Geometry-limited Diffusion in the Characterization of Geometric Patches in Images

TR91-039

September, 1991

Ross T. Whitaker

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

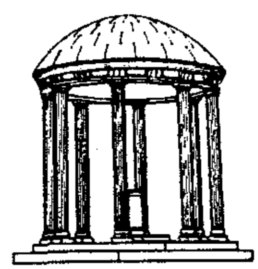

The research was supported by NIH grant number P01 CA47982.
Submitted to CVGIP: Image Understanding.

Geometry-limited Diffusion In The Characterization of Geometric Patches In Images

Ross T. Whitaker

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

Abstract

Much of image processing and artificial vision has focused on the detection of edges - particularly, edges that are measured by the gradient magnitude. Higher order geometry can provide a richer variety of information about objects within images and can also yield useful measurements which are invariant to certain kinds of intensity transformations. However, analyzing higher order geometry can be difficult because of the sensitivity of higher order filters to noise. Low pass filters can alleviate the effects of high frequency noise but tend to distort the geometry in ways that make the resulting measurements less useful.

This paper suggests a generalization of anisotropic diffusion as a mechanism for making reliable and precise geometric measurements in the presence of blurring and noise. This mechanism is a generalized form of edge-affected diffusion that applies to multi-valued functions. We pursue the interpretation of multi-valued descriptors as positions in a feature space and describe how this premise yields a natural form for a set of coupled anisotropic diffusion equations that depend on one's choice of distance in the resulting feature space. The appropriate choice of distance allows one to measure areas of the image where the feature positions are changing rapidly and vary the conductance in the diffusion equation accordingly. These features can be the outputs of some multi-valued imaging device, or measurements made (via filters) on a single valued image.

Feature spaces that consist of measurements made on single-valued images can reflect geometric properties of the local intensity surface. The anisotropic diffusion can be used to segment images into patches that share local geometric properties so that the boundaries of these patches are geometrically and visually interesting. The appropriate choice of distance in such feature spaces can yield meaningful geometric information. One such geometric feature space consists of the first order derivatives. This paper presents a distance measure in this space that results in a process for reliable and accurate detection of 'creases' and 'corners'. These ideas can be generalized to other features, including higher order derivatives. The appropriate choice of distance in such feature spaces could yield meaningful higher order geometric information.

Introduction

Koenderink [1, 2, 3] has argued that important information about an image can be discerned not only from the intensity value at every point in the image but from the local geometry of the intensity surface. He also shows that the appropriate procedure for measuring such geometry is convolution of the image with filters from within a family of receptive fields that resemble derivatives of Gaussians. Ter Haar Romeny and Florack [4, 5] show that these derivatives can be combined to create certain invariant (independent of choice of coordinate systems) geometrical measures and that these invariants can be combined to form measurements that have visual significance, i.e. edges, corners, etc. Furthermore, gaussian blurring can improve the signal to noise ratio of higher order derivative filters.

Lowpass filters such as the gaussian can have adverse effects on the characterization of objects whose shapes depend on high frequency information. Anisotropic diffusion has been proposed [6, 7, 8] as an alternative to the isotropic scale space, which corresponds to convolution with gaussian kernels. More specifically, edge-effected diffusion, in which conductance varies according to local gradient information, has been shown to reduce unwanted noise and preserve, or even enhance, edges. One can think of such edge-effected diffusion as a regularization process [8] which preserves areas of rapid change in intensity (large gradient magnitude).

These authors propose edge-affected diffusion in the form of the equation,

$$\nabla \cdot g(|\nabla f|)\nabla f = \frac{\partial f}{\partial t} \tag{1}$$

where g, the conductance modulating term, is some bounded, positive, decreasing function of $|\nabla f|$, and t is the time or evolution parameter.

Whitaker [9] has shown that isotropic scale can be incorporated into the conductance term in order to account for the unreliability of local gradient measurements in the presence of correlated and uncorrelated noise. Furthermore, by decreasing the scale at which the gradient is measured over time, one can obtain boundary information that reflects both small scale and large scale gradient information. This results in the multi-scale anisotropic diffusion equation,

$$\nabla \cdot g(|\nabla G(s) * f|)\nabla f = \frac{\partial f}{\partial t} \tag{2}$$

in which "G(s)*" denotes convolution with a gaussian kernel of a particular size s(t), which generally decreases as the process evolves.

The result of such processing is typically a set of smoothly varying regions that have relatively sharp boundaries - a kind of piecewise continuity. These images lend themselves to segmentation schemes that rely on local gradient measurements. The appropriate choices of s(t) can control the scale of the resulting smooth regions, and yet maintain accurate information at the boundaries of these regions, preserving their overall shapes.

Eqs. (1) and (2) implicitly assume that a single intensity function is the appropriate measure to characterize points in the image. If we relax this notion and admit that there are many possible measures which could characterize points in the image, then we might expect there to be a regularization process which captures not only a single value in the image, but any number of values that are relevant to the task at hand.

Multi-valued diffusion

A conventional digital image (an array of pixels) can be viewed as a discrete sampling of some continuous function $f:\Re^n \rightarrow \Re$. The domain is the space in which one typically views the image. For a photograph: n = 2, and the domain is the flat surface on which one views the picture. We call this the image space. The range is a number that characterizes each point in that picture - the intensity, or lightness. We call this the feature space. For digital images, the range is usually measured (or computed) and recorded at discrete locations in the domain on a rectilinear grid. The exact nature of this sampling is not directly important for the following discussion.

A multi-valued image is a discrete sampling of the function $\mathbf{f}:\Re^n \rightarrow \Re^m$, in which the bold **f** designates it as multi-valued, and each point in the domain is characterized by its position in the feature space, recorded as a finite array of measurements at each pixel. At times it will be important to remember that the set of numbers used to record the position of a pixel in the range is not indicative of the shape or nature of that space. More specifically, one cannot assume that an appropriate measure of distance in the range space follows directly from the array of measurements that one uses.

Perhaps the most apparent examples of multi-valued data are the images that result from certain imaging devices that measure several physical quantities at the same location in space. Multi-echo MRI or multi-spectral Landsat data are two examples. One could treat each physical measurement in these datasets as a separate image, and perform the above edge-affected diffusion process separately for each measurement. However, if the goal of the edge-effected diffusion process is to capture dissimilarities between pixels across a number of these measurements, then the regularization process should reflect this.

The diffusion process introduces a time or evolution parameter, t, into the function $\mathbf{f}: \Re^n \times \Re^+ \rightarrow \Re^m$ so that there is a multi-valued function at each point in time, or each level of processing. The multi-valued diffusion equation is

$$\nabla \cdot g(\mathcal{D}(G(s) * \mathbf{f}))D\mathbf{f} = \frac{\partial \mathbf{f}}{\partial t} \tag{3}$$

where $\mathcal{D}: \Re^m \rightarrow \Re$ is a *dissimilarity* operator. The convolution, $G(s)*\mathbf{f}$, incorporates the above notion (Eq. (2)) of time varying, isotropic scale. $D\mathbf{f}$ is the derivative of $\mathbf{f}$ in the form of a matrix. The conductance, g, is a scaler, and the operator "$\nabla \cdot$" is a vector that is applied to the matrix $gD\mathbf{f}$ using the standard convention of matrix multiplication. Eq. (3) is a system of separate single valued diffusion processes, evolving simultaneously, and sharing a common conductance modulating term.

$$\begin{aligned} \nabla \cdot c(\mathcal{D}(G(s) * \mathbf{f}))\nabla f_1 &= \frac{\partial f_1}{\partial t} \\ &\vdots \\ \nabla \cdot c(\mathcal{D}(G(s) * \mathbf{f}))\nabla f_m &= \frac{\partial f_m}{\partial t} \end{aligned} \tag{4}$$

Eqs. (3) and (4) describe a smoothing process that respects boundaries and operates on a set of images simultaneously. The boundaries are not defined on any one image, but are shared among (and possible dependent on) all of the images in the set, as shown in Fig. 1.

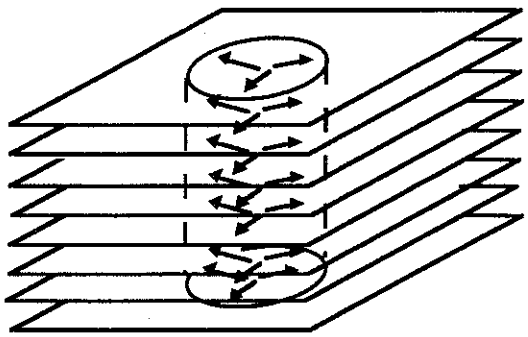

Fig. 1. Mult-valued diffusion is a system of diffusion processes that are coupled through the conductance function.

Diffusion in a Feature Space

The behavior of the system (3) is clearly dependent on the choice of the dissimilarity operator, $\mathcal{D}$. In the single feature case, the gradient magnitude proves to be an appropriate measure. That is, $\mathcal{D}f = (\nabla f \cdot \nabla f)^{1/2}$. For higher dimensions it is helpful to consider the range of $\mathbf{f}$ as a feature space and to evaluate dissimilarity based on distances between pixels in this space.

Applications of statistical pattern recognition to image segmentation are available in the literature [11, 12, 13, 14, 15]. Such approaches apply multi-variate statistics in order to find clusters of pixels in these feature spaces. Boundaries between clusters can be defined in a way that optimizes certain metrics. Pixels can be classified on the basis of their distance to groups of pixels that are nearby in the feature space. Effective use of such techniques requires an appropriate measure of distance. Typically, linear transformations are performed on the feature space in order to produce meaningful clusters. These transformations can be determined by statistical measures made on the available data. Coggins [16] has shown that multiple measurements of single valued images, obtained through sets of linear filters, can be treated using these same principles.

The dissimilarity operator is constructed to capture the manner in which neighborhoods in the image space (n dimensional) map into the feature space. If $\mathbf{f}$ is a well behaved function (Lipschitz continuous for example), then for a point $\mathbf{x}_0 \in \Re^n$ in the domain of $\mathbf{f}$, and a neighborhood of $\mathbf{x}_0$, there is a point $\mathbf{y}_0 = \mathbf{f}(\mathbf{x}_0)$ in the feature space and a corresponding neighborhood of $\mathbf{y}_0$. The dissimilarity at $\mathbf{x}_0$ is a measure of the *density* of space in the neighborhood of $\mathbf{x}_0$ after it is mapped to $\mathbf{y}_0$.

If the resulting space is dense, it will indicate that the neighborhood of $\mathbf{x}_0$ has low dissimilarity as shown in Fig. 2. There appear to be a number of ways to measure this notion of similarity on a multi-valued function. We propose a dissimilarity that is computed by trace of the square of the Jacobian. If J is the Jacobian of **f**, then the dissimilarity is

$$\mathcal{D}\,\mathbf{f} = \mathrm{Tr}(J \cdot J^t) \tag{5}$$

In the case of m = 1, this expression is the gradient magnitude squared, as in edge-affected diffusion. It can be shown that the square root of this expression is a norm in the precise sense, and therefore can be thought of as a distance measure that measures the "length" of the matrix J.

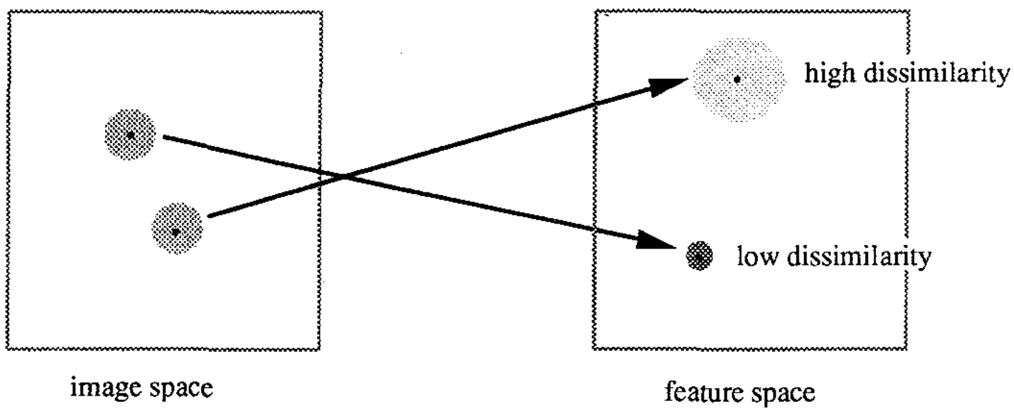

Fig. 2. Every neighborhood in the image space maps onto a neighborhood in the feature space. The 'size' of the neighborhood in the feature space indicates dissimilarity.

This approach has several advantages over statistical pattern recognition approaches mentioned above. First, because the feature space is used in conjunction with the diffusion equation, the process combines information about the image space with position in the feature space. As the process evolves pixels that are nearby in image space will be drawn together in the feature space, except in cases where the dissimilarity is large enough to reduce the conductance. The result is a clustering that is sensitive to the spatial cohesiveness of the image space. Gerig [9] has shown experimentally that a pair of coupled nonuniform diffusion equations when applied to dual-echo MRI, can improve the distinction between clusters of pixels in the two dimensional feature space. Mult-valued diffusion as described here constructs a generalized framework for such systems of coupled diffusion equations.

The second advantage is that the dissimilarity depends only on the differential structure of **f**, and therefore does not rely on the global structure of the feature space. The dependence on the local structure of the feature space allows for a great deal of flexibility for transformations within that space. The dissimilarity measure can be generalized to allow for local coordinate transformations that are not required to preserve global notions of distance. If one considers coordinate transformations (rotations and rescaling of axis) to be changes in the relative importance of features in the calculation of distance, then the dissimilarity measure could allow the relative importance of various features to vary with position in the feature space. If $\Phi(\mathbf{y})$ is the local coordinate transformation, then the generalized dissimilarity becomes:

$$\mathcal{D}(\Phi)\mathbf{f} = \mathrm{Tr}[(\Phi \cdot \mathrm{J}) \cdot (\Phi \cdot \mathrm{J})^{t}] \tag{6}$$

The practical implications of this can imagined by considering a multi-valued data set gathered by some imaging device. Suppose that the relative importance of each feature depends not only on it's own value but the values of the other features. Such behavior could be accounted for by choosing the proper Φ. At this point it is still not evident how one would go about choosing Φ given set of image data. It is conceivable that Φ should depend on the statistical properties (local and global) of the feature space, the nature of the imaging device, and the task at hand. These are issues for future research. Instead, the following discussion will examine how these ideas can be applied to multi-valued data derived from a single-valued intensity function.

Geometry Limited Diffusion

The goal is characterize the shape of the local intensity surface of a single valued function at every point in the image, and represent this shape through a finite set of scalar values. A very convenient way to do this is to use the power series expansion of the image intensity in Cartesian coordinates. If $I(x,y)$ is the original image ($n = 2$), then the local surface at each point in the image can be represented by a the function

$$f(x', y', x, y) = I|_{x,y} + I_{x'}|_{x,y}\, x' + I_{y'}|_{x,y}\, y' + I_{x'y'}|_{x,y}\, x'y' + (1/2)I_{x'x'}|_{x,y}\, x'^2 + (1/2)I_{y'y'}|_{x,y}\, y'^2 + \dots \tag{7}$$

The x′ and y′ are local coordinates that exist at each choice of x and y. The function, f(x′, y′, x, y), is the same function at every x and y, except that it is expressed in local coordinates. Thus f(0, 0, x, y) = I(x,y) as can be seen in Eq. (7). In practice, the series can be truncated in order to produce a local approximation to the surface, and to produce a finite number of terms. Truncated versions of this function are no longer the same at each point in space - at each point there is a distinct local approximation to the surface structure. The coefficients of the truncated expansion can be considered a multi-valued function of the original image space.

$$\mathbf{f}(x,y) = \{I|_{x,y}, I_{x'}|_{x,y}, I_{y'}|_{x,y}, I_{x'y'}|_{x,y}, (1/2)I_{x'x'}|_{x,y}, (1/2)I_{y'y'}|_{x,y}, \dots\} \tag{8}$$

If the local coordinates x′ and y′ are aligned with the image coordinates, then **f**(x,y) is the set of partial derivatives (to within a constant factor) of I(x, y).

$$\mathbf{f}(x,y) = \{I(x,y), I_x(x,y), I_y(x,y), I_{xy}(x,y), (1/2)I_{xx}(x,y), (1/2)I_{yy}(x,y), \dots\} \tag{9}$$

The range of this function is a feature space. Distances in the this feature space are measurements of the similarity of the underlying truncated power series expansion. In the case of a digital image, these functions will all be sampled on a grid. From this point of view, methods which attempt to segment images by measuring the gradient magnitude, are using only the first term in the series expansion.

The strategy adopted in this paper is to make a set of measurements on the original image in order to create a multi-valued function. These measurements are obtained by applying derivatives of Gaussians [3] at a particular scale that is chosen so as not to compromise interesting visual features. The resulting function is that of Eq. (9) with a finite number of terms that correspond to the set of initial measurements made on the image. The measurements are treated as feature vectors in the multi-valued diffusion equation, with a choice of Φ that captures geometric properties that are of interest and a choice of scale in the diffusion equation that eliminates unwanted fluctuations in the feature measurements. After some appropriate amount of processing, further measurements can be made on the resulting resulting set of features, and decisions about presence of certain geometric quantities can be made on the basis of those measurements. The goal is to divide the image into distinct regions that share a set of geometric properties (features) so that the boundaries between these regions form visually interesting areas. This strategy is depicted in Fig. 3.

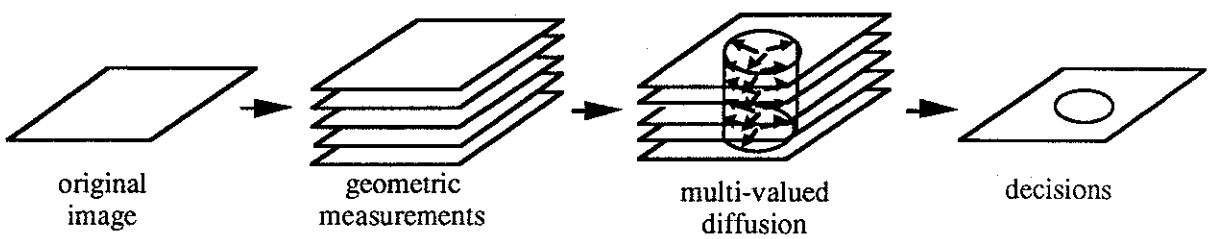

Fig. 3. Geometry-limited diffusion in which features are geometric measurements made on a single valued image.

Creases and First Order Geometry

Because this approach has already been studied in the zero order case (edge effected diffusion based on gradient magnitude), a logical next step is to include first order derivatives as features. The resulting feature space is three dimensional. The remainder of this section will investigate the use of first order geometric features for the purpose of finding 'creases' (sometimes called 'ridges').

Pizer [10, 18, 19] has shown that creases and the corresponding flank regions are very powerful for forming hierarchical segmentations of medical images. Although the notion of creases has no one precise definition (there are several) there is a reliable intuition that one can test via a simple experiment. Imagine an image as an intensity surface, and suppose that the intensity axis is aligned with a gravitation field so that the direction of increasing intensity is 'up'. If one places drops of water on the surface, then creases are places where a drop of water tends to split and run in two or more different directions. Likewise, one could turn the surface over and repeat the experiment. The results in the first case are 'ridges' and in the second case 'valleys'. Of course if the surface is continuous, every place on the image except local extrema has exactly one downhill direction. However, if one considers the higher order structure of the surface (a larger drop of water) then one can imagine that creases are places that have a large disparity of gradient directions in an immediate neighborhood. If one characterizes flanks as contiguous regions that have the same (or nearly the same) uphill direction, then creases can be thought of as places where flank regions meet.

This definition of crease is explicitly constructed to be independent of image intensity. For this reason, it is appropriate and convenient for this discussion to not include the zero order term as one of the features - including it would have no effect on the results. The resulting space is two dimensional, and consists of the features I_x and I_y. The local coordinate transformation, Φ, which determines the way local distances are measured in the feature space, is chosen specifically to capture the notion of creases. The appropriate choice of Φ can best be understood by considering the feature space $\{I_x, I_y\}$ as it is described by polar coordinates (Fig. 4). Each position, (x, y), in the image has a position in this feature space. The polar coordinates of each point in this feature space have a geometric interpretation. The magnitude of the gradient vector at that point in the image is ρ, and θ is the direction of that vector.

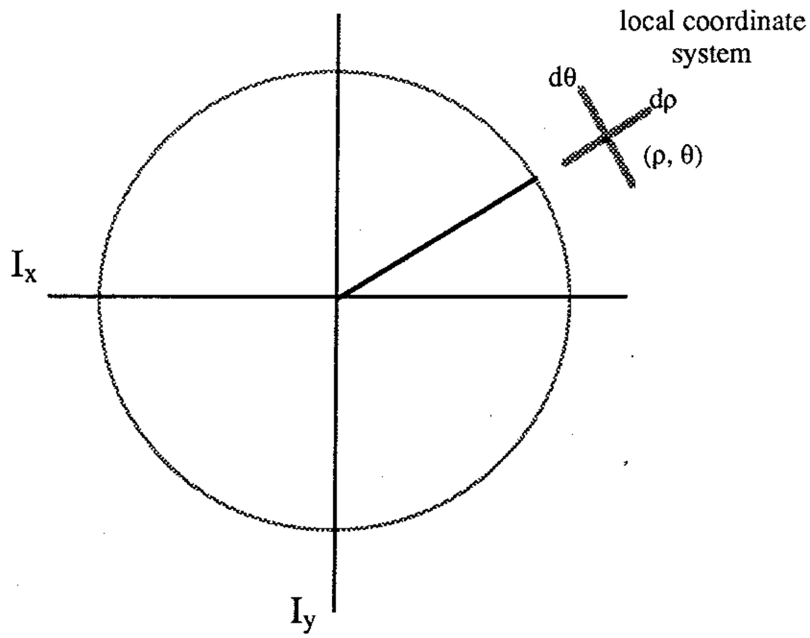

Fig. 4. First order feature space as represented in polar coordinates. Every point in this feature space has a local coordinate system that is aligned with the dq and dr directions.

It is convenient to express Φ as a rotation, R, followed by a scaling S. The matrix R is a rotation of the local feature space coordinates into a coordinate system which is aligned with $d\rho$ and $d\theta$ in the polar coordinates. The matrix S controls the scaling, and thereby the relative effects, of each of these properties. This allows for control of the amount of influence that the gradient magnitude and gradient direction have on the dissimilarity operator. The notion of creases described above does not depend on the values of gradient magnitudes, only on direction. Therefore differences between pixels in the radial direction will not contribute to the dissimilarity. This has the effect of locally collapsing the feature space onto a circle with its center at the origin. The perpendicular component, $d\theta$, will be scaled by the inverse of the radius so that small changes in the in the $d\theta$ direction capture changes in the angle. The result is

$$\Phi = S \cdot R = \begin{bmatrix} \frac{1}{\rho} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} -\sin\theta & \cos\theta \\ \cos\theta & \sin\theta \end{bmatrix} = \begin{bmatrix} \frac{-\sin\theta}{\rho} & \frac{\cos\theta}{\rho} \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} \frac{-\tilde{I}_x}{(\tilde{I}_x^2+\tilde{I}_y^2)} & \frac{\tilde{I}_y}{(\tilde{I}_x^2+\tilde{I}_y^2)} \\ 0 & 0 \end{bmatrix} \tag{10}$$

The tilde over the the terms in this expression is to denote that these are not the derivatives of the image, but the values of these features after they have undergone nonuniform diffusion. Because the diffusion is non-linear, the partial derivatives do not commute with the operators in the process. It us understood that $\tilde{I}_x$ and $\tilde{I}_y$ are from the same time slice. It is not clear that $\tilde{I}_x$ and $\tilde{I}_y$ represent the derivatives of any image because after some amount of blurring the mixed partials are not necessarily equal.

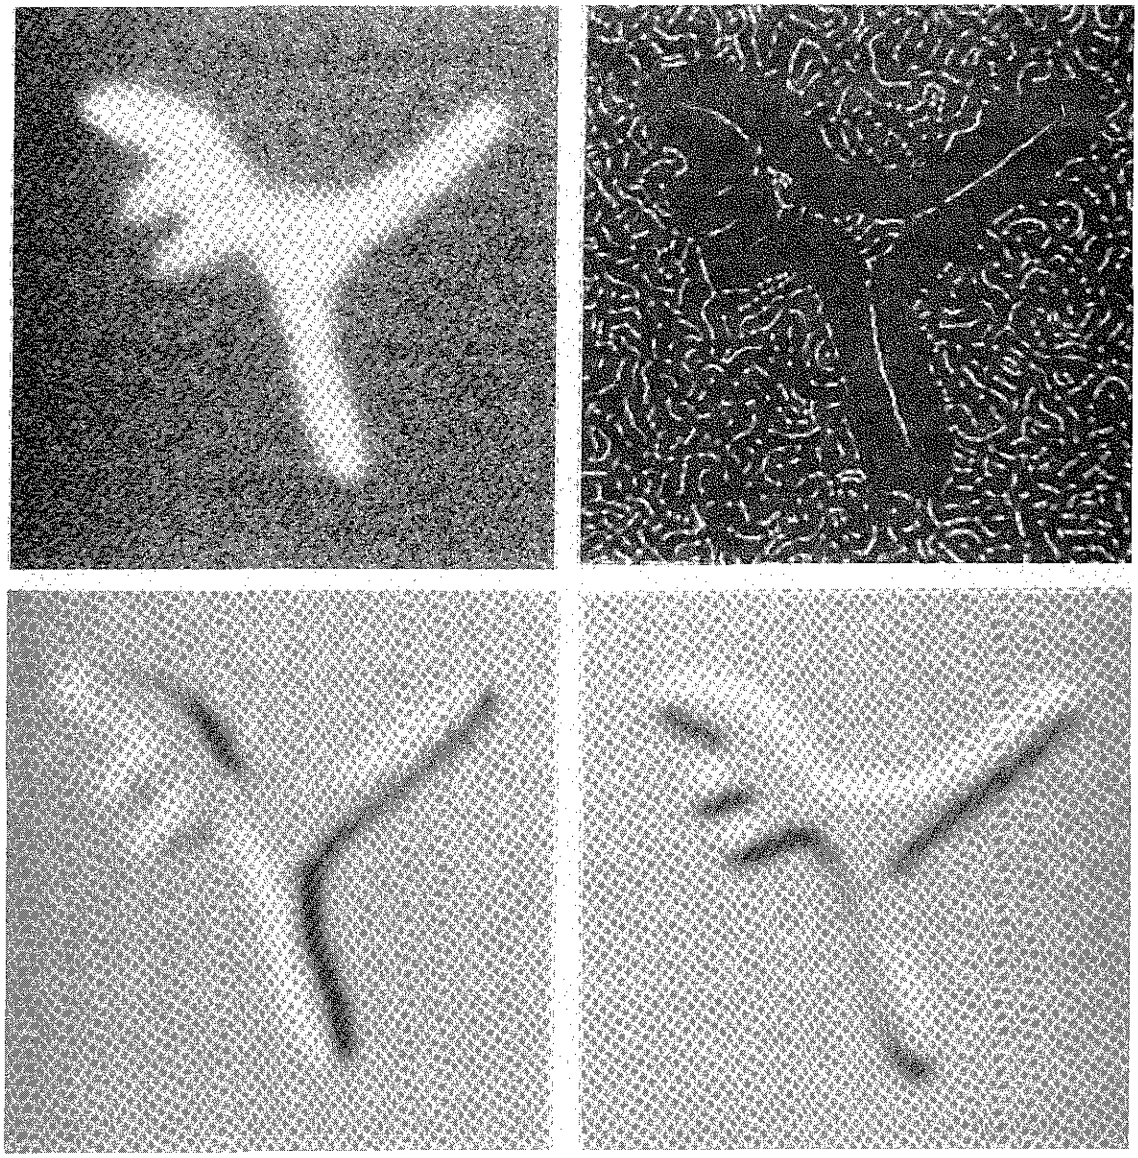

Fig. 5. Clockwise from the upper left: (a)A noisy blob is created by blurring a white on black figure and then adding uniformly distributed random noise. (b) The initial values of the first order features I_x and I_y and (c) the dissimilarity measure on these features.

This form for Φ when applied to the Jacobian is precisely the same as mapping all the points in this feature space to the unit circle, and then computing the Jacobian on the resulting functions. It has several desirable properties. First, it can be computed on discrete images by first dividing each pair of features by the length of the vector that they form, and then computing the local derivatives. Second, the $\mathcal{D}$ that results from the Φ describes above is invariant to any monotonic intensity transformation on the original image. This has the effect of not discriminating between bright and dim features in the original image. Third, when using finite differences, this measure is bounded. The farthest that one pixel can be from a nearby neighbor is on the opposite side of the unit circle (this happens only at critical points of intensity in the original image), a distance of 2. Because this measure is bounded one can construct a conductance function that behaves properly for any number of images regardless of the units one uses to express intensity or the range intensity values within an image. The choice of scale at which to make the geometric measurements, and the range of scales over which to run the diffusion will depend on the image, the properties of the noise, and the size of the structures that one wishes to characterize.

The test image in Fig. 5a is created by drawing a white figure on a black background, and then blurring the result. Uncorrelated random noise is added to the image so that the range of the noise is half the overall intensity of the foreground. Figs. 5b-d show the initial values of the features, I_x and I_y, and the dissimilarity measure at start of the diffusion process. Because the dissimilarity measure normalizes features with respect to the gradient magnitude, areas of the image that were initially flat are susceptible to noise. Fig. 6 shows the features and the dissimilarity measure after processing. The areas of high dissimilarity indicate boundaries between flank regions. The pixels within flank regions have been regularized so that they have virtually the same gradient values.

Treating the resulting features as x and y derivatives of an image allows one to make measurements that indicate second order geometry. Consider the 'diffused' Hessian.

$$\tilde{H} = \begin{bmatrix} \tilde{I}_{xx} \tilde{I}_{yx} \\ \tilde{I}_{xy} \tilde{I}_{yy} \end{bmatrix} \qquad (11)$$

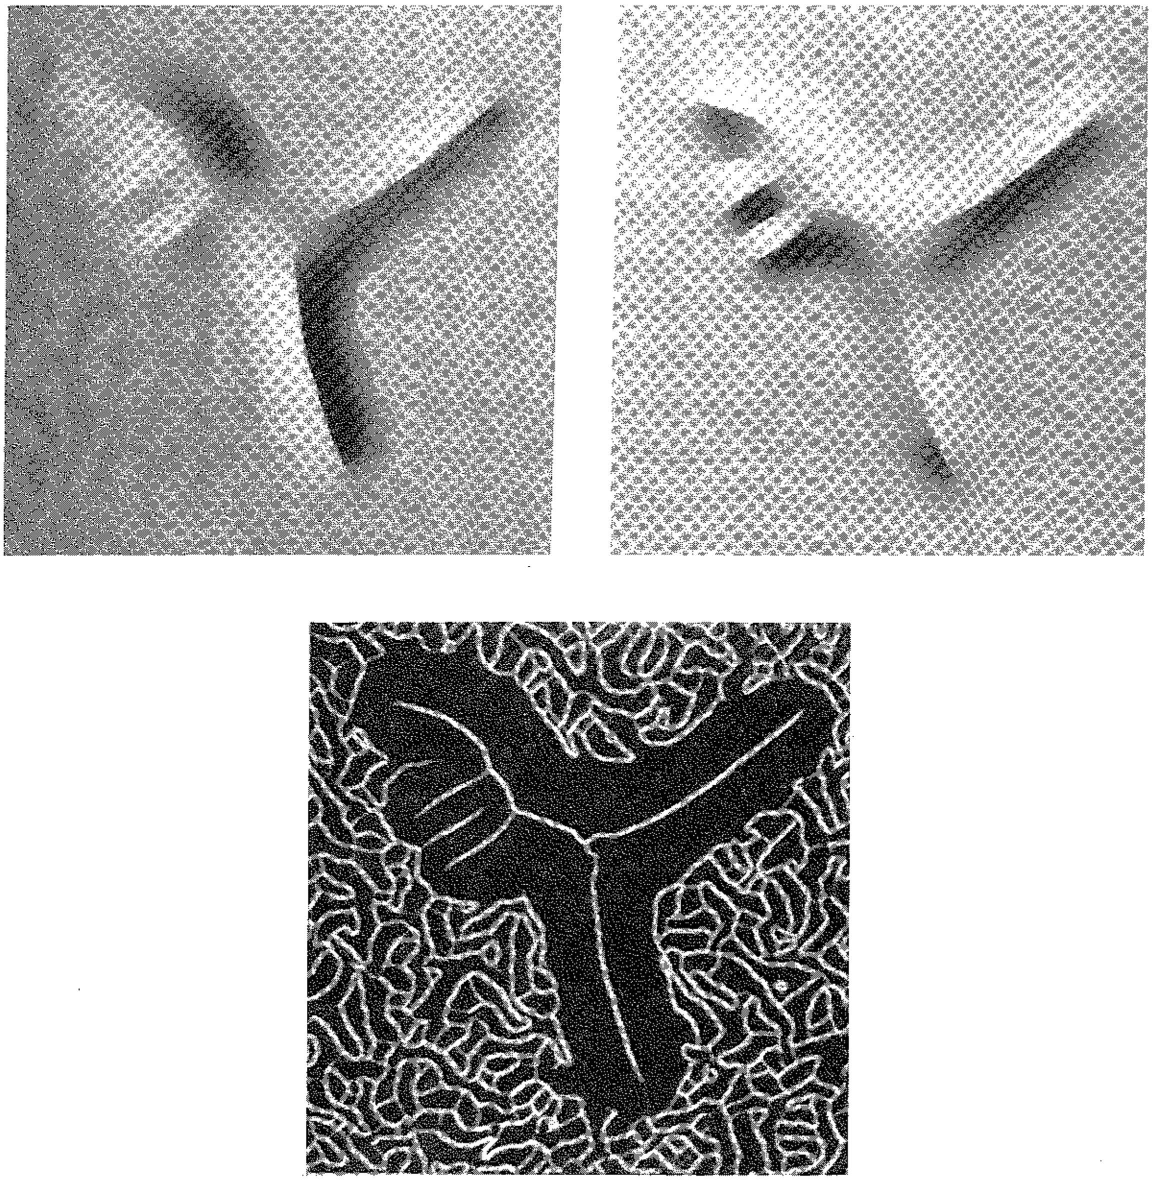

Fig. 6. Clockwise from the upper left: The result values of the first order features (a) I_x and (b) I_y after 150 iterations of the geometry-limited diffusion process and (c) the dissimilarity measure on these features.

Fig. 7 shows the images that result from computing $Tr[\tilde{H}\tilde{H}]$ and $Tr[\tilde{H}]$ respectively. While $Tr[\tilde{H}\tilde{H}]$ indicates the deviation from flatness, a possible quantitative measure of 'creaseness', $Tr[\tilde{H}]$ (also the mean curvature) can be used to decide between between ridges and valleys. It is worth noting that these quantities could be computed on the original image, or some gaussian blurred version of that image. However, such measurements on noisy or blurred images do not offer an immediate and reliable means of making decisions about the presence of second order features. By regularizing the first order information and confining second order variation to a limited number of boundaries by nonuniform diffusion, one can make a number of accurate decisions about the existence of second order features.

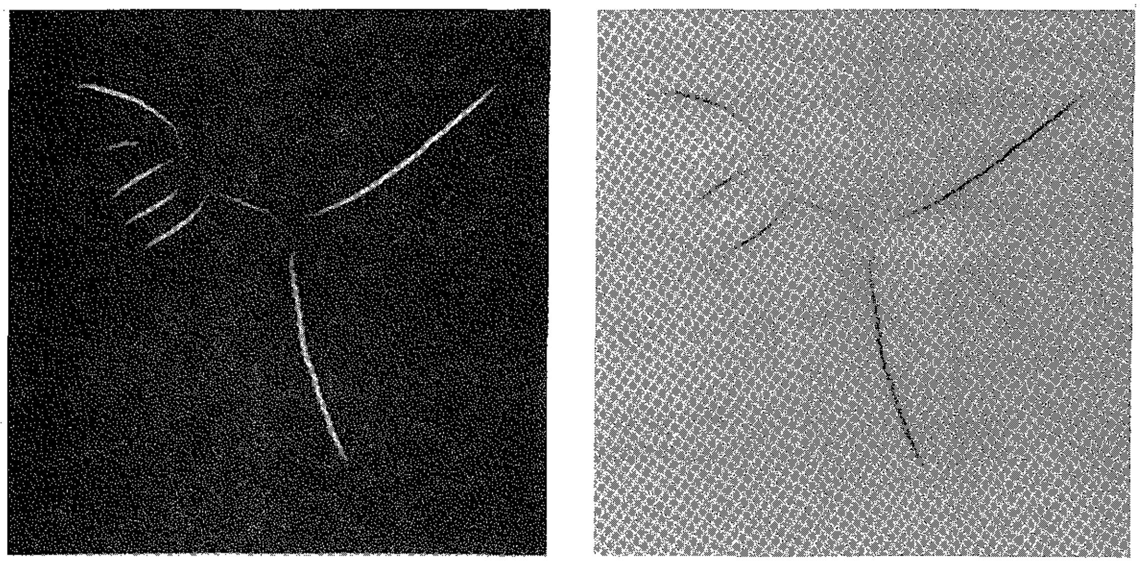

Fig. 7. From the features shown in Fig. 6 the (a) trace of the square of the diffused Hessian and (b) the trace of the Hessian.

Combining Information of Zero and First Orders - Corners

In the example above the image intensity was not used as one of the features. It is conceivable that there exist a three dimensional feature space and an associated Φ that would allow segmentation on the basis of image intensity and gradient information in order to produce visually interesting regions and boundaries that do include image intensity.

Consider the visual feature called a 'corner'. The difficulty of a finding corners using geometry-limited diffusion is that corners in 2D are zero order sets - they are points. Therefore, they cannot form boundaries between geometric patches. If one were to define a dissimilarity measure to capture corners, features in the surrounding space would most likely flow around these points and undermine the local changes in geometry which are indicative of the corner. An alternate strategy is to compute and process zero and first order feature spaces separately and then combine the information that results from each of these processes. A reasonable description of a corner is a place on the edge of an object where the boundary turns very sharply. In order for boundaries to turn suddenly, the gradients along those boundary must twist suddenly, which is precisely the definition used for creases above. A corner is a place in the image that is both a crease and an edge. This suggests that one could obtain corners by combining the results of the crease calculations with edges discerned from zero order diffusion. Because a corner must have both a high gradient (derivative of zero order quantity) and a rapid rapid change in gradient direction (derivative of a first order quantity) it is natural to multiply the measurements that result from zero order diffusion with those of the first order diffusion described above. This strategy is depicted in Fig. 8.

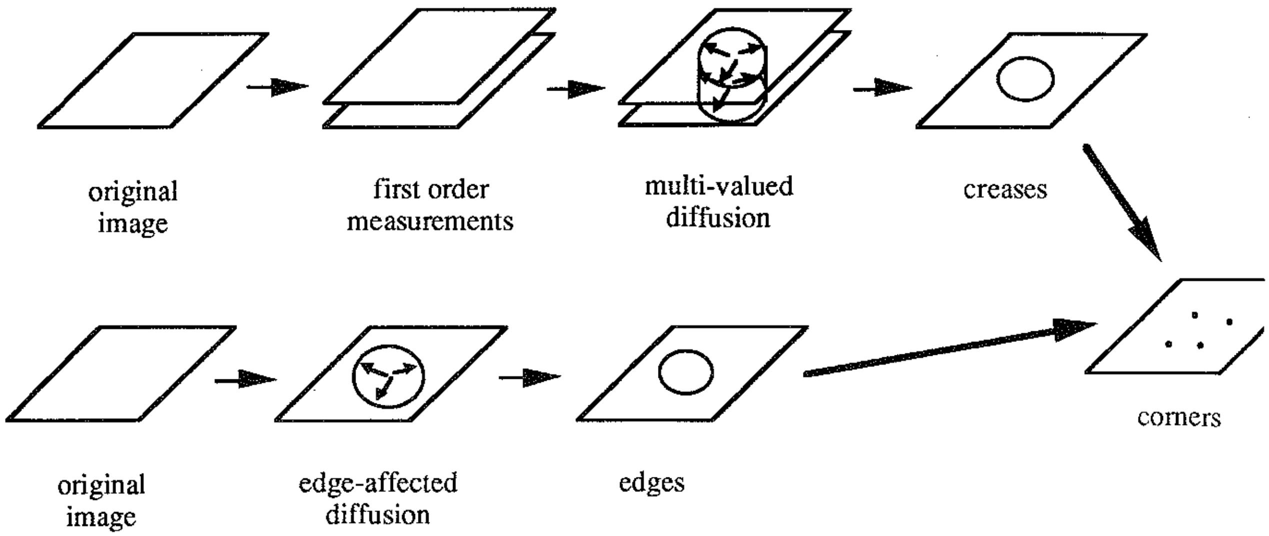

Fig. 8. The strategy for finding corners is to isolate the zero order and first order patches separately and then combine the results.

The sample image in Fig. 9a is a white hexagon one a black background. Random noise as described above has been added to the image. The value of $Tr[\tilde{H}\tilde{H}]$ that results from the first order diffusion processing and the value of $\nabla\tilde{I} \cdot \nabla\tilde{I}$ (or $Tr[\tilde{J}\tilde{J}]$) that results from the zero order processing are shown in Figs. 9b-c. The image of Fig. 9d is the result of multiplying the pixel values of the images in Figs. 9b and 9c. The corners that result from a simple threshold of Fig. 9d are accurate to within 3 pixels.

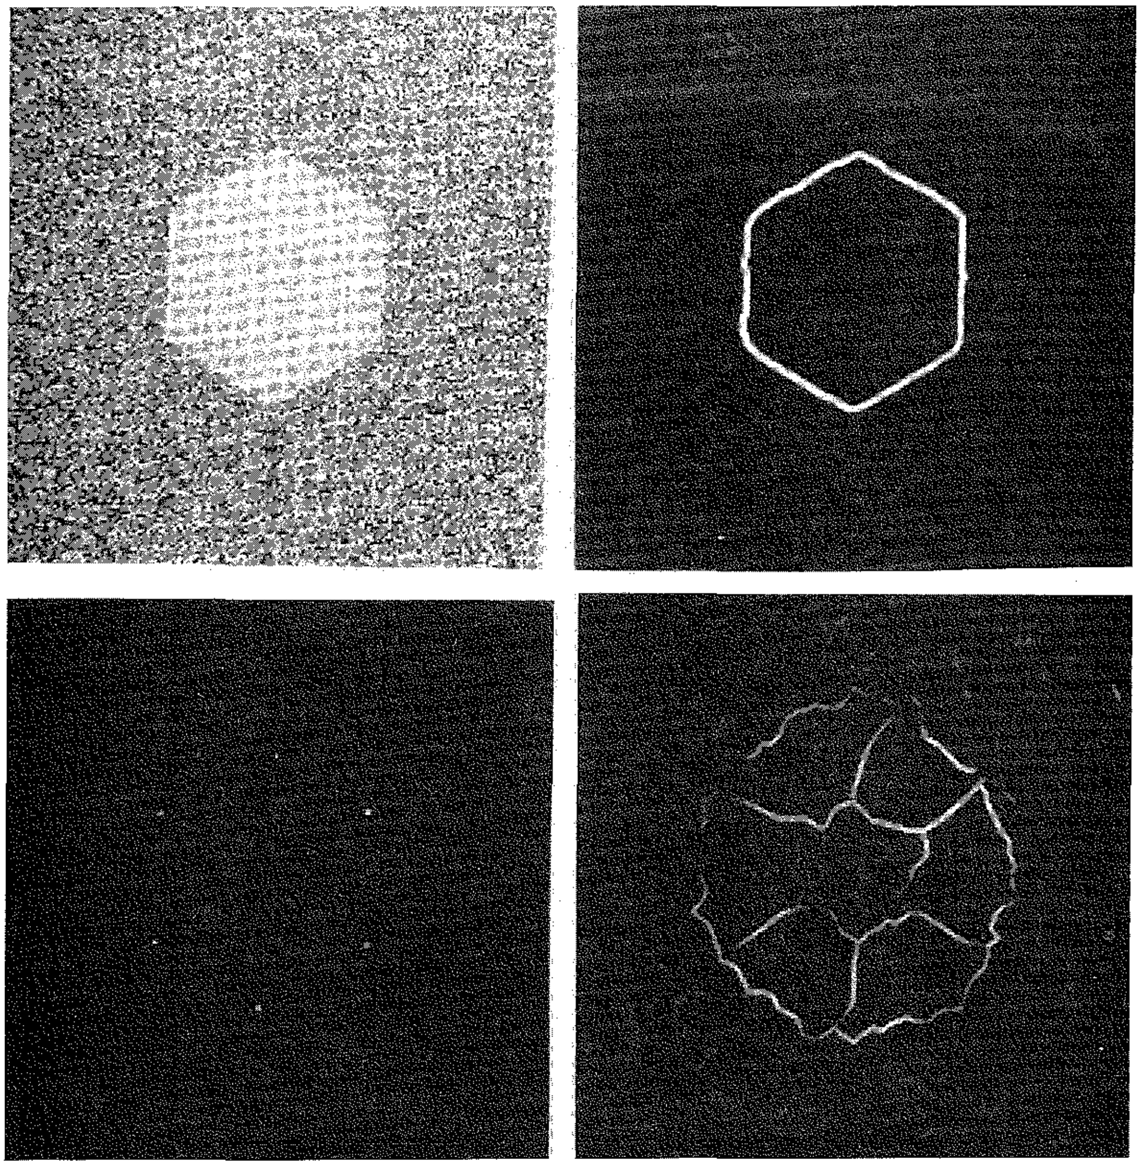

Fig. 9. Clockwise from the upper left: (a) A test image composed of white hexagon on a black background with additive random noise. (b) The gradient squared after zero order diffusion. (c) The creases that result from first order diffusion. (d) The result of multiplying the edge and crease measures in order to detect corners.

Conclusions

Edge-affected diffusion can be generalized to capture regularities in images across multiple features simultaneously. Interpreting multiple scalar values at each point in the image as a position in a feature space provides a generalized dissimilarity operator that controls the conductance in the diffusion of these features. The gradient magnitude is an example of this operator in the case when there is only one feature. The dissimilarity operator incorporates a flexible notion of distance in the feature space, and it allows for local transformations the feature space that do not require global interpretations of distance.

Geometric measurements made on single valued images can be treated as a multiple features in the multi-valued diffusion. In particular, measurements that consist of mixed partial derivatives of the image can be used to characterize the local surface structure as described by a power series expansion. An appropriate choice for local distance in the feature space can produce a dissimilarity measure which captures boundaries between regions that share interesting geometric properties. Nonuniform diffusion blurs these measurements within boundaries, but enhances the boundaries themselves. The result is a distinct set of patches that have well defined boundaries. The richness of higher order data allows one to calculate a variety of measurements on the patches, and their boundaries.

Creases can be described as boundaries between regions that have similar gradient directions. Application of multi-valued nonuniform diffusion to a first order (two dimensional) feature space provides a robust and accurate means of locating creases. Combining crease information with changes in zero order information (intensity) provides a means of accurately detecting corners.

The richness of higher order information holds a great deal of opportunity for this approach to processing image. The difficulty lies in interpreting higher order information in a manner which produces useful boundaries. Future work will explore the use of second order information in order to locate the boundaries of objects ("edges"). This work will also explore conceptual frameworks which could systematically produce useful dissimilarity operators for higher order feature spaces.

Acknowledgements

I would like to thank Dr. Stephen Pizer for his help and guidance throughout this work. Special thanks also go to Dr. James Coggins, Dr. James Damon, Dr. David Eberly, and Dr. Robert Gardner for their useful input.

This research has been supported by NIH grant #P01 CA47982

References

1. J. J. Koenderink, The Structure of Images, *Biol. Cybern.* **50**, 1984, 363-370.

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

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

4. 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.

5. 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.

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

7. 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.

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

9. G. Gerig et. al., Nonlinear Anistropic Filtering of MRI Data, Technical Report BIWI-TR-124, Institute for Communication Technology, Image Science Division, ETH-Zentrum, Surich, Switzerland, 1991.

10. S. M. Pizer, J. M. Coggins, and C. A. Burbeck, Formation of Image Objects in Human Vision, to appear in *Proceedings 1991 Conference on Computer Assisted Radiology*, Springer-Verlag, Berlin, 1991

11. A. K. Jain and R.C. Dubes, *Algorithms for Clustering Data*, Prentice-Hall, New Jersey, 1988.

12. J. R. Jensen, *Introductory Digital Image Processing: A Remote Sensing Perspective*, Prentice-Hall, 1988.

13. R. O. Duda and P. E. Hart, in *Pattern Classification and Scene Analysis*, John Wiley and Sons, New York, 1973.

15. G. B. Coleman and H. C. Andrews, Image Segmentation by Clustering, in *Proceedings of the IEEE* **67** (5), 1979, pp. 773-785.

16. J. M. Coggins and A. K. Jain, A spatial filtering approach to texture analysis, *Pattern Recognition Letters* **3**, 1985, pp. 195-203.

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

18. S. M. Pizer et. al., Descriptions of Intensity Structure via Scale and Symmetry, *Proceedings First Conference on Visualization in Boimedical Computing,* IEE Cat. #90TH0311-1, 1990, pp. 94-101.

19. L. M. Lipshitz and S. M. Pizer, A Multiresolution Hierarchical Approach to Image Segmentation Based on Intensity Extrema, *IEEE Trans. PAMI* **12** (6), 1990, pp. 529-540.