

Natural scene sampling reveals reliable coarse-scale orientation tuning in human V1

Zvi N. Roth^{1*}, Kendrick Kay^{2,#}, Elisha P. Merriam^{1,#}

*Corresponding author: zvi.roth@nih.gov

#These authors jointly supervised this work

¹Laboratory of Brain and Cognition, National Institute of Mental Health, NIH, Bethesda, MD, USA

²Center for Magnetic Resonance Research (CMRR), Department of Radiology, University of Minnesota, Minneapolis, MN, USA

Short Title: Coarse-scale orientation tuning in human V1

Abstract:

Orientation selectivity in primate visual cortex is organized into cortical columns. Since cortical columns are at a finer spatial scale than the sampling resolution of standard BOLD fMRI measurements, analysis approaches have been proposed to peer past these spatial resolution limitations. It was recently found that these methods are predominantly sensitive to stimulus vignetting - a form of selectivity arising from an interaction of the oriented stimulus with the aperture edge. Beyond vignetting, it is not clear whether orientation-selective neural responses are detectable in BOLD measurements. Here, we leverage a dataset of visual cortical responses measured using high-field 7T fMRI. Fitting these responses using image-computable models, we compensate for vignetting and nonetheless find reliable tuning for orientation. Results further reveal a coarse-scale map of orientation preference that may constitute the neural basis for known perceptual anisotropies. These findings settle a long-standing debate in human neuroscience, and provide insights into functional organization principles of visual cortex.

Keywords: visual cortex, orientation, fMRI, radial bias, cardinal bias

Introduction:

Neurons in human visual cortex are organized according to their functional selectivity. A number of stimulus features are organized in coarse-scale cortical maps. For example, retinotopic visual cortex is organized as a log-transform of the visual field, with the polar dimensions of visual space (angle and eccentricity) corresponding to the Cartesian dimensions of the cortical surface^{1, 2, 3}. Additionally, receptive field size^{4, 5} and preferred spatial frequency⁶ increase monotonically with visual eccentricity. Selectivity to other visual features, such as ocular dominance and temporal frequency, are organized at a fine spatial scale, often forming columns through the cortical depth^{7, 8, 9, 10}. As in primates and carnivores¹¹, orientation selectivity in humans has been shown to also be organized at a fine spatial scale, in cortical columns that are approximately 0.7-0.8 mm in width across the cortical surface of primary visual cortex¹².

In addition to the fine-scale columnar architecture, fMRI studies have offered evidence for a coarse-scale organization of orientation preference^{13, 14, 15, 16}. Because of the relatively low spatial resolution of standard fMRI measurements (typically around 2 x 2 x 2 mm), these studies do not reveal the fine-scale columnar architecture. Rather, these studies have leveraged the broad spatial coverage afforded by fMRI to reveal a radial bias of orientation preference: voxels respond more strongly to orientations that point from the receptive field center toward fixation. This radial bias was originally assumed to reflect a physiological map of orientation selectivity¹³, analogous to maps of receptive field location and spatial frequency selectivity. However, this radial bias map, it turns out, does not necessarily reflect solely a physiological map. Instead, it was shown to likely be, in large part, the result of stimulus vignetting, an interaction between the edges of the stimulus aperture and the spatial frequency envelope of the stimulus^{17, 18}. The principle underlying vignetting is that the Fourier spectrum changes in the vicinity of a change in contrast, such as an edge. In typical orientation mapping experiments, the assumption is that each condition contains a single orientation. But because of vignetting, this isn't true: different regions in the image contain different orientations and different Fourier power. Specifically, at the stimulus edge, there is more power for the radial orientation than for other orientations. These observations challenge the interpretation of a large body of studies over the past 20 years that were presumed to measure orientation-selective responses in humans¹⁹.

If previous fMRI studies do not provide clear evidence regarding a physiological map of orientation selectivity, how is orientation actually represented in human visual cortex? Does the human brain contain a map for orientation selectivity at a coarse spatial scale, distinct from columnar architecture? Alternatively, is the presumed fine-scale columnar map the sole organizational principle for orientation in visual cortex? The answer to these questions has been obfuscated by stimulus vignetting: if there is indeed a coarse-scale map for orientation, it may be entirely eclipsed by stimulus vignetting.

To overcome these challenges, we apply a computational framework for studying orientation selectivity that explicitly models the effects of stimulus vignetting in order to access orientation-selective signals that would be otherwise obscured. To this end, we leverage a massive 7T fMRI dataset, the Natural Scenes Dataset (NSD), consisting of extensive sampling of responses to natural scene stimuli in a small number of intensively-studied participants^{20, 21, 22}. The large number of measurements and unique stimuli, combined with the high signal-to-noise ratio (SNR) of the fMRI measurements, enables us to robustly fit models that include dozens of parameters per voxel.

Image-computable models have been used to study a wide range of questions in visual neuroscience^{17, 23, 24, 25}. To assess orientation selectivity, our modeling approach exploits two

image-computable models based on the steerable pyramid²⁶. The constrained model, which includes both visual field position tuning and spatial frequency tuning, but pools equally across orientation-selective filters, is sensitive to the effects of stimulus vignetting¹⁷. This model is based on the model used previously to demonstrate vignetting¹⁷, and it can fully explain responses modulated by total Fourier power. However, because the orientation filters at each level are pooled before the filter responses are computed, the model cannot capture any information about orientation. In other words, any apparent orientation selectivity in the model output is entirely due to stimulus vignetting. The full model, on the other hand, which allows unequal contributions from orientation-selective filters, is sensitive to orientation-selectivity beyond the effects of stimulus vignetting (Fig 1). Combined, the pair of models enables us to assess voxel-wise orientation selectivity while simultaneously accounting for the impact of stimulus vignetting. The modeling results reveal the existence of physiological maps for orientation preference.

Results:

Both constrained and full models explain roughly similar amounts of variance

The NSD dataset contains measurements of 7T BOLD fMRI responses from 8 participants who each viewed 9,000–10,000 distinct color natural scenes (22,000–30,000 trials) over the course of 30–40 scan sessions²². We fit two models characterizing V1 voxel responses to the natural scene stimuli. Both models fit voxel responses as a weighted sum of steerable pyramid filters. The constrained model pools across orientation, and is therefore effectively composed of a range of spatial frequency filters across the visual field. In contrast, the full model includes flexible weights for both spatial frequency and orientation tuned filters (Fig 1). In order to evaluate how well each model fits the data, we cross-validated both models and measured the variance explained by each model on out-of-sample data.




Fig 1. Analysis pipeline for a single voxel. Top, full model. Bottom, constrained model. Natural scene images were converted to grayscale and passed as input to a steerable pyramid (shown here with only 4 orientations and 3 spatial frequencies for visualization purposes). Each

filter yielded an energy response that was sampled by the voxel pRF, resulting in a scalar output for each filter. Linear regression was performed on the response amplitudes observed for each voxel with filter output values as predictors. This procedure yielded a set of model weights, which were subsequently used to simulate responses to a range of gratings in order to determine the voxel's preferred orientation and spatial frequency. The constrained model (bottom) was identical to the full model except for an additional step of summing model outputs across orientations. As a result, the constrained model involves a weight for each spatial frequency filter but enforces equal contribution across orientation. Example images shown here were created by the authors for illustration only and were not used in the study.

The constrained and the full models explained a similar portion of variance in the BOLD measurements (mean R^2 : full 0.0311, constrained 0.0310)(Fig 2A). If BOLD activity reflects orientation selectivity, we would expect the constrained model to be unable to account for modulations driven by local orientation differences between images. The full model, however, includes orientation tuning, and therefore should be able to account for this additional response variability, assuming that the parameters that characterize orientation tuning can be reliably estimated. But we found that both models performed comparably, explaining similar amounts of variance, with the constrained model slightly but significantly outperforming the full model (median R^2 : full 0.0204, constrained 0.0212; two-sided Wilcoxon signed-rank test p -value $< 10^{-10}$). Moreover, cross-validated R^2 was highly correlated between the two models ($r=0.9835$, $p < 10^{-30}$), likely reflecting gross differences in signal-to-noise ratio across voxels (Fig 2B).

At first glance, these results suggest that there is no reliable orientation tuning to be modeled. However, since the full model has many more free parameters than the constrained model (57 vs. 8 parameters), we expect that, for voxels with low SNR, the full model would likely result in overfitting. For voxels with high SNR, on the other hand, the full model may be able to capture orientation tuning and lead to higher cross-validated R^2 values. Therefore, inspecting model performance only using summary statistics (e.g. median or mean) does not provide a complete picture, and it is necessary to analyze how model performance varies across voxels.




Fig 2. Both constrained and full models fit voxel responses. A. Distribution of R^2 values for full (pink) and constrained (blue) models. Distributions for both models were similar, with the constrained model R^2 median slightly but significantly higher ($p=7.4 \times 10^{-11}$, two-sided Wilcoxon rank sum test). B. R^2 values for full model as a function of R^2 values for constrained model. C. R^2 values for constrained model as a function of pRF R^2 . D. R^2 values for full model as a function of pRF R^2 . E. R^2 values for full model minus R^2 for constrained model as function of pRF R^2 . F. R^2 values for full model minus R^2 for constrained model as function of pRF eccentricity. G. pRF R^2 values as function of pRF eccentricity. Solid lines in C-G indicate a running mean computed over 20 bins.

Model performance depends on voxel SNR

For each individual voxel, model filter outputs were sampled using a population receptive field (pRF) estimated for that voxel from an independent pRF-mapping experiment that was conducted as part of the NSD dataset. In both the constrained and full models, the estimated pRF determined the portion of stimulus from which the model output is sampled (see *Methods: pRF sampling*). Hence, any error in estimating the size or location of the pRF would propagate forward, adversely affecting the model fits. Specifically, if the pRF estimate is inaccurate, the model would attempt to explain BOLD activity based on the portion of the natural scene image dictated by the pRF estimate, while the voxel would in fact be driven by another portion of the image. Therefore, it is likely that model performance depends on the quality of the pRF estimate.

Sorting voxels according to pRF R^2 , we found that this was indeed the case. R^2 of both models were correlated with pRF R^2 (Constrained: $r=0.5859$, $p<10^{-30}$; Full: $r=0.5728$, $p<10^{-30}$). Our interpretation is that pRF R^2 is a good proxy for voxelwise SNR, and that voxels with high SNR will tend to have both high-quality pRF estimates and good model performance on the natural scene responses (Fig 2C,D). We furthermore found that the amount of additional variance that the full model explained beyond the constrained model was also correlated with voxel pRF R^2 ($r=0.2418$, $p<10^{-30}$) (Fig 2E). This result implies that in cases of low SNR, including orientation tuning in an encoding model will likely result in overfitting, and that doing so is unlikely to reveal reliable orientation selectivity. However, for high SNR voxels in the NSD dataset used here, we are able to estimate reliable orientation selectivity that improves generalization on out-of-sample data. Additional analyses (Fig 2G) indicate that voxels with high SNR tend to lie away from the fovea, and it is for these more peripheral voxels that we can reliably estimate orientation selectivity.

Coarse-scale map of orientation selectivity

Explicitly modeling voxel responses enabled us to capture robust orientation selectivity that is not due to stimulus vignetting. What is the source of this orientation tuning? Although the voxel size in the NSD dataset ($1.8 \times 1.8 \times 1.8$ mm) is much larger than the size of orientation columns, fMRI studies using multivariate decoding methods have suggested that even with large fMRI voxels ($3 \times 3 \times 3$ mm), V1 voxel responses might exhibit robust orientation biases originating from random sampling of cortical columns^{27, 28, 29}. Therefore, we ask: is orientation preference scattered in a salt-and-pepper fashion, suggesting a random bias from orientation column sampling, or is orientation preference organized in a systematic map across the cortical surface?

200 To assess the potential organization of orientation selectivity, we plotted voxel orientation
201 preference in visual space (Fig 3A). We observe a clear coarse-scale map of orientation:
202 preferred orientation varies smoothly across visual angle. This observation implies that voxels
203 with nearby pRF locations have similar orientation preference, unlike the expectation based on
204 salt-and-pepper organization. Indeed, when we visualize cortical maps of estimated orientation
205 preference, we see in V1 a clear progression of orientation that mimics the well-known
206 organization of visual field angle (Fig 3B).
207




Fig 3. Orientation preference changes smoothly across visual space and V1. A, Orientation preferences plotted in visual space. Each line represents a single voxel and is positioned at the voxel's pRF center. Hue and orientation of the line indicate preferred orientation. Line length, width, and scale reflect the amount of variance (R^2) explained by the constrained model. Solid square at ± 4.2 deg indicates the size of the natural-scene stimuli. See also individual subject plots in Supplemental Fig S1. B, Orientation and pRF polar angle maps overlaid on left and right inflated 'fsaverage' surfaces. For angle map and unthresholded orientation map, all vertices in V1, V2, V3, and V4 are plotted. For the thresholded orientation map, only vertices with the top 50% full model R^2 are plotted. See also individual subject maps in Supplemental Fig S2.

The gradual progression of orientation preference across the visual field and cortical surface provides evidence against a random-sampling bias of orientation columns. Instead, this observation reveals an organizational principle for visual cortex, namely, a coarse-scale spatial map for orientation preference.

Radial map of orientation selectivity

The orientation map appears to resemble a radial map (Fig 3A), but quantitative analysis is necessary to test this hypothesis. We quantified the similarity of the orientation map to a radial map by computing the angular deviation of each voxel's orientation preference from the

preference predicted by a perfectly radial map (Fig 4, right). Deviation from radial as a function of pRF eccentricity resembled an inverted U (Fig 5A): deviation was maximal at the fovea and at high (>5 deg) eccentricities, and lowest at intermediate (2-4 deg) eccentricities. We also compared the measured orientation map to two alternative possibilities: a vertical map and a cardinal map (Fig 4; Fig 5A). Deviation from the ideal radial map was lower than for either the ideal vertical or the ideal cardinal maps (Fig 5B, C), indicating that orientation selectivity in V1 is approximately organized in a radial map.




Fig 4. Schematic of three ideal orientation maps. Ideal maps of orientation preference plotted in visual space. Each line represents a single ideal voxel and is positioned at the voxel's pRF center. Orientation of the line indicates preferred orientation. Left, vertical map. Center, cardinal map. Right, radial map.

Deviation from radial was not uniform across pRF polar angle (Fig 5A, right), but was lowest at the horizontal and vertical meridians (Fig 5A, orange), similar to the deviation from cardinal (Fig 5A, green). Confirming this observation, deviation from radial was linearly correlated with angular distance from the closest meridian ($r=0.141$, $p<10^{-38}$). However, voxel-wise SNR is a potential limiting factor in this analysis, since voxel-wise SNR (as indexed by pRF R^2) correlates slightly negatively with distance from the meridians ($r=-0.055$, $p<10^{-6}$). Nevertheless, when controlling for the effect of pRF R^2 , the partial correlation between radial deviation and distance from meridian remained strong ($r=0.133$, $p<10^{-34}$). We interpret this to mean that modulation of the radial bias by distance from the meridian is a feature of the coarse-scale orientation map, rather than a trivial result of anisotropies in pRF SNR across the visual field.




Fig 5. Coarse-scale bias is mostly radial in organization. A. Deviation of preferred orientation from ideal vertical (blue), cardinal (green), and radial (orange) maps, as function of pRF eccentricity (left) and pRF angle (right). B. Deviation from vertical minus deviation from radial, as function of pRF eccentricity (left) and pRF angle (right). Gray lines, individual subjects. Orange points, significantly lower deviation from radial than from vertical ($p < 0.05$, one-way t-test, 7 degrees of freedom, no correction for multiple comparisons). C. Deviation from cardinal minus deviation from radial as function of pRF eccentricity (left) and pRF angle (right). Gray lines, individual subjects. Orange points, significantly lower deviation from radial than from cardinal ($p < 0.05$, one-way t-test, 7 degrees of freedom, no correction for multiple comparisons). D. Deviation from radial as function of angular distance from closest meridian. E. pRF R^2 as function of angular distance from meridian.

Ideal radial and vertical maps are identical at the vertical meridian, and differ maximally at the horizontal meridian (Fig 4). Therefore, comparing the empirical orientation map with radial and vertical maps around the vertical meridian will not be particularly informative. Around the

horizontal meridian, however, radial and vertical maps differ maximally. We therefore expected deviation from vertical to be roughly equal to deviation from radial around the vertical meridian. Around the horizontal meridian, on the other hand, we expected deviation from radial to differ maximally from deviation from vertical. Consistent with this expectation, deviation from radial was significantly lower than deviation from vertical around the horizontal meridian, but not at the vertical meridian (Fig 5B, right).

We found that radial and cardinal maps are identical at the horizontal and vertical meridians, and differ maximally at oblique (diagonal) angles (Fig 4). Therefore, deviation from radial and cardinal should differ maximally around oblique angles and minimally around the meridians. Again, this was indeed the case: deviation from radial was significantly lower than from cardinal only at oblique angles of the visual field (Fig 5C, right).

Deviation from radial was significantly lower than deviation from cardinal at intermediate eccentricities (Fig 5C, left). This may be related to lower pRF R^2 around the fovea (Fig 2G). Therefore, we cannot determine whether the radial bias is weakest at low eccentricities, or whether the larger deviation is entirely due to less accurate pRF estimates, while the strength of the radial map is in fact constant across eccentricities.

We conclude that the radial map is strongest around the meridians (i.e. at cardinal angles), and is weakest around oblique angles.

Controlling for analysis pipeline

We have assumed thus far that orientation selectivity in the full model reflects variance that cannot be explained by the constrained model. We wondered if this result could somehow be an artifact of our modeling procedures, given the complexity of the model fitting pipeline. For example, it is conceivable that orientation selectivity in the full model (i.e. the weights the model assigns to different orientation filters) in fact reflects variance that *can* be explained by the constrained model as well, but with other combinations of regressors. One example of this type of confound is stimulus vignetting¹⁷, where apparent orientation selectivity in fact reflects spatial frequency tuning and not orientation tuning. Perhaps stimulus vignetting or other sources of response variance that do not genuinely reflect orientation tuning effectively masquerade as the orientation selectivity observed in the full model.

To test this possibility directly, we first regressed out all variance that was explained by the constrained model from the experimental data. We then fit the full model to these residual data. If orientation selectivity of the full model reflects variance that can be explained by a combination of parameters in the constrained model, we would expect that regressing out that variance would leave the full model with no orientation-selective variance left to fit. If, on the other hand, orientation selectivity in the full model reflects only variance that cannot be explained by the constrained model (i.e. true orientation tuning) then regressing out variance explained by the constrained model should have no effect on the full model orientation selectivity.

Consistent with the second scenario, the map derived from analyzing the residuals (Fig 6A-B, left) was nearly identical to the original map (Fig 3). Complementarily, fitting the full model to the variance explained by the constrained model (i.e., the output predicted by the constrained model) resulted in a random map of orientation preference (Fig 6A-B, right). This indicates that the orientation selectivity in the full model cannot be explained by the constrained model.




Fig 6. Control analyses demonstrate that the full model indeed captures orientation selectivity that cannot be derived from the constrained model. A, Schematic Venn diagram illustrating the control analyses. First, the constrained model was fit. The full model was then fit either to the residual (left) or the output (right) of the constrained model. B, Orientation map in visual space derived from the two analyses. Same format as Fig 3A. The orientation map on the left is nearly identical to the map estimates from the original data (see Fig 3A), indicating that the full model capitalizes on variance not modeled by the constrained model. In contrast, the orientation map on the right has no discernable organization, demonstrating that the constrained model cannot explain variance reflecting orientation selectivity.

Discussion:

Summary

By analyzing a unique, extensive visual fMRI dataset using a model-based framework, we obtained robust estimates of orientation selectivity in human visual cortex. Unlike previous fMRI studies in which measurements of orientation selectivity were potentially confounded by interactions between the oriented stimuli and the stimulus aperture, the orientation selectivity that we report here cannot be attributed to stimulus vignetting. We uncovered a radial bias of orientation selectivity that is coarse-scale and widespread throughout the extent of V1. The radial bias that we describe here is distinct from the spatial pattern expected by the sampling of cortical columns. This coarse-scale orientation bias may be a fundamental organizational principle of human V1, providing a physiological basis for well-documented behavioral biases in orientation judgements.

Multiple scales of stimulus representation in human visual cortex

The human brain likely contains multiple representations of stimulus orientation at different spatial scales, co-existing within the same cortical visual area but arising from distinct neural computations. The first, and most familiar, is the fine-scale, columnar organization. The orientations of visual features are represented in an orderly pinwheel-like progression within each hypercolumn across the cortical surface^{30, 31, 32}. While orientation columns have not been measured directly in humans using electrophysiological methods, based on postmortem measurements of ocular dominance columns in humans^{33, 34}, orientation columns are likely to be less than a millimeter wide along the cortical surface. In addition to this fine-scale columnar structure, a number of fMRI studies have also reported a second, coarse-scale map-like organization for orientation selectivity that is at the scale of the retinotopic organization of V1^{13, 14}, spanning tens of centimeters, orders of magnitude larger than cortical columns. For the purposes of this discussion, we define a pattern of cortical activity to be 'fine-scale' if it has features that are smaller than the point-spread function of a conventional fMRI voxel (~2×2×2mm). According to this definition, the radial bias that we report here is clearly coarse-scale.

The presence of both fine- and coarse-scale patterns of orientation selectivity raises two fundamental questions. The first regards the scale of information leveraged by multi-voxel pattern analysis (MVPA) methods to decode orientation. In a landmark study, Kamitani and Tong²⁷ demonstrated that it is possible to use a linear classifier to decode the orientation of a grating presented to the subject on an individual trial. Cortical columns are irregularly organized with respect to the rectilinear voxel grid, which, it was posited, could lead to small biases in voxel responses that are decodable with MVPA. The conjecture that MVPA methods are sensitive to fine-scale signals had a profound impact on fMRI research, well beyond visual neuroscience, because it implied the feasibility of studying neural representations in the human brain that are instantiated at a spatial scale smaller than an fMRI voxel²⁸. However, it has been surprisingly difficult to fully support or fully refute this claim, especially in light of an alternative account suggesting that MVPA is primarily sensitive to the coarse-scale orientation bias¹³. This uncertainty has engendered ongoing and unresolved debate^{35, 36, 37, 38, 39, 40, 41, 42, 43, 44}. The current study does not directly bear on this debate. Here, we quantified orientation information in fMRI BOLD responses using a neurally-inspired image-computable model. While our results reaffirm the existence of a prominent coarse-scale organization for orientation (i.e., a radial bias), we did not use MVPA, nor did we explicitly model the scale of orientation information with a spatial model of the cortical surface. One fruitful approach may be to build spatial models of V1^{3, 45, 46}, which could provide a computational platform for testing hypotheses regarding spatial

scale of information and explicitly test the sensitivity of MVPA methods to different spatial scales of orientation information.

The second question concerns the neural response properties that give rise to the coarse-scale orientation bias. While the cortical architecture giving rise to fine-scale orientation columns is relatively well understood^{47, 48, 49}, much less is known about the origin of the coarse-scale bias. It was initially presumed that the coarse-scale bias stems from a pattern of cortical organization in which neurons with similar patterns of orientation selectivity form gradients over large swaths of cortex¹³. This idea was challenged by the insightful work of Carlson¹⁸, who demonstrated through simulation that the radial bias could arise from properties of the stimulus aperture or 'vignette', rather than from the pattern of orientation selectivity in the brain. This proposal suggested that stimuli with the same underlying orientation could be shown through different vignettes and this could reverse the observed orientation selectivity. These predictions have been confirmed empirically¹⁷, suggesting that orientation selectivity measured with fMRI is indeed influenced to some degree by stimulus vignetting.

Is stimulus vignetting the primary driver of the coarse-scale orientation biases reported in previous fMRI studies? The magnitude of stimulus vignetting reported in Roth, Heeger¹⁷ was roughly commensurate with the magnitude of the radial bias reported in previous studies of orientation selectivity¹³, suggesting that the radial bias could, in theory, be due to stimulus vignetting. fMRI studies have reported a radial bias using a wide range of stimulus parameters, i.e. different spatial frequencies and aperture sizes. Further theoretical simulations (Supplementary Fig 4) suggest that each of these configurations could produce stimulus vignetting, depending on the location of the edge in the retinotopic map, the spatial frequency channels that contribute most to the voxel's response, and the form of response normalization assumed. On the other hand, there have been reports of orientation decoding away from the stimulus edge, suggesting orientation information in fMRI measurements that are not due to vignetting (Kamitani and Tong²⁷ Supplementary Figure 3; Wardle, Ritchie⁴¹). These considerations have, up to now, left the field at an impasse, with no clear way to characterize true coarse-scale orientation selectivity in the face of a potential stimulus confound. Here, we turned to natural scene stimuli and an image-computable modelling approach to overcome these challenges. By explicitly modelling the spatial frequency and orientation content of each image, we account for the presence of a stimulus aperture, enabling the accurate characterization of orientation selectivity in human visual cortex.

The characterization of a coarse-scale orientation bias does not preclude a contribution from cortical columns. It is likely the case that orientation-selective signals measured with fMRI are multi-scale^{38, 42}, reflecting a contribution of both fine- and coarse-scale orientation signals, with the relative contribution dependent on critical experimental parameters, such as fMRI voxel size, acquisition method (gradient echo, spin echo, or VASO⁵⁰), proximity of a voxel to veins^{51, 52}, as well the stimulus protocol itself⁵³. These considerations imply that various reports of orientation selectivity over the years may have differentially emphasized fine- and/or coarse-scale components of the signal depending on which combination of these parameters were used. Further work, perhaps using an extension of the image-computable model that we have developed here, may be able to tease apart the relative contribution of each of these orientation-selective signals.

Multiple coarse-scale biases

Here we report a radial bias of orientation preference in visual cortex, extending a number of prior BOLD fMRI studies^{13, 14, 15, 16, 54}. But while the radial bias in previous fMRI studies may have been entirely a result of stimulus vignetting, the radial bias we report here is not

explainable by the constrained model (Fig 6), and therefore reflects a coarse-scale orientation map distinct from stimulus vignetting. In addition to a radial bias that was most pronounced in the periphery, a previous study identified a vertical bias at mid-eccentricities, closer to the fovea¹⁶. We tested for such a vertical bias, but found no evidence to support this possibility. At intermediate eccentricities, orientation preference was significantly closer to radial than to vertical, although close to fixation, the radial preference was not significant (Fig 5B). This may be a result of lower SNR around fixation, as reflected in the lower pRF R^2 values (Fig 2G). Therefore, we cannot definitively rule out a vertical bias at fixation, as reported by Freeman, Heeger¹⁶.

A cardinal bias has been identified previously in human visual cortex⁵⁵, in other primates^{56, 57, 58}, and in carnivores such as ferrets^{59, 60, 61} and cats^{62, 63, 64, 65}. The cardinal bias is typically described as a stronger response to vertical and horizontal orientations compared to oblique orientations. To enable a comparison to the radial map, we defined a cardinal map as a stronger response to the meridian closest to the pRF center, which is radial only for pRFs along the meridian (Fig 4, center). This definition describes an orientation map, relating orientation preference to retinotopic preference. The cardinal map as we defined it entails a cardinal bias, since when averaging across all of V1⁵⁵, cardinal orientations evoke the strongest responses.

Although the orientation map was closer to a radial map than to the cardinal map, we did find that the radial bias was strongest around the vertical and horizontal meridians. This pattern can be described as cardinal modulation of a radial map, and is consistent with previous fMRI findings¹⁵.

What physiological factors underly the cardinal modulation of the radial map? It is possible that the radial bias is stronger around the meridians because of stronger SNR in those regions, or because of small artifacts away from the meridians caused by cortical unfolding and veins^{66, 67, 68}. However, we believe such artifacts and SNR differences should manifest similarly in the pRF data, yet we found that the lower pRF R^2 (Fig 5E) could not fully explain the weaker radial bias away from the meridians. Similarly, a fine-scale columnar bias could potentially cause voxel preference to deviate from the radial orientation, but we would expect such an effect to take place uniformly across all polar angles. Instead, we believe it more likely that the cardinal modulation reflects the true nature of the radial bias: it is possible there are two co-existing biases, a cardinal and a radial, or that neurons around the meridians show a stronger preference for the radial orientation.

Source of coarse-scale orientation bias

How does the coarse-scale orientation map form? The mechanism is likely related to the source of orientation selectivity itself, which is still debated. When orientation-selective neurons were first discovered in cat visual cortex, their tuning properties were proposed to arise from convergence of center-surround neurons in lateral geniculate nucleus (LGN) that were themselves not orientation selective⁶⁹. Local interactions between V1 neurons have also been shown to amplify orientation selectivity⁷⁰, suggesting that orientation selectivity arises from both the convergence of feedforward input and local circuit interactions. But it has also been suggested that orientation selectivity is computed earlier in the visual pathway, and that it is to some degree inherited by V1 neurons. Multiple lines of evidence point towards orientation selectivity being present already in some LGN neurons^{71, 72, 73, 74, 75, 76} and even in retinal ganglion cells^{77, 78, 79, 80, 81}, raising the possibility that orientation selectivity in V1 reflects computations at earlier processing stages. Consistent with this possibility, retinal size relative to V1 size predicts across species whether orientation preference will be arranged in cortical columns or scattered in a salt-and-pepper fashion⁸². The coarse-scale orientation map may be

the result of the same mechanisms that form orientation selectivity, or it may involve other unique factors. Determining the source of the coarse-scale orientation bias, and whether it differs from sources of fine-scale selectivity, will require additional research involving measurements at a range of spatial scales.

A number of distinct neurophysiological mechanisms could, in theory, give rise to the coarse-scale bias uncovered here. For example, the radial bias could reflect a higher number of neurons preferring the radial orientation. Alternatively, the radial bias could reflect a higher firing rate for neurons preferring the radial orientation. A third possibility is that neurons preferring the radial orientation could have a narrower tuning bandwidth. All three of these scenarios presume a higher mean population firing rate in response to radial stimulus orientations, which would presumably translate to larger BOLD fMRI responses^{83, 84, 85, 86}. Since BOLD fMRI measures a hemodynamic signal and is an indirect measure of neural activity, it is also possible that the orientation tuning we measured here reflects synaptic inputs from either feedforward or feedback projections, or local field potentials^{87, 88}. Future electrophysiology studies in humans and non-human primates may shed more light on the relationship between the coarse-scale orientation bias and the underlying neurophysiology.

An alternate possibility is that coarse-scale orientation biases are an emergent property of visual cortex. When the visual system is modeled with large, unconstrained models, certain anisotropies and biases emerge, including a cardinal bias⁸⁹ and a radial bias^{90, 91}. These biases are due to statistics of the images used to train the models. Image statistics may underlie orientation biases in human visual cortex as well^{92, 93}, although it is unclear whether biases evident in neural networks account for similar biases we have observed in human visual cortex.

Natural scenes vs. oriented gratings

Most prior fMRI studies of orientation selectivity have relied on oriented gratings. This approach stems from a long and successful history in visual neurophysiology dating back to Enroth-Cugell and Robson⁹⁴. Such ‘synthetic’ grating stimuli are optimal for driving individual V1 neurons because they can be presented at full-contrast, and because the parameters of the stimuli (size, position, and spatial frequency) can be carefully tailored to the individual neuron being recorded. However, such gratings may be less appropriate when studying large neural populations, as with fMRI, since a single voxel reflects the pooled activity of many neurons with a wide range of selectivities, and therefore no single grating will be optimal for every neuron contributing to the voxel’s response¹⁹. Natural scene stimuli are inherently broadband along multiple dimensions, and hence may be more appropriate for studying population responses. However, natural scene stimuli do have, on average, lower contrast than gratings, and will not drive individual V1 neurons at their maximal firing rates. Thus, natural scenes might not be the most efficient set of stimuli for estimating the voxel-wise encoding models used in the current study. Nonetheless, this loss of efficiency is counteracted by the massive number of trials in the dataset.

Image-computable models and V1

A previous study⁹⁵ decoded natural images from BOLD responses in visual cortex by fitting encoding models to individual voxels. One finding was that including orientation tuning did not improve decoding beyond the accuracy obtained with only spatial frequency tuning^{95; Supplemental Figure 8}. This result seems to be at odd with our findings here. However, direct comparison between these studies is somewhat complicated. First, in a decoding approach, performance is sensitive to responses across multiple voxels and how they jointly encode stimuli. Such an approach yields results that are more difficult to interpret compared to a more straightforward approach in which properties of individual voxels are examined. Second, although modeling orientation tuning in addition to spatial frequency tuning did not lead to improved decoding

accuracy, when spatial frequency was absent or assumed to be identical across the entire region-of-interest, including orientation tuning did improve decoding. Thus, the results of the present study are not necessarily inconsistent with the previous study. Finally, it is important to note that estimation of model parameters for the Gabor wavelet encoding model used in the previous study was performed using gradient descent with early stopping. This type of regularization (early stopping) reduces variance at the expense of introducing bias, and the exact nature of this bias is dependent, in a complex way, on the statistics of the model inputs (e.g., orientation statistics in natural images). A major advantage of the approach used in the present study is the use of unregularized ordinary least-squares for parameter estimation, which was made possible by the combination of the sheer size of the NSD dataset and the pRF constraints incorporated into our models. This approach avoids complications associated with regularization and facilitates accurate interpretation of voxel selectivity.

The image-computable model that we used here was based on the steerable pyramid²⁶, a sub-band image transform that decomposes an image into orientation and spatial frequency channels (see *Methods: Steerable Pyramid*). In our previous study of stimulus vignetting¹⁷, we made two simplifying assumptions. First, because most of the power in the stimulus was at a single spatial frequency, we only analyzed the response of the model at a single spatial frequency channel centered at the spatial frequency of the stimulus (Fig S4). This approach provided good qualitative fits to the fMRI data, and we further confirmed that an alternative approach of averaging across all the channels did not change the model predictions (Fig S4). In the present study, such an approach is not feasible since the naturalistic images are broadband in spatial frequency. Instead, our modeling approach enabled fitting weights to all model channels, essentially estimating a spatial frequency tuning curve for each voxel. An alternative approach would be to weigh the different channels according to independent measures of spatial frequency tuning for each voxel^{6, 96}. The Natural Scenes Dataset could then be used to fit only the weights on orientation channels, which may result in more accurate estimates of orientation preference because of the smaller number of free parameters.

The second simplifying assumption involves scaling of the channel outputs. For each orientation and spatial frequency, the pyramid includes a quadrature pair: two RFs with different phases. The sum of the squares of the responses of the two RFs is typically taken, yielding an ‘energy’ response, which uniformly tiles all orientations and spatial frequencies^{97, 98}. This energy response is often nonlinearly scaled in order to better match the contrast-response function of V1 neurons. However, determining the model architecture and normalization pool appropriate for an fMRI voxel is not trivial and very much an area of active investigation^{99, 100, 101, 102}. We acknowledge that the form of scaling could have an impact on the size and spatial extent of vignetting. Determining the most appropriate scaling is an important issue that remains unresolved. In the context of the current study, we think it unlikely that the main results are sensitive to the particular regime of scaling employed.

Behavioral correlates of the coarse-scale bias

Understanding the stimulus selectivity of neurons and their organization is fundamental for understanding how neural computations lead to visual perception. In particular, coarse-scale organizations are likely critical elements of neural accounts of behavior, since large populations of neurons are likely to contribute to the final behavioral readout.

A behavioral radial bias has been reported by several psychophysics studies: sensitivity is higher to radial orientations than to other orientations^{14, 103, 104, 105}. Another well-known bias is the oblique effect: across the visual field sensitivity is higher to cardinal orientations (vertical or

horizontal) than to oblique orientations (diagonal)^{104, 105, 106, 107, 108}. The physiological orientation selectivity measured here may underlie both behavioral effects. The radial map may be the source of the behavioral radial bias, while the cardinal modulation may underlie the behavioral oblique effect. It has been suggested that fMRI response amplitudes reflect the neural SNR, which in turn determines the perceptual performance¹⁰⁹. In that case, a stronger fMRI response to the radial orientation should correspond to higher perceptual performance for stimuli with radial orientations.

A major endeavor in neuroscience is to link brain properties with behavioral readout, and visual neuroscience has made significant progress toward this goal. Recently it has been shown that V1 size and cortical magnification in individual subjects is correlated with contrast sensitivity^{110, 111}. Similarly, the extent of cortical magnification in V1 corresponds to orientation discrimination performance in individual participants¹¹². If the coarse-scale map revealed here constitutes the neural basis for behavioral anisotropies, we hypothesize that individual differences in the orientation map that we report here are related to individual differences in perception. Successfully demonstrating such a correspondence would provide a crucial link between brain and behavior.

Methods:

Natural Scenes Dataset:

The fMRI data analyzed here is from the Natural Scenes Dataset (NSD; <http://naturalscenesdataset.org>)²⁰. The NSD dataset contains measurements of fMRI responses from 8 participants who each viewed 9,000–10,000 distinct color natural scenes (22,000–30,000 trials) over the course of 30–40 scan sessions. Scanning was conducted at 7T using whole-brain gradient-echo EPI at 1.8-mm resolution and 1.6-s repetition time. Images were taken from the Microsoft Common Objects in Context (COCO) database¹¹³, square cropped, and presented at a size of 8.4° x 8.4°. A special set of 1,000 images were shared across subjects; the remaining images were mutually exclusive across subjects. Images were presented for 3 s with 1-s gaps in between images. Subjects fixated centrally and performed a long-term continuous recognition task on the images. The fMRI data were pre-processed by performing one temporal interpolation (to correct for slice time differences) and one spatial interpolation (to correct for head motion). A general linear model was then used to estimate single-trial beta weights. Cortical surface reconstructions were generated using FreeSurfer, and both volume- and surface-based versions of the beta weights were created. In this study, we used the 1.8-mm volume preparation of the NSD data and version 3 of the NSD single-trial betas in percent signal change units (betas_fithrf_GLMdenoise_RR). The results in this study are based on data from all NSD scan sessions, from all 8 subjects who participated in the NSD study.

Stimuli:

NSD images were originally 425 × 425 pixels, and were then upsampled for display purposes to 714 × 714 pixels. We reproduced this upsampling in our stimulus preparation, and padded the images with a gray border on all four sides (mimicking the scanner display environment), resulting in a final image dimension of 1024 × 1024 pixels (12.05° x 12.05°). A semitransparent red fixation point was added at the center to simulate the actual stimulation experienced by the subjects during the experiment. Images were converted to grayscale by averaging across the 3 color channels. To speed up subsequent computations, the images were then downsampled to 512 × 512 pixels. To enable cross-validation, the set of 10,000 images assigned to each subject was randomly divided into 2 partitions of 5,000 images each. For subjects who completed fewer than 40 sessions, only the viewed images were used, which resulted in a slightly different number of images included in each partition.

Steerable Pyramid:

We built two models based on the steerable pyramid²⁶: a full model and a constrained model. The full model simulates each neuron in V1 with a receptive field that is tuned for both spatial frequency and orientation, and then allows for variable weighting of these model neurons. The constrained model also simulates populations of V1 neurons, but enforces equal weighting of model neurons across orientation by summing across orientation subbands of the pyramid. It is possible to create steerable pyramid models with a wide range of parameters, each instantiating different hypotheses regarding the tuning properties of individual neurons.

We used a steerable pyramid with 8 orientations, 7 spatial frequency levels, and a spatial frequency bandwidth of 1 octave, resulting in tuning profiles that resemble those of individual V1 neurons¹¹⁴. The pyramid, and the full model, had a total of 56 filters. After summing across the 8 orientation filters, the constrained model consisted of 7 filters. The number of spatial frequency levels was determined by the size of the image (512 × 512) and the spatial frequency tuning bandwidth (1 octave). This results in 7 filters, with preferred spatial frequencies of 128,

64, 32, 16, 8, 4, and 2 cycles per image. These values were then converted to cycles per degree given the size of the image in degrees (12.05°): 21.85, 10.93, 5.46, 2.73, 1.37, 0.68, and 0.34 cycles/degree (cpd) (Fig S3). In our previous work with the steerable pyramid we used only the level corresponding to the stimulus spatial frequency. In this study all levels were fit to the data. We chose to have 8 orientations, 2 more than in our previous study, in order to increase the accuracy of the estimated orientation preference, while maintaining tuning width that was comparable to those measured in primate electrophysiological recordings¹¹⁴. For each orientation and spatial frequency, the pyramid includes a quadrature pair: two RFs with different phases. We take the sum of the squares of the responses of the two RFs, yielding an ‘energy’ response, which uniformly tiles all orientations and spatial frequencies^{97, 98}.

pRF modeling:

pRF estimates are included in the NSD, where full details are found²⁰. Briefly, pRFs were estimated based on a single session (6 runs, 300 s each) of a pRF mapping experiment. Stimuli consisted of slowly moving apertures (bars, wedges, and rings) filled with a dynamic colorful texture, that appeared within a circular region of 8.4 deg diameter. Subjects performed a color change detection task at fixation. pRFs were estimated using the Compressive Spatial Summation (CSS) model¹⁰¹.

Regions of interest:

Regions of interest V1, V2, V3, hV4 were defined in the NSD dataset based on the pRF maps. In this study we analyzed all 4 regions but focused on V1 where orientation selectivity has been studied most extensively. Results are presented for V1 only, except for the surface maps (Fig 3, Supplemental Fig 2) which show all regions.

pRF sampling:

The output of each filter in the steerable pyramid was sampled by each voxel’s pRF by multiplying the 2D pRF with the filter output. The pRF was modeled as a 2D isotropic (circular) Gaussian, using the ‘size’ parameter as the Gaussian’s standard deviation. (Note that the ‘size’ parameter, as estimated as part of NSD, reflects the response of the modeled pRF to point stimuli and takes into account the exponent used in the CSS model.) For filter k of image j ($F_i^{j,k}$), the sampled output for voxel i with a pRF centered at (x_i, y_i) and standard deviation of σ_i , is computed as dot product between the pRF and the filter:

$$(1) f_i^{j,k} = \sum_{x,y} F_{x,y}^{j,k} \cdot e^{-\frac{(x_i-x)^2 + (y_i-y)^2}{2\sigma_i^2}}$$

The full model had 56 sampled outputs per image, for each voxel. For the constrained model, sampled outputs were summed across orientations. Thus, the constrained model had 7 sampled outputs per image, for each voxel.

Multiple regression:

We modeled the responses of voxel i , y_i , as a linear combination of the sampled filter outputs plus noise:

$$(2) y_i = f_i \cdot \beta_i + \varepsilon_i$$

Here f_i is a matrix consisting of voxel i ’s sampled outputs for all filters of all images and a constant term (images x filters+1). β_i is a vector of beta weights (filters+1 x 1), and ε_i is a set of residuals (images x 1).

Beta weights were estimated using ordinary least-squares:

$$(3) \hat{\beta}_i = (f_i^T f_i)^{-1} f_i^T y_i$$

Note that each voxel not only had different beta weights but also different predictors due to the incorporation of each voxel's unique pRF, thus distinguishing this regression from a general linear model analysis of the voxel responses.

To assess model accuracy, we performed cross-validation. After estimating model parameters on one half of the data, the regression prediction was calculated as:

$$(4) \tilde{y}_i^{\text{pred}} = f_i \cdot \tilde{\beta}_i$$

where f_i is constructed for the other half of the data, and $\tilde{\beta}_i$ are the betas weights estimated using the other partition. The residual of this prediction is given by

$$(5) \tilde{y}_i^{\text{resid}} = y_i - \tilde{y}_i^{\text{pred}} = y_i - f_i \cdot \tilde{\beta}_i$$

Cross-validated R^2 is then computed as

$$(6) R_i^2 = 1 - \frac{SS(\tilde{y}_i^{\text{resid}})}{SS(y_i - \bar{y}_i)}$$

where $\bar{y}_i$ is the mean response across images, and SS denotes the sum of squares.

Regression was performed separately for the full model and for the constrained model on each of the 2 partitions. Regression coefficients and R^2 values were then averaged across partitions.

Inferring preferred orientation and spatial frequency:

After estimating the optimal weights for each voxel, we simulated an electrophysiology experiment for quantifying neural tuning, by probing the model with gratings at different orientations and spatial frequencies and measuring its predicted response. Full contrast gratings were 512×512 pixels at 30 spatial frequencies ranging from a single (horizontal) cycle in the image (0.083) to the Nyquist frequency (21.25 cpd), spaced exponentially. For each spatial frequency, gratings were oriented at 30 different angles, spaced uniformly between 0 and pi. All gratings were then passed through the steerable pyramid, and each filter's outputs were summed. Voxel responses to the gratings were simulated by multiplying the model outputs for each grating with the voxel's filter weights. Preferred spatial frequency of a voxel was estimated by first averaging simulated responses across orientations and then computing the mean frequency, weighted by response amplitudes. Similarly, preferred orientation of a voxel (for the full model) was estimated by averaging across frequencies and then computing the circular mean, weighted by response amplitudes. For all weighted means, the minimal weight was first subtracted from all weights to eliminate any negative weights.

Cortical surface maps

In order to create a group map on the cortical surface, each subject's data in volume space was transformed to surface space using nearest-neighbor interpolation, using the mrTools toolbox in Matlab¹¹⁵: each vertex was assigned the value of a single voxel, and multiple vertices could inherit values from the same voxel. Next, all subjects' surface data were transformed to a single cortical space, FreeSurfer's 'fsaverage' space. For the group map, we computed the circular mean across subjects for each vertex in V1-V4, weighted by the full model R^2 values. The resulting map of mean orientation preference was displayed on the 'fsaverage' inflated cortical surface (Fig 3B).

Quantifying coarse-scale biases

We quantified the strength of the coarse-scale orientation map by comparing it to an ideal, perfectly radial map, as well as ideal vertical and cardinal maps (see Fig 4). For each voxel, we computed the circular distance from the preferred orientation to the predicted preferred orientation. For the radial map the predicted preferred orientation was the radial orientation, according to the voxel's pRF angle (Fig 4, right). For the cardinal map the predicted orientation was vertical for pRF angles closer to the vertical meridian than to the horizontal meridian, and horizontal for pRF angles closer to the horizontal meridian than to the vertical meridian (Fig 4, center). For the vertical map the predicted orientation was vertical for all voxels (Fig 4, left). To average across voxels, we divided voxels into 20 bins according to voxel pRF eccentricity, pRF angle, and pRF R^2 . For eccentricity binning, bin width increased exponentially with eccentricity. Voxels with pRF R^2 values below zero were excluded from binning. To compare between the strength of the radial bias and other alternative biases, we averaged across voxels within each bin separately for each subject, and subject means were then submitted to a paired-sample t-test, with 7 degrees of freedom.

Control analysis: analyzing regression residuals

For this analysis, after performing multiple regression with the constrained model as predictors, we took the residuals from the same partition used to estimate the beta weights. We then performed regression on the residuals, this time using the full model.

Control analysis: analyzing regression prediction

For this analysis after performing multiple regression with the constrained model as predictors, we multiplied the regression coefficients of the same partition with the predictors to get the regression prediction. We then performed regression on the prediction, using the full model as predictors.

Data availability:

The NSD dataset is freely available at <http://naturalscenesdataset.org>. Images used for NSD were taken from the Common Objects in Context database (<https://cocodataset.org>). Source data are provided with this paper.

Code availability:

Code for analyzing the data and generating the figures is available at: <https://github.com/elimerriam/nsdOtopy>¹¹⁶.

References

1. Grill-Spector K, Malach R. The human visual cortex. *Annu Rev Neurosci* **27**, 649-677 (2004).
2. Schwartz EL. Computational anatomy and functional architecture of striate cortex: a spatial mapping approach to perceptual coding. *Vision Res* **20**, 645-669 (1980).
3. Benson NC, Butt OH, Datta R, Radoeva PD, Brainard DH, Aguirre GK. The retinotopic organization of striate cortex is well predicted by surface topology. *Current biology : CB* **22**, 2081-2085 (2012).
4. Dumoulin SO, Wandell BA. Population receptive field estimates in human visual cortex. *NeuroImage* **39**, 647-660 (2008).
5. Amano K, Wandell BA, Dumoulin SO. Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of neurophysiology* **102**, 2704-2718 (2009).
6. Aghajari S, Vinke LN, Ling S. Population spatial frequency tuning in human early visual cortex. *Journal of neurophysiology* **123**, 773-785 (2020).
7. Yacoub E, Shmuel A, Logothetis N, Uğurbil K. Robust detection of ocular dominance columns in humans using Hahn Spin Echo BOLD functional MRI at 7 Tesla. *NeuroImage* **37**, 1161-1177 (2007).
8. de Hollander G, van der Zwaag W, Qian C, Zhang P, Knapen T. Ultra-high field fMRI reveals origins of feedforward and feedback activity within laminae of human ocular dominance columns. *NeuroImage* **228**, 117683 (2021).
9. Cheng K, Waggoner RA, Tanaka K. Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron* **32**, 359-374 (2001).
10. Sun P, Ueno K, Waggoner RA, Gardner JL, Tanaka K, Cheng K. A temporal frequency-dependent functional architecture in human V1 revealed by high-resolution fMRI. *Nature neuroscience* **10**, 1404-1406 (2007).
11. Kaschube M, Schnabel M, Löwel S, Coppola DM, White LE, Wolf F. Universality in the evolution of orientation columns in the visual cortex. *Science* **330**, 1113-1116 (2010).
12. Yacoub E, Harel N, Ugurbil K. High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences of the United States of America* **105**, 10607-10612 (2008).

13. Freeman J, Brouwer GJ, Heeger DJ, Merriam EP. Orientation decoding depends on maps, not columns. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **31**, 4792-4804 (2011).
14. Sasaki Y, Rajimehr R, Kim BW, Ekstrom LB, Vanduffel W, Tootell RB. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron* **51**, 661-670 (2006).
15. Sun P, *et al.* 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 (2013).
16. Freeman J, Heeger DJ, Merriam EP. 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 (2013).
17. Roth ZN, Heeger DJ, Merriam EP. Stimulus vignetting and orientation selectivity in human visual cortex. *eLife* **7**, (2018).
18. Carlson TA. 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 (2014).
19. Gardner JL, Merriam EP. Population Models, Not Analyses, of Human Neuroscience Measurements. *Annu Rev Vis Sci* **7**, 225-255 (2021).
20. Allen EJ, *et al.* A massive 7T fMRI dataset to bridge cognitive and computational neuroscience. *bioRxiv*, (2021).
21. Naselaris T, Allen E, Kay K. Extensive sampling for complete models of individual brains. *Current Opinion in Behavioral Sciences* **40**, 45-51 (2021).
22. Allen EJ, *et al.* A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature neuroscience*, (2021).
23. Hermes D, Petridou N, Kay KN, Winawer J. An image-computable model for the stimulus selectivity of gamma oscillations. *eLife* **8**, (2019).
24. Jigo M, Heeger DJ, Carrasco M. An image-computable model of how endogenous and exogenous attention differentially alter visual perception. *Proceedings of the National Academy of Sciences of the United States of America* **118**, (2021).
25. Morgenstern Y, *et al.* An image-computable model of human visual shape similarity. *PLoS computational biology* **17**, e1008981 (2021).

- 881
882 26. Simoncelli EP, Freeman WT, Adelson EH, Heeger DJ. Shiftable multiscale transforms.
883 *IEEE transactions on Information Theory* **38**, 587-607 (1992).
884
885 27. Kamitani Y, Tong F. Decoding the visual and subjective contents of the human brain.
886 *Nature neuroscience* **8**, 679-685 (2005).
887
888 28. Boynton GM. Imaging orientation selectivity: decoding conscious perception in V1.
889 *Nature neuroscience* **8**, 541-542 (2005).
890
891 29. Haynes JD, Rees G. Predicting the orientation of invisible stimuli from activity in human
892 primary visual cortex. *Nature neuroscience* **8**, 686-691 (2005).
893
894 30. Das A, Gilbert CD. Distortions of visuotopic map match orientation singularities in
895 primary visual cortex. *Nature* **387**, 594-598 (1997).
896
897 31. Grinvald A, Lieke E, Frostig RD, Gilbert CD, Wiesel TN. Functional architecture of cortex
898 revealed by optical imaging of intrinsic signals. *Nature* **324**, 361-364 (1986).
899
900 32. Ohki K, Chung S, Kara P, Hübener M, Bonhoeffer T, Reid RC. Highly ordered arrangement
901 of single neurons in orientation pinwheels. *Nature* **442**, 925-928 (2006).
902
903 33. Adams DL, Sincich LC, Horton JC. Complete pattern of ocular dominance columns in
904 human primary visual cortex. *The Journal of neuroscience : the official journal of the*
905 *Society for Neuroscience* **27**, 10391-10403 (2007).
906
907 34. Horton JC, Hedley-Whyte ET. Mapping of cytochrome oxidase patches and ocular
908 dominance columns in human visual cortex. *Philos Trans R Soc Lond B Biol Sci* **304**, 255-
909 272 (1984).
910
911 35. Kamitani Y, Sawahata Y. Spatial smoothing hurts localization but not information: pitfalls
912 for brain mappers. *NeuroImage* **49**, 1949-1952 (2010).
913
914 36. Op de Beeck HP. Against hyperacuity in brain reading: spatial smoothing does not hurt
915 multivariate fMRI analyses? *NeuroImage* **49**, 1943-1948 (2010).
916
917 37. Op de Beeck HP. Probing the mysterious underpinnings of multi-voxel fMRI analyses.
918 *NeuroImage* **50**, 567-571 (2010).
919
920 38. Swisher JD, *et al.* Multiscale pattern analysis of orientation-selective activity in the
921 primary visual cortex. *The Journal of neuroscience : the official journal of the Society for*
922 *Neuroscience* **30**, 325-330 (2010).
923

39. Alink A, Walther A, Krugliak A, Kriegeskorte N. Local opposite orientation preferences in V1: fMRI sensitivity to fine-grained pattern information. *Scientific reports* **7**, 7128 (2017).
40. Carlson TA, Wardle SG. Sensible decoding. *NeuroImage* **110**, 217-218 (2015).
41. Wardle SG, Ritchie JB, Seymour K, Carlson TA. 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 (2017).
42. Alink A, Krugliak A, Walther A, Kriegeskorte N. fMRI orientation decoding in V1 does not require global maps or globally coherent orientation stimuli. *Front Psychol* **4**, 493 (2013).
43. Chaimow D, Yacoub E, Ugurbil K, Shmuel A. Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns. *NeuroImage* **56**, 627-642 (2011).
44. Tong F, Harrison SA, Dewey JA, Kamitani Y. Relationship between BOLD amplitude and pattern classification of orientation-selective activity in the human visual cortex. *NeuroImage* **63**, 1212-1222 (2012).
45. Schira MM, Wade AR, Tyler CW. Two-dimensional mapping of the central and parafoveal visual field to human visual cortex. *Journal of neurophysiology* **97**, 4284-4295 (2007).
46. Benson NC, Winawer J. Bayesian analysis of retinotopic maps. *eLife* **7**, (2018).
47. Jin J, Wang Y, Swadlow HA, Alonso JM. Population receptive fields of ON and OFF thalamic inputs to an orientation column in visual cortex. *Nature neuroscience* **14**, 232-238 (2011).
48. Paik SB, Ringach DL. Retinal origin of orientation maps in visual cortex. *Nature neuroscience* **14**, 919-925 (2011).
49. Vidyasagar TR, Eysel UT. Origins of feature selectivities and maps in the mammalian primary visual cortex. *Trends Neurosci* **38**, 475-485 (2015).
50. Huber L, Uludağ K, Möller HE. Non-BOLD contrast for laminar fMRI in humans: CBF, CBV, and CMR(O2). *NeuroImage* **197**, 742-760 (2019).
51. Gardner JL. Is cortical vasculature functionally organized? *NeuroImage* **49**, 1953-1956 (2010).

52. Kriegeskorte N, Cusack R, Bandettini P. How does an fMRI voxel sample the neuronal activity pattern: compact-kernel or complex spatiotemporal filter? *NeuroImage* **49**, 1965-1976 (2010).
53. Pratte MS, Sy JL, Swisher JD, Tong F. Radial bias is not necessary for orientation decoding. *NeuroImage* **127**, 23-33 (2016).
54. Mannion DJ, McDonald JS, Clifford CW. Orientation anisotropies in human visual cortex. *Journal of neurophysiology* **103**, 3465-3471 (2010).
55. Furmanski CS, Engel SA. An oblique effect in human primary visual cortex. *Nature neuroscience* **3**, 535-536 (2000).
56. Xu X, Collins CE, Khaytin I, Kaas JH, Casagrande VA. Unequal representation of cardinal vs. oblique orientations in the middle temporal visual area. *Proceedings of the National Academy of Sciences of the United States of America* **103**, 17490-17495 (2006).
57. Mansfield RJ. Neural basis of orientation perception in primate vision. *Science* **186**, 1133-1135 (1974).
58. Fang C, Cai X, Lu HD. Orientation anisotropies in macaque visual areas. *Proceedings of the National Academy of Sciences of the United States of America* **119**, e2113407119 (2022).
59. Grabska-Barwińska A, Distler C, Hoffmann KP, Jancke D. Contrast independence of cardinal preference: stable oblique effect in orientation maps of ferret visual cortex. *Eur J Neurosci* **29**, 1258-1270 (2009).
60. Coppola DM, White LE, Fitzpatrick D, Purves D. Unequal representation of cardinal and oblique contours in ferret visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* **95**, 2621-2623 (1998).
61. Chapman B, Bonhoeffer T. Overrepresentation of horizontal and vertical orientation preferences in developing ferret area 17. *Proceedings of the National Academy of Sciences of the United States of America* **95**, 2609-2614 (1998).
62. Berman NE, Wilkes ME, Payne BR. Organization of orientation and direction selectivity in areas 17 and 18 of cat cerebral cortex. *Journal of neurophysiology* **58**, 676-699 (1987).
63. Huang L, Shou T, Chen X, Yu H, Sun C, Liang Z. Slab-like functional architecture of higher order cortical area 21a showing oblique effect of orientation preference in the cat. *NeuroImage* **32**, 1365-1374 (2006).

64. Wang G, Ding S, Yunokuchi K. Difference in the representation of cardinal and oblique contours in cat visual cortex. *Neurosci Lett* **338**, 77-81 (2003).
65. Li B, Peterson MR, Freeman RD. Oblique effect: a neural basis in the visual cortex. *Journal of neurophysiology* **90**, 204-217 (2003).
66. Kay K, Jamison KW, Vizioli L, Zhang R, Margalit E, Ugurbil K. A critical assessment of data quality and venous effects in sub-millimeter fMRI. *NeuroImage* **189**, 847-869 (2019).
67. Kurzwaski JW, Gulban OF, Jamison KW, Winawer J, Kay K. The influence of non-neural factors on BOLD signal magnitude. *bioRxiv*, (2021).
68. Kurzwaski JW, Gulban OF, Jamison K, Winawer J, Kay K. Non-neural factors influencing BOLD response magnitudes within individual subjects. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **42**, 7256-7266 (2022).
69. Hubel DH, Wiesel TN. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* **160**, 106-154 (1962).
70. Lien AD, Scanziani M. Tuned thalamic excitation is amplified by visual cortical circuits. *Nature neuroscience* **16**, 1315-1323 (2013).
71. Sun W, Tan Z, Mensh BD, Ji N. Thalamus provides layer 4 of primary visual cortex with orientation- and direction-tuned inputs. *Nature neuroscience* **19**, 308-315 (2016).
72. Ling S, Pratte MS, Tong F. Attention alters orientation processing in the human lateral geniculate nucleus. *Nature neuroscience* **18**, 496-498 (2015).
73. Cheong SK, Tailby C, Solomon SG, Martin PR. Cortical-like receptive fields in the lateral geniculate nucleus of marmoset monkeys. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **33**, 6864-6876 (2013).
74. Zhao X, Chen H, Liu X, Cang J. Orientation-selective responses in the mouse lateral geniculate nucleus. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **33**, 12751-12763 (2013).
75. Piscopo DM, El-Danaf RN, Huberman AD, Niell CM. Diverse visual features encoded in mouse lateral geniculate nucleus. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **33**, 4642-4656 (2013).
76. Scholl B, Tan AY, Corey J, Priebe NJ. Emergence of orientation selectivity in the Mammalian visual pathway. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **33**, 10616-10624 (2013).

77. Levick WR, Thibos LN. Analysis of orientation bias in cat retina. *J Physiol* **329**, 243-261 (1982).
78. Antinucci P, Suleyman O, Monfries C, Hindges R. Neural Mechanisms Generating Orientation Selectivity in the Retina. *Current biology : CB* **26**, 1802-1815 (2016).
79. Schall JD, Perry VH, Leventhal AG. Retinal ganglion cell dendritic fields in old-world monkeys are oriented radially. *Brain Res* **368**, 18-23 (1986).
80. Schall JD, Vitek DJ, Leventhal AG. Retinal constraints on orientation specificity in cat visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* **6**, 823-836 (1986).
81. Thibos LN, Levick WR. Astigmatic visual and deprivation in cat: behavioral, optical and retinophysiological consequences. *Vision Res* **22**, 43-53 (1982).
82. Jang J, Song M, Paik SB. Retino-Cortical Mapping Ratio Predicts Columnar and Salt-and-Pepper Organization in Mammalian Visual Cortex. *Cell Rep* **30**, 3270-3279.e3273 (2020).
83. Mukamel R, Gelbard H, Arieli A, Hasson U, Fried I, Malach R. Coupling between neuronal firing, field potentials, and fMRI in human auditory cortex. *Science* **309**, 951-954 (2005).
84. Nir Y, *et al.* Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to interneuronal correlations. *Current biology : CB* **17**, 1275-1285 (2007).
85. Heeger DJ, Ress D. What does fMRI tell us about neuronal activity? *Nature reviews Neuroscience* **3**, 142-151 (2002).
86. Rees G, Friston K, Koch C. A direct quantitative relationship between the functional properties of human and macaque V5. *Nature neuroscience* **3**, 716-723 (2000).
87. Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. *Nature* **412**, 150-157 (2001).
88. Yen CC, Fukuda M, Kim SG. BOLD responses to different temporal frequency stimuli in the lateral geniculate nucleus and visual cortex: insights into the neural basis of fMRI. *NeuroImage* **58**, 82-90 (2011).
89. Henderson M, Serences JT. Biased orientation representations can be explained by experience with nonuniform training set statistics. *J Vis* **21**, 10 (2021).
90. Philips RT, Chakravarthy VS. A Global Orientation Map in the Primary Visual Cortex (V1): Could a Self Organizing Model Reveal Its Hidden Bias? *Front Neural Circuits* **10**, 109 (2016).

91. da Costa D, Goebel R, Senden M. ConvNets Develop Organizational Principles of the Visual Cortex when using Ganglion Cell-Based Sampling. *bioRxiv*, (2021).
92. Bruce ND, Tsotsos JK. A statistical basis for visual field anisotropies. *Neurocomputing* **69**, 1301-1304 (2006).
93. Rothkopf CA, Weisswange TH, Triesch J. Learning independent causes in natural images explains the spacevariant oblique effect. In: *2009 IEEE 8th International Conference on Development and Learning*). IEEE (2009).
94. Enroth-Cugell C, Robson JG. The contrast sensitivity of retinal ganglion cells of the cat. *J Physiol* **187**, 517-552 (1966).
95. Kay KN, Naselaris T, Prenger RJ, Gallant JL. Identifying natural images from human brain activity. *Nature* **452**, 352-355 (2008).
96. Broderick WF, Simoncelli EP, Winawer J. Mapping spatial frequency preferences across human primary visual cortex. *J Vis* **22**, 3 (2022).
97. Adelson EH, Bergen JR. Spatiotemporal energy models for the perception of motion. *Josa a* **2**, 284-299 (1985).
98. Heeger DJ. Half-squaring in responses of cat striate cells. *Vis Neurosci* **9**, 427-443 (1992).
99. Klímová M, Bloem IM, Ling S. The specificity of orientation-tuned normalization within human early visual cortex. *Journal of neurophysiology* **126**, 1536-1546 (2021).
100. Aqil M, Knapen T, Dumoulin SO. Divisive normalization unifies disparate response signatures throughout the human visual hierarchy. *Proceedings of the National Academy of Sciences of the United States of America* **118**, (2021).
101. Kay KN, Winawer J, Mezer A, Wandell BA. Compressive spatial summation in human visual cortex. *Journal of neurophysiology* **110**, 481-494 (2013).
102. Winawer J, Fang Z, Ma WJ. Normalization by the variance across orientation channels in human V1–V3. *Journal of Vision* **19**, 43b-43b (2019).
103. Temme LA, Malcus L, Noell WK. Peripheral visual field is radially organized. *Am J Optom Physiol Opt* **62**, 545-554 (1985).
104. Westheimer G. The distribution of preferred orientations in the peripheral visual field. *Vision Res* **43**, 53-57 (2003).

- 1141 105. Westheimer G. Anisotropies in peripheral vernier acuity. *Spat Vis* **18**, 159-167 (2005).
1142
- 1143 106. Girshick AR, Landy MS, Simoncelli EP. Cardinal rules: visual orientation perception
1144 reflects knowledge of environmental statistics. *Nature neuroscience* **14**, 926-932 (2011).
1145
- 1146 107. Coppola DM, Purves HR, McCoy AN, Purves D. The distribution of oriented contours in
1147 the real world. *Proceedings of the National Academy of Sciences of the United States of*
1148 *America* **95**, 4002-4006 (1998).
1149
- 1150 108. Appelle S. Perception and discrimination as a function of stimulus orientation: the
1151 "oblique effect" in man and animals. *Psychological bulletin* **78**, 266-278 (1972).
1152
- 1153 109. Pestilli F, Ling S, Carrasco M. A population-coding model of attention's influence on
1154 contrast response: Estimating neural effects from psychophysical data. *Vision Res* **49**,
1155 1144-1153 (2009).
1156
- 1157 110. Himmelberg M, Winawer J, Carrasco M. Linking individual differences in human V1 to
1158 perception around the visual field. (2021).
1159
- 1160 111. Himmelberg MM, Winawer J, Carrasco M. Linking individual differences in human
1161 primary visual cortex to contrast sensitivity around the visual field. *Nature*
1162 *communications* **13**, 3309 (2022).
1163
- 1164 112. Benson NC, Kupers ER, Barbot A, Carrasco M, Winawer J. Cortical magnification in
1165 human visual cortex parallels task performance around the visual field. *eLife* **10**, (2021).
1166
- 1167 113. Lin T-Y, et al. Microsoft coco: Common objects in context. In: *European conference on*
1168 *computer vision*). Springer (2014).
1169
- 1170 114. Ringach DL, Shapley RM, Hawken MJ. Orientation selectivity in macaque V1: diversity
1171 and laminar dependence. *The Journal of neuroscience : the official journal of the Society*
1172 *for Neuroscience* **22**, 5639-5651 (2002).
1173
- 1174 115. Gardner JL, Merriam EP, Schluppeck D, Besle J, Heeger DJ. mrTools: Analysis and
1175 visualization package for functional magnetic resonance imaging data, version 4.7.
1176 (2018).
1177
- 1178 116. Roth ZN, Kay K, Merriam EP. nsdOtopy: Natural scene sampling reveals reliable coarse-
1179 scale orientation tuning in human V1. (2022).
1180
1181
1182

Acknowledgements

This research was supported by the Intramural Research Program of the NIH (ZIAMH002966) to E.P.M. Collection of the NSD dataset was supported by NSF IIS-1822683 to K.K. and NSF IIS-1822929.

Author contributions:

Conceptualization, Methodology, Investigation, Validation, Writing – Review & Editing: Z.N.R., K.K., & E.P.M.
Software, Data Curation: Z.N.R. & E.P.M.
Resources, Supervision: K.K. & E.P.M.
Project Administration, Funding Acquisition: E.P.M.
Formal Analysis, Investigation, Writing – Original Draft, Visualization: Z.N.R.

Competing interests: The authors declare no competing interests.


Supplementary Materials for
**Natural scene sampling reveals reliable coarse-scale orientation tuning in
human V1**

Zvi N. Roth*, Kendrick Kay, Elisha P. Merriam,


*Corresponding author. Email: zvi.roth@nih.gov

This PDF file includes:


Figs. S1 to S4



Supplemental Fig S1. Orientation maps are consistent across participants. Same color legend and scaling as Fig 3A. Orientation preferences plotted in visual space. Each line represents a single voxel and is positioned at the voxel's pRF center. Hue and orientation of the line indicate preferred orientation. Line length, width, and scale reflect the amount of variance (R^2) explained by the constrained model. Solid square at +/- 4.2 deg indicates the size of the natural-scene stimuli.



Supplemental Fig S2. Orientation surface maps are consistent across individuals. Same color legend as Fig 3B. Individual subject orientation maps (top, thresholded; middle, unthresholded) and pRF polar angle maps (bottom) overlaid on left and right inflated ‘fsaverage’ surfaces. For angle map and unthresholded orientation map, all vertices in V1, V2, V3, and V4 are plotted. For thresholded orientation map only vertices with the top 50% full model R^2 are plotted.



Supplemental Fig S3. Full model and constrained model filters in Fourier space. The full model allows unequal contributions from orientation-selective filters, and is therefore sensitive to orientation-selectivity beyond the effects of stimulus vignetting. The constrained model pools equally across orientation-selective filters, and therefore has no orientation selectivity, but due to spatial frequency tuning it accounts for effects of stimulus vignetting.



Supplemental Fig S4. Response of constrained model filters to circular apertured gratings at various spatial frequencies. Responses to vertical orientation minus responses to horizontal orientation are shown for all 7 filters, for 5 different grating spatial frequencies. Stimulus is 512×512 pixels, here simulated to span 10×10 degrees of visual angle. For each spatial frequency, responses are averaged across 8 different phases. Unscaled responses (middle row for each spatial frequency) shows responses scaled from the minimum (black) to maximum (white) for each individual filter. This enables one to see how the spatial patterns change across the different filters, but doesn't convey a sense of the relative magnitude. In the bottom row of each spatial frequency the gray scale for all panels ranges from the minimum to maximum across all filters ('scaled'). The filter response corresponding to the stimulus frequency has the largest response, ranging from black to white, but other filters have smaller responses, which are therefore washed out and almost invisible. The model filter that corresponds to the grating spatial frequency (marked with a red box) exhibits a radial bias, while the adjacent filters show an opposite effect. However, when scaling all filters to the same maximal value it becomes apparent that the effects in adjacent filters are much weaker than that of the filter with the maximal response. Therefore, as long as the spatial frequency of the stimulus is matched to the preferred spatial frequency tuning of a local population of neurons in a voxel, and assuming a circular aperture, vignetting should result in a radial bias.