A Multi-scale Approach to Nonuniform Diffusion

TR91-040
September, 1991

Ross T. Whitaker
Stephen M. Pizer

Medical Image Display Group
Department of Computer Science
Department of Radiation Oncology
The University of North Carolina
Chapel Hill, NC 27599-3175

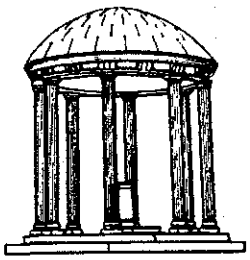

The research was supported by NIH grant number P01 CA47982.
Submitted to CVGIP: Image Understanding.

A Multi-scale Approach to Nonuniform Diffusion

Ross T. Whitaker and Stephen M. Pizer

Department of Computer Science, University of North Carolina
Chapel Hill, NC, 27514

Abstract

This paper examines a new method of image processing that combines information at multiple scales in order to locate boundaries. This method employs a technique of edge-affected diffusion, where blurring is limited by the presence of edges as measured at the scale of interest. By repeating such processing and measuring gradients at successively smaller scales one is able to trace a 'path' through scale space which can preserve accurate information about boundaries of objects, and yet selectively remove objects that fall below a scale of interest. This method is compared with the edge-affected diffusion technique described by Perona and Malik, which depends only on the local gradient of intensity of the processed image. This paper shows some examples which indicate that this method could be useful for boundary detection in the presence of blurring and noise and is also capable of performing grouping of distinct objects at various scales. This paper also examines the sensitivity of this process with respect to ones choice of parameters.

Introduction

Given the task of dividing an discretely sampled image into a finite number of meaningful regions, there is a wide range of possible strategies. One reasonable approach is to group nearby pixels on the basis of similarity in intensity values. For such an approach, boundaries between regions are areas where pixels are dissimilar from pixels nearby. One measure of dissimilarity is the gradient vector. For a two dimensional image, I(x, y), the gradient vector is

$$\nabla I(x, y) = \left\langle \frac{\partial I}{\partial x}, \frac{\partial I}{\partial y} \right\rangle \qquad (1)$$

The magnitude of this vector is the rate of change of intensity in the direction of maximal change. Using the gradient, there are a number of options for identifying divisions or boundaries between regions. Perhaps the simplest is to threshold the gradient magnitude. Pixels that coincide with gradients higher than a specified threshold serve as boundaries to separate other regions. More complex approaches include the Canny edge [1, 2] which identifies boundaries as local maxima of gradient magnitude in the gradient direction. In continuous images these Canny edges form closed curves. These approaches suffer from a common problem. The gradient is not a reliable metric when measured on the presence of unwanted, small scale, luminance fluctuations. That is, the gradient is susceptible to noise.

One response to this problem is to compute the gradient of the image by a sampling process which includes some weighted averaging over a local neighborhood, or equivalently computing local derivatives on blurred versions of the original image. Koenderink [3, 4] argues that in the continuous case, with a fairly natural set of constraints, this sampling function is uniquely defined as a Gaussian that falls off with spatial distance. These Gaussian kernels incorporate a continuous scale parameter that allows one to create a dense set of blurred versions of the original image, a 'scale space'. Ter Haar Romeny and Florack [5, 6] show that in the case of Gaussian noise, the application of Gaussian blurring can improve the signal to noise ratio of image measurements that consist of derivatives of the image intensity. Lindeberg [12] applies a similar set of constraints to discrete images and shows that the necessary blur kernels are modified Bessel functions.

In order to understand the limitations of uniform Gaussian blurring, consider the example in Fig. 1a. It was constructed by drawing two overlapping equiluminant circles on a black background and then adding uniformly distributed random noise with

a range that is twice the intensity difference between the circles and the background. The image is noisy, and yet the outline of the shape formed by the overlapping circles is easily distinguishable. This visual boundary includes not only the smooth outline of the original circles, but the sharp cusp that is formed by their intersection. Fig. 1b shows a display of the gradient magnitude of this image, that was determined by computing differences between nearest neighbor pixels. Although there are some areas of particularly high values that coincide with the borders of the original circles, these areas are incomplete and difficult to reliably characterize.

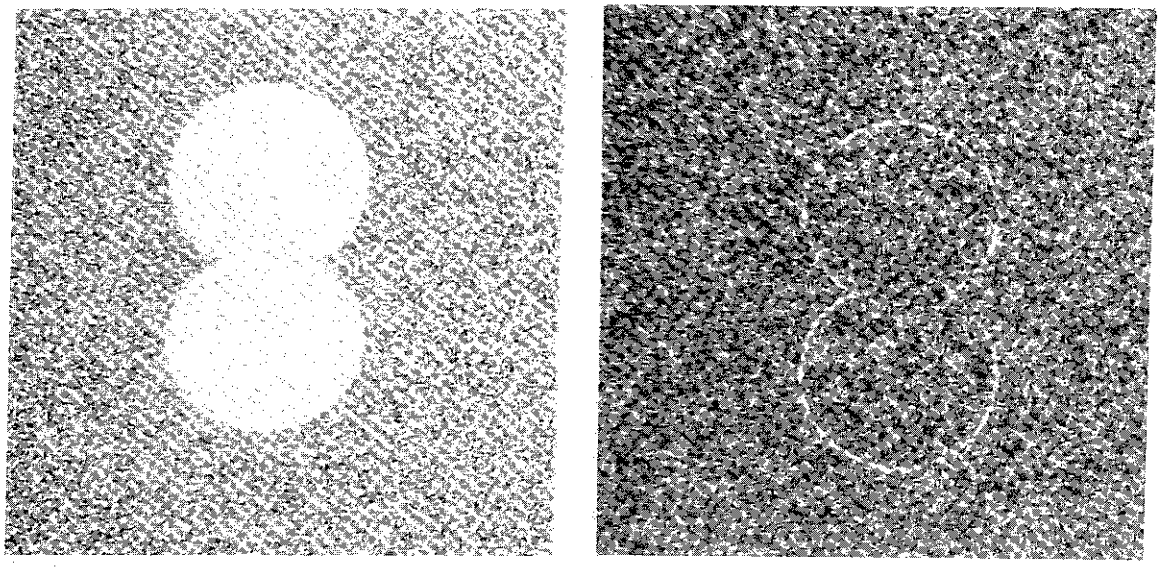

Fig. 1. (a) A white figure on a dark background with additive noise that is uniformly distributed and has a range of twice the intensity difference of the foreground and background. (b) The gradient magnitude of the same image.

Fig. 2 shows the effects of two levels of Gaussian blurring in the image and the gradient magnitude of these blurred images. The effects of the random noise have been dramatically reduced, and the gradient images are more coherent. However, the gradients have formed relatively wide bright bands that could introduce some ambiguity about the precise location of the boundaries in question.

Once could interpret the 'fuzzy' response of the gradient measure at a pixel for a given scale as a graded membership function that indicates the likelihood of a boundary at that pixel. The width of these bands in the gradient image suggests a tolerance on the locations of boundaries that result from these measurements. Since the stated goal is a discrete segmentation of the image, this interpretation begs the question of how to decide the most likely location of boundary within an area of relatively high gradient measure. The approach of Canny is to choose as boundaries the set of points that are local maxima

of gradient in the direction of the gradient as boundaries. Effectively, this approach applies higher order information (directional second derivative) in order to resolve the ambiguity associated with smoothly varying intensity functions or intensity functions that are measured at finite scales.

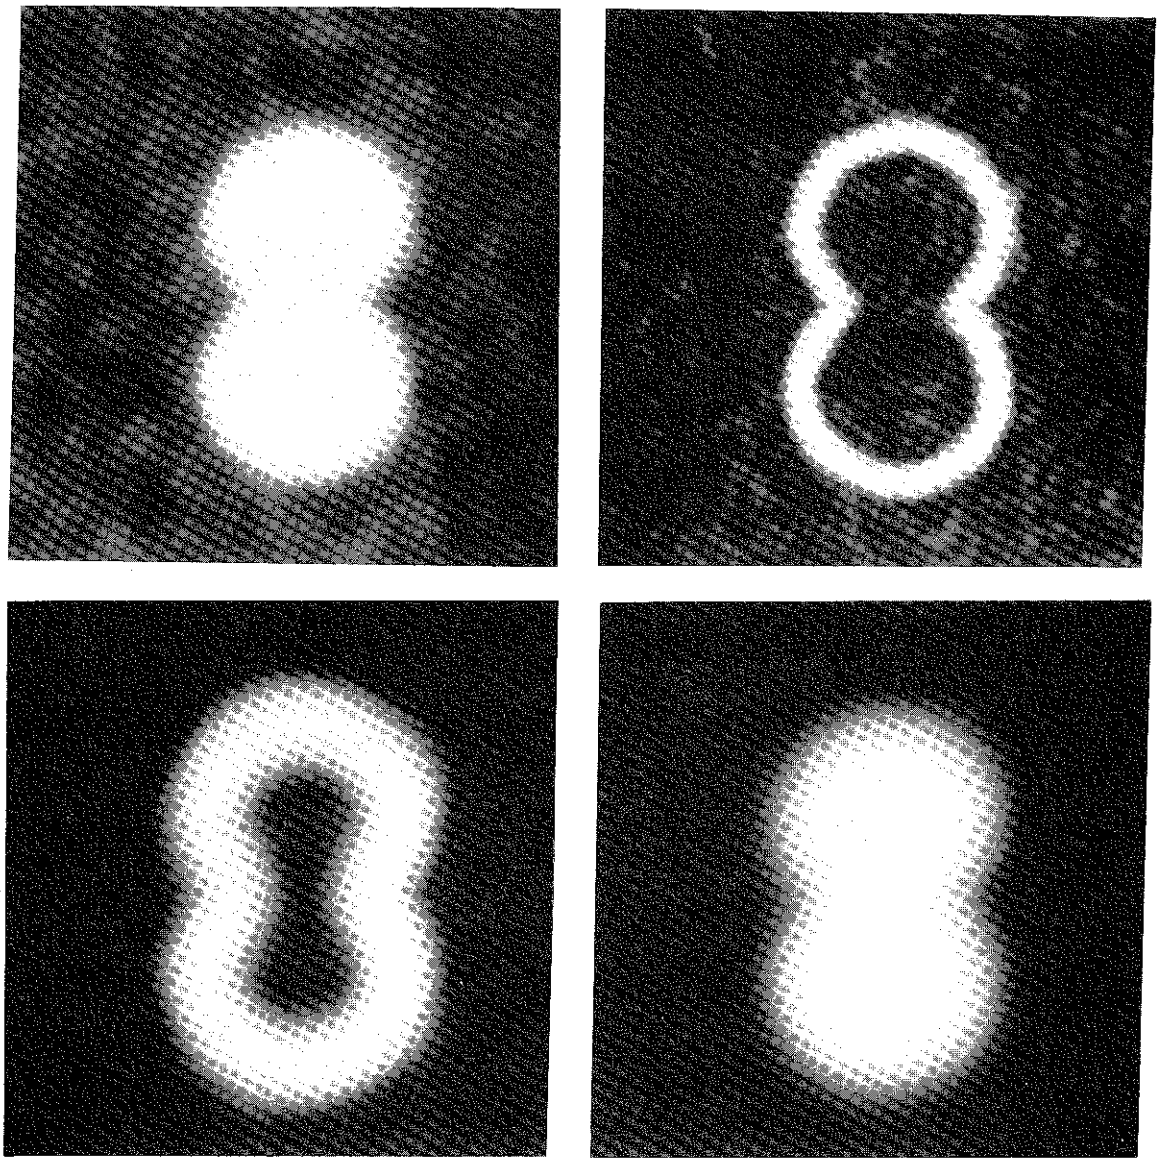

Fig. 2. Clockwise from the upper left: The image of Fig. 1a and the gradient magnitude of that image after blurring with Gaussian kernels. (a) standard deviation of 0.02 (relative to the width of the image) and (b) the gradient magnitude. (c) A kernel with standard deviation of 0.05 and (d) the gradient magnitude.

Such an approach does not address all the difficulties of Gaussian blurring. Referring to Fig. 2, one can see that the sharp cusps where the circles meet which were apparent in the original image are much less distinct. The sharp points associated with these cusps degrade very rapidly with Gaussian blurring, even though they might be

important to the overall shape created by the two circles. The indiscriminate reduction of small scale information associated with uniform blurring can remove important information about features that have both small scale and large scale components. Ideally, one would like to improve the tolerance of gradient measurements at places in the image where the large tolerances associated with larger Gaussian kernels lead to poor decisions about boundary location.

One possibility for making use of Gaussian blurred images is to combine information gathered at multiple scales. It is possible to imagine that large scale (large Gaussian kernel) information could be used to indicate the presence, in an imprecise way, of interesting boundaries, while smaller scale information could be used to determine more precisely the location of these boundaries. Baxter and Coggins [7] have combined information at multiple scales, by associating with each pixel measurements made at a number of discrete scales (samples in scale space) and treating the resulting measurements as positions in a feature space. He then applies some techniques of statistical pattern recognition to the resulting feature space in order to find pixels that have similar behavior through various levels of blurring.

This paper explores another approach to combining information at multiple scales. This approach combines information from a nearly continuous range of scales via an nonuniform diffusion process.

Edge-affected Diffusion

Convolution with the Gaussian kernel is one time slice of the solution to the uniform diffusion equation (or heat equation) with the original image as the initial condition (Eq. (1)).

$$\text{Uniform diffusion} \quad . \qquad \nabla \cdot c \nabla I = \frac{\partial I}{\partial t} \qquad (1)$$

The constant, c, is the conductance and controls the rate of blurring with respect to the time parameter. In most expressions of this equation, units are chosen so that c = 1, so that it does not appear explicitly. The full solution, I(x, y, t), is the continuous scale space mentioned above and each time slice is a version of the original image that has undergone some amount of blurring. This allows one to view Gaussian blurring as a continuous process that evolves from the original image. Nonuniform diffusion, as described by Grossberg [9] as well as Perona and Malik [10], uses a variation on the heat equation which allows the conductance to vary over space.

Nonuniform diffusion $$\nabla \cdot c(x, y, t)\nabla I = \frac{\partial I}{\partial t} \quad (2)$$

More specifically, the above authors describe a type of edge-affected diffusion in which conductance varies in response to the presence of edges in the processed image (Eq. (2)).

Edge-affected diffusion $$\nabla \cdot g(|\nabla I|)\nabla I = \frac{\partial I}{\partial t} \quad (3)$$

In this case the conductance, $g(|\nabla I|)$ depends explicitly on the gradient magnitude of the function itself. Thus Eq. (3), unlike (1) and (2), is nonlinear and its behavior is somewhat different than applying a different size Gaussian kernel at every point in the image. The solutions that result from (3) are not characterized by any explicit convolution with a kernel. Perona and Malik have shown that (3) can have the effect of limiting blurring near edges, and also increasing the steepness or gradient of edges that are sufficiently steep in the initial image.

One way to describe Eq. 3 is the uniform diffusion equation applied to the image as it undergoes local warping of space and time. Let x' and t' be local coordinate systems that vary as functions of the local gradient and let $x' = (1/g(|\nabla I|))x$, $t' = (1/g(|\nabla I|))t$. The primed operators are computed with respect to these local coordinates. Then the uniform diffusion equation in the local coordinates becomes:

$$\frac{d}{dx'} = g(|\nabla I|)\frac{d}{dx'}, \quad \frac{d}{dt'} = g(|\nabla I|)\frac{d}{dt'}, \quad \nabla' \cdot \nabla' I = \frac{dI}{dt'}$$

$$\Rightarrow \quad \nabla \cdot g(|\nabla I|)\nabla I = \frac{dI}{dt}$$

This local "stretching" of space and time near edges can reverse the sign of the laplacian operator (Fig. 3) and allow the 'sharpening' or 'enhancement' of edges near places of high gradient [10,11].

Obtaining the desired behavior from Eq. (3) depends on an appropriate choice of $g(|\nabla I|)$. Some constraints that yield behaved solutions are that $g(|\nabla I|)$ be positive and bounded, and in order to limit diffusion at edges, $g(|\nabla I|)$ should be monotonically decreasing. Perona and Malik provide further constraints on $g(|\nabla I|)$ that enable (3) to exhibit the 'edge enhancing' behavior, and propose

$$g(|\nabla I|) = e^{-(|\nabla I|^2/k^2)}$$

as one possibility. This conductance function introduces a parameter, k. This parameter controls the effect a given gradient value will have on the conductance. Ideally, k should be chosen to reflect the range of gradients in the image, or possibly the gradients in a local neighborhood of every point. For the examples in this paper k will be

expressed as some multiple of the root mean squared (rms) of the gradients at each pixel in image. This has the effect of making the process independent of the particular units in which one chooses to express intensity and thereby allows a small range of k values to apply to a wide variety of images.

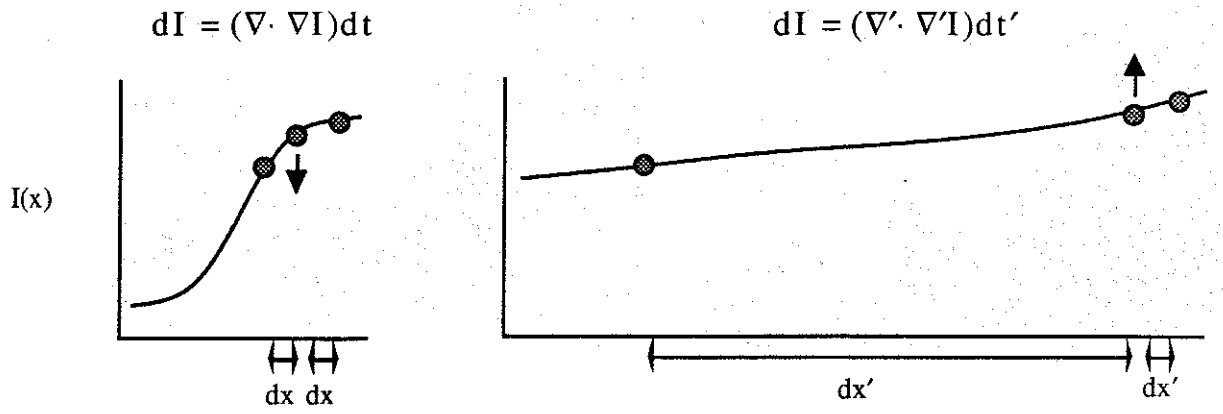

Fig. 3. The local stretching of space and time causes the laplacian operator to switch signs resulting in the sharpening of edges.

Fig. 4 shows the effects of allowing the process to run on a blurred disk. The algorithm employs the method of finite differences. It starts with the initial image and proceeds through discrete time steps. The gradient magnitudes are computed using nearest neighbor differences in orthogonal directions. Derivatives are computed on both a horizontal-vertical and diagonal grid in order to improve the overall behavior of the system. The boundary conditions are chosen to be adiabatic, so that the derivative of intensity is zero in the direction perpendicular to the boundary. The time increments are chosen in order to maintain numerical stability [10, 13], and the total elapsed time is chosen in order to accentuate the edge enhancing effects of the algorithm.

The diffusion process described above can blur images, reduce unwanted noise, and also preserve (or even enhance) boundaries. The images that result tend to have steep, distinct boundaries that adhere to the shapes of objects in the original image. Solutions with adequate values of t offer more reliable gradient measures then in the original image, and facilitate the use of Canny edges or even gradient thresholding as reliable indicators of boundaries.

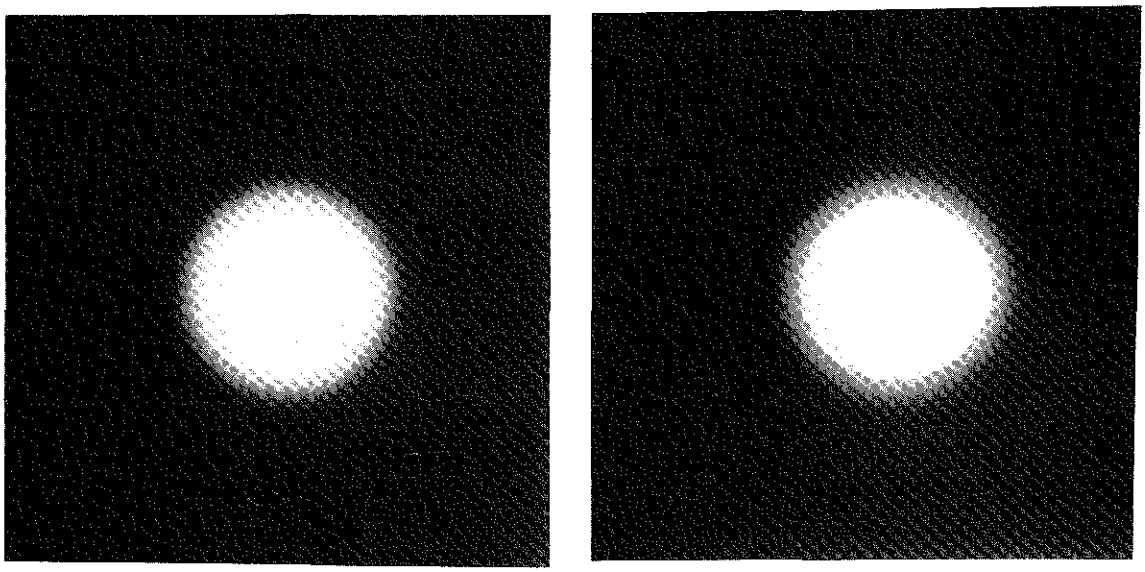

Fig. 4. The result of nonuniform diffusion on a blurred disk. (a) The initial image. (b) The resulting image after 100 iterations with a k value of 5.5 rms VI.

Scale Within Nonuniform Diffusion

Despite the fact that solutions to (3) offer an improvement over Gaussian blurring, they still suffer from the problem of the unreliability of the local gradient measures that are used in calculating the conductance. Local gradients can be so poor in situations where the noise is substantial, that they can drive the process to undesirable results. Consider how solutions behave in the case of the blurred disk in Fig. 4 with the addition of uniformly distributed random noise that has a range of one half of the intensity of the original disk. Fig. 5 shows the original image and samples of the solutions for two distinct values of k.

The algorithm described above has arbitrarily assumed that the scale of individual pixels (sometimes referred to as 'inner scale' [4]) is a meaningful scale at which to make measurements of the gradient. The additive noise introduced into the blurred disk makes the pixel scale a particularly poor scale for such measurements. The result is that the process is unable to capture the larger scale regularity in the image.

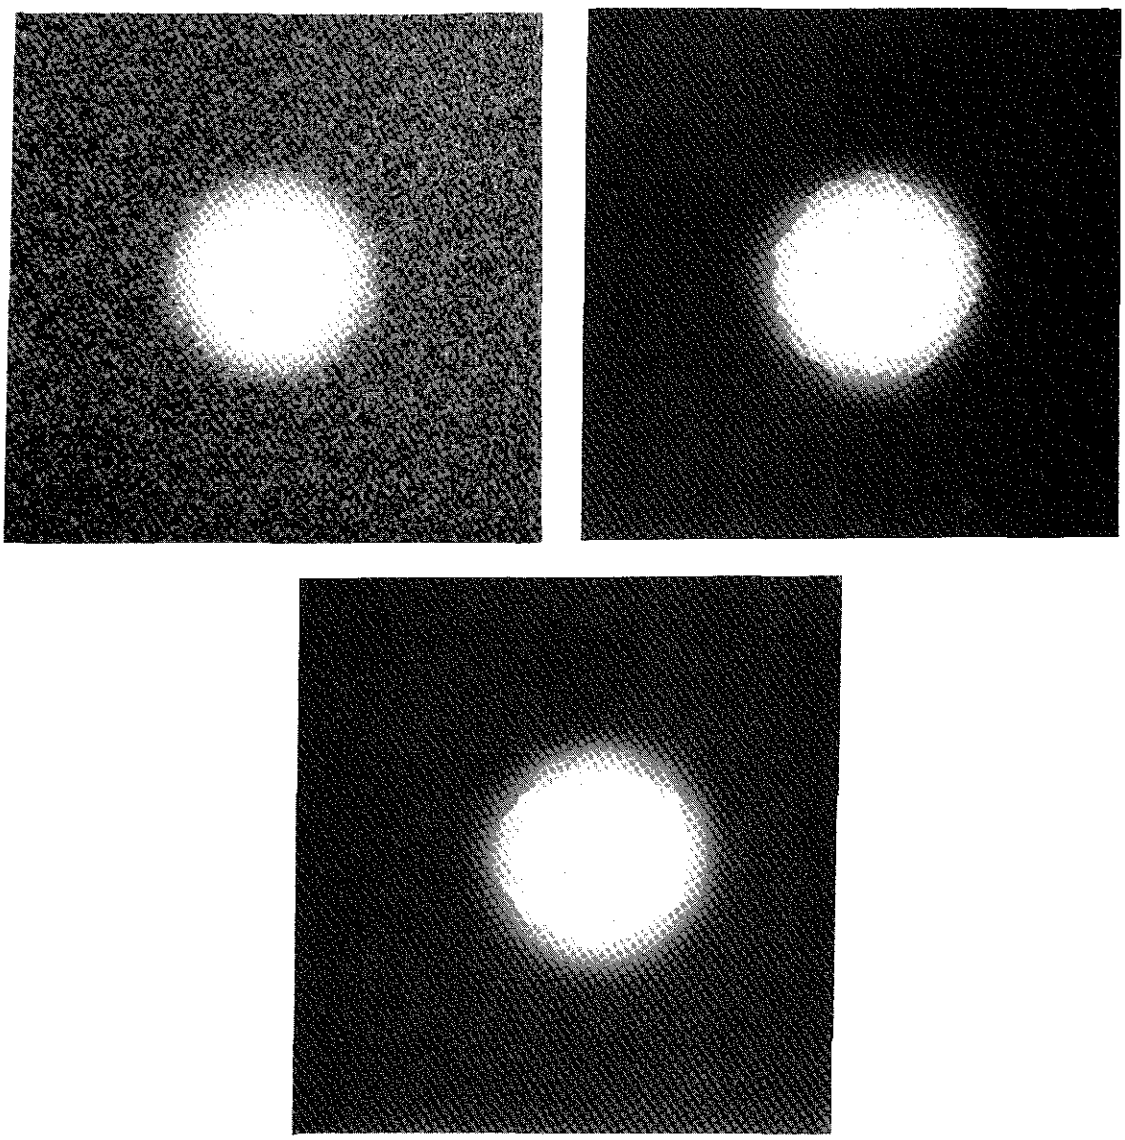

Fig. 5. Clockwise from the upper left: (a) A blurred disk with additive uniformly distributed noise. (b) The result after 100 iterations with a k value of 3.0 rms VI and (c) a k value of 5.5 rms VI.

The edge-affected diffusion process can produce more reliable measures of image gradient for the purpose of making decisions about boundaries. However, in order to carry out this process one needs a best estimate of the image gradient at each point in the image and at each point in time. The problem appears to be circular. In the absence of any a priori information about the image, the best estimate of the gradient (for the conductance determination) that one could hope for [3] is the gradient of the Gaussian blurred image. This suggests that (3) should contain a scale parameter, s, as follows:

$$\nabla \cdot g(|\nabla G(s) * I(x, y, t)|)\nabla I = \frac{\partial I}{\partial t} \qquad (4)$$

where G(s)*I(x, y, t) denotes a convolution (over x and y) of the image at time t with a Gaussian kernel of scale s. It is important at this point to distinguish t and s in (4). The 'evolution' or 'time' parameter, t, characterizes a particular level of nonuniform blurring. The scale parameter, s, describes a level of uniform blurring used to make a gradient measurement of the image at some value of t. The uniform blurring associated with s is not so much a transformation applied to the image, but a description of the another (uniform) diffusion process used to measure the gradient at each point in a single time slice.

Scale as a Function of Evolution

Edge-effected diffusion is an iterative process which uses tentative estimates of the gradient to make incremental changes to the image. Because of unwanted luminance fluctuations these tentative estimates are measured at some scale s. If the process works as we hope, then the unwanted luminance fluctuations should diminish more rapidly than the signal that we wish to ultimately characterize. That is, gradient measurements should become more reliable as the process evolves. This suggests that one should not measure gradients at a single scale throughout the process, but should decrease the scale parameter s to reflect an increasing confidence subsequent versions of the image. The argument above prescribes a diffusion equation of the form

$$\nabla \cdot g(|\nabla G(s(t)) * I(x, y, t)|)\nabla I = \frac{\partial I}{\partial t} \qquad (5)$$

and suggests that s(t) should be some decreasing function of the evolution parameter as in Fig. 6.

The result is a process which gradually 'narrows in' on edges that belong to objects which do not blur away at some chosen scale. If we consider the wide bands in the gradient image of Fig. 2, then we can see that as the process continues, and scale decreases, the bands will become progressively narrower and the nonuniform blurring will continue to smooth closer to the boundary of the desired object. Thus, it is possible to retain very accurate descriptions of edges that belong to objects which are sufficiently large in the initial image.

It is important that s(t) not decrease too quickly, or else edges which are not sufficiently large at some scale $s(t_n)$ might 'resurface' at some scale $s(t_{n+1})$. Likewise, if

s(t) decreases slowly enough, then edges which are sufficiently small at one scale $s(t_n)$, will be adequately smoothed so that they cannot reappear in measurements made at some later time. On the other hand, reducing scale too slowly could allow important information about large scale objects to be lost. Unfortunately, the above arguments do not prescribe a precise form for s(t). It is conceivable that s(t) might not be the same, or even have the same form, for every image, but could vary according to some smoothness measurement made on the process itself. This analysis is an area of future work.

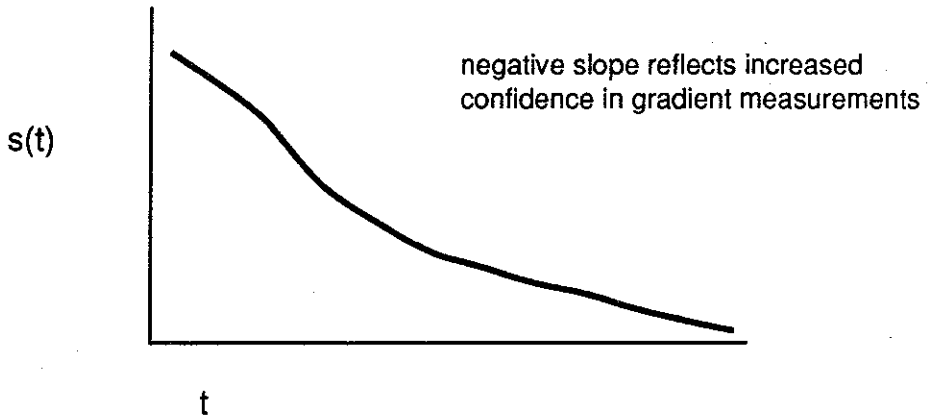

Fig. 6. A typical scale 'path' for a multi-scale nonuniform diffusion process.

For the following examples s(t) has been chosen so that the standard deviation of the Gaussian kernel that is applied to make the gradient measurement decreases linearly with evolution. The slopes and intercepts of these functions are determined empirically. The strategy is to choose an initial value for the scale function, $s(t_0)$, that will screen out features that are not large enough to be interesting. An addition one must choose an appropriate value of k, and and some stopping value for the evolution parameter.

Fig. 7 shows the results of this approach on the blurred disk and overlapping circles with additive noise from Figs. 1a and 5a. The images shown offer a dramatic improvement over Gaussian blurring and edge-affected diffusion with conductance measured at the inner scale. In both cases, the gradients for the conductance function are measured initially with a Gaussian kernel with a radius of about 6 pixels. As the process evolves the size of this kernel is slowly decreased, so that measurements are made at the inner scale only in the final iterations. In the case of Fig. 7b, this process was able to preserve the sharp cusps where the circles meet.

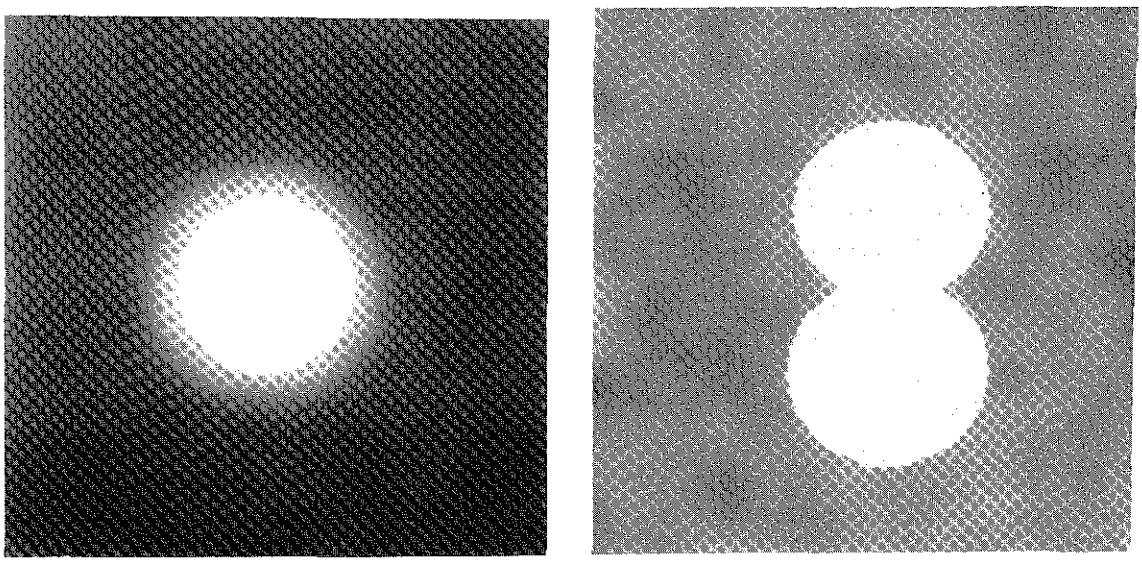

Fig. 7. (a) The result of multi-scale nonuniform diffusion on the noisy blurred disk from Fig. 5a after 200 iterations with a k value of 4.0 rms ∇I and Gaussian scale starting at 0.025 of the image width and decreasing linearly. (b) The overlapping circles from Fig. 1a after 130 iterations with a k value of 3.0 rms ∇I and Gaussian scale starting at 0.005 of the image width and decreasing linearly.

Grouping and Subjective Boundaries

The choices of the initial value and shape of s(t) determine a lower limit on the size and intensity of objects that will appear in later stages of the diffusion process. This has a significant impact on the results of this type of processing. It's reasonable that $s(t_0)$ should depend not only on the properties of the image, but the task one wishes to perform. The reasoning described above does not make any assumptions about what is considered noise, or what form that noise should take; it only assumes that there is an improvement in the signal to noise ratio with applications of Gaussian kernels. This provides a great deal of flexibility in the kinds of features that can qualify as noise. In particular it is possible to choose very large Gaussian kernels in order to prevent relatively large structures (of course they must be smaller than the objects one wishes to characterize) from contributing to the gradient measure. The only requirement is that regions of interest must have measurable intensity differences at some scale.

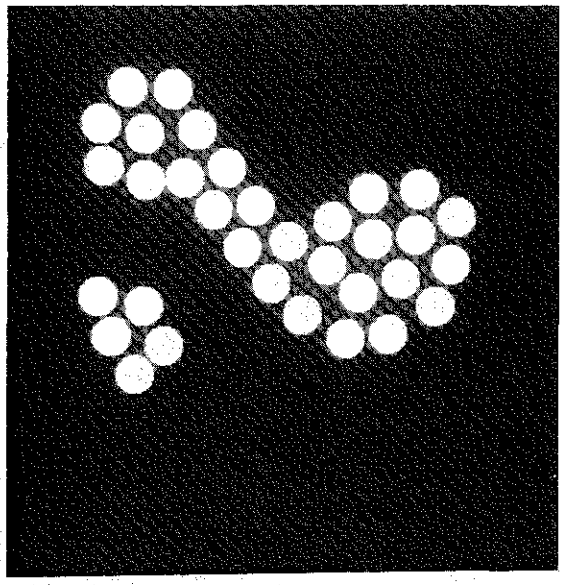

Fig. 8. An image composed of two distinct groups of dots and exhibiting subjective boundaries.

Consider the example in Fig. 8. It consists of 33 white dots on a black background. The dots are readily distinguishable from the background by any of a number of very simple boundary criteria. The picture also has another interpretation: It consists of two light forms, one small and the other large, on a dark background. In the context of this interpretation, the white dots can be considered a kind of highly structured noise which serves to obscure the two light forms. It is even possible to ascribe boundaries to these larger forms, but these boundaries would inevitably pass through black areas of the image that have no luminance variation. We refer to these as subjective boundaries.

By choosing an $s(t_0)$ that has a kernel size larger than the white dots one can use the edge-effected diffusion to 'fill in' luminance across the larger forms and complete the subjective boundaries. Fig. 9 shows several time slices of the dotted image at different times in the process. These images show that the smoothing begins near the center of the object and flows outward toward the boundaries.

The process has groups dots based on their averaged intensities as measured as some large scale. The result in Fig. 9d is a pair of objects that have virtually flat luminance functions with well defined boundaries. As in the earlier examples (Fig. 7), the presence of noise (dots) has influenced the shape of the boundary. The framework for diffusion described in this paper offers some flexibility in the amount of information retained at the boundaries. Allowing s(t) to decrease more slowly, choosing larger values of k, or beginning the process with a Gaussian blurred version of the original image, can all result in smoother boundaries. Of course, these measures do not

discriminate between the effects of noise along the boundary and desirable features of objects.

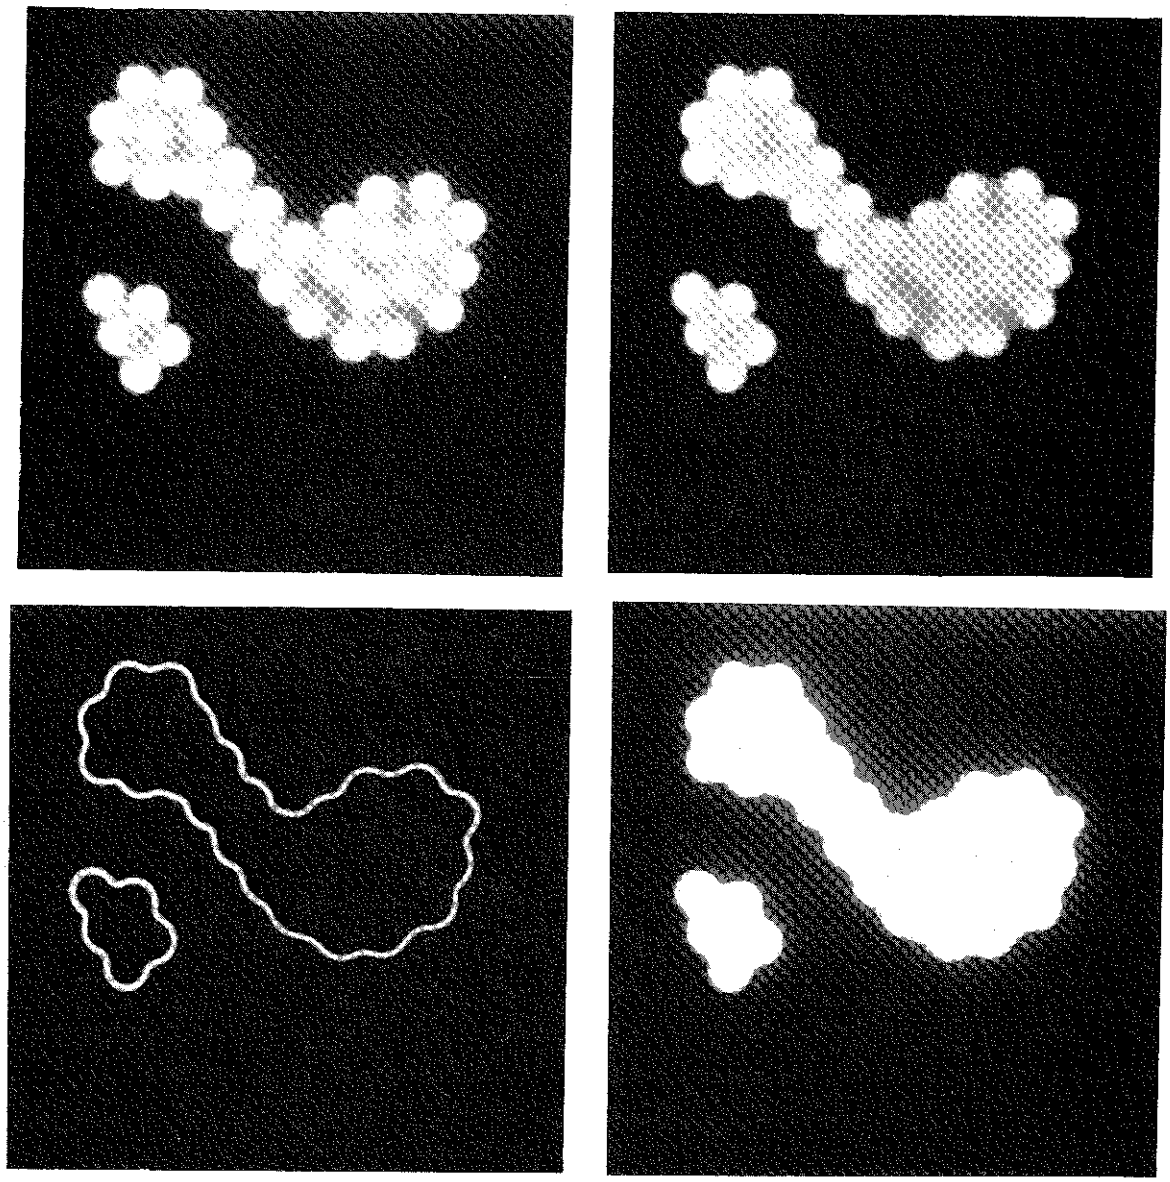

Fig. 9. Clockwise from the upper left: (a) - (c) Consecutive time slices that demonstrate 'filling in' phenomenon associated with as the process evolves on the image from Fig. 8 and (d) the resulting gradient magnitude.

Sensitivity to The Conductance Parameter

The behavior described is sensitive to one's choice of parameters. In particular, the results of this processing are dependent on appropriate choices for the values of the conductance parameter, k, and the time, t, which measures the extent to which the process evolves.

Not all edges in an image in an image are enhanced - smaller edges are blurred away. Choices of k provide a lower limit on the gradients of edges that are enhanced. Define (for the one dimensional case) the flow, $\phi(I_x) = g(|I_x|)I_x$ as in [10]. From Eq. (3) the change in image gradient over time becomes

$$\frac{\partial I_x}{\partial t} = \phi'' I_{xx}^2 + \phi' I_{xxx} \tag{6}$$

The sign of the right hand side of Eq. (6) indicates whether gradients are increasing or decreasing. For a local maximum of gradient, $I_{xx} = 0$ and $I_{xxx} < 0$. Thus, edges that are local maxima of gradient are becoming more steep over time ($\partial I_x/\partial t > 0$) provided that $\phi' < 0$. Perona and Malik [10] suggest choices of $g(|I_x|)$ so that $\phi'(I_x)$ has the property that there is some α such that

$$\phi'(I_x) \begin{cases} > 0 & \text{if } I_x < \alpha \\ < 0 & \text{if } I_x > \alpha \end{cases} \tag{7}$$

The conductance function $g(|\nabla I|) = e^{-(|\nabla I|^2/k^2)}$ has this property, and $\alpha = (1/2)^{1/2}k$. This suggests that one could choose k to reflect the lower bound on the steepness of edges that will be enhanced; edges below α will be blurred away while edges above α will be enhanced. Unfortunately, there are further constraints on the choice of k. Fig. 10 shows the edge-affected diffusion applied to the blurred disk from Fig. 4a. For this analysis k was purposely chosen to be low; it is two times the root mean squared of the image gradient. The result is a 'staircasing' effect. Instead of isolating a single steep portion of the disk boundary as in the examples above, the process has broken the boundary into many discrete steps.

In order to better understand this phenomenon we will analyze a particular case of the one dimension edge and then try to generalize the results. Consider an edge I(x), where

$$I(x) = \int_{-\infty}^{x} A e^{-s^2/2\sigma} ds \tag{8}$$

This is equivalent to a Gaussian blurred step function. The advantage of using this function is that it has closed form analytical expressions for all of it's derivatives. The derivative of this function,

$$I_x(x) = A e^{-x^2/2\sigma},$$

is the familiar bell shaped Gaussian and I(x) has a maximum gradient at $x_0 = 0$. The gradient at the most steep point, x_0, is increasing for appropriate choices of the conductance coefficient (Eq. (6)). However, this says nothing about the way I_x changes in the neighborhood of x_0. In order for this edge to become we would expect that the gradient at x_0 to increase *more quickly* than the gradients in a local neighborhood of x_0. If the local neighborhood of x_0 increases it's slope more quickly than $I_x(x_0)$ then the edge, I(x), becomes steeper, but also 'flatter' as shown in Fig. 11.

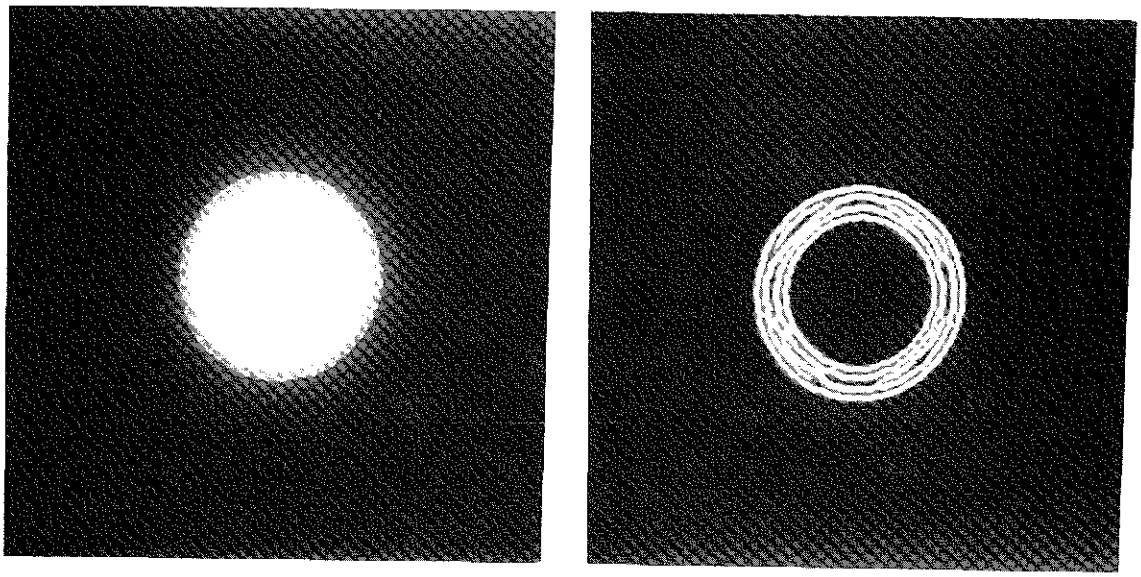

Fig. 10 The 'staircasing' effect as exhibited on the blurred disk of Fig. 4 resulting from choices of k that are too low.

Characterizing this behavior requires analyzing at the manner in which higher order terms change over time. From (6) and (7) and expanding the exponential in a Taylor series

$$\frac{\partial I_x}{\partial t} = \left[\phi''\frac{x^2}{\sigma^2}Ae^{-x^2/2\sigma} + \phi'\left(\frac{x^2}{\sigma^2} - \frac{1}{\sigma}\right)\right] Ae^{-x^2/2\sigma}$$

$$= -\phi'\frac{A}{\sigma} + \left(\phi''A + \frac{3}{2}\phi'\right) A\frac{x^2}{\sigma^2} + O(x^4) \qquad (9)$$

The local behavior of $\partial I_x/\partial t$ is dependent on the sign of the expression $(\phi''A + (3/2)\phi')$. For $(\phi''A + (3/2)\phi') < 0$, $\partial I_{xxx}/\partial t$ is positive and edge becomes sharper in the neighborhood of x_0, while $(\phi''A + (3/2)\phi') > 0$ creates the 'flattening' effect shown in figure 12. Because $\phi' < 0$, as required for the sharpening process, this flattening can only happen when $\phi'' > 0$. Because the flow, ϕ, is strictly positive and has a negative slope for sufficiently large I_x, any acceptable choice of $g(I_x)$ will result in a positive ϕ'' for I_x sufficiently large.

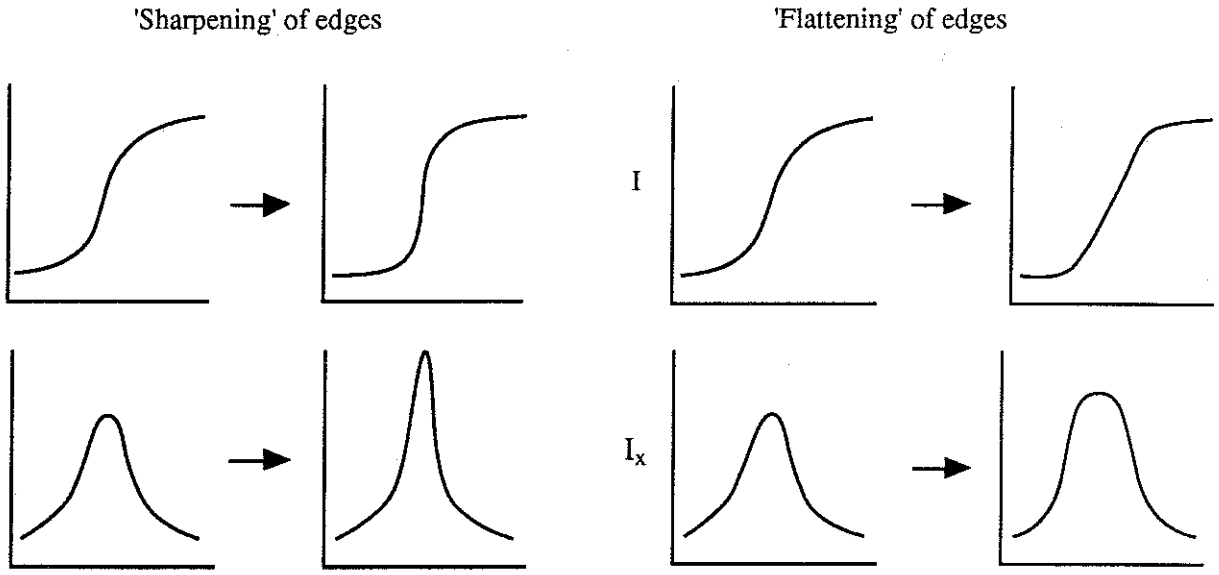

Fig. 11. 'Sharpening' is distinguished from 'flattening' by examining the manner in which the first derivative changes in a local neighborhood.

In generalizing this analysis to edges other than step functions, it is helpful to define a new function

$$e(I_x,\ I_{xx},\ I_{xxx}) = \partial I_x/\partial t = \phi''(I_x)I_{xx}^2 + \phi'(I_x)I_{xxx}$$

which is Eq. (6) with a change in formal parameters. This function describes the rate of change of the first derivative at any location in the image, given the local properties of up to third order. In order for the process to accentuate the most steep portions of a smooth edge and reduce the steepness of other portions of the same edge, the function $e(I_x, I_{xx}, I_{xxx})$ should penalize locations that are not local maxima of gradient. We will consider its behavior in a neighborhood of x_0 where $I_x(x_0)$ is a local maxima. If we fix I_x and I_{xxx} in this neighborhood then this function is a quadratic in $I_{xx}(x)$.

$$e(I_{xx}) = \phi''I_{xx}(x)^2 + \phi'I_{xxx}$$

This quadratic has a local maxima at $I_{xx} = 0$ only if $A = \phi'' < 0$. First and third order properties being equal, local maxima of intensity gradient ($I_{xx} = 0$ and $I_{xxx} < 0$) increase maximally only if $\phi'' < 0$. For $\phi'' > 0$ places that have high curvature and only slightly lower gradients can increase their gradients more quickly than those places that are locally the most steep in the same smooth edge. For the conductance term $g(|\nabla I|) = e^{-(|\nabla I|^2/k^2)}$ that was used in these examples, $\phi'' > 0$ only for $I_x < (3/2)^{1/2}k$. In the event that there are smooth edges with gradients larger than $(3/2)^{1/2}k$, the staircasing described above can result.

This analysis suggests that for a given conductance function there is a limited range of gradient values that will produce edge enhancement without the staircasing of smooth edges as shown in Fig. 12. The choice of k will be very difficult for images that have a great variation in gradient values. For such images it might be impossible to characterize both dim and bright objects. Fortunately, experiments have shown that this staircasing tends to occur only on wide smooth edges as in the blurred disk. For a wide range of images that have only very narrow boundary regions, this phenomenon was not widespread.

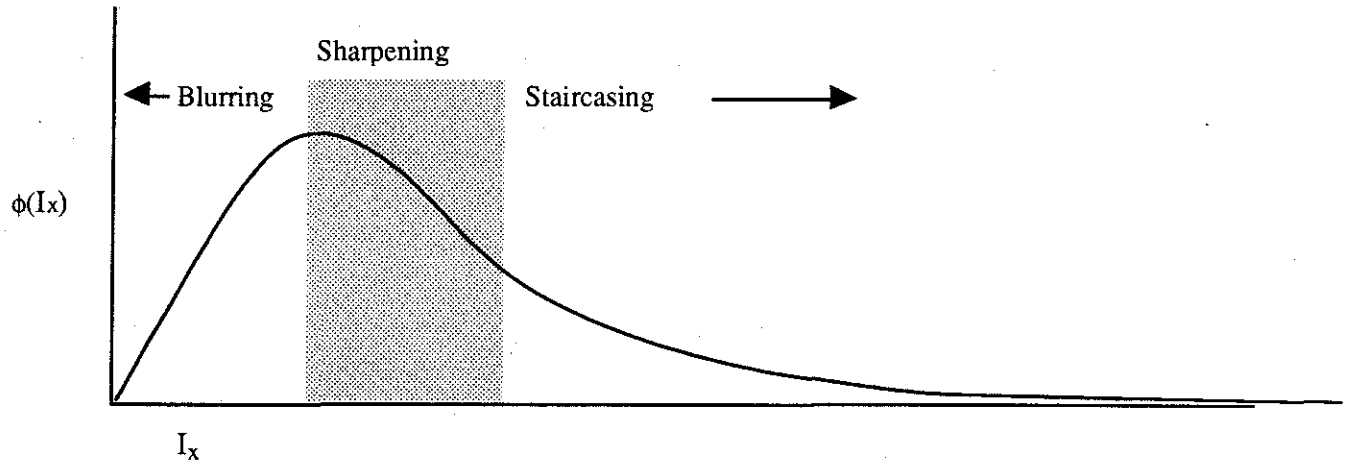

Fig. 12. Because of staircasing there is a limited range of gradient values which produce desirable results.

Sensitivity to Total Elapsed Time

The amount of processing that should occur on an image requires another important decision. In principle one would like to run this process for some period of time and use the result to make decisions about boundaries in the original image. If the process converged to some useful result, then one could carry out the process until it reached a point where changes were inconsequential. In the discrete algorithm described this is not the case. The conductance can never drop to zero because the gradients (as computed by finite differences) are bounded by the difference in minimum and maximum intensity values in the original image. Given enough time, the solutions as computed by finite differences will converge to a single value. This is best understood by considering solutions with $dI/dt = 0$ in two dimensions.

$$\frac{dI}{dt} = \nabla \cdot e^{-(\nabla I \cdot \nabla I)/k^2} \nabla I$$

$$= e^{-(\nabla I \cdot \nabla I)/k^2}\left(-\frac{2}{k^2}\nabla I \cdot \nabla \nabla I \cdot \nabla I + \nabla \cdot \nabla I\right)$$

$$= e^{-I_v^2/k^2}\left[\left(1 - \frac{2}{k^2}I_v^2\right)I_{vv} + I_{uu}\right] = 0$$

The final expression is shown in gauge coordinates, and the subscripts w and v represent derivatives in the direction of the gradient and the direction tangent to the level set, respectively. This expression can hold if either the term in square brackets is zero (or approaches zero in the limit) or the exponential term approaches zero.

One can argue that the term in brackets has only a trivial solution for some very simple examples. Consider an image which is half black and half white with a border running vertically down the center of the image. For such an image I_{vv} is zero (isophotes are straight) and I_{ww} is zero only at the most steep point in the boundary between black and white and the flat areas on either side. There must be areas near that point that have non zero I_{ww} . Therefore, with the boundary conditions discussed above, the term in brackets approaches zero for the entire image only as I becomes flat.

It is conceivable that the exponential term could become zero within the precision of the numerical representation of the discrete image. A simple analysis shows that for floating point representation this happens only for images that have a large number of samples. For this analysis it is best to express k in terms of the expected value of the gradient magnitude squared over the entire image: $k = f\langle I_w^2\rangle$, where f is a constant and $\langle I_w^2\rangle$ is the root of the mean of the gradient magnitude squared. If r is the number of bits in the mantissa and b is the base used, then the condition for exponential term to have no effect on the image is

$$e^{-I_v^2/f\langle I_v^2\rangle} \le b^{-2^{r-1}}$$

$$\frac{I_v^2}{\langle I_v^2\rangle} \ge 2^{r-1}\, f\, \ln|b|$$

Generally speaking, $I_w^2/\langle I_w^2\rangle$ is proportional to the square root of the number of samples in the image. For the black on white image discussed earlier, $I_w^2/\langle I_w^2\rangle$ is equal to the width of the image. For the this image and r = 16, b = 10, and f = 2 the image would need to be over 150,000 pixels wide in order for the numerical discretization in the exponential term to allow the process to stop changing.

Edges, except for those that meet the above criteria, will continually 'leak' as the process evolves. This leaking is not like the blurring associated with Gaussian blurring - the edge remains distinct - but the grey levels on either side of the edge will slowly

approach a middle value. This analysis, combined with edges enhancing behavior discussed earlier, suggests that edges undergo an evolution of increasing gradient followed by a period of slow decay, until the gradient reaches the lower limit on the enhancing behavior ($I_x < \alpha$), after which it decays rapidly.

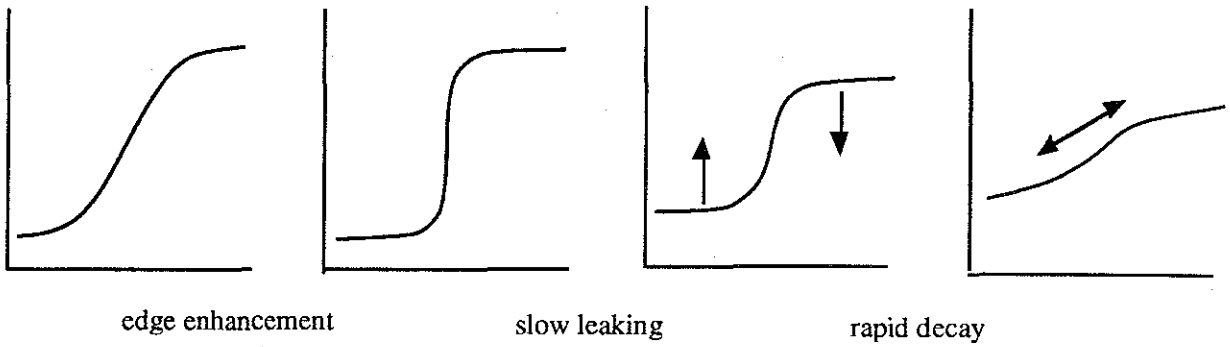

Fig. 13. Edges undergo a period of enhancement followed by a period of slow decay due to the bounded nature of derivatives as measured by finite differences.

Experiments have indicated that the leakage happens relatively slowly compared with the edge enhancement. The leaking, however, is greater for edges of lower contrast. In images that have a great disparity of contrasts, edges of less contrast could begin deteriorating while greater edges are still increasing in steepness. For the examples in this paper, the amount of time that the images evolved is chosen experimentally, in order to provide the best qualitative results in terms of distinctive edges and object separation.

Nordstrom [11] shows that edge-affected diffusion can be described as regularization process and can be thought of as an iterative approximation to an optimal tradeoff between associated 'smoothness' and 'edgeness' cost functions. Nordstrom also introduces a variation on (3) that enables the process to have a nontrivial optimal (steady state) solution.

Conclusions

The edge-affected diffusion equation offers an attractive alternative to Gaussian blurring in detecting the boundaries of objects in the presence of limited amounts of noise. However, the edge-affected diffusion process requires requires measurements of image gradient that must be made at some scale. Previous work has implicitly assumed that the inner scale of the image is the appropriate scale to make such measurements, but in the case of large amounts of noise this is not a reasonable assumption and the process produces less than adequate results.

One should make gradient measurements at each step in the process using Gaussian kernels, and choose the size of these kernels based on the nature of noise in the image. In addition, allowing the size of these kernels to decrease over time in response to the improved noise characteristics of the processed image produces a diffusion that gradually narrows in on edges. This diffusion process can reduce unwanted noise at a range of scales, and yet provides accurate boundaries of objects that are of sufficient size. This approach allows for the reduction highly correlated noise, so that larger objects can be formed from groups of smaller objects when the smaller objects are viewed as noise and the appropriately sized Gaussians are chosen.

The results of this processing are sensitive to the conductance parameter and the amount of time allowed to evolve. This is especially true in images that have a wide range of contrasts. Future work will concentrate on developing conductance functions that are less sensitive to intensity transformations, so that dark and light objects in the same image will be treated equally.

Acknowledgements

We would like to thank Dr. David Eberly for his help in editing this paper as well as Dr. James Coggins, Dr. James Damon, and Dr. Robert Gardner for their input on this work.

This research has been supported by NIH grant #P01 CA47982

References

1. J. Canny, Finding edges and lines in images, Technical Report 720, MIT, Articficial Intellegence Laboratory, 1983.

2. J. Canny, A computational approach to edge detection. *IEEE PAMI* **8** (6), 1987, pp. 679-698.

3. J. J. Koenderink and A. J. Van Doorn, Receptive Field Families, *Biol. Cybern.* **63**, 1990, pp. 291-298.

4. J. J. Koenderink, *Solid Shape*, MIT press, Cambridge, Mass., 1990.

5. T. H. Romeny and L. Florack, A Multiscale Geometric Approach to Human Vision, in *Perception of Visual Information* (B. Hendee and P. N. T. Wells, Ed.), Springer Verlag, Berlin, 1991.

6. T. H. Romeny and L. Florack, Scale Space; its natural operators and differential invariants, Lecture Notes in Computer Science **511** (A. C. F. Colchester and D. J. Hawkes Eds.), Springer-Verlag, 1991, pp. 239-255.

7. L. C. Baxter and J. M. Coggins, Supervised Pixel Classification Using a Feature Space Derived from an Artificial Visual System, in *Proceedings SPIE Conference on Intellegent Robots and Computer Vision IX: Algorithms and Techniques, Boston, November 5-7, 1990*, pp. 459-469.

9. S. Grossberg, Neural Dynamics of Brightness Perception: Features, Boundaries, Diffusion, and Resonance, *Percpetion and Psychophysics* **36** (5), 1984, pp. 428 - 456.

10. P. Perona and J. Malik, Scale-space and Edge Detection Using Anisotropic Diffusion, *IEEE Transactions on Pattern Analysis Machine Intelligence* **12**, pp. 429-439, 1990.

11. N. Nordstrom, Biased Anisotropic Diffusion - A unified Regularization and Diffusion Approach to Edge Detection, *Image and Visual .Comp.* **8** (4), pp. 318-327, 1990.

12. T. Lindeberg, A Scale-space for descrete signals, *IEEE PAMI* **12** (3), 1990, pp. 234-245.

13. I. G. Petrovskii, *Partial Differential Equations*, W. B. Saunders Company, Philadelphia, 1967.