

Coarse-to-Fine Segmentation and Tracking Using Sobolev Active Contours

Ganesh Sundaramoorthi, Anthony Yezzi, *Senior Member, IEEE*, and Andrea C. Mennucci

Abstract—Recently proposed Sobolev active contours introduced a new paradigm for minimizing energies defined on curves by changing the traditional cost of perturbing a curve and thereby redefining gradients associated to these energies. Sobolev active contours evolve more globally and are less attracted to certain intermediate local minima than traditional active contours, and it is based on a well-structured Riemannian metric, which is important for shape analysis and shape priors. In this paper, we analyze Sobolev active contours using scale-space analysis in order to understand their evolution across different scales. This analysis shows an extremely important and useful behavior of Sobolev contours, namely, that they move successively from coarse to increasingly finer scale motions in a continuous manner. This property illustrates that one justification for using the Sobolev technique is for applications where coarse-scale deformations are preferred over fine-scale deformations. Along with other properties to be discussed, the coarse-to-fine observation reveals that Sobolev active contours are, in particular, ideally suited for tracking algorithms that use active contours. We will also justify our assertion that the Sobolev metric should be used over the traditional metric for active contours in tracking problems by experimentally showing how a variety of active-contour-based tracking methods can be significantly improved merely by evolving the active contour according to the Sobolev method.

Index Terms—Active contours, tracking, coarse-to-fine evolutions, Sobolev gradients.

1 INTRODUCTION

TRACKING objects in video sequences with active contours has been an active research area ever since the introduction of *snakes* in [2] (see [3] for a survey). This is often a two-step procedure. The first step is *detection*. Here, an initial estimate of the object boundary being tracked in a particular image (video frame) is given, and the goal is to evolve this initial contour toward the object of interest in that particular frame. A wide variety of different energy-based schemes have been proposed [2], [4], [5], [6], [7], [8], [9], [10], [11], [12]. The second step is *prediction*, where the objective is to predict the object's boundary in the upcoming image based on the presently detected contour, as well as contours detected in previous images. Measured (or assumed) dynamics are then extrapolated forward to predict the upcoming contour. Many times, the result from prediction is then averaged in an appropriate manner with the result of detection on the prediction to form an *estimate* of the contour (for example, see [13], [14]). A trivial approach, which we call the *naive tracker*, assumes no change in dynamics and therefore uses the contour detected in the current frame as the prediction (initial contour) for the next frame. More sophisticated prediction steps may be found in [15], [13], [16], [17] for parametric snakes and, more recently, [18], [14], [19] for geometric active contours.

The prediction step in many contour tracking algorithms is needed because the detection step is too sensitive to initial contour placement, thereby rendering the *naive tracker* inadequate. Indeed, if we had a robust detection scheme that could operate in real time, then the prediction step could be eliminated, and the *naive tracker* would suffice. This sensitivity of active contour models comes in part due to a lack of inherent smoothness in the way the active contours evolve or deform.

Typically, an object being tracked deforms rather smoothly from frame to frame; otherwise, a prediction would make no sense. Note that we are referring to the smoothness of the contour *deformation* (along the contour), not the contour itself. Active contour energies, through the use of regularizers, may easily be adapted to favor smoothness in the final detected contour. However, in tracking, it makes sense to ensure the smoothness of the deformation of the contour in each frame to the next, regardless of how smooth we want the contour to be. Most current and previous active contour algorithms allow an initial contour to deform in very complex ways, as it flows toward an energy minimum. Even if the final contour in each frame has the exact same shape as the initial contour up to translation, the intermediate contours attained during the evolution may vary immensely from the initial and final shapes. This nonpreferential *freedom* of the contour to undergo arbitrarily complicated deformations as it flows can attract the contour to undesirable intermediate local minima before it reaches the desired object boundary.

It would thus be beneficial, when tracking with active contours, to evolve the initial contour, whether or not it was obtained by the naive tracker or by a prediction step, toward its final configuration in a manner that mimics the evolution behavior of the objects we wish to track. In particular, it would be ideal if the evolution first favored rigid motions that did not change the actual shape of the evolving contour and then gave preferential treatment to

• G. Sundaramoorthi and A. Yezzi are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, 777 Atlantic Drive (TSRB Room 427), Atlanta, GA 30332.

E-mail: {ganeshs, ayezzi}@ece.gatech.edu.
 • A.C. Mennucci is with Scuola Normale Superiore, P.za dei cavalieri, 7, I-56126 Pisa, Italy. E-mail: a.mennucci@sns.it.

Manuscript received 6 Feb. 2007; revised 13 May 2007; accepted 17 May 2007; published online 26 July 2007.

Recommended for acceptance by G. Sapiro.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0078-0207. Digital Object Identifier no. 10.1109/TPAMI.2007.70751.

coarser scale or more global deformations, resorting only at the end to finer scale deformations when necessary.

Recently, *Sobolev Active Contours* [20], [21] introduced a new paradigm for minimizing energies defined on curves (see also [22], [23]). This yields a completely new way to evolve active contours by exploiting the fact that the gradient flow used to evolve a contour is influenced not only by the energy it minimizes but also by how we measure the cost of perturbing the curve. The works [24], [25] revealed many undesirable properties associated with the usual cost (H^0) inherent in all previous geometric active contour models. Accordingly, the authors in [20] and [22] considered using other norms for perturbing active contours based on Sobolev spaces. Sobolev active contours evolve more globally and are less attracted to certain intermediate local minima than traditional active contours. In contrast to the usual strategy of substituting simple energies with more complex (and costly) energies exhibiting fewer local minima, Sobolev active contours minimize the same energy but follow an entirely different deformation to reach their steady-state configuration, thereby avoiding many local minima that would otherwise have been encountered along the way.

Applying Sobolev norms to variational problems has been done in areas other than active contours to gain many of the same advantages that are gained in active contour problems. For example, the book [26] (see also references within) presents the theory of Sobolev gradients and applies it to various physical problems. The key difference between those metrics and the ones we are considering is that we are defining *geometrized* Sobolev norms; these are to be thought of as the metrics for a (yet to be completely studied) Riemannian manifold of curves, equivalent up to reparameterization. The usual Sobolev norms are associated to Hilbert (or Banach) vector spaces of *parameterized* curves. A recent work that uses the Sobolev gradient for boundary reconstruction is presented in [27]. Geometrized Sobolev metrics have been used for defining shape spaces for *shape analysis*, where the objective is to be able to perform statistical operations on *shapes*, which may be contours. For example, the authors in [28], [29], [30] consider Sobolev metrics on the space of plane contours for shape analysis, and other variants are considered in [31].

Although Sobolev active contours was introduced earlier and tracking is mentioned in both [21], [23], our contribution is to expand on those ideas and give a detailed new analysis, which among other things explains the compelling reasons for using the Sobolev metric in tracking situations. Indeed, we examine Sobolev active contours using a scale-space-type analysis that shows, along with other properties to be discussed, that these active contours are quite naturally suited for tracking problems, performing (given the exact same energy functional) significantly better than the corresponding traditional active contour. This makes the generic tracking algorithm less dependent on its prediction step as the initial contour does not need to be placed within as narrow an attraction basin in order to reach the desired minimum. In fact, we will see that Sobolev active contours often allows even the *naïve tracker* to perform well with simple energies that are otherwise plagued by undesirable local minima problems. For a more detailed discussion and analysis of the benefits of Sobolev active contours for tracking, see Section 4.

1.1 Related Works

We now discuss previous techniques that have been explicitly designed to obtain multiscale and global motion properties that are naturally inherent in Sobolev active contours and discuss the advantages of Sobolev active contours in relation to these techniques.

Many active contour works have explicitly incorporated information from successive *scales* of an image to perform a systematic segmentation that matches image data at both coarse and fine scales. For example, in [32], the image is downsampled to a coarse scale, and an active contour [2] is evolved until convergence. The resulting active contour is upsampled to a finer scale of the image, and the process is continued on successively finer scale representations of the image until the active contour is evolved in the image itself. The method in [32] makes it less likely that the active contour becomes stuck in irrelevant local minimum of the underlying energy caused by fine-scale features of the image. There is also a computational advantage of these methods since the algorithm works with downsampled images.

One problem with this method is knowing which and how many scales of the image to use. Ideally, one would like to use a continuum of scales that have gradually more information added at each successive scale. However, this is practically not possible in the framework of the work in [32]. When an arbitrary discrete sampling of scales is chosen, there is a greater possibility of the evolving contour becoming trapped in artificial local minima. Choosing a large number of scales reduces the chance of being trapped in a local minimum; however, the computational cost increases.

One limitation of multiscale techniques like [32] is that there is a limit to how much the domain of the image may be coarsened whenever complicated geometrical objects are present in the image, and it is not trivial to obtain the limit. If the domain of the image is coarsened or downsampled enough, the object to be segmented in the coarsened representation of the image may have a different topology than the same object in the original image (for example, if we coarsen the domain of the image in Fig. 6 enough, then the object will become two circles—a different topology than the original object). At each instant of time of the contour evolution, multiscale methods give a possibly different topology of the object to be segmented (at different scales). This is clearly not a desirable property. In general, this problem is exhibited whenever thin structures are present in the image. We note that Sobolev active contours have no such limitation in that the object to be segmented stays a consistent topology at each instant of time during the contour evolution (even when complicated geometrical objects and thin structures are present).

As shown in [21] (see also Section 2), Sobolev active contours are *global* flows in that they incorporate image and curve information from the entire curve in order to evolve a single point on the curve. Typical multiscale methods that are used for curve evolutions also use global information in order to deform the contour, but the notion of “global” is different than for Sobolev active contours. Indeed, in the coarse-scale evolution of a multiscale algorithm, a point $c(s)$ of the curve uses image information from the neighborhood $\{x \in \mathbb{R}^2 : \|x - c(s)\| < R\}$, where $R > 0$ depends on the amount of smoothing/downsampling performed. Although this gives some advantages over traditional active contours including that the evolution incorporates coarse-scale

information before finer details, the evolution of the curve is not coarse to fine, and there is no preference to coarse-scale deformations before finer deformations because the traditional metric is used. Our scale-space-type analysis shows that Sobolev active contours have inherent multiscale behavior. Indeed, the Sobolev technique is an automatic multiscale evolution that incorporates information from all scales of the image.

Numerous approaches have been applied to tracking, which deal with the problem of traditional active contours being attracted to irrelevant local features of an image resulting from the flexibility of the contour to undergo arbitrary deformations. The approach taken by some to avoid unlikely deformations in tracking applications has been to restrict the degrees of freedom of the active contour so that more global deformations of the contour are only possible. For example, this approach has been taken in [33], [34], [35]. In [33], a deforming contour is represented by B-splines, and [34], [35] uses polygons. Since there are fewer control points for a B-spline and polygon than a typical parametric snake, this results in more global deformations of the contour. An advantage of the approaches used in [34], [35] over traditional active contours is that information is *integrated* over adjacent edges of the polygon in order to move the corresponding vertex and, hence, this adds the robustness of the polygon to noise and other local features. Other related methods use a *finite* number of Fourier or wavelet coefficients to represent the evolving contour (for example, see [32]). A disadvantage of the approaches used in [34], [35] is that topological changes become hard to handle (see [36]). Another disadvantage of these approaches is that the motion of the contour is restricted and, therefore, it becomes impossible to detect fine features of the image when they are needed. Moreover, these methods do not generally evolve in a coarse-to-fine fashion, which is an advantage of Sobolev active contours.

Techniques have been designed to *force* a coarse-to-fine evolution of active contours in order to avoid undesirable local minima of energies in tracking applications. In [37], [38], the authors propose to optimize energies that are defined on both the set of curves and a set of *global* group motions. In the simplest case, these energies are defined as $E_{\text{new}}(c, g) := E(g \circ c)$, where E is an active contour energy, c is a curve, and g is a global group action (for example, affine motion). For tracking, it is beneficial to optimize with respect to g first while keeping c fixed since the global motion is the most important and then to optimize (using the traditional metric) with respect to c to obtain fine-scale changes of the object being tracked.

More recently, the authors of [22], [23] (see a related idea in [39]) have proposed to optimize energies using “spatially coherent” flows to achieve the effect of [37], [38] by constructing norms on the space of perturbations of a curve that favor various group motions (for example, affine motions). These norms are equivalent (unlike Sobolev-type norms) to the traditional H^0 norm, and thus, the notion of “locality” is the same for both H^0 and “spatially coherent” norms (see Section 2.2 for a consequence of this fact). This approach is useful for tracking if one has a prior assumption that the object of interest is moving according to an affine motion. In many tracking situations (for example, see the real sequences in Section 5), it is not necessarily the case that the object of interest is moving according to an affine motion, but we can generally say that the object is moving according to

“coarse” deformations with respect to its boundary that cannot be captured with simple group motions. The advantage of using Sobolev active contours for tracking is the inherent coarse-to-fine behavior, which the inner products based on group motions do not have, and the fact that explicit groups do not need to be chosen.

2 INTRODUCTION TO SOBOLEV ACTIVE CONTOURS

Sobolev active contours were introduced in [20], [21]. We give a brief review of the theory. Let M denote the set of immersed curves in \mathbb{R}^2 , which is a differentiable manifold. For a curve $c \in M$, we denote by $T_c M$ the tangent space of M at c , which is isomorphic to the set of smooth perturbations of the form $h : S^1 \rightarrow \mathbb{R}^2$, where S^1 denotes the circle. We denote by $E : M \rightarrow \mathbb{R}$ an energy functional on M .

Definition 2.1. Let $E : M \rightarrow \mathbb{R}$. If $c \in M$ and $h \in T_c M$, then the *variation of E* is $dE(c) \cdot h = \frac{d}{dt} E(c + th)|_{t=0}$, where $(c + th)(\theta) := c(\theta) + th(\theta)$, and $\theta \in S^1$.

Assume that $\langle \cdot, \cdot \rangle_c$ is an inner product on $T_c M$. The *gradient of E* is a vector field $\nabla E(c) \in T_c M$ that satisfies $dE(c) \cdot h = \langle h, \nabla E(c) \rangle_c$ for all $h \in T_c M$.

One can interpret the gradient as the most efficient perturbation; that is, the gradient maximizes the change in energy per cost of perturbing the curve. The following proposition justifies the previous statement.

Proposition 2.1. Let $\| \cdot \|_c$ be the norm induced from the inner product $\langle \cdot, \cdot \rangle_c$ on $T_c M$. Suppose that $dE(c) \neq 0$, and $\nabla E(c)$ exists; then, the problem

$$\sup_{\{h \in T_c M, \|h\|_c=1\}} dE(c) \cdot h = \sup_{\{k \in T_c M, k \neq 0\}} \frac{dE(c) \cdot k}{\|k\|_c}$$

has a unique solution up to a multiplicative constant, $k = \nabla E(c) \in T_c M$, $h = k/\|k\|_c$.

We review the new inner products on $T_c M$ proposed in [20], [21], which are based on inner products in Sobolev spaces.

Definition 2.2. Let $c \in M$, L be the length of c , and $h, k \in T_c M$. Let $\lambda > 0$. We assume that h and k are parameterized by the arc-length parameter of c :

1. $\langle h, k \rangle_{H^0} := \frac{1}{L} \int_c h(s) \cdot k(s) ds$,
2. $\langle h, k \rangle_{H^n} := \langle h, k \rangle_{H^0} + \lambda L^{2n} \langle h^{(n)}, k^{(n)} \rangle_{H^0}$, and
3. $\langle h, k \rangle_{\tilde{H}^n} := \text{avg}(h) \cdot \text{avg}(k) + \lambda L^{2n} \langle h^{(n)}, k^{(n)} \rangle_{H^0}$,

where $\text{avg}(h) := \frac{1}{L} \int_c h(s) ds$, and $h^{(n)}$ is the n th derivative of h with respect to the arc length.

Note that the length-dependent scale factors give the above inner products and corresponding norms invariance under rescaling of the curve (for example, when the domain of the curve, that is, the image domain, is scaled). Also, it should be noted that the general definition of the Sobolev inner product of order n contains the H^0 inner product of all lower than order n derivatives; however, as shown in [20], [21], all these definitions are topologically equivalent to the definitions we present, and moreover, the qualitative behavior depends on the leading derivative. As noted in [20], [21] (see also Section 5), for implementation purposes, there is an algorithm for the \tilde{H}^n gradient flow that is independent of the parameter λ .

Changing the Riemannian metric associated with the space of curves regularizes the minimizing flows associated with active contour energies without requiring the addition of regularization penalties in the original active contour energies. The change of metric does not affect the global minima of the energy, but it completely changes the notions of gradient and “neighborhood of a curve.” As a result of the change of “locality,” Sobolev active contours are much more robust to the local minima that strongly influence standard active contours (for example, local minima due to noise). See Section 2.2 for more details.

Sobolev gradients are related to H^0 gradients by linear ordinary differential equations (ODEs), and the solution to these ODEs results in the following expressions:

$$\nabla_{H^n} E = \nabla_{H^0} E * K_{\lambda,n}, \text{ and } \nabla_{\tilde{H}^n} E = \nabla_{H^0} E * \tilde{K}_{\lambda,n}, \quad (1)$$

where $*$ denotes convolution on S^1 , $K_{\lambda,n}$, and $\tilde{K}_{\lambda,n}$ are appropriate kernels. It was noted that gradient flows with respect to H^n and \tilde{H}^n norms have the same qualitative properties and that they have similar geometric properties. The advantage of using the \tilde{H}^n gradient is that the convolution formula need not be used; $\nabla_{\tilde{H}^n} E$ at all points of the curve can be solved from $\nabla_{H^0} E$ by computing a couple of integrals around the contour. This means that the computational costs of computing the H^0 and \tilde{H}^n gradients are nearly the same; indeed, computing both gradients have the same order of computational complexity. If m is the number of sample points of the curve and we assume that computing $\nabla_{H^0} E$ is $\mathcal{O}(m)$ (which is typical), then to compute $\nabla_{\tilde{H}^n} E$ is $\mathcal{O}(m + nm) = \mathcal{O}(m)$, where \mathcal{O} denotes Big-Oh notation. On the other hand, to compute $\nabla_{H^n} E$ is $\mathcal{O}(m^2)$, which is much more expensive.

2.1 Motivation: Consistent Theory of Shape

One of the main motivations of our work on Sobolev metrics in the space of curves is to obtain a *consistent* way of doing both shape analysis (for example, computing statistics of shapes) and shape optimization (for example, optimizing active contour energies) in a *consistent* manner. To date, this is the only work that offers such a consistent theory. Indeed, it has recently been proven that the Riemannian metric arising from the Sobolev metrics we consider in the space of curves yield well-defined, that is, nonzero, distances and that the metrics are complete with respect to Lipschitz curves [29]. Also, recently, Michor and Mumford [30] computed the geodesic equation for the H^n metric and proved the existence of geodesics for small times and smooth initial data. In the past, there have been many metrics proposed for doing shape analysis (for example, [40], [41], [38], [31]), but the optimization procedure for energies defined on the space of shapes is inconsistent with these metrics and the geometry of the space and, therefore, artificial.

There are many advantages of the Riemannian metric approach for shape analysis, for example, being able to define a principal component analysis (PCA) of a set of shapes: We may perform a PCA in the tangent space to the mean shape on the vectors that point in the geodesic direction of each individual shape. A consistent approach for optimization on the space of curves is desired for many reasons. For example, consider the simple application of incorporating prior information into segmentation using active contours (for example, see [42], [40], [43], [44], [45]). In the simplest case, we may have prior information that the object to be detected from

an image is close in the sense of our metric to the shape c_0 . Our active contour energy may be the following:

$$E(c) = E_{\text{image}}(c) + d(c, c_0),$$

where c is a shape (for example, curve), d is the Riemannian metric (shape metric) on the space of curves, and E_{image} is an image-based term. To minimize the energy, one can consider calculating the gradient descent flow, which depends on the metric used to define gradient. If we choose the same Riemannian metric for d and the gradient, then the gradient of $d(\cdot, c_0)$ becomes simply the vector pointing in the direction of the geodesic, which is quite natural. On the other hand, if we make an inconsistent choice, then the gradient of $d(\cdot, c_0)$ is the direction that maximizes the change in $d(c, c_0)$ while also minimizing the (inconsistent) cost of perturbing c , which is quite artificial. Moreover, when minimizing E with respect to a different metric than d , we end up far away (in the sense of d) from the initial curve as we step in the gradient direction, which may have detrimental effects for tracking applications where the initial curve is usually quite close to the target curve. Even when computing the average of two shapes, which is a typical computation for shape analysis, a typical procedure is to use a gradient descent to find the average. Although using the same metric to define the gradient leads to a gradient descent that is intimately tied to geodesics and, therefore, to the geometry of the shape space, an inconsistent metric has no such natural connection to the geometry of the space.

2.2 Motivation: Fewer Numerical Local Minimizers

A very important motivation for using Sobolev metrics to define gradient flows is the fact that, as mentioned earlier, the notion of “local” in the space of curves changes. Indeed, since the Sobolev norms are not equivalent to the H^0 norm, the notion of “local” completely changes. Local minimizers for the same energy in H^0 may not be local minimizers when considering Sobolev norms (they may change to critical points that are not local minimizers). However, a local minimizer with respect to Sobolev norms is a local minimizer with respect to H^0 since Sobolev norms dominate the H^0 norm (that is, $\|h\|_{H^0} \leq \|h\|_{\text{Sobolev}}$). Another consequence of the change of locality in Sobolev norms is that numerically many local minimizers of an energy due to noise vanish. Indeed, curves that are local to a local minimizer due to noise in H^0 are “pushed” so far away from the “local” minimizer when considering Sobolev metrics that, numerically, the local minimum no longer exists. These previous points are not true for the “spatially coherent” norms where rigid motions are favored (considered in [23]) since these norms are equivalent to the H^0 norm.

To demonstrate the numerical vanishing of local minima due to noise, we conduct the following experiment:

1. Initialize contour in a noisy image.
2. Run H^0 gradient flow on the energy

$$E(c) = E_{cv}(c) + L(c), \quad (2)$$

where E_{cv} is the Chan-Vese energy, and L is the length functional. Call the converged contour c_0 .

3. Adjust c_0 at **one sample point** by **one pixel**; call modification \hat{c}_0 (\hat{c}_0 is an H^0 local perturbation of c_0).
4. Run and compare H^0 , translation-favored H^0 , and Sobolev gradient flows initialized with \hat{c}_0 .

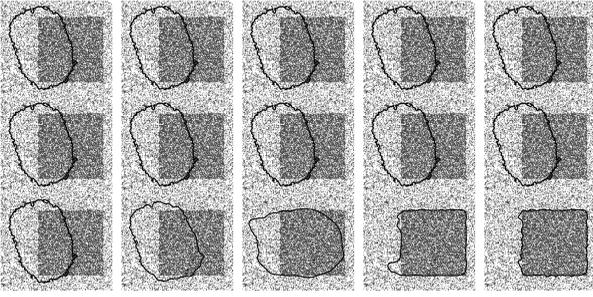


Fig. 1. Top row: H^0 , middle: translation-favored H^0 , and bottom: Sobolev flows for (2). The initialization is a “local” minimum with respect to the H^0 of the energy (2) (this is \hat{c}_0 as noted in the text). Noise: salt and pepper (density = 0.6).

The results are given in Fig. 1. The H^0 active contour flows back to the local minimizer c_0 , since that is local to \hat{c}_0 with respect to H^0 . The same is true for the translation-favored H^0 flow since the corresponding norm is equivalent to the H^0 norm. On the other hand, the Sobolev active contour is not attracted to the “local” minimizer c_0 since \hat{c}_0 is no longer local to c_0 . Note that identical results are achieved for this experiment even if the local minimum c_0 is not perturbed!! This is because, numerically, a local minimum cannot be represented exactly due to finite precision issues, and these finite precision perturbations are local with respect to the H^0 norm but not with respect to Sobolev norms. Therefore, **numerically, this local minimum due to noise no longer exists.**

3 SCALE-SPACE TYPE ANALYSIS OF SOBOLEV ACTIVE CONTOURS

In this section, through Fourier analysis, we show that Sobolev active contours favor more coarse-scale motions than regular active contours and that they generally first undergo coarse-scale motions before resorting to fine-scale deformations in optimizing the chosen energy.

3.1 Sobolev Norms in Frequency Domain

Notice that since any $h \in T_c M$ is smooth on S^1 , it follows that $h \in L^2(S^1)$. Thus, we may write h as a Fourier series:

$$h(s) = \sum_{l \in \mathbb{Z}} \hat{h}(l) \exp\left(\frac{2\pi i}{L} ls\right) \quad (3)$$

with convergence in $L^2(S^1)$ (and in fact, pointwise since h is smooth), where $\hat{h} \in \ell^2(\mathbb{Z})$ is defined by

$$\hat{h}(l) = \frac{1}{L} \int_c h(s) \exp\left(-\frac{2\pi i}{L} ls\right) ds. \quad (4)$$

It should be noted that (3) decomposes the perturbation into the orthonormal basis of exponentials. This allows us to write Definition 2.2 in the frequency domain. By Parseval’s theorem

$$\int_0^L h(s) \cdot k(s) ds = L \sum_{l \in \mathbb{Z}} \hat{h}(l) \cdot \overline{\hat{k}(l)},$$

where $\bar{\cdot}$ denotes complex conjugation. We also have that

$$\int_0^L h^{(n)}(s) \cdot k^{(n)}(s) ds = L \sum_{l \in \mathbb{Z}} \left(\frac{2\pi i}{L} l\right)^{2n} \hat{h}(l) \cdot \overline{\hat{k}(l)},$$

therefore,

Proposition 3.1. *If $h, k \in T_c M$, L is the length of c , and $\hat{h}, \hat{k} : \mathbb{Z} \rightarrow \mathbb{C}$ are defined by (4), then*

$$\langle h, k \rangle_{H^n} = \sum_{l \in \mathbb{Z}} (1 + \lambda(2\pi l)^{2n}) \hat{h}(l) \cdot \overline{\hat{k}(l)}, \quad (5)$$

$$\langle h, k \rangle_{\tilde{H}^n} = \hat{h}(0) \cdot \overline{\hat{k}(0)} + \sum_{l \in \mathbb{Z}} \lambda(2\pi l)^{2n} \hat{h}(l) \cdot \overline{\hat{k}(l)}, \quad (6)$$

and the corresponding norms are

$$\|h\|_{H^n}^2 = \sum_{l \in \mathbb{Z}} (1 + \lambda(2\pi l)^{2n}) |\hat{h}(l)|^2, \quad (7)$$

$$\|h\|_{\tilde{H}^n}^2 = |\hat{h}(0)|^2 + \sum_{l \in \mathbb{Z}} \lambda(2\pi l)^{2n} |\hat{h}(l)|^2. \quad (8)$$

Notice that Proposition 3.1 allows us to define the H^n and \tilde{H}^n inner products for n that is any real number greater than zero. These inner products are defined the same way as in (5) and (6). It is easy to verify in this case too that the definitions are indeed inner products. Unfortunately, for an n that is not an integer, the inner products (therefore, norms) are not local, that is, they cannot be written as integrals of derivatives of the curves, but given $r \in \mathbb{R}^+$, we can represent them for n integer $n > r + 1/4$ as $\langle h, k \rangle_{\tilde{H}^r} = \text{avg}(h) \cdot \text{avg}(k) + \lambda L^{2n} \langle h^{(n)}, K * k^{(n)} \rangle_{H^0}$ for a kernel K with Fourier coefficients:

$$\hat{K}(l) = \begin{cases} 1 & \text{if } l = 0 \\ (2\pi l)^{2r-2n} & \text{if } l \neq 0. \end{cases}$$

The norms shown in (7) and (8) measure the perturbation magnitude in terms of its Fourier coefficients, which are the weights of its corresponding frequency components. We see that for both H^n and \tilde{H}^n norms, high-frequency components of the perturbation contribute increasingly to the norm of the perturbation.

3.2 Sobolev Gradients in the Frequency Domain

We now calculate Sobolev gradients of an arbitrary energy in the frequency domain. By Definition 2.1, if the H^0 and H^n gradients of an energy $E : M \rightarrow \mathbb{R}$ exist, then it follows that

$$dE(c) \cdot h = \langle \nabla_{H^0} E(c), h \rangle_{H^0} = \langle \nabla_{H^n} E(c), h \rangle_{H^n}$$

for all $h \in T_c M$. Using Parseval’s theorem, the last expression becomes

$$\sum_{l \in \mathbb{Z}} (1 + \lambda(2\pi l)^{2n}) \widehat{\nabla_{H^n} E}(l) \cdot \hat{h}(l) = \sum_{l \in \mathbb{Z}} \widehat{\nabla_{H^0} E}(l) \cdot \hat{h}(l).$$

Since the last expression holds for all $h \in T_c M$, we have

$$\widehat{\nabla_{H^n} E}(l) = (1 + \lambda(2\pi l)^{2n})^{-1} \widehat{\nabla_{H^0} E}(l) \quad \text{for } l \in \mathbb{Z}. \quad (9)$$

Using a similar argument, we see that

$$\widehat{\nabla_{\tilde{H}^n} E}(l) = \begin{cases} \widehat{\nabla_{H^0} E}(0) & l = 0 \\ (\lambda(2\pi l)^{2n})^{-1} \widehat{\nabla_{H^0} E}(l) & l \in \mathbb{Z} \setminus \{0\}. \end{cases}$$

It is clear from the previous expressions that high-frequency components of $\nabla_{H^0} E$ are less pronounced in the various forms of Sobolev gradients when compared with the H^0 gradient.

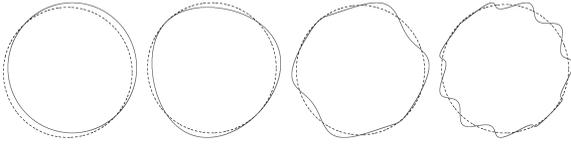


Fig. 2. Increasingly higher frequency perturbations applied to a circle (left to right, $l = 0, 2, 5, 10$).

Higher order Sobolev gradients damp high-frequency components with faster decay rates.

3.3 Coarse-to-Fine Motion of Sobolev Contours

We now discuss the implications of the results in the previous sections. Note that the Fourier basis of the perturbations of a curve decomposes $T_c M$ from global perturbations (low-frequency perturbations) to increasingly more finer perturbations (high-frequency perturbations). Indeed, the zero-frequency perturbation is a simple translation of the curve, which is completely global. See Fig. 2. Therefore, by (9) and comments in the previous section, it is apparent that Sobolev gradients yield perturbations with more pronounced global components than the standard H^0 gradient. Although H^0 gradients give equal weighting across all scales, Sobolev gradients give less weight to finer scales. However, this does not mean that very fine scale deformations of the curve are restricted in Sobolev gradient flows. It just means that if there exists a low-order perturbation (a more global motion) that increases the given energy just as would a higher order perturbation (a more finer deformation), then the low-order perturbation will be preferred in the Sobolev gradient, as shown by Proposition 2.1. Also, if no perturbations in G_m , given by

$$G_m = \left\{ \sum_{|l| \leq m} a_l \exp\left(\frac{2\pi i}{L} l \cdot\right) : a_l \in \mathbb{C}, a_{-l} = \overline{a_l} \right\}$$

can increase the energy E ; that is, $dE(c) \cdot h \leq 0$ for all $h \in G_m$, then by Definition 2.1, we must have that $\widehat{\nabla_{H^0} E}(l) = 0$ for $l \leq m$ and, therefore, we can write

$$\widehat{\nabla_{H^n} E}(l) = \frac{1}{\lambda(m+1)^{2n}} \begin{cases} 0, & |l| \leq m \\ \frac{1}{(2\pi l/(m+1))^{2n}} \widehat{\nabla_{H^0} E}(l), & |l| > m. \end{cases}$$

We see that since the gradient flow does not geometrically depend on a scale factor, the Sobolev gradient automatically has the weights on high-order perturbations of the gradient readjusted (so that perturbations near $|l| = m+1$ become more pronounced). This means that the Sobolev gradient flow at this particular instant of the evolution changes the finer scale structure of the curve. Thus, with Sobolev active contours, this implies, at least locally, a progression from coarse-scale motion to finer scale motion, unlike the standard H^0 active contour. These comments are illustrated in Fig. 3, which shows the tracking of a noisy square image using both H^0 and H^1 active contours. Notice that with the H^0 active contour, the fine structure of the curve is changed immediately, whereas the H^1 active contour gradually changes finer scale features of the curve after changing coarse-scale features.

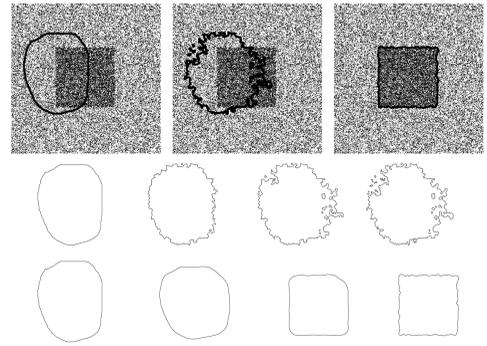


Fig. 3. Standard H^0 active contour (second row) alters the fine structure of the curve immediately; Sobolev (H^1) active contour (bottom) moves from coarse to finer scale motions. Both use the same energy. Top row: initialization, final H^0 and H^1 segmentations.

3.4 Analytic Examples of Coarse-to-Fine Motion

In this section, we give two analytic examples of Sobolev gradient flows and show explicitly that as the artificial time parameter of the evolution increases, the gradient generally moves from coarse-scale to finer scale perturbations. This will verify that the curve initially deforms in a coarse manner before resorting to finer deformations to decrease the energy of interest. Since geometric energies make use of the arc-length parameterization, which lead to nonlinear equations that are difficult to analyze analytically, we perform an analysis on related parametric energies. However, the parametric analysis we perform is relevant in inferring the qualitative behaviors, particularly the coarse-to-fine properties, of the geometric Sobolev gradient flows as we shall justify throughout this section by illustrating the similarities between the parametric and geometric evolutions in the cases where the geometric behavior is known through analysis from other methods.

The inner products that we use are the parametric equivalent of those given in Definition 2.2; that is, if $h, k : S^1 \rightarrow \mathbb{R}^2$ are perturbations of a curve $c : S^1 \rightarrow \mathbb{R}^2$, then we define

$$\begin{aligned} \langle h, k \rangle_{H^0} &:= \int_{S^1} h(u) \cdot k(u) du \\ \langle h, k \rangle_{\tilde{H}^n} &:= \text{avg}(h) \cdot \text{avg}(k) + \lambda \langle h^{(n)}, k^{(n)} \rangle_{H^0}, \end{aligned}$$

where $\text{avg}(h) := \int_{S^1} h(u) du$, and $h^{(n)}$ is the n th derivative with respect to the parameter u . Note that we use \tilde{H}^n instead of H^n for simplicity, but similar conclusions hold for the H^n inner product.

The first energy we consider is the following curve-matching energy, $E_m : C^\infty(S^1, \mathbb{R}^2) \rightarrow \mathbb{R}^+$:

$$E_m(c) = \frac{1}{2} \int_{S^1} |c(u) - c_0(u)|^2 du, \quad (10)$$

where $c_0 : S^1 \rightarrow \mathbb{R}^2$ is a prespecified target curve, which is the known *data*. This energy is a representative of data-based energies, for example, possible energies that depend on image data such as the image-based term of the Chan-Vese energy [11].

Remark 3.1. A geometric version of the energy (10) might be the following:

$$E(c) = \inf_{b \in [0, L_0]} \frac{1}{2} \int_c |c(s) - c_0(L_0/L_c \cdot s + b)|^2 ds,$$

where s is the arc-length parameter of c , L_0 is the length of c_0 , and L_c is the length of c . An optimization of this energy would require a joint evolution of the parameter b and the curve c . The H^0 gradient of this energy with respect to c is

$$\begin{aligned} \nabla_{H^0} E(c)(s) &= (c(s) - c_0(L_0/L_c \cdot s + b)) \cdot \mathcal{N}(s) \\ &\quad - |c(s) - c_0(L_0/L_c \cdot s + b)|^2 c_{ss}(s). \end{aligned}$$

As we will see, the first term is the normal component of the gradient descent of E_m , and we will show later when we consider the energy E_s below that the second term in the Sobolev domain has the coarse-to-fine property we seek to show. Thus, our parametric analysis gives us a good idea of the behavior of the geometric energy also.

Obviously, the optimum curve for (10) is $c = c_0$, which is the global minimizer. However, we are interested in *how* the curve evolves in the frequency domain to attain the minimum. To see this, let us note our definition of Fourier transform:

$$\widehat{f}(k) = \int_{S^1} f(u) \exp(-2\pi i k u) du,$$

where $f : S^1 \rightarrow \mathbb{R}$. Note that

$$dE_m(c) \cdot h = \int_{S^1} (c(u) - c_0(u)) \cdot h(u) du$$

and, so, $\nabla_{H^0} E(c) = c - c_0$. Note that

$$\nabla_{\widehat{H}^n} \widehat{E}_m(l) = \begin{cases} \nabla_{H^0} \widehat{E}_m(0), & l = 0 \\ \frac{\nabla_{H^0} \widehat{E}_m(l)}{\lambda(2\pi l)^{2n}}, & l \in \mathbb{Z} \setminus \{0\}. \end{cases}$$

Let us now consider the \widehat{H}^n gradient flow.

$$\partial_t c(u, t) = -\nabla_{\widehat{H}^n} E_m(c).$$

We may write c as a Fourier series:

$$\begin{aligned} c(u, t) &= \sum_{l \in \mathbb{Z}} \widehat{c}(l, t) \exp(2\pi i l u) \\ \nabla_{\widehat{H}^n} E_m(u, t) &= \nabla_{H^0} \widehat{E}_m(0, t) \\ &\quad + \sum_{l \in \mathbb{Z} \setminus \{0\}} \frac{\nabla_{H^0} \widehat{E}_m(l, t)}{\lambda(2\pi l)^{2n}} \exp(2\pi i l u) \\ &= \widehat{c}(0, t) - \widehat{c}_0(0) \\ &\quad + \sum_{l \in \mathbb{Z} \setminus \{0\}} \frac{\widehat{c}(l, t) - \widehat{c}_0(l)}{\lambda(2\pi l)^{2n}} \exp(2\pi i l u). \end{aligned}$$

Assuming that $\widehat{c}(\cdot, t)$ is uniformly bounded by an $\ell^1(\mathbb{Z})$ function, the flow becomes

$$\begin{aligned} \sum_{l \in \mathbb{Z}} \partial_t \widehat{c}(l, t) \exp(2\pi i l u) &= \\ -(\widehat{c}(0, t) - \widehat{c}_0(0)) - \sum_{l \in \mathbb{Z} \setminus \{0\}} \frac{\widehat{c}(l, t) - \widehat{c}_0(l)}{\lambda(2\pi l)^{2n}} \exp(2\pi i l u). \end{aligned}$$

Since $\{\exp(2\pi i l \cdot)\}_{l \in \mathbb{Z}}$ is an orthogonal basis, we have that

$$\begin{aligned} \partial_t \widehat{c}(0, t) &= -(\widehat{c}(0, t) - \widehat{c}_0(0)) \quad \text{and} \\ \partial_t \widehat{c}(l, t) &= -\frac{\widehat{c}(l, t) - \widehat{c}_0(l)}{\lambda(2\pi l)^{2n}} \quad \text{for } l \in \mathbb{Z} \setminus \{0\} \end{aligned}$$

or

$$\begin{aligned} \partial_t \widehat{c}(0, t) + \widehat{c}(0, t) &= \widehat{c}_0(0) \quad \text{and} \\ \partial_t \widehat{c}(l, t) + \frac{1}{\lambda(2\pi l)^{2n}} \widehat{c}(l, t) &= \frac{\widehat{c}_0(l)}{\lambda(2\pi l)^{2n}} \quad \text{for } l \in \mathbb{Z} \setminus \{0\}. \end{aligned}$$

Solving the previous equations yield

$$\begin{aligned} \widehat{c}(0, t) &= \exp(-t) \widehat{c}(0, 0) + \widehat{c}_0(0) [1 - \exp(-t)] \\ \widehat{c}(l, t) &= \exp\left(-\frac{t}{\lambda(2\pi l)^{2n}}\right) \widehat{c}(l, 0) \\ &\quad + \widehat{c}_0(l) \left[1 - \exp\left(-\frac{t}{\lambda(2\pi l)^{2n}}\right)\right]. \end{aligned}$$

Therefore,

$$\begin{aligned} \nabla_{\widehat{H}^n} \widehat{E}_m(0, t) &= \widehat{c}(0, t) - \widehat{c}_0(0) = \exp(-t) (\widehat{c}(0, 0) - \widehat{c}_0(0)) \\ \nabla_{\widehat{H}^n} \widehat{E}_m(l, t) &= \frac{\widehat{c}(l, t) - \widehat{c}_0(l)}{\lambda(2\pi l)^{2n}} \\ &= \frac{1}{\lambda(2\pi l)^{2n}} \exp\left(-\frac{t}{\lambda(2\pi l)^{2n}}\right) (\widehat{c}(l, 0) - \widehat{c}_0(l)). \end{aligned}$$

For convenience, define

$$g_n(l, t) := \frac{1}{\lambda(2\pi l)^{2n}} \exp\left(-\frac{t}{\lambda(2\pi l)^{2n}}\right).$$

Note that

- $\lim_{t \rightarrow +\infty} g_n(l, t) = 0$ for all $l \in \mathbb{Z}$, and also, the rate of convergence to zero as $|l|$ becomes larger is slower (when $n > 0$). Thus, we find that the coarse-scale frequency components ($|l|$ small) of c converge to the global minimum much faster than the fine-scale frequencies.
- $\lim_{|l| \rightarrow +\infty} g_n(l, t) = 0$.
- We want to show that as time increases, the low-frequency components of $\nabla_{\widehat{H}^n} \widehat{E}_m$ decrease to zero monotonically for a gradually increasing number of low-frequency components. Note that $\|\nabla_{\widehat{H}^n} \widehat{E}_m\|_{\ell^\infty(\mathbb{Z})} \rightarrow 0$ as $t \rightarrow +\infty$. Because a common scale factor for $\nabla_{\widehat{H}^n} \widehat{E}_m$ does not have any effect on the geometry of the curve evolution, we scale $\nabla_{\widehat{H}^n} \widehat{E}_m$ by its maximum value in l to show the convergence to zero of low-frequency components relative to the rest of the components of $\nabla_{\widehat{H}^n} \widehat{E}_m$. Consider

$$\tilde{g}_n(l, t) = \frac{g_n(l, t)}{1/(et)} = etg_n(l, t).$$

Note that

$$\frac{\partial \tilde{g}_n}{\partial t} = \frac{e}{\lambda(2\pi l)^{2n}} \exp\left(-\frac{t}{\lambda(2\pi l)^{2n}}\right) \left(1 - \frac{t}{\lambda(2\pi l)^{2n}}\right)$$

is less than zero when $l < \frac{1}{2\pi} (t/\lambda)^{1/(2n)}$. Note also that $\lim_{t \rightarrow +\infty} \tilde{g}_n(l, t) = 0$ for all $l \in \mathbb{Z}$. These facts verify our assertion.

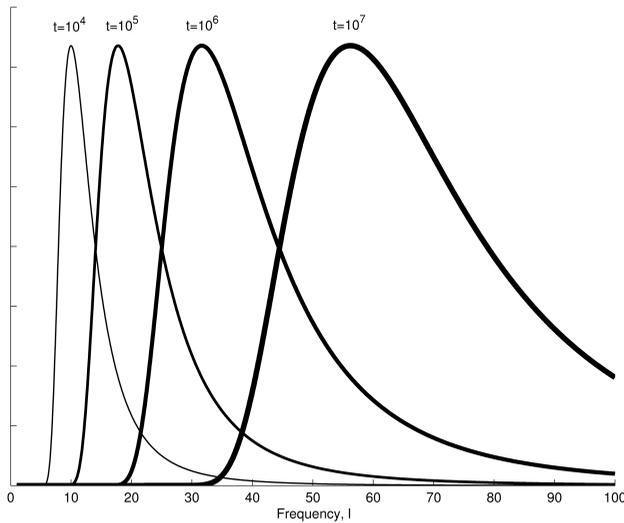


Fig. 4. Plot of $\tilde{g}_n(\cdot, t)$ with $n = 2$ for various t , which shows the coarse-to-fine behavior and that the evolution transitions through *all* possible frequencies.

- For a fixed time, we approximate the frequency component that is changing the most. To do this, we suppose that l is a real number (even though it is an integer) and compute the derivative:

$$\frac{\partial g_n}{\partial l} = \frac{2}{\lambda(2\pi)^{2n} l^{2n+1}} \exp\left(-\frac{t}{\lambda(2\pi l)^{2n}}\right) \left[-1 + \frac{t}{l^{2n}}\right]$$

and, so, $l_s \approx \pm t^{\frac{1}{2n}}$ is the frequency component being changed the most at a particular time instant. In particular, this shows a successive transition through *all* possible frequencies in a coarse-to-fine manner (see the plot in Fig. 4).

In Fig. 4, we show a plot of \tilde{g}_n to visually show the movement of $\nabla_{\tilde{H}^n} E_m$ from coarse to finer scale perturbations. The figure shows the case when $n = 2$; however, other $n > 0$ have the same qualitative behavior. A larger n shows a slower and more pronounced progression from coarse to fine. Note that when $n = 0$, $\tilde{g}_n(l, t) = \frac{\epsilon t}{\lambda} \exp(-\frac{t}{\lambda})$, which is constant in l for a fixed time and, so, the plot in Fig. 4 would simply be a horizontal line for all time. The plot also shows that the Sobolev active contour transitions through all possible frequencies as time increases.

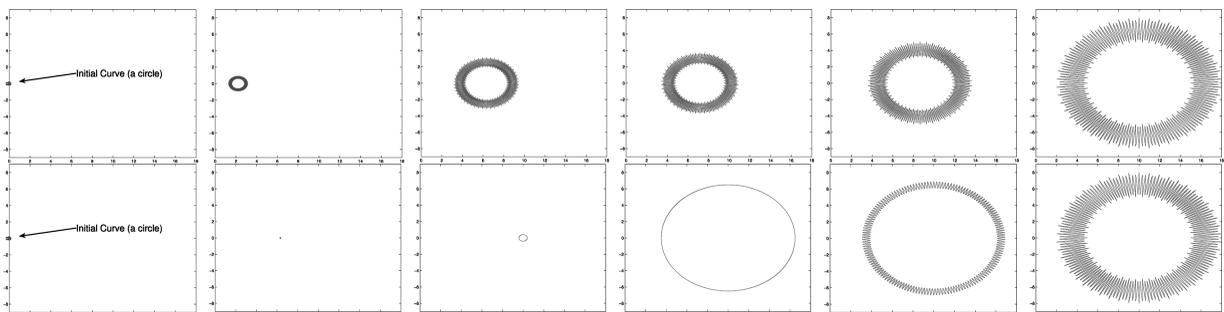


Fig. 5. Minimizing E_m using a “spiked circle” (which is shown on the right of each row) for c_0 , and the initialized curve is a circle: $c(u, 0) = \epsilon(\cos 2\pi u, \sin 2\pi u)$, $\epsilon = 0.001$ for $u \in S^1$ (left of each row, enlarged for visibility). Top: Snapshots of the H^0 evolution. Bottom: Snapshots of the \tilde{H}^2 evolution. The right of each row shows that both flows converge to the desired “spiked” circle but take quite different paths.

The evolutions (H^0 and \tilde{H}^2) minimizing E_m are shown in Fig. 5. Notice that the H^0 -evolution favors all frequency components equally, and so, the coarse and fine features of the target curve are both detected at the same times. On the other hand, the Sobolev evolution changes in a coarse manner, detecting the coarse structure of the target curve, before finally detecting the fine “spiked” structure and deforming in a fine-scale manner.

Remark 3.2. For real images, as long as the image data does not undergo drastic changes with respect to movements of the curve, the coarse-to-fine motion shown in this example also applies to real image-based evolutions (in particular, for many tracking applications). In the general case, we would expect this coarse-to-fine deformation to repeat during successive intervals of time.

The next energy we consider is the following smoothing energy, $E_s : C^\infty(S^1, \mathbb{R}^2) \rightarrow \mathbb{R}^+$:

$$E_s(c) = \frac{1}{2} \int_{S^1} |c^{(m)}(u)|^2 du, \quad (11)$$

where $m \geq 1$.

Remark 3.3. The geometric equivalent to (11) is

$$E(c) = \frac{1}{2} \int_c |D_s^m c(s)|^2 ds.$$

In the case that $m = 1$, the energy is half of the arc length, and in the case that $m = 2$, the energy is the elastic energy. The Sobolev flows for these energies were calculated in [20], [21]. As we progress, we will illustrate the similar qualitative behavior between the parametric and geometric energies, which justifies our use of a parametric analysis to infer some qualitative behaviors in the geometric case.

Note that $\nabla_{H^0} E_s(c) = (-1)^m c^{(2m)}$; therefore,

$$-\nabla_{H^0} E_s(l) = -(-1)^m (2\pi i l)^{2m} \hat{c} = -(2\pi l)^{2m} \hat{c}(l)$$

and also

$$-\nabla_{\tilde{H}^n} E_s(l) = \begin{cases} 0 & \text{for } l = 0 \\ -\lambda^{-1} (2\pi l)^{2(m-n)} \hat{c}(l) & \text{for } l \neq 0. \end{cases}$$

We would like to compute the gradients as a function of time in the Fourier domain. Note that $\partial_t \hat{c}(l, t) = -\nabla_{\tilde{H}^n} E_s(l)$, so

$$\partial_t \widehat{c}(l, t) = \begin{cases} 0 & l = 0 \\ -\lambda^{-1}(2\pi l)^{2(m-n)} \widehat{c}(l) & l \neq 0, \end{cases}$$

which yields a solution of

$$\widehat{c}(l, t) = \begin{cases} \widehat{c}(0, 0) & l = 0 \\ \exp(-\lambda^{-1}(2\pi l)^{2(m-n)} t) \widehat{c}(l, 0) & l \neq 0. \end{cases}$$

Hence, we see

$$\nabla_{\widehat{H}^n} \widehat{E}_s(l, t) = \begin{cases} 0 & l = 0 \\ \lambda^{-1}(2\pi l)^{2(m-n)} \exp(-\lambda^{-1}(2\pi l)^{2(m-n)} t) \widehat{c}(l, 0) & l \neq 0. \end{cases}$$

We make the following observations:

- For $n < m$, the rate of convergence of $\nabla_{\widehat{H}^n} \widehat{E}_s(l, \cdot)$ to zero (also $\widehat{c}(l, \cdot)$ to zero) increases as $|l|$ increases. Using similar arguments as for E_m , one can see that this means that $\nabla_{\widehat{H}^n} \widehat{E}_s$ moves from fine to coarse-scale perturbations. Note that this case includes the case when $m = 1$ and $n = 0$, which is a linear heat flow. This flow is well known to have a smoothing effect on the curve and removes fine-scale curve information before removing coarser scale information (that is, $\nabla_{\widehat{H}^n} \widehat{E}_s$ moves from fine to coarse). Similarly, in the geometric case when $m = 1$ (that is, the energy is the length of the curve) and $n = 0$, the evolution is a nonlinear geometric heat flow [46], [47], which is well known to have a fine-to-coarse smoothing effect.
- For $n = m$, the rate of convergence of all frequency components of $\nabla_{\widehat{H}^n} \widehat{E}_s(l, \cdot)$ to zero is the same for all l . The curve evolution is a simple rescaling of the contour about its parametric centroid. It is also apparent in this case that the curve evolution exists when the gradient ascent flow is considered. Similarly, as verified in [20], [21] through direct computations, in the case of the geometric energy when $m = 1$ and $n = 1$, the evolution is a simple rescaling of the curve about the geometric centroid and, therefore, also stable for the ascent.
- For $n > m$, the rate of convergence of $\nabla_{\widehat{H}^n} \widehat{E}_s(l, \cdot)$ to zero (also $\widehat{c}(l, \cdot)$ to zero) decreases as $|l|$ increases. As in the case of the energy E_m , we can show that $\nabla_{\widehat{H}^n} \widehat{E}_s$ moves from coarse to finer scale motions.
- From the above statements, we see that the Sobolev gradient flows move in a more coarse-to-fine way than the H^0 gradient flow, and as the order of the Sobolev gradient increases, this coarse-to-fine motion is more pronounced.

4 BENEFITS OF SOBOLEV CONTOURS FOR TRACKING

The scale-space analysis of Sobolev active contours performed in Section 3 that shows a coarse-to-fine evolution of the contour also shows why Sobolev active contours are ideal for tracking. The fact that H^0 gradient flows change the fine structure of the curve immediately when energetically favorable and, hence, are easily attracted by undesirable local minima is one reason for predicting the motion and dynamics of the object being tracked. By predicting the motion and dynamics of the moving object, a better estimate of the object's upcoming position can be attained, thereby

placing the initial guess hopefully closer to its desired final position. Many prediction schemes apply low-dimensional global motions to the contour. Thus, the initial global motion followed by an H^0 flow is less likely than the naive tracker to get caught in an intermediate undesirable local minimum of the energy. Notice that since Sobolev gradient flows naturally move from coarse to successively finer motions, the contour is less likely to be trapped by intermediate local minima caused by local features of the image and is therefore likely to be less dependent on the prediction of motion and dynamics of the object. We also wish to emphasize that the transition from coarse to increasingly finer motions is automatic and continuous in comparison to other works (for example, [38]) where the global motions must be deliberately specified, and the transition from the global motion to more local deformation is not continuous. Indeed, even discrete attempts to deliberately graduate from more global to more local motions are not trivial as one typically starts from translations, then rotations, and then scale, but beyond this, it becomes less clear and natural how to progress to finer scale deformations.

Another advantage of using Sobolev active contours for tracking is the speed of convergence compared to standard H^0 active contours. Although computing the \widehat{H}^n gradient is slightly more computationally costly than computing the H^0 gradient, though both have the same order of complexity, we point out that without accurate prediction, the number of iterations in typical contour tracking applications required to update the active contour from frame to frame is usually much smaller with Sobolev active contours. Therefore, the total computational time for processing between frames is significantly lower with Sobolev active contours. The reason is that the frame-to-frame motion of the object to be tracked is, as mentioned previously, usually dominated by more global motions: translations, scaling, and coarse-scale deformations. Accordingly, a Sobolev active contour needs only a few iterations to lock onto the object in the next frame because the Sobolev gradient moves globally at first, preferring coarse-scale motions in the first few iterations before proceeding to fine-scale motions in later iterations. In contrast, standard H^0 active contours requires many more iterations since they are immediately deformed by local motions, significantly changing their initial shape (often to meaningless intermediate shapes), before deforming back to only slightly deformed, translated, and scaled versions of their initial shape and that is assuming they do not first get trapped into intermediate local minima!

We now illustrate the advantages discussed in the previous paragraphs with a simple synthetic image sequence (Fig. 6) in which we employ the naive tracker using the energy functional for geodesic active contours [4], [5]:

$$E_{geo}(c) = \int_c \phi(c(s)) ds, \quad \text{where } \phi = \frac{1}{1 + \|\nabla I\|^2}. \quad (12)$$

Fig. 6 shows the tracking for both the H^0 gradient flow and the \widehat{H}^1 gradient flow. The flows are run until convergence in each frame. Note that the H^0 active contour deforms its initial shape greatly to react to local information. Hence, the contour changes shape and must deform back to its initial shape. However, the contour gets trapped in an undesirable local minimum. The Sobolev active contour, on the other hand, only changes shape slightly while moving in an overall translation. This means that the number of iterations until convergence for the H^0 active contour is

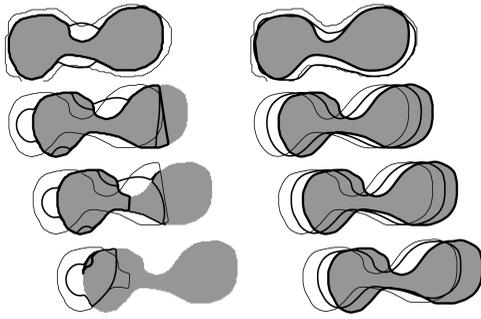


Fig. 6. Simple tracking using geodesic active contours: The standard (H^0) active contour (left column) deforms the initialized contour greatly and is stuck in local minima, and the Sobolev active contour (right column) moves in a global manner only slightly changing shape. In each frame, the initial contour (given by the contour detected in the previous frame) is the thinnest contour, the intermediate contour is the next thinnest, and the final detected contour is the thickest contour.

much greater than the Sobolev active contour, and therefore, the computational time is also much greater. See Fig. 7 for a simple quantitative analysis of computational times. In this simulation, we segment the object shown in Fig. 6 when the initial contour is a translated and slightly deformed version of the object. We quantify the difference by using the set symmetric difference (SSD) between the desired object and the initial contour. From the graph in Fig. 7, we see that the number of iterations and the computational time are significantly lower for the \tilde{H}^1 active contour.

5 EXPERIMENTS

We now demonstrate significant performance gains by replacing standard H^0 active contours with their Sobolev counterparts for the *exact same* detection energy in a variety of tracking scenarios on real videos, both when using the *naive tracker* and when tracking with a predictor. These experiments give evidence to support our claim that the Sobolev metric rather than the traditional metric should be used in tracking applications that make use of active contours.

In the next experiments, we use the \tilde{H}^1 active contour and the algorithm described in [21] that is independent of the parameter λ in the definition of the \tilde{H}^1 inner product. The algorithm evolves by the translation component of the \tilde{H}^1 gradient until this term becomes zero followed by one

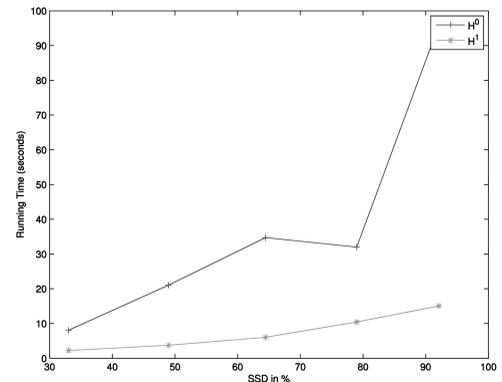


Fig. 7. Graph showing the total time for convergence versus the SSD of initial region and desired object in percent. The total time (and, thus, number of iterations) is lower for the Sobolev active contour.

iteration of the deformation component, which is geometrically independent of λ , and the process is iterated.

Fig. 8 shows the results for a sequence in which a man is walking on a street. The sequence is heavily corrupted by noise (Gaussian noise: $\mu = 0$, $\sigma^2 = 0.3$). The tracking is done using the naive tracker (no prediction) where the detection energy is the Chan-Vese energy [11]

$$E_{cv}(c) = \int_{c_{in}} (I - u)^2 dA + \int_{c_{out}} (I - v)^2 dA + \alpha L(c), \quad (13)$$

where u and v are mean intensities inside and outside c , respectively, and $\alpha \geq 0$ specifies a penalty on the length (used for curve regularity of the H^0 active contour) $L(\cdot)$ of the curve. The top row shows the standard H^0 ($\alpha = 5,000$) active contour, and the bottom row shows the Sobolev \tilde{H}^1 ($\alpha = 0$) active contour. Sobolev active contours favor translations. In order to show that the translation-favoring property of Sobolev active contours is not solely responsible for the pleasing tracking results but, more generally, the coarse-to-fine property, we also show the results of tracking where the energy minimization is performed using an H^0 inner product that has a heavily favored translation component in the middle of Fig. 8. Note that the advantages of the Sobolev technique over such explicitly favored group motions were discussed at the end of Section 1.1). A similar result to the translation-favored H^0 flow is obtained for an affine-favored H^0 flow. We have used an alternating algorithm between a

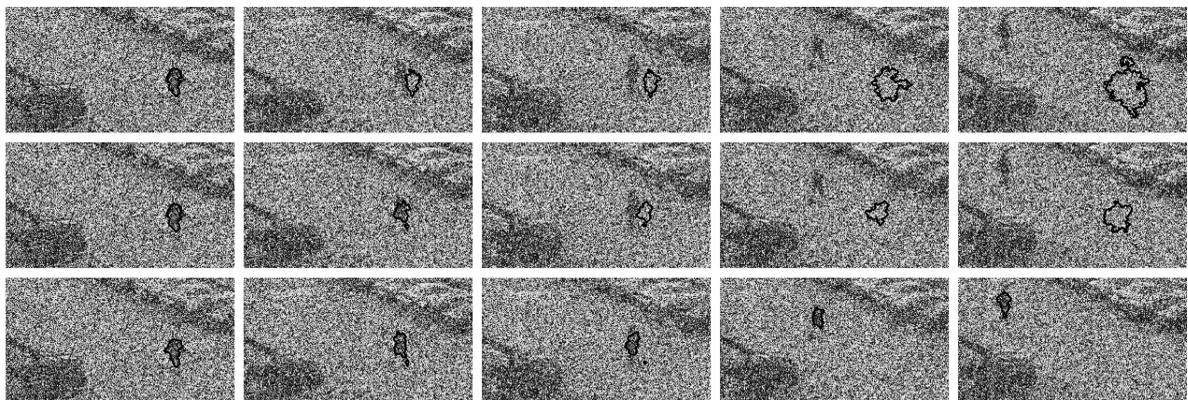


Fig. 8. Tracking of a person in a noisy image sequence with a region-based (Chan-Vese) energy with H^0 (top row), with H^0 translation favored (middle row), and with \tilde{H}^1 (bottom row) active contours.

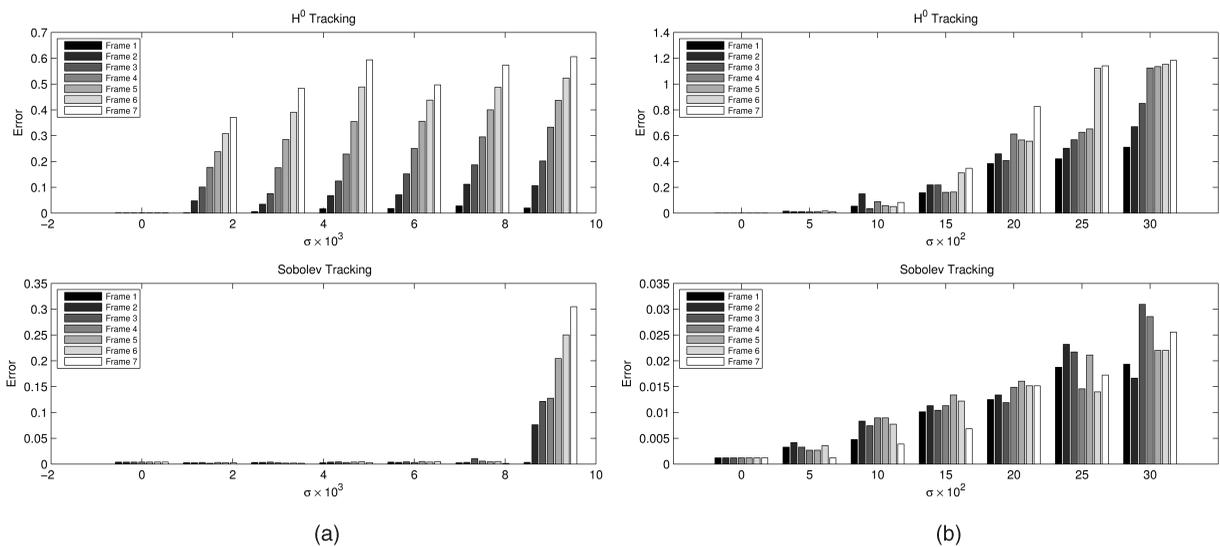


Fig. 9. Plots of error for tracking a square that is translating and slightly changing its area with various degrees of noise (Gaussian mean 0, standard deviation σ). (a) Using geodesic active contours. (b) using the Chan-Vese model. Note the difference in the scales of each plot; in particular, the plots show that the results of using Sobolev active contours is vastly better than the corresponding H^0 active contour.

translation and the H^0 gradient minus the translation to avoid picking the weight on the translation, but the result is similar to that in Fig. 8 (middle row). The contours are evolved until convergence between frames. After a few frames, the H^0 active contour gets stuck in noise and loses track of the person. The translation-favored H^0 contour initially does better than the H^0 active contour but soon loses track of the person, becoming stuck in noise. Note that the translation, in the first frames, initially pushes the contour in the vicinity of the person, but then, the active contour immediately detects the fine-scale noise since H^0 minus the translation does not favor coarser motions and cannot more accurately detect the person. The Sobolev active contour, because of its more global initial motions (translations and other coarse motions), avoids the intermediate local minima caused by noise and keeps tracking the person. Due to the high noise level, however, the precise shape of the person is not captured in any of the cases.

To quantify the robustness to noise of the Sobolev active contour versus the traditional H^0 active contour as seen in the previous experiment, we have conducted experiments with a synthetically generated image sequence (so that the ground truth is known) in which various degrees of noise are added. The sequence is binary images in which a square is translating (to represent motion) and changing its area slightly (to represent deformation). In the first experiment, we use the naive tracker and the detection energy (12). In this case, the square is translated by three pixels, and the length of the side is randomly changed by ± 2 pixels when compared with the square in the previous frame. The segmentation error for various degrees of Gaussian noise ($\mu = 0$, and standard deviation σ specified) using both H^0 and Sobolev \tilde{H}^1 active contours is shown in Fig. 9a. Similar experiments are done using the energy (13), but the square is translated by 17 pixels, and the length is adjusted by a random ± 5 pixels. Results are shown in Fig. 9b. The results for the H^0 active contour are shown for the best value of α chosen for the given noise levels. Note that our measure of error is half of the number of false positive classified pixels plus false negative classified pixels divided by the ground-truth number of pixels of the object. In

both cases of the detection energies chosen, the Sobolev active contour does significantly better than the corresponding H^0 active contour.

In the next experiment (Fig. 10), we demonstrate that the Sobolev active contour is useful not only for noisy situations but also in other cases where one is trying to track an object in a cluttered or textured environment where the object shares some visual characteristics with the background. In this experiment, we track a sea creature at the bottom of an ocean using the naive tracker and the detection energy

$$E_{cv+\sigma}(c) = (1 - \beta)E_{cv}(c) - \frac{1}{2}\beta(\sigma_u^2 - \sigma_v^2)^2,$$

where σ_u^2 and σ_v^2 are the average mean-square errors inside and outside c , respectively. Since the mean values of some regions in the background are closer to the mean value inside the creature rather than to other regions of the background, a first-order Chan-Vese energy is not enough to capture the object and, thus, we incorporate second-order information. For this experiment, we chose $\beta = 0.6$, although different β 's produced similar results. In Fig. 10, we see that the H^0 active contour tracks the object for some time, but when the object's statistics look closer to the light part of the background than the dark part, the contour leaks into the background. On the other hand, because the Sobolev active contour moves globally before gradually changing its fine structure, the contour is able to avoid the distracting features of the background and get a rough approximation of the object before detecting finer features of the object, thereby locking into a more desirable local minimum than the H^0 active contour.

In the next experiment (in Fig. 11), we show that the Sobolev active contour can offer improvements over the traditional metric for tracking an object through partial occlusions. In particular, we track a car that moves under a lamp post. The energy functional used for the active contours is the Mumford-Shah functional [7]

$$E_{ms}(c, f, g) = \int_{c_{in}} (I - f)^2 dA + \int_{c_{out}} (I - g)^2 dA, \quad (14)$$

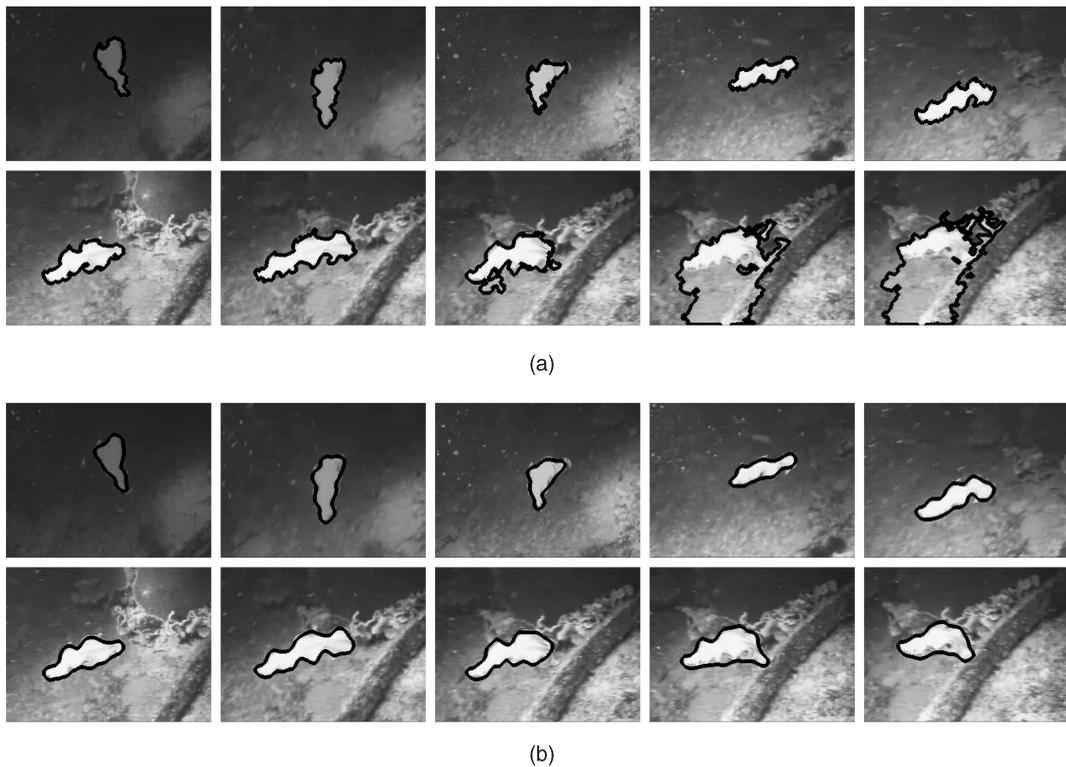


Fig. 10. Tracking of a sea creature at the sea bottom using an energy that incorporates the mean intensity and variance information inside and outside the contour. (a) Tracking using the H^0 active contour. (b) Tracking using the Sobolev (\tilde{H}^1) active contour.

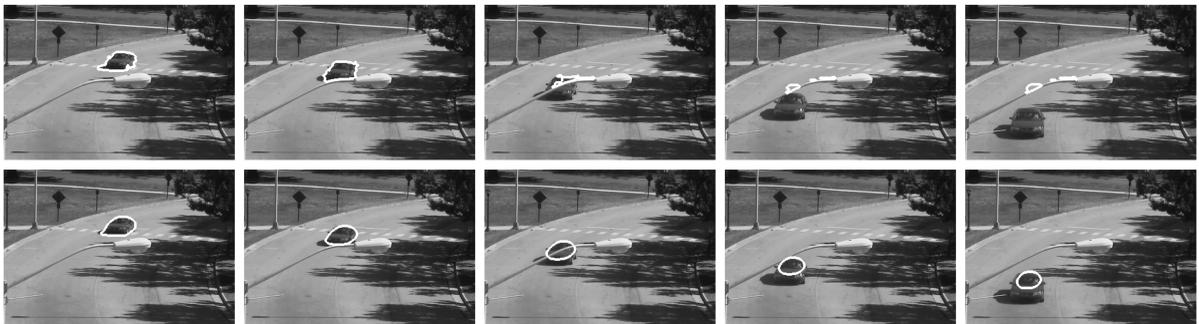


Fig. 11. Tracking of a car under an occlusion using the Mumford-Shah energy with H^0 (top) and \tilde{H}^1 (bottom) active contours.

where f and g are smooth functions defined inside (respectively, outside) the curve. The functional is minimized jointly in c , f , and g (see [48], [33], [49] for implementation details). A fixed number of iterations (300) are used to evolve the curve at each frame. The top row shows the H^0 active contour, which is thrown off as soon as the contour hits the lamp post. This is because each point of the H^0 active contour moves in a direction independent from the other points. Hence, the points close to the lamp post do not want to move past the post. On the other hand, the Sobolev \tilde{H}^1 flow moves globally first and hence does not get stuck on the lamp post and continues to track the car, although at the end, the contour misses the outer parts of the car.

In the last experiment, we illustrate that the Sobolev active contour can improve the traditional active contour even when a prediction step is used to obtain the overall global motion (indeed an affine motion is predicted). The experiment (in Fig. 12) tries to address the problem with the previous experiment by using a predictor and observer/estimator. We

use the detection/prediction algorithm considered in [14]. The detection step involves a simultaneous segmentation and rigid registration (that is, affine) of three consecutive frames using the Mumford-Shah functional (14) and a fixed (300) number of iterations. For the prediction, a constant acceleration model is assumed for the parameters of the rigid registration. The measurements that the estimator uses to estimate the contour and its registrations are the results of the detection step. A Kalman gain is used to determine if more weight is put on the measured contour versus the model prediction. As can be seen in Fig. 12, the H^0 active contour prefers coarse-scale and fine-scale perturbations equally, and therefore, the contour immediately becomes distracted by the pole, which is detected by fine-scale perturbations, and the estimator/predictor is of little help. Note that a higher regularization penalty may be used; however, like the previous experiments, the length penalty, in addition to restricting the deformation into the pole, also shrinks the curve, and a significant portion of the car is not detected. On

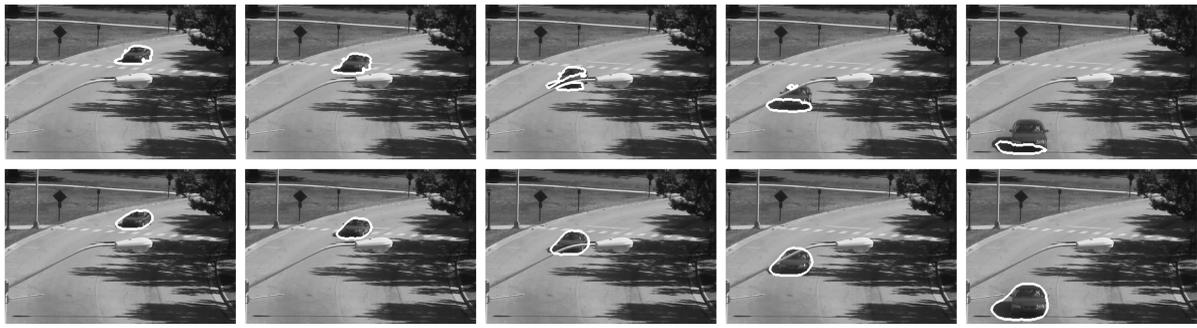


Fig. 12. Tracking a car under an occlusion using estimation with Mumford-Shah energy functional for the detection. H^0 (top) and \tilde{H}^1 (bottom) active contours.

the other hand, for the Sobolev \tilde{H}^1 , the estimator/predictor greatly improves the result, as the shape is more accurately captured. Note that if the detection step is iterated for a very large number of iterations, the Sobolev active contour will also be distracted by the pole in this experiment. However, in real-time tracking applications, it is often the case that the detection is not run until convergence; therefore, it is nice to know that with a limited number of iterations, the Sobolev active contour detects the coarse deformations, which are more essential than the fine deformations. In both cases (with and without the predictor/estimator), it is clear that simply replacing the standard H^0 active contour with the Sobolev active contour greatly improves the tracking performance.

6 CONCLUSION

We have shown that Sobolev active contours move successively from coarse to fine-scale motions in a continuous manner through a scale-space-type analysis. This property gives more justification for using the Sobolev framework. We have shown that this property, along with others, makes Sobolev active contours natural for tracking, and experiments have shown that the Sobolev technique is beneficial over the standard technique both when tracking with or without a predictor. The property of coarse-to-fine motion, as we saw, implies that Sobolev active contours take fewer iterations (and also less time) to converge to the desired local minimum than H^0 active contours. This is important for real-time tracking systems where a more efficient detection scheme with better accuracy is beneficial. Note that existing tracking algorithms, which use active contours, need not be modified, nor does the energy functional for the active contour, just a simple addition of a procedure to compute the Sobolev gradient is necessary, which is straightforward to obtain from the original gradient.

In this paper, we have analyzed Sobolev gradient flows for curves and showed important properties of these flows, which are quite useful for tracking applications. The next step is to extend these ideas to surface evolutions, where there are many applications such as tracking of the heart in ultrasound sequences. Although the Sobolev method extends to surfaces, there are no equivalent convolution formulas (and in particular, no simple integral solution as in \tilde{H}^n) for Sobolev gradients with respect to surfaces, and they are computationally expensive to compute. Future work is

to formulate computationally feasible methods for computing Sobolev gradients for surfaces.

ACKNOWLEDGMENTS

Note that a preliminary version of this manuscript has appeared as a conference proceeding [1]. This work was supported by US NSF grant CCF-0728911.

REFERENCES

- [1] G. Sundaramoorthi, J.D. Jackson, A.J. Yezzi, and A. Mennucci, "Tracking with Sobolev Active Contours," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 674-680, 2006.
- [2] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active Contour Models," *Int'l J. Computer Vision*, vol. 1, pp. 321-331, 1987.
- [3] A. Blake and M. Isard, *Active Contours*. Springer, 1998.
- [4] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic Active Contours," *Proc. Fifth Int'l Conf. Computer Vision*, pp. 694-699, June 1995.
- [5] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi, "Gradient Flows and Geometric Active Contour Models," *Proc. Fifth Int'l Conf. Computer Vision*, pp. 810-815, 1995.
- [6] R. Goldenberg, R. Kimmel, E. Rivlin, and M. Rudzsky, "Fast Geodesic Active Contours," *IEEE Trans. Image Processing*, vol. 10, no. 10, pp. 1467-1475, 2001.
- [7] D. Mumford and J. Shah, "Optimal Approximations by Piecewise Smooth Functions and Associated Variational Problems," *Comm. Pure and Applied Math.*, vol. 42, pp. 577-685, 1989.
- [8] S.C. Zhu, T.S. Lee, and A.L. Yuille, "Region Competition: Unifying Snakes, Region Growing, Energy/Bayes/MDL for Multi-Band Image Segmentation," *Proc. Fifth Int'l Conf. Computer Vision*, pp. 416-423, 1995.
- [9] A. Yezzi, A. Tsai, and A. Willsky, "A Statistical Approach to Snakes for Bimodal and Trimodal Imagery," *Proc. Int'l Conf. Computer Vision*, pp. 898-903, Oct. 1999.
- [10] N. Paragios and R. Deriche, "Geodesic Active Contours and Level Sets for the Detection and Tracking of Moving Objects," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 3, pp. 266-280, Mar. 2000.
- [11] T. Chan and L. Vese, "Active Contours without Edges," *IEEE Trans. Image Processing*, vol. 10, no. 2, pp. 266-277, Feb. 2001.
- [12] N. Paragios and R. Deriche, "Geodesic Active Regions: A New Paradigm to Deal with Frame Partition Problems in Computer Vision," *Int'l J. Visual Comm. and Image Representation*, special issue on partial differential equations in image processing, computer vision, and computer graphics, vol. 13, no. 2, pp. 249-268, June 2002.
- [13] D. Terzopoulos and R. Szeliski, "Tracking with Kalman Snakes," *Active Vision*, MIT Press, 1992.
- [14] J. Jackson, A. Yezzi, and S. Soatto, "Tracking Deformable Moving Objects under Severe Occlusions," *Proc. 43rd IEEE Conf. Decision and Control*, Dec. 2004.
- [15] A. Blake and R. Brockett, "On Snakes and Estimation Theory," *Proc. 33rd IEEE Conf. Decision and Control*, 1994.
- [16] M. Isard and A. Blake, "Condensation—Conditional Density Propagation for Visual Tracking," *Int'l J. Computer Vision*, vol. 1, no. 29, pp. 5-28, 1998.

- [17] N. Peterfreund, "Robust Tracking of Position and Velocity with Kalman Snakes," *IEEE Trans. Pattern Analysis Machine Intelligence*, vol. 21, no. 6, pp. 564-569, June 1999.
- [18] M. Niethammer and A. Tannenbaum, "Dynamic Geodesic Snakes for Visual Tracking," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 660-667, 2004.
- [19] Y. Rathi, N. Vaswani, A. Tannenbaum, and A. Yezzi, "Particle Filtering for Geometric Active Contours and Application to Tracking Deforming Objects," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, 2005.
- [20] G. Sundaramoorthi, A. Yezzi, and A. Mennucci, "Sobolev Active Contours," *Proc. Third Int'l Workshop Variational, Geometric, and Level Set Methods in Computer Vision*, pp. 109-120, 2005.
- [21] G. Sundaramoorthi, A. Yezzi, and A. Mennucci, "Sobolev Active Contours," *Int'l J. Computer Vision*, vol. 73, no. 3, pp. 345-366, July 2007.
- [22] G. Charpiat, R. Keriven, J.-P. Pons, and O.D. Faugeras, "Designing Spatially Coherent Minimizing Flows for Variational Problems Based on Active Contours," *Proc. 10th Int'l Conf. Computer Vision*, pp. 1403-1408, 2005.
- [23] G. Charpiat, R. Keriven, J.-P. Pons, and O.D. Faugeras, "Generalized Gradients: Priors on Minimization Flows," *Int'l J. Computer Vision*, vol. 73, no. 3, pp. 325-344, July 2007.
- [24] P. Michor and D. Mumford, *Riemannian Geometries on the Space of Plane Curves*, ESI Preprint 1425, arXiv:math.DG/0312384, Dec. 2003.
- [25] A. Yezzi and A. Mennucci, *Metrics in the Space of Curves*, Preprint, arXiv:math.DG/0412454, May 2005.
- [26] J.W. Neuberger, "Sobolev Gradients and Differential Equations," *Lecture Notes in Math.*, vol. 1670, 1997.
- [27] M. Burger, "A Framework for the Construction of Level Set Methods for Shape Optimization and Reconstruction," *Interfaces and Free Boundaries*, vol. 5, pp. 301-329, 2003.
- [28] L. Younes, "Computable Elastic Distances between Shapes," *SIAM J. Applied Math.*, vol. 58, no. 2, pp. 565-586, 1998.
- [29] A.C.G. Mennucci, A. Yezzi, and G. Sundaramoorthi, *Properties of Sobolev-Type Metrics in the Space of Curves*, Preprint, arXiv:math.DG/0605017, Apr. 2006.
- [30] P.W. Michor and D. Mumford, *An Overview of the Riemannian Metrics on Spaces of Curves Using the Hamiltonian Approach, Applied and Computational Harmonic Analysis (ACHA)*, arXiv:math.DG/0605009, 2006.
- [31] W. Mio and A. Srivastava, "Elastic-String Models for Representation and Analysis of Planar Shapes," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, vol. 2, pp. 10-15, 2004.
- [32] B. Leroy, I. Herlin, and L. Cohen, "Multi-Resolution Algorithms for Active Contour Models," *Proc. 12th Int'l Conf. Analysis and Optimization of Systems: Images, Wavelets and PDE's*, 1996.
- [33] D. Cremers, F. Tischhäuser, J. Weickert, and C. Schnörr, "Diffusion Snakes: Introducing Statistical Shape Knowledge into the Mumford-Shah Functional," *Int'l J. Computer Vision*, vol. 50, no. 3, pp. 295-313, 2002.
- [34] G.B. Unal, A.J. Yezzi, and H. Krim, "Information-Theoretic Active Polygons for Unsupervised Texture Segmentation," *Int'l J. Computer Vision*, vol. 62, no. 3, pp. 199-220, 2004.
- [35] G.B. Unal, H. Krim, and A.J. Yezzi, "Fast Incorporation of Optical Flow into Active Polygons," *IEEE Trans. Image Processing*, vol. 14, no. 6, pp. 745-759, 2005.
- [36] T. McInerney and D. Terzopoulos, "Topologically Adaptable Snakes," *Proc. Fifth Int'l Conf. Computer Vision*, pp. 840-845, 1995.
- [37] A. Yezzi, L. Zollei, and T. Kapur, "A Variational Framework for Joint Segmentation and Registration," *Proc. IEEE Workshop Math. Methods Biomedical Image Analysis*, 2001.
- [38] S. Soatto and A.J. Yezzi, "DEFORMOTION: Deforming Motion, Shape Average and the Joint Registration and Segmentation of Images," *Proc. Seventh European Conf. Computer Vision*, vol. 3, pp. 32-57, 2002.
- [39] A.-R. Mansouri, D.P. Mukherjee, and S.T. Acton, "Constraining Active Contour Evolution via Lie Groups of Transformation," *IEEE Trans. Image Processing*, vol. 13, no. 6, pp. 853-863, 2004.
- [40] A. Tsai, A.J. Yezzi, C. Tempny, D. Tucker, A. Fan, W.E.L. Grimson, and A.S. Willsky, "Model-Based Curve Evolution Technique for Image Segmentation," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 463-468, 2001.
- [41] G. Charpiat, O.D. Faugeras, and R. Keriven, "Approximations of Shape Metrics and Application to Shape Warping and Empirical Shape Statistics," *Foundations of Computational Math.*, vol. 5, no. 1, pp. 1-58, 2005.
- [42] M. Leventon, E. Grimson, and O. Faugeras, "Statistical Shape Influence in Geodesic Active Contours," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 316-323, 2000.
- [43] M. Rousson and N. Paragios, "Shape Priors for Level Set Representations," *Proc. Seventh European Conf. Computer Vision*, vol. 2, pp. 78-93, 2002.
- [44] Y. Chen, H. Tagare, S. Thiruvankadam, F. Huang, D. Wilson, K. Gopinath, R. Briggs, and E. Geiser, "Using Prior Shapes in Geometric Active Contours in a Variational Framework," *Int'l J. Computer Vision*, vol. 50, no. 3, pp. 315-328, Dec. 2002.
- [45] D. Cremers, S.J. Osher, and S. Soatto, "Kernel Density Estimation and Intrinsic Alignment for Shape Priors in Level Set Segmentation," *Int'l J. Computer Vision*, vol. 69, no. 3, pp. 335-351, Sept. 2006.
- [46] M. Grayson, "The Heat Equation Shrinks Embedded Planes Curves to Round Points," *J. Differential Geometry*, no. 26, pp. 285-314, 1987.
- [47] M. Gage and R.S. Hamilton, "The Heat Equation Shrinking Convex Plane Curves," *J. Differential Geometry*, vol. 23, pp. 69-96, 1986.
- [48] A. Tsai, A.J.Y. Jr., and A.S. Willsky, "Curve Evolution Implementation of the Mumford-Shah Functional for Image Segmentation, Denoising, Interpolation, and Magnification," *IEEE Trans. Image Processing*, vol. 10, no. 8, pp. 1169-1186, 2001.
- [49] L.A. Vese and T.F. Chan, "A Multiphase Level Set Framework for Image Segmentation Using the Mumford and Shah Model," *Int'l J. Computer Vision*, vol. 50, no. 3, pp. 271-293, 2002.



Ganesh Sundaramoorthi received the MS degree in mathematics and the PhD degree in electrical and computer engineering from the Georgia Institute of Technology in 2005 and 2007, respectively. His research spans computer vision and image processing with a particular interest in applications of partial differential equation (PDE) and differential geometry to these areas.



Anthony Yezzi received the PhD degree from the Department of Electrical Engineering, University of Minnesota, in 1997. After completing a postdoctoral research position in the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology, he joined the faculty of the School of Electrical and Computer Engineering, Georgia Institute of Technology, in 1999, where he currently holds the position of associate professor. He has also

consulted for a number of medical imaging companies, including GE, Picker, and VTI. His research lies primarily within the fields of image processing and computer vision. He has worked on a variety of problems including image denoising, edge detection, segmentation and grouping, shape analysis, multiframe stereo reconstruction, tracking, and registration. Some central themes of his research include curve and surface evolution theory, differential geometry, and partial differential equations. He has been senior member of the IEEE since 1999.



Andrea C. Mennucci received the Laurea degree in math from the University of Pisa, Italy, the Diploma in math from Scuola Normale Superiore, Pisa, Italy, and the master of science in electrical engineering from California Institute of Technology (Caltech) in 1993. In 1995, he left the PhD program to join the Scuola Normale Superiore, where he currently holds the position of ricercatore. His research spans Hamilton-Jacobi equations, metric spaces, and computer vision (with a particular interest in shape theory).

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.