

Local Parallel Computation of Stochastic Completion Fields

Lance R. Williams

David W. Jacobs

NEC Research Institute, Princeton, NJ 08540, USA

We describe a local parallel method for computing the stochastic completion field introduced in the previous article (Williams and Jacobs, 1997). The stochastic completion field represents the likelihood that a completion joining two contour fragments passes through any given position and orientation in the image plane. It is based on the assumption that the prior probability distribution of completion shape can be modeled as a random walk in a lattice of discrete positions and orientations. The local parallel method can be interpreted as a stable finite difference scheme for solving the underlying Fokker-Planck equation identified by Mumford (1994). The resulting algorithm is significantly faster than the previously employed method, which relied on convolution with large-kernel filters computed by Monte Carlo simulation. The complexity of the new method is $O(n^3m)$, while that of the previous algorithm was $O(n^4m^2)$ (for an $n \times n$ image with m discrete orientations). Perhaps most significant, the use of a local method allows us to model the probability distribution of completion shape using stochastic processes that are neither homogeneous nor isotropic. For example, it is possible to modulate particle decay rate by a directional function of local image brightnesses (i.e., anisotropic decay). The effect is that illusory contours can be made to respect the local image brightness structure. Finally, we note that the new method is more plausible as a neural model since (1) unlike the previous method, it can be computed in a sparse, locally connected network, and (2) the network dynamics are consistent with psychophysical measurements of the time course of illusory contour formation.

1 Introduction

In recent years, several different researchers (Grossberg & Mingolla, 1985; Guy & Medioni, 1996; Heitger & von der Heydt, 1993; Williams & Jacobs, 1997) have formulated computational models of perceptual completion and illusory contours based on large-kernel convolution.¹ Although the function computed in each case is different, each method requires computing

¹ See "Stochastic Completion Fields" elsewhere in this journal for a detailed literature review.

one or more convolutions with filters represented by kernels of size equal to the largest gaps to be bridged (usually taken to be the size of the image). From the standpoint of computer vision, it should be clear that computing large numbers of convolutions with kernels of size equal to the image size is prohibitively expensive. These methods also have deficiencies as models of human visual processing, since implementation in vivo would require pairs of neurons separated by large amounts of visual arc to be densely interconnected. Although this possibility cannot be ruled out,² a model based on local connections clearly makes weaker demands on neuroanatomy (see Figure 1).

Data from human psychophysics are also inconsistent with methods requiring large-kernel convolution. First, experiments by Rock and Anson (1979) (and by Ramachandran, Ruskin, Cobb, Rogers-Ramachandran, & Tyler, 1994) demonstrate (see Figure 2) that illusory contours can be suppressed by the presence of texture or other figural elements along the path the completion would follow if these elements were absent. That is, the shape, salience, and sharpness of the completion are not solely a function of characteristics of the pair of elements it bridges but are also a function of the intervening pattern of image brightnesses. Second, recent studies of the time course of illusory contour formation (Rubin, Shapley, & Nakayama, 1995) show that the time required for an illusory contour to form is at least partly a function of the size of the gap to be bridged.³ Again, this is consistent with a local-iterative process since a global process would require an amount of time independent of the size of the gap to be bridged.

In the previous article (Williams and Jacobs, 1997), we presented a theory of the shape, salience, and sharpness of illusory contours and other perceptual completions. This theory was based on the assumption that the visual system computes the probability that an object boundary (possibly occluded) exists at each position and orientation in the visual field. This representation was termed a *stochastic completion field*. The inputs to this computation are measurements derived from the image about locations where completions might begin and end and a model for the probability distribution of completion shapes. Like Mumford (1994), we proposed that the prior probability distribution of completion shapes can be modeled by

² Absence of evidence is not evidence of absence.

³ This result is due to Rubin, Shapley, and Nakayama (1995), who used a forced-choice discrimination task together with masking stimulus to estimate the time required for illusory contour formation. They first showed that the time required for an illusory contour to form increases in direct proportion to distance when the "pacmen" in a Kanizsa display are "pulled apart." Somewhat surprisingly, they then showed that the time required for an illusory contour to form remains approximately constant when the entire figure is uniformly scaled. This somewhat paradoxical result is consistent with a stochastic completion field "pyramid" computed by repeated small-kernel convolution.




Figure 1: Two possible neural networks for computing stochastic completion fields. (a) Convolution with large-kernel filters requires that neurons separated by large amounts of visual arc be densely interconnected. (b) Repeated small-kernel convolution makes weaker demands on neuroanatomy since it can be implemented in a network with only local connections.

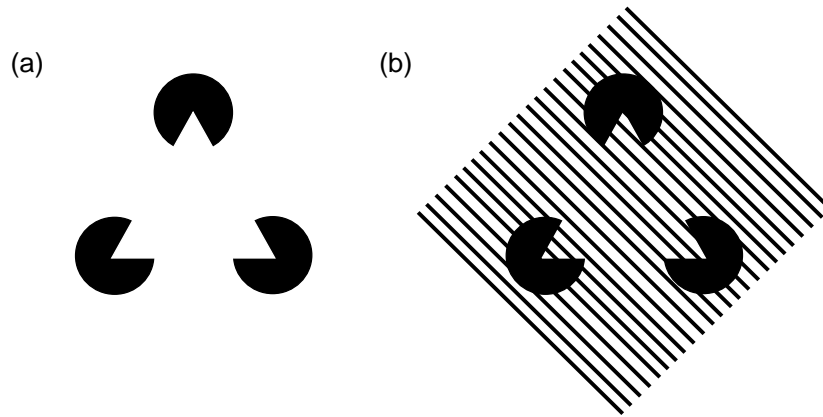


Figure 2: Rock and Anson (1979) show that illusory contours in the Kanizsa triangle can be suppressed by the presence of a background texture. (a) Kanizsa triangle. (b) Kanizsa triangle with background texture.

a particle undergoing a stochastic motion (i.e., a directional random walk).⁴ Unlike the familiar two-dimensional isotropic random walk, where a particle's state is simply its position in the plane, the particles of the directional random walk possess both position and orientation. The particle moves in the plane with constant speed in the direction specified by its orientation. Change in orientation ($\dot{\theta}$) is a normally distributed random variable with zero mean and variance equal to σ^2 so that the particle's orientation (θ) is a Brownian motion. The result is that particles tend to travel in straight lines but deviate from straightness by an amount that depends on σ^2 . This reflects a prior expectation of smoothness. In addition, a constant fraction of particles decay per unit time ($1 - e^{-\frac{1}{\tau}}$, where τ is the half-life of the particle), and this reflects a prior expectation of shortness.

If there exist two subsets of measurements, P and Q (representing the beginning and ending points of a set of boundary fragments), then the stochastic completion field, $C(u, v, \phi)$, is defined as the probability that a particle, initially at (x_p, y_p, θ_p) , for some $p \in P$, will, in the course of a di-

⁴ The important thing is not the particles themselves but the distribution of occluded shapes that the device of stochastic particle motion is intended to characterize. A particle follows a trajectory that represents one of the many possible shapes that an occluded object's boundary might have in the region where it is not directly visible. The probability that a particle follows a path of a particular shape is presumed to equal the frequency with which that shape occurs in the world.

rectional random walk, pass through (u, v, ϕ) , on its way to (x_q, y_q, θ_q) , for some $q \in Q$. This definition emphasizes the input and output of the computation that the visual system is performing (the function). It describes what information from the image and generic knowledge of the world should be combined to estimate the shape of occluded object boundaries. Using Marr's terminology (Marr, 1982), this is a *computational theory*—level model. We feel that our work demonstrates the value of Marr's approach for understanding human vision. This approach distinguishes between theories at three levels: (1) computational theories, (2) algorithm- and representation-level theories, and (3) implementation-level theories. Identical functions might be computable by different algorithms, and identical algorithms can be implemented (in the brain) in different ways. Many previous theories of figural completion have been specified only at the level of algorithm and representation. These models use an abstraction of biological neural networks as a kind of programming language. Unfortunately, it is difficult to understand exactly what problem these programs solve and what assumptions they make. Consequently the activity of units in these models and the computations they perform are often impossible to interpret. In contrast, our computational theory has allowed us to formulate algorithms based on operations that are meaningful for distributions of contours.⁵

In this article, we show how the stochastic completion field can be computed efficiently in a local parallel network. The network computation is based on a finite difference scheme for solving the Fokker-Planck equation identified by Mumford (1994). This algorithm avoids computationally prohibitive large-kernel convolutions and is more consistent with known psychophysics and neuroanatomy.

2 Review of Previous Method

Because the random walk is defined by a Markov process, $C(u, v, \phi)$ is proportional to the product of the probability that a particle beginning in a source will reach (u, v, ϕ) before it decays (the *source field*) and the probability that a particle beginning at (u, v, ϕ) will reach a sink before it decays (the *sink field*). Furthermore, because the Markov process is translation and rotation invariant, the probability that a particle will reach (u, v, ϕ) can be computed by "convolving" the source distribution, $p(x, y, \theta; 0)$, with the Green's function representing the probability that a particle at $(0, 0, 0)$ at time zero will reach (u', v', ϕ') ,

$$p'(u, v, \phi) = \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy \int_{-\pi}^{\pi} d\theta G'(u', v', \phi') p(x, y, \theta; 0),$$

⁵ It has also made obvious the relationship between the output of our algorithms and curves of least energy (see Mumford, 1994; "Stochastic Completion Fields," elsewhere in this journal). This relationship is totally opaque in previous models.

where (u', v', ϕ') is (u, v, ϕ) rotated by $-\theta$ about (x, y) , so that $u' = (u - x) \cos \theta + (v - y) \sin \theta$, $v' = -(u - x) \sin \theta + (v - y) \cos \theta$ and $\phi' = \phi - \theta$. The Green's function, G' , was originally computed by a Monte Carlo method that involved simulating the random walks of 1.0×10^6 particles on a 256×256 grid with 36 fixed orientations. The probability that a particle beginning at $(0, 0, 0)$ reaches (x, y, θ) before it decays was approximated by the fraction of simulated trajectories beginning at $(0, 0, 0)$ that intersect the region $(x \pm 1.0, y \pm 1.0, \theta \pm \pi/72)$. The sink field, $q'(u, v, \phi)$, which represents the probability that a particle leaving (u, v, ϕ) will reach a sink state before it decays, can be computed in a similar fashion,

$$q'(u, v, \phi) = \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy \int_{-\pi}^{\pi} d\theta G'(x', y', \theta') q(x, y, \theta; 0),$$

where (x', y', θ') is (x, y, θ) rotated by $-\phi$ about (u, v) , so that $x' = (x - u) \cos \phi + (y - v) \sin \phi$, $y' = -(x - u) \sin \phi + (y - v) \cos \phi$, and $\theta' = \theta - \phi$. Finally, the stochastic completion field, $C(u, v, \phi)$, which represents the relative likelihood that a particle leaving a source state will pass through (u, v, ϕ) and enter a sink state before it decays, equals the product of the source and sink fields:

$$C(u, v, \phi) = p'(u, v, \phi) \cdot q'(u, v, \phi).$$

Clearly, the run-time complexity (neglecting the time required to compute G) of the above process is dominated by the convolution with the Green's function G' . The overall run-time complexity is therefore $O(n^4 m^2)$ for an $n \times n$ image with m discrete orientations.

3 Local Parallel Computation

In this section we show how to compute the completion field in a local parallel network. Fundamentally, the computation is the same as before:

1. Compute the source field.
2. Compute the sink field.
3. Compute their product.

The difference is in the way the source and sink fields are computed. In the new algorithm, the source field is computed by integrating the Fokker-Planck equation for the stochastic process identified by Mumford (1994). The probability density function representing a particle's position at time t' is given by

$$p(x, y, \theta; t') = p(x, y, \theta; 0) + \int_0^{t'} \frac{\partial p(x, y, \theta; t)}{\partial t}$$

and

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

where P is $p(x, y, \theta; t)$. The Fokker-Planck equation for the stochastic process is best understood by considering each of the four terms separately. The first two terms taken together are the so-called *advection equation*:

$$\frac{\partial P}{\partial t} = -\cos \theta \frac{\partial P}{\partial x} - \sin \theta \frac{\partial P}{\partial y}.$$

Solutions of this equation have the form

$$p(x, y, \theta; t) = p(x + \cos \theta \Delta t, y + \sin \theta \Delta t, \theta; t + \Delta t)$$

for small Δt . That is, the probability density function for constant θ at times t and $t + \Delta t$ is related through a translation by the vector $[\cos \theta \Delta t, \sin \theta \Delta t]$. Intuitively, for constant θ , probability density is being transported in the θ direction with unit speed. The third term is the classical *diffusion equation*:

$$\frac{\partial P}{\partial t} = \sigma^2/2 \frac{\partial^2 P}{\partial \theta^2}.$$

Solutions of the diffusion equation are given by gaussian convolutions. Although, strictly speaking, the gaussian is not defined for angular quantities, for small Δt and variance σ^2 the following will approximately hold:

$$p(x, y, \theta; t + \Delta t) = \frac{1}{\sigma \sqrt{2\pi \Delta t}} \int_{-\pi}^{\pi} d\phi p(x, y, \theta - \phi; t) \cdot \exp(-\phi^2/2\sigma^2 \Delta t).$$

The last term implements the exponential decay:

$$\frac{\partial P}{\partial t} = -1/\tau P.$$

It is possible to interpret the Fokker-Planck equation in this case as describing a set of independent advection equations (one for each possible value of θ) coupled in the θ dimension by the diffusion equation. Taken together, the terms comprising the Fokker-Planck equation faithfully model the evolution in time of the probability density function representing a particle's position and orientation.

The value of the source field will also depend on the initial conditions (the probability density function [p.d.f.] describing the particle's position at time

zero). In our implementation, the p.d.f. defining the initial conditions consist of oriented impulses located at corners detected in the intensity image. These impulses represent the starting (or ending) points of possible completions. Strictly speaking, the Fokker-Planck equation can model our diffusion process only if we assume smooth initial conditions. Unfortunately, because the p.d.f.s defining the source and sink distributions are discontinuous, their partial derivatives are not defined. Moreover, the discrete method that we describe for solving the Fokker-Planck equation will be accurate only when a first-order approximation holds well over the length of scales that we discretely represent. For now, we will assume that the initial conditions are smooth and are well approximated by first-order expressions. We then show how to solve this partial differential equation iteratively. In the next section, we consider the accuracy of this method and describe a stochastic interpretation of the local method that does not require continuous initial conditions.

A standard technique for solving partial differential equations on a grid is to use the Taylor series expansion to write the partial derivatives at grid points as functions of the values of nearby grid points and of higher-order terms. Then by neglecting higher-order terms, we obtain an equation that expresses the value of a grid point at time $t + 1$ as a function of nearby grid points at time t , and which is accurate to first order (Ames, 1992; Ghez, 1988). This technique leads to the following iterative method for solving the Fokker-Planck equation:

$$\begin{aligned} \text{Step 1: } p_{x,y,\theta}^{t+1/4} &= p_{x,y,\theta}^t - \cos \theta \cdot \begin{cases} p_{x,y,\theta}^t - p_{x-1,y,\theta}^t & \text{if } \cos \theta > 0 \\ p_{x+1,y,\theta}^t - p_{x,y,\theta}^t & \text{if } \cos \theta < 0 \end{cases} \\ \text{Step 2: } p_{x,y,\theta}^{t+1/2} &= p_{x,y,\theta}^{t+1/4} - \sin \theta \cdot \begin{cases} p_{x,y,\theta}^{t+1/4} - p_{x,y-1,\theta}^{t+1/4} & \text{if } \sin \theta > 0 \\ p_{x,y+1,\theta}^{t+1/4} - p_{x,y,\theta}^{t+1/4} & \text{if } \sin \theta < 0 \end{cases} \\ \text{Step 3: } p_{x,y,\theta}^{t+3/4} &= \lambda p_{x,y,\theta-\Delta\theta}^{t+1/2} + (1 - 2\lambda) p_{x,y,\theta}^{t+1/2} + \lambda p_{x,y,\theta+\Delta\theta}^{t+1/2} \\ \text{Step 4: } p_{x,y,\theta}^{t+1} &= e^{-\frac{1}{\tau}} \cdot p_{x,y,\theta}^{t+3/4} \end{aligned}$$

where $\lambda = \sigma^2/2(\Delta\theta)^2$. These equations require a few comments. First we have split the computation into four separate steps. This standard fractional method allows us to compute $p(x, y, \theta; t + 1)$ with a series of separate convolutions in each of the three spatial directions. Second, notice that the first two steps of the computation depend on the sign of $\cos \theta$ and $\sin \theta$. To understand why, consider a particle traveling in the positive x direction. In this case, we wish to ensure that our approximation of $\partial p(x, y, \theta; t)/\partial x$ depends on $p(x - 1, y, \theta; t)$ but not on $p(x + 1, y, \theta; t)$ because $p(x, y, \theta; t + 1)$ depends on the first value but not the second. This method of ensuring that the

local difference used to approximate derivatives depends on the direction relevant to the underlying physical process is called *upwind differencing* and is necessary for stability.

Finally, we can see that the recurrence relationship given by these equations can be computed through small-kernel convolution. For example, we can compute the first recurrence using a 1×3 kernel centered on the pixel being convolved. This kernel is $[\cos \theta, 1 - \cos \theta, 0]$ for $\cos \theta > 0$ and $[0, 1 + \cos \theta, -\cos \theta]$ for $\cos \theta < 0$.

The source field, $p'(x, y, \theta)$, which represents the probability that a particle beginning at a source will reach (x, y, θ) before it decays, is computed by integrating the probability density function representing the particle's position over time:

$$p'(x, y, \theta) = \int_0^{\infty} dt p(x, y, \theta; t).$$

We can approximate the integral of $p(x, y, \theta; t)$ over all times less than some fixed time t' as follows,

$$p'(x, y, \theta; t') \approx \sum_{t=0}^{t'} p(x, y, \theta; t),$$

which can be computed by the following recurrence:

$$p'(x, y, \theta; t+1) = p'(x, y, \theta; t) + p(x, y, \theta; t+1).$$

Since $O(n)$ iterations are required to bridge the largest gaps that might be found in an image of size $n \times n$ and each iteration requires $O(n^2 m)$ time, the run-time complexity of the small-kernel method is $O(n^3 m)$. In our experiments, the image size is typically 256×256 with 36 discrete orientations. This represents an effective speedup on the order of one thousand times when compared against the previous algorithm. We note that since the new algorithm is local and parallel, the communication overhead on a SIMD computer would be negligible, so that the parallel time complexity would be $O(n)$ on a machine with $O(n^2 m)$ processors. This is consistent with the estimated time required for illusory contours to form in human vision (Rubin et al., 1995).

4 Stability and Accuracy Considerations

We have presented a method of computing the stochastic completion field using repeated small-kernel convolutions. We now consider how faithfully this new method models the underlying stochastic process. First, we identify the conditions under which our finite-difference scheme converges to the underlying partial differential equation that it models. We then present a stochastic interpretation of the small-kernel method. This interpretation embodies our original (i.e., physical) intuition and applies even in the case

of discontinuous initial conditions. Finally, we compare the accuracy of our new method to our previous approach, which used a single large convolution mask.

We can use standard techniques to determine that our method will be stable, provided that the following three conditions are met:

1. $\lambda = (\sigma^2 \Delta t) / 2(\Delta \theta)^2 \leq \frac{1}{2}$.
2. $\cos \theta \frac{\Delta t}{\Delta x} \leq 1$.
3. $\sin \theta \frac{\Delta t}{\Delta y} \leq 1$.

Informally, stability is necessary to ensure that the solution of the partial differential equation at each time step incorporates all the possibly relevant information from the initial conditions and to ensure convergence to the correct solution as space and time are discretized more finely. In the experiments that we report in this article, we assume that $\Delta x = \Delta y = \Delta t = 1$. Therefore, the last two conditions will always be met. However, should we choose to discretize space more finely (to achieve greater accuracy), then we must also discretize time proportionately. Regarding the first condition, in all of our experiments we set $\Delta \theta = \pi/36$. With this level of discretization, our method will be stable only when $\sigma \leq \pi/36$. All of our experiments use values of σ much less than this limit.

We now present a stochastic interpretation of repeated small-kernel convolution. This discussion serves two purposes. First, the stochastic interpretation more directly embodies our original (physical) intuition that the prior probability distribution of completion shape can be modeled as a random walk in a lattice of discrete positions and orientations. Second, our stochastic interpretation applies even in the case of discontinuous initial conditions, where interpreting the small-kernel method as a first-order approximation to a partial differential equation breaks down.

The small-kernel method can be viewed as computing the probability distribution of the position of a particle undergoing a random walk in a lattice in $\mathbb{R}^2 \times S^1$. Suppose there is a particle in this lattice, with position (x, y, θ) , which at the first time step:

- Moves to the neighboring site in the x direction with probability $\max(0, \cos \theta)$.
- Moves to the neighboring site in the $-x$ direction with probability $\max(0, -\cos \theta)$.
- Remains at the current site with probability $(1 - |\cos \theta|)$.

See Figure 3a. Clearly the convolution defined by the Step 1 equation updates the probability distribution on the lattice consistent with this motion. In a similar fashion, the change in the probability distribution of the particle's y -position is updated in the second time step. This update is described

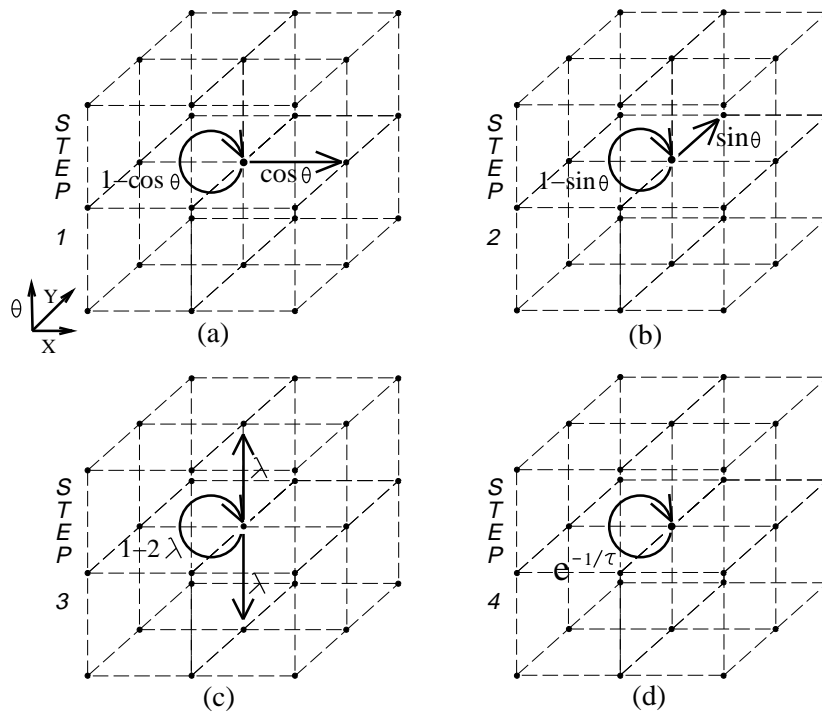


Figure 3: The small-kernel method can be interpreted as a random walk in a lattice of discrete positions and orientations. (a) Step 1: Advection in x direction (for $\cos \theta \geq 0$). (b) Step 2: Advection in y direction (for $\sin \theta \geq 0$). (c) Step 3: Diffusion in θ . (d) Step 4: Exponential decay.

by the Step 2 equation (see fig. 3b). In the third time step, the particle:

- Moves to the neighboring site in the θ direction with probability λ .
- Moves to the neighboring site in the $-\theta$ direction with probability λ .
- Remains at the current site with probability $(1 - 2\lambda)$.

See Figure 3c. The Step 3 equation updates the probability distribution on the lattice consistent with this motion. Finally, the Step 4 equation implements the exponential decay of the particle (see Figure 3d). After t iterations of these four steps, we have computed the distribution of the position of a set of particles at time t , given initial conditions that specify a distribution of their positions at time zero. This stochastic process is an exact model of

what our small-kernel method computes and is valid even when the initial conditions are discontinuous.

We now estimate the difference in magnitude between the outputs of the small- and large-kernel methods. Previously we constructed a large-kernel filter by simulating the paths of a large number of particles. Each particle starts its stochastic motion at the origin, with a direction of $\theta = 0$. Time was divided into discrete units, and at each time step, $\Delta\theta$ was drawn from a continuous, gaussian distribution. The x and y positions of the particle at each time step were represented as real numbers. The magnitude of the large-kernel filter at any given point equaled the fraction of trajectories intersecting a small volume of $x - y - \theta$ space. It is straightforward to show that the distribution of a particle's x position after t time steps is $\sum_{i=1}^t c_i$ (where $c_i = \cos \theta_i$ and $s_i = \sin \theta_i$).

We now contrast this with the small-kernel method. If we let x_i be a random variable that is 1 with probability c_i and 0 with probability $1 - c_i$, then repeated small-kernel convolution produces values in the x direction that are a probability distribution for the random variable $\sum_{i=1}^t x_i$. Since the expected value of x_i is c_i , it follows that the distribution of $\sum_{i=1}^t x_i$ has mean $\sum_{i=1}^t c_i$ and variance $\sum_{i=1}^t (x_i - c_i)^2 = \sum_{i=1}^t (c_i - c_i^2)$. This variance can range from 0 (when all $c_i = 0$ or all $c_i = 1$) to as much as $t/4$ when all c_i equal $1/2$ (when $\theta = \pi/3$). Moreover, it follows from Lindeburg's theorem that if this distribution is normalized to be of mean 0 and unit variance, it will converge to a normal distribution as $t \rightarrow \infty$. This depends on the variables satisfying the Lindeburg condition, which states roughly that each term in the sum becomes, in the limit, a vanishingly small component of the overall sum (see Feller, 1971).

Similar reasoning tells us that the distribution of a particle's position in the y direction will have mean $\sum_{i=1}^t s_i$ and variance $\sum_{i=1}^t (y_i - s_i)^2 = \sum_{i=1}^t (s_i - s_i^2)$, which can also range from 0 to $t/4$. So, for example, with $t = 100$ the standard deviation of this "spread" in the position of a particle is bounded by five pixels.

This effect is not isotropic. The above expressions show that particles traveling in directions parallel to the x and y axis will have zero variance, while those traveling in the $\pi/4$ direction will have maximum variance.⁶ Also, the source field can be computed with greater accuracy (if needed) at the cost of additional computation. Suppose that we divide each unit in space, angle, and time into d subunits and perform the same computations on these subunits. Similar calculations show that the variance in each direction will be bounded by $t/4d$ pixels (i.e., the standard deviation drops with

⁶ The lack of any detectable global phase coherence in the orientation preference structure of visual cortex (Niebur & Wörgötter, 1994) suggests that (unlike the obvious computer implementation on a rectangular grid) systematic long-range errors due to local anisotropy will be minimal in vivo.

the square root of the number of discrete units used, as might be expected). We must, however, perform td convolutions on a field of size $nd \times nd \times m$, rather than t convolutions on a field of size $n \times n \times m$, increasing the run time by a factor of d^3 .

5 Anisotropic Decay

The fact that previous computation took the form of a convolution with a large kernel was a simple consequence of the translational and rotational invariance of the random process. That is, the likelihood of a particle's taking a random walk of a given shape is assumed to be independent of the random walk's starting point and orientation. Because the new method is based on repeated small-kernel convolution, we are free to consider stochastic processes that are not the same everywhere, that is, processes that are neither homogeneous nor isotropic. There is some precedent for this in computer vision, where different methods lumped together under the term "anisotropic diffusion" are becoming increasingly popular for image enhancement (Perona & Malik, 1990; Nitzberg & Shiota, 1992; Weickert, 1995). The analogy breaks down in the sense that the quantity being "diffused" in the one case is image brightness, and in our case, the probability density of a particle's position and orientation.

Conceivably each of the two parameters defining the stochastic process (σ^2 and τ) could be modulated by local directional functions of the image brightness. Interpreted within a Bayesian framework, this could be viewed as augmenting the prior model with the additional information provided by the local image brightnesses to estimate better the parameters of the stochastic process modeling completion shape within the local neighborhood. In this article, we take only a first step in this direction by modulating the half-life of the particle by a directional function of the local image brightness. Using this anisotropic decay mechanism, we can ensure that illusory contours do not cross large brightness discontinuities.

The first requirement is that in regions where brightness is changing slowly, a particle's half-life should remain constant. This will allow illusory contours to form in these regions. The second requirement is that along edges, the half-life should be very short in directions leading away from the edge but very long in directions leading toward the edge. This will have the effect of channeling particles toward the edges. Finally, we require that the half-life should be very long in the direction tangent to the edge. This will encourage particles to continue traveling along the edge. All of these requirements can be satisfied by a function based on the directional second derivative,

$$\tau(\theta) = \begin{cases} \tau \cdot e^{D_\theta^2 I} & \text{if } D_{\theta+\pi/2}^2 I \neq 0 \text{ and } D_\theta^1 I > 0 \\ \tau \cdot e^{-D_\theta^2 I} & \text{if } D_{\theta+\pi/2}^2 I \neq 0 \text{ and } D_\theta^1 I < 0 \\ \tau + |\nabla I| & \text{otherwise,} \end{cases}$$

where $D_\theta^1 I$ and $D_\theta^2 I$ are the first and second directional derivatives of image brightness in the direction of the particle's motion and $D_{\theta+\pi/2}^2 I$ is the second directional derivative in the direction normal to the particle's motion. The condition $D_{\theta+\pi/2}^2 I \neq 0$ is simply used to decide whether the particle is traveling along an edge (a zero crossing of the second directional derivative in the direction normal to the particle's motion). If so, the half-life is taken to be $\tau + |\nabla I|$, which becomes τ in homogeneous areas as required, irrespective of whether there are zero crossings. If the particle is not traveling in the direction of an edge, then the half-life is $\tau \cdot \exp(D_\theta^2 I)$ or $\tau \cdot \exp(-D_\theta^2 I)$, depending on the sign of $D_\theta^1 I$. This expression is large or small depending on whether the particle is moving toward or away from an edge, and it evaluates to τ in homogenous areas.

6 Experimental Results

Through a sequence of pictures shown in Figure 4, we first demonstrate the propagation of two probability density "waves" emitted from a pair of sources located on a line parallel to the x -axis and with orientations equal to $\pi/6$ and $5\pi/6$. The pictures show the probability density ($p(x, y, \theta; t)$) summed over all orientations (the marginal distribution in θ). The time interval between successive pictures in the sequence, Δt , equals 5, the variance, σ^2 , equals 0.005, and the half-life, τ , equals 100. Error due to the inherent anisotropy of the finite-difference scheme is especially pronounced in the first five frames (frames 1–5) where it manifests itself as a flattening of the face of the wave along the vertical direction. This effect diminishes as the size of the wave increases, the result of greater effective resolution. In frames 11–20 the waves from the different sources can be seen to pass through each other. Because they are traveling in different "orientation planes," there is no "collision." Finally, we observe that there is a visible decrease in brightness between the first and last frame of the sequence. Although some of this decrease in brightness can be attributed to the diffusion in space, the larger part can be attributed to the exponential decay of the particles.

A second sequence of images, shown in Figure 5, depicts the formation of the completion field as a function of time. To simplify this demonstration, p.d.f.s representing the source and sink distributions at time zero were assumed to be equal. Frames 1–10 are totally dark, since insufficient time has elapsed for particles traveling from the two sources to reach each other. The completion field first appears in frame 11, at the point midway between both sources. In frames 12–20 the completion field rapidly spreads outward in both directions away from the midpoint and back toward the two sources.⁷

⁷ Assuming that the neural locus of the stochastic completion field is the population of neurons in V2 identified by von der Heydt, Peterhans, and Baumgartner (1984) (see "Stochastic Completion Fields" elsewhere in this journal), then this prediction could be verified by measuring the latency between stimulus presentation and the onset of the




Figure 4: Sequence of images depicting probability density “waves” emitted by a pair of sources located on a line parallel to the x -axis and with orientations equal to $\pi/6$ and $5\pi/6$. The pictures show the probability density $(p(x, y, \theta; t))$ summed over all orientations (the marginal distribution in θ). The time interval between successive pictures in the sequence, Δt , equals 5; the variance, σ^2 , equals 0.005; and the half-life, τ , equals 100.




Figure 5: Sequence of images depicting the formation of the completion field ($C(u, v, \phi; t)$) as a function of time. Frames 1–10 are totally dark, since insufficient time has elapsed for particles traveling from the two sources to reach each other. The completion field first appears in frame 11, at the point midway between both sources. In frames 12–20 the completion field rapidly spreads outward in both directions away from the midpoint and back toward the two sources.




Figure 6: (a) Kanizsa triangle. (b) Stochastic completion field computed by large-kernel convolution with filter generated by Monte Carlo simulation (see “Stochastic Completion Fields” elsewhere in this journal) (summed over all orientations and superimposed on the brightness gradient magnitude image). (c) Stochastic completion field computed by large-kernel convolution with filter generated by analytic expression (Thornber & Williams, 1996). (d) Stochastic completion field computed by repeated small-kernel convolution.

The third demonstration is the well-known Kanizsa triangle (Kanizsa, 1979) (see Figure 6a). The sources and sinks were identified using the method

neuron’s increased firing rate. Because the wave has to arrive from both flanking elements, the function describing this latency should have the general form, $\max(d_1, d_2)$ (where d_1 and d_2 are distances from the center of the neuron’s receptive field to each flanking element).




Figure 7: Numerical artifacts caused by anisotropy inherent in finite difference solution of the advection equation. (a) Completion field in neighborhood of “pacman” computed by large-kernel convolution. (b) Completion field in same neighborhood computed by repeated small-kernel convolution.

described in Williams and Jacobs (1997). Figure 6b shows the stochastic completion field computed using large-kernel convolution with a filter generated by Monte Carlo simulation (see our other article in this issue). The completion field is summed over all orientations and superimposed on the brightness gradient magnitude image for illustrative purposes. Figure 6c shows the same, but this result was computed with a filter generated by numerical integration of the analytic expression for G derived by Thornber and Williams (1996). This is by far the most accurate method of computing the stochastic completion field. Figure 6d shows the stochastic completion field computed using repeated small-kernel convolution. There are small but noticeable numerical artifacts due to the inherent anisotropy of the finite difference scheme for solving the advection equation (see Figure 7). The similarity between Figures 6b–d demonstrates that our characterization of the stochastic completion field is distinct from the algorithm (and neural network) that computes it.

Although space considerations preclude our showing them here, we have rerun all of the experiments from our other article in this issue using the small-kernel algorithm and achieved results of similar quality. Finally, we have tested the small-kernel method on a picture of a dinosaur. The stochastic completion field, summed over all orientations and superimposed on an edge image for illustration purposes, is shown in Figure 8. The body of the dinosaur has been completed where it is occluded by the legs.

In Figure 9 we demonstrate the effectiveness of the anisotropic decay




Figure 8: (a) Dinosaur image (*Apatosaurus*). (b) The body of the dinosaur has been completed where it is occluded by the legs.

function described in the previous section. This figure also illustrates the strong connection between the stochastic completion field and the active energy minimizing contours (“snakes”) of Kass, Witkin, and Terzopolous (1987). The stochastic completion field can be thought of as a probability distribution representing a family of snakes with given end conditions. This connection is strengthened by the inclusion of the anisotropic decay mechanism, since this device plays a role similar to the “external” energy term of the snake formulation. Figure 9a shows two probability density “waves” emitted from a pair of sources located on the dinosaur’s back. The source and sink fields were assumed equal and the time, $t = 32$. Figure 9b shows the stochastic completion field computed for this pair of sources. The mode of the stochastic completion field (the curve of least energy) lies significantly below the back of the dinosaur. Furthermore, the variance of the distribution is quite large. Figure 9c shows the probability density “waves” emitted from this same pair of sources but with particle half-life determined by the anisotropic decay function. The probability density is maximum where the wave intersects the back of the dinosaur. Figure 9d shows that the mode of the resulting completion field conforms closely to the back of the dinosaur and that the variance of the distribution is significantly reduced.

In our last experiment, we demonstrate the effect of the anisotropic decay function using three variations of the Kanizsa square figure (Ramachandran et al., 1994). Although we compare the results of our simulation against the human percepts in each case, we make no special claims for the specific anisotropic decay mechanism described in this article. Our intention was not to propose a faithful model of human visual processing but to illustrate




Figure 9: The stochastic completion field can be thought of as a probability distribution representing a family of snakes with given end conditions. (a) Probability density at $t = 32$ ($p(x, y, \theta; 32)$) for a pair of sources placed on the back of the dinosaur. This was computed by the small-kernel method without anisotropic decay. (b) The mode of the stochastic completion field (the curve of least energy) lies significantly below the back of the dinosaur. (c) Probability density at $t = 32$ for same sources computed with anisotropic decay. The probability density is maximum where the wave intersects the back of the dinosaur. (d) The mode of the resulting completion field conforms closely to the back of the dinosaur.

that the local-parallel algorithm, unlike the large-kernel algorithm, allows illusory contour formation to be modulated by local image brightnesses.

The first variation of the Kanizsa square, shown in Figure 10a, consists of four “pacmen” on a uniform white background. This figure is perceived as an illusory square occluding four black discs. The second, shown in Figure

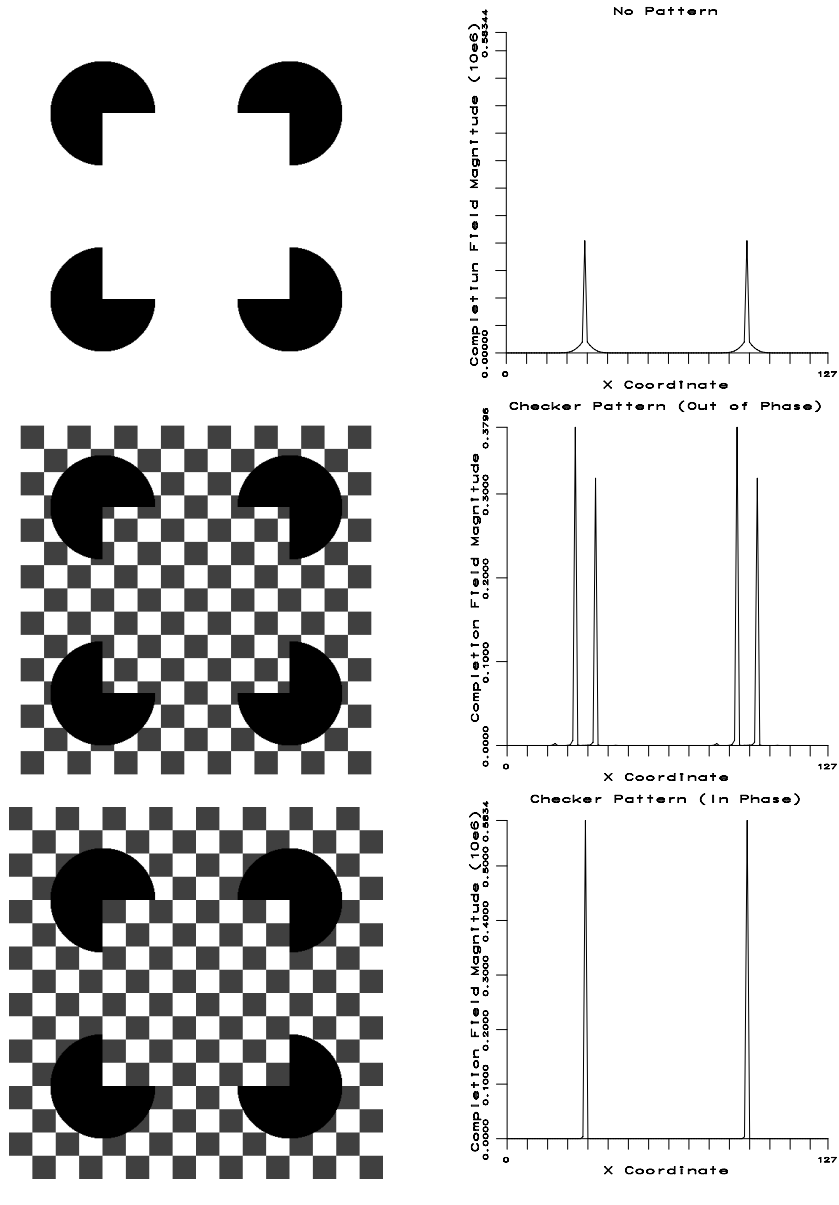


Figure 10: (a) Kanizsa square. (b) Stochastic completion field magnitude along the line $y = 64$. (c) Kanizsa square with out-of-phase checkered background (see Ramachandran et al., 1994). (d) The enhancement of the nearest contrast edge is not noticeable unless the magnitudes are multiplied by a very large factor (approximately 1.0×10^6). (e) Kanizsa square with in-phase checkered background. (f) The peak completion field magnitude is increased almost threefold, and the distribution has been significantly sharpened.

10c, is the same except that a checkered background has been added. The phase of the checkered pattern is such that illusory contours joining the “pacmen” must traverse large brightness gradients. Ramachandran et al. report not only that the illusory square is suppressed by the addition of the out-of-phase checkered background but also that the edges of the checker pattern nearest the old illusory square are enhanced. The third variation, shown in Figure 10e, also places the “pacmen” on a checkered background, but this time the phase of the checkered pattern is such that the illusory contours encounter no large brightness gradients but instead coincide with the borders of the squares. Ramachandran et al. report that the illusory square is enhanced by the addition of the in-phase checkered background.

In our simulations, the images were of size 128×128 with 12 discrete orientations. The positions and orientations of the sources and sinks (the corners of the “pacmen”) were defined manually. The diffusivity, $\sigma^2 = 0.01$ and the particle half-life, $\tau = 100$. The magnitude of the stochastic completion field along the line $y = 64$ is displayed to the right of each of the Kanizsa squares. Figure 10b and 10f are scaled by the same factor, so the results may be directly compared. We note that the peak completion field magnitude is increased almost threefold by the addition of the in-phase checkered background and that the distribution has been significantly sharpened. The effect of adding the out-of-phase checkered background is that the stochastic completion field from Figure 10b is almost totally suppressed. To a first approximation, this is consistent with Ramachandran et al.’s observation. However, the enhancement of the nearest contrast edge, also reported in Ramachandran et al. (1994), is not noticeable unless the completion field magnitudes are multiplied by a very large factor (approximately 1.0×10^6). See Figure 10d. To summarize, the results of our simulations are largely consistent with the observations reported in Ramachandran et al. (1994), but the enhancement effect observed in the case of the out-of-phase checkered background is far smaller.

7 Conclusion

We have shown how to model illusory contour formation using repeated, small-kernel convolutions. This method directly improves on several published models because of its greater efficiency (by several orders of magnitude) and because it allows edges of objects to be connected in a way that takes account of all relevant image intensities, forming illusory contours only when they are consistent with these intensities. Our work therefore helps to bridge the gap between work on understanding human performance on illusory contours and work in computer vision (e.g., snakes) on finding contours that are smooth yet faithful to the image intensities.

References

- Ames, W. (1992). *Numerical methods for partial differential equations*. San Diego: Academic Press.
- Feller, W. (1971). *An introduction to probability theory and its applications* (Vol. 2). New York: Wiley.
- Ghez, R. (1988). *A primer of diffusion problems*. New York: Wiley.
- Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, *92*, 173–211.
- Guy, G., & Medioni, G. (1996). Inferring global perceptual contours from local features. *International Journal of Computer Vision*, *20*, 113–133.
- Heitger, R., & von der Heydt, R. (1993). A computational model of neural contour processing, figure-ground and illusory contours. *Proc. of 4th Intl. Conf. on Computer Vision*, Berlin, Germany.
- Kanizsa, G. (1979). *Organization in vision*. New York: Praeger.
- Kass, M., Witkin, A., & Terzopolous, D. (1987). Snakes: Active minimum energy seeking contours. *Proc. of the First Intl. Conf. on Computer Vision*, London, 259–268.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Mumford, D. (1994). Elastica and computer vision. In C. Bajaj (Ed.), *Algebraic Geometry and Its Applications*. New York: Springer-Verlag.
- Niebur, E., & Wörgötter, F. (1994). Design principles of columnar organization in visual cortex. *Neural Computation*, *6*, 602–614.
- Nitzberg, M., & Shiota, T. (1992). Nonlinear image filtering with edge and corner enhancement. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, *14*, 826–833.
- Perona, P., & Malik, J. (1990). Scale space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, *12*, 629–639.
- Ramachandran, V. S., Ruskin, D., Cobb, S., Rogers-Ramachandran, D., & Tyler, C. W. (1994). On the perception of illusory contours. *Vision Research*, *34*(23), 3145–3152.
- Rock, I., & Anson, R. (1979). Illusory contours as the solution to a problem. *Perception*, *8*, 665–681.
- Rubin, N., Shapley, R., & Nakayama, K. (1995). Rapid propagation speed of signals triggering illusory contours. *Investigative Ophthalmology and Visual Science (ARVO)*, *36*(4).
- Thorner, K. K., & Williams, L. R. (1996). Analytic solution of stochastic completion fields. *Biological Cybernetics*, *75*, 141–151.
- von der Heydt, R., Peterhans, E., & Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science*, *224*, 1260–1262.
- Weickert, J. (1995). Multi-scale texture enhancement. *Conf. on Computer Analysis of Image and Patterns, Prague, Czech Republic*, 230–237.
- Williams, L. R., & Jacobs, D. W. (1997). Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation*, *9*, 837–858.