

STIMULUS VIGNETTING AND ORIENTATION SELECTIVITY IN HUMAN VISUAL CORTEX

(Abbreviated title: Stimulus vignetting)

Zvi N. Roth¹, David J. Heeger^{2,3}, Elisha P. Merriam¹

¹ Laboratory of Brain and Cognition
National Institute of Mental Health, National Institutes of Health,
Building 10, Rm 4C-432B
10 Center Dr
Bethesda, MD 20892 USA

² Department of Psychology and ³ Center for Neural Science
New York University
4-6 Washington Place
New York, NY 10003 USA

Correspondence should be addressed to Z.N.R. (zvi.roth@nih.gov)

Number of Figures: 7

Number of pages: 25

Research Article

Keywords: orientation; multivariate classification; primary visual cortex; fMRI

Acknowledgements: Supported by NIH grants R01-EY025673, and by the Intramural Research Program of the National Institutes of Health (ZIA-MH-002909) - National Institute of Mental Health Clinical Study Protocol 93-M-0170, NCT00001360. ZNR was supported by The ELSC Postdoctoral Fellowship Abroad, Hebrew University of Jerusalem. Special thanks to Eero Simoncelli and to Peter Bandettini for helpful comments.

Abstract

Neural selectivity to orientation is one of the simplest and most thoroughly-studied cortical sensory features. Here, we show that a large body of research that purported to measure orientation tuning may have in fact been inadvertently measuring sensitivity to second-order changes in luminance, a phenomenon we term 'vignetting'. Using a computational model of neural responses in primary visual cortex (V1), we demonstrate the impact of vignetting on simulated V1 responses. We then used the model to generate a set of predictions, which we confirmed with functional MRI experiments in human observers. Our results demonstrate that stimulus vignetting can wholly determine the orientation selectivity of responses in visual cortex measured at a macroscopic scale, and suggest a reinterpretation of a well-established literature on orientation processing in visual cortex.

Introduction

Primary visual cortex (V1) is likely the best studied sensory cortical area, and is a model for understanding broad principles of cortical processing. Similarly, orientation in V1 is likely one of the simplest and best studied cortical sensory features. Orientation is used as a model for understanding more complex feature processing in other cortical areas, and oriented V1-like receptive fields play an important role in successful computational models of vision. Yet even something as basic as the map of orientation in V1 is inadequately understood. At a fine scale, the map shows an orderly periodic structure with pinwheels in hypercolumns, which have a periodicity of about 1 mm in monkeys (Hubel and Wiesel, 1963; Blasdel and Salama, 1986; Grinvald et al., 1986; Das and Gilbert, 1997; Maldonado et al., 1997; Ohki et al., 2006), and which are thought to be slightly larger in humans (Adams et al., 2007; Yacoub et al., 2008). While the fine-scale structure has been studied extensively, the coarse-scale structure of orientation selectivity is poorly established.

We and others, using fMRI, discovered a coarse-scale orientation bias in human V1; each voxel exhibits an orientation preference that depends on the region of space that it represents (Furmanski and Engel, 2000; Sasaki et al., 2006; Mannion et al., 2010; Freeman et al., 2011; Freeman et al., 2013; Larsson et al., 2017). We observed a radial bias in the peripheral representation of V1: voxels that responded to peripheral locations near the vertical meridian tended to respond most strongly to vertical orientations; voxels along the peripheral horizontal meridian responded most strongly to horizontal orientations; likewise for oblique orientations. This phenomenon had gone mostly unnoticed previously. We discovered this striking phenomenon with fMRI because fMRI covers the entire retinotopic map in visual cortex, making it an ideal method for characterizing such coarse-scale representations.

The observation of a coarse-scale orientation bias raises two questions. First, what is the relative contribution of coarse- and fine-scale biases to fMRI measurements of orientation selectivity? A leading conjecture is that the orientation preferences in fMRI measurements arise primarily from random spatial irregularities in the fine-scale columnar architecture (Boynton, 2005; Haynes and Rees, 2005; Kamitani and Tong, 2005). This conjecture has inspired a large literature on the use of multivariate statistical analyses to exploit these small biases, both to study orientation selectivity in V1 and to study a wide range of other functions throughout the brain. On the other hand, we have argued that the coarse-scale orientation bias is the predominant orientation-selective signal measured with fMRI, and that multivariate decoding analysis methods are successful because of it (Freeman et al., 2011; Freeman et al., 2013). This conjecture remains controversial because the notion that fMRI is sensitive to fine-scale neural activity is highly attractive, even though it has been proven difficult to validate empirically (Alink et al., 2013; Pratte et al., 2016; Alink et al., 2017).

Second, while the neural architecture that gives rise to orientation selectivity is well understood, the neural basis for the coarse-scale orientation bias remains unknown. It has been proposed that the coarse-scale bias is a byproduct of the edge of the stimulus

(Carlson, 2014; Clifford and Mannion, 2015; But see Wardle et al., 2017). By this account, the coarse-scale bias does not directly reflect enhanced neural responses for particular orientations, but rather properties of the stimuli. This account, if correct, would require the reinterpretation of a vast number of previous studies, since in nearly all experiments on orientation, the stimuli are confined by the edges of a finite aperture, or, at very least, by the edge of the display screen.

We tested the hypothesis that stimulus vignetting is the primary underlying source of the coarse-scale orientation bias. Stimulus vignetting refers to the potential influence of the edge of the aperture. We use this term to emphasize that it is not the edge per se, but rather an interaction between the orientation of the stimulus and a second feature of the display (the aperture or “vignette”) that bounds the stimulus. We developed a theoretical account of stimulus vignetting using a computational model of V1 activity. We used this computational model to generate predictions, which we then tested with empirical data from fMRI experiments. Our results provide a framework for reinterpreting a wide-range of findings in the visual system.

Results

Coarse-scale bias and stimulus vignetting: Model simulations

We studied the impact of the stimulus vignetting on coarse-scale orientation preference maps. The principle underlying stimulus vignetting is that the vignette spreads Fourier energy. The direction and amplitude of this spread is determined by the underlying stimulus spectrum and the aperture geometry, in much the same way that the Fourier transform of a sinusoid windowed by a rectangle consists of the sinusoid frequency convolved with a sinc function. Due to this spectral spread at the aperture edge, neurons with receptive fields straddling the aperture edge respond to a combination of spatial frequency and orientation components, different from neurons with receptive fields far from the edge (Figure 1).

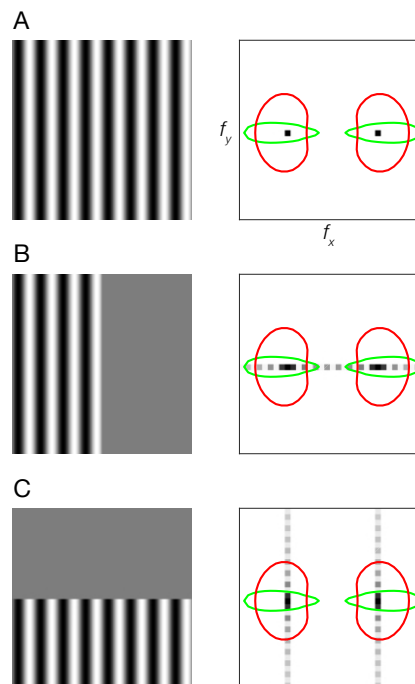


Figure 1. Stimulus vignetting. (A) Stimulus image (left) and Fourier amplitude (right) for sinusoidal grating that is infinite in extent. Black dots, Fourier amplitude restricted to a single spatial frequency and orientation component. Colored ovals, orientation tuning (polar angle cross-section through the oval) and spatial frequency tuning (radial cross-section) of two hypothetical V1 neurons. Red oval, neuron with broad orientation tuning and narrow spatial frequency tuning. Green oval, neuron with narrow orientation tuning and broad spatial frequency tuning. (B) Sinusoidal grating vignettted by vertical aperture. Light gray dots, spread of Fourier energy because of the aperture. (C) Sinusoidal grating vignettted by horizontal aperture.

We began by simulating the results of our previous fMRI experiment (Freeman et al., 2011) using a computational model of V1 (Simoncelli et al., 1992) (see Methods: Theoretical Model). The inputs to the model were identical to the stimuli used by Freeman et al. (2011). Specifically, the stimuli consisted of 16 evenly spaced orientation

angles between 0° and 180° . All stimuli were vignetted by a 4° annulus (inner radius, 5° ; outer radius, 9°). Both the inner and outer edges of the annulus were blurred with a 1° raise-cosine transition (centered on the inner and outer edges). We measured the model's responses to each image separately. Each pixel of the model's output can be thought of as a simulated neuron in the retinotopic map of V1. To simulate an fMRI voxel's response to the stimuli used in Freeman et al., we summed the model responses across all orientation channels. This procedure yielded a spatial array of simulated voxel responses with different population receptive fields (pRFs). The response of a single simulated fMRI voxel corresponded to a single spatial location on the map of neural responses (e.g., with a pRF straddling the aperture edge). This yielded a set of 16 simulated voxel maps, one for each stimulus orientation (Figure 2). Importantly, the model's response should not exhibit any orientation tuning, since we summed across all of the orientation channels. But the simulated neural responses exhibited a clear coarse-scale bias for orientation (Figure 2). The coarse-scale orientation bias matched Freeman et al. (2011) in that the largest responses were observed for radial orientations (Figure 2).

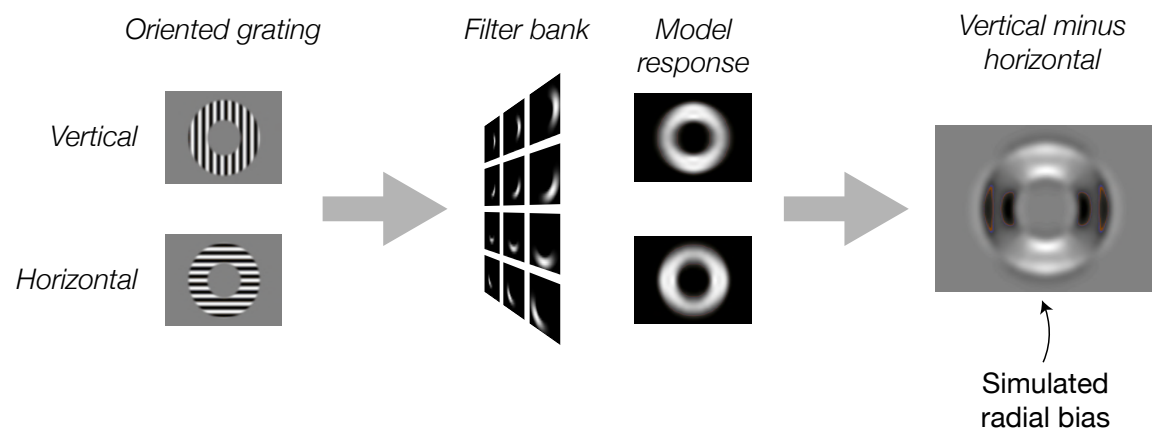


Figure 2. Simulated neural responses illustrating the impact of stimulus vignetting. Oriented gratings were vignetted by inner and outer annuli, using stimulus parameters identical to those from Freeman et al. (2012). Next, the gratings were used as input to the model. Model responses were computed separately for vertical (top) and horizontal (bottom) gratings. Finally, model output was computed as the responses to vertical minus response to horizontal gratings. Model exhibits a preference for horizontal gratings along the horizontal meridian and a preference for vertical gratings along the vertical meridian (i.e., a radial bias).

Having established that stimulus vignetting can produce a radial coarse-scale bias, we next asked whether the computational model predicts the influence of vignetting on responses to novel stimuli. The simulation results suggest that orientation-selective responses measured with fMRI could reflect, at least in part, the shape of the vignette, rather than the underlying orientation tuning of individual neurons. To test this possibility, we generated a novel set of stimuli with vignettes that were orthogonal to one another. If the vignette affects orientation bias, changing the vignette should have a predictable

effect on orientation bias.

The novel stimuli for this simulation were created as follows. Vertical and horizontal Cartesian gratings (the carrier) were multiplied by a radial or an angular polar grating (the modulator). These compound stimuli were then passed through the model. Subtracting horizontal from vertical carrier grating outputs resulted in an image of orientation bias for the model.

The model predicted that the coarse-scale orientation bias should depend on the modulator. The radial modulator evoked a coarse-scale bias that was radial (Figure 3). In other words, at regions around the vertical meridian there is a greater response for the vertical orientation than for the horizontal orientation (light regions in Figure 3), and vice versa around the horizontal meridian (dark regions in Figure 3). The model predicted that changing the orientation of the modulator should affect the orientation preference to the carrier grating. We found that the angular modulator evoked a coarse-scale bias that was tangential. Thus, at the vertical meridian the horizontal grating yielded higher energy, and at the horizontal meridian the vertical grating evokes higher energy responses.

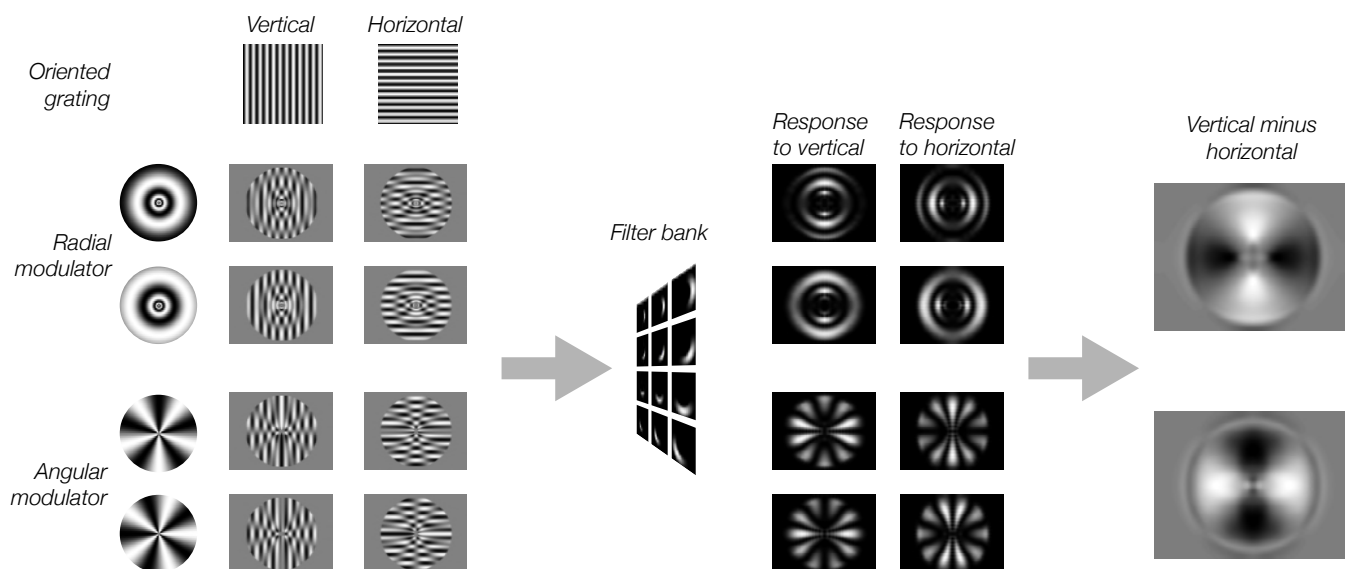


Figure 3. Model predictions for modulated gratings. Each of the eight stimuli were created by multiplying a vertical or horizontal grating by a radial or angular modulator. These stimuli were used as input to the model. Model responses were computed as in Figure 2. For radial modulated gratings (top two rows), the model exhibits a radial preference: larger responses to horizontal gratings along the horizontal meridian, larger responses to vertical gratings along the vertical meridian. However, for angular modulated gratings (bottom two rows), the orientation preference is tangential: larger responses for horizontal gratings along the vertical meridian and larger responses for vertical gratings along the horizontal meridian.

The stimuli used in Freeman et al. (2011), and in our simulations in Figure 2, had vignettes with gradual contrast changes, indicating that the impact of vignetting on orientation bias is not dependent on the presence of sharp edges. But is it possible that a more gradual edge would have mitigated vignetting effects? We answered this questions by using a sinusoidal modulation that lacked any edges (Figure 3, left). We found similar similar results for both sinusoidal modulators and modulators with hard edges (i.e. square-wave), as expected, because vignetting arises from any contrast change in the image, rather than an “edge” per se (see Discussion).

Coarse-scale bias and stimulus vignetting: fMRI experiments

We found evidence of stimulus vignetting at both 3T and 7T field strengths, at different spatial resolutions, and for both square wave and sinusoidal modulators, as predicted by the theoretical model. The radial modulator evoked a radial bias (i.e., orientation preferences pointing inward toward the fovea). The angular modulator evoked a tangential bias (i.e., orientation preferences that were rotated by 90 deg from radial) (Figure 4). Furthermore, on a voxel-by-voxel basis, preferred orientations were shifted by 90 deg across modulators (Figure 5).

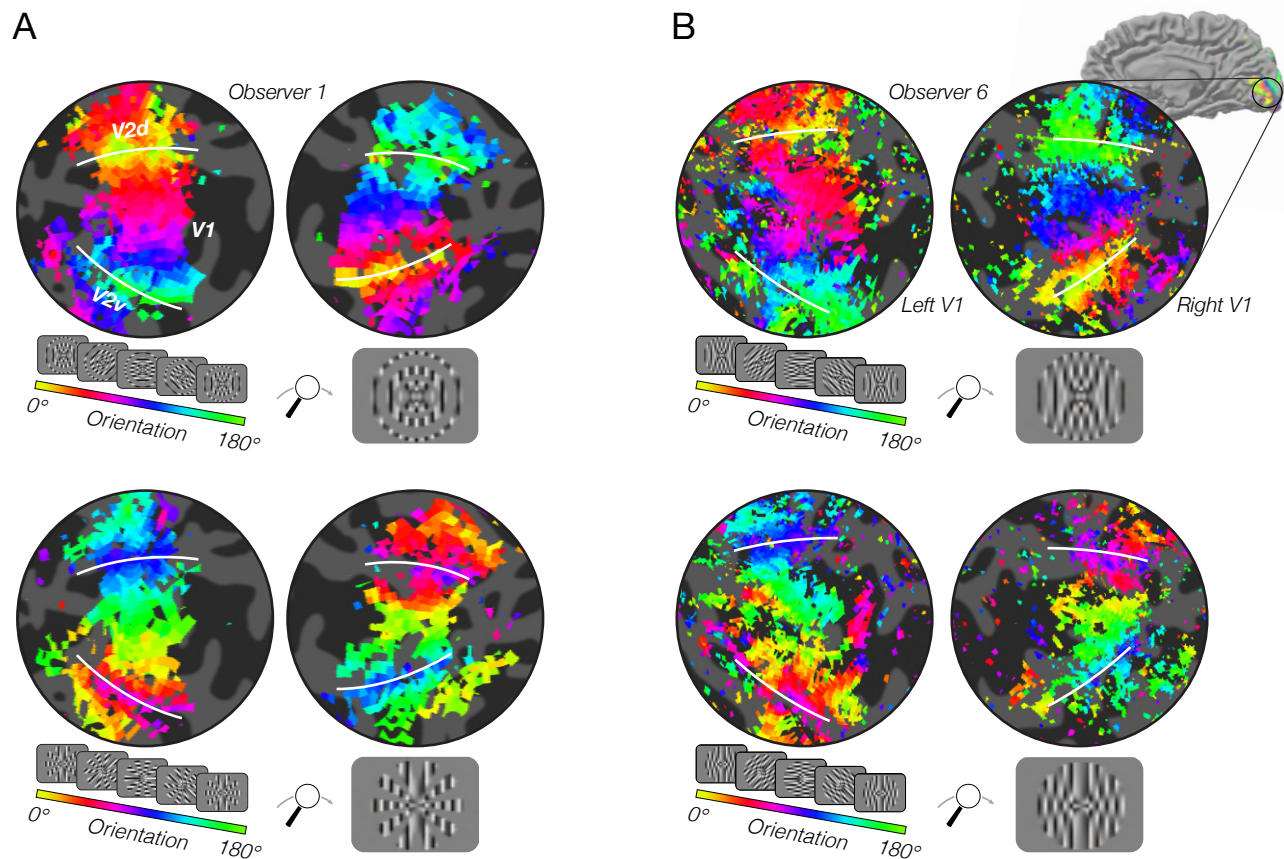


Figure 4. fMRI measurements of orientation bias depend on stimulus vignetting. (A) Top: Responses to phase-encoded oriented gratings multiplied by a static angular modulator (shown

in inset). The oriented grating cycled through 16 steps of orientation ranging from 0° to 180° every 24 s. The angular modulator was constant during the entire experiment. Map thresholded at coherence of 0.2. Hue indicates phase of the best fitting sinusoid. White lines indicate V1/V2 boundaries. Bottom: Responses to the same oriented grating as in A, but here the grating is multiplied by a radial modulator. As predicted by the model, the radial modulator gave rise to a radial orientation bias, while the angular modulator gave rise to a tangential orientation bias. (B) High-resolution, high field strength measurements of orientation preference for radial and angular modulators. Stimuli and conventions same as for A, except the modulators were radial and angular sinusoids.

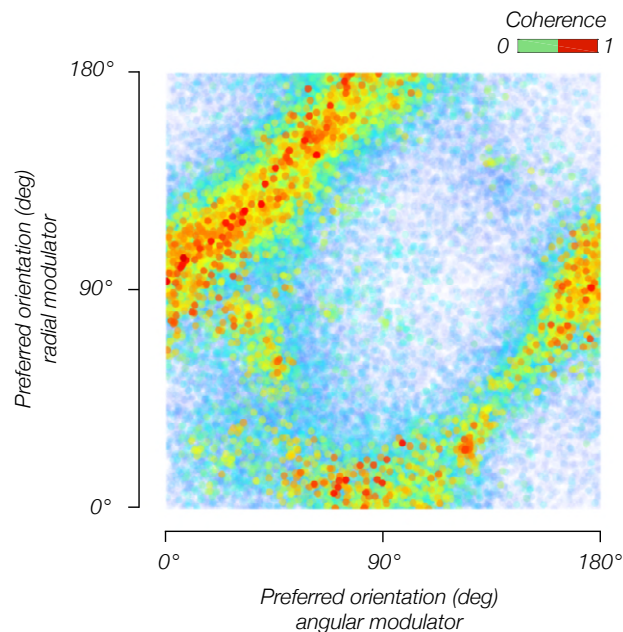


Figure 5. Orientation bias depends on the modulator. Orientation bias for V1 voxels pooled across 12 subjects for the angular modulator (*y*-axis) and radial modulator (*x*-axis).

Comparing model predictions to data from individual voxels, we observed a clear correspondence between predicted and measured orientation preference (Figure 6). Circular correlations between predicted and measured orientation preferences were 0.26 and 0.20 for radial and angular modulators, respectively ($p < 0.0001$ for both). Note that many sources of noise in the measurements that affect this analysis: in estimating the location and size of each voxel's pRF, in measuring each voxel's orientation preference, and in coregistering each voxel's response across multiple fMRI sessions. Moreover, the model we use is highly simplified; for example, it does not take into account changes in spatial frequency tuning at greater eccentricities. Yet, despite the multiple sources of noise and the simplified assumptions of the model, the correspondence between the model's prediction and the empirical measurements are highly statistically significant. From this, we conclude that stimulus vignetting is a primary source of the course scale bias.

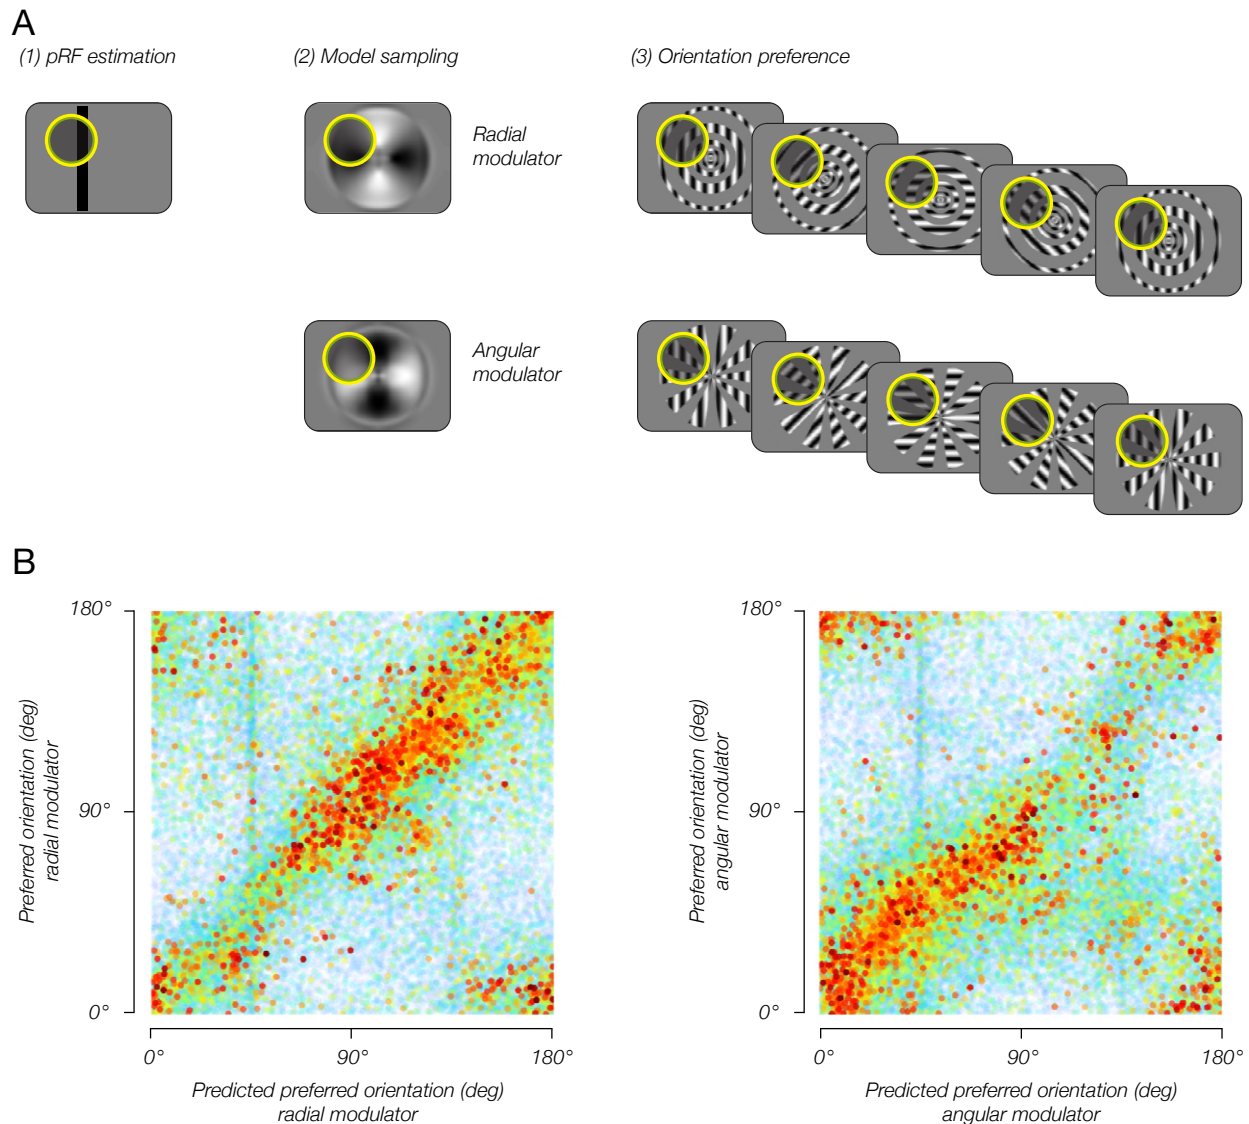


Figure 6. Combining the model predictions with pRF estimates enables reliable predictions of single-voxel orientation preference. (A) pRF sampling method. (1) Each voxel's pRF was estimated based on an independent pRF-mapping scanning session. (2) Model output was sampled by the voxel pRF, resulting in an estimated preferred orientation for each modulator. (3) Estimated preferred orientation was compared to measured orientation preference, for each modulator. (B) Estimated orientation preference plotted against measured orientation preference. Sampling the model output with estimated pRFs yielded accurate predictions for each voxels' orientation preferences, indicating that the model of stimulus vignetting provides a good account for the fMRI measurements of orientation preference.

We hypothesized that if stimulus vignetting accounts for the radial bias, and changing the direction of the modulator flips voxels' orientation preferences by 90 deg, then the population pattern of activity should also shift by 90 deg. We tested this hypothesis by training a classifier on fMRI responses to the radial modulator, and testing how

accurately it decodes responses to the angular modulator, and vice versa. We compared between-modulator decoding accuracy to accuracy obtained from training and testing with data from the same modulator (within-modulator). Within-modulator decoding yielded high accuracies ($p < 0.0001$, tested individually for each modulator with a non-parametric permutation test). Across-modulator decoding accuracy was below chance ($p < 0.0001$ for each)(Figure 7A). These results indicate that the population response to oriented gratings was dependent on the modulator — responses to an oriented grating multiplied by one modulator (e.g. radial) were entirely different from responses to the same oriented grating when multiplied by the other modulator (e.g. by the angular modulator).

Does changing the modulator (between radial and angular) shift the population response pattern by 90 deg, as predicted by the model? To test this possibility, we repeated the between-modulator decoding analysis described above, but shifted the orientation labels by 90°. After shifting the orientation labels, decoding accuracy returned to levels significantly above chance ($p < 0.0001$ for both modulators), only slightly lower than within-modulator decoding accuracy on the unshifted data. Note that the 2 modulators share inner and outer radial edges. Thus, not all voxels are expected to shift by 90 deg. These common edges may be the cause of the slightly lower decoding accuracy across modulators.

We performed a complimentary analysis, based on multidimensional scaling, to visualize the responses for both modulators in a single 2D plot. In Figure 7B, each line represents a population response to a single stimulus orientation. The color and the line orientation in the plot both indicate the orientation of the stimulus. The distances between the lines reflect the dissimilarity between the corresponding population responses: similar responses are near to each other, and dissimilar responses are further apart. For each modulator, we found the responses to different orientations to be organized continuously in a ring. We found the two rings, corresponding to the two modulators, to be misaligned with each other, reflecting the change in the population responses resulting from the change in modulator (Figure 7B). Specifically, responses to identical orientations (i.e. identical colored lines) are distant, while responses to dissimilar orientations are proximal. However, after rotating the labels for responses to the angular modulator by 90 deg, the two sets of responses are well aligned. This demonstrates once again that changing modulators completely flips voxels' orientation bias, and this is reflected at the level of entire voxel populations.

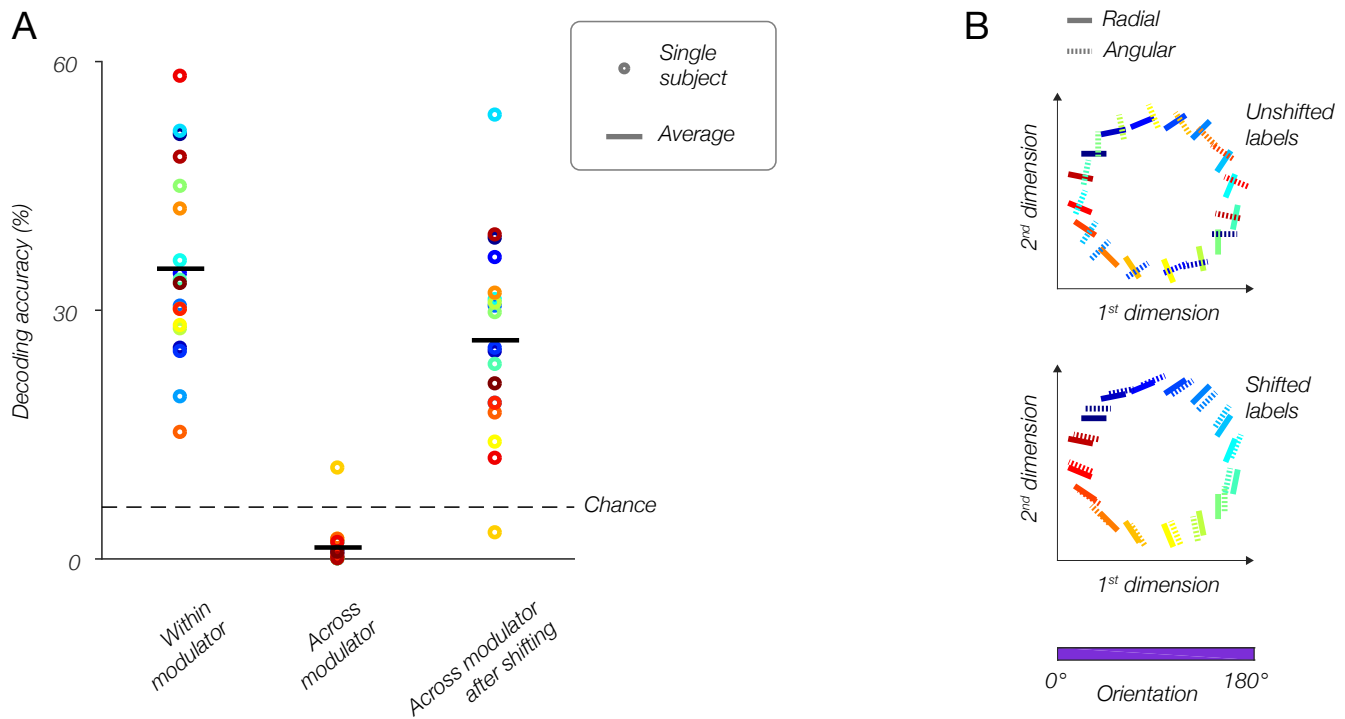


Figure 7. Population responses reflect the modulator, not stimulus orientation. (A) fMRI decoding. Within-modulator decoding: classifier was trained and tested on data from each modulator separately, in a leave-one-run-out cross-validation approach. Across modulator decoding: classifier was trained on data from one modulator and tested on data from the other modulator (and vice-versa). Across modulator decoding after shifting: data for one modulator was shifted by 90°, then the classifier was trained on that modulator and tested on the other (and vice-versa). (B) Multidimensional scaling. Top: population responses to all orientations for both modulators (solid lines, radial modulator; dotted lines, angular modulator) are plotted in 2D, such that distances between conditions reflect how different (i.e. dissimilarity = 1- r , where r is the correlation coefficient) the responses are. For the original data, orthogonal orientations from the two modulators have similar responses and are plotted near each other, while identical orientations taken from the two modulators evoked very different responses, and lay distant from each other. Bottom: after shifting the labels for the angular modulator by 90 deg, the two datasets align well, and similar orientations correspond to similar population responses regardless of the modulator.

Discussion

Our results show that stimulus vignetting strongly affects voxel orientation bias. This suggests that results from many fMRI studies attempting to measure neural orientation tuning properties may actually reflect stimulus vignetting. Since stimulus vignetting masquerades as orientation selectivity amongst voxels that correspond to the stimulus edge, and because edges tend to be continuous over a large spatial extent, vignetting will generally result in a coarse-scale map of orientation bias. In many fMRI studies, gratings are presented within a circular annulus (Haynes and Rees, 2005; Kamitani and Tong, 2005; Harrison and Tong, 2009; Freeman et al., 2011; Sun et al., 2013; Larsson et al., 2017), and in these cases the model predicts a radial coarse-scale bias. However, in some cases, other stimulus vignetting exists, resulting in more subtle coarse-scale patterns of orientation bias (Swisher et al., 2010; Cichy et al., 2015; Sengupta et al., 2017). One study (Alink et al., 2017) used inner and outer circular annuli, but added additional angular edges, the result of which should be a combination of radial and tangential biases. Indeed, this study reported that voxels had a mixed pattern of selectivity, with a considerable number of voxels reliably preferring tangential gratings, and other voxels reliably favoring radial orientations.

Stimulus vignetting is a general issue of concern in visual neuroscience (Carter and Henning, 1971; Cichy et al., 2015; Tang et al., 2018). According to the theoretical model that we have applied here, whenever a neuron's receptive field overlaps a stimulus edge or a change in contrast, stimulus vignetting will spread the Fourier power and affect the neuron's response. As a result, such a neuron may exhibit ostensible orientation tuning even when it is not orientation selective. Furthermore, even if a neuron is truly orientation selective, measurements of its orientation tuning may be affected by the aperture shape.

A major question concerning the neural origin of orientation selectivity is whether and to what degree lateral geniculate nucleus (LGN) responses are tuned to orientation. Recently, this question was addressed using fMRI. One group found that orientation can be decoded from LGN, suggesting that thalamic responses are tuned to orientation (Ling et al., 2015), perhaps inheriting their selectivity from the retina (Schall et al., 1986). Another group, using calcium imaging to measure bouton responses in mice, found extensive orientation tuning as well (Sun et al., 2016). But in both of these studies, however, stimulus vignetting may have played a critical role. In the former study (Ling et al., 2015), gratings were presented within a circular annulus, likely yielding a radial bias due to stimulus vignetting. In the latter study (Sun et al., 2016), the grating covered the entire visual field, so at first glance vignetting shouldn't have been a dominant source of orientation selectivity. However, mouse neuronal receptive fields are remarkably large, at least an order of magnitude larger than in primates (Van den Bergh et al., 2010; Tang et al., 2016). Therefore, a large proportion of neuron RFs measured in Sun et al. (2016) would have overlapped the edges of the screen. Stimulus vignetting may well account for the dominant vertical bias and less prominent horizontal bias that they observed, since such a bias would reflect the rectangular shape of the stimulus screen.

Do voxel responses contain any orientation preferences that are not the result of stimulus vignetting? We found no evidence for any other orientation information present in the data. Based on the current study we cannot rule out the possibility of other factors affecting fMRI orientation preferences, at both coarse and fine spatial scales. For example, fine-scale signals originating from a columnar organization may, theoretically, affect fMRI signals (Boynton, 2005). But these columnar signals may be entirely eclipsed by stimulus vignetting. Similarly, coarse-scale signals such as a physiological radial bias may indeed exist, but be relatively small, and hence also eclipsed by stimulus vignetting. Our previous results, describing a vertical bias near the fovea (Freeman et al., 2013), in addition to the radial bias in the periphery (Freeman et al., 2011), lend support to the possibility that there are multiple sources of coarse-scale bias.

Do fMRI voxels reflect any fine-scale orientation signals originating from a columnar organization, in addition to an orientation bias resulting from vignetting? If each voxel reflects the activity of a large number of neurons, signals from multiple orientation columns will be summed within a voxel, thus attenuating any columnar signal. If this is the case, decreasing the voxel size should lower the number of different columns summed within each voxel, and amplify columnar orientation selectivity. On the other hand, orientation bias resulting from vignetting should not depend on scanning resolution. Nearby neurons respond to nearby locations in the visual field, and therefore will respond similarly to the vignette. Thus, increasing the acquisition resolution will augment columnar signals but not vignetting signals. Our results at 2 mm voxels demonstrate primarily vignetting with no evidence of a columnar signal. Increasing the acquisition resolution to 0.9 mm did not alter this result. From this we conclude that the coarse-scale bias due to vignetting is the predominant source of orientation information measured with fMRI.

Stimulus vignetting has previously been referred to as an “edge effect”, and we demonstrated the intuition of stimulus vignetting using examples of hard edges (Figure 1). Several studies have used gratings with smoothed edges, and some have smoothed the edges explicitly to eliminate any effects related to the edge (Freeman et al., 2011; Warren et al., 2014). According to the model, however, vignetting does not occur exclusively in the vicinity of edges. Rather, a contrast change, even a gradual change, is all that is necessary to produce stimulus vignetting. Smoothing the stimulus edge only acts to spatially spread the affect on the Fourier spectrum across a larger portion of the visual field. Indeed, we found that both square-wave modulators and sinusoidal modulators evoked similar levels of vignetting, empirically confirming the model’s prediction.

Here, we demonstrate that stimulus vignetting can result in a radial bias. By this account, the coarse-scale bias does not directly reflect greater neural responses for particular orientations, but rather primarily properties of the stimuli. However, other possible sources for the coarse-scale bias may exist as well.

First, the coarse-scale orientation bias may reflect, in addition to stimulus vignetting, a neuronal “gain map” (Philips and Chakravarthy, 2017). Neurons in visual cortex may have different response gains depending on their preferred orientations and receptive field locations. For example, the responses of a population of neurons tuned for vertical might be highest for those neurons with receptive fields near the vertical meridian and smallest for those along the horizontal meridian, and vice versa for neurons tuned for horizontal, conferring a cardinal bias.

A gain map could be instantiated in visual cortex through a number of different physiological mechanisms. For example, in a population of neurons along the horizontal meridian, a gain map might be instantiated by 1) a larger proportion of neurons that prefer horizontal orientations, or 2) higher firing rates for the sub-population of neurons that prefer horizontal orientations, or 3) wider tuning bandwidth for sub-population of neurons that prefer orientations close to vertical. Each of these three mechanisms would give rise to a larger population response for horizontal gratings in fMRI voxels that lay along the horizontal meridian.

Additionally, the coarse-scale orientation bias might reflect “asymmetric surround suppression”, an imbalance in suppression between different parts of the region surrounding a neuron’s receptive field. The responsiveness of a V1 neuron is suppressed by stimuli outside its classical RF (Sceniak et al., 1999; Sceniak et al., 2001; Cavanaugh et al., 2002a, b; Zenger-Landolt and Heeger, 2003). Suppression is stronger for neurons with RFs near the center of a stimulus aperture, less so for neurons with RFs near the aperture edge, because there is less stimulation surrounding those RFs. For many neurons in V1, surround suppression is asymmetric, such that there is more suppression at the “ends” of the receptive field (i.e., along the preferred orientation) than the “sides” (Walker et al., 1999; Cavanaugh et al., 2002a, b; Tanaka and Ohzawa, 2009; Hallum and Movshon, 2014). Because of this asymmetry, the release from suppression may be greater for aperture edges that are orthogonal to each neuron’s preferred orientation, such that a circular aperture centered at fixation will confer a radial bias.

Regardless of whether a gain map and/or asymmetric surround suppression contribute to coarse-scale orientation bias, stimulus vignetting shows a strong and robust effect on measured orientation preference. Orientation is used as a model for understanding more complex feature processing in other cortical areas, and oriented V1-like receptive fields play an important role in successful computational models of vision. Yet, the computations that give rise to orientation tuning in V1 remain inadequately understood. Our results provide a framework for reinterpreting a wide-range of findings on orientation selectivity, measured with both fMRI in human subjects and in single units.

Materials and Methods

Theoretical model.

The model is based on the steerable pyramid (Simoncelli et al., 1992), a subband image transform that decomposes an image into orientation and spatial frequency channels (Figure 1). The pyramid simulates the responses of a large number of linear receptive fields (RFs), each of which computes a weighted sum of the stimulus image; the weights determine the orientation and spatial frequency tuning. RFs with the same orientation and spatial-frequency tuning, but shifted to different locations, are called a “channel”. In the model, the number of spatial frequency channels, orientation channels, and orientation bandwidth are model parameters that can each be chosen flexibly. We used an instantiation of the model with 6 orientations, which resulted in a bandwidth comparable to empirical tuning curves measured in primate electrophysiological recordings (Ringach et al., 2002). Using 4 or more bands with correspondingly broader or narrower tuning curves yielded similar results supporting the same conclusions. The number of spatial frequency channels was determined by the size of the input images and the spatial frequency bandwidth. We used images that were 768 x 1024 pixels, and a spatial frequency bandwidth of 0.5 octaves, and so the model had 15 different spatial frequency channels. The RFs cover all orientations and spatial frequencies evenly (i.e., the sum of the squares of the tuning curves is exactly equal to one for all orientations and spatial frequencies). For each orientation and spatial frequency, the pyramid includes RFs with two different phases, analogous to odd- and even-phase Gabor filters. The sum of the squares of the responses of two such RFs computes what has been called an “energy” response (Adelson and Bergen, 1985; Heeger, 1992) because it depends on the local spectral energy within a spatial region of the stimulus, for a particular orientation and spatial frequency. The energy responses are also arranged in channels. The linear RFs in the model are hypothesized to be a basis set of orientation and spatial frequency tuning curves of V1 neurons; any invertible linear transform of the basis set can be substituted (Freeman and Adelson, 1991; Simoncelli et al., 1992; Simoncelli, 1993) and different such transforms can be substituted at each spatial location, consistent with the diversity in the tuning properties of V1 neurons, without changing the nature of the representation.

We used the theoretical model to predict fMRI responses to a wide range of stimuli. We then chose stimuli that were predicted to have opposite patterns of orientation bias when measured with fMRI. Predicted responses were calculated by finding the scale with maximal responses, summing the energy responses for the channels at that scale, and averaging across stimulus grating phases (See: fMRI experiments, Stimuli).

fMRI experiments.

fMRI experiments were conducted at two sites. Experiments at 3T were conducted at NYU Center for Brain Imaging. Experiments at 7T were conducted at Functional Magnetic Resonance Imaging Core Facility at NIH. Methods were similar between the two sites except where noted.

Stimuli. Stimuli consisted of two gratings (a carrier and a modulator) multiplied by one another (Figure 3). The carrier grating consisted of a large, oriented sinusoidal Cartesian grating (spatial frequency, NYU: 0.5 cycles per degree; NIH: 1.4 cycles per degree) presented within an annulus (NYU: inner radius, 0.5°; outer radius 9.5°; NIH: inner radius 0.75°; outer radius 9°). The spatial phase of the carrier grating was randomized every 250 ms from a predefined set of 16 phases uniformly distributed between 0 and 2π . The orientation of the carrier grating cycled through 16 evenly-spaced angles between 0° and 180° (1.5 s per orientation). The carrier grating completed 10.5 cycles in each run. Each cycle was 24 s long, producing a run length of 4 min, 12 s. The orientation of the carrier grating cycled clockwise in half of the runs and counter-clockwise in the other half.

The modulator grating was a polar-transformed grating. In the “square wave modulator” experiments (at NYU and NIH), the modulator grating was square, with hard edges. In the “sinusoidal modulator” grating experiment (NIH), the modulator did not have any edges. On half the runs the modulator had a radial orientation, and an angular orientation in the other half. The modulator grating produced either a set of rings emanating from the fovea (radial orientation) or a set of inward-pointing wedges encircling the fovea (angular orientation). In the NIH experiments, the rings were scaled with eccentricity. On half the runs, the modulator grating was sine phase and cosine phase in the other half. We tested every combination of clockwise/counter-clockwise cycling carrier gratings, radial/angular modulator gratings, and sine/cosine phase modulator gratings. This produced a total of 8 conditions. Each of the 8 conditions was tested in two independent runs, producing a total of 16 runs per scanning session.

Because the only stimulus feature that varied within each run was the orientation of the carrier grating, any fMRI activity measured within each run could be attributed to either the orientation of the carrier grating, or an interaction between the orientation of the carrier grating and the static modulator grating. The modulator itself did not evoke any modulation in the fMRI responses because it did not change during the course of a scanning run.

Observers. Fourteen observers (8 females, aged 22-27 years) with normal or corrected-to-normal vision participated in the study. Observers provided written informed consent. The consent and experimental protocol were in compliance with the safety guidelines for MRI research, and were approved by both the University Committee on Activities Involving Human Subjects at New York University, and the Institutional Review Board at National Institutes of Health.

Each observer participated in multiple scanning sessions. Five observers (O1–O5) each participated in one session of the “square wave modulator” experiment on the 3T scanner at NYU. Nine observers (O6–O14) each participated in two sessions at NIH: one session for the “square wave modulator” experiment and one session for the “sine wave modulator” experiment. Data from a single session were discarded due to severe ghosting in the BOLD images. In addition, each of the fourteen observers participated in

one session to obtain a set of high-resolution anatomical volumes, used for cortical segmentation, and an additional session for retinotopic mapping.

Orientation mapping experiments. Stimuli were generated using Matlab (MathWorks, MA) and MGL (<http://gru.stanford.edu/doku.php/mgl/overview>) on a Macintosh computer. Stimuli were displayed via an LCD projector (3T scanner: Eiki LC-XG250; resolution: 1024 × 768 pixels; refresh rate: 60 Hz. 7T scanner: PLUS U2-1200; resolution: 800 × 600 pixels; refresh rate: 60 Hz) onto a back-projection screen in the bore of the magnet. Observers viewed the display through an angled mirror at a viewing distance of approximately 58 cm (field of view in the 3T at NYU: 31.6 deg × 23.7 deg; field of view in the 7T at NIH: 20.5 deg × 16.1 deg).

The novel stimuli consisted of two gratings (a carrier and a modulator) multiplied by one another (Figure 3). The carrier grating consisted of a large, oriented sinusoidal Cartesian grating presented within an annulus. The spatial phase of the carrier grating was randomized every 250 ms from a predefined set of 16 phases uniformly distributed between 0 and 2π . The orientation of the carrier grating cycled through 16 evenly-spaced angles between 0° and 180° (1.5 s per orientation). The carrier grating completed 10.5 cycles in each run. Each cycle was 24 s long, producing a run length of 4 min, 12 s. The orientation of the carrier grating cycled clockwise in half of the runs and counter-clockwise in the other half.

Behavioral task. Throughout each run, observers continuously performed a demanding two-interval, forced-choice task to maintain a consistent behavioral state and stable fixation, and to divert attention from the main experimental stimuli. In each trial of the task, the fixation cross (a 0.4-deg crosshair) dimmed twice (for 400 ms at a time), and the observer indicated with a button press the interval (1 or 2) in which it was darker. The observer had 1 s to respond and received feedback through a change in the color of the fixation cross (correct green, incorrect red). Each contrast decrement and response period was preceded by a 800 ms interval, such that each trial lasted 4.2 s. The fixation task was out of phase with the main experimental stimulus presentation. Contrast decrements were presented using an adaptive, 1-up-2-down staircase procedure (Levitt, 1971) to maintain performance at approximately 70% correct.

MRI data were acquired on two different scanners at two different sites. At NYU, data were acquired for observers O1-O5 on Siemens (Erlangen, Germany) 3T Allegra head-only scanner using a transmit head coil (NM-011, Nova Medical, Wakefield, MA), and an eight-channel phased-array surface receive coil (NMSC-071, Nova Medical). Functional imaging was conducted with 24 slices oriented perpendicular to the calcarine sulcus and positioned with the most posterior slice at the occipital pole (TR: 1500 ms; TE: 30 ms; FA: 72°; voxel size: 2 × 2 × 2.5 mm; grid size: 104 × 80 voxels).

At NIH, data were acquired from observers O6-O14 on a research-dedicated Siemens 7T Magnetom scanner using a 32-channel head coil, located in the Clinical Research Center on the National Institutes of Health campus (Bethesda, MD). Functional imaging was conducted with 54 slices oriented perpendicular to the calcarine sulcus covering

the poster half of the brain (TR: 1500 ms; TE 23 ms; FA: 55°; voxel size: 1.2 × 1.2 × 1.2 mm or 0.9 × 0.9 × 0.9 mm with 0% or 10% gap between slices, respectively; grid size: 160 x 160 voxels. Multiband factor 2 or 3, GRAPPA/iPAT factor 3).

For each observer, a high-resolution anatomy of the entire brain was acquired by co-registering and averaging between 2 and 8 T1-weighted anatomical volumes (magnetization-prepared rapid-acquisition gradient echo, or MPRAGE; TR: 2500 ms; TE: 3.93 ms; FA: 8°; voxel size: 1 × 1 × 1 mm; grid size: 256 × 256 voxels). The averaged anatomical volume was used for co-registration across scanning sessions and for gray-matter segmentation and cortical flattening. Functional scans were acquired using T2*-weighted, gradient recalled echo-planar imaging to measure blood oxygen level-dependent (BOLD) changes in image intensity (Ogawa et al., 1990). At NYU, an MPRAGE anatomical volume with the same slice prescription as the functional images (“inplane”) was also acquired at the beginning of each scanning session (TR: 1530 ms; TE: 3.8 ms; FA: 8°; voxel size: 1 × 1 × 2 mm with 0.5 mm gap between slices; grid size: 256 × 160 voxels). The inplane anatomical was aligned to the high-resolution anatomical volume using a robust image registration algorithm (Nestares and Heeger, 2000).

At NIH, prior to the first experimental functional run of each session, 30 volumes were acquired with identical scanning parameters and slice prescription as the subsequent functional runs, except for the phase encoding direction which was reversed. This scan was used for correcting spatial distortions (see next section).

fMRI time series preprocessing. The single reverse phase-encoded run was used to estimate the susceptibility-induced off-resonance field using a method similar to that described in (Andersson et al., 2003) as implemented in FSL (Smith et al., 2004). This estimate was then used to correct the spatial distortions in each subsequent run in the session. Data from the first half cycle (8 volumes) at the beginning of each functional run were discarded to minimize the effect of transient magnetic saturation and to allow hemodynamic response to reach steady state: the first half cycle. Functional data were compensated for head movement within and across runs (Nestares and Heeger, 2000), linearly detrended, and high-pass filtered (cutoff: 0.01 Hz) to remove low-frequency noise and drift (Smith et al., 1999). The time series for each voxel was divided by its mean to convert from arbitrary intensity units to percent change in image intensity. Time series data from each run were shifted back in time by 3 volumes (4.5 s) to compensate approximately for hemodynamic lag. Time series for the clockwise runs were time reversed to match the sequence of the counterclockwise runs. Time series were then averaged within modulator type (the matrix of 8 conditions: cosine and sine modulator crossed with clockwise and counterclockwise orientation). Each individual voxel’s time course was fitted to a cosine, and the orientation corresponding to the phase of the cosine was assigned as that voxel’s preferred orientation. It should be noted that shifting, inverting and averaging minimizes differences in hemodynamic responses between voxels, and results in a smooth time series that fits well to a cosine. Without this stage, the time series would have less Fourier power at the fundamental frequency

and much of the power would instead shift to harmonic frequencies (see (Pratte et al., 2016) for examples of these effects).

Retinotopic maps. Retinotopy was measured using non-periodic traveling bar stimuli and analyzed using the population receptive field (pRF) method (Dumoulin and Wandell, 2008). Bars were 3 deg wide and traversed the field of view in sweeps lasting 24 s. Eight different bar configurations (4 orientations and 2 traversal directions) were presented. The pRF of each voxel was estimated using standard fitting procedures (Dumoulin and Wandell, 2008), implemented in Matlab using mrTools (<http://gru.stanford.edu/doku.php/mrtools/overview>). Visual area boundaries were drawn by hand on the flat maps, following published conventions (Engel et al., 1997; Larsson and Heeger, 2006; Wandell et al., 2007).

pRF sampling analysis. To test the model's predictions for the responses of individual fMRI voxels, we estimated pRFs and sampled the model's output at the corresponding region of the visual field. Specifically, each voxel's Gaussian pRF was used to weight the model output for each grating orientation, averaged across grating phase and modulator phase, resulting in a single predicted response value for each orientation. We then fitted a cosine to the responses, with the phase reflecting the voxel's predicted preferred orientation. This was done separately for each of the two modulators. Finally, we compared the predicted preferred orientation to the measured preferred orientation (see Figure 6, top).

fMRI MVPA analysis. Classification was performed on fMRI time series, in a leave-one-run-out manner, with a naive Bayes classifier (using Matlab function 'classify' with 'diagLinear' option). First, for each voxel, we averaged across all 10 cycles within a run, yielding 16 time points per voxel, which reflected the amplitude of the response to each stimulus orientation. We refer to these 16 vectors across voxels as population responses, a single response vector per time point, or equivalently, per orientation.

For within-modulator decoding, the classifier was trained on data (response vectors and corresponding orientation labels) from 7 runs, and tested on the 8th run (leave-one-run-out decoding). If there were consistent differences in the responses to different orientations the classifier should decode correctly at an accuracy rate that is above chance level. We then cross-validated, by repeating this procedure 8 times, each time leaving out a different run for testing. This was repeated for both modulators, and results were averaged across them.

For across-modulator decoding, the classifier was trained on data from one modulator, then tested on data from the other modulator. We then swapped the training and testing datasets and repeated the analysis. For shifted across-modulator decoding, we first shifted the orientation labels for the angular modulator by 90 deg, then repeated the across-modulator decoding procedure, and averaged the accuracy. Statistical significance of decoding accuracy was determined by permuting the orientation labels 10,000 times and comparing the mean accuracy across subjects to the mean accuracy

without permuting. Decoding accuracy X was deemed significant at $p < a$ if the fraction of permutations that yielded decoding accuracies higher than X was smaller than a .

While decoding enabled us to measure orientation information present in population responses, it provided only an indirect view of the shift the responses undergo, as a result of changing the modulator. To visualize the responses in a more direct fashion, we used representational similarity analysis. Each condition (i.e., 16 orientations for each of 2 modulators), represented by a population response vector averaged across cycles and runs, was correlated with all other conditions, yielding a correlation matrix. The matrix was transformed into a dissimilarity matrix by subtracting from 1 (i.e., dissimilarity = $1-r$, where r is the correlation coefficient). We then averaged across subjects and used multidimensional scaling (Matlab function 'mdscale') to reduce the dimensionality to 2. Results are plotted both with the original orientation labels, and after shifting the angular modulator's orientation labels by 90 deg.

Bibliography

Adams DL, Sincich LC, Horton JC (2007) Complete pattern of ocular dominance columns in human primary visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 27:10391-10403.

Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. *JOSA A* 2:284-299.

Alink A, Krugliak A, Walther A, Kriegeskorte N (2013) fMRI orientation decoding in V1 does not require global maps or globally coherent orientation stimuli. *Frontiers in psychology* 4:493.

Alink A, Walther A, Krugliak A, Kriegeskorte N (2017) Local opposite orientation preferences in V1: fMRI sensitivity to fine-grained pattern information. *Scientific reports* 7:7128.

Andersson JL, Skare S, Ashburner J (2003) How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage* 20:870-888.

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. *Nature neuroscience* 8:541-542.

Carlson TA (2014) Orientation decoding in human visual cortex: new insights from an unbiased perspective. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 34:8373-8383.

Carter BE, Henning GB (1971) The detection of gratings in narrow-band visual noise. *The Journal of physiology* 219:355-365.

Cavanaugh JR, Bair W, Movshon JA (2002a) Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *Journal of neurophysiology* 88:2530-2546.

Cavanaugh JR, Bair W, Movshon JA (2002b) Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *Journal of neurophysiology* 88:2547-2556.

Cichy RM, Ramirez FM, Pantazis D (2015) Can visual information encoded in cortical columns be decoded from magnetoencephalography data in humans? *NeuroImage* 121:193-204.

Clifford CW, Mannion DJ (2015) Orientation decoding: Sense in spirals? *NeuroImage* 110:219-222.

Das A, Gilbert CD (1997) Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature* 387:594-598.

Dumoulin SO, Wandell BA (2008) Population receptive field estimates in human visual cortex. *NeuroImage* 39:647-660.

Engel SA, Glover GH, Wandell BA (1997) Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral cortex (New York, NY: 1991)* 7:181-192.

Freeman J, Heeger DJ, Merriam EP (2013) Coarse-scale biases for spirals and orientation in human visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 33:19695-19703.

- Freeman J, Brouwer GJ, Heeger DJ, Merriam EP (2011) Orientation decoding depends on maps, not columns. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 31:4792-4804.
- Freeman WT, Adelson EH (1991) The design and use of steerable filters. *IEEE Transactions on Pattern analysis and machine intelligence* 13:891-906.
- Furmanski CS, Engel SA (2000) An oblique effect in human primary visual cortex. *Nature neuroscience* 3:535-536.
- 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.
- Hallum LE, Movshon JA (2014) Surround suppression supports second-order feature encoding by macaque V1 and V2 neurons. *Vision research* 104:24-35.
- 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. *Nature neuroscience* 8:686-691.
- Heeger DJ (1992) Half-squaring in responses of cat striate cells. *Visual neuroscience* 9:427-443.
- Hubel DH, Wiesel TN (1963) Shape and arrangement of columns in cat's striate cortex. *The Journal of physiology* 165:559-568.
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nature neuroscience* 8:679-685.
- Larsson J, Heeger DJ (2006) Two retinotopic visual areas in human lateral occipital cortex. *Journal of Neuroscience* 26:13128-13142.
- Larsson J, Harrison C, Jackson J, Oh SM, Zeringyte V (2017) Spatial scale and distribution of neurovascular signals underlying decoding of orientation and eye of origin from fMRI data. *Journal of neurophysiology* 117:818-835.
- Levitt H (1971) Transformed up-down methods in psychoacoustics. *The Journal of the Acoustical society of America* 49:467-477.
- Ling S, Pratte MS, Tong F (2015) Attention alters orientation processing in the human lateral geniculate nucleus. *Nature neuroscience* 18:496-498.
- Maldonado PE, Godecke I, Gray CM, Bonhoeffer T (1997) Orientation selectivity in pinwheel centers in cat striate cortex. *Science (New York, NY)* 276:1551-1555.
- Mannion DJ, McDonald JS, Clifford CW (2010) Orientation anisotropies in human visual cortex. *Journal of neurophysiology* 103:3465-3471.
- Nestares O, Heeger DJ (2000) Robust multiresolution alignment of MRI brain volumes. *Magnetic Resonance in Medicine* 43:705-715.
- Ogawa S, Lee TM, Nayak AS, Glynn P (1990) Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magnetic resonance in medicine* 14:68-78.

- Ohki K, Chung S, Kara P, Hubener M, Bonhoeffer T, Reid RC (2006) Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442:925-928.
- Philips RT, Chakravarthy VS (2017) A global orientation map in the primary visual cortex (V1): Could a self organizing model reveal its hidden bias? *Frontiers in neural circuits* 10:109.
- Pratte MS, Sy JL, Swisher JD, Tong F (2016) Radial bias is not necessary for orientation decoding. *NeuroImage* 127:23-33.
- Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: diversity and laminar dependence. *Journal of Neuroscience* 22:5639-5651.
- 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.
- Sceniak MP, Hawken MJ, Shapley R (2001) Visual spatial characterization of macaque V1 neurons. *Journal of neurophysiology* 85:1873-1887.
- Sceniak MP, Ringach DL, Hawken MJ, Shapley R (1999) Contrast's effect on spatial summation by macaque V1 neurons. *Nature neuroscience* 2:733-739.
- Sengupta A, Yakupov R, Speck O, Pollmann S, Hanke M (2017) The effect of acquisition resolution on orientation decoding from V1 BOLD fMRI at 7T. *NeuroImage* 148:64-76.
- Simoncelli EP (1993) Distributed representation and analysis of visual motion. In: Massachusetts Institute of Technology.
- Simoncelli EP, Freeman WT, Adelson EH, Heeger DJ (1992) Shiftable multiscale transforms. *IEEE transactions on Information Theory* 38:587-607.
- Smith AM, Lewis BK, Ruttimann UE, Frank QY, Sinnwell TM, Yang Y, Duyn JH, Frank JA (1999) Investigation of low frequency drift in fMRI signal. *NeuroImage* 9:526-533.
- Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE (2004) Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage* 23:S208-S219.
- Sun P, Gardner JL, Costagli M, Ueno K, Waggoner RA, Tanaka K, Cheng K (2013) Demonstration of tuning to stimulus orientation in the human visual cortex: a high-resolution fMRI study with a novel continuous and periodic stimulation paradigm. *Cerebral cortex (New York, NY : 1991)* 23:1618-1629.
- Sun W, Tan Z, Mensh BD, Ji N (2016) Thalamus provides layer 4 of primary visual cortex with orientation-and direction-tuned inputs. *Nature neuroscience* 19:308-315.
- 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. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 30:325-330.
- Tanaka H, Ohzawa I (2009) Surround suppression of V1 neurons mediates orientation-based representation of high-order visual features. *Journal of neurophysiology* 101:1444-1462.
- Tang J, Jimenez SCA, Chakraborty S, Schultz SR (2016) Visual receptive field properties of neurons in the mouse lateral geniculate nucleus. *PLoS one* 11:e0146017.

Tang S, Lee TS, Li M, Zhang Y, Xu Y, Liu F, Teo B, Jiang H (2018) Complex Pattern Selectivity in Macaque Primary Visual Cortex Revealed by Large-Scale Two-Photon Imaging. *Current Biology* 28:38-48. e33.

Van den Bergh G, Zhang B, Arckens L, Chino YM (2010) Receptive-field properties of V1 and V2 neurons in mice and macaque monkeys. *Journal of Comparative Neurology* 518:2051-2070.

Walker GA, Ohzawa I, Freeman RD (1999) Asymmetric suppression outside the classical receptive field of the visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 19:10536-10553.

Wandell BA, Dumoulin SO, Brewer AA (2007) Visual field maps in human cortex. *Neuron* 56:366-383.

Wardle SG, Ritchie JB, Seymour K, Carlson TA (2017) Edge-Related Activity Is Not Necessary to Explain Orientation Decoding in Human Visual Cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 37:1187-1196.

Warren S, Yacoub E, Ghose G (2014) Featural and temporal attention selectively enhance task-appropriate representations in human V1. *Nature communications* 5:5643.

Yacoub E, Harel N, Ugurbil K (2008) High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences of the United States of America* 105:10607-10612.

Zenger-Landolt B, Heeger DJ (2003) Response suppression in v1 agrees with psychophysics of surround masking. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 23:6884-6893.