

Origin and Function of Tuning Diversity in Macaque Visual Cortex

Highlights

- Visual cortex neurons differ in the precision of their orientation selectivity
- Using novel stimuli and a model, we dissect mechanisms that create tuning diversity
- Variations in linear receptive fields and in nonlinear sharpening create diversity
- Tuning diversity makes neural populations more efficient encoders of natural images

Authors

Robbe L.T. Goris, Eero P. Simoncelli,
J. Anthony Movshon

Correspondence

Robbe.Goris@nyu.edu (R.L.T.G.),
movshon@nyu.edu (J.A.M.)

In Brief

Neurons in primary visual cortex are selective for stimulus orientation, but the precision of this selectivity differs across neurons. Using novel stimuli and a computational “cascade” model, Goris et al. shed light on the origins and purpose of this diversity.



Origin and Function of Tuning Diversity in Macaque Visual Cortex

Robbe L.T. Goris,^{1,2,*} Eero P. Simoncelli,^{1,2} and J. Anthony Movshon^{1,*}

¹Center for Neural Science

²Howard Hughes Medical Institute

New York University, 4 Washington Place, Room 809, New York, NY 10003, USA

*Correspondence: Robbe.Goris@nyu.edu (R.L.T.G.), movshon@nyu.edu (J.A.M.)

<http://dx.doi.org/10.1016/j.neuron.2015.10.009>

SUMMARY

Neurons in visual cortex vary in their orientation selectivity. We measured responses of V1 and V2 cells to orientation mixtures and fit them with a model whose stimulus selectivity arises from the combined effects of filtering, suppression, and response nonlinearity. The model explains the diversity of orientation selectivity with neuron-to-neuron variability in all three mechanisms, of which variability in the orientation bandwidth of linear filtering is the most important. The model also accounts for the cells' diversity of spatial frequency selectivity. Tuning diversity is matched to the needs of visual encoding. The orientation content found in natural scenes is diverse, and neurons with different selectivities are adapted to different stimulus configurations. Single orientations are better encoded by highly selective neurons, while orientation mixtures are better encoded by less selective neurons. A diverse population of neurons therefore provides better overall discrimination capabilities for natural images than any homogeneous population.

INTRODUCTION

Orientation selectivity is the best-known property of neurons in primary visual cortex (Hubel and Wiesel, 1959). But the precision of this selectivity is diverse: some neurons are highly selective for orientation, while others are more broadly tuned (Rose and Blakemore, 1974; Schiller et al., 1976; Ringach et al., 2002; Scholl et al., 2013). This diversity is maintained in downstream cortical areas (Albright 1984; Levitt et al., 1994). Many studies have examined the origins of orientation selectivity, and many theories have been proposed to explain its functional role, but the origin and function of its diversity has received relatively little attention.

Orientation selectivity arises from the organization of excitatory LGN input (Hubel and Wiesel, 1962; Chapman et al., 1991; Reid and Alonso, 1995; Jin et al., 2011). It is further sculpted by intracortical inhibitory circuits (Sillito, 1975; Ringach et al., 1997; Xing et al., 2011) and by the nonlinear transduction of membrane voltage to spiking activity (Carandini and Ferster, 2000). To dissociate the contributions of these different mecha-

nisms, we measured responses to mixture stimuli that vary in the distribution of their orientation content. We fit these responses with a linear–nonlinear–linear–nonlinear (LN-LN) cascade model, which has its roots in the work of Hubel and Wiesel (1962) and has been refined by many others (Movshon et al., 1978; Heeger, 1992; Carandini et al., 1997; Touryan et al., 2002; Sharpee et al., 2004; Rust et al., 2005; Vintch et al., 2015). The initial linear filtering stage creates basic orientation tuning, which is then modified by subsequent nonlinearities that capture the effects of untuned suppression and spike threshold. The model provides a good account of the diverse orientation selectivity of V1 and V2 cells. Diversity arises from neuron-to-neuron variability in all three model components, but the primary source is the initial filtering stage.

Functionally, this diversity permits the efficient encoding of different kinds of stimulus information. Narrowly tuned neurons transmit more information about uniquely oriented stimuli, while broadly tuned neurons transmit more information about orientation mixtures. This specialization may be an adaptation to the diverse orientation content found in the natural environment. Specifically, we show that a population of neurons with mixed orientation selectivity provides better discrimination capabilities for a set of natural images than homogeneous populations of any particular selectivity. These results suggest that tuning diversity originates from variations in the organization of excitatory inputs and improves the ability of the cortex to encode the visual information in nature.

RESULTS

Diversity of Orientation Selectivity in V1 and V2

The diversity of orientation selectivity in visual cortex can be appreciated by considering the responses of V1 cells to sinusoidal gratings drifting in different directions, as shown for three example neurons in Figure 1A. For each example, the mean response varied with drift direction. Some neurons responded exclusively to a narrow range of stimulus orientations, but others were less selective. The tuning may be conveniently summarized with an orientation selectivity index (OSI)

$$\text{OSI} = \frac{\left| \sum_{j=1}^J R_j e^{i2\theta_j} \right|}{\sum_{j=1}^J |R_j|},$$

where θ_j is the orientation and R_j is the response magnitude (with baseline subtracted) for the j^{th} stimulus (Leventhal et al., 1995).

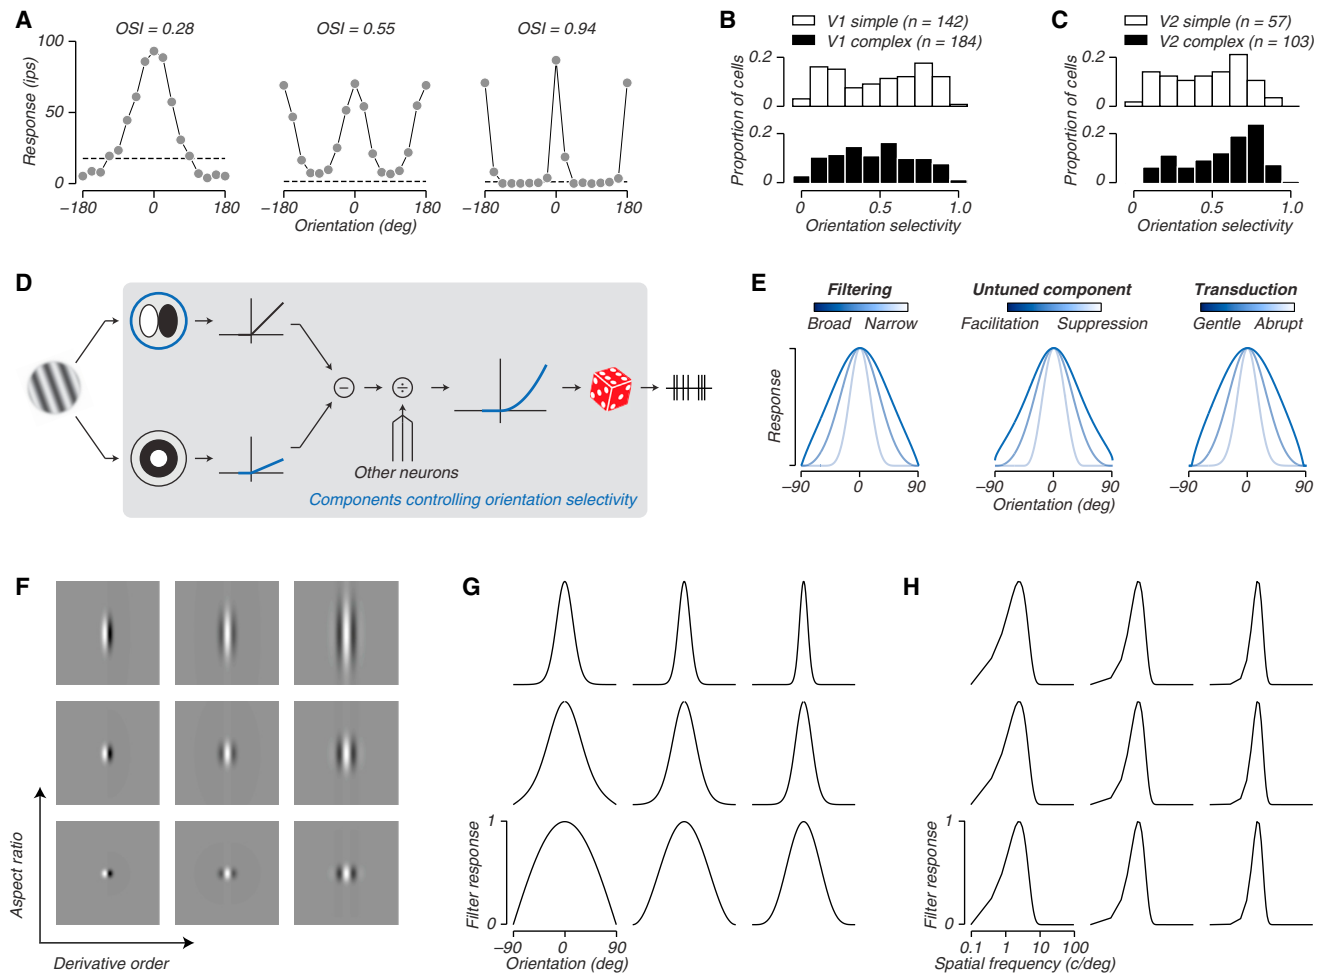


Figure 1. Diversity in Orientation Selectivity

(A) Responses of three example V1 neurons, stimulated with gratings drifting in different directions. Responses were computed by counting spikes in a 1,000 ms window following response onset. Dashed lines show baseline activity.

(B) Distribution of OSI (one minus normalized circular variance of tuning curve) for a population of V1 simple (white) and complex (black) cells. Cells were classified based on the temporal modulation of the response to the preferred stimulus.

(C) Distribution of orientation selectivity for cells in V2.

(D) The LN-LN model. The stimulus is processed in two parallel channels: one orientation-selective (top left) and one untuned (bottom left). Within each channel, a linear filtering stage is followed by rectification. The channel responses are combined linearly and modified by a divisive gain control set by the pooled response of other neurons. The normalized response is passed through a nonlinearity to obtain a firing rate, which drives an inhomogeneous Poisson process.

(E) Orientation tuning for three families of model neurons. Within each family, one model component that controls orientation selectivity was varied (left: filter selectivity; middle: sign and strength of the untuned component; right: nonlinearity exponent).

(F) Example Gaussian derivative filters, with different derivative order (left to right: 1, 2, 4) and aspect ratio (bottom to top: 1, 2, 4).

(G and H) Orientation tuning and spatial frequency tuning of the corresponding example filters shown in (F).

This index varies between 0 and 1 and does not depend on direction selectivity. Orientation selectivity is broadly distributed over the range from 0 to 1 for both simple and complex cells (Figures 1B and 1C), with a slightly larger median in V2 than in V1 ($p = 0.047$, Wilcoxon rank-sum test).

To examine the computations underlying this diversity of orientation selectivity, we used an LN-LN cascade model (Figure 1D). In the model, the stimulus is processed in two channels, each consisting of a linear filter—one orientation selective and one untuned—followed by rectification. The oriented filters were spatial derivatives of a 2D Gaussian (Figure 1F; see also

Figure S4). The channel responses are subtracted, subjected to divisive normalization, and transformed with an instantaneous response nonlinearity whose output drives the firing rate of an inhomogeneous Poisson process.

In this model, tuning for orientation arises from the combined action of three different model components. To illustrate this, we simulated tuning curves under changes of each of these components (Figure 1E).

First, changes in the orientation selectivity of the linear filter can cause corresponding changes in the overall model orientation tuning (Figure 1E, left). The linear filters in our model are

derivatives of a Gaussian and are parameterized by the aspect ratio of the Gaussian and the derivative order (Figure 1F). Both parameters have an effect on the orientation selectivity (Figure 1G). The latter parameter also governs spatial frequency selectivity—the more cycles of modulation within the receptive field, the narrower the spatial frequency selectivity (Figure 1H) (Tolhurst and Thompson, 1981).

Second, the nonlinear function that is applied to the untuned component can either increase or decrease overall orientation tuning specificity relative to that of the linear filter (Figure 1E, middle) without changing receptive field structure. Specifically, overall tuning width depends on the baseline value of the untuned component, which (depending on its sign) either facilitates or suppresses responses to non-optimal stimuli. Untuned suppression can sharpen orientation tuning (Xing et al., 2005) while untuned facilitation can broaden it.

Third, changes in the final response (transduction) nonlinearity have a direct effect on the model orientation tuning (Figure 1E, right), again without changing the receptive field structure. For our model, which uses a power law nonlinearity, exponents larger than one yield an accelerating nonlinearity that sharpens model orientation tuning, whereas smaller exponents broaden it. Since intracellular recordings indicate that spike threshold exclusively sharpens tuning (Carandini and Ferster, 2000; Priebe et al., 2004), we constrain our model fits to use exponents greater than one.

Finally, it is worth noting that divisive normalization in our model does not alter stimulus selectivity but simply rescales responses, since we assume a normalization signal that is untuned for orientation (Heeger, 1992).

Modeling Responses to Orientation Mixtures

How can the contributions of these mechanisms to orientation selectivity be isolated? The initial filter is linear and must obey the principle of superposition: the response to a sum of stimuli is equal to the sum of responses to the individual stimuli. The other two model components are nonlinear and will cause deviations from superposition. Responses to linear mixtures of oriented stimuli therefore offer a way to separate the effects of linear and nonlinear mechanisms. To distinguish the effects of the nonlinear model components, we consider their operation in detail. Inhibition and nonlinear transduction both refine selectivity by vetoing weak filter responses. But a crucial difference is that the vetoing power of inhibitory mechanisms scales with stimulus contrast, while a threshold imposes a contrast-independent veto. Responses to mixture stimuli that vary in contrast therefore offer a way to disentangle the specific contributions of excitation, inhibition, and transduction to orientation selectivity.

For a subset of our cells, we measured responses to mixtures of incoherently drifting sinusoidal gratings whose orientations were spaced at 20° intervals and whose orientation-dependent contrasts followed a circular Gaussian profile centered on a particular orientation (spread 0°–55°, Figure 2A). We presented each stimulus at 16 center orientations and at high and low total contrast (100% and 33%, respectively). To fully identify the linear filter (Figure 1F), we also characterized spatial frequency selectivity by measuring responses to a collection of high-contrast single gratings of different spatial frequency presented at the

preferred orientation. For each cell, we fit the LN-LN model to the full set of responses to the orientation mixtures and single component gratings, adjusting model parameters to maximize the likelihood under a modulated Poisson spiking model (Goris et al., 2014) (see Experimental Procedures). We verified that this set of measurements yields data that are sufficiently rich to unambiguously identify the different model components (see Figure S5).

The mixture stimuli revealed differences between cells that were not evident from grating responses alone. Consider the actual and predicted responses for two example V1 cells (Figure 2B). These neurons had nearly identical tuning when stimulated with single gratings (OSI = 0.94 for neuron 1 and 0.95 for neuron 2). But their responses to mixture stimuli differed markedly. Neuron 1 responded selectively to all high-contrast mixture stimuli, though its response magnitude fell with stimulus spread. In contrast, neuron 2 responded only to stimuli with spread of 20° or less (Figure 2B). Moreover, reducing the contrast of a single grating to 33% rescaled the responses of neuron 1 but silenced neuron 2 (Figure 2B). The fitted model (solid colored lines) accurately captures the behaviors of both neurons and explains why they differ (Figures 2C and 2D): neuron 1's model has a more selective filter than neuron 2's (filter OSI: 0.86 versus 0.47), but weaker inhibition (gain untuned channel: 0.05 versus -0.48) and a slightly milder response nonlinearity (power law exponent: 3.8 versus 4.6). These differences in mechanism are also reflected in the spatial frequency tuning, which was narrow for neuron 1 but broad for neuron 2 (Figure 2E). Neuron 1's model fit is rather extreme in the selectivity of its filter; neuron 2's fit is rather extreme in the strength of untuned suppression. Additional example neurons and model fits are shown in Figure S1.

The fitted model predicts tuning curves for the different mixture stimuli that are in excellent agreement with the measured responses in V1 ($r = 0.93$, median correlation between the predicted and measured mean spike counts for hold-out data in a 10-fold cross-validation analysis; see Experimental Procedures) and V2 ($r = 0.88$). Although small, this difference is significant ($p = 0.0013$, Wilcoxon rank-sum test), indicating that the model better captures V1 responses than V2 responses. To assess the model's goodness of fit more explicitly, we compared the log-probability of the data with a simulated reference distribution (see Experimental Procedures). This test indicates the degree to which the model-predicted spike count distributions capture the observed responses on a trial-by-trial basis. For 73 of 80 neurons, the model fit cannot be rejected at a significance level of 5%. We elected not to remove the remaining cells from the data-set, although doing so does not affect any of the following results.

Contributions of Different Mechanisms to Selectivity

The full model explains 86% (V1) and 55% (V2) of the diversity of orientation selectivity across the population (Figure 3A). Our measure of explained diversity is analogous to explained variance in a regression model, but it captures both the overall correlation between observed and predicted values of OSI and the systematic errors in the prediction (see Experimental Procedures). We wondered how much each recovered model component contributed to the tuning diversity of the population. We

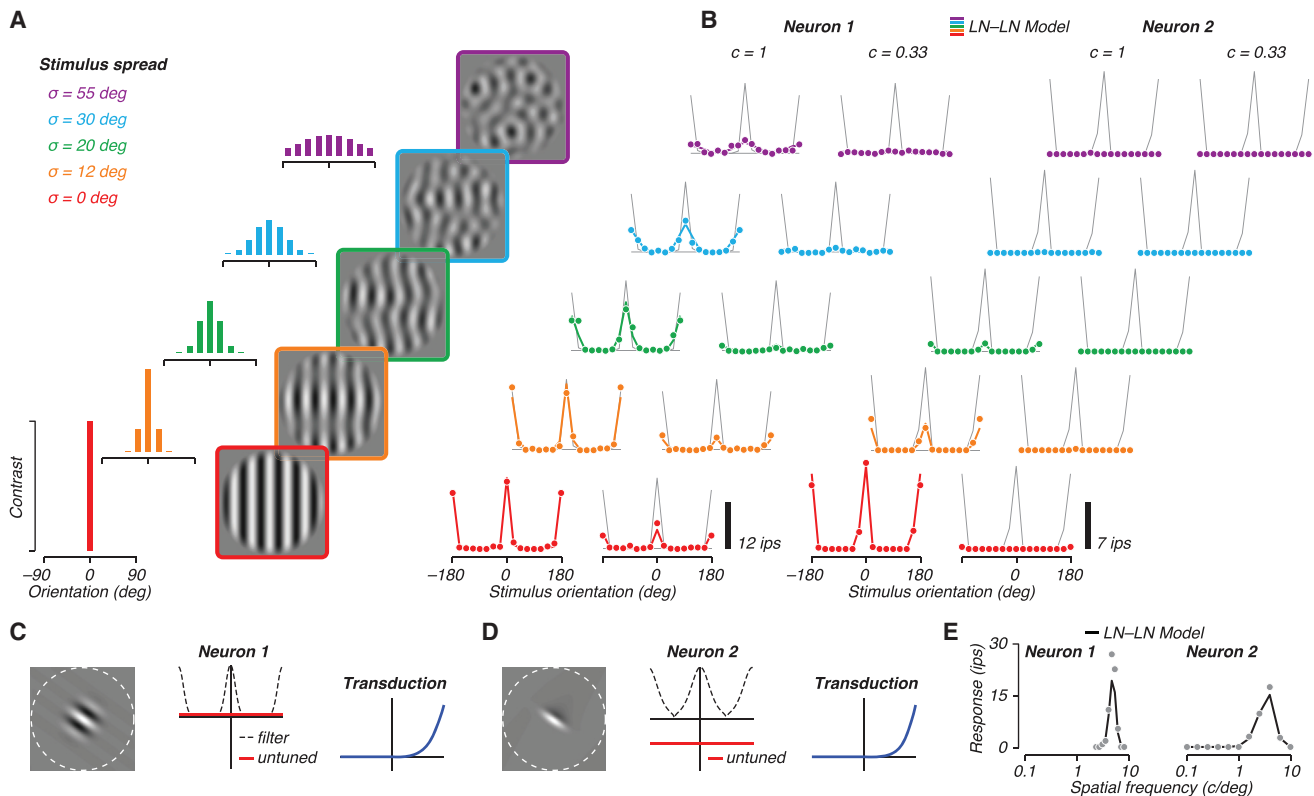


Figure 2. Characterization of the LN-LN Cascade Model

(A) Each stimulus consisted of a windowed sum of drifting sinusoidal gratings, with drift directions drawn from one of five unimodal distributions (colored histograms), centered around a randomly selected primary direction. All gratings had the same spatial frequency, optimized for the cell. An example movie frame is shown, corresponding to each distribution.

(B) Measured and predicted responses of two example V1 neurons plotted against modal stimulus direction for high- and low-contrast stimuli. Colored points indicate mean spike counts, colored lines indicate predicted responses of the fitted model. As a reference, thin gray lines reproduce the model responses to a single high-contrast grating.

(C and D) Recovered model components for the two example neurons. Left: best-fitting Gaussian derivative filter (white dashed line illustrates the stimulus size). Middle: orientation tuning of the filter (dashed black line) and gain of the untuned component (thick red line). Right: response nonlinearity.

(E) Comparison of model-predicted spatial frequency tuning curves to data, for both example neurons.

examined the relationship between each of the three components—the selectivity of the filter, the gain of the untuned component, and the steepness of transduction—and the observed orientation selectivity.

Recall Figure 1E, in which we illustrated three families of tuning curves created by changing different mechanisms. In Figure 3B, we plot three parametric families showing the effects of these manipulations. The upper curves each pass through points corresponding to the OSI values of the tuning curves in Figure 1E. The other curves in each panel of Figure 3B show the effects of changing the values of the unvaried mechanisms (over the range of values fit to the data), revealing the complex interdependence of the effects of these parameters on the OSI.

Now compare these patterns with the relationships between the three model components and the observed OSI (Figure 3C). In V1 (upper panels), filter selectivity and orientation selectivity are well correlated ($r = 0.59$, $p < 0.001$, Spearman correlation). But the other two model components also correlate with OSI: $r = -0.55$ ($p < 0.001$) for the untuned component, and $r = 0.46$ ($p < 0.001$) for the response exponent. Results from V2

(lower panels) largely recapitulate those from V1: variation in the filter stage has the strongest association with tuning diversity ($r = 0.52$, $p < 0.001$), followed by the untuned component ($r = -0.45$, $p < 0.01$), and the response nonlinearity ($r = 0.32$, $p = 0.051$).

Though the three components have similar correlations with OSI in V1 and V2, their actions are not all identical. In V1, the untuned component excites most neurons and therefore reduces orientation selectivity (0.067 ± 0.033 , mean gain of untuned component \pm SEM); in V2, it mostly inhibits responses (-0.10 ± 0.044) and consequently sharpens tuning. This difference remains significant when controlling for the different distributions of orientation selectivity in these samples of V1 and V2 neurons (ANCOVA, $F_{1,76} = 6.5$, $p = 0.013$).

To compare the explanatory power of each component, we fit restricted versions of the model in which only one parameter was allowed to vary, while the two other parameters were held fixed at the median of values estimated across all neurons under the full model (filter aspect ratio = 2.1, filter derivative order = 0.75, gain untuned component = 0, response exponent = 2.15).

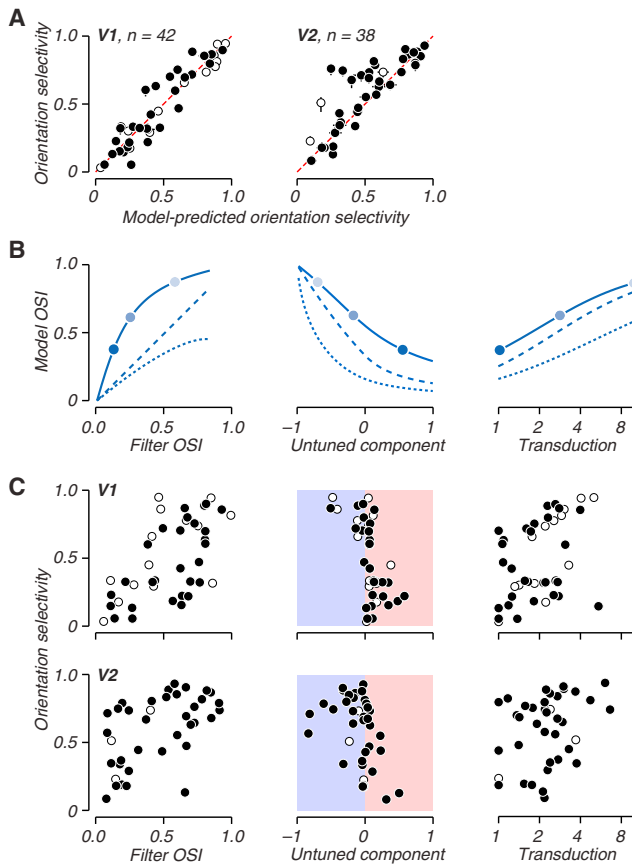


Figure 3. Analysis of Recovered Model Components

(A) Comparison of measured and predicted orientation selectivity for V1 (left) and V2 (right) neurons. Confidence intervals, where visible, illustrate the inter-quartile-range.

(B) Orientation selectivity for three families of model neurons. Within each family, one model component that controls orientation selectivity varies. Different line types correspond to different values for the non-varied components. The left panel summarizes responses for a range of filter OSI values with an untuned component and response exponent of $(-0.2, 3)$ (solid line), $(0, 1)$ (dashed line), and $(0.2, 1)$ (dotted line). The middle panel summarizes responses for a range of untuned component values with a filter OSI and response exponent of $(0.25, 3)$ (solid line), $(0.33, 1)$ (dashed line), and $(0.15, 1)$ (dotted line). The right panel summarizes responses for a range of response exponents with a filter OSI and untuned component of $(0.25, -0.2)$ (solid line), $(0.33, 0.2)$ (dashed line), and $(0.15, 0)$ (dotted line). Points correspond to the color-matched tuning curves shown in Figure 2B.

(C) Orientation selectivity plotted against filter selectivity (left), the untuned component (middle), and the response exponent (right) for a population of V1 (top) and V2 (bottom) neurons. Open circles indicate simple cells, filled circles complex cells. The colored background in the middle column refers to untuned excitation (red) versus inhibition (blue).

Consider the goodness of fit of the different models (Figure 4A). The full model (ordinate) predicts the measured tuning curves for the different mixture stimuli better than any of the restricted models (top row). This means that tuning diversity in V1 and V2 originate from multiple sources. Among the restricted models, the one in which the filter width is allowed to vary (abscissa) has the highest predictive accuracy (left column). This model is also the one that best explains the measured diversity in orienta-

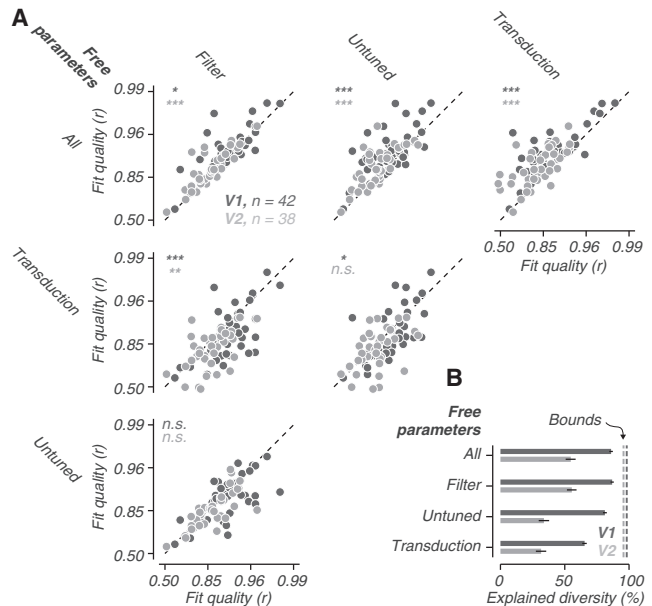


Figure 4. Comparison of Restricted Models

(A) Comparison of fit quality of the full model against various restricted models, over populations of V1 (dark gray) and V2 (light gray) cells. Fit quality is captured with the correlation of model-predicted and measured responses for a held-out set of data and is plotted on Fisher transformed axes. For restricted models, only one model component is free to vary across neurons to account for variable orientation selectivity.

(B) Fraction of tuning diversity in V1 (dark gray) and V2 (light gray) populations that is explained by full and restricted models. Explained diversity expresses the fraction of across-neuron variance in OSI that is accounted for by the model. It is bounded below 100% because neural responses are variable (see Experimental Procedures). Confidence intervals, where visible, illustrate the inter-quartile-range of the fraction of explained diversity. n.s., not significant. * $p < 0.05$; ** $p < 0.01$, *** $p < 0.001$ (Wilcoxon signed-rank test).

tion selectivity (Figure 4B), which suggests that variability in the filtering stage is the main source of tuning diversity. The restricted model in which the untuned component is allowed to vary comes close in performance to the best restricted model in V1 (81.6% versus 87% explained diversity) but not in V2 (33.7% versus 55.5%). This likely results from the different actions of the untuned component in the two areas (Figure 3C). We conclude that the tuning diversity in visual cortex is due to multiple factors, of which the most important is variability in the filtering stage.

Model-Based Analysis of Spatial Frequency Selectivity

In the LN-LN model, selectivity for orientation and spatial frequency both emerge from a linear filter (Figures 1F–1H), and both are modified by the untuned component and the response nonlinearity. Cortical neurons differ in the sharpness of their spatial frequency tuning (Figure 2E) (De Valois et al., 1982; Xing et al., 2004). Our experiments primarily probed orientation tuning, but we also measured spatial frequency tuning, and we wondered whether our model could also account for its diversity. For single gratings, the spatial frequency tuning of the unrestricted model was strongly correlated with the measured tuning




Figure 5. The LN-LN Model Explains Diversity in Spatial Frequency Tuning

(A) Distribution over V1 (top) and V2 (bottom) cells of the correlation of the model-generated and the measured spatial frequency tuning curves, plotted on a Fisher transformed axis.

(B) Measured spatial frequency bandwidth (full bandwidth at half height) is plotted against model bandwidth for a population of V1 (left) and V2 (right) neurons. Confidence intervals, where visible, give the inter-quartile-range.

(C) Comparison of explained bandwidth diversity for a family of restricted models in a population of V1 (dark gray) and V2 (light gray) neurons. Confidence intervals, where visible, illustrate the inter-quartile-range.

in V1 ($r = 0.95$, median correlation between the predicted and measured mean spike counts) and V2 ($r = 0.92$, Figure 5A). The spatial frequency bandwidth at half height computed from the unrestricted model therefore reasonably approximates the observed bandwidth in V1 (Figure 5B, 61.5% of diversity explained). The model gives a poorer account of V2 diversity (30% explained), in part due to noisy data (the level of explainable diversity in V2 is low, as indicated by the reduced bound in Figure 5C). By analogy with the procedure we described earlier for orientation, we also asked how well restricted versions of the model account for spatial frequency tuning diversity. In V1 and V2, the restricted models all explain a similar fraction of the measured diversity (Figure 5C). In V1, restricted models explain less diversity than the full model (Figure 5C), suggesting that differences in spatial frequency tuning arise from multiple sources. With the caveat that we only made measurements with single gratings and not spatial frequency mixtures, we conclude that spatial frequency tuning diversity in visual cortex arises from the same factors that create diversity in the orientation domain.

As others have previously reported, selectivity for both orientation and spatial frequency are also related to other elements of receptive field structure, such as preferred spatial frequency, size, and phase sensitivity (“simple” versus “complex” cells) (De Valois et al., 1982; Xing et al., 2004). In our sample, the strongest of these effects—and the only one to achieve statistical significance—was the relationship between orientation selectivity and preferred spatial frequency. Interestingly, the variation in orientation selectivity with preferred spatial frequency seems to result from two distinct mechanisms, one captured in the model by its linear filter and one captured by the transduction nonlinearity (Figure S2).

Quantifying Orientation Information Carried by Individual Neurons

Neurons with different degrees of selectivity may be adapted to signal information from particular types of image content. Consider the responses of a broadly tuned example neuron to stimuli with different orientation dispersion (Figure 6A). This cell responded less selectively but more robustly to orientation mixtures than the narrowly tuned example cells shown in Figure 2B. This suggests that there exists a trade-off between selectivity and robustness that might enable less selective neurons to transmit more information about orientation mixtures.

To quantify how much orientation information individual neurons transmit, we estimated Fisher information from the

observed responses (Figure 6A; see Experimental Procedures). This statistic expresses the accuracy with which stimulus orientations are represented and provides a lower bound on discrimination thresholds of an observer whose decisions are based solely on this neuron (Serriès et al., 2009). We computed how much information each neuron carries, taking into account the number of neurons needed to cover the entire orientation domain (see Experimental Procedures). We refer to this statistic as Fisher information per neuron (FI/N).

The relationship between FI/N and orientation selectivity is shown in Figure 6B for five values of orientation dispersion. For simple gratings (left, red), as one might expect, FI/N increases with selectivity, so that more selective neurons convey more information about orientation ($r = 0.41$, $p < 0.001$). For broadly dispersed stimuli (right, purple), however, the relationship is reversed—the less selective neurons carry more information about orientation ($r = -0.30$, $p = 0.006$). The reasons are twofold. First, the response of highly selective neurons falls when the orientation distribution is broad (Figure 2B). Less selective neurons, on the other hand, remain responsive and more or less equally selective for all values of orientation dispersion (Figure 6A). Second, narrow populations require more cells to cover the orientation domain than broad populations. For broadly dispersed stimuli, the Fisher information for narrow and broad populations is similar, but the FI/N decreases with orientation selectivity. We conclude that narrowly tuned neurons transmit more information about the orientation of simple oriented stimuli, while broadly tuned neurons transmit more information about orientation mixtures.

Diversity of Orientation Information in the Natural Environment

What perceptual role might tuning diversity serve? Consider the distribution of image energy across the Fourier spectrum in small image patches sampled from artificial and natural stimuli (Figures 7A and 7B). For a sine wave grating, all energy is concentrated around a single orientation; for a random noise sample, the energy is typically distributed across many orientations (Figure 7B). In the natural environment, local image energy is sometimes concentrated at a single orientation (e.g., at the edge of an object) but is often more broadly distributed (e.g., at a corner or an isolated bright or dark spot). To characterize this with a single number, we computed an orientation concentration index (OCI) using the same formula we used to measure orientation selectivity for cortical neurons. Across a large number of local patches

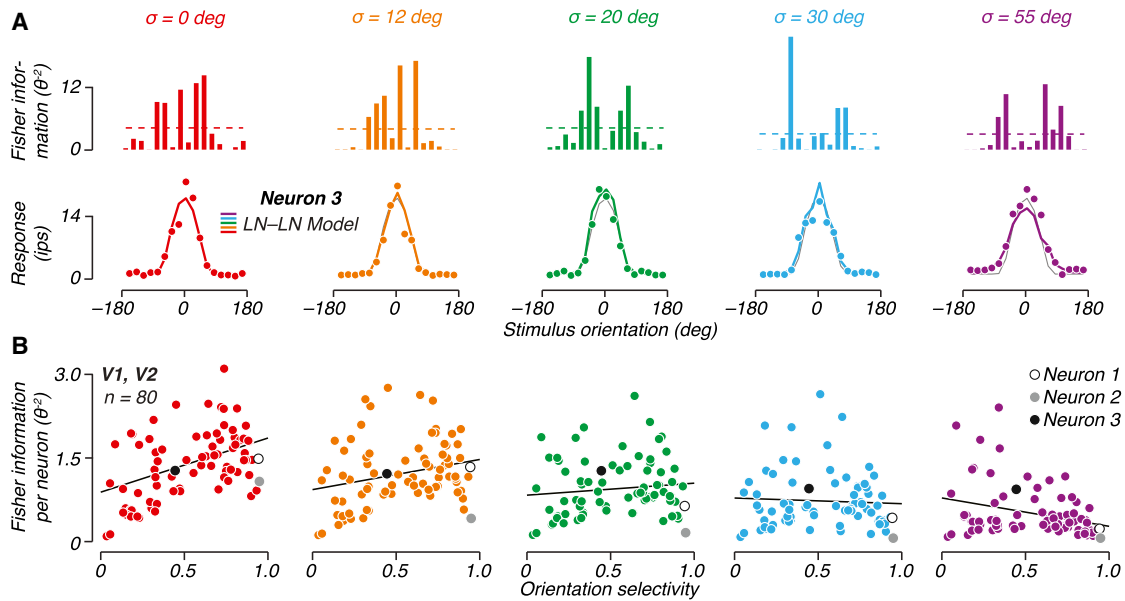


Figure 6. Tuning Diversity and Orientation Coding

(A) Firing rate (bottom) and Fisher information (top) are plotted against stimulus orientation for an example V2 neuron, for stimulus ensembles with five different dispersion values (σ). Only data from the high-contrast conditions are shown. Top: dashed lines indicate mean Fisher information. Bottom: colored points are data; colored lines are model-predicted responses. For reference, the thin gray line reproduces the model responses to a single grating.

(B) Fisher information per neuron is plotted against grating orientation selectivity for all V1 and V2 neurons. Black lines show the best-fitting linear relationship. The three example cells from Figures 2 and 6A are highlighted.

drawn from natural images, orientation concentration is approximately uniformly distributed and strikingly similar to the distribution of orientation selectivity in V1 (Figure 7C).

Tuning Diversity and Visual Coding

Is the similarity between the diversity of the visual environment and the tuning of neurons in primary visual cortex mere coincidence? Theoretical studies have shown that natural scenes can be described with sets of elementary functions whose linear representation of image information is statistically independent (Bell and Sejnowski, 1997) and sparse (Olshausen and Field, 1996); these functions are often oriented and resemble cortical simple cell receptive fields. These analyses do not, however, predict filters with the same kind of diversity evident in biological measurements (Ringach, 2002). Perhaps tuning diversity gives a coding advantage that is not revealed by maximizing independence or sparseness but relates instead to the perceptual usefulness of the image representation.

We asked how well differently composed populations of model neurons represent differences in local image patches drawn from the same three stimulus families (Figure 8A). Our synthetic populations were composed of LNP neurons and equated with respect to the number of neurons (45 neurons), total spike rate (average of ten spikes per neuron per image), the distribution of preferred orientation (uniform), and the response nonlinearity (a power law with an exponent of 3.3, the mean value found in a study on the nonlinear properties of spike threshold) (Priebe et al., 2004). The populations consisted of neurons that varied in orientation bandwidth; we created this variation by changing the linear filter (Figure 8B). We considered three homogeneous

populations of 45 broad, medium, or narrowly tuned filters and one diverse population with 15 broad, 15 medium, and 15 narrowly tuned filters. The images to be discriminated were a randomly sampled patch from a particular image family (gratings, noise, natural images) and its nearest neighbor in the population of random samples of that family ($n \sim 10,000$). We estimated discriminability between these stimuli from the square root of Fisher information, a generalization of the d' statistic used in signal detection theory (see Experimental Procedures).

As expected, the allocation of resources (neurons and spikes) affects the efficiency of visual coding for different stimulus families. For example, discrimination of sine wave gratings is best supported by a homogeneous population of narrowly tuned neurons (Figure 8B). This makes sense given that the tuning properties of this population best match the high degree of orientation concentration characteristic of grating stimuli (Figure 7C). Likewise, discrimination of noise images is better supported by a homogeneous population of broadly tuned neurons than by a population of narrowly tuned neurons (Figure 8B). For the latter task, the performance of the diverse population is similar to the best homogeneous population (Figure 8B). But for natural images, the diverse population outperforms all homogeneous populations (Figure 8B).

DISCUSSION

Cortical neurons differ substantially in the precision of their orientation tuning, and we have undertaken a systematic exploration of the origins and consequences of this diversity. Our results are expressed in terms of a functional model of a kind that has

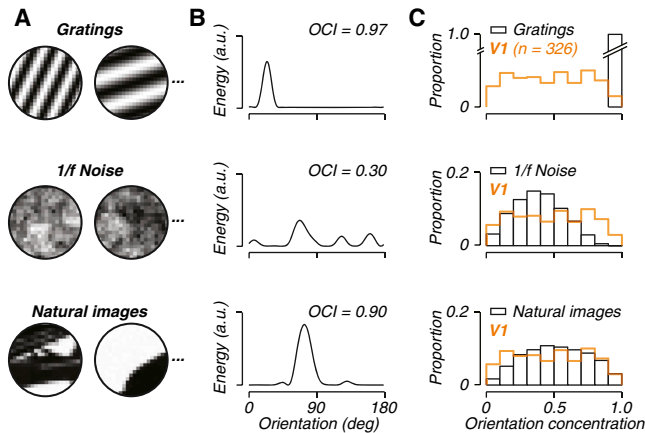


Figure 7. Distribution of Local Orientations for Three Image Ensembles

(A) Example image patches taken from an ensemble of grating stimuli (top), 1/f noise images (middle), and natural photographic images (bottom).

(B) The distribution of image energy in the orientation domain for the left-hand patches in (A).

(C) The distribution of OCI (one minus circular variance) across a large number of patches drawn from each ensemble. Orange distributions recapitulate the distribution of orientation selectivity for our population of V1 cells.

proved capable of explaining much of the response behavior of cortical cells (Priebe and Ferster, 2012), consisting of a tuned linear filter, an untuned input, and an output nonlinearity. Interpreting neuronal responses with this model leads to the conclusion that tuning diversity arises from neuron-to-neuron variability in each of these components, with the most important role ascribed to variability in the linear filter.

All models are approximate (Box, 1976). Moreover, the mapping from model components to underlying biophysical mechanism is uncertain, and speculations about specific circuit elements are necessarily tentative. There is some reason to believe that the filter stage of our model is instantiated by the aggregated feedforward, excitatory input (Chapman et al., 1991; Reid and Alonso, 1995; Jin et al., 2011). More complex interpretations are possible, because recurrent cortical networks may also act as linear filters (Wieland et al., 2001; Tao et al., 2004), and there is some evidence that the required pattern of specific cortico-cortical connections exists (Nauhaus et al., 2009; Cossell et al., 2015). As we have noted, the mechanistic interpretation of the untuned component depends on its sign, because it can signify either a broad pool of excitatory input or an untuned suppression of the kind described by Xing et al. (2005). We must also acknowledge that our decision to make this model component perfectly untuned is open to question, especially in light of experimental evidence for tuned suppression (Xing et al., 2005). We elected to use untuned suppression in part because the added complexity was not needed for the model to fit our data (73 of our 80 cells pass an absolute goodness of fit test; see Results) and in part because the key elements of our findings are not likely to be sensitive to the exact choice of model architecture. The essence of a functional model is to find an economical formulation that explains the data, and our model meets that test. Perhaps most straightforward, the output nonlin-

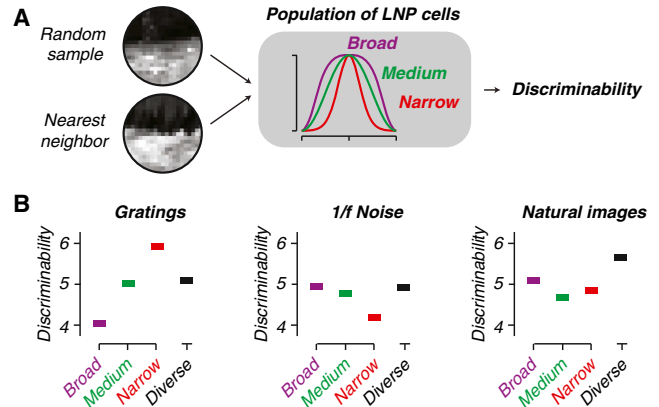


Figure 8. Tuning Diversity Is Beneficial for Visual Coding

(A) Population coding analysis. Each population of LNP neurons was used to estimate the discriminability of sample image patches and their nearest neighbors drawn from a large ensemble of patches ($n \sim 10,000$). Three different orientation tuning curves for the linear filter are illustrated in the middle panel. (B) Median discriminability of gratings (left), 1/f noise (middle), and natural images (right) for three homogeneous populations of neurons—exclusively consisting of broad, medium, or narrow filters—and one diverse population of neurons. The precise values of discriminability depend on the parameters chosen for the simulation (neural population size, firing rate, image contrast and size, and number of image samples), but the relative performance does not.

earity presumably captures the final relationship between excitation and firing rate (Priebe and Ferster, 2012).

The origins of orientation selectivity remain a topic of vigorous debate. At issue is the question of whether cortical neurons derive their selectivity from simple feedforward filtering or from a more complex feedback process that encompasses intracortical inhibition. Our work offers a novel perspective by asking, in functional terms, what causes the selectivity of cells to differ from one another? We find that the dominant source of variation is the linear filter in our model. Indeed, a restricted model in which only the linear filter varied gave a good account of tuning diversity. Our results thus support the view that neurons primarily derive their stimulus selectivity from the way they gather their inputs.

The model was designed to account for orientation selectivity but also provides a natural account of variations from neuron to neuron in spatial frequency tuning. Orientation selectivity correlates with spatial frequency selectivity and with spatial frequency preference (De Valois et al., 1982; Xing et al., 2004) (see also Figure S2). Our experiments were designed to probe orientation selectivity. Nevertheless, our results and our model suggest that the linear filter and the output nonlinearity jointly drive the relationship between preferred spatial frequency and orientation selectivity (see Supplemental Information). Future work could investigate this in more detail by fitting the model to responses elicited by both orientation and spatial frequency mixtures.

Tuning diversity is not limited to V1 (Albright 1984; Levitt et al., 1994). It seems likely that V2 largely inherits this diversity from its V1 inputs. We compared the recovered distribution of model

components in V1 and V2 and found them similar, except for the untuned component (Figure 3C). In V1, the untuned component on average has an excitatory role, reflecting the fact that many neurons respond above baseline to all orientations. In V2, the untuned component is on average inhibitory. A straightforward interpretation is that the untuned component of early cortical response is gradually removed by inhibitory cortical mechanisms, which also explains the slight increase in orientation selectivity from V1 to V2.

Tuning diversity is not an inescapable consequence of selectivity. For example, directional selectivity in area MT is remarkably uniform (Zeki, 1974; Maunsell and Van Essen, 1983; Albright, 1984). This raises the question of why diversity exists. Previous work suggests that diversity can be used to overcome the harmful effects of collective noise fluctuations in neural populations, an advantage for stimulus encoding by neural populations (Shamir and Sompolinsky, 2006). Other constraints and objectives may also contribute to the diversity of V1 selectivity. For example, V1 computations are presumably constrained by the need to minimize cortical wiring and metabolic cost (Attwell and Laughlin, 2001; Chklovskii et al., 2002). In addition, V1 neurons encode variables other than orientation, such as spatial frequency, speed, spatial position, phase, and size. Simultaneous encoding of these attributes might interact with the representation of orientation information and might thus contribute to its diversity.

Here we have presented evidence that tuning diversity in visual cortex supports efficient sensory coding for a variety of inputs. Stimuli containing a single orientation are better encoded by highly selective neurons, while orientation mixtures are better encoded by less selective neurons (Figure 6B). Our synthetic stimuli mirror a neglected feature of the visual environment—small patches of natural images exhibit notable diversity in orientation content (Figure 7). We therefore asked whether the diversity we measured in our neural populations confers any advantage and discovered that populations of neurons with dispersed tuning indeed convey more information about the local content of natural images than uniform populations (Figure 8C). Our results suggest that diverse sensory mechanisms are the brain's solution to the challenge of processing information from a diverse sensory world.

There is an ongoing debate about the use of natural images to probe mechanism in visual neuroscience. Our understanding of visual function is expressed in the models we formulate, and developing those models in an efficient and principled manner requires carefully designed and well-controlled stimuli, yet the ultimate test of such models is their ability to predict responses in the natural world (Felsen and Dan, 2005; Rust and Movshon, 2005). Here, we used well-specified artificial stimuli, tailored to each cell, to answer a specific question about neural computation—something we could not have achieved easily with natural stimuli. Then we used natural stimuli to determine how this neural computation plays against the richness of the natural world—something we could not have done with artificial stimuli. This combination of artificial and natural stimuli offers an effective way to study the brain in its natural context through the dynamic interplay of theory and experiment.

EXPERIMENTAL PROCEDURES

Surgical Preparation

We recorded from six anesthetized, paralyzed, adult macaque monkeys (one female *Macaca nemestrina* and five male *M. cynomolgus*). Surgical preparation of animals and single-unit recordings have been reported in detail previously (Cavanaugh et al., 2002). Briefly, experiments typically lasted 5 to 6 days, during which we maintained anesthesia with infusion of sufentanil citrate (6–30 $\mu\text{g kg}^{-1} \text{h}^{-1}$) and paralysis with infusion of vecuronium bromide (Norcuron; 0.1 $\text{mg kg}^{-1} \text{h}^{-1}$) in isotonic dextrose-Normosol solution. We monitored vital signs (heart rate, blood pressure, lung pressure, end-tidal pCO_2 , EEG, body temperature, urine flow, and osmolarity) and maintained them within appropriate physiological ranges. Pupils were dilated with topical atropine. The eyes were protected with gas-permeable contact lenses and refracted with supplementary lenses chosen through direct ophthalmoscopy. All procedures were conducted in compliance with the *NIH Guide for the Care and Use of Laboratory Animals* and with the approval of the New York University Animal Welfare Committee.

Unit Recording

Extracellular recordings were made with quartz-platinum-tungsten microelectrodes (Thomas Recording), advanced mechanically into the brain through a craniotomy and small durotomy. We distinguished V1 from V2 on the basis of depth from the cortical surface and changes in the receptive field location of the recorded units. We made recordings from every single unit with a spike waveform that rose sufficiently above noise to be isolated. Stimuli were presented in random order. Data are reported from every unit for which isolation could be maintained for at least three stimulus repetitions (median: ten repetitions).

Visual Stimulation

We presented visual stimuli on a gamma-corrected CRT monitor (Eizo T966; mean luminance, 33 cd/m^2) at a resolution of 1,280 \times 960 with a refresh rate of 120 Hz. Stimuli were presented using *Expo* software (<http://corevision.cns.nyu.edu>). For each isolated unit, we first determined eye dominance and occluded the non-preferred eye. We then presented suitably vignetted sinusoidal grating stimuli to map each cell's receptive field and determined its preferred size and drift rate. We characterized neuronal selectivity for spatial frequency. Thereafter, we presented the mixture stimuli at optimal spatial frequency in a window of the preferred size. Stimuli were presented in random order, each for 1,000 ms interleaved with a blank screen for 500 ms.

We generated mixture stimuli by summing nine sinusoidal gratings whose orientations were spaced at 20° intervals and whose contrasts followed a circular Gaussian profile (illustrated in Figure 2A). The drift rate of each stimulus component was selected at random from a Gaussian distribution centered on the preferred rate, with an SD equal to 1/5 of this value, resulting in an incoherently drifting mixture. Single-frame images of example stimuli are shown in Figure 2A. The spatiotemporal luminance profile of the stimulus is fully specified by spatial frequency, ω ; mean temporal frequency, $\mu_{\omega t}$; and peak, dispersion, and amplitude of the contrast profile in the orientation domain (μ_{θ} , σ_{θ} , and C_{θ} , respectively):

$$S(x, y, t; \omega, \mu_{\omega t}, \mu_{\theta}, \sigma_{\theta}, C_{\theta}) = W(x, y) \sum_{k=-4}^4 c_{\theta} f\left(\frac{\pi k}{9\sigma}\right) \cos\left(\omega\left(x \cos\left(\mu_{\theta} - \frac{\pi k}{9}\right) + y \sin\left(\mu_{\theta} - \frac{\pi k}{9}\right) + \mu_{\omega t} \left(1 + \frac{n_k}{5}\right) t\right)\right)$$

Here, x and y are spatial coordinates, t is time, $w(x, y)$ is a soft-edged window that tapers contrast over 1/8 of the aperture diameter, f is a circular Gaussian weighting function, and n_k are random samples drawn from a normal distribution of unit variance. The complete stimulus set contained all combinations of 16 peak orientations ($\mu_{\theta} = \pi k/8$, $k \in \{0, 1, \dots, 15\}$), five orientation dispersions ($\sigma_{\theta} \in \{0.05, 0.1, 0.2, 0.35, 1\}$), and two contrasts ($C_{\theta} \in \{0.33, 1\}$).

Analysis of Neuronal Response

We counted spikes in a 1,000-ms window following response onset. We chose a latency for each cell by maximizing the stimulus-associated response

variance (Smith et al., 2005). For each cell, we computed an OSI, defined in Results. We estimated the variability of the OSI through a bootstrap analysis. For each iteration of the bootstrap, we computed the OSI for a randomly resampled dataset. For model-predicted responses, we computed the OSI from simulated spike counts, sampled from a modulated Poisson model (Goris et al., 2014). The number of simulated spike counts was matched to the number acquired experimentally. Any bias in the OSI resulting from trial-to-trial variability is thus approximately matched for actual and model-predicted estimates.

We computed the full bandwidth at half height of the spatial frequency tuning curve from a suitable descriptive function fit to the mean response. We estimated the variability of the bandwidth estimate through a bootstrap analysis, using the same procedure as described for the OSI.

LN-LN Model

We fit a simplified two-stage feedforward model to data from individual V1 and V2 neurons (Figure 1D). The model describes how a mixture stimulus composed of superimposed drifting gratings is transformed into the firing rate of a cortical cell. Stimuli are processed in two parallel channels, each composed of a linear filter whose output is half wave rectified. Both channels are tuned for temporal frequency, but only one channel is tuned for orientation and spatial frequency.

The linear filter's spatial profile is given by a derivative of a 2D Gaussian function (specific examples shown in Figure 1F; construction of the filter is explained in detail in the Supplemental Information). At the preferred orientation, the spatial frequency selectivity of the filter solely depends on the order of the derivative, b (Figure 1H):

$$r_{\omega}(\omega; \omega_0, b) \propto \left[(\omega/\omega_0) e^{-\frac{1}{2}(\omega/\omega_0)^2} \right]^b,$$

where ω is stimulus spatial frequency and ω_0 the filter's preferred spatial frequency. At the preferred spatial frequency, the orientation selectivity of this filter depends on the aspect ratio of the Gaussian, α ; the order of the derivative, b ; and the directional selectivity, d (Figure 1G):

$$r_{\theta}(\theta; \theta_0, \alpha, b, d) \propto \left(1 + \frac{d}{2} (\text{sgn}(\cos(\theta - \theta_0)) - 1) \right) \cdot \left[\cos(\theta - \theta_0) \times e^{-\frac{1}{2}(1-\alpha^2)\cos^2(\theta - \theta_0)} \right]^b,$$

where θ is stimulus orientation, θ_0 the filter's preferred orientation, and parameter $d \in [0, 1]$ determines the direction selectivity. The function $\text{sgn}(\cdot)$, computes the sign of its argument (returning -1 for negative and $+1$ for positive), and the initial parenthesized expression serves to multiply the half of the tuning curve in the non-preferred direction by $(1 - d)$.

If the Gaussian kernel is circularly symmetric ($\alpha = 1$), orientation tuning is simply described with a cosine function raised to a power; such functions can perfectly tile the orientation domain. We exploited this property to compute the normalization signal. The untuned channel has a linear filter whose amplitude is uniform over orientation and spatial frequency. For both channels, temporal frequency tuning is described with a function of the same form as the spatial frequency tuning function and is six octaves wide at half height. The time-varying responses of the tuned and untuned channel, $C_T(t)$ and $C_U(t)$, are obtained by computing the half-rectified linear filter response for every movie frame.

The second stage of the model consists of a linear combination of the two channel responses, followed by divisive normalization:

$$N(t) = \frac{C_T(t) - wC_U(t)}{\sigma + \frac{1}{k} \sum_k C_k(t)},$$

where w controls the gain of untuned channel. The normalization signal consists of the sum of a stimulus-independent constant, σ , and the pooled activity of a diverse set of neurons with spatially distributed receptive fields, whose responses, $C_k(t)$, are computed in the same way as the tuned channel responses. The normalization pool includes both orientation- and direction-selective cells, drawn from five bandwidth families (25° , 40° , 60° , 90° , and 120° at half height),

each tiling the orientation domain. The normalization pool was fixed for all model cells and was not fit to the data.

Finally, the normalized response is transformed into a firing rate. This requires the inclusion of a spontaneous discharge. For some cells, spontaneous activity can be suppressed by stimuli that fail to excite the neuron. But for other cells, the spontaneous discharge simply adds to the stimulus-driven response. We therefore included two forms of spontaneous discharge, governed by two separate parameters (ε_1 and ε_2). The final firing rate is computed using a power-law nonlinearity:

$$R(t_i) = \varepsilon_1 + \beta \left[\frac{\varepsilon_2 + C_T(t_i) - wC_U(t_i)}{\sigma + \frac{1}{k} \sum_k C_k(t_i)} \right]^q.$$

Estimating Model Parameters for Individual Cells

In total, the LN-LN model has 11 free parameters: five filter parameters (orientation preference θ_0 , spatial frequency preference ω_0 , spatial aspect ratio α , derivative order b , and directional selectivity d), one untuned channel parameter (gain w), one parameter for the nonlinearity (exponent q), and four parameters controlling response range and amplitude (semi-saturation constant σ , response scale β , and maintained discharge ε_1 and ε_2). The exponent q captures the translation from excitation to firing rate. Priebe et al. (2004) estimated exponents from intracellular recordings from cat V1 and reported values between 1.5 to 6.5 (mean: 3.3, SD: 0.9). We incorporated this knowledge in the form of a log-normal prior distribution on the exponent (mean: 3.3, SD: 1). For each cell, we optimized the model parameters by maximizing the product of the prior on the exponent and the likelihood over the observed data, assuming spike counts arise from a modulated Poisson model. One additional free parameter, σ_G^2 , describes across-trial fluctuations in neural response gain (Goris et al., 2014). We computed the model response for every presented movie frame and took the average over the one second stimulus movie as its predicted firing rate. We used a simplex algorithm and multi-start fitting procedure with randomized starting values to find the best fitting parameters. The distributions of the fitted model parameters that directly control visual response properties are shown in Figure S3.

The model contains three mechanisms (governed by four parameters) that control orientation selectivity (the spatial aspect ratio α , the derivative order b , the gain of the untuned channel w , and the response exponent q). To test the explanatory power of the corresponding mechanisms directly, we fit restricted versions of the model to the data (Figures 4 and 5). In this analysis, two of the three selectivity mechanisms were not fit to the data but were assigned a value that reflected the median estimate under the full model. The third component was fit to the data. Changing the nonlinear selectivity parameters typically affects the model's response scale. To maximize the model's ability to capture the data, we simultaneously optimized the parameters controlling maintained discharge and response scale. The values of preferred direction of motion and the variance of the gain were not refit—instead we used the values fit under the full model.

Goodness of Fit

We assessed model performance in three different ways. First, we evaluated the model's absolute goodness of fit. We fit the unrestricted LN-LN model to the full set of responses and compared the log-probability of the observed data under the fitted model with a reference distribution, obtained by computing the log-probability of 1,000 synthetic datasets, sampled from the fitted model (Goris et al., 2014). Each synthetic dataset contained the same number of responses as the experimentally acquired data. The reference distribution thus characterizes the expected distribution of log-probability if the model correctly describes the cell's response characteristics.

Second, we evaluated the models' relative goodness of fit. We compared the unrestricted and restricted versions of the model in a 10-fold cross-validation analysis. We fit each model to a subset of the data (all the trials that make up the tuning curves for nine out of ten mixture stimuli) and predicted the tuning curve for the remaining stimulus type. We repeated this procedure ten times, once for each stimulus type, to obtain a full set of predictions. The training

set always included the spatial frequency tuning curve. We expressed model performance as the correlation between the predicted and observed mean spike counts (Figure 4A).

Third, we evaluated how well the models explain diversity in orientation selectivity (Figures 3A and 4B). We used the predicted tuning curve for the high-contrast, single-component grating obtained in the cross-validation analysis and computed a statistic (explained diversity) that expresses the fraction of across-neuron variance in OSI that is accounted for by the model. Explained diversity, σ_E^2 , is defined as

$$\sigma_E^2 = 1 - \frac{SS_R}{SS_T},$$

where SS_R is the residual sum of squares of the predicted single grating OSI and SS_T is the total sum of squares of the observed OSI. If the model prediction were perfect for every cell, this diversity measure would equal 100%. To obtain confidence intervals, we computed explained diversity for 1,000 randomly resampled and model-simulated datasets. Because the observed OSI is variable, not all diversity is explainable. We estimated the explainable diversity by calculating how much diversity is explained by an "oracle" model in which each neuron's OSI is predicted from a set of randomly sampled responses generated by that neuron (the number of sampled responses matched the number acquired experimentally). We used the same procedure to evaluate how well the model explains diversity in spatial frequency bandwidth (Figure 5C).

Estimation of Fisher Information for V1 and V2 Neurons

In the orientation domain, Fisher information can be simply written as a function of the measured tuning curve, $h(s)$:

$$I_f(s) = \frac{h'(s)^2}{h(s)},$$

where $h'(s)$ is the derivative of the tuning curve (Seung and Sompolinsky, 1993). We estimated how much information each neuron transmits in a hypothetical population of neurons composed of shifted copies of the recorded neuron, each responding independently. To make it easier to compare the effect of selectivity across neurons, we normalized responses by the largest observed response; we verified that there was no systematic relationship between selectivity and peak firing rate ($r = -0.16$, $p = 0.15$). We computed the average Fisher information across stimulus orientations and divided this average by the number of neurons needed to tile the orientation domain. To estimate tiling density, we derived the relation between the circular variance of the stimulus-driven response (the observed response with the weakest stimulus-response subtracted), V , and the full bandwidth of the tuning curve at half height, B , under a simple descriptive model of orientation tuning. In particular, for a tuning curve given by $h = [\cos(\theta - \theta_p)]^{2b}$, the circular variance of the responses only depends on the exponent: $V = 1/(1 + b)$. This is also true of the full bandwidth at half height: $B = 2\cos^{-1}(0.5^{1/2b})$. Combining these gives a simple formula relating bandwidth to circular variance: $B = 2\cos^{-1}(0.5^{V/(2 - 2V)})$. For every cell, we estimated V from the raw data, used this formula to obtain a value for B , from which we defined the tiling density to be $2\pi/B$.

Orientation Distribution in Images

We measured the orientation distribution in families of natural and synthetic images. We drew natural images from a publicly available database (Olmos and Kingdom, 2004) containing 653 2,560 × 1,920 TIF photographs of natural scenes, with a pixel size of 0.0028 deg. We linearized the natural images in RGB color space using camera parameters provided with the database, averaged the three color channels, and normalized the image by the mean luminance. We randomly selected 20 patches of 21 × 21 pixels from each image and excluded a small number of patches whose contrast was too low to evoke reliable filter responses.

These natural image patches constituted one family. We created two other families: sinusoidal gratings that varied in orientation and spatial frequency, and 1/f noise. We calculated the inner product of each image with 180 odd-symmetric Gaussian derivative filters (derivative order: 3, aspect ratio: 4, peak spatial frequency: 6 c/deg), which perfectly tiled the orientation domain.

We computed image energy as a function of orientation by taking the squared filter responses (Figure 7B) and defined an OCI as one minus the normalized circular variance of the energy profile.

Discriminability of Images for a Population of Independent LNP Neurons

We measured discriminability of pairs of similar images for populations of LNP neurons designed to mirror the properties we measured in V1 and V2 (Figure 8). Our populations contained 45 neurons whose preferred orientations covered the entire domain evenly. The filter was a Gaussian derivative (order 2), with preferred spatial frequency of 6 c/deg. Its aspect ratio controlled orientation selectivity and was set to 0.25, 1, or 2. The nonlinearity was an exponent of 3.3. The images to be discriminated were a randomly selected patch of 21 × 21 pixels and its nearest neighbor (in terms of Euclidean norm of the pixel values) in the population of random samples drawn either from a collection of gratings, 1/f noise images, or the natural image data base. All populations were equated with respect to "metabolic cost," taken as the average firing rate per image, which was set to 10 spikes/image. We capture the average discriminability of pairs of similar images using a metric based on the Fisher information matrix:

$$D(\vec{s}_1, \vec{s}_2) = \sqrt{(\vec{s}_1 - \vec{s}_2)^T \left[\frac{I_F(\vec{s}_1) + I_F(\vec{s}_2)}{2} \right] (\vec{s}_1 - \vec{s}_2)},$$

where the $\{\vec{s}_1, \vec{s}_2\}$ are two stimuli (images) and the Fisher information matrix, $I_F(\vec{s})$, is computed as the expected value of the covariance of the gradient of the log likelihood of the population response

$$I_F(\vec{s}) = E_r \left\{ \vec{\nabla}_s \log p(\vec{r} | \vec{s}) \cdot \vec{\nabla}_s^T \log p(\vec{r} | \vec{s}) \right\}, \quad (1)$$

where $E_r\{\cdot\}$ is the expectation over the vector \vec{r} , representing the response of the neural population (i.e., a vector of spike counts). For an LNP model with independent responses, the log likelihood function is

$$\log p(\vec{r} | \vec{s}) = \sum_k [r_k \log f_k(\vec{s}) - f_k(\vec{s}) - \log(r_k!)], \quad (2)$$

where the firing rate function is

$$f_k(\vec{s}) = (\vec{w}_k^T \vec{s})^{\rho_k} H(\vec{w}_k^T \vec{s}), \quad (3)$$

with w_k the receptive field (weight vector) for the k th neuron, ρ_k the exponent describing its nonlinear response, and $H(\cdot)$ the Heaviside step function (equal to one when its argument is positive, zero otherwise). Combining Equations 2 and 3 and taking the gradient gives

$$\vec{\nabla}_s \log p(\vec{r} | \vec{s}) = \sum_k \frac{\rho_k}{\vec{w}_k^T \vec{s}} [r_k - f_k(\vec{s})] \vec{w}_k,$$

where we have assumed that the actual responses r_k are zero whenever $\vec{w}_k^T \vec{s}$ is negative. Plugging this into Equation 1 gives the Fisher information:

$$\begin{aligned} I_F(\vec{s}) &= \sum_{j,k} \frac{\rho_j \rho_k}{(\vec{w}_j^T \vec{s}) (\vec{w}_k^T \vec{s})} E_r \{ [r_j - f_j(\vec{s})] [r_k - f_k(\vec{s})] \} \vec{w}_j \vec{w}_k^T \\ &= \sum_k \frac{\rho_k^2}{(\vec{w}_k^T \vec{s})^2} E_r \{ [r_k - f_k(\vec{s})]^2 \} \vec{w}_k \vec{w}_k^T \\ &= \sum_k \frac{\rho_k^2}{(\vec{w}_k^T \vec{s})^2} f_k(\vec{s}) \vec{w}_k \vec{w}_k^T, \end{aligned}$$

where the $j \neq k$ terms have vanished because the responses are assumed independent and the $j = k$ terms have been simplified by noting that the response variance of a Poisson process is equal to its mean, $f_k(\vec{s})$.

SUPPLEMENTAL INFORMATION

Supplemental Information includes five figures and Supplemental Experimental Procedures and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2015.10.009>.

AUTHOR CONTRIBUTIONS

R.L.T.G., E.P.S., and J.A.M. designed the research; R.L.T.G. collected and analyzed the data; and R.L.T.G., E.P.S., and J.A.M. wrote the paper.

ACKNOWLEDGMENTS

This work was supported by NIH grants EY04440 and EY022428, the Howard Hughes Medical Institute, and postdoctoral fellowships from the Fund for Scientific Research of Flanders and the Belgian American Educational Foundation awarded to R.L.T.G. We are grateful to Robert M. Shapley for helpful discussions and to Nicholas Priebe for sharing his data with us.

Received: April 3, 2015

Revised: July 14, 2015

Accepted: September 30, 2015

Published: November 5, 2015

REFERENCES

- Albright, T.D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *J. Neurophysiol.* *52*, 1106–1130.
- Attwell, D., and Laughlin, S.B. (2001). An energy budget for signaling in the grey matter of the brain. *J. Cereb. Blood Flow Metab.* *21*, 1133–1145.
- Bell, A.J., and Sejnowski, T.J. (1997). The “independent components” of natural scenes are edge filters. *Vision Res.* *37*, 3327–3338.
- Box, G.E.P. (1976). Science and Statistics. *J. Am. Stat. Assoc.* *71*, 791–799.
- Carandini, M., and Ferster, D. (2000). Membrane potential and firing rate in cat primary visual cortex. *J. Neurosci.* *20*, 470–484.
- Carandini, M., Heeger, D.J., and Movshon, J.A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* *17*, 8621–8644.
- Cavanaugh, J.R., Bair, W., and Movshon, J.A. (2002). Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *J. Neurophysiol.* *88*, 2547–2556.
- Chapman, B., Zaksas, K.R., and Stryker, M.P. (1991). Relation of cortical cell orientation selectivity to alignment of receptive fields of the geniculocortical afferents that arborize within a single orientation column in ferret visual cortex. *J. Neurosci.* *11*, 1347–1358.
- Chklovskii, D.B., Schikorski, T., and Stevens, C.F. (2002). Wiring optimization in cortical circuits. *Neuron* *34*, 341–347.
- Cossell, L., Iacaruso, M.F., Muir, D.R., Houlton, R., Sader, E.N., Ko, H., Hofer, S.B., and Mrsic-Flogel, T.D. (2015). Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* *518*, 399–403.
- De Valois, R.L., Albrecht, D.G., and Thorell, L.G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Res.* *22*, 545–559.
- Felsen, G., and Dan, Y. (2005). A natural approach to studying vision. *Nat. Neurosci.* *8*, 1643–1646.
- Goris, R.L.T., Movshon, J.A., and Simoncelli, E.P. (2014). Partitioning neuronal variability. *Nat. Neurosci.* *17*, 858–865.
- Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* *9*, 181–197.
- Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurons in the cat's striate cortex. *J. Physiol.* *148*, 574–591.
- Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* *160*, 106–154.
- Jin, J., Wang, Y., Swadlow, H.A., and Alonso, J.M. (2011). Population receptive fields of ON and OFF thalamic inputs to an orientation column in visual cortex. *Nat. Neurosci.* *14*, 232–238.
- Leventhal, A.G., Thompson, K.G., Liu, D., Zhou, Y., and Ault, S.J. (1995). Concomitant sensitivity to orientation, direction, and color of cells in layers 2, 3, and 4 of monkey striate cortex. *J. Neurosci.* *15*, 1808–1818.
- Levitt, J.B., Kiper, D.C., and Movshon, J.A. (1994). Receptive fields and functional architecture of macaque V2. *J. Neurophysiol.* *71*, 2517–2542.
- Maunsell, J.H.R., and Van Essen, D.C. (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J. Neurophysiol.* *49*, 1127–1147.
- Movshon, J.A., Thompson, I.D., and Tolhurst, D.J. (1978). Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J. Physiol.* *283*, 53–77.
- Nauhaus, I., Busse, L., Carandini, M., and Ringach, D.L. (2009). Stimulus contrast modulates functional connectivity in visual cortex. *Nat. Neurosci.* *12*, 70–76.
- Olmos, A., and Kingdom, F.A.A. (2004). A biologically inspired algorithm for the recovery of shading and reflectance images. *Perception* *33*, 1463–1473.
- Olshausen, B.A., and Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* *381*, 607–609.
- Priebe, N.J., and Ferster, D. (2012). Mechanisms of neuronal computation in mammalian visual cortex. *Neuron* *75*, 194–208.
- Priebe, N.J., Mechler, F., Carandini, M., and Ferster, D. (2004). The contribution of spike threshold to the dichotomy of cortical simple and complex cells. *Nat. Neurosci.* *7*, 1113–1122.
- Reid, R.C., and Alonso, J.M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* *378*, 281–284.
- Ringach, D.L. (2002). Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *J. Neurophysiol.* *88*, 455–463.
- Ringach, D.L., Hawken, M.J., and Shapley, R. (1997). Dynamics of orientation tuning in macaque primary visual cortex. *Nature* *387*, 281–284.
- Ringach, D.L., Shapley, R.M., and Hawken, M.J. (2002). Orientation selectivity in macaque V1: diversity and laminar dependence. *J. Neurosci.* *22*, 5639–5651.
- Rose, D., and Blakemore, C. (1974). An analysis of orientation selectivity in the cat's visual cortex. *Exp. Brain Res.* *20*, 1–17.
- Rust, N.C., and Movshon, J.A. (2005). In praise of artifice. *Nat. Neurosci.* *8*, 1647–1650.
- Rust, N.C., Schwartz, O., Movshon, J.A., and Simoncelli, E.P. (2005). Spatiotemporal elements of macaque V1 receptive fields. *Neuron* *46*, 945–956.
- Schiller, P.H., Finlay, B.L., and Volman, S.F. (1976). Quantitative studies of single-cell properties in monkey striate cortex. II. Orientation specificity and ocular dominance. *J. Neurophysiol.* *39*, 1320–1333.
- Scholl, B., Tan, A.Y., Corey, J., and Priebe, N.J. (2013). Emergence of orientation selectivity in the Mammalian visual pathway. *J. Neurosci.* *33*, 10616–10624.
- Seriès, P., Stocker, A.A., and Simoncelli, E.P. (2009). Is the homunculus “aware” of sensory adaptation? *Neural Comput.* *21*, 3271–3304.
- Seung, H.S., and Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proc. Natl. Acad. Sci. USA* *90*, 10749–10753.
- Shamir, M., and Sompolinsky, H. (2006). Implications of neuronal diversity on population coding. *Neural Comput.* *18*, 1951–1986.
- Sharpee, T., Rust, N.C., and Bialek, W. (2004). Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput.* *16*, 223–250.
- Sillito, A.M. (1975). The contribution of inhibitory mechanisms to the receptive field properties of neurons in the striate cortex of the cat. *J. Physiol.* *250*, 305–329.
- Smith, M.A., Majaj, N.J., and Movshon, J.A. (2005). Dynamics of motion signaling by neurons in macaque area MT. *Nat. Neurosci.* *8*, 220–228.
- Tao, L., Shelley, M., McLaughlin, D., and Shapley, R. (2004). An egalitarian network model for the emergence of simple and complex cells in visual cortex. *Proc. Natl. Acad. Sci. USA* *101*, 366–371.

- Tolhurst, D.J., and Thompson, I.D. (1981). On the variety of spatial frequency selectivities shown by neurons in area 17 of the cat. *Proc. R. Soc. Lond. B Biol. Sci.* 213, 183–199.
- Touryan, J., Lau, B., and Dan, Y. (2002). Isolation of relevant visual features from random stimuli for cortical complex cells. *J. Neurosci.* 22, 10811–10818.
- Vintch, B., Movshon, J.A., and Simoncelli, E.P. (2015). A convolutional subunit model for neuronal responses in macaque V1. *J. Neurosci.* 35, <http://dx.doi.org/10.1523/JNEUROSCI.2815-13.2015>.
- Wieland, D.J., Shelley, M., McLaughlin, D., and Shapley, R. (2001). How simple cells are made in a nonlinear network model of the visual cortex. *J. Neurosci.* 21, 5203–5211.
- Xing, D., Ringach, D.L., Shapley, R., and Hawken, M.J. (2004). Correlation of local and global orientation and spatial frequency tuning in macaque V1. *J. Physiol.* 557, 923–933.
- Xing, D., Shapley, R.M., Hawken, M.J., and Ringach, D.L. (2005). Effect of stimulus size on the dynamics of orientation selectivity in Macaque V1. *J. Neurophysiol.* 94, 799–812.
- Xing, D., Ringach, D.L., Hawken, M.J., and Shapley, R.M. (2011). Untuned suppression makes a major contribution to the enhancement of orientation selectivity in macaque v1. *J. Neurosci.* 31, 15972–15982.
- Zeki, S.M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *J. Physiol.* 236, 549–573.