NeuroImage 44 (2009) 223–231

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions

John T. Serences [a,*,1], Sameer Saproo [a,1], Miranda Scolari [a,1], Tiffany Ho [a,1], L. Tugan Muftuler [b]

[a] Department of Cognitive Sciences, University of California, Irvine, CA 92697-5100, USA
[b] Radiological Sciences, School of Medicine, University of California, Irvine, CA 92697-5100, USA

ARTICLE INFO

Article history:
Received 5 May 2008
Revised 18 July 2008
Accepted 20 July 2008
Available online 5 August 2008

Keywords:
Attention
Multivariate analysis
fMRI
Mutual information
Vision
Perception
Orientation

ABSTRACT

In order to form stable perceptual representations, populations of sensory neurons must pool their output to overcome physiological noise; selective attention is then required to ensure that behaviorally relevant stimuli dominate these 'population codes' to gain access to awareness. However, the role that attention plays in shaping population response profiles has received little direct investigation, in part because most traditional neurophysiological methods cannot simultaneously assess changes in activity across large populations of sensory neurons. Based on single-unit recording studies, current theories hold that attending to a relevant feature sharpens the population response profile and improves the signal-to-noise ratio of the resulting perceptual representation. Here, we test this hypothesis using fMRI and an analysis approach that is able to estimate the influence of feature-based attentional modulations on population response profiles. We first derive orientation tuning functions for single voxels in human primary visual cortex, and then use these tuning functions to sort voxels according to their orientation preference. We then show that selective attention systematically biases population response profiles so that behaviorally relevant stimuli are represented in the visual system at the expense of behaviorally irrelevant stimuli. Collectively, the present results (1) provide a new approach for precisely characterizing feature-selective responses in human sensory cortices and (2) reveal how behavioral goals can shape population response profiles to support the formation of coherent perceptual representations.

Every perception is thought to be based on a pattern of electrical activity distributed across a population of feature-selective sensory neurons. Without loss of generality, the problem of 'decoding' these patterns of activity can be framed by considering an experiment in which the subject's goal is to determine the orientation of a stimulus in a visual display. Suppose that the subject has only one sensory neuron with a Gaussian tuning function centered on a particular orientation (Fig. 1A). In the absence of noise, or variability in the response of this single neuron, the problem of inferring the stimulus orientation is relatively straightforward because a specific firing rate is associated with only one or two possible orientations. However, real sensory neurons – like all biological units – are noisy and the firing rate evoked by the same stimulus will vary across repeated presentations. As a result, the activity level recorded on a single trial is consistent with any number of different orientations and carries little useful information about the true nature of the stimulus (Fig. 1B, Pouget et al., 2000; Pouget et al., 2003; Seung and Sompolinsky, 1993).

Now consider this experiment performed on a subject with many orientation-selective neurons, all with tuning functions of the same shape but with preferred orientations distributed uniformly over a range from 0° to 180° (Fig. 1C). When this subject views a stimulus, a noisy response will be recorded from each neuron. However, using relatively simple estimation procedures, the distribution of responses across neurons (or the *population response profile*) can be used to efficiently recover the approximate orientation of the stimulus even in the presence of physiological noise (Georgopoulos et al., 1982; Jazayeri and Movshon, 2006; Ma et al., 2006; Pouget et al., 2003; Sanger, 1996; Seung and Sompolinsky, 1993; Shadlen et al., 1996). While the details vary, most estimation techniques capitalize on the fact that combining the information provided by multiple sensory neurons results in a more stable prediction compared to estimates based only on the firing rate of the most active neuron(s).

In addition to using population coding schemes to overcome internal neural noise, perceptual mechanisms must also deal with crosstalk in the sensory input stream that can arise when many objects are simultaneously present in the visual field. Multiple stimuli pose a fundamental problem for perception because the brain has a limited processing capacity, so inputs must compete for representation in cortex (Desimone and Duncan, 1995; Reynolds and Desimone, 1999; Tsotsos, 1997). For example, you can either choose to be aware of

* Corresponding author. Department of Psychology, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0109, USA.
E-mail address: jserences@ucsd.edu (J.T. Serences).
[1] Present address: Department of Psychology, University of California, San Diego, CA 92093-0109, USA.

1053-8119/$ – see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2008.07.043

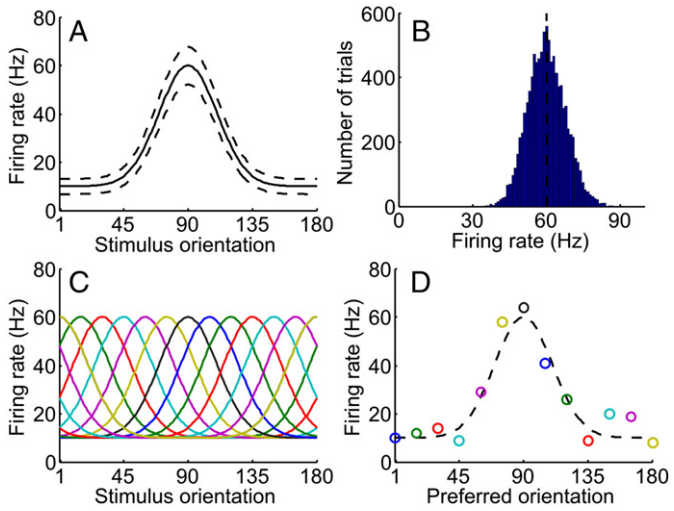

Fig. 1. Inference based on single units and population response profiles. (A) The simulated firing rate of a single V1 neuron that responds maximally to a 90° stimulus (±1 standard deviation with Poisson noise) (B) Distribution of responses from neuron in panel A to a 90° stimulus on 10,000 simulated trials. Although the average response is given by the height of the tuning function – 60 Hz in this case – there is considerable variability in the response on a trial-by-trial basis. (C) Population of 13 V1 neurons that are tuned to different orientations. (D) Response of each neuron in panel C on a single trial to a 90° stimulus plotted as a function of the neuron's preferred orientation. Although each neuron produces a noisy response, the population response profile is approximately centered on the correct orientation (dashed black line).

the pressure of the chair on your back or the banter of coworkers in the hall — but probably not both at precisely the same time. Countless single-unit recording studies have established that such goal directed (or *top-down*) deployments of attention can enhance the firing rate of single neurons that encode behaviorally relevant stimuli (e.g. Desimone and Duncan, 1995; Martinez-Trujillo and Treue, 2004; McAdams and Maunsell, 1999; Reynolds et al., 2000). In turn, the signal-to-noise ratio of single neurons should improve as firing rates increase, thereby improving the reliability of information conveyed by population codes about stimuli in the environment (this is true under reasonable assumptions about the nature of physiological noise, Borst and Theunissen, 1999; Paradiso, 1988; Pouget et al., 2001; Shadlen et al., 1996).

In sum, theoretical studies suggest that perception is based on population codes, and empirical studies have shown that top-down attentional factors are necessary to disambiguate competing sensory inputs so that relevant stimuli win representation at the expense of irrelevant stimuli. However, little work has been done to directly examine the impact of attentional modulations on the information content of population codes. Multiunit recording techniques are still being refined and functional magnetic resonance imaging (fMRI) studies typically lack the ability to clearly link modulations in the blood oxygenation level dependent (BOLD) signal with population level dynamics. To address this issue, sophisticated multivoxel pattern analysis (MVPA) algorithms have recently been developed by Kamitani, Tong and others to discriminate activation patterns within human visual cortex that are associated with different attentional or perceptual states (Kamitani and Tong, 2005, 2006; Haxby et al., 2001; Norman et al., 2006; Haynes and Rees, 2005). These methods discriminate activation patterns by combining weak feature-selective responses from all voxels in a visual area; however, this pooling of information can partially obscure the tuning properties of the individual voxels (and by inference the neurons) that are modulated by attentional or perceptual factors.

Here, we demonstrate that complex pattern classification algorithms are not always needed to make inferences about changes in population response profiles. Instead, robust feature-tuning functions can be derived from individual voxels and used to more directly quantify the influence of cognitive factors on population codes. This approach allowed us to sort voxels on the basis of their tuning preference and to show that attention biases population response profiles in favor of voxels that represent behaviorally relevant features. This attention-mediated bias in the population code presumably supports the formation of stable perceptual representations by influencing the signal-to-noise ratio of neurons that encode behaviorally relevant stimuli.

Methods

Subjects

Four neurologically healthy subjects were recruited from the University of California, Irvine (UCI) community. Each subject gave written informed consent as per Institutional Review Board requirements at UCI, and completed 1 h of training outside the scanner, and two 1.5 h scanning sessions held on separate days. Compensation for participation was $10/h for training and $20/h for scanning.

Tuning function scans

Subjects were instructed to maintain fixation on a central white fixation point (subtending 0.5°) that remained on the screen for the duration of each scan (where a 'scan' refers to an entire fMRI data collection run lasting approximately 5 min, see below). On each 12 s trial, a full contrast grayscale sinusoidal grating (0.5 cycles/°) was flickered at 2 Hz (250 ms on, 250 ms off) in one of 8 possible orientations, evenly spaced across 180° (0°, 22.5°, 45°, 67.5°, 90°, 112.5°, 135°, and 157.5°, where 0° is defined as horizontal, see Fig. 2A). The spatial phase of the grating was randomly selected from a set of four possible phases on each presentation to attenuate the perception of apparent motion. The oriented stimulus subtended 20° visual angle with a small circular aperture removed around fixation (2.5° radius). The order of orientations was randomized on each scan with the constraint that the same orientation could not be presented on successive trials. Subjects were instructed to respond when the contrast of the stimulus decreased slightly, which is henceforth referred to as a target event. The contrast reduction that defined a target was titrated on an individual basis so that the hit rate remained at approximately 80% over the course of all scanning sessions (the projector used to the display the stimuli had a linearized output). Each target was presented for a single 250 ms frame, and there were 5 targets in each trial. The timing of each target was pseudorandomly determined with the

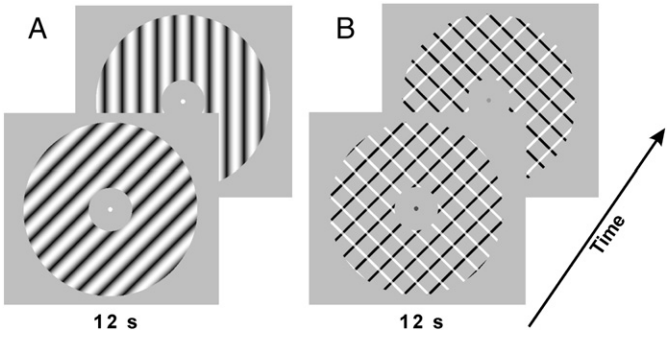

Fig. 2. Task design. (A) Schematic of the stimulus paradigm used to evaluate orientation selectivity in human visual cortex (see Figs. 3–5). There were eight possible stimulus orientations, and the subject's task was to press a button whenever the contrast of the grating dimmed slightly. (B) Schematic of the stimulus paradigm used to evaluate attentional modulations in visual cortex. The color of the central fixation point instructed observers to either monitor the 45° oriented bars or the 135° oriented bars; they were to press a button whenever the width of the attended lines increased slightly. See text for more details.

following constraints: each target was separated from the previous one by at least 1.5 s, and targets were restricted to a temporal window of 2–11 s following the onset of the trial. Button-press responses made within 1 s of target onset were considered correct. Each trial was separated by a blank 500 ms inter-trial-interval before the next oriented stimulus was presented. Observers completed five tuning function scans per 1.5 h scanning session; each scan lasted for 300 s and contained 3 presentations of each stimulus orientation for a total of 24 trials/scan.

Attention scans

On attention scans, the stimulus was comprised of two superimposed squarewave gratings; one grating was always oriented at 45° and the other oriented at 135° (Fig. 2B). The 45° and 135° lines were always rendered in a different color (black or white); however, the color of the 45° and 135° lines was perfectly counterbalanced within each scan. 500 ms prior to the beginning of each trial, participants were instructed to attend to one grating and ignore the other via the color of the fixation point (red and green were assigned to each orientation and counterbalanced across subjects). As in the tuning function scans, subjects were instructed to maintain gaze on the central fixation point for the duration of each scan. On each 12 s trial, the gratings flickered at a rate of 2 Hz, and subjects had to press a button whenever the width of the attended lines increased slightly on a single 250 ms frame, while ignoring equally probable changes in the width of the unattended lines. There were 5 target/distractor events on each trial, with an equal number of targets and distractors across each scan. The temporal sequence of targets/distractors was governed by the same parameters that dictated target presentations in the tuning function scans. Button-press responses made within 1 s of target onset were considered correct; button presses made within 1 s of a distractor were considered false alarms. Observers completed five scans of this procedure during each session, and each scan consisted of 24 trials (12 trials in each attention condition, for a scan duration of 300 s). Attention scans were interleaved with the tuning function scans over the course of an experimental session.

Functional localizer scans

Independent functional localizer scans were run to identify voxels within each retinotopically organized visual area that responded to the spatial position occupied by the stimulus aperture in the tuning function and attention scans. The sequence of events on the localizer scans was nearly identical to the tuning function scans; however, there was a blank inter-trial-interval of 12 s and each of the 8 orientations was presented only once during the scan (yielding a scan duration of 192 s).

Retinotopic mapping scans

Retinotopic mapping data were obtained in one to two scans per observer using a checkerboard stimulus and standard presentation parameters (stimulus flickering at 8 Hz and subtending 60° of polar angle Engel et al., 1994; Sereno et al., 1995). This procedure was used to identify ventral visual areas V1, V2v, V3v, and V4v. Our high-resolution scanning protocol did not provide sufficient coverage to acquire data from dorsal occipital areas V2d, V3d and V3a or visual areas in parietal and frontal cortex. To aid in the visualization of early visual cortical areas, we projected the retinotopic mapping data onto a computationally inflated representation of each observer's gray/white matter boundary. All retinotopic mapping data was aligned with the tuning function/attention scan data (as opposed to moving all functional data into register with the retinotopy data) to minimize the number of spatial transformations applied to the tuning function/attention scan data.

fMRI data acquisition and analysis

MRI scanning was carried out on a Philips Achieva 3-Tesla scanner equipped with an 8-channel SENSE head coil at the John Tu and Thomas Yuen Center for Functional Onco Imaging, University of California, Irvine. Anatomical images were acquired using a MPRAGE T1-weighted sequence that yielded images with a 1 mm^3 resolution (TR/TE = 11/3.3 ms, TI = 1100 ms, 150 slices, flip angle = 18° with no SENSE acceleration). Functional images were acquired using a gradient echo planar imaging (EPI) pulse sequence, which covered the occipital lobe with 25 oblique transverse slices. Slices were acquired in sequential order with 1.5 mm thickness and 0.5 mm gap to avoid slice crosstalk; therefore a 50 mm thick slab was acquired (TR = 2000 ms, TE = 25 ms, flip angle = 70°, image matrix = 120 (AP) × 92 (RL), with FOV = 240 mm (AP) × 180 mm (RL), SENSE factor = 2, voxel size = 2 mm × 2 mm × 1.5 mm).

Data analysis was performed using BrainVoyager QX (v 1.86; Brain Innovation, Maastricht, The Netherlands) and custom timeseries analysis routines written in Matlab (version 7.1; The Math Works, Natick, Massachusetts). Data from the tuning function and attention experiments were collected in 10 scans per subject (5 scans for each task type per session). All EPI images were slice-time corrected, motion-corrected (both within and between scans) and high pass filtered (3 cycles/run) to remove low frequency temporal components from the timeseries.

Region of interest selection procedure

To identify voxels that responded to the retinotopic position of the stimulus aperture, data from the functional localizer scans were analyzed using a GLM that contained a regressor marking each 12 s stimulus epochs (a boxcar convolved with a gamma function, Boynton et al., 1996). Voxels within each visual area were included in all subsequent analyses if they passed a threshold of $p < .01$, corrected for multiple comparisons using the False Discovery Rate (FDR) algorithm implemented in Brain Voyager. See Table 1 for the number of voxels within each visual area that passed this threshold.

Mulitvoxel pattern analysis of orientation selectivity (MVPA, Fig. 3)

First, we extracted the raw timeseries from each voxel within each ROI (e.g. left V1) during a time period extending from 6 s to 18 s after the presentation of each stimulus (i.e. the timeseries was shifted by 6 s to account for the hemodynamic lag). The epoched timeseries from each voxel was then normalized on a scan-by-scan basis using a z-transform. Temporal epochs from all but one scan were extracted to form a 'training' data set for the classification analysis; epochs from the remaining scan were defined as a 'test' set (note that we use the term "scan" to refer to an entire 310 s data collection sequence so the training and test data were always independent). We then trained a Support Vector Machine (the OSU-SVM implementation, downloaded from http://sourceforge.net/projects/svm/) based only on the training data and then used it to classify the orientation of the stimulus on each trial from the test scan (Kamitani and Tong, 2005). This procedure was repeated using a 'hold-one-scan-out' cross-validation approach so that epochs from every scan were used as a test set in turn. Since each observer completed 5 scans/session, the overall classification accuracy

Table 1
Size of each ROI in voxels

Visual area	Min size (voxels)	Max size (voxels)	Mean (voxels, ±STD)
V1	280	703	454 (128)
V2v	76	336	217 (75)
V3v	48	314	146 (67)
V4v	27	349	162 (88)

was defined as the average classification accuracy across all 5 possible permutations of holding one scan out as the test set and using the remaining scans as a training set. Classification accuracy was computed separately for each ROI based on all voxels within that ROI that exhibited a significant response to the stimulus used in the functional localizer scan (Table 1); accuracy data from complementary ROIs in each hemisphere were then averaged because no significant differences were observed between left and right visual areas. Data from each session were analyzed separately and then averaged to generate an overall estimate for each observer.

Tuning functions (Figs. 4 and 5)

Before computing tuning functions or attention effects (see next section), the timeseries from every voxel was z-normalized on a scan-by-scan basis to remove differences in mean signal intensity. Each voxel from a visual area was then assigned to an orientation-preference bin based on data from all scans except one; the preference of each voxel was determined based on the orientation that evoked the largest mean response over a temporal epoch extending from 6 s to 18 s post-stimulus after removing the mean response across all voxels at each orientation (to correct for main effects that had a common influence on the response of every voxel). Then, using only data from the remaining scan, we computed the response of voxels in each bin to every stimulus orientation. As described in the MVPA section above, this hold-one-scan cross-validation procedure was repeated 5 times to establish internal reliability and to avoid selection procedures that biased or predetermined experimental outcomes (i.e. issues of circularity). Data from each visual area, hemisphere, and experimental session were analyzed separately and then combined appropriately to form an average tuning function for each visual area of each subject. The tuning functions were then characterized using a circular Gaussian (von Mises) function

$$\delta = b + a * e^{[k*\cos(x-\mu)-1]} \quad (1)$$

where a is the response amplitude, b is the baseline, and k is the concentration parameter. Although care was taken to avoid circularity in this analysis procedure, we also verified that passing in a data set comprised of independent identically distributed (IID) noise did not yield any meaningful results (i.e. we found flat tuning functions because there was no consistent signal or internal reliability in the noise data set).

In Supplemental Fig. 1, we show the results of this tuning function analysis without removing the mean activation level across all voxels at each orientation (as we describe above in the main analysis). The shape of the tuning functions is similar in the absence of this normalization procedure; however, the distribution of preferred orientations across voxels is shifted towards near vertical angles (67.5°, see Supplemental Fig. 2). The shift in orientation preferences occurs because this orientation evoked a larger BOLD response across nearly *all* voxels in a visual area. We factored out these common responses because we believe that it is unlikely that information shared by all voxels (or more importantly by all neurons) contributes in a meaningful way to the perceptual discrimination of the current orientation when the visual system reads out the population response profile. Thus, by subtracting out these common responses, we focus on differential patterns of activation across the population, which is more likely to reflect the information that supports perceptual inferences about the nature of the sensory stimulus. However, even if these assumptions are shown to be incorrect or are otherwise ill-conceived, similar tuning functions and attentional modulations (see below) are obtained.

Mutual information measures

To examine the variability of voxel-based tuning functions across each visual area, we sorted voxels based on the amount of information that they conveyed about stimulus orientation. First, we computed the entropy of the BOLD responses in each voxel, which is a measure of response uncertainty across all stimulus orientations

$$H(B) = -\sum_{b\in B} p(b)\log_2 p(b) \quad (2)$$

Since the standard definition of entropy holds for both continuous and discrete variables, we converted the continuous BOLD response into a discrete variable (B) by dividing the range of responses into a set of equidistant bins (b) of sufficiently small size (Cover and Thomas, 1991). In this formulation, $p(b)$ is the frequency with which a response falls into bin b divided by the total number of responses measured from a given voxel. The bins were defined based on the range of responses across all voxels within a visual area from a subject in a scanning session. Before discretizing the BOLD responses, we z-normalized the timeseries from each voxel to have zero mean and unit variance (Fuhrmann Alpert et al., 2007).

Next, we computed conditional entropy $p(b|\theta)$, which yields a measure of response uncertainty given knowledge of the stimulus orientation. If there is a dependence between stimulus orientation and the observed BOLD responses, the introduction of Θ as a conditional random variable should reduce the uncertainty and hence the entropy of the random variable B

$$H(B|\Theta) = -\sum_{\theta\in\Theta} p(\theta)\sum_{b\in B} p(b|\theta)\log_2 p(b|\theta) \quad (3)$$

The information content carried by each voxel can then be defined as the reduction in uncertainty for each voxel's BOLD response given the stimulus orientation, or the mutual information (MI). Given that we are using logarithm to the base 2 in our calculations, the unit of measure is *bit*.

$$\mathrm{MI}(B;\Theta) = H(B) - H(B|\Theta) \quad (4)$$

Bayes' rule ensures that the equation is symmetric, so that $p(B|\theta)$ could also be used to infer $p(\theta|B)$, or the reduction in uncertainty about stimulus orientation given a distribution of BOLD responses (under the assumption of flat priors).

Therefore, if there is a strong mutual dependence between the distribution of BOLD responses (B) and stimulus orientation (θ) for a particular voxel, then that voxel will produce a high MI value. If they are completely dependent – such that knowing one gives complete information about the other – then MI would be equal to the entropy of either one of them (because MI is symmetric $\mathrm{MI}(B;\Theta) = \mathrm{MI}(\Theta;B)$). On the other hand, if the distribution of BOLD responses and stimulus orientation are completely independent, then that voxel will yield a MI value of zero.

For an informative voxel, entropy $H(B)$ should be high, implying that the magnitude of the BOLD response changes substantially as a function of stimulus orientation. Thus a voxel that has a relatively 'peaked' tuning function (up to a point) would be more informative and hence have higher entropy than a voxel that has a flat tuning function. Conditional entropy $H(B|\Theta)$ for an informative voxel should be relatively low (with respect to the entropy) because that voxel should yield similar BOLD responses on each successive presentation of a particular orientation. Thus, a high MI value indexes the information conveyed by a voxel about stimulus orientation, and provides a principled method for feature selection (Peng et al., 2005). Moreover, this metric has the advantage that it makes no a priori assumptions about the shape of the underlying voxel-based tuning functions (see also Fuhrmann Alpert et al., 2007).

Feature-based attentional modulations (Fig. 7)

To assess the influence of attention on population response profiles, we first assigned each voxel into one of eight bins based on

the orientation that evoked the largest response during the tuning function scans (see previous section). Then, we computed the average evoked response separately for all attend 45° trials and all attend 135° trials (see Fig. 2B) for voxels in each bin over a temporal epoch extending from 6 s to 18 s post-stimulus.

Results

Orientation selectivity in primary visual cortex

Recently, several groups have used fMRI and multivariate pattern analysis (MVPA) methods to estimate feature-selective neural modulations within small regions of human visual cortex (Kamitani and Tong, 2005, 2006; Cox and Savoy, 2003; Haxby et al., 2001; Haynes and Rees, 2005; Norman et al., 2006; Peelen and Downing, 2007; Serences and Boynton, 2007a, b). The basic MVPA approach is based on the assumption that feature-selective subregions of visual cortex such as V1 contain submillimeter columns of neurons that are selective for different stimulus orientations. Functional MRI voxels are quite large in comparison (2×2×1.5 mm in the present study, but more often 3×3×3 mm). However, if a slight preponderance of neurons preferring a particular orientation happens to be sampled within a voxel, then that voxel will exhibit a small but detectable orientation-selective response bias (Kamitani and Tong, 2005). By examining the distributed voxel-by-voxel activation pattern across a visual area such as V1, inferences can be made about changes in the underlying population response profile, and simple classification algorithms can be used to predict the specific feature that an observer is attending or perceiving (Kamitani and Tong, 2005, 2006; Haynes and Rees, 2005; Serences and Boynton, 2007a, b).

Several assumptions must be made in order to use this general MVPA approach to infer changes in population response profiles. First, even though each voxel contains neurons tuned to many different orientations, the final activation level will be dominated by the tuning preference of the most prevalent type of neuron within that voxel. Linear classifiers then pool the information from many weakly selective voxels to discriminate stimulus features such as orientation (see e.g. Fig. 1b from Kamitani and Tong, 2005). Second, since BOLD fMRI fundamentally measures changes in the magnetic properties of blood, it is assumed that the vasculature in early visual areas is sensitive to anisotropies in the distribution of feature-selective neurons.

To evaluate the ability of MVPA methods to discriminate differences in population response profiles, we first assessed sensitivity to fine-grained changes in activation profiles across subregions of early visual cortex. Specifically, patterns evoked by similar stimuli should be more alike than patterns evoked by highly dissimilar stimuli (just as one would expect of the underlying neuronal population response profile, see Fig. 1D). We adapted a task from Kamitani and Tong (2005) in which observers viewed full contrast sinusoidal gratings rendered in one of eight different orientations (in 22.5° increments across 180°) while in the MR scanner (Fig. 2A). Each trial consisted of one oriented grating flickering at 2 Hz (250 ms on, 250 ms off) for 12 s; the slow flicker rate and a randomly selected phase on each presentation cycle attenuated the perception of apparent motion (Kahneman and Wolman, 1970). The stimuli were centered at fixation and subtended 20° visual angle (diameter) with a circular region around fixation masked out (2.5° radius). The subject was to press a button whenever the contrast of the grating dimmed slightly, which occurred at five pseudorandomly determined times during the course of each trial. The magnitude of the target-defining change in contrast was titrated as necessary during scanning to maintain a constant level of task difficulty (mean contrast decrement across observers, ±S.E.M.: 15.75%± 5.9%, mean accuracy, ±S.E.M.: 81.2%±4.7%).

A Support Vector Machine (SVM) was used to classify the stimulus orientation on each trial based on the pattern of activation across all voxels within left and right visual areas that responded to the region of space occupied by the stimulus aperture (as identified using independent functional localizer scans). Data from each hemisphere were always analyzed separately and then averaged because no systematic differences between the reliability of activation patterns in left and right hemispheres were observed. All active voxels from each visual area were included in the analyses (see Table 1).

To perform the classification, we first split the data into two independent sets: all scans except one comprised a 'training' set, and the remaining scans comprised a 'test' set. We then compared the multivoxel activation pattern on each trial in the test set with the activation patterns associated with each of the 8 orientations collapsed across all training scans; each test pattern was then assigned to the stimulus category that it most closely resembled (see Mulitvoxel pattern analysis of orientation selectivity section of Methods). The overall classification accuracy for each observer was defined as the average prediction accuracy across all possible permutations of holding one scan out for use as a test set (which is referred to as the 'hold-one-scan-out cross-validated' estimate of the mean classification accuracy). Since we collected 5 scans, we were able to obtain 5 estimates of classification accuracy for each visual area in each subject. This approach ensures that test and training data sets were always independent, thereby avoiding issues of circularity (Kamitani and Tong, 2005, 2006; Haynes and Rees, 2005; Serences and Boynton, 2007a, b).

Replicating Kamitani and Tong (2005), the classifier correctly identified the stimulus orientation on approximately 43% of trials based on activation patterns in V1, and classification errors most often occurred for orientations adjacent to the true stimulus orientation (Fig. 3A). The Gaussian shape of the classification function demonstrates that stimulus similarity determines the similarity of activation patterns, since activation patterns associated with adjacent stimuli are more easily confusable. Similar patterns of data were also observed in other retinotopically organized visual areas; however, overall classification accuracy was generally lower, perhaps because fewer voxels were identified in these regions (Figs. 3B–D, see Table 1). Accuracy is below chance in the sidebands because the area under each curve in Fig. 3 must sum to 1 (or 100% of all trials). This constraint arises because the classifier only makes one 'guess' per trial so if the algorithm is above chance at guessing the correct orientation, then it must by necessarily pick some other orientation offsets with a relatively low probability.

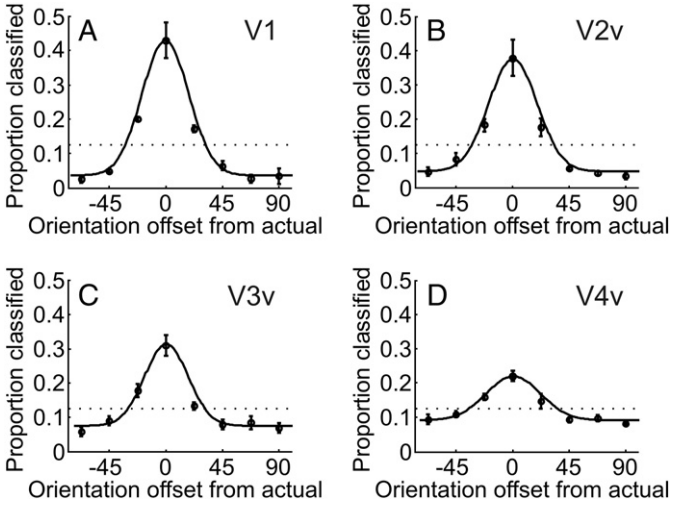

Fig. 3. MVPA of orientation selectivity. (A) Proportion of trials classified based on activation patterns in V1 as a function of the orientation offset from the actual stimulus orientation (black circles), averaged across all subjects (chance = 12.5%). Solid black line is best fitting circular Gaussian. When the classifier made an error, it most often reported an orientation that was adjacent to the actual orientation (see also Kamitani and Tong, 2005). (B–D) Analogous data from areas V2v, V3v, and V4v. All error bars ±S.E.M. across subjects.

The SVM classification approach infers stimulus orientation by computing a weighted sum of inputs from all voxels in a visual area, where the weight assigned to each voxel reflects the degree to which that voxel discriminates a particular orientation. On the one hand, this decoding method is statistically efficient because it accumulates information across all voxels and can powerfully discriminate between activation patterns that are associated with different attentional states. On the other hand, the process of pooling weighted activation levels from each voxel means that one cannot easily make inferences about the specific tuning characteristics of the voxels (and by implication the specific neurons) that are modulated by an attentional manipulation. Thus, using a linear classifier is an extremely powerful method for determining *if there are differences* between activation patterns associated with different attentional states. However, pooling information from many feature-selective voxels makes it difficult to determine exactly *how* feature-selective voxels are modulated by attention or other cognitive factors. In the next section we develop a method for directly examining the tuning functions of individual voxels in an effort to more precisely characterize the influence of attention on population response profiles.

Voxel-based feature-tuning functions

The results shown in Fig. 3 (and those reported by Kamitani and Tong, 2005) demonstrate that activation patterns in regions of early visual cortex carry information about stimulus orientation. However, to make more precise inferences about exactly how the activation patterns changed, we tested to see if voxels exhibited graded activation profiles in response to different levels of a particular feature, just as single neurons in early areas of visual cortex are tuned to specific feature values (e.g. Fig. 1A). Based on data taken from four out of five orientation scans, we sorted each of the voxels in V1 into eight separate bins based on the orientation that maximally stimulated each voxel. Next, we computed the response of the voxels in each bin to the eight oriented gratings presented during the fifth scan; this hold-one-scan-out procedure was then repeated five times and the results from each unique permutation were averaged. As shown in Fig. 4, the orientation preference exhibited by a voxel on four of the scans reliably predicted the orientation preference on the remaining scan, revealing a series of Gaussian-shaped tuning functions representing the response profile of subpopulations of voxels within V1 that respond maximally to each of the different orientations. Fig. 5 shows the average tuning functions across all voxels in each visual area after re-centering all tuning functions on their preferred orientation (solid black lines). Although all regions exhibited reliable voxel-based tuning functions, the amplitude of the tuning functions was generally lower in regions outside of V1, which is consistent with the physiology of monkey striate and extrastriate visual cortex (Vanduffel et al., 2002).

To explore the variability of these tuning functions, we used a mutual information (MI) metric to estimate the extent to which measuring the

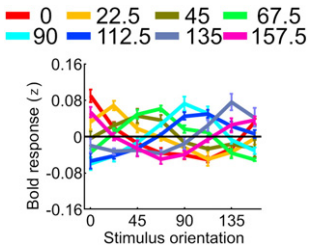

Fig. 4. Voxel-based tuning functions in V1. Each colored line represents the response of a group of V1 voxels to all orientations during one scan after the voxels were binned based on their orientation preference during the other four scans. For instance, the red line shows the response in voxels that respond maximally to 0° when stimulated by all possible orientations.

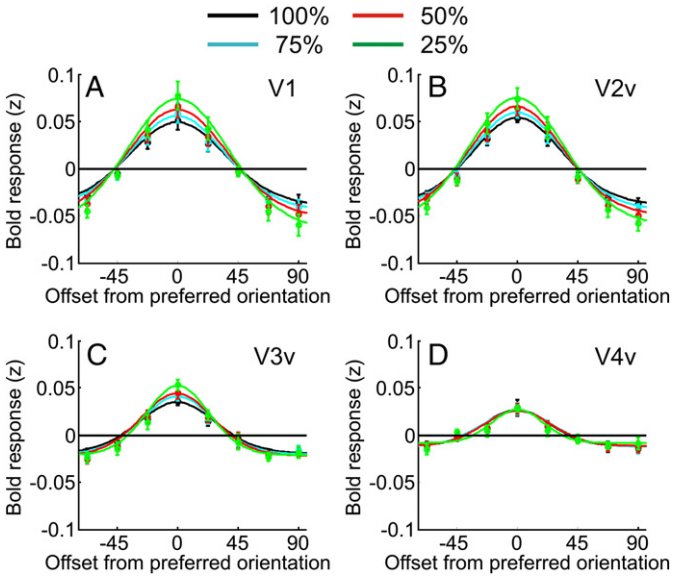

Fig. 5. Mean voxel-based tuning functions. (A) Mean voxel-based tuning function in V1 after centering each curve in Fig. 4 on the preferred orientation (black solid line is best fitting circular Gaussian). (B–D) Analogous plots for areas V2v–V4v. Each color of line represents voxels selected from one of four equally populated bins based on the MI estimates. Thus, the black line represents data using all voxels, and at the other extreme, the green line represents data only from the 25% of voxels with the highest MI. Error bars are ±S.E.M. across subjects.

tuning function of each voxel reduced uncertainty about the orientation of the stimulus being presented. MI was computed using only four scans, and then voxels preferring each orientation were sorted into four equally populated bins based on their MI value (this procedure was then permuted five times using the cross-validation procedure described above). A two-way repeated-measures ANOVA on response amplitude with visual area and MI bin as factors revealed that more informative voxels showed sharper tuning functions compared to voxels that were less informative ($F(3,9)=23.03$, $p<.001$). However, the influence of MI was more pronounced in earlier regions, as voxels in V4v showed uniformly poor orientation tuning functions ($F(9,27)=3.21$, $p<.01$). Together, these data reveal (1) reliable voxel-based orientation tuning curves in regions of early visual cortex, and (2) that the amount of information carried by a voxel about orientation is stable across scans. Note also that the data presented in Figs. 4 and 5 cannot be attributed to random noise because the hold-one-scan-out cross-validation procedure ensures that the observed influence of stimulus orientation and MI on the tuning functions is internally reliable. Finally, a previous report using a similar paradigm did not observer such robust voxel-based tuning functions in early visual cortex (Kamitani and Tong, 2005). The lack of strong voxel tuning instead led those authors to employ a linear classifier to pool information across large groups of voxels in order to maximize their ability to discriminate activation patterns (as in Fig. 3 above). We highlight several potential reasons for this apparent discrepancy between our robust voxel-based tuning functions and previous results in the Discussion section.

We also measured the distribution of orientation tuning preferences across voxels within each visual area. After removing the mean response across all voxel to each orientation (see Methods), a higher proportion of voxels responded maximally to horizontal (0°) stimuli compared to any other orientation; this effect was more pronounced in V1, V2v, and V3v compared to V4v (Fig. 6), and suggests that more neurons within early visual cortex are dedicated to representing the horizontal meridian (although see Supplemental Fig. 2). The over-representation of voxels that respond relatively more to horizontal stimuli is consistent with the organization of primate V1, as well as psychophysical studies that reveal superior performance for detecting

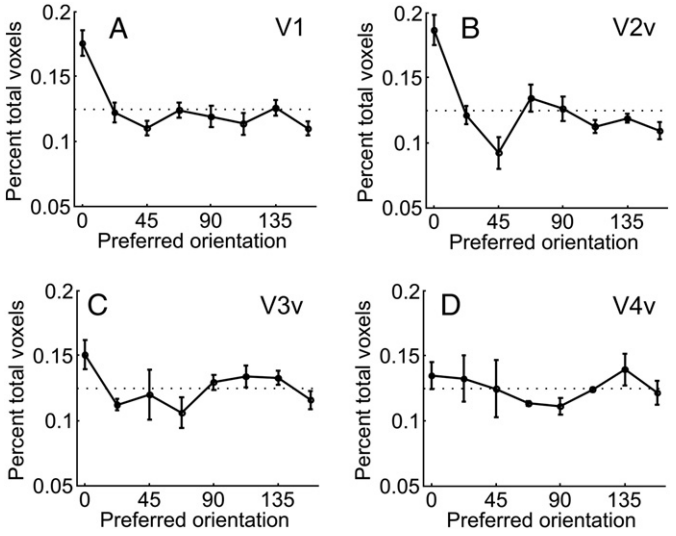

Fig. 6. Orientation tuning histograms. (A–D) Histograms showing the proportion of voxels in each visual area that respond maximally to each stimulus orientation.

stimuli presented on the horizontal meridian compared to oblique orientations (the 'oblique effect', Campbell et al., 1966; Furmanski and Engel, 2000; McMahon and MacLeod, 2003; Orban et al., 1984).

Attentional modulations of population response profiles

In the same scanning session, each of our subjects participated in five 'attention' scans that were interleaved with the orientation tuning scans discussed above. In these attention scans, subjects viewed a grating stimulus comprised of overlapping 45° and 135° lines on each 12 s trial (Fig. 2B). The color of the central fixation point instructed the subjects to attend one set of oriented lines and to press a button whenever the width of the attended lines increased slightly; subjects were required to ignore changes in the width of the unattended lines that occurred with equal frequency. The change in line-width that defined a target was titrated on a scan-by-scan basis to ensure that the task remained challenging (mean change in width ±S.E.M.: 0.1° ±0.02°, mean d-prime ±S.E.M.: 2.75±0.37, 2.84±0.1 for attend 45° and attend 135° trials, respectively). We analyzed the activation level of voxels that responded maximally to each of the eight orientations during epochs of attention to 45° and 135° (where the tuning preference was determined using the independent orientation tuning scans, see Figs. 4 and 5). As shown in Fig. 7A, voxels in primary visual cortex that maximally responded to 45° were most active when the 45° oriented lines were attended, and voxels that maximally responded to 135° were most active when the 135° lines were attended. Voxels with intermediate orientation preferences showed appropriately graded responses. A repeated-measures ANOVA with visual area, attended orientation, and voxel orientation preference as factors revealed a significant interaction between orientation preference and the attended feature across all visual areas ($F(7,21)=2.67$, $p<0.05$). However, this interaction was primarily driven by the crossover modulations in V1 and V4v (individual two-way repeated-measures ANOVA: $F(7,21)=3.37$, $p<0.02$, $F(7,21)=2.5$, $p<0.05$, respectively). The attention effect in area V2v approached significance ($p=0.15$), but the data from area V3v were noisy and no systematic pattern of modulation was present (as evidenced by the very poor fits). Finally, we computed attention effects as a function of MI in each visual area, where MI was determined based on the tuning function scans. However, only data in areas V1 and V4v were stable across all MI bins, most likely because the number of trials in each bin was quite low (Supplemental Fig. 3).

The color of the oriented lines that made up each grating (black or white) was counterbalanced over the course of each scan, so sensory differences in the display cannot account for the attentional modulations. In addition, the specificity of the attention effects is especially robust because the orientation preference of each voxel was assigned using independent scans – thereby avoiding any issue of circularity – and different stimulus sets were used to evaluate orientation preferences and attentional modulations (sinusoidal gratings versus simple oriented lines). One remaining puzzle is the observation that robust attention effects are largely restricted to V1 and V4v, especially given that orientation tuning is more pronounced in areas V2v and V3v compared to V4v (Fig. 5).

Discussion

Stable perceptual representations are thought to be based on pooled activity across small populations of sensory neurons; selective attention is required to ensure that these population codes are biased in favor of behaviorally relevant stimuli. However, few studies have directly examined the influence of attention on the shape of population response profiles, primarily due to technical limitations associated with both electrophysiological recording and fMRI methods. Recently, MVPA methods have demonstrated that attention systematically modulates activation patterns within visual cortex, although these previous reports did not specify precisely how activation patterns were altered (Kamitani and Tong, 2005, 2006; Serences and Boynton, 2007a). The primary limitation of these previous efforts is that classification algorithms (e.g. SVMs) obscure the relationship between changes in global activation patterns and changes in the activation level of voxels tuned to specific features. Here we address this limitation by demonstrating that voxels show graded response profiles that are similar in form to the tuning functions that are commonly reported in the single-unit recording literature. Then in separate scans we use these voxel-based tuning functions to explicitly demonstrate that attention biases population response profiles in favor of a relevant feature. This increase in activation presumably reflects an increase in the signal-to-noise ratio of neurons tuned to the attended stimulus attribute, ensuring that the relevant stimulus is represented more robustly in visual cortex (Boynton, 2005; Desimone and Duncan, 1995; Martinez-Trujillo and Treue, 2004; Pouget et al., 2001).

In a previous report that used a similar stimulation protocol, very weak orientation tuning functions were observed for single voxels

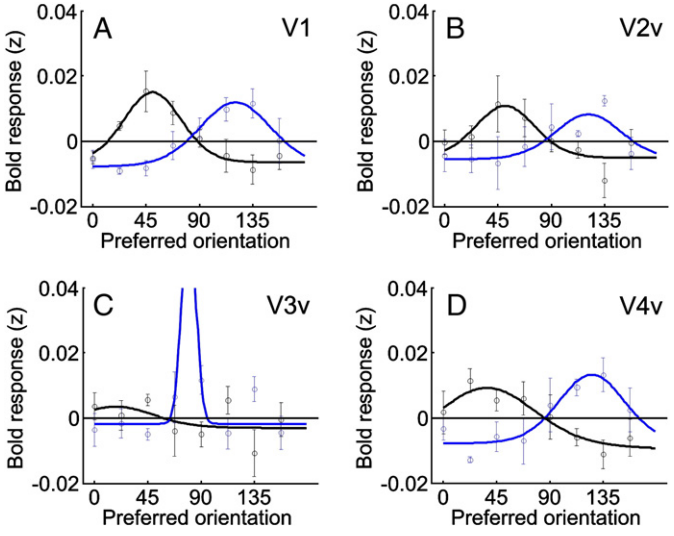

Fig. 7. Feature-based attentional modulations of population response profiles. (A) BOLD response plotted as a function of the attended orientation (attend 45°=black, attend 135°=blue) and the preferred orientation of each voxel in V1. Solid lines are best fitting circular Gaussians. (B–D) Analogous plots for areas V2v–V4v. Error bars ±S.E.M. across subjects.

(see Kamitani and Tong, 2005, their Supplemental Fig. 2). Due to this weak orientation tuning, the authors used a linear classifier to pool information across all voxels to maximize their ability to discriminate the currently viewed orientation. In contrast, the present voxel-based tuning functions provide a more direct method to examine the role of attention in modulating population response profiles. While the exact reason for the discrepancy between our results and previous efforts is not entirely clear, two potential contributing factors should be considered. First, we used higher-resolution imaging (2×2×1.5 mm voxels) compared to the 3×3×3 mm voxels used by Kamitani and Tong (2005). The use of higher-resolution imaging may better isolate signal from orientation columns in areas like V1, which may give rise to more pronounced feature tuning at the voxel level. In addition, smaller voxels may increase BOLD SNR because less of the volume in each cortical voxel is occupied by white matter, which does not yield as robust a signal (Hyde et al., 2001). Second, we minimized the amount of spatial smoothing that was performed on the fMRI data in our study. While Kamitani and Tong (2005) did not explicitly smooth their data with a Gaussian kernel as is often done, their functional images were aligned with retinotopic mapping data collected in a separate session, translated into the ACPC plane, and then converted into Talairach space. Each time an image matrix is transformed, rotated, or scaled, the data are resampled and thus some spatial smoothing is introduced. We avoided all three of these transformation steps in our analysis by aligning our retinotopic mapping data to our experimental data and by performing all analyses in native scanner space (i.e. no ACPC or Talairach transformations). This issue of spatial smoothing may turn out to be critical because smoothing should effectively average the response of neighboring voxels, which should in turn attenuate orientation tuning functions.

Treue and Maunsell (1999) and Martinez-Trujillo and Treue (2004) demonstrated that feature-based attention boosts the gain of single neurons tuned to an attended feature. By extension, these single-unit results suggest that feature-based attention should bias the entire population response profile in favor of neurons tuned to an attended feature, and that attentional gain should fall off in a roughly Gaussian manner as a function of the distance between the preferred feature of a neuron and the current focus of attention. However, this predicted influence of attention on the shape of population codes has not yet been directly demonstrated: the present data provide this missing link by estimating changes in the response profile across groups of orientation-selective voxels.

Attention may influence cortical activity in a number of ways. First, attention might enhance the firing rate of all neurons that are tuned to the attended feature, even before a stimulus is presented (analogous to spatial 'baseline shifts', Kastner et al., 1999; Luck et al., 1997). Second, attention has been shown to induce a multiplicative increase in the stimulus evoked response (e.g. Martinez-Trujillo and Treue, 2004; McAdams and Maunsell, 1999; Treue and Martinez-Trujillo, 1999; Williford and Maunsell, 2006). In the present report, we cannot distinguish these accounts because we did not measure pre-stimulus attentional modulations timelocked to the presentation of the attention cue, and we did not map out the effects of attention on the full voxel-based tuning functions. However, future studies might adopt our voxel-based tuning function approach to inform this long standing debate about the nature of attentional modulations on neural activity in visual cortex.

The feature-selective attention effects shown in Fig. 7 are also reminiscent of attentional modulations previously reported within anatomically segregated subregions of ventral temporal cortex that are selectively responsive to different object types (e.g. faces and houses, O'Craven et al., 1999; Serences et al., 2004). However, the feature-based attentional modulations reported here reflect changes in the response profile within a single population of feature-selective voxels. Thus, with the advent of MVPA methods and voxel-based tuning functions, it is now possible to move beyond distinguishing stimulus categories (e.g. faces versus houses). Instead, precise inferences can be made about how population response profiles change in response to specific tuning function properties, stimulus features, or category exemplars (e.g. individual faces, natural scenes, etc. Dumoulin and Wandell, 2008; Kamitani and Tong, 2005, 2006; Kay et al., 2008; Williams et al., 2007).

Consistent with the known orientation selectivity of V1 neurons in non-human primates, activation patterns in area V1 discriminated orientations better than activation patterns in other areas (Fig. 3A) and also had sharper tuning functions (Fig. 5A). However, the observation of relatively large attentional modulations in V1 contradicts some single-unit recording studies that show modest (or absent) attention effects in V1, including a recent study that recorded spiking activity directly on the cortical surface of V1 in awake behaving humans (Yoshor et al., 2007). The reason for this discrepancy is not entirely clear; Yoshor et al. suggest that behavioral demands may alter the linearity of the relationship between neural activity and the BOLD response, thereby inflating the estimated magnitude of attentional modulations in human visual cortex. However, it also seems plausible that we observed robust effects in V1 simply because V1 voxels were highly selective for the stimulus attribute that was being attended (see Fig. 3–5). We cannot distinguish these possibilities here, but examining the magnitude of attention effects in visual areas that are associated with processing specific features – such as motion selective area MT – might reveal a systematic relationship between attentional modulations and the underlying feature-selectivity of a visual area.

The attentional modulations presented in Fig. 7 are consistent with the two leading models of feature-based attention: the *feature-similarity gain* model proposed by Martinez-Trujillo and Treue (2004) and Treue and Martinez Trujillo (1999), and the *optimal gain model* proposed by Navalpakkam and Itti (2007). According to both accounts, the firing rates of neurons tuned to the currently attended feature should be enhanced – consistent with the data presented in Fig. 7 – because these neurons most effectively discriminate the highly dissimilar (orthogonal) target and distractor that were used in the present study (Butts and Goldman, 2006; Hol and Treue, 2001; Jazayeri and Movshon, 2006, 2007; Navalpakkam and Itti, 2007; Regan and Beverley, 1985). However, if a fine perceptual discrimination is required (e.g. distinguish a 95° target from a 90° distractor), then boosting the gain of neurons tuned to the target orientation would be suboptimal because those neurons would respond about equally well to the target and to the distractor (Regan and Beverley, 1985; Jazayeri and Movshon, 2006; Butts and Goldman, 2006). Instead, the activation peaks in Fig. 7 should be shifted towards orientations tuned slightly away from the target because those neurons are more informative, a prediction of the optimal gain model that can be tested in future studies using voxel-based tuning functions.

The data presented here demonstrate that attention biases population response profiles in favor of an attended stimulus; however, we cannot make any statements about how a subject 'reads out' this population code to infer the physical properties of the display. On the one hand, subjects may base their interpretation of the scene on the firing rate and the tuning function of the maximally active sensory neurons (e.g. Zhaoping and May, 2007). While applying such a 'max' rule is computationally inexpensive, the resulting estimate might be relatively unreliable given that the firing rate of sensory neurons is highly variable on a trial-by-trial basis (see Fig. 1B). On the other hand, a statistically optimal maximum likelihood estimate can be derived by computing the probability that an orientation was present given the observed vector of responses recorded from each neuron in the population (Jazayeri and Movshon, 2006; Pouget et al., 2003; Seung and Sompolinsky, 1993). While this optimal read-out rule will be, on average, much more precise than a max rule, it is also potentially more cumbersome to compute. We speculate that the exact read-out rule might depend critically on the nature of the behavioral task being performed, and our analysis approach might eventually yield insights

into how population codes are influenced by attention and read-out to achieve stable and coherent perceptual representations of the surrounding environment.

Acknowledgments

We thank Alyssa Brewer and Emily Grossman for their helpful comments.

Author contributions: J.S., T.H., and M.S. designed the experiment; J.S. and S.S. analyzed the data; S.S. wrote code for mutual information analysis; L.T.M. designed high-resolution pulse sequences; and J.S., M.S., S.S., and T.H. wrote the paper.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2008.07.043.

References

Borst, A., Theunissen, F.E., 1999. Information theory and neural coding. Nat. Neurosci. 2, 947–957.
Boynton, G.M., 2005. Attention and visual perception. Curr. Opin. Neurobiol. 15, 465–469.
Boynton, G.M., Engel, S.A., Glover, G.H., Heeger, D.J., 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. J. Neurosci. 16, 4207–4221.
Butts, D.A., Goldman, M.S., 2006. Tuning curves, neuronal variability, and sensory coding. PLoS Biol. 4, e92.
Campbell, F.W., Kulikowski, J.J., Levinson, J., 1966. The effect of orientation on the visual resolution of gratings. J. Physiol. 187, 427–436.
Cover, T.M., Thomas, J.A., 1991. Elements of Information Theory. Wiley, New York.
Cox, D.D., Savoy, R.L., 2003. Functional magnetic resonance imaging (fMRI) "brain reading": detecting and classifying distributed patterns of fMRI activity in human visual cortex. Neuroimage 19, 261–270.
Desimone, R., Duncan, J., 1995. Neural mechanisms of selective visual attention. Annu. Rev. Neurosci. 18, 193–222.
Dumoulin, S.O., Wandell, B.A., 2008. Population receptive field estimates in human visual cortex. Neuroimage 39, 647–660.
Engel, S.A., Rumelhart, D.E., Wandell, B.A., Lee, A.T., Glover, G.H., Chichilnisky, E.J., Shadlen, M.N., 1994. fMRI of human visual cortex. Nature 369, 525.
Fuhrmann Alpert, G., Sun, F.T., Handwerker, D., D'Esposito, M., Knight, R.T., 2007. Spatio-temporal information analysis of event-related BOLD responses. Neuroimage 34, 1545–1561.
Furmanski, C.S., Engel, S.A., 2000. An oblique effect in human primary visual cortex. Nat. Neurosci. 3, 535–536.
Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., Massey, J.T., 1982. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J. Neurosci. 2, 1527–1537.
Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., Pietrini, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science 293, 2425–2430.
Haynes, J.D., Rees, G., 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. Nat. Neurosci. 8, 686–691.
Hol, K., Treue, S., 2001. Different populations of neurons contribute to the detection and discrimination of visual motion. Vision Res. 41, 685–689.
Hyde, J.S., Biswal, B.B., Jesmanowicz, A., 2001. High-resolution fMRI using multislice partial k-space GR-EPI with cubic voxels. Magn. Reson. Med. 46, 114–125.
Jazayeri, M., Movshon, J.A., 2006. Optimal representation of sensory information by neural populations. Nat. Neurosci. 9, 690–696.
Jazayeri, M., Movshon, J.A., 2007. A new perceptual illusion reveals mechanisms of sensory decoding. Nature 446, 912–915.
Kahneman, D., Wolman, R.E., 1970. Stroboscopic motion: effects of duration and interval. Percept. Psychophys. 8, 161–164.
Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. Nat. Neurosci. 8, 679–685.
Kamitani, Y., Tong, F., 2006. Decoding seen and attended motion directions from activity in the human visual cortex. Curr. Biol. 16, 1096–1102.
Kay, N.K., Naselaris, T., Prenger, R.J., and Gallant, J.L., 2008. Identifying natural images from human brain activity. Nature 452 (7185), 352–355.
Kastner, S., Pinsk, M.A., De Weerd, P., Desimone, R., Ungerleider, L.G., 1999. Increased activity in human visual cortex during directed attention in the absence of visual stimulation. Neuron 22, 751–761.
Luck, S.J., Chelazzi, L., Hillyard, S.A., Desimone, R., 1997. Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. J. Neurophysiol. 77, 24–42.
Ma, W.J., Beck, J.M., Latham, P.E., Pouget, A., 2006. Bayesian inference with probabilistic population codes. Nat. Neurosci. 9, 1432–1438.
Martinez-Trujillo, J.C., Treue, S., 2004. Feature-based attention increases the selectivity of population responses in primate visual cortex. Curr. Biol. 14, 744–751.
McAdams, C.J., Maunsell, J.H., 1999. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. J. Neurosci. 19, 431–441.
McMahon, M.J., MacLeod, D.I., 2003. The origin of the oblique effect examined with pattern adaptation and masking. J. Vis. 3, 230–239.
Navalpakkam, V., Itti, L., 2007. Search goal tunes visual features optimally. Neuron 53, 605–617.
Norman, K.A., Polyn, S.M., Detre, G.J., Haxby, J.V., 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends. Cogn. Sci. 10, 424–430.
O'Craven, K.M., Downing, P.E., Kanwisher, N., 1999. fMRI evidence for objects as the units of attentional selection. Nature 401, 584–587.
Orban, G.A., Vandenbussche, E., Vogels, R., 1984. Human orientation discrimination tested with long stimuli. Vision Res. 24, 121–128.
Paradiso, M.A., 1988. A theory for the use of visual orientation information which exploits the columnar structure of striate cortex. Biol. Cybern. 58, 35–49.
Peelen, M.V., Downing, P.E., 2007. Using multi-voxel pattern analysis of fMRI data to interpret overlapping functional activations. Trends. Cogn. Sci. 11, 4–5.
Peng, H., Long, F., Ding, C., 2005. Feature selection based on mutual information: criteria of max-dependency, max-relevance, and min-redundancy. IEEE Trans. Pattern Anal. Mach. Intell. 27, 1226–1238.
Pouget, A., Dayan, P., Zemel, R., 2000. Information processing with population codes. Nat. Rev. Neurosci. 1, 125–132.
Pouget, A., Dayan, P., Zemel, R.S., 2003. Inference and computation with population codes. Annu. Rev. Neurosci. 26, 381–410.
Pouget, A., Deneve, S., Latham, P.E., 2001. The relevance of Fisher Information for theories of cortical computation and attention. In: Braun, J., Koch, C., Davis, J.L. (Eds.), In Visual attention and cortical circuits. MIT Press, Cambridge.
Regan, D., Beverley, K.I., 1985. Postadaptation orientation discrimination. J. Opt. Soc. Am. A. 2, 147–155.
Reynolds, J.H., Desimone, R., 1999. The role of neural mechanisms of attention in solving the binding problem. Neuron 24, 19–29 111–125.
Reynolds, J.H., Pasternak, T., Desimone, R., 2000. Attention increases sensitivity of V4 neurons. Neuron 26, 703–714.
Sanger, T.D., 1996. Probability density estimation for the interpretation of neural population codes. J. Neurophysiol. 76, 2790–2793.
Serences, J.T., Boynton, G.M., 2007a. Feature-based attentional modulations in the absence of direct visual stimulation. Neuron 55, 301–312.
Serences, J.T., Boynton, G.M., 2007b. The representation of behavioral choice for motion in human visual cortex. J. Neurosci. 27, 12893–12899.
Serences, J.T., Schwarzbach, J., Courtney, S.M., Golay, X., Yantis, S., 2004. Control of object-based attention in human cortex. Cereb. Cortex 14, 1346–1357.
Sereno, M.I., Dale, A.M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., Tootell, R.B., 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science 268, 889–893.
Seung, H.S., Sompolinsky, H., 1993. Simple models for reading neuronal population codes. Proc. Natl. Acad. Sci. U. S. A. 90, 10749–10753.
Shadlen, M.N., Britten, K.H., Newsome, W.T., Movshon, J.A., 1996. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J. Neurosci. 16, 1486–1510.
Treue, S., Martinez Trujillo, J.C., 1999. Feature-based attention influences motion processing gain in macaque visual cortex. Nature 399, 575–579.
Tsotsos, J.K., 1997. Limited Capacity of Any Realizable Perceptual System Is a Sufficient Reason for Attentive Behavior. Conscious Cogn. 6, 429–436.
Vanduffel, W., Tootell, R.B., Schoups, A.A., Orban, G.A., 2002. The organization of orientation selectivity throughout macaque visual cortex. Cereb. Cortex 12, 647–662.
Williams, M.A., Dang, S., Kanwisher, N.G., 2007. Only some spatial patterns of fMRI response are read out in task performance. Nat. Neurosci. 10, 685–686.
Williford, T., Maunsell, J.H., 2006. Effects of spatial attention on contrast response functions in macaque area V4. J. Neurophysiol. 96, 40–54.
Yoshor, D., Ghose, G.M., Bosking, W.H., Sun, P., Maunsell, J.H., 2007. Spatial attention does not strongly modulate neuronal responses in early human visual cortex. J. Neurosci. 27, 13205–13209.
Zhaoping, L., May, K.A., 2007. Psychophysical tests of the hypothesis of a bottom-up saliency map in primary visual cortex. PLoS Comput. Biol. 3, e62.