Scale-space and edge detection using anisotropic diffusion

Pietro Perona and Jitendra Malik
Electrical Engineering and Computer Science department
University of California at Berkeley
Berkeley, CA 94720
e-mail: perona@united.berkeley.edu

December 20, 1988

Abstract

The scale-space technique introduced by Witkin involves generating coarser resolution images by convolving the original image with a Gaussian kernel, or equivalently by using the original image as the initial condition of a diffusion process. This approach has a major drawback: it is difficult to obtain accurately the locations of the 'semantically meaningful' edges at coarse scales. In this paper we suggest a new definition of scale-space, and introduce a class of algorithms that realize it using a diffusion process. The diffusion coefficient is chosen to vary spatially in such a way as to encourage intra-region smoothing in preference to inter-region smoothing. It is shown that the 'No new maxima should be generated at coarse scales' property of conventional scale space is preserved. As the region boundaries in our approach remain sharp, we obtain a high quality edge detector which successfully exploits global information. Experimental results are shown on a number of images. The algorithm involves simple, local operations replicated over the image making parallel hardware implementation feasible.

1 Introduction

The importance of multi–scale descriptions of images has been recognized from the early days of computer vision e.g. Rosenfeld and Thurston [18]. A clean formalism for this problem is the idea of scale-space filtering introduced by Witkin [19] and further developed in Koenderink[10], Babaud, Duda and Witkin[1], Yuille and Poggio[20], and Hummel[7].

The essential idea of this approach is quite simple: embed the original image in a family of derived images $I(x, y, t)$ obtained by convolving the original image $I_0(x, y)$ with a gaussian kernel $G(x, y; t)$ of variance t:

$$I(x, y, t) = I_0(x, y) * G(x, y; t) \tag{1}$$

Larger values of t, the scale-space parameter, correspond to images at coarser resolutions. See Fig. 1.

As pointed out by Koenderink [10], and Hummel [7] this one parameter family of derived images may equivalently be viewed as the solution of the heat conduction or diffusion equation

$$I_t = c \cdot \triangle I = c \cdot (I_{xx} + I_{yy}) \tag{2}$$

with the initial condition $I(x, y, 0) = I_0(x, y)$, the original image.

Koenderink motivates the diffusion equation formulation by stating two criteria :

1. *Causality :* Any feature at a coarse level of resolution is required to possess a (not necessarily unique) "cause" at a finer level of resolution although the reverse need not be true. In other words, no spurious detail should be generated when the resolution is diminished.

2. *Homogeneity and Isotropy :* The blurring is required to be space invariant.

These criteria lead naturally to the diffusion equation formulation. It may be noted that the second criterion is only stated for the sake of simplicity. We will have more to say on this later. In fact the major theme of this paper is to replace this criterion by something more useful.

It should also be noted that the causality criterion does not force uniquely the choice of a Gaussian to do the blurring, though it is perhaps the simplest. Hummel [7] has made the important observation that a version of the maximum principle from the theory of parabolic differential equations is equivalent to causality. We will discuss this further in section 4.1.

This paper is organized as follows: section 2 critiques the standard scale space paradigm and presents an additional set of criteria for obtaining 'semantically meaningful' multiple scale descriptions. In section 3, we show that by allowing the

diffusion coefficient to be anisotropic, one can satisfy these criteria. In section 4.1, the maximum principle is reviewed and used to show how the causality criterion is still satisfied by our scheme. In section 5 some experimental results are presented. In section 6 we compare our scheme with other edge detection schemes. Section 7 presents some concluding remarks.

2 Weaknesses of the standard scale-space paradigm

We now examine the adequacy of the standard scale–space paradigm for vision tasks which need 'semantically meaningful' multiple scale descriptions. Surfaces in nature usually [12] have a hierarchical organization composed of a small discrete number of levels. At the finest level, a tree is composed of leaves with an intricate structure of veins. At the next level, each leaf is replaced by a single region, and at the highest level there is a single blob corresponding to the treetop. There is a natural range of resolutions (intervals of the scale-space parameter) corresponding to each of these levels of description. Furthermore at each level of description, the regions (leaves, treetops or forests) have well-defined boundaries.

In the standard scale-space paradigm the true location of a boundary at a coarse scale is not directly available at the coarse scale image. This can be seen clearly in the 1-D example in Figure 2. The locations of the edges at the coarse level t_1 are shifted from their true locations. In 2D images there is the additional problem that edge junctions, which contain much of the spatial information of the edge drawing, are destroyed. The only way to obtain the true location of the edges that have been detected at a coarse scale is by tracking across the scale space to their position in the original image. This technique proves to be complicated and expensive [5].

The reason for this spatial distortion is quite obvious–Gaussian blurring does not 'respect' the natural boundaries of objects. Suppose we have the picture of a treetop with the sky as background. The Gaussian blurring process would result in the green of the leaves getting 'mixed' with the blue of the sky, long before the treetop emerges as a feature (after the leaves have been blurred together). Figure 3 show a sequence of coarser images obtained by Gaussian blurring which illustrates this phenomenon. It may also be noted that the region boundaries are generally quite diffuse instead of being sharp.

With this as motivation, we enunciate [17] the criteria which we believe any candidate paradigm for generating multi-scale 'semantically meaningful' descriptions of images must satisfy:

1. **Causality** : As pointed out by Witkin and Koenderink, a scale-space representation should have the property that no spurious detail should be generated passing from finer to coarser scales.

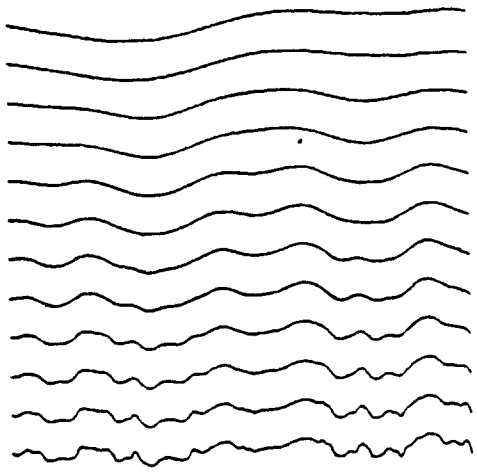

Figure 1: A family of 1-D signals, I(x,t), obtained by convolving the original one (bottom) with Gaussian kernels whose variance increases from bottom to top. (Adapted from Witkin [19]).

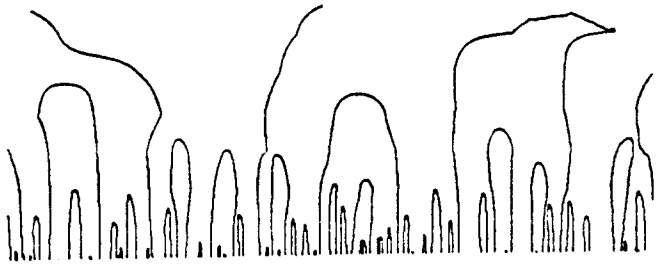

Figure 2: Position of the edges (zeros of the Laplacian with respect to x) through the linear scale space of fig. 1 (Adapted from Witkin [19]).

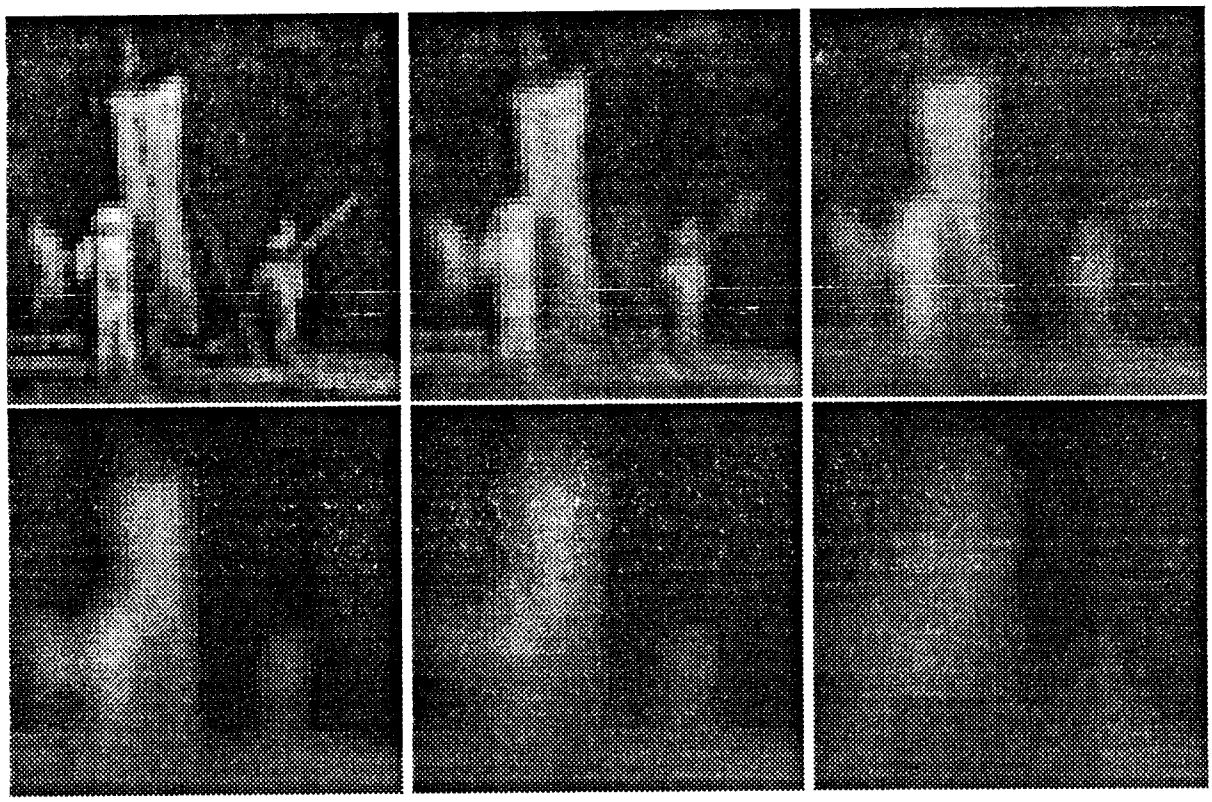

Figure 3: Scale-space (scale parameter increasing from top to bottom, and from left to right) produced by isotropic diffusion. Compare with figure 11.

2. **Immediate Localization** : At each resolution, the region boundaries should be sharp and coincide with the semantically meaningful boundaries at that resolution.

3. **Piecewise Smoothing**: At all scales, intra–region smoothing should occur preferentially over inter–region smoothing. In the tree example mentioned earlier, the leaf regions should be collapsed to a treetop *before* being merged with the sky background.

In the next section we will outline a framework for achieving these objectives.

3 Anisotropic diffusion

There is a simple way of modifying the scale–space paradigm to achieve the objectives that we have put forth in the previous section. In the diffusion equation framework of looking at scale–space, the diffusion coefficient c is assumed to be a constant independent of the space location. There is no fundamental reason why this must be so. To quote Koenderink [10], (pg. 364, left column, l. 19 from the bottom) "... I do not permit space variant blurring. Clearly this is not essential to the issue, but it simplifies the analysis greatly". We will show how a suitable choice of $c(x, y, t)$ will enable us to satisfy the second and third criteria listed in the previous section. Furthermore this can be done without sacrificing the causality criterion.

Consider the anisotropic diffusion equation

$$I_t = div(c(x, y, t)\nabla I) = c(x, y, t)\triangle I + \nabla c \cdot \nabla I \tag{3}$$

where we indicate with *div* the divergence operator, and with ∇, and $\triangle$ the gradient, and laplacian operators with respect to the space variables. It reduces to the isotropic heat diffusion equation $I_t = c\triangle I$ if $c(x, y, t)$ is a constant. Suppose at the time (scale) t, we knew the locations of the region boundaries appropriate for that scale. We would want to encourage smoothing *within* a region in preference to smoothing *across* the boundaries. This could be achieved by setting the conduction coefficient to be 1 in the interior of each region and 0 at the boundaries. The blurring would then take place separately in each region with no interaction between regions. The region boundaries would remain sharp.

Of course, we do *not* know in advance the region boundaries at each scale (If we did the problem would already have been solved!). What can be computed is a current best estimate of the location of the boundaries (edges) appropriate to that scale.

Let $\mathbf{E}(x, y, t)$ be such an estimate: a vector valued function defined on the image which ideally should have the following properties:

1. $\mathbf{E}(x, y, t) = \mathbf{0}$ in the interior of each region.

2. $\mathbf{E}(x,y,t) = K\mathbf{e}(x,y,t)$ at each edge point, where $\mathbf{e}$ is a unit vector normal to the edge at the point, and K is the local contrast (difference in the image intensities on the left and right) of the edge.

Note that the word *edge* as used above has not been formally defined–we mean here the perceptual subjective notion of an edge as a region boundary. A completely satisfactory formal definition is likely to be part of the solution, rather than the problem definition!

If an estimate $\mathbf{E}(x,y,t)$ is available, the conduction coefficient $c(x,y,t)$ can be chosen to be a function $c = g(\|\mathbf{E}\|)$ of the magnitude of $\mathbf{E}$. According to the previously stated strategy $g(\cdot)$ has to be a nonnegative monotonically decreasing function with $g(0) = 1$. This way the diffusion process will mainly take place in the interior of regions, and it will not affect the region boundaries where the magnitude of $\mathbf{E}$ is large.

It is intuitive that the success of the diffusion process in satisfying the three scale-space goals of section 2 will greatly depend on how accurate the estimate $\mathbf{E}$ is as a "guess" of the position of the edges. Accuracy though is computationally expensive and requires complicated algorithms. We are able to show that fortunately the simplest estimate of the edge positions, obtained by looking at the gradient of the brightness function i.e. setting $\mathbf{E}(x,y,t) = \nabla I(x,y,t)$, gives excellent results.

There are many possible choices for $g(\cdot)$, the simplest being a binary valued function. In the next section we show that in case we use the edge estimate $\mathbf{E}(x,y,t) = \nabla I(x,y,t)$ the choice of $g(\cdot)$ is restricted to a subclass of the monotonically decreasing functions.

4 Properties of anisotropic diffusion

We start in section 4.1 by recalling a general result of the partial differential equation theory that guarantees that the causality criterion of scale-space is satisfied by anisotropic diffusion.

In section 4.2 we will show that a diffusion in which the conduction coefficient is chosen locally as a function of the magnitude of the gradient of the brightness function, i.e.

$$c(x,y,t) = g(\|\nabla I(x,y,t)\|) \tag{4}$$

will not only preserve, but also sharpen, the brightness edges if the function $g(\cdot)$ is chosen properly.

4.1 The maximum principle

In this section we show that the solutions to the anisotropic diffusion equation respect the maximum principle, and argue that this implies compliance with the causality requirement.

Call Ω an open bounded set of $\mathbf{R}^n$ (in our case Ω is the plane of the image, a rectangle of $\mathbf{R}^2$), and $\mathbf{T} = (a, b)$ an interval of $\mathbf{R}$. Let $\mathbf{D}$ be the open cylinder of $\mathbf{R}^{n+1}$ formed by the product $\mathbf{D} = \Omega \times \mathbf{T} = \{(x,t) : x \in \Omega, t \in \mathbf{T}\}$. Call $\partial\mathbf{D}$ the boundary of $\mathbf{D}$, $\overline{\mathbf{D}}$ it's closure, and $\partial'\mathbf{D}$, and $\partial''\mathbf{D}$ the side-bottom, and the top portions of $\partial\mathbf{D}$:

$$\partial'\mathbf{D} = \{(x,t) : x \in \Omega, t = a\} \bigcup \{(x,t) : x \in \partial\Omega, t \in \mathbf{T}\}$$

$$\partial''\mathbf{D} = \{(x,t) : x \in \Omega, t = b\}$$

The following theorem holds:

Theorem.

Consider a function $U : \mathbf{R}^{n+1} \longrightarrow \mathbf{R}$ that is continuous on $\overline{\mathbf{D}}$, and twice differentiable on $\mathbf{D} \bigcup \partial''\mathbf{D}$. If U satisfies the differential inequality

$$C(x,t)U_t - c(x,t)\triangle U - \nabla c \cdot \nabla U \leq 0 \tag{5}$$

on $\mathbf{D}$, with : $C : \mathbf{R}^{n+1} \longrightarrow \mathbf{R}^+$ continuous on $\overline{\mathbf{D}}$, and $c : \mathbf{R}^{n+1} \longrightarrow \mathbf{R}^+$ continuous on $\overline{\mathbf{D}}$, and differentiable on $\mathbf{D} \bigcup \partial''\mathbf{D}$, then it obeys the maximum principle, i.e.

$$\max_{\overline{\mathbf{D}}} U = \max_{\partial'\mathbf{D}} U$$

The following proof is adapted from John [9].

Proof.

First consider U satisfying the stricter condition

$$C(x,t)U_t - c(x,t)\triangle U - \nabla c \cdot \nabla U < 0 \tag{6}$$

By hypothesis U is continuous on $\overline{\mathbf{D}}$, a compact set, hence it has a maximum in it. Call $p = (y, \tau)$ this maximum.

Suppose that $p \in \mathbf{D}$. Since U is twice continuously differentiable in $\mathbf{D}$ all the first derivatives of U are equal to zero at p, and the second derivatives in the laplacian are negative, or equal to zero (it is easy to convince oneself that this is true by looking at the Taylor expansion of U about p). This would mean that at p

$$C(x,t)U_t - c(p)\triangle U - \nabla c \cdot \nabla U \geq 0$$

contradicting the hypothesis.

Similarly, if $p \in \partial''\mathbf{D}$ the first derivative with respect to t of U could only be positive or equal to zero, while the first derivatives with respect to the x variables would have to be equal to zero, and the second derivatives with respect to the x variables could only be equal to zero or negative, giving the same inequality at p as above. This would again contradict the hypothesis. So, if U satisfies equation (6), then it obeys the maximum principle.

If U satisfies

$$C(x,t)U_t - c(x,t)\triangle U - \nabla c \cdot \nabla U \leq 0$$

the function $V = U - \lambda(t-a)$ satisfies equation (6), and hence the maximum principle, for any $\lambda > 0$. Observe that $U = V + \lambda(t-a) \leq V + \lambda(b-a)$ on $\overline{\mathbf{D}}$, and because of this

$$\max_{\overline{\mathbf{D}}} U \leq \max_{\overline{\mathbf{D}}} (V + \lambda(b-a)) =$$

$$= \max_{\partial' \mathbf{D}} (V + \lambda(b-a)) \leq \max_{\partial' \mathbf{D}} (U + \lambda(b-a))$$

Letting $\lambda \to 0$ we obtain the thesis. □

Notice that the maximum principle also guarantees that there are no local maxima of U in $\mathbf{D} \bigcup \partial'' \mathbf{D}$. One can in fact use the same technique as in the proof to see that the existence of one at $p \in \mathbf{D} \bigcup \partial'' \mathbf{D}$ would violate the differential inequality.

If a function U satisfies the differential equation

$$C(x,t)U_t - c(x,t)\triangle U - \nabla c \cdot \nabla U = 0 \tag{7}$$

with the hypotheses already stated on the functions $C(\cdot)$ and $c(\cdot)$, the arguments above can be run for U and $W = -U$ proving that both a maximum and minimum principle have to be satisfied.

The diffusion equation (3) is a special case of equation (7) (set $C(x,t) = 1$, and $U = I$), hence the scale-space brightness function $I(x,y,t)$ obeys the the maximum principle provided that the conduction coefficient c never takes negative value (in fact the condition that c does not take negative value where U has a maximum is sufficient) and is differentiable. This is relevant to our discussion because if we identify "features" in the images with "blobs" of the brightness function $I(x,y,t)$ for different values of the scale parameter, then the maximum principle guarantees that no new features are created for increasing t. In fact the birth of a new "blob" would imply the creation of a maximum, or a minimum, that would have to belong either to the interior or the top face $\partial'' \mathbf{D}$, of the domain of interest in the scale space.

Solutions U of equation (3) have an additional property if the conduction coefficient is constant along the space axes: $c = c(t)$. Any spatial derivative of U is then a solution of equation (3) too and thus satisfies the hypotheses of the maximum principle. So the causality criterion is satisfied by all such functions: the components of the gradient, the Laplacian etc. . It is important to notice that this is not true in general for solutions of equation (3) when the conduction coefficient varies in space. We will show in the next section that in fact anisotropic diffusion can increase the contrast (i.e. the magnitude of the gradient) of edges in the image.

4.2 Edge Enhancement

With conventional low-pass filtering and diffusion the price payed for eliminating the noise, and for performing scale space, is the blurring of edges. This causes their detection and localization to be difficult. An analysis of this problem is presented in [4].

Edge enhancement and reconstruction of blurry images can be achieved by high-pass filtering or running the diffusion equation backwards in time. This is an ill-posed problem, and gives rise to numerically unstable computational methods, unless proper regularization is employed [8].

We will show here that if the conduction coefficient is chosen to be an appropriate function of the image gradient we can make the anisotropic diffusion enhance edges while running *forward* in time, thus enjoying the stability of diffusions which is guaranteed by the maximum principle.

We model an edge as a step function convolved with a Gaussian. Without loss of generality, assume that the edge is aligned with the y axis.

The expression for the divergence operator simplifies to:

$$div(c(x,y,t)\nabla I) = \frac{\partial}{\partial x}(c(x,y,t)I_x)$$

We choose c to be a function of the gradient of I: $c(x,y,t) = g(I_x(x,y,t))$ as suggested by equation (4). Let $\phi(\cdot)$ denote the product $c \cdot I_x : \phi(I_x) \equiv g(I_x) \cdot I_x$.

Then the 1-D version of the diffusion equation (3) becomes

$$I_t = \frac{\partial}{\partial x}\phi(I_x) \tag{8}$$

We are interested in looking at the variation in time of the slope of the edge: $\frac{\partial}{\partial t}(I_x)$. If $c(\cdot) > 0$ the function $I(\cdot)$ is smooth, and the order of differentiation may be inverted:

$$\frac{\partial}{\partial t}(I_x) = \frac{\partial}{\partial x}(I_t) = \frac{\partial}{\partial x}(\frac{\partial}{\partial x}\phi(I_x)) = \phi'' \cdot {I_{xx}}^2 + \phi' \cdot I_{xxx} \tag{9}$$

Suppose the edge is oriented in such a way that $I_x > 0$ (see Fig.4). At the point of inflection $I_{xx} = 0$, and $I_{xxx} \ll 0$ since the point of inflection corresponds, or is very close, to the point with maximum slope (see figure 4). Then in a neighbourhood of the point of inflection $\frac{\partial}{\partial t}(I_x)$ has sign opposite to $\phi'(I_x)$. If $\phi'(I_x) > 0$ the slope of the edge will decrease with time; if, on the contrary $\phi'(I_x) < 0$ the slope will increase with time.

Notice that this increase in slope cannot be caused by a scaling of the edge, because this would violate the maximum principle. The edge becomes sharper.

There are several possible choices for $\phi(\cdot)$, for example, $g(I_x) = \frac{C}{1+(\frac{I_x}{K})^{1+\alpha}}$ with $\alpha > 0$ (see fig. 5). Then there exists a certain threshold value related to K, and

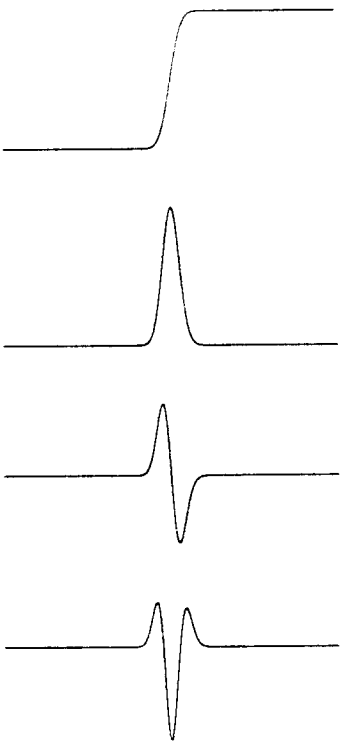

Figure 4: A mollified step edge and its derivatives.

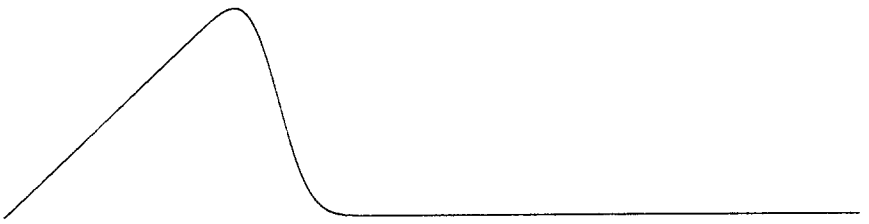

Figure 5: A choice of the function $\phi(\cdot)$ that leads to edge enhancement.

α, below which $\phi(\cdot)$ is monotonically increasing, and beyond which $\phi(\cdot)$ is monotonically decreasing, giving the desirable result of blurring small discontinuities *and* sharpening edges.

5 Experimental Results

Our anisotropic diffusion, scale-space, and edge detection ideas were tested using a simple numerical scheme that is described in this section.

Equation (3) can be discretized on a square lattice, with brightness values associated to the vertices, and conduction coefficients to the arcs. A 4-nearest-neighbours discretization of the Laplacian operator was used:

$$I_{i,j}^{t+1} = I_{i,j}^{t} + \lambda[c_N \cdot \nabla_N I + c_S \cdot \nabla_S I + c_E \cdot \nabla_E I + c_W \cdot \nabla_W I]_{i,j}^{t} \qquad (10)$$

where $\lambda \leq \frac{1}{4}$ for the numerical scheme to be stable, N,S,E,W are the mnemonic subscripts for North, South, East, West, the superscript and subscripts on the square bracket are applied to all the terms it encloses, and the symbol ∇ indicates nearest-neighbour differences:

$$\begin{aligned} \nabla_N I_{i,j} &\equiv I_{i-1,j} - I_{i,j} \\ \nabla_S I_{i,j} &\equiv I_{i+1,j} - I_{i,j} \\ \nabla_E I_{i,j} &\equiv I_{i,j+1} - I_{i,j} \\ \nabla_W I_{i,j} &\equiv I_{i,j-1} - I_{i,j} \end{aligned} \qquad (11)$$

The conduction coefficients are updated at every iteration as a function of the brightness gradient (equation (4)):

$$\begin{aligned} c_{N_{i,j}}^{t} &= g(\|(\nabla I)_{i+\frac{1}{2},j}^{t}\|) \\ c_{S_{i,j}}^{t} &= g(\|(\nabla I)_{i-\frac{1}{2},j}^{t}\|) \\ c_{E_{i,j}}^{t} &= g(\|(\nabla I)_{i,j+\frac{1}{2}}^{t}\|) \\ c_{W_{i,j}}^{t} &= g(\|(\nabla I)_{i,j-\frac{1}{2}}^{t}\|) \end{aligned} \qquad (12)$$

The value of the gradient can be computed on different neighbourhood structures achieving different compromises between accuracy and locality. The simplest choice consists in approximating the norm of the gradient at each arc location with the absolute value of its projection along the direction of the arc:

$$\begin{aligned} c_{N_{i,j}}^{t} &= g(|\nabla_N I_{i,j}^{t}|) \\ c_{S_{i,j}}^{t} &= g(|\nabla_S I_{i,j}^{t}|) \\ c_{E_{i,j}}^{t} &= g(|\nabla_E I_{i,j}^{t}|) \\ c_{W_{i,j}}^{t} &= g(|\nabla_W I_{i,j}^{t}|) \end{aligned} \qquad (13)$$

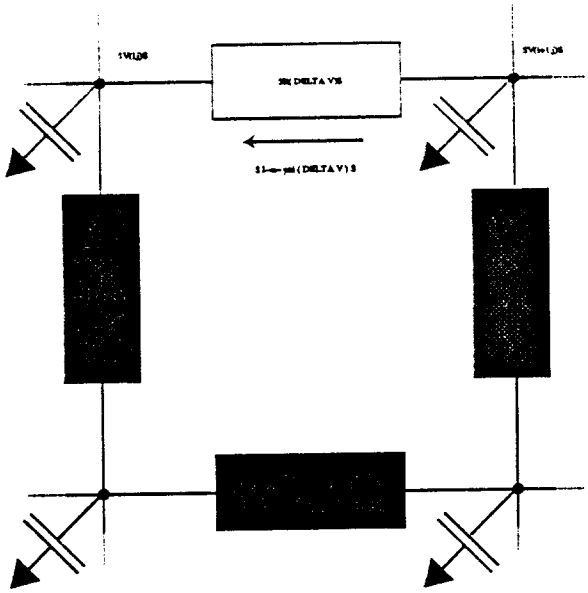

Figure 6: The structure of a network realizing the implementation of anisotropic diffusion described in section 5, and more in detail in [16]. The charge on the capacitor at each node of the network represents the brightness of the image at a pixel. Linear resistors produce isotropic linear diffusion. Resistors with an I-V characteristic as in Figure 5 produce anisotropic diffusion.

This scheme is not the exact discretization of equation (3), but of a similar diffusion equation in which the conduction tensor is diagonal with entries $g(\mid I_x \mid)$ and $g(\mid I_y \mid)$ instead of $g(\|\nabla I\|)$ and $g(\|\nabla I\|)$. This discretization scheme preserves the property of the continuous equation (3) that the total amount of brightness in the image is preserved. Additionally the "flux" of brightness through each arc of the lattice only depends on the values of the brightness at the two nodes defining it, which makes the scheme a natural choice for analog VLSI implementations [16]. See fig. 6. Less crude approximations of the gradient yielded perceptually similar results at the price of increased computational complexity.

The numerical scheme used to obtain the pictures in this paper is the one given by equations (10), (11), (13), using the original image as the initial condition, and adiabatic boundary conditions, i.e. setting the conduction coefficient to zero at the boundaries of the image. A constant value for the conduction coefficient c (i.e. $g(\cdot) \equiv 1$) leads to Gaussian blurring (see figure 3).

Different functions were used for $g(\cdot)$ giving perceptually similar results. The images in this paper were obtained using

$$g(\nabla I) = e^{(-(\frac{\|\nabla I\|}{K})^2)}$$

(Fig. 8), and

$$g(\nabla I) = \frac{1}{1 + (\frac{\|\nabla I\|}{K})^2}$$

(Fig. 11,13). The scale-space generated by these two functions is different; This point will be analyzed in a forthcoming paper.

The constant K was fixed either by hand at some fixed value (See Fig. 7, 8, 13), or using using the "noise estimator" described by Canny [4]: a histogram of the absolute values of the gradient throughout the image was computed, and K was set equal to the 90% value of its integral at every iteration (see Fig. 11(b)).

The computational scheme described in this section has been chosen for its simplicity. Numerically efficient implementations of anisotropic diffusion are under study and will be described in a future paper.

6 Comparison with other edge detection schemes

This section is devoted to comparing the anisotropic diffusion scheme that we present in this paper with previous work on edge detection, image segmentation and image restoration.

We will divide edge detectors in two classes: fixed-neighbourhood edge detectors, and energy/probability "global" schemes.

6.1 Fixed neighbourhood detectors

This class of detectors makes use of local information only–they typically examine a small window of the image and try to be clever about deciding if and where there is an edge. This decision is ambiguous and difficult. In our approach, the image is being blurred in such a way that the pixel values at an edge come to reflect the values of their near and distant neighbors from the *same* region.

We pick Canny's scheme [4] as a representative of this class of detectors. The image is convolved with directional derivatives of a Gaussian – the idea is to do smoothing parallel to the edge and thus reduce noise without blurring the edge too much. There are two major difficulties: (a) the inevitable tradeoff [4] between localization accuracy and detectability that comes from using linear filtering; (b) the complexity of combining outputs of filters at multiple scales. Anisotropic diffusion is a nonlinear process, hence in principle is not subject to limitation (a). The complication of multiple scale, multiple orientation filters is avoided by locally adaptive smoothing.

We can thus summarize the advantages of the scheme we propose over linear fixed-neighbourhood edge detectors:

Locality - The shape and size of the neighbourhood where smoothing occurs are determined locally by the brightness pattern of the image, and adapt to the

a)

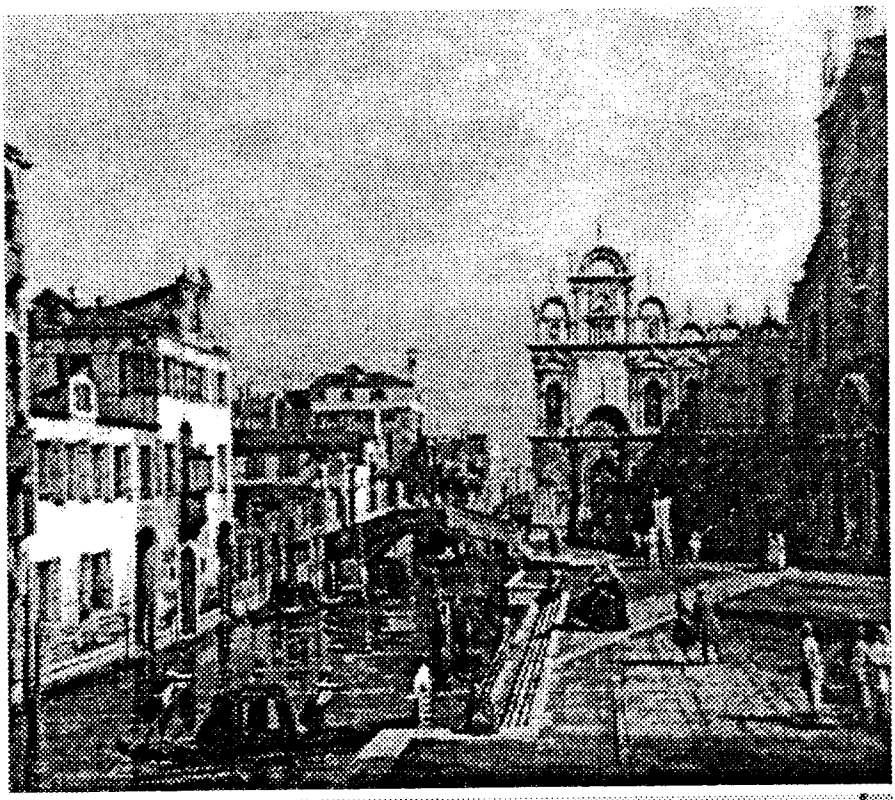

b)

Figure 7: Effect of anisotropic diffusion (b) on the Canaletto image (a) [3].

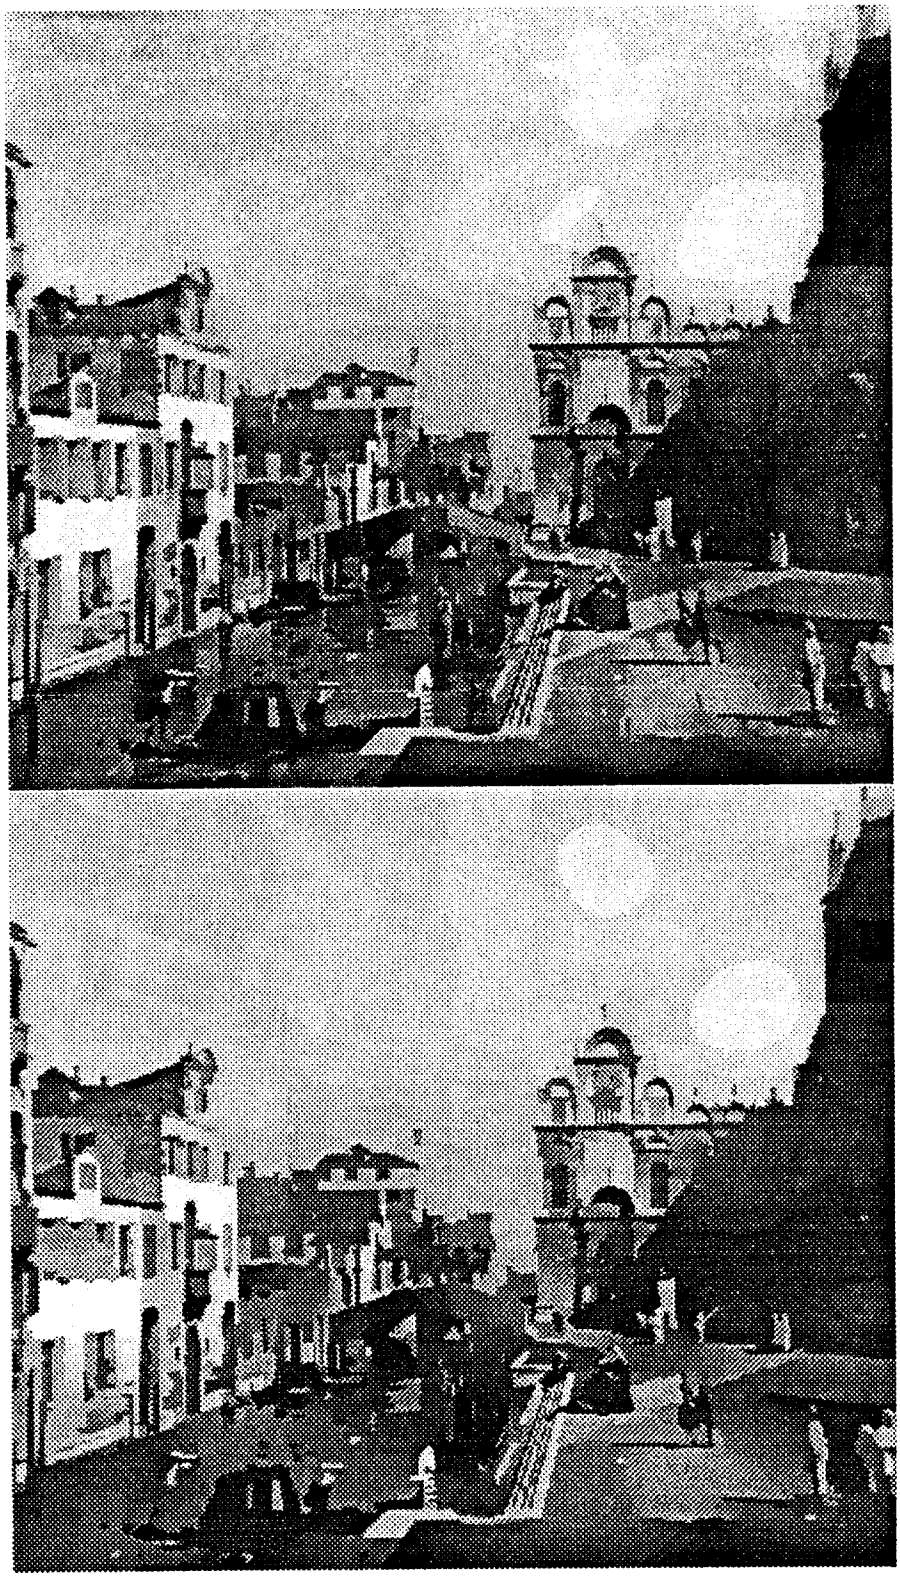

Figure 8: Two intermediate diffusion stages on the Canaletto image.

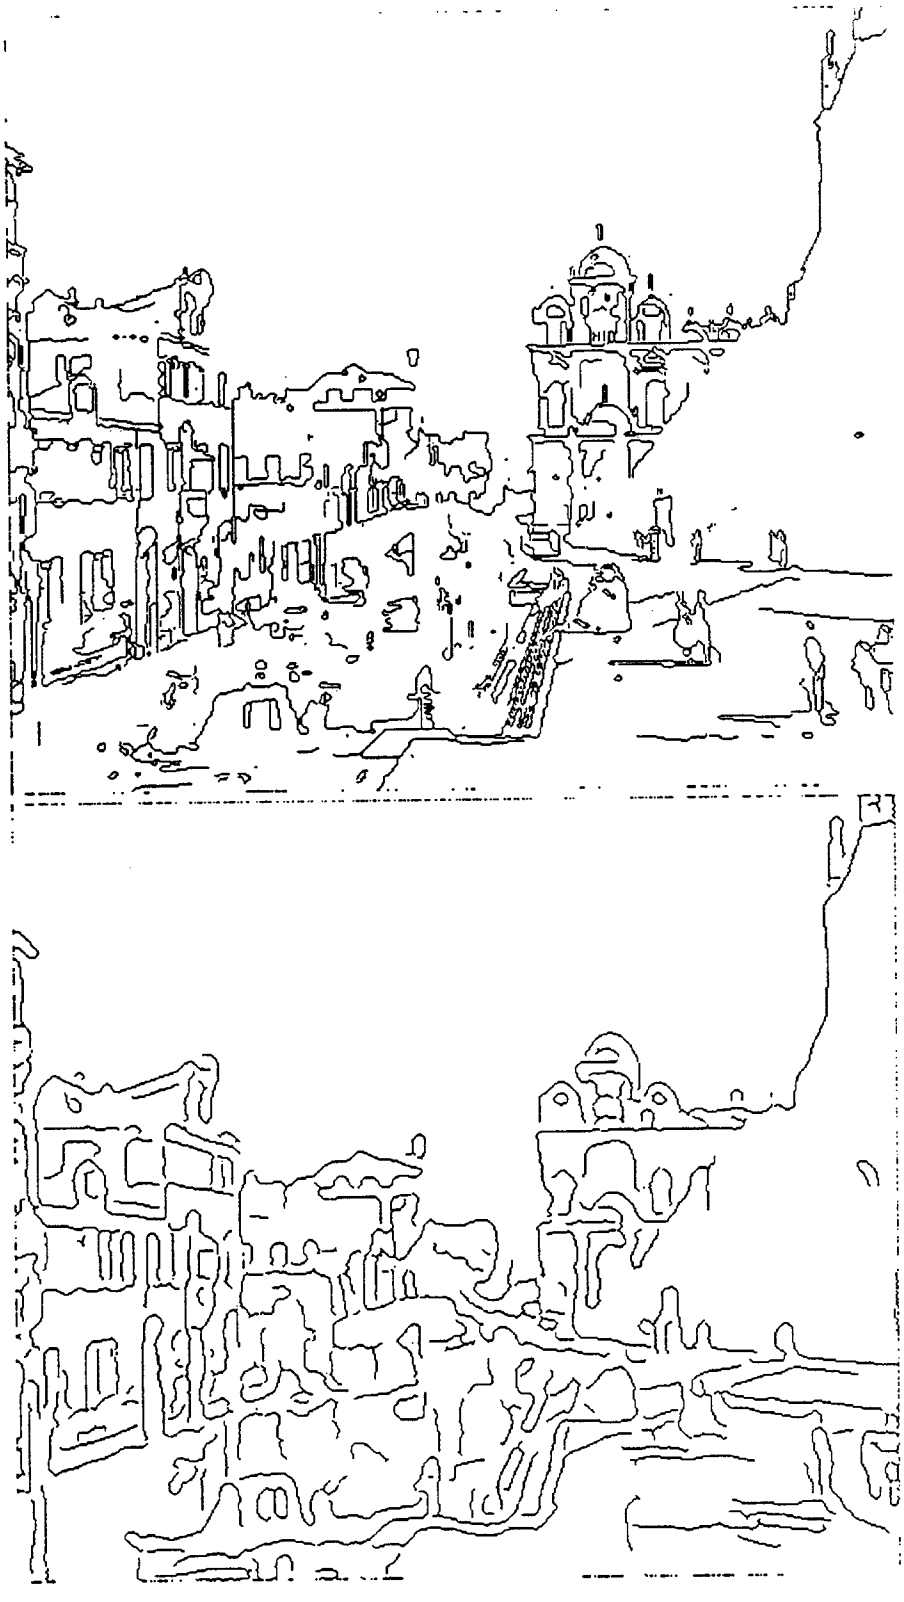

Figure 9: Edges detected using (a) anisotropic diffusion and (b) Gaussian smoothing (Canny detector).

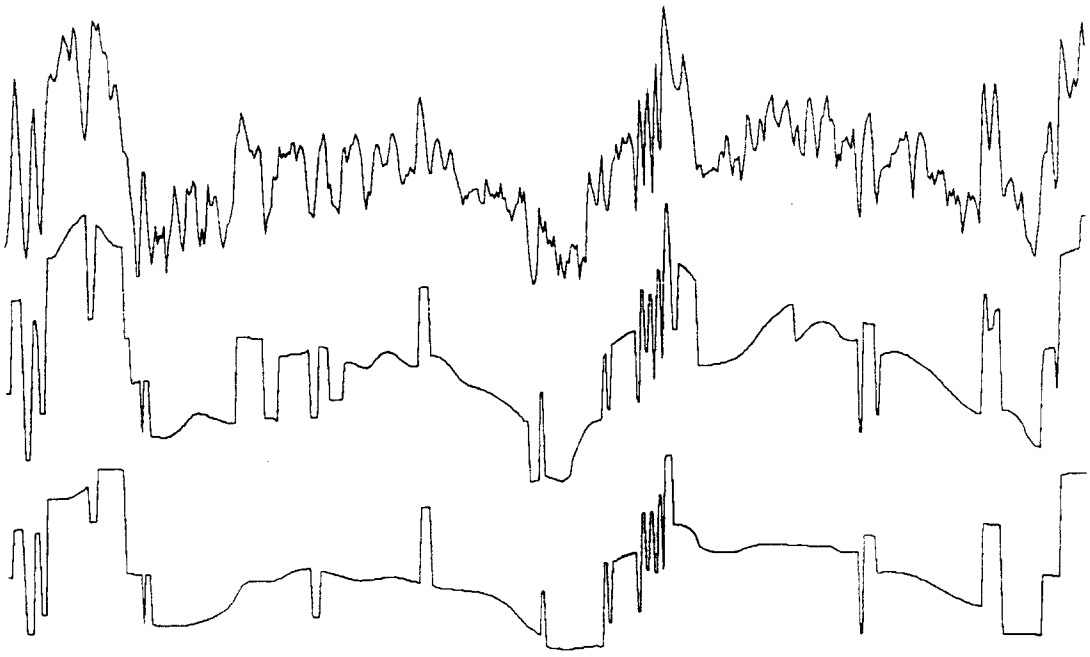

Figure 10: Scale space obtained with anisotropic diffusion. The diffusion was performed in 2D on the Canaletto image of which one line (the horizontal line number 400 out of 480 - just above the gondola) is shown. Notice that the edges remain sharp until their disappearance.

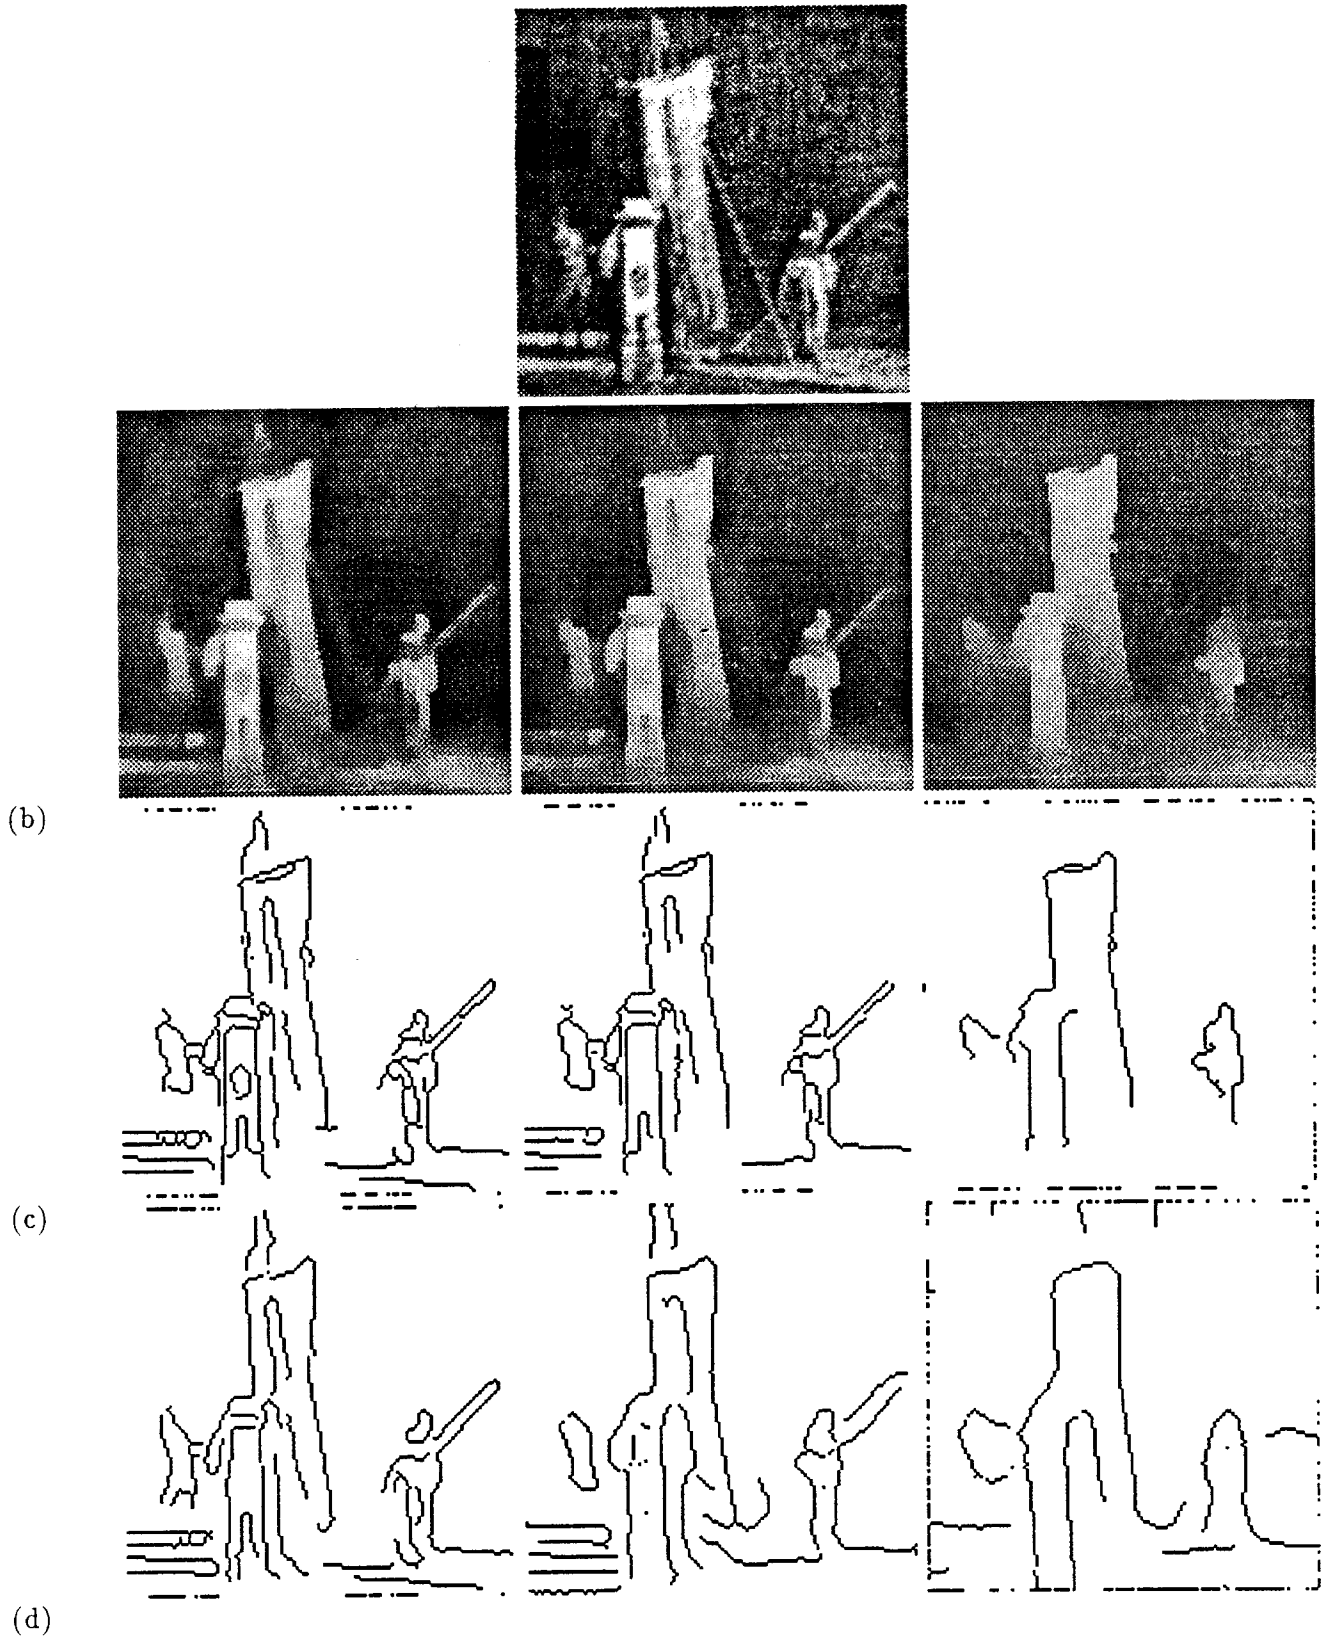

Figure 11: From top to bottom: (a) original image, (b) scale-space using anisotropic diffusion (10, 20, 80 iterations), (c) edges of the same, (d) edges at comparable scales detected using the Canny detector (convolution kernels of variance 1, 2, 4 pixels).

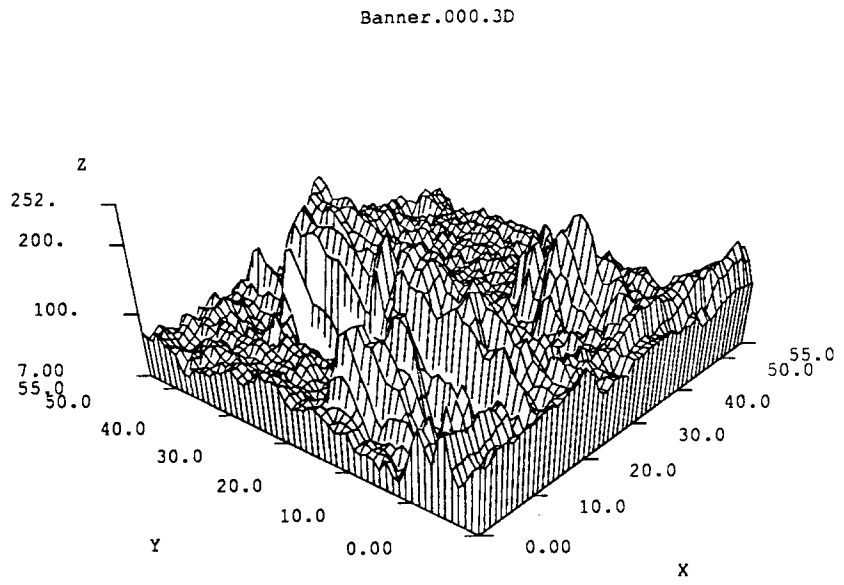

(a)

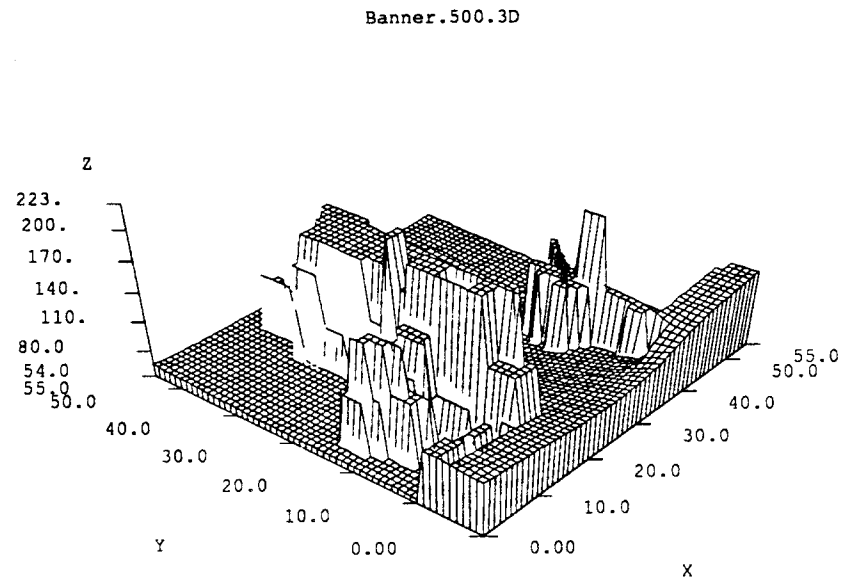

(b)

Figure 12: Scale-space using anisotropic diffusion. Three dimensional plot of the brightness in figure 11.

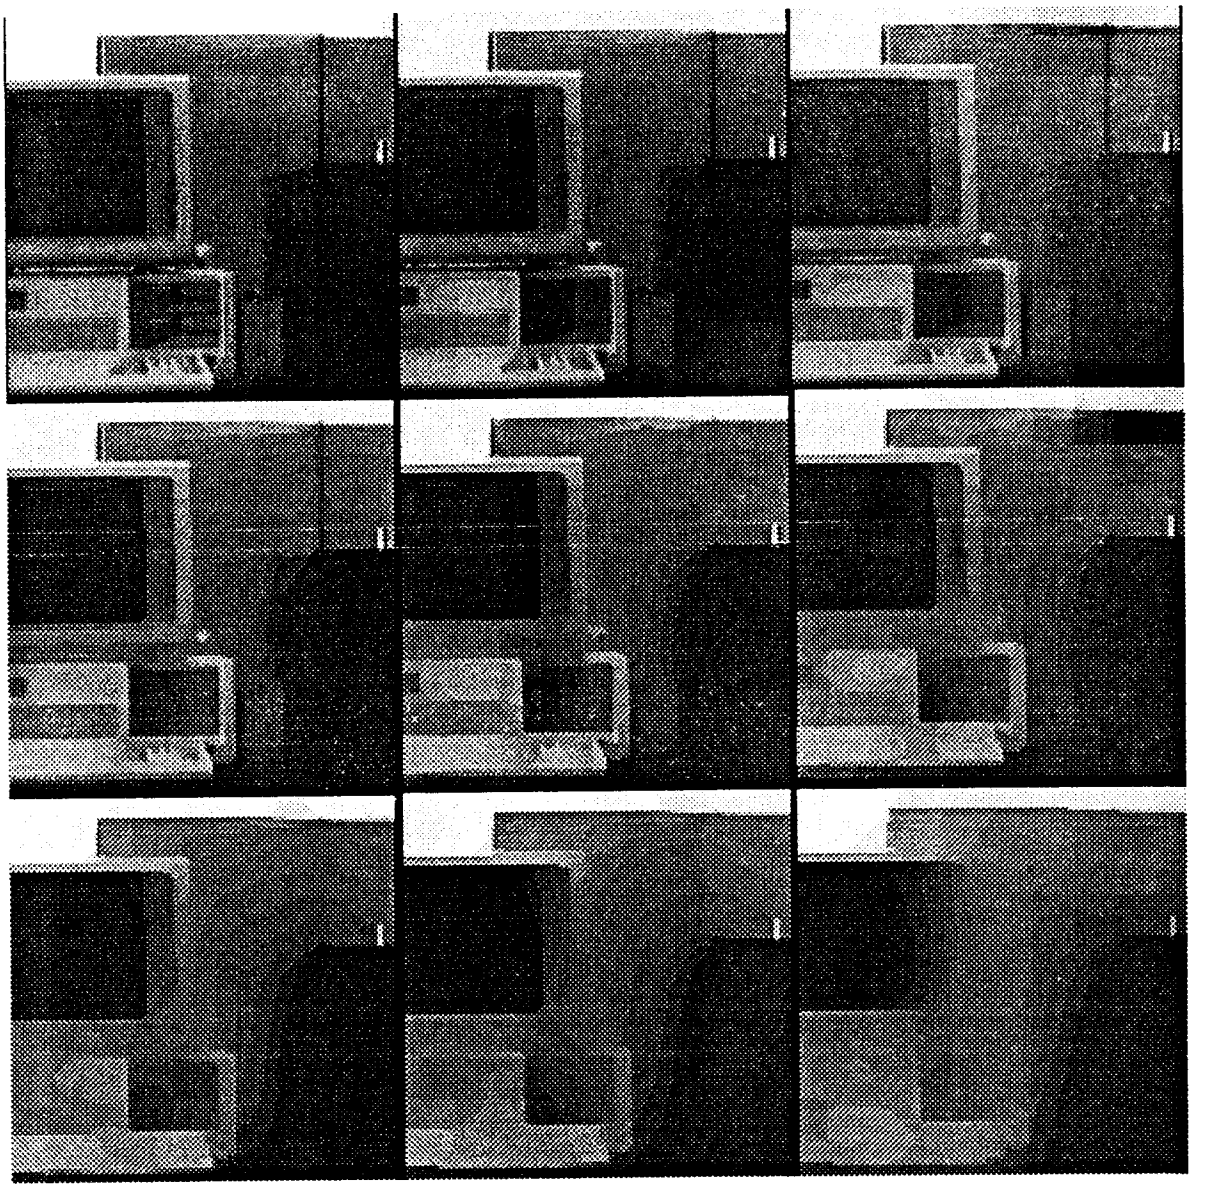

Figure 13: Scale-space using anisotropic diffusion. Original image (top left) and coarser scale images after (left to right, top to bottom) 20, 60, 120, 160, 220, 280, 320, 400 iterations.

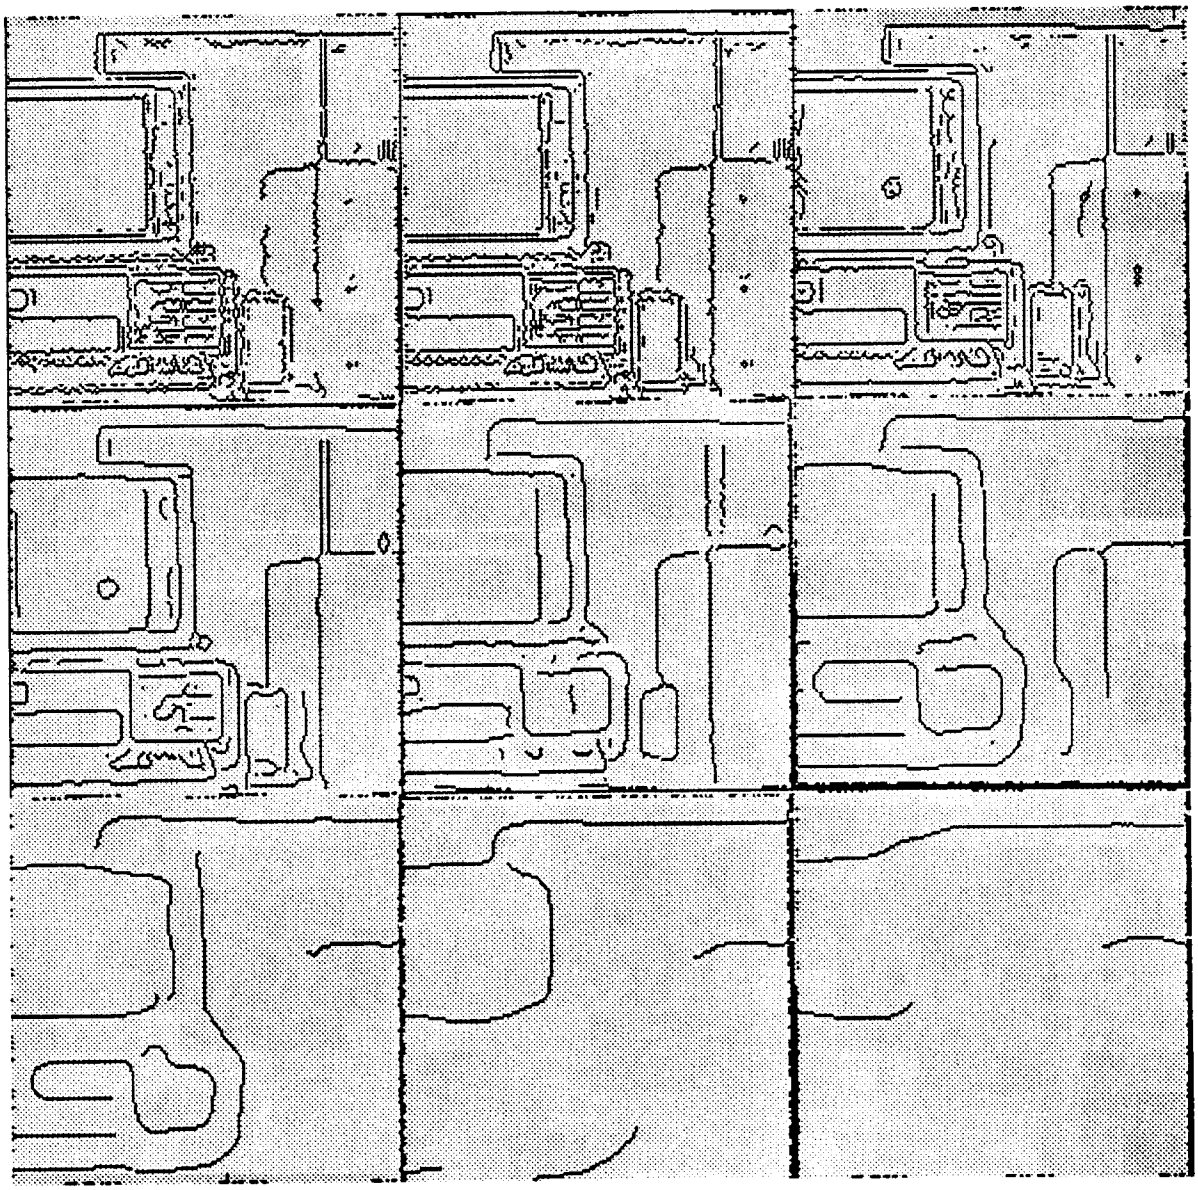

Figure 14: Scale-space using linear convolution. The edges are distorted and the junctions disappear. Images generated using the Canny detector and smoothing Gaussian kernels of variance (top left to bottom right) 1/2, 1, 2, 4, 8, 16 pixels.

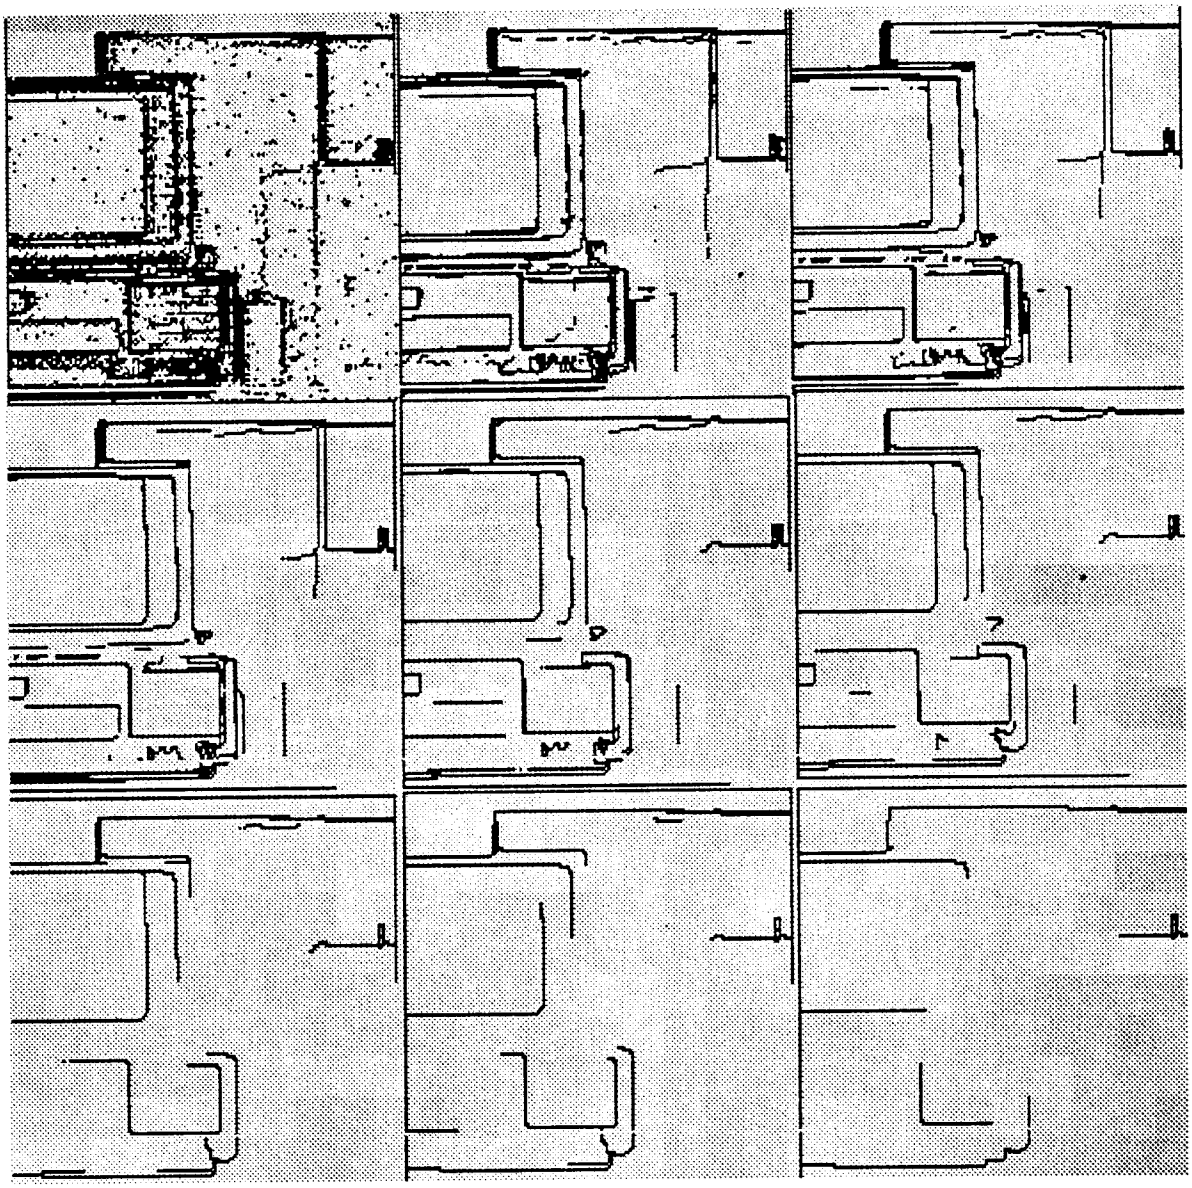

Figure 15: Edges detected by thresholding the gradient on figure 13. Linking is not necessary. Thinning is only for the finer scales. Compare with figure 16 where thinning and linking have been used.

Figure 16: Edges detected on figure 13 using a thinning and linking stage [4].

shape and size of the regions within which smoothing is required. In schemes based on linear smoothing or fixed-neighbourhood processing the shape and size of the areas where smoothing occurs are constant throughout the image. This causes distortions in the shape of the meaningful regions, and in the loss of structures like edge junctions (see Fig. 9(b), 11(d),14) which contain much of the information that allows a three dimensional interpretation of the edge line-drawing [11].

Simplicity - The algorithm consists in identical nearest-neighbour operations (4-8 differences, a function evaluation or a table look-up, and 4-8 sums) iterated over the nodes of a 4 (8) connected square lattice. By comparison the Canny detector requires a number of convolutions (each involving large neighbourhoods at a time) as a pre-processing stage, and a stage of cross-scales matching. Moreover with our algorithm the edges are made sharp by the diffusion process discussed in section 4.2, so that edge thinning and linking are almost unnecessary, especially at coarse scales of resolution (compare Fig. 16 with Fig. 15). For edge detectors based on convolution this is an essential, delicate, and expensive step since linear low-pass filtering has the effect of blurring the edges. The simplicity of the computations involved in anisotropic diffusion makes it a good candidate for digital hardware implementations.

Parallelism- The structure of the algorithm is parallel which makes it very cheap to run on arrays of simple parallel processors.

On sequential machines, anisotropic diffusion is computationally more expensive than convolution-based detectors. This is because in the diffusion process a continuum of scales are generated instead of a small fixed number.

6.2 Energy-based methods for image reconstruction and segmentation

A number of methods have appeared in the literature where the edge detection/image segmentation process is performed by the minimization of an energy function of type

$$U(\vec{z}) = \sum_{i \in \mathbf{I}, j \in \mathbf{N}(i)} V(z_i, z_j) + \sum_{i \in \mathbf{I}} W_i(z_i) \tag{14}$$

with $\mathbf{I}$ indicating the set of the nodes of a lattice, $\mathbf{N}(i) \in \mathbf{I}$ indicating the nodes neighbouring node i, and z a function defined on the lattice, typically the brightness function [6,15,2]. An equivalent formulation is based on finding maxima of a Markov probability distribution function defined on the space of all images:

$$\mathbf{P_Z}(\vec{z}) = \frac{1}{K} e^{-U(\vec{z})} \tag{15}$$

where the function $U(\cdot)$ has the form of equation (14) [6.14]. Because the exponential function is monotonic the maxima of the probability distribution and the minima of the energy function coincide, and we can limit our attention to the schemes based on minimizing the energy.

The energy function (14) is the sum of two terms: the "a priori" term (the sum of the "clique" functions V containing the "a priori" knowledge about the image space – see any one of [6,15,2] for a complete discussion), and a term depending on the data available (the sum of the functions W_i). $V(\cdot,\cdot)$ is typically an even function depending only on the value of the difference of its arguments (with abuse of notation $V(z_i, z_j) = V(z_i - z_j)$). It has minimum at zero and it is monotonic on the positive and negative semilines assigning higher energy ($\Leftrightarrow$ lower probability) to the pairs i, j of lattice nodes whose brightness difference $\|z_i - z_j\|$ is bigger. We will show that the simple approximation of anisotropic diffusion that we suggest in section 5 may be seen as a gradient descent of the "a priori" part of the energy function

$$U(\vec{z}) = \sum_{i \in \mathbf{I}, j \in \mathbf{N}(i)} V(z_i, z_j) \tag{16}$$

The steepest descent strategy for finding minima of a function consists of starting from some initial state, and then changing iteratively the state following the opposite of the gradient vector. The gradient of the energy function, which may be computed from equation (16) differentiating with respect to the z_i, is the vector of components

$$\nabla U_i(z) = 2 \sum_{j \in \mathbf{N}(i)} \dot{V}(z_i - z_j) \tag{17}$$

therefore the gradient descent algorithm is

$$\frac{\partial}{\partial t} z_i = -A \cdot \sum_{j \in \mathbf{N}(i)} \dot{V}(z_i - z_j) \tag{18}$$

where A is some "speed" factor.

Suppose that $V(\cdot)$ is differentiable in the origin and define $\phi(\cdot) = -\dot{V}$. Since $V(\cdot)$ is even, $\phi(\cdot)$ is an odd function and $\phi(0) = 0$. Then we may write $\phi(s) = s \cdot c(s)$ for some function $c(\cdot)$ even and positive. Substituting into equation (18) we obtain:

$$\frac{\partial}{\partial t} z_i = A \cdot \sum_{j \in \mathbf{N}(i)} c(z_j - z_i) \cdot (z_j - z_i) \tag{19}$$

which is exactly the anisotropic diffusion algorithm defined by equations (10), (11), (13) if the neighbourhood structure is given by natural nearest-neighbour cliques of a square lattice. The flux functions obtained by differentiating the local energy functions $V(\cdot)$ of [6,13,2] are very similar to the shape of flux function that the analysis in section 4.2 suggests. See figure 17.

To summarize: anisotropic diffusion may be seen as a gradient descent on some energy function. The data (the original image) are used as the initial condition. In the energy-based methods [6,15,2] the closedness of the solution to the data is imposed by a term in the energy function. This makes the energy function non-convex and more complicated optimization algorithms than gradient descent are necessary. Most of the algorithms that have been proposed (simulated annealing for example) appear too slow for vision applications. Perhaps the only exception is the GNC algorithm proposed by Blake and Zisserman [2] which does not guarantee to find the global optimum for generic images, but appears to be a good compromise between speed and accuracy.

7 Conclusion

We have introduced a tool, anisotropic diffusion, that we believe will prove useful in many tasks of early vision. Diffusion based algorithms involve very simple, local, identical computations over the entire image lattice. Implementations on massively parallel architectures like the connection machine would be almost trivial. Implementations using hybrid analog-digital networks also seem feasible.

We have shown that a very simple version of anisotropic diffusion can be applied with success to multiscale image segmentation. As a pre-processing step it makes thinning and linking of the edges trivial to realize, it preserves the edge junctions, and it does not require complicated comparison of images at different scales since shape and position are preserved at every single scale.

In images in which the brightness gradient generated by the noise is greater than that of the edges, and the level of the noise varies significantly across the image the scheme that we have described proves insufficient to obtain a correct multi-scale segmentation. This is because a global noise estimate does not provide an accurate local estimate, and the local value of the gradient provides too partial a piece of information for distingushing noise-related and edge-related gradients. Moreover the abscissa K of the peak of the flux function $\phi(\cdot)$ has to be set according to the typical contrast value. If this changes considerably through the image the value of K has to be locally adapted.

In a forthcoming paper we study methods to overcome these two difficulties. These methods are still based on the diffusion equation, and the algorithms we suggest preserve the inherent parallelism of anisotropic diffusion.

Acknowledgements

This work has been supported by grant 82-11-008 of the Semiconductor Research Corporation (P.P.), an IBM faculty developement award (J.M.) and by DARPA contract N00039-88-C-0292.

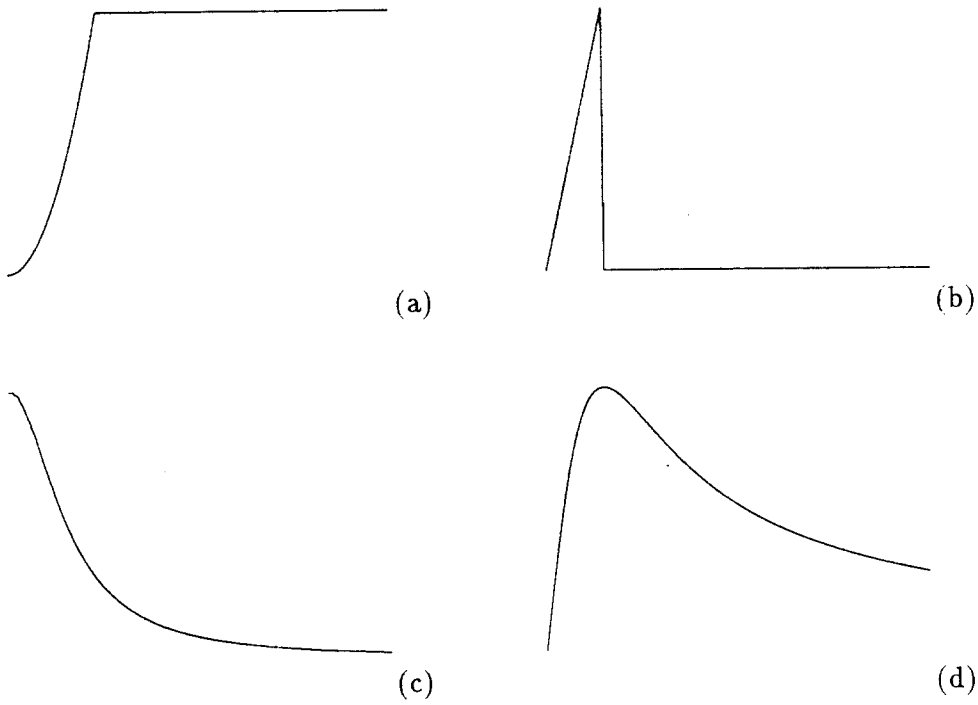

Figure 17: (a) The local energy function proposed by [6,2,15] typically is equal to the square of the nearest-neighbour brightness difference, and saturates at some threshold value. (b) The first derivative of the energy function (a). (c) (d) The anisotropic diffusion conduction coefficient and flux function as a function of the brightness gradient magnitude, proportional to the nearest neighbour brightness difference in the discrete case. (b) and (d) have the same role (see text in section 6.2).

References

[1] J. Babaud, A. Witkin, Baudin, and R.Duda. Uniqueness of the gaussian kernel for scale-space filtering. *IEEE Trans. on PAMI*, 8, January 1986.

[2] A. Blake and A. Zisserman. *Visual reconstruction.* MIT press, 1987.

[3] Canaletto. View in Venice. Circa 1740.

[4] J. Canny. *Finding edges and lines in images.* Technical Report 720, MIT, Artificial Intelligence Laboratory, 1983.

[5] J. Clark. Singularity theory and phantom edges in scale space. *IEEE trans. on pattern analysis and machine intelligence*, 10(5):720–727, 1988.

[6] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on PAMI*, 6:721–741, November 1984.

[7] R. Hummel. *Representations based on Zero-crossings in Scale-Space.* Technical Report 225, New York University, Courant Institute of Mathematical Sciences, 1986.

[8] R. Hummel, B. Kimia, and S.Zucker. Deblurring gaussian blur. *Computer vision, graphics and image processing*, 38:66–80, 1987.

[9] F. John. *Partial Differential Equations.* Springer Verlag , New York, 1982.

[10] J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.

[11] J. Malik. Interpreting line drawings of curved objects. *International Journal of Computer Vision*, 1(1):73–103, 1987.

[12] D. Marr. *Vision.* W.H.Freeman & Co., 1982.

[13] J. Marroquin. *Probabilistic solution of inverse problems.* AI-TR 860, M.I.T. - Artificial Intelligence Laboratory, 1985.

[14] J. Marroquin. *Probabilistic solution of inverse problems.* PhD thesis, MIT, 1985.

[15] D. Mumford and J. Shah. Boundary detection by minimizing functionals. *Unpublished manuscript*, 1986.

[16] P. Perona and J. Malik. A network for edge detection and scale space. In *IEEE International Symposium on Circuits and Systems - Helsinki*, pages 2565–2568, June 1988.

[17] P. Perona and J. Malik. Scale space,and edge detection using anisotropic diffusion. In *IEEE Computer Society Workshop on Computer Vision - Miami*, pages 16–22, 1987.

[18] A. Rosenfeld and M. Thurston. Edge and curve detection for visual scene analysis. *IEEE Trans. on Computers*, C-20:562–569, May 1971.

[19] A. Witkin. Scale-space filtering. In *International Joint Conference on Artificial Intelligence*, pages 1019–1021, 1983. Karlsruhe.

[20] A. Yuille and T. Poggio. Scaling theorems for zero crossings. *IEEE Transactions on Pattern Analysis, and Machine Intelligence*, 8, January 1986.