Computing Stochastic Completion Fields in Linear-Time Using a Resolution Pyramid

Lance Williams* John Zweck**
Tairan Wang† Karvel Thornber‡

Apr 28 1999

*Dept. of Computer Science
University of New Mexico
Albuquerque, NM 87131

†Dept. of Physics
M.I.T.
Cambridge, MA 02139

‡NEC Research Institute
4 Independence Way
Princeton, NJ 08540

Abstract

We describe a linear-time algorithm for computing the likelihood that a completion joining two contour fragments passes through any given position and orientation in the image plane. Our algorithm is a resolution pyramid based method for solving a partial differential equation (PDE) characterizing a distribution of short, smooth completion shapes. The PDE consists of a set of independent advection equations in (x, y) coupled in the θ dimension by the diffusion equation. A previously described algorithm used a first-order, explicit finite difference scheme implemented on a rectangular grid. This algorithm required $O(n^3m)$ time for a grid of size $n \times n$ with m discrete orientations. Unfortunately, systematic error in solving the advection equations produced unwanted anisotropic smoothing in the (x, y) dimension. This resulted in visible artifacts in the completion fields. The amount of error and its dependence on θ has been previously characterized. We observe that by careful addition of extra spatial smoothing, the error can be made totally isotropic. The combined effect of this error and of intrinsic smoothness due to diffusion in the θ dimension is that the solution becomes smoother with increasing time, i.e., the high spatial frequencies drop out. By increasing Δx and Δt on a regular schedule, and using a second-order, implicit scheme for the diffusion term, it is possible to solve the modified PDE in $O(n^2m)$ time, i.e., time linear in the problem size. Using current hardware and for problems of typical size, this means that a solution which previously took one hour to compute can now be computed in about two minutes.

1 Introduction

The problem of computing the shape of the contour joining a pair of boundary fragments was first examined by Ullman in 1976. It has since become conventional wisdom that the

*J.Z. was supported (in part) by the Albuquerque High Performance Computing Center.

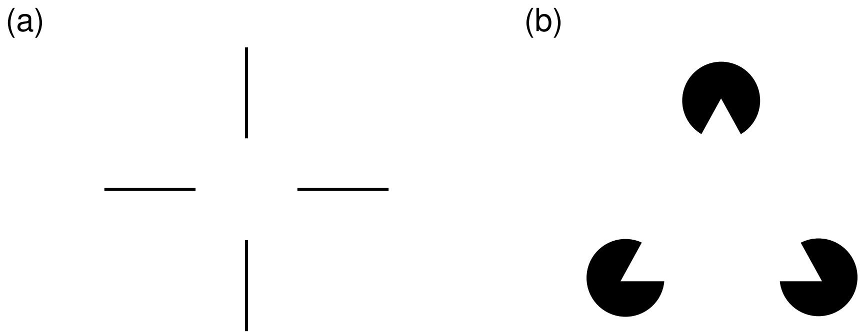

Figure 1: **(a)** Ehrenstein Figure. **(b)** Kanizsa Triangle.

shape of this contour is described by the curve of least energy, i.e., the curve which minimizes a functional of the form, $E = \int_\Gamma (\alpha\kappa^2(s) + \beta)\, ds$ (see Horn[5]). The curve of least energy can be regarded as a maximum likelihood estimate of the shape of the completion. However, humans do not experience a sharply defined, well localized illusory contour when presented with a stimulus such as the Ehrenstein figure or Kanizsa Triangle (see Figure 1). In a recent paper, Williams and Jacobs[11] argue that our perception more closely resembles the distribution of possible shapes and not simply the most likely shape. According to this view, the degree of sharpness is related to the variance of this distribution. In an earlier paper, Mumford[6] proposed that the distribution of completion shapes could be modeled as the set of paths followed by particles traveling with constant speed in directions described by Brownian motions. He shows that the maximum likelihood paths followed by such particles are curves of least energy and gives a partial differential equation (PDE) which describes the evolution in time of the probability density function describing a particle's position and direction, i.e., a *Fokker-Planck equation.* Williams and Jacobs[12] subsequently proposed a neural model[1] of illusory contour shape, salience and sharpness based on a finite difference scheme for integrating this PDE. Although the dynamics of this model are consistent with known human visual psychophysics, the algorithm is fairly slow, requiring over an hour for a problem of typical size on a modern workstation. Being able to solve this PDE efficiently will allow it to be applied profitably to problems in computer vision, e.g., to the problem of identifying smooth, closed shapes amidst background clutter.

[1]Other models of illusory contour formation are described by Grossberg and Mingolla[2], Guy and Medioni[3], Heitger and von der Heydt[4] and Geiger, Kumaran and Parida[1].

2 Stochastic Completion Fields

Given 1) a set of position and orientation constraints (also known as *keypoints*)[2] representing the beginning and ending points of a set of contour fragments; and 2) a probability distribution of completion shapes; the magnitude of the *stochastic completion field* at (x, y, θ) is the probability that a completion from this distribution will pass through (x, y, θ) on a path joining two of the contour fragments. The probability distribution of completion shapes is modeled as the set of paths followed by particles traveling with constant speed in directions described by Brownian motions. Williams and Jacobs[11] showed that the stochastic completion field could be factored into a *source field* and a *sink field*. The source field, $p'(x, y, \theta)$, represents the probability that a contour beginning at a keypoint (x_p, y_p, θ_p) will pass through (x, y, θ) and the sink field, $q'(x, y, \theta)$, represents the probability that a contour beginning at (x, y, θ) will reach a keypoint (x_q, y_q, θ_q).

Given a probability density function describing a set of particles' positions and directions at time zero, $p(x, y, \theta\,; 0)$, the probability density function describing their positions and directions at time t is computed by integrating the Fokker-Planck equation described by Mumford[6]:

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

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

where $P = p(x, y, \theta\,;\, t)$. This PDE can be viewed as a set of independent *advection* equations in (x, y) (the first and second terms) coupled in the θ dimension by the *diffusion* equation (the third term). The effect of the advection equations is to translate probability mass in the θ direction with unit speed. The diffusion term models the Brownian motion in direction. Finally, the effect of the fourth term is that a constant fraction of particles decay per unit time. This represents our prior expectation on the length of gaps—most are quite short.

[2]We adopt this term used in [4]. A keypoint can represent either of the two orientations at a corner or the normal orientation at a line termination. These are points where it is likely that one surface occludes the boundary of another.

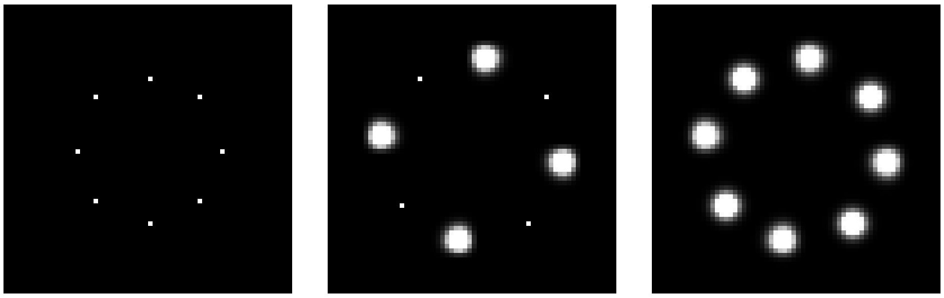

Figure 2: Demonstration of the anisotropic nature of the advection error and its correction. Left: The initial condition consists of eight impulses uniformly spaced around the circumference of a circle and in directions tangent to the circle. For illustration purposes, there is no diffusion in θ and no decay. After 12 time-steps, the mass traveling in the $45°$, $135°$, $225°$ and $315°$ directions is noticeably dispersed. The mass traveling in the $0°$, $90°$, $180°$ and $270°$ directions remains concentrated. Right: Advection error after correction.

The algorithm for computing the source field described by Williams and Jacobs is based on a first order, explicit scheme for integrating the Fokker-Planck equation:

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

where $\lambda = \sigma^2/2(\Delta\theta)^2 < 1/2$, and $\Delta t = \Delta x = \Delta y = 1$. The four steps correspond to the four terms of the PDE. The first two steps employ *upwind differencing* to ensure stability (see [7]). The third step is stable when $\lambda < 0.5$ and the fourth step is unconditionally stable. By repeating the above four steps, $p(x, y, \theta; t)$ can be computed for increasing values of t. The source field, $p'(x, y, \theta) = \int_0^\infty dt\, p(x, y, \theta\,;\, t)$, is computed using the recurrence equation

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

Since the amount of remaining probability mass, $\int\int\int dx\, dy\, d\theta\, p(x, y, \theta; t) = e^{-t/\tau}$, is typically sufficiently small for $t > n$, the time complexity of this method is $O(n^3 m)$.

Due to the way in which the advection equations are finite-differenced on a rectangular grid, the above method introduces additional non-isotropic, spatial smoothing (see Figure 2). It is straightforward to show that after one time-step, the expected values and variances of a particle's position (with respect to its previous position) are given by

$$< x(\theta) >= \cos\theta, \qquad < y(\theta) >= \sin\theta$$

$$\sigma^2_{xerr}(\theta) \;=\; |\cos\theta|\;(1-|\cos\theta|), \qquad \sigma^2_{yerr}(\theta) \;=\; |\sin\theta|\;(1-|\sin\theta|).$$

Note that this error is highly non-isotropic—σ^2_{xerr} has a minimum value of 0 when $\theta = 0°$ and a maximum value of 0.25 when $\theta = 60°$. This means that the PDE which Williams and Jacobs actually solve more closely resembles

$$\frac{\partial P}{\partial t} = -\cos\theta\frac{\partial P}{\partial x} - \sin\theta\frac{\partial P}{\partial y} + \frac{\sigma^2}{2}\frac{\partial^2 P}{\partial\theta^2} - \frac{1}{\tau}P + \frac{\sigma^2_{xerr}(\theta)}{2}\frac{\partial^2 P}{\partial x^2} + \frac{\sigma^2_{yerr}(\theta)}{2}\frac{\partial^2 P}{\partial y^2}.$$

This non-isotropic advection error in the source and sink fields leads to visible artifacts in the completion fields (see Figures 4 and 5 below).

3 A Linear-Time Method

The basic idea underlying our new algorithm is to selectively increase spatial smoothing to make the advection error isotropic and then to increase Δt and Δx on a regular schedule as the high spatial frequencies drop out of the solution. Undersampling in this manner will lead to a linear-time algorithm.

We begin by describing how the PDE and finite-differencing scheme are modified to make the advection error isotropic. Then we describe the undersampling procedure used to pass the probability density function (p.d.f.) up the resolution pyramid from a finer to a coarser grid, and the interpolation procedure used to recover the source field on the finer grid from that on the coarser grid. The question of how many time steps one should evolve the p.d.f. on a fine grid before passing it up to a coarser grid is answered, with the aid of the Shannon-Whittaker Sampling Theorem, in Section 4 below.

The modified PDE is

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

where $\sigma_x^2 = \sigma_{xerr}^2(\theta) + \sigma_{xcorr}^2(\theta) = \frac{1}{4}$ and $\sigma_y^2 = \sigma_{yerr}^2(\theta) + \sigma_{ycorr}^2(\theta) = \frac{1}{4}$. The variances, $\sigma_{xcorr}^2(\theta)$ and $\sigma_{ycorr}^2(\theta)$, are the correction factors needed to make the advection error isotropic. Their values are given by the expressions

$$\sigma_{xcorr}^2(\theta) = \frac{1}{4} - |\cos\theta|(1 - |\cos\theta|), \quad \sigma_{ycorr}^2(\theta) = \frac{1}{4} - |\sin\theta|(1 - |\sin\theta|),$$

which were derived by linear interpolation of the variances for the maximum and minimum error directions. The modified PDE can be solved using the following finite-differencing scheme:

$$\text{STEP 1:} \quad p_{x,y,\theta}^{t+1/6} = p_{x,y,\theta}^{t} - \cos\theta \cdot \begin{cases} p_{x,y,\theta}^{t} - p_{x-\Delta x,y,\theta}^{t} & \text{if } \cos\theta > 0 \\ p_{x+\Delta x,y,\theta}^{t} - p_{x,y,\theta}^{t} & \text{if } \cos\theta < 0 \end{cases}$$

$$\text{STEP 2:} \quad p_{x,y,\theta}^{t+2/6} = p_{x,y,\theta}^{t+1/6} - \sin\theta \cdot \begin{cases} p_{x,y,\theta}^{t+1/6} - p_{x,y-\Delta y,\theta}^{t+1/6} & \text{if } \sin\theta > 0 \\ p_{x,y+\Delta y,\theta}^{t+1/6} - p_{x,y,\theta}^{t+1/6} & \text{if } \sin\theta < 0 \end{cases}$$

$$\text{STEP 3:} \quad p_{x,y,\theta}^{t+3/6} = \lambda_x \, p_{x-\Delta x,y,\theta}^{t+2/6} + (1 - 2\lambda_x)\, p_{x,y,\theta}^{t+2/6} + \lambda_x \, p_{x+\Delta x,y,\theta}^{t+2/6}$$

$$\text{STEP 4:} \quad p_{x,y,\theta}^{t+4/6} = \lambda_y \, p_{x,y-\Delta y,\theta}^{t+3/6} + (1 - 2\lambda_y)\, p_{x,y,\theta}^{t+3/6} + \lambda_y \, p_{x,y+\Delta y,\theta}^{t+3/6}$$

$$\text{STEP 5:} \quad p_{x,y,\theta}^{t+5/6} = \lambda \, p_{x,y,\theta-\Delta\theta}^{t+4/6} + (1 - 2\lambda)\, p_{x,y,\theta}^{t+4/6} + \lambda \, p_{x,y,\theta+\Delta\theta}^{t+4/6} \quad \text{if } \lambda < \tfrac{1}{2},$$

$$-\lambda \, p_{x,y,\theta+\Delta\theta}^{t+5/6} + (1 + 2\lambda)\, p_{x,y,\theta}^{t+5/6} - \lambda \, p_{x,y,\theta-\Delta\theta}^{t+5/6} =$$

$$\lambda \, p_{x,y,\theta-\Delta\theta}^{t+4/6} + (1 - 2\lambda)\, p_{x,y,\theta}^{t+4/6} + \lambda \, p_{x,y,\theta+\Delta\theta}^{t+4/6} \quad \text{if } \lambda \geq \tfrac{1}{2}$$

$$\text{STEP 6:} \quad p_{x,y,\theta}^{t+6/6} = e^{-\Delta t/\tau} \cdot p_{x,y,\theta}^{t+5/6},$$

where $\lambda = \frac{1}{2}\sigma^2 \Delta t/(\Delta\theta)^2$, $\lambda_x = \frac{1}{2}\sigma_{xcorr}^2(\theta)$ and $\lambda_y = \frac{1}{2}\sigma_{ycorr}^2(\theta)$. (Note that λ_x is independent of Δt and Δx since, for a grid spacing of Δx and a time step size of Δt, the variance of a particle's position is given by $\sigma_{xerr}^2(\theta)(\Delta x)^2/\Delta t$.) As long as $\Delta x/\Delta t = \Delta y/\Delta t = 1$, the first two steps will be stable. Since our undersampling schedule described below ensures that $\Delta x = \Delta t$, this condition will always hold. Because the maximum value of $\sigma_{xcorr}^2(\theta)$ and $\sigma_{ycorr}^2(\theta)$ is 0.25, it follows that $\lambda_x < 0.5$ and $\lambda_y < 0.5$. Consequently the third and fourth steps will always be stable. More care is required to implement the θ-diffusion since Δt is increased each time we undersample, and so it is possible that λ could exceed $\frac{1}{2}$ in which case the first order explicit scheme described in Section 2 will become unstable. To overcome this problem, in the fifth step, when $\lambda \geq \frac{1}{2}$ we switch to a second order implicit scheme, i.e., the Crank-Nicholson method (see [7]), which is stable for all values of λ. Although the Crank-Nicholson method is more expensive than the explicit method, in practice this is often not an issue. For example, the completion fields in this paper are of size 256×256 with 36 discrete orientations and have a variance of order 10^{-3}. They were produced with 18 time steps per level and required a pyramid with 5 levels. So it is only on the last level or two that λ might exceed $\frac{1}{2}$.

After c time steps are performed at a given resolution the p.d.f., p, is spatially undersampled with a matching increase in the speed of evolution of p. This is accomplished by doubling Δx, Δy and Δt and by means of the *reduction operation*

$$p_{i+1}(x, y, \theta\,; 1) \;\leftarrow\; p_i(2x, 2y, \theta\,; c),$$

where i and $i+1$ are successive levels in a resolution pyramid. Using the above strategy we can efficiently simulate the evolution of the PDE for a sufficient length of time. However, to compute the source field, we must still compute the integral of the probability over all time. Within a given level, the probability can be accumulated as before, using the recurrence equation

$$p_i'(x, y, \theta) \;\leftarrow\; p_i'(x, y, \theta) + p_i(x, y, \theta\,; t).$$

After c time-steps, $p_i'(x, y, \theta)$ will hold the partial sum for the i-th level:

$$p_i'(x, y, \theta) \;=\; \sum_{t=1}^{c} p_i(x, y, \theta\,; t).$$

The only thing that remains is to combine the partial sums for each level. This is accomplished using repeated *projection operations* to push the partial sums down the resolution pyramid:
for $i = L$ downto 1 do

$$p_{i-1}'(x, y, \theta) \;\leftarrow\; p_{i-1}'(x, y, \theta) \;+\; \mathrm{Proj}(p_i')\,(x, y, \theta)$$

endo

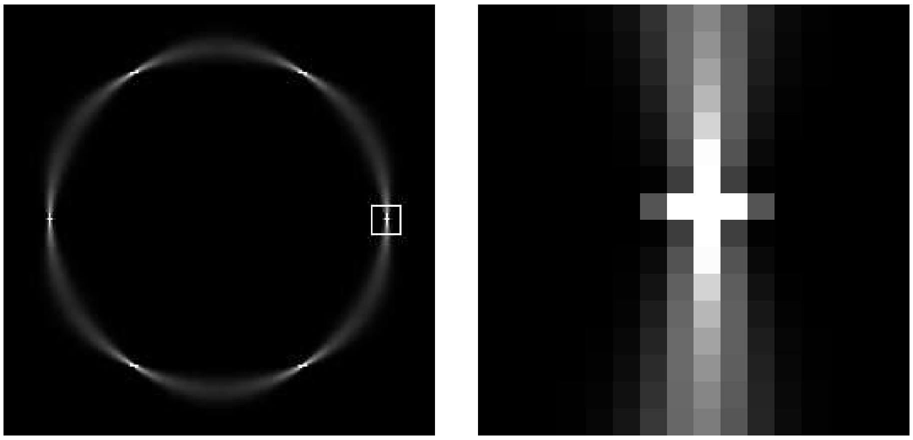

Figure 3: Demonstration of the moonwalking artifact. Left: The completion field due to a six stick Ehrenstein initial stimulus computed using the linear-time method of this paper. The moonwalking artifact is visible in a neighborhood of each of the keypoints, (c.f. Figure 5). Right: The region inside the box, magnified 16×.

where

$$\mathrm{Proj}(p'_i)\,(x,y,\theta) \;\leftarrow\; \begin{cases} p'_i(\frac{x}{2},\frac{y}{2},\theta) & \text{if } x,y \text{ even,}\\ \frac{1}{2}\{p'_i(\frac{x}{2},\frac{y-1}{2},\theta)+p'_i(\frac{x}{2},\frac{y+1}{2},\theta)\} & \text{if } x \text{ even, } y \text{ odd,}\\ \frac{1}{2}\{p'_i(\frac{x-1}{2},\frac{y}{2},\theta)+p'_i(\frac{x+1}{2},\frac{y}{2},\theta)\} & \text{if } x \text{ odd, } y \text{ even,}\\ \frac{1}{4}\{p'_i(\frac{x-1}{2},\frac{y-1}{2},\theta)+p'_i(\frac{x-1}{2},\frac{y+1}{2},\theta) & \\ \quad +\; p'_i(\frac{x+1}{2},\frac{y-1}{2},\theta)+p'_i(\frac{x+1}{2},\frac{y+1}{2},\theta)\} & \text{if } x,y \text{ odd.} \end{cases}$$

The addition of extra spatial smoothing in Steps 3 and 4 introduces an undesirable artifact in the completion field which we call "moonwalking" (so called because it is due to probability mass moving backwards), see Figure 3. This artifact is characterized by the property that near where the initial p.d.f. is non-zero the completion field is greater than expected in directions roughly perpendicular to the direction of motion. For example consider the case where the initial p.d.f. is non-zero only at $(0,0,\frac{\pi}{2})$. As expected the completion field is greatest along the y axis. However it is also quite strong in a small interval about the origin on the x axis. This is because the combined effect of Steps 3 and 4 is to move a small but significant fraction of the initial probability mass around small loops, thereby propagating mass backwards. Consequently the source and sink fields, and hence the completion field, are greater than expected on the x axis. Put succinctly the moonwalking artifact is due to the short term interaction between the source and sink fields. The simplest way to remove this artifact is to subtract the product of short term source and sink fields from the completion field. More precisely, the source field $p'(x,y,\theta)$ is decomposed as $p' = p'_S + p'_L$ where the *short term source field*, p'_S, is given by integrating the p.d.f, p, up to a time T_S, i.e. $p'_S = \int_0^{T_S} p\,dt$. Then the stochastic completion field is redefined to be

$$C \;=\; p'q' \;-\; p'_S q'_S \;=\; p'_S q'_L \;+\; p'_L q'_S \;+\; p'_L q'_L, \tag{3.1}$$

where $q' = q'_S + q'_L$ denotes the corresponding decomposition of the sink field. In the examples presented in Section 6 below we needed the short term source and sink fields to be about eight time steps long to completely remove the moonwalking artifact from the completion field.

4 An Analysis of the Resolution Pyramid Algorithm

The aim of this section is to analyze how well the multi-level resolution pyramid algorithm presented in the previous section computes the completion field by comparing it to the completion field obtained by evolving the same six step finite-differencing scheme on the bottom level of the pyramid. In particular given a tolerance level, ϵ, we will estimate the number of time steps, c, required to evolve the p.d.f., p, on each level before passing p up to the next level of the pyramid in order that the relative error between the multilevel and single level solutions is less than the tolerance level ϵ. We will show that c depends primarily on the tolerance level ϵ, i.e. it is essentially independent of the grid size and the parameters in the PDE.

To state the problem more precisely, let $p(x,y,\theta\,;t)$ be the p.d.f. obtained by evolving an initial p.d.f. on the bottom level of the pyramid for time t, and let $\mathrm{PR}\,(p)(x,y,\theta\,;t)$ be the p.d.f. on the base-level grid obtained (by interpolation) from the undersampled values of p on a grid which is twice as coarse in the spatial variables x and y, i.e. $\mathrm{PR}\,(p)(x,y,\theta\,;t)$ is the projection back to the base level of the reduction of $p(x,y,\theta\,;t)$ to the next level of the pyramid. As time increases the high frequency content of p decreases (due to the extra spatial diffusion introduced in Steps 3 and 4 of the finite-differencing scheme) and so $\mathrm{PR}\,(p)$ is a better approximation to p as time increases. The degree to which $\mathrm{PR}\,(p)$ approximates p can be estimated using the Shannon-Whittaker Sampling Theorem [8]. The condition we use to estimate the number of time steps per level, c, is that by time c, the probability density functions $\mathrm{PR}\,(p)$ and p are relatively close in the L^2-norm, in that

$$\frac{\|\mathrm{PR}\,(p) \;-\; p\|_2}{\|p\|_2} \;<\; \epsilon, \tag{4.1}$$

where ϵ is the given tolerance level and where $\|f\|_2 := \int\limits_{-\infty}^{\infty}\int\limits_{-\infty}^{\infty} |f(x,y)|^2\,dxdy$ denotes the L^2-norm in the spatial variables.

To make the calculation of c tractable we make two simplifying assumptions. Firstly we assume that the interpolation function used in the projection operation is the full sinc function rather than the small kernel approximation for it which we actually used. This assumption allows us to easily apply the Sampling Theorem. Secondly we assume that the p.d.f. $p(x,y,\theta\,;t)$ is simply a function $p(x,y\,;t)$ evolving according to the spatial diffusion equation

$$\frac{\partial P}{\partial t} \;=\; \frac{1}{8}\left(\frac{\partial^2 P}{\partial x^2} \;+\; \frac{\partial^2 P}{\partial y^2}\right), \tag{4.2}$$

that is, we ignore the advection, θ-diffusion and decay terms since, as numerical experiments confirm (see below), these terms do not increase the spatial frequency content of the p.d.f. The finite-differencing scheme used to implement equation (4.2) is

$$\begin{aligned} \textsf{STEP 3}':\quad p_{n,m}^{t+1/2} \;&=\; \tfrac{1}{8}\,p_{n-1,m}^{t} \;+\; \tfrac{3}{4}\,p_{n,m}^{t} \;+\; \tfrac{1}{8}\,p_{n+1,m}^{t} \\ \textsf{STEP 4}':\quad p_{n,m}^{t+1} \;&=\; \tfrac{1}{8}\,p_{n,m-1}^{t+1/2} \;+\; \tfrac{3}{4}\,p_{n,m}^{t+1/2} \;+\; \tfrac{1}{8}\,p_{n,m+1}^{t+1/2} \end{aligned} \tag{4.3}$$

where n and m are integers.

Given these assumptions we have the following result, which is proved in the Appendix.

Theorem 4.1 *Suppose that a p.d.f. $p(x,y\,;t)$ evolves on an $N \times N$ grid according to the finite-differencing scheme (4.3). Then the number of time steps per level, c, required for equation (4.1) to hold with relative error ϵ is determined by the inequality*

$$\frac{\sum\limits_{|n|\le N/4} \left(\frac{1}{4}\cos\frac{2\pi n}{N} \;+\; \frac{3}{4}\right)^{2c}}{\sum\limits_{n=-N/2}^{N/2-1} \left(\frac{1}{4}\cos\frac{2\pi n}{N} \;+\; \frac{3}{4}\right)^{2c}} \;>\; \sqrt{1-\frac{\epsilon^2}{2}}. \tag{4.4}$$

Table 4.1							
Relative L^2 error, ϵ	11%	8%	5 %	3%	2%	1%	0.5 %
Number of time steps per level, c	8	9	11	12	14	16	18

For large N the left hand side of this inequality is a good approximation to the ratio of two integrals and so is essentially independent of N. Table 4.1 shows the relationship between the number of time steps per level c and the reconstruction error ϵ for a grid of size $N = 256$.

In the experimental results in Section 6 we used $c = 18$ time steps per level, which corresponds to a relative L^2-error of $\epsilon = 0.5\%$. We required such a small relative error because of the simplifying assumptions made in Theorem 4.1. For instance, the theorem does not account for the fact that the projection operation was implemented with a small kernel of width three, rather than with a large kernel approximation to the sinc function. However the following numerical experiment suggests that our simplifying assumptions produce a model which approximates the frequency evolution of the p.d.f. very well. The initial p.d.f. for our experiment was a six stick Ehrenstein figure on a grid of size 256×256 with 36 discrete orientations. The p.d.f. was evolved for 16 time steps using the six step finite-differencing scheme in Section 3. Using the small width-three kernel approximation to the sinc function, we found that the relative L^2 error in reconstructing p on the fine grid from it's values on the coarse grid was 5%. We also observed that the projection of the p.d.f. from the coarse grid back to the fine grid was more spatially dispersed than the p.d.f. on the fine grid. (A wide kernel approximation to the sinc function gives a relative L^2 error of $< 1\%$. However we did not use such a kernel to compute completion fields since, if the conditions of the Sampling Theorem are not met quite accurately enough, the lack of locality in the wide kernel can result in a severe degradation of the completion field.)

Of course even if the p.d.f. on a given level could be perfectly reconstructed from its reduction to the next level, the finite-differencing scheme used to further evolve the p.d.f. on the next level would produce a different completion field than if the evolution had continued for an equivalent amount of time on the finer grid. We observed that evolving the p.d.f. on the coarser grid and then projecting back to the finer grid resulted in some additional spatial dispersion of the p.d.f., and hence of the completion field, (see Figures 6 and 7 below).

5 Time-Complexity

Given an $n \times n$ image with m discrete orientations, i.e., a problem of size n^2m, the number of levels in the resolution pyramid is:

$$L = 1 + \log_2(n/c)$$

and the cost at the i-th level to perform c time-steps is:

$$c \cdot \frac{n^2m}{4^{i-1}}$$

It follows that the total cost of running the PDE forward in time is:

$$c \cdot (1 + \frac{1}{4} + ... + \frac{1}{4^{L-1}}) \cdot n^2m$$

The total cost of pushing the partial sums for each level down the resolution pyramid is:

$$(\frac{1}{4^{L-1}} + ... + \frac{1}{4} + 1) \cdot n^2m$$

Consequently, the total amount of time required by both stages is:

$$(c+1) \cdot (1 + \frac{1}{4} + ... + \frac{1}{4^{L-1}}) \cdot n^2m < \frac{4}{3}(c+1) \cdot n^2m$$

so that the time-complexity of the new algorithm is $O(n^2m)$, i.e., linear in the problem size. In practice, we have observed that problems which previously took over an hour to finish ($n = 256$ and $m = 36$), can now be computed in about 2 minutes.

6 Experimental Results

In this section we present examples of stochastic completion fields constructed using the linear-time algorithm described in Section 3 and compare them to completion fields constructed using the algorithm of Williams and Jacobs, recalled in Section 2.

In each example the completion field is represented by an array of size $256 \times 256 \times 36$. In the figures brightness encodes the sum over all orientations. In each figure completion

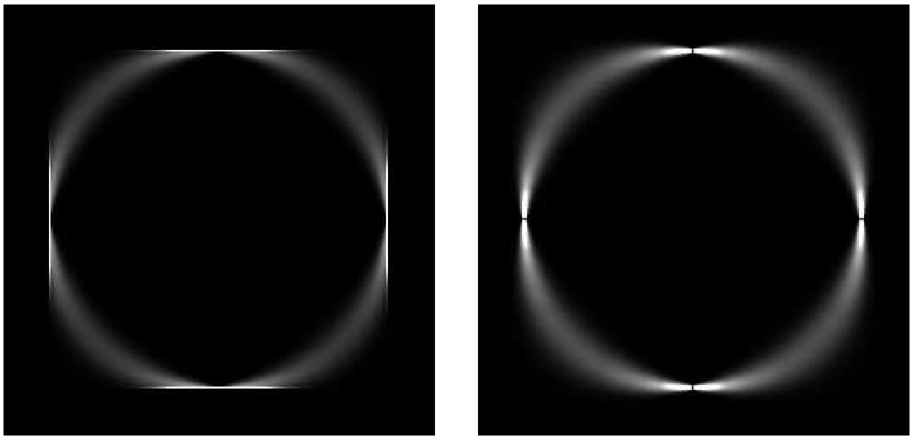

Figure 4: Left: Stochastic completion field for a four stick Ehrenstein figure computed using the method of Williams and Jacobs [11]. Right: The same, but computed using the linear-time method described in this paper.

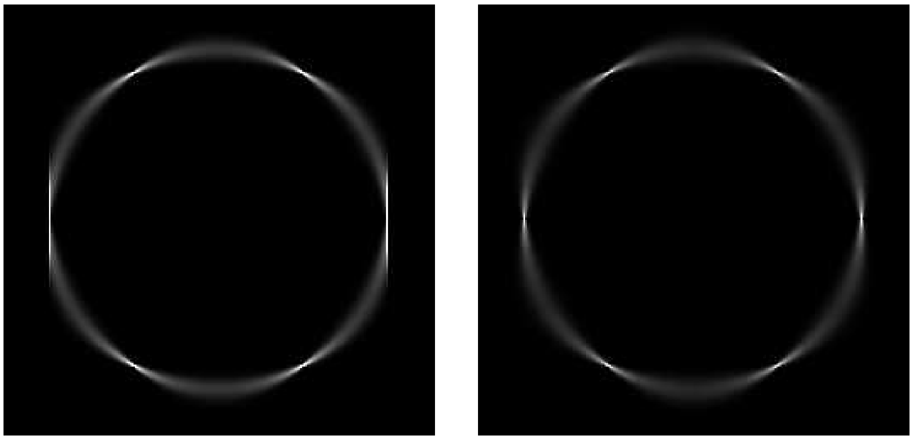

Figure 5: Left: Stochastic completion field for a six stick Ehrenstein figure computed using the method of Williams and Jacobs [11]. Right: The same, but computed using the linear-time method described in this paper.

fields on the left were computed using the method of Williams and Jacobs, while those on the right were obtained using the linear-time algorithm in this paper. In the multi-level resolution pyramid algorithm 18 time steps were performed at each level of the pyramid and (unless otherwise stated) eight time steps were used in the short term source to ensure that the moonwalking artifact was removed.

As a first demonstration we consider the stochastic completion field due to an Ehrenstein stimulus consisting of four sticks, (c.f. [11]). The variance is $\sigma^2 = 0.0004/\gamma^2$ and the decay constant is $\tau = 160\gamma$, where the speed of propagation is $\gamma = 1$. Figure 4 (left) shows the completion field computed using the method of Williams and Jacobs, both clipped above

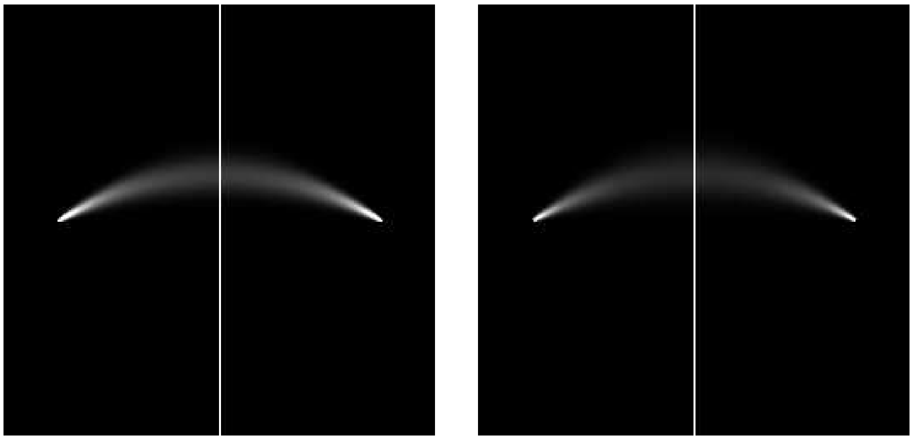

Figure 6: Stochastic completion fields for a two-point initial distribution with $p_1 = (32, 128, 30°)$ and $p_2 = (224, 128, -30°)$.

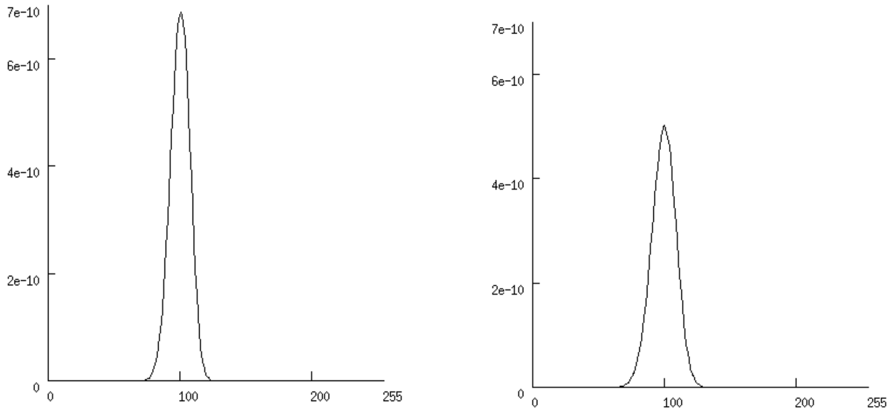

Figure 7: Sections of the completion fields in Figure 6. The horizontal axes of these graphs corresponds to the vertical lines $x = 128$ shown in Figure 6. Notice that the section on the right is somewhat more dispersed since by the time probability mass reaches the midpoint of the arc the Fokker-Planck equation is being solved on a very coarse grid.

at 10^{-9}. The horizontal and vertical straight edges are due to the anisotropic nature of the advection error. Figure 4 (right) show the completion field computed using the linear-time algorithm, scaled by 5.0×10^{10}. In our second demonstration the initial stimulus is an Ehrenstein figure with six sticks. The variance σ^2 and decay rate τ are the same as for the four stick Ehrenstein initial stimulus. The completion fields are shown in Figure 5, both clipped above at 3×10^{-7}. Notice that the completion field on the right is a little duller and is more dispersed half way between each stimulus in comparison with that on the left.

Finally we consider the completion field due to an initial distribution consisting of two points, at $(32, 128, 30°)$ and $(224, 128, -30°)$. Figure 6 shows the completion fields, which are concentrated along the arc of a circle. In this case $\sigma^2 = 0.0004/\gamma^2$ and $\tau = 160\gamma$ where the speed was chosen to be $\gamma = 1.92$ since the radius of the circular arc is 1.92 times the radius of the Ehrenstein initial stimuli discussed above. Twelve time steps were required to remove the moonwalking artifact. Both completion fields are clipped above at 3×10^{-9}. The interesting thing to note here is that even with such a large distance between the two initial points the completion field constructed with only 18 steps per level looks quite good.

7 Conclusion

In two recent papers, Williams and Jacobs[11, 12] have described a representation of visible and occluded image contours called a stochastic completion field. The stochastic completion field is based on the idea that the distribution of contour shapes can be modeled by the paths of particles traveling with constant speed in directions described by Brownian motions. The algorithm described in the more recent paper[12] is based on a simple finite-differencing scheme for integrating the partial differential equation given by Mumford[6]. The time-complexity of the algorithm described in [12] is $O(n^3 m)$ (for an $n \times n$ image with m discrete orientations). Their solution also introduced unwanted, anisotropic smoothing which resulted in noticeable artifacts in the completion fields. In our approach, we carefully add smoothing so as to make the error isotropic. We then decrease the spatial and temporal sampling rate as the high spatial frequencies drop out of the evolving solution. The result is an $O(n^2 m)$ method. In practical terms, the previous algorithm took over an hour to produce an answer (for a problem of size $256 \times 256 \times 36$) and the new method takes about 2 minutes.

8 Appendix

In this appendix we prove Theorem 4.1. The calculation of $\|\mathrm{PR}\,(p) \; - \; p\|_2$ proceeds as follows. First recall that $\|\mathrm{PR}\,(p) \; - \; p\|_2 = \|\widehat{\mathrm{PR}\,(p)} \; - \; \widehat{p}\|_2$ by Plancherel's Theorem. Second, by the Sampling Theorem,

$$\mathrm{PR}\,(p)(x, y\,; t) \;=\; \sum_{n,m=-\infty}^{\infty} p(2n, 2m\,; t)\ \mathrm{sinc}((x - 2n)\tfrac{\pi}{2})\ \mathrm{sinc}((y - 2m)\tfrac{\pi}{2}),$$

where $\operatorname{sinc}(x) := \frac{\sin x}{x}$. The (continuous) Fourier transform of $\mathrm{PR}\,(p)$ is

$$\widehat{\mathrm{PR}\,(p)}(\omega,\eta\,;t) \;=\; \chi(\omega,\eta) \sum_{n,m=-\infty}^{\infty} \widehat{p}(\omega - n\pi, \eta - m\pi\,; t), \tag{8.1}$$

where $\chi(\omega,\eta)$ is the indicator function of the square $|\omega| \le \frac{\pi}{2}$, $|\eta| \le \frac{\pi}{2}$. Since the high frequencies drop out as time increases it is we may assume that

$$\widehat{p}(\omega,\eta) \text{ is negligible for } |\omega| \text{ or } |\eta| > \frac{3\pi}{2}. \tag{8.2}$$

(A calculation of $\widehat{p}$ shows that this assumption is reasonable.) This assumption implies that $\mathrm{PR}\,(p)$ is well approximated by

$$\widehat{\mathrm{PR}\,(p)} \;\approx\; \chi \sum_{n,m=-1}^{1} \widehat{p}_{n,m},$$

where $\widehat{p}_{n,m}(\omega,\eta\,;t) = \widehat{p}(\omega - n\pi, \eta - m\pi\,; t)$, and so

$$\widehat{\mathrm{PR}\,(p)} \;-\; \widehat{p} \;\approx\; (1-\chi)\widehat{p} \;+\; \chi \sum_{(n,m)\in I} \widehat{p}_{n,m}, \tag{8.3}$$

where the index set is $I = \{(n,m) \;:\; |n|, |m| \le 1, (n,m) \ne (0,0)\}$.

A second application of the assumption (8.2), together with the fact that the $\widehat{p}_{n,m}$ are translates of $\widehat{p}$, implies that the two terms in equation (8.3) are approximately equal, and so

$$\widehat{\mathrm{PR}\,(p)} \;-\; \widehat{p} \;\approx\; 2(1-\chi)\widehat{p}.$$

Consequently our condition (4.1) for c can be reformulated as

$$\frac{\|(1-\chi)\widehat{p}\|_2}{\|\widehat{p}\|_2} \;<\; \frac{\epsilon}{\sqrt{2}}. \tag{8.4}$$

To calculate the time evolution of $\widehat{p}$, let p evolve on an $N \times N$ grid according to the finite-differencing scheme given by Steps $3'$ and $4'$ in equation (4.3), and let $\widehat{p}(n,m\,;t)$ denote the discrete Fourier transform of p, i.e. $\widehat{p}(n,m\,;t)$ is the component of p with x-frequency $\frac{2\pi n}{N}$ and y-frequency $\frac{2\pi m}{N}$. Here N is assumed to be even and n, m are integers in the range $-\frac{N}{2} \le n, m \le \frac{N}{2} - 1$. A calculation of $\widehat{p}(n,m\,;t)$ presented below shows that, if the p.d.f. p is initially concentrated at a single point of the grid, then condition (8.4) is equivalent to the inequality (4.4) given in Theorem 4.1.

The discrete Fourier transform $\widehat{p}(n, m, \,; t)$ of the p.d.f., $p(n, m\,; t)$, is calculated as follows. Let $p_x(n\,; t)$ and $p_y(m\,; t)$ evolve according to Steps $3'$ and $4'$ of the finite-differencing scheme (4.3) respectively. Since Steps $3'$ and $4'$ commute it follows that

$$\widehat{p}(n, m\,; t) \;=\; \widehat{p}_x(n\,; t)\,\widehat{p}_y(m\,; t),$$

where

$$\widehat{p}_x(n\,; t) \;=\; \frac{1}{N}\sum_{k=-N/2}^{N/2-1} p_x(k\,; t)\, e^{-\frac{2\pi i\, nk}{N}}$$

is the discrete Fourier transform of $p_x(n\,; t)$. Let T_x be the linear operator given by $p_x(\cdot\,; t) = T_x p_x(\cdot\,; t-1)$ and let $h_x = T_x\delta$ be the response of T_x to the impulse at the origin, δ. Since T_x is a shift-invariant operator it follows that $\widehat{p}_x(n\,; t) = (\widehat{h}_x(n))^t\,\widehat{p}_x(n\,; 0)$. If $p(\cdot\,; t) = \delta$ then a short calculation shows that

$$\widehat{p}_x(n\,; t) \;=\; \frac{1}{N^{t+1}}\left(\tfrac{1}{4}\cos\tfrac{2\pi n}{N} \;+\; \tfrac{3}{4}\right)^t.$$

A similar formula holds for $\widehat{p}_y(m\,; t)$. The required inequality (4.4) now follows from Condition (8.4) using elementary algebra.

Acknowledgments The authors wish to thank Achi Brandt for suggesting the use of a multi-resolution method. Joachim Weickert also offered helpful comments.

References

[1] Geiger, D., Kumaran, K. and Parida, P., Visual Organization for Figure/Ground Separation, *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR '96)*, San Francisco, CA, 1996.

[2] Grossberg, S., and Mingolla, E., Neural Dynamics of Form Perception: Boundary Completion, Illusory Figures, and Neon Color Spreading, *Psychological Review* **92**, pp. 173-211, 1985.

[3] Guy, G., and Medioni G., Inferring Global Perceptual Contours from Local Features, *Intl. Journal of Computer Vision* **20**, pp. 113-133, 1996.

[4] Heitger, R. and von der Heydt, R., A Computational Model of Neural Contour Processing, Figure-ground and Illusory Contours, *Proc. of 4th Intl. Conf. on Computer Vision*, Berlin, Germany, 1993.

[5] Horn, B.K.P., The Curve of Least Energy, MIT AI Lab Memo No. 612, MIT, Cambridge, Mass., 1981.

[6] Mumford, D., Elastica and Computer Vision, *Algebraic Geometry and Its Applications*, Chandrajit Bajaj (ed.), Springer-Verlag, New York, 1994.

[7] Press, W.H., Flannery, B.P., Teukolsky, S.A., and W.T. Vetterling, *Numerical Recipes in C*, Cambridge University Press, 1988.

[8] Shannon, C.E., Communications in the presence of noise. *Proc. of the I.R.E.*, **37**, pp 10-21, January 1949.

[9] Strang, G., *Introduction to Applied Mathematics*, Wellesley-Cambridge Press, Cambridge, MA, 1986.

[10] Ullman, S., Filling-In the Gaps: The Shape of Subjective Contours and a Model for their Generation, *Biological Cybernetics* **21**, pp. 1-6, 1976.

[11] Williams, L.R., and Jacobs, D.W., Stochastic Completion Fields: A Neural Model of Illusory Contour Shape and Salience, *Neural Computation* **9**(4), pp. 837-858, 1997.

[12] Williams, L.R., and Jacobs, D.W., Local Parallel Computation of Stochastic Completion Fields, *Neural Computation* **9**(4), pp. 859-881, 1997.