nature neuroscience

Predicting the orientation of invisible stimuli from activity in human primary visual cortex

John-Dylan Haynes[1,2] & Geraint Rees[1,2]

Humans can experience aftereffects from oriented stimuli that are not consciously perceived, suggesting that such stimuli receive cortical processing. Determining the physiological substrate of such effects has proven elusive owing to the low spatial resolution of conventional human neuroimaging techniques compared to the size of orientation columns in visual cortex. Here we show that even at conventional resolutions it is possible to use fMRI to obtain a direct measure of orientation-selective processing in V1. We found that many parts of V1 show subtle but reproducible biases to oriented stimuli, and that we could accumulate this information across the whole of V1 using multivariate pattern recognition. Using this information, we could then successfully predict which one of two oriented stimuli a participant was viewing, even when masking rendered that stimulus invisible. Our findings show that conventional fMRI can be used to reveal feature-selective processing in human cortex, even for invisible stimuli.

Orientation-selective aftereffects can result from exposure to grating stimuli that are too fine to be consciously perceived, suggesting that orientation-selective but unconscious activation of visual cortex is possible[1–3]. However, direct physiological measurement of such unconscious feature-selective processing in human V1 has proven elusive due to the relatively low spatial resolution of functional neuroimaging methods. Typically, human fMRI has a spatial resolution of a few millimeters. However, in the primary visual cortex of primates, neurons with different orientation preferences are systematically mapped in a pinwheel pattern, with regions of similar orientation selectivity separated by around 500 microns[4,5]. Thus, the topographic representation of orientation selectivity is below the spatial resolution of functional MRI, so direct comparison of activity evoked by differently oriented stimuli is generally considered unrevealing. Indirect methods such as prolonged selective adaptation[6] can be used to assess orientation processing in human V1, but such effects are generally weak[7] and rely on prolonged exposure to adapting stimuli.

Most functional neuroimaging studies use analysis techniques in which brain signals are averaged across space and time to improve signal quality, combining samples from different spatial positions (voxels) and averaging across many individual measurements. Averaging across space and time in this way might obscure any information that is present in the spatial pattern of responses in individual samples of brain activity. Multivariate pattern recognition[8] provides a way of taking into account any such information about stimulus orientation that may be contained in patterns of activity across the whole of primary visual cortex. Encouraged by recent reports that multivariate analyses can be used to successfully classify fMRI measurements of brain activity evoked by different object categories[9–11], and even stimulus orientation (Y. Kamitani & F. Tong, *J. Vis.* **4**, 186a, 2004), we set out to test the hypothesis that an unconscious representation of orientation might exist in human V1. We found that this analytical method could successfully distinguish different patterns of activity evoked in V1 by a brief presentation of orthogonally oriented gratings. Moreover, when the gratings were masked so that participants were unable to report their orientation correctly, we could still use brief measurements of V1 activity from individual participants to predict stimulus orientation.

RESULTS

Experiment 1

In the first experiment, we investigated whether any representation of stimulus orientation could be detected reliably in human V1 using high-field functional magnetic resonance imaging (fMRI) at a conventional spatial resolution. We measured brain activity while four participants passively viewed a clearly visible, obliquely oriented grating that was presented in an annulus around the fixation point (Methods; **Fig. 1a**). In separate blocks, the grating was oriented either 45° to the right or orthogonally, 45° to the left. For the 100 voxels with the strongest responses to both categories, we first computed how much they preferred stimuli of either orientation. Most voxels showed a slight orientation preference (**Fig. 1b**). Approximately 30% of voxels showed a significant orientation preference at an uncorrected threshold (**Fig. 1b**), and between 5% and 28% were significant even at a conservative Bonferroni-corrected level. Next we investigated the reproducibility of these weak orientation preferences. To achieve this, we divided the dataset in half and computed the orientation bias for each voxel separately for each half. Notably, a large number of voxels

[1]Wellcome Department of Imaging Neuroscience, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3BG, UK. [2]Institute of Cognitive Neuroscience, University College London, Alexandra House, 17 Queen Square, London WC1N 3AR, UK. Correspondence should be addressed to J.D.H. (haynes@fil.ion.ucl.ac.uk).

Published online 24 April 2005; doi:10.1038/nn1445

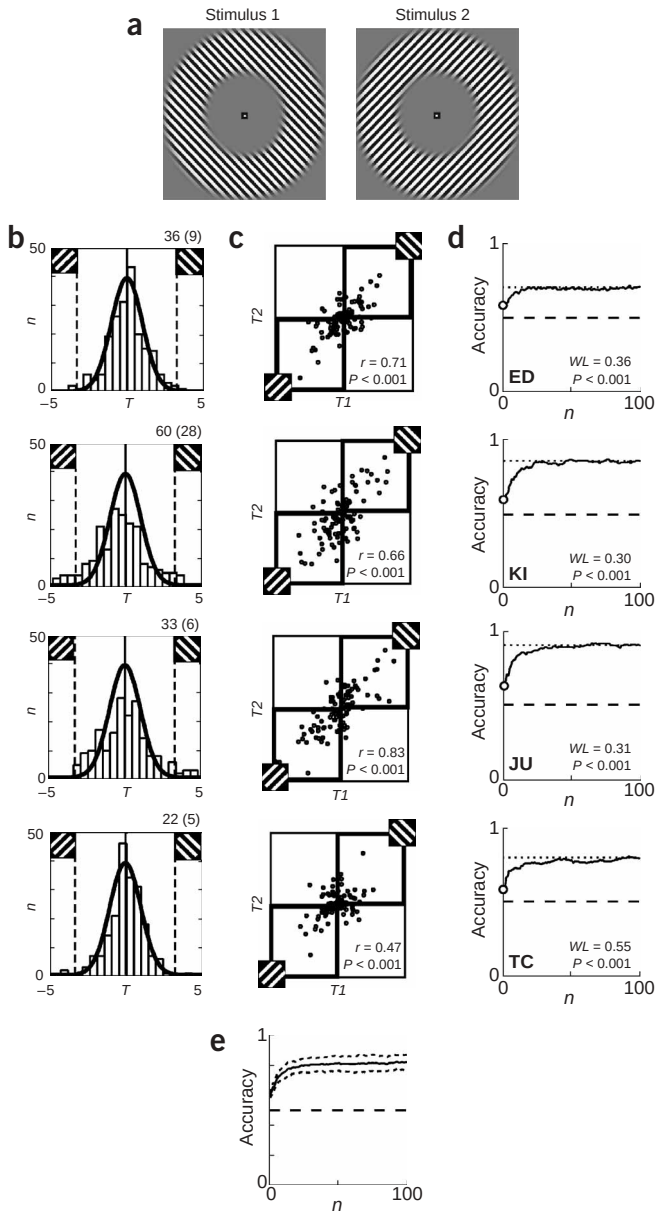

Figure 1 Orientation selectivity of fMRI responses in V1. (**a**) The stimuli used for experiment 1 were two contrast-reversing gratings with orthogonal orientation. (**b**) Orientation preference. *T*-values measuring the difference in activation under both stimulation conditions were computed as indices of voxel orientation bias for 100 voxels in V1 (Methods). The solid curve shows the expected distribution of these *T*-values if voxels showed no orientation bias. The numbers above each figure indicate the percentage of voxels exceeding an uncorrected (and in parentheses a Bonferroni-corrected) statistical threshold of $P = 0.05$. Several voxels show an orientation bias that even exceeds the conservative Bonferroni threshold (vertical dashed lines). (**c**) Orientation bias for data acquired in the first four runs (*T1*) and in the second four runs (*T2*). The positive correlation between these two indices indicates that the orientation bias of many voxels is reproducible across time. (**d**) Performance of a pattern classifier based on linear discriminant analysis. Accuracy is plotted against number of voxels *n* included to train the classifier. Performance based on one voxel (circle) is above chance (lower dashed line) but increases significantly with increasing number of voxels. A multivariate test statistic, Wilk's lambda[37] (*WL*), is also reported for each subject and confirms the significant difference between the multivariate responses for each subject under the two stimulation conditions. The upper dotted line shows asymptotic performance with increasing number of voxels. (**e**) Accuracy as in **d** averaged across the four participants. Dashed lines, s.e.m.

reproduced their bias from the first half in the second half (**Fig. 1c**), thus pointing toward a high stability of their slightly orientation-biased responses. Additional simulations using published data on the geometry of orientation pinwheeels[5] confirmed that a moderate orientation bias could indeed be expected at the spatial scale of our fMRI protocol (**Supplementary Fig. 1** online). Further analyses of the spatial distribution of orientation bias revealed no apparent spatial pattern and no differences between radial and tangential sectors of the stimulus (**Supplementary Figs. 2** and **3** online).

Having established weak but reproducible orientation-selective responses, we characterized how well the response of the entire 'population' of voxels could be used to predict the orientation of a stimulus from single measurements of brain activity. To do this, we jointly analyzed responses from multiple concurrently recorded voxels using multivariate pattern recognition methods. Pattern classifiers based on linear discriminant analysis were trained (for each participant) to distinguish between the population responses to the two stimuli using data from seven of the eight recorded sessions of imaging data (Methods; **Supplementary Fig. 4** online). These trained classifiers were then applied to the independently acquired test volumes from the eighth session to see how well the orientation of stimuli during acquisition of these volumes could be predicted. Data used for training and testing were independent time series collected at different times (Methods).

We found that single volumes from the test session could indeed be classified with high accuracy (**Fig. 1d,e**). Moreover, classification accuracy improved as patterns of activity across larger numbers of voxels in V1 were taken into account. When based only on the single voxel that showed the highest orientation bias, prediction accuracy was around 60%. However, accuracy increased markedly as more voxels were taken into account; it reached an asymptote at around 80% accuracy for 20 to 50 voxels (**Fig. 1d,e**). Prediction accuracy was equally good, and significantly better than chance, in every subject. After training, our classifier predicted stimulus orientation based on single measurements of V1 activity (Methods). Thus, a single 2-s measurement of activity in V1 was sufficient to predict with 80% accuracy which of two orthogonally orientated gratings was presented visually. This is in marked contrast to typical fMRI data analysis, in which many minutes of data across many participants are averaged to reveal slight differences in activation. Indeed, averaging activity over V1 voxels in the present experiment (as for typical fMRI data analysis, based on voxel-wise or region-of-interest analyses) did not reveal any differential activation comparing the orthogonally oriented gratings (**Supplementary Fig. 5** online). Single functional MRI images therefore contain far more information than is frequently appreciated[12].

Experiment 2

In the second experiment, we tested whether our pattern classification approach could be used to predict the orientation of a masked and invisible stimulus based on response patterns in V1. We again presented four participants with one of two orthogonally orientated gratings, but now masked using a variant of the 'standing wave of invisibility'[13], in which a single target bar is repeatedly alternated with two flanking masks that share an overlapping contour. This produces reliable and prolonged invisibility of the target bar, due to a combination of forward and backward masking, and strongly modulates target-specific responses in primate V1 (ref. 13). We adapted this protocol (Methods; **Fig. 2**) to produce prolonged masking of our oriented

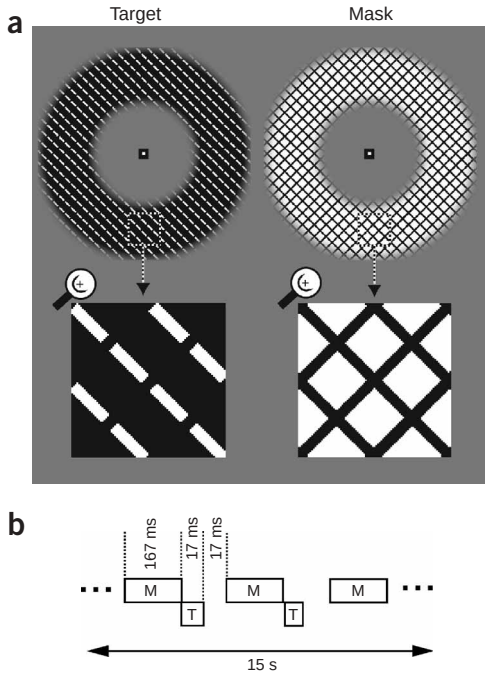

Figure 2 Procedures for experiment 2. (**a**) Target and mask stimuli used for experiment 2. Targets were bright dashed lines (with either –45° or +45° orientation) presented in an annular window on a dark background. The mask was a bright annulus (190 Cd m^{-2}) into which dark lines of two orthogonal orientations were cut. The lines of the masks and targets were in the same position, and the breaks in the target lines coincided with the intersections of the two orthogonal lines patterns of the mask (bottom, enlarged view). (**b**) During each 15-s trial, masks were presented for 167 ms, followed by a 34-ms gap during which the target was briefly presented for 17 ms. Then after 17 ms, a new mask-target cycle followed. This creates a 'standing wave' of orientation invisibility[13], in which the orientation of the target is invisible, even when the target-mask cycle is presented repeatedly for extended periods of time.

grating annulus. In separate blocks, participants viewed masked gratings oriented either to the right, or orthogonally to the left, repeatedly alternating with the mask stimulus (Methods; **Fig. 2**). At the end of each block, participants indicated with a button press which masked orientation they thought had been presented. The masking protocol was explained to four participants before the experiment. As the same oriented grating was presented throughout a 15 s block, even if participants had only a fleeting impression of its orientation, they would be able to accurately discriminate its orientation. However, despite these instructions, participants were completely unaware of the orientation of the masked gratings and were at chance performance (50.3 $\pm$ 0.4%, s.e.m.) in discriminating their orientation.

Pattern classifiers were trained for each participant using data from a subset of trials (Methods), as for the consciously perceived gratings. Then the classifiers were applied to an independent set of fMRI data (the remaining trials) to predict the orientation of the unconscious gratings based on single fMRI image volumes (that is, a single measurement for each voxel in V1). Again, we found that single volumes could be classified with an accuracy that was significantly above chance for each of the four participants (**Fig. 3**). Prediction accuracy was at chance when only a single voxel was taken into account, but improved substantially with increasing number of voxels. When all 100 voxels were used, prediction was significantly above chance for every participant. Thus, even when participants' conscious reports indicated that they themselves could not distinguish the orientation of a masked grating, their brain state contained information that could permit such discrimination.

Prediction accuracy for the masked and invisible gratings, though significantly better than chance, was nevertheless lower than for the consciously perceived grating classification. One possible interpretation is that the representation of masked orientation is substantially weaker in V1, either because of the masking procedure itself or simply because the conscious and unconscious stimuli differed in their precise physical characteristics. Taken together, the ability of participants' V1 to outperform their conscious perception in prediction accuracy indicates the presence of some information about orientation in V1 even for the unconscious targets. This therefore represents the first direct evidence that human V1 is sensitive to stimulus orientations outside conscious awareness. We also investigated orientation prediction based on signals from visual areas V2 and V3. Prediction accuracy for both visible and invisible gratings dropped substantially from V1 to V2, and still further from V2 to V3 (**Fig. 4**). For invisible gratings, we were only able to reliably reconstruct orientation from V1.

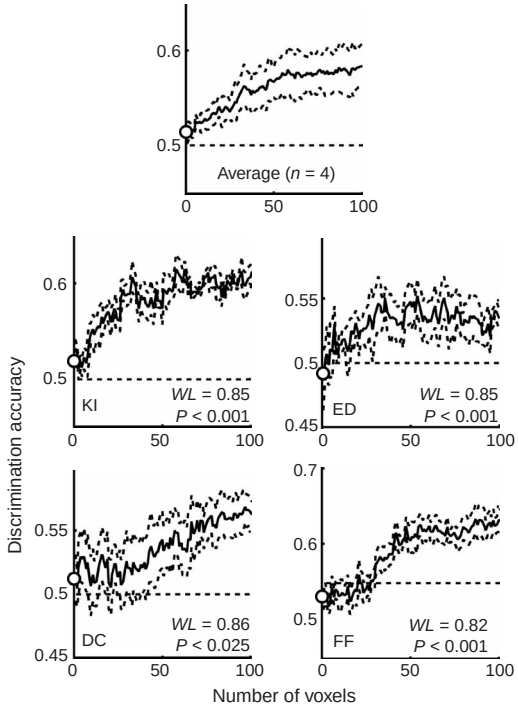

Figure 3 Discrimination accuracy for prediction of orientation of the invisible target stimuli from single samples of V1 activity. Discrimination is at chance when prediction is based in a single voxel (circle), but increases substantially when more voxels are included, reaching on average around 57%. Top, data collapsed across 4 participants (dashed lines, s.e.m.). Horizontal dashed line, chance performance. Bottom, individual subject data (here dashed lines indicate the s.e.m. for averaging across the five possible partitionings of the data into training and test datasets). It is clear that even for individual participants, the prediction of the orientation of the invisible grating is above chance when all 100 voxels are taken into account. The most conventional multivariate test statistic, Wilk's lambda[37] (*WL*), is also reported for each subject and confirms the significant difference between the multivariate responses for the two invisible gratings.

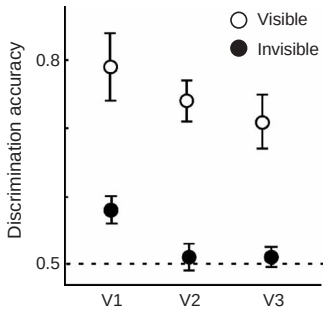

Figure 4 Discrimination accuracy compared across visual areas (averaged across participants, error bars, s.e.m.). For both visible and invisible orientation stimuli, the prediction is significantly better in V1 than in V2 and V3. For invisible gratings, reliable above-chance prediction is possible only from signals recorded in V1. The accuracy in V2 and V3 for invisible stimuli is much lower (compared with visible stimuli) than it is in V1. (Predictive accuracy for unseen stimuli as a proportion of accuracy for seen stimuli is 27.6% in V1, 4.3% in V2 and 5.0% in V3, using $p_{observed} = p_{pred} + (1 - p_{pred})\, p_{guess}$ to adjust for guess rate, where $p_{observed}$ is the observed accuracy, p_{pred} is the corrected accuracy and p_{guess} is 0.5). However, because of the large differences in cortical architecture and size between V1, V2 and V3, we cannot exclude the possibility that V2 and V3 are weakly activated by orientation stimuli.

DISCUSSION

Our findings demonstrate that individual voxels in human V1 show small but reliable orientation biases, even though the size of orientation columns in humans is likely to be much smaller than the size of a single voxel. The information present in these biased signals could be accumulated across many voxels using multivariate pattern recognition algorithms, yielding a powerful measure of orientation-selective processing. This enabled us to infer, from single fMRI measurements of V1 activity, which of two oriented stimuli was being viewed by our participants. We were able to make such predictions even when the stimuli were completely invisible to the participants. Thus, human V1 can represent information about the orientation of visual stimuli that cannot be used by participants to make a simple behavioral discrimination.

Several possibilities may account for the weak orientation bias that we found in individual voxels, despite the relatively low spatial resolution of our fMRI measurements relative to that of orientation pinwheels[4,5]. First, anisotropies in orientation processing exist for horizontal and vertical versus oblique stimuli[14], and for radial versus tangential stimuli[15,16]. Our obliquely oriented stimuli were deliberately chosen to avoid potential differences in V1 activity resulting from these anisotropies. However, our stimuli do have some regions that are predominantly radial and others that are predominantly tangential (or concentric) relative to fixation. We examined activity in these regions to determine whether such regional differences in orientation might cause the orientation bias (**Supplementary Fig. 3**). However, we found no differences in orientation bias between sectors with radial versus tangential orientations, and no other apparent pattern in the spatial distribution of orientation bias (**Supplementary Figs. 2** and **3**).

A second possibility is that the orientation biases result from anisotropies in the orientation map itself. The orientation profile sometimes does not change for up to 1–2 mm along the cortical sheet, and there are also fractures and discontinuities[5] in the pattern of orientation pinwheels. Thus, any randomly placed neuroimaging voxel will sample an anisotropic distribution of orientation preferences. Our simulations using published data on the geometry of orientation pinwheels in monkeys[5] (**Supplementary Fig. 1**) confirmed that such a moderate orientation bias of individual voxels could be expected at the spatial resolution of fMRI. In addition, the convoluted nature of the cortical surface means that individual voxels will not be collinear with the cortical surface, leading to further subtle anisotropies of the number of orientation-specific columns in each voxel. Although such a macroscopic orientation-biased signal cannot be used to study the detailed tuning of individual orientation columns, it can be used (as here) as a direct indicator of orientation-selective processing, rather than using indirect measures such as selective adaptation[1–3,6,7].

Despite complete and continuous perceptual suppression, responses of human primary visual cortex nevertheless still showed orientation selectivity. This has important implications for the role of V1 in visual awareness[17,18], which has been controversial. Although some studies have shown that activity in V1 can be closely correlated[13,18–22] with conscious visual perception, others have found no such correlation[23,24]. However, such observations do not establish whether activity in primary visual cortex necessarily leads to conscious perception. To determine if activity in V1 necessarily leads to awareness, it is necessary to investigate whether or not responses in V1 can occur in the complete absence of conscious perception. Evidence that human V1 processes unconscious stimuli has previously been indirect. Psychophysical studies have demonstrated selective adaptation to invisible orientation stimuli[1–3], suggesting that these stimuli are processed in V1. Also, temporal, parietal and frontal cortical areas can be activated by unperceived stimuli[25–29]. However, given the evidence for direct thalamic input into extrastriate cortex[30] and the substantial activation of extrastriate cortex without input from V1 (ref. 31), it is not evident that this relies on the pathway through V1. Furthermore, many of these studies have relied on types of behavioral judgments that have been criticized as too lenient for accurately judging the absence of awareness[32]. The only widely accepted method for ensuring that a stimulus was not consciously perceived is to show that a subject's performance is at chance level when forced to perform a discrimination on the stimulus[32], as in the present study. Our study thus provides direct evidence that human primary visual cortex can process the orientation of completely invisible stimuli, suggesting that selective processing in V1 is not sufficient to cause visual awareness. Whether to be represented in conscious experience information has to cross a threshold level of activity, or perhaps needs to be relayed to another region of the brain, is an intriguing question for further research.

In summary, we have shown that individual voxels in primary visual cortex show weak but reliable orientation preferences. Using multivariate pattern recognition, we have been able to accumulate this weak information across many voxels, yielding a powerful and direct index of orientation-selective processing. We found that even completely invisible oriented stimuli were processed in human V1 in an orientation-selective fashion. Thus, feature-specific representation in V1 of humans may be necessary but is not sufficient for consciousness.

METHODS

Participants and experimental design. Six healthy, right-handed volunteers with normal vision (age 25–33 years) gave written informed consent to participate in the study, which was approved by the University College London ethics committee. Two people participated in experiment 1 alone, two participated in experiment 2 only, and two participated in both experiments. In the first experiment, stimuli were tilted gratings (spatial frequency 2 cpd; Michelson contrast 90%) presented within a smoothed annular window that subtended from 4° to 8° eccentricity (**Fig. 1**). Stimuli were presented in alternating blocks of 30 s, during which gratings with one of two possible orientations (either –45° or +45°) were selected and continuously contrast

reversed at a frequency of 4 Hz. Participants viewed the gratings passively while monitoring a color change to the central fixation spot. Mean luminance of the gratings was equal to background luminance (100 Cd m^{-2}). During each of the eight scanning runs, four such blocks were presented, two with gratings of each orientation. Before scanning, participants practiced to ensure they would be able to perform the task and maintain stable fixation.

In the second experiment, participants viewed stimuli similar to those used in the first experiment but now rendered invisible by metacontrast masking[13]. Target stimuli were 'dashed' white gratings (spatial frequency 2 cpd; luminance 145 Cd m^{-2}; duty cycle 10%) with an orientation of either –45° or +45° that were presented within an annulus (diameter 4° to 8°) on a black background (**Fig. 2**). Targets were repeatedly presented briefly for 17 ms every 200 ms, interleaved by long presentations of the mask stimuli (190 Cd m^{-2}) for 167 ms. This creates a so-called 'standing wave of invisibility'[13], which optimally combines forward and backward metacontrast masking and renders the orientation of the target invisible, even for prolonged periods of presentation. As in previous studies[13], pilot experiments showed that maximum masking was achieved when the mask followed the target by a small delay and the target followed the mask immediately. During each trial, one target orientation was randomly chosen, and the target-mask cycle was presented for 15 s. Participants were required to maintain fixation on a central fixation spot and to try and identify the orientation of the target. After offset of the target, participants were given 500 ms to judge the orientation of the target. The performance of all four participants was at chance level. Before scanning, participants practiced to ensure they would be able to perform the task and maintain stable fixation.

fMRI acquisition. A Siemens Allegra 3T scanner with Nova Medical occipital surface coil was used to acquire functional MRI volumes (20 slices; TR, 1.3 s) at a conventional resolution of 3 × 3 × 3 mm. Eight runs of 125 functional MRI volumes per participant were acquired in experiment 1. In experiment 2, acquisition was modified and arranged into 90 runs, each of which comprised 13 volumes (corresponding to one trial). Runs were self-paced with a pause of around 5 s between each and with long breaks after every thirtieth run. In both experiments, a T1-weighted structural image was also acquired, together with 2–3 retinotopic mapping runs of 165 volumes each, during which participants viewed standard stimuli that mapped the horizontal and vertical meridians.

Data analysis. Data were preprocessed using SPM2 (http://www.fil.ion.ucl.ac.uk/spm). After discarding the first three images of each scanning run to allow for magnetic saturation effects, we realigned and coregistered the remaining images to the individual participants' structural scans. To identify stimulus-driven cortical regions, the data were modeled voxel-wise, using a general linear model that included the two experimental conditions[33]. To extract activity from primary visual cortex, we created a mask volume defining V1. This was obtained using the meridian localizers from the retinotopic mapping sessions following standard definitions of V1 (ref. 34) together with segmentation and cortical flattening in MrGray[35,36].

Voxel orientation bias. As a measure of orientation bias, we computed the difference in activation under both stimulation conditions as $T = (\bar{x}_1 - \bar{x}_2)/\sqrt{(s_1^2/N_1 + s_2^2/N_2)}$ separately for the 100 voxels in V1 that showed the strongest stimulus-driven responses. Here $\bar{x}_i$ and s_i denote the univariate means and standard deviations for one voxel when stimulated with orientation i. If signals in these voxels were not orientation biased, the distribution of T values across voxels should follow a T-distribution with 542 degrees of freedom (solid lines in **Fig. 1b**).

Pattern classification. In the first experiment, 17 recorded images were extracted for each of the 32 stimulation blocks (eight runs × four blocks per run; delayed by three volumes to account for the delay of the hemodynamic response function), yielding a total of 544 volumes (details in **Supplementary Fig. 4**). Classification performance was assessed using linear discriminant analysis with m-fold cross-validation[8]. The images were split into eight groups, each including the two blocks of each stimulation condition from one run (**Supplementary Fig. 4**). Images from one of these groups were assigned to a test dataset, and the images for the remaining seven groups were assigned to a training dataset yielding 238 training volumes and 34 test volumes for each of the two conditions. (Training and test datasets were acquired in independent blocks.) In the second experiment, ten recorded images were extracted for each of the 90 trials (delayed by three volumes to account for the delay of the hemodynamic response function and T1 saturation), yielding a total of 900 volumes (90 runs × 10 images per run). The images were split into five groups, each including nine runs of each stimulation condition. Images from one of these groups were assigned to a test dataset, and the images for the remaining runs were assigned to a training dataset, yielding 90 training volumes and 360 test volumes for each of the two conditions. The sequence of training and test was repeated eight times for experiment 1 and five times for experiment 2 (each time assigning a different group of images to the independent test dataset).

For each training and test volume, we computed first a raw activation vector $\mathbf{a}$ that was obtained by extracting fMRI signal intensity from the 100 voxels of that subject's V1 that showed a maximal main effect of stimulation across all conditions. This was transformed to a normalized activation vector $\mathbf{x}$ with unit length following $\mathbf{x} = \mathbf{a}/||\mathbf{a}||$. To achieve optimal (minimum error-rate) classification, a sample with response vector $\mathbf{x}$ from the test dataset is assigned to the stimulus category s_i for which the posterior probability distribution

$$p(s_i|\mathbf{x}) = p(\mathbf{x}|s_i)\,p(s_i)/p(\mathbf{x}) \quad (1)$$

estimated from the training dataset is maximal[8]. Under flat (unbiased) priors, $p(s_i)$ as here is sufficient to maximize $p(\mathbf{x}|s_i)$. When the responses of both categories follow a multivariate normal distribution with equal covariance matrices $\mathbf{\Sigma} = \mathbf{\Sigma}_1 = \mathbf{\Sigma}_2$, the maximization of $p(\mathbf{x}|s_i)$ can be achieved by computing the Mahalanobis distance between a given test sample $\mathbf{x}$ and the two sample means from the training dataset, $\bar{\mathbf{x}}_1$ and $\bar{\mathbf{x}}_2$, (ignoring constants)[8]:

$$MD(i) = (\mathbf{x} - \bar{\mathbf{x}}_i)^T S^{-1} (\mathbf{x} - \bar{\mathbf{x}}_i) \quad (2)$$

where S is the pooled covariance matrix estimated from the sample. The vector is then assigned to the stimulus category for which $MD(i)$ is minimal. The parameters S, $\bar{\mathbf{x}}_1$ and $\bar{\mathbf{x}}_2$ are obtained using only the training dataset. In our case, the pattern classification was done on an increasing number of voxels (between 1 and 100), for which the voxels were rank ordered according to their T-value (for the difference between the two orientations), again computed from the training dataset only.

Additionally, we trained a non-parametric classification algorithm that estimates $p(\mathbf{x}|s_i)$ based on Parzen windows[8]. Both parametric and non-parametric classifiers[8] performed equivalently and significantly above chance in classifying both conscious and unconscious oriented gratings (**Supplementary Fig. 5**). Thus, the precise choice of discriminant function did not affect the reliability of our findings. In contrast, prediction was much worse when based on a conventional region of interest (ROI) analysis or on the single best voxel (**Supplementary Fig. 5**). As a further validation of our results, we computed the most conventional multivariate test statistic, Wilk's lambda, which tests for significant differences between multivariate means[37]. This has the advantage that it does not rely on a training-test cycle. This confirmed the significant differences between the response patterns (**Figs. 1** and **3**).

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank K. Friston and J. Driver for comments on the manuscript. The Wellcome Trust funded this work.

AUTHOR CONTRIBUTIONS

J.D.H. and G.R. conceived and carried out the experiment. J.D.H. analyzed the data. G.R. and J.D.H. co-wrote the paper.

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Received 16 February; accepted 29 March 2005
Published online at http://www.nature.com/natureneuroscience/

1. He, S., Cavanagh, P. & Intriligator, J. Attentional resolution and the locus of visual awareness. *Nature* **383**, 334–337 (1996).
2. He, S. & MacLeod, D.I. Orientation-selective adaptation and tilt after-effect from invisible patterns. *Nature* **411**, 473–476 (2001).
3. Rajimehr, R. Unconscious orientation processing. *Neuron* **41**, 663–673 (2004).

4. Bartfeld, E. & Grinvald, A. Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. *Proc. Natl. Acad. Sci. USA* **89**, 11905–11909 (1992).
5. Obermayer, K. & Blasdel, G.G. Geometry of orientation and ocular dominance columns in monkey striate cortex. *J. Neurosci.* **13**, 4114–4129 (1993).
6. Tootell, R.B. *et al.* Functional analysis of primary visual cortex (V1) in humans. *Proc. Natl. Acad. Sci. USA* **95**, 811–817 (1998).
7. Boynton, G.M. & Finney, E.M. Orientation-specific adaptation in human visual cortex. *J. Neurosci.* **23**, 8781–8787 (2003).
8. Duda, O.R., Hart, P.E. & Stork, D.G. *Pattern Classification* (Wiley, New York, 2001).
9. Haxby, J.V. *et al.* Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* **293**, 2425–2430 (2001).
10. Cox, D.D. & Savoy, R.L. Functional magnetic resonance imaging (fMRI) 'brain reading': detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* **19**, 261–270 (2003).
11. Carlson, T.A., Schrater, P. & He, S. Patterns of activity in the categorical representation of objects. *J. Cogn. Neurosci.* **15**, 704–717 (2003).
12. Dehaene, S. *et al.* Inferring behavior from functional brain images. *Nat. Neurosci.* **1**, 549–550 (1998).
13. Macknik, S.L. & Livingstone, M.S. Neuronal correlates of visibility and invisibility in the primate visual system. *Nat. Neurosci.* **1**, 144–149 (1998).
14. Furmanski, C.S. & Engel, S.A. An oblique effect in human primary visual cortex. *Nat. Neurosci.* **3**, 535–536 (2000).
15. Westheimer, G. The distribution of preferred orientations in the peripheral visual field. *Vision Res.* **43**, 53–57 (2003).
16. Wilkinson, F. *et al.* An fMRI study of the selective activation of human extrastriate form vision areas by radial and concentric gratings. *Curr. Biol.* **10**, 1455–1458 (2000).
17. Crick, F. & Koch, C. Are we aware of neural activity in primary visual cortex? *Nature* **375**, 121–123 (1995).
18. Tong, F. Primary visual cortex and visual awareness. *Nat. Rev. Neurosci.* **4**, 219–229 (2003).
19. Ress, D. & Heeger, D.J. Neuronal correlates of perception in early visual cortex. *Nat. Neurosci.* **6**, 414–420 (2003).
20. Polonsky, A., Blake, R., Braun, J. & Heeger, D.J. Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. *Nat. Neurosci.* **3**, 1153–1159 (2000).
21. Tong, F. & Engel, S.A. Interocular rivalry revealed in the human cortical blind-spot representation. *Nature* **411**, 195–199 (2001).
22. Super, H., Spekreijse, H. & Lamme, V.A. Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nat. Neurosci.* **4**, 304–310 (2001).
23. Gur, M. & Snodderly, D.M. A dissociation between brain activity and perception: chromatically opponent cortical neurons signal chromatic flicker that is not perceived. *Vision Res.* **37**, 377–382 (1997).
24. Leopold, D.A. & Logothetis, N.K. Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* **379**, 549–553 (1996).
25. Moutoussis, K. & Zeki, S. The relationship between cortical activation and perception investigated with invisible stimuli. *Proc. Natl. Acad. Sci. USA* **99**, 9527–9532 (2002).
26. Rees, G. *et al.* Unconscious activation of visual cortex in the damaged right hemisphere of a parietal patient with extinction. *Brain* **123**, 1624–1633 (2000).
27. Vuilleumier, P. *et al.* Neural fate of seen and unseen faces in visuospatial neglect: a combined event-related functional MRI and event-related potential study. *Proc. Natl. Acad. Sci. USA* **98**, 3495–3500 (2001).
28. Dehaene, S. *et al.* Imaging unconscious semantic priming. *Nature* **395**, 597–600 (1998).
29. Luck, S.J., Vogel, E.K. & Shapiro, K.L. Word meanings can be accessed but not reported during the attentional blink. *Nature* **383**, 616–618 (1996).
30. Yukie, M. & Iwai, E. Direct projection from the dorsal lateral geniculate nucleus to the prestriate cortex in macaque monkeys. *J. Comp. Neurol.* **201**, 81–97 (1981).
31. Goebel, R., Muckli, L., Zanella, F.E., Singer, W. & Stoerig, P. Sustained extrastriate cortical activation without visual awareness revealed by fMRI studies of hemianopic patients. *Vision Res.* **41**, 1459–1474 (2001).
32. Hannula, D.E., Simons, D.J. & Cohen, N.J. Imaging implicit perception: promise and pitfalls. *Nat. Rev. Neurosci.* **6**, 247–255 (2005).
33. Friston, K.J. *et al.* Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* **2**, 189–210 (1995).
34. Sereno, M.I. *et al.* Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* **268**, 889–893 (1995).
35. Teo, P.C., Sapiro, G. & Wandell, B.A. Creating connected representations of cortical gray matter for functional MRI visualization. *IEEE Trans. Med. Imaging* **16**, 852–863 (1997).
36. Wandell, B.A., Chial, S. & Backus, B.T. Visualization and measurement of the cortical surface. *J. Cogn. Neurosci.* **12**, 739–752 (2000).
37. Everitt, B.S. & Dunn, G. *Applied Multivariate Data Analysis* (Edward Arnold, London, 1991).