

Orientation Decoding Depends on Maps, Not Columns

Jeremy Freeman,¹ Gijs Joost Brouwer,^{1,2} David J. Heeger,^{1,2} and Elisha P. Merriam^{1,2}

¹Center for Neural Science and ²Department of Psychology, New York University, New York, New York 10003

The representation of orientation in primary visual cortex (V1) has been examined at a fine spatial scale corresponding to the columnar architecture. We present functional magnetic resonance imaging (fMRI) measurements providing evidence for a topographic map of orientation preference in human V1 at a much coarser scale, in register with the angular-position component of the retinotopic map of V1. This coarse-scale orientation map provides a parsimonious explanation for why multivariate pattern analysis methods succeed in decoding stimulus orientation from fMRI measurements, challenging the widely held assumption that decoding results reflect sampling of spatial irregularities in the fine-scale columnar architecture. Decoding stimulus attributes and cognitive states from fMRI measurements has proven useful for a number of applications, but our results demonstrate that the interpretation cannot assume decoding reflects or exploits columnar organization.

Introduction

The orientations of visual features are represented in an orderly pinwheel-like progression within each hypercolumn ($\sim 1 \times 1$ mm in monkey and $\sim 2 \times 2$ mm in human) of primary visual cortex (V1) (Hubel and Wiesel, 1963; Blasdel and Salama, 1986; Grinvald et al., 1986; Das and Gilbert, 1997; Maldonado et al., 1997; Ohki et al., 2006; Adams et al., 2007). Less is known, however, about representations of orientation at a coarse spatial scale.

Several studies have shown that it is possible to assess orientation selectivity in human V1 using functional magnetic resonance imaging (fMRI). fMRI pools cortical activity over several millimeters, so the underlying orientation-selective neural responses might cancel out, resulting in fMRI responses that are not orientation selective. Rather, it has been demonstrated, using multivariate statistical analyses (e.g., using a classifier to decode stimulus orientation from the spatial pattern of activity across many fMRI voxels), that individual fMRI voxels exhibit small but reliable orientation preferences (Haynes and Rees, 2005; Kamitani and Tong, 2005; Kay et al., 2008). The source of these orientation preferences, however, remains controversial.

A leading conjecture is that the orientation preferences in fMRI measurements arise from random spatial irregularities in the fine-scale columnar architecture (Boynton, 2005; Haynes and Rees, 2005; Kamitani and Tong, 2005; Op de Beeck et al., 2008). According to this account, multivariate statistical analyses, which are exceptionally sensitive, exploit these small biases to reveal orientation-specific signals, thereby providing a window into subvoxel columnar structure. It has been difficult to confirm or

reject this conjecture (Mannion et al., 2009; Gardner, 2010; Kamitani and Sawahata, 2010; Kriegeskorte et al., 2010; Op de Beeck, 2010; Shmuel et al., 2010; Swisher et al., 2010).

An alternative possibility is that decoding accuracy, at least for orientation in V1, reflects a coarse-scale organization, such as a global preference for cardinal orientations (Furmanski and Engel, 2000; Serences et al., 2009) or for radial orientations (i.e., orientations that point toward the fovea) (Sasaki et al., 2006; Mannion et al., 2010).

Motivated by the multivariate statistical analysis of orientation selectivity in human V1, a wealth of fMRI studies have used similar methods to probe selectivity for various stimulus features and cognitive processes throughout the brain (e.g., Haynes and Rees, 2006; Norman et al., 2006; Op de Beeck et al., 2008). Resolving the source of the orientation preferences in fMRI measurements from V1 would help guide the interpretation of the rapidly growing number of results based on multivariate statistical analyses in other brain areas.

We used fMRI to investigate the coarse-scale organization of orientation in human V1 and found a topographic map of radial orientations, confirming and extending previous reports of a radial bias (Sasaki et al., 2006; Mannion et al., 2010). The orientation map was tightly colocalized with the angular-position component of the retinotopic map in V1 (Engel et al., 1994; Sereno et al., 1995; DeYoe et al., 1996). This coarse-scale map of orientation was both necessary and sufficient for orientation decoding. Hence, our results do not support the conjecture that orientation decoding, at conventional spatial sampling resolutions ($\geq 2 \times 2 \times 2$ mm voxels), reflects the fine-scale columnar architecture. Rather, the coarse-scale orientation map provides a parsimonious, but sobering, explanation for fMRI classification results: orientation decoding in V1 is not an example of how fMRI might probe representations at fine (i.e., subvoxel) spatial scales.

Materials and Methods

Subjects. Data were acquired from four healthy subjects with normal or corrected-to-normal vision (three males; age range, 24–34 years). Three

Received Sept. 30, 2010; revised Feb. 10, 2011; accepted Feb. 11, 2011.

Author contributions: J.F., G.J.B., D.J.H., and E.P.M. designed research; J.F., G.J.B., and E.P.M. performed research; J.F., G.J.B., and E.P.M. analyzed data; J.F., D.J.H., and E.P.M. wrote the paper.

This work was supported by National Institutes of Health Grants R01-EY016752 (D.J.H.) and R01-NS047493 (E.M. and D.J.H.) and by a National Science Foundation Graduate Student Fellowship (J.F.). We thank Michael Landy, Clay Reid, Christopher Genovesi, E. J. Chichilnisky, and Yan Karklin for helpful suggestions.

Correspondence should be addressed to either Jeremy Freeman or Elisha P. Merriam, New York University, 6 Washington Place, New York, NY 10003. E-mail: freeman@cns.nyu.edu or eli@cns.nyu.edu.

DOI:10.1523/JNEUROSCI.5160-10.2011

Copyright © 2011 the authors 0270-6474/11/314792-13\$15.00/0

subjects were authors (S1–S3). Experiments were conducted with the written consent of each subject and in accordance with the safety guidelines for fMRI research, as approved by the University Committee on Activities Involving Human Subjects at New York University. Each subject participated in at least three scanning sessions: one session to obtain a set of high-resolution anatomical volumes, one session for standard retinotopic mapping (single-wedge, angular-position, and expanding-ring eccentricity), and one double-wedge, angular-position retinotopic mapping session for comparison with orientation maps. Three subjects (S1–S3) participated in the main orientation mapping experiment with annular stimulus. Three subjects (S1–S3) participated in the blurred-edge orientation mapping control experiment. Three subjects (S1, S2, and S4) participated in the full-field sinusoidal orientation mapping control experiment. Two subjects (S1 and S3) participated in the square-wave orientation mapping control experiment. Two subjects (S1 and S3) participated in a second double-wedge, angular-position retinotopic mapping session. One subject (S1) participated in a repeat of the main orientation mapping experiment with annular stimuli. One subject (S2) participated in an event-related orientation mapping experiment.

Stimuli. Stimuli were generated using Matlab (MathWorks) and MGL (available at <http://justingardner.net/mgl>) on a Macintosh computer. Stimuli were displayed via an LCD projector onto a back-projection screen in the bore of the magnet. Subjects were supine and viewed the projected stimuli through an angled mirror (maximum eccentricity of 12° of visual angle).

Orientation mapping experiments. The stimulus was a large, oriented sinusoidal grating (spatial frequency, 0.5 cycles per degree) presented within a 5° peripheral annulus (inner radius, 4.5°; outer radius, 9.5°). Both the inner and outer edges of the annulus were blurred with a 1° raised cosine transition (centered on the inner and outer edges) from 100 to 0% contrast. The spatial phase of the grating was randomized every 250 ms from a predefined set of 16 phases uniformly distributed between 0 and 2π . Regions outside the annulus were a uniform gray, equal to the mean luminance of the gratings (526 cd/m²). The orientation of the grating cycled through 16 evenly spaced angles between 0° and 180° (1.5 s per orientation). The stimulus completed 10.5 cycles in each run. Each cycle was 24 s long. Subjects completed 14–18 runs in each scanning session. The stimuli cycled clockwise in half of the runs and counterclockwise in the other half.

Orientation mapping control experiments. The first control experiment tested whether the orientation map was due to a potential confound related to the visible edge of the stimulus. The stimulus consisted of the same annular oriented sinusoidal grating as in the main experiment except that the stimulus edge was blurred over a much larger area so that there was no visible edge. Both edges were blurred with a raised cosine transition region (from 100 to 0% contrast). The transition region was 5° for the inner edge (extending from 2° to 7° eccentricity) and 5° for the outer edge (extending from 7° to 12°), such that the stimulus was 0% contrast at 2°, 100% contrast at 7°, and 0% contrast at the outermost edge of the screen (12° eccentricity). The grating thus achieved 100% contrast only at a single eccentricity (7°), and the region of blurring (5°) was more than two times the spatial period of the sinusoid (2°).

The second control experiment tested whether the map varied as a function of eccentricity by using a stimulus that included the central part of the visual field (9° annulus; inner radius, 0.5°; outer radius, 9.5°). The stimulus was otherwise identical to that in the main experiment. The third control experiment tested whether the map generalized to the stimulus parameters used in other fMRI studies (Kamitani and Tong, 2005; Swisher et al., 2010). The stimulus was a square-wave grating (rather than sinusoidal) with a higher spatial frequency of 1.4 cycles per degree [matched to the stimulus used by Kamitani and Tong (2005)]. This stimulus also included the central part of the visual field (9° annulus; inner radius, 0.5°; outer radius, 9.5°).

In a fourth control experiment, we measured orientation responses using an event-related protocol. The stimulus was a contrast-reversing

oriented sinusoidal grating (spatial frequency, 1 cycle per degree; temporal frequency, 1.33 cycles per degree), presented within a 7.5° annulus (inner radius, 0.5°; outer radius, 8°). Stimuli were presented at six different orientations (0°, 30°, 60°, 90°, 120°, and 150°) in randomized order. Each stimulus presentation was 1.5 s in duration, interleaved with interstimulus intervals that ranged from 3 to 6 s, in steps of 1.5 s. All six orientations were presented eight times in each run, along with eight blank trials. Standard methods were used to estimate the hemodynamic impulse response function and to estimate a response amplitude to each orientation, separately for each voxel [see Brouwer and Heeger (2009) for a description of a similar analysis]. To estimate the preferred orientation of each voxel, we created six unit vectors with angles equally distributed between 0 and 2π . These six vectors were multiplied by the six response amplitudes to each presented orientation and were averaged to a single vector. The direction of the resulting vector represented the preferred orientation in radians (between 0 and 2π). Multiplying by $180/2\pi$ converted the preferred orientation to degrees (between 0° and 180°).

Angular-position mapping experiments. For angular-position mapping, the stimulus was a black-and-white, high-contrast radial checkerboard pattern presented within a pair of wedges on opposite sides of fixation. Each wedge occupied 45° of angular position, and the wedges were restricted to a 5° peripheral annulus (inner radius, 4.5°; outer radius, 9.5°). No blurring was applied to the edges of the wedges. We used two wedges to equate the phase range of the orientation and angular-position maps; in both cases, one full cycle covered 180° of the stimulus variable (angular position or orientation). The angular position of the wedges rotated through 16 evenly spaced angular positions between 0° and 180° (1.5 s per angular position). Each radial strip of the checkerboard pattern within the wedge moved randomly inward or outward on each stimulus frame at a speed of 2°/s, giving rise to vivid motion boundaries between adjacent strips. Regions outside the wedges were a uniform gray, equal to the mean luminance of the checkerboard patterns within the wedges. As in the orientation mapping experiments, the stimulus completed 10.5 cycles in each run, each cycle was 24 s long, subjects completed 14–18 runs in each session, and the stimuli cycled clockwise in half of the runs and counterclockwise in the other half.

Task. Observers performed a demanding two-back detection task continuously throughout each run to maintain a consistent behavioral state, encourage stable fixation, and divert attention from the peripheral stimulus. A continuous sequence of digits (0 to 9) was displayed at fixation, changing every 400 ms. The observer's task was to indicate, by means of a button press, whether the current digit matched the digit from two steps earlier.

MRI acquisition. MRI data were acquired on a Siemens 3T Allegra head-only scanner using a head coil (NM-011; Nova Medical) for transmitting and an eight-channel phased array surface coil (Nova Medical) for receiving. Functional scans were acquired with gradient recalled echo-planar imaging to measure blood oxygen level-dependent changes in image intensity (Ogawa et al., 1990). Functional imaging was conducted with 24 slices oriented perpendicular to the calcarine sulcus and positioned with the most posterior slice at the occipital pole (1500 ms repetition time; 30 ms echo time; 72° flip angle; $2 \times 2 \times 2$ mm voxel size; 104×80 voxel grid). A T1-weighted magnetization-prepared rapid gradient echo anatomical volume (MPRAGE) was acquired in each scanning session with the same slice prescriptions as the functional images (1530 ms repetition time; 3.8 ms echo time; 8° flip angle; $1 \times 1 \times 2.5$ mm voxel size; 256×160 voxel grid). A high-resolution anatomical volume, acquired in a separate session, was the average of three MPRAGE scans that were aligned to one another and averaged (2500 ms repetition time; 3.93 ms echo time; 8° flip angle; $1 \times 1 \times 1$ mm voxel size; 256×256 voxel grid). This high-resolution anatomical scan was used both for registration across scanning sessions and for gray matter segmentation and cortical flattening.

Eye tracking. In one orientation mapping session, we measured gaze positions during the fMRI experiment. The eye tracker (EyeLink

1000, SR Research) infrared camera was placed in the bore, recording the pupil centroid and corneal reflection. Raw gaze positions were calibrated and converted to degrees of visual angle (Stampe, 1993). Preprocessing involved only blink removal: all samples during and shortly (100 ms) preceding and following blinks were excluded from analysis. Eye position data were divided into segments corresponding to each stimulus orientation. Within each of these segments, average gaze position was computed over each period of stable fixation (defined as periods in which gaze velocity was $<22^\circ/\text{s}$ and gaze drift was $<1^\circ$). The angular positions of averaged gaze positions were rotated so that their range matched the range of stimulus orientations. Finally, circular correlation, Equation 5, was used to evaluate whether gaze position angular position varied as a function of stimulus orientation (supplemental Fig. 1, available at www.jneurosci.org as supplemental material).

Defining retinotopic regions of interest. Each subject participated in a standard retinotopic mapping experiment, explained in great detail previously (Larsson and Heeger, 2006; Gardner et al., 2008). Standard traveling wave methods were used to identify meridian representations corresponding to the borders between retinotopically organized visual areas V1 and V2. For each subject, we defined an annular subregion of area V1, including voxels responding preferentially to stimuli between 5° and 10° . These regions of interest (ROIs) were defined with data that were statistically independent of those from the main experiment (i.e., measured in a separate scanning session). The annular V1 ROI was used in the correlation analyses shown in Figure 3A and in the classification analyses shown in Figure 5. We also defined a rectangular ROI to examine the effect of filtering on classification (see below, Spatial filtering). This ROI was the minimum rectangular volume that included the entire annular V1 ROI. Because the cortical surface is folded, this volume necessarily included additional voxels in the fovea and far periphery of V1, along with voxels from other visual areas (e.g., V2) and other tissue types (e.g., white matter). This rectangular ROI was used only for the classification analysis shown in Figure 6.

Preprocessing. The anatomical volume acquired in each scanning session was aligned to the high-resolution anatomical volume of the same subject's brain, using a robust image registration algorithm (Nestares and Heeger, 2000). Data from the first half cycle (eight frames) of each functional run were discarded to minimize the effect of transient magnetic saturation and allow the hemodynamic response to reach steady state. Head movement within and across scans was compensated for using standard procedures (Nestares and Heeger, 2000). The time series from each voxel was divided by its mean to convert from arbitrary intensity units to percentage modulation and was high-pass filtered (cutoff, 0.01 Hz) to remove low-frequency noise and drift (Smith et al., 1999). Data acquired with different stimulus rotation directions (clockwise or counterclockwise) were combined to estimate the response phase, independent of the lag caused by the hemodynamic delay. Time series data from each scan were shifted back in time by three frames. The time series for the counterclockwise runs were then time reversed. Averaging clockwise and time-reversed counterclockwise runs cancelled the residual hemodynamic lag (Engel et al., 1997; Kalatsky and Stryker, 2003).

Maps of orientation and angular-position preference. The average time series of each voxel was fit with a sinusoid with period matched to the period of stimulus rotations (24 s). The phase of the best-fitting sinusoid indicated the angular-position (or orientation) preference of the voxel. The responses of most voxels were very well fit by a sinusoidal model, and there was no evidence for systematic deviation from a sinusoid (e.g., multiple peaks in the responses) for either angular position or orientation. If there was any such nonsinusoidal component in the responses for some voxels, the phase value would not on its own sufficiently capture that voxel's orientation/angular-position preference in the maps; the classification analysis (see below, Classification) considered the full time series of responses to each stimulus, not just the phase.

We visualized orientation (and angular-position) preference on computationally flattened representations of each subject's occipital lobe. The signal-to-noise ratio of the responses was quantified as the coherence between the time series and the best-fitting sinusoid, separately for each voxel (Engel et al., 1997). For each subject, the orientation maps were thresholded by displaying only voxels exceeding a coherence of 0.3. The coherence threshold for the angular-position maps was higher to reflect the fivefold larger signal amplitude (Fig. 3B). Coherence is the ratio of the Fourier power (squared amplitude) at the signal frequency (S) to the sum of power across all frequencies, including both the signal frequency and nonsignal, [i.e., noise frequencies (N)]:

$$\frac{S}{S+N} > x, \quad (1)$$

where x represents the coherence threshold. If the signal increases by a factor of m and the noise remains constant, the threshold should be increased correspondingly so that the probability of exceeding the threshold is unchanged. To do so, we used Equation 1 to derive a new coherence threshold, y , that accounted for a signal amplitude that was larger by a factor of m . First, rewrite Equation 1 so that signal amplitude appears alone on one side:

$$S > \left(\frac{x}{1-x} \right) N. \quad (2)$$

If the signal amplitude is scaled by a factor of m but the noise is unchanged, we seek a new threshold, y , such that the probability of the signal exceeding the threshold remains the same. Without making any assumptions about the distribution of S , the following equality holds:

$$P\left(S > \left(\frac{x}{1-x} \right) N\right) = P\left(mS > \left(\frac{y}{1-y} \right) N\right), \quad (3)$$

if the new threshold, y , is as follows:

$$y = \frac{mx}{1+x(m-1)}. \quad (4)$$

In our case, the threshold for orientation preference was 0.3 and the signal amplitude for angular position was five times higher than for orientation, so we used a threshold of 0.68 for the angular-position maps, computed using Equation 4 with $x = 0.3$ and $m = 5$.

Map similarity. Circular correlation was used to characterize the similarity between orientation and angular-position maps (Jammalamadaka and Sengupta, 2001). For each voxel, phase values from orientation (o) and angular-position (p) mapping were modeled as being equal up to a phase shift:

$$p = o + \Delta_\theta \text{ modulo } 2\pi. \quad (5)$$

The circular correlation coefficient, r_c , is as follows:

$$r_c = \frac{R(o-p) - R(o+p)}{2\sqrt{\sum_{j=1}^n \sin^2(o_j - \delta) \sum_{j=1}^n \sin^2(p_j - \hat{p})}}, \quad (6)$$

where j indexes voxels, n is the number of voxels, (\wedge) indicates the circular mean of all phase values for the corresponding measurement (either orientation or angular position), and $R(o \pm p)$ gives the concentration of the angular difference (or sum):

$$R(o \pm p) = \left| \sum_{j=1}^n e^{j(o_j \pm p_j)} \right|, \quad (7)$$

where $||$ is the magnitude of a complex number. When two angular variables are related by Equation 5, their difference will be tightly concentrated around Δ_θ , so $R(o-p)$ will be close to n and $R(o+p)$ will be

close to 0. If the relationship is of opposite sign, the reverse will be true. Thus, the relative magnitude of the two terms in the numerator of Equation 6 determines the degree of correlation (positive or negative), and the denominator serves to normalize the coefficient. The value of r_c ranges from -1 to 1 , and the closer it is to 1 , the greater the extent to which Equation 5 explains the relationship between o and p (Jammalamadaka and Sengupta, 2001). From here on, we refer to circular correlation simply as “correlation.” For all experiments, correlations were computed using Equation 6 on a subset of voxels from an independently defined annular region of interest within V1 (see above, Defining retinotopic regions of interest). The phase values computed from the time series data for the two measurements (orientation and angular position) spanned the range 0 to 2π . These values were used to compute correlation in Equation 5, but in Figure 3, these values were scaled to match the actual stimulus orientations (and angular positions) tested, which ranged from 0 to π (or 0 to 180°). A randomization test was used to assess the statistical significance of the correlations. Orientation and angular-position phase values were shuffled (i.e., randomly reassigned to different voxels) to simulate the null hypothesis that the two maps are unrelated, and correlations were recomputed. The process was repeated 10,000 times to obtain a null distribution of correlation values. A p value was then computed as the proportion of samples in the null distribution that were less than the observed value.

An additional analysis was used to characterize the phase shift between the two maps. Specifically, the best-fitting value of the parameter Δ_θ in the model from Equation 5 was found by minimizing the following:

$$\sum_{j=1}^n (\cos(p_j) - \cos(o_j - \Delta_\theta))^2 + (\sin(p_j) - \sin(o_j - \Delta_\theta))^2, \quad (8)$$

where j indexes voxels and n is the total number of voxels. This corresponds to the norm of the difference between two complex numbers and is analogous to least squares but for circular variables. Differentiating Equation 8 and setting it equal to 0 yields the following estimate of Δ_θ :

$$\hat{\Delta}_\theta = \tan^{-1} \left(\frac{\sum_{j=1}^n \sin(o_j - p_j)}{\sum_{j=1}^n \cos(o_j - p_j)} \right). \quad (9)$$

Phase shift was computed using Equation 9 on the same independently defined subset of voxels from the annular region of interest within V1 used in the correlation analysis.

A final analysis was used to characterize the correspondence between orientation and angular-position preference on a voxel-by-voxel basis. This analysis assessed how well the angular-position map explained the response time courses of individual voxels in the main orientation experiment. For each voxel, we randomly divided the orientation runs into “fitting” and “testing” halves. The fitting data were used to compute the amplitude and phase of the best-fitting sinusoid. We then computed the correlation between that sinusoid and the averaged time series from the testing half. This correlation value indicated the reliability of orientation preference across halves of the data within a session. In the second stage of the analysis, we used the fitting data to estimate only the best-fitting amplitude of a sinusoid having phase equal to that voxel’s phase in the independently measured angular-position experiment. We then computed the correlation between this sinusoid and the time series from the test data. This correlation value indicated how well the angular-position preference predicted the responses in the main orientation mapping experiment. We report the median correlation across voxels for the two versions of the analysis. To ensure that these correlations were not artifacts of the fitting, we repeated the fitting procedure using a random phase for each voxel to construct a null distribution. Correlations for both orientation and angular-position preference well exceeded the 95th percentile of this null distribution.

Classification. In multivariate classification analysis of fMRI data, each condition is represented by a set of points in a multidimensional space,

where each point corresponds to a repeated measurement and where the dimensionality is equal to the number of voxels. Accurate decoding is possible when the responses corresponding to different conditions form distinct clusters within this high-dimensional space. We took the average of all 10 cycles in each run as providing measurements of the responses to each of the 16 orientations (or angular positions). Thus, the 16 time points of the cycle-averaged run were the “categories” to be classified. Before averaging, the counterclockwise runs were reversed and time shifted to match the clockwise runs (as described above in Preprocessing). The cycle-averaged time courses were stacked across runs, forming an $m \times n$ matrix, with m being the number of voxels in the V1 region of interest and n being the number of repeated measurements in the session (n ranged from 224 to 288: 1 cycle-averaged time course per run \times 16 orientations \times between 14 and 18 runs per session).

Decoding was performed with a maximum likelihood classifier, using the Matlab function “classify” with the option “diagLinear.” The maximum likelihood classifier optimally separated responses to each of the 16 orientations (or angular positions) if the response variability in each voxel was normally distributed and statistically independent across voxels. Because the number of voxels, m , was large relative to the number of repeated measurements, n , the computed covariance matrix would have been a poor estimate of the true covariance. This would have made the performance of the classifier unstable, as it relied on inversion of this covariance matrix. We therefore ignored covariances between voxels and modeled the responses as being statistically independent across voxels. Although noise in nearby voxels was likely correlated, the independence assumption was conservative; including accurate estimates of the covariances, if available, would have improved the decoding accuracies.

Decoding accuracy was computed using split-halves cross-validation. The $m \times n$ data matrix was randomly partitioned along the n dimension (repeated measurements) into training and testing sets, each containing an equal number of runs. Data in the training and testing sets were drawn from different runs in the same session and were thus statistically independent. The training set was used to estimate the parameters (multivariate means and variances) of the maximum-likelihood classifier. The testing set was then used for decoding. Decoding accuracy was determined as the proportion of the 16 test examples that the classifier was able to correctly assign to one of the 16 orientations (or angular positions). Recall that the 16 different test examples corresponded to the 16 time points of the cycle-averaged run, each of which yielded a particular pattern of voxel responses. The data were repartitioned into testing and training halves, and decoding accuracy was recalculated. The median of many repeated cross-validations provided a stable estimate of decoding accuracy. A nonparametric permutation test was used to determine the significance of decoding accuracy. Specifically, we constructed a distribution of accuracies expected under the null hypothesis that there is no relationship between the presented orientation and the corresponding time point of the cycle-averaged run. To generate this null distribution, each cycle-averaged run in the training and testing data was phase-randomized and accuracy was recalculated. Phase-randomizing preserved the temporal autocorrelation (and power spectrum) of the responses, but still randomized the relationship between time points and orientations. Repeating this randomization 10,000 times yielded a distribution of accuracies expected under the null hypothesis. Accuracies computed using the unrandomized training data were then considered statistically significant when decoding accuracy was higher than the 95th percentile of the null distribution ($p < 0.05$, one-tailed permutation test). We used unaveraged (Fig. 5A) and unfiltered (Fig. 6A,B) data to compute the null distributions.

Sufficiency: angular-position-based averaging. This analysis determined whether the coarse-scale topographic map of orientation was sufficient for orientation decoding. Statistically independent angular-position measurements were used to assign each voxel to one of a fixed set of bins, each corresponding to a different range (bin width)

of angular positions. We then averaged the time series from the orientation experiment of all voxels within each bin to yield a small number of “supervoxels” and repeated orientation classification. This procedure was performed for bins of different widths (ranging from 0.26° to 60°). For comparison, we also did the same amount of averaging while assigning voxels to bins randomly rather than based on the angular-position map.

Necessity: angular-position removal. This analysis determined whether the coarse-scale topographic map was necessary for orientation decoding. The angular-position map was “removed” from the orientation mapping data as follows. First, for each voxel, we computed the phase of the sinusoid that best fit the time series, averaged across all runs of the angular-position mapping experiment. These phase values were used to project the corresponding sinusoids out of the time series measured during orientation mapping. Specifically, let y be the time series measured during orientation mapping and let x be a sinusoid with frequency equal to the stimulus frequency during orientation mapping and phase equal to the best-fitting phase from the angular-position mapping experiment. We computed a residual time series, r , as follows:

$$\vec{r} = \vec{y} - \frac{\vec{y} \cdot \vec{x}}{\vec{x} \cdot \vec{x}} \vec{x}. \quad (10)$$

Removal by projection ensured that the residual time series, r , was orthogonal to the removed component, x . Finally, classification analysis was performed on the residual time series (as described above in Classification).

A confidence interval was determined for the residual decoding accuracy after map removal, according to the hypothesis that the residual decoding depended only on imperfections (e.g., due to measurement noise and registration across experiments) in our ability to completely remove the map. Two subjects repeated the double-wedge, angular-position mapping experiment. For these subjects, the first angular-position mapping experiment was used to remove the map from the second angular-position mapping experiment. Classification analysis was performed on data from the second angular-position mapping experiment before and after removing the map. The classification of angular position was identical to the classification of orientation, except that the 16 time points corresponded to different angular positions rather than different orientations. Signal-to-noise ratio was much higher for angular-position than for orientation measurements. Gaussian noise was added to the second angular-position dataset (from which the map was removed) so that baseline decoding accuracy (before map removal) was the same as for orientation mapping. Specifically, noise was added to the cycle-averaged time series of each run, the amplitude of which was slightly different for each of the two subjects (SD of 6 for one subject, 6.5 for the other). Resampling was used to obtain a confidence interval for decoding accuracy. On each resampling, a new sample of noise was added to the data, and runs were randomly repartitioned into training and testing halves. Repeating this 10,000 times yielded a distribution of decoding accuracy. The median and 16th and 84th percentiles of this distribution are reported (Fig. 5B, gray horizontal lines) (percentiles were averaged across the two subjects). Orientation decoding accuracy (after map removal) fell within this confidence interval. The noise amplitude was set to match baseline decoding accuracy between the two data sets before map removal, but our conclusions depend on the fact that the decoding accuracies were indistinguishable after map removal.

The projection analysis could degrade decoding accuracy even if the phase to be removed was randomly chosen. This could occur if, by chance for some voxels, the phase to be removed matched the voxel’s orientation preference. An additional test was performed to ensure that the degradation in classification was specifically due to removing the map. Specifically, the orientation and angular-position phase values were shuffled (i.e., randomly reassigned to different voxels) to project out sinusoids with

random phase. Pairings were reshuffled 10,000 times to yield a distribution of decoding accuracies, and the median is reported (Fig. 5B, random component removed).

Spatial filtering. fMRI time series data were spatially filtered using a volumetric filtering procedure. The procedure was similar to the approach used in related studies (Op de Beeck, 2010; Swisher et al., 2010). After motion correction, each frame of the time series was filtered separately in each of its three dimensions. A one-dimensional filter was constructed by multiplying the inverse Fourier transform of an ideal low-pass filter by a Hamming window in the spatial domain. In the Fourier domain, this filter that has a nearly flat passband followed by a transition region (centered on the specified frequency cutoff) and a stopband with minimal ripple. The filter achieves better frequency isolation than a Gaussian filter while minimizing artifacts due to ringing. For consistency with previous reports (Swisher et al., 2010), we report the size of each low-pass filter as the full-width at half-maximum (in the spatial domain) of a Gaussian that achieves half power at the same frequency cutoff (in the Fourier domain). Filtering was performed on a fixed rectangular region of interest that was cropped from the full volume to avoid artifacts associated with extracting the data from ROIs after filtering (Kamitani and Sawahata, 2010). Filtering was performed in the spatial domain, separately along each dimension, using reflective boundary handling. For each filter size, low-pass filtering was performed by filtering the data with the corresponding filter, and high-pass filtering was performed by subtracting the results of low-pass filtering from the unfiltered data. Many different versions of the analysis (e.g., sharp-cutoff “ideal” filter, nonrectangular region of interest) yielded a qualitatively similar result: classification monotonically decreased with low-pass filtering and increased with high-pass filtering, for both orientation and angular position. Like Swisher et al. (2010), we also found qualitatively similar results when performing filtering on the surface (instead of in the volume).

The response amplitudes in the angular-position mapping sessions were five times larger than the response amplitudes in the orientation mapping sessions (Fig. 3B,C). The higher signal-to-noise ratio in the angular-position measurements resulted in near-perfect decoding accuracy with any kind of spatial filtering, which we interpreted as a ceiling effect. To avoid the ceiling effect, only a single cycle of angular-position data was included in the analysis, rather than the average of all 10 cycles. In each iteration of the cross-validation procedure, a single cycle of the angular-position data was randomly selected for each run. This approximately compensated for the difference in the signal-to-noise ratio between the two datasets and avoided a ceiling effect without directly modifying the time series data.

Two different procedures were adopted to reduce the signal-to-noise of angular-position measurements. In the projection analysis (see above, Necessity: angular-position removal), the goal was to precisely equate baseline classification accuracy, and this was best accomplished by adding noise to the time series. In the spatial filtering analysis, we simply wanted to avoid a ceiling effect; thus, we used the simplest procedure that modified the data the least (i.e., using only a single cycle of the time series per run).

Results

Topographic map of orientation in human V1

We studied the representation of orientation using periodic, phase-encoded stimuli similar to those used in studies of retinotopic organization. In typical retinotopic mapping studies, a high-contrast visual stimulus moves periodically through the visual field (e.g., to measure the angular-position component of the retinotopic map, the stimulus rotates around the fixation point), creating a traveling wave of activity in cortex (Engel et al., 1994; Sereno et al., 1995; DeYoe et al., 1996). Following similar reasoning, to “map” orientation, we used a large oriented sinusoidal grating (0.5 cycles per degree), the orientation of which cycled through 16 evenly spaced angles between 0° and 180°, 1.5 s per orientation (Fig. 1A). The grating filled a peripheral annulus with a smoothed edge (inner radius, 4.5°; outer radius, 9.5°). The spa-

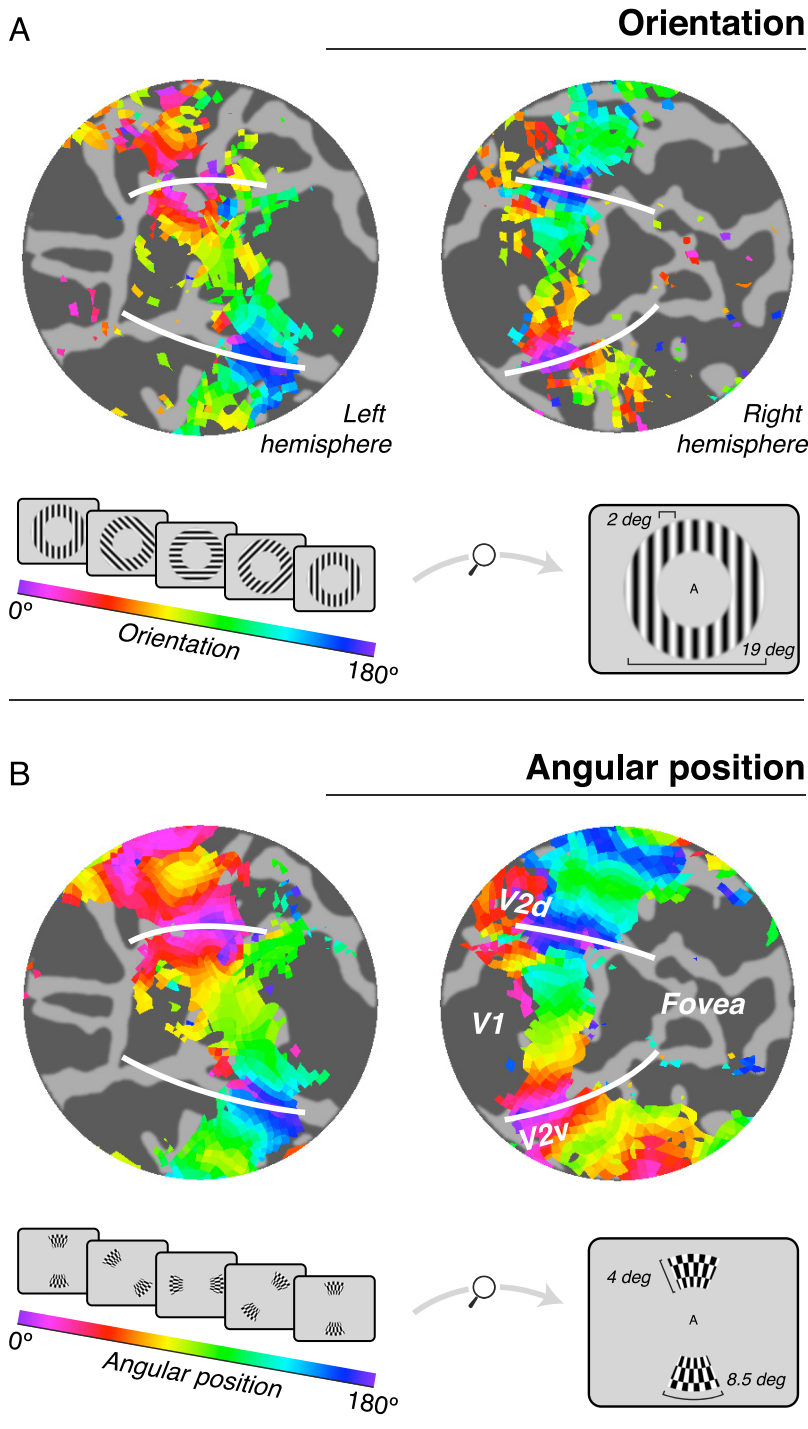


Figure 1. Orientation and angular-position topographic maps for a single subject shown on a flattened representation of the occipital lobe. **A**, Responses to phase-encoded oriented gratings (shown in inset). The stimulus cycled through 16 steps of orientation, ranging from 0° to 180° every 24 s. The map is thresholded at a coherence of 0.3. **B**, Responses to double-wedge retinotopy stimulus (shown in inset). The stimulus cycled through 16 steps of angular position, ranging from 0° to 180° every 24 s. The map is thresholded at a coherence of 0.68 to account for differences in signal-to-noise ratio between the two experiments (see Materials and Methods). Color indicates phase of best-fitting sinusoid; white lines indicate the V1/V2 boundaries.

tial phase of the grating was randomized every 250 ms. Subjects performed a demanding rapid serial visual presentation two-back task at fixation to control the locus of attention and divert it from the peripheral stimulus. The responses of each voxel were fit to a sinusoid with the period of the stimulus as is typically done with retinotopic mapping data. The phase of the best-fitting sinu-

oid indicated the preferred orientation of the voxel. The coherence between the best-fitting sinusoid and the measured response time courses provided a measure of the signal-to-noise ratio of the responses. Orientation preferences were visualized on a flattened representation of visual cortex (Figs. 1, 2; see also Fig. 4).

Orientation preferences exhibited a smooth progression across V1. For each location within V1, the responses exhibited a preference for radial orientations, matching the angular-position component of the retinotopic map (Figs. 1, 2). The angular-position component of the retinotopic map was measured in the same subjects using standard methods (Engel et al., 1997; Larsson and Heeger, 2006; Wandell et al., 2007). Stimuli were high-contrast flickering checkerboard patterns presented within two collinear wedges shown within the same annulus used for the oriented gratings (Fig. 1B). We used two wedges (instead of one, which is more typical for retinotopic mapping) to equate the phase-range of the maps; in both cases, one full cycle covered 180° of the stimulus variable (angular position or orientation). The angular-position and orientation maps were qualitatively similar (Figs. 1, 2). We use the term “orientation map” to refer to the fact that orientation responses are topographically organized in a radial bias, in register with the angular-position map. The similarity between the orientation map and the angular-position map was quantified by correlating the orientation preference and angular-position preference of each voxel (Fig. 3A) (most voxels fall on or near the line of equality). Correlation between the maps (r_c , circular correlation) was highly statistically significant (Fig. 3A) ($r_c = 0.52$, $p < 0.0001$; randomization test, combining data across $n = 3$ subjects). Correlations were also significant, and of similar magnitude, within each individual subject (S1: $r_c = 0.53$, $p < 0.0001$; S2: $r_c = 0.40$, $p < 0.0001$; S3: $r_c = 0.65$, $p < 0.0001$; randomization test). Voxels for which the maps differed substantially (farther from the line of equality in Fig. 3A) appeared to cover a wide range of retinotopic locations and tended to have low coherence values. Computing correlations only using high-coherence voxels (those exceeding the 50th percentile) yielded an even higher correlation ($r_c = 0.78$; $n = 3$), but we prefer to report correlations without thresholding to avoid any potential statistical bias (Kriegeskorte et al., 2009; Vul et al., 2009). Circular correlations are invariant to phase offsets, so we also computed the best-fitting phase shift (Δ_θ) between orientation and angular-position preference. The phase

soid indicated the preferred orientation of the voxel. The coherence between the best-fitting sinusoid and the measured response time courses provided a measure of the signal-to-noise ratio of the responses. Orientation preferences were visualized on a flattened representation of visual cortex (Figs. 1, 2; see also Fig. 4).

Orientation preferences exhibited a smooth progression across V1. For each location within V1, the responses exhibited a preference for radial orientations, matching the angular-position component of the retinotopic map (Figs. 1, 2). The angular-position component of the retinotopic map was measured in the same subjects using standard methods (Engel et al., 1997; Larsson and Heeger, 2006; Wandell et al., 2007). Stimuli were high-contrast flickering checkerboard patterns presented within two collinear wedges shown within the same annulus used for the oriented gratings (Fig. 1B). We used two wedges (instead of one, which is more typical for retinotopic mapping) to equate the phase-range of the maps; in both cases, one full cycle covered 180° of the stimulus variable (angular position or orientation). The angular-position and orientation maps were qualitatively similar (Figs. 1, 2). We use the term “orientation map” to refer to the fact that orientation responses are topographically organized in a radial bias, in register with the angular-position map. The similarity between the orientation map and the angular-position map was quantified by correlating the orientation preference and angular-position preference of each voxel (Fig. 3A) (most voxels fall on or near the line of equality). Correlation between the maps (r_c , circular correlation) was highly statistically significant (Fig. 3A) ($r_c = 0.52$, $p < 0.0001$; randomization test, combining data across $n = 3$ subjects). Correlations were also significant, and of similar magnitude, within each individual subject (S1: $r_c = 0.53$, $p < 0.0001$; S2: $r_c = 0.40$, $p < 0.0001$; S3: $r_c = 0.65$, $p < 0.0001$; randomization test). Voxels for which the maps differed substantially (farther from the line of equality in Fig. 3A) appeared to cover a wide range of retinotopic locations and tended to have low coherence values. Computing correlations only using high-coherence voxels (those exceeding the

shifts were near 0° and deviations from 0° were not systematic across subjects (S1: $\Delta_\theta = 2^\circ$; S2: $\Delta_\theta = -8^\circ$; S3: $\Delta_\theta = 9^\circ$), indicating that the two maps were in register with one another. An additional visualization was used to summarize the structure of the map (Fig. 3C), which shows a strong and consistent bias toward radial orientations across all angular positions within the peripheral annulus (Fig. 3C). As a complementary assessment of the degree to which the radial bias explained the orientation preferences, we computed, for each voxel, the correlation between the orientation responses and two different sinusoids: one with phase estimated from orientation mapping and the other with phase estimated from angular-position mapping. The median value of this correlation across voxels was 0.31 ± 0.035 (for orientation) and 0.33 ± 0.037 (for angular position) (mean \pm SE, $n = 3$ subjects), confirming that a substantial fraction of the orientation responses is explained by angular position.

The orientation map cannot be explained by eye movements. One potential eye-movement-related confound could occur if subjects made involuntary eye movements that were correlated with the orientation of the stimulus. This pattern would shift the retinotopic position of the stimulus in the direction of the eye movements. For example, moving the eyes upward for a vertical grating would shift the stimulus down, but not to the left or right. Shifts in the retinotopic position of the stimulus would co-occur with changes in orientation, which could produce a map that appeared to show a radial bias. To test this possibility, we measured gaze position for a single subject during orientation mapping (supplemental Fig. 1, available at www.jneurosci.org as supplemental material). Nearly all fixations (99%) fell within 1.5° of the center of the screen, indicating the retinotopic position of the stimulus was nearly constant for the duration of the experiment. Furthermore, there was no evidence of correlation between fixation location and stimulus orientation ($r_c = -0.027$, $p = 0.1220$).

The orientation map was robust to a variety of stimulus manipulations. First, we wondered whether the map was driven by responses to the inner and outer edges of the annulus. We repeated the experiment using an annular grating with a much blurrier edge. The spatial extent of blurring (5°) was more than two times the period of the sinusoidal grating (2°). We observed a qualitatively similar map (Fig. 4A), with significant correlations between angular-position and orientation maps (Fig. 4A) ($r_c = 0.36$, $p < 0.0001$, $\Delta_\theta = 0^\circ$; $n = 3$). Correlation magnitude was lower than in the main experiment, but so was the signal-to-noise ratio, likely due to lower overall stimulus contrast. Second, the experiment was repeated with a sinusoidal grating that included the central part of the visual field (outer radius, 9.5° ; inner radius, 0.5°), to determine whether the map remained for larger stimuli and whether the map might vary as a function of eccentricity. In the periphery, the map obtained using this stimulus was qualitatively similar (Fig. 4B) and correlated with the angular-position map (Fig. 4B) ($r_c = 0.54$, $p < 0.0001$; $\Delta_\theta = 0^\circ$; $n = 3$, computed within a peripheral ROI) (see Materials and Methods). Qualita-

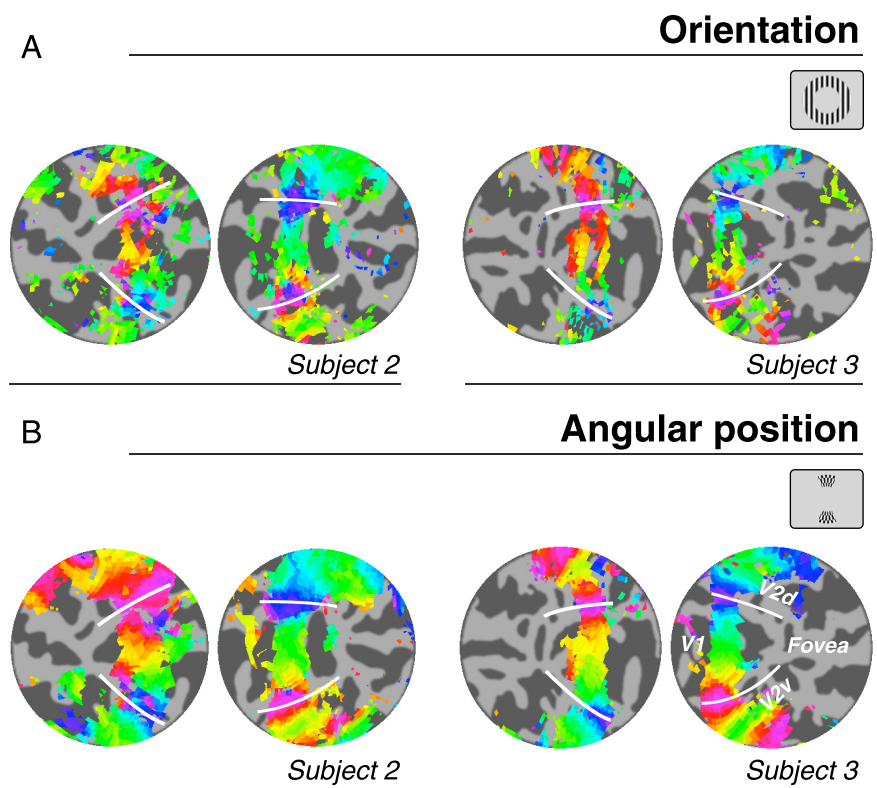


Figure 2. Orientation and angular-position maps for two additional subjects. **A**, Orientation maps for subjects 2 and 3. **B**, Angular-position maps for subjects 2 and 3. The same conventions are used as in Figure 1.

tively, we did not observe as clear a map near the fovea. Finally, we measured responses to higher-spatial-frequency (1.4 cycles per degree) phase-randomized, square-wave-oriented gratings matched to that used in other studies (Kamitani and Tong, 2005; Swisher et al., 2010) and obtained similar results in the periphery (Fig. 4C) ($r_c = 0.28$, $p < 0.0001$; $\Delta_\theta = -4^\circ$; $n = 2$). The correlation magnitudes for the square-wave gratings were lower than for other stimuli, possibly because the high-spatial-frequency components of the square waves were poorly matched to peripheral receptive fields. For the square-wave stimulus, which included the central visual field, the map was again not evident near the fovea.

The response amplitudes were five times smaller during orientation than angular-position mapping (Fig. 3B). We suspect that we were able to uncover a robust orientation map despite this smaller signal amplitude because we used a phase-encoded stimulus protocol that is optimal for topographic mapping. For comparison, we measured orientation preferences in one of our subjects using an event-related protocol with similar stimuli (see Materials and Methods). The correlation between angular-position map and the orientation map measured with the event-related protocol was significant (S2: $r_c = 0.20$, $p < 0.0001$), demonstrating that the orientation map is robust to the different stimulus protocols. As expected, the structure of the orientation map was qualitatively less clear in the event-related experiment, confirming the relative advantage of the phase-encoded protocol.

The coarse-scale orientation map is sufficient and necessary for orientation decoding

The observed coarse-scale orientation map was sufficient for multivariate decoding of orientation. If the map is sufficient for decoding of orientation, then it should be possible to group voxels with similar angular-position preferences and

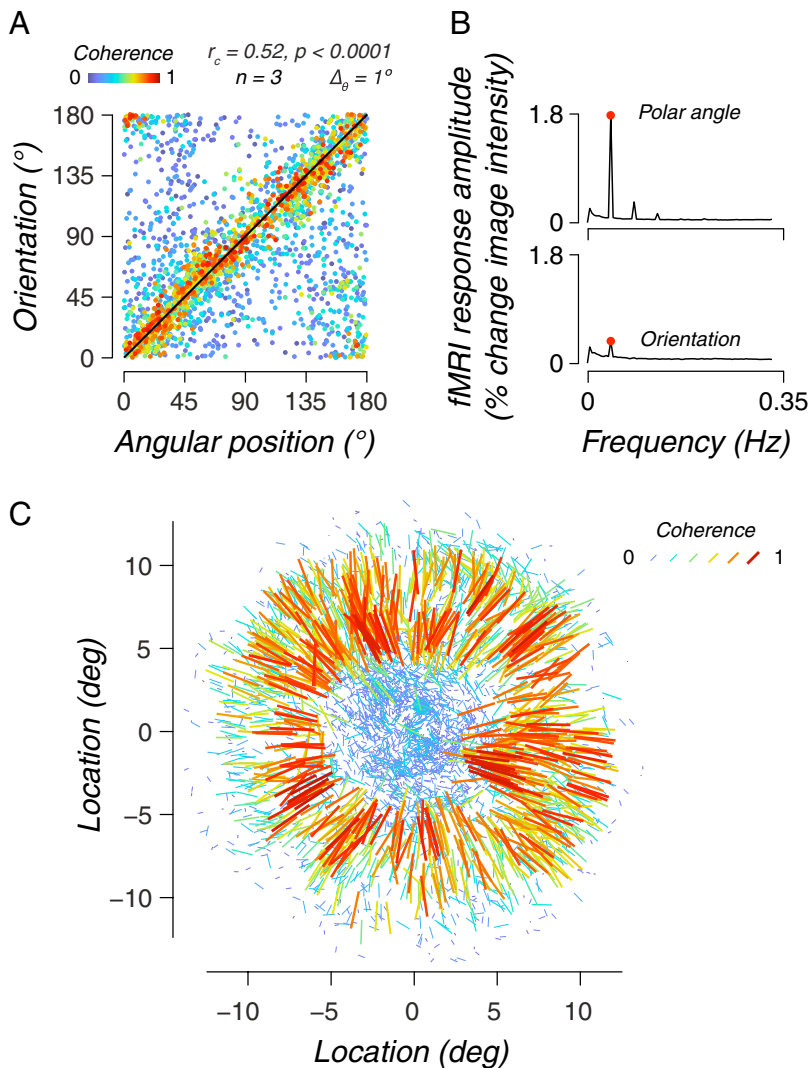


Figure 3. Orientation map matches angular-position map. **A**, Circular correlation between orientation and angular-position maps. Preferred orientation is plotted against preferred angular position. Each dot corresponds to a voxel from an annular region of interest in V1 defined using data from an independent eccentricity mapping experiment. Data were combined across subjects ($n = 3$). Dot color indicates coherence from the orientation mapping experiment. **B**, Response amplitude (percentage change in image intensity), averaged across voxels and across subjects ($n = 3$), as a function of temporal frequency for angular-position (top) and orientation mapping (bottom). Red dot indicates stimulus frequency. **C**, Line vector plot showing radial bias of orientation preference in the visual field. Each line corresponds to a voxel. The center of each line is the voxel's retinotopic position (angular position and eccentricity), and the angle of each line indicates the voxel's orientation preference. Line color and size indicate coherence from the orientation mapping experiment. Data were combined across $n = 3$ subjects.

average their responses (within each group) and still decode orientation based on these averaged signals. To test this prediction, the angular-position measurements were used to assign each voxel to one of a fixed set of angular-position bins, each corresponding to a different range of angular positions, specified by the width of each bin (in degrees). The time series of all voxels within each bin were averaged to yield a small number of “supervoxels,” and orientation decoding was performed on these averaged supervoxel responses. Specifically, we treated the 16 steps of orientation as 16 different stimulus categories and estimated the responses to each of the 16 categories. We then split the data in half to train and test a linear classifier. The classifier used the multivariate pattern of supervoxel responses in the training data to predict the stimulus orientation from the test data (see Materials and Methods). For comparison, we did the same amount of averaging while

assigning voxels to random phase bins that were not based on the angular-position map. For both random and angular-position-based averaging, the number of the angular-position bins was varied. For a large number of bins, there was very little averaging, so decoding accuracy was high for both random and angular-position-based averaging. As the number of bins decreased, decoding accuracy for random averaging degraded to chance (Fig. 5A); however, for angular-position averaging, decoding accuracy remained well above chance, even when using a very small number of angular-position bins, each averaging over a wide range of angular positions (Fig. 5A). Thus classification accuracy remained high despite substantial averaging, so long as the averaging was based on the structure of the coarse-scale map.

A complementary analysis was performed to determine whether the coarse-scale orientation map was necessary for orientation decoding. Orientation decoding may rely on multiple factors at multiple spatial scales. First, decoding may depend on the systematic relationship between orientation preference and the angular-position map. Indeed, the sufficiency analysis showed that this relationship alone was capable of supporting accurate decoding. However, the responses of each voxel may additionally depend on other factors unrelated to the coarse-scale map (e.g., voxel-to-voxel variability in orientation preference arising from irregular sampling of columns). We hypothesize that the coarse-scale (radial bias) map is necessary for decoding, which implies a negligible contribution from the other factors. To test this hypothesis, for each voxel, the angular-position data were used to remove the map-based component of the orientation data. Specifically, for each voxel in the orientation mapping experiment, we removed (by regression, i.e., linear projection) a sinusoid

having phase equal to the angular-position preference of that voxel. Thus, any signal component at the stimulus frequency with the angular-position phase was completely removed from the orientation data. A classification analysis was performed before and after angular-position removal. Even if the map was necessary for decoding, this angular-position removal procedure was unlikely to reduce decoding accuracy to chance. Any mismatch in the maps between the two experiments (e.g., due to measurement noise) would have prevented complete removal of the map-based component from one experiment based on data from other. Small mismatches in the map likely occurred because estimates of preference (for angular position or orientation) in both experiments were corrupted by noise and by imperfect coregistration across the two experiments (which were on different days). A confidence interval was determined for the residual

decoding accuracy after angular-position removal, according to the hypothesis that any residual decoding depended only on these imperfections in our ability to completely remove the angular-position map. Specifically, two subjects participated in an additional session of angular-position mapping, and a similar procedure was used to remove the angular-position map measured in the first session from that measured in the second session (see Materials and Methods). When removing angular position from angular position, any residual decoding accuracy after map removal reflected only imprecision in matching phase estimates across the two experiments.

The coarse-scale orientation map was necessary for orientation decoding. Baseline decoding accuracy for orientation before map removal was 54% (average of $n = 3$ observers) (Fig. 5B). Removing the angular-position map from the orientation data degraded decoding by more than a factor of 2, down to 20%. Orientation decoding accuracy after map removal was within the confidence interval for the expected degradation in decoding accuracy (Fig. 5B) (see Materials and Methods). An additional test was performed to ensure that the degradation in decoding was due to removing the map, rather than a mere consequence of the projection procedure. Voxel pairings were shuffled to project out sinusoids with random phase, and residual decoding accuracy remained high (43%). The projection procedure removed only the smooth, single-peaked component of each voxel's responses (because it removed a sinusoid). If there was a reliable nonsinusoidal (e.g., multi-peaked) component of the responses, not predicted by angular-position preference, it would have remained after the projection. Thus, the necessity of the coarse-scale map implies that any such components of the responses were negligible.

Previous studies have used spatial filtering to determine the scale of orientation information in fMRI data (Op de Beeck, 2010; Swisher et al., 2010). In particular, Swisher et al. (2010) found that orientation decoding degraded with low-pass filtering (i.e., smoothing) and concluded that orientation decoding reflects irregular spatial arrangements of orientation columns and their supporting vasculature. In a complementary analysis, they showed that high-pass filtering, which removes low-spatial-frequency information, still allowed for decoding well above chance, and concluded that decoding does not require a coarse-scale signal like a radial bias. The logic that links both results to the source of information—either columns or a radial bias—crucially assumes that signals arising from a radial bias are exclusively in low spatial frequencies, such that high-pass filtering removes them and low-pass filtering leaves only them. This reasoning predicts that the decoding of a topographically organized variable with no hypothesized contribution from columnar irregularities, such as angular-position retinotopy, should not degrade with low-pass filtering or remain after high-pass filtering.

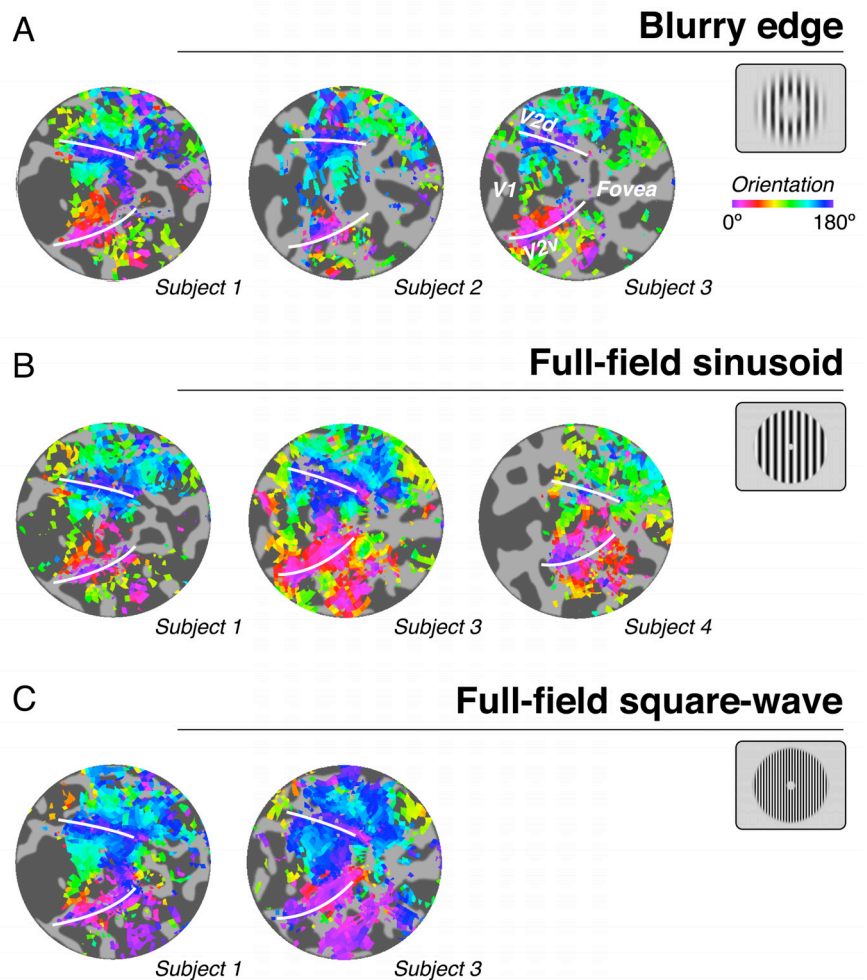


Figure 4. Orientation maps are robust to a variety of stimulus manipulations. Conventions and threshold are the same as in Figures 1A and 2A. **A**, Orientation maps using a sinusoidal grating with a much blurrier edge. **B**, Orientation maps using a sinusoidal grating that included the central visual field. **C**, Orientation maps using a high-frequency square-wave grating that also included the center of the visual field.

We measured decoding accuracy with different degrees of spatial filtering (both low-pass and high-pass), for orientation, and also for angular position. Decoding accuracy for orientation increased slightly with a small amount of low-pass filtering (presumably because local averaging increased the signal-to-noise ratio of the measurements), but decoding accuracy decreased as the low-pass filter attenuated a broader range of frequencies (Fig. 6A), replicating previous reports (Swisher et al., 2010). Decoding was well above chance after high-pass filtering (Fig. 6A), also replicating previous reports (Swisher et al., 2010). However, we found nearly identical results for decoding angular position (Fig. 6B). Of course, decoding angular position (i.e., retinotopic location) does not depend on the sampling of spatial irregularities in the fine-scale columnar architecture. Obtaining such similar results for orientation and angular position suggests that the spatial filtering analysis does not distinguish well between biases derived from maps or columns.

Rather, both sets of results are consistent with the fact that when responses are organized in large-scale maps (orientation or angular-position), there are signals at multiple spatial frequencies that carry information about the stimulus, even without any contribution from the columnar architecture. The orientation (or angular-position) map is approximately a triangle wave (a ramp within V1 with reversals at the V1/V2 boundary) and hence

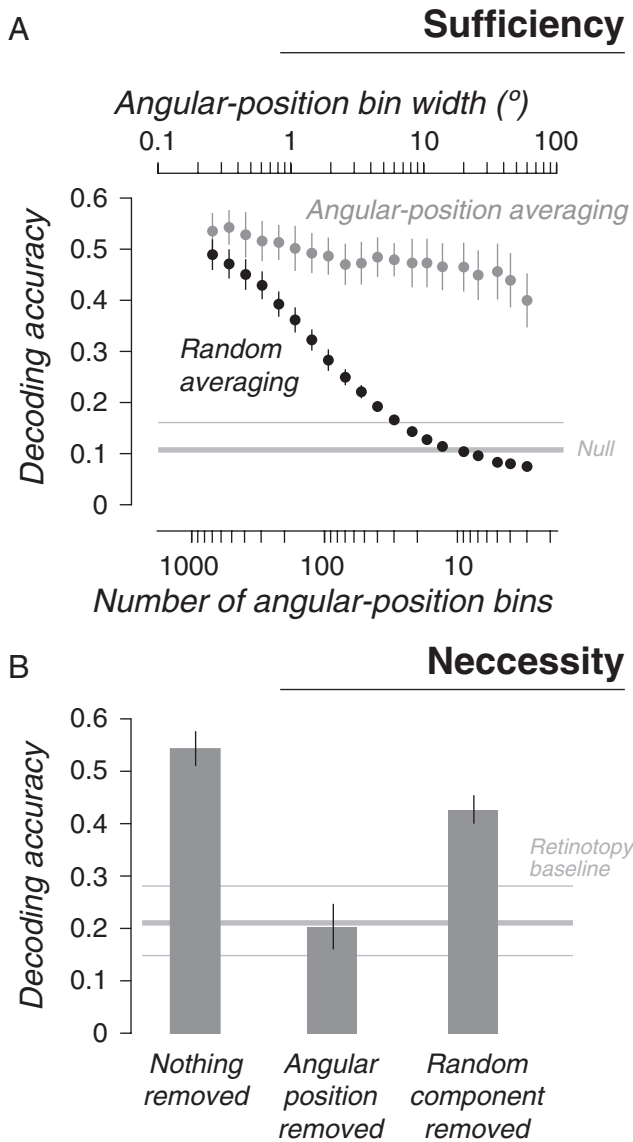


Figure 5. Orientation map necessary and sufficient for classification. **A**, Sufficiency. Gray symbols, The angular-position measurements were used to assign each voxel to one of a fixed number of angular-position bins, each corresponding to a different range of angular positions (“angular-position bin width”). Time series of all voxels within each bin were averaged, and classification was performed on the resulting averaged time series. Black symbols, Voxels were assigned to bins randomly, not based on the angular-position map. Error bars indicate SEM across subjects ($n = 3$). Thick and thin gray lines, The median and fifth percentile of a one-sided null distribution (randomization test) (see Materials and Methods). **B**, Necessity. Removing the coarse-scale map degraded classification of orientation. Left bar, Baseline orientation decoding accuracy without removing any component from the map. Middle bar, Decoding accuracy after projecting out, from each voxel’s response, a sinusoid having phase equal to the angular position preference of that voxel as measured in the angular-position mapping experiment. Right bar, Decoding accuracy after removing a sinusoid with random phase. Error bars indicate SEM across subjects ($n = 3$). Gray horizontal lines, The median (thick line) and 68% confidence interval (thin lines) of a null distribution for the residual decoding accuracy expected after map removal if decoding was driven entirely by the map. Specifically, this distribution was computed by removing the angular-position phases from a separate angular-position mapping experiment before measuring decoding accuracy for angular position (see Materials and Methods).

is broadband, containing both low and high spatial frequencies. The details and irregularities of its projection onto the cortex add additional high-frequency components that carry information. Analyses, like ours, that either explicitly average the data in accordance with the map (“sufficiency”) or remove the map (“neces-

sity”) enable a more precise link between the radial bias and classification. Conceptually, the angular-position-based averaging analysis described above (“sufficiency”) is a form of low-pass filtering, because data from multiple voxels are combined into a single signal. However, unlike spatial low-pass filtering, angular-position-based averaging was constrained by the angular-position map and did not necessarily combine signals locally or uniformly in space (whether performed on the surface or in the volume). Our results for spatial filtering of orientation responses agree with previous work, but our results for spatial filtering of angular-position responses compel a different interpretation; the effect of filtering on orientation decoding is consistent with the conjecture that orientation decoding is primarily dependent on a coarse-scale map that follows the same topographic organization as the angular-position map.

Discussion

We observed a coarse-scale topographic map of orientation preference in human V1. The map was tightly colocalized with the well known angular-position component of the retinotopic map. The orientation map was robust to a variety of stimulus parameters and was not due to either attention or eye movements. Our mapping results are consistent with fMRI studies in both human and macaque (Sasaki et al., 2006; Mannion et al., 2010) and electrophysiological recordings in cat V1 (Leventhal, 1983; Schall et al., 1986). Most population-level physiological measurements in primary visual cortex (e.g., those using calcium imaging) have not been at a spatial scale sufficient to measure this coarse-scale radial orientation bias.

The map of radial orientation bias in human V1 may be partly inherited from the properties of the early visual system. Dendritic fields of peripheral retinal ganglion cells are elongated toward the fovea in macaque (Schall et al., 1986) and human (Rodieck et al., 1985). There are radial biases in the orientation preference of cat retinal ganglion cells, especially in the periphery (Levick and Thibos, 1980, 1982), and radial orientation biases have been identified in the lateral geniculate nucleus (LGN) of both cat (Shou et al., 1986) and macaque (Smith et al., 1990). However, the exact contribution of these early radial biases to the organization of orientation tuning in V1 depends on the mechanisms by which orientation tuning in V1 is constructed (Schall et al., 1986; Ringach, 2007). The topographic map of radial orientation bias in human V1 may thus depend on a variety of factors, including the structure and organization of feedforward inputs from LGN to V1 and intracortical connectivity within V1.

In both macaque and cat, the radial bias (in either the retina, LGN, or V1) has been reported as predominantly in the periphery, which is consistent with our finding in humans that the radial orientation map was especially pronounced in the periphery (Fig. 4). However, our failure to observe a robust radial orientation map closer to the fovea may only reflect a methodological limitation. Signal-to-noise ratio is typically lower closer to the fovea, and even retinotopic mapping with fMRI is difficult in this part of visual cortex (Schira et al., 2009). Thus, we do not claim that there is no map near the fovea, only that we failed to observe one. In some hemispheres, there was a weak tendency at mid-eccentricities for vertical and near-vertical orientations to dominate. This may be linked to the well known “oblique effect,” in which observers are better able to detect and discriminate cardinal orientations. Psychophysical measurements show that contrast sensitivity and orientation discrimination for gratings exhibit cardinal biases in the central part of the visual field and radial biases in the periphery (Berkley et al., 1975; Rovamo et al., 1982; Westheimer, 2003), consistent with

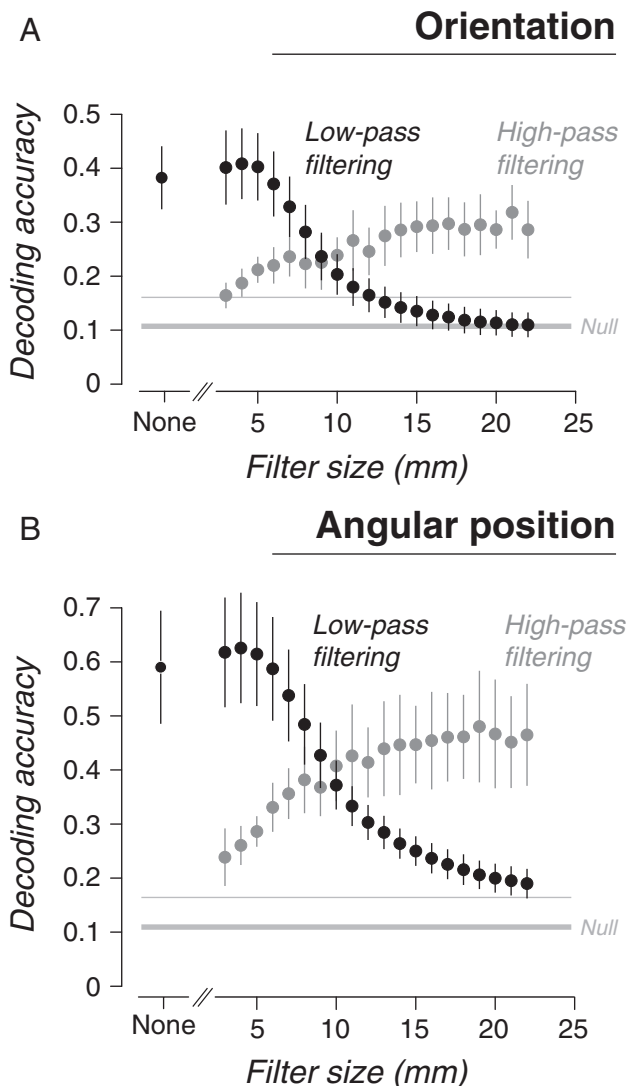


Figure 6. Decoding accuracy for both orientation and angular position decreases with low-pass filtering (i.e., spatial smoothing) and increases with high-pass filtering. *A*, Decoding accuracy for orientation using unfiltered data (None) and for several different spatial filtering kernel sizes (see Materials and Methods). Black symbols, Decoding low-pass filtered data. Gray symbols, Decoding high-pass filtered data. Error bars indicate SEM across subjects ($n = 3$). Thick and thin gray lines, The median and fifth percentile of a one-sided null distribution (randomization test) (see Materials and Methods). *B*, Decoding accuracy for angular position. Conventions are the same as in *A*.

the peripheral radial bias in our measurements. These biases in the representation of orientation, neural and psychophysical, foveal and peripheral, may reflect an efficient representation of image statistics (Ganguli and Simoncelli, 2011), which show a cardinal bias at image locations that are likely to be fixated and radial biases in the periphery (Rothkopf et al., 2009).

Several previous studies have reported orientation biases using both fMRI and electrophysiology, with ostensibly conflicting results (Mansfield, 1974; Wilson and Sherman, 1976; De Valois et al., 1982; Furmanski and Engel, 2000; Li et al., 2003; Sasaki et al., 2006; Mannion et al., 2010; Swisher et al., 2010). Some studies found predominantly cardinal biases (related to the oblique effect), whereas others found radial biases instead of, or in addition to, the cardinal biases. In these studies, however, responses were typically averaged over an entire visual area or the portion of a visual area corresponding to an entire quadrant of the visual field

(fMRI) or were localized exclusively to one sector of the visual field (electrophysiology). The discrepancies between these results could, therefore, be reconciled if the orientation biases depended systematically on eccentricity.

The radially oriented component of the retinotopic mapping stimulus that we used may have contributed to the measurements of the angular-position maps. The checkerboard stimulus combined radial motion, radial orientations, and high contrast. However, we have measured angular-position maps in these same subjects using a population receptive field method with moving horizontal and vertical bars showing natural image stimuli that were not radially orientated (Dumoulin and Wandell, 2008). These methods yielded nearly identical angular-position maps, confirming that the angular-position maps were not driven by the radial orientations of the checkerboard stimulus.

Our results provide a parsimonious explanation for the accuracy of orientation decoding from fMRI measurements. Voxels show a small but reliable bias for radial orientations organized in a topographic map, and multivariate methods exploit these biases to decode orientation. Previous studies have used spatial filtering to argue that above-chance decoding arises from sampling of the underlying irregular columnar architecture, rather than from coarse-scale biases (Swisher et al., 2010; but see Op de Beeck, 2010). We found, however, that the effects of spatial filtering on decoding were similar for orientation and angular-position retinotopy, the latter clearly not reflecting the columnar architecture, demonstrating that spatial filtering does not distinguish well between maps and columns. Our analyses showed that averaging the responses based on the angular-position map did not prevent accurate orientation decoding, proving that the coarse-scale orientation map was sufficient for decoding. Also decoding accuracy was degraded by removing the angular-position map from the responses to different orientations, proving that the coarse-scale orientation map was necessary for orientation decoding.

This explanation helps elucidate the surprising fact that orientation decoding can be performed across experiments performed on different days (Kamitani and Tong, 2005; Gardner, 2010). This finding alone makes the columnar hypothesis unlikely but leaves open the question of why decoding works. Our results provide a mechanistic explanation: Responses based on a topographic map would be nearly consistent across different sessions, but small biases due to sampling spatial irregularities in the fine-scale columnar architecture would not. In fact, one of our subjects repeated the orientation mapping experiment in a second scanning session, and orientation preference was highly consistent across sessions ($r_c = 0.69$, $p < 0.0001$; $\Delta_\theta = -3^\circ$).

Several studies of orientation decoding in humans successfully classified orientation, but did not report maps like those reported here (Haynes and Rees, 2005; Kamitani and Tong, 2005; Serences et al., 2009). This may seem surprising given the robustness of our maps, but the discrepancy likely reflects the substantial differences in experimental protocol. The periodic stimulation protocol that we adopted for orientation mapping was based on an approach that has been optimized for topographic mapping (Kalatsky and Stryker, 2003; Wandell et al., 2007). Specifically, the power in the response time course is concentrated at a single temporal frequency, selected to obtain a reasonable trade-off between the signal attenuation at high frequencies due to the sluggishness of the hemodynamics (Boynton et al., 1996) and the noise and drift that dominate fMRI signals at low frequencies (Smith et al., 1999). These aspects of the periodic design increase signal-to-noise ratio compared with event-related designs; we found weaker, but statistically significant, evidence for the

radial orientation map using an event-related experiment. The studies cited above were able to classify orientation, but perhaps they failed to see a map because the signal-to-noise ratio required to observe a clear topographic map is higher than that required for accurate decoding. Multivariate classification methods are extraordinarily sensitive (Kriegeskorte et al., 2006) and could exploit radial biases even in data where the map itself is not clear.

The presence of an orientation map, and its necessity and sufficiency for orientation decoding, casts doubt on the conjecture that fMRI classification techniques, when applied to orientation-specific responses in V1, primarily reflect differential signals that arise from sampling spatial irregularities in the fine-scale columnar architecture. Clearly, the mere finding of above-chance decoding of fMRI signals in any given brain region does not indicate the existence of fine-scale columnar architecture for neurons representing the corresponding stimulus, task, or cognitive state. However, orientation decoding in V1 has become a well established example of how fMRI, at conventional spatial sampling resolutions ($\geq 2 \times 2 \times 2$ mm voxels), can ostensibly exploit irregular columnar structure and probe subvoxel representations. This columnar conjecture appears to be false in one case where it was thought to be true. Thus, we caution the extension of this widely held conjecture to other sensory and cognitive representations in other brain areas. At the same time, the finding of a coarse-scale orientation map exemplifies how fMRI is well suited to reveal topographic structure in the human brain, particularly when using experimental protocols optimized for the temporal frequency response of the fMRI signal.

Numerous studies that have applied classification and other multivariate methods to fMRI data have reached conclusions that do not depend on whether or not differential signals arise from columnar scale architecture. In general, multivariate methods are statistically more sensitive than univariate methods (Kriegeskorte et al., 2006). In fMRI studies, that increased sensitivity is most advantageous when differential signals are robust but present in a small subset of voxels or not robust but distributed across many voxels. Furthermore, multivariate methods, especially those based on encoding and decoding models, can be used to investigate whether neural representations generalize across stimulus conditions and cognitive states (Kamitani and Tong, 2005, 2006; Dinstein et al., 2008; Kay et al., 2008; Brouwer and Heeger, 2009; Harrison and Tong, 2009; Kay and Gallant, 2009). For example, the initial report of orientation decoding (Kamitani and Tong, 2005) showed that a classifier trained while subjects viewed single oriented gratings could be used to predict the orientation that subjects attended to while viewing two superimposed gratings, thus demonstrating that feature-based attention can bias neural responses toward those of the attended orientation. This particular conclusion does not depend on the source of the underlying orientation preference exploited by the classifier. Examples like this show that multivariate methods can be used to draw inferences about neural representation, without relying on the assumption that decoding accuracy reflects or exploits a fine-scale (e.g., subvoxel or columnar) functional architecture.

References

- Adams DL, Sincich LC, Horton JC (2007) Complete pattern of ocular dominance columns in human primary visual cortex. *J Neurosci* 27:10391–10403.
- Berkley MA, Kitterle F, Watkins DW (1975) Grating visibility as a function of orientation and retinal eccentricity. *Vision Res* 15:239–244.
- Blasdel GG, Salama G (1986) Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature* 321:579–585.
- Boynton GM (2005) Imaging orientation selectivity: decoding conscious perception in V1. *Nat Neurosci* 8:541–542.
- Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16:4207–4221.
- Brouwer GJ, Heeger DJ (2009) Decoding and reconstructing color from responses in human visual cortex. *J Neurosci* 29:13992–14003.
- Das A, Gilbert CD (1997) Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature* 387:594–598.
- De Valois RL, Yund EW, Hepler N (1982) The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res* 22:531–544.
- DeYoe EA, Carman GJ, Bandettini P, Glickman S, Wieser J, Cox R, Miller D, Neitz J (1996) Mapping striate and extrastriate visual areas in human cerebral cortex. *Proc Natl Acad Sci U S A* 93:2382–2386.
- Dinstein I, Gardner JL, Jazayeri M, Heeger DJ (2008) Executed and observed movements have different distributed representations in human aIPS. *J Neurosci* 28:11231–11239.
- Dumoulin SO, Wandell BA (2008) Population receptive field estimates in human visual cortex. *Neuroimage* 39:647–660.
- Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN (1994) fMRI of human visual cortex. *Nature* 369:525.
- Engel SA, Glover GH, Wandell BA (1997) Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7:181–192.
- Furmanski CS, Engel SA (2000) An oblique effect in human primary visual cortex. *Nat Neurosci* 3:535–536.
- Ganguli D, Simoncelli E (2011) Implicit encoding of prior probabilities in optimal neural populations. In: *Advances in Neural Information Processing Systems (NIPS)*, Vol 23 (Lafferty J, Williams C, Shawe-Taylor J, Zemel RS, Culotta A, eds), pp 658–666. Cambridge, MA: MIT.
- Gardner JL (2010) Is cortical vasculature functionally organized? *Neuroimage* 49:1953–1956.
- Gardner JL, Merriam EP, Movshon JA, Heeger DJ (2008) Maps of visual space in human occipital cortex are retinotopic, not spatiotopic. *J Neurosci* 28:3988–3999.
- Grinvald A, Lieke E, Frostig RD, Gilbert CD, Wiesel TN (1986) Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature* 324:361–364.
- Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632–635.
- Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
- Haynes JD, Rees G (2006) Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7:523–534.
- Hubel DH, Wiesel TN (1963) Shape and arrangement of columns in cat's striate cortex. *J Physiol* 165:559–568.
- Jammalamadaka SR, Sengupta A (2001) *Topics in circular statistics*. River Edge, NJ: World Scientific.
- Kalatsky VA, Stryker MP (2003) New paradigm for optical imaging: temporally encoded maps of intrinsic signal. *Neuron* 38:529–545.
- Kamitani Y, Sawahata Y (2010) Spatial smoothing hurts localization but not information: pitfalls for brain mappers. *Neuroimage* 49:1949–1952.
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685.
- Kamitani Y, Tong F (2006) Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16:1096–1102.
- Kay KN, Gallant JL (2009) I can see what you see. *Nat Neurosci* 12:245.
- Kay KN, Naselaris T, Prenger RJ, Gallant JL (2008) Identifying natural images from human brain activity. *Nature* 452:352–355.
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci U S A* 103:3863–3868.
- Kriegeskorte N, Simmons WK, Bellgowan PS, Baker CI (2009) Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci* 12:535–540.
- Kriegeskorte N, Cusack R, Bandettini P (2010) How does an fMRI voxel sample the neuronal activity pattern: compact-kernel or complex spatio-temporal filter? *Neuroimage* 49:1965–1976.
- Larsson J, Heeger DJ (2006) Two retinotopic visual areas in human lateral occipital cortex. *J Neurosci* 26:13128–13142.
- Leventhal AG (1983) Relationship between preferred orientation and receptive field position of neurons in cat striate cortex. *J Comp Neurol* 220:476–483.

- Levick WR, Thibos LN (1980) Orientation bias of cat retinal ganglion cells. *Nature* 286:389–390.
- Levick WR, Thibos LN (1982) Analysis of orientation bias in cat retina. *J Physiol* 329:243–261.
- Li B, Peterson MR, Freeman RD (2003) Oblique effect: a neural basis in the visual cortex. *J Neurophysiol* 90:204–217.
- Maldonado PE, Gödecke I, Gray CM, Bonhoeffer T (1997) Orientation selectivity in pinwheel centers in cat striate cortex. *Science* 276:1551–1555.
- Mannion DJ, McDonald JS, Clifford CW (2009) Discrimination of the local orientation structure of spiral Glass patterns early in human visual cortex. *Neuroimage* 46:511–515.
- Mannion DJ, McDonald JS, Clifford CW (2010) Orientation anisotropies in human visual cortex. *J Neurophysiol* 103:3465–3471.
- Mansfield RJ (1974) Neural basis of orientation perception in primate vision. *Science* 186:1133–1135.
- Nestares O, Heeger DJ (2000) Robust multiresolution alignment of MRI brain volumes. *Magn Reson Med* 43:705–715.
- Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424–430.
- Ogawa S, Lee TM, Kay AR, Tank DW (1990) Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci U S A* 87:9868–9872.
- Ohki K, Chung S, Kara P, Hübener M, Bonhoeffer T, Reid RC (2006) Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442:925–928.
- Op de Beeck HP (2010) Against hyperacuity in brain reading: spatial smoothing does not hurt multivariate fMRI analyses? *Neuroimage* 49:1943–1948.
- Op de Beeck HP, Haushofer J, Kanwisher NG (2008) Interpreting fMRI data: maps, modules and dimensions. *Nat Rev Neurosci* 9:123–135.
- Ringach DL (2007) On the origin of the functional architecture of the cortex. *PLoS ONE* 2:e251.
- Rodieck RW, Binmoeller KF, Dineen J (1985) Parasol and midget ganglion cells of the human retina. *J Comp Neurol* 233:115–132.
- Rothkopf CA, Weisswange TH, Triesch J (2009) Learning independent causes in natural images explains the spacevariant oblique effect. Paper presented at the IEEE 8th International Conference on Development and Learning, Shanghai, China, June 5–7, 2009.
- Rovamo J, Virsu V, Laurinen P, Hyvärinen L (1982) Resolution of gratings oriented along and across meridians in peripheral vision. *Invest Ophthalmol Vis Sci* 23:666–670.
- Sasaki Y, Rajimehr R, Kim BW, Ekstrom LB, Vanduffel W, Tootell RB (2006) The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron* 51:661–670.
- Schall JD, Vitek DJ, Leventhal AG (1986) Retinal constraints on orientation specificity in cat visual cortex. *J Neurosci* 6:823–836.
- Schira MM, Tyler CW, Breakspear M, Spehar B (2009) The foveal confluence in human visual cortex. *J Neurosci* 29:9050–9058.
- Serences JT, Saproo S, Scolari M, Ho T, Muftuler LT (2009) Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *Neuroimage* 44:223–231.
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–893.
- Shmuel A, Chaimow D, Raddatz G, Ugurbil K, Yacoub E (2010) Mechanisms underlying decoding at 7 T: ocular dominance columns, broad structures, and macroscopic blood vessels in V1 convey information on the stimulated eye. *Neuroimage* 49:1957–1964.
- Shou T, Ruan D, Zhou Y (1986) The orientation bias of LGN neurons shows topographic relation to area centralis in the cat retina. *Exp Brain Res* 64:233–236.
- Smith AM, Lewis BK, Rüttemann UE, Ye FQ, Sinnwell TM, Yang Y, Duyn JH, Frank JA (1999) Investigation of low frequency drift in fMRI signal. *Neuroimage* 9:526–533.
- Smith EL 3rd, Chino YM, Ridder WH 3rd, Kitagawa K, Langston A (1990) Orientation bias of neurons in the lateral geniculate nucleus of macaque monkeys. *Vis Neurosci* 5:525–545.
- Stampe D (1993) Heuristic filtering and reliable calibration methods for video-based pupil-tracking systems. *Behav Res Methods Instrum Comput* 25:137–142.
- Swisher JD, Gatenby JC, Gore JC, Wolfe BA, Moon CH, Kim SG, Tong F (2010) Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. *J Neurosci* 30:325–330.
- Vul E, Harris C, Winkielman P, Pashler H (2009) Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspect Psychol Sci* 4:274–290.
- Wandell BA, Dumoulin SO, Brewer AA (2007) Visual field maps in human cortex. *Neuron* 56:366–383.
- Westheimer G (2003) The distribution of preferred orientations in the peripheral visual field. *Vision Res* 43:53–57.
- Wilson JR, Sherman SM (1976) Receptive-field characteristics of neurons in cat striate cortex: changes with visual field eccentricity. *J Neurophysiol* 39:512–533.