

Geometrical Computations Explain Projection Patterns of Long-Range Horizontal Connections in Visual Cortex

Ohad Ben-Shahar and Steven W. Zucker

Department of Computer Science and the Interdepartmental Neuroscience Program

Yale University, New Haven, Connecticut 06520, USA

Neurons in primary visual cortex respond selectively to oriented stimuli such as edges and lines. The long-range horizontal connections between them are thought to facilitate contour integration. While many physiological and psychophysical findings suggest that colinear or association-field models of good continuation dictate particular projection patterns of horizontal connections to guide this integration process, significant evidence of interactions inconsistent with these hypotheses is accumulating. We first show that natural random variations around the colinear and association-field models cannot account for these inconsistencies, a fact that motivates the search for more principled explanations. We then develop a model of long-range projection fields that formalizes good continuation based on differential geometry. The analysis implicates curvature(s) in a fundamental way and the resulting model explains both consistent data and apparent outliers. It quantitatively predicts the (typically ignored) spread in projection distribution, its non-monotonic variance, and the differences found between individual neurons. Surprisingly, and for the first time, this model also indicates that texture (and shading) continuation can serve as alternative and complimentary functional explanations to contour integration. Because current anatomical data support both (curve and texture) integration models equally, and because both are important computationally, new testable predictions are derived to allow their differentiation and identification.

1 Introduction

The receptive fields (RFs) of neurons in visual cortex characterize their response to patterns of light in the visual field. In primary visual cortex this response is often selective for stimulus orientation in a small region (Hubel and Wiesel, 1977). The clustered long-range horizontal connections between such cells (Rockland and Lund, 1982) link those with non-overlapping RFs and are thought to facilitate contour integration (Field *et al.*, 1993). However, there is no direct physiological evidence that these connections *only* support curve integration while, on the other hand, there also remains much ambiguity about the precise connections required to support the integration of curves. Our goal in this paper is to address both of these concerns.

1.1 Biological data and integration models

The argument that associates long-range horizontal connections with curve integration begins with the realization that the finite spatial extent of RFs and their broad orientation tuning lead to significant uncertainties in the position and the local orientation measured from visual stimuli. This causes a further uncertainty in determining which of the many nearby RFs signal the next section of a curve (Fig. 1a).

[Figure 1 about here]

All these uncertainties underlying curve integration can be reduced by interactions between neurons whose RFs are close in retinotopic coordinates. Starting with Mitchison and Crick and their hypothesis about interactions between iso-oriented RFs (Mitchison and Crick, 1982), physiological and anatomical findings have been accumulating to suggest a roughly *colinear* interaction. The main evidence supporting this conclusion is based on the distribution of angular differences between preferred orientations of connected cells. These distributions are computed by taking the orientation difference between a target cell and every other cell it is connected to with a long-range horizontal

connection. Indeed, as is exemplified in Fig. 1b, these distributions have been shown to be unimodal on average, with maximal interaction between iso-oriented RFs (Ts'o *et al.*, 1986; Gilbert and Wiesel, 1989; Weliky *et al.*, 1995; Schmidt *et al.*, 1997; Buzás *et al.*, 1998; Bosking *et al.*, 1997; Malach *et al.*, 1993; Sincich and Blasdel, 2001; Schmidt and Löwel, 2002). Furthermore, direct anatomical studies reveal long-range interactions between co-axial cells (Bosking *et al.*, 1997; Schmidt *et al.*, 1997) and indirect psychophysical experiments report a general *association field* (Field *et al.*, 1993; Kapadia *et al.*, 1995, 2000) which emphasizes straight or slowly varying continuations while allowing some support for more rapidly varying continuations as well (Fig. 2a).

[Figure 2 about here]

With the accumulation of these data, however, are a growing number of observations that are difficult to reconcile with the intuition that neural spatial integration is based on colinearity, or that it only serves curve integration. Facilitory interaction between cells of significant orientation difference (Kapadia *et al.*, 1995), short range co-axial *inhibition* (Polat and Sagi, 1993), iso-orientation *side* facilitation (Adini *et al.*, 1997), and strong correlations between iso-oriented, non-overlapping, and *parallel* receptive fields (Ts'o *et al.*, 1986, see bottom cell on page 1163 or top cell on page 1164) are functionally inconsistent. Evidence of cross-orientation (Matsubara *et al.*, 1985; Kisvárdy *et al.*, 1997) and non-axial (Gilbert and Wiesel, 1989) connections, plus roughly isotropic retinotopic extent (Malach *et al.*, 1993; Sincich and Blasdel, 2001), suggest anatomical inconsistencies.

These inconsistencies prompt a closer examination of the interactions within visual cortex and their population statistics. As the evidence suggests, individual cells, or small collections of adjacent cells captured in tracer injections, may have qualitatively different connectivity distributions (Bosking *et al.*, 1997): some are narrow and high while others are very wide, as is illustrated in Fig 1c. When averaged, the pooled distribution of long-range connections (e.g., those extending beyond $500\mu\text{m}$ in Bosking *et al.*, 1997)

is

- unimodal,
- peaks at zero orientation offset,
- indicates a non-negligible fraction of connections linking cells of significantly different orientation preferences (Malach *et al.*, 1993; Kisvárdy *et al.*, 1994; Kisvárdy *et al.*, 1997; Bosking *et al.*, 1997),
- crosses the uniform distribution at approximately ± 40 degrees ¹, and
- has a *non-monotonically* changing variance as the orientation difference increases (Malach *et al.*, 1993; Bosking *et al.*, 1997).

Neither colinearity nor association field models predict all of these features. While both models imply unimodal pooled distributions over orientation differences (Fig. 2b), they also suggest a *fixed* projection field and thus neither predicts *any* variance for the pooled distribution, let alone a non-monotonic one. Furthermore, colinearity is clearly at odds with the significant spread in the distribution of connections over orientation differences, whether it is measured via extracellular injections (e.g., Bosking *et al.*, 1997) or the more elaborate intracellular protocol (Buzás *et al.*, 1998).

The data in Bosking *et al.* (1997) contain one injection site of possibly different connection distribution which may substantially contribute to the non-monotonic nature of the variance. Since the variance will become central to this paper, we examined whether or not this statistical feature depends critically on this one, possibly outlier, measurement. We re-analyzed the data from Bosking *et al.* (1997) after removing the data from this injection site and calculating the statistical properties of the rest. We further examined the robustness of the non-monotonicity by running two additional analyses: one in

¹This crossing point provides a reference for the bias of projection patterns toward particular orientations; considering the offsets where the connection distribution crosses the uniform line quantifies this bias in way independent of scale or quantization level.

which we removed the sample points (one from each orientation bin) that contribute the most to the variance, and another in which we removed those sample points (again, one from each bin) that maximized local changes in the variance. In all these tests, including the last one which flattens the variance the most, the tri-modal non-monotonicity, and the two local minima at $\pm 30^\circ$, were preserved. All these findings suggest that the non-monotonicity of the variance is a critical feature that deserves attention both from biologists and from modelers.

1.2 Integration models and random physiological variations

It is tempting to explain the apparent anomalies and inconsistencies between the predicted and measured distributions of long-range horizontal connections as random physiological variations, for example by asserting that anatomy only *approximates* the correct connections. We tested this explanation by applying different noise models to the colinearity and association field connectivity distributions from Fig. 2, and checked whether the resultant pooled distributions possess the properties listed above. The results of the most natural noise model are illustrated in Fig. 3b. Under this model, each long-range horizontal connection, ideally designated to connect cells of orientation difference $\Delta\theta$, is shifted to connect cells of orientation difference $\Delta\theta + \epsilon_\sigma$, where ϵ_σ is a wrapped Gaussian (i.e., normally distributed and wrapped on \mathcal{S}^1) random variable with zero mean and variance σ (see appendix for details). As the figure shows, it takes an overwhelming amount of noise (s.d. $\geq 35^\circ$) to transform the colinear distribution to one that resembles the measured data in terms of spread and peak height, but the non-monotonic behavior of the variance is *never* reproduced (for space considerations we omit the results of other connection-based noise models, or the noisy distributions based on the association field model, all of which were even less reminiscent of the measured physiological data).

A second possible source for the inconsistencies between the predicted and mea-

sured distributions may be the extracellular injection protocol commonly in use by physiologists to trace long-range horizontal connections (e.g., Gilbert and Wiesel, 1989; Malach *et al.*, 1993; Kisvárdy *et al.*, 1994; Kisvárdy *et al.*, 1997; Bosking *et al.*, 1997; Schmidt *et al.*, 1997; Sincich and Blasdel, 2001). Due to the site-selection procedure used, cells stained by these injections are likely to have similar orientation preferences (e.g. Bosking *et al.* (1997), page 2113; or Schmidt *et al.* (1997), page 1084). However, their orientation tuning may nevertheless be different, sometimes significantly (e.g., note such a cell in Bosking *et al.*, 1997, Fig. 4B). Consequently, the distribution of presynaptic terminals (boutons) traced from the injection site may incorporate an artificial, random spread relative to the single orientation typically assumed at the injection site. Preliminary evidence from a recently developed single-cell protocol (Buzás *et al.*, 1998) suggests that leakage in the injection site cannot bridge the gap between the predicted colinear distribution and those measured anatomically. However, we also examined this possibility computationally by modeling the leakage in the injection site as a wrapped Gaussian random variable of predefined variance². The base distribution (colinear or association field) of the computational cells selected by this process were then summed up and normalized, and the resultant (random) distribution was attributed to the original cell representing the injection site. Repeating this process many times yielded a collection of (different) distributions, for which we calculated an average and variance (see appendix for details). The results are illustrated in Fig. 3c. Similar to random variations at the level of individual connections, here too it takes an overwhelming amount of noise (s.d. $\geq 35^\circ$) to transform the colinear distribution to one that resembles the measured data in terms of spread and peak height, but the non-monotonic behavior of the variance is *never* reproduced.

[Figure 3 about here]

²A wrapped Gaussian model was particularly suitable here due to the injection site selection protocol typically used in the extracellular injection protocol; see appendix).

The thinking around long-range horizontal connections has been dominated by their first-order statistics and its peak at zero orientation offset. However, the non-monotonicity of the variance was first reported almost a decade ago (Fig. 3d in Malach *et al.*, 1993) and we have further confirmed it from the more detailed measurements in Bosking *et al.* (1997), as was illustrated in Fig. 3a. Since neither colinearity nor association field models can explain this aspect of the physiological data, even if much noise is allowed, it is necessary to consider whether this and the other subtle properties of the pooled data reflect genuine functional properties of long-range horizontal connections. We therefore developed a geometric model of projection patterns, and examined quantitatively both pooled connection statistics and connectivity patterns of individual cells generated by it. Since many findings suggest that long-range horizontal connections are primarily excitatory, especially those extending beyond one hypercolumn (Ts'o *et al.*, 1986; Gilbert and Wiesel, 1989; Kapadia *et al.*, 1995; Kisvárdy *et al.*, 1997; Buzás *et al.*, 1998; Sincich and Blasdel, 2001), our model concentrates on this class of connections.

2 From differential geometry to integration models

Curve integration, the hypothesized functional role ascribed to long-range horizontal connections, is naturally based in differential geometry. The tangent, or the local linear approximation to a curve, abstracts orientation preference, and the collection of all possible tangents at each (retinotopic) position can be identified with the orientation hypercolumn (Hubel and Wiesel, 1977). Formally, since position takes values in the plane \mathbb{R}^2 (think of image coordinates x, y) and orientation in the circle S^1 (think of an angle θ varying between 0 and 2π), the primary visual cortex can be abstracted as the product space $\mathbb{R}^2 \times S^1$ (Fig. 4). Points in this space represent both position and orientation to abstract visual edges of given orientation at a particular spatial (i.e., retinotopic) position. It is in this space that our modeling takes place.

[Figure 4 about here]

Since any single tangent is the limit of any smooth curve passing through a given (retinotopic) point in a given direction, the question of curve integration becomes one of determining how two tangents at nearby positions are related. (Colinearity, for example, asserts that the tangent orientation hardly changes for small displacements along the curve). In general terms, the angular difference between RFs captures only one part of the relationship between nearby tangents; their relative spatial offset also must be considered. Thus, in the mathematical abstraction, relationships between tangents correspond to relationships between points in $\mathbb{R}^2 \times S^1$. Physiologically, these relationships are carried by the long-range horizontal connections, with variation in retinotopic position corresponding to \mathbb{R}^2 , and variation along orientation hypercolumns corresponding to S^1 (Fig. A). Determining them amounts, in mathematical terms, to determining what is called a connection structure. As we discuss in the rest of this paper, the relationship between these two types of “connections”, the mathematical and the physiological, is more than linguistic.

[Figure 5 about here]

2.1 The geometry of orientation in the retinal plane

Orientation in the 2D (retinal) plane is best represented as a unit length tangent vector $\hat{\mathbf{E}}(\vec{\mathbf{q}})$ attached to point of interest $\vec{\mathbf{q}} = (x, y) \in \mathbb{R}^2$. Having such a tangent vector attached to every point of an object of interest (e.g., a smooth curve or oriented texture) results in a unit length *vector field* (O’Neill, 1966). Assuming good continuation (Wertheimer, 1955), a small translation $\vec{\mathbf{V}}$ from the point $\vec{\mathbf{q}}$ results in a small change (i.e., rotation) in the vector $\hat{\mathbf{E}}(\vec{\mathbf{q}})$. To apply techniques from differential geometry, a suitable coordinate frame $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$ is placed at the point $\vec{\mathbf{q}}$ and the basis vector $\hat{\mathbf{E}}_T$ is identified with $\hat{\mathbf{E}}(\vec{\mathbf{q}})$ – the tangent vector at $\vec{\mathbf{q}}$ (Fig. 6). Note that $\hat{\mathbf{E}}_T$ is drawn at an angle θ – the local orientation measured relative to the horizontal axis in retinotopic coordinates – such that $(\vec{\mathbf{q}}, \theta) \in \mathbb{R}^2 \times S^1$. Nearby tangents are displaced both in po-

sition and orientation according to the *covariant derivatives* of the underlying pattern. These covariant derivatives, $\nabla_{\vec{v}} \hat{\mathbf{E}}_T$ and $\nabla_{\vec{v}} \hat{\mathbf{E}}_N$, are naturally represented as vectors in the basis $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$ itself:

$$\begin{pmatrix} \nabla_{\vec{v}} \hat{\mathbf{E}}_T \\ \nabla_{\vec{v}} \hat{\mathbf{E}}_N \end{pmatrix} = \begin{bmatrix} w_{11}(\vec{V}) & w_{12}(\vec{V}) \\ w_{21}(\vec{V}) & w_{22}(\vec{V}) \end{bmatrix} \begin{pmatrix} \hat{\mathbf{E}}_T \\ \hat{\mathbf{E}}_N \end{pmatrix}. \quad (1)$$

The coefficients $w_{ij}(\vec{V})$, known as *1-forms*, are functions of the displacement direction vector \vec{V} , and since the basis $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$ is orthonormal, they are skew-symmetric $w_{ij}(\vec{V}) = -w_{ji}(\vec{V})$. Thus $w_{11}(\vec{V}) = w_{22}(\vec{V}) = 0$ and the system reduces to:

$$\begin{pmatrix} \nabla_{\vec{v}} \hat{\mathbf{E}}_T \\ \nabla_{\vec{v}} \hat{\mathbf{E}}_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(\vec{V}) \\ -w_{12}(\vec{V}) & 0 \end{bmatrix} \begin{pmatrix} \hat{\mathbf{E}}_T \\ \hat{\mathbf{E}}_N \end{pmatrix}. \quad (2)$$

This last system is known as Cartan's *connection equation* (O'Neill, 1966) and $w_{12}(\vec{V})$ is called the *connection form*. Since $w_{12}(\vec{V})$ is linear in \vec{V} , it can be represented in terms of $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$:

$$w_{12}(\vec{V}) = w_{12}(a \hat{\mathbf{E}}_T + b \hat{\mathbf{E}}_N) = a w_{12}(\hat{\mathbf{E}}_T) + b w_{12}(\hat{\mathbf{E}}_N) .$$

The relationship between nearby tangents is thus governed by two scalars at each point. We define them as follows

$$\begin{aligned} \kappa_T &\triangleq w_{12}(\hat{\mathbf{E}}_T) \\ \kappa_N &\triangleq w_{12}(\hat{\mathbf{E}}_N) \end{aligned} \quad (3)$$

and interpret them as *tangential* (κ_T) and *normal* (κ_N) curvatures, since they represent a directional rate of change of orientation in the tangential and normal directions, respectively.

[Figure 6 about here]

While the connection equation describes the local behavior of orientation for the general two dimensional case, it is equally useful for the one-dimensional case of curves.

Now, only $\nabla_{\hat{\mathbf{E}}_T}$ is relevant and Equation 2 simplifies to

$$\begin{pmatrix} \nabla_{\hat{\mathbf{E}}_T} \hat{\mathbf{E}}_T \\ \nabla_{\hat{\mathbf{E}}_T} \hat{\mathbf{E}}_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(\hat{\mathbf{E}}_T) \\ -w_{12}(\hat{\mathbf{E}}_T) & 0 \end{bmatrix} \begin{pmatrix} \hat{\mathbf{E}}_T \\ \hat{\mathbf{E}}_N \end{pmatrix}. \quad (4)$$

In its more familiar form, where T, N , and κ replace $\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N$, and κ_T , respectively, this is the classical *Frenet equation* (O’Neill, 1966) (primes denote derivatives by arclength):

$$\begin{pmatrix} T' \\ N' \end{pmatrix} = \begin{bmatrix} 0 & \kappa \\ -\kappa & 0 \end{bmatrix} \begin{pmatrix} T \\ N \end{pmatrix}. \quad (5)$$

2.2 Integration models and projection patterns of horizontal connections

The geometrical analysis discussed above and illustrated in Fig 6 shows how the relationship between nearby tangents depends on the covariant derivative: for *curves* the connection is dictated by one curvature; for *texture flows* – or oriented 2-dimensional patterns – two curvatures are required. By estimating these quantities at a given retinal point $\vec{\mathbf{q}}$ it is possible to approximate the underlying geometrical object, and thus a coherent distribution of tangents, around $\vec{\mathbf{q}}$. This, in turn, can be used to model the set of horizontal connections that are required to facilitate the response of a cell if its RF is embedded in a visual context that reflects good continuation. Naturally, to describe such a local approximation and to use it for building projection patterns, the appropriate domain of integration must be determined. However, since RF measurements provide only the tangent, possibly curvature (Dobbins *et al.*, 1987; Versavel *et al.*, 1990), but not whether the stimulus pattern is a curve (1D) or a texture (2D), it is necessary to consider continuations for both.

Since estimates of curvature at point $\vec{\mathbf{q}}$ hold in a neighborhood containing the tangent, the discrete continuation for a *curve* is commonly obtained by approximating it locally by its osculating circle (do Carmo, 1976) and quantizing. This relationship, which is based on the constancy of curvature around $\vec{\mathbf{q}}$, is known as *co-circularity* (Parent and

Zucker, 1989; Zucker *et al.*, 1989; Sigman *et al.*, 2001; Geisler *et al.*, 2001), and in $\mathbb{R}^2 \times S^1$ it takes the form of a *helix* (Fig. 7a-b). Different estimates of curvature give rise to different helices whose points define both the spatial position and the local orientation of nearby RFs that are compatible with the estimate at \vec{q} (Fig. 7c). Together, all these compatible cells induce a curvature-based field of long-range horizontal connections (Fig. 7a-c and Fig. 8a-d). While different curvatures induce different projection fields, the “sum” over curvatures gives an association field (Fig. 8e) reminiscent of recent psychophysical findings (Field *et al.*, 1993). Note, however, that as a psychophysical entity, the association field is not necessarily a one-to-one reflection of connectivity patterns in the visual cortex. In fact, representing a “cognitive union” across displays of different continuations, the association field is unlikely to characterize any single cell.

[Figure 7 about here]

Similar considerations can be applied toward the local approximation of texture flows, although now the construction of a rigorous local model is slightly more challenging. Unlike curves, this model must depend upon the estimate of *two* curvatures at the point \vec{q} , $K_T = \kappa_T(\vec{q})$ and $K_N = \kappa_N(\vec{q})$, but more importantly, these estimates cannot be held constant in the neighborhood of \vec{q} , however small; they must covary for the pattern to be feasible (Ben-Shahar and Zucker, 2003b). Nevertheless, invariances between the curvatures do exist and formal considerations of good continuation have been shown to yield a unique approximation which, in $\mathbb{R}^2 \times S^1$, takes the form of a *right helicoid* (Fig. 7c-d) and whose orientation function has the following expression

$$\theta(x, y) = \tan^{-1}\left(\frac{K_T x + K_N y}{1 + K_N x - K_T y}\right). \quad (6)$$

The most unique property of this object is that it induces an identical covariation of the two curvature functions κ_T and κ_N in the neighborhood of the point \vec{q} . The osculating helicoid is the formal 2D analogue of the osculating circle and, as with co-circularity for curves, the fields of connections between neurons that this model generates (Fig. 8f-j)

depend intrinsically on curvature(s). Such connectivity structures can be used to compute coherent texture and shading flows in a “neural”, distributed fashion (Ben-Shahar and Zucker, 2003b) and two examples are shown in Fig. 9.

[Figure 8 about here]

[Figure 9 about here]

3 Results

The computational connection fields generated above contain all the geometrical information needed for predictions about long-range horizontal connections of individual cells (or after some averaging, that of tracer injection sites) in visual cortex. Thus we now turn to the central question: how well do these connectivity maps match the available data about projection fields in visual cortex? In particular, do they make better predictions than those arising from colinearity or association field models (Section 1)?

To answer these questions, we focused on anatomical studies that report population statistics (Malach *et al.*, 1993; Bosking *et al.*, 1997) and compared their data to predictions produced by performing “computational anatomy” on our model³. We randomly sampled the population of model-generated fields analogously to the way anatomists sample cells, or injection sites, in neural tissue and computed both individual and population statistics of their connection distributions. To generate robust predictions we repeated these sampling procedures many times and calculated the expected values and standard errors of the frequency distribution.

³Anatomical studies such as Bosking *et al.* (1997) and Malach *et al.* (1993) were preferred to psychophysical or electrophysiological studies because the latter two typically contribute no population statistics and are generally more difficult to interpret directly in terms of the structure of horizontal connections.

3.1 Computational anatomy predicts biological measurements

Fig. 10 illustrates the main results computed from our models, and compares them to the corresponding anatomical data reported in the literature (Malach *et al.*, 1993; Bosking *et al.*, 1997). The agreement of the computational process to the biological data is striking, both qualitatively and quantitatively. As with the association field, our model correctly predicts the spread of the pooled distribution with similar peak height (approximately 11% for orientation resolution of 10 degrees) and a similar orientation offset at which it crosses the uniform distribution (approximately ± 40 degrees). Unlike colinearity and association field models, however, ours predict qualitative differences between distributions of individual neurons, or injection sites, similar to findings in the literature (Fig 10c). Most importantly, our model predicts the consistently non-monotonic standard deviation. At orientation resolution of 10 degrees, both the anatomical data and the computational models exhibit variance local minima at approximately ± 30 degrees. This property holds both for a random sample of cells (Fig. 11), and for the computational population as a whole (not shown for space consideration).

[Figure 10 about here]

[Figure 11 about here]

3.2 Curvature quantization and population statistics

The geometrical model discussed in this paper must be quantized both in orientation and curvature before projection patterns can be computed and computational predictions can be made. We fixed the orientation quantization to the same level used in Bosking *et al.* (1997). Curvature quantization, however, is not addressed in the physiological literature and thus it is necessary to examine its effect on the resultant connectivity distributions. We note that even with orientation represented to hyper-acuity levels there are sufficient number of cells to represent such quantization (Miller and Zucker, 1999).

Broad orientation tuning implies discrete orientation quantization and suggests even more discrete curvature quantization. The results presented in Figs 10 and 11 are based on quantizing curvature into 5 classes⁴. This is a likely upper bound, given the broad band-pass tuning of cortical neurons that have been observed (Dobbins *et al.*, 1987; Versavel *et al.*, 1990) and modeled (Dobbins *et al.*, 1989). However, to study the effect of curvature quantization, we repeated the entire set of computations with both a smaller (3) and a larger (7) number of curvature classes. 3 is clearly the lower limit, which may correspond to the tree shrew (Bosking *et al.*, 1997) or other simple mammals, and 7 is more than required computationally (Ben-Shahar and Zucker, 2003b). We found that all of the properties predicted initially remain invariant under these changes. In particular, regardless of quantization level, the pooled distribution remains unimodal, it peaks at zero orientation difference with approximately 11%, it crosses the uniform distribution at $\pm 40^\circ$, and it has non-monotonic variance with local minima at $\pm 30^\circ$ (with somewhat increased variance around zero orientation for higher quantization levels). Qualitative differences between individual neurons are predicted regardless of the number of curvature classes. All these results are illustrated in Fig. 12.

[Figure 12 about here]

3.3 Relationship between cells' distribution and connections' distribution

Since both anatomical and computational studies must sample the population of (biological or computational) cells to measure the distribution of their horizontal connections, an important consideration is whether the underlying distribution of *cells* (based on their curvature tuning) can affect the *pooled* distribution of connections. For example, if most

⁴In the context of curves these 5 classes may be labeled as straight, slowly curving to the left, slowly curving to the right, rapidly curving to the left, and rapidly curving to the right.

cells in the population are tuned for zero or very small curvature, the pooled connection distribution may differ from that of a population dominated by high curvature cells.

The results presented in Figs 10 and 11 are based on the assumption that cells of different curvature tuning (or put differently, of different connectivity patterns) are distributed uniformly. Such an assumption follows naturally from the mathematical abstraction that allocates the same number of computational units to equal portions of $\mathbb{R}^2 \times S^1$. However, if this assumption were wrong, would a bias in the distribution of cells affect significantly the predictions made from our models?

Unfortunately, few data about such distributions are available, partially because anatomists need not assume any particular cells' distribution for their measurements of projection fields, and partially because curvature tuning is rarely considered. Some data available on the distribution of end-stopped cells (Kato *et al.*, 1978; Orban, 1984), in conjunction with the functional equivalence of end-stopping with curvature selectivity (Orban *et al.*, 1987; Dobbins *et al.*, 1987, 1989; Versavel *et al.*, 1990), suggest that cells are distributed *bimodally* in the curvature domain, with peaks at both zero and high curvature tuning. Alternatively, statistical studies of edge correlations in natural images (Dimitriv and Cowan, 1998; August and Zucker, 2000; Sigman *et al.*, 2001; Kaschube *et al.*, 2001; Geisler *et al.*, 2001; Pedersen and Lee, 2002) show that colinear co-occurrences are more frequent than others. Although these co-occurrence measurements neither depend on curvature nor do they necessarily indicate any particular distribution of cells at the computational level, implicitly they may suggest that cells are distributed *unimodally* in the curvature domain, with peak at zero curvature only.

Since our model raises the possibility of a curvature bias effect, we thus redistributed the population of our computational cells by one or the other of these non uniform (bimodal and unimodal) distributions, and then repeated the computational anatomy process described in Section 3.1. All computations were done on the more general 2D (texture) model. The bimodal distribution was modeled as a radial two-Gaussian mixture model (GMM) parametrized by the total curvature $\kappa = \sqrt{\kappa_T^2 + \kappa_N^2}$ and parameters

$\mu_0 = 0$, σ_0 , μ_1 , and σ_1 . The unimodal distribution was modeled as a two dimensional Gaussian of zero mean and variances σ_T and σ_N in the κ_T and κ_N dimensions, respectively.

Fig 13 illustrates one example of the resultant statistical measures for the *bimodal* cell distribution. In this example $\sigma_0 = 0.04$ and $\sigma_1 = 0.05$, where the slight difference accounts for corresponding differences in the two modes as reported in Kato *et al.* (1978) and Orban (1984). As is shown, this non-uniform distribution hardly changes the expected median, while further emphasizing the non-monotonic nature of the variance (compared to the statistics obtained with the same number of curvature classes and uniform cell distribution; Fig. 12d). Similar results were obtained with other values of σ_0 and σ_1 , and for the curvature quantized to 5 classes as well⁵.

[Figure 13 about here]

Fig 14 illustrates another example, this time using the *unimodal* cell distribution mentioned above. In this example $\sigma_T = \sigma_N = 0.15$ such that cells with zero curvature tuning are 8 times more frequent than cells tuned to the maximum value of curvatures. As expected, this strongly non-uniform distribution slightly elevated the peak of the population statistics, but otherwise, all other features that were predicted from the uniform cell distribution, and in particular the non-monotonic variance, were fully preserved. Similar results were obtained with other values of σ_T and σ_N , and for all quantization levels of the curvatures (as in Section 3.2).

[Figure 14 about here]

In summary, we have shown that even if cells in primary visual cortex were distributed non-uniformly in their curvature tuning, the pooled distribution of long-range

⁵Quantization of curvature to 3 classes was irrelevant in this case because the bimodality of the distribution could not be expressed using too few samples.

horizontal contentions in the orientation domain would preserve its fundamental properties, and in particular its wide spread and non-monotonic variance. Thus, our conclusions are not biased by an implicit assumption about curvature-dependent distribution of cells.

4 Discussion

The findings presented from our computational anatomy support the functional identification of the long-range horizontal connections with those obtained mathematically. However, the question of why the texture model is necessary becomes unavoidable, and we believe this issue is more than just formal mathematics. Certain physiological and psychophysical findings, such as iso-orientation side facilitation (Adini *et al.*, 1997), functional and anatomical connections between retinotopically *parallel* receptive fields (Ts'o *et al.*, 1986; Gilbert and Wiesel, 1989), and roughly isotropic retinotopic extent of projection fields (Malach *et al.*, 1993; Sincich and Blasdel, 2001), suggest the perceptual integration of texture flows rather than curves. Although this class of patterns may seem less important than curves as a factor in perceptual organization, their perceptual significance has been established (Glass, 1969; Kanizsa, 1979; Todd and Reichel, 1990). Furthermore, recent computational vision research implicates them in the analysis of visual shading (Lehky and Sejnowski, 1988; Huggins *et al.*, 2001), as was demonstrated in Fig. 9), and even color (Ben-Shahar and Zucker, 2003a).

Whether or not projection patterns of cells in primary visual cortex come in different flavors (i.e., curve vs. texture/shading integration) is an open question. To answer it one is likely to exploit the many physiologically measurable differences between these classes of projection patterns as suggested by Fig. 8. Unfortunately, the statistical data obtained so far do not distinguish between the two (curve and texture) integration models; without a spatial dimension, the statistical differences between the two models in the orientation domain are too fine to measure relative to the accuracy of current laboratory

techniques.

Until full spatio-angular data are obtained, however, the inclusion of even *weak* spatial information is sufficient to generate further testable predictions. In particular, incorporating the retinotopic *distance* between linked cells into the statistics (or estimating it from their cortical distance) can produce predictions regarding the dependency of the distribution's spread and shape on the integration distance, as illustrated in Fig. 15. Some verification for these predictions can be seen in the measurements of Kisvárday *et al.* (Kisvárday *et al.*, 1997) (top row in their Fig. 9 shows developing peaks resembling the ones in Fig 15b,c). Similar annular analyses, which focus on sectors of the annuli in directions other than parallel to the RF's preferred orientation, provide measurable differences between curve and texture projection fields.

[Figure 15 about here]

In summary, we have presented mathematical analysis and computational models that predict both the pooled distribution of long-range horizontal connections and the distributions of individual cells/injection sites. For the first time the modeling goes beyond the unimodal first-order data and falsifies a common conclusion from it. In particular, while co-aligned facilitation entails the pooled first-order data, the converse is not necessarily true: these data are also consistent with curvature-dependent connections. The second-order (variance) data, however, remain consistent with curvature-dependent connections but not with co-aligned facilitation.

The explanatory force of our model derives from sensory integration, and we observed in the Introduction that most researchers limit this to curve integration via co-linearity. We, however, conclude in an enlarged context: differential geometry provides a foundation for connections in visual cortex that predicts both dependency on curvature(s) and an expanded functional capability, including curve, texture, and shading integration. Since the same geometrical analysis applies to many other domains in which orientation and direction fields are fundamental features, coherent motion pro-

cessing (Series *et al.*, 2002) and coherent color perception (Ben-Shahar and Zucker, 2003a) might also have been included. Since all follow from the geometry, and all are important for vision, more targeted experiments are required to articulate their neural realization.

Acknowledgments

We are grateful to Allan Dobbins, David Fitzpatrick, Kathleen Rockland, Terry Sejnowski, and Michael Stryker, for reviewing this manuscript and for providing valuable comments, to Lee Iverson for the curve compatibility fields in Fig. 8, and to an anonymous reviewer for pointing out a possible artifact in the data (Sec. 1.1). This research was supported by AFOSR, ONR, and DARPA.

A Noise models

Two basic noise models are used in this paper to examine whether variations of the basic colinear distribution (Fig. 2) can produce a pooled statistics with similar properties to the biological one. This appendix describes both procedures in detail.

To examine natural random variations at the level of individual connections, each long-range horizontal connection, ideally designated to connect cells of orientation difference $\Delta\theta$, is shifted to connect cells of orientation difference $\Delta\theta + \epsilon_\sigma$, where ϵ_σ is a wrapped Gaussian noise with zero mean and variance σ . To do this computation, the base distribution (colinear or association field) from Fig. 2, initially given as probabilities over 18 orientation bins of 10° each, was normalized and quantized to a connection histogram in the range $[0, N]$, where N represents the total number of connections a cell makes. To each such connection to orientation difference $\Delta\theta$ we then added a wrapped Gaussian noise ϵ_σ of zero mean and variance σ , and the new (random) connection was accumulated at the bin $\Delta\theta + \epsilon_\sigma$ of the resultant histogram. This process was repeated

200 times to produce 200 different perturbations, from which both expected distribution and variance were computed bin-wise. The parameter σ was set to the value that produced an expected distribution of peak height and spread similar to the biological one. Since different anatomical studies, and protocols, indicate a different number of total connections (e.g., hundreds in Schmidt *et al.* (1997), approximately 3500 in Buzás *et al.* (1998), and up to 20K for injection sites of approximately 20 cells in Bosking *et al.* (1997)), we repeated this statistical test for normalizations in different ranges. As expected, changing N only scaled the variance uniformly across the expected distribution but did not affect its mean. Thus, for better clarity of its monotonicity, the result in Fig. 3b reflects a smaller number of total connections ($N = 200$), as in, e.g., Schmidt *et al.* (1997).

To examine random variations due to “leakage” of tracer from an injection site of preferred orientation θ_0 to nearby orientation columns, we modeled such leakage by selecting $i = 1 \dots M$ cells of preferred orientation $\theta_i = \theta_0 + \Delta\theta_\sigma$, where $\Delta\theta_\sigma$ is a wrapped Gaussian random variable of zero mean and variance σ . A normalized and quantized based distribution (colinear or association field) was then centered around each of θ_i and all were summed up and normalized to yield a resultant (random) distribution of connections for the injection site at θ_0 . As before, we repeated this generation process 200 times to produced 200 different perturbations, from which both expected distribution and variance were computed bin-wise. The parameter σ was again set to that value that produced an expected distribution of peak height and spread similar to the biological one. The number of cells in an injection was set to $M = 20$, approximately as reported in Bosking *et al.* (1997). Unlike random variations at the level of individual connections, the range parameter N had no effect on the variance of the expected distribution.

References


- Adini, Y., Sagi, D., and Tsodyks, M. (1997). Excitatory-inhibitory network in the visual cortex: Psychophysical evidence. *Proc. Natl. Acad. Sci. U.S.A.*, **94**, 10426–10431.
- August, J. and Zucker, S. (2000). The curve indicator random field: Curve organization via edge correlation. In K. Boyer and S. Sarkar, editors, *Perceptual Organization for Artificial Vision Systems*. Kluwer Academic Publishers.
- Ben-Shahar, O. and Zucker, S. (2003a). Hue fields and color curvatures: A perceptual organization approach to color image denoising. In *Proc. Computer Vision and Pattern Recognition*, pages 713–720.
- Ben-Shahar, O. and Zucker, S. (2003b). The perceptual organization of texture flows: A contextual inference approach. *IEEE Trans. Pattern Anal. Machine Intell.*, **25**(4), 401–417.
- Bosking, W., Zhang, Y., B., S., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in the tree shrew striate cortex. *J. Neurosci.*, **17**(6), 2112–2127.
- Breton, P. and Zucker, S. (1996). Shadows and shading flow fields. In *Proc. Computer Vision and Pattern Recognition*, pages 782–789.
- Buzás, P., Eysel, U., and Kisvárdy, Z. (1998). Functional topography of single cortical cells: an intracellular approach combined with optical imaging. *Brain Res. Prot.*, **3**, 199–208.
- Dimitriv, A. and Cowan, J. (1998). Spatial decorrelation in orientation-selective cortical cells. *Neural Comput.*, **10**, 1779–1795.
- do Carmo, M. (1976). *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Inc.

- Dobbins, A., Zucker, S., and Cynader, M. (1987). Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature*, **329**(6138), 438–441.
- Dobbins, A., Zucker, S., and Cynader, M. (1989). Endstopping and curvature. *Vision Res.*, **29**(10), 1371–1387.
- Field, D., Hayes, A., and Hess, R. (1993). Contour integration by the human visual system: Evidence for a local “association field”. *Vision Res.*, **33**(2), 173–193.
- Geisler, W., Perry, J., Super, B., and Gallogly, D. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res.*, **41**(6), 711–724.
- Gilbert, C. and Wiesel, T. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J. Neurosci.*, **9**(7), 2432–2442.
- Glass, L. (1969). Moiré effect from random dots. *Nature*, **223**(5206), 578–580.
- Hubel, D. and Wiesel, T. (1977). Functional architecture of macaque monkey visual cortex. In *Proc. R. Soc. London Ser. B*, volume 198, pages 1–59.
- Huggins, P., Chen, H., Belhumeur, P., and Zucker, S. (2001). Finding folds: On the appearance and identification of occlusion. In *Proc. Computer Vision and Pattern Recognition*, pages 718–725.
- Kanizsa, G. (1979). *Organization in Vision: Essays on Gestalt Perception*. Praeger Publishers.
- Kapadia, M., Ito, M., Gilbert, C., and Westheimer, G. (1995). Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron*, **15**, 843–856.
- Kapadia, M., Westheimer, G., , and Gilbert, C. (2000). Spatial distribution of contextual interactions in primary visual cortex and in visual perception. *J. Neurophysiol.*, **84**, 2048–2062.

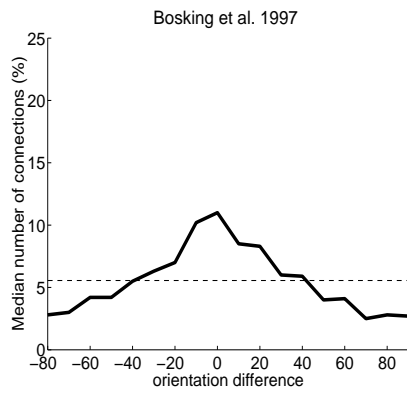
- Kaschube, M., Wolf, F., Geisel, T., and Löwel, S. (2001). The prevalence of colinear contours in the real world. *Neurocomputing*, **38-40**, 1335–1339.
- Kato, H., Bishop, P., and Orban, G. (1978). Hypercomplex and the simple/complex cell classification in cat striate cortex. *J. Neurophysiol.*, **41**, 1071–1095.
- Kisvárdy, Z., Kim, D., Eysel, U., and Bonhoeffer, T. (1994). Relationship between lateral inhibition connections and the topography of the orientation map in cat visual cortex. *J. Neurosci.*, **6**, 1619–1632.
- Kisvárdy, Z., Tóth, É., Rausch, M., and Eysel, U. (1997). Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat. *Cereb. Cortex*, **7**, 605–618.
- Lehky, S. and Sejnowski, T. (1988). Network model of shape-from-shading: neural function arises from both receptive and projective fields. *Nature*, **333**, 452–454.
- Malach, R., Amir, Y., Harel, M., and Grinvald, A. (1993). Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proc. Natl. Acad. Sci. U.S.A.*, **90**, 10469–10473.
- Matsubara, J., Cynader, M., Swindale, N., and Stryker, M. (1985). Intrinsic projections within visual cortex: Evidence for orientation specific local connections. *Proc. Natl. Acad. Sci. U.S.A.*, **82**, 935–939.
- Miller, D. and Zucker, S. (1999). Computing with self-excitatory cliques: A model and application to hyperacuity-scale computation in visual cortex. *Neural Comput.*, **11**, 21–66.
- Mitchison, G. and Crick, F. (1982). Long axons within the striate cortex: Their distribution, orientation, and patterns of connections. *Proc. Natl. Acad. Sci. U.S.A.*, **79**, 3661–3665.

- O'Neill, B. (1966). *Elementary Differential Geometry*. Academic Press.
- Orban, G. (1984). Receptive field organization in areas 17,18 and 19 of the cat. In *Neural Operations in the Visual Cortex*, chapter 4, pages 87–124. Springer-Verlag.
- Orban, G., Versavel, M., and Lagae, L. (1987). How do striate neurons represent curved stimuli. In *Abstracts of the Society for Neuroscience*, volume XIII, page 404.10.
- Parent, P. and Zucker, S. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Anal. Machine Intell.*, **11**(8), 823–839.
- Pedersen, K. and Lee, A. (2002). Toward a full probability model of edges in natural images. APPTS Technical Report 02-1, Division of Applied Mathematics, Brown University.
- Polat, U. and Sagi, D. (1993). Lateral interactions between spatial channels: Suppression and facilitation revealed by lateral masking experiments. *Vision Res.*, **33**(7), 993–999.
- Rockland, K. and Lund, J. (1982). Widespread periodic intrinsic connections in the tree shrew visual cortex. *Science*, **215**(4941), 1532–1534.
- Schmidt, K. and Löwel, S. (2002). Long-range intrinsic connections in cat primary visual cortex. In B. Payne and A. Peters, editors, *The Cat Primary Visual Cortex*, pages 387–426. Academic Press.
- Schmidt, K., Goebel, R., Löwel, S., and Singer, W. (1997). The perceptual grouping criterion of colinearity is reflected by anisotropies in the primary visual cortex. *Eur. J. Neurosci.*, **9**, 1083–1089.
- Series, P., Georges, S., Lorenceau, J., and Fregnac, Y. (2002). Orientation dependent modulation of apparent speed: a model based on the dynamics of feed-forwards and horizontal connectivity in v1 cortex. *Vision Res.*, **42**, 2781 – 2797.

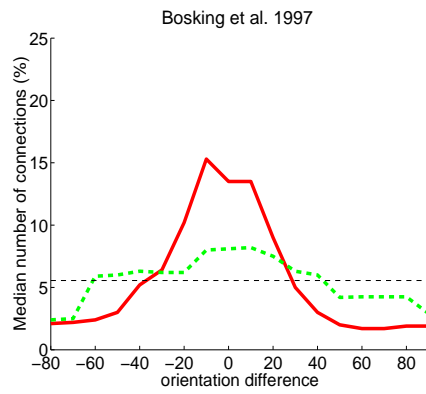
- Sigman, M., Cecchi, G., Gilbert, C., and Magnasco, M. (2001). On a common circle: Natural scenes and gestalt rules. *Proc. Natl. Acad. Sci. U.S.A.*, **98**(4), 1935–1940.
- Sincich, L. and Blasdel, G. (2001). Oriented axon projections in primary visual cortex of the monkey. *J. Neurosci.*, **21**(12), 4416–4426.
- Todd, J. and Reichel, F. (1990). Visual perception of smoothly curved surfaces from double-projected contour patterns. *J. Exp. Psych.: Human Perception and Performance*, **16**(3), 665–674.
- Ts'o, D., Gilbert, C., and Wiesel, T. (1986). Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.*, **6**(4), 1160–1170.
- Versavel, M., Orban, G., and Lagae, L. (1990). Responses of visual cortical neurons to curved stimuli and chevrons. *Vision Res.*, **30**(2), 235–248.
- Weliky, M., Kandler, K., Fitzpatrick, D., and Katz, L. (1995). Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to oriented columns. *Neuron*, **15**, 541–552.
- Wertheimer, M. (1955). Laws of organization in perceptual forms. In W. Ellis, editor, *A source book of Gestalt Psych.*, pages 71–88. Routledge & Kegan Paul Ltd.
- Zucker, S., Dobbins, A., and Iverson, L. (1989). Two stages of curve detection suggest two styles of visual computation. *Neural Comput.*, **1**, 68–81. See also, Iverson, L.A., PhD thesis, McGill University, 1994.



a



b



c

Figure 1: **Ben-Shahar/#2773.**

Caption of Figure 1: Visual integration and the distribution of long-range projections. **(a)** Broad tuning in orientation and position introduce uncertainty in curve integration even if a single curve model (thick, red curve) is assumed through the RF. Determining which nearby RF the curve continues through can be facilitated by interaction between neurons with mutually aligned, retinotopically close RFs. **(b)** A fundamental measurable property of long-range connection is their distribution in the orientation domain, i.e., the percentage of connections between interconnected neurons as a function of preferred orientation (angular) difference. This graph shows the median distribution of lateral connections (distance $> 500\mu\text{m}$) of 7 cell clusters in primary visual cortex of tree shrew (redrawn from Bosking *et al.*, 1997, their Fig. 6c). Qualitatively similar (though coarser) measurements are available on primates as well (Malach *et al.*, 1993). **(c)** Connectivity distribution of individual cell clusters reveals significant variability and qualitative differences between them. Shown here are distributions from two injection sites from Bosking *et al.* (1997).

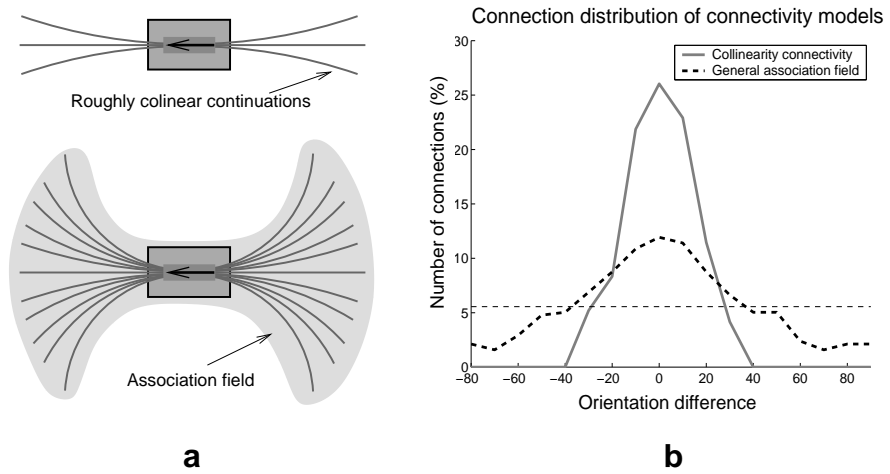




Figure 2: **Ben-Shahar/#2773.**

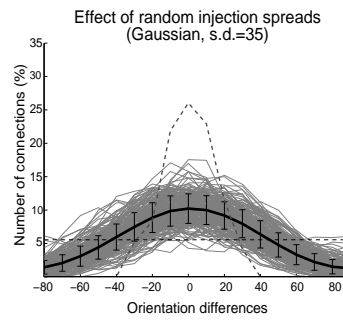
Caption of Figure 2: Colinear facilitation, association fields, and their predicted distribution of connections. **(a)** Informally, two visual integration, or continuation models are typically considered in the physiological and psychophysical literatures. Collinearity, the predominant model, predicts only few possible curve continuations (top). On the other hand, many possible continuations reveal an association field (bottom), similar to those observed psychophysically (Field *et al.*, 1993). **(b)** The corresponding distribution derived from the collinearity and association field models. Observe that collinearity predicts a very narrow distribution which is clearly at odds with the significant spread frequently measured anatomically or electrophysiologically (compare to Fig. 1b). The association field leads to a wider spread, but like collinearity, it predicts a *fixed* distribution for all cells, a hypothesis refuted in recent studies (see text). The collinearity distribution (solid) was calculated from the field depicted in Fig 8a while the association field distribution (dashed) was calculated from the field in Fig 8e. The dashed horizontal line depicts the uniform distribution.



a




b




c

Figure 3: **Ben-Shahar/#2773.**


Caption of Figure 3: Results of a statistical perturbation of colinear connectivity distribution. **(a)** Mean connection distribution computed from the data in Bosking *et al.* (Bosking *et al.*, 1997), shown here for reference. Error bars are ± 1 standard deviation. Note the unimodal distribution that peaks at approximately 11%, the wide spread, the crossing of the uniform distribution (dashed horizontal line) around $\pm 40^\circ$, and the non-monotonic variance. Can all these features be replicated by applying noise to the base distribution induced by the standard collinearity model? **(b)** Result of simulating physiological deviation at the individual connection level. Dashed line is the base colinear distribution. Gray region is the superposition of individual applications of the noise model to the base distribution. Solid graph is the expected distribution and error bars are ± 1 standard deviation. Permitting large enough developmental variations (shown here is the result of wrapped Gaussian i.i.d. noise of s.d. = 35°) in the connections to model the first-order statistics significantly violates the underlying connectivity principle of good continuation but still cannot model the second-order statistics. **(c)** Results of simulating measurement errors due to leakage in the injection site. All parts are coded as in panel b. Again, permitting large enough injection spread to model the first-order statistics (shown here is the result of Gaussian noise of s.d. = 35° and assuming 20 cells per injection site (Bosking *et al.*, 1997)) cannot model the second-order statistics.




a



b



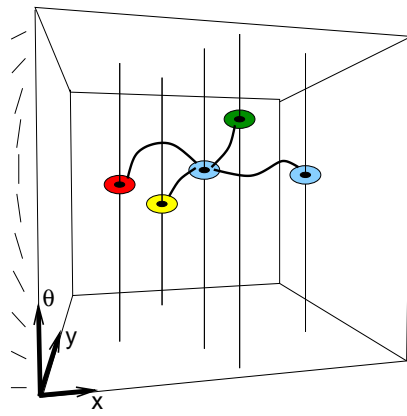
c




d

Figure 4: **Ben-Shahar/#2773.**


Caption of Figure 4: Abstracting the primary visual cortex as $\mathbb{R}^2 \times \mathcal{S}^1$, or *position* \times *orientation* space. **(a)** The “ice cube” cartoon of visual cortex (Hubel and Wiesel, 1977) (cytochrome-oxidase blobs and distortions due to cortical magnification factor are not shown). A tangential penetration in the superficial layers reveals an orientation hypercolumn of cells whose RFs have similar spatial (retinotopic) coordinates. With cells of similar orientation tuning grouped by color, the hypercolumn is cartooned as a horizontal cylinder. **(b)** With ocular dominance columns omitted, the superficial layers of the primary visual cortex can be viewed as a collection of (horizontally arranged) orientation hypercolumns. **(c)** Drawing the cylinders vertically emphasizes that RFs of cells within a column overlap in retinotopic coordinates (x, y) and makes explicit this aspect of their organization. **(d)** Since different hypercolumns correspond to different retinotopic positions, the set of all hypercolumns abstracts the visible subspace of $\mathbb{R}^2 \times \mathcal{S}^1$, with each column corresponding to a different vertical fiber in that space. The θ axis in this space corresponds to a tangential penetration within V1 hypercolumns (colors within the column represent different orientation tunings), and the XY plane corresponds to retinotopic coordinates.




a



b



c



d

Figure 5: **Ben-Shahar/#2773.**

Caption of Figure 5: Abstracting long-range horizontal connections as relationships between points in $\mathbb{R}^2 \times \mathcal{S}^1$. **(a)** Since visual integration must involve not only the relative orientation between RFs, but their spatial offset as well, it is more fully abstracted by relationships between points in $\mathbb{R}^2 \times \mathcal{S}^1$. The exact nature of these relationships is determined by the underlying integration model. **(b)** Redrawing $\mathbb{R}^2 \times \mathcal{S}^1$ fibers as orientation hypercolumns in V1 reveals the connection between the integration model in $\mathbb{R}^2 \times \mathcal{S}^1$ and the distribution of long-range horizontal connections between the hypercolumns. **(c)** Collapsing the $\mathbb{R}^2 \times \mathcal{S}^1$ abstraction to a cortical orientation map (i.e., flattening each orientation cylinder and re-distributing its orientation-selective parts as orientation columns in the superficial cortical layers), the integration model implies a particular set of long-range horizontal connections between orientation domains (colors represent orientation tuning similar to panels **a,b** and Fig. 4). Such links have been identified and measured through optical imaging and anatomical tracing (e.g., (Malach *et al.*, 1993; Bosking *et al.*, 1997; Buzás *et al.*, 1998)) and thus can be compared to the model's predictions. **(d)** A real counterpart to the schematic in panel **c**. Reproduced from Bosking *et al.* (1997), this image shows an optical image of intrinsic signals combined with long-range horizontal connections traced through extracellular injection of biocytin. The white dots at the upper left corner represent the injection site while the black dots represent labeled boutons. The white bar in the inset represents the orientation preference at the injection site.

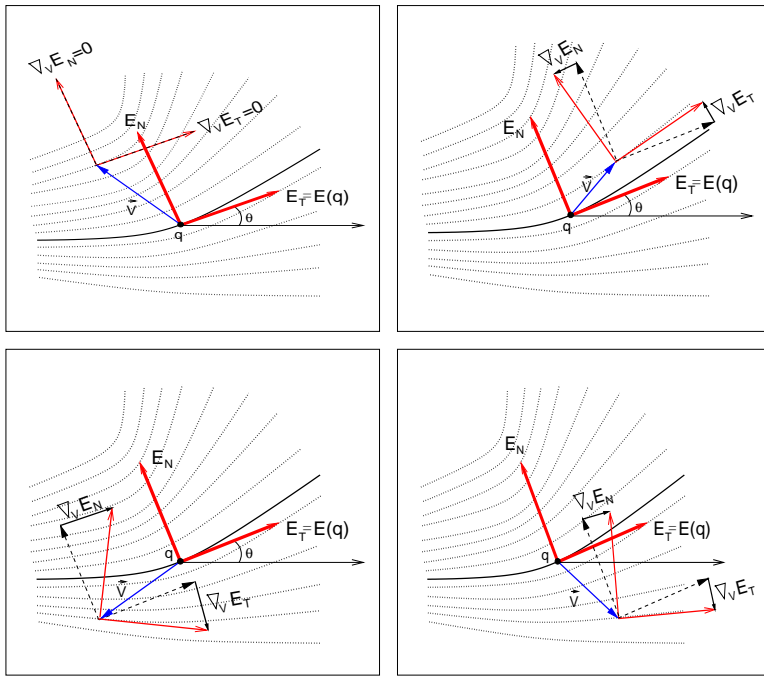




Figure 6: **Ben-Shahar/#2773.**


Caption of Figure 6: Visual integration under good continuation involves the question of how a measurement of orientation at one retinal position relates to another measurement of orientation at a nearby retinal position. Formally this amounts to specifying how a tangent (orientation measurement) at position \vec{q} relates to another nearby tangent displaced by a vector \vec{V} . This tangent displacement amounts to rotation, and as shown above this rotation can differ for different displacements. Formally, the rotation is specified locally by the covariant derivative $\nabla_{\vec{V}}$, and the mathematical analysis is facilitated by defining an appropriate coordinate frame. Shown is the Frenet basis $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$, where $\hat{\mathbf{E}}_T$ corresponds to a unit vector in the orientation's tangential direction and $\hat{\mathbf{E}}_N$ corresponds to a unit vector in the normal direction. Associated with this frame is an angle θ defined relative to external fixed coordinate frame (black horizontal line). The covariant derivative specifies the frame's *initial* rate of rotation for any direction vector \vec{V} . The four different cases in this figure illustrate how this rotation depends on \vec{V} both quantitatively (i.e., different magnitudes of rotation) and qualitatively (i.e., clockwise, counter-clockwise, or zero rotation). Since displacement is a two-dimensional vector and $\nabla_{\vec{V}}$ is linear, two numbers are required to fully specify the covariant derivative. These two numbers describe the initial rate of rotation in two independent displacement directions. Using the Frenet basis once again, two natural directions emerge. A pure displacement in the tangential direction ($\hat{\mathbf{E}}_T$) specifies one rotation component, and a pure displacement in the normal direction ($\hat{\mathbf{E}}_N$) specifies the other component. We call them the *tangential* curvature (κ_T) and the *normal* curvature (κ_N), respectively. If visual integration based on good continuation relates to two dimensional patterns of orientation, then both of these curvatures are required. For good continuation along individual curves only the tangential curvature is required since displacement is possible only in the tangential direction (that is, along the curve only).




a




b




c



d



e



f

Figure 7: Ben-Shahar/#2773.

Caption of Figure 7: Differential geometry, integration models, and horizontal connections between compatible RFs. **(a)** Estimate of tangent (light blue vector) and curvature at a point \vec{q} permits modeling a curve with the osculating circle as a good-continuation approximation in its neighborhood. Given the approximation, compatible (green) and incompatible (pink) tangents at nearby locations can be explicitly derived. **(b)** With height representing orientation (see scale along the θ axis), the osculating circle lifts to a helix in $\mathbb{R}^2 \times S^1$ whose points define both the spatial location and orientation of compatible nearby tangents. Color-coded as in panel **a**, the green point is compatible with the blue one while the pink points are incompatible with it. **(c)** The consistent structure in panels **a,b** illustrated as RFs and their spatial arrangement. As an abstraction for visual integration, the ideal geometrical model – the osculating circle – induces a discrete set of RFs which can facilitate the response of the central cell. Shown here is an example for one particular curvature tuning at the central cell. **(d)** For textures, determination of good continuation requires two curvatures at a point. Based on these curvatures, a local model of good continuation can determine both the position, orientation, and curvatures of (spatially) nearby coherent points. Given these two curvatures at a point, there exists a unique model of good continuation that guarantees identical covariation of the curvature functions. Given the approximation, compatible (green) and incompatible (pink) flow patches at nearby locations can be explicitly derived. **(e)** In $\mathbb{R}^2 \times S^1$, our model for 2D orientation good continuation lifts to a right helicoid, whose points define both the spatial location and orientation of compatible (green) nearby flow tangents. **(f)** As an abstraction for visual integration, the ideal geometrical model – the right helicoid – induces a discrete set of RFs which can facilitate the response of the central cell. Shown here is an example for one particular curvature tuning at the central cell. Note that broad RF tuning means that both the helix and the helicoid must be dilated appropriately, thus resulting in compatible “volumes” in $\mathbb{R}^2 \times S^1$ and possibly multiple compatible orientations at given spatial positions. This dilation should be reflected in the set of compatible RFs and the horizontal links to them, but to avoid clutter, we omit it from this figure. The effect of this dilation is illustrated in Fig. 8 and consequently in all our calculations.

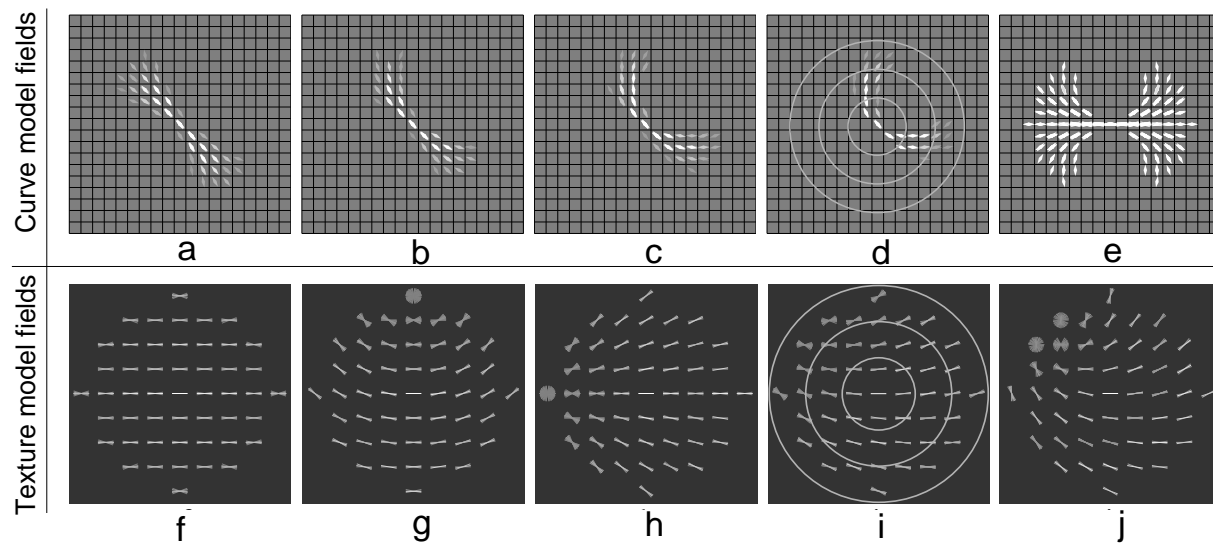


Figure 8: **Ben-Shahar/#2773**.

Caption of Figure 8: Illustration of connection fields for curves (top, based on co-circularity (Parent and Zucker, 1989)) and textures (bottom, based on right helicoidal model (Ben-Shahar and Zucker, 2003b)). Each position in these fields represents one orientation hypercolumn while individual bars represent the orientation preference of single neurons, all of which are connected to the central cell in each field. Multiple bars at any given point represent multiple neurons in the same hypercolumn that are connected to the central cell, a result of the dilation of the compatible structure due to broad RF tuning (see caption of Fig. 7). All fields assume that orientation tuning is quantized to 10 degrees and their radius of influence is set to 4-5 non-overlapping hypercolumns to reflect a 6-8mm cortical range of horizontal connections (Gilbert and Wiesel, 1989) and hypercolumn diameter of 1.5mm (to account for ocular dominance domains). **(a-d)** Examples of co-circularity projection fields (Parent and Zucker, 1989) for cells with orientation preference of 150 degrees (center bars) and different values of curvature tuning (based on the implementation by Iverson (Zucker *et al.*, 1989)). **(a)** $\kappa = 0.0$ (curvature in units of pixels^{-1}). **(b)** $\kappa = 0.08$. **(c)** $\kappa = 0.16$. **(d)** $\kappa = 0.24$. **(e)** The union of all projection fields of all cells with same orientation preference (0 degrees in this case) but different curvature tuning. Note the similarity to the schematic association field in Fig. 6b. **(f-j)** Examples of the texture flow projection fields (Ben-Shahar and Zucker, 2003b) for cells with horizontal orientation preference (center bars) and different curvature tuning. Note the intrinsic dependency on curvatures and the qualitatively different connectivity patterns that they induce. **(f)** $(\kappa_T, \kappa_N) = (0.0, 0.0)$. **(g)** $(\kappa_T, \kappa_N) = (0.2, 0.0)$. **(h)** $(\kappa_T, \kappa_N) = (0.0, 0.2)$. **(i)** $(\kappa_T, \kappa_N) = (0.1, 0.1)$. **(j)** $(\kappa_T, \kappa_N) = (0.2, 0.2)$. Note that while the majority of connections link roughly co-linear cells, some connect cells of large orientation differences. The fields shown are just a few examples sampled from the models, both of which contain similar (rotated) connection fields for each of the possible orientation preferences in the central hypercolumn. The circles superimposed on panels **d,i** are used to characterize retinotopic distance zones for the predictions made in Fig. 15.




Figure 9: **Ben-Shahar/#2773.**

Caption of Figure 9: Example of coherent texture (a-d) and shading (e-g) flow computation based on contextual facilitation with right helicoidal connectivity patterns (Ben-Shahar and Zucker, 2003b). **(a)** Natural image of a tree stump with perceptual texture flow. **(b)** A manually drawn flow structure as perceived by a typical observer. **(c)** Noisy orientation field reminiscent of RF responses. The computed measurements are based on the direction of the image intensity gradient. **(d)** The outcome of applying a contextual and distributed computation (Ben-Shahar and Zucker, 2003b) which facilitates the response of individual cells based on their interaction with nearby cells through the connectivity structures in Fig. 8. Compare to panel **b** and note how the measurements in the area of the knot, where no RF is embedded in a coherent context, were rejected altogether. **(e)** An image of a plane. **(f)** Measured shading flow field (white) and edges (black). In biological terms edges are measured by RFs of particular orientation preferences tuned to high spatial frequencies. The shading field may be measured by cells tuned to low frequencies. **(g)** Applying the right helicoidal-based computation on the shading information results in a coherent shading field on the plane's nose and a complete rejection of the incoherent shading information on the textured background. Such an outcome can be used to segment smoothly curved surfaces in the scene (Ben-Shahar and Zucker, 2003b), to resolve their shape (Lehky and Sejnowski, 1988), to identify shadows (Breton and Zucker, 1996), and to determine occlusion relationship underlying edge classification (Huggins *et al.*, 2001).

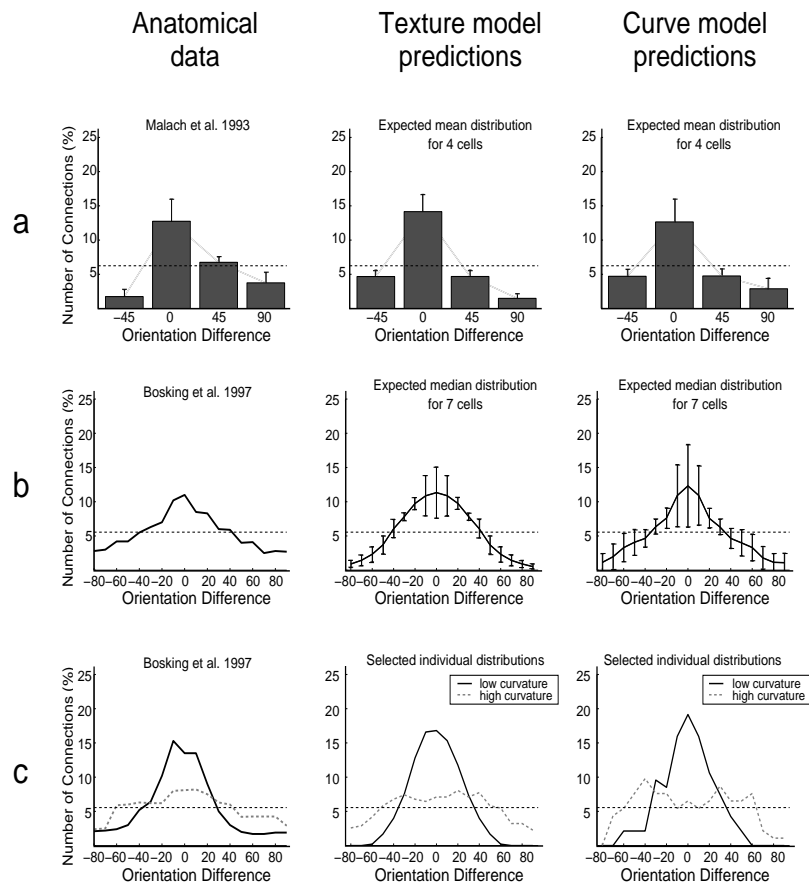




Figure 10: Ben-Shahar/#2773.

Caption of Figure 10: Comparison of anatomical data and model predictions for the distribution of long-range horizontal connections in the orientation domain. In all graphs, dashed horizontal lines represent the uniform distribution and error bars represent ± 1 standard error. **(a)** Mean connection distribution of 4 injection sites from Malach *et al.* (1993) vs. the computational prediction from our models (expected mean, $N=4$, 100 repetitions). Note the dominant peak around zero orientation difference and the considerable width of the histogram. The asymmetry in the pooled distribution measured by Malach *et al.* (1993) likely derives from a bias at the injection site (see their Fig. 4D) rather than being intrinsic. **(b)** Median distribution of 7 injection sites from Bosking *et al.* (1997) against the computational prediction from our models (expected median, $N=7$, 100 repetitions). Note in particular the similarity in peaks' height and in the orientation offset at which the graphs cross the uniform distribution, and the strongly non-monotonic behavior of the variance. **(c)** Two individual injection sites with qualitatively different connection distributions reproduced from Bosking *et al.* (1997). The counterpart computational instances are sampled from our models. Solid graphs correspond to the fields in Fig. 8b,i. Dashed graphs correspond to the fields in Fig. 8c,j.

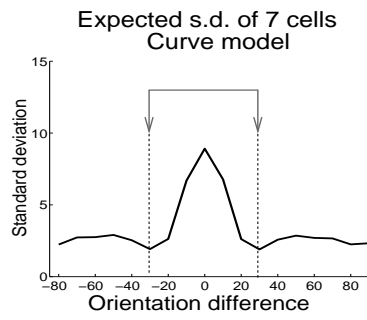
Mean and s.d. of 7 injection sites
Bosking et al. 1997



a



b



c

Figure 11: Ben-Shahar/#2773.

Caption of Figure 11: Although both the computational and the physiologically measured distributions of the mean are monotonically decreasing, their standard deviation is consistently non-monotonic. **(a)** While Bosking *et al.* (1997) used the population *median*, we further analyzed their published data (from 7 injection sites) to find its *mean* and *standard deviation*. It is evident that the s.d. is non-monotonic, with two local minima at ± 30 degrees (marked with red arrows). **(b)** Expected standard deviation for the texture model. **(c)** Expected standard deviation for the curve model. Both graphs depict the expected s.d. for 7 randomly selected cells ($N=7$, 100 repetitions). and both show a similar non-monotonic behavior with pronounced s.d. local minima at approximately ± 30 degrees. Note how the computational local minima coincide with the anatomical ones (arrows are copied from panel **a** and overlaid on the computational graphs). Compare also to the s.d. on the median graphs in Fig. 10b. Note that as with the distributions themselves, both computational models produce quantitatively similar s.d. results.

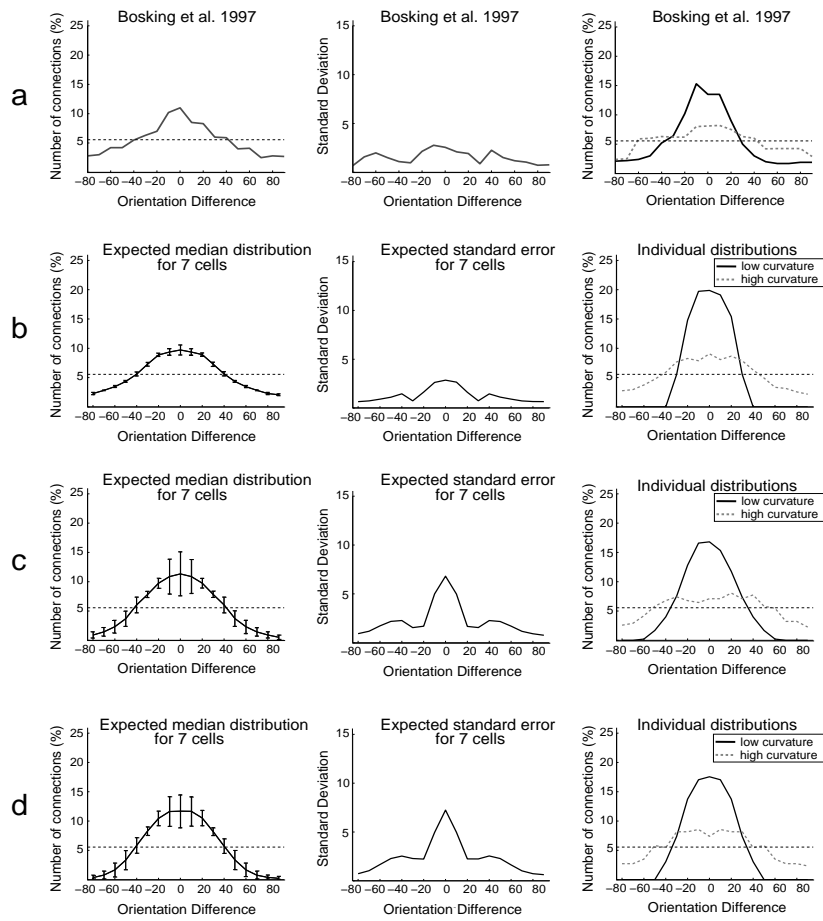
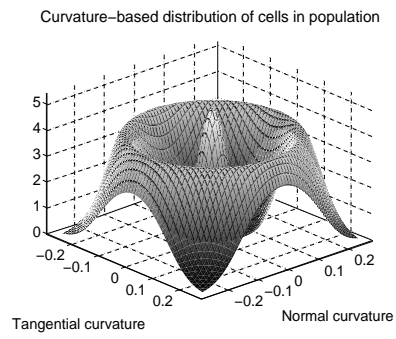
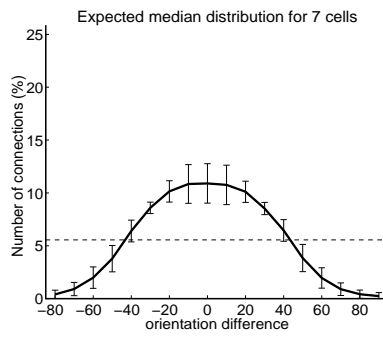


Figure 12: Ben-Shahar/#2773.

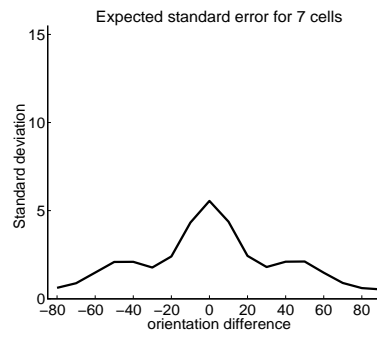
Caption of Figure 12: Different quantization levels of curvature tuning have little effect on the expected median distribution and its standard deviation. **(a)** Anatomical data from Bosking *et al.* (1997) shown for comparison with the computational predictions. **(b)** Computational predictions with 3 curvature classes. **(c)** Computational predictions with 5 curvature classes. **(d)** Computational predictions with 7 curvature classes. In all cases, left column depicts the expected median for 7 cells (bars are 1 s.d.), middle column depicts the expected standard deviation for 7 cells, and the right column shows two qualitatively different distributions from two different cells. For space considerations we show the results from the texture model only.



a



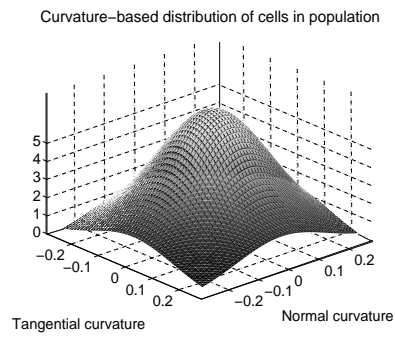
b




c

Figure 13: Ben-Shahar/#2773.


Caption of Figure 13: A statistical confirmation that the properties of our models persist even when the population of cells is distributed bimodally (Kato *et al.*, 1978; Orban, 1984). Illustrated here is the result from a distribution modeled as a GMM with $\mu_0 = 0$, $\sigma_0 = 0.04$, $\mu_1 = 0.2$, and $\sigma_1 = 0.05$. Since the bimodal nature of the distribution is best represented with higher number of curvature classes, presented here is the case of the texture model with curvatures quantized to 7 classes each. **(a)** The (radially) bimodal distribution of cells in the curvature domain normalized for number of cells. X and Y axes represent tangential and normal curvature tuning, respectively, and the Z axis represents the number of such curvature-tuned cells for any given orientation tuning. **(b)** Expected median of 7 cells. Error bars are ± 1 standard deviation. **(c)** Expected standard error for 7 cells. Compare both graphs to Fig. 12d.



a



b



c

Figure 14: Ben-Shahar/#2773.

Caption of Figure 14: A statistical confirmation that the properties of our models persist even when the population of cells is distributed normally (i.e., *unimodally*). Such distribution is implicitly suggested by statistics of edges in natural images (Dimitriv and Cowan, 1998; August and Zucker, 2000; Sigman *et al.*, 2001; Kaschube *et al.*, 2001; Geisler *et al.*, 2001; Pedersen and Lee, 2002) in which colinear edges are much more frequent. The case presented here ($\sigma_T = \sigma_N = 0.15$) induces a distribution in which cells of zero curvature tuning are 8 times more frequent than those of maximal curvature tuning. The graphs in this figure correspond to the texture model with curvatures quantized to 3 classes each. Similar results were obtained with other quantization levels as well. **(a)** The normal distribution of cells in the curvature domain normalized for number of cells. X and Y axes represent tangential and normal curvature tuning, respectively, and the Z axis represents the number of such curvature-tuned cells for any given orientation tuning. **(b)** Expected median of 7 cells. Error bars are ± 1 standard deviation. **(c)** Expected standard error for 7 cells. Compare both graphs to Fig. 12b.

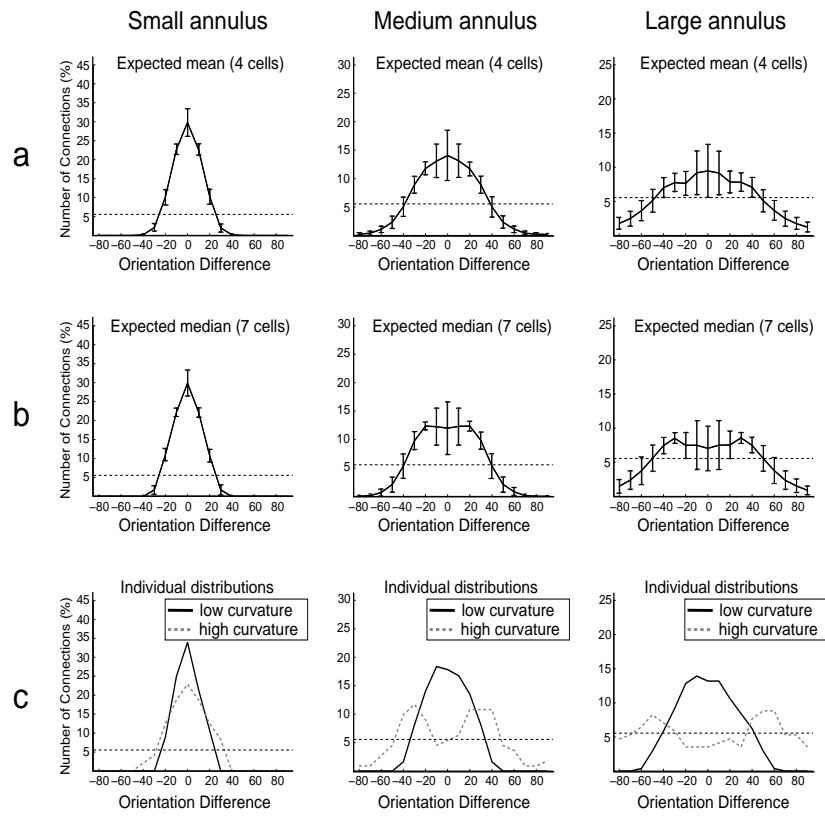


Figure 15: Ben-Shahar/#2773.

Caption of Figure 15: Model predictions of connection distributions by retinotopic annulus. Left, middle, and right columns correspond to predictions based on small, medium, and large annuli, respectively (circles in Fig. 8d,i). All annuli refer to distances *beyond* one orientation hypercolumn thus the small annulus should *not* be confused with distances less than the diameter of one hypercolumn. In a,b the same sampling procedure and the same sample set sizes described in Fig. 10 were repeated. For lack of space we omit the very similar graphs of the mean and median of the entire population, and present predictions from the texture model only. **(a)** Expected mean distribution and standard error (N=4, 100 repetitions). Note the spread with increased retinotopic distance. **(b)** Expected median distribution and standard error (N=7, 100 repetitions). Note developing symmetric peaks that further depart from the iso-orientation domain as the spatial distance increases. The correspondence of these peaks to the minima of the standard error is remarkable, thus designating them as statistical “anchors” suitable for empirical verification. **(c)** Individually tuned cells show the qualitative difference between distributions of high and low curvature cells. In particular, note how the distributions of high curvature cells (dashed graphs) are the ones that develop the peaks mentioned in panel **b** above.