

Contextual modulation of V1 receptive fields depends on their spatial symmetry

Tatyana O. Sharpee · Jonathan D. Victor

Received: 20 April 2008 / Revised: 9 June 2008 / Accepted: 10 June 2008 / Published online: 5 August 2008
© Springer Science + Business Media, LLC 2008

Abstract The apparent receptive field characteristics of sensory neurons depend on the statistics of the stimulus ensemble—a nonlinear phenomenon often called contextual modulation. Since visual cortical receptive fields determined from simple stimuli typically do not predict responses to complex stimuli, understanding contextual modulation is crucial to understanding responses to natural scenes. To analyze contextual modulation, we examined how apparent receptive fields differ for two stimulus ensembles that are matched in first- and second-order statistics, but differ in their feature content: one ensemble is enriched in elongated contours. To identify systematic trends across the neural population, we used a multidimensional scaling method, the Procrustes transformation. We found that contextual modulation of receptive field components increases with their spatial extent. More surprisingly, we also found that odd-symmetric components change

systematically, but even-symmetric components do not. This symmetry dependence suggests that contextual modulation is driven by oriented On/Off dyads, i.e., modulation of the strength of intracortically-generated signals.

Keywords Primary visual cortex · Plasticity · Linear–nonlinear model · Reverse correlation

1 Introduction

Neurons in primary visual cortex are often described as detectors of oriented features (Hubel and Wiesel 1959, 1968), or as oriented filters (De Valois et al. 1982a, b; De Valois and Thorell 1988). Such spatial processing can be adequately characterized by stimuli that vary along only one dimension, such as bars, edges, and gratings. Moreover, implicit (or sometimes explicit) in this view is the notion that the main qualitative features of a V1 neuron's response can be described in terms of a linear spatial filter, or a linear spatial filter followed by a simple nonlinearity such as a threshold, or an energy operation.

However, there is a large and increasing body of evidence that this simple picture is incomplete. First, there is the intriguing role of the surround of the receptive field. Increasing stimulation of the surround increases sparseness (Vinje and Gallant 2000) and efficiency of information transmission (Vinje and Gallant 2002) during natural vision. Second, even simple visual stimuli placed in the surround strongly influence a cell's responses. These modulatory effects may arise at several levels. Some may reflect processes that have already occurred in the retina (Shapley and Victor 1978) and lateral geniculate nucleus (Ohzawa et al. 1985). Some can be explained by the nonlinearities of processing within a single cortical cell

Electronic supplementary material The online version of this article (doi:10.1007/s10827-008-0107-5) contains supplementary material, which is available to authorized users.

Action Editor: Matthew Wiener

T. O. Sharpee (✉)
Computational Neurobiology Laboratory,
The Salk Institute for Biological Studies,
La Jolla, CA 92037, USA
e-mail: sharpee@salk.edu

T. O. Sharpee
Center for Theoretical Biological Physics,
University of California, San Diego,
La Jolla, CA 92093, USA

J. D. Victor
Weill Cornell Medical College,
New York, NY 10065, USA
e-mail: jdvicto@med.cornell.edu

(Priebe and Ferster 2006) or synapse (Carandini et al. 2002; Freeman et al. 2002). Others appear to require cortical network mechanisms (Das and Gilbert 1999) including suppressive connections between neurons of different orientations and spatial frequencies (Bonds 1989; Carandini et al. 1997; Heeger 1992).

Here we examine how the shape of the apparent receptive field sensitivity profile depends on context, and focus on aspects of context that are not likely to trigger simple gain control mechanisms. We do this by analyzing responses of neurons to two sets of two-dimensional Hermite functions (Fig. 1). The two stimulus sets are matched in spatial extent, contrast, and second-order statistics, but differ in their higher-order statistics. These differences are visually obvious and lead to qualitatively different two-dimensional patterns: one set has Cartesian symmetry and contains Gabor-like patterns, checkerboard-like patterns, and elongated contours; the other set has polar symmetry and contains target-like and pinwheel patterns. If the neural response were described by a linear function of the stimulus (perhaps followed by a nonlinearity), then both stimulus ensembles would provide identical reconstructions of the neural receptive field. This is because individual stimuli within one stimulus ensemble can be expressed as linear combinations of stimuli from the other stimulus ensemble. Moreover, even if the spatial sensitivity were modulated by mechanisms sensitive to the “context” of first- and second-order statistics (such as mean luminance, contrast, and power spectrum), the two ensembles would still provide identical reconstructions—because their statistics are matched at first and second order.

The focus on context-dependence induced by statistics of order greater than two might at first seem like a mathematical curiosity, but it is actually of central impor-

tance. This is because it is these high-order statistics (phase correlations) that lead to visually salient structure such as lines and edges, both in artificial stimuli (Morrone and Burr 1988) and in natural scenes (Oppenheim and Lim 1981; Ruderman and Bialek 1994; Simoncelli and Olshausen 2001). Moreover, standard models for V1 neuronal responses tend to be inaccurate specifically for natural scene stimuli in which such high-order correlations are present (Felsen et al. 2005). Thus, delineation of the influence of high-order statistics on V1 neurons is crucial to advance our understanding of the computations that they carry out.

However, studying context-dependence induced by high-order statistics via responses to natural scenes is difficult, because high-order statistics of natural scenes are complex and incompletely characterized. This, specifically, motivated our study of context-dependence driven by Hermite functions, which are analytically simple, and have fully-characterized statistics. We recently showed (Victor et al. 2006) that Hermite functions induced significant context-dependent changes in a large fraction of V1 cat and macaque neurons. However, it was not clear how those changes in apparent receptive field shape could be succinctly characterized, or what features of these changes were systematic. For example, across the population, we found no consistent shifts in either spatial frequency or orientation tuning. Therefore, we sought a more data-driven, non-parametric approach to characterize these changes.

The present study carries out such an approach. By considering the population of V1 neurons as a whole, it succeeds in identifying two systematic aspects of the context-dependence of their receptive fields. First, receptive field components covering a wider area show stronger context-dependent changes than those that cover a narrower area. Second, we find that consistency of context-dependent changes across the population of cells depends on symmetry properties of receptive field components: in particular, components that invert their contrast under 180-degree spatial rotation change in a more consistent manner than components that remain unchanged with this rotation.

This latter conclusion is unexpected, and neither conclusion was apparent from casual inspection of the dataset. Rather, they emerged from a novel method that we used to study context-dependent changes, which we believe is applicable far beyond the present scenario. The approach is based on a version of the Procrustes transformation of multi-dimensional scaling (Kabsch 1976). The main idea of this method is to find the best transformation from a set of receptive fields obtained in one context to a set of receptive fields from the same neurons recorded in a second context.

A distinctive feature of this approach is that it seeks to identify a single transformation that applies to all measured

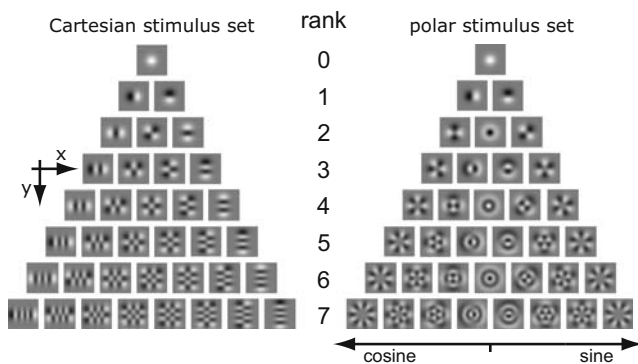


Fig. 1 Stimulus patterns used in these experiments. Stimuli were two-dimensional Hermite functions (Victor et al. 2006). Cartesian stimuli (*left panel*) are products of one-dimensional Hermite functions in vertical and horizontal coordinates. Each *row* contains all the Cartesian stimuli of a given rank. Linear combinations of patterns of the same rank can be constructed to be separable in radial and angular coordinates. These are the polar stimuli (*right panel*). Modified with permission from Victor et al. 2006

receptive fields within a population. Thus, it can identify patterns that would not be apparent if each neuron's receptive field were studied independently. As such, this is an approach that is likely to be widely applicable to the problem of characterizing how receptive field shapes depend on some manipulation, such as contrast, attention, or a pharmacologic perturbation.

2 Methods

Electrophysiological recordings The data presented here were collected for a previous study (Victor et al. 2006), and the experimental preparation has been detailed there. Briefly, neural responses were collected using tetrodes from the primary visual cortex of anesthetized, paralyzed cats and macaques. Single units were tentatively identified during recording using real-time spike-sorting software (Datawave Technologies); quantitative analysis of receptive fields was carried out off-line following spike sorting with custom software (Reich 2001).

At each recording site, one well-isolated single unit with signal-to-noise of the spike waveform $>2:1$, and usually $3:1$, was selected as the “target” neuron. Orientation tuning of the target neuron was determined from responses to drifting gratings at orientations spaced in steps of 22.5° (or, for narrowly tuned units, 11.25°), with spatial and temporal frequency determined by initial assessment. Next, the spatial frequency tuning was determined using drifting gratings at an 8-to 16-fold range of spatial frequencies straddling the value determined by auditory assessment. Temporal tuning was determined from responses to 1-, 2-, 4-, 8-, and 16 Hz drifting gratings of optimal orientation and spatial frequency. Finally, a contrast response function was determined from responses to drifting gratings at contrast of 0, 0.0625, 0.125, 0.25, 0.5, and 1.0, with orientation, spatial and temporal frequencies previously determined.

The center of the receptive field was determined from the response to either a bright or a dark bar, moving slowly (≤ 1 deg/s) and symmetrically in both directions along the preferred axis. The center of the receptive field along the preferred axis was determined as the position corresponding to the mean time of the peak responses elicited by bars moving slowly in the preferred and anti-preferred directions. The center of the receptive field in the orthogonal direction was taken as a halfway point between the upper and lower edges of the receptive field, which in turn were determined by the appearance of a response to slowly swept patches along multiple trajectories parallel to the preferred axis.

Once centered, the size of the classical receptive field was determined from response to a drifting grating (all

parameters optimized) presented in disks of increasing diameter and in a series of annuli. The effective diameter D of the receptive field was taken to be the smallest inner diameter of an annulus that did not produce a statistically significant response above zero. The set of annuli were chosen so that D was determined to within $1/2$ deg or, for smaller receptive fields, $1/4$ deg. The effective diameter D was used (below) to set the width of the zero-rank two-dimensional Hermite function, a Gaussian. Higher-rank functions were scaled in proportion to this Gaussian, as in Fig. 1.

Stimulus presentation Neurons were probed with visual patterns derived from two-dimensional Hermite functions (see Fig. 1). The contrast profiles of these patterns are polynomials multiplied by a Gaussian envelope. Stimuli were rotated so that the x -axis was along the target neuron's preferred orientation and the positive y -axis was the preferred direction for drifting gratings. We set the spatial parameter σ of the Hermite functions (Eqs. A1 through A5 of Victor et al. 2006) at $\sigma = D/10$, where D is the diameter of the classical receptive field of the target neuron as determined by responses to disks and annuli containing optimal drifting gratings (see above and Victor et al. 2006). This choice of σ creates Hermite function stimuli that have one, two, or three oscillations within a region of space that covers the receptive fields. Each stimulus was presented as illustrated in Fig. 1, and also with luminance polarity reversed.

Computation of receptive fields for Cartesian and polar stimuli For the reverse-correlation modeling approach (L-filters), receptive fields were computed in a manner similar to a spike-triggered average. The contribution of each basis function to the linear component of the estimated receptive field (L-filter) was determined from the *difference* between the response to that basis function, and to its contrast-inverse. This yields the desired filter shapes because individual stimulus patterns are orthogonal to each other. See (Victor et al. 2006) and Supplementary Information for further details, including characterization of the nonlinear component of the receptive field.

For the MID modeling approach (M-filters), receptive fields were computed using the method of (Sharpee et al. 2004). Any candidate receptive field shape can be interpreted as a direction in stimulus space. Stimuli can be projected along that direction to yield a set of scalars. The receptive field determined by this method is the direction for which the mutual information between these scalars and the firing rate is maximized. Advantages of this method are that (a) it does not place requirements on the correlation structure of the stimulus set, and (b) it does not postulate a specific form for the relationship between the projection

and the response—this can be an arbitrary nonlinear function.

Details of the computational algorithm have been previously described in (Sharpee et al. 2004) and in the supplementary information of (Sharpee et al. 2006). Briefly, the spike-triggered average was chosen as an initial guess. Then, we used a combination of gradient ascent and simulated annealing to maximize information; 1000 line optimizations were used. Each line optimization was carried out along the gradient of information evaluated along the current candidate relevant dimension. No regularization by the performance on the test set was implemented. Thus, computation of jackknife estimates (by setting aside one presentation, see below) was completely independent from the data in the omitted presentation.

Debiasing and standard error estimation for correlation coefficients We used correlation coefficients to compare receptive field shapes. Two kinds of comparisons were made: receptive fields derived from the two modeling approaches (reverse correlation vs. MID) for a single stimulus set (Cartesian or polar), and receptive fields derived from two stimulus sets but according to a single modeling approach. In each case, we used a jackknife method to debias the estimates of the correlation coefficients, and to obtain estimates of their standard error.

We are interested in the absolute value of the correlation coefficient $cc = |\vec{v} \cdot \vec{u}|$ between two receptive field estimates, \vec{v} and \vec{u} . The absolute value is used because for MID, receptive fields are defined up to a scaling factor, which could be negative.

To carry out the jackknife procedure (Efron and Tibshirani 1998), we computed receptive fields from all stimulus presentations \vec{v}_{full} and \vec{u}_{full} , and also the “drop one” jackknife estimates \vec{v}_i and \vec{u}_i computed by leaving out the i th stimulus repetition from the analysis.

The jackknife estimate of the bias of a statistical estimator θ that has an approximately Gaussian distribution is

$$0_{\infty} = \theta_{full} - (N - 1)(\langle \theta_i \rangle - \theta_{full}), \quad (1)$$

where θ_{full} is estimated taking all of the data into account, and $\langle \theta_i \rangle$ is the mean of the drop-one estimates, obtained by successively omitting one of the N repetitions (in this dataset N ranged from 8 to 16). Correlation coefficients are limited to the range between -1 and 1 , and their distribution is not well approximated by a Gaussian. To normalize their distribution and make the jackknife procedure applicable, we first applied the Fisher z-transform to the correlation coefficients (Efron and Tibshirani 1998), i.e., $\theta = \text{arctanh}(cc)$. After the z-transformed correlation coefficients were debiased according to Eq. (1), we applied the inverse transformation, i.e., $cc = \tanh(\theta)$, to obtain debiased correla-

tion coefficients. We note that even though correlation coefficients computed from jackknife estimates and the full dataset are always positive, negative values can appear after debiasing.

The standard error of jackknife estimates was also first computed on the z-transformed variables θ : $\Delta\theta = \sqrt{\frac{N-1}{N} \sum_{i=1}^N (\theta_i - \theta_{full})^2}$ and then transformed back to obtain standard error for the correlation coefficients:

$$\sigma_{\text{jack},\pm} = \tanh(\theta_{\infty} \pm \Delta\theta) - \tanh(\theta_{\infty}), \quad (2)$$

where θ_{∞} was computed according to Eq. (1). This results in an asymmetric confidence interval. We quote the larger of the two values when describing correlation coefficients in figure legends.

3 Results

3.1 Context-dependence of responses to Cartesian and polar stimuli

Our analysis is based on $n=51$ single neurons ($n=34$ from cat visual cortex and $n=17$ from monkey visual cortex) previously recorded (Victor et al. 2006). This study found differences between the receptive fields measured with the Cartesian and polar stimulus sets showed in Fig. 1, but was unable to characterize the nature of this context-induced change—which is the goal of the present study.

To pursue this goal, the first step is to characterize the responses to each stimulus set in a concise fashion—i.e., by determining the effective receptive field in each context. This characterization necessarily entails assumptions (a receptive field model) and approximations (which aspects of the response are incorporated into the model). Therefore, before drawing conclusions about the differences between Cartesian and polar stimulus sets, a prerequisite is to show that our characterization is robust. To do this, we will use two complementary strategies to estimate the receptive fields from each context, and show that these two modeling approaches lead to the same estimate.

The basic element of our receptive field model is a linear filter followed by a static nonlinearity (an “LN” model, de Boer and Kuyper 1968; Meister and Berry 1999; Schwartz et al. 2006; Victor and Shapley 1980). One way of fitting this model consists essentially of reverse correlation (Victor et al. 2006). The key assumption of this approach is that the nonlinearity can be decomposed into a linear component and an even-symmetric component. In the second approach, known as the maximally informative dimension (MID) method (Sharpee et al. 2004), the best LN approximation is determined according to an information-theoretic criterion.

This allows us to avoid making assumptions about the shape of the nonlinearity. Another difference between these two approaches is that in the reverse-correlation method, the estimate of the L-filter is unaffected by the spatial distribution of On/Off responses (instead, the On/Off responses are characterized by a second LN pathway), while in the MID method as implemented here, a single LN pathway is used to model all responses. These approaches also differ in their sensitivity to response variability and limited amounts of data. This is because the MID method relies on a nonlinear optimization, while the reverse correlation method does not.

We will refer to the linear component of the LN model extracted by reverse correlation as L-filters, and those computed as MIDs as M-filters.

As we now show, despite these differences in how the L-filters and M-filters are estimated, they have similar shapes when estimated from responses to the same basis set (i.e., in the same context), and show parallel changes in shape when context is altered (Fig. 2 and Supplementary Figs. 1 and 2). Thus, the LN model (as determined by either method) is a meaningful way to characterize the response to either basis set, and how this characterization changes with context.

Figure 2 shows several example cells that cover the observed range of context dependence. In all cases, the L-filters and the M-filters were similar to each other when obtained with the same stimulus set (Cartesian or polar stimuli). Both characterizations (L-filters and M-filters) also had a similar dependence on context (Cartesian vs. polar stimuli). In Fig. 2(a) we show examples of cells that did not show contextual modulation. In these cells, the L-filters and the M-filters (compare two columns shown for each cell) are similar, as are the filters obtained from Cartesian and polar stimulus sets (compare the two rows shown for each cell). The example cells of Fig. 2(b,c) show a context effect: the L-filters and the M-filters determined from the same stimulus set (either Cartesian or polar) are similar, but the filters depend on which stimulus set (Cartesian vs. polar) is used. For the cells of Fig. 2(b), the context-dependence was statistically significant ($p < 0.05$, see Methods) for both the L-filters and the M-filters. For the cells of Fig. 2(c), the context-dependence was statistically significant only for the L-filters. For the M-filters, the differences between the Cartesian and polar M-filters are qualitatively similar to the differences seen for the L-filters, but did not reach statistical significance.

To quantify the magnitude of context-dependent changes at the population level, we compute the distribution of correlation coefficients between normalized L-filters computed from Cartesian and polar stimulus set [Fig. 3(a)]. This quantity measures the extent to which the receptive field map, as determined by reverse correlation, depends on the stimulus set. In agreement with the previous report (Victor

et al. 2006), more than half (39/51) of the cells in our dataset had significant differences in receptive fields of the linear model between Cartesian and polar stimuli, with 28 cells showing differences significant at $p < 0.01$, and 11 cells showing differences significant at $0.01 < p < 0.05$. For the remaining 12 cells, there was no significant difference in receptive fields determined from Cartesian and polar stimuli. Figure 3(b) shows a parallel analysis using the correlation coefficient determined from the M-filters: 20 out of 51 cells showed significant context-dependence (17 cells at $p < 0.01$ and three cells at $0.01 < p < 0.05$). Fewer cells with significant differences between receptive fields for Cartesian and polar stimuli were found for the M-filters (20) than for the L-filters (39). Most likely this is due to larger uncertainty in the M-filter profiles than for the L-filters, which is in turn a consequence of the absence of assumptions on the shape of the nonlinear gain function in the MID method (see Supplementary Fig. 3). As examples in Fig. 2(c) illustrate, the context-dependence revealed by the two methods is qualitatively similar, even if it does not reach statistical significance with one of the approaches.

3.2 Identifying systematic changes in receptive fields across a population

Having found that approximately half of the neurons exhibit changes in receptive field profiles when probed with Cartesian versus polar stimuli, we now seek to determine the extent to which these changes are systematic. As previously reported in this population of cells (Victor et al. 2006), the change in effective receptive field as a function of Cartesian vs. polar context was not associated with a systematic change in sensitivity, spatial frequency tuning, or orientation tuning. This motivated us to pursue a more data-driven, nonparametric approach to this problem. As a first step, we make use of the fact that each receptive field profile (e.g., the profiles shown in Fig. 2) can be considered to be members of a vector space of moderate dimension. This is because each filter shape is a linear combination of the stimuli. Moreover, since each set of stimuli is a linear combination of the other one, we can use either stimulus set as a basis for describing the filter shapes. For definiteness, we choose the polar ones. With this convention, the filter shape of a single cell is described by a vector of 36 numbers—the coefficients in the linear combination of the 36 polar functions.

Extending this idea, the set of filters obtained from all cells with Cartesian stimuli can be thought of as a set of vectors in a 36-dimensional space of polar filter shapes. Similarly, the set of filters obtained from the same cells with the polar stimuli can be thought of as a corresponding set of vectors in the same 36-dimensional space.

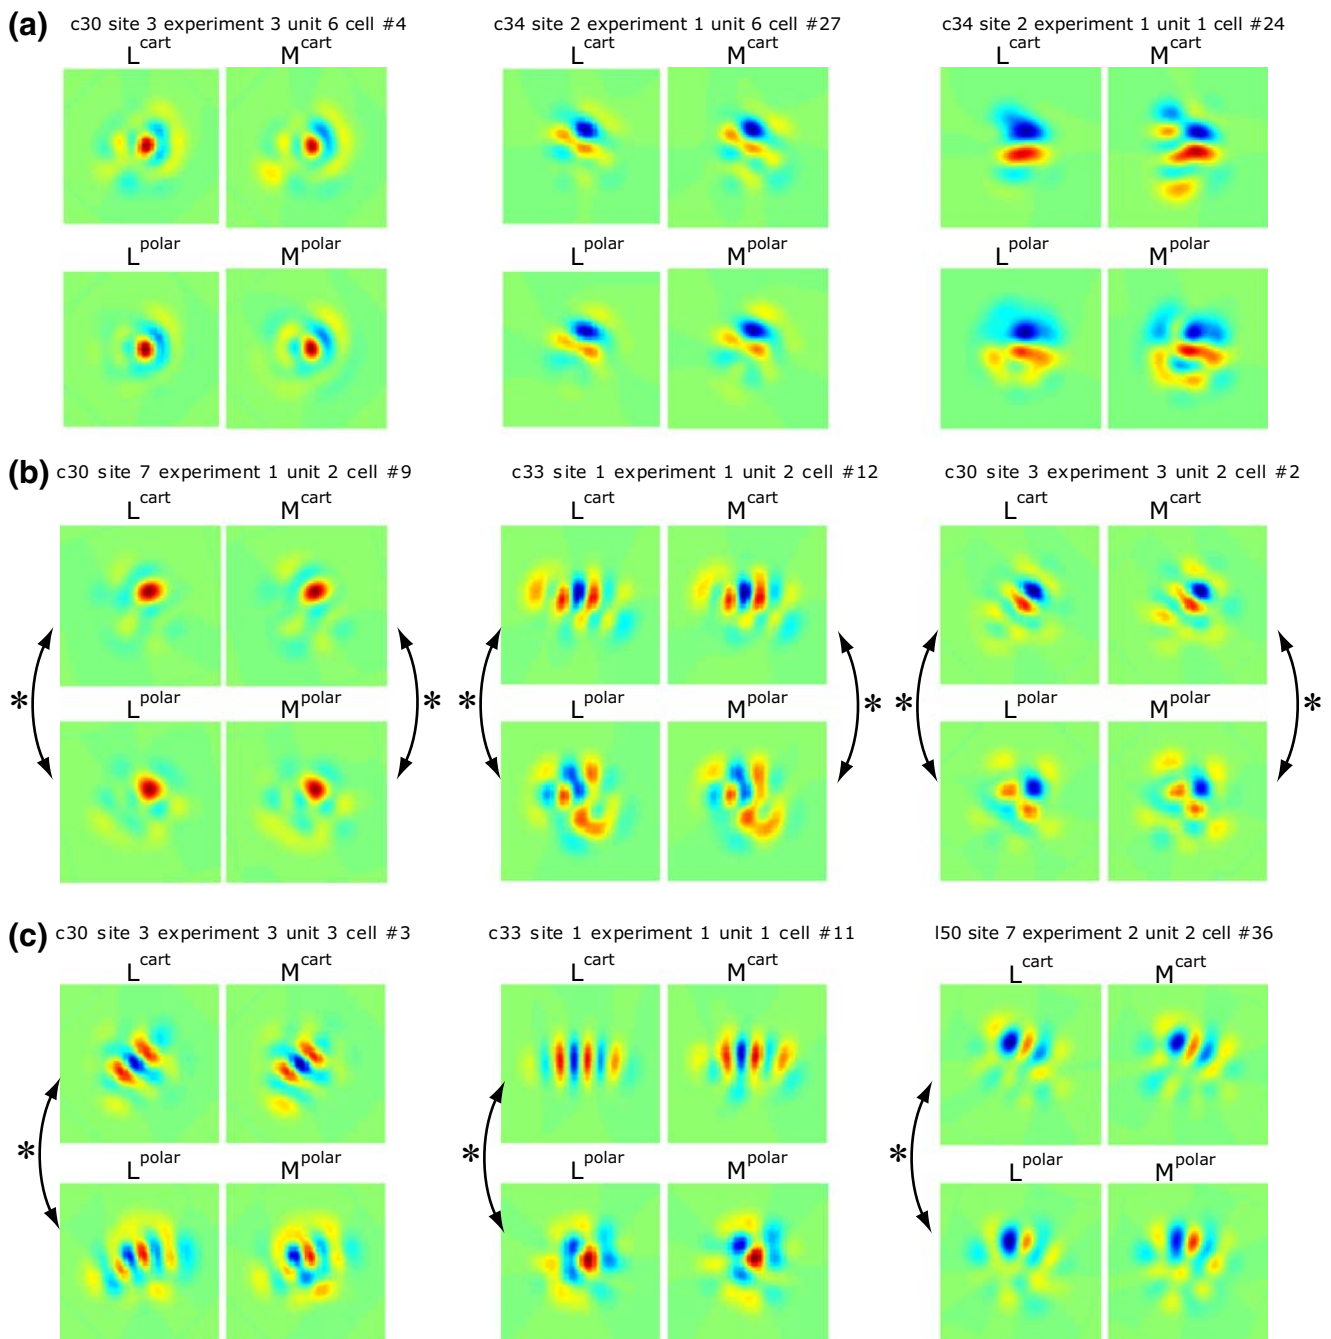


Fig. 2 Comparison of receptive fields for Cartesian and polar stimuli computed within two models. **(a)** Cells with no consistent differences between receptive fields determined from Cartesian or polar basis sets by reverse correlation (L-filters) or MID (M-filters). L^{cart} : Linear filter derived from reverse correlation of Cartesian responses, M^{cart} : filter derived from MID analysis of Cartesian responses, L^{polar} : linear filter derived from reverse correlation of polar responses, M^{polar} : filter derived from MID analysis of polar responses. Correlation coefficients of filters determined from the two stimulus contexts (L^{cart} vs. L^{polar} , M^{cart} vs. M^{polar}) were not significantly different from 1. L-filters (from left to right): 0.89 ± 0.09 , 0.97 ± 0.02 , 0.92 ± 0.05 (all $p > 0.05$); M-filters: 0.92 ± 0.07 , 0.991 ± 0.009 , 0.99 ± 0.01 (all $p > 0.05$). **(b)** Cells with context-dependent receptive fields as determined by reverse correlation (L-filters) and MID (M-filters). Correlation coefficients of filters determined from the two stimulus contexts were all significantly

different from 1 (marked by arrows). L-filters: 0.89 ± 0.03 ($p < 0.01$), 0.70 ± 0.10 ($p < 0.01$), 0.73 ± 0.08 ($p < 0.05$); M-filters: 0.70 ± 0.10 ($p < 0.05$), 0.92 ± 0.04 ($p < 0.05$), 0.60 ± 0.10 ($p < 0.01$). **(c)** Cells with context-dependence of receptive fields as determined by reverse correlation (L-filters) but not by MID (M-filters). Correlation coefficients of filters determined from the two stimulus contexts were significantly different from 1 (marked by arrows) for L-filters: 0.40 ± 0.15 , 0.50 ± 0.10 , $cc = 0.91 \pm 0.03$ (all $p < 0.01$) but not for M-filters: 0.91 ± 0.16 , 0.98 ± 0.03 , 0.90 ± 0.10 (all $p > 0.05$). None of these nine cells had consistent differences between L- and M-filters for either Cartesian or polar basis sets. Color scale is arbitrary, but is the same for all of the four filters pertaining to a neuron. For each neuron, the color scale covers the range from the minimal to the maximal value across the four filters




Fig. 3 Distribution of correlation coefficients between receptive fields determined from Cartesian and polar stimulus sets. *Left* Stimulus-dependent changes in filters calculated by reverse correlation, *right* stimulus-dependent changes in filters calculated by MID. For both kinds of models, significant context-dependent changes (correlation

coefficients <1) are prevalent. Cells with no significant changes ($p > 0.05$) are shown in white, those with significant changes in gray ($0.01 < p < 0.05$) and black ($p < 0.01$). The *abscissa* labels indicate the filters whose profiles are compared. Debiasing (see [Methods and Supplementary Material](#)) can result in estimated correlation coefficients <0

It is then natural to ask what linear transformation best accounts for differences between the two sets of receptive fields. That is, can we find a transformation of the 36-dimensional space that maps the Cartesian filter shapes into the corresponding polar filter shapes?

There are many linear transformations that will transform one vector (e.g., the filter from Cartesian stimuli for a particular cell) into a second vector (e.g., the filter for that cell from polar stimulus set). However, since each cell within the population has its own set of filters, the problem of finding a *single* transformation that works for all cells is highly constrained. Indeed, since we have more cells (51) than coordinates (36), it would be surprising if a single transformation would be able to map each cell's filter obtained from Cartesian stimuli into each cell's filter obtained from polar stimuli. Since we cannot expect to account for all of the variance with a single linear transformation, instead we seek the linear transformation that accounts for the largest possible fraction of the variance. This linear transformation indicates the differences between the Cartesian and polar filters that are systematic across the population.

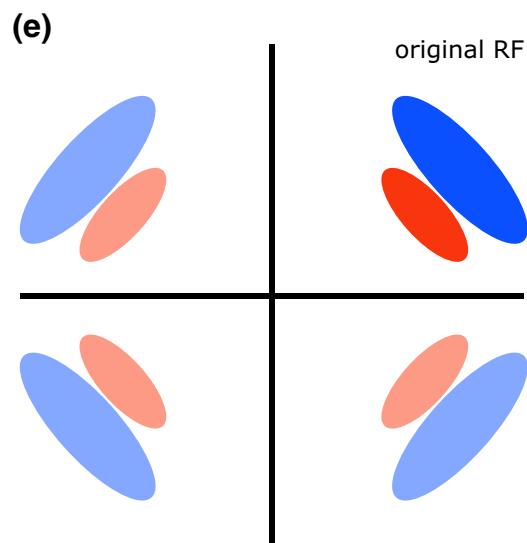
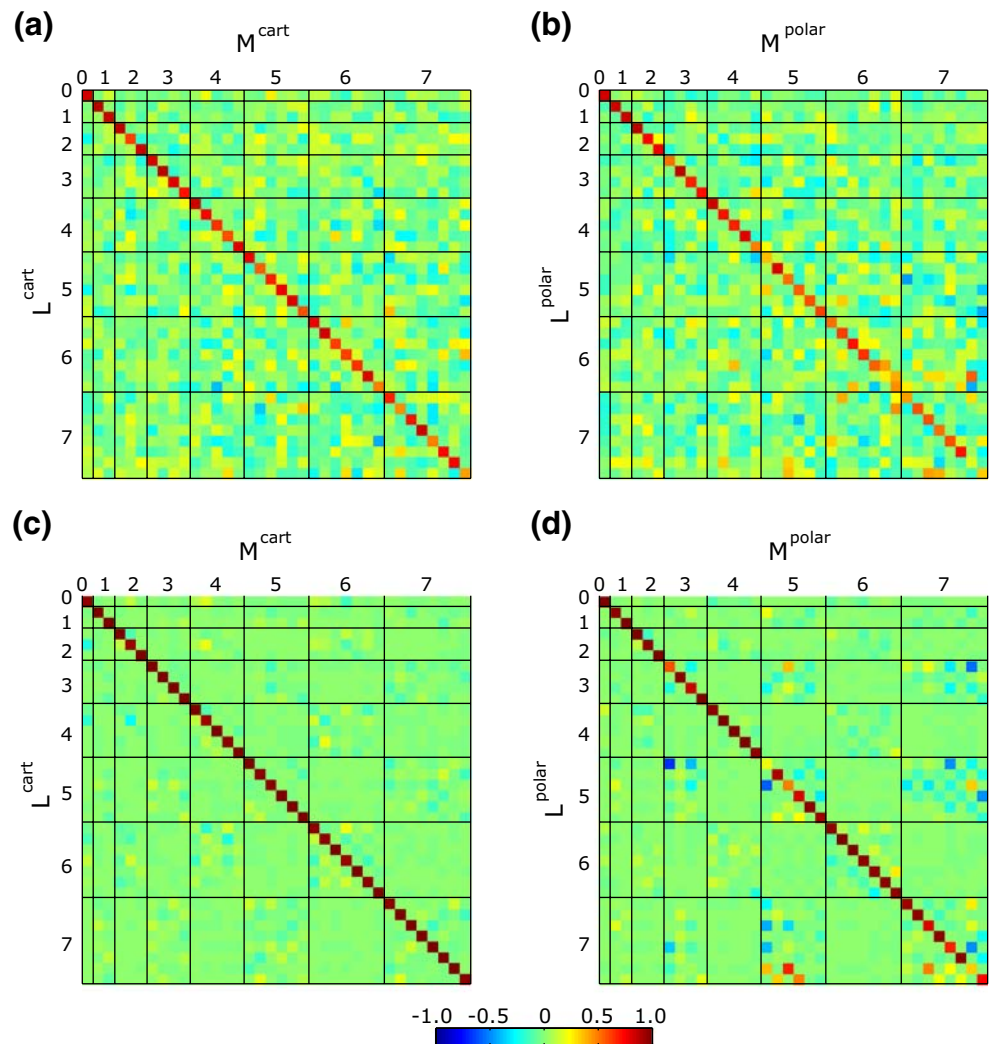
This is generally known as the Procrustes problem, often encountered in multidimensional scaling (Cox and Cox 2000). Because we are interested in how the shapes, not the amplitudes, of receptive fields depend on context, we normalize filters to unit length, and look for the best rotation that can transform one set of vectors into another with minimum least square error. Note that the rotations we seek are not literal rotations in space, but abstract rotations within the 36-dimensional shape space defined by the stimulus patterns. For example, a rotation might specify that neurons whose sensitivity profiles had had a large component of the rank-0 (Gaussian) basis function when studied with Cartesian stimuli tended to have a large component of the rank-2 target-like basis function when studied with polar stimuli.

3.3 No systematic differences between L-filters and M-filters

As a test of this approach, we first compare receptive fields obtained from Cartesian stimuli via reverse correlation (L^{cart}) and with those obtained from the same stimulus set by MID (M^{cart}). Since we have already seen that the L-filters and the M-filters were similar (Fig. 2), we expect that the best rotation matrix is close to the identity—and this is what we found [Fig. 4(a)]. Similarly, the best rotation matrix that transforms the receptive fields obtained from polar stimuli via reverse correlation (L^{polar}) into the receptive fields obtained from the same stimulus set by MID (M^{polar}) is also close to the identity matrix [Fig. 4(b)]. This supports the validity of the approach.

Because the rotation matrices identified by this procedure are not *precisely* the identity, the possibility remains that there are systematic differences between the receptive fields determined by the two methods that not evident from casual inspection of Fig. 4(a,b). To determine whether this is the case, we exploit the symmetry of the Hermite functions (Fig. 1). Note, for example, that reflection of the coordinate system across the horizontal axis inverts the polarity of some stimulus patterns, but leaves others preserved. Thus, when a given receptive field shape is represented as a vector, the sign assigned to some coordinates depends on the (arbitrary) choice of which horizontal direction is “positive”, while the sign assigned to other coordinates does not have this dependence. Consequently, any apparent linear relationship between the inverting and non-inverting coordinates must be due to chance. A similar argument holds for reflection across the vertical axis. Thus, we can restrict the search for rotations to those that do not violate these parity rules. This eliminates rotations that mix basis functions of even and odd ranks, or that mix basis functions in even and odd positions within ranks (Fig. 1). This procedure has another

Fig. 4 Comparison of receptive fields computed by the two models. Panels (a)–(b) pseudo-color display of the rotation matrix that is the optimal transformation between receptive fields estimated by two methods (reverse correlation (L-filters) and MID (M-filters)) from one set of stimuli. (a) Cartesian stimuli. (b) Polar stimuli. Panels (c) and (d) the same calculations, augmented by adding receptive fields reflected around horizontal and vertical axes (panel e). The heavy lines in panels (a)–(d) separate the ranks, shown increasing from 0 to 7. Within each rank, basis elements are ordered from most centrally-weighted (middle of pyramid of Fig. 1, right) to most peripherally-weighted (edges of pyramid of Fig. 1, right). Color scale covers the interval $[-1, 1]$, with green indicating 0, red-brown indicating 1, and blue indicating -1 . The matrices are all similar to the identity matrix, indicating close correspondence between receptive fields derived by the two methods



interpretation [Fig. 4(e)]. It is equivalent to adding cells to the dataset whose receptive fields are related by mirror symmetry to the cells actually recorded. The above analysis corresponds to assuming that since these added cells have the same relationship to the center and axes of the stimuli as real recorded neurons, their responses will be the same.

As seen in Fig. 4(c,d), this refinement increases the similarity of the inferred rotation to the identity matrix. Thus, the results of Fig. 2 are not only corroborated, but also extended: we do not identify any systematic difference between the L-filters and the M-filters obtained with the same stimulus set (i.e., in the same context).

3.4 Context-dependence of a component depends on its size and symmetry

A qualitatively different result emerges when we use this approach to compare receptive fields measured in different

contexts (Cartesian vs. polar). Comparisons of the L-filters derived from neural responses to Cartesian and polar stimulus sets are shown in Fig. 5(a), while those based on the M-filters are shown in Fig. 5(b). Because there is typically some correlation between receptive field profiles derived from Cartesian and polar stimulation, we expect that a component of the identity matrix will be present. This is indeed the case, but it is strongest for low-rank coefficients, and becomes less pronounced as the rank increases. Thus, higher-rank components, which have larger spatial extent (Fig. 1), are more subject to systematic stimulus-dependent changes than the lower-rank components.

The deviations of diagonal elements from unity measures the size of the systematic change of the corresponding coordinates. In Fig. 5(e) we plot diagonal elements of the best rotation matrices derived for L-filters (magenta) and M-filters (blue). In both cases, the amplitude of diagonal elements decreases with rank, indicating increasing devia-

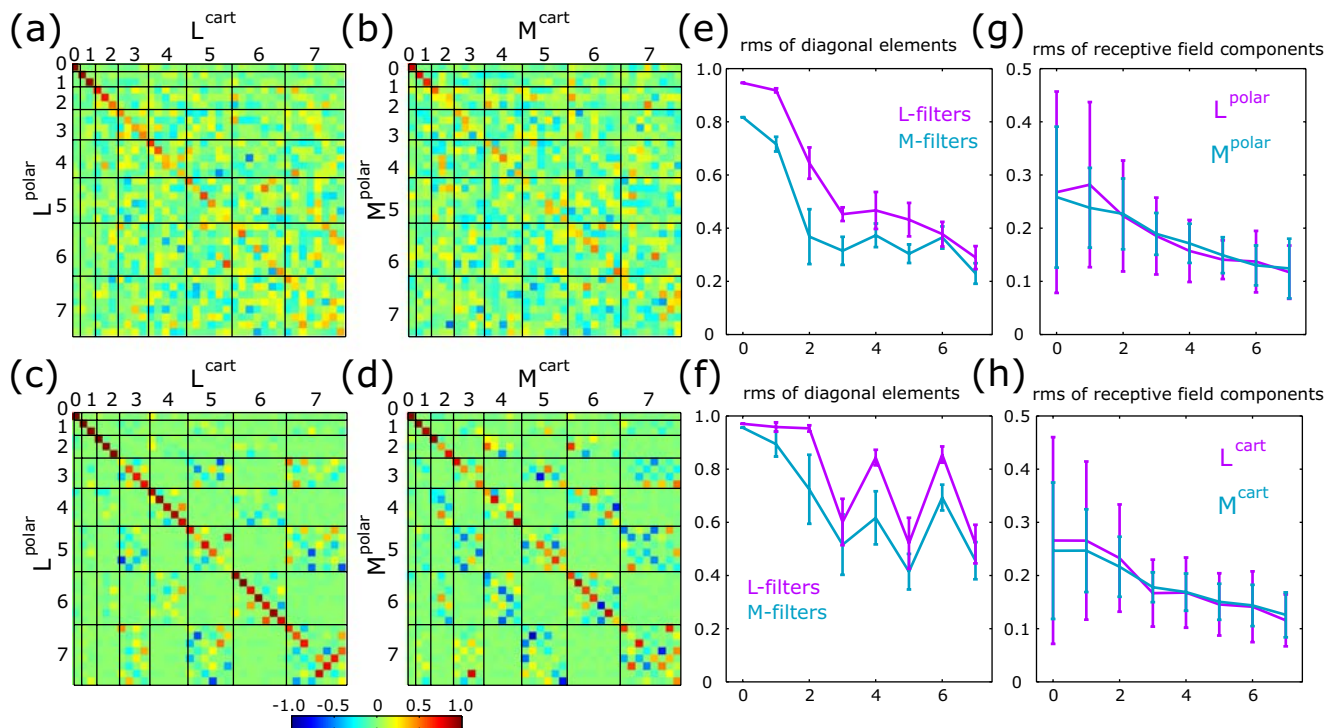


Fig. 5 Comparison of receptive fields derived from Cartesian and polar stimuli. Panels (a)–(b) pseudocolor display of the rotation matrix that is the optimal transformation between a set of receptive fields determined from responses to the two contexts (Cartesian vs. polar) with either modeling method: L-filters (panel a) or M-filters (panel b). Panels (c) and (d): The same calculations, augmented by adding receptive fields reflected around the vertical and horizontal axes. For both models, rotation matrices show that odd-ranked coefficients change more between Cartesian and polar stimuli than the even-ranked coefficients. Display details as in Fig. 4. (e) root-mean-square (rms) of rotation matrix diagonal elements at each rank, derived from L-filters (magenta, based on panel a) and M-filters (blue, based on panel b). (f) Shows the analogous rms values computed using the

symmetry-augmented datasets, taken from panels (c) and (d). Error bars in panels (e) and (f) show standard errors of the mean within each rank. The decrease in the amplitude of diagonal values is strongly non-monotonic: odd ranks are further from the identity than even ranks, showing that odd ranks have a greater systematic context-dependence than even ranks. Panels (g) and (h) show the overall magnitude of the receptive field components derived from responses to Cartesian and polar stimuli. Error bars in panels (g) and (h) show standard deviations across different cells and components within a given rank. The difference between even and odd rank receptive field components in (e) and (f) is not explained by the difference in their magnitude or estimation error, which changes monotonically

tions from the identity matrix. In other words, the higher the rank of the receptive field component, the more it depends on context.

Applying the above symmetry argument removes some of the off-diagonal points of the matrix, but a significant discrepancy between the transformation matrix and the identity persists [Fig. 5(c,d)]. The decrease in amplitude of the on-diagonal elements with increasing rank remains evident [Fig. 5(f)]. Intriguingly, the odd-rank diagonal elements (ranks 3, 5, and 7) are significantly smaller than the even-rank elements (ranks 2, 4, and 6). Since the Procrustes matrix, by definition, is orthogonal, these reductions in the on-diagonal components are accompanied by nonzero components off the diagonal, in the corresponding rows and columns. This, in turn, implies that when odd-rank components are present in one context (e.g., Cartesian), there is a significant tendency for *other* odd-rank components to be present in the alternative context (e.g., polar). That is, the odd-rank components of the receptive fields profiles are modified in a systematic way by context.

No comparable changes are seen in the even-rank components. (For ranks 0 and 1, a comparison is not meaningful, since Cartesian and polar stimuli are identical at these ranks.)

These two findings—a decline in the responses on the diagonal as rank increases, and smaller on-diagonal values for odd ranks than even ranks—are present both for the analysis based on reverse correlation [Fig. 5(c)] and MID [Fig. 5(d)]. The diagonal elements are generally smaller in the MID analysis [Fig. 5(f)]; this is likely due to greater

uncertainty in determination of the MID receptive field profiles (Supplementary Fig. 3).

There are two ways that the above difference between odd and even components can be interpreted: One possibility is that *only* the odd-order components depend on context. The second possibility is that the even-order components also depend on context, but, this dependence is not systematic—so that the best prediction for the size of an even-order component in one context is its size in the other context. Our data indicate that the second alternative is correct.

To see this, we first examine the difference between the responses in each of the two contexts, broken down by parity. As Fig. 6(a) shows, there is approximately the same amount of context-dependent change in the odd- and the even-order components. However, when we compare the actual responses in the polar context with the response predicted by the Procrustes transformation of the Cartesian response, we see a significant difference between odd- and even-order responses. For the even-order responses, virtually none of the context-dependent change is accounted for by the Procrustes transformation [Fig. 6(b)], while for the odd-order responses, approximately half of the context-dependent change can be accounted for in this fashion [Fig. 6(c)]. After the Procrustes transformation, the mean square difference per component was significantly smaller for odd-rank than for even-rank components ($p=0.008$, Wilcoxon paired test). This was also evident from smaller slopes of the regression line between variance values before and after the Procrustes transformation in the case of odd components (0.64) vs. even components (0.86).

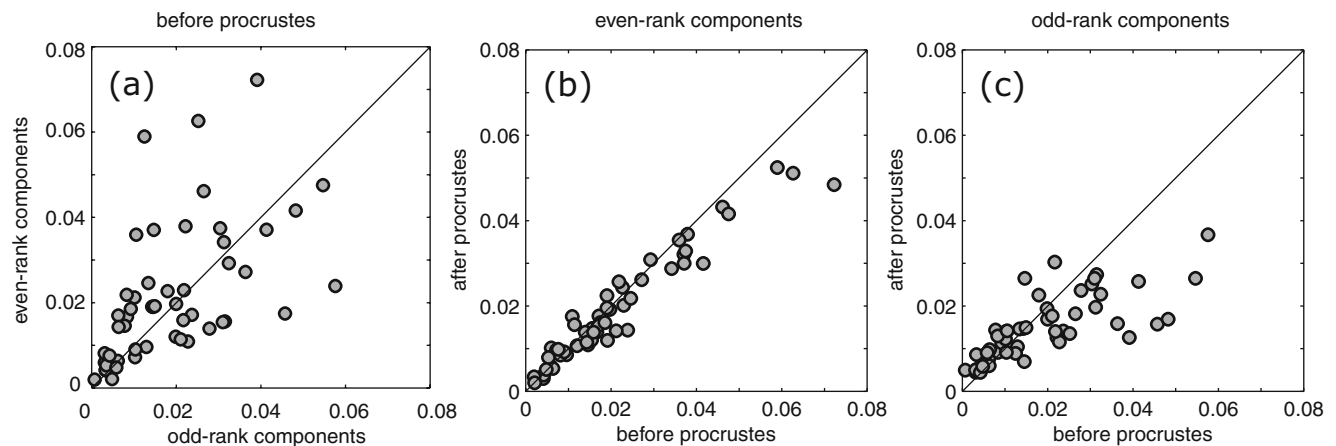


Fig. 6 Analysis of context-dependent changes in receptive fields according to symmetry. Panel (a) shows the overall context-dependent change in receptive fields broken down by the parity of the basis element, and quantified as the mean square difference per basis element computed between components of L-filters for Cartesian and polar stimuli. Even and odd-parity components show an equal amount of overall context dependence. Panels (b) and (c) compare the overall context-dependent change (*abscissa*) to the context-dependent change

that is not accounted for by the Procrustes transformation. Comparison is based on the mean square difference between components of L-filters for polar stimuli and those of L-filters for Cartesian stimuli after the Procrustes transformation. For even-parity components (panel b), the Procrustes transformation does not account for a substantial fraction of the variance, but for the odd-parity components (panel c), it accounts for approximately half of the variance

To emphasize that the Procrustes analysis has identified an overall change in the population that would not be evident from analyzing individual neurons, we show examples of two typical cells (Fig. 7). For each cell, the left column shows the L-filter obtained with Cartesian stimuli, and its decomposition into odd-symmetric and even-symmetric components. The right column shows the analogous filters obtained with polar stimuli. While it is clear that there are changes in receptive field shape [20% of the variance per component in Fig. 7(a), 50% in Fig. 7(b)], the systematic aspect of these changes is not at all apparent.

The middle column of each panel shows the prediction of the polar filters from the Procrustes transformation determined from the entire dataset. For the cell of panel (a), the Procrustes transformation correctly predicts the blobs above and below the center of the receptive field in the odd-rank component of the polar filter (middle row, last two plots). For the cell of panel (b), the Procrustes transformation correctly predicts a loss in the number of lobes in the odd-rank component of the polar filter (middle row, last two plots). The even-rank components also show changes in the effective receptive field for Cartesian and polar plots, but these are not predicted by the Procrustes transformation. These observations are quantified by the average mean square difference per component accounted for by the Procrustes transformation. For the odd ranks, the average mean square difference was decreased by 43% [Fig. 7(a)] and 65% [Fig. 7(b)]; for the even ranks, it was decreased only by 3% [Fig. 7(a)] and 8% [Fig. 7(b)].

3.5 Exclusion of potential confounds

Finally, we consider (and exclude) three potential confounds in this analysis. First, because higher-rank Cartesian and polar stimuli tend to elicit smaller responses than lower-rank stimuli, estimation of higher ranks' contribution to receptive field shape would be subject to greater measurement error, and the similarity of the inferred rotation matrix to the identity would be artifactually reduced. However, this cannot account for our findings. As shown in Fig. 5(g,h), the average magnitude of receptive field components decreases *monotonically* and gradually as the rank increases. In contrast, the difference between the Procrustes matrix and the identity increases abruptly at rank 3, and the high-order odd ranks are more discrepant than adjacent even ranks [Fig. 5(f)]. Second, we consider the effect of imperfect alignment of the center of stimulus patterns relative to the center of the receptive field. While such imperfect alignment was likely present, it would not result in a difference between odd- and even ranks, but would instead tend to *dilute* such differences. This is because a shift in position of an *n*th-rank pattern




Fig. 7 Procrustes transformations for two example cells. *Left column* shows L-filters obtained from Cartesian stimulus set; *middle column* shows the effect of Procrustes transformation [Fig. 5(c)] on these three different profiles, and *right column* shows L-filter obtained from polar stimulus set. The *top row* is the complete L-filter; the *middle row* is the odd-rank component, and the *bottom row* is the even-rank component. Consistent with the analysis of Fig. 6, the Procrustes transformation of the Cartesian L-filter recovers some of the features of the polar L-filter for odd ranks (*middle row*), but not for even ranks (*bottom row*)

leads, in a first approximation, to addition of a small amount of a pattern of rank $n+1$. Thus, any phenomenon that would be confined to either even or odd ranks (if all cells were perfectly centered) would spread to the other parity by misalignment. Thirdly, the fact that odd-rank components show more stimulus-dependent changes is not due to the dataset extension by symmetry (Supplementary Fig. 4).

3.6 Hypothesis: context-dependence arises from independent local perturbations with anti-symmetric profiles

Above we have shown that lability (context-dependence) of receptive field components (a) increases with their rank, i.e. spatial extent, and (b) is more systematic for the odd-symmetric components, than for the even-symmetric ones. We now turn to a consideration of possible explanations for these findings. To begin, we make the general hypothesis that the observed changes in the receptive field profile occur as the net result of independent local perturbations at different positions within the receptive field. The overall change in the receptive field profile is the combined result from all of these separate perturbations. At a circuit level, these local perturbations could correspond to perturbations affecting the strength of synapses from neurons that are presynaptic to the neuron under study, or the intrinsic sensitivities of these presynaptic neurons. (Note we use the term “perturbation” is to refer to an elementary physiologic process that affects the receptive field, but the term “receptive field component” to a basis function used to describe or measure the receptive field profile.) As we now show, the observation that the changes in the odd-symmetric components are more systematic than the even ones has implications both for the spatial profiles of these perturbations, and their distribution. Our reasoning is qualitative, and relies on an observation concerning the interpretation of the Procrustes matrix elements. This observation, which is formalized in the [Appendix](#), is as follows: for a sufficiently diverse set of receptive fields, a matrix element of the Procrustes transformation will deviate from the identity if, on average across the population, perturbations of individual neurons’ receptive fields are spatially correlated with their expression of that basis element.

As a simple example, consider a scenario in which perturbations are Gaussian blobs placed randomly within the receptive field, and of random polarity (ON or OFF). Because of their random nature, they will not be correlated with the receptive field. Thus, they will not lead to a deviation of the Procrustes matrix from the identity. Put another way, although these perturbations might strongly affect each cell’s receptive field, there is no systematic way of predicting how they will change the receptive field from

the Cartesian to the polar context—because they are random. So the best prediction of the polar receptive field is that it is equal to the Cartesian receptive field, i.e., the Procrustes transformation is the identity.

Now suppose that these Gaussian blobs tend to be centered in the receptive field. We further suppose that the sign of the perturbation is correlated with the sign of the receptive field at that location—e.g., a localized ON-perturbation tends to lie in an ON-portion of the receptive field. The even-rank basis elements would in turn tend to be correlated with these perturbations, since they either peak in the receptive center or are zero in the receptive field center but have adjacent flanks of the same sign. Thus, the even-rank Procrustes matrix elements would be expected to deviate from the identity. The odd-rank receptive field components would not be expected to change, because the dot-product of an odd-rank (antisymmetric) component with an even-rank (symmetric) perturbation is zero. A similar argument holds if the perturbations were some other even-symmetric profile, such as a circular center-surround patch, or an oriented bar with antagonistic flankers, or a cosine Gabor. So in the scenario of even-symmetric, centrally-located perturbations, we would expect that only the even-rank receptive field components would change, resulting in even-rank matrix elements of the Procrustes transformation that deviate from the identity. This is contrary to our observations: we find that both even and odd-rank components are labile, and it is the *odd* components, not the even components, that change systematically.

This motivates us to consider a contrasting situation, in which the perturbations are odd-symmetric. Examples of odd-symmetric perturbations include a dyad of adjacent on and off regions, or a sine Gabor, or an idealized simple cell receptive field. The symmetry argument now implies that only the odd rank components will change. Moreover, provided that these perturbations are correlated with the receptive field profile, the changes will be systematic. The result will be deviations of the odd-rank elements of the Procrustes matrix from the identity, as we observe.

In addition to the systematic lability of the odd-rank components, we also need to account for non-systematic lability of the even-rank components [Fig. 6(a)]. There are at least three factors that may contribute to changes in the receptive field profiles that are not identified as systematic by the Procrustes analysis. First, if the On/Off perturbations are not perfectly centered, then their effects will “leak” into the even ranks. Since the sign of the effect will depend on exactly where the perturbations are placed in the receptive field (i.e., how the lobes of the perturbation line up with the lobes of the basis element), this contribution will not be systematic. A scatter of perturbations throughout the receptive field would also account for the greater lability of the higher ranks, compared to the lower ranks. Second, if

the perturbations themselves have a random even-symmetric component, this will lead to random changes in the even-rank basis elements. And finally, any measurement error would contribute to an apparent non-systematic variation of the receptive field components. But measurement error can be no more than the variability of the odd-rank components that is not explained by the Procrustes transformation [Fig. 6(c)].

We sum up the main points of our analysis. First, deviation of the Procrustes matrix from the identity implies that the perturbations of the receptive field are correlated with specific receptive field components. Second, assuming that receptive field lability is a consequence of spatially localized perturbations, our observation that primarily the odd-rank elements of the Procrustes matrix deviate from the identity implies that these local perturbations tend to be odd-symmetric (like On/Off dyads).

4 Discussion

Stimulus-dependent changes in receptive fields have been observed in a variety of species and in both auditory (Rieke et al. 1995; Theunissen et al. 2000, 2001; Woolley et al. 2006) and visual (Baccus and Meister 2002; Chander and Chichilnisky 2001; David et al. 2004; Hosoya et al. 2005; Kim and Rieke 2001; Sharpee et al. 2006) neurons. It is often difficult to go beyond stating that the receptive fields depend on the stimulus set and to provide a parametric description of how receptive fields change between two different stimulus sets. Recent reports point to changes in inhibitory components of visual receptive fields (David et al. 2004), with the strongest effects at low-to-medium spatial frequencies (Sharpee et al. 2006). However, a unified picture of how neural receptive fields change between different stimulus ensembles remains elusive.

4.1 Summary of findings

Here we examine stimulus-dependent changes in V1 receptive fields that occur when high-order statistics change, but first- and second-order statistics (mean luminance, contrast, spatial frequency composition) are held constant. We chose to study the effects of high-order statistics, since their presence distinguishes natural scenes from structureless noise. There are two ingredients in our analysis.

First, we used two complementary methods of receptive field analysis to show that our measures of contextual modulation did not depend on assumptions made in computing receptive fields. Consistent results were obtained with two different models. In one approach, we assumed that the gain function that relates stimulus

components along the receptive field to firing rate was described as a combination of a linear and full-rectifying function (Victor et al. 2006). In the other approach, we used a nonlinear optimization method of maximally informative dimensions without introducing a parametric description for the gain function. We found that both approaches produced receptive fields that were consistent and exhibited similar context-dependent changes.

Second, we introduce a new approach to analyzing these stimulus-dependent changes that seeks to identify consistent changes across a neural population. This analysis yields two findings: the lability of receptive field components increases with their spatial extent, and, perhaps unexpectedly, that odd-symmetric components of the receptive field (i.e., those that are anti-symmetric with respect to 180-degree rotation) change in a more consistent manner across the neural population than even-symmetric ones.

4.2 A mechanistic hypothesis

To explain these findings, we suggest a hypothesis based on the idea that global changes in receptive fields arise as a result of independent, local perturbations. Arguing primarily from symmetry considerations, our observation that odd-rank components change in a more systematic fashion than even-rank ones can be explained if we postulate that the individual perturbation profiles are predominantly odd-symmetric, such as sine-Gabor profiles or a pair of adjacent On/Off subregions (Ringach 2004, 2007; Soodak 1987). The difference between odd- and even-rank components will depend on many factors, including the degree to which the perturbation profiles are clustered at the center of the receptive field, the extent to which they are purely odd-symmetric, whether their occurrence is correlated with the sensitivity profile of the unperturbed receptive field, and the accuracy with which the receptive field profiles can be measured. Our data are not sufficient to analyze the contributions of these factors. However, point-like profiles, circularly-symmetric center-surround profiles, cosine Gabors, and random perturbation profiles will not account for our observation—since they would predict that even-rank receptive field profiles are systematically modified, opposite to what we observe.

The proposition that a localized profile with a pair of adjacent On/Off subregions can serve as a unit for receptive field changes is a natural one, given the known properties of cortical circuitry. Recent theoretical models of primary visual cortex have shown that On/Off receptive field profiles can arise during development simply as a result of “haphazard sampling” from the discrete mosaic of retinal ganglion cells (RGC; Ringach 2004, 2007; Soodak 1987). This argument works because nearest neighbors in the X-

cell RGC mosaic tend to be of similar size but opposite sign (Wassle et al. 1981): pooling inputs from a small area in the visual field is likely to result in two displaced Gaussians of opposite contrast. This “haphazard sampling” argument can account for many known experimental results on the development of visual cortex (Ringach 2007), including the early emergence of simple cells in layer 4 without a prior stage of substantial On/Off overlap and that blocking ON-center RGC cells precludes the development of orientation tuning (Chapman and Godecke 2000). Here we propose that in the adult cortex, the On/Off dyads that are postulated to underlie orientation tuning also underlie contextual modulation—that is, contextual modulation of receptive fields involves adjustments at the level of the output of an On/Off dyad, rather than separate adjustments of its components.

Our analysis focused on changes in the spatial sensitivity profile, and used only a specific, highly non-natural, set of stimuli. But if our hypothesis is correct, then it should generalize. In particular, we would predict that perturbations similar to On/Off dyads, localized in space and time, would account for context-dependence of the *temporal* structure of receptive fields (Chander and Chichilnisky 2001; Hsu et al. 2004; Sharpee et al. 2006; Smirnakis et al. 1997), since these temporal context-dependent changes are likely governed by the same principles. We would also predict that the same proposed mechanism of context-dependent changes could also explain receptive field changes under more natural conditions (David et al. 2004; Sharpee et al. 2006; Theunissen et al. 2000, 2001; Woolley et al. 2006), for example between natural stimuli and correlated Gaussian noise with the same second-order structure (Felsen et al. 2005).

4.3 A method for studying receptive field transformations

Our analysis rests on a novel method for studying how receptive fields change between two conditions. Although our two conditions (Cartesian and polar stimulus patches) are defined by highly structured statistics, this is by no means a prerequisite for the approach. The approach can be applied without regard to the nature of the two conditions, which may include a change in the statistics of sensory inputs, a behavioral manipulation such as modulation of attention, or even a pharmacologic manipulation. The main advantages of this approach are that it (a) does not rely on preconceived notions of what types of changes may be present and (b) can identify patterns that would not be apparent if each receptive field were studied independently.

The main prerequisites of the method are that the same neurons are studied under both conditions, and that their

receptive fields can be represented as vectors in a finite-dimensional space. Here, we described receptive fields by 36-dimensional vectors of weights with respect to Hermite functions, but other basis sets could be used, including spatiotemporal basis functions or Fourier components. Once the basis set is chosen, the receptive fields measured under each condition (context) correspond to a set of vectors, one vector for each neuron. The final, crucial step is to find a rotation matrix that most closely transforms the set of vectors measured under one set of conditions into the set of vectors measured under the other set of conditions. To do so, we use an algorithm that was first developed in crystallography (Kabsch 1976).

At this stage, it is important that the number of neurons (or receptive-field pairs under the two different conditions) is larger than the number of dimensions in which the best transformation is sought. Otherwise, the problem of identifying the optimal rotation would be underdetermined.

Since we were interested in changes in shape but not overall sensitivity (which, here, would be absorbed into the nonlinear stage, and was previously analyzed, Victor et al. 2006), we did not allow for a dilation following the rotation. However, the method can be readily extended to allow for dilations that follow the rotation matrix, or to look for general linear transformations between the two sets of vectors—thus providing an analysis of increases or decreases in sensitivity.

5 Conclusion

Our analysis of how V1 receptive fields respond to changes in higher-order stimulus statistics shows that stability of receptive field components depends on both their spatial extent and symmetry properties. Qualitatively, this finding can be explained by a simple, physiologically plausible mechanism based on the idea that receptive fields change as a result of independent local perturbation events. To account for our data, local perturbation profiles must be predominantly anti-symmetric, e.g., constructed from a pair of adjacent On/Off subregions. Recent theoretical models of cortical development argue that such a profile can result just from local pooling in space of a small number of RGC cells and is sufficient for seeding the formation of cortical circuits. Perhaps the mechanisms based on On/Off pairs that guide the development of visual cortex continue to shape neural receptive fields in the adulthood. Our approach can be used to determine whether similar mechanisms underlie changes in spatiotemporal receptive fields, and adaptation to context driven by natural scene statistics.

Note added in proof: An analysis of an additional 57 macaque V1 neurons showed results comparable to those presented here.

Acknowledgements This research was supported by the Center for Theoretical Biological Physics (NSF PHY-0822283), a grant from the Swartz Foundation and the Mentored Quantitative Career Development Award MH068904 to TS from NIMH, and EY9314 to JDV. Computing resources were provided by the National Science Foundation through Partnerships for Advanced Computational Infrastructure at the Distributed Terascale Facility and Terascale Extensions.

Appendix

Relationship between Procrustes matrix and receptive field differences

In order to understand possible mechanisms that might explain the observed difference in context-dependent changes for even- and odd-rank components, we analyze the relationship between the elements of the Procrustes matrix and receptive field changes. The Procrustes transformation implemented in Figs. 4, 5, 6, 7 determines the best rotation matrix between the two sets of receptive fields obtained in the two stimulus contexts. We will work in the approximation where changes in the stimulus components are small compared to the magnitude of stimulus components. This approximation implies that the best rotation matrix will also be close to the best linear transformation that includes both rotation and scaling, because scaling factor will be close to 1. This significantly simplifies analysis, making the relationship between elements of the Procrustes matrix and changes in receptive field components more transparent. In this approximation, the elements C_{nk} of the matrix given by the difference between the Procrustes matrix and the unit matrix can be found by minimizing mean square error:

$$\sum_{n,i} \left(\delta a_n^{(i)} - C_{nk} a_k^{(i)} \right)^2$$

where $a_k^{(i)}$ represents the k th component of receptive field of the i th neuron obtained from the first, for example Cartesian, stimulus set. We denote the difference between receptive field coefficients for the two stimulus sets by $\delta a_k^{(i)}$, which we assume to be small, so that the sum of identity matrix and C is nearly orthogonal. The optimal matrix elements C_{nk} can be found from the following equation:

$$\sum_k C_{nk} \left\langle a_k^{(i)} a_j^{(i)} \right\rangle_i = \left\langle \delta a_n^{(i)} a_j^{(i)} \right\rangle_i. \tag{3}$$

Here all averages are computed across the population of cells.

If we first assume that receptive fields are sufficiently diverse so that receptive field components are uncorrelated, the expression for matrix C can be simplified:

$$C_{nk} = \frac{\left\langle \delta a_n^{(i)} a_k^{(i)} \right\rangle_i}{\left\langle a_k^{(i)2} \right\rangle_i}. \tag{4}$$

Thus, for small context dependent changes $\delta a_k^{(i)}$ in receptive fields, the deviations of the Procrustes matrix from the unit matrix are determined the average product between k th receptive field components and the change in the n th receptive field component, normalized by the average magnitude of k th component across all receptive fields. The assumption of uncorrelation is clearly a strong one, but a weaker assumption—namely, that there are no correlations between even and odd components see [Fig. 4(e)]—suffices to support the analysis in the main text. That is, Eq. (3) shows that odd-parity elements of the Procrustes matrix imply correlations of the perturbations with odd-order basis elements.

References

- Baccus, S. A., & Meister, M. (2002). Fast and slow contrast adaptation in retinal circuitry. *Neuron*, 36, 909–919. doi:10.1016/S0896-6273(02)01050-4.
- Bonds, A. B. (1989). Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience*, 2, 41–55.
- Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *The Journal of Neuroscience*, 17, 8621–8644.
- Carandini, M., Heeger, D. J., & Senn, W. (2002). A synaptic explanation of suppression in visual cortex. *The Journal of Neuroscience*, 22, 10053–10065.
- Chander, D., & Chichilnisky, E. J. (2001). Adaptation to temporal contrast in primate and salamander retina. *The Journal of Neuroscience*, 21, 9904–9916.
- Chapman, B., & Godecke, I. (2000). Cortical cell orientation selectivity fails to develop in the absence of ON-center retinal ganglion cell activity. *The Journal of Neuroscience*, 20, 1922–1930.
- Cox, T. F., & Cox, M. A. A. (2000). *Multidimensional scaling*, Second edn: Chapman and Hall).
- Das, A., & Gilbert, C. D. (1999). Topography of contextual modulations mediated by short-range interactions in primary visual cortex. *Nature*, 399, 655–661. doi:10.1038/21371.
- David, S. V., Vinje, W. E., & Gallant, J. L. (2004). Natural stimulus statistics alter the receptive field structure of V1 neurons. *The Journal of Neuroscience*, 24, 6991–7006. doi:10.1523/JNEUROSCI.1422-04.2004.
- de Boer, E., & Kuypers, P. (1968). Triggered correlation. *IEEE Transactions on Bio-Medical Engineering*, 15, 169–179. doi:10.1109/TBME.1968.4502561.
- De Valois, R. L., Albrecht, D. G., & Thorell, L. G. (1982a). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, 22, 545–559. doi:10.1016/0042-6989(82)90113-4.

- De Valois, R. L., & Thorell, L. G. (1988). Spatial vision. New York: Oxford University Press.
- De Valois, R. L., Yund, E. W., & Helper, N. (1982b). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, 22, 531–544. doi:10.1016/0042-6989(82)90112-2.
- Efron, B., & Tibshirani, R. J. (1998). An introduction to the bootstrap. London: Chapman and Hall/CRC.
- Felsen, G., Touryan, J., Han, F., & Dan, Y. (2005). Cortical sensitivity to visual features in natural scenes. *PLoS Biology*, 3, e342. doi:10.1371/journal.pbio.0030342.
- Freeman, T. C., Durand, S., Kiper, D. C., & Carandini, M. (2002). Suppression without inhibition in visual cortex. *Neuron*, 35, 759–771. doi:10.1016/S0896-6273(02)00819-X.
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9, 181–198.
- Hosoya, T., Baccus, S. A., & Meister, M. (2005). Dynamic predictive coding by the retina. *Nature*, 436, 71–77. doi:10.1038/nature03689.
- Hsu, A., Woolley, S. M., Fremouw, T. E., & Theunissen, F. E. (2004). Modulation power and phase spectrum of natural sounds enhance neural encoding performed by single auditory neurons. *The Journal of Neuroscience*, 24, 9201–9211. doi:10.1523/JNEUROSCI.2449-04.2004.
- Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurons in the cat's striate cortex. *The Journal of Physiology*, 148, 574–591.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, 195, 215–243.
- Kabsch, W. (1976). A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica. Section A*, 32, 922–923. doi:10.1107/S0567739476001873.
- Kim, K. J., & Rieke, F. (2001). Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. *The Journal of Neuroscience*, 21, 287–299.
- Meister, M., & Berry, M. J. (1999). The neural code of the retina. *Neuron*, 22, 435–450. doi:10.1016/S0896-6273(00)80700-X.
- Morrone, M. C., & Burr, D. C. (1988). Feature detection in human vision: a phase-dependent energy model. *Proceedings of the Royal Society of London—Series B: Biological Sciences*, 235, 221–245.
- Ohzawa, I., Sclar, G., & Freeman, R. D. (1985). Contrast gain control in the cat's visual system. *Journal of Neurophysiology*, 54, 651–667.
- Oppenheim, A. V., & Lim, J. S. (1981). The importance of phase in signals. *Proceedings of the IEEE*, 69, 529–541.
- Priebe, N. J., & Ferster, D. (2006). Mechanisms underlying cross-orientation suppression in cat visual cortex. *Nature Neuroscience*, 9, 552–561. doi:10.1038/nn1660.
- Reich, D. S. (2001) PhD Thesis: Information encoding by individual neurons and groups of neurons in primary visual cortex, The Rockefeller University, New York.
- Rieke, F., Bodnar, D. A., & Bialek, W. (1995). Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory afferents. *Proceedings of the Royal Society of London—Series B: Biological Sciences*, 262, 259–265. doi:10.1098/rspb.1995.0204.
- Ringach, D. L. (2004). Haphazard wiring of simple receptive fields and orientation columns in visual cortex. *Journal of Neurophysiology*, 92, 468–476. doi:10.1152/jn.01202.2003.
- Ringach, D. L. (2007). On the origin of the functional architecture of the cortex. *PLoS ONE*, 2, e251. doi:10.1371/journal.pone.0000251.
- Ruderman, D. L., & Bialek, W. (1994). Statistics of natural images: Scaling in the woods. *Physical Review Letters*, 73, 814–817. doi:10.1103/PhysRevLett.73.814.
- Schwartz, O., Pillow, J. W., Rust, N. C., & Simoncelli, E. P. (2006). Spike-triggered neural characterization. *Journal of Vision (Charlottesville, Va.)*, 6, 484–507. doi:10.1167/6.4.13.
- Shapley, R. M., & Victor, J. D. (1978). The effect of contrast on the transfer properties of cat retinal ganglion cells. *The Journal of Physiology*, 285, 275–298.
- Sharpee, T., Rust, N. C., & Bialek, W. (2004). Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Computation*, 16, 223–250. doi:10.1162/089976604322742010.
- Sharpee, T. O., Sugihara, H., Kurgansky, A. V., Rebrik, S. P., Stryker, M. P., & Miller, K. D. (2006). Adaptive filtering enhances information transmission in visual cortex. *Nature*, 439, 936–942. doi:10.1038/nature04519.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24, 1193–1216. doi:10.1146/annurev.neuro.24.1.1193.
- Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W., & Meister, M. (1997). Adaptation of retinal processing to image contrast and spatial scale. *Nature*, 386, 69–73. doi:10.1038/386069a0.
- Soodak, R. E. (1987). The retinal ganglion cell mosaic defines orientation columns in striate cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 84, 3936–3940. doi:10.1073/pnas.84.11.3936.
- Theunissen, F. E., David, S. V., Singh, N. C., Hsu, A., Vinje, W. E., & Gallant, J. L. (2001). Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network*, 12, 289–316.
- Theunissen, F. E., Sen, K., & Doupe, A. J. (2000). Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *The Journal of Neuroscience*, 20, 2315–2331.
- Victor, J., & Shapley, R. (1980). A method of nonlinear analysis in the frequency domain. *Biophysical Journal*, 29, 459–483.
- Victor, J. D., Mechler, F., Repucci, M. A., Purpura, K. P., & Sharpee, T. (2006). Responses of V1 neurons to two-dimensional hermite functions. *Journal of Neurophysiology*, 95, 379–400. doi:10.1152/jn.00498.2005.
- Vinje, W. E., & Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287, 1273–1276. doi:10.1126/science.287.5456.1273.
- Vinje, W. E., & Gallant, J. L. (2002). Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. *The Journal of Neuroscience*, 22, 2904–2915.
- Wassle, H., Boycott, B. B., & Illing, R. B. (1981). Morphology and mosaic of on- and off-beta cells in the cat retina and some functional considerations. *Proceedings of the Royal Society of London—Series B: Biological Sciences*, 212, 177–195.
- Woolley, S. M., Gill, P. R., & Theunissen, F. E. (2006). Stimulus-dependent auditory tuning results in synchronous population coding of vocalizations in the songbird midbrain. *The Journal of Neuroscience*, 26, 2499–2512. doi:10.1523/JNEUROSCI.3731-05.2006.