Neuroscience Abstracts, 1:126

SHEFFIELD, F.D. (1966). Induction Theory of Reinforcement. In Current Research in Motivation. Edited by R.N. Haber. Holt, Rinehart and Winston, N.Y.

SHEFFIELD, F.D. and CAMPBELL, B.A. (1954). The role of experience in the "spontaneous" activity of hungry rats. J. Comp. Physiol. Psychol., 47:97-100

SHIK, M.L. and ORLOVSKY, G.N. (1976). Neurophysiology of locomotor automatism. Physiological Reviews, 56, July

SHOLL, D.A. (1956). The organisation of the cerebral cortex. Methuen

SIGG, E.B. (1971). Central sympathetic nervous system. In Emotions: their Parameters and Measurement. Edited by L. Levi. Raven Press, N.Y.

SINGER, W. and TRETTER, F. (1976). Receptive field properties and neuronal connectivity in striate and parastriate cortex of contour deprived cats. J. Neurophysiol. Vol. 39, No.3, May

SKINNER, B.F. (1938). The behaviour of organisms. Appleton-Century-Crofts, N.Y.

SOKOLOV, E.N. (1960). Neuronal models and the orienting reflex. In The Central Nervous System and Behaviour. Third Conf. Edited by M.A.B. Brazier. Josiah Macey Jr., pp.187-212

SPERRY, R.W. (1958). Physiological plasticity and brain circuit theory. In Biological and Biochemical Bases of Behaviour. Edited by Harlow and Woolsey. Univ. of Wisconsin.

SPERRY, R.W. (1959). The Growth of Nerve Circuits. Scientific American, Nov.

SPERRY, R.W. and MINER, N. (1955). Pattern perception following insertion of mica plates into visual cortex. J. Comp. Physiol. Psychol., 48:463-469

SPERRY, R.W., MINER, N. and MYERS, R.E. (1955). Visual pattern perception following sub-pial splicing and tantalum wire implantations. J. Comp. Physiol. Psychol., 48:50-58

SPIGEL, I.M. (1962). Contour absence in the inhibition of the decay of a movement after-effect. J. Psychol., 54:221-228

SPIGEL, I.E. (1965). Readings in the study of visually perceived movement. Harper and Row, N.Y.

SPINELLI, D.N. (1970). OCCAM, a context addressable memory model for the brain. In The Biology of Memory. Edited by K.H. Pribram and D. Broadbent. Academic Press, N.Y.

SPRAGUE, J.M., BELLUCCHI, G. and RIZZOLATTI, G. (1973). The role of the Superior Colliculus and Pretectum in vision and visually guided behaviour. Chapter 2 in Central Processing of Visual Information, Vol. VII/3 B of Handbook of Sensory Physiology. Edited by R. Jung. Springer-Verlag, pp.27-101

STANLEY, J.C. and KILMER, W.L. (1975). A wave model of temporal sequence learning. Int. J. Man Machine Studies, 7:395-412

STEIN, B.E., MAGALHAES-CASTRO, B. and KRUGER, L. (1976). Relationship between visual and tactile representations in Cat Superior Colliculus. J. Neurophysiol., 39:No.2, March, 401-479

STEVENS, C.F. (1979). The neuron. Scientific American. Sept. Vol.241 No.3, 48-59

SUTHERLAND, N.S. (1957). Visual discrimination of orientation and shape by Octopus. Nature, 179:11-13

SUTHERLAND, N.S. (1958). Visual discrimination of shape by Octopus, squares and triangles. Quart. J. Exper. Psychol., 10:40-47

SUTHERLAND, N.S. (1959). Stimulus analysing mechanisms. In Proc. Symp. Mechanisation of Thought Processes. National Physical Laboratory, HMSO

SUTHERLAND, N.S. (1961). Discrimination of Horizontal and Vertical extents by Octopus. J. Comp. Physiol. Psychol., 54:43-48

SUTHERLAND, N.S. (1961). Figural after-effects and apparent size. Quart. J. Exper. Psychol., 13:222-228

SUTHERLAND, N.S. (1968). Outlines of a theory of visual pattern recognition in animals and man. Proc. Roy. Soc. B, 171:297-317

SZENTAGOTHAI, J. (1967). The anatomy of complex integrative units in the nervous system. In Results in Neuroanatomy, Neurochemistry, Neuropharmacology and Neurophysiology. Recent developments of neurobiology in Hungary I. Budapest, Akademiai, Kiado, pp.9-45

SZENTAGOTHAI, J. (1973). Synaptology of the Visual Cortex. In Central Processing of Visual Information, Vol. VII/3 B of Handbook of Sensory Physiology. Edited by R. Jung. Springer-Verlag, pp.269-324

SZENTAGOTHAI, J. (1975). The "module concept" in cerebral cortex architecture. Brain Res., 95:475-496

SZENTAGOTHAI, J. (1976). Basic circuitry of the Neocortex. In Afferent and Intrinsic Organisation of Laminated Structures in the Brain. Edited by O. Creutzfeldt. Max-Planck Institut fur Geschichte. Springer-Verlag, pp.282-287

SZENTAGOTHAI, J. (1978). The local neuronal apparatus of the cerebral cortex. In Cerebral Correlates of Conscious Experience. Edited by P.A. Buser and A. Rougeul-Buser. Elsevier/North Holland Biomedical Press, pp.131-138

TALBOT, S.A. (1942). cited in Marshall W.H. and Talbot, S.A. Recent evidence for neural mechanisms in vision leading to a general theory of sensory acuity. In Visual Mechanisms. Edited by H. Kluver, Lancaster: Jacques Cattell, pp.117-139. Reprinted in Pattern Recognition. Edited by L. Uhr, John Wiley, 1966, pp.195-211

THORPE, W.H. (1963). Learning and Instinct in Animals. Methuen

TINBERGEN, N. (1951). The Study of Instinct. Oxford Univ. Press

TOLMAN, E.C. (1932). Purposive Behaviour in Animals and Man. Appleton-Century-Crofts.

TOLMAN, E.C. (1959). Principles of purposive behaviour. In Psychology, a study of a science. Edited by S. Koch. McGraw-Hill, N.Y.

TOWE, A.L. (1973). Motor Cortex and the pyramidal system. In Efferent Organisation and the Integration of Behaviour. Edited by J. Maser. Academic Press, pp.67-99

UHR, L. (1966). Pattern Recognition. John Wiley, N.Y.

UHR, L. (1973). Pattern Recognition, Learning and Thought. Prentice-Hall, N.Y., USA

UTTAL, W.R. (1973). The Psychobiology of Sensory Coding. Harper and Row, N.Y.

UTTLEY, A.M. (1954). The classification of signals in the nervous system. Electroencephalogr. clin. Neurophysiol.,6:479

UTTLEY, A.M. (1956a). Conditional probability machines and conditioned reflexes. In Automata Studies. Edited by Shannon and McCarthy. No.34. Princeton Univ. Press, pp.253-274

UTTLEY, A.M. (1956b). Temporal and spatial patterns in a Conditional Probability Machine. In Automata Studies. Edited by Shannon and McCarthy. No.34. Princeton Univ. Press, pp.277-285

UTTLEY, A.M. (1959). Conditional Probability Computing in a nervous system. In Symp. Mechanisation of Thought Processes. National Physical Laboratory, HMSO

UTTLEY, A.M. (1961). The engineering approach to the problem of neural integration. In Progress in Biophysics and Biopnysical Chemistry 11. Edited by J.A.V. Butler, B. Katz and R.E. Zirkle, Pergamon Press, pp.25-32

VALENSTEIN, E.S. (1970). Stability and plasticity of motivation systems. In The Neurosciences: Second Study Program. Edited by G.C. Quarton, T. Melnechuk and G. Adelman. Editor-in-chief, F.O. Schmitt. Rockefeller Univ. Press

VALVERDE, F. (1971). Short axon neuronal sub-systems in the visual cortex of monkey. International J. Neuroscience, 1:181-197

VAN DER LOOS, H. and WOOLSEY, T.A. (1973). Somatosensory cortex: structural alterations following early injury to sense organs. Science, 179:395-397

VANDERWOLF, C.H., BLAND, B.H. and WHISHAW, I.Q. (1973). Diencephalic, Hippocampal and Neocortical mechanisms in voluntary movement. In Efferent Organisation and the Integration of Behaviour. Edited by J.D. Maser. Academic Press, pp.229-263

VARGA, M.Y. and PRESSMAN, J.M. (1963) (cited in Konorski, 1967). Some forms of relationship between two temporarily connected reflexes. In Proc. Conf. Liblice, May 1961. Publ. House of Czech. Acad. Sci. Edited by E. Gurman, pp.279-284

VERZEANO, M. (1970). Evoked responses and network dynamics. In The Neural Control of Behaviour. Edited by R. Whalen, R.F. Thompson, M. Verzeano and N.M. Weinberger. Academic Press, pp.27-54

VERZEANO, M. (1977). The activity of neuronal networks in memory consolidation. In The neurobiology of sleep and memory. Edited by R. Drucker-Colin and J.L. McGaugh. Academic Press, pp.75-97

WALLEY, R.E. and WEIDEN, T.D. (1973). Lateral inhibition and cognitive masking: a neurophysiological theory of attention. Psychol. Review, 80, No.4:284-302

WALTER, G.W. (1953). Theoretical properties of diffuse projection systems in relation to behaviour and consciousness. In Brain Mechanisms and Consciousness. Edited by J.F. Delafresnaye. Blackwell Oxford

WALTER, G. (1956). The Living Brain. Penguin Books, Reprinted, 1971

WALTER, G. (1969). Can "Attention" be defined in physiological terms? In Attention in Neurophysiology. Edited by C.R. Evans and T.B. Mulholland. Butterworths, pp.27-37

WELT, C., ASCHOFF, J.C., KAMEDA, K. and BROOKS, V.B. (1967). Intracortical organisation of cat's motorsensory neurons. In Neurophysiological Basis of Normal and Abnormal Motor Activities. Edited by M.D. Yahr and D.P. Purpura. Raven Press, N.Y., pp.255-294

WERNER, G. (1970). The topology of the body representation in the somatic efferent pathways. In The Neurosciences: Second Study Program. Edited by G.C. Quarton, T. Melnechuk, and G. Adelman. Editor-in-chief, F.O. Schmitt. Rockefeller Univ. Press, N.Y., pp.605-616

WERBER, G., WHITSEL, B.L. and PETRUCELLI, L.M. (1972). Data structure and algorithms in the primate somatosensory cortex. In Brain and Human Behaviour. Edited by A.G. Karczmar and J.C. Eccles. Springer-Verlag, pp.164-186

WIENER, N. (1949). Cybernetics. John Wiley, N.Y.

WIESEL, T.N., HUBEL, D.H. and LAMM, D. (1974). Autoradiographic demonstration of ocular dominance columns in the monkey striate cortex by means of trans-synaptic transport. Brain Res., 79:273-279

WILLSHAW, D.J., BUNEMAN, O.P. and LONGUET-HIGGINS, H.C. (1969). Non-holographic association memory. Nature, 222:960-962

WINOGRAD, T. (1972). Understanding natural language. Edinburgh Univ. Press

WOHLGEMUTH, A. (1911). On the after-effect of seen movement. Brit. J. Psychol., Monogr. Suppl. 1

WOOLSEY, T.A. and Van Der LOOS, H. (1970). The structural organisation of layer IV in the somatosensory region (S1) of Mouse Cerebral Cortex. Brain Res., 17:205-242

WURTZ, R.H. and MOHLER, C.W. (1976). Organisation of monkey superior colliculus: enhanced visual response of superficial layer cells. J. Neurophysiol., 39:745-765

WURTZ, R.H. and MOHLER, C.W. (1976). Enhancement of visual responses in monkey striate cortex and frontal eye fields. J. Neurophysiol., 39:766-772

YARBUS, A.L. (1967). Eye movements and vision. Plenum Press

YOUNG, J.Z. (1964). A model of the brain. Oxford Univ. Press

ZEKI, S.M. (1978a). The cortical projections of foveal striate cortex in the rhesus monkey. J. Physiol., 277:227-244

ZEKI, S.M. (1978b). The third visual complex of rhesus monkey. J. Physiol., 277:245-272

ZEKI, S.M. (1978c). Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex. J. Physiol., 277:273-290

ZUSNE, L, (1970). Visual perception of form. Academic Press

Appendix A

Monitor Printout of a computer run from an extended version of the model ON MICROFICHE

I General

This microfiche is presented as an authentication that the work described in the thesis has actually been carried out. It does not form part of the presentation of the model and need not be referred to except for purposes of authentication.

II General layout of Fiche

A. Fortran Listing of program with FTN R3 Maps	Fiche 1
	Fiche 2 (part).
B. Update Listing	Fiche 2 (part).
C. Monitor printouts during run	Fiche 2 (part).
	Fiche 3,4 & 5.

III Printout at each cycle of program

Each main cycle of the program representing an Instant consisted of two sub-cycles. The printout for each sub-cycle begins with a statement of the Instant No., the position of the notional organism in the Environmental Maze, the Trial No. and the Input Pattern No.

```
TIME   405.012
 INSTANT NO      520
MAZE:          37    9   43   37POS:   4 ALONG,   4 UP, TRIAL   NO.  10
MAZE:           0   37   37   37POS:   4 ALONG,   4 UP, TRIAL   NO.  10
IMOVE  3JMOVE  1
PRESENT POSITION:ACROSS:  4  UP:  4
IFACE=  1
 IGL10   ITR  13    PATTERN NO   47
```

A large number of variables are printed out which need not concern the reader. In particular, at each programming sub-unit to process a particular Neuronal Array, some key variables relating to tht Array are reported. These are headed Unit 1, Part 1, etc.

These Units are related to the Neuronal Arrays of the model as follows:

Sub-cortical	Sensory		Units 1 and 2
	Motor	Consummatory	Unit 3, Part 1
		Instrumental	Unit 3, Part 2
		Preconsummatory	Unit 4, Part 1
		Protective	Unit 4, Part 2
	Motivational		Units 3 and 4, *Groups 1 and 6.
Cortical	Sensory	Projection 1	Units 5 and 6
		Projection 2	Units 9 and 10
	Motor	Consummatory	Unit 7, Part 1
		Instrumental	Unit 7. Part 2
		Preconsummatory	Unit 8, Part 1
		Protective	Unit 8, Part 2
	Motivational		Units 7 and 8, *Groups 1 and 6

* A Group in the computer program approximated to a Neuronal Pool.

As illustrated below, variables reported upon include:**ALEVEL**- the feed-back criterion for Excitation control; AVEN - the mean of Excitation Counts; **DIFFAV** - the Diffuse Volley Value; AVRIB - the mean of Inhibition Counts; DINHIB - the diffuse Inhibition increment; and a number of adjusting variables associated with these functions.

```
                          ALEVEL     AVEN      DIFFAV     ADDIFF     SLOSS      AVRIB
UNIT  1  PART  1          880.00   1066.32   1300.00     200.00    1400.00    177.73   201
 ICALIN  0  IX  4  IZ  4
UNIT  1  PART  2          880.00    989.09   1000.00     200.00    1400.00       .10   163
 ICALIN  0  IX  4  IZ  3
UNIT  2  PART  1          880.00    946.88    900.00     200.00    1400.00     66.15   193
 ICALIN  0  IX  4  IZ  3
```

Response Diagnosis: Under Unit 6, Part 2, are reported the states of a number of counting variables. These indicate the notional location of signal volleys in delayed pathways. They enable the state of the motor system to be assessed and the cause of emission of any motor response to be diagnosed. The relevant printout begins DIAG.

```
UNIT  6  PART  2     880.00    854.03   1000.00    200.00    1000.00        .10   10
 AFF  6  2  0  2       1832.00    600.00   1000.00
 ORSP     2.     33.     33.     33.ODOMR     0.INMO3    1RSET     3.     0.     0.
 DIAG BW 872    2    0    1    1    0  533    6    2   13   16    5
 ICALIN  1  IX  8  IZ  0
```

State of Integrating Neurons. Once during every computer run, sample portions of the computer arrays representing Neuronal Arrays were printed out (middle of Fiche 2). Each row reports the variables in each Integrating Neuron: Constellation Label, Excitation Count, Linking Control Constant, Linking Control Variable, Threshold,

Inhibition Count, Ouput address, if any, Afferent Input. The printout of these samples from Neuronal Arrays may be identified from Figure A.1.

Constellation report. At each cycle a table was printed out similar to that described in Chapter 8 in respect of the initial set of simulations. The table was composed by sorting the Excitation Counts of all the Integrating Neurons in the system in relation to the Constellation Label of the Neuronsto which they belonged. Each row of the table represents the state, at the time of printout, of the Linked Constellation formed at the Instant shown in the first column of the table. The remaining columns show: (2) the aggregate of Excitation Counts of the Integrating Neurons belonging to that Constellation; (3) the number of Integrating Neurons remaining in that Constellation; (4) the mean Excitation Count. See below:

246	1.03	1.03	247.18	509	59690.71
247	854.89	3.14	272.08	510	84921.58
248	911.92	4.14	220.16	511	92370.45
249	656.57	2.14	306.52	512	211158.23
250	5900.93	16.14	365.56	513	153954.89
251	9419.95	25.14	374.67	514	55799.33
252	1719.99	5.14	334.50	515	55812.36
253	1009.99	3.14	321.45	516	61005.68
254	593.67	1.14	582.63	517	44868.88
255	5790.02	9.14	633.34	518	49952.82
256	3992.25	7.14	558.98	519	47389.55
257	867.30	2.14	404.90	520	38315.17
258	1966.20	6.14	320.12	521	100946.49
259	1.03	1.03	488.03	522	54469.71
260	3428.12	5.14	666.69	523	37935.28
261	5534.14	8.14	679.70	524	39622.32
262	609.10	1.14	538.00	525	30884.68
263	414.67	1.14	450.36	526	8809.30

DISCHARGE MATRIX FOR INSTANT NO 526

Discharge Matrix. At each cycle a table was printed out showing the aggregate of Spikes emitted by Neuron within each Neuronal Pool of the model. These were equal to the Spike Frequencies of the respective Efferent Fibres:

1.11	929.31			524	19897.89	20.13	985.88		

INSTANT NO 524

1.0	0.0	0.0	0.0	0.0	2506.7	1502.4	8279.6	UNIT	1	SEC 1
5.1	6140.3	7088.9	0.0	0.0	0.0	0.0	0.0	UNIT	1	SEC 2
1.0	6432.7	1963.2	0.0	0.0	1280.5	0.0	0.0	UNIT	2	SEC 1
1.2	4928.5	0.0	0.0	2020.6	0.0	0.0	2020.5	UNIT	2	SEC 2
1.0	6.1	.0	0.0	0.0	0.0	0.0	0.0	UNIT	3	SEC 1
1.0	1922.4	.1	0.0	.1	.1	.1	.1	UNIT	3	SEC 2
1.5	0.0	.0	0.0	0.0	0.0	0.0	0.0	UNIT	4	SEC 1
1.0	0.0	.0	1195.2	.0	.0	.0	.0	UNIT	4	SEC 2
1.0	0.0	0.0	0.0	0.0	19580.4	19575.4	19291.3	UNIT	5	SEC 1
5.4	27106.9	24600.4	0.0	176.6	0.0	373.9	103.0	UNIT	5	SEC 2
1.5	24523.5	27088.5	9.1	71.4	23694.7	18.0	88.1	UNIT	6	SEC 1
1.3	25253.8	97.1	4.0	23467.3	5.4	96.0	24635.4	UNIT	6	SEC 2
1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	UNIT	7	SEC 1
1.0	0.0	0.0	0.0	19.3	1884.2	0.0	0.0	UNIT	7	SEC 2
1.5	366.5	366.5	0.0	699.2	366.5	366.5	366.5	UNIT	8	SEC 1
1.0	3.8	0.0	0.0	0.0	0.0	0.0	0.0	UNIT	8	SEC 2
1.0	0.0	0.0	0.0	0.0	17659.2	17908.5	18299.4	UNIT	9	SEC 1
1.0	18567.3	15668.4	0.0	0.0	0.0	0.0	0.0	UNIT	9	SEC 2
1.0	17285.7	19865.2	0.0	0.0	15542.3	0.0	0.0	UNIT	10	SEC 1

Control Centre (Thalamus). At each cycle the notional Active state or Inactive state of each Fibre leaving the Control Centre for a Neuronal Pool in one or other of the Neuronal Arrays was reported in the following matrix.

MDDIN 150.00 IPOS TTX 712 0

AM	0	1	0	1	1	0	0	1	0	0	BM	0	0	1	0	0	0	1	0	1	1
AM	0	0	0	0	0	0	1	1	1	1	BM	0	0	0	0	1	0	1	1	1	0
AM	0	0	0	0	0	0	0	0	0	0	i	0	0	0	0	0	0	0	0	0	0
AM	0	0	0	0	0	0	0	0	0	0	BM	0	0	0	0	0	2	0	0	0	0
.M	0	0	1	1	0	0	1	0	0	1	BM	0	1	0	1	1	0	0	1	0	0
AM	0	1	0	1	0	0	0	1	0	1	BM	0	0	0	1	0	0	1	1	0	1
AM	0	0	0	0	0	0	0	0	0	0	BM	0	0	0	0	0	0	0	0	0	0
AM	0	0	0	0	0	0	0	0	0	0	BM	0	0	0	0	0	0	0	0	0	0
AM	0	0	1	1	0	0	1	0	0	1	BM	0	1	0	1	1	0	0	1	0	0
AM	0	1	0	1	0	0	0	1	0	1	BM	0	0	0	1	0	0	1	1	0	1
AM	0	0	0	0	0	0	0	0	0	0	BM	0	0	0	0	0	0	0	0	0	0
AM	0	0	0	0	0	0	0	0	0	0	BM	0	0	0	0	0	0	0	0	0	0
ICH	2																				

IV Summary reports.

At the end of each 15 Instants of the run, a summary report was output. (Figure A.2). The chief rows indicate at each successive half-cycle the occurrence of key Active Volleys. The key to the respective rows is : LL PATTERN - Input Pattern to Subcortical Sensory Array; UPP PATT - Input Pattern to delayed projection Cortical Sensory Array; POSITION - location in maze; INSTRUMENT - Identifying number of Neuronal Pool in Cortical Instrumental Array to which an Active Volley was sent. CONSUMM.- ditto in respect of Consummatory Array. PRECONSUM - ditto in respect of Preconsummatory Array; NOT DO - ditto in respect of Instrumental Suppression Array; NOTPREC - ditto in respect of Preconsummatory Suppression Array; NOT CONS - ditto in respect of Consummatory Suppression Array; STANDSTILL - transmission of Active Volley to Standstill response Neuronal Pool. IREL - a counting control variable. ICH - the state of the Motor Release Mechanism; LEARNED PREC RESP. etc. The Anticipatory Response Indication detected in each of the Motor Cortical Arrays indicated.

\- - - - - - - - - - - - - - - - - -

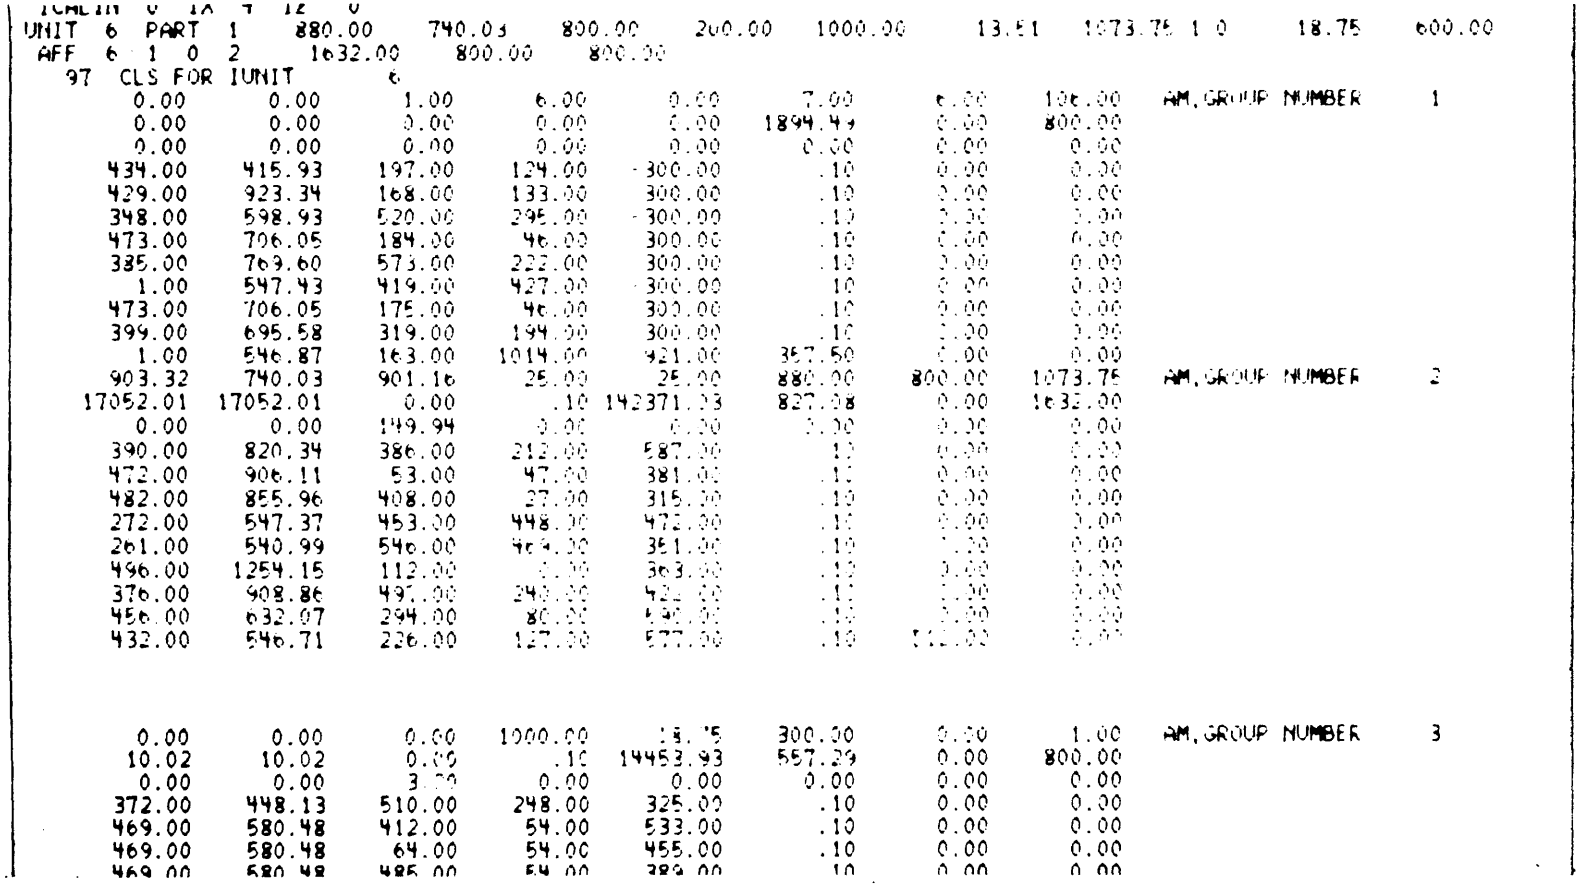

```
 ICHLIN  0  1A  7  12  0
UNIT  6  PART  1     880.00     740.03     800.00     200.00    1000.00      13.51    1073.75 1 0      18.75     600.00
 AFF  6  1  0  2      1632.00     800.00     800.00
   97  CLS FOR IUNIT        6
        0.00       0.00       1.00       6.00       0.00       7.00       0.00     106.00  AM,GROUP NUMBER      1
        0.00       0.00       0.00       0.00       0.00    1894.49       0.00     800.00
        0.00       0.00       0.00       0.00       0.00       0.00       0.00       0.00
      434.00     415.93     197.00     124.00    -300.00        .10       0.00       0.00
      429.00     923.34     168.00     133.00     300.00        .10       0.00       0.00
      348.00     598.93     520.00     295.00    -300.00        .10       0.00       0.00
      473.00     706.05     184.00      46.00     300.00        .10       0.00       0.00
      385.00     769.60     573.00     222.00     300.00        .10       0.00       0.00
        1.00     547.43     419.00     427.00    -300.00        .10       0.00       0.00
      473.00     706.05     175.00      46.00     300.00        .10       0.00       0.00
      399.00     695.58     319.00     194.00     300.00        .10       0.00       0.00
        1.00     546.87     163.00    1014.00     421.00     357.50       0.00       0.00
      903.32     740.03     901.16      25.00      25.00     880.00     800.00    1073.75  AM,GROUP NUMBER      2
    17052.01   17052.01       0.00        .10  142371.33     827.08       0.00    1632.00
        0.00       0.00     149.94       0.00       0.00       0.00       0.00       0.00
      390.00     820.34     386.00     212.00     587.00        .10       0.00       0.00
      472.00     906.11      53.00      47.00     381.00        .10       0.00       0.00
      482.00     855.96     408.00      27.00     315.00        .10       0.00       0.00
      272.00     547.37     453.00     448.00     472.00        .10       0.00       0.00
      261.00     540.99     546.00     464.00     351.00        .10       0.00       0.00
      496.00    1254.15     112.00       0.00     363.00        .10       0.00       0.00
      376.00     908.86     497.00     240.00     422.00        .10       0.00       0.00
      456.00     632.07     294.00      80.00     490.00        .10       0.00       0.00
      432.00     546.71     226.00     127.00     577.00        .10     112.00       0.00



        0.00       0.00       0.00    1000.00      15.75     300.00       0.00       1.00  AM,GROUP NUMBER      3
       10.02      10.02       0.00        .10   14453.93     557.29       0.00     800.00
        0.00       0.00       3.00       0.00       0.00       0.00       0.00       0.00
      372.00     448.13     510.00     248.00     325.00        .10       0.00       0.00
      469.00     580.48     412.00      54.00     533.00        .10       0.00       0.00
      469.00     580.48      64.00      54.00     455.00        .10       0.00       0.00
      469.00     580.48     485.00      54.00     389.00        .10       0.00       0.00
```

Figure A.1 Part of the monitor printout showing the variables within Neuronal Arrays representing Integrating Neurons.

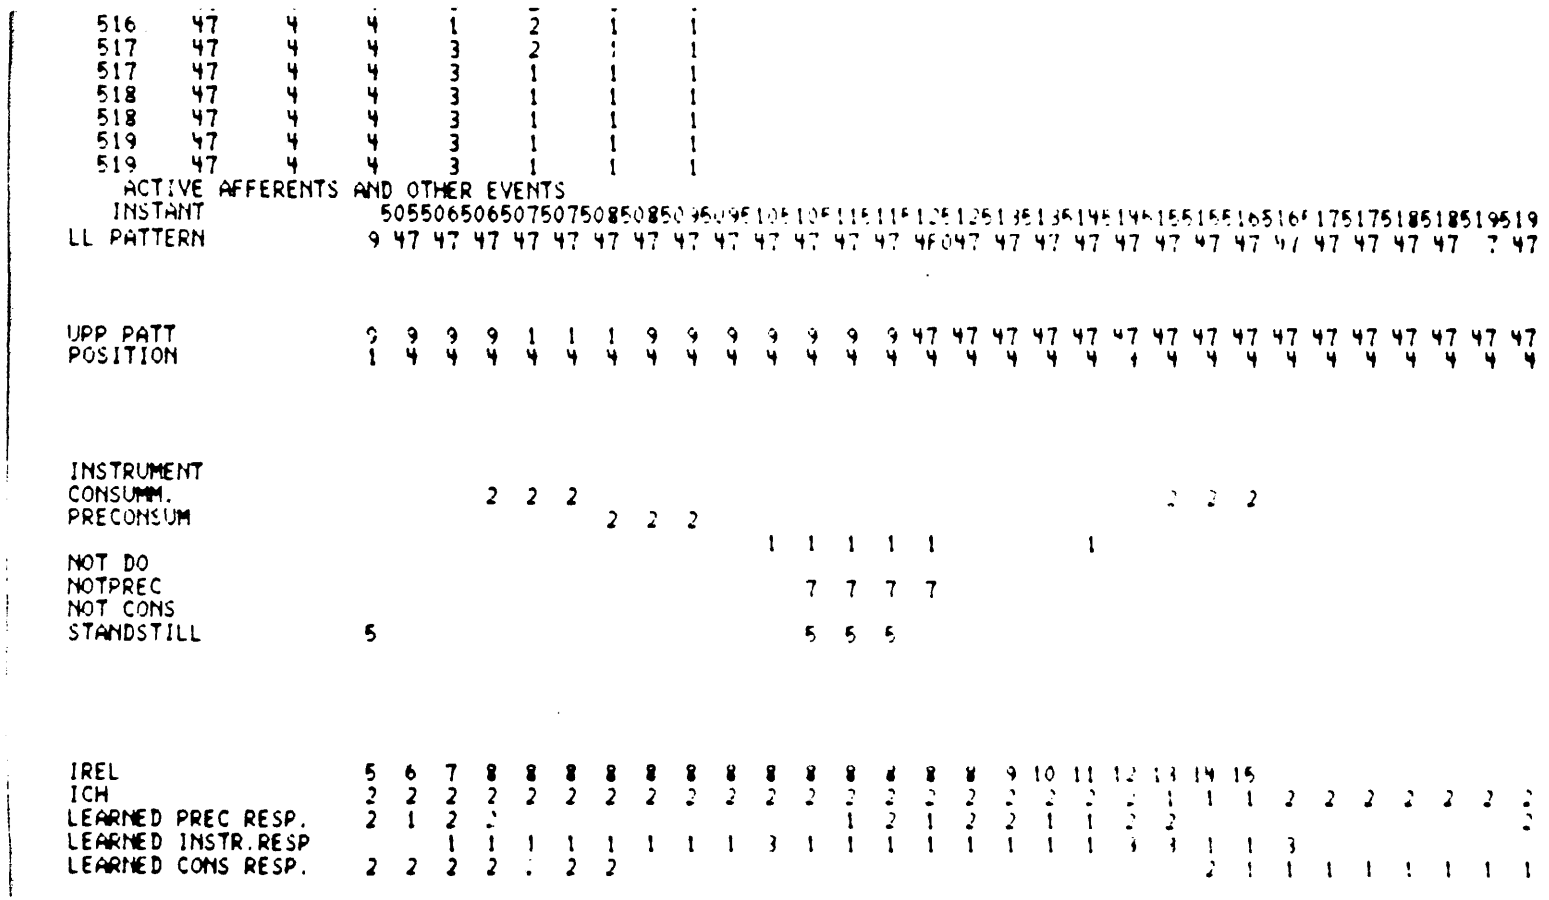

Figure A.2 Summary of performance printed out every 15 Instants. The main information shown is the transmission of Afferent Volleys between the Control Centre (Thalamus) and key Cortical Arrays. The Anticipatory Motor Response Indications are also shown.

Appendix B

Monitor Printout of a computer run illustrating the separate simulations investigating mediation of Linked Constellations by orthodox synaptic connections.

I. General

As in the case of the Microfiche presented in Appendix A, the Microfiche presented here is intended only to authenticate that the work described in the thesis was actually carried out. The monitor printout for the simulations investigating the mediation of Linked Constellations by orthodox synapses was much simpler than that for the main simulations.

II General layout

A. Fortran Listing of program

B. Monitor printouts.

III Main output at each sub-cycle.

In these programs in order to improve the fineness of the grain of iterative representation of the physiology of the model each Instant consisted of upwards of ten subcycles of the main loop of the program. At each sub-cycle the main printout consists of a matrix (illustrate in Figures 16.1 and 16.2) reporting the Spikes emitted by sample neurons in each Neuronal Pool of the model. The effects of the processing of Excitation and Inhibition variables may be followed by considering these output tables at successive sub-cycles of successive Instants.

Further Monitor Printouts from various programs discussed in the text

I. General

The microfiche included in this Appendix are taken from seven different experimental runs. The general layout and information furnished in the microfiche is as described in Appendix A. In addition to Monitor Printouts relating to the runs reported in Tables in the text, further printouts are included which contain programming code relating to mechanisms discussed more informally in the thesis.

II List of sets of microfiche and contents

C1 Monitoring part of the experimental run shown in Table 14.1, p.320. This run showed increased frequency of a reinforced Response. The Sub-Cortical emission mechanism was a variety of Maturational Linked Constellation, the code for which is set out in microfiche C6 below. (See p.302).

C2 Monitoring part of the experimental run UACTO1H referred to in Table 14.4, representing discrimination learning and utilising a 'time since last emission' Sub-Cortical Motor Release Mechanism. (See p. 296 and p.321)

C3 Monitoring part of the experimental run UACTOG1 referred to in Table 14.4, p.321

C4 Monitoring part of the experimental run UACTOTA shown in Table 14.7 , p.338. This run included the Motivational Mechanisms discussed in Section IV of Chapter 12 (page 284) and further discussed at page 323 and 337.

C5 Part of the experimental run shown in Table 14.6, page 335. The program was similar to that of the run to which C4 relates, except that in this run the model acted under the influence of two different Motivational conditions, notional Hunger and Thirst, either acting at the same time or at different times (see page 334).

C6 A Fortran Source Program and related monitor printout of a run, in which the routines for a number of mechanisms experimented with earlier are still present in some form. The program includes Fortran code for a version of Response initiation by Maturational Linked Constellations (page 302), and for Protective Responses Bump and Hurt (page 259), for lateral and recurrent Inhibition (page 139).

C7 A Fortran Source Program and related monitor printout of a version of the model run recently during the proposed further developments in the model discussed in Section III, Chapter 17 (page 390). The program provided for the 'labelling' of Linked Constellations with a marker indicating their temporal association with a Preconsummatory Response and with a TOTE test of the Anticipatory Preconsummatory Response as the criterion of the positive (notional) Consolidation of the marker. This is the program which always produced a 'confirmation' regardless of validity (see page 395), the remedy of which is discussed in immediately following pages.