

Basing Perceptual Decisions on the Most Informative Sensory Neurons

Miranda Scolari and John T. Serences

Department of Psychology, University of California, San Diego

Submitted 17 March 2010; accepted in final form 13 July 2010

Scolari M, Serences JT. Basing perceptual decisions on the most informative sensory neurons. *J Neurophysiol* 104: 2266–2273, 2010. First published July 14, 2010; doi:10.1152/jn.00273.2010. Single unit recording studies show that perceptual decisions are often based on the output of sensory neurons that are maximally responsive (or “tuned”) to relevant stimulus features. However, when performing a difficult discrimination between two highly similar stimuli, perceptual decisions should instead be based on the activity of neurons tuned away from the relevant feature (*off-channel* neurons) as these neurons undergo a larger firing rate change and are thus more informative. To test this hypothesis, we measured feature-selective responses in human primary visual cortex (V1) using functional magnetic resonance imaging and show that the degree of off-channel activation predicts performance on a difficult visual discrimination task. Moreover, this predictive relationship between off-channel activation and perceptual acuity is not simply the result of extensive practice with a specific stimulus feature (as in studies of perceptual learning). Instead, relying on the output of the most informative sensory neurons may represent a general, and optimal, strategy for efficiently computing perceptual decisions.

INTRODUCTION

Early sensory areas such as primary visual cortex (V1) play a critical role in encoding information about the low-level features of behaviorally relevant stimuli in the environment (e.g., information about object edges, colors, motion, etc.). This low-level featural information must then be interpreted— or decoded—by downstream brain regions involved in decision making and action planning (Gold and Shadlen 2007). In a typical perceptual decision-making experiment, observers view a noisy stimulus (usually a field of moving dots) and indicate a stimulus attribute (e.g., whether a subset of dots is moving to the left or right) by making a saccade to a prespecified location (Gold and Shadlen 2000; Newsome and Pare 1988; Newsome et al. 1989; Roitman and Shadlen 2002). The amount of evidence that supports each decision outcome is proportional to the firing rates of feature-selective neurons in early visual cortex that are tuned to each stimulus alternative (Britten et al. 1996; Ditterich et al. 2003; Gold and Shadlen 2001, 2007; Mazurek et al. 2003; Newsome et al. 1989; Salzman et al. 1992; Shadlen and Newsome 2001; Shadlen et al. 1996). This sensory information is then accumulated by oculomotor neurons in higher-order areas such as the lateral intraparietal area (LIP), dorsal-lateral prefrontal cortex (DLPFC), frontal eye field (FEF), and the superior colliculus (SC) until a decision is reached and an eye movement is executed (Gold and Shadlen 2001, 2003, 2007; Hanes and Schall 1996; Horwitz et al. 2004; Leon and Shadlen 1999; Roitman and Shadlen 2002; Schall 2001; Shadlen and Newsome 2001). Recent evidence also

suggests that similar accumulation processes occur when other response modalities are used to indicate the outcome of the perceptual decision (e.g., manual movements: Heekeren et al. 2004, 2006; Ho et al. 2009; Romo and Salinas 2003).

Here we focus on understanding the optimality with which sensory evidence is represented in early visual cortex during the decision making process. Sensory neurons are typically corrupted by Poisson noise (or near-Poisson noise) (Mitchell et al. 2007; Shadlen and Newsome 1994), so increasing the gain of a cell should generally increase the signal-to-noise ratio (SNR) and thus the capacity to accurately represent information about relevant sensory features. Accordingly, increasing the physical salience of a stimulus leads to higher spiking rates in sensory neurons and to a corresponding increase in both the speed and accuracy of perceptual decisions (Newsome et al. 1989; Roitman and Shadlen 2002). In addition, using microstimulation to directly increase the gain of sensory neurons biases decisions in favor of the feature that drives a maximal response in the stimulated neurons (Ditterich et al. 2003; Salzman et al. 1992). In combination with computational studies (Beck et al. 2008; Gold and Shadlen 2002; Mazurek et al. 2003), these data collectively suggest that the rate of evidence accumulation is directly tied to the relative spiking rates of sensory neurons tuned to each possible stimulus alternative.

While existing studies are *consistent* with the notion that perceptual decisions are based on the gain of neurons tuned to the relevant sensory feature, these empirical studies typically required discriminations between highly dissimilar stimuli such as opposite directions of motion. When performing such *coarse* discriminations, basing decisions on the maximally responsive neurons is indeed optimal because these neurons best signal which of the two stimuli are present (Fig. 1A). However, in other instances, this *on-channel* strategy is sub-optimal, particularly when an observer is faced with a challenging discrimination that requires distinguishing small differences between two visual stimuli (e.g., a radiologist discriminating between cancerous and normal tissue in a noisy X-ray). In such situations, decisions should counterintuitively be guided by neurons tuned away from the critical stimulus feature because these *off-channel* neurons undergo a larger firing rate change in response to the different stimulus alternatives (Fig. 1B) (Beck et al. 2008; Hol and Treue 2001; Jazayeri and Movshon 2006, 2007a,b; Law and Gold 2009; Navalpakkam and Itti 2007; Purushothaman and Bradley 2005; Regan and Beverley 1985; Scolari and Serences 2009; Schoups et al. 2001). In particular, neurons for which the relevant stimulus alternatives fall near the steepest part of the tuning function—or the point where the derivative of the tuning function is highest—undergo the largest change in firing rate and therefore provide the most useful information for perceptual

Address for reprint requests and other correspondence: M. Scolari, University of California, 9500 Gilman Dr. #0109, La Jolla, CA, 92093-0109 (E-mail: mscolari@ucsd.edu).

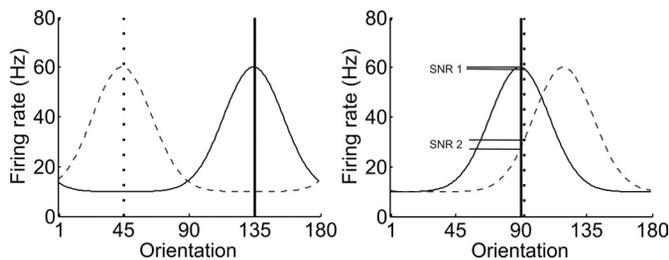


FIG. 1. *A*: relying on the most responsive neurons is optimal during a coarse discrimination task because they respond maximally to the relevant feature and minimally to the irrelevant feature(s), thus resulting in a high signal-to-noise ratio (SNR). *B*: when performing a difficult fine discrimination, a neuron tuned to the target feature (—) does not discriminate stimulus alternatives very well (SNR-1). However, a neuron tuned to a flanking off-channel orientation (---) undergoes a large change in firing rate because its tuning function has a steeper slope at the sample orientation (SNR-2). ···, the target (90°) and the distractor(s) (92°). Adapted with permission from Fig. 4 of Navalpakkam and Itti (2007).

decision-making. Thus the nature of the evidence on which decisions are based should adaptively vary with task demands: a coarse discrimination should primarily be based on the activity of on-channel neurons, whereas a difficult fine discrimination should be based on the activity of off-channel neurons. Previous studies that used difficult visual discrimination tasks generally support this theoretical framework (Hol and Treue 2001; Jazayeri and Movshon 2007a,b; Purushothaman and Bradley 2005; Navalpakkam and Itti 2007; Regan and Beverley 1985; Schoups et al. 2001; Scolari and Serences 2009).

To test the hypothesis that the relative activation level of off-channel neural populations in early visual cortex determines performance on a difficult delayed-match-to-sample (DMTS) orientation-discrimination task, we used human observers, fMRI, and feature-selective voxel tuning functions (VTFs) (Serences et al. 2009; see also: Kamitani and Tong 2005; Kay et al. 2008). We focused on feature-selective response profiles in V1 given the high density of orientation-selective neurons and prior evidence of pronounced orientation selectivity as measured using fMRI (Haynes and Rees 2005; Kamitani and Tong 2005; Kay et al. 2008; Serences et al. 2009). We reasoned that increasing the gain of informative off-channel neurons should increase the slope of their respective tuning functions, rendering these neurons more sensitive to detect small changes between the two discriminants in the DMTS task. Thus we predicted that the probability of a successful discrimination should be correlated with the magnitude of gain in these off-channel neural populations.

METHODS

Subjects

Nine neurologically healthy subjects (3 males and 6 females) with an age range of 22–31 yr [25.2 ± 2.9 (SD)] were recruited from the University of California, San Diego (UCSD) community. Each subject gave written informed consent as per Institutional Review Board requirements at UCSD and completed 1 h of training outside the scanner prior to one 2 h scanning session. Compensation for participation was \$10/h for training and \$20/h for scanning. The design and execution of this study conformed to the ethical guidelines of The American Physiological Society.

Training session

All subjects were trained on the fine discrimination task for five to eight blocks outside of the scanner (outlined in the following text). Subjects received auditory feedback in the form of beeps on each trial to aid in training (the beeps were not presented during the scan session where feedback was only given at the end of each scan). The offset between the sample and test was titrated based on performance via a staircasing procedure so that an appropriate offset could be chosen for the scan session. Subjects also practiced one block of the contrast dimming task used during the functional localizer scan (described in the following text). The stimuli used in both the training and scanning sessions were generated using the Matlab programming language (v7.6, Natick, MA) with the Psychophysics Toolbox (version 3) (Brainard 1997; Pelli 1997), and the stimulus used for retinotopic mapping (see following text) was generated using custom libraries and the C++ programming language.

Fine discrimination scans

Subjects were instructed to maintain fixation on a central green fixation point (subtending 0.4°) that remained on the screen for the duration of each scan. On each fine discrimination trial, a full contrast grayscale square-wave grating (1.2 cycle/°, and subtending 10° diam with a 1.25° diam cutout around fixation) was flickered at 5 Hz (100 ms on, 100 ms off) for 2 s in 1 of 10 possible sample orientations that were evenly distributed across 180° (0, 18, 36, 54, 72, 90, 108, 126, 144, and 162° ; this first stimulus will henceforth be referred to as the *sample stimulus*). A small and randomly selected jitter of between 0 and 3° was then either added or subtracted from the selected sample orientation. The order of sample orientations was randomized on each scan with the constraint that the same orientation could not be presented on successive trials. Following a 400 ms blank period, a second square-wave grating (termed the *test stimulus*) onset with the same physical and temporal parameters listed in the preceding text except it was rotated approximately $\pm 2.2^\circ$ from the sample stimulus (on average across subjects; the offset was titrated on an individual basis so that performance remained at $\sim 65\%$ correct over the course of the scanning session; see RESULTS). Subjects indicated whether the test stimulus was rotated clock- or counterclockwise from the sample; button-press responses made within 400 ms of test offset were considered valid, yielding a response window of 2.4 s because the subjects could also respond when the test stimulus was present on the screen. The spatial phases of the sample and test gratings were randomly selected to ensure that subjects could not perform the discrimination task simply by monitoring a single image pixel (or small region of the display). Subjects completed 50 fine discrimination trials per scan (5 per orientation), and each trial was separated by a blank 200 ms intertrial interval after the response window expired. Eight null trials were randomly presented during a single scan in which only the fixation point was presented for 5 s (the total length of a single fine discrimination trial). Subjects completed eight fine discrimination scans, each of which lasted 300 s (including a 10 s passive fixation period at the end of each scan).

Functional localizer scans

Independent functional localizer scans were run on each subject to identify voxels within retinotopically organized visual areas that responded to the spatial position occupied by the stimulus aperture in the fine discrimination scans. A full contrast checkerboard stimulus flickered for 10 s at 5 Hz. As in the fine discrimination scans, the stimulus subtended 10° visual angle with a small circular cutout (1.25°) around a white fixation point (0.4° diam). Subjects were instructed to make a button-press response when the contrast of the stimulus decreased slightly, hereby referred to as a target. The contrast reduction that defined a target was titrated on an individual basis so

that performance remained above chance and below ceiling (~75% hit rate). Each target was presented for two video frames (33.33 ms), and there were three targets presented on each trial. The timing of each target was pseudorandomly determined with the following constraints: each target was separated from the previous one by ≥ 1 s and targets were restricted to a window of 1–9 s following the onset of the trial. Button-press responses made within 1 s of target onset were considered correct. Each trial was followed by a blank 10 s inter-trial-interval. Observers completed two localizer scans during the scan session; each scan contained 15 trials and lasted for 312 s.

Retinotopic mapping procedures

Retinotopic mapping data were obtained in one or two scans per subject 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 hV4 (we did not have adequate coverage of dorsal visual areas in some subjects due to the fMRI data acquisition parameters, see following text). To aid in the visualization of early visual cortical areas, we projected the retinotopic mapping data onto a computationally inflated representation of each subject's gray/white matter boundary.

fMRI data acquisition and analysis

Magnetic resonance imaging (MRI) scanning was carried out on a GE Excite HDx 3-Tesla scanner equipped with an eight-channel head coil at the Keck Center for Functional MRI, University of California, San Diego. Anatomical images were acquired using a T1-weighted sequence that yielded images with a 1 mm^3 resolution (TR/TE = 11/3.3 ms, TI = 1,100 ms, 163 slices, flip angle = 18°). Functional images were acquired using a gradient echo EPI pulse sequence which covered the occipital lobe with either 24 or 28 transverse slices. Slices were acquired in ascending interleaved order with 3 mm thickness [TR = 1,500 ms, TE = 30 ms, flip angle = 90° , image matrix = 64 (AP) \times 64 (RL), with FOV = 192 mm (AP) \times 192 mm (RL), ASSET factor = 2, voxel size = $3 \times 3 \times 3$ mm].

Data analysis was performed using BrainVoyager QX (v 1.91; Brain Innovation, Maastricht, The Netherlands) and custom time series analysis routines written in Matlab (version 7.1; The Math Works, Natick, MA). EPI images were slice-time corrected, motion-corrected (both within and between scans), and high-pass filtered (3 cycle/run) to remove low frequency temporal components from the time series. To identify voxels that responded to the retinotopic position of the stimulus aperture, data from the functional localizer scan were analyzed using a GLM that contained a regressor marking each 10 s stimulus epoch (a boxcar convolved with a gamma function as implemented in Brain Voyager). Voxels within each visual area were included in all subsequent analyses if they passed a threshold of $P < 0.01$, corrected for multiple comparisons using the false discovery rate (FDR) algorithm implemented in Brain Voyager.

Voxel-based tuning functions

The general methods that we used have also been described elsewhere in more detail (Serences et al. 2009). Before computing tuning functions, the time series from every voxel was z -normalized on a scan-by-scan basis to have zero mean and unit variance. The magnitude of the response in each voxel was then estimated using a GLM with a separate regressor for each orientation (each regressor was formed by convolving a boxcar model of the stimulus presentation sequence with a gamma function as implemented in Brain Voyager). Based on data from all scans except one, each voxel from a visual area was assigned to an orientation-preference bin based on the orientation that evoked the largest response after removing the mean response across all voxels to correct for main effects that had a common

influence on the response of every voxel (Haxby et al. 2001; Serences et al. 2009). Then, using only data from the remaining scan, we computed the response of voxels in each orientation-preference bin to each of the possible sample stimulus orientations. This hold-one scan cross-validation procedure was repeated across all eight unique permutations (given that there were 8 scans/subject) to avoid issues of circularity that arise when the same data set is used to both assign orientation preference and evaluate the shape of the feature-selective response profile. Finally, the response profile of each voxel was re-centered so that their preferred orientation was assigned to 0° . Data for each hemisphere were analyzed separately and then averaged to form an estimate of the *population response profile* across neural subpopulations tuned to different angles with respect to the sample stimulus (e.g., Figs. 3 and 4).

Characterizing the shape of VTFs with half-Gaussians

Before fitting a function to characterize the shape of the VTF in each subject, we re-centered tuning functions from voxels selective for different orientations, averaged them together (collapsing across correct and incorrect trials), and then averaged across equivalent offsets on either side of 0° to improve the signal-to-noise ratio (i.e., averaging data from ± 18 , ± 36 , etc., see Fig. 3B). We then fit the data from each subject with a half-Gaussian of the form

$$B(\Theta) = \alpha + \beta * f(\kappa, \Theta) \quad (1A)$$

where

$$f(\kappa, \Theta) = \exp(-\{[(\pi/2 + \Theta - \Theta_{\text{pref}}) \bmod \pi] - \pi/2\}^2 / (2 * \sigma^2)), x \geq 0 \quad (1B)$$

where $f(\cdot)$ is a half-Gaussian function (only evaluated for positive values of x), α is the baseline offset parameter, β is the multiplicative gain parameter, and σ is the SD. To avoid local minima during the fitting procedure, we first performed a global fit using a series of GLMs with a regressor based on a half-Gaussian function derived using the equation in the preceding text and with σ varying across a range from 0.1 to 180° in steps of 1° (i.e., 1 GLM was run for each level of σ). Each GLM produced an estimate of the amplitude and baseline offset (via the regression coefficients for the half-Gaussian and the constant term, respectively) for the half-Gaussian with the specified value of σ . The parameters that resulted in the best fit to the data (lowest root mean square error) across all tested GLMs were then used to seed a gradient descent algorithm implemented in the Matlab Optimization Toolbox to fine tune the model parameters to achieve the best possible fit (built around the “fminsearch” function that employs Nelder-Mead Simplex Direct Search) (Lagarias et al. 1998). Once the multiplicative, baseline, and SD parameters were estimated, the full width at half-maximum (FWHM) estimate of the best fitting half-Gaussian was calculated as

$$\text{FWHM} = 2\sqrt{2 \ln 2} \sigma \quad (2)$$

Note that the FWHM of a half-Gaussian is technically undefined, but because we collapsed across negative and positive values simply to improve the SNR of the data, the FWHM still provides a useful characterization of the width of the curve.

Relating changes in VTF gain with changes in sensitivity

The degree to which a particular neural population supports fine discriminations should depend on at least two factors: the slope of their tuning functions at the target orientation and the degree of variability, or noise, in their tuning functions at the target orientation. Intuitively, steep slopes should translate into large changes in activation levels in response to the sample and test stimuli (see Fig. 1B). On the other hand, high variance should undermine this sensitivity, even if the slope is relatively steep. This relationship between slope and

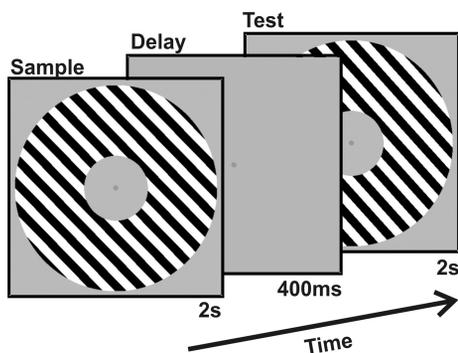


FIG. 2. Schematic of the experimental paradigm. Subjects indicated whether the test grating was rotated clockwise or counterclockwise from the sample grating. The sample stimulus was rendered at an orientation selected from a set of 10 values evenly spanning 180° , plus an additional offset randomly selected from a range of $\pm 3^\circ$. This jitter encouraged subjects to learn a method of discrimination that could be applied to any arbitrary orientation, rather than learning to discriminate a single particular orientation (as in typical studies of perceptual learning; see DISCUSSION for more details).

variance is captured by the Fisher information (FI) metric, which in this case takes the form

$$FI_i = \frac{f'_i(\theta)^2}{n_i(\theta)} \quad (3)$$

where $f'_i(\theta)$ is the slope of the tuning function at orientation θ , and $n_i(\theta)$ is the variance of the tuning function at that point (defined here as the variance across all trials). To estimate FI at each point along each individual subject's VTF, we used the best fitting half-Gaussian (see previous section) to estimate the slope of the response function. We then squared the slope and weighted it by the variance at each point. To compute the variance at each point along this smooth curve, we first estimated the variance of each of the 10 measured data points in the VTF across all trials, and then used linear interpolation to generate an estimate of the variance at each point along the half-Gaussian.

RESULTS

Behavior

Subjects performed the DMTS task shown in Fig. 2. Task difficulty was maintained by adjusting the orientation offset between the sample and test stimuli so as to produce 65% correct responses during the fMRI scanning session (average offset across subjects: $\pm 2.2 \pm 1.3^\circ$, range: $1\text{--}4.75^\circ$). This relatively low level of accuracy was selected to ensure that there were enough data to later compare evoked BOLD responses on correct and incorrect trials.

Voxel-based estimates of population response profiles

To estimate the relative activation level in different populations of feature selective neurons while subjects performed the fine discrimination task, we used an analysis based on "voxel tuning functions" (VTFs) (Kay et al. 2008; Serences et al. 2009). This analysis rests on the assumption that the distribution of feature-selective neurons—in this case, the distribution of orientation-selective columns in V1—is not uniform across a given visual area (Boynton 2005; Kamitani and Tong 2005). In turn, single voxels should exhibit a differential response when the subject views stimuli rendered in different orienta-

tions due to a nonuniform sampling of orientation columns within a voxel, giving rise to small but detectable feature-selective biases in the BOLD response (Kamitani and Tong 2005; Kay et al. 2008; Miyawaki et al. 2008; Serences et al. 2009; Swisher et al. 2010). As long as the magnitude of the BOLD response is monotonically related to changes in the magnitude of neural activity (Logothetis 2003; Logothetis et al. 2001), then feature-selective modulations observed at the voxel level should indirectly index feature-selective changes within the sampled pool of neurons. Note that we used the VTF approach in the present study instead of the more common "multi-voxel pattern classifiers" because these pattern classifiers pool information from *all* voxels within a given cortical area and are therefore designed to establish that an area is encoding some form of information about stimulus features (e.g., Haxby et al. 2001; Haynes and Rees 2005; Kamitani and Tong 2005; Norman et al. 2006; Peelen and Downing 2007; Serences and Boynton 2007). In contrast, directly examining modulations within specific subsets of feature-selective voxels can in principle provide information about *how* a particular experimental manipulation selectively influences population response profiles in early visual cortex (Serences et al. 2009).

Evaluating off-channel gain in primary visual cortex (V1)

We first used independent functional localizer scans to select voxels in V1 that responded to the spatial position of the stimulus aperture used in the fine discrimination scans (see METHODS). Next, data from seven of the eight fine discrimination scans were used to sort each voxel into 1 of 10 bins based on the stimulus orientation that produced the largest response. Then using only data from the eighth scan, we computed the response of voxels in each orientation-selective bin as a function of their offset from the sample stimulus on each trial (where the orientation of the sample stimulus on each trial was always defined as 0° by convention, see METHODS and Fig. 3).

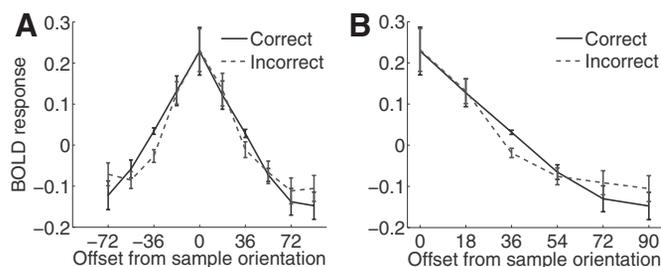


FIG. 3. Average blood-oxygen-level-dependent (BOLD) response for correct (solid line) and incorrect (dashed line) trials in V1. The x axis indicates the offset of a voxel's tuning preference from the sample orientation (which is always defined as 0° by convention). *A*: the BOLD response is displayed for clockwise (CW; negative values on the x axis) and counterclockwise (CCW; positive values on the x axis) offsets from sample separately. At offsets of $\pm 36^\circ$, the BOLD signal is significantly greater for correct trials. *B*: given that the BOLD response is essentially symmetrical $\sim 0^\circ$, we averaged across analogous CW and CCW rotations. Error bars represent ± 1 SE across subjects. Note that the between-subject error is smallest at the 36° offset. This compression happened because the BOLD data have a mean of 0 (as a consequence of z -transforming), so the BOLD response corresponding to the condition that is closest to the mean of the data set for a given subject will have an amplitude near 0. It turns out that the 36° point was close to the mean for each subject, thus giving rise to low between-subject variability. However, because all statistics are based on within-subjects variance (repeated-measures), this compression of the between-subjects error does not influence our conclusions.

This procedure was repeated across all permutations of holding-one-scan out, and data were then averaged across all permutations. The end result of this procedure was an estimate of the activation level in groups of voxels that were tuned to all orientations surrounding the sample stimulus, which provides an estimate of the population response profile. As expected, the BOLD response was highest in voxels that were tuned to the orientation of the sample stimulus (0° offset), as the stimulus evoked a large sensory response. Voxels tuned progressively farther away from the sample orientation were less active (see Fig. 3A), consistent with a Gaussian selectivity of the underlying feature-tuned response profile (see also Serences et al. 2009).

We next estimated the shape of the population response profile separately on correct and incorrect trials (Fig. 3A). There was no difference in the activation level of voxels tuned to the sample orientation (0° offset) on correct and incorrect trials in V1 [$t(8) = -0.09$, $P = 0.9$]. Because the shape of the population response profile was nearly symmetric about the 0° point, we collapsed across data points that were offset from 0° by an equivalent amount in either direction (i.e., data were averaged across voxels tuned to ± 18 , ± 36 etc. away from the sample, Fig. 3B). A two-way repeated measures ANOVA with orientation offset from sample (6 levels) and response correct/incorrect (2 levels) was run to evaluate the influence of discrimination accuracy on the shape of the VTFs. As noted above, the magnitude of the BOLD response monotonically decreased in voxels tuned progressively further from 0° on both correct and incorrect trials [$F(5,40) = 17.53$, $P < 0.001$; this was also true for V2v, V3v, and hV4, all $P_s < 0.001$, see Fig. 4]. In addition, there was a significant interaction between accuracy and orientation offset in V1 [$F(5,40) = 2.57$, $P = 0.042$] such that responses were higher on correct trials in voxels tuned $\pm 36^\circ$ from the sample (Wilcoxon signed-rank test on data from only the 36° offset: $P < 0.001$, all 9 subjects

showed larger responses on correct compared with incorrect trials). The difference between activation levels on correct and incorrect trials was not significant at any other offset (all $P_s > 0.43$). Finally, logistic regression revealed that larger responses in voxels tuned 36° from the target predicted correct discriminations on a trial-by-trial basis [mean fit coefficient \pm SE: 0.1 ± 0.04 ; $t(8) = 2.75$, $P = 0.025$; no significant effects were observed for any other orientation offset]. No differences in off-channel gain were observed in areas V2v, V3v, and hV4 (Fig. 4).

Because the stimuli used on correct and incorrect trials were identical (i.e., same rotational offset between sample and test), this off-channel modulation of the BOLD response in V1 was not driven by differences in the sensory input. We therefore conclude that the relative magnitude of activity in neural populations tuned away from the target—as indirectly indexed using feature-selective BOLD responses—predicts success on a difficult visual discrimination task.

Evaluating change in slope of voxels tuned 36° from sample

The utility of a voxel (and by inference a large population of neurons) should depend on the slope of the VTF at the sample orientation (Fig. 1B). Thus it is important to establish that increasing the gain of voxels tuned 36° from the sample significantly increases the slope of their respective VTFs. To evaluate this possibility, we first computed the FWHM of the response profiles: the mean FWHM (\pm SE) across observers was $59 \pm 6.2^\circ$, which translates to an offset of $\sim 30^\circ$ from the peak on either side and is consistent with a high VTF slope at the sample orientation (see also: Gur et al. 2005; Ringach et al. 2002; Snowden 1992). However, the slope of the observed VTFs is maximal just inside the FWHM point along the curve. Therefore we computed a more formal metric of informativeness based on a ratio of VTF slope and the variance of the VTF at each point (see METHODS). We predicted that increasing the gain of off-channel VTFs by the amount that we observed at the 36° offset (Fig. 3B) should lead to an increase in the slope of these VTFs around the sample orientation (0°) and hence an increase in information to support the fine-discrimination (Fig. 5). More specifically, the slope of voxels tuned 36° from the sample should peak at 0° if off-channel gain was indeed optimally deployed to the most informative neural populations. The slope of the VTFs centered 36° from the sample actually peaked $11 \pm 1.2^\circ$ (mean ± 1 SE across subjects) away from the sample orientation (Fig. 5E). However, note that an 11° offset is quite close to 0° given that we sampled the VTFs in step sizes of 18° . Moreover, the magnitude of the gain increase on correct trials in voxels tuned 36° from the sample gives rise to a significant increase in the amount of information available to support the fine perceptual discrimination [i.e., a significant increase in slope at the sample orientation, $t(8) = 4.18$, $P = 0.003$, Fig. 5F, all 9 subjects showed the effect].

An alternative account holds that the increased activation in voxels tuned 36° from the sample was not due to an amplitude gain in the VTFs (as we assumed in Fig. 5) but instead due to a selective increase in the bandwidth of the VTFs. If this account is correct, then the response in voxels tuned 36° from the sample would indeed be higher, but this increase in bandwidth should *decrease* the slope of the tuning functions and in turn reduce the sensitivity of these neural populations to

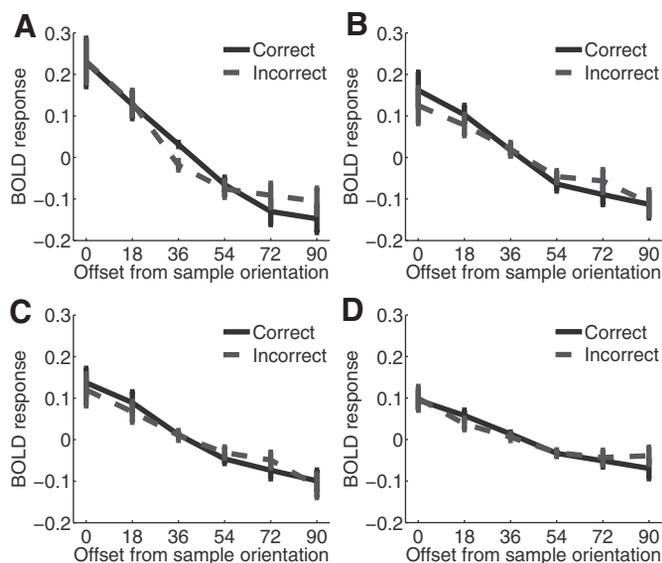


FIG. 4. Average BOLD response for correct (solid line) and incorrect (dashed line) trials in V1 (A), V2v (B), V3v (C), and hV4 (D). As in Fig. 3B, data were collapsed across analogous CW and CCW rotations in all panels. The significant difference at the 36° offset observed in V1 is not replicated across the remaining three visual areas (all $P_s > 0.43$). Error bars represent ± 1 SE across subjects (and see note in Fig. 2 legend about the compression of between-subject error at the 36° offset).

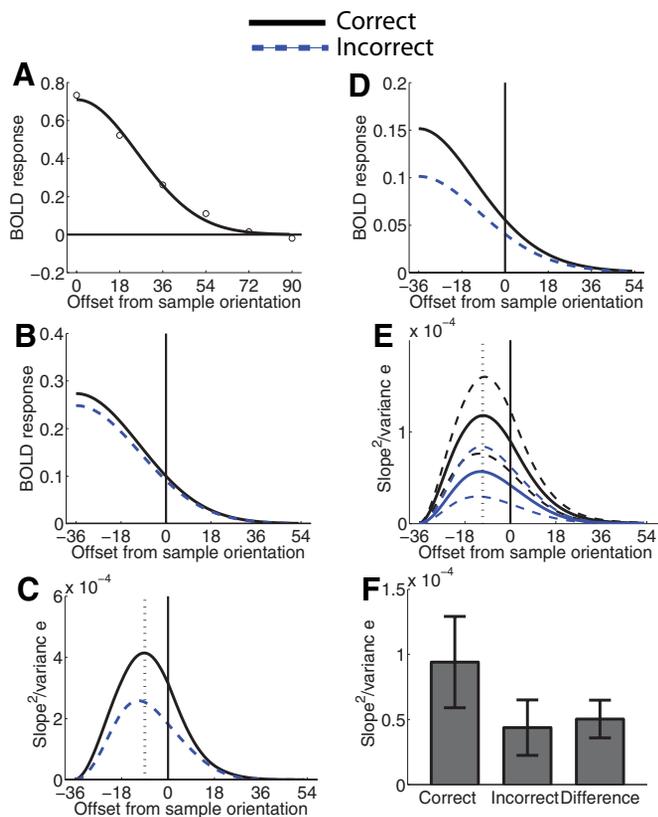


FIG. 5. *A*: best fitting half-Gaussian to the data from a single subject (*AA*), collapsed across correct and incorrect trials (see METHODS). *B*: estimated shape of the half-Gaussian in the same subject on correct (black) and incorrect (blue) trials in voxels tuned 36° away from the target. The additive baseline offset was removed from each curve, and the amplitudes of the black and blue curves were scaled by the observed difference in the response amplitude of voxels tuned 36° from the target on correct and incorrect trials, respectively. *C*: the squared slope divided by the within-subject trial-by-trial variance for each of the curves shown in *B*. Solid vertical lines indicate the sample orientation, and the vertical dashed lines indicate where the slope²/variance function peaks (8° from the sample in this subject). Note that because these values are based on voxels tuned to $\pm 36^\circ$ from sample, the predicted “optimal” off-channel gain would be at 0° . *D*: same conventions as in *B* and *C* but averaged across all subjects. Note that the difference in the peaks of these curves matches the difference between response amplitudes on correct and incorrect trials in voxels offset 36° from the sample, as shown in Fig. 3*B*. *E*: the color of the lines signifies correct (black) and incorrect (blue) trials; solid lines indicate the squared slope divided by within-subject trial-by-trial variance as in *B–D*, and dashed lines indicate the ± 1 SE across subjects. Here the curve peaks at an offset of 11° from the sample, where again the prediction of optimal off-channel gain would be at 0° . *F*: cross-subject mean slope²/variance metric on correct and incorrect trials at the sample orientation. Error bars represent ± 1 SE across subjects.

discriminate the sample and test stimuli (see Supplemental Fig. S1).¹ Thus we argue against this account because the increased activation in off-channel voxels was associated with improved perceptual performance, not a decrease in perceptual performance as would be predicted by a change in bandwidth.

DISCUSSION

Many previous studies suggest that simple perceptual decisions are based on the activity level of the sensory neurons that are tuned to the relevant stimulus feature (see Gold and Shadlen 2007 for a review). However, when faced with a

difficult decision between two highly similar stimuli, off-channel neurons most reliably signal differences between adjacent stimulus features and should therefore be relied on to provide the “evidence” that is used to inform a decision (Jazayeri and Movshon 2006; Law and Gold 2009). Consistent with this account, we found that the magnitude of off-channel gain in primary visual cortex predicts performance on a difficult discrimination task. This functionally significant off-channel gain in voxels tuned $\pm 36^\circ$ away from the sample is also consistent with the prediction that these neurons should indeed be highly informative given that V1 tuning functions generally have a bandwidth of $\sim 30\text{--}60^\circ$ (Gur et al. 2005; Ringach et al. 2002; Snowden 1992).

Several previous reports have shown a similar correlation between the activity of off-channel sensory neurons and performance on a fine-discrimination task (e.g., Purushothaman and Bradley 2005; Schoups et al. 2001; Yang and Maunsell 2004). However, these prior studies employed an experimental procedure that involved extensive training with a specific feature value [e.g., discriminate 90° motion from a 93° motion as in Purushothaman and Bradley (2005)]. In contrast, our subjects viewed a large set of possible stimulus features (10 in total) that were evenly distributed across the span of possible orientation values. Thus our subjects were not learning to rely on the gain of a *specific set* of sensory neurons (e.g., neurons tuned to an absolute orientation of $\sim 60^\circ$ when repeatedly discriminating a 90° from a 93° stimulus) but rather were dynamically evaluating the gain of off-channel neurons with respect to the sample orientation.

We predicted that off-channel gain during a difficult orientation discrimination task would occur in V1 given the high density of orientation-selective neurons compared with other visual areas (Hubel and Weisel 1962). Indeed the BOLD signal in off-channel V1 voxels predicted performance on the fine discrimination task, whereas no such relationship was observed in any other visual area (V2v–hV4; Fig. 4). The failure to observe off-channel gain related to task performance in other visual areas may also be due to a lack of sensitivity and signal; V1 is large and relatively flat compared with other visual areas, and thus may lend itself to sampling via voxel-based fMRI methods. Furthermore, vascular density is thought to be higher in both human and nonhuman primate striate cortex compared with extrastriate cortex (Duvernoy et al. 1981; Weber et al. 2008), which might lead to better overall BOLD SNR in V1 than in V2v, V3v, and hV4. Finally, it is possible that if we used another set of visual features—such as colored disks or translating motion patterns—then off-channel gain would be magnified in other visual areas that contain a higher proportion of neurons tuned to these other feature domains. Thus in more real-world situations in which one must discriminate between two similar but more complex stimuli (e.g., a radiologist discriminating between regions of a CT image), it is quite possible that the sensory evidence driving perceptual decisions would originate from the most informative neurons in whatever visual areas were most relevant for performing the required task.

One unresolved question is whether the observed gain changes in off-channel neural populations are the result of top-down attentional processes or whether they reflect trial-by-trial fluctuations in neural activity that incidentally increase the amount of information available to perform the discrimination.

¹ The online version of this article contains supplemental data.

Consistent with the top-down attentional account, Navalpakkam and Itti (2007) and Scolari and Serences (2009) provided behavioral evidence that attention modulates the influence of informative off-channel neurons during fine-discrimination tasks. However, neither study could determine if attention directly influenced the gain of off-channel neural populations in early visual cortex or if attention instead acted to optimize the read-out of information from early visual cortex during decision-making (Law and Gold 2008; Palmer et al. 2000). While the present observation of modulations in V1 is consistent with an attentional account, Britten et al. (1996) demonstrated that seemingly random fluctuations in the spiking activity of sensory neurons tuned to a particular direction of motion predicted choice behavior even when the physical stimulus was ambiguous and contained no diagnostic information. A similar phenomenon might be operating here as well: when the firing rate of informative off-channel neurons happens to be high on a given trial, then these neurons will be more sensitive to signal small orientation changes between the sample and the test stimuli. Unfortunately, because we post-sorted the trials based on behavioral performance, we cannot directly rule in favor of either a top-down attentional account or a “random fluctuation” account of our data. However, in future studies we will disambiguate these alternatives to determine if off-channel gain in low-level sensory neurons is under the direct control of voluntary attentional mechanisms.

GRANTS

This work was supported by National Mental Health Institute Grant R21-MH-083902 to J. T. Serences.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

REFERENCES

- Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A. Probabilistic population codes for Bayesian decision making. *Neuron* 60: 1142–1152, 2008.
- Boynton GM. Attention and visual perception. *Curr Opin Neurobiol* 15: 465–469, 2005.
- Brainard DH. The Psychophysics Toolbox. *Spat Vision* 10: 433–436, 1997.
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13: 87–100, 1996.
- Ditterich J, Mazurek ME, Shadlen MN. Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat Neurosci* 6: 891–898, 2003.
- Duvernoy HM, Delon S, Vannson JL. Cortical blood vessels of the human brain. *Brain Res Bull* 7: 519–579, 1981.
- Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN. fMRI of human visual cortex. *Nature* 369: 525, 1994.
- Gold JI, Shadlen MN. Representation of a perceptual decision in developing oculomotor commands. *Nature* 404: 390–394, 2000.
- Gold JI, Shadlen MN. Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5: 10–16, 2001.
- Gold JI, Shadlen MN. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36: 299–308, 2002.
- Gold JI, Shadlen MN. The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J Neurosci* 23: 632–651, 2003.
- Gold JI, Shadlen MN. The neural basis of decision making. *Annu Rev Neurosci* 30: 535–574, 2007.
- Gur M, Kagan I, Snodderly DM. Orientation and direction selectivity of neurons in V1 of alert monkeys: functional relationships and laminar distributions. *Cereb Cortex* 15: 1207–21, 2005.
- Hanes DP, Schall JD. Neural control of voluntary movement initiation. *Science* 274: 427–430, 1996.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293: 2425–2430, 2001.
- Haynes JD, Rees G. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8: 686–691, 2005.
- Heekeren HR, Marrett S, Bandettini PA, Ungerleider LG. A general mechanism for perceptual decision-making in the human brain. *Nature* 431: 859–862, 2004.
- Heekeren HR, Marrett S, Ruff DA, Bandettini PA, Ungerleider LG. Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. *Proc Natl Acad Sci USA* 103: 10023–10028, 2006.
- Ho TC, Brown S, Serences JT. Domain general mechanisms of perceptual decision making in human cortex. *J Neurosci* 29: 8675–8687, 2009.
- Hol K, Treue S. Different populations of neurons contribute to the detection and discrimination of visual motion. *Vision Res* 41: 685–689, 2001.
- Horowitz GD, Batista AP, Newsome WT. Representation of an abstract perceptual decision in macaque superior colliculus. *J Neurophysiol* 91: 2281–2296, 2004.
- Hubel DH, Wiesel TN. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J Physiol* 160: 106–154, 1962.
- Jazayeri M, Movshon JA. Optimal representation of sensory information by neural populations. *Nat Neurosci* 9: 690–696, 2006.
- Jazayeri M, Movshon JA. A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* 446: 912–915, 2007a.
- Jazayeri M, Movshon JA. Integration of sensory evidence in motion discrimination. *J Vision* 7: 1–7, 2007b.
- Kamitani Y, Tong F. Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8: 679–685, 2005.
- Kay NK, Naselaris T, Prenger RJ, Gallant JL. Identifying natural images from human brain activity. *Nature* 452: 352–355, 2008.
- Lagarias JC, Reeds JA, Wright MH, Wright PE. Convergence properties of the Nelder-Mead simplex method in low dimensions. *SIAM J Optimization* 9: 112–147, 1998.
- Law CT, Gold JI. Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat Neurosci* 11: 505–513, 2008.
- Law CT, Gold JI. Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nat Neurosci* 12: 655–663, 2009.
- Leon MI, Shadlen MN. Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron* 24: 415–425, 1999.
- Logothetis NK. The underpinnings of the BOLD functional magnetic resonance imaging signal. *J Neurosci* 23: 3963–3971, 2003.
- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412: 150–157, 2001.
- Martinez-Trujillo JC, Treue S. Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14: 744–751, 2004.
- Maunsell JH, Treue S. Feature-based attention in visual cortex. *Trends Neurosci* 29: 317–322, 2006.
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN. A role for neural integrators in perceptual decision making. *Cereb Cortex* 13: 1257–1269, 2003.
- Mitchell JF, Sundberg KA, Reynolds JH. Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* 55: 131–141, 2007.
- Miyawaki Y, Uchida H, Yamashita O, Sato M, Morito Y, Tanabe H, Sadato N, Kamitani Y. Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron* 60: 915–929, 2008.
- Navalpakkam V, Itti L. Search goal tunes visual features optimally. *Neuron* 53: 605–617, 2007.
- Newsome WT, Britten KH, Movshon JA. Neuronal correlates of a perceptual decision. *Nature* 341: 52–54, 1989.
- Newsome WT, Pare EB. A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J Neurosci* 8: 2201–2211, 1988.
- Norman KA, Polyn SM, Detre GJ, Haxby JV. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10: 424–430, 2006.

- Palmer J, Verghese P, Pavel M.** The psychophysics of visual search. *Vision Res* 40: 1227–1268, 2000.
- Peelen MV, Downing PE.** Using multi-voxel pattern analysis of fMRI data to interpret overlapping functional activations. *Trends Cogn Sci* 11: 4–5, 2007.
- Pelli DG.** The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision* 10: 437–442, 1997.
- Pouget A, Deneve S, Latham PE.** The relevance of Fisher Information for theories of cortical computation and attention. In: *Visual Attention and Cortical Circuits*, edited by Braun J, Koch C, Davis JL. Cambridge, MA: MIT Press, 2001, p. 265–284.
- Purushothaman G, Bradley DC.** Neural population code for fine perceptual decisions in area MT. *Nat Neurosci* 8: 99–106, 2005.
- Regan D, Beverley KI.** Postadaptation orientation discrimination. *J Opt Soc Am* 2: 147–155, 1985.
- Ringach DL, Shapley RM, Hawken MJ.** Orientation selectivity in macaque V1: diversity and laminar dependence. *J Neurosci* 22: 5639–5651, 2002.
- Roitman JD, Shadlen MN.** Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22: 9475–9489, 2002.
- Romo R, Salinas E.** Flutter discrimination: neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4: 203–218, 2003.
- Salzman CD, Murasugi CM, Britten KH, Newsome WT.** Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* 12: 2331–2355, 1992.
- Schall JD.** Neural basis of deciding, choosing and acting. *Nat Rev Neurosci* 2: 33–42, 2001.
- Schoups A, Vogels R, Qian N, Orban G.** Practising orientation identification improves orientation coding in V1 neurons. *Nature* 412: 549–553, 2001.
- Scolari M, Serences JT.** Adaptive allocation of attentional gain. *J Neurosci* 29: 11933–11942, 2009.
- Serences JT, Boynton GM.** Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55: 301–312, 2007.
- Serences JT, Saproo S, Scolari M, Ho T, Muftuler LT.** Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *NeuroImage* 44: 223–231, 2009.
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB.** Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268: 889–893, 1995.
- Shadlen MN, Britten KH, Newsome WT, Movshon JA.** A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16: 1486–1510, 1996.
- Shadlen MN, Newsome WT.** Noise, neural codes, and cortical organization. *Curr Opin Neurobiol* 4: 569–579, 1994.
- Shadlen MN, Newsome WT.** Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86: 1916–1936, 2001.
- Snowden RJ.** Orientation bandwidth: the effects of spatial and temporal frequency. *Vision Res* 32: 1965–1974, 1992.
- Swisher JD, Gatensby JC, Gore JC, Wolfe BA, Moon C-H, Kim S-G, Tong F.** Multiscale pattern analysis of orientation-selective activity in primary visual cortex. *J Neurosci* 30: 325–330, 2010.
- Treue S, Martinez-Trujillo JC.** Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399: 575–579, 1999.
- Weber B, Keller AL, Reichold J, Logothetis NK.** The microvascular system of the striate and extrastriate visual cortex of the macaque. *Cereb Cortex* 18: 2318–2330, 2008.
- Yang T, Maunsell JHR.** The effect of perceptual learning on neuronal responses in monkey visual area V4. *J Neurosci* 24: 1617–1626, 2004.