November, 1996

LIDS- P 2368

Research Supported By:

AFOSR grant F49620-95-1-0083
ONR grant N00014-91-J-1004
BU. GC123919NGD
AFOSR Multidisciplinary Research

Image Segmentation and Edge Enhancement with Stabilized Inverse Diffusion Equations

Pollak, I.
Willsky, A.S.
Krim, H.

LIDS-P-2368

Image Segmentation and Edge Enhancement with Stabilized Inverse Diffusion Equations.[1]

Ilya Pollak[2]

e-mail: ipollak@mit.edu, tel: 617-253-3816, fax: 617-258-8553, Room 35-427

Alan S. Willsky

Hamid Krim

EDICS numbers: IP 1.6, IP 1.5, IP 1.3.

Abstract.[3] We introduce a family of first-order multi-dimensional ordinary differential equations (ODEs) with discontinuous right-hand sides and demonstrate their applicability in image processing. An equation belonging to this family is an inverse diffusion everywhere except at local extrema, where some stabilization is introduced. For this reason, we call these equations "stabilized inverse diffusion equations" ("SIDEs"). Existence and uniqueness of solutions, as well as stability, are proven for SIDEs. A SIDE in one spatial dimension may be interpreted as a limiting case of a semi-discretized Perona-Malik equation [16, 17]. In an experimental section, SIDEs are shown to suppress noise while sharpening edges present in the input signal. Their application to image segmentation is demonstrated.

[1]The work of the authors was supported in part by AFOSR grant F49620-95-1-0083, ONR grant N00014-91-J-1004, and by subcontract GC123919NGD from Boston University under the AFOSR Multidisciplinary Research Program on Reduced Signature Target Recognition.

[2]The authors are affiliated with the Department of Electrical Engineering and Computer Science and the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139.

[3]Permission to publish this abstract separately is granted.

1 Introduction

In this paper we introduce, analyze, and apply a new class of nonlinear image processing algorithms. These algorithms are motivated by the great recent interest in using evolutions specified by partial differential equations (PDE's) as image processing procedures for tasks such as edge enhancement and segmentation, among others. In a sense that we will make precise, our algorithms can be viewed as a logical extension of one line of thought involving such evolution equations and, in some sense, as a limiting case, which has both some important mathematical properties as well as what we believe is considerable promise for edge enhancement and segmentation, especially in the presence of noise.

To understand the conceptual basis for our approach, it is useful to briefly review one of the lines of thought that has spurred work in evolution-based methods for image analysis. In [26] Witkin proposed filtering an original image $u_0(x, y)$ with Gaussian kernels of variance t, to result in a one-parameter family of images $u(x, y, t)$ he referred to as "a scale space". This filtering technique has both a very important interpretation and a number of significant limitations that inspired the search for alternative scale spaces that are better adapted to edge detection and image segmentation. In particular, a major limitation is that linear Gaussian smoothing blurs and displaces edges, merges boundaries of objects that are close to each other, and removes edge junctions [22]. However, the important insight found, for example, in [10], is that the family of images $u(x, y, t)$ is the solution of the linear heat equation with $u_0(x, y)$ as the initial data. This insight led to the pursuit and development of a new paradigm for processing images via the evolution of nonlinear PDEs [16, 17, 15, 1] which effectively lift the limitations of the linear heat equation. In addition, thanks to the interpretation of the heat equation as the steepest descent equation for the functional $\int |\nabla u|^2 \, dx \, dy$, there has been a great deal of activity in defining functionals adapted to various problems in image processing [14, 2, 21, 19, 18, 3, 13]. One such formulation is that of deformable contours and surfaces [8, 23, 4, 12, 5, 9, 20] which provides another framework for PDE-based segmentation.

While the analysis of the techniques mentioned above is most often performed in the continuous setting, where an image is identified with a function of two continuous spatial variables, the implementation of such equations generally involves their discrete approximation. As a consequence, as Weickert pointed out in [25], "a scale-space representation cannot perform better than its discrete realization". Following this suggestion, we concentrate in this paper on semi-discrete scale spaces (i.e., continuous in scale (or time) and discrete in space). More specifically, the main contribution of this paper is a new family of semi-discrete evolution equations which stably sharpen edges and suppress noise. The starting point for the development of these equations is a discrete interpretation of anisotropic diffusions such as that used by Perona-Malik [16,17]. One motivation for such equations is precisely that of achieving both noise removal and edge enhancement through the use of a diffusion-like equation which in essence acts as an unstable inverse diffusion near edges and as a stable linear-heat-equation-like diffusion in homogeneous regions without edges. In a sense that we will make both precise and conceptually clear, the evolutions that we introduce may be viewed as a conceptually limiting case of such diffusions. These evolutions have discontinuous right-hand sides and act as inverse diffusions "almost everywhere" with stabilization resulting from the presence of the discontinuities in the vector field defined by the evolution. As we will see, the scale space of such an equation is a family of segmentations of the original image, with larger values of the scale parameter t corresponding to segmentations at coarser resolutions. Moreover, in contrast to continuous evolutions, the ones introduced here naturally define a sequence of logical "stopping times", i.e. points along the evolution fraught with useful information one may wish to extract, and corresponding to times at which the evolution hits a discontinuity surface of its solution field. These times are data-adaptive, i.e., they depend on the initial image, and result in a sequence of images at increasingly coarser resolutions, where the resolutions are adapted to the image being analyzed.

In the next section we begin by describing a convenient mechanical analog for the visualization of many spatially-discrete evolution equations, including discretized linear or nonlinear diffusions such as that of Perona-Malik, as well as the discontinuous equations that we introduce in Section 3.

Because of the discontinuous right-hand side, some care must be taken in defining solutions, but as we show in Section 4, once this is done, the resulting evolutions have a number of important properties. Moreover, as we have indicated, they lead to very effective algorithms for edge enhancement and segmentation, something that we demonstrate in Section 5. In particular, as we will see, they can produce sharp enhancement of edges in high noise as well as accurate segmentations of very noisy imagery such as synthetic aperture radar (SAR) imagery subject to severe speckle.

2 A Spring-Mass Model for Certain Evolution Equations

As we indicated in the introduction, the focus of this paper is on discrete-space, temporally-continuous evolutions of the following general form

$$\begin{aligned} \dot{\mathbf{u}}(t) &= \mathcal{F}(\mathbf{u})(t), \\ \mathbf{u}(0) &= \mathbf{u}_0, \end{aligned} \tag{1}$$

where $\mathbf{u}$ is either a discretized signal, i.e., an N-point discrete sequence ($\mathbf{u} = (u_1, ..., u_N)^T \in \mathbb{R}^N$), or an N-by-N image whose j-th entry in the i-th row is u_{ij} ($\mathbf{u} \in \mathbb{R}^{N^2}$). The initial condition $\mathbf{u}_0$ corresponds to the original signal or image to be processed, and $\mathbf{u}(t)$ then represents the evolution of this signal/image at time (scale) t, resulting in a scale-space family for $0 \leq t < \infty$.

The nonlinear operators $\mathcal{F}$ of interest in this paper can be conveniently visualized through the following simple mechanical model. For the sake of simplicity in visualization, let us first suppose that $\mathbf{u} \in \mathbb{R}^N$ is a one-dimensional (1-D) sequence, and interpret $\mathbf{u}(t) = (u_1(t), ..., u_N(t))^T$ in (1) as the vector of vertical positions of the N particles of masses $M_1, ..., M_N$, depicted in Figure 1. The particles are forced to move along N vertical lines. Each particle is connected by springs to its two neighbors (except the first and last particles, which are only connected to one neighbor.) Every spring whose vertical extent is v has energy $E(v)$, i.e., the energy of the spring between the n-th and $(n+1)$-st particles is $E(u_{n+1} - u_n)$. We impose the usual requirements for an energy function:

$$
\begin{aligned}
E(v) &\geq 0, \qquad E(0) = 0, \\
E'(v) &\geq 0 \text{ for } v > 0, \\
E(v) &= E(-v).
\end{aligned} \tag{2}
$$

Then the derivative of $E(v)$, which we refer to as "the force function" and denote by $F(v)$, satisfies

$$
\begin{aligned}
F(0) &= 0, \qquad F(v) \geq 0 \text{ for } v > 0, \\
F(v) &= -F(-v).
\end{aligned} \tag{3}
$$

We make the movement of the particles non-conservative by stopping it after a small period of time Δt and re-starting with zero velocity. We assume that during one such step, the total force $F_n = -F(u_n - u_{n+1}) - F(u_n - u_{n-1})$, acting on the n-th particle, stays approximately constant. The displacement during one iteration is equal to the product of acceleration and the square of the time interval, divided by two:

$$
u_n(t + \Delta t) - u_n(t) = \frac{(\Delta t)^2}{2} \frac{F_n}{M_n}.
$$

Letting $\Delta t \to 0$, while fixing $\frac{2M_n}{\Delta t} = m_n$, where m_n is a positive constant, leads to

$$
\dot{u}_n = \frac{1}{m_n}(F(u_{n+1} - u_n) - F(u_n - u_{n-1})), \quad n = 1, 2, ..., N, \tag{4}
$$

with the conventions $u_0 = u_1$ and $u_{N+1} = u_N$ imposed by the absence of springs to the left of the first particle and to the right of the last particle. We will refer to m_n as "the mass of the n-th particle" in the remainder of the paper. In the three examples below, we set $m_n = 1$.

Example 1. A linear force function $F(v) = v$ leads to the semi-discrete linear heat equation

$$
\dot{u}_n = u_{n+1} - 2u_n + u_{n-1}.
$$

This corresponds to a simple discretization of the 1-D linear heat equation and results in evolutions which produce increasingly low-pass filtered and smoothed versions of the original signal $\mathbf{u}_0$. ■

In general, we call $F(v)$ a "diffusion force" if, in addition to (3), it is monotonously increasing:

$$v_1 < v_2 \quad \Rightarrow \quad F(v_1) < F(v_2), \tag{5}$$

which is illustrated in Figure 2. We shall call the corresponding energy a "diffusion energy" and the corresponding evolution (4) a "diffusion". The evolution in Example 1 is clearly a diffusion.

Example 2. The force function $F(v) = -v$ results in the linear inverse diffusion equation

$$\dot{u}_n = -u_{n+1} + 2u_n - u_{n-1}.$$

In contrast to the linear diffusion in Example 1, this evolution tends to accentuate and increase variations in $\mathbf{u}$ rather than blur them. It would thus appear that such an inverse diffusion might have the capacity to enhance edges. Note, however, that the resulting evolution is unstable. ■

In general, we shall call a monotonously decreasing force $F(v)$ an "inverse diffusion force":

$$v_1 < v_2 \quad \Rightarrow \quad F(v_1) > F(v_2), \tag{6}$$

as displayed in Figure 3. We shall call the corresponding energy an "inverse diffusion energy" and the corresponding evolution (4) an "inverse diffusion". As in Example 2, inverse diffusions have the characteristic of enhancing abrupt differences in $\mathbf{u}$ corresponding to "edges" in the 1-D sequence. As also seen in the example however, such pure inverse diffusions lead to unstable evolutions. The following example, which is prototypical of the examples considered by Perona and Malik, defines a stable evolution that captures at least some of the edge enhancing characteristics of inverse diffusions.

Example 3. Taking $F(v) = v \exp\left(-(\frac{v}{K})^2\right)$, as illustrated in Figure 4, yields a 1-D semi-discrete (continuous in scale and discrete in space) version of the Perona-Malik equation (see equations (3.3), (3.4), and (3.12) in [17]). In general, given a positive constant K, we shall call a force $F(v)$ a "Perona-Malik force of thickness K" if, in addition to (3), it satisfies the following conditions:

$$F(v) \ \text{ has a unique maximum at } \ v = K, \tag{7}$$

$$F(v_1) = F(v_2) \quad \Rightarrow \quad (|v_1| - K)(|v_2| - K) < 0.$$

We shall call the corresponding energy a "Perona-Malik energy" and the corresponding evolution a "Perona-Malik evolution of thickness K". As Perona and Malik demonstrate (and as can also be inferred from our results), evolutions with such a force function act like inverse diffusions in the regions of high gradient and like usual diffusions elsewhere. They are stable and capable of achieving some level of edge enhancement depending on the exact form of $F(v)$. ■

Finally, to extend our mechanical model of Figure 1 to images, we simply replace the sequence of vertical lines along which the particles move with an N-by-N square grid of such lines. The particle at location (i, j) is connected by springs to its four neighbors: $(i-1, j)$, $(i, j+1)$, $(i+1, j)$, $(i, j-1)$, except for the particles in the four corners of the square (which only have two neighbors each), and the rest of the particles on the boundary of the square (which have three neighbors). The view from above of this arrangement is depicted in Figure 5. It is reminiscent of (and, in fact, was suggested by) the resistive network of Figure 8 in [16]. The analog of the equation (4) for images is then:

$$\dot{u}_{ij} = \frac{1}{m_{ij}}(F(u_{i+1,j} - u_{ij}) - F(u_{ij} - u_{i-1,j}) + F(u_{i,j+1} - u_{ij}) - F(u_{ij} - u_{i,j-1})), \tag{8}$$

with $i = 1, 2, ..., N,\ \ j = 1, 2, ..., N$, and the conventions $u_{0,j} = u_{1,j}$, $u_{N+1,j} = u_{N,j}$, $u_{i,0} = u_{i,1}$ and $u_{i,N+1} = u_{i,N}$ imposed by the absence of springs outside of $1 \leq i \leq N$, $1 \leq j \leq N$.

3 Stabilized Inverse Diffusion Equations (SIDEs): the Definition

In this section, we introduce a discontinuous force function, resulting in a system (4) that has discontinuous right-hand side (RHS). Such equations received much attention in control theory because of the wide usage of relay switches in automatic control systems. More recently, deliberate introduction of discontinuities has been used in control applications to drive the state vector onto lower-dimensional surfaces in the state space [24]. As we will see, this objective of driving a trajectory onto a lower-dimensional surface also has value in image analysis and in particular in image segmentation. Segmenting a signal or image, represented as a high-dimensional vector $\mathbf{u}$, consists of evolving it so that it is driven onto a comparatively low-dimensional subspace which corresponds to a segmentation of the signal or image domain into a small number of regions.

The type of force function of interest to us here is illustrated in Figure 6. More precisely, we wish to consider force functions $F(v)$ which, in addition to (3), satisfy the following properties:

$$\begin{aligned} F'(v) &\leq 0 \quad \text{for} \quad v \neq 0, \\ F(0^+) &> 0, \\ F(v_1) = F(v_2) \quad &\Leftrightarrow \quad v_1 = v_2. \end{aligned} \tag{9}$$

Contrasting this form of a force function to the Perona-Malik function in Figure 4, we see that in a sense one can view the discontinuous force function as a limiting form of the continuous force function in Figure 4. In essence this new force function acts as an inverse diffusion operator as long as its argument is not zero. This would appear, at first, to lead to potential problems, since the way in which Perona-Malik-type equations achieve stability is through the positive diffusion effects resulting from the behavior of $F(v)$ for $v \in [-K, K]$. More fundamentally, because of the discontinuity at the origin of the force function in Figure 6, there is a question of how one defines solutions of the equation (4) for such a force function. Indeed, if the equation (4) evolves toward a point of discontinuity of its RHS, the value of the RHS of (4) apparently depends on the direction from which this point is approached (because $F(0^+) \neq F(0^-)$), making further evolution non-unique. We therefore need a special definition of how the trajectory of our evolution proceeds at these discontinuity points.[4] For this definition to be useful, the resulting evolution must satisfy well-posedness properties: the existence and uniqueness of solutions, as well as stability of solutions with respect to the initial data. In the rest of this section we describe how we define solutions to (4) for force functions (9). Assuming the resulting evolutions to be well-posed, we demonstrate that they have the qualitative properties we desire, namely that they both are stable and also act as inverse diffusions and hence enhance edges. We address the issue of well-posedness and other properties in the next section.

Consider the evolution (4) with $F(v)$ as in Figure 6 and Eq. (9) and with all of the masses m_n equal to 1. Notice that the RHS of (4) has a discontinuity at a point $\mathbf{u}$ if and only if $u_i = u_{i+1}$

[4] Having such a definition is crucial because, as we will show in the next section, equation (4) will reach a discontinuity point of its RHS in finite time, starting with any initial condition.

for some i between 1 and $N-1$. It is when a trajectory reaches such a point $\mathbf{u}$ that we need the following definition. In terms of our spring-mass model of Figure 1, once the vertical positions u_i and u_{i+1} of two neighboring particles become equal, the spring connecting them is replaced by a rigid link. In other words, the two particles are simply merged into a single particle which is twice as heavy (see Figure 7), yielding the following modification of (4) for $n=i$ and $n=i+1$:

$$\dot{u}_i = \dot{u}_{i+1} = \frac{1}{2}((F(u_{i+2}-u_{i+1}) - F(u_i - u_{i-1})).$$

(The differential equations for $n \neq i, i+1$ do not change.) Similarly, if m consecutive particles reach equal vertical position, they are merged into one particle of mass m $(1 \leq m \leq N)$:

$$\dot{u}_n = \ldots = \dot{u}_{n+m-1} = \frac{1}{m}(F(u_{n+m} - u_{n+m-1}) - F(u_n - u_{n-1})) \tag{10}$$
$$\text{if} \quad u_{n-1} \neq u_n = u_{n+1} = \ldots = u_{n+m-2} = u_{n+m-1} \neq u_{n+m}.$$

Notice that this system is the same as (4), but with possibly unequal masses. It is convenient to re-write this equation so as to explicitly indicate the reduction in the number of state variables:

$$\dot{u}_{n_i} = \tfrac{1}{m_{n_i}}(F(u_{n_{i+1}} - u_{n_i}) - F(u_{n_i} - u_{n_{i-1}})), \tag{11}$$
$$u_{n_i} = u_{n_i+1} = \ldots = u_{n_i+m_{n_i}-1},$$
$$\text{where } i = 1, \ \ldots, \ p,$$
$$1 = n_1 < n_2 < \ldots < n_{p-1} < n_p \leq N,$$
$$n_{i+1} = n_i + m_{n_i}.$$

The compound particle described by the vertical position u_{n_i} and mass m_{n_i} consists of m_{n_i} unit-mass particles u_{n_i}, u_{n_i+1}, ..., $u_{n_i+m_{n_i}-1}$ that have been merged, as shown in Figure 7. The evolution can then naturally be thought of as a sequence of stages: during each stage, the right-hand side of (11) is continuous. Once the solution hits a discontinuity surface of the right-hand side, the state reduction and re-assignment of m_{n_i}'s, described above, takes place. The solution then proceeds according to the modified equation until it hits the next discontinuity surface, etc.

Notice that such an evolution automatically produces a multiscale segmentation of the original signal if we view each compound particle as a region of the signal. Viewed as a segmentation algorithm, our evolution can be summarized as follows:

1. Start with the trivial initial segmentation: each sample is a distinct region.
2. Evolve (11) until the values in two or more neighboring regions become equal.
3. Merge the neighboring regions whose values are equal.
4. Go to step 2.

In addition, the evolution also naturally defines a sequence of times at which it may be stopped or which may be used in order to characterize features of the image being processed. Specifically, the successive hitting times at which the evolution hits one of the hyperplanes of discontinuity (and at which the evolution changes in the manner we have just described) play such a role. We will have more to say about these in subsequent sections.

The same algorithm can be used for 2-D images, which immediately follows upon re-writing Equation (11) as follows:

$$\dot{u}_{n_i} = \frac{1}{m_{n_i}} \sum_{n_j \in A_{n_i}} F(u_{n_j} - u_{n_i}) p_{ij}, \tag{12}$$

where

m_{n_i} is again the mass of the compound particle n_i (= the number of pixels in the region n_i);

A_{n_i} is the set of the indices of all the neighbors of n_i, i.e., all the compound particles that are connected to n_i by springs;

p_{ij} is the number of springs connecting regions n_i and n_j (this is always equal to one in 1-D, but can be larger in 2-D).

Just as in 1-D, two neighboring regions n_1 and n_2 are merged by replacing them with one region n of mass $m_n = m_{n_1} + m_{n_2}$ and the set of neighbors $A_n = A_{n_1} \cup A_{n_2} \backslash \{n_1, n_2\}$.

We close this section by describing one of the basic and most important properties of these evolutions, namely that the evolution is stable but nevertheless behaves like an inverse diffusion. Notice that a force function $F(v)$ satisfying (9) can be represented as the sum of an inverse diffusion force $F_{id}(v)$ and a positive multiple of sign(v):

$$F(v) = F_{id}(v) + C\,\mathrm{sign}(v),$$

where $C = F(0^+)$ and $F_{id}(v)$ satisfies (3) and (6). Therefore, if $u_{n_{i+1}} - u_{n_i}$ and $u_{n_i} - u_{n_{i-1}}$ are of the same sign (which means that u_{n_i} is not a local extremum of the sequence $(u_{n_1}, ..., u_{n_p})$), then (11) can be written as

$$\dot{u}_{n_i} = \frac{1}{m_{n_i}}(F_{id}(u_{n_{i+1}} - u_{n_i}) - F_{id}(u_{n_i} - u_{n_{i-1}})). \tag{13}$$

If $u_{n_i} > u_{n_{i+1}}$ and $u_{n_i} > u_{n_{i-1}}$ (i.e., u_{n_i} is a local maximum), then (11) is

$$\dot{u}_{n_i} = \frac{1}{m_{n_i}}(F_{id}(u_{n_{i+1}} - u_{n_i}) - F_{id}(u_{n_i} - u_{n_{i-1}}) - 2C). \tag{14}$$

If $u_{n_i} < u_{n_{i+1}}$ and $u_{n_i} < u_{n_{i-1}}$ (i.e., u_{n_i} is a local minimum), then (11) is

$$\dot{u}_{n_i} = \frac{1}{m_{n_i}}(F_{id}(u_{n_{i+1}} - u_{n_i}) - F_{id}(u_{n_i} - u_{n_{i-1}}) + 2C). \tag{15}$$

Equation (13) says that the evolution is a pure inverse diffusion at the points which are not local extrema. It is not, however, a *global* inverse diffusion, since pure inverse diffusions drive local maxima to $+\infty$ and local minima to $-\infty$ and thus are unstable. In contrast, equations (14) and (15) show that at local extrema, our evolution is an inverse diffusion plus a stabilizing term which guarantees that the local maxima do not increase and the local minima do not decrease. For this reason, we call the new evolution (11), (12) a "stabilized inverse diffusion equation" ("SIDE"), a force function satisfying (9) a "SIDE force", and the corresponding energy a "SIDE energy".

4 Properties of SIDEs

The SIDEs described in the previous section enjoy a number of interesting properties which validate and explain their adaptability to segmentation problems. We first examine the SIDEs in one spatial dimension for which we can make the strongest statements.

We define the n_i-th discontinuity hyperplane of a SIDE (11) by $S_{n_i} = \{\mathbf{u} \in \mathbb{R}^p : u_{n_i} = u_{n_{i+1}}\}$, $i = 1, ..., p-1$. Sometimes it is more convenient to work with the vector $\mathbf{v} = (v_{n_1}, \ ..., \ v_{n_{p-1}})^T \in \mathbb{R}^{p-1}$

of the first differences of $\mathbf{u}$: $v_{n_i} = u_{n_{i+1}} - u_{n_i}$, for $i = 1, ..., p-1$. We abuse notation by also denoting $S_{n_i} = \{\mathbf{v} \in \mathbb{R}^{p-1} : v_{n_i} = 0\}$.

On such hyperplanes, we defined the solution of a SIDE as the solution to a modified, lower-dimensional, equation whose RHS is continuous on S_{n_i}. In what follows, we will assume that the SIDE force function $F(v)$ is sufficiently regular away from zero, so that the ODE (11), restricted to the domain of continuity of its RHS, is well-posed. As a result, existence and uniqueness of solutions of SIDEs immediately follow from the existence and uniqueness of solutions of ODEs with continuous RHS. Continuous dependence on the initial data is also guaranteed for a trajectory segment lying inside a region of continuity of the RHS. In order to show, however, that the solutions that we have defined are continuous with respect to initial conditions over *arbitrary* time intervals, we must take into account the presence of discontinuities on the RHS. In particular, what we must show is that trajectories that start very near a discontinuity surface remain close to one that starts on the surface. More precisely, we need to be able to show that a trajectory whose initial point is very close to S_{n_i} will, in fact, hit S_{n_i}. In the literature on differential equations and control theory [7, 24], the behavior that our differential equations exhibit is referred to as "sliding modes". Specifically, as proven in Appendix A, the behavior of our evolution near discontinuity hyperplanes satisfies the following:

Lemma on Sliding. *Let σ be a permutation of $(n_1, \ ..., \ n_{p-1})$, and m an integer between 1 and $p-1$, and let*

$$S = \bigcap_{q=1}^{m} S_{\sigma(q)} \backslash (\bigcup_{q=m+1}^{p-1} S_{\sigma(q)}).$$

Then, as $\mathbf{v}$ approaches S from any quadrant,[5] $\lim(\dot{v}_{\sigma(q)} \, sign(v_{\sigma(q)})) \leq 0$ *for $q = 1, \ ..., \ m$, and for at least one q this inequality is strict.* ■

Intuitively, and as illustrated in Figure 8, this lemma states that the solution field of our

[5] In $\mathbb{R}^{p-1}$, a quadrant containing a vector $\mathbf{a} = (a_1, \ ..., \ a_{p-1})^T$ such that $a_i \neq 0$ for $i = 1, \ ..., \ p-1$ is the set $Q = \{\mathbf{b} \in \mathbb{R}^{p-1} : b_i a_i > 0 \text{ for } i = 1, \ ..., \ p-1\}$.

equation near any discontinuity surface points toward that surface. As a consequence, a trajectory which hits such a surface may be continuously extended to "slide" along the surface, as shown in [7, 24]. For this reason the discontinuity surfaces are commonly referred to as "sliding surfaces" and the associated trajectories as "sliding modes". In our case, a simple calculation verifies that the dynamics along such a surface, obtained through any of the three classical definitions in [7, 24], correspond exactly to the definition given in the preceding section.

The Lemma on Sliding, together with the well-posedness of SIDEs inside their continuity regions, directly implies the overall well-posedness of 1-D SIDEs: for finite T, the trajectory from $t = 0$ to $t = T$ depends continuously on its initial point. As shown in Property 2 to follow, a SIDE reaches a steady state in finite time, which establishes its well-posedness for infinite time intervals.

We call u_{n_i}, with $i \in \{2, ..., p-1\}$ a local maximum (minimum) of the sequence $(u_{n_1}, ..., u_{n_p})$ if $u_{n_i} > u_{n_{i\pm 1}}$ ($u_{n_i} < u_{n_{i\pm 1}}$). The point u_{n_1} is a local maximum (minimum) if $u_{n_1} > u_{n_2}$ ($u_{n_1} < u_{n_2}$); u_{n_p} is a local maximum (minimum) if $u_{n_p} > u_{n_{p-1}}$ ($u_{n_p} < u_{n_{p-1}}$). Similarly, a region of a 2-D image is a local maximum (minimum) if its value is larger (smaller) than the values of its neighbors. Re-phrasing this definition in terms of our spring-mass model, a maximum (minimum) is a particle with all its attached springs directed downward (upward). Therefore, we immediately have (as we saw in Equations (14), (15)) that the maxima (minima) are always pulled up (down):

Property 1 (maximum principle) *Every local maximum is decreased and every local minimum is increased by a SIDE. Therefore,*

$$|u_i(t)| < \max_n |u_n(0)| \quad for \quad t > 0 \tag{16}$$

Using this result, we can prove the following:

Property 2 (finite evolution time) *A SIDE, started at* $\mathbf{u}_0 = (u_{0,1}, ..., u_{0,N})^T$, *reaches its equilibrium (i.e., the point* $\mathbf{u} = (u_1, ..., u_N)^T$ *where* $u_1 = ... = u_N = \frac{1}{N}\sum_{i=1}^{N} u_{0,i}$*) in finite time.*

Proof. The sum of the vertical positions of all unit-mass particles is equal to the sum of the vertical positions of the compound particles, weighted by their masses: $\sum_{n=1}^{N} u_n = \sum_{i=1}^{p} u_{n_i} m_{n_i}$. The time derivative of this quantity is zero, as verified by summing up the right-hand sides of (11). Therefore, the mean vertical position $\frac{1}{N}\sum_{n=1}^{N} u_n$ is constant throughout the evolution. Writing (11) for $i = 1$,

$$\dot{u}_{n_1} = \frac{1}{m_{n_1}} F(u_{n_2} - u_{n_1}),$$

we see that the leftmost compound particle is stationary only if $p = 1$, i.e., if all unit-mass particles have the same vertical position: $u_{n_1} = u_1 = u_2 = \ldots = u_N$. Since the mean is conserved, the unique steady state is $u_1 = \ldots = u_N = \frac{1}{N}\sum_{i=1}^{N} u_{0,i}$. To prove that it is reached in finite time, we use the fact that a SIDE force function assigns larger force to shorter springs. If we put $L = 2 \max_n |u_n(0)|$, then the maximum principle implies that in our system there cannot exist a spring with vertical extent larger than L at any time during the evolution. Therefore, the rate of decrease of the absolute maximum, according to the equation (11), is at least $F(L)/N$ (because $F(L)$ is the smallest force possible in the system, and N is the largest mass). Similarly, the absolute minimum always increases at least as quickly. They will meet no later than at $t = \frac{LN}{2F(L)}$, at which point the sequence $\mathbf{u}(t)$ must be a constant sequence. ∎

The above property allows us immediately to state the well-posedness results as follows:

Property 3 (well-posedness) *For any initial condition $\mathbf{u}_0^*$, a SIDE has a unique solution $\mathbf{u}^*(t)$ satisfying $\mathbf{u}^*(0) = \mathbf{u}_0^*$. Moreover, for any $\varepsilon > 0$ there exists a $\delta > 0$ such that $|\mathbf{u}_0 - \mathbf{u}_0^*| \leq \delta$ implies $|\mathbf{u}(t) - \mathbf{u}^*(t)| \leq \varepsilon$ for $t \geq 0$, where $\mathbf{u}(t)$ is the solution of the SIDE with the initial condition $\mathbf{u}_0$.* ∎

As we pointed out in the previous section, a SIDE evolution defines a natural set of hitting times which intuitively should be of use in characterizing features in an image. For this to be true, however, we would need some type of continuity of this hitting time sequence. Specifically, let $t_n(\mathbf{u}_0)$ denote the "n-th hit time", i.e., the time when the solution starting at $\mathbf{u}_0$ reaches the sliding hyperplane S_n. By Property 2, this is a finite number. Let $\mathbf{u}(t)$ be "a typical solution" if it never reaches two different sliding hyperplanes at the same time: $t_i(\mathbf{u}(0)) \neq t_j(\mathbf{u}(0))$ if $i \neq j$. One of the consequences of the Lemma on Sliding is that a trajectory that hits a single hyperplane S_n does so transversally (that is, cannot be tangent to it). Since trajectories vary continuously, this means that nearby solutions also hit S_n. Therefore, for typical solutions the following holds:

Property 4 (stability of the succession of hit times) *If* $\mathbf{u}(t)$ *is a typical solution, all solutions with initial data sufficiently close to* $\mathbf{u}(0)$ *get onto surfaces* S_n *in the same order as* $\mathbf{u}(t)$. ■

The sequence in which a trajectory hits surfaces S_n is an important characteristic of the solution. Property 4 says that, for a typical solution $\mathbf{u}(t)$, the (strict) ordering of hit times $t_n(\mathbf{u}(0))$ is stable with respect to small disturbances in $\mathbf{u}(0)$:

$$t_{n_1}(\mathbf{u}(0)) < t_{n_2}(\mathbf{u}(0)) < \dots < t_{n_{N-1}}(\mathbf{u}(0)), \tag{17}$$

where $(n_1, \dots, n_{N-1})$ is a permutation of $(1, \dots, N-1)$. For the purposes of segmentation and edge detection, the only interesting output occurs at these $N-1$ time points, since they are the only instances when the segmentation of the initial signal changes (i.e., when regions are merged and edges are erased). While a thorough investigation of how to use these hitting times and in particular how to stop a SIDE so as to obtain the best segmentation is an open one, the fact that our choice is limited to a finite set of time points provides us with both a natural sequence of segmentations of increasing granularity and with, at the very least, some simple stopping rules. For example, if the number of "useful" regions, r, is known or bounded a priori, a natural candidate for a stopping time would be $t_{n_{N-r-1}}$, i.e., the time when exactly r regions remain. In the next section we illustrate the effectiveness of such a rule in the simplest case, namely when $r = 2$ so that we are seeking a partition of the field of interest into two regions. These results together with the properties described here provide ample motivation for a more detailed examination of the properties of the sequence of segmentations produced by a SIDE flow. Such an investigation is currently ongoing.

We already mentioned that our definition of solutions on sliding surfaces for SIDEs in one spatial dimension coincides with all three classical definitions of solutions for a general equation with discontinuous right-hand side, which are presented on pages 50-56 of Filippov's book [7]. We use a result on page 95 of [7] to infer the following:

Property 5 (continuous dependence on the RHS) *Let $F_S(v)$ be a SIDE force function, and let $p_K(v)$ be a smoothing kernel of width K:*

$$p_K(v) \geq 0, \quad supp(p_K) = [-K; K], \quad \int p_k(v)\, dv = 1.$$

Let $F_K(v) = \int F_S(w) p_K(v-w)\, dw$ be a regularized version of $F_S(v)$. Consider system (4) with $m_n = 1$ and $F(v) = F_K(v)$. Then for any ε, there is a K such that the solution of this system stays closer than ε to the solution of the SIDE with the same initial condition and force $F_S(v)$. ■

We note that if the smoothing kernel $p_k(v)$ is appropriately chosen, then the resulting $F_K(v)$ will be a Perona-Malik force function of thickness K. (For example, one easy choice for $p_k(v)$ is a multiple of the indicator function of the interval [-K;K].) Thus, semi-discrete Perona-Malik evolutions with small K are regularizations of SIDEs, and consequently a SIDE in 1-D can be viewed as a limiting case of a Perona-Malik-type evolution. However, as we will see in the next section, SIDE evolutions have behavior that appears to have some advantages over such regularized evolutions even in 1-D.

As we stated at the start of this section, the analysis and properties we have just derived have focused on SIDEs for 1-D signals. Let us close this section by commenting on the properties of SIDEs in 2-D. The existence and uniqueness of solutions again follow easily from our construction of solutions. Property 1 (the maximum principle) is easily inferred from the spring-mass model of Figure 5. Property 2 (finite evolution time) also carries over, with the same proof. There is, however, no analog of the Lemma on Sliding in 2-D: it is easy to show that the solutions in the vicinity of a discontinuity hyperplane of (12) do not necessarily slide onto that hyperplane. Therefore, there is no global continuous dependence on the initial data. In particular, the sequence of hitting times and associated discontinuity planes does not depend continuously on initial conditions, and our SIDE evolution does not correspond to a limiting form of a Perona-Malik evolution in 2-D but in fact represents a decidedly different type of evolutionary behavior. Several factors, however, indicate the value of this new evolution and also suggest that a weaker stability result can be proven. First of all, as shown in the experimental results in the next section, SIDEs can produce excellent segmentations in 2-D images even in the presence of considerable noise. Moreover, thanks to the maximum principle, excessively wild behavior of solutions is impossible, something that is again confirmed by the experiments of the next section. Consequently, the sequence of hit times (17) does not seem to be very sensitive to the initial condition in that the presence of noise, while perhaps perturbing the ordering of hitting times and the sliding planes that are hit, seem to introduce perturbations that are, in some sense, "small". We are currently working on defining an appropriate metric on such hitting plane/time sequences that captures this behavior and that allows us to characterize the qualitatively stable behavior that SIDE evolutions display in the experiments described next.

5 Experiments

Our computer implementation of SIDEs uses the algorithm given in Section 3, with a slight modification imposed by the numerical precision, leading to the replacement of the equality condition $u_i = u_{i+1}$ with the condition $|u_{i+1} - u_i| < \varepsilon$, where ε is a small number. In other words, two neighboring regions are merged if their values differ by less than ε. The examples below are generated with the following SIDE force function:

$$F(v) = 1 - \frac{v}{L} \text{ if } v > 0$$
$$F(v) = -1 - \frac{v}{L} \text{ if } v < 0,$$

where $L/2$ is the maximum of the absolute value of the initial condition.

5.1 Experiment 1: 1-D Unit Step in High Noise Environment

We first test this SIDE on a unit step function corrupted by additive white Gaussian noise whose standard deviation is equal to the amplitude of the step. The noise-free unit step is shown in Figure 9(a), while the noise-corrupted measurement of the step is depicted in Figure 9(b). The remaining parts of this figure display snapshots of the SIDE evolution starting with the noisy data in Figure 9(b), i.e., they correspond to the evolution at a selected set of hitting times. The particular members of the scale space which are illustrated are labeled according to the number of remaining regions. Note that the last remaining edge, i.e., the edge in Figure 9(f) for the hitting time at which there are only two regions left, is located between samples 96 and 97, which is quite close to the position of the original edge (between the 100-th and 101-st samples). In this example, the step in Figure 9(f) also has amplitude that is close to that of the original unit step. In general, thanks to the stability of SIDEs, the sizes of discontinuities will be diminished through such an evolution, much as they are in other evolution equations. However, from the perspective of segmentation this is irrelevant–i.e., the focus of attention is on detecting and locating the edge, not on estimating its amplitude–and that is the aspect on which we wish to focus here.

This example also provides us with the opportunity to contrast the behavior of a SIDE evolution with a Perona-Malik evolution and in fact to describe the behavior that originally motivated our work. Specifically, as we noted in the discussion of Property 5 of the previous section, a SIDE in 1-D can be approximated with a Perona-Malik equation of a small thickness K. Observe that a Perona-Malik equation of a large thickness K will diffuse the edge before removing all the noise. Consequently, if the objective is segmentation, the desire is to use as small a value of K as possible. Following the procedure prescribed by Perona, Shiota, and Malik in [17], we computed the histogram of the absolute values of the gradient throughout the initial signal, and fixed K at 10% of its integral. The resulting evolution is shown in Figure 10. In addition to its good denoising performance, it also blurs the edge, which is clearly undesirable if the objective is a sharp segmentation. The comparison of Figures 9 and 10 strongly suggests that the smaller K the better. It was precisely this observation that originally motivated the development of SIDEs. However, while in 1-D a SIDE evolution can be viewed precisely as a limit of a Perona-Malik evolution as K goes to 0, there is still an advantage to using the form of the evolution that we have described rather than a Perona-Malik evolution with a very small value of K. Specifically, the presence of explicit reductions in dimensionality during the evolution makes a SIDE implementation more efficient than that described in [17]. Even for this simple example the Perona-Malik evolution that produced the result comparable to that in Figure 9 evolved approximately 5 times more slowly than our SIDE evolution. Although a SIDE in 2-D cannot be viewed as a limit of Perona-Malik evolutions, the same comparison in speed of evolution is still true, although in this case the difference in computation time can be orders of magnitude.

5.2 Experiment 2: Deblurring in 1-D

Our second one-dimensional example shows that SIDEs can stably de-blur signals. The staircase signal in the upper left-hand corner of Figure 11 was convolved with a Gaussian and corrupted by additive noise. The evolution was stopped when there were only four regions (three edges) left. The locations of the edges are very close to those in the original signal.

5.3 Experiment 3: SIDE Evolutions in 2-D

This example illustrates SIDEs in 2-D and confirms their similar behavior to that in 1-D. Figure 12 shows that a 2-D SIDE eliminates the less important edges first: the finer textures disappear while the silhouette of the ballerina changes very little. We also see that the boundary between two neighboring regions is always sharp, until it is erased.

5.4 Experiments 4 and 5: SIDE Evolutions in 2-D (continued)

The sharpness of boundaries is also evident in the second and third image experiments. In the second one, shown in Figure 13, we see that if allowed to evolve until exactly two regions are left, the SIDE produces the most important boundary in the image, namely between the phone and everything else. This property is also evident and used to advantage in segmenting a SAR image in which only two textures are present (forest and trees). The initial SAR image and the scale space are shown in Figure 14, and the resulting boundary is superimposed onto the original image in Figure 15. SAR imagery, such as the example shown here, are subject to the phenomenon known as speckle, which is present in any coherent imaging system and which leads to the large amplitude variations and noise evident in the original image. Consequently, the accurate segmentation of such imagery can be quite challenging and in particular cannot be accomplished using standard edge detection algorithms. In contrast, the two-region segmentation displayed in Figure 15 is extremely accurate.

Finally we note, that, as mentioned in Experiment 1, the SIDE evolutions require far less computation time than Perona-Malik-type evolutions. Since in 2-D a SIDE evolution is not a limiting form of a Perona-Malik evolution, the comparison is not quite as simple. However, in experiments that we have performed in which we have devised Perona-Malik evolutions that produce results as qualitatively similar to those in Figure 14 as possible, we have found that the resulting computational effort is roughly 130 times slower for this (201×201) image than our SIDE evolution.

6 Conclusion

In this paper we have presented a new approach to edge enhancement and segmentation and demonstrated its successful application to signals and images with very high levels of noise, as well as to blurry signals. Our approach is based on a new class of evolution equations for the processing of imagery and signals which we have termed stabilized inverse diffusion equations or SIDEs. These evolutions, which have discontinuous right-hand sides, have conceptual and mathematical links to other evolution-based methods in signal and image processing, but they also have their own unique qualitative characteristics and properties that, together with the promising results presented here, suggest the merit of several further lines of investigation.

First, while we have described stability results for SIDEs in 1-D, we have also pointed out that there is an open question in terms of identifying the appropriate notion of stability in 2-D. More to the point, even in 1-D and in light of some of our results (Properties 1, 3, and 4 of Section 4), and our experiments, we have reasons to believe that stronger results can be obtained. For example, Property 4 implies that if the output of a SIDE is a step edge at location n (i.e., between the n-th and $(n+1)$-st samples), then perturbing the input by a small amount will not change that. That is, recalling our definition of hitting times and surfaces, we know that small perturbations in the input data will not change the sequence of hitting surfaces, resulting in locally stable edge location estimates. However, if the noise is large enough to change the order in which the surfaces are hit, we have seen from our experimental results that the consequence is that the surfaces that are hit are perturbed to what we can naturally think of as "nearby" surfaces, i.e., sliding surfaces corresponding to edge estimate locations near the correct edge location. This suggests that there should be a natural topology on the set of sliding surfaces and hitting times that will allow us to define a more global characterization of performance in noise than the local stability result in Property 5 and, moreover, a characterization that applies in 2-D as well as 1-D. Furthermore, the accuracy of the segmentations shown in Section 5 suggests that the results of such a characterization will confirm the noise-insensitivity of SIDEs.

At the same time as we think about noise suppression properties of SIDEs, we must also decide on what exactly we consider to be the output of a SIDE. For the purposes of the examples 1, 2, and 5 in the preceding section, prior knowledge of the number of segments (two for Experiments 1 and 5 and four for Experiment 2) enabled us to simply stop the evolution at the point where the dimensionality of the SIDE had been reduced to the number of desired regions. More generally, however, SIDEs hold promise for more adaptive ways in which to extract information from the evolution. In particular, the value of the evolution $\mathbf{u}(t)$ at some or all of the hitting times–and the sequence of hitting times and planes themselves–form natural candidates for outputs or features. For example, suppose we characterize a segmented region in terms of its size (area in 2-D, length in 1-D) and contrast (amplitude relative to surrounding regions). Can we devise a SIDE that can robustly determine the number, sizes, and shapes of regions with size greater than some minimum value and contrast greater than a second specified minimum? Questions such as this require an investigation not only of the information present in the hitting times and planes and in the evolved images at these points but also of the role played by the detailed form of the force function $F(v)$.

As we also mentioned in Section 4, while SIDEs may do an excellent job of segmentation and location of edges, their estimation of the values within regions or between edges could be improved. We saw this, for example, in Experiments 1 and 2 in which the locations of the discontinuities were determined accurately while the amplitude of each edge had some error and in fact was generally reduced in magnitude. There are two natural ways to address this limitation. The first is simply to use the SIDE for segmentation and then to use optimal linear estimation or filtering within each segmented region in order to get both accurate edge estimates as well as denoising within each so-identified region. Alternatively, at least some of the bias in estimating amplitudes within each region is completely predictable based on the SIDE evolution and could be removed directly. In particular, given knowledge of the hitting time at which a particular edge has been located and knowing the form of the SIDE evolution, we can estimate the amount by which the amplitude has evolved toward the ultimate steady-state value corresponding to the overall average of the signal or image. Using this estimate we can then simply re-scale the amplitude to correct for this effect. Investigations of both of these ideas are currently underway as well.

Finally, once one admits the concept of sliding surfaces for signal or image evolution, the question immediately arises as to whether one can design other sliding surfaces than those used here. In particular, the sliding surfaces used here, corresponding to the enforced equality of neighboring points or pixels, correspond directly to piecewise-constant approximations of signals, and the resulting SIDE evolution in essence produces an adapted sequence of staircase approximations to a signal or image. It is also possible to produce a sequence of linear spline approximations to signals and images by appropriately defining the sliding surfaces. In essence the SIDE evolution decides where the knots in such approximations should be placed at a sequence of levels of granularity, allowing one to identify the most significant knots as the ones that persist longest in the evolution. In fact, one can perform similar approximations for a variety of different choices for the class of approximating basis functions that one chooses. Research expanding on this idea and making connections with topics such as wavelet shrinkage [6, 11] is also ongoing.

A Proof of Lemma on Sliding

To simplify notation, we replace n_i with i in (11):

$$\dot{u}_i = \frac{1}{m_i}(F(u_{i+1} - u_i) - F(u_i - u_{i-1})), \tag{18}$$

with i running from 1 to p, or, in terms of $v_i = u_{i+1} - u_i$,

$$\dot{v}_i = \frac{1}{m_{i+1}}(F(v_{i+1}) - F(v_i)) - \frac{1}{m_i}(F(v_i) - F(v_{i-1})). \tag{19}$$

We need to prove that if $(i_1, ..., i_{p-1})$ is any permutation of $(1, ..., p-1)$, then, as $\mathbf{v}$ approaches $S = \cap_{q=1}^{m} S_{i_q} \backslash (\cup_{q=m+1}^{p-1} S_{i_q})$, $\lim(\dot{v}_{i_q} \text{sign}(v_{i_q})) \leq 0$ for $q = 1, ..., m$, and for at least one q the inequality is strict (i.e., the trajectories enter S transversally).

Fix $v_{i_{m+1}}, ..., v_{i_{p-1}}$ at non-zero values, call $\varepsilon = \frac{1}{2}\min_{m+1 \leq j \leq p-1} |v_{i_j}|$, set initially $|v_{i_1}| = ... = |v_{i_m}| = \delta = \varepsilon$, and drive $\mathbf{v}$ towards S by letting δ go to zero. If $m < p-1$, then there is a j between 1 and m such that at least one of the two neighbors of v_{i_j} is not going to zero: $|v_{i_{j+1}}| > \varepsilon$ or $|v_{i_{j-1}}| > \varepsilon$. Without loss of generality, suppose it is the left neighbor: $|v_{i_{j-1}}| > \varepsilon$. If $m = p-1$, define $j = 1$. Supposing $v_{i_q} = \delta \to 0$, we have $v_{i_q} \leq |v_{i_{q\pm1}}|$, implying $F(v_{i_q}) \geq F(v_{i_{q\pm1}})$, which makes the RHS of (19) for $i = i_q$ non-positive:

$$\lim_{\delta \to 0} \dot{v}_{i_q} \leq 0$$

Moreover, for the particular j described above (supposing $v_{i_j} = \delta$), $F(v_{i_j}) - F(v_{i_{j-1}}) > F(v_{i_j}) - F(\varepsilon)$, and hence (19) for $i = i_j$ has a strictly negative limit:

$$\lim_{\delta \to 0} \dot{v}_{i_j} < 0.$$

Similar reasoning for the cases $v_{i_q} = -\delta$ $v_{i_j} = -\delta$ leads to:

$$\lim_{\delta \to 0} \dot{v}_{i_q} \geq 0 \quad \text{and} \quad \lim_{\delta \to 0} \dot{v}_{i_j} > 0. \qquad \blacksquare$$

References

[1] L. Alvarez, P.L. Lions, and J.M. Morel. Image selective smoothing and edge detection by nonlinear diffusion, II. *SIAM J. Num. Anal.*, 29(3), 1992.

[2] L. Ambrosio and V.M. Tortorelli. Approximation of functionals depending on jumps by elliptic functionals via Γ-convergence. *Communications on Pure and Applied Mathematics*, 43, 1990.

[3] C. Ballester, V. Caselles, and M. Gonzalez. Affine invariant segmentation by variational method. *SIAM J. Appl. Math.*, 56(1), 1996.

[4] V. Caselles, F. Catté, T. Coll, and F. Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66, 1993.

[5] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proc. ICCV*, pages 694–699, Cambridge, 1995.

[6] D.L. Donoho and I.M. Johnstone. Adapting to unknown smoothness via wavelet shrinkage. *JASA*, 90:1200–1223, 1995.

[7] A.F. Filippov. *Differential Equations with Discontinuous Righthand Sides*. Kluwer Academic Publishers, 1988.

[8] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1988.

[9] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *Proc. ICCV*, pages 810–815, Cambridge, 1995.

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

[11] H. Krim and J-C Pesquet. *On the Statistics of Best Bases Criteria*, volume Wavelets in Statistics of *Lecture Notes in Statistics*. Springer-Verlag, 1995.

[12] R. Malladi, J.A. Sethian, and B.C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. on PAMI*, 17(2), 1995.

[13] J.-M. Morel and S. Solimini. *Variational Methods in Image Segmentation*. Birkhauser, 1995.