Biol. Cybern. 88, 2–10 (2003)
DOI 10.1007/s00422-002-0370-x

Biological Cybernetics

A rotation and translation invariant discrete saliency network

Lance R. Williams[1], **John W. Zweck**[2]

[1] Department of Computer Science, University of New Mexico, Albuquerque, NM 87110, USA
[2] Department of CS and EE, University of Maryland, Baltimore County, MD 21250, USA

Received: 11 July 2002 / Accepted in revised form: 25 October 2002

Abstract. We describe a neural network that enhances and completes salient closed contours in images. Our work is different from all previous work in three important ways. First, like the input provided to primary visual cortex (V1) by the lateral geniculate nucleus (LGN), the input to our computation is isotropic. That is, it is composed of spots, not edges. Second, our network computes a well-defined function of the input based on a distribution of closed contours characterized by a random process. Third, even though our computation is implemented in a discrete network, its output is invariant to continuous rotations and translations of the input image.

1 Introduction

Vision, whether that of humans or machines, must solve certain fundamental information-processing problems. One such problem is robustness to rotations and translations of the input image. Ideally, the output of a visual computation should vary continuously under continuous transformation of its input. We believe that an understanding of how the human visual system achieves this invariance will lead to significant improvements in the performance of machine vision systems. Conversely, by studying the fundamental information processing problems common to both human and machine vision we can achieve a deeper understanding of the structure and function of the human visual cortex.

There is a long history of research on neural networks inspired by the structure of visual cortex whose functions have been described as contour completion, saliency enhancement, orientation sharpening, or segmentation, e.g., August (2000), Li (1998), Yen and Finkel (1997), Guy and Medioni (1996), Iverson (1993), Heitger and von der Heydt (1993), Parent and Zucker (1989), Sha'ashua and Ullman (1988), and Grossberg and Mingolla (1985). In this paper we describe a neural network that enhances and completes salient closed contours in images. Our work is different from all previous work in three important ways. First, like the input provided to primary visual cortex (V1) by the lateral geniculate nucleus (LGN), the input to our computation is isotropic. That is, the input is composed of spots, not edges.[1] Second, our network computes a well-defined function of the input based on a distribution of closed contours characterized by a random process. Third, even though our computation is implemented in a discrete network, its output is invariant to continuous rotations and translations of the input image.[2]

There are two important properties a computation must possess if it is to be invariant to rotations and translations, i.e., Euclidean invariant. First, the input, the output, and all intermediate representations must be Euclidean invariant. Second, all transformations of these representations must also be Euclidean invariant. Previous models are not Euclidean invariant, first and foremost because their input representations are not Euclidean invariant. That is, not all rotations and translations of the input can be represented equally well. This problem is often skirted by choosing input patterns that match particular choices of sampling rate and phase. For example, Li (1998) used only six samples in orientation (including 0°) and Heitger and von der Heydt (1993) only twelve (including 0°, 60° and 120°). Li's first test pattern was a dashed line of orientation, 0°, while Heitger and von der Heydt (1993) used a Kanizsa Triangle with sides of 0°, 60°,

Correspondence to: L.R. Williams
(e-mail: Williams@cs.unm.edu)

[1] Like our model, Guy and Medioni's (1996) can also be applied to input patterns consisting of spots.

[2] Guy and Medioni (1996) represent a distribution on $\mathbf{R}^2 \times S^1$ as an $N \times N$ array of 2×2 scatter matrices. Ignoring the issue of the biological plausibility of such an encoding scheme, we observe that although the representation of S^1 is clearly continuous, the rectangular grid used to sample $\mathbf{R}^2$ is not. It follows that their computation is not Euclidean invariant.

and 120° orientation.[3] There is no reason to believe that the experimental results shown in these papers would be similar if the input images were rotated by as little as 5°. To our knowledge, until now no one has ever commented on this problem.

2 A continuum formulation of the saliency problem

The following section reviews the continuum formulation of the contour completion and saliency problem as described in Williams and Thornber (2001).

2.1 Shape distribution

Mumford (1994) observed that the probability distribution of object boundary shapes could be modeled by a *Fokker-Planck* equation of the following form,

$$\frac{\partial p}{\partial t} = -\cos\theta\frac{\partial p}{\partial x} - \sin\theta\frac{\partial p}{\partial y} + \frac{\sigma^2}{2}\frac{\partial^2 p}{\partial \theta^2} - \frac{1}{\tau}p \ , \tag{1}$$

where $p(\boldsymbol{x}, \theta; t)$ is the probability that a particle is located at position $\boldsymbol{x} = (x, y)$ and is moving in direction θ at time t. This partial differential equation can be viewed as a set of independent *advection* equations in x and y (the first and second terms) coupled in the θ dimension by the *diffusion* equation (the third term). The advection equations translate probability mass in direction, θ, with unit speed, while the diffusion term implements a Brownian motion in direction, with *diffusion parameter* σ. The combined effect of these three terms is that particles tend to travel in straight lines, but over time they drift to the left or right by an amount proportional to σ^2. Finally, the effect of the fourth term is that particles decay over time, with a half-life given by the decay constant τ.

2.2 The propagators

The Green's function, $G(\boldsymbol{x}, \theta; t_1 | \boldsymbol{u}, \phi; t_0)$, gives the probability that a particle observed at position $\boldsymbol{u}$ and direction ϕ at time t_0 will later be observed at position $\boldsymbol{x}$ and direction θ at time t_1. It is the solution, $p(\boldsymbol{x}, \theta; t_1)$, of the Fokker-Planck initial value problem with initial value $p(\boldsymbol{x}, \theta; t_0) = \delta(\boldsymbol{x} - \boldsymbol{u})\delta(\theta - \phi)$ where δ is the Dirac delta function. Unlike Williams and Thornber (2001), who assumed that the location of input edges could be specified with arbitrary precision, we explicitly model the limited spatial acuity of the edge-detection process. We consider two edges to be indistinguishable if they are separated by a distance smaller than some scale, Δ. The Green's function is used to define two *propagators*. The long-time propagator:

$$P_0(\boldsymbol{x}, \theta | \boldsymbol{u}, \phi) = \int_0^\infty \mathrm{d}t \chi(t)\, G(\boldsymbol{x}, \theta; t | \boldsymbol{u}, \phi; 0) \tag{2}$$

gives the probability that $(\boldsymbol{x}, \theta)$ and $(\boldsymbol{u}, \phi)$ are distinguishable edges from the boundary of a single object. The short-time propagator:

$$P_1(\boldsymbol{x}, \theta | \boldsymbol{u}, \phi) = \int_0^\infty \mathrm{d}t[1 - \chi(t)]\, G(\boldsymbol{x}, \theta; t | \boldsymbol{u}, \phi; 0) \tag{3}$$

gives the probability that $(\boldsymbol{x}, \theta)$ and $(\boldsymbol{u}, \phi)$ are from the boundary of a single object but are indistinguishable. In both of these propagators, $\chi(.)$ is a cut-off function with $\chi(0) = 0$ and $\lim_{t\to\infty}\chi(t) = 1$:

$$\chi(t) = \frac{1}{2}\left[1 + \frac{2}{\pi}\mathrm{atan}\left(\mu\left[\frac{t}{\Delta} - \alpha\right]\right)\right] \ . \tag{4}$$

The cutoff function is characterized by three parameters – α, μ, and Δ. The parameter Δ is the scale of the edge-detection process. Relative to this scale, the parameter α specifies the location of the cutoff, and μ specifies its hardness.

2.3 Eigenfunctions

We use the long-time propagator, P_0, to define an integral linear operator, $Q(.)$, which combines three sources of information: (1) the probability that two edges belong to the same object, (2) the probability that the two edges are distinguishable, and (3) the probability that the two edges exist. It is defined as follows:

$$Q(\boldsymbol{x}, \theta | \boldsymbol{u}, \phi) = b(\boldsymbol{x})^{\frac{1}{2}} P_0(\boldsymbol{x}, \theta | \boldsymbol{u}, \phi) b(\boldsymbol{u})^{\frac{1}{2}} \tag{5}$$

where the *input bias function*, $b(\boldsymbol{x})$, gives the probability that an edge exists at $\boldsymbol{x}$. Note that, as defined, $b(\boldsymbol{x})$ depends only on the position, $\boldsymbol{x}$, and not the direction, θ, of the edge. This reflects the fact that the input to the computation is isotropic.[4]

As described in Williams and Thornber (2001), the right and left eigenfunctions, $s(.)$ and $\bar{s}(.)$, of $Q(.)$ with the largest positive real eigenvalue, λ, play a central role in the computation of saliency:

$$\lambda s(\boldsymbol{x}, \theta) = \int\int\int_{\mathbf{R}^2\times S^1} \mathrm{d}\boldsymbol{u}\mathrm{d}\phi Q(\boldsymbol{x}, \theta | \boldsymbol{u}, \phi) s(\boldsymbol{u}, \phi) \tag{6}$$

$$\lambda \bar{s}(\boldsymbol{x}, \theta) = \int\int\int_{\mathbf{R}^2\times S^1} \mathrm{d}\boldsymbol{u}\mathrm{d}\phi \bar{s}(\boldsymbol{u}, \phi) Q(\boldsymbol{u}, \phi | \boldsymbol{x}, \theta) \ . \tag{7}$$

Because $Q(.)$ is invariant under the transformation

$$Q(\boldsymbol{x}, \theta | \boldsymbol{u}, \phi) = Q(\boldsymbol{u}, \phi + \pi | \boldsymbol{x}, \theta + \pi) \ , \tag{8}$$

which reverses the order and direction of its arguments, the right and left eigenfunctions are related by

$$\bar{s}(\boldsymbol{x}, \theta) = s(\boldsymbol{x}, \theta + \pi) \ . \tag{9}$$

[3] Nor are we blameless in this regard. Williams and Jacobs (1997) used 36 discrete orientations including 0°, 60°, and 120°, and also show results on a conveniently oriented Kanizsa Triangle.

[4] The model could easily be extended to include nonisotropic input by making $b(.)$ a function of both $\boldsymbol{x}$ and θ. This would allow us to model any orientation bias present in feedforward connections from LGN to V1.

2.4 Stochastic completion field

The fact that the distribution of closed contours, or the *stochastic completion field*, can be factored into a *source field* derived from $s(.)$ and a *sink field* derived from $\bar{s}(.)$ is discussed in detail in Williams and Thornber (2001). The stochastic completion field, $c(\boldsymbol{u}, \phi)$, gives the probability that a closed contour containing any subset of the edges exists at $(\boldsymbol{u}, \phi)$. It is the sum of three terms:

$$c(\boldsymbol{u}, \phi) = \frac{p_0(\boldsymbol{u}, \phi)\bar{p}_0(\boldsymbol{u}, \phi) + p_0(\boldsymbol{u}, \phi)\bar{p}_1(\boldsymbol{u}, \phi) + p_1(\boldsymbol{u}, \phi)\bar{p}_0(\boldsymbol{u}, \phi)}{\lambda \int\int\int_{\mathbf{R}^2 \times S^1} \mathrm{d}\boldsymbol{x}\mathrm{d}\theta\, s(\boldsymbol{x}, \theta)\bar{s}(\boldsymbol{x}, \theta)} \tag{10}$$

where $p_m(\boldsymbol{u}, \phi)$ is a source field and $\bar{p}_m(\boldsymbol{u}, \phi)$ is a sink field:

$$p_m(\boldsymbol{u}, \phi) = \int\int\int_{\mathbf{R}^2 \times S^1} \mathrm{d}\boldsymbol{x}\mathrm{d}\theta\, P_m(\boldsymbol{u}, \phi|\boldsymbol{x}, \theta) b(\boldsymbol{x})^{\frac{1}{2}} s(\boldsymbol{x}, \theta)$$

$$\bar{p}_m(\boldsymbol{u}, \phi) = \int\int\int_{\mathbf{R}^2 \times S^1} \mathrm{d}\boldsymbol{x}\mathrm{d}\theta\, \bar{s}(\boldsymbol{x}, \theta) b(\boldsymbol{x})^{\frac{1}{2}} P_m(\boldsymbol{x}, \theta|\boldsymbol{u}, \phi) \ .$$

The purpose of writing $c(\boldsymbol{u}, \phi)$ in this way is to remove the contribution, $p_1(\boldsymbol{u}, \phi)\bar{p}_1(\boldsymbol{u}, \phi)$, of closed contours at scales smaller than Δ that would otherwise dominate the completion field. Given the above expression for the completion field, it is clear that the key problem is computing the eigenfunction, $s(.)$, of $Q(.)$ with the largest positive real eigenvalue. To accomplish this, we can use the well-known power method. See Golub and Van Loan (1996). In this case, the power method involves repeated application of the linear operator, $Q(.)$, to the function, $s(.)$, followed by normalization:

$$s^{(m+1)}(\boldsymbol{x}, \theta) = \frac{\int\int\int_{\mathbf{R}^2 \times S^1} \mathrm{d}\boldsymbol{u}\mathrm{d}\phi Q(\boldsymbol{x}, \theta|\boldsymbol{u}, \phi) s^{(m)}(\boldsymbol{u}, \phi)}{\int\int\int_{\mathbf{R}^2 \times S^1} \int\int\int_{\mathbf{R}^2 \times S^1} \mathrm{d}\boldsymbol{x}\mathrm{d}\theta\mathrm{d}\boldsymbol{u}\mathrm{d}\phi\, Q(\boldsymbol{x}, \theta|\boldsymbol{u}, \phi) s^{(m)}(\boldsymbol{u}, \phi)} \tag{11}$$

In the limit, as m gets very large, $s^{(m+1)}(\boldsymbol{x}, \theta)$ converges to the eigenfunction of $Q(.)$, with the largest positive real eigenvalue.[5] We observe that the above computation can be considered a continuous-state, discrete-time, recurrent neural network.

3 A neural implementation

The fixed point of the neural network we describe is the eigenfunction with the largest positive real eigenvalue of the linear integral operator, $Q(.)$, formed by composing the input-independent linear operator, $P_0(.)$, and the input-dependent linear operator, $b(.)$. The dynamics of the neural network are derived from the update equation for the standard power method for computing eigenvectors. It is useful to draw an analogy between our neural network for contour completion and the models for the emergence of orientation selectivity in V1 described by Somers et al. (1995) and Ben-Yishai et al. (1995). See Fig. 1. First, we can identify $s(.)$ with simple cells in V1 and the input-bias function, $b(.)$, which modulates $s(.)$ in the numerator of the update equation, with the feedforward excitatory connections from the LGN. Second, we can identify $P_0(.)$ with the intracortical excitatory connections, which Somers et al. hypothesize are primarily responsible for the emergence of orientation selectivity in V1. As in the model of Somers et al., these connections are highly specific and target mainly cells of similar orientation preference. Third, we identify the denominator of the update equation with the orientation-nonspecific intracortical inhibitory connections, which Somers et al. hypothesize keep the level of activity within bounds. Finally, we identify $c(.)$, the stochastic completion field, with the population of cells in V2 described by von der Heydt et al. (1984).

4 A discrete implementation of the continuum formulation

Following Zweck and Williams (2000), the continuous functions comprising the state of the computation are represented as weighted sums of a finite set of *shiftable-twistable* basis functions. The weights form the coefficient vectors for the functions. The computation we describe is biologically plausible in the sense that all transformations of state are effected by linear transformations (or other vector parallel operations) on the coefficient vectors.

4.1 Shiftable-twistable bases

The input and output of the above computation are functions defined on the continuous space, $\mathbf{R}^2 \times S^1$, of

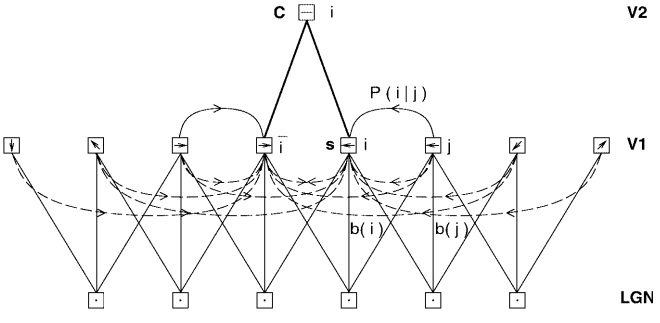

Fig. 1. *Thin solid lines* indicate feedforward connections from LGN, which provide a weak orientation bias, i.e., $b(.)$, to simple cells in V1, i.e., $s(.)$. *Solid lines with arrows* indicate orientation-specific intracortical excitatory connections, i.e., $P_0(.)$. *Dashed lines with arrows* indicate orientation-nonspecific intracortical inhibitory connections. *Thick solid lines* indicate feedforward connections between V1 and V2, i.e., $c(.)$

[5] It is only in the limit of large m, i.e., after the power-method iteration has converged, that $c(\boldsymbol{x}, \theta)$ represents a distribution of **closed** contours through location $(\boldsymbol{x}, \theta)$. Prior to convergence, $s^{(m)}(\boldsymbol{x}, \theta)\bar{s}^{(m)}(\boldsymbol{x}, \theta)$ can be considered a distribution of nonclosed contours of length $2m + 1$ edges centered at location $(\boldsymbol{x}, \theta)$.

positions in the plane, $\mathbf{R}^2$, and directions in the circle, S^1. For such computations, the important symmetry is determined by those transformations, $T_{\boldsymbol{x}_0,\theta_0}$, of $\mathbf{R}^2 \times S^1$, that perform a shift in $\mathbf{R}^2$ by $\boldsymbol{x}_0$, followed by a twist in $\mathbf{R}^2 \times S^1$ through an angle, θ_0. A *twist* through an angle, θ_0, consists of two parts: (1) a rotation, R_{θ_0}, of $\mathbf{R}^2$ and (2) a translation in S^1, both by θ_0. The symmetry $T_{\boldsymbol{x}_0,\theta_0}$, which is called a *shift-twist transformation*, is given by the formula

$$T_{(\boldsymbol{x}_0,\theta_0)}(\boldsymbol{x},\theta) = (R_{\theta_0}(\boldsymbol{x}-\boldsymbol{x}_0), \theta-\theta_0) \ . \tag{12}$$

The relationship between continuous shift-twist transformations and computations in primary visual cortex has been hypothesized by Williams and Jacobs (1997) and by Bressloff et al. (2001). The continuous shift-twist group characterizes the symmetry of the Green's function of the Fokker-Planck equation described by Mumford (1994):

$$G(\boldsymbol{x},\theta;t_1|\boldsymbol{u},\phi;t_0) = G(R_\phi(\boldsymbol{x}-\boldsymbol{u}),\theta-\phi;t_1-t_0|0,0;0) \ . \tag{13}$$

The *continuous* shift-twist symmetry of the Green's function can be understood by considering the analogous *discrete* symmetry of highest order. An idealized model of V1 might be based on a set of hypercolumns arranged in a hex lattice and with six discrete orientation preferences (see Fig. 2). By virtue of the shift-twist symmetry of the interconnection matrix, the long-range, orientation-specific, intracortical excitatory connections labeled A and B have equal weight. However, because this lattice is not invariant under continuous shift-twist transformations, the output of such a model (unlike ours) would not be Euclidean invariant. A visual computation, C, on $\mathbf{R}^2 \times S^1$ is called *shift-twist invariant* if, for *all* $(\boldsymbol{x}_0,\theta_0) \in \mathbf{R}^2 \times S^1$, a shift-twist of the input by $(\boldsymbol{x}_0,\theta_0)$ produces an identical shift-twist of the output. This property can be depicted in the following commutative diagram:

$$\begin{array}{ccc} b(\boldsymbol{x},\theta) & \xrightarrow{C} & c(\boldsymbol{x},\theta) \\ \downarrow T_{\boldsymbol{x}_0,\theta_0} & & \downarrow T_{\boldsymbol{x}_0,\theta_0} \\ b(R_{\theta_0}(\boldsymbol{x}-\boldsymbol{x}_0),\theta-\theta_0) & \xrightarrow{C} & c(R_{\theta_0}(\boldsymbol{x}-\boldsymbol{x}_0),\theta-\theta_0), \end{array}$$

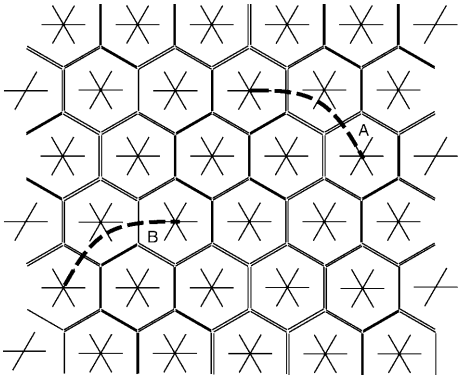

Fig. 2. The *continuous* shift-twist symmetry of the Green's function may be understood by considering the analogous *discrete* symmetry of highest order

where $b(.)$ is the input, $c(.)$ is the output, $\xrightarrow{C}$ is the computation, and $\xrightarrow{T_{\boldsymbol{x}_0,\theta_0}}$ is the shift-twist transformation. Correspondingly, we define a *shiftable-twistable basis*[6] of functions on $\mathbf{R}^2 \times S^1$ to be a set of functions on $\mathbf{R}^2 \times S^1$ with the property that whenever a function, $f(\boldsymbol{x},\theta)$, is in their span, then $f(T_{\boldsymbol{x}_0,\theta_0}(\boldsymbol{x},\theta))$ is as well for every choice of $(\boldsymbol{x}_0,\theta_0)$ in $\mathbf{R}^2 \times S^1$. As such, the notion of a shiftable-twistable basis on $\mathbf{R}^2 \times S^1$ generalizes that of a shiftable-steerable basis on $\mathbf{R}^2$ introduced by Freeman and Adelson (1991) and Simoncelli et al. (1992).

Shiftable-twistable bases can be constructed as follows. Let $\Psi(\boldsymbol{x},\theta)$ be a function on $\mathbf{R}^2 \times S^1$ that is periodic (with period X) in both spatial variables, $\boldsymbol{x}$. In analogy with the definition of a shiftable-steerable function on $\mathbf{R}^2$, we say that Ψ is *shiftable-twistable* on $\mathbf{R}^2 \times S^1$ if there are integers, K and M, and *interpolation functions*, $a_{\boldsymbol{k},m}(\boldsymbol{x}_0,\theta_0)$, such that for each $(\boldsymbol{x}_0,\theta_0) \in \mathbf{R}^2 \times S^1$, the shift-twist of Ψ by $(\boldsymbol{x}_0,\theta_0)$ is a linear combination of a finite number of basic shift-twists of Ψ by amounts $(\boldsymbol{k}\Delta, m\Delta_\theta)$, i.e., if

$$\Psi(T_{\boldsymbol{x}_0,\theta_0}(\boldsymbol{x},\theta)) = \sum_{\boldsymbol{k},m} a_{\boldsymbol{k},m}(\boldsymbol{x}_0,\theta_0)\Psi(T_{\boldsymbol{k}\Delta,m\Delta_\theta}(\boldsymbol{x},\theta)) \ . \tag{14}$$

Here $\Delta = X/K$ is the *basic shift amount* and $\Delta_\theta = 2\pi/M$ is the *basic twist amount*. The sum in Eq. (14) is taken over all pairs of integers, $\boldsymbol{k} = (k_x,k_y)$, in the range $0 \le k_x,k_y < K$, and all integers, m, in the range $0 \le m < M$. As in Simoncelli et al. (1992) the interpolation functions, $a_{\boldsymbol{k},m}$, are sinc functions.

The Gaussian-Fourier basis is the product of a shiftable-steerable basis of Gaussians in $\boldsymbol{x}$ and a Fourier series basis in θ. For the experiments in this paper, the standard deviation of the Gaussian basis function, $g(\boldsymbol{x}) = \frac{1}{\Delta}e^{-\|\boldsymbol{x}\|^2/2\Delta^2}$, equals the basic shift amount, Δ. We regard $g(\boldsymbol{x})$ as a periodic function of period, X, which is chosen to be much larger than Δ, so that $g(X/2,X/2)$ and its derivatives are essentially zero. For each frequency, ω, and shift amount, Δ (where $K = X/\Delta$ is an integer), we define the *Gaussian-Fourier basis functions*, $\Psi_{\boldsymbol{k},\omega}$, by

$$\Psi_{\boldsymbol{k},\omega}(\boldsymbol{x},\theta) = g(\boldsymbol{x}-\boldsymbol{k}\Delta)e^{i\omega\theta} \ . \tag{15}$$

Zweck and Williams (2000) showed that the Gaussian-Fourier basis is shiftable-twistable.

4.2 Power method update formula

Suppose that $Q(.)$'s approximate eigenfunction, $s^{(m)}(\boldsymbol{x},\theta)$, can be represented in the Gaussian-Fourier basis as

$$s^{(m)}(\boldsymbol{x},\theta) = \sum_{\boldsymbol{k},\omega} s^{(m)}_{\boldsymbol{k},\omega}\Psi_{\boldsymbol{k},\omega}(\boldsymbol{x},\theta) \ . \tag{16}$$

The vector, $\mathbf{s}^{(m)}$, with components, $s^{(m)}_{\boldsymbol{k},\omega}$, will be called the *coefficient vector* of $s^{(m)}(\boldsymbol{x},\theta)$. In the next two sections,

[6] We use this terminology even though the basis functions need not be linearly independent.

we demonstrate how the following integral linear transform:

$$s^{(m+1)}(\boldsymbol{x},\theta) = \int\int\int_{\mathbf{R}^2\times S^1} \mathrm{d}\boldsymbol{u}\mathrm{d}\phi\, P_0(\boldsymbol{x},\theta|\boldsymbol{u},\phi) b(\boldsymbol{u}) s^{(m)}(\boldsymbol{u},\phi) \quad (17)$$

(i.e., the basic step in the power method) can be implemented as a discrete linear transform in a Gaussian-Fourier shiftable-twistable basis:

$$\mathbf{s}^{(m+1)} = \mathbf{PBs}^{(m)} \ . \quad (18)$$

4.3 The propagation operator **P**

In practice, we do not explicitly represent the matrix, **P**. This is because it is quite large, and explicitly computing the above matrix-vector product using the matrix would be prohibitively expensive. Instead, we compute the matrix-vector product using the advection-diffusion-decay operator in the Gaussian-Fourier shiftable-twistable basis, $\mathbf{A}\circ\mathbf{D}$, which is derived and described in detail in Zweck and Williams (2000). The coefficient vector, $\mathbf{s}^{(m+1)}$, is a weighted average of solutions, $\mathbf{q}^{(m,n)}$, of a Fokker-Planck initial-value problem. The average is over all positive times, $t = n\Delta t$. That is,

$$\mathbf{s}^{(m+1)} = \sum_{n=0}^{\infty} \chi(n\Delta t)\mathbf{q}^{(m,n)} \quad (19)$$

where $\chi(.)$ is the cutoff function, and the initial value for $\mathbf{q}$ is the product of the bias operator, **B**, and the coefficient vector, $\mathbf{s}^{(m)}$,

$$\mathbf{q}^{(m,0)} = \mathbf{Bs}^{(m)} \ . \quad (20)$$

For convenience, we introduce a variable, $\mathbf{p}^{(m,n)}$, to represent the partial sum of $\mathbf{q}^{(m,n)}$ from $t = 0$ to $t = n\Delta t$. This leads to the following update equations:

$$\mathbf{q}^{(m,n+1)} = (\mathbf{A}\circ\mathbf{D})\mathbf{q}^{(m,n)} \quad (21)$$

$$\mathbf{p}^{(m,n+1)} = \mathbf{p}^{(m,n)} + \chi(n\Delta t)\mathbf{q}^{(m,n+1)} \ , \quad (22)$$

where **A** is the the advection operator and **D** is the diffusion-decay operator. In the shiftable-twistable basis, the advection operator, **A**, is a discrete convolution:

$$q^{(m,n+\frac{1}{2})}_{\boldsymbol{\ell},\boldsymbol{\eta}} = \sum_{\boldsymbol{k},\omega} \hat{a}_{\boldsymbol{\ell}-\boldsymbol{k},\eta-\omega}(\Delta t)\, q^{(m,n)}_{\boldsymbol{k},\omega} \quad (23)$$

with the kernel,

$$\hat{a}_{\boldsymbol{k},\eta}(\Delta t) = \frac{1}{2\pi}\int_0^{2\pi} a_{\boldsymbol{k}}(\Delta t[\cos\theta,\sin\theta]^{\mathrm{T}})\exp(-i\eta\theta)\,\mathrm{d}\theta \quad (24)$$

where the $a_{\boldsymbol{k}}$ are interpolation functions. Let N be the number of Fourier-series frequencies, ω, used in the shiftable-twistable basis, and let $\Delta\theta = 2\pi/N$. The diffusion-decay operator, **D**, is given by the diagonal matrix-vector product:

$$q^{(m,n+1)}_{\boldsymbol{k},\omega} = \mathrm{e}^{-\Delta t/\tau}(\lambda\mathrm{e}^{-i\omega\Delta\theta} + (1-2\lambda) + \lambda\mathrm{e}^{i\omega\Delta\theta})\, q^{(m,n+\frac{1}{2})}_{\boldsymbol{k},\omega} \ , \quad (25)$$

where $\lambda = \frac{\sigma^2}{2}\frac{\Delta t}{(\Delta\theta)^2}$.

4.4 The bias operator **B**

In the continuum, the bias operator multiplies the function, $s(.)$, by the input-bias function, $b(.)$. Our aim is to identify an equivalent linear operator in the shiftable-twistable basis. Suppose that both $s(.)$ and $b(.)$ are represented in a Gaussian basis, $g_{\boldsymbol{k}}(\boldsymbol{x}) = g(\boldsymbol{x}-\boldsymbol{k}\Delta)$. Their product is:

$$b(\boldsymbol{x})s(\boldsymbol{x}) = \sum_{\boldsymbol{k}} s_{\boldsymbol{k}} g_{\boldsymbol{k}}(\boldsymbol{x}) \cdot \sum_{\boldsymbol{\ell}} b_{\boldsymbol{\ell}} g_{\boldsymbol{\ell}}(\boldsymbol{x}) \quad (26)$$

$$= \sum_{\boldsymbol{k}}\sum_{\boldsymbol{\ell}} s_{\boldsymbol{k}} b_{\boldsymbol{\ell}} g_{\boldsymbol{k}}(\boldsymbol{x}) g_{\boldsymbol{\ell}}(\boldsymbol{x}) \ . \quad (27)$$

Now, the product of two Gaussian basis functions, $g_{\boldsymbol{k}}$ and $g_{\boldsymbol{\ell}}$, is a Gaussian of smaller variance (with higher-frequency content) that cannot be represented in the Gaussian basis, $g_{\boldsymbol{k}}$. Because $b(\boldsymbol{x})s(\boldsymbol{x})$ is a linear combination of the products of pairs of Gaussian-basis functions, it cannot be represented in the Gaussian basis, either. However, we observe that the convolution of $b(\boldsymbol{x})s(\boldsymbol{x})$ and a Gaussian, $h(\boldsymbol{x}) * [b(\boldsymbol{x})s(\boldsymbol{x})]$, where

$$h(\boldsymbol{x}) = \frac{1}{\Delta^2\pi}\mathrm{e}^{-||\boldsymbol{x}||^2/\Delta^2} \quad (28)$$

can be represented in the Gaussian basis. It follows that there exists a matrix, **B**, such that

$$h(\boldsymbol{x}) * [b(\boldsymbol{x})s(\boldsymbol{x})] = \sum_{\boldsymbol{k}} [\mathbf{Bs}]_{\boldsymbol{k}} g_{\boldsymbol{k}}(\boldsymbol{x}) \ . \quad (29)$$

The formula for the matrix, **B**, is derived by first completing the square in the exponent of the product of two Gaussians to obtain:

$$g_{\mathbf{k}}(\boldsymbol{x}) g_{\boldsymbol{\ell}}(\boldsymbol{x}) = g(\sqrt{2}(\boldsymbol{x} - \frac{\Delta}{2}(\boldsymbol{k}+\boldsymbol{\ell})))g(\frac{\Delta}{\sqrt{2}}(\boldsymbol{k}-\boldsymbol{\ell})) \ . \quad (30)$$

This product is then convolved with h to obtain a function, $f(\boldsymbol{x})$, which is a shift of the Gaussian basis function, $g(\boldsymbol{x})$. Finally, we use the shiftability formula

$$g(\boldsymbol{x}-\boldsymbol{x}_0) = \sum_{\boldsymbol{k}} a_{\boldsymbol{k}}(\boldsymbol{x}_0) g_{\boldsymbol{k}}(\boldsymbol{x}) \ , \quad (31)$$

where $a_{\boldsymbol{k}}$ are the interpolation functions, to express $f(\boldsymbol{x})$ in the Gaussian basis. The result is

$$B_{\boldsymbol{k},\boldsymbol{\ell}} = \sum_{\boldsymbol{i}} b_{\boldsymbol{i}} \exp(-||\boldsymbol{i}-\boldsymbol{\ell}||^2/4) a_{\boldsymbol{k}}(\Delta(\boldsymbol{i}+\boldsymbol{\ell})/2) \ . \quad (32)$$

Although this formula is very general, using it to compute the matrix-vector product would be prohibitively expensive due to the nonsparseness of the matrix **B**. To overcome this problem we have developed an

alternative representation for the bias operator in the Gaussian-Fourier basis. For this representation, we assume that the input-bias function is of the form

$$b(\boldsymbol{x}) = \sum_{j=1}^{N} g(\boldsymbol{x} - \boldsymbol{x}_j) \ , \tag{33}$$

where the centers, $\boldsymbol{x}_j$, of the Gaussian spots are assumed to be far apart from each other relative to their widths. In this case, the matrix $\mathbf{B}$ representing the bias operator can be expressed as a sum of N outer products of vectors representing values of the Gaussian-basis functions g_ℓ and values of the sinc interpolation functions a_k at the locations $\boldsymbol{x}_j$ of the Gaussian spots:

$$B_{k,\ell} = \sum_{j=1}^{N} g_\ell(\boldsymbol{x}_j) a_k(\boldsymbol{x}_j) \ . \tag{34}$$

Consequently, the action of the bias operator on a coefficient vector is given by:

$$[\mathbf{Bs}]_k = \sum_{j=1}^{N} \left[\sum_\ell s_\ell g_\ell(\boldsymbol{x}_j) \right] \cdot a_k(\boldsymbol{x}_j) \ . \tag{35}$$

In this form, the bias operator is relatively inexpensive to compute. For this reason, it was used in the experimental results described below. However, in the continuum, this representation of the bias operator does not agree with the representation given in Eq. 32. To understand how the second representation is related to the first, we now describe how it was derived. Assume that the input-bias function is simply $b(\boldsymbol{x}) = g(\boldsymbol{x} - \boldsymbol{x}_0)$ and that $s(\boldsymbol{x}) = \sum_\ell s_\ell g_\ell(\boldsymbol{x})$. Then

$$\sum_\ell [\mathbf{Bs}]_\ell g_\ell(\boldsymbol{x}) = b(\boldsymbol{x}) s(\boldsymbol{x}) \sum_\ell s_\ell g(\boldsymbol{x} - \boldsymbol{x}_0) g_\ell(\boldsymbol{x}) \tag{36}$$

$$\simeq \left[\sum_\ell s_\ell g_\ell(\boldsymbol{x}_0) \right] \cdot g(\boldsymbol{x} - \boldsymbol{x}_0) \ , \tag{37}$$

where the final equality holds exactly in the limit that $g(\boldsymbol{x} - \boldsymbol{x}_0) \to \delta(\boldsymbol{x} - \boldsymbol{x}_0)$, where δ denotes the Dirac delta function. The second representation of the bias operator in the basis then follows by expressing $g(\boldsymbol{x} - \boldsymbol{x}_0)$ in the Gaussian basis using the shiftability formula.

5 Experimental results

In our experiments the Gaussian-Fourier basis consisted of $K = 192$ translates (in each spatial dimension) of a Gaussian (of period, $X = 70.0$) and $N = 92$ harmonic signals in the orientation dimension. The standard deviation of the Gaussian was set equal to the shift amount, $\Delta = X/K$. For illustration purposes, all functions were rendered at a resolution of 256×256. The diffusion parameter, σ, equaled 0.1473, and the decay constant, τ, equaled 12.5. The time step, Δt, used to solve the Fokker-Planck equation in the basis equaled $\Delta/2$. The parameters for the cutoff function used to eliminate self-loops were $\alpha = 4$ and $\mu = 15$.

In the first experiment, the input-bias function, $b(\boldsymbol{x})$, consists of eight Gaussian spots spaced equally on a circle. The positions are real valued, i.e., they do not lie on the grid of basis functions. The stochastic completion field, $\int_{S^1} c(\boldsymbol{u}, \phi) \mathrm{d}\phi$, after five iterations of the power method, is shown in Fig. 3. Although the local orientations of the circle at the locations of the spots were not specified in the input-bias function, they are an emergent property of the completion field.

In the second experiment, the input-bias function was created by adding twenty randomly positioned Gaussian spots to twenty spots on the boundary of an avocado. See Fig. 4 (left). The stochastic completion field computed using 32 iterations of the power method is shown in Fig. 4 (right).

In the last experiment, the input-bias function from the first experiment was rotated by 45° and translated by half the distance between the centers of adjacent basis functions, $b(R_{45^\circ}(\boldsymbol{x} - [\frac{\Delta}{2}, \frac{\Delta}{2}]^{\mathrm{T}}))$. See Fig. 5 (left). The stochastic completion field is identical (up to rotation and translation) to the one computed in the first experiment. See Fig. 5 (right). Finally, we plot the estimate of the largest positive real eigenvalue, λ, as a function of m, the power-method iteration for both the rotated and nonrotated input patterns. See Fig. 6. The final value of λ

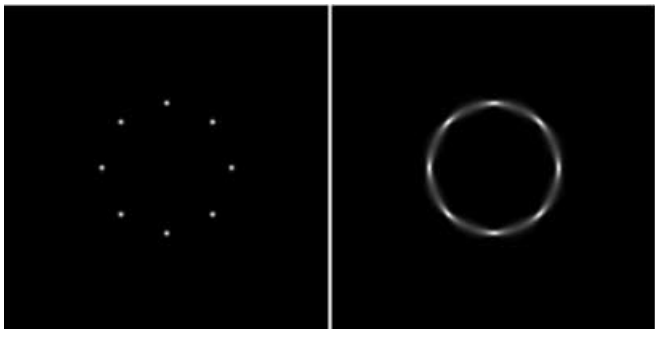

Fig. 3. *Left*: The input-bias function, $b(\boldsymbol{x})$, consists of eight Gaussian spots spaced equally on a circle. The positions are real valued, i.e., they do not lie on the grid of basis functions. *Right*: The stochastic completion field, $\int_{S^1} c(\boldsymbol{u}, \phi)\, \mathrm{d}\phi$, computed using $192 \times 192 \times 92$ basis functions

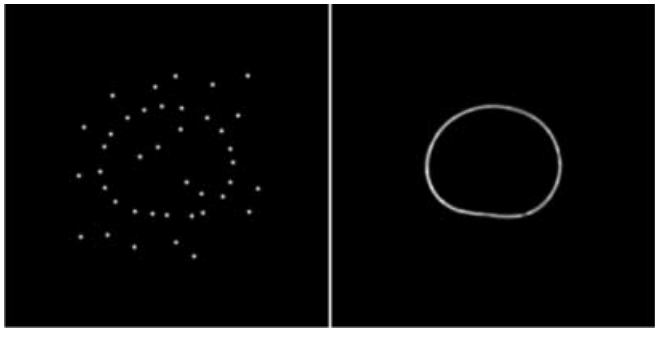

Fig. 4. *Left*: The input-bias function was created by adding 20 randomly positioned Gaussian spots to 20 spots on the boundary of an avocado. The positions are real valued, i.e., they do not lie on the grid of basis functions. *Right*: The stochastic completion field, $\int_{S^1} c(\boldsymbol{u}, \phi)\, \mathrm{d}\phi$, computed using $192 \times 192 \times 92$ basis functions

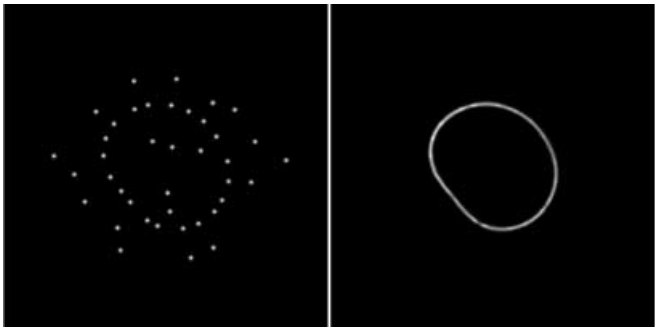

Fig. 5. *Left*: The input-bias function from Fig. 4, rotated by 45° and translated by half the distance between the centers of adjacent basis functions, $b(R_{45^\circ}(\boldsymbol{x} - [\frac{\Delta}{2}, \frac{\Delta}{2}]^{\mathrm{T}}))$. *Right*: The stochastic completion field is identical (up to rotation and translation) to the one shown in Fig. 4. This demonstrates the Euclidean invariance of the computation

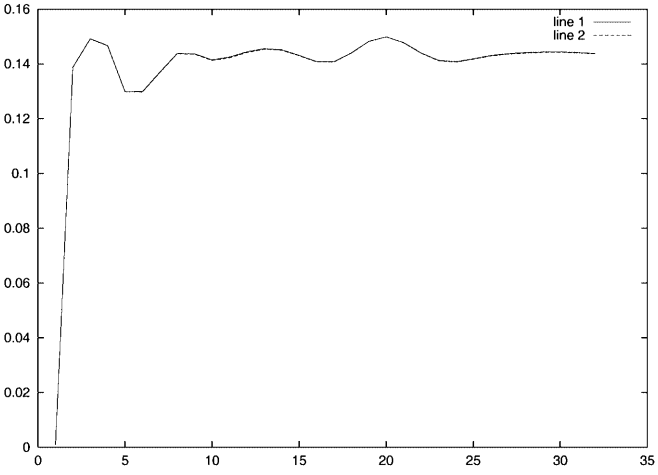

Fig. 6. The estimate of the largest positive real eigenvalue, λ, as a function of m, the power-method iteration for the rotated (*line 2*) and nonrotated (*line 1*) input patterns. The final value of λ (and all intermediate values) is identical in both cases; hence line 1 and line 2 overlap completely

(and all intermediate values) is identical in both cases. It follows that the $Q(.)$ operator has the same eigenvalues for the rotated and nonrotated inputs, i.e., the $Q(.)$'s are related by a similarity transform. This vividly demonstrates the Euclidean invariance of the computation.

6 Discussion

In this section we discuss the relationship of our model to the known structure and function of areas V1 and V2 of the visual cortex.

6.1 Isotropy vs. anisotropy

Given the physical anisotropy of our environment (a consequence of the gravitational field), the observation of a corresponding anisotropy in the statistics of natural scenes is not surprising. However, it does not follow that this environmental anisotropy is necessarily manifest in the computations performed by V1, or in its structure, especially since it is not manifest in the circularly symmetric shapes of retinal and LGN receptive fields–areas that immediately precede V1 in the visual pathway.

As far as the representation of orientation in V1 is concerned, the evidence from neuroscience is inconclusive. On the one hand, there is evidence for the overrepresentation of horizontal and vertical orientations in the primary visual cortex of ferrets. See Chapman et al. (1998). On the other hand, we are aware of no corresponding result for primates, and many authors have described the distribution of orientation preference in area V1 of monkeys as uniform, e.g., Bartfeld and Grinvald (1992), Hubel and Wiesel (1974).

6.2 Continuity vs. isotropy

A computation is continuous if and only if a continuous transformation of the input produces a continuous (not necessarily identical) transformation of the output. It follows that continuity with respect to rotation is a weaker property than isotropy. Indeed, there are well-known examples of visual computations that are continuous with respect to rotation but not isotropic, e.g., the dependence of shape-from-shading on illumination direction. See Sun and Perona (1998). Admittedly, in the model we describe here we have achieved the first (weaker) property by enforcing the second (stronger) property. However, to the best of our knowledge, it is an open question whether contour computations in V1 are isotropic or are merely continuous with respect to rotation.

In any case, it would be a mistake to conclude that, because human vision is not completely isotropic, there is a lack of continuity with respect to rotation in contour computations in V1. Yet apart from our own model, no existing model of contour completion has output that varies continuously under continuous rotation of its input. This is due to the nonshiftability-twistability of the representations of position and orientation used in these models.

6.3 Receptive fields vs. basis functions

In neuroscience, a receptive field characterizes the response of a neuron to a stimulus that varies in position, size, phase, and/or orientation. In contrast, a basis function is a concept from linear algebra. We hypothesize that the functions that comprise the states of visual computations are represented by expansions in bases consisting of a family of functions related by translation, dilation, phase-shift, and/or rotation. In other words, basis functions are the building blocks of visual intermediate representations.

Under the assumption that the retina, LGN, and V1 can be modeled as linear shift-invariant systems, the notion of a receptive field can be given a precise meaning. That is, it can be regarded as the *impulse response function* of the system. Under this assumption, the relationship of receptive fields to basis functions can also

be formally characterized as a *frame* and its *dual*. Recently, Salinas and Abbott (2000) have argued that simple-cell receptive fields do not form a *tight frame*. If they are correct, then the corresponding basis functions, which are only implicitly defined, do not have the same shape as the receptive fields. Consequently, even under the strong assumption of linearity, and even when the dimensionalities of the input and the intermediate representation are the same, it is a mistake to equate receptive fields and basis functions.

With regard to V1 and V2, this is a moot point, since the assumption of linearity is known to be unrealistic. Many researchers have demonstrated that the response of simple cells in V1 is modulated by stimuli outside of the classical (i.e., linear) receptive field. Furthermore, the dynamics of this modulatory activity is consistent with a computation implemented in a recurrent network, not a simple feedforward convolution of the input from LGN. The assumption of linearity is even less justified in V2, where von der Heydt et al. (1984) have identified neurons that "respond" to illusory contours of specific position and orientation. Since these neurons do not respond to an isolated spot, the concept of impulse response function is not useful.

6.4 Relative resolution of V1 and V2

The variance of the Gaussian-basis functions required to represent the completion field is half of that required to represent the source and sink fields. See Zweck and Williams (2000). It follows that our model would be supported if the spatial resolution of V2 were twice as high as that of V1. However, the fact that V2 receptive fields are, on average, larger than those of V1 does not mean that our model is wrong. Given the nonlinearity of the proposed computation, it would be a mistake to equate the size of receptive fields with the size of the basis functions used to represent the completion field. Indeed, given the spatial support of V2 receptive fields in V1, one would *expect* them to be significantly larger. A better measure of spatial resolution would be average distance (in visual angle) between orientation "pinwheel" centers.

However, it may simply be the case that the completion field is not computed at its theoretically maximum resolution relative to the resolutions of the source and sink fields, i.e., V2 represents a lowpass, downsampled completion field. We note that V1 directly projects to brain regions other than V2, primarily area MT but also back to LGN. It is possible that the computational requirements of MT (or perhaps another computation within V2) constrains the resolution of V1 more strongly than does contour completion.

6.5 Locality in orientation

One significant difference between our computer model and the visual cortex is the exact form of the basis functions that span $\mathbf{R}^2 \times S^1$. In a previous section, we proposed that observed long-range, intracortical, excitatory connections within V1 were the neural embodiment of the $P_0(.)$ operator. The fact that these connections are orientation-specific implies that the basis functions that represent the source field are localized in S^1. However, our implementation uses harmonic signals, which have zero locality, to represent the orientation dimension. The problem of identifying a basis of shiftable-twistable functions that are localized in orientation and specifying the corresponding $P_0(.)$ operator is a topic for future research.

7 Conclusion

We described a neural network that enhances and completes salient closed contours. Even though the computation is implemented in a discrete network, its output is invariant under continuous rotations and translations of the input.

Acknowledgements. L.R.W. was supported in part by Los Alamos National Laboratory. J.W.Z. was supported in part by the Albuquerque High Performance Computing Center. We wish to thank Jonas August and Steve Zucker for their insightful comments.

References

August J (2000) The curve indicator random field. PhD dissertation, Department of Computer Science, Yale University

Bartfeld E, Grinvald A (1992) Relationships between orientation preference pinwheels, cytochrome oxidase blobs and ocular dominance columns in primate striate cortex. Proc Nat Acad Sci USA 89: 1195–1199

Ben-Yishai R, Lev Bar-Or R, Sompolinksy H (1995) Theory of orientation tuning in visual cortex. Proc Nat Acad Sci USA 92: 3844–3848

Bressloff PC, Cowan JD, Golubitsky M, Thomas P, Wiener M (2001) Geometric visual hallucinations, Euclidean symmetry, and the functional architecture of striate cortex. Phil Trans Roy Soc (Lond) B 356: 1–32

Chapman B, Bonhoeffer T (1998) Overrepresentation of horizontal and vertical orientation preferences in developing ferret area 17. Proc Nat Acad Sci USA 95: 2609–2614

Freeman W, Adelson E (1991) The design and use of steerable filters. IEEE Trans Pattern Anal Mach Intell 13(9): 891–906

Golub GH, Van Loan CF (1996) Matrix Computations. Johns Hopkins University Press, Baltimore, MD

Grossberg S, Mingolla E (1985) Neural dynamics of form perception: boundary completion, illusory contours and neon color spreading. Psychol Rev 92: 173–211

Guy G, Medioni G (1996) Inferring global perceptual contours from local features. Int J Comput Vis 20: 113–133

Heitger R, von der Heydt R (1993) A computational model of neural contour processing, figure-ground and illusory contours. In: Proceedings of the 4th international conference on computer vision, Berlin

Hubel DH, Wiesel TN (1974) Sequence regularity and geometry of orientation columns in the monkey striate cortex. J Comp Neurol 158: 267–294

Iverson L (1993) Toward discrete geometric models for early vision. PhD dissertation, McGill University, Montreal

Li Z (1998) A neural model of contour integration in primary visual cortex. Neural Comput 10(4): 903–940

Mumford D (1994) Elastica and computer vision. In: Chandrajit B (ed) Algebraic geometry and its applications. Springer, Berlin Heidelberg New York, pp 491–506

Parent P, Zucker SW (1989) Trace inference, curvature consistency and curve detection. IEEE Trans Pattern Analy Mach Intell 11: 823–889

Salinas E, Abbott LF (2000) Do simple cells in primary visual cortex form a tight frame? Neural Comput 12(2): 313–335

Sha'ashua A, Ullman S (1988) Structural saliency: the detection of globally salient structures using a locally connected network. In: Proceedings of the 2nd international conference on computer vision. Clearwater, FL, pp 321–327

Simoncelli E, Freeman W, Adelson E, Heeger D (1992) Shiftable multiscale transforms. IEEE Trans Inf Theory 38(2): 587–607

Somers DC, Nelson SB, Sur M (1995) An emergent model of orientation selectivity in cat visual cortical cells. J Neurosci 15: 5448–5465

Sompolinsky H, Shapley R (1997) New perspectives on the mechanisms for orientation selectivity. Curr Opinion Neurobiol 7: 514–522

Sun J, Perona P (1998) Where is the sun? Nature Neurosci 1: 183–184

Von der Heydt R, Peterhans E, Baumgartner G (1984) Illusory contours and cortical neuron responses. Science 224: 1260–1262

Williams LR, Jacobs DW (1997) Stochastic completion fields: a neural model of illusory contour shape and salience. Neural Comput 9(4): 837–858

Williams LR, Thornber KK (2001) Orientation, scale, and discontinuity as emergent properties of illusory contour shape. Neural Comput 13(8): 1683–1711

Yen S, Finkel L (1997) Salient contour extraction by temporal binding in a cortically-based network. Adv Neural Inf Process Syst 10: 915–921

Zweck JW, Williams LR (2000) Euclidean group invariant computation of stochastic completion fields using shiftable-twistable functions. In: Proceedings of European conference on computer vision (ECCV '00), Dublin, Ireland