Technical Report OSU-CISRC-7/98-TR22
Department of Computer and Information Science
The Ohio State University
Anonymous ftp: ftp.cis.ohio-state.edu
File: /pub/tech-report/1998/TR22.ps.gz
WWW: http://www.cis.ohio-state.edu/Research/tech-report.html

Boundary Detection by Contextual Nonlinear Smoothing

Xiuwen Liu[1,2] DeLiang Wang[1,3] J. Raul Ramirez[2]

[1]*Department of Computer and Information Science*

[2]*Center for Mapping* [3]*Center for Cognitive Science*

The Ohio State University, Columbus, Ohio 43210, USA

{liux, dwang}@cis.ohio-state.edu {xiuliu, raul}@cfm.ohio-state.edu

Abstract

In this paper we present a two-step boundary detection algorithm. The first step is a nonlinear smoothing algorithm which is based on an orientation-sensitive probability measure. By incorporating geometrical constraints through the coupling structure, we obtain a robust nonlinear smoothing algorithm, where many nonlinear algorithms can be derived as special cases. Even when noise is substantial, the proposed smoothing algorithm can still preserve salient boundaries. Compared with anisotropic diffusive approaches, the proposed nonlinear algorithm not only performs better in preserving boundaries but also has a non-uniform stable state, whereby reliable results are available within a fixed number of iterations independent of images. The second step is simply a Sobel edge detection algorithm without non-maximum suppression and hysteresis tracking. Due to the proposed nonlinear smoothing, salient boundaries are extracted effectively. Experimental results using synthetic and real images are provided.

Keywords: Nonlinear smoothing, Contextual information, Anisotropic diffusion, Edge detection, Boundary detection, Orientation-sensitive probability

1. Introduction

One of the fundamental tasks in low-level machine vision is to locate discontinuities in images corresponding to physical boundaries between a number of regions. A common practice is to identify local maxima in local gradients of images - collectively known as edge detection algorithms. The Sobel edge detector[1] consists of two 3 x 3 convolution kernels, which respond maximally to vertical and horizontal edges respectively. Local gradients are estimated by convolving the images with the two kernels, and thresholding is then applied to get rid of noisy responses. The Sobel edge detector is computationally efficient but sensitive to noise. To make the estimation of gradients more reliable, the image can be convolved with a low-pass filter before estimation and two influential methods are due to Marr and Hildreth[2] and Canny [3]. By convolving the image with a Laplacian of Gaussian kernel, the resulting local maxima, which are assumed to correspond to meaningful edge points, are zero-crossings in the filtered image[2]. Canny[3] derived an optimal step edge detector using variational techniques starting from some optimal criteria and used the first derivative of a Gaussian as a good approximation of the derived detector. Edge points are then identified using a non-maximum suppresion and hysteresis thresholding for better continuity of edges. As noticed by Marr and Hildreth[2], edges detected at a fixed scale are not sufficient and multiple scales are essentially needed in order to obtain good results. By formalizing the multiple scale approach, Witkin[4] and Koenderink[5] proposed Gaussian scale space. The original image is embedded in a family of gradually smoothing images controlled by a single parameter, which is equivalent to solving a heat equation with input as the initial condition[5]. While Gaussian scale space has nice properties and is widely used in machine vision[6], a major limitation is that Gaussian smoothing inevitably blurs edges and other important features due to its low-pass nature. To overcome the limitation, non-uniform diffusion, which was proposed by Cohen and Grossberg[7] in modeling the primary visual cortex, was formulated by Perona and Malik[8]:

$$\frac{\partial I}{\partial t} = div\,(g\,(\|\nabla I\|)\,\nabla I) \tag{1}$$

Here *div* is the divergence operator, and g is a nonlinear monotonically decreasing function and ∇I denotes the gradient. By making the diffusion conductance dependent explicitly on local gradients, the non-uniform diffusion prefers intra-region smoothing over inter-region smoothing, resulting immediate localization while noise is reduced[8]. Because it produces visually impressive results, non-uniform diffusion generates much theoretical as well as practical interest (see reference 9 for a recent review). While many improvements have been proposed, including spatial regularization[10] and edge-enhancing anisotropic diffusion[11], the general framework remains the same. As shown by You et al[12], anisotropic diffusion given by (1) is the steepest gradient descent minimizer of the following energy function:

$$E(I) = \int_{\Omega} f(\|\nabla I\|)\, d\Omega \tag{2}$$

with

$$g(\|\nabla I\|) = \frac{f'(\|\nabla I\|)}{\|\nabla I\|}$$

Because the energy function given by (2) has a trivial global minimum, where the image is constant everywhere, interesting results exist only within a certain period of time of diffusion. An immediate problem is how to determine the termination time, which we refer to as the *termination problem*. While there are some heuristic rules on how to choose the stop time[10-11], in general it corresponds to the open problem of automatic scale selection. As in Gaussian scale space, a fixed time would not be sufficient to obtain good results. Another problem of anisotropic diffusion is that diffusion conductance is a deterministic function of local gradients, which, similar to non-maximum suppresion in edge detection algorithms, makes an implicit assumption that larger gradients are due to true boundaries. When noise is substantial and gradients due to noise and boundaries cannot be distinguished based on magnitudes, the approach tends to fail to preserve meaningful region boundaries.

To illustrate the problems, Figure 1(a) shows a noise-free image, where the gradient magnitudes along the central square change considerably. Figure 1(b) shows a noisy version of Fig. 1(a) by adding Gaussian noise with zero mean and $\sigma = 40$, and Fig. 1(c) shows its local gradient magnitude obtained using Sobel operators[1]. While the three major regions in Fig. 1(b) may be perceived, Fig. 1(c) is very noisy and the strong boundary fragment is barely visible. Figures 1(d)-(f) show the smoothed images by an anisotropic diffusive algorithm[13] with specified numbers of iterations. Figures 1(g)-(i) show the edge maps of Fig. 1(d)-(f), respectively, using the Sobel edge detection algorithm. While at the 50th iteration the result is still noisy, the result becomes meaningless at the 1000th iteration. Even though the result at the 100th iteration is visually good, the boundaries are still fragmented and it is not clear how to identify a “good” number of iterations automatically.

(Figure 1 should be inserted around here)

These problems to a large extent are due to the assumption that local maxima in gradient images are good edge points. In other words, due to noise, responses from true boundaries and those from noise are not distinguishable based on magnitude. To overcome these problems, contextual information, i.e., responses from neighboring pixels, should be incorporated in order to reduce ambiguity as in relaxation labeling and related methods[14-17]. In general, relaxation labeling methods use pair-wise compatibility measure, which is determined based on a priori models associated with labels, and convergence is not known and often very slow in numerical simulations[18]. In this paper, by using an orientation-sensitive probability measure, we incorporate contextual information through the geometrical constraints on the coupling structure. Numerical simulations show that the resulting nonlinear algorithm has a non-uniform stable state and good results can be obtained within a fixed number of iterations independent of input images. Also, the oriented probability measure is defined on input data, and thus no a priori models need to be assumed.

In Section 2, we formalize our contextual nonlinear algorithm and show that many nonlinear algorithms can be treated as special cases. Section 3 gives some theoretical results as well as numerical simulations regarding to the stability and convergence of the algorithm. Section 4 provides experimental results using synthetic and real images. Section 5 concludes the paper with further discussions.

2. Contextual nonlinear algorithm

2.1 Design of the algorithm

To design a statistical algorithm, with no prior knowledge, we assume a Gaussian distribution within each region. That is, given a pixel (i_0, j_0) and a window $R_{(i_0, j_0)}$ at pixel (i_0, j_0), consisting of a set of pixel locations, we assume that:

$$P(I_{(i_0, j_0)}, R) = \frac{1}{\sqrt{2\pi}\sigma_R} \exp\left\{-\frac{(I_{(i_0, j_0)} - \mu_R)^2}{2\sigma_R^2}\right\} \tag{3}$$

where $I_{(i, j)}$ is the intensity value at pixel location (i, j). To simplify notation, without confusion, we use R to stand for $R_{(i_0, j_0)}$. Intuitively, $P(I_{(i_0, j_0)}, R)$ is a measure of compatibility between intensity value at pixel (i_0, j_0) and statistical distribution in window R. To estimate the unknown parameters of μ_R and σ_R, consider the pixels in R as n realizations of (3), where $n = |R|$. The likelihood function of μ_R and σ_R is(19):

$$L(R; \mu_R, \sigma_R) = \left(\frac{1}{\sqrt{2\pi}\sigma_R}\right)^n \exp\left(-\frac{1}{2\sigma_R^2} \sum_{(i, j) \in R}^{n} (I_{(i, j)} - \mu_R)^2\right) \tag{4}$$

By maximizing (4), we get the maximum likelihood estimators for μ_R and σ_R:

$$\hat{\mu}_R = \frac{1}{n} \sum_{(i, j) \in R} I_{(i, j)} \tag{5a}$$

$$\hat{\sigma}_R = \frac{1}{\sqrt{n}} \sqrt{\sum_{(i, j) \in R} (I_{(i, j)} - \hat{\mu}_R)^2} \tag{5b}$$

To do a nonlinear smoothing, similar to selective smoothing filters(20-21), suppose that there are M windows $R^{(m)}$, where $1 \le m \le M$, around a central pixel (i_0, j_0). Here these $R^{(m)}$'s can be generated from one or several basis windows through rotation, which are motivated by the experimental findings of orientation selectivity in the visual cortex(22). Simple examples are rectangular windows, which are used throughout this paper for synthetic and real images. The probability that pixel (i_0, j_0) belongs to $R^{(m)}$ can be estimated from equations (3) and (5). By assuming that the weight of each $R^{(m)}$ should be proportional to the probability, as in relaxation labeling(14-15), we obtain an iterative nonlinear smoothing filter:

$$I^{t+1}_{(i_0, j_0)} = \frac{\sum_m P\left(I^t_{(i_0, j_0)}, R^{(m)}\right) \hat{\mu}^t_{R^{(m)}}}{\sum_m P\left(I^t_{(i_0, j_0)}, R^{(m)}\right)} \tag{6}$$

A problem with this filter is that it is not sensitive to weak edges due to the linear combination. To generate more semantically meaningful results and increase the sensitivity even to weak edges, we apply a nonlinear function on weights, which is essentially same to anisotropic diffusion[8]:

$$I^{t+1}_{(i_0, j_0)} = \frac{\sum_m g\left(P\left(I^t_{(i_0, j_0)}, R^{(m)}\right)\right)\hat{\mu}^t_{R^{(m)}}}{\sum_m g\left(P\left(I^t_{(i_0, j_0)}, R^{(m)}\right)\right)} \tag{7}$$

Here g is a nonlinear function. A good choice for g is an exponential function, which is widely used in nonlinear smoothing anisotropic diffusion approaches:

$$g(x) = \exp\left(-x^2 / K\right) \tag{8}$$

Here parameter K controls the sensitivity to edges[23]. Equation (7) provides a generic model for a wide range of nonlinear algorithms, the behavior of which largely depends on the sensitivity parameter K. When K is large, (7) reduces to the equally weighted average smoothing filter. When K is around 0.3, g is close to a linear function in [0,1] and (7) then reduces to (6). When K is a small positive number, (7) will be sensitive to all discontinuities. No matter how small the weight of one window can be, theoretically speaking, if it is nonzero, when $t \to \infty$, the system will reach a uniform stable state. Similar to anisotropic diffusion approaches, the desired results will be time-dependent and the termination problem becomes a critical issue for autonomous solutions. To overcome this limitation, we restrict smoothing only within the window with the highest probability similar to selective smoothing[20-21]:

$$m^* = \max_{1 \le m \le M} \left(P\left(I_{(i_0, j_0)}, R^{(m)}\right)\right) \tag{9}$$

The nonlinear smoothing through (9) is desirable in regions that are close to edges. By using appropriate $R^{(m)}$'s, (9) encodes discontinuities implicitly. But in homogenous regions, (9) may produce artificial block effects due to intensity variations. Under the proposed statistical formulation, there is an adaptive method to detect homogeneity. Based on the assumption that there are M $R^{(m)}$ windows around a central pixel (i_0, j_0), where each window has a Gaussian distribution, consider the mean in each window as a new random variable:

$$\mu^{(m)} = \frac{1}{\left|R^{(m)}\right|} \sum_{(i,j) \in R^{(m)}} I_{(i,j)} \tag{10}$$

Because $\mu^{(m)}$ is a linear combination of random variables with a Gaussian distribution, $\mu^{(m)}$ has also a Gaussian distribution with the same mean and a standard deviation given by:

$$\sigma_{\mu^{(m)}} = \frac{1}{\sqrt{\left|R^{(m)}\right|}} \sigma_{R^{(m)}} \tag{11}$$

This provides a probability measure of how likely that the M windows are sampled from one homogenous region. Given a confidence level α, for each pair of windows $R^{(m1)}$ and $R^{(m2)}$ we have:

$$\left|\mu^{(m1)} - \mu^{(m2)}\right| \le min\left(\sqrt{\frac{\log(1/\alpha)}{\left|R^{(m1)}\right|}}\hat{\sigma}_{R^{(m1)}}, \sqrt{\frac{\log(1/\alpha)}{\left|R^{(m2)}\right|}}\hat{\sigma}_{R^{(m2)}}\right) \tag{12}$$

If all the pairs satisfy (12), the M windows are likely from one homogenous region with confidence α. Intuitively, under the assumption of a Gaussian distribution, when we have more samples, i.e., the window $R^{(m)}$ is larger, the estimation of the mean is more precise and so the threshold should be smaller. In a region with a larger standard deviation, the threshold should be larger because larger variations are allowed.

The nonlinear smoothing algorithm outlined above works well when noise is not very large. In cases when signal to noise ratio is very low, the probability measure given in (3) would be unreliable because pixel values change considerably. This problem can be alleviated by using the mean value of pixels sampled from R which are close to the central pixel (i_0, j_0), or along a certain direction to make the algorithm more orientation sensitive.

To summarize, we obtain a nonlinear smoothing algorithm. We define M oriented windows which can be obtained by rotating one or more basis windows. At each pixel, we estimate parameters using (5). If all the M windows belong to a homogenous region according to (12), we do the smoothing within all the M windows. Otherwise, the smoothing is done only within the most compatible window given by (9).

2.2 A generic nonlinear smoothing framework

In this section we will show how to derive several widely used nonlinear algorithms from the statistical nonlinear algorithm outlined above. Several early nonlinear filters[20-21] do the smoothing in a window where the standard deviation is the smallest. These filters can be obtained by simplifying (3) to

$$P(I_{(I_0, j_0)}, \hat{\mu}, \hat{\sigma}) = \frac{1}{\sqrt{2\pi}\hat{\sigma}} C \tag{13}$$

where C is a constant. Then the solution to (9) is the window with the smallest deviation. This also justifies why those filters produce good results.

Another nonlinear smoothing filter is the gradient-inverse filter[24]. Suppose that there is one window, i.e., $M = 1$, consisting of the central pixel (i_0, j_0) itself only, the estimated deviation for a given pixel (i, j) in (5b) now becomes

$$\hat{\sigma} = \left| I_{(i,j)} - I_{(i_0, j_0)} \right| \tag{14}$$

Equation (14) is a popular way to estimate local gradients. Using (13) as the probability measure, (6) becomes exactly the gradient inverse nonlinear smoothing filter[24].

SUSAN nonlinear smoothing filter[25] is proposed based on the SUSAN (Smallest Univalue Segment Assimilating Nucleus) principle. The SUSAN smoothing filter is:

$$I^{t+1}_{(i_0, j_0)} = \frac{\sum_{(\delta i, \delta j) \neq (0,0)} I^t_{(i_0 + \delta i, j_0 + \delta j)} W(i_0, j_0, \delta i, \delta j)}{\sum_{(\delta i, \delta j) \neq (0,0)} W(i_0, j_0, \delta i, \delta j)} \tag{15}$$

where

$$W(i_0, j_0, \delta i, \delta j) = \exp\left(-\frac{r^2}{2\sigma^2} - \frac{\left(I^t_{(i_0 + \delta i, j_0 + \delta j)} - I^t_{(i_0, j_0)} \right)^2}{T^2} \right)$$

Here (i_0, j_0) is the central pixel under consideration, and $(\delta i, \delta j)$ defines a local neighborhood. Essentially, it integrates Gaussian smoothing in spatial and brightness domains. The parameter T is a threshold for intensity values. It is easy to see from (15) that the weights are derived based on pair-wise intensity value differences. It would be expected that the SUSAN filter performs well when images consists of relatively homogenous regions and within each region noise is smaller than T. When noise is substantial, it fails to preserve structures due to the pair-wise difference calculation, where no geometrical constraints are incorporated. This is consistent with the experimental results, which will be discussed later. To get the SUSAN filter, we define one window including the central pixel itself only. For a given pixel (i, j) in its neighborhood, (3) can be simplified to:

$$P(I_{(i,j)}, R) = C \cdot \exp\left\{-\frac{(I_{(i,j)} - \hat{\mu}_R)^2}{T^2}\right\} \tag{16}$$

where C is a scaling factor. Because now $\hat{\mu}_R$ is $I_{(i_0, j_0)}$, (6) with the probability measure given by (16) is equivalent to Gaussian smoothing in the brightness domain in (15).

Now consider the anisotropic diffusion given by (1). By discretizing (1) in image domain with four nearest neighbor coupling[13] and rearranging the terms, we have:

$$I^{t+1}_{(i,j)} = \eta^t_{(i,j)} I^t_{(i,j)} + \lambda \sum_m g\left(P\left(I^t_{(i,j)}, R^{(m)} \right)\right) \hat{\mu}^t_{R^{(m)}} \tag{17}$$

If we have four singleton regions, (17) is essentially a simplified version of (7) with an adaptive learning rate.

3. Analysis

3.1 Theoretical results

One of the distinctive characteristics of the proposed algorithm is that it requires spatial constraints among responses from neighboring locations through coupling structure as opposed to pair-wise coupling structure. Figure 2 illustrates the concept using a manually constructed example. Figure 2(a) shows the oriented windows in a 3 x 3 neighborhood, and Fig. 2(b) shows the coupling structure if we apply the proposed algorithm to a small image patch. The directed graph is constructed as follows. There is a directed edge from (i_1, j_1) to (i_0, j_0) if and only if (i_1, j_1) contributes to the smoothing of (i_0, j_0) according to equations (12) and (9). By doing so, the coupling structure is represented as a directed graph as shown in Fig. 2(b). Connected components and strongly connected components[26] of the directed graph can be used to analyze the temporal behavior of the proposed algorithm. A strongly connected component is a set of vertices, or pixels here, where there is a directed path from any vertex to all the other vertices in the set. We obtain a connected component if we do not consider the direction of edges along a path. In the example shown in Fig. 2(b), all the black pixels form a strongly connected component and so do all the white pixels. Also, there are obviously two connected components.

(Figure 2 should be around here)

Essentially our nonlinear smoothing algorithm can be viewed as a discrete dynamic system, the behavior of which is complex due to spatial constraints imposed by coupling windows and adaptive coupling structure by probabilistic grouping. We now prove that a constant region satisfying certain geometrical constraints is a stable state of the smoothing algorithm.

Theorem. If a region S of a given image I satisfies:

$$(i_1, j_1) \in S \text{ and } (i_2, j_2) \in S \Rightarrow I_{(i_1, j_1)} = I_{(i_2, j_2)} \tag{18a}$$

$$\forall (i, j) \in S \Rightarrow \exists m \ \ R_{(i, j)}^{(m)} \subseteq S \tag{18b}$$

Then S is stable with respect to the proposed algorithm.

Proof. Condition (18a) states that S is a constant region and the standard deviation is zero if $R^{(m)}$ is within S according to equation (5b). Consider a pixel (i_0, j_0) in S. Inequality (12) is satisfied only when all $R^{(m)}$'s are within S. In this case, the smoothing algorithm does not change the intensity value at (i_0, j_0). Otherwise, $R^{(m^*)}$ according to equation (9) must be within S because there exists at least one such window according to (18b) and thus the smoothing algorithm does not change the intensity value at (i_0, j_0) also. So S is stable. Q.E.D.

A maximum connected component of the constructed graph is stable when its pixels are constant and thus maximum connected components of the constructed graph are a piecewise constant stable solution of the proposed algorithm. For the image patch given in Fig 2(b), for example, a stable solution is that pixels in each of the two connected components are constant. The noise-free image in Fig. 1(a) is also a stable solution by itself, as we will demonstrate through numerical simulations later on.

It is easy to see from the proof that any region which satisfies conditions (18a) and (18b) during temporal evolution will stay unchanged. In addition, due to the smoothing nature of the algorithm, a local maximum at iteration t cannot increase according to the smoothing kernel by equation (12) or (9), and similarly, a local minimum cannot decrease. We conjecture that any given image approaches an image that is almost covered by homogenous regions. Due to the spatial constraints given by (18b), it is not clear if the entire image converges to a piece-wise constant stable state. Within each resulting homogenous region, (18b) is satisfied and thus the region becomes stable. For pixels near boundaries, corners, and junctions, it is possible that (18b) is not uniquely satisfied within one constant region, and small changes may persist. The whole image in this case attains a quasi-equilibrium state. This is supported by the following numerical simulations using synthetic and real images. While there are pixels which do not converge within 1000 iterations, the smoothed image as a whole does not change noticeably at all.

The two maximum strongly connected components in Fig. 2(b) satisfy condition (18b). Both of them are actually uniform regions and thus are stable. Gray pixels would be grouped into one of the two stable regions according to pixel value similarity and spatial constraints.

3.2 Numerical simulations

Because it is difficult to derive the speed of convergence analytically, we use numerical simulations to demonstrate the temporal behavior of the proposed algorithm. Since smoothing is achieved using equally weighted average within selected windows, the algorithm should converge rather quickly in homogenous regions. To obtain quantitative estimations, we define two measures similar to variance. For synthetic images, where a noise-free image is available, we define the deviation from the ground truth image as:

$$D_{(I)} = \sqrt{\frac{\sum_i \sum_j \left(I_{(i,j)} - I^{gt}{}_{(i,j)} \right)^2}{|I|}} \tag{19}$$

Here I is the image to be measured and I^{gt} is the ground truth image. The deviation gives an objective measure of how good the smoothed image is with respect to the true image. To measure the convergence, we define relative variance for image I at time t:

$$V^t_{(I)} = \sqrt{\frac{\sum_i \sum_j \left(I^t_{(i,j)} - I^{t-1}_{(i,j)} \right)^2}{|I|}} \tag{20}$$

We have applied the proposed algorithm on the noise-free image shown in Fig. 1(a) and six noisy images generated from it by adding zero-mean Gaussian noise with σ from 5 to 60. Figure 3 shows the deviation from the ground truth image with iterations, and Figure 4 shows the relative variance of the noise-free image and four selected noisy images to make the figure more readable. As we can see from Fig. 3, the noise-free image is a stable solution by itself, where the deviation is always zero. For the noisy images, the deviation from true image is stabilized within a few number of iterations independent of the amount of noise. Figure 4 shows that relative variance is bounded with a small upper limit after 10 iterations. This variance is due to the pixels close to boundaries, corners and junctions that do not belong to any resulting constant region. As discussed before, because the spatial constraints cannot be satisfied within one homogenous region, these pixels have connections from pixels belonging to different homogenous regions, and thus fluctuate. These pixels are a small fraction of the input image in general, and thus the fluctuations do not affect the quality of the smoothed images noticeably. As shown in Fig. 3, the deviation is stabilized quickly.

(Figure 3 should be around here)

(Figure 4 should be around here)

Real images are generally more complicated than synthetic images statistically and structurally, and we have also applied our algorithm to the four real images shown in Fig. 9-12 which include a texture image. Figure 5 shows the relative variance in 100 iterations, where the variance is bounded after 10 iterations independent of images. This indicates that the proposed algorithm behaves similarly for synthetic and real images.

(Figure 5 should be around here)

4. Experimental results

4.1 Results of the proposed algorithm

The nonlinear smoothing algorithm formalized in this paper integrates discontinuity and homogeneity through the orientation-sensitive probability framework. Equation (9) represents discontinuity implicitly and (12) encodes homogeneity explicitly. Because of the probability measure, the initial errors for choosing smoothing windows due to noise can be overcome by the coupling structure. Essentially only when majority of the pixels in one window make a wrong decision, the final result would be affected. As illustrated in Fig. 2, the coupling structure is robust.

The oriented windows used throughout the experiments in this section are shown in Fig. 6, which are generated by rotating two rectangular basis windows. The preferred orientation of each window is consistent with orientation sensitivity of cell responses in the visual cortex[22]. Asymmetric window shapes are used so that 2-D features such as corners and junctions can be preserved.

(Figure 6 should be around here)

As is evident from numerous simulations, the proposed algorithm generates stable results around 10 iterations regardless of input images. Thus, all the boundaries from the proposed algorithm are generated using smoothed images at the 11th iteration. As stated above, boundaries are detected using the Sobel edge detector due to its efficiency.

Figure 7 shows the results by applying the proposed algorithm on a set of noisy images obtained from the noise-free image shown in Fig. 1(a) by adding Gaussian noise with σ of 10, 40, and 60 respectively. Same parameters for smoothing are used for the three images. When noise is relative small, the proposed algorithm preserves boundaries accurately as well as corners and junctions, as shown in Fig. 7(a). When noise is substantial, due to the coupling structure, the proposed algorithm is robust to noise and salient boundaries are well preserved. Because only local information is used in the system, it would be expected that the boundaries are less accurate when noise is larger. This uncertainty is an intrinsic property of the proposed algorithm because reliable estimation gets more difficult when noise gets larger as shown in Fig. 7(b) and (c). The results seem consistent with our perceptual experience.

(Figure 7 should be around here)

Figure 8 shows the result for another synthetic image, which was extensively used by Sarkar and Boyer[27]. As shown in Fig. 8(b), noise is reduced greatly and boundaries as well as corners are well preserved. Even using the simple Sobel edge detector, the result is better than the best result from the optimal infinite impulse responses filters[27] obtained using several parameter combinations with hysteresis thresholding. This is because their edge detector does not consider the responses from neighboring pixels, but rather assumes the local maxima as good edge points.

(Figure 8 should be around here)

Figure 9 shows an image of a grocery store advertisement which was used throughout the book by Nitzberg, Mumford, and Shiota[(28)]. In order to get good boundaries, they first applied an edge detector and then several heuristic algorithms to close gaps and delete noise edges. In our system, the details and noise are smoothed out due to the coupling structure and the salient boundaries, corners and junctions are preserved. The result shown in Fig. 9(c) is comparable with the result after several post-processing steps shown on page 43 of the book.

(Figure 9 should be around here)

Figure 10 shows a high resolution satellite image of a natural scene, consisting of a river, soil land, and a forest. As shown in Fig. 10(b), the river boundary which is partially occluded by the forest is delineated. The textured forest is smoothed out into a homogenous region. The major boundaries between different types of features are detected correctly.

(Figure 10 should be around here)

Figure 11 shows an image of a woman which includes detail features and shading effects, the color version of which was used by Zhu and Yuille[(24)]. In their region competition algorithm, Zhu and Yuille[(29)] used a mixture of Gaussian model. A nonconvex energy function consisting of several constraint terms was formulated under Bayesian framework. The algorithm, derived using variational principles, is guaranteed to converge to only a local minimum. For our nonlinear algorithm, as shown in Fig. 11(b), the details are smoothed out while important boundaries are preserved. The final result in Fig. 11(c) is comparable with the result from the region competition algorithm[(29)] applied on the color version after 130 iterations. Compared with the region competition algorithm, the main advantage of our approach is that local statistical properties are extracted and utilized effectively in the oriented probabilistic framework instead of fitting the image into a global model which, in general, cannot be guaranteed to fit the given data well.

(Figure 11 should be around here)

To further demonstrate the effectiveness of the proposed algorithm, we have also applied it to a texture image as shown in Fig. 12(a). As shown in Fig. 12(b), the boundaries between different textures are preserved while most of detail features are smoothed out. Figure 12(c) shows the detected boundaries by the Sobel edge detector. While there are some noisy responses due to the texture patterns, the main detected boundaries are connected. A simple region growing algorithm would segment the smoothed image into four regions. While this example is not intended to show that our algorithm can process texture images, it demonstrates that the proposed algorithm can be generalized to handle distributions that are not Gaussian, which was assumed when formalizing the algorithm.

(Figure 12 should be around here)

4.2 Comparison with nonlinear smoothing algorithms

In order to evaluate the performance of the proposed algorithm relative to existing nonlinear smoothing methods, we have conducted a comparison with three recent methods. The SUSAN nonlinear filter(25) has been claimed to be the best by integrating smoothing both in spatial and brightness domains. The original anisotropic model by Perona and Malik(8) is still widely used and studied. The edge-enhancing anisotropic diffusion model proposed by Weickert(11)(9) incorporates true anisotropy using a diffusion tensor calculated from a Gaussian kernel, and is probably by far the most sophisticated diffusion-based smoothing algorithm.

To do an objective comparison using real images is difficult because there is no universally accepted ground truth. Here we use synthetic images where the ground truth is known and the deviation calculated by (19) gives an objective measure of the quality of smoothed images. We have also tuned parameters to achieve best possible results for the methods to be compared. For the SUSAN algorithm, we have used several different values for the critical parameter T in (15). For the Perona and Malik model, we have tried different nonlinear functions g in (1) with different parameters. For the Weickert model, we have chosen a good set of parameters for diffusion tensor estimation. We in addition choose their best results in terms of deviation from the ground truth, which are then used for boundary detection.

Because the three methods and proposed algorithm all can be applied iteratively, first we compare their temporal behavior. We apply each of them to the image shown in Fig. 7(b) for 1000 iterations and calculate the deviation and relative variance with respect to the number of iterations using (19) and (20). Figure 13 shows the deviation from the ground-truth image. The SUSAN filter, which quickly reaches a best state, and converges quickly also to a uniform state due to the

Gaussian smoothing term in the filter (see equation (15)). The temporal behavior of the Perona-Malik model and the Weickert model is quite similar while the Weickert model converges more rapidly to and stay longer in good results. The proposed algorithm converges and stabilizes quickly to a non-uniform state, and thus the smoothing can be terminated after several iterations.

(Figure 13 should be around here)

Figure 14 shows the relative variance of the four methods along the iterations. Because the SUSAN algorithm converges to a uniform stable state, the relative variance goes to zero after a number of iterations. The relative variance of Perona-Malik model is closely related to the g function in (1). Due to the spatial regularization using a Gaussian kernel, Weickert model changes continuously and the diffusion lasts much longer, which accounts for the fact why good results exist for a longer period of time than Perona-Malik model. As shown in Fig. 4 and 5, the proposed algorithm generates bounded small ripples in the relative variance measure. Those ripples do not affect smoothing results noticeably as the deviation from the ground truth, shown in Fig. 13, is stabilized quickly.

(Figure 14 should be around here)

Now we compare the effectiveness of the four methods in preserving meaningful boundaries. For the proposed algorithm, we use the result at the 11th iteration because the proposed algorithm converges within several iterations. As mentioned before, for the other three methods, we tune critical parameters and choose the smoothed images with the smallest deviation. Figure 15 shows the smoothed images along with the detected boundaries using the Sobel edge detector, for the image shown in Fig. 11(a), where added noise is Gaussian with zero mean and $\sigma = 10$. All of the four methods perform well. The boundary of the square is preserved accurately. For the central oval, the proposed algorithm gives a better connected boundary while the other three have gaps. Also the proposed algorithm generated the sharpest edges while most edges from Weickert model are blurred.

(Figure 15 should be around here)

Figure 16 shows the result for the image in Fig. 7(b), where noise is substantial. As shown in Fig. 16(a), the SUSAN filter tends to fail to preserve boundaries, resulting in noisy boundary fragments. The Perona-Malik model produces good but fragmented boundaries. The Weickert model produces good boundaries for strong segments but weak segments are blurred considerably. The proposed algorithm preserves the connected boundary of the square and partially fragmented boundaries of the central oval also. As shown in Fig. 13, the smoothed image of the Weickert model has a smaller deviation than the result from our algorithm, but the detected boundaries are fragmented. This is because our algorithm produces sharp boundaries, which induce larger penalties according to (19) when not accurately marked.

(Figure 16 should be around here)

We have also applied the four methods on the natural satellite image shown in Fig. 10. The result from the proposed algorithm is at the 11th iteration as already shown in Fig. 10. The results from the other three methods are picked up manually for best possible results. Due to the termination problem, results from most nonlinear smoothing algorithms have to be chosen manually, making them difficult to be used automatically. The results from the other three methods are similar, and boundaries between different regions are not formed. In contrast, our algorithm generated connected boundaries separating major different regions.

(Figure 17 should be around here)

5. Conclusions

In this paper we have presented a two-step robust boundary detection algorithm. The first step is a nonlinear smoothing algorithm based on an orientation sensitive probability measure. This algorithm is motivated by the orientation sensitivity of cells in the visual cortex(22). By incorporating geometrical constraints through the coupling structure, the algorithm is robust to noise while preserving meaningful boundaries. Even though the algorithm was formulated based on Gaussian distribution, it performs well for real and even textured images, showing the generalization capability of the algorithm. It is also easy to see that the formalization of the proposed algorithm would extend to other known distribution by changing equations (3)-(5) accordingly. One such an extension would be to use a mixture of Gaussian distributions(30) so that the model may be able to describe arbitrary probability distribution.

Compared with recent anisotropic diffusion methods, our algorithm approaches a non-uniform stable state and reliable results can be obtained after a fixed number of iterations. In other words, it provides a solution for the termination problem. When noise is substantial, our algorithm preserves meaningful boundaries better than the diffusion-based methods, because the coupling structure employed is more robust than pair-wise coupling structure.

Scale is an intrinsic parameter in machine vision as interesting features may exist only in a limited range of scales. Scale spaces based on linear and nonlinear smoothing kernels do not represent semantically meaningful structures explicitly(31). A solution to the problem could be to use parameter K in equation (8) as a control parameter(23), which is essentially a threshold in gray values. Under this formalization, (12) could offer an adaptive parameter selection. With the robust coupling structure, our algorithm with adaptive parameter selection may be able to provide a robust multiscale boundary detection method.

Another advantage of the probability measure framework is that there is no need to assume a priori knowledge about each region, which is necessary in relaxation labeling(14-15) and the comparison across windows with different sizes and shapes is feasible. This could lead to an adaptive window selection that preserves small but important features which cannot be handled well by the current implementation.

There is one intrinsic limitation common to many smoothing approaches including our proposed one. After smoothing, the available feature is the average gray value, resulting in loss of information for further processing. One way to overcome this problem is to apply the smoothing in feature spaces derived from input images(32). Another disadvantage of the proposed algorithm is relatively intensive computation due to the use of oriented windows. Each oriented window takes roughly as long in one iteration as the edge-enhancing diffusion method(11). On the other hand, because our algorithm is entirely local and parallel, computation time would not be a problem on parallel and distributed hardware. Computation on serial computers could be reduced dramatically by decomposing the oriented filters hierarchically so that oriented windows would be used only around discontinuities rather than in homogenous regions. The decomposition techniques for steerable and scalable filters(33) could also help to reduce the number of necessary convolution kernels.

Acknowledgments. The authors would like to thank Dr. Ke Chen and Erdogan Cesmeli for their useful discussions and Dr. Ke Chen for letting us use his implementation of the edge-enhancing anisotropic diffusion algorithm(11). This work was partially supported by an NSF grant (IRI-9423312) and an ONR Young Investigator Award (N00014-96-1-0676) to DLW. The authors also thank National Imagery and Mapping Agency and the Ohio State University Center for Mapping for supporting this research. Thanks to K. Boyer, S. Zhu, and M. Nitzberg for making their test images available on-line.

References

1. K. R. Castleman, *Digital Image Processing*, Prentice Hall, Englewood Cliffs, NJ (1996).
2. J. Canny, A computational approach to edge detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**, 679-698 (1986).
3. D. Marr and E. Hildreth, Theory of edge detection, *Proceedings of Royal Society London B* **207**, 187-217 (1980).
4. A. Witkin, Scale-space filtering, In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 1023-1026 (1983).
5. J. Koenderink, The structure of images, *Biological Cybernetics* **50**, 363-370 (1984).
6. T. Linderberg, *Scale-Space Theory in Computer Vision*, Kluwer Academic Publishers, Dordrecht, Netherlands (1994).
7. M. A. Cohen and S. Grossberg, Neural dynamics of brightness perception: features, boundaries, diffusion and resonance, *Perception and Psychophysics* **36**, 428-456 (1984).
8. P. Perona and J. Malik, Scale space and edge detection using anisotropic diffusion, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**, 16-27 (1990).
9. J. Weickert, A review of nonlinear diffusion filtering, In *Proceedings of the First International Conference on Scale-Space*, 3-28 (1997).
10. F. Catte, P.-L. Lions, J.-M. Morel, and T. Coll, Image selective smoothing and edge detection by nonlinear diffusion, *SIAM Journal on Numerical Analysis* **29**, 182-193 (1992).
11. J. Weickert, Theoretical foundations of anisotropic diffusion in image processing, In *Theoretical Foundations of Computer Vision*, also *Computing Supplement* **11**, W. Kropatsch, R. Klette, and F. Solina, eds, Springer-Verlag, Wien, Austria, 221-236 (1996).
12. Y. L. You, W. Xu, A. Tannenbaum, and M. Kaveh, Behavioral analysis of anisotropic diffusion in image processing, *IEEE Transactions on Image Processing* **5**, 1539-1553 (1996).
13. P. Perona, T. Shiota, and J. Malik, Anisotropic diffusion, In *Geometry-Driven Diffusion in Computer Vision*, B. M. ter Haar Romeny, ed, Kluwer Academic Publishers, Dordrecht, Netherlands, 73-92 (1994).
14. A. Rosenfeld, R. A. Hummel, and S. W. Zucker, Scene labeling by relaxation operations, *IEEE Transactions on Systems, Man, and* Cybernetics **6**, 420-433 (1976).
15. R. A. Hummel and S. W. Zucker, On the foundations of relaxation labeling processes, *IEEE*

Transactions on Pattern Analysis and Machine Intelligence **5**, 267-287 (1983).

16. C. H. Li and C. K. Lee, Image smoothing using parametric relaxation, *Graphical Models and Image Processing* **57**, 161-174 (1995).
17. M. W. Hansen and W. E. Higgins, Relaxation methods for supervised image segmentation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **19**, 949-962 (1997).
18. D. Marr, *Vision*, W. H. Freeman, San Francisco (1982).
19. H. Stark and J. W. Woods, *Probability, Random Processes and Estimation Theory for Engineers*, Prentice-Hall, Englewood Cliffs, NJ (1994).
20. F. Tomita and S. Tsuji, Extraction of multiple regions by smoothing in selected neighborhoods, *IEEE Transactions on Systems, Man, and Cybernetics* **7**, 107-109 (1977).
21. M. Nagao and T. Matsuyama, Edge preserving smoothing, *Computer Graphics and Image Processing* **9**, 394-407 (1979).
22. D. H. Hubel, *Eye, Brain, and Vision*, W. H. Freeman and Company, New York (1988).
23. P. Saint-Marc, J.-S. Chen, and G. Medioni, Adaptive smoothing: a general tool for early vision, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**, 514-529 (1991).
24. D. C. C. Wang, A. H. Vagnucci, and C. C. Li, Gradient inverse weighted smoothing scheme and the evaluation of its performance, *Computer Graphics and Image Processing* **15**, 167-181 (1981).
25. S. M. Smith and J. M. Brady, SUSAN - a new approach to low level image processing, *International Journal of Computer Vision* **23**, 45-78 (1997).
26. T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*, MIT Press, Cambridge, MA (1997).
27. S. Sarkar and K. L. Boyer, On optimal infinite impulse response edge detection filters, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**, 1154-1171 (1991).
28. M. Nitzberg, D. Mumford, and T. Shiota, *Filtering, Segmentation, and Depth*, Springer-Verlag, Berlin (1993).
29. S. C. Zhu and A. Yuille, Region competition: unifying snakes, region growing, and Bayes/MDL for multiband image segmentation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **18**, 884-900 (1996).
30. B. S. Everitt and D. J. Hand, *Finite Mixture Distributions*, Chapman and Hall, London (1981).

31. M. Tabb and N. Ahuja, Multiscale image segmentation by integrated edge and region detection, *IEEE Transactions on Image Processing* **6**, 642-655 (1997).

32. R. Whitaker and G. Gerig, Vector-valued diffusion, In *Geometry-Driven Diffusion in Computer Vision*, B. M. ter Haar Romeny, ed, Kluwer Academic Publishers, Dordrecht, Netherlands, 93-134 (1994).

33. P. Perona, Deformable kernels for early vision, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**, 488-499 (1995).

Figure Caption

Figure 1. An example with non-uniform boundary gradients and substantial noise. (a) A noise-free synthetic image. (b) A noisy version of (a) with Gaussian noise of $\sigma = 40$. (c) Local gradient map of (b) using the Sobel operators. (d)-(f) Smoothed images from an anisotropic diffusion algorithm[13] at 50, 100, and 1000 iterations. (g)-(i) Corresponding edge maps of (d)-(f) respectively using the Sobel edge detector.

Figure 2. Illustration of the coupling structure of the proposed algorithm. (a) Eight oriented windows and a fully connected window defined on a 3 x 3 neighborhood. (b) The resulting coupling structure. There is a directed edge from (i_1, j_1) to a neighbor (i_0, j_0) if and only if (i_1, j_1) contributes to the smoothing of (i_0, j_0) according to equations (12) and (9). Filled circles represent black pixels, empty circles represent white pixels, and hatched circles represent gray pixels. Ties in (9) are broken according to left-right and top-down preference of the oriented windows in (a).

Figure 3. Temporal behavior of the proposed algorithm with respect to the amount of noise. Six noisy images are obtained by adding zero-mean Gaussian noise with σ of 5, 10, 20, 30, 40 and 60, respectively, to the noise-free image shown in Fig. 1(a). The plot shows the deviation from the ground truth image with respect to iterations of the noise-free image and six noisy images.

Figure 4. Relative variance of the proposed algorithm for the noise-free image shown in Fig. 1(a) and four noisy images with Gaussian noise of zero-mean and σ of 5, 20, 40 and 60, respectively.

Figure 5. Relative variance of the proposed algorithm for real images shown in Fig. 9-12.

Figure 6. The oriented bar-like windows used throughout this paper for synthetic and real images.

Figure 7. The smoothed images at the 11th iteration and detected boundaries for three synthetic images with different amount of noise. Top row shows the input images, middle the smoothed image at the 11th iteration, and bottom the detected boundaries using the Sobel edge detector.

Figure 8. The smoothed image at the 11th iteration and detected boundaries for a synthetic image with corners.

Figure 9. The smoothed image at the 11th iteration and detected boundaries for a grocery store advertisement. Details are smoothed out while major boundaries and junctions are preserved accurately.

Figure 10. The smoothed image at the 11th iteration and detected boundaries for a natural satellite image with several land use patterns. The boundaries between different regions are formed from noisy segments due to the coupling structure.

Figure 11. The smoothed image at the 11th iteration and detected boundaries for a woman image. While the boundaries between large features are preserved and detected, detail features such as facial features are smoothed out.

Figure 12. The smoothed image at the 11th iteration and detected boundaries for a texture image. The boundaries between different textured regions are formed while details due to textures are smoothed out.

Figure 13. Deviations from the ground truth image for the four nonlinear smoothing methods. Dashed line: The SUSAN filter(25); Dotted line: The Perona-Malik model(8); Dash-dotted line: The Weickert model of edge enhancing anisotropic diffusion(11); Solid line: The proposed algorithm.

Figure 14. Relative variance of the four nonlinear smoothing methods. Dashed line: The SUSAN filter(25); Dotted line: The Perona-Malik diffusion model(8); Dash-dotted line: The Weickert model(11); Solid line: The proposed algorithm.

Figure 15. Smoothing results and detected boundaries of the four nonlinear methods for a synthetic image shown in Fig. 7(a). Here noise is not large and all of the methods perform well in preserving boundaries.

Figure 16. Smoothing results and detected boundaries of the four nonlinear methods for a synthetic image with substantial noise shown in Fig. 7(b). The proposed algorithm generates sharper and better connected boundaries than the other three methods.

Figure 17. Smoothing results and detected boundaries of a natural scene satellite image shown in Fig. 10. Smoothed image of the proposed algorithm is at the 11th iteration while smoothed images of the other three methods are chosen manually. While the other three methods generate similar fragmented boundaries, the proposed algorithm forms the boundaries between different regions due to its coupling structure.

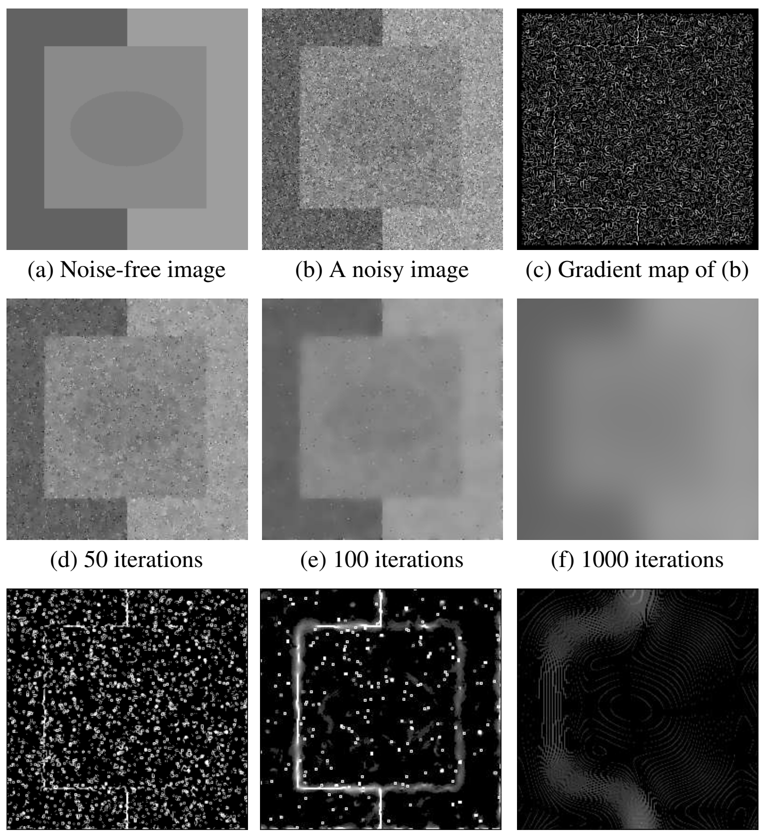

(a) Noise-free image (b) A noisy image (c) Gradient map of (b)

(d) 50 iterations (e) 100 iterations (f) 1000 iterations

(g) Edge map of (d) (h) Edge map of (e) (i) Edge map of (f)

Figure 1

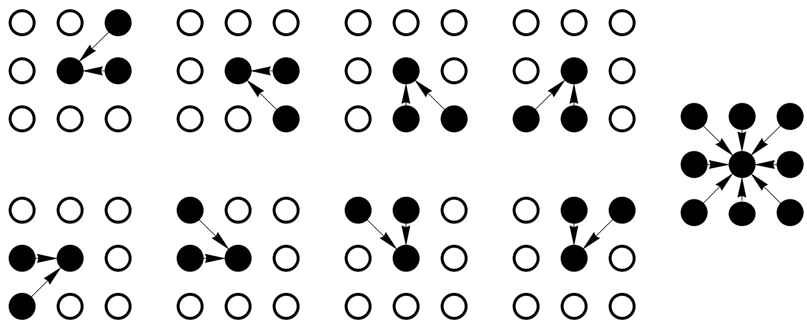

(a) Oriented windows

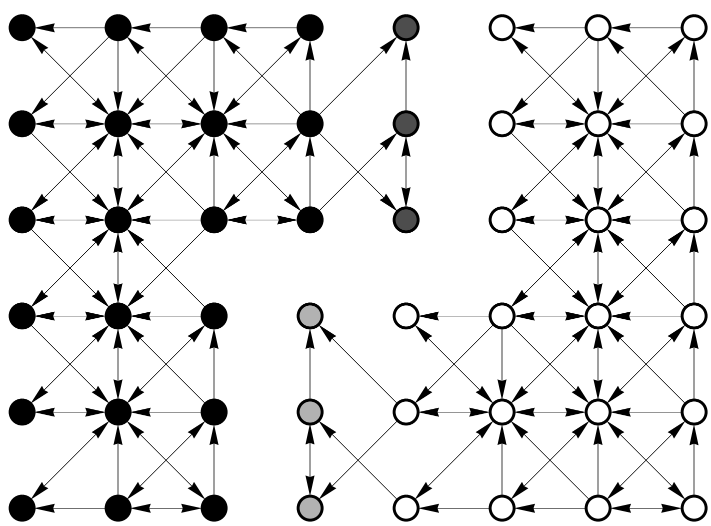

(b) Coupling structure of a small image

Figure 2

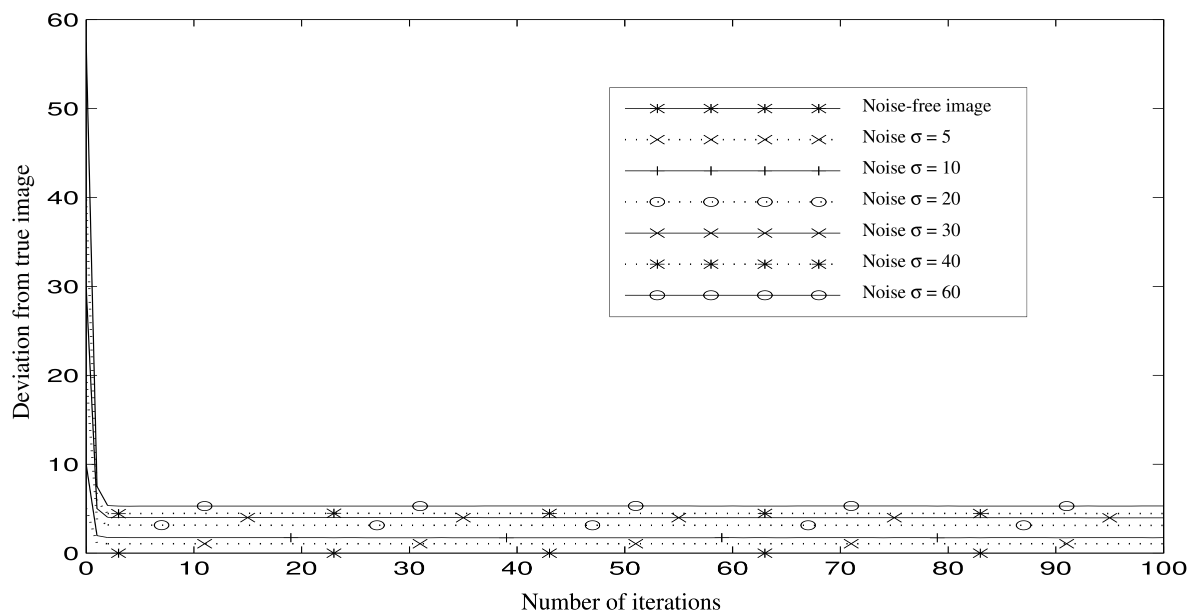

Noise-free image
Noise σ = 5
Noise σ = 10
Noise σ = 20
Noise σ = 30
Noise σ = 40
Noise σ = 60
Deviation from true image
Number of iterations
0
10
20
30
40
50
60
70
80
90
100

Figure 3

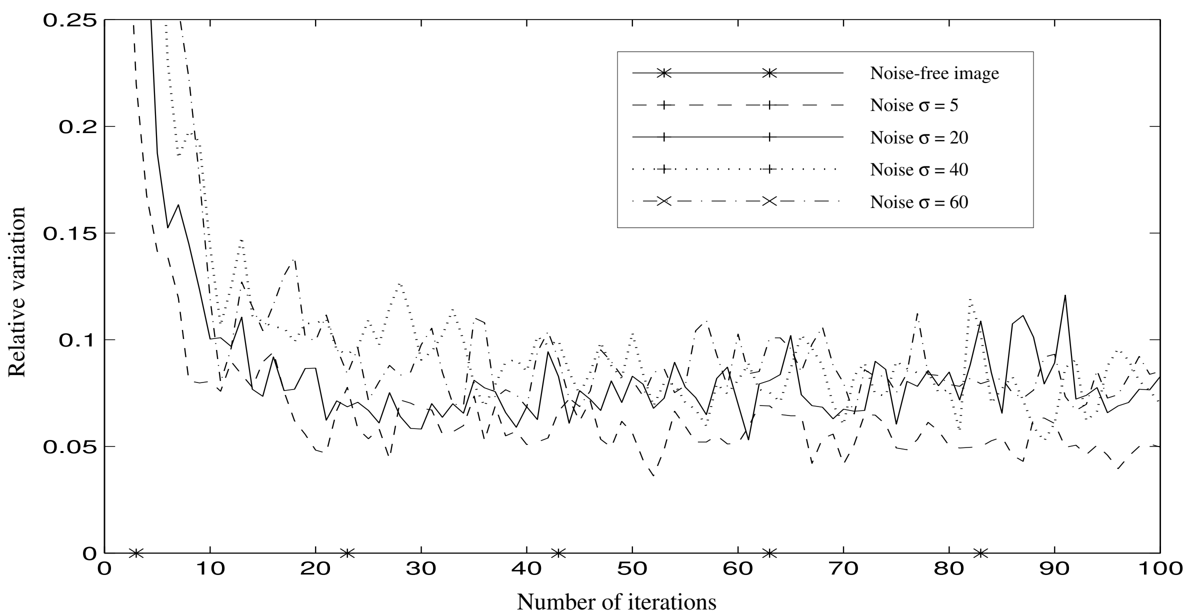
0.25
0.2
0.15
0.1
0.05
0
Relative variation
0
10
20
30
40
50
60
70
80
90
100
Number of iterations
Noise-free image
Noise σ = 5
Noise σ = 20
Noise σ = 40
Noise σ = 60

Figure 4

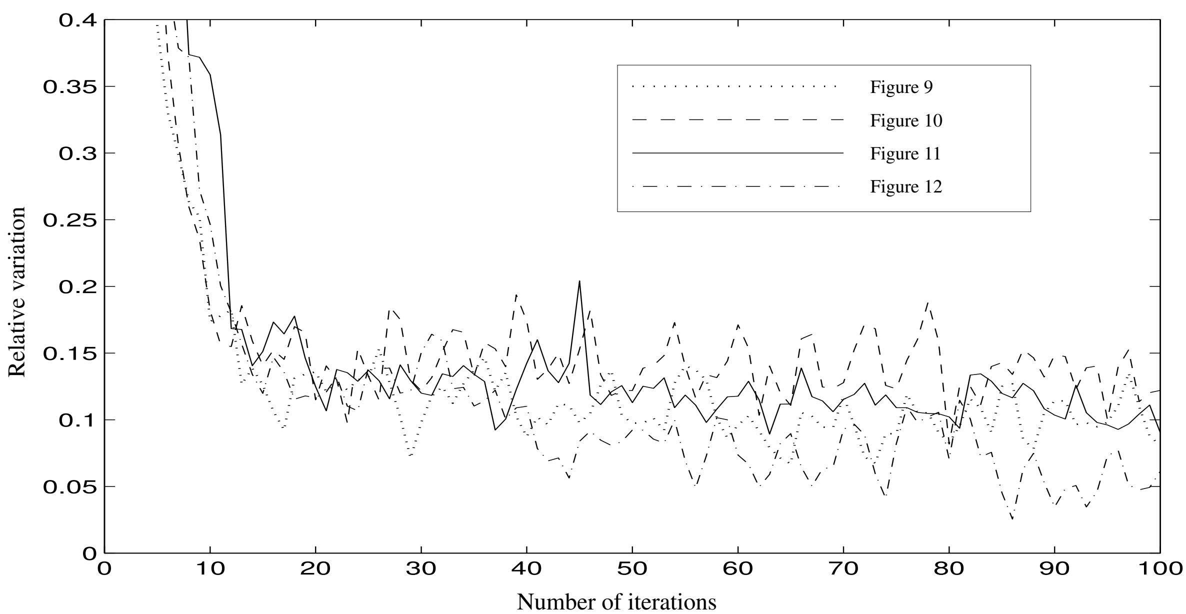
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
Relative variation
0
10
20
30
40
50
60
70
80
90
100
Number of iterations
Figure 9
Figure 10
Figure 11
Figure 12

Figure 5

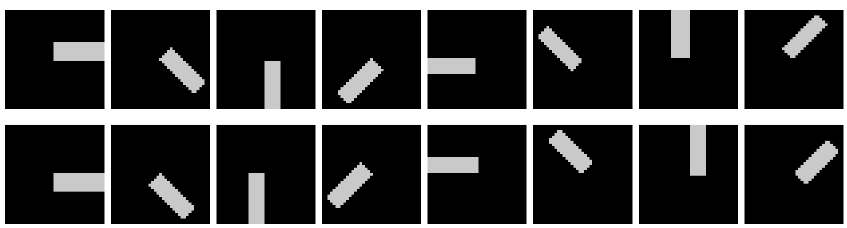

Figure 6

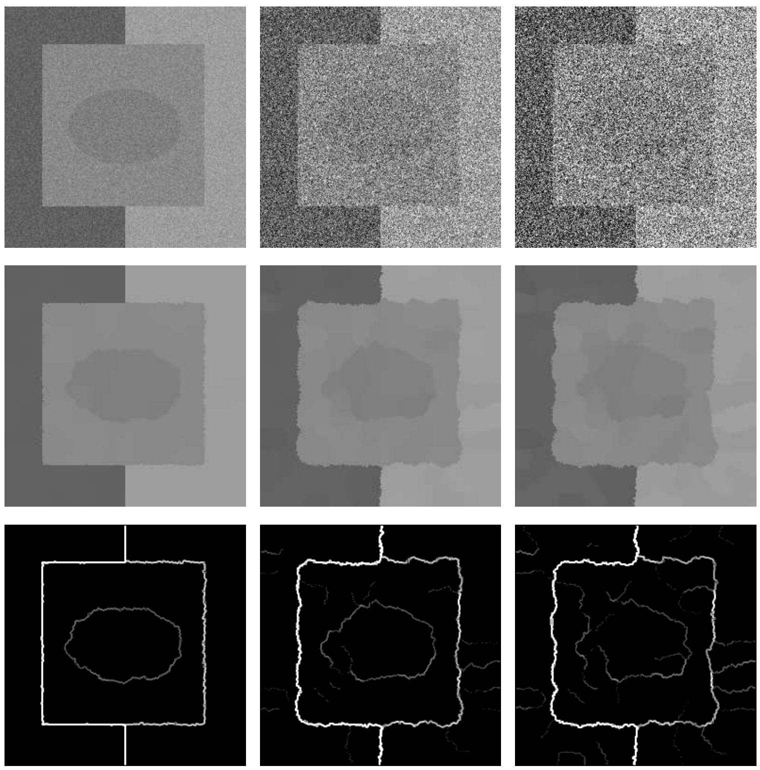

(a) Noise σ = 10 (b) Noise σ = 40 (c) Noise σ = 60

Figure 7

(a) Input image (b) Smoothed image (c) Boundaries detected

Figure 8

(a) Input image (b) Smoothed image (c) Boundaries detected

Figure 9

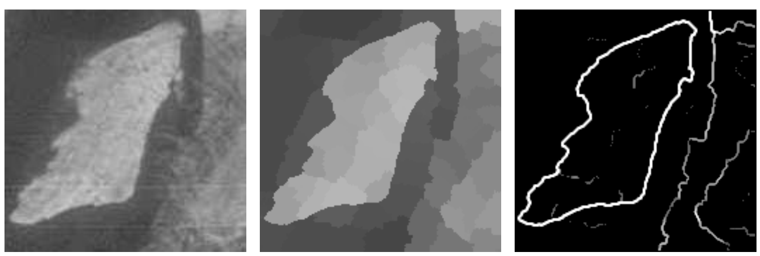

(a) Input image (b) Smoothed image (c) Boundaries detected

Figure 10

(a) Input image (b) Smoothed image (c) Boundaries detected

Figure 11

(a) Input image (b) Smoothed image (c) Boundaries detected

Figure 12

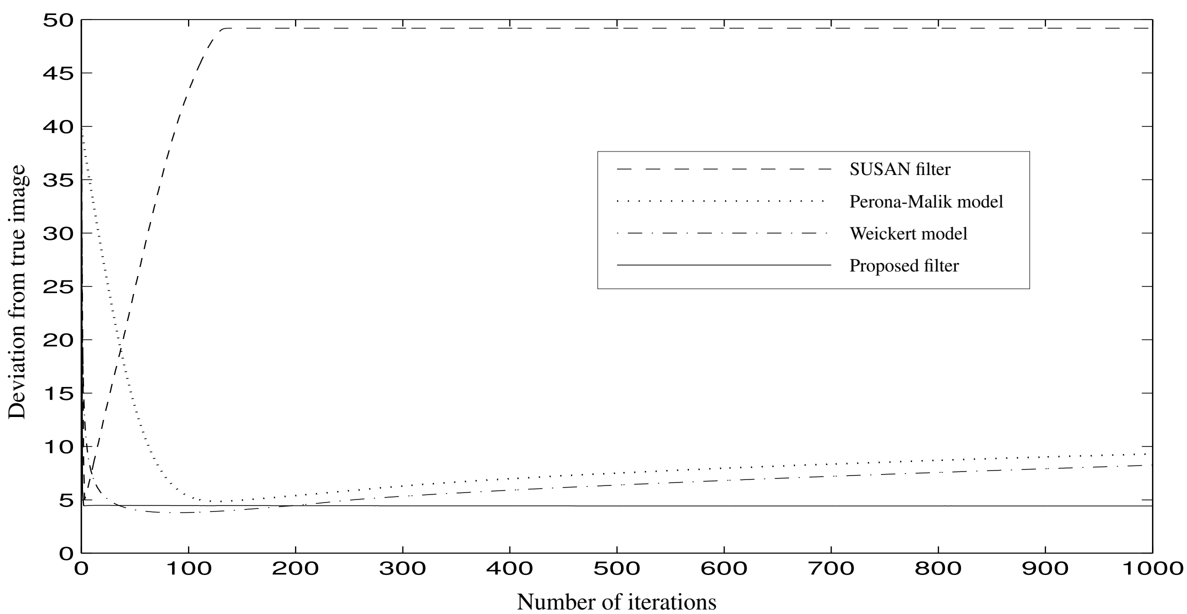

50
45
40
35
30
25
20
15
10
5
0
0
100
200
300
400
500
600
700
800
900
1000
Deviation from true image
Number of iterations
SUSAN filter
Perona-Malik model
Weickert model
Proposed filter

Figure 13

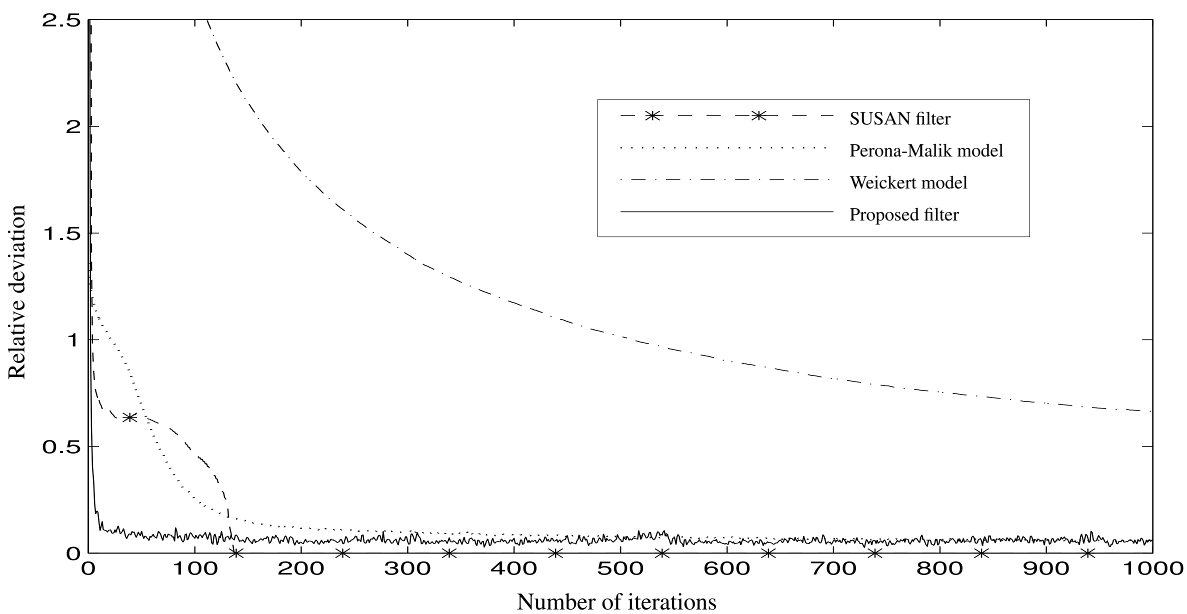

2.5
2
1.5
1
0.5
0
Relative deviation
0
100
200
300
400
500
600
700
800
900
1000
Number of iterations
SUSAN filter
Perona-Malik model
Weickert model
Proposed filter

Figure 14

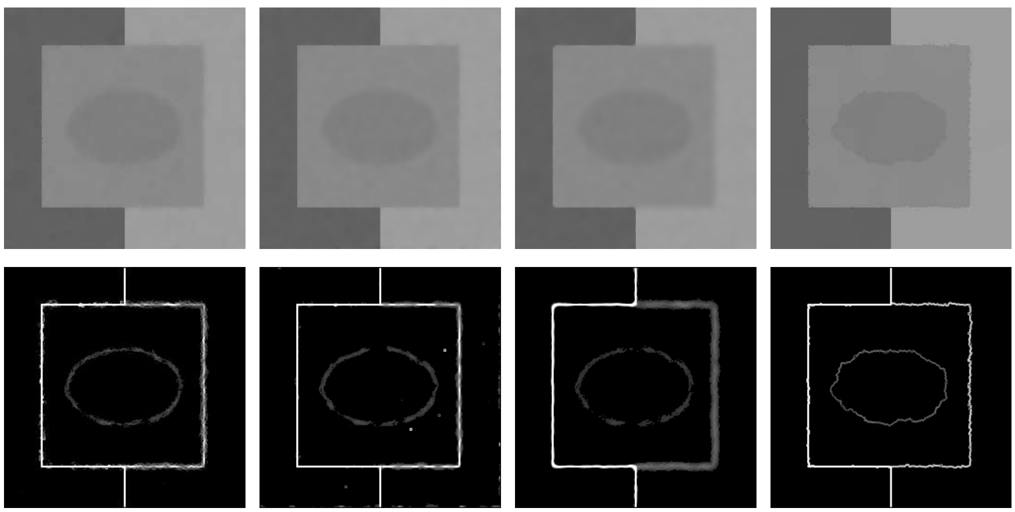

(a) SUSAN (b) Perona-Malik model (c) Weickert model (d) Proposed filter

Figure 15

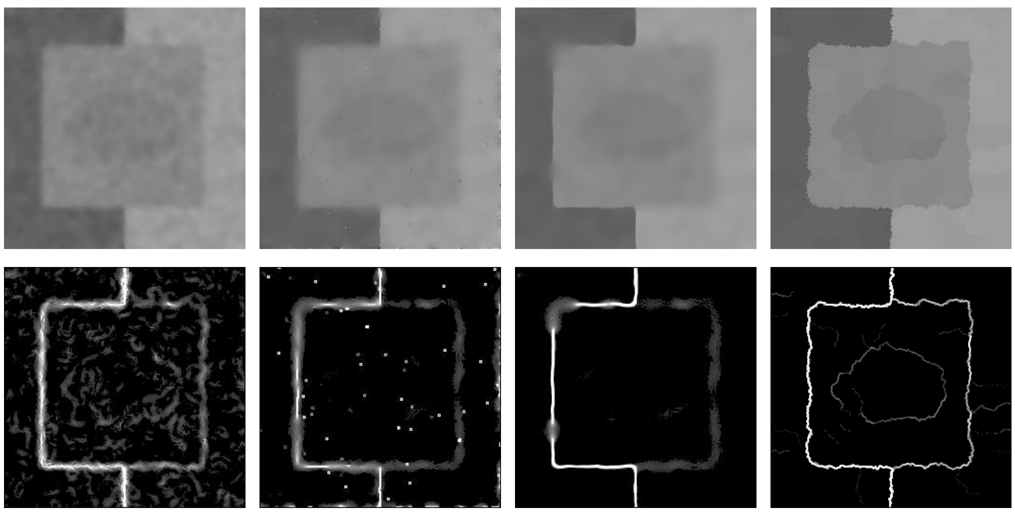

(a) SUSAN (b) Perona-Malik model (c) Weickert model (d) Proposed filter

Figure 16

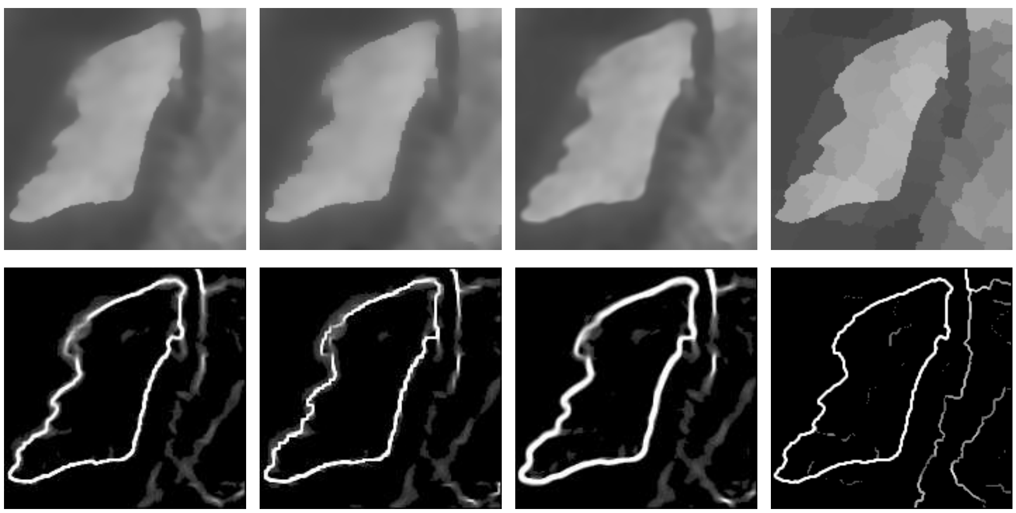

(a) SUSAN (b) Perona-Malik Model (c) Weickert model (d) Proposed filter

Figure 17