

An Exploration of Stochastic Completion Fields

Trần-Quân Luong

School of Computer Science
McGill University
ID no: 119731007
email: Tran-Quan.Luong@mail.mcgill.ca

December 22, 2003

abstract

We present in this paper a method originally proposed by Williams and Jacobs to determine the shape of illusory contours. The theory on stochastic completion fields is based on the assumption that the prior probability distribution of completion shapes can be modelled by the random walks of particles. Two different algorithms to compute completion fields are presented. The first is based on the convolution of a large kernel determined by Monte-Carlo simulation and the other on the integration of the Fokker-Planck equation by repeated convolutions of small kernels. Experimental results are then presented to demonstrate the effectiveness of stochastic completion fields. We conclude with a discussion that suggests when the use of completion fields may fail.

1 Introduction

A significant challenge in the field of computer vision is to infer the shape of 3D structures given the information that is provided by their projection on a 2D plane. If there are no overlapping surfaces on the image plane, then one is assured that there is a correspondence between a neighboring patch in the image and its counterpart in the 3D world. On the other hand, if a surface occludes another one, then this 1-to-1 correspondence is lost and the task of inferring shape of the occluded structures becomes much more complicated. Hence, most research in the field of “Shape-from-X” has focussed on the recovery of depth and surface orientation of visible surfaces but few attempts have been made to recover information of occluded regions.

Williams suggests in [7] that the human visual system infers geometry not only for visible surfaces, but also for occluded ones. The argument he provides is of introspective nature, but a convincing one. Imagine lifting a cue-ball on a pool table. How surprising would it be if upon picking the ball, there lay a whole of the exact same size. The surprise Williams argues, must be due to the fact that our visual system had inferred characteristics about the underlying surface. Thus, in order to grasp a full understanding of the human visual system, it is necessary to determine the boundary of occluded surfaces. As demonstrated in [6], this task is equivalent to building a Huffman-labeled figure that represents the boundaries of visible and occluded surfaces.

Central to the problem of figural completion is the necessity to determine the underlying paths that connect a set of boundary fragments. Williams and Jacobs introduced in [8] a novel approach to this problem, what they called *stochastic completion fields*. At the heart of their model is the theory that the prior probability distribution of shape boundary completions can be determined by the paths of particles that undergo a random walk in the image plane. Their belief that this approach is the correct one is reinforced by the observation that neuronal activity in receptive fields can be viewed as the probability distribution of a particle’s position. Given a subset of states P that contains the beginning points of boundary fragments and a set Q that includes the corresponding ending points, we would like to determine the probability that a particles leaving a point $p \in P$ will through the course of a random walk pass through (u, v, ϕ) before finally reaching a point $q \in Q$. This distribution over all locations (u, v) and orientation ϕ is what’s referred as the *stochastic completion field* of connected boundary fragments.

In this paper, we explore the initial formulation of stochastic completion fields and several subsequent improvements. The paper is organized as follows. We begin in section 2 with the motives that gave birth to the theory of stochastic completion fields. In section 3, its formal mathematical foundations are explained followed in section 4 by several improvements over the original formulation; these include a local approach to the computation of completion fields and the introduction of anisotropic decay. An implementation of the proposed method is then presented in section 5. Finally, we discuss the strengths and weaknesses of stochastic completion fields in section 6.

2 Motivation

Prior to the introduction of stochastic completion fields, a rich literature on the problem of figural completion already existed. Following observations of the behavior and preferences expressed by the visual cortex, several of these algorithms were based on similar assumptions such as translation and rotational invariance. Most solutions proposed can be categorized into one of three groups. The first class of strategy relates the shape of boundary completion to cocircularity [2]. The second group of solution is based on curves of least energy and the last class on the prior distribution of completion shapes modeled as random walks [3].

As it is well known, the problem of determining the boundary of occluded surfaces is ill-posed. It is impossible to determine with exactitude the location of the underlying contour since many such solutions exist. The general consensus in the literature developed so far was that given a set of hypothesized completion shapes, the visual system chose the most likely one. When introducing stochastic completion fields, Williams and Jacobs suggested that significant psychophysical properties of the shape were not preserved when using this approach. They asserted for instance that the sharpness of illusory contour is related to the variance of the distribution. Rather than selecting the completion shape with maximum likelihood, they proposed instead that the visual system derives a distribution of possible completion shapes.

3 Stochastic Completion Fields

In the initial incarnation of stochastic completion fields, the authors suggest that the distribution of boundary completion shapes can be modelled by the paths of particles that undergo a random walk in the image plane with lattices of discrete position and orientation. Using this approach similar to the one proposed by Mumford [3] provides several advantages. First, Brownian motion embodies the Gestalt principles of *proximity* and *good continuation* since the paths of particles are smooth. More importantly however is the Markov property of random walks. By the Markov assumption, it becomes possible to separate the stochastic completion field into two components, its source and sink fields.

A particle of random walk is defined as the three element vector that includes position (x, y) and orientation θ . Following each time step, the particle's position and orientation are updated such that its new components \dot{x} , \dot{y} and $\dot{\theta}$ are respectively defined as $\cos \theta$, $\sin \theta$ and $\hat{\kappa}(0, \sigma^2; t)$. The change in orientation $\dot{\theta}$ is drawn from a normally distributed random variable with zero mean and variance σ^2 (referred as the *diffusivity* factor). Additionally, a particle also possesses an assigned half-life that controls the length of its trajectory. The decay mechanism ensures that longer paths are exponentially less likely. Williams and Jacobs suggest that using a particle defined over the space $\mathbb{R}^2 \times \mathbb{S}^1$ is natural since the two dimensional surface of the visual cortex is simply a mapping of $\mathbb{R}^2 \times \mathbb{S}^1$ onto \mathbb{R}^2 .

The probability that a random walk of length t will start in state (x, y, θ) and end in (u, v, ϕ) is defined by the Green's function $G(u, v, \phi, x, y, \theta; t)$. Given that $p(x, y, \theta; 0)$ is the probability density function of the particle's initial position and orientation, the probability

density function of a particle with position (u, v) and orientation ϕ at time t is described by

$$p(u, v, \phi; t) = \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy \int_{-\pi}^{\pi} d\theta G(u, v, \phi, x, y, \theta; t) p(x, y, \theta; 0) \cdot e^{-\frac{t}{\tau}}.$$

Since random walks are invariant to rotation and translation, the probability distribution of a random path is independent of the particle's initial position and orientation. That is, any six order tensor G can be mapped to a three order tensor $G(u', v', \phi'; t)$ where the initial state is $(0, 0, 0)$. We can then restate $p(u, v, \phi; t)$ as

$$p(u, v, \phi; t) = \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy \int_{-\pi}^{\pi} d\theta G(u', v', \phi'; t) p(x, y, \theta; 0) \cdot e^{-\frac{t}{\tau}}$$

where $u' = (u - x) \cos \theta + (v - y) \sin \theta$, $v' = -(u - x) \sin \theta + (v - y) \cos \theta$, and $\phi' = \phi - \theta$.

We would now like to compute the probability that a particle will begin in a source keypoint and by its random walk pass through (u, v, ϕ) before it decays. This can be provided by what the authors call the *stochastic source field* $p'(u, v, \phi)$ and is determined by taking the integral of the position and orientation probability density function over time.

$$p'(u, v, \phi) = \int_0^{\infty} dt p(u, v, \phi; t)$$

By defining a new Green's function G' that incorporates the decay factor, it is possible to rewrite $p'(u, v, \phi)$ as the convolution of the source probability density function with G'

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

Similarly, the probability that a particle leaving a state (u, v, ϕ) will reach a source point (x, y, θ) before it decays, the *sink field* $q'(u, v, \phi)$, can be determined by

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

where $x' = (x - u) \cos \phi + (y - v) \sin \phi$, $y' = -(x - u) \sin \phi + (y - v) \cos \phi$ and $\theta' = \theta - \phi$.

The stochastic field $C(u, v, \phi)$ denoting the relative likelihood that a particle will pass through (u, v, ϕ) in its path from a source to a sink before it decays can finally be found by taking the product of the source and sink fields.

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

Provided a set P of sources and a set Q of sinks, the stochastic completion field can therefore be computed at every position in the image to reveal the contour of boundary fragments.

It was mentioned earlier that work on the problem of figural completion generally falls into one of three broad classes. The first based on cocircularity of two boundary fragments, the second on the assumption that the most likely path connecting two edges is a curve of least energy and the final category on the assumption that shape completion should be modelled as probability distribution. As can be evidently observed, the theory on stochastic completion fields falls into the latter category. The appeal of the approach proposed by Williams and Jacobs becomes much more significant however when one realises that there is a close relationship between the stochastic motion of particles and curves of least energy. It can be demonstrated¹ that the path of maximum likelihood of the distribution of completion

¹See appendix of [8].

shapes is in fact a curve of least energy!

4 Local Parallel Computation

The original formulation of stochastic completion fields by Williams and Jacobs is a novel departure from the classical approach that only seeks the most likely completion shape. Their method that makes use of the random walk of a particle to model the prior probability distribution of shape boundaries is deceptively simple when compared to the complex strategies that require numerical relaxation. The authors claim that their method depicts a plausible neural model but given the large kernel filter they use, it seems improbable. As can be observed upon inspection of the original method, in order to bridge the largest gap that might separate a source and sink point, the kernel filter has to be at least as large as the maximal distance that separates boundary fragments. In general, the size of the filter is therefore taken to be the size of the entire image. Early research by Hubel and Wiesel [1] strongly suggests that the visual system works in a highly local fashion. Hence, a method as proposed by Williams and Jacobs that necessitates a pair of spatially distant neurons to be interconnected appears unlikely.

In order to account for this discrepancy and improve in significant ways their original model of stochastic completion fields, Williams and Jacobs introduced in [9] a method based on local parallel computation². Based on the same assumption that the prior probability distribution of completion shape can be modelled by the random walk of a particle in a lattice of discrete positions and orientations, the new approach no longer relies on the convolution of a large kernel filter. Instead, the local parallel method computes the source and sink fields by integrating the Fokker-Planck equation for the stochastic motion of a particle. This new approach improves over the previous method in several ways. The most significant of which is that one can now model the distribution of completion shape using stochastic processes that are neither homogenous nor isotropic. Furthermore, the local parallel approach seems more plausible from a neuroanatomy standpoint since the network only needs to be connected within neighboring regions. Finally, an important speed improvement is gained from the local computation approach.

4.1 Fokker-Planck Equation

In the local parallel method, the basic strategy to compute the completion fields remains essentially the same. As before, we must determine the source and sink fields. Once these values have been computed over the entire image, the completion field can be determined by taking the product of the two vector-fields. The novelty proposed in this new method is in the introduction of the Fokker-Planck equation which is now integrated to obtain the source and sink fields. The Fokker-Planck equation describes the time evolution of the probability density function of a travelling particle. Its function allows us to tell the distribution of

²Interestingly, the original paper and the subsequent one were both presented back to back in the same issue of *Neural Computation*. One has to wonder why the first paper was even presented if an improved formulation of stochastic completion fields already existed. Still, it remains interesting to compare the two methods which are quite different from one another.

particles at a given location and velocity at any instant in time. This new approach is a significant improvement over the previous method that evaluated the Brownian motion of particles by Monte-Carlo simulation, a computationally expensive solution. Assuming a velocity of constant unit speed and the particle's initial position at $t = 0$ is known, the Fokker-Planck equation describes the probability density function of a particle's position at time t' as follows:

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

and

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

where P is $p(x, y, \theta; t)$, the probability that a particle is at position (x, y) with orientation θ at time t . As described, the Fokker-Planck equation can be separated into four terms that clearly demonstrate the motion of a particle. Combined together, the first two terms can be referred as the *advection equation*:

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

One can imagine that the first term represents the advection in the x direction while the second term advection in the y direction. If time interval Δt is sufficiently small, solutions to this equation have the form:

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

Thus, given a particle at location (x, y) at time t , at time $t + \Delta t$ the particle will have been displaced by a vector $(\cos \theta \Delta t, \sin \theta \Delta t)$. The third term of the Fokker-Planck equation is the classical *diffusion equation*:

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

which is typically found through Gaussian convolution. Since the Gaussian is not defined for angular quantities, Williams and Jacobs use the following approximation:

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

Finally, as in the original method, exponential decay is introduced to exponentially decrease the likelihood of long paths:

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

In effect, for a fixed orientation θ , the advection term of the Fokker-Planck equation can be thought of as representing the probability that a particle is at position (x, y) at time t . If we simultaneously maintain a set of m such equations (for each orientation θ) and let them interact through the diffusion process, we obtain a model that describes the probability density function of a particle's position and orientation evolving in time.

4.2 Fractional Computation

Evaluating the Fokker-Planck equation requires an initially smooth condition such that it is possible to determine the necessary partial derivatives. This assumption is unfortunately not held in the discrete case where the probability density function that describes the particle's position consists initially of impulses located at keypoints on the image. A standard approach to solve partial differential equations in a discrete image is to evaluate the partial derivatives at a grid point as a functions of its neighbors. By taking the Taylor series expansion and neglecting higher order terms, we can obtain a first-order approximation. The fractional method next presented allows us to iteratively compute the Fokker-Planck equation in four time steps:

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

where $\lambda = \sigma^2/2(\Delta\theta)^2$. As can be observed, the value of $p(x, y, \theta; t + 1)$ can be computed by successively evaluating each of the four term of the Fokker-Planck equation, building every time on the previous step. The reader will also observe that the first and second step vary according to the sign of $\cos \theta$ and $\sin \theta$. This is necessary because the local differencing scheme used to approximate the partial derivatives requires that we only consider locations where the probability distribution function has already been assigned a numerical value. This precaution, known as *upwind differencing*, is necessary to ensure numerical stability of the solution. Furthermore, the following conditions must be met in order to guarantee convergence of the finite-difference scheme:

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

In the experiments performed in this paper, $\Delta t = \Delta x = \Delta y = 1$ which satisfies the last two conditions. Orientation is discretized with a value of $\Delta\theta = \frac{2\pi}{16}$ and so for all values of σ^2 used, the first condition is also respected.

An important advantage brought by the evaluation of the Fokker-Planck equation by fractional computation is that the first three steps can be conceived as small kernel convolutions. Unlike the prohibitively large kernel of the previous method, the new kernels used in

the local parallel method are only of size 3×1 . In the first step where $\cos \theta > 0$, the kernel can be derived as follows:

$$\begin{aligned}
p_{x,y,\theta}^{t+1/4} &= p_{x,y,\theta}^t - \cos \theta (p_{x,y,\theta}^t - p_{x-1,y,\theta}^t) \\
&= p_{x,y,\theta}^t - \cos \theta \cdot p_{x,y,\theta}^t + \cos \theta \cdot p_{x-1,y,\theta}^t \\
&= (1 - \cos \theta) p_{x,y,\theta}^t + (\cos \theta) p_{x-1,y,\theta}^t \\
&= (\cos \theta) p_{x-1,y,\theta}^t + (1 - \cos \theta) p_{x,y,\theta}^t + (0) p_{x+1,y,\theta}^t \\
&= K * p_{x,y,\theta}^t
\end{aligned}$$

where the kernel K is defined as $[\cos \theta, 1 - \cos \theta, 0]$. In a similar fashion, kernels for the other spatial directions can easily be derived. Imagine a particle at position (x, y) with orientation θ at time t . At the next time step, the probabilities of its motion are stated as follows:

- With probability $\max(0, \cos \theta)$, the particle will move to the neighboring pixel to its right.
- With probability $\max(0, -\cos \theta)$, the particle will move to the neighboring pixel to its left.
- With probability $(1 - |\cos \theta|)$ it will remain at its current location.

The motion of the particle can be computed similarly in the y direction. As can be seen, the convolution kernels, update the probability distributions at each iteration. The diffusion of the particle to neighboring orientations can be described as follows:

- With probability λ , the particle increases its orientation by $\Delta\theta$.
- With probability λ , the particle decreases its orientation by $\Delta\theta$.
- With probability $(1 - 2\lambda)$, the particle maintains its orientation.

Finally, to obtain the desired source field $p'(x, y, \theta)$, we integrate the probability density function of the particle's position over time:

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

In the discrete domain, the integration can be approximated by the summation of $p(x, y, \theta; t)$ starting at $t = 0$ up to a fixed value t' :

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

using the recurrence relation

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

The sink field q' is computed in an analogous fashion. As performed in the original method, the stochastic completion field $C(x, y, \theta)$ is determined by taking the product of the source field p' and sink field q' .

4.3 Anisotropic Decay

It is suggested in [5] that the illusory contours perceived on a Kanizsa square can be greatly enhanced when superimposed over a checkerboard pattern. If the squares of the checkerboard are “in-phase” with the edges suggested by the pacmen (i.e. no large brightness gradients are crossed), then the appearance of illusory contour becomes more vivid. On the other hand, if the squares of the checkerboard are misaligned, then they will interfere with the appearance of the subjective contour. There is therefore strong evidence that the presence of textures can greatly enhance the path of boundary completions. As witnessed, the completion shape depends not only on the location of connected boundary fragments but also on the patterns of image brightness across the gap. It should therefore prove desirable when modeling the motion of particles to consider image intensities along their paths.

In the original method by Williams and Jacobs, the large kernel was derived based on the translational and rotational invariance of the stochastic motion. Using a single kernel simplified the task of determining the source and sink fields but the process was purely isotropic. In other words, what permitted this simplification also prohibited the use of local information along the path of particles. In the new local parallel approach based on repeated convolution of small kernels, it becomes possible to modulate the stochastic process according to image brightness. This anisotropic procedure is similar to the one proposed by Perona and Malik [4] in which the diffusion process varies locally to favor smoothing within a region but not across it. However, whereas their diffusion method works on image intensities, the anisotropic process utilised in our case is applied to the probability density of a particle’s position and orientation.

An interesting modification one can apply to the diffusion process is to vary the particle’s half-life according to image brightness. The anisotropic decay we seek should favor the motion of particles that move towards an edge and particularly those that follow the path tangent to the edge. The half-life of particles that move towards an edge should therefore be increased whereas the half-life of particles that move away from the edge should be shortened. Finally, where brightness varies slowly, the decay should remain constant. The desired behavior can be modelled according to the following function:

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

where D_{θ}^1 and D_{θ}^2 represent the first and second derivatives of image intensity taken in direction θ . When a particle is travelling exactly along an edge (case 3), then the second derivative of the direction normal to the particle’s motion will exhibit a zero-crossing (i.e. $D_{\theta+\pi/2}^2 = 0$) and the half-life τ is maintained. On the other hand, if there are no zero-crossings (case 1 and 2), then the particle is straying away from an edge. Depending on whether it moves closer or away, the half-life increases or decreases accordingly.

It should be noted that although Williams and Jacobs only propose methods to modulate the decay, similar steps could be taken to modify the diffusivity term σ^2 according to local characteristics of the image.

5 Experimental Results

We demonstrate in this section experimental results obtained using stochastic completion fields as proposed by Williams and Jacobs. In order to validate the effectiveness of the local parallel method elaborated in [9], we have recreated selected experiments the authors have performed. In all the simulations that follow, the image consisted of a 100 x 100 grid space where 16 distinct directions were selected (i.e. $\Delta\theta = 2\pi/16$).

The first sequence of images in Figure 1 shows the evolution in time of two probability density “waves” (as the authors call them) summed over all orientations. These correspond in our mathematical derivation to $p(x, y, \theta; t)$ and $q(x, y, \theta; t)$. A source point was placed at (25, 25) with direction $\frac{\pi}{4}$ and a sink point at (75, 25) with direction $\frac{3\pi}{4}$. As selected by Williams and Jacobs, the diffusivity of particles was taken to be $\sigma^2 = 0.005$ with a half-life $\tau = 100$. Each successive picture represents an evolution of 10 time steps over the previous frame. An effect that is unfortunately not apparent in this sequence is the decrease in brightness that should occur as time progresses due to the effect of decay and to a lesser extend of spatial diffusion. Thus, although not visible, there is a decrease by three orders of magnitude in the maximal probability between the first and last frame. Moreover, although Williams and Jacobs do not mention it in their paper, the image intensities had to be clamped to only show the highest probabilities. Finally, although the waves appear to collide with one another, they do not interact since each is computed independently.

Using results just presented, we can compute the stochastic completion field by taking the product of the source and sink fields. Results are shown in the sequence of Figure 2. The first three frames are empty since the probability density waves of Figure 1 have not had time to meet yet. In the fourth frame however, one can see the formation of the completion shape midway between the source and sink points. As time progresses, the completion field propagates back to the originating keypoints as expected. Figure 3 shows two more completion fields to demonstrate the effectiveness of the local parallel method. The image to the left demonstrates a distribution that possesses an inflection point and the one to the right shows the subjective contour perceived on a Kanizsa triangle (pacmen were left out for clarity).

Due to space constraints imposed on this paper, we did not discuss a significant failing of the finite differencing scheme used to approximate the advection term of the Fokker-Planck equation. Because of the evaluation of the advection equation on a rectangular grid, an anisotropic smoothing is unfortunately introduced. This effect is most noticeable in the horizontal and vertical components. Figure 4 shows (left) the initial location of impulses placed to simulate an 8-stick Ehrenstein figure and the result produced (middle) with the non-isotropic error. As can be observed, the problem is quite severe and significantly flattens the completion shape. Fortunately, a solution is proposed in [10] by introducing additional terms which render the smoothing isotropic. Necessary modifications were implemented and results are presented in Figure 4 (right). As can be seen, the completion shape no longer exhibits a flattening in the horizontal and vertical direction and produces much more satisfying results.

It should be noted that the local parallel computation of stochastic completion fields is vastly more rapid than the original method by large kernel convolution. Whereas the previous algorithm had a running time of $O(n^4m^2)$, the new method runs in $O(n^3m)$ (for




Figure 1: Propagation of probability density waves. Pictures show the probability density $p(x, y, \theta; t)$ and $q(x, y, \theta; t)$ summed over all orientations. Time interval between each frame is 10 time steps. Diffusivity $\sigma^2 = 0.005$ and half-life $\tau = 100$.

an $n \times n$ image with m discrete orientations). The speed improvements are quite significant and the new approach can compute the completion fields of an image within minutes. As currently implemented however, the algorithm's performance drastically decreases when isotropic smoothing and anisotropic decay are activated.

6 Discussion

As witnessed in the previous section, the model proposed by Williams *et al* to the figural completion problem appears quite effective. We have demonstrated that stochastic completion fields can easily be used to determine subjective contours as perceived in the Kanizsa or Ehrenstein figure. The local parallel method that uses repeated small kernel convolution offers several advantages over the original formulation, including a more plausible neural model, anisotropic decay and significant speed improvement. Furthermore, as suggested by the authors, the diffusivity σ^2 can easily be modified to change the perceived sharpness of completion boundaries. This flexibility however, leads to a natural question. Given an im-




Figure 2: Formation of completion fields by product of source and sink fields. Time interval between each frame is 10 time steps.




Figure 3: Diffusivity $\sigma^2 = 0.01$ and half-life $\tau = 100$. Left: Completion field with an inflection point. Right: Kanizsa triangle without pacmen figures.

age, how exactly should diffusivity be set? Similar inquiries may be made on the half-life of particles. Psychological experiments have shown that when a Kanizsa triangle is uniformly scaled, the time of formation of illusory contour remains constant. According to the local parallel computation method, the time required to join two boundary fragments is




Figure 4: Left: Initial placement of impulses to simulate an 8-stick Ehrenstein figure. Diffusivity $\sigma^2 = 0.01$ and half-life $\tau = 100$. Middle: Completion field with anisotropic smoothing error. Right: Corrected completion field with isotropic smoothing.

directly proportional to the distance that separates the keypoints. There is thus a discrepancy between results obtained through psychological experiments and the method proposed by Williams. It therefore remains uncertain whether stochastic completion fields truthfully model the way the human visual system detects illusory contours.

Determining the correct location and orientation of keypoints in an image is crucial to the method proposed by Williams *et al.* Unfortunately, the authors spend little time discussing this critical issue short of saying that they use a steerable one-sided filter. In practice, we suspect that they probably perform this operation manually. In almost all cases, if source and sink points are not placed correctly, then the wrong completion shape will be revealed. Unless strict rules are set to determine admissible keypoints, it unfortunately becomes possible to create any illusory contour. What greatly complicates this task is the fact that different people perceive different contours. Although most “see” a circle when presented with the four stick Ehrenstein figure, certain perceive a square. Another manifestation of this problem is noticeable with Koffka crosses of increasing width (Figure 5) where the shape of the illusory contour changes from a circle to a square. The appearance of corners also seems problematic since the stochastic completion fields model (as currently defined) cannot account for sharp contours. Finally, in the Ehrenstein figure for instance, even if a stick is removed, most would still perceive a circle. Using completion fields however, even assuming perfect placement of source and sink points on the three remaining sticks, the emerging boundary shape will not be circular. This is by the nature of completion fields that usually attempt to create closed form figures. Put in other words, the diffusion process relies heavily on the principle of *closure*, possibly too much.

It is also important to keep in mind that the initial intent of completion fields is to determine the boundaries of occluded surfaces. Although optical illusions such as the Kanizsa triangle are used to demonstrate subjective contour, these figures only happen to be extreme cases where the occluding surface has the same intensity as its background. The problem of figural completion must therefore be solved at several levels according to the depth of the figure for which we would like to retrieve the boundaries. That is, each surface must have its own source and sink points that must not interact with keypoints of other figures. This is a subtlety that is never mentioned by Williams and might lead one to wrongfully believe that the diffusion process can be applied uniformly over an image. Figure 6 (left) for example




Figure 5: Koffka crosses with increasing arm width. The perceived completion shape changes from a circle to a square.

is an illustration of two triangles superimposed one over the other. In order to obtain the correct boundaries of each triangle (middle), source and sink points of the triangle pointing up should not form completion fields with the triangle pointing down. When this constraint is not respected, incorrect boundaries (right) become visible. It remains unclear however how one would determine which keypoint belongs to which figure. Even assuming L, T, X and Y-junctions could be identified correctly, how can corners due to natural discontinuities be told apart from corners created by occlusions? To detect boundaries of occluded surfaces, only the latter type is of interest to us. Finally, assuming an ideal scenario where source and sink points are perfectly detected and keypoints are correctly assigned to their respective figure, the theoretical model at hand dictates an increase in computational cost as more figures are added. Through introspective observation, this does not seem to correspond to the way our visual system functions. Adding more figures does not appear to slow down the speed at which one can determine the boundary of occluded surfaces.




Figure 6: Left: Two superimposed triangles. Middle: Desired completion shape obtained when diffusion process occurs independently at each depth level. Right: Incorrect completion shape obtained when source and sink points of the two triangles interact.

7 Conclusion

We have presented in this paper a method originally introduced by Williams and Jacobs to determine the shape of illusory contours. The theory on stochastic completion fields is based on the assumption that the prior probability distribution of completion shapes can be

modelled by the random walks of particles. Two different algorithms to compute the completion fields were presented one based on the convolution of a large kernel determined by Monte-Carlo simulation and the other based on the integration of the Fokker-Planck equation via repeated small kernel convolutions. An implementation of the proposed method was presented with results that convincingly demonstrate the effectiveness of stochastic completion fields. Key issues unfortunately not addressed by the authors were finally discussed to suggest when the use of completion fields may fail.

References

- [1] D.H. Hubel and T.N. Wiesel. “Brain Mechanisms of Vision”. *Scientific American*, 241, pp. 150-162, 1979.
- [2] P. Parent and S.W. Zucker. “Trace Inference, Curvature Consistency, and Curve Detection”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8)823–839, 1989
- [3] D. Mumford, “Elastica and Computer Vision”. *Algebraic Geometry and Its Applications*, Chandrajit Bajaj (ed.), Springer-Verlag, New York, 1994.
- [4] P. Perona and J. Malik, “Scale Space and Edge Detection Using Anisotropic Diffusion”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7), pp. 629-639, 1990.
- [5] V.S. Ramachandran, D. Ruskin, S. Cobb, D. Rogers-Ramachandran and C. W. Tyler, “On the Perception of Illusory Contours”, *Vision Research*, 34(23), pp. 3145-3152, 1994.
- [6] L.R. Williams, “Perceptual Organization of Occluded Surfaces”. Ph.D. Dissertation, Dept. of Computer Science, University of Massachusetts at Amherst, Amherst, Mass., 1994.
- [7] L.R. Williams and A.R. Hanson, “Perceptual Completion of Occluded Surfaces”, *Computer Vision and Image Understanding*, 64, No. 1, pp. 1-20, 1996.
- [8] L.R. Williams and D.W. Jacobs, “Stochastic Completion Fields: A Neural Model of Illusory Contour Shape and Saliency”. *Neural Computation*, 9(4), pp. 837-858, 1997.
- [9] L.R. Williams and D.W. Jacobs, “Local Parallel Computation of Stochastic Completion Fields”. *Neural Computation*, 9(4), pp. 859-881, 1997.
- [10] L.R. Williams, J.W. Zweck, T. Wang and K.K. Thornber, “Computing Stochastic Completion Fields in Linear-Time Using a Resolution Pyramid”, *Computer Vision and Image Understanding*, 76(3), pp. 289-297, 1999.
- [11] L.R. Williams and K.K. Thornber, “A Comparison of Measures for Detecting Natural Shapes in Cluttered Backgrounds”, *Intl. Journal of Computer Vision*, 34 (2/3), pp. 81-96, 1999.