The confounding effect of response amplitude on MVPA performance measures

AT Smith[1], P Kosillo[1] and AL Williams[2]

[1]Department of Psychology, Royal Holloway, University of London, Egham TW20 0EX, UK

[2] Centre for Cognition & Neuroimaging, Brunel University, Uxbridge, Middlesex UB8 3PH, UK

Corresponding author:

Prof. Andy Smith
Address as above
email: a.t.smith@rhul.ac.uk
Tel: +44 (0)1784 443717
Fax: +44 (0)1784 434347

ABSTRACT

Multi-voxel pattern analysis (MVPA) is proving very powerful in the analysis of fMRI time-series data, yielding surprising sensitivity, in many different contexts, to the response characteristics of neurons in a given brain region. However, MVPA yields a metric (classification performance) that does not readily lend itself to quantitative comparisons across experimental conditions, brain regions or people. This is because performance is influenced by a number of factors other than the sensitivity of neurons to the experimental manipulation. One such factor that varies widely but has been largely ignored in MVPA studies is the amplitude of the response being decoded. In a noisy system, it is expected that measured classification performance will decline with declining response amplitude, even if the underlying neuronal specificity is constant. We document the relationship between response amplitude and classification performance in the context of orientation decoding in the visual cortex. Flickering sine gratings were presented at each of two orthogonal orientations in a block design (multivariate experiment) or an event-related design (univariate experiment). Response amplitude was manipulated by varying stimulus contrast. Orientation classification performance in retinotopically defined occipital area V1 increased approximately linearly with the logarithm of stimulus contrast. As expected, univariate response amplitude also increased with contrast. Similar results were obtained in V2, V3 and V3A. Plotting classification performance against response amplitude gave a function with a compressive non-linearity that was well fit by a power function. Knowledge of this function potentially allows adjustment of classification performance to take account of the effect of response size, making comparisons across brain areas, categories or people more meaningful.

INTRODUCTION

Multivoxel pattern analysis (MVPA) has recently been applied, in conjunction with functional magnetic resonance imaging (fMRI), to a wide variety of different issues in human sensory and cognitive neuroscience. Numerous successful studies have identified specificities of neural response properties that, at least in some cases, were not evident from univariate fMRI studies. The sensitivity of the technique is impressive in relation to traditional univariate analysis, in which sensitivity is limited by the restricted quantity of information obtainable from a single voxel.

However, MVPA yields a metric (classifier performance) that does not readily lend itself to numerical comparisons across experimental conditions, sessions, brain regions or subjects. This is because decoding performance is influenced by a number of factors other than the sensitivity of neurons to the experimental manipulation. These include some factors that can readily be controlled, such as the number of voxels used for pattern analysis and the number of stimulus repetitions in the experiment, but also others that cannot. One such factor is the amplitude of the response being decoded. In a noisy system, it is expected that measured decoding performance will increase with response amplitude (or 'effect size') if all other factors are held constant. The expected effect of amplitude is large since, in the extreme, performance must decline to near chance for very small activations, where signals are dominated by noise. Yet this factor has received little attention in the literature. The problem cannot be circumvented by assuming that response amplitude is correlated with response specificity. For example, it is easy to imagine one population of neurons that is highly responsive to visual stimuli but relatively unselective for colour, and another that responds in a highly colour-specific way but not particularly strongly. Clearly an index of colour specificity should favour the latter, but MVPA might well yield greater decoding efficiency for the former.

A number of recent studies have compared numerical decoding efficiencies across brain regions and have implicitly or explicitly taken high performance to indicate high neuronal specificity. Most such studies have conducted separate MVPA analyses in two or more different brain regions and have confined analysis and conclusions to

statements of whether performance is or is not significantly above chance in each area, with no explicit comparison across areas (Brouwer and Ee, 2007; Etzel et al., 2008; Fu et al., 2008; Haynes et al., 2007; Li et al., 2007; Mannion et al., 2009; Preston et al., 2008; Serences and Boynton, 2007; Sterzer et al., 2008). In such cases, any indication that one brain area has greater specificity for the experimental variable than another is implicit, or at least not endorsed by statistical comparisons. A few studies have gone further and made quantitative comparisons among classification efficiencies from different brain regions. For example, Ostwald et al. (2008) used MVPA to document the ability to decode the global structure of Glass patterns in various visual areas. They found a progressive increase in classification performance from V1 to LOC, the reliability of which they tested with an ANOVA. Similarly, Eger et al. (2008) presented participants with pictures of objects and compared numerical classifier performance between posterior and anterior parts of LOC, finding statistically significant superior performance for posterior regions. Beauchamp et al. (2009) decoded the anatomical location of a somatosensory stimulus in S1, S2 and MST/STP. They obtained different classification performances in different areas which they compared in order to demonstrate differences.

Comparing classifier accuracies in this way raises the potential problem that the cause of the observed difference may lie elsewhere than in differences of neuronal selectivity. Any statement that one brain region is more sensitive than another to a particular stimulus attribute assumes that equally sensitive measurements have been obtained in both areas. This may not be the case if one area responds more strongly than the other. Quantitative comparisons across different classification pairs within the same brain region (e.g. Reddy and Kanwisher, 2007) or indeed within the whole brain (e.g. Shinkareva et al., 2008) may be somewhat safer, since at least the brain tissue included will be invariant across the results compared, but again results may be confounded by differences in response amplitude. Similar considerations apply to quantitative comparisons across subject populations. For example Yoon et al. (2008) measured the ability to decode faces, objects, scenes and scrambled images in a large, object-sensitive region of the ventral occipital cortex. They compared decoding performance in schizophrenics with that in healthy controls and found a statistically reliable difference. The difference may well be real, but the interpretation of such differences may not be straightforward.

It would be highly desirable to move towards a metric that was derived from classification performance but allowed greater scope for quantitative comparisons of the kind discussed above. In this paper, we introduce the concept of amplitude-weighted classification performance. In principle, if the relationship between response amplitude and classification performance were known, it would be possible to adjust observed performance values to account for the influence of amplitude and so to make it possible to compare decoding performance across stimuli, brain regions and subject groups in a more meaningful and reliable way. As a first step, we measure this relationship in the context of responses in human visual cortex to simple visual stimuli. It has been shown (Haynes and Rees, 2005; Kamitani and Tong, 2005) that the orientation of a grating stimulus can be decoded from the pattern of responses elicited across voxels in the visual cortex. We use a similar paradigm but systematically vary stimulus contrast so as to vary response amplitude and we document and quantify the effect of this manipulation on decoding efficiency.

METHODS

Participants

Five healthy volunteers (mean age 25 years) participated. All had normal or corrected-to-normal vision and were screened according to standard MRI exclusion criteria. Local research ethics approval and written informed consent were obtained.

Data Acquisition

MRI images were obtained with a 3-Tesla Siemens Magnetom Trio scanner and either a standard Siemens 8-channel array head coil (anatomical scans) or a custom 8-channel posterior-head array coil (Stark Contrast, Erlangen, Germany) that gives improved SNR in occipital cortex (functional scans). For each participant, a high-resolution T1-weighted 3D anatomical image was acquired (modified driven-equilibrium Fourier transform, MDEFT (Deichmann et al., 2004), 176 axial slices, in-plane resolution 256 x 256, 1 mm isotropic voxels, TR = 7.92 ms, TE = 2.45 ms, flip angle = 16°, bandwidth = 195 Hz/pixel). This anatomical image was used as a

reference to which all the functional images from all experiments were co-registered. It was also used to generate flattened cortical representations of occipital cortex for use in ROI definition. MDEFT was chosen in place of standard 3D anatomical sequences because of its improved contrast between grey matter and white matter, which is beneficial for segmentation. The functional data were acquired with a gradient echo, echoplanar sequence (TR = 2000ms, 28 contiguous axial slices covering the occipital cortex, interleaved acquisition order, 3 mm isotropic voxels, FoV 192x192 mm, flip angle = 80°, TE = 32 ms, bandwidth = 1396 Hz/pixel).

Stimuli and design

Computer-generated visual stimuli were projected by an LCD projector onto a rear-projection screen located at the head end of the scanner bore. This could be seen by participants via a mirror attached to the headcoil. The mean luminance of the stimuli was approximately 1500 $cd.m^{-2}$. Three separate experiments were conducted on different days: one to allow estimation of decoding efficiency, one to allow estimation of response magnitude and one to provide an independent region of interest (ROI) for use in the analysis.

Multivariate experiment (decoding orientation)

This experiment was similar to previous studies of orientation decoding (Haynes and Rees, 2005; Kamitani and Tong, 2005) except that stimulus contrast was varied. The stimuli were counterphasing sine gratings (2 cycles/deg, 4 Hz) presented in one of two orthogonal orientations (±45 deg from vertical). Gratings were presented in a large circular window (diameter 24 deg visual angle) so as to stimulate most of the primary visual cortex. A central fixation cross was present throughout but there was no task other than fixation. Five contrasts were used, ranging from 1% to 100% in equal log steps. Contrast is defined as $(L_{max} - L_{min})/(L_{max} + L_{min})$ where L_{max} and L_{min} are the grating luminances occurring at the peak and trough respectively. The range chosen spans the visible range (absolute detection thresholds are in the region of 0.5% contrast).

A block design was employed. The block duration was 16s and the two orientations alternated between blocks, with no blank intervals. Each scan run had 31 blocks and lasted 8 min 16s. The first block was discarded, leaving 30 blocks. Stimulus contrast was constant throughout a given run. Five such scan runs were performed, one for each of the five contrasts, separated by short rest breaks. A different random order of contrasts was used for each participant. The data from the five runs were analysed separately. Each run can be regarded as a standard orientation decoding experiment performed at one contrast, and the five runs as independent repetitions of the experiment at different contrasts.

Univariate Experiment (estimating response amplitude)

Accurate estimation of response amplitude is difficult using a block design of the type described above. Estimation benefits greatly from the inclusion of baseline blocks that contain no stimulus. In order to make it possible to independently decode five pairs of stimuli (five contrasts) with data from one scan, we maximized the number of exemplars by alternating the two orientations with no baseline blocks. Thus the timecourses were essentially flat. We verified empirically that it was not possible to obtain meaningful amplitude measures with a standard univariate analysis. We therefore conducted a separate experiment to estimate the effect of cntrast on amplitude. This had an event-related design because, in our view, several factors militate towards using event-related designs for studying the effect of an independent variable on response amplitude. First, block designs are afflicted by confounding effects of adaptation during the block, whereas adaptation is minimal for brief events. Second, event-related designs more readily allow five different stimuli to be interleaved. Third, block designs are more vulnerable to contamination by low-frequency noise near the frequency of the block cycle.

The same participants were therefore scanned again using an event-related design. The same set of stimuli (5 contrasts, 2 orientations) was used but each stimulus was presented for only 2s. Stimulus events were separated by a variable inter-trial interval (range 7-13s), selected at random on each trial with a rectangular probability distribution. During the ITI, the screen was uniform apart from the fixation cross. Its luminance was the same as the space-average luminance of the gratings. Trials of the

five contrasts were intermixed within each scan run, in an order that was pseudo-random with the constraint that each contrast was presented the same number of times (six, to give 30 trials per scan run). For each trial, grating orientation was selected randomly from the two possible values (±45 deg), again with the constraint that equal numbers of each were presented within a run. Each scan run lasted an average of 6min 10s, depending on the ITIs. Eight such scans were performed, separated by short rest breaks.

Retinotopic mapping

Retinotopic mapping was performed in the same participants on a third occasion. In most cases, this had already been done in connection with other projects and so this scan was the first of the three scans. Standard methods were used to map visual field polar angle in terms of temporal response phase (Sereno et al., 1995). A counterphasing checkerboard ‘‘wedge’’ stimulus (a 24º sector) rotated clockwise at a rate of 64 s/cycle. Eight rotation cycles were completed per run. The counterphase frequency was 8 Hz. The rotating wedge covered an area 24º visual angle in diameter, as in the main experiments. Check size was scaled by eccentricity in approximate accordance with the cortical magnification factor. Participants fixated a central fixation spot throughout.

Data analysis

Data analysis was carried out with BrainVoyager QX (version 2.0.7, Brain Innovation, The Netherlands). The data for each participant was analysed separately and the final results averaged across participants. The data were first corrected for head motion by aligning each functional volume to a template volume acquired at the beginning of the session. Each timeseries was filtered with a high-pass temporal filter with cut-off at 0.01 Hz. The functonal data were accurately co-registered with the 3D anatomy. Co-registration accuracy was checked visually. A flattened representation of each hemisphere was created by segmenting and reconstructing the border between grey and white matter within each hemisphere of the MDEFT scan. The resulting surfaces were smoothed, inflated, and cut along the calcarine sulcus. Finally, the

surface was flattened and corrected for linear distortions. Functional data could then be viewed either on the sliced 3D anatomy or on the flatmap.

For retinotopic mapping, a model was fitted to the timecourse obtained with the rotating wedge stimulus. This consisted of a rectangular wave of duty cycle 24/360, reflecting the duration of stimulation at any portion of the visual field, convolved with a canonical haemodynamic response function (HRF). The phase of the fitted response was taken as an index of visual-field location, in terms of polar angle. Phases were projected onto the flattened surface as a colour overlay. Reversals of the direction of phase change across the cortical surface were taken as boundaries of visual areas. The boundary of visual areas V1-V3A were drawn by eye in each hemisphere and the ROI created was projected back onto the participant's reference anatomy, to generate a voxel cluster. The left and right V1 clusters were then combined to provide a single ROI corresponding to bilateral V1 for use in the multivariate and univariate experiments. The mean number of voxels in this combined ROI across the five participants was 505 (SD 37). Figure 1 shows a typical V1 ROI and its derivation. The analysis focussed on V1 but additional analyses were performed using data from V2, V3 and V3A for comparison.

For the multivariate experiment, the core analysis used a linear support vector machine (SVM). In *BrainVoyager*, a model, consisting of the event time convolved with a canonical HRF, is fitted separately to each event or block in a general linear model (GLM) regression analysis. The resulting beta value (effect size) is then taken as the exemplar for that trial or block. Thus, each scan run yielded 15 exemplars of each orientation. For each scan run, the first 8 volumes (one block) were discarded and the remainder of the timeseries was divided into five sections, each containing six consecutive blocks. Four of these sub-runs (80% of the data) were used for training and the fifth for testing. The multivariate analysis was performed using the bilateral V1 ROI defined in the same participant. All voxels in the ROI were included in the analysis, irrespective of whether significant activity was present in the multivariate experiment. The analysis looked at each 16s test block in turn and established whether the pattern of activity across V1 voxels better matched one orientation or the other, based on the patterns established by the training runs. Performance was defined as

percent correct decisions. The analysis was repeated five times, using a different sub-run for testing each time, and the results were averaged.

For the univariate experiment, a standard GLM analysis was conducted with one regressor for each contrast. Instances of the two orientations were treated as a single event type in each case. Thus both orientations were represented but they were not distinguished. Six additional regressors (three translation, three rotation) derived from the head movement data were also included, as was a "session regressor" modelling the baseline activity in each run. The data from all eight runs were included to give a single parameter estimate (beta) for each stimulus contrast. This was then converted to percent signal change. An estimate of baseline activity was derived from the session regressor (mean signal over the whole timecourse after modelling out the responses and effects of head movement). Stimulus-related signal change was then divided by the baseline estimate.

RESULTS

The results of the multivariate experiment are shown in Figure 2(a). This shows mean orientation classification performance in V1 as a function of stimulus contrast, which is plotted on a logarithmic scale. In line with previous studies, stimulus orientation could readily be decoded from the data, at least at high contrast. Performance improves monotonically with contrast and a straight line (in log contrast space) provides an acceptable fit to the data. Figure 2(b) shows the results of the univariate analysis for the same ROI in the same participants. Mean response amplitude is shown, expressed as percent signal change from baseline derived from the GLM, as a function of stimulus contrast. As previously shown (Buracas and Boynton, 2007), the response increases monotonically with contrast. The response is a saturating function of contrast on a linear contrast axis but is approximately a linear function of log contrast (our data show a modest expansive non-linearity).

Since it would be desirable to obtain estimates of amplitude and decoding performance from the same dataset, we attempted to decode orientation using the data from the event-related experiment. However, performance was at chance due to an

insufficiency of exemplars at each contrast. In our hands, orientation can be decoded for a single contrast based on a one-hour event-related scan, but it cannot be done five times over in a single scan unless a block design is used.

The orientation tuning bandwidth of neurons in V1 is essentially invariant with contrast (Sclar and Freeman, 1982; Skottun et al., 1987). Therefore the effect of contrast on classification performance shown in Fig 2(a) probably does not reflect variations in orientation specificity in the brain with contrast, but instead reflects a measurement problem: the confounding effect of response amplitude. In short, orientation specificity is invariant but our estimate of it varies widely with response amplitude.

Figure 3(a) plots the results of the two experiments against each other, to show the effect of response amplitude on classification performance. This function shows that classifier performance increases with response amplitude, sharply at low amplitude and then more gradually. A good fit is provided by a power function of the form:

$$P_{raw} = k * A_{resp}^{0.2}$$

where P_{raw} is uncorrected classifier performance, A_{resp} is response amplitude and k is a scaling constant. For our data, the fitted value of k is 102 but this can be expected to vary among studies. The exponent 0.2 may be more consistent. Figure 3(b) shows the same data together with the corresponding results for V2, V3 and V3A. The results are similar in all areas and the fit to the pooled data is similar to that for V1 alone.

In principle, this function enables us to adjust classification performance. If we assume that the orientation tuning bandwidth of neurons in V1 is invariant with contrast, then a transformed version of the function in Figure 2(a) that accounts for the contaminating effect of amplitude on classifier performance will be a straight line of zero slope. We could therefore establish a transformation that is based on the function in Figure 3 and transforms Figure 2(a) with such a result. The same transformation could then be applied to other data, where the underlying neuronal specificity is not constant, in order better to estimate it. However, developing a

generic correction that can be applied in other contexts would require solving several problems (see Discussion) and would need to be tested in other contexts before its use could be advocated.

Finally, the estimates of amplitude and classification performance to be compared should be derived from the same dataset. We attempted to obtain classification performance measures from the event-related data used for the amplitude estimates. However, classification performance was around chance levels for all contrasts. We attribute this to having insufficient data from each participant. MVPA with an event-related design has been implemented successfully in a few studies but it requires quite large numbers of trials (e.g. Beauchamp et al, 2009). It is not feasible to decode orientation independently for five contrasts based on data from a single one-hour scan, and it is undesirable to combine data across repeated scans because of the difficulty of ensuring that the voxels are placed in exactly the same location.

DISCUSSION

Our results show that MVPA orientation classification performance is strongly dependent on the mean amplitude of the responses being decoded. If it is assumed that orientation specificity does not change with stimulus contrast then our results indicate that classification performance cannot be taken as a straightforward index of stimulus specificity. A possible problem with this assumption is that even though the orientation bandwidth of individual neurons may be unchanged, the number of responsive neurons may increase with contrast. Different neurons have different contrast thresholds (e.g. Albrecht and Hamilton, 1982), although most are responsive over most of the range we used, and there is psychophysical evidence for high-pass contrast "channels" (Georgeson, 1985). Thus, the amount of information about orientation in a voxel could increase with no change in tuning bandwidth, simply because of recruitment of an increasing proportion of neurons as contrast increases. Ultimately, it is the amount of information in a neural population that is estimated with MVPA, not the tuning properties of the neurons. Variable recruitment of this type is very plausible at low contrasts, however it becomes less so at high contrasts, where essentially all neurons are expected to be active. Such a mechanism therefore

predicts a large effect of contrast on classification performance at low contrasts but little or none at high contrasts. In fact we see a large effect of contrast even when only high contrasts are considered (Figure 2a). Contrast-related changes in neuron recruitment could contribute to our result but are unlikely to explain it fully.

In light of our results, we argue that it is not safe to rely on precise numerical comparisons of classifier performance across stimuli, stimulus categories, brain regions or people unless either (i) it is known that the mean univariate response amplitudes are similar across the instances compared or (ii) a correction is applied to the classification performance data to take account of differences in response amplitude. This is not to say that such comparisons are completely meaningless, only that they may be inaccurate or misleading. We have outlined a general strategy for calculating a correction that could be used to give an improved index of neural specificity that can be characterised as amplitude-weighted classification performance. We have restricted ourselves to outlining a principle and we have not developed a full method for calculating adjusted performance. Although we have fitted curves to our plots, our point is fundamentally a qualitative one.

In order to implement amplitude-weighted classification performance, it is necessary to have a good estimate of amplitude. In our study, we obtained this from a separate experiment. The multivariate and univariate experiments used designs optimized for multivariate and univariate analysis, respectively. This might be unsatisfactory as a routine method of performing both types of analysis: it is both more desirable and more efficient to obtain both measures from the same dataset. This requires a design that is amenable to both types of analysis but still provides acceptable measurement efficiency for each. This is challenging because block designs pose problems for amplitude measurement (see Results) and event-related designs require many trials for successful MVPA analysis. We have found the use of a single experiment not to be feasible when attempting to decode five different pairs of stimuli in one scan, but it should be feasible when decoding only one pair.

In addition to genuine variations of response amplitude, it is also necessary to consider confounding methodological factors relating to the way in which amplitude is calculated. For instance, increasing or decreasing the size of an ROI could lead to a

change in the mean amplitude, depending on the proportion of included voxels that are strongly active. This is likely to have a bigger effect on amplitude than on classification performance, changing the quantitative relationship between the two. The problem should not be severe in our case because we used retinotopically defined visual areas. Within these, response amplitude will be broadly uniform and the mean will be broadly independent of ROI size. However, it becomes more severe if an ROI is defined by thresholding functional data, because of the arbitrary choice of statistical threshold and because statistical significance depends on the number of trials as well as the amplitude of the response to each. The way in which responses are modelled is also important, since estimated amplitude is highly dependent on the fit of the model. It may appear that this is not a major factor when comparing brain regions or participants within a single study that uses a consistent method, but variability in the waveform of the BOLD response should also be considered. There have been clear demonstrations of differences in temporal characteristics between brain regions and between subjects (e.g. Miezin et al., 2000). For a variety of reasons, developing a reliable universal method for compensating for effects of response amplitude on MVPA performance is challenging.

The use of separate experiments with different designs means that the function relating decoding performance to amplitude (Fig, 3) may not be quantitatively accurate or typical. However, this is largely a matter of scaling; the form of the function is expected to be similar however the two variables are estimated. Thus, our conclusion can be expected to generalize to other MVPA studies qualitatively, even if not quantitatively.

It is difficult to know the extent to which failure to correct for response amplitude has already led to misleading conclusions. Univariate amplitudes are generally not stated in MVPA papers and even if they were, it would be difficult to know whether the level of variability was sufficient to have caused serious distortions. Moreover, the effect of a given amplitude difference will be greater when the responses decoded are strong than when they are weak (because of the non-linearity in Figure 3). Relatively few studies involving numerical comparisons have so far been conducted and it may be that all their conclusions are valid. However, as the MVPA literature burgeons, it is likely that false conclusions will sometimes be drawn if no correction for amplitude is

applied. The danger is perhaps illustrated by existing studies of orientation selectivity in human visual cortex. Both the original studies in this field (Haynes and Rees, 2005; Kamitani and Tong, 2005) compared orientation classification performance across visual areas V1, V2 and V3 and report that performance is greatest in V1 and least in V3, the difference being marked in one case and subtle in the other. A possible interpretation (although the authors do not state it in either study) is that orientation coding is strongest in V1 and degrades as response properties become more complex in later areas. Such a phenomenon would not be in line with macaque physiology, which shows that orientation selectivity is well preserved in V2 and V3. A better explanation comes from other fMRI data. In our hands (see for example Fig. 1c of Wall and Smith, 2008), visual stimuli consistently yield large activations in human V1 while response amplitude progressively diminishes as we progress from V1 to V2, V2 toV3 and V3 toV4. The typical univariate response in V3 is little more than half that in V1. We suggest that this may be what causes the decline in orientation classification accuracy from V1 to V3 as measured with MVPA and that, in fact, underlying neuronal tuning may be unchanging, or changing in some different way through this progression. Indeed, Haynes and Rees (2005) are careful on this point, saying "we cannot exclude the possibility that V2 and V3 are weakly activated by orientation stimuli."

In conclusion, we have highlighted the clear confounding effects that neuronal response amplitudes can have on MVPA performance measures. Studies that utilise these approaches should exercise some caution in their interpretation, particularly when comparisons between performance measures are being made.

REFERENCES

Albrecht, D.G., Hamilton, D.B., 1982. Striate cortex of monkey and cat: contrast response function. Journal of Neurophysiology 48, 217-237.

Beauchamp, M.S., LaConte, S., Yasar, N., 2009. Distributed Representation of Single Touches in Somatosensory and Visual Cortex Human Brain Mapping 30, 3163-3171.

Brouwer, G.J., Ee, R.v., 2007. Visual Cortex Allows Prediction of Perceptual States during Ambiguous Structure-From-Motion Journal of Neuroscience 27, 1015–1023.

Buracas, G.T., Boynton, G.M., 2007. The Effect of Spatial Attention on Contrast Response Functions in Human Visual Cortex. Journal of Neuroscience 27, 93-97.

Deichmann, R., Schwarzbauer, C., Turner, R., 2004. Optimisation of the 3D MDEFT sequence for anatomical brain imaging: technical implications at 1.5 and 3 T. NeuroImage 21, 757-767.

Eger, E., Ashburner, J., Haynes, J.-D., Dolan, R.J., Rees, G., 2008. fMRI Activity Patterns in Human LOC Carry Information about Object Exemplars within Category Journal of Cognitive Neuroscience 20, 356-370.

Etzel, J.A., Gazzola, V., Keysers, C., 2008. Testing Simulation Theory with Cross-Modal Multivariate Classification of fMRI Data PLOS One 3, e3690.

Fu, C.H.Y., Mourao-Miranda, J., Costafreda, S.G., Khanna, A., Andre F. Marquand, Williams, S.C.R., Brammer, M.J., 2008. Pattern Classification of Sad Facial Processing: Toward the Development of Neurobiological Markers in Depression Biological Psychiatry 63, 656-662.

Georgeson, M.A., 1985. The effect of spatial adaptation on perceived contrast. Spatial Vision 1, 103-112.

Haynes, J.-D., Rees, G., 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. Nature Neuroscience 8, 686 - 691.

Haynes, J.-D., Sakai, K., Rees, G., Gilbert, S., Frith, C., Passingham, R.E., 2007. Reading Hidden Intentions in the Human Brain Current Biology 17, 323–328.

Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. Nature Neuroscience 8, 679 - 685.

Li, S., Ostwald, D., Giese, M., Kourtzi, Z., 2007. Flexible Coding for Categorical Decisions in the Human Brain Journal of Neuroscience 27, 12321–12330.

Mannion, D.J., McDonald, J.S., Clifford, C.W.G., 2009. Discrimination of the local orientation structure of spiral Glass patterns early in human visual cortex NeuroImage 46, 511-515.

Miezin, F.M., Maccotta, L., Ollinger, J.M., Petersen, S.E., Buckner, R.L., 2000. Characterizing the Hemodynamic Response: Effects of Presentation Rate, Sampling Procedure, and the Possibility of Ordering Brain Activity Based on Relative Timing. NeuroImage 11, 735-759.

Ostwald, D., Lam, J.M., Li, S., Kourtzi, Z., 2008. Neural Coding of Global Form in the Human Visual Cortex. Journal of Neurophysiology 99, 2456-2469.

Preston, T.J., Li, S., Kourtzi, Z., Welchman, A.E., 2008. Multivoxel Pattern Selectivity for Perceptually Relevant Binocular Disparities in the Human Brain. Journal of Neuroscience 28, 11315-11346.

Reddy, L., Kanwisher, N., 2007. Category Selectivity in the Ventral Visual Pathway Confers Robustness to Clutter and Diverted Attention. Current Biology 17, 2067–2072.

Sclar, G., Freeman, R.D., 1982. Orientation Selectivity in the cat's striatal cortex is invariant with stimulus contrast. Experimental Brain Research 46, 457-461.

Serences, J.T., Boynton, G.M., 2007. The Representation of Behavioral Choice for Motion in Human Visual Cortex. Journal of Neuroscience 27, 12893-12899.

Sereno, M.I., Dale, A.M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., Tootell, R.B.H., 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science 268, 889-893.

Shinkareva, S.V., Mason, R.A., Malave, V.L., Wang, W., Mitchell, T.M., Just, M.A., 2008. Using fMRI Brain Activation to Identify Cognitive States Associated with Perception of Tools and Dwellings PLOS One 3, e1394.

Skottun, B.C., Bradley, A., Sclar, G., Ohzawa, I., Freeman, R.D., 1987. The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behaviour. Journal of Neurophysiology 57, 773-785.

Sterzer, P., Haynes, J.-D., Rees, G., 2008. Fine-scale activity patterns in high-level visual areas encode the category of invisible objects. Journal of Vision 8(15):10, 1-12.

Wall, M., Smith, A., 2008. The representation of egomotion in the human brain. Current Biology 18, 191-194.

Yoon, J.H., Tamir, D., Minzenberg, M.J., Ragland, J.D., Ursu, S., Carter, C.S., 2008. Multivariate Pattern Analysis of Functional Magnetic Resonance Imaging Data Reveals Deficits in Distributed Representations in Schizophrenia Biological Psychiatry 64, 1035-1041.

FIGURE LEGENDS

Figure 1

Illustration of the V1 region of interest used in the analysis. **(a)** Inflated view of the segmented grey matter surface of the occipital cortex (medial view) in the right hemisphere of one participant. Response phases from the retinotopic mapping experiment are overlaid in colour (see key to visual field location, inset) and the location of V1 is shown by a white outline. The location of the fundus of the calcarine sulcus is marked with a pink line. **(b)** A sagittal slice through the occipital cortex showing significant ($p<0.001$ unc.) activity from the 100% contrast trials of the univariate experiment (orange and yellow). The green overlay shows voxels in the slice that fall within the V1 region of interest, outlined in white.

Figure 2

Results of the two experiments. **(a)** Multivariate orientation classification performance in primary visual cortex (V1), averaged across participants, as a function of stimulus contrast. **(b)** Univariate response amplitude, averaged across all voxels in V1 and across participants. Error bars show ±1 SEM based on the means for participants (n=5).

Figure 3

(a) Relationship between orientation classification performance and response amplitude in V1, derived from the results in Figure 2. The curve fit is a power function with exponent 0.2. **(b)** Results for three additional visual areas in the same format. The curve is fitted to the pooled data.

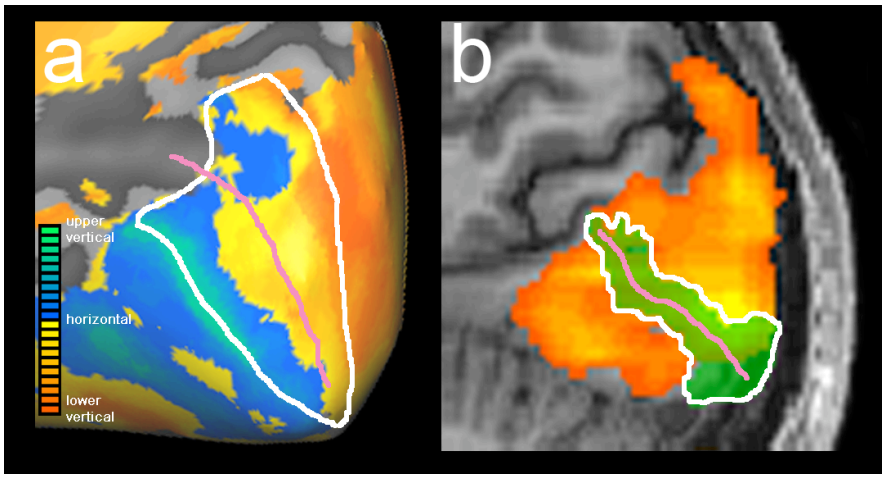

Figure1

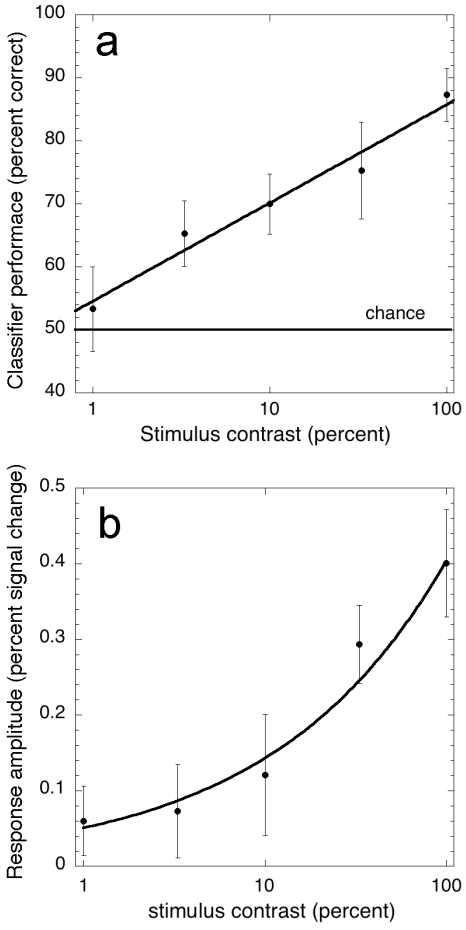

Figure 2

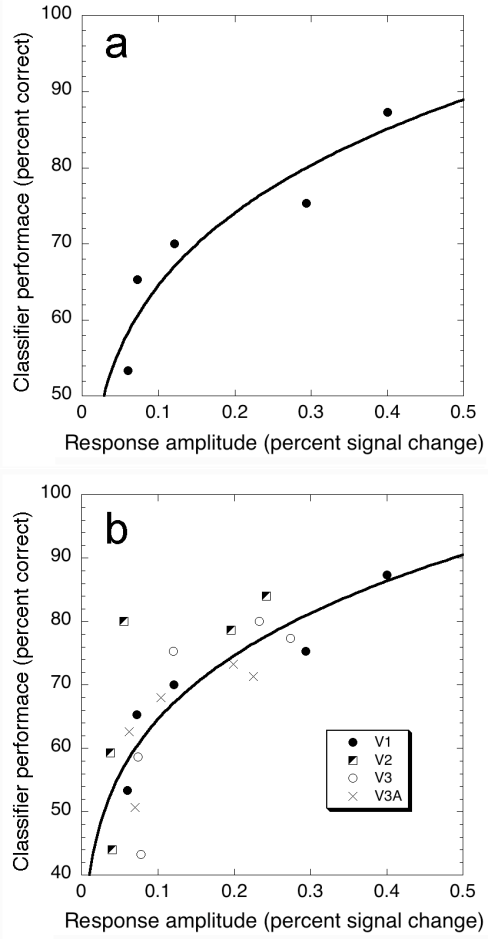

Figure 3