

Tracking With Sobolev Active Contours

Ganesh Sundaramoorthi, Jeremy D. Jackson, Anthony Yezzi
School of Electrical and Computer Engineering
Georgia Institute of Technology
{ganeshs, gtg120d, ayezzi}@ece.gatech.edu

Andrea C. Mennucci
Scuola Normale Superiore
a.mennucci@sns.it

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 their gradients. Sobolev active contours evolve more globally and are less attracted to certain intermediate local minima than traditional active contours. In this paper we analyze Sobolev active contours in the Fourier domain in order to understand their evolution across different scales. This analysis shows an important and useful behavior of Sobolev contours, namely, that they move successively from coarse to increasingly finer scale motions in a continuous manner. Along with other properties, the previous observation reveals that Sobolev active contours are ideally suited for tracking problems that use active contours. Our purpose in this work is to show how a variety of active contour based tracking methods can be significantly improved merely by evolving the active contours according to the Sobolev method.

1. Introduction

Tracking objects in video sequences with active contours has been an active research area ever since the introduction of *snakes* in [8] (see [2] 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 to do this, including both edge-based [3, 9] and region-based [12, 4, 14] active contours. The second step 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 estimate the upcoming contour. A trivial approach, which we call the *naive tracker*, assumes no change and therefore uses the contour detected in the current frame as the prediction (ini-

tial contour) for the next frame. More sophisticated prediction steps may be found in [1, 18, 6] for parametric snakes and more recently [13, 7, 15] for geometric active contours.

The prediction step in many contour tracking algorithm 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 inherent smoothness in the way they 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 smoothness of the contour *deformation*, 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 smoothness of the deformation of the contour from one 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 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 non-preferential *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 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 coarser or more global deformations, resorting only at the end to finer or more local deformations when necessary.

Recently, *Sobolev Active Contours* [17] introduced a new paradigm for minimizing energies defined on curves (see also [5]). This yields a completely new way to evolve active contours by exploiting the fact that the gradient flow used to evolve a contour is not only influenced by the energy it minimizes but also by how we measure the cost of perturbing the curve. The works [11, 19] revealed many undesirable properties associated with the usual cost (H^0) inherent in all previous geometric active contour models. Accordingly, [17] and [5] 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.

In this paper, we examine Sobolev active contours using a scale-space type analysis which 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 Sobolev active contours often allows even the *naive tracker* to perform well with simple energies that are otherwise plagued by undesirable local minima problems.

Finally, the scale space analysis we carry out shows an additional pragmatic benefit, again ideal for tracking, which may be exploited not only for performance gains but also for speed considerations. Namely, as we show that Sobolev active contours first undergo coarse scale deformations before yielding to finer scale changes, it therefore becomes more justified when using Sobolev active contours to truncate the number of evolution steps when updating the contour between video frames than with traditional active contours. Such truncation is often required in practical algorithms for the sake of speed. It is nice to know that given a limited number of iterations, that the coarser scale changes of the object will be captured first as such changes are typically more important and useful when tracking.

2. Review of Sobolev Active Contours

Sobolev active contours were introduced in [17] (see related work in [5]). 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 [10]. For a curve $c \in M$, we denote by

T_cM 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 also denote by $E : M \rightarrow \mathbb{R}$ an energy functional on M , which is known.

Definition 1 Let $E : M \rightarrow \mathbb{R}$.

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

We formalize a comment briefly alluded to in [17] with a proposition, which interprets the gradient as the most efficient perturbation; that is, the gradient maximizes the change in energy per cost of perturbing the curve.

Proposition 1 Let $\| \cdot \|_c$ be the norm induced from the inner product $\langle \cdot, \cdot \rangle_c$ on T_cM . Suppose $dE(c) \neq 0$; then the problem

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

has a unique solution, $k = \nabla E(c) \in T_cM, h = k/\|k\|$.

In [17, 5], it was noted that all previous geometric active contour models that have been formulated as gradient flows of various energies use the same L^2 -type inner product (aka H^0) to define the notion of gradient. We review the new inner products on T_cM proposed in [17], based on inner products in Sobolev spaces.

Definition 2 Let $c \in M$, L be the length of c , and $h, k \in T_cM$. Let $\lambda > 0$. We assume h and k are parameterized by the arclength parameter of c .

1. $\langle h, k \rangle_{H^0} := \frac{1}{L} \int_0^L h(s) \cdot k(s) ds$
2. $\langle h, k \rangle_{H^n} := \langle h, k \rangle_{H^0} + \lambda L^{2n} \langle h^{(n)}, k^{(n)} \rangle_{H^0}$
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 $\text{avg}(h) := \frac{1}{L} \int_0^L h(s) ds$, and the derivatives are with respect to arclength.

Note that the length dependent scale factors give the above inner products and corresponding norms invariance under rescaling of the curve.

It was shown in [17] that Sobolev gradients can be expressed as a convolution on S^1 of the traditional H^0 gradient and appropriate kernels. It was noted that gradient flows from H^n and \tilde{H}^n 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 single integral 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 computational complexity.

Whereas the focus in active contour literature for the past two decades has been on changing the energy in dealing with the problem of local minima, the idea of changing the inner product to define the notion of gradient in [17, 5] allows one to use the same energy with less susceptibility to local minima. The usefulness of Sobolev active contours was demonstrated in [17]. It was shown that Sobolev flows are smooth in the space of curves, are not as dependent on local image information as H^0 flows, are more global flows than H^0 , and reduce the order of the evolution PDE in comparison to H^0 .

3. Fourier Analysis of Sobolev Active Contours

In this section we study Sobolev inner-products and the corresponding gradient flows in the Fourier domain, allowing us to analyze the evolution of a Sobolev active contour across different scales. In particular, we will see that coarse scale evolution components are weighted more heavily.

3.1. Sobolev Norms in Frequency Domain

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

$$h(s) = \sum_{l \in \mathbb{Z}} \hat{h}(l) \exp\left(\frac{2\pi i}{L} l s\right) \quad (1)$$

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

$$\hat{h}(l) = \frac{1}{L} \int_0^L h(s) \exp\left(-\frac{2\pi i}{L} l s\right) ds. \quad (2)$$

It should be noted that (1) decomposes the perturbation into the orthonormal basis of exponentials. This allows us to write Definition 2 in the frequency domain. By Parseval's theorem,

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

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

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

therefore,

Proposition 2 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 (2). Then,

$$\langle h, k \rangle_{H^n} = \sum_{l \in \mathbb{Z}} (1 + \lambda(2\pi l)^{2n}) \hat{h}(l) \cdot \overline{\hat{k}(l)} \quad (3)$$

$$\langle h, k \rangle_{\tilde{H}^n} = \hat{h}(0) \cdot \overline{\hat{k}(0)} + \sum_{l \in \mathbb{Z}} \lambda(2\pi l)^{2n} \hat{h}(l) \cdot \overline{\hat{k}(l)}. \quad (4)$$

and the corresponding norms are

$$\|h\|_{H^n}^2 = \sum_{l \in \mathbb{Z}} (1 + \lambda(2\pi l)^{2n}) |\hat{h}(l)|^2 \quad (5)$$

$$\|h\|_{\tilde{H}^n}^2 = |\hat{h}(0)|^2 + \sum_{l \in \mathbb{Z}} \lambda(2\pi l)^{2n} |\hat{h}(l)|^2. \quad (6)$$

Notice that Proposition 2 allows us to define the H^n and \tilde{H}^n inner products for n that is any real number greater than 0. These inner products are defined the same way as in (3) and (4). It is easy to verify in this case too, the definitions are indeed inner products.

The norms shown in (5) and (6) 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. Indeed, the norm of a frequency component increases with frequency, and the growth rate of the weights of the frequency coefficients for both H^n and \tilde{H}^n are the same.

3.2. Sobolev Gradients in Frequency Domain

We now calculate Sobolev gradients of an arbitrary energy in the frequency domain. By Definition 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

$$\widehat{\nabla_{H^n} E}(l) = (1 + \lambda(2\pi l)^{2n})^{-1} \widehat{\nabla_{H^0} E}(l) \quad \text{for } l \in \mathbb{Z} \quad (7)$$

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 the H^n gradients, with higher order Sobolev gradients damping high frequency components with faster decay rates. We should also remark that these expressions give a decomposition of the gradients onto an orthogonal basis of simple motions starting from translations to higher order trigonometric motions. We see that with the Sobolev gradients, these high order motions do not contribute to the motion of the curve as much as the H^0 gradient. In fact, these high order motions decay at a much quicker rate than in H^0 . It should be noted that in $\widehat{\nabla_{H^0} E}(l)$ decays to zero as $|l| \rightarrow +\infty$, but can there can be an arbitrary large frequency component of $\widehat{\nabla_{H^0} E}$. The Sobolev gradients' frequency components decay rate at a much faster rate than H^0 : the larger the frequency, the more it will be killed.

3.3. Coarse-to-Fine Motion of Sobolev Contours

We now discuss the implications of the analysis of Sobolev active contours in the Fourier domain. We note that the Fourier basis of the perturbations of a curve decomposes $T_c M$ from global perturbations (low frequency perturbations) to increasingly more local perturbations (high frequency perturbations). Indeed the zero frequency perturbation is a simple translation of the curve, which is completely global. See Figure 1. Therefore, by (7), 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. While H^0 gradients give equal weighting across all scales, Sobolev gradients give less weight to finer scales. However, this does not mean that very local (fine scale) deformations of the curve are restricted from 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 local motion), then the low order perturbation will be preferred in the Sobolev gradient, as shown by Proposition 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 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

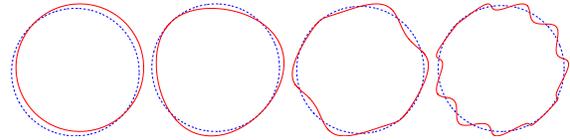


Figure 1. Increasingly higher frequency perturbations applied to a circle (left to right, $l = 0, 2, 5, 10$).

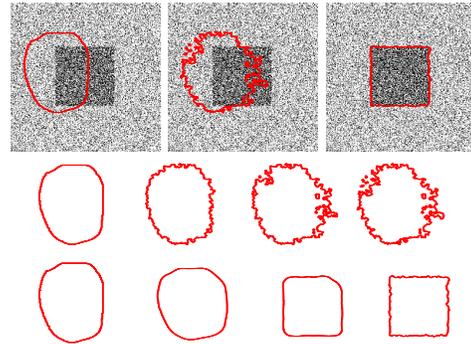


Figure 2. Standard H^0 active contour (2nd row) alters fine structure of the curve immediately; Sobolev (H^1) active contour (bottom) moves from coarse to finer scale motions. Both use same energy. Top row: initialization, final H^0 and H^1 segmentations.

readjusted (so that perturbations near $|l| = m + 1$ become more pronounced). This means the Sobolev gradient flow at this particular instant of the evolution changes the fine scale structure of the curve. Thus, with Sobolev active contours, we see a progression from coarse scale motion to finer scale motion, much more so than the standard H^0 active contour. Figure 2 shows the tracking of a noisy square image using both H^0 and H^1 active contours, which illustrates the ideas of the previous comments. Notice that with the H^0 active contour, the fine structure of the curve is changed immediately, while the H^1 active contour gradually changes finer scale features of the curve after changing coarse-scale features.

We comment that the effect of using higher order (n large) Sobolev gradients is higher favorability to lower order perturbations in the flow.

4. Benefits of Sobolev Contours for Tracking

In this section, we outline the benefits of switching from the standard H^0 active contour evolution to a Sobolev active contour in tracking algorithms that use active contours.

We note that typically an object that is being tracked, during a small period of time, is moving globally according to a translation and locally according to a small deformation. This is assumed in many tracking algorithms that use active contours (for example [16]). Sobolev active contours are ideally suited for this typical tracking assumption. For λ large, by expression (7), we see that most of the motion of the Sobolev active contour is given by a translation, but

there is still a small deformation of the curve. This may lead to the question of how large to choose λ . For the particular case of \tilde{H}^n , as noted in [17], we can implement the curve evolution without a choice of λ and have the same behavior as λ large. This is done by iteratively evolving by the translation component of the gradient until this term becomes zero followed by the deformation component of the \tilde{H}^n gradient, from which we may clearly omit the factor λ .

The Fourier 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 fine structure of the curve immediately when energetically favorable, and hence are easily attracted by undesirable local minima, is one reason for predicting motion and dynamics of the object being tracked. By predicting 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, 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 (*e.g.*, [16]) 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, 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 speed of convergence compared to standard H^0 active contours. While computing the \tilde{H}^n gradient is slightly more computationally costly than computing the H^0 gradient, though both have the same 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

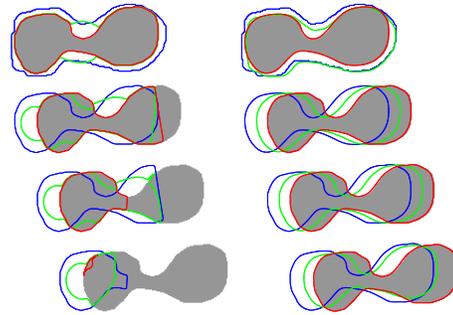


Figure 3. Simple tracking using geodesic active contours: Standard (H^0) active contour (left column) deforms the initialized contour greatly and is stuck in local minima, and Sobolev active contour (right column) moves in a global manner only slightly changing shape. In each frame, the initial curve (given by the contour detected in the previous frame) is blue, the intermediate curve is green, and the final detected curve is red.

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 immediately deform 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 don't first get trapped into intermediate local minima!

We now illustrate the advantages discussed in the previous paragraphs with a simple synthetic image sequence (Figure 3) in which we employ the naive tracker using the energy functional for geodesic active contours [3, 9]. Figure 3 shows the tracking for both the H^0 gradient flow and the \tilde{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 re-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 much greater than the Sobolev active contour, and therefore the computational time is also much greater. See Figure 4 for a simple quantitative analysis of the number of iterations and computational times. In this simulation, we segment the object shown in Figure 3 when the initial contour is a translated and a slightly deformed version of the object. We quantify the difference by using the set symmetric difference between the desired object and the initial contour. From the graph in Figure 4, we see that the number of iterations and the computational time is significantly lower for

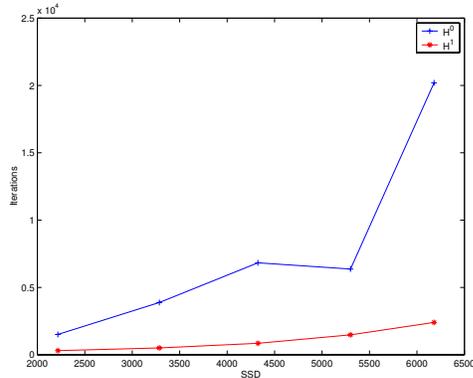


Figure 4. Graph showing number of iterations to converge versus set symmetric difference (SSD) of initial region and desired object.

the \tilde{H}^1 active contour.

5. Experiments

We now demonstrated improved performance by replacing standard H^0 active contours with their Sobolev counterparts in a variety of tracking scenarios on real images, both in *naive tracking* as well as in tracking with prediction.

Figure 5 shows the results for a sequence that contains a man walking. The sequence is heavily corrupted by noise. The tracking is done using the naive tracker (no prediction) with the Chan-Vese energy functional [4]. The left column shows the standard H^0 active contour, and the right shows the Sobolev \tilde{H}^1 active contour. 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 Sobolev active contour, because of its more global initial motions, skips over the intermediate noise and keeps tracking the person. Due to the high noise level, however, the precise shape of the person is not captured in either of the cases.

Figure 6 shows the tracking of a car that passes through an occlusion using the naive tracker (no prediction). The energy functional used for the active contours is the Mumford-Shah functional [12]. A fixed number of iterations 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 \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.

Figure 7 tries to address the problem with the previous experiment with a predictor and estimator. In this experiment we compare the behavior of using an estimator with H^0 active contours and using it with Sobolev active con-

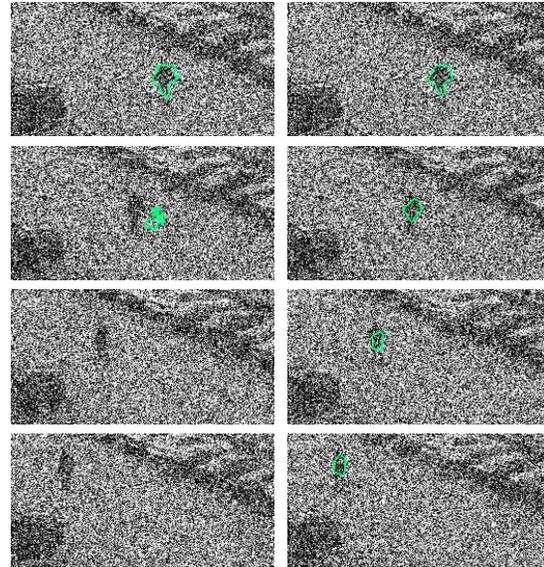


Figure 5. Tracking of a person in a noisy image sequence using a region-based energy with H^0 (left) and \tilde{H}^1 (right) active contours.

tours. Using an estimator in the H^0 case was done in [7]. The measurements that the estimator uses to estimate the contour and its registrations are just the output of a simultaneous flow that finds a segmentation and registration between frames. A gain is used to determine if more weight is put on the measured contour versus the model. The problem with using the H^0 flow for the measured contour is that if one uses more iterations to get the results of the registration/segmentation (i.e the measurements), then as it passes by an occlusion it has more opportunity to get distracted by it. On the other hand, the for, \tilde{H}^1 , the estimator greatly improves the result, as the shape is more accurately captured. 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 through Fourier analysis. Therefore, Sobolev active contours are robust to local minima when compared to H^0 active contours derived from the *same* energy functional. We have shown that this property, along with others, makes Sobolev active contours natural for tracking. The property of this 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 ac-

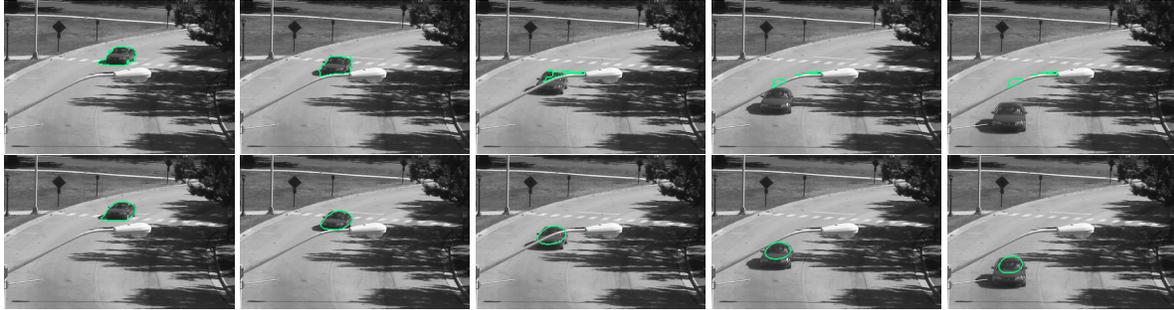


Figure 6. Tracking of a car under an occlusion using the Mumford-Shah energy with H^0 (top) and \tilde{H}^1 active contours.

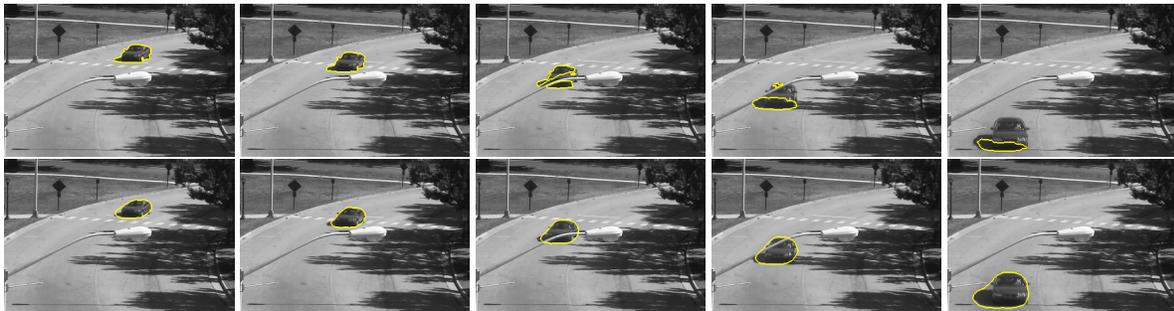


Figure 7. Tracking a car under an occlusion using estimation with Mumford-Shah energy functional for the detection. H^0 (top) and \tilde{H}^1 (bottom) active contours.

tive 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 active contours is necessary, which is straight forward to obtain from the original active contour.

References

- [1] A. Blake and R. Brockett. On snakes and estimation theory. In *IEEE CDC*, 1994.
- [2] A. Blake and M. Isard. *Active Contours*. Springer Verlag, 1998.
- [3] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proceedings of the IEEE Int. Conf. on Computer Vision*, pages 694–699, Cambridge, MA, USA, June 1995.
- [4] T. Chan and L. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, February 2001.
- [5] G. Charpiat, R. Keriven, J. Pons, and O. Faugeras. Designing spatially coherent minimizing flows for variational problems based on active contours. In *ICCV*, 2005.
- [6] M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking. *IJCV*, 1(29):5–28, 1998.
- [7] J. Jackson, A. Yezzi, and S. Soatto. Tracking deformable moving objects under severe occlusions. In *IEEE Conference on Decision and Control*, Dec. 2004.
- [8] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1987.
- [9] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *Proceedings of the IEEE Int. Conf. on Computer Vision*, pages 810–815, 1995.
- [10] A. Kriegel and P. Michor. *The Convenient Setting of Global Analysis*, volume 53 of *Mathematical Surveys and Monographs*. American Mathematical Society, 1997.
- [11] P. Michor and D. Mumford. Riemannian geometries on the space of plane curves. *ESI Preprint 1425*, *arXiv:math.DG/0312384*, Dec. 2003.
- [12] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989.
- [13] M. Niethammer and A. Tannenbaum. Dynamic geodesic snakes for visual tracking. In *CVPR*, volume 1, pages 660–667, 2004.
- [14] N. Paragios and R. Deriche. Geodesic active regions: a new paradigm to deal with frame partition problems in computer vision. *International Journal of Visual Communication and Image Representation, Special Issue on Partial Differential Equations in Image Processing, Computer Vision and Computer Graphics*, 13(2):249–268, June 2002.
- [15] Y. Rathi, N. Vaswani, A. Tannenbaum, and A. Yezzi. Particle filtering for geometric active contours and application to tracking deforming objects. In *IEEE CVPR*, 2005.
- [16] S. Soatto and A. J. Yezzi. DEFORMOTION: Deforming motion, shape average and the joint registration and segmentation of images. In *ECCV (3)*, pages 32–57, 2002.
- [17] G. Sundaramoorthi, A. Yezzi, and A. Mennucci. Sobolev active contours. In *VLSM*, pages 109–120, 2005.
- [18] D. Terzopoulos and R. Szeliski. *Active Vision*, chapter Tracking with Kalman Snakes. MIT Press, 1992.
- [19] A. Yezzi and A. Mennucci. Metrics in the space of curves. *Preprint, arXiv:math.DG/0412454*, May 2005.