

Relationship between BOLD amplitude and pattern classification of orientation-selective activity in the human visual cortex

Frank Tong^{a,*}, Stephenie A. Harrison^a, John A. Dewey^a, Yukiyasu Kamitani^b

^a Psychology Department and Vanderbilt Vision Research Center, Vanderbilt University, Nashville, TN, USA

^b ATR Computational Neuroscience Laboratories, Keihanna Science City, Kyoto, Japan

ARTICLE INFO

Article history:

Accepted 5 August 2012

Available online 10 August 2012

Keywords:

fMRI

Decoding

Multivoxel pattern analysis

V1

ABSTRACT

Orientation-selective responses can be decoded from fMRI activity patterns in the human visual cortex, using multivariate pattern analysis (MVPA). To what extent do these feature-selective activity patterns depend on the strength and quality of the sensory input, and might the reliability of these activity patterns be predicted by the gross amplitude of the stimulus-driven BOLD response? Observers viewed oriented gratings that varied in luminance contrast (4, 20 or 100%) or spatial frequency (0.25, 1.0 or 4.0 cpd). As predicted, activity patterns in early visual areas led to better discrimination of orientations presented at high than low contrast, with greater effects of contrast found in area V1 than in V3. A second experiment revealed generally better decoding of orientations at low or moderate as compared to high spatial frequencies. Interestingly however, V1 exhibited a relative advantage at discriminating high spatial frequency orientations, consistent with the finer scale of representation in the primary visual cortex. In both experiments, the reliability of these orientation-selective activity patterns was well predicted by the average BOLD amplitude in each region of interest, as indicated by correlation analyses, as well as decoding applied to a simple model of voxel responses to simulated orientation columns. Moreover, individual differences in decoding accuracy could be predicted by the signal-to-noise ratio of an individual's BOLD response. Our results indicate that decoding accuracy can be well predicted by incorporating the amplitude of the BOLD response into simple simulation models of cortical selectivity; such models could prove useful in future applications of fMRI pattern classification.

© 2012 Elsevier Inc. All rights reserved.

Introduction

Orientation is a fundamental feature that provides the basis by which the visual system encodes the local contours of visual objects. Psychophysical studies indicate that observers are exquisitely sensitive to visual orientation. For example, humans can discriminate between gratings that differ by less than one degree (Regan and Beverley, 1985; Skottun et al., 1987) or detect a colinear arrangement of oriented patches embedded in a field of random orientations (Field et al., 1993; Kovacs and Julesz, 1993). Neurophysiological studies in animals have revealed the prevalence of orientation-selective neurons and cortical columns throughout the visual cortex, with stronger selectivity and columnar organization found in lower tier visual areas and broader tuning found in higher visual areas (Blasdel and Salama, 1986; Bonhoeffer and Grinvald, 1991; Desimone and Schein, 1987; Felleman and Van Essen, 1987; Gegenfurtner et al., 1997; Hubel and Wiesel, 1968; Levitt et al., 1994; Ohki et al., 2006; Vanduffel et al., 2002).

Advances in functional magnetic resonance imaging (fMRI) have also led to promising approaches for investigating orientation selectivity in the human visual system. These include amplitude measures of orientation-selective adaptation (Boynton and Finney, 2003; Fang et al., 2005; Murray et al., 2006; Tootell et al., 1998b), recent efforts to isolate orientation columns using high-resolution functional imaging (Yacoub et al., 2008), and the decoding of orientation-selective responses from cortical activity patterns (Kamitani and Tong, 2005). Here, we focus on the orientation decoding approach, as our previous work has revealed that this method is both robust and flexible, allowing for the detection of orientation signals from multiple visual areas when scanning at standard fMRI resolutions.

In the study by Kamitani and Tong (2005), the authors applied pattern classification algorithms to determine whether reliable orientation information might be present in the ensemble patterns of fMRI activity found in the human visual cortex. To their initial surprise, they found that pattern classifiers could predict which of eight possible orientations a subject was viewing on individual stimulus blocks with remarkable accuracy. Given that the fMRI signals were sampled at a standard spatial resolution ($3 \times 3 \times 3$ mm) that greatly exceeded the presumed size of human orientation columns, how was this possible? Analyses and simulations

* Corresponding author at: Psychology Department, Vanderbilt University, 301 Wilson Hall, 111 21st Avenue South, Nashville, TN 37240, USA.

E-mail address: frank.tong@vanderbilt.edu (F. Tong).

indicated that random variations in the local distribution of orientation columns could lead to weak but reliable orientation biases in individual voxels, even across small shifts in head position. By pooling the orientation information available from many voxels, it was possible to extract considerable orientation information from these patterns of cortical activity. High-resolution functional imaging of the visual cortex in cats and humans has shown that orientation-selective information can be found at multiple spatial scales, ranging from the scale of columns up to several millimeters, thereby allowing for reliable decoding at standard fMRI resolutions (Swisher et al., 2010).

Since its advent, this approach for measuring feature-selective responses has attracted considerable interest (Haynes and Rees, 2006; Kriegeskorte and Bandettini, 2007; Norman et al., 2006; Tong and Pratte, 2012). Applications of this approach have led to successful decoding of the contents of feature-based attention (Jehee et al., 2011; Kamitani and Tong, 2005, 2006; Serences and Boynton, 2007a), conscious perception (Brouwer and van Ee, 2007; Haynes and Rees, 2005b; Serences and Boynton, 2007b), subliminal processing (Haynes and Rees, 2005a), and even the contents of visual working memory (Harrison and Tong, 2009; Serences et al., 2009). Given that top-down processes such as voluntary attention or working memory can have such a powerful influence on these measures of feature-selective activity, it is important to determine whether the reliability of these feature-selective responses would also scale with strength of the stimulus-driven response. In recent studies, we have found that top-down visual processes of selective attention or working memory can lead to dissociable effects on these orientation-selective activity patterns and the overall BOLD response (Harrison and Tong, 2009; Jehee et al., 2011). For example, one can accurately predict the orientation being retained in working memory even when the overall activity in the visual cortex falls to near-baseline levels (Harrison and Tong, 2009). Thus, the relationship between response amplitude and reliability of the orientation-selective activity pattern is not entirely straightforward.

The goal of the present study was to characterize how these orientation-selective activity patterns vary as a function of the strength and quality of the visual input. Specifically, we investigated how systematic manipulations of stimulus contrast (Experiment 1) and spatial frequency (Experiment 2) affected these fMRI responses, using the accuracy of orientation decoding as an index of the reliability of the orientation-selective activity patterns in the visual cortex. We also conducted more in-depth analyses to gain insight into the relationship between the overall amplitude of stimulus-driven activity and the amount of orientation information contained in the cortical activity patterns. Specifically, we determined whether fMRI amplitudes could effectively predict the accuracy of orientation decoding using a simple simulation model, by incorporating these amplitude values to adjust the response strength of simulated orientation columns. This model was then used to predict the overall accuracy of orientation decoding performance across changes in stimulus contrast and spatial frequency, based on the fMRI response amplitudes observed in those conditions. We also performed a variety of correlational analyses to investigate the relationship between fMRI response amplitudes and orientation decoding performance, both at a group level and when inspecting individual data.

Based on previous studies of the visual system, we predicted that the reliability of orientation-selective activity patterns should increase as a function of stimulus contrast (Lu and Roe, 2007; Smith et al., 2011). Moreover, we predicted that the effects of contrast should be more pronounced in V1 than in higher visual areas, due to the gradual increase in contrast invariance that occurs at progressively higher levels of the visual hierarchy (Avidan et al., 2002; Boynton et al., 1999; Kastner et al., 2004; Tootell et al., 1995). We further predicted that high spatial frequency gratings would be better discriminated with central than peripheral presentation, and that V1 might show an advantage over extrastriate areas in discriminating fine orientations.

Experiment 1

In Experiment 1, we determined how well the activity patterns in individual visual areas (V1–V4) could discriminate between orthogonal grating patterns presented at varying levels of stimulus contrast (4, 20 or 100%). Sine-wave gratings of independent orientation (45° or 135°) were presented in the left and right visual fields while subjects performed a simple visual detection task (Fig. 1). Activity patterns obtained from the contralateral visual cortex were used to classify the orientation seen in each hemifield; control analyses were performed using voxels sampled from the ipsilateral visual cortex. We expected that contralateral visual areas should show better orientation discrimination at higher stimulus contrasts, reflecting the strength of incoming visual signals. Moreover, we predicted that manipulations of stimulus contrast should have a greater impact on orientation discrimination performance in lower-tier visual areas, especially V1, as these areas are thought to be more contrast-dependent (Avidan et al., 2002; Boynton et al., 1999; Gegenfurtner et al., 1997; Kastner et al., 2004; Levitt et al., 1994; Sclar et al., 1990; Tootell et al., 1995).

To investigate the relationship between orientation-selective responses and the overall BOLD amplitude found in each region of interest, we performed both correlational analyses and simulation analyses. For our simulation analysis, we generated a 1-dimensional array of simulated orientation columns, following the work of Kamitani and Tong (2005). A modest degree of random spatial jitter was used to generate the preferred orientations for the repeating cycle of columns, which led to small local anisotropies in orientation preference (see Materials and methods for details). The response of each column was specified by a Gaussian-shaped orientation-tuning curve centered around its preferred orientation. From this fine-scale columnar array, we sampled large voxel-scale activity patterns (with random independent noise added to each voxel's response), and submitted these to our linear classifier. To determine whether decoding of these simulated activity patterns might account for our fMRI decoding results in a given visual area, we adjusted the sharpness of the columnar orientation tuning curve to match the mean level of classification accuracy across the 3 contrast levels; this was the only free parameter in our model. The amplitude (height) of the Gaussian-tuned response curve was scaled according to the mean




Fig. 1. Examples of the visual orientation displays from Experiment 1. On each 16-s stimulus block, sine-wave gratings of independent orientation (45° or 135°) were presented in the two hemifields at 1–9° eccentricity. Gratings flashed on and off at a rate of 2 Hz.

fMRI response amplitude observed in the corresponding stimulus condition. For this simulation, the decoding accuracy across changes in stimulus contrast provided a measure of how well the mean stimulus-driven BOLD response could account for the reliability of the fMRI orientation-selective activity patterns.

Materials and methods

Participants

Ten healthy adults (eight male, two female), ages 23–35, with normal or corrected-to-normal visual acuity participated in this experiment. All subjects gave written informed consent. The study was approved by the Institutional Review Board at Vanderbilt University.

Experimental design and stimuli

Visual displays were rear-projected onto a screen in the scanner using a luminance-calibrated MR-compatible LED projector (Avotec, Inc.). All stimuli were generated by a Macintosh G4 computer running Matlab and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997).

The stimulus display consisted of oriented sine-wave gratings presented in the left and right visual fields (Fig. 1). The gratings were displayed at an intermediate spatial frequency of 1 cycle per degree (cpd) within an annular region extending from 1 to 9° eccentricity. Along the vertical meridian, a 30°-sector region was removed to avoid potential spillover of visual input to ipsilateral regions of visual cortex. In addition, a Gaussian mask was applied to each grating to reduce visual responses to the edges of the stimulus.

The lateralized gratings flashed on and off every 250 ms (2 Hz), appearing at a randomized spatial phase with each presentation. The extended blank interval between presentations of the gratings minimized the appearance of apparent motion (Kahneman and Wolman, 1970). The participant's task was to maintain fixation on a central bull's-eye throughout each fMRI run and to report rare occasions in which the gratings failed to appear. This simple monitoring task was chosen because it could be performed well at all contrast levels and ensured that subjects had to attend to the display. The orientation of each of the two gratings was varied independently on each stimulus block, by presenting an equal number of blocks with left/right gratings of 45°/45°, 45°/135°, 135°/45°, and 135°/135°. These four possible display conditions occurred in a randomized order on each fMRI run. Each run consisted of a series of 16-s blocks: an initial fixation-rest block, 12 consecutive stimulus blocks with no intervening rest periods, and a final fixation block.

The contrast of the gratings was randomly varied across runs (4%, 20%, or 100%) while the spatial frequency was held constant at 1 cpd. Participants performed a total of 12–15 runs, consisting of 4–5 runs at each contrast level. Thus, within a given hemifield, participants viewed a total of 24–30 stimulus blocks of each orientation by contrast condition. The fMRI activity collected during each stimulus block served as an independent sample for classification analysis.

On separate runs within each experimental session, subjects viewed a reference stimulus to localize retinotopic regions corresponding to the stimulus locations in the main experiment. This “visual field localizer” consisted of dynamic random dots (dot size 0.2°, 10 random-dot images per second) presented within a contrast-modulated Gaussian envelope that matched that of the orientation display, with the exception that only a single hemifield was stimulated at a time. Participants performed 2 localizer runs in a session, which consisted of 10 cycles of stimulus presentation alternating between the left and right visual fields.

Retinotopic mapping of visual areas

All subjects participated in a separate retinotopic mapping session. Subjects viewed rotating and expanding checkerboard stimuli, which were used to generate polar angle and eccentricity maps, respectively. This allowed us to identify the boundaries between visual

areas on flattened cortical representations, using previously described methods (Engel et al., 1997; Sereno et al., 1995).

MRI acquisition

Scanning was performed on a 3.0-Tesla Philips Intera Achieva scanner using a standard 8-channel head coil at the Vanderbilt University Institute for Imaging Sciences. A high-resolution anatomical T1-weighted scan was acquired from each participant (FOV 256 × 256, 1 × 1 × 1 mm resolution). To measure BOLD contrast, standard gradient-echo echoplanar T2*-weighted imaging was used to collect 28 slices perpendicular to the calcarine sulcus, which covered the entire occipital lobe as well as the posterior parietal and temporal cortices (TR, 2000 ms; TE = 35 ms; flip angle, 80°; FOV 192 × 192; slice thickness, 3 mm, no gap; in-plane resolution, 3 × 3 mm). Participants used a custom-made bite bar system to stabilize head position and minimize motion.

Functional MRI data preprocessing

All fMRI data underwent three-dimensional (3D) motion correction using automated image registration software, followed by linear trend removal to eliminate slow drifts in signal intensity. No spatial or temporal smoothing was applied. The fMRI data were aligned to the retinotopic mapping data collected from the separate session, using Brain Voyager software (Brain Innovation). All automated alignment was subjected to careful visual inspection and manual fine-tuning to correct for potential residual misalignment. Rigid-body transformations were performed to align fMRI data to the within-session 3D anatomical scan, and then to the retinotopy data. After across-session alignment, fMRI data underwent Talairach transformation and re-interpolation using 3-mm isotropic voxels. This procedure allowed us to delineate individual visual areas on flattened cortical representations and to restrict the selection of voxels to the gray matter regions of the cortex.

Voxels used for orientation decoding analysis were selected from the cortical surface of areas V1 through V4. First, voxels near the gray–white matter boundary were identified within each visual area using retinotopic maps delineated on a flattened cortical surface representation. Then, the voxels were sorted according to the strength of their responses to the visual field localizer. We used 60 voxels from each hemisphere for each of areas V1, V2, V3, and V3A/V4 combined, selecting the most activated voxels as assessed by their t-value. The t-values of all selected voxels in areas V1–V3 were highly significant, typically greater than 10 and ranging upward to 30 or more in many subjects. The lowest t-values observed in any subject or voxel were 5.20, 7.32, and 4.22 ($p < 0.00005$) for areas V1, V2, and V3, respectively. Data from V3A and V4 were combined due to the smaller number of visually active voxels found in these regions. Note that the classification analyses performed separately on V3A and V4 did not reveal any reliable statistical differences in orientation decoding, so pooling of the data from these regions would not have affected the overall pattern of results.

The data samples used for orientation classification analysis were created by shifting the fMRI time series by 4 s to account for the hemodynamic delay of the BOLD response, and then averaging the MRI signal intensity of each voxel for each 16-s stimulus block. Response amplitudes of individual voxels were normalized relative to the average of all stimulus blocks within the run to minimize baseline differences across runs. The resulting activity patterns were labeled according to their corresponding stimulus orientation to serve as input to the orientation classifier.

Classification analysis

fMRI activity patterns from individual visual areas were analyzed using a linear classifier to predict the orientation shown in each hemifield. Each fMRI data sample could be described as a single point in a multidimensional space, where each dimension served to represent the response amplitude of a specific voxel in the activity pattern. We

used linear support vector machines (SVM) to obtain a linear discriminant function that could separate the fMRI data samples according to their orientation category (Vapnik, 1998). SVM is a powerful classification algorithm that aims to minimize generalization error by finding the hyperplane that maximizes the distance (or margin) between the most potentially confusable samples from each of the two categories.

Mathematically, deviations from this hyperplane can be described by a linear discriminant function: $g(\mathbf{x}) = \mathbf{w} \cdot \mathbf{x} + w_0$, where \mathbf{x} is a vector specifying the fMRI amplitude of each voxel in the activity pattern, \mathbf{w} is a vector specifying the corresponding weight of each voxel that determines the orientation of the hyperplane, and w_0 is an overall bias term to allow for shifts from the origin. For a given training data set, linear SVM finds optimal weights and bias for the discriminant function. If the training data samples are linearly separable, then the output of the discriminant function will be positive for activity patterns induced by one stimulus orientation and negative for those induced by the other orientation. This discriminant function can then be applied to classify the orientation of independent test samples. We have previously described methods for extending this approach to multi-class decoding (Kamitani and Tong, 2005).

To evaluate orientation classification performance, we performed an N -fold cross-validation procedure using independent samples for training and testing. This involved dividing the data set into N pairs of 45° and 135° blocks, training the classifier using data from $N - 1$ pairs, and then testing the decoder on the remaining pair of blocks. We performed this validation procedure repeatedly until all pairs were tested, to obtain a measure of classification accuracy for each visual area, visual hemifield, stimulus condition, and subject.

Statistical analyses

Orientation classification performance was calculated and plotted for the following visual areas: V1, V2, V3, and V3A/V4 combined. Next, we tested for differences in classification accuracy across stimulus contrasts and visual areas, focusing on areas V1–V3 because of their consistently strong performance. (Combined area V3A/V4 was excluded from these statistical analyses, because the low overall performance found in this region would otherwise lead to floor effects that could distort the analytic results.) Our measures of orientation classification performance for each contrast level (4, 20, 100%), hemifield (left or right), visual area (V1, V2, V3), and subject, served as data for within-subjects analysis of variance and planned contrasts. Since both hemispheres were stimulated independently, data from the two hemispheres were treated as independent observations. Interaction effects were further investigated using paired contrasts and general trend analysis. In addition, classification performance of each visual area was compared to chance-level performance of 50% using one-sample t -tests.

To ensure the statistical validity of our experimental design and analyses, we performed a permutation test by randomizing the labels for the activity patterns of participants in Experiment 1. The result of 60,000 iterations of this procedure revealed a mean classification accuracy of 50.3% (chance level 50%). We calculated the false positive rate for analyzing group data ($N = 10$), by applying a t -test (with alpha of 0.05) to each of the 6000 sets of data. We observed a false positive rate of 5.10%, indicating that the application of t -tests and other parametric statistical tests would lead to minimal inflation in the likelihood of committing a Type I statistical error.

Simulation analyses

To determine whether fMRI response amplitudes could account for the accuracy of orientation decoding, we applied the same classification analysis to voxel-scale activity patterns that were sampled from a 1-dimensional array of simulated orientation columns. Previously, we have shown that when random spatial jitter is applied to a regular array of orientation columns, individual voxels develop a weak orientation bias such that the resulting pattern of activity across

many voxels can allow for reliable orientation decoding (Kamitani and Tong, 2005). It should be noted that the effects of spatial irregularity are similar for 1D and 2D simulations, with the degree of irregularity determining the amount of bias evident at more coarse spatial scales. (For simplicity, in the present analysis we excluded the simulation of head motion jitter.)

Parameters for the simulation were as follows. We assumed a voxel width of 3 mm, and determined the average spacing between neighboring iso-orientation columns based on previously reported values in the macaque monkey (Vanduffel et al., 2002). Average iso-orientation spacings were 760, 950, and 970 μm for V1, V2, and V3 respectively, with 20 individual columns used to span this cortical distance. A modest degree of random spatial jitter was used to generate the preferred orientations for the repeating cycle of columns, allowing each column to deviate by 0.3 SD units relative to the average rate of change in orientation preference across the cortex. The orientation-tuned response function of each column was specified by a Gaussian function of fixed width, centered at the preferred orientation of the individual column. The height of the Gaussian was scaled according to the mean fMRI response amplitude observed in the corresponding stimulus condition and visual area. The width of the Gaussian tuning function was the only free parameter in our model, which was determined separately for each visual area to fit the mean level of classification accuracy found across the 3 contrast levels in Experiment 1. The orientation tuning widths obtained for each visual area in Experiment 1 were then used to fit the data in Experiment 2.

For classification of the simulated data, we generated the same number of voxels (i.e., features) and data samples as was used in the fMRI study. An array of 60 voxels was used to coarsely sample activity from the orientation columns. The response of each voxel was calculated using a boxcar function to obtain the mean level of activity from the relevant portion of the columnar array. Independent Gaussian-distributed noise was added to each voxel's response, to simulate the physiological noise sources inherent to fMRI measures of brain activity. The same level of random noise was used for all stimulus conditions and visual areas. After randomly generating an array of orientation columns for a particular visual area, we generated 30 samples of voxel activity patterns (with random independent noise) for each orientation (45° or 135°) and each of the 3 stimulus conditions, using the fMRI response amplitude in the corresponding condition to scale the amplitude of the columnar orientation responses. The classifier was trained to discriminate stimulus orientation using $N - 1$ pairs of data samples, and tested on the remaining pair of orientations using a leave-one-pair-out cross-validation procedure. We performed 1000 simulations for each visual area, and calculated the average level of classification accuracy.

Results

fMRI classification analyses

We found that the activity patterns in all retinotopic visual areas could discriminate between orientations presented in the contralateral hemifield at above-chance levels, with higher stimulus contrasts leading to better performance (Fig. 2A). Overall, orientation classification performance was strongest in V1 and V2, moderate in V3, and much worse in V3A/V4. The decline in orientation decoding performance when ascending the visual hierarchy replicates our previous findings with high-contrast square-wave gratings (Kamitani and Tong, 2005). The results are generally consistent with reports of weaker columnar organization and broader orientation selectivity in higher extrastriate areas (Desimone and Schein, 1987; Vanduffel et al., 2002). Because orientation classification in V3A/V4 was less reliable and could potentially lead to floor effects, we focused on V1–V3 in our analytic comparisons of visual areas.

Statistical analyses confirmed that orientation classification performance was highly dependent on stimulus contrast ($F = 28.36$,

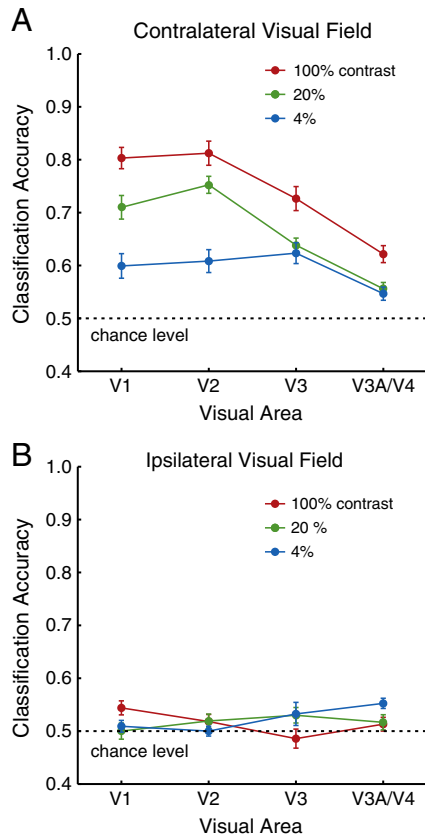


Fig. 2. Effects of stimulus contrast on orientation classification performance, plotted by visual area for Experiment 1. Areas V1–V3 performed significantly better than chance level of 50% at all contrast levels for gratings presented to the contralateral visual field (A), but could not reliably discriminate orientations presented in the ipsilateral visual field (B). Higher stimulus contrasts led to better discrimination of contralateral orientations, especially in early areas such as V1. Error bars indicate ± 1 S.E.M.

$p < 0.00001$). Orientation classification improved as the contrast of the grating increased from 4% to 20% ($t = 4.0$, $p < 0.001$), and from 20% to 100% ($t = 4.47$, $p < 0.001$). Thus, the reliability of these ensemble orientation-selective responses depended on the strength of the incoming visual signals. Visual areas also differed in classification performance ($F = 6.6$, $p < 0.005$), with areas V1 and V2 showing stronger orientation selectivity than V3 ($t = 3.35$, $p < 0.005$).




Fig. 3. Orientation classification performance plotted as a function of number of voxels used for pattern analysis for areas V1, V2 and V3. The most visually active voxels were first selected for analysis, based on t -values obtained from independent visual localizer runs. The pattern of results was stable over a wide range of voxel numbers, and classification performance reached near-asymptotic levels with the selection criterion of 60 voxels used for all main analyses.

Of particular interest was whether visual areas differed in their ability to discriminate between orientations as a function of contrast. We observed a significant interaction effect between visual area and contrast level ($F = 4.66$, $p < 0.005$), indicating a reliable difference in contrast sensitivity between visual areas. V1 and V2 showed larger differences in classification performance between the high- and low-contrast conditions than did area V3 ($t = 2.95$, $p < 0.005$), indicating that lower-tier areas are more contrast dependent and V3 is relatively more “contrast invariant”. Admittedly, when stimuli were presented at full strength (i.e., 100% contrast), our decoding method was less sensitive at detecting orientation-selective responses from V3 than from V1 or V2. Nonetheless, V3 performed as well as V1/V2 in the low-contrast condition, indicating a relative boost in sensitivity, specifically in the low-contrast range.

Control analyses confirmed that visual areas in each hemisphere could only predict the orientation of stimuli in the corresponding visual field. All visual areas could discriminate between orientations shown in the contralateral visual field (Fig. 2A) but not those shown in the ipsilateral visual field (Fig. 2B), indicating the location specificity of this orientation information. Also, the pattern of results remained consistent and stable when a sufficient numbers of voxels (~ 40 or more) was used for classification analysis (Fig. 3). Orientation classification accuracy was near asymptote at the level of 60 voxels used in the main analysis, suggesting that most of the available information had been extracted from the activated region.

To what extent can these orientation decoding results be understood in terms of the overall amplitude of activity in the visual cortex, which is known to increase monotonically with stimulus contrast and to approach asymptotic levels more quickly in higher visual areas? Previous fMRI studies have found that response amplitudes gradually become more invariant to contrast at progressively higher levels of the visual pathway, although reports have been mixed regarding whether early areas V1–V3 might differ in contrast sensitivity (Avidan et al., 2002; Boynton et al., 1999; Kastner et al., 2004; Tootell et al., 1995). In Fig. 4A, we show the overall amplitude of fMRI activity across the visual hierarchy for each level of stimulus contrast. Here, one can observe much stronger responses at higher contrasts and also a gradual increase in contrast invariance when ascending the visual hierarchy; area V1 showed a significantly greater difference in amplitude between high and low contrast conditions than did area V3 ($F = 11.98$, $p < 0.00001$). This pattern of results corresponds well with the orientation decoding results. Mean decoding accuracy was highly correlated with the mean amplitude of activity found in areas V1 through V3 across all contrast levels ($R = 0.71$, $p < 0.01$), indicating a strong positive relationship.

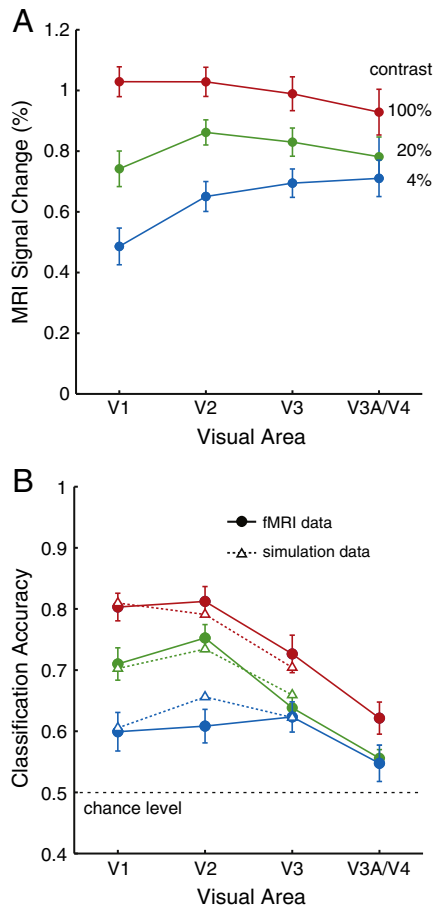


Fig. 4. Analysis of fMRI response amplitudes and their relationship to classification accuracy. A. Mean response amplitudes for each stimulus contrast and visual area. B. Model fits of orientation classification accuracy based on the mean fMRI response amplitudes observed in Experiment 1. Solid lines, original fMRI data; dashed lines, model fitted data. Mean fMRI responses were used to scale the response amplitude of simulated orientation columns, from which voxel-scale activity patterns were sampled and analyzed by a linear classifier. Simulation results provided a good fit of fMRI classification performance for areas V1–V3 ($R^2 = 0.938$, $p < 0.0001$). Error bars indicate ± 1 S.E.M.

Why should orientation decoding improve with increased response amplitudes as measured with fMRI? Single-unit recording studies have found that higher firing rates are accompanied by a proportional increase in variance, such that increases in stimulus contrast lead to relatively small improvements in the signal-to-noise ratio and transmission of orientation information (Skottun et al., 1987). However, fMRI reflects the averaged activity of large populations of neurons, thereby minimizing the impact of variability in individual neuronal firing rates. Noise levels in fMRI are largely dominated by other physiological sources, which include respiratory, vascular, and metabolic fluctuations and also pulsatile head motion (Kruger and Glover, 2001). Gratings of higher contrast should therefore evoke stronger orientation-selective activity patterns without increasing the variability of these fMRI responses, thereby increasing the signal-to-noise ratio. Consistent with this view, we found that fMRI noise levels remained stable across changes in stimulus contrast, based on estimates of the variability of mean amplitudes when measured across subjects (see error bars in Fig. 4A) and also across blocks within each subject (average standard deviation at 4%, 20% and 100% contrast, 0.37, 0.35 and 0.38 percent signal change, respectively).

Simulation analyses: comparison with fMRI decoding

We performed a simulation analysis to determine whether the mean amplitude of fMRI activity found in individual visual areas

could provide a good account of the accuracy of orientation decoding. Previously, we have shown that when random spatial jitter is applied to a regular array of orientation columns, it becomes possible to decode orientation-selective activity patterns at voxel resolutions because these local anisotropies lead to weak biases in orientation preference in the individual voxels (Kamitani and Tong, 2005). Here, we constructed a simple 1D model simulation of orientation columns, from which coarse-scale activity patterns were sampled using an array of 60 voxels (see Materials and methods for details). Our orientation classifier was then trained and tested on the pattern of responses from these 60 voxels (plus independent Gaussian noise) to repeated presentations of 45° and 135° orientations. The goal of this simulation was to determine whether scaling of the response amplitude of simulated orientation columns could effectively predict the reliability of fMRI orientation decoding.

In constructing this simulation, we attempted to use a minimum set of assumptions and free parameters. The same parameter values were used for all visual areas and contrast levels (e.g., number of voxels, number of samples, jitter in orientation preference, fMRI noise level), with the exception of two fixed parameters (columnar spacing and response amplitude) and one free parameter (columnar orientation tuning). The spacings between iso-orientation columns were determined individually for V1, V2 and V3 (760, 950 and 970 μm , respectively) based on previously reported values obtained from the monkey visual cortex (Vanduffel et al., 2002). The response amplitude of the simulated orientation columns was scaled according to the mean fMRI amplitude found in that contrast condition and visual area. The only free parameter consisted of the orientation tuning width of the cortical columns, which was adjusted separately for each visual area to obtain a good fit of the mean decoding accuracy averaged across all contrast levels. Thus, manipulations of this parameter cannot account for changes in performance as a function of stimulus contrast; such differences must instead reflect the change in response amplitude.

We obtained the following fitted values for simulated columns in areas V1, V2, and V3; orientation tuning was 38°, 58° and 64° respectively, based on the standard deviation of the Gaussian tuning curve. Such a decline in columnar orientation tuning when ascending the visual hierarchy is in rough agreement with the known properties of early visual areas. Results from Vanduffel et al. (2002) suggested that the orientation tuning width of cortical columns are about 50% more broad in areas V2/V3 than in area V1. Fig. 4B shows the results of decoding the simulated voxel responses to these orientation columns. The model provided a very good fit of the orientation decoding data by incorporating the mean amplitude of fMRI responses at each contrast level ($R^2 = 0.938$, $p < 0.0001$). From these results, we can conclude that the reliability of ensemble orientation-selective responses is highly dependent on the amplitude of stimulus-driven activity in the visual cortex, and that this relationship can be well modeled by scaling the amplitude of orientation-selective responses.

It should be noted that this model provides an excellent fit to the fMRI data over a range of initial parameters. This is because decoding accuracy depends on the signal-to-noise ratio of the differential activity, which is determined by the difference of Gaussians for columns tuned to different orientations. Thus, changes to one fixed parameter, such as iso-orientation spacing, can be compensated for by adjustments to the single free parameter of orientation tuning width to obtain differential patterns of about equal discriminability. For example, if we arbitrarily specify that areas V1–V3 should instead share the same iso-orientation spacing of 760 μm , then different estimates of columnar orientation tuning are obtained (V1–V3: 38°, 42°, 54° respectively). Nonetheless, with these parameters the model still provides an excellent fit of fMRI decoding performance for visual areas across changes in stimulus contrast ($R^2 = 0.931$, $p < 0.0001$). In the analyses described here, we report the results using columnar spacing values of 760, 950 and 970 μm for areas V1, V2 and V3, respectively.

Additional correlational analyses

We performed a variety of correlational analyses for more exploratory purposes, to determine whether the mean BOLD response amplitude in a given visual area might be predictive of the accuracy of orientation decoding performance for a given individual. The goal of these analyses was not to test for the effects of stimulus manipulations, such as the effects of contrast, but rather, to investigate whether individual differences in one measure might be predictive of another. For example, perhaps individuals who exhibit relatively greater response amplitudes in area V2, as compared to V1, also exhibit better decoding in V2, when tested with a common set of stimuli. Although the mean BOLD response does not contain reliable orientation-selective information per se, aspects of this gross signal, such as its amplitude or reliability, might be predictive of the reliability of orientation-selective patterns for an individual. With a sample size of 10 subjects, it was possible to explore such questions pertaining to individual differences in decoding performance.

As a first-pass analysis, we evaluated whether fMRI amplitudes might be predictive of decoding accuracy when pairs of data points were separately entered for each contrast level and visual area (V1–V3) of every subject. We observed a modest correlation ($R=0.36$, $p<10^{-6}$), but much of this apparent relationship was driven by the effect of stimulus contrast on fMRI amplitudes and decoding accuracy. Correlations were no longer reliable when tested within each contrast level (R values of 0.098, 0.106 and 0.077 for 100%, 20% and 4% contrast, respectively, $p>0.40$ in all cases).

We also explored the possibility that decoding performance for an individual might be related to the signal-to-noise ratio of their mean fMRI response. Instead of using response amplitudes, we used the ratio resulting from the mean response amplitude divided by the standard deviation to test for correlations. This signal-to-noise measure was quite strongly correlated with decoding accuracy when considering paired data points for every contrast level and visual area ($R=0.51$, $p<10^{-12}$). More importantly, we observed a reliable relationship within each contrast condition (R values of 0.48, 0.35, and 0.27 for 100%, 20% and 4% contrast respectively, $p<0.05$ in all cases). These results indicate that reliability of orientation-selective activity patterns can be partly predicted by the signal-to-noise ratio of fMRI responses for a given individual. These individual differences could potentially reflect differences at a neural level, such as differences in the stability of low-level visual responses or high-level cognitive effects of attention/arousal. Alternatively, it could reflect noise factors resulting from other physiological sources (e.g., head motion, stability of breathing and heart rate, strength of blood flow response) or even non-physiological sources such as the quality of the shim or stability of the MRI signal for that scanning session. The results of Experiment 1 provide compelling evidence of the relationship between orientation decoding performance and response amplitude across changes in stimulus contrast. We also find suggestive evidence of a relationship that might account for individual differences in decoding accuracy within a stimulus condition.

Experiment 2

In this experiment, we determined the extent to which fMRI activity patterns in individual visual areas could discriminate between orthogonal gratings shown at various spatial frequencies (0.25, 1 or 4 cpd). Higher spatial frequencies (above 5 cpd) were avoided because optical factors can impair sensitivity in this range (Banks et al., 1987). In the main experiment (Experiment 2A), we measured orientation classification performance for separate gratings presented in the left and right visual fields (cf. Fig. 1). Because these stimuli were presented in the periphery (1–9° eccentricity), we expected that contralateral visual areas would show better discrimination at moderate than high spatial frequencies. A follow-up study (Experiment 2B) was conducted to determine whether centrally presented stimuli (0–4° eccentricity) would lead to better orientation discrimination of high spatial frequency

gratings. Such differences in orientation discrimination would be expected due to the decline in spatial resolution found with increasing eccentricity in both the retina and cortex (De Valois et al., 1982; Enroth-Cugell and Robson, 1966; Hilz and Cavonius, 1974).

This experiment also allowed us to test for differences in spatial frequency sensitivity across early visual areas. We expected that V1 might show a relative advantage at discriminating orientations of high spatial frequency, in comparison to extrastriate areas. Previous studies have found evidence of a shift in tuning toward lower spatial frequencies when ascending the visual pathway (Foster et al., 1985; Gegenfurtner et al., 1997; Henriksson et al., 2008; Sasaki et al., 2001; Singh et al., 2000), suggestive of a gradual loss of fine detail information.

We further tested whether fMRI response amplitudes could predict the accuracy of orientation decoding, by applying both correlational and simulation analyses. Critically, for the simulation analysis, we used the same orientation tuning parameter values as were obtained from the fMRI data of Experiment 1, to predict fMRI classification accuracy in Experiment 2 for stimuli presented in a common visual location. A good fit across such changes in stimulus manipulation (i.e., stimulus contrast and spatial frequency) would provide additional support for the general validity of this simple modeling approach.

Materials and methods

All aspects of the experimental design, stimuli, MRI scanning procedures, and analysis were identical to those of Experiment 1, except as noted below.

Participants

Five participants from Experiment 1 (four male, one female) participated in Experiment 2A involving lateralized gratings. Four of these participants (three male, one female) went on to participate in Experiment 2B in which they viewed central presentations of a single grating. All subjects provided written informed consent. The study was approved by the Institutional Review Board at Vanderbilt University.

Experimental design and stimuli

In the main experiment, lateralized gratings of varying orientation (45° or 135°) were presented in the left and right visual fields (Fig. 1). In Experiment 2B, a single grating of varying orientation (45° or 135°) was centrally presented within a Gaussian envelope centered at the fovea (0–4° eccentricity). All gratings were presented at 100% contrast, with spatial frequency varied across runs (0.25, 1, or 4 cpd). For the foveal presentation experiment, voxels were selected from both left and right visual areas for classification analysis. Because retinotopic boundaries are not well defined in the foveal region, the regions of interest consisted of parafoveal and peripheral regions of V1–V4 that responded to the localizer stimulus.

Results

Fig. 5 shows orientation classification performance at each spatial frequency, plotted by visual area. Areas V1–V3 could discriminate the orientation of contralateral stimuli at above-chance levels for all spatial frequencies (min $t=3.16$, $p<0.05$), whereas discrimination of ipsilateral stimuli was very poor and at chance level. Again, orientation discrimination was quite poor in V3A/V4, with above-chance performance found in the 0.25 and 1 cpd conditions, but chance-level performance was found in the 4 cpd condition.

Focused analyses of areas V1–V3 indicated that orientation decoding was highly dependent on spatial frequency ($F=21.32$, $p<0.0001$), with much poorer discrimination at 4 cpd than at lower spatial frequencies ($t=8.57$, $p<0.001$). Overall classification performance also differed across visual areas ($F=9.98$, $p<0.001$), with better performance found for V1/V2 than V3 ($t=3.80$, $p<0.005$).

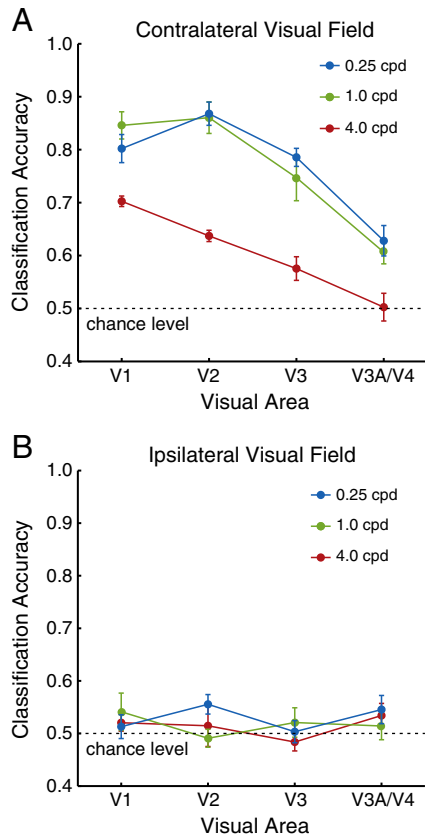


Fig. 5. Effects of spatial frequency on orientation classification performance, plotted by visual area for Experiment 2A. A, contralateral stimulation; B, ipsilateral stimulation. Gratings were presented in the left and right visual fields at 0.25, 1.0 or 4.0 cycles per degree (cpd). For contralateral regions, areas V1–V3 showed strong classification performance for orientations of low and moderate spatial frequencies, but poorer performance for high spatial frequencies. Nonetheless, V1 showed a relative advantage at discriminating 4 cpd gratings in this condition.

Of greater interest, the orientation classification accuracy found across visual areas seemed to depend on the spatial frequency content of the gratings, as indicated by a significant interaction effect ($F = 2.98, p < 0.05$). For high spatial frequency gratings of 4 cpd, orientation discrimination performance declined when progressing along the visual hierarchy from V1 to V3 (significant linear trend, $F = 11.828, p < 0.005$). At 1 cpd, performance was equally good in V1 and V2 ($t = 0.71, n.s.$) and worse in V3 (significant linear trend, $F = 9.73, p < 0.05$). The results obtained with 1-cpd gratings were very similar to those of Experiment 1 when the same stimulus conditions were tested (see Fig. 2, red curve). In the lowest spatial frequency condition of 0.25 cpd, orientation decoding was quite comparable across visual areas, with V3 performing as well as V1 ($t = 0.42, n.s.$), and V2 showing a relative advantage in performance (significant quadratic trend, $F = 6.62, p < 0.05$). Overall, the results indicate that V1 can discriminate high spatial frequency patterns better than higher extrastriate areas, whereas low spatial frequency patterns can be discriminated about equally well by areas V1–V3. Area V3 exhibited a relative advantage at discriminating orientations discriminating orientations with lower spatial frequencies.

We investigated the relationship between fMRI response amplitudes and orientation decoding accuracy. Plots of response amplitudes indicated that V1 responded most strongly to the 1 cpd gratings, whereas extrastriate areas responded more strongly to the low spatial frequency gratings (Fig. 6A). Mean decoding accuracy was highly correlated with the mean amplitude of activity found in areas V1 through V3 across the three spatial frequency conditions ($R = 0.767, p < 0.005$), indicating a positive relationship similar to that found in Experiment 1. This positive relationship was further supported by our decoding analysis of modeled

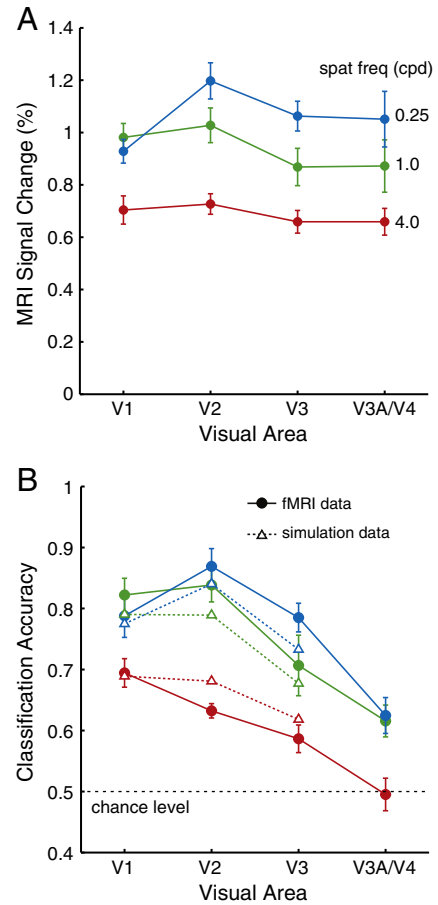


Fig. 6. Analysis of fMRI response amplitudes for Experiment 2A. A, mean response amplitudes for each spatial frequency condition and visual area. B, model fits of orientation classification accuracy based on the mean fMRI response amplitudes observed in Experiment 2A. Solid lines, original fMRI data; dashed lines, model fitted data. Mean fMRI responses were used to scale the response amplitude of simulated orientation columns, from which voxel-scale activity patterns were sampled and analyzed by a linear classifier. Simulation results provided a good fit of fMRI classification performance for areas V1–V3 ($R^2 = 0.925, p < 0.0001$).

voxel responses sampled from simulated orientation columns. Using the exact same parameter values and orientation tuning widths from Experiment 1, we obtained a good fit of orientation decoding performance based on the mean fMRI amplitudes found in each spatial frequency condition and visual area (Fig. 6B, $R^2 = 0.925, p < 0.0001$). These results indicate the predictive power of this rather simple model, and how the reliability of these orientation-selective activity patterns is highly dependent on the strength of the stimulus-driven activity. This proved true across manipulations of both stimulus contrast and spatial frequency.

We also tested whether our measures of ensemble orientation selectivity might reveal evidence of the bandpass tuning characteristics of visual areas. This was done by evaluating generalization performance across changes in spatial frequency. Fig. 7 shows generalization performance plotted by visual area, in which the decoder was trained on orientations of one spatial frequency and then tested on orientations of another spatial frequency. As expected, orientation classification performance was somewhat poorer across changes in spatial frequencies than when training and testing with the same spatial frequency (cf. Figs. 7 and 5A, respectively). Nonetheless, generalization performance for areas V1–V3 exceeded chance levels, with the exception of V3 in the 4 and 1 cpd conditions ($t = 1.6, n.s.$). V1 showed a relative advantage at generalizing between moderate and high spatial frequencies of 1 and 4 cpd (significant linear trend across V1–V3 in the 4 and 1 cpd condition, $F = 8.86, p < 0.05$), consistent with the notion that this region is sensitive to a higher range of




Fig. 7. Generalization of the orientation classifier across changes in spatial frequency for contralateral regions in areas V1 through V4. The classifier was trained on orientations of one spatial frequency and then tested on orientations of the other spatial frequency to evaluate generalization performance.

spatial frequencies. By contrast, extrastriate visual areas exhibited better generalization between the 0.25 and 1 cpd conditions, presumably because of their greater sensitivity to orientation at moderate and low spatial frequencies than to the high frequency 4 cpd condition. Generalization was significantly better for gratings differing by two octaves than for those differing by four octaves in the 0.25 and 4 cpd analyses ($t = 4.58, p < 0.001$).

These results indicate that more similar spatial frequencies allow for better generalization of orientation preference, as one would expect, but further demonstrate that some degree of orientation generalization is possible over a fairly broad spatial range spanning 4 octaves. Such generalization might partly reflect the bandwidth of individual orientation-selective neurons, as well as the tendency for neurons with similar orientation preference to form clusters or columns. Neurophysiological studies suggest that orientation-selective neurons in V1 are sensitive to a range of spatial frequencies, with bandwidths ranging from 0.5 to 5 octaves (Xing et al., 2004). However, recent optical imaging studies of the monkey visual cortex have found that orientation columns remain quite stable across large changes in spatial frequency of up to 8–16 octaves (Lu and Roe, 2007). The extensive bandwidth of the orientation map implies a high degree of clustering of neurons with a similar orientation preference across differences in spatial frequency preference.

Results for Experiment 2B: foveal gratings

In a follow-up experiment, we investigated whether gratings centered at the fovea would lead to stronger orientation-selective responses at high spatial frequencies. Previous studies have shown that spatial acuity is highest at the fovea and declines steadily toward the periphery, with estimates of peak sensitivity in the foveal region ranging from 3 to 8 cpd, depending on the study (Banks et al., 1987; De Valois et al., 1974; Hilz and Cavonius, 1974). A large circular grating was centrally presented, extending from 0–4° eccentricity. For our decoding analysis, we focused on the contributions of the parafovea and near periphery, since the boundaries between visual areas are difficult to determine in the foveal region.

Central presentation of the gratings led to a dramatic boost in the discrimination of high spatial frequency orientations throughout areas V1–V3. These areas showed excellent orientation classification performance across all spatial frequencies tested, including the highest 4-cpd condition (Fig. 8). Classification accuracy did not reliably differ as a function of spatial frequency ($F = 0.3$, n.s.), nor was there a significant interaction between spatial frequency and visual area ($F = 0.5$, n.s.). Thus, it appears that the cutoff sensitivity of areas V1–V3 exceeds 4 cpd in central vision.

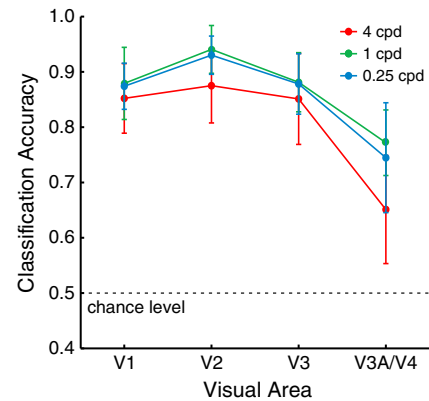


Fig. 8. Orientation classification performance for centrally presented gratings in Experiment 2B. Areas V1–V3 exhibited similarly high levels of orientation discrimination across all spatial frequencies tested, with much better performance for high spatial frequency orientations than was found for peripheral gratings in Experiment 2A.

Discussion

This study demonstrates that orientation decoding depends greatly on the strength and quality of the visual input. Early visual areas, especially V1, showed much better discrimination of high- than low-contrast oriented gratings, even though all stimuli were above perceptual threshold. This basic finding indicates that perception of orientation alone is not sufficient to evoke strong orientation-selective activity patterns and that these orientation-selective responses are not simply due to the perception of a global oriented pattern. Manipulations of spatial frequency provided further support for this view. Although the gratings were always suprathreshold orientation-selective responses were quite poor for high spatial frequency gratings presented in the near periphery but considerably better with central presentation. Thus, orientation decoding depends on the quality of the visual input and the spatial resolution available at the given location in the visual field.

More detailed analyses revealed that orientation classification performance was highly dependent on the amplitude of the stimulus-driven activity. Compelling evidence of this relationship was revealed by both correlational analyses and decoding applied to simulated voxel responses to orientation columns. Correlations between mean response amplitudes and decoding accuracy were reasonably high ($R > 0.70$), when considered across stimulus conditions and visual areas. An even better fit of mean decoding performance was obtained by incorporating response amplitudes in a simple simulation model of voxel responses to spatially irregular arrays of orientation columns. When fMRI response amplitudes were used to determine the response strength of the simulated orientation columns, decoding of the resulting voxel responses provided a very good fit of the actual fMRI results. Model predictions accounted for more than 90% of the variance observed across experimental conditions for areas V1–V3. This proved true even when the parameter values obtained from Experiment 1 were used to predict decoding accuracy based on the response amplitudes observed in Experiment 2. Our results indicate that the reliability of the voxel-scale activity patterns can be reasonably modeled by incorporating changes in the response amplitude of simulated orientation columns. These findings provide compelling evidence of the close relationship between fMRI response amplitudes and the strength of these feature-selective responses under stimulus-driven conditions.

Previous work has shown that voxel response amplitudes are an important consideration for effective decoding. For example, the classification of whole-brain activity patterns can be improved considerably by implementing *feature selection* to restrict the analysis to those voxels that are most active across all of the experimental conditions to be discriminated (Mitchell et al., 2004). In our previous (Kamitani and Tong,

2005, 2006) and present works, we preselected voxels in individual visual areas based on their responses to independent visual localizer runs to ensure statistical independence of our methods for feature selection and pattern classification (Kriegeskorte et al., 2009). This approach allows us to identify voxels that correspond to the retinotopic location of the orientation stimulus, and typically we find that decoding accuracy steadily improves as more voxels are added (Fig. 4), until the region of interest begins to exceed the boundaries of the cortical region that receives direct stimulus input. Since we analyzed the same set of voxels across manipulations of stimulus contrast or spatial frequency, any changes in decoding performance that we observed must be attributable to changes in the strength of the underlying orientation-selective activity patterns within the prespecified set of voxels.

Of considerable relevance to the current work is a recent fMRI study that reported monotonic improvements in orientation decoding as a function of stimulus contrast (Smith et al., 2011). The authors found that classification performance in early visual areas improved with stimulus-driven increases in BOLD response amplitude, approximately following a power law with a 0.2 exponent due to the compressive non-linearity that occurred at higher amplitudes. In the present study, we conducted in-depth analyses to better understand the relationship between response amplitude and orientation decoding accuracy, and demonstrated that decoding of orientation responses can be well predicted by incorporating fMRI response amplitudes into a simple simulation model. Using the parameter settings obtained from the first experiment, we were able to extend this approach to predict orientation classification performance for gratings of varying spatial frequency in a second experiment. These findings indicate the predictive power of our modeling approach to account for orientation decoding performance.

Our study suggests that the effects of response amplitude on orientation decoding can be understood by incorporating principles of visual cortical organization. Whereas Smith et al. interpreted the effects of response amplitude on orientation decoding as a confounding variable that limits the ability to compare pattern classification results across visual areas, we consider response amplitude to be one of several factors that can impact the accuracy of orientation decoding for a given visual area. As is indicated by our model, multiple factors can readily affect decoding performance for a given visual area, including columnar size, iso-orientation spacing, spatial variability in the distribution of columns, sharpness of columnar orientation tuning, and the amount of neural or fMRI noise present in a visual area. Given that these factors should remain constant throughout an experiment, it should be possible to investigate whether visual areas are differentially affected by systematic manipulations of a stimulus.

Our fMRI decoding approach was sufficiently sensitive to detect differences between visual areas across stimulus manipulations. In Experiment 1, we found that orientation-selective responses in V1 showed greater dependence on stimulus contrast than higher extrastriate areas, such as V3. These findings are generally consistent with previous reports that fMRI response amplitudes tend to become more “contrast-invariant” at progressively higher levels of the visual hierarchy. However, evidence of a difference in contrast sensitivity between early visual areas, such as V1 and V3, has been mixed, with several studies unable to find a reliable difference between these early visual areas (Avidan et al., 2002; Boynton et al., 1999; Gardner et al., 2005; Tootell et al., 1995). In Experiment 2, we found evidence of a selective attenuation of high spatial frequency information as signals propagated up the visual hierarchy. Areas V1–V3 showed equally good orientation discrimination of low frequency patterns presented at mid-eccentricities, whereas V1 showed a relative advantage at discriminating high frequency patterns. These results are consistent with previous single-unit studies in monkeys, which have reported differences in spatial frequency tuning between V1 and V2 (Foster et al., 1985; Levitt et al., 1994), and also between V2 and V3 (Gegenfurtner et al., 1997). Previous neuroimaging studies of response amplitudes have found that spatial frequency preference is most strongly influenced by visual eccentricity (Tootell et al., 1998a), with some studies

reporting evidence of a modest shift toward lower spatial frequencies when ascending the visual hierarchy (Henriksson et al., 2008; Kay et al., 2008; Sasaki et al., 2001; Singh et al., 2000).

Although we observe a strong relationship between orientation decoding accuracy and BOLD response amplitude under these stimulus-driven conditions, it is important to note that this relationship is unlikely to prevail under conditions of pure top-down processing. In a recent study of visual working memory, we discovered that activity patterns in early visual areas contain reliable orientation information even after the overall amplitude of activity has fallen to baseline levels (Harrison and Tong, 2009). It is conceivable that inhibitory mechanisms, elicited by some forms of top-down feedback, might lead to a decoupling between overall response amplitudes and the reliability of the orientation-selective activity patterns when observers must retain a grating in working memory.

It is also interesting to note that our measures of orientation-selective responses in the human visual cortex are in generally good agreement with a recent optical imaging study of the macaque monkey (Lu and Roe, 2007). These authors examined the effects of contrast and spatial frequency on the response of orientation columns in areas V1 and V2. Higher stimulus contrasts led to stronger orientation responses in both areas, though V2 responses tended to saturate at high contrasts indicating relatively greater contrast invariance. V1 also preferred somewhat higher spatial frequencies than V2. Orientation columns in V1 were most evident at moderate spatial frequencies of 1.7 cpd, weak but visible at 3.4 cpd, and no longer reliable at very high spatial frequencies of 6.7 cpd. These results are consistent with our findings that human V1 exhibited more modest orientation selectivity at high spatial frequencies of 4 cpd, but still exhibited a relative advantage in decoding performance when compared to extrastriate areas V2 and V3. Interestingly, orientation columns in monkey V1 remained quite stable over a wide range of spatial frequencies (0.2–3.4 cpd). Such generalization suggests that neurons with similar orientation preference (but different spatial frequency preferences) tend to cluster in cortical columns. This might account for the above-chance generalization performance we observed across a 4-octave difference in spatial frequency in our study. The fact that fMRI orientation decoding proved sufficiently sensitive to detect differences in contrast and spatial frequency sensitivity across visual areas indicates that this approach could provide an effective tool for investigating the feature-selective response properties of the human visual system. A related approach could be to estimate the conjoint orientation and spatial frequency tuning of individual voxels within a visual area (Kay et al., 2008), or to assess the amount of orientation information available in high-resolution fMRI activity patterns (Swisher et al., 2010).

In recent years, multivariate pattern analysis has become an increasingly popular method for characterizing fMRI activity in studies of human perception and cognition (Tong and Pratte, 2012). Although most studies have used these methods to investigate the neural representation of sensory stimuli or mental states, one might also ask whether decoding is likely to be more successful in some individuals than others. Here, we found evidence that the signal-to-noise ratio of the BOLD response is modestly predictive of individual differences in decoding accuracy. It might be interesting for future studies to explore what other factors contribute to individual differences in the reliability of classification performance. In the case of orientation decoding, some factors to explore might include the size or cortical magnification properties of the visual area in question (Duncan and Boynton, 2003), the fine-scale orientation structure in these regions (Swisher et al., 2010; Yacoub et al., 2008), as well as the degree to which global orientation biases might be present. Although a bias for radial orientations can indeed be found in the primary visual cortex (Freeman et al., 2011; Sasaki et al., 2006), spiral patterns that lack any such bias nonetheless allow for reliable orientation decoding (Mannion et al., 2009). It could prove useful to characterize how local and more global orientation biases tend to vary from one individual to another, and to investigate whether these factors might account for individual differences in the reliability of fMRI orientation decoding.

Acknowledgments

This research was supported by a grant from the National Eye Institute of the U.S. National Institutes of Health (NEI R01 EY017082) to F.T. a graduate scholarship from the Natural Sciences and Engineering Research Council of Canada to S.H. and a grant from the Ministry of Education, Culture, Sports, Science and Technology in Japan (Strategic Research Program for Brain Science) to Y.K. We thank D. K. Brady for technical assistance and the Vanderbilt University Institute for Imaging Science for MRI support.

References

- Avidan, G., Harel, M., Hendler, T., Ben-Bashat, D., Zohary, E., Malach, R., 2002. Contrast sensitivity in human visual areas and its relationship to object recognition. *J. Neurophysiol.* 87, 3102–3116.
- Banks, M.S., Geisler, W.S., Bennett, P.J., 1987. The physical limits of grating visibility. *Vision Res.* 27, 1915–1924.
- Blasdel, G.G., Salama, G., 1986. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature* 321, 579–585.
- Bonhoeffer, T., Grinvald, A., 1991. Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* 353, 429–431.
- Boynton, G.M., Finney, E.M., 2003. Orientation-specific adaptation in human visual cortex. *J. Neurosci.* 23, 8781–8787.
- Boynton, G.M., Demb, J.B., Glover, G.H., Heeger, D.J., 1999. Neuronal basis of contrast discrimination. *Vision Res.* 39, 257–269.
- Brainard, D.H., 1997. The psychophysics toolbox. *Spat. Vis.* 10, 433–436.
- Brouwer, G.J., van Ee, R., 2007. Visual cortex allows prediction of perceptual states during ambiguous structure-from-motion. *J. Neurosci.* 27, 1015–1023.
- De Valois, R.L., Morgan, H., Snodderly, D.M., 1974. Psychophysical studies of monkey vision-III. Spatial luminance contrast sensitivity tests of macaque and human observers. *Vision Res.* 14, 75–81.
- De Valois, R.L., Albrecht, D.G., Thorell, L.G., 1982. Spatial frequency selectivity of cells in macaque visual cortex. *Vision Res.* 22, 545–559.
- Desimone, R., Schein, S.J., 1987. Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *J. Neurophysiol.* 57, 835–868.
- Duncan, R.O., Boynton, G.M., 2003. Cortical magnification within human primary visual cortex correlates with acuity thresholds. *Neuron* 38, 659–671.
- Engel, S.A., Glover, G.H., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex* 7, 181–192.
- Enroth-Cugell, C., Robson, J.G., 1966. The contrast sensitivity of retinal ganglion cells of the cat. *J. Physiol.* 187, 517–552.
- Fang, F., Murray, S.O., Kersten, D., He, S., 2005. Orientation-tuned fMRI adaptation in human visual cortex. *J. Neurophysiol.* 94, 4188–4195.
- Felleman, D.J., Van Essen, D.C., 1987. Receptive field properties of neurons in area V3 of macaque monkey extrastriate cortex. *J. Neurophysiol.* 57, 889–920.
- Field, D.J., Hayes, A., Hess, R.F., 1993. Contour integration by the human visual system: evidence for a local “association field”. *Vision Res.* 33, 173–193.
- Foster, K.H., Gaska, J.P., Nagler, M., Pollen, D.A., 1985. Spatial and temporal frequency selectivity of neurons in visual cortical areas V1 and V2 of the macaque monkey. *J. Physiol.* 365, 331–363.
- Freeman, J., Brouwer, G.J., Heeger, D.J., Merriam, E.P., 2011. Orientation decoding depends on maps, not columns. *J. Neurosci.* 31, 4792–4804.
- Gardner, J.L., Sun, P., Waggoner, R.A., Ueno, K., Tanaka, K., Cheng, K., 2005. Contrast adaptation and representation in human early visual cortex. *Neuron* 47, 607–620.
- Gegenfurtner, K.R., Kiper, D.C., Levitt, J.B., 1997. Functional properties of neurons in macaque area V3. *J. Neurophysiol.* 77, 1906–1923.
- Harrison, S.A., Tong, F., 2009. Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458, 632–635.
- Haynes, J.D., Rees, G., 2005a. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat. Neurosci.* 8, 686–691.
- Haynes, J.D., Rees, G., 2005b. Predicting the stream of consciousness from activity in human visual cortex. *Curr. Biol.* 15, 1301–1307.
- Haynes, J.D., Rees, G., 2006. Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* 7, 523–534.
- Henriksson, L., Nurminen, L., Hyvarinen, A., Vanni, S., 2008. Spatial frequency tuning in human retinotopic visual areas. *J. Vis.* 8 (5), 1–13.
- Hilz, R., Cavonius, C.R., 1974. Functional organization of the peripheral retina: sensitivity to periodic stimuli. *Vision Res.* 14, 1333–1337.
- Hubel, D.H., Wiesel, T.N., 1968. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* 195, 215–243.
- Jehee, J.F., Brady, D.K., Tong, F., 2011. Attention improves encoding of task-relevant features in the human visual cortex. *J. Neurosci.* 31, 8210–8219.
- 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.
- Kastner, S., O'Connor, D.H., Fukui, M.M., Fehd, H.M., Herwig, U., Pinsk, M.A., 2004. Functional imaging of the human lateral geniculate nucleus and pulvinar. *J. Neurophysiol.* 91, 438–448.
- Kay, K.N., Naselaris, T., Prenger, R.J., Gallant, J.L., 2008. Identifying natural images from human brain activity. *Nature* 452, 352–355.
- Kovacs, I., Julesz, B., 1993. A closed curve is much more than an incomplete one: effect of closure in figure-ground segmentation. *Proc. Natl. Acad. Sci. U. S. A.* 90, 7495–7497.
- Kriegeskorte, N., Bandettini, P., 2007. Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage* 38, 649–662.
- Kriegeskorte, N., Simmons, W.K., Bellgowan, P.S., Baker, C.I., 2009. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat. Neurosci.* 12, 535–540.
- Kruger, G., Glover, G.H., 2001. Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magn. Reson. Med.* 46, 631–637.
- Levitt, J.B., Kiper, D.C., Movshon, J.A., 1994. Receptive fields and functional architecture of macaque V2. *J. Neurophysiol.* 71, 2517–2542.
- Lu, H.D., Roe, A.W., 2007. Optical imaging of contrast response in macaque monkey V1 and V2. *Cereb. Cortex* 17, 2675–2695.
- Mannion, D.J., McDonald, J.S., Clifford, C.W., 2009. Discrimination of the local orientation structure of spiral glass patterns early in human visual cortex. *Neuroimage* 46, 511–515.
- Mitchell, T.M., Hutchinson, R., Niculescu, R.S., Pereira, F., Wang, X.R., Just, M., Newman, S., 2004. Learning to decode cognitive states from brain images. *Mach. Learn.* 57, 145–175.
- Murray, S.O., Olman, C.A., Kersten, D., 2006. Spatially specific fMRI repetition effects in human visual cortex. *J. Neurophysiol.* 95, 2439–2445.
- 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.
- Ohki, K., Chung, S., Kara, P., Hubener, M., Bonhoeffer, T., Reid, R.C., 2006. Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442, 925–928.
- Pelli, D.G., 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* 10, 437–442.
- Regan, D., Beverley, K.I., 1985. Postadaptation orientation discrimination. *J. Opt. Soc. Am. A* 2, 147–155.
- Sasaki, Y., Hadjikhani, N., Fischl, B., Liu, A.K., Marrett, S., Dale, A.M., Tootell, R.B., 2001. Local and global attention are mapped retinotopically in human occipital cortex. *Proc. Natl. Acad. Sci. U. S. A.* 98, 2077–2082.
- Sasaki, Y., Rajimehr, R., Kim, B.W., Ekstrom, L.B., Vanduffel, W., Tootell, R.B., 2006. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron* 51, 661–670.
- Scal, G., Maunsell, J.H., Lennie, P., 1990. Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res.* 30, 1–10.
- 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., Ester, E.F., Vogel, E.K., Awh, E., 2009. Stimulus-specific delay activity in human primary visual cortex. *Psychol. Sci.* 20, 207–214.
- 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.
- Singh, K.D., Smith, A.T., Greenlee, M.W., 2000. Spatiotemporal frequency and direction sensitivities of human visual areas measured using fMRI. *Neuroimage* 12, 550–564.
- Skottun, B.C., Bradley, A., Scal, G., Ohzawa, I., Freeman, R.D., 1987. The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *J. Neurophysiol.* 57, 773–786.
- Smith, A.T., Kossilo, P., Williams, A.L., 2011. The confounding effect of response amplitude on MVPA performance measures. *Neuroimage* 56, 525–530.
- Swisher, J.D., Gatenby, J.C., Gore, J.C., Wolfe, B.A., Moon, C.H., Kim, S.G., Tong, F., 2010. Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. *J. Neurosci.* 30, 325–330.
- Tong, F., Pratte, M.S., 2012. Decoding patterns of human brain activity. *Annu. Rev. Psychol.* 63, 483–509.
- Tootell, R.B., Reppas, J.B., Kwong, K.K., Malach, R., Born, R.T., Brady, T.J., Rosen, B.R., Belliveau, J.W., 1995. Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *J. Neurosci.* 15, 3215–3230.
- Tootell, R.B., Hadjikhani, N.K., Mendola, J.D., Marrett, S., Dale, A.M., 1998a. From retinotopy to recognition: fMRI in human visual cortex. *Trends Cogn. Sci.* 2, 174–183.
- Tootell, R.B., Hadjikhani, N.K., Vanduffel, W., Liu, A.K., Mendola, J.D., Sereno, M.I., Dale, A.M., 1998b. Functional analysis of primary visual cortex (V1) in humans. *Proc. Natl. Acad. Sci. U. S. A.* 95, 811–817.
- 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.
- Vapnik, V.N., 1998. *Statistical Learning Theory*. Wiley, New York.
- Xing, D., Ringach, D.L., Shapley, R., Hawken, M.J., 2004. Correlation of local and global orientation and spatial frequency tuning in macaque V1. *J. Physiol.* 557, 923–933.
- Yacoub, E., Harel, N., Ugurbil, K., 2008. High-field fMRI unveils orientation columns in humans. *Proc. Natl. Acad. Sci. U. S. A.* 105, 10607–10612.