Towards a Local Kalman Filter for Visual Tracking

Ibrahima J. Ndiour and Patricio A. Vela
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0250

***Abstract*— This paper considers the task of closed curve filtering for visual tracking. Segmentation-based visual tracking strategies provide the closed curve measurements to filter. This paper discusses the derivation of a local, linear description for planar curve variation and curve uncertainty. It consists of a family of non-intersecting trajectories transverse to a given curve. Along one of the single-dimensional transverse trajectories, linear curve operations are feasible. Using the linear operation, a simple, locally optimal filtering procedure is derived. The filtering procedure is then used to define an observer for segmentation-based visual tracking. Experiments conducted validate the proposed method and resulting observer design.**

I. Introduction

Segmentation-based tracking algorithms in computer vision involve the tracking of a target boundary. When there is imaging noise, the segmentation model is incorrect, or other similar perturbations exist, the bounding contour is incorrect and corrupted. Over time, poor measurements lead to loss of track. Filtering of the segmenting contour provides a means to correct this problem [1]. This paper sketches initial progress towards a strategy for optimally filtering planar, smooth, closed curves, given a temporal sequence of such curves as generated by a segmentation algorithm. Due to the infinite-dimensional nature of the space of smooth, closed curves, this is but one potential approach [10]. The contents advance recent work in this area [12] and attempt to integrate related work [8].

Related work regarding curve filtering strategies includes [2], [11], [15]. Some works consider the problem to be a nonlinear finite-dimensional system [4], [5], [18]. Others consider the optimal filtering problem using the entire spatio-temporal volume [13]. For additional references on related work, we refer to [12].

Currently, most of the recent filtering designs proposed for tracking consider a fixed gain, manually specified. High gains are chosen when the image sequences are free of visual disturbances or the measurement model easily distinguishes target from background. Low gains are chosen in the opposing case. While the gains do relate to qualitatively observable aspects of the image sequence, they have not been formally tied to quantifiable aspects of the image sequence, such as uncertainty of the image data, of the segmentation, or of the segmentation models.

Contribution: The principal contribution of the paper is the derivation of a mechanism for estimating the optimal gain associated to the curve filtering process for planar curves, given quantitative uncertainty levels. To achieve this, we define a family of transverse curves to the curve in question. The family of transverse curves forms a collection of coordinate frames within which curve operations are linear. Linearity is only valid within a neighborhood of the curve in question. Once the curve filtering equations are derived, they are placed within the greater context of observer design for the estimation of a curve as it evolves in the plane. The observer states include the rigid pose and the deforming curve states. The correction of the two states is decoupled. The work can be extended to the 3D case involving surfaces.

The paper is organized as follows. In Section II, we describe the transverse coordinate system for closed curves. Section III describes the formulation of the optimal filtering gain. Section IV discusses the extension of the work to the problem of visual tracking. Experimental validation follows in Section V. Section VI concludes the paper. While this paper attempts to be self-contained, understanding of [8], [12] would be beneficial.

II. Transverse Curve Coordinates

In [12], a curve filtering strategy was proposed relying on curve correspondences derived from a combined Laplace-Poisson equation approach. Given two smooth curves, a scalar function defined on the curve domain was derived. The scalar function generated an associated error vector field. In this section, we describe how the setup is modified to obtain a characteristic vector field (which agrees with the error vector field modulo sign). The characteristic vector field will define the family of 1D transverse coordinate systems for a given smooth curve local to the two smooth curves.

A. The Characteristic Vector Field

The error vector field associated to two smooth curves $\mathcal{C}_0$ and $\mathcal{C}_1$ and their local region is obtained from the solution to a Laplace-Poisson partial differential equation (PDE), for which the induced flow is a diffeomorphism.

The variational formulation of the Laplace-Poisson problem is

$$\min_u \int \|\nabla u\|^2 \, d\Omega, \tag{1}$$

such that $trace(\mathcal{C}_0) = u^{-1}(0)$ and $trace(\mathcal{C}_1) = u^{-1}(1)$. Its solution requires careful construction of the interior and boundary conditions. The source curve and the target curve define the following solution domain decomposition of the total space Ω, $R := \overline{int(\mathcal{C}_0) \ominus int(\mathcal{C}_1)}$, $R_{pi} := int(\mathcal{C}_0) \cap int(\mathcal{C}_1)$, and $R_{lo} := \Omega \setminus \left(\overline{R \cup R_{pi}}\right)$, where $int(\mathcal{C})$ denotes the interior of the curve $\mathcal{C}$ and $\ominus$ is the set-symmetric difference; see Figure 1(a) for a depiction of the domains.

Instead of the defined boundary conditions above, set the boundary conditions to 0 for the interior curve portions $(\partial\overline{R_{pi}} \setminus (\mathcal{C}_0 \cap \mathcal{C}_1))$ and to 1 for the exterior curve portions $(\partial\left(R \cup R_{pi}\right))$. The exterior and the interior curve parts may comprise of subsets of $\mathcal{C}_0$ *and* $\mathcal{C}_1$ if $\mathcal{C}_0$ and $\mathcal{C}_1$ intersect. Via the calculus of variations, a solution to (1) in the domain enclosed by the source and target curves *with modified boundary conditions* satisfies

$$\Delta u_s(\boldsymbol{x}) = 0, \quad \boldsymbol{x} \in R, \tag{2}$$

where $\Delta = \nabla^2$, with the boundary conditions

$$\begin{aligned} u_s(\boldsymbol{x}) &= 0, \ \boldsymbol{x} \in \partial\overline{R_{pi}} \setminus (\mathcal{C}_0 \cap \mathcal{C}_1), \\ u_s(\boldsymbol{x}) &= 1, \ \boldsymbol{x} \in \partial\left(R \cup R_{pi}\right), \end{aligned} \tag{3}$$

which is a simple reformulation of the minimization problem (1) based on the domain decomposition depicted in Figure 1. On R_{pi}:

$$\Delta u_{pi}(\boldsymbol{x}) = c \quad \boldsymbol{x} \in \overline{R_{pi}}, \ c > 0 \tag{4}$$

$$\Delta u_{lo}(\boldsymbol{x}) = 0, \quad \boldsymbol{x} \in \overline{R_{lo}}, \tag{5}$$

with boundary conditions

$$\begin{aligned} &u_{pi}(\boldsymbol{x}) = 0, \ \boldsymbol{x} \in \partial\overline{R_{pi}}, \\ &u_{lo}(\boldsymbol{x}) = 1, \ \boldsymbol{x} \in \partial\left(R \cup R_{pi}\right), \\ &u_{lo}(\boldsymbol{x}) = 2, \ \boldsymbol{x} \in \partial\Omega. \end{aligned} \tag{6}$$

The combined solution

$$u(\boldsymbol{x}) = \begin{cases} u_{lo}(\boldsymbol{x}), & \boldsymbol{x} \in R_{lo}, \\ u_{pi}(\boldsymbol{x}), & \boldsymbol{x} \in R_{pi}, \\ u_s(\boldsymbol{x}), & \boldsymbol{x} \in R \end{cases} \tag{7}$$

defines the characteristic vector field X_{char} on Ω,

$$X_{char}(\boldsymbol{x}) := \begin{cases} \nabla u / \|\nabla u\|, & \boldsymbol{x} \in R, \\ \nabla u_o / \|\nabla u_o\|, & \boldsymbol{x} \in R_{lo}, \\ \nabla u_i / \|\nabla u_i\|, & \boldsymbol{x} \in R_{pi} \end{cases} \tag{8}$$

via the normalized gradient. Figure 1(b) depicts the distance map generated by the characteristic vector field arising from two curves. Following the distance characteristics forwards and backwards starting at a curve point defines the local transverse coordinate system. Figure 1(c) depicts a few transverse coordinate systems associated to the blue curve. The white lines indicate a range of $[-5, 5]$ in the local 1D coordinate system for five different blue curve points intersected by the white lines.

The described scheme is fast and parametrization-free. It is not invariant to translations, rotations or scale. Therefore, we will assume that the two curves are or have been registered with respect to rotation and translation.

B. The Family of 1D Tranverse Coordinate Frames.

A family of tranverse 1D curves will be defined for a closed curve $\mathcal{C}$ that lies within a local region of the curve $\mathcal{C}_0$ and $\mathcal{C}_1$ (it may well be equal to $\mathcal{C}_0$ or $\mathcal{C}_1$). Given the field X_{char} and a particle $\boldsymbol{x}_0 \in \mathcal{C}$, its traveling distance, d, at position $\boldsymbol{x}$ along the characteristic of X_{char} through $\boldsymbol{x_0}$ is defined as the arc-length of the characteristic curve connecting $\boldsymbol{x}_0$ and $\boldsymbol{x}$. To measure these traveling distances from a complete set of initial locations, as specified by $d^{-1}(\cdot, 0)$, solve

$$\begin{cases} d(\cdot, 0) = 0, \\ d_\tau + X_{char}^T \nabla_x d = 1, \end{cases} \tag{9}$$

where τ is an artificial time parameter for the PDE equation and $d : \mathbb{R}^2 \times \mathbb{R}^+ \to \mathbb{R}$. From (8), X_{char} has unit norm.

For a given curve, $\mathcal{C}$, the travelling distance map and its characteristics define the 1D family of transverse coordinate frames. Let s be the arc-length parameter of the curves in question, and assume that $\mathcal{C}_0(0)$, $\mathcal{C}_1(0)$ and $\mathcal{C}(0)$ lie on the same characteristic. Then at $\mathcal{C}(s)$ the associated transverse curve intersects, for example, the curves $\mathcal{C}_0$ and $\mathcal{C}_1$, at the same arc-length parameter s. The coordinate location on the transverse curve of the curve point $\mathcal{C}_0(s)$ is given by

$$sd(\mathcal{C}_0(s); \mathcal{C}(s)) = d(\mathcal{C}_0(s), 0), \quad \text{where } \mathcal{C} = d^{-1}(\cdot, 0),$$

where $sd(p; \mathcal{C})$ describes the signed travelling distance from a point $\boldsymbol{x}$ to its corresponding point on $\mathcal{C}$. The distance is negative when $\boldsymbol{x}$ lies interior to $\mathcal{C}$ and positive exterior to $\mathcal{C}$.

To be more concrete about the fact that these transverse curves have their own coordinates, we will specify the coordinates on the transverse line of a particular curve point $\hat{\mathcal{C}}(s)$ by $\hat{x}(s)$, whose value is given by $\hat{x}(s) = sd(\hat{\mathcal{C}}(s); \mathcal{C}(s))$. We will call this the *point notation* for a curve. The point notation will be used to define linear operations on curves that are locally close, and to define curve variances.

III. Curve Filtering and the Optimal Gain

This section describes the derivation of a filtering scheme on the space of smooth, closed curves. The scheme itself relies on three fundamental assumptions:

Assumption 1: The curves have been approximately registered.

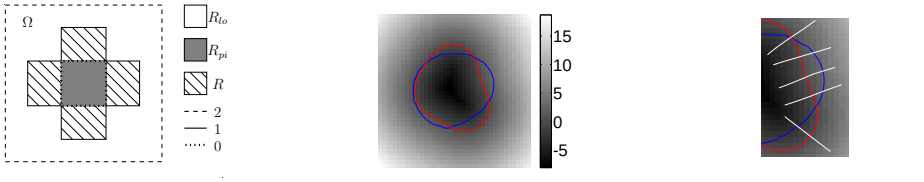

(a) Solution domain decomposition. (b) Distance map using X_{char}. (c) Transverse coordinates.

Fig. 1. The topology and geometry of curve comparison and transverse curve coordinates.

Assumption 2: Given the level of uncertainty about the closed curve, the error is local. In particular, local curve metrics suffice to quantify the error.

Assumption 3: Curve variance and uncertainty is given by a smooth scalar function along the curve.

These assumptions exist to make the curve filter feasible in practice. The rest of this section details the consequences of the assumptions.

Because the contents of this section will ultimately serve to temporally filter a curve, we will utilize conventional estimation notation (involving hats). Consider two curves, $\hat{\mathcal{C}}$ and $\hat{\mathcal{C}}^-$, which are both estimates for the true curve $\mathcal{C}$. Further, consider an additional curve $\mathcal{C}^m$ representing a noisy measurement of the true curve $\mathcal{C}$.

The curve error between $\hat{\mathcal{C}}$ and $\mathcal{C}$, and $\hat{\mathcal{C}}^-$ and $\mathcal{C}$ can be described using the distance function derived from X_{char}. We will use the curve $\mathcal{C}$ as the zero curve, e.g., $\mathcal{C} = d^{-1}(\cdot; 0)$. In point notation, the curve errors of $\hat{\mathcal{C}}^-$ and $\hat{\mathcal{C}}$ with respect to the curve $\mathcal{C}(s)$ are given by

$$\hat{e}^-(s) = \hat{x}^-(s) - x(s), \text{ and}$$
$$\hat{e}(s) = \hat{x}(s) - x(s)).$$

The variance associated with the error is

$$P^-(s) = E\left(\left(\hat{x}^-(s) - x(s)\right)^2\right) > 0, \text{ and}$$
$$P(s) = E\left(\left(\hat{x}(s) - x(s)\right)^2\right) > 0.$$

Furthermore, presume that the measurement error variance, computed in a similar fashion is $R(s) > 0$ and varies smoothly with s, and that the measurements are independent of $\hat{C}^-$, $\mathrm{Cov}(\mathcal{C}^{\mathrm{m}}, \hat{\mathcal{C}}^-) = 0$.

What we seek is an optimal selection of K so that the covariance $P(s)$ is minimized under the update

$$\hat{x}(s) = \hat{x}^-(s) + K(\hat{x}^m(s) - \hat{x}^-(s)), \tag{10}$$

given in point notation,

Note that this setup essentially reduces the problem of finding the optimal selection of K to a one-dimensional problem. In particular, it reduces to the Kalman gain for a 1D system [3]. Thus the optimal choice of K is

$$K = P^-\left(P^- + R\right)^{-1}, \tag{11}$$

and the associated error variance is

$$P = P^-\left(1 - K\right). \tag{12}$$

Since all quantitites are smooth and bounded away from zero as needed for denominators, the above operations are smooth. Therefore, the resulting curve will be smooth. What remains to be shown is: (1) that the selection of Equation (10) is unbiased, and (2) what the second-order filtering strategy should be. We leave this for future work.

IV. Extension to Visual Tracking

The derived filtering strategy is applied to visual tracking in this section. Contour-based visual tracking will require the curve to be locally registered from frame to frame, therefore we will consider the problem to be one of identifying both the rigid motion and the non-rigid motion associated to the track target [19]. Rigid motion can be given by the Euclidean linear group, $E(2)$, or the special Euclidean group, $SE(2)$, depending on one's needs. Non-rigid motion will be determined by a curve and its normal motion.

Combining the curve filtering strategy discussed earlier, with a group filtering strategy, plus a dynamical model for the system will result in an observer for visual tracking. The observer components are described below in the ordering: dynamic model, sensor measurement model, system measurement model, and correction. Because the second-order dynamical model is nonlinear, the gain calculation (11) will only be used on the shape state, whereas the curve velocities will have static gains. The observer is continuous-discrete, e.g., continuous time dynamics with discrete time measurements.

A. Dynamical Model.

The dynamical model describes the target's evolution. For the rigid group states, we presume a constant velocity model,

$$\dot{g} = g \cdot \xi, \qquad \dot{\xi} = 0.$$

For the contour, model prediction will be done using the dynamic elastic prior, which is a dynamic active contour [11], minimizing the action integral

$$\mathcal{L} = \int_{t=t_0}^{t_1} \int_0^1 \frac{1}{2}\mu \|\hat{\mathcal{C}}_t\|^2 \|\hat{\mathcal{C}}_p\| dp dt, \tag{13}$$

with μ a mass constant. When restricted to normal curve propagation, the resulting dynamic elastic prior is

$$\mathcal{C}_t = \beta \mathcal{N}, \quad \beta_t = \frac{1}{2}\beta^2 \kappa.$$

where s denotes arc-length, κ is the curvature, $\mathcal{T}$ is the unit tangent vector to $\hat{\mathcal{C}}$, and β is the normal speed in the direction of $\mathcal{N}$ [11]. Additional priors are described in [12]. The curve covariance is propagated along with the curve $\mathcal{C}_t$ and increases by $Q(s)$, a smooth scalar function. Further the noise process leading to $Q(s)$ is independent of the curve $\mathcal{C}(s)$ and the measurement $\mathcal{C}^m$.

The prediction will generate $(\hat{g}, \hat{\xi}, \hat{\mathcal{C}}, \hat{\beta})$, in addition to propagating forward any covariances (group) or variances (shape) associated with the system. If desired, the system can be reduced to a single-order static model.

B. Sensor Measurement

Measurement of the target can be achieved through any segmentation algorithm applied to the current image. The segmentation can then be converted into a curve description if it is not already of that form. Candidate algorithms include Bayesian segmentation [6], active contours [16], graph cuts [9], etc.

In a classical observer the measurements would be completely independent of the observer states, however image analysis of video has the nature of not explicitly providing the necessary signal. Instead it must be extracted from the image using an image processing or computer vision algorithm. The measurement procedure may not completely determine the necessary target state measurements (due to non-uniqueness of the group + shape decomposition). Consequently, a registration step is required to describe the predicted and measured shape with respect to the same coordinate frame.

Once segmentation and localization are performed on the current image, a registration procedure is applied to match the measured contour with the predicted contour, yielding a measurement g_m for the group motion and the measurement $\mathcal{C}_m$ for the shape. If desired, the velocity field β_m can be measured by computing the optical flow [7] between two subsequent aligned images and projecting it onto the measured curve normals. In practice, the group velocity ξ is not directly measurable.

C. Model Measurement

The model measurement is obtained by extracting the components of the internal state model that are equivalent to those obtained from the sensor measurement.

D. Correction

From a prediction on the observer state and a measurement on the observable part of the state, the update step updates the estimate of the observer internal state.

a) Group correction: Since the group motion is either a translation or a composition of a rotation and translation, it can be represented by a three-state vector. Depending on the prediction motion model chosen (linear vs. nonlinear), finite-dimensional filtering such as Kalman, extended or unscented Kalman are appropriate to use in order to correct the measured group.

b) Shape correction: Correction on the shape requires the construction of the transverse coordinates. The first step is to generate the characteristic vector field using the predicted and measured curves. While the actual curve state was used to define the origin of the transverse curve coordinates in Section II, any curve within the local neighborhood can be used as the reference curve for the transverse coordinate frames. Invariance to the reference curve is due to the fact that the filter equations rely on differences and that the characteristic vector field is independent of the chosen curve. Here we use the predicted curve.

The correction applied to the curve, in point notation, is then

$$\hat{x}(s) = \hat{x}^-(s) + K(x^m(s) - \hat{x}^-) = Kx^m(s),$$

since the predicted curve is at the origin. The curve state $\hat{\mathcal{C}}(s)$ is uniquely defined by $\hat{x}(s)$. The gain chosen is the optimal gain selection from Equation (11), and the covariance update is as per Equation (12).

The shape velocity lives in the tangent space of the shape manifold. Since it is a vector space, the velocity field is relatively simple to correct. Through errors between the measured and predicted shape and between the predicted and measured shape velocities, we can induce correction on the shape velocity:

$$\hat{\beta}(s) = \hat{\beta}^- + K_{vx}(x^m(s) - \hat{x}^-(s)) + K_{vv}(\beta^m(s) - \hat{\beta}^-(s))$$

The parameters K_{vx} and K_{vv} are constants in the range $(0, 1)$, chosen by the user according to expected velocity measurement noise.

V. Experiments and Results

c) Error Measurement: One of the fundamental assumptions underlying the method is that curve error is localized. Being so, a locally valid linear description is feasible. To verify when such locality occurs, we generated noisy grayscale images to segment with known ground truth. Then we computed the Laplace error of the system as a function of the level of noise. The level of noise is measurable through the Bhattacharyya coefficient. Figure 2 depicts the expected error as a function of the noise level. At the level 0.6 and above, the image starts to get sufficiently corrupted that the segmentations become nonsensical (the error variance spikes; not depicted). Thus, for reasonable imagery, we can expect the error to be sufficiently local.

d) Static Kalman Filtering: In order to verify that the behavior as derived is occuring, we performed tests on a static image corrupted with noise. A series of segmentations were generated and then locally Kalman filtered using the derived equations. A gain sweep from 0.05 to 0.95 in 0.05 gain increments was performed to verify if the Kalman gain converged to is indeed optimal. Several

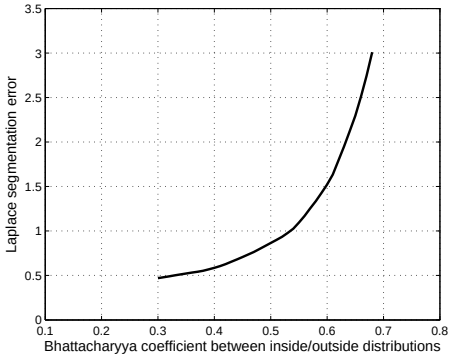

Fig. 2. Expected Laplace error versus Bhattacharya coefficient.

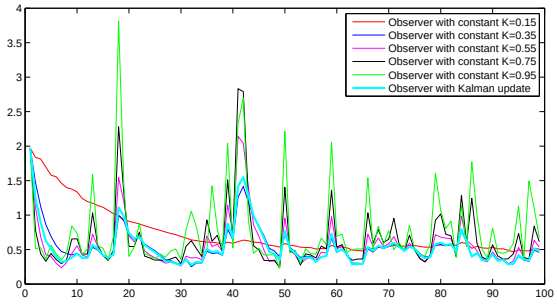

Fig. 3. Gain sweep test.

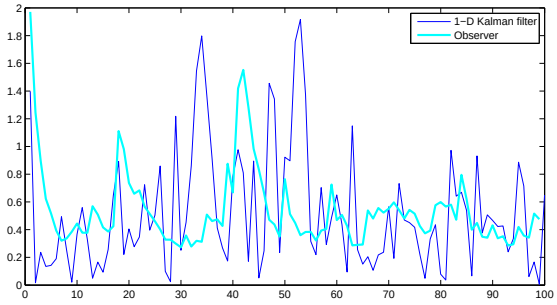

Fig. 4. Visual comparison against a true 1D system.

sweeps were run, one of which is depicted in Figure 3 for a limited set of gains (to avoid figure clutter). The red curve corresponds to a low gain, and little update of the observer state. The cyan curve represents the optimal gain based on the derived mathematics.

Additionally, several 1D Kalman filter simulations were run to compare the Laplace simulations against a true 1D system. Figure 4 depicts the evolution of the error for a 1D system and for a simulated static tracking scenario. The converged error variance of the two systems is of the same order.

For this experiment and those that follow, the gain was selected to be $K(s,0) = K_0$, and the variances $P(s) = P_0$, $Q(s) = Q_0$, and $R(s) = R_0$. Thus, as a function $K(s,k) = K_0(k)$.

e) Tracking Experiments: Tracking experiments were performed on a variety of image sequences. The first image sequence is extracted from a low SNR grayscale video featuring the deformation of a biomembrane. The remaining color sequences are obtained from a construction database; they are difficult to track because of poor resolution, clutter and partial occlusions of the targets.

The local Kalman filter algorithm and three other contour-based tracking algorithms were tested. The three algorithms were: Bayesian active contours [16], the deformotion filter [8], and a shape-based filter [4]. Implementation of the contour equations was done in the level-set formulation, which describes curves implicitly using a signed-distance function [17]. For details on how the filtering implementation works using the implicit contour description, see [12]. For all these tracking techniques, we apply Kalman filtering on the rigid group state with the same measurement and noise covariance matrices.

The shape database consisted of 67 sample shapes obtained from the construction imagery with the top 10 eigenvectors kept. For the deformotion tracking technique [8], a gain was selected to the best of our understanding. For the proposed method, we defined the observation and process noises, Q and R, for both the group and the shape spaces, plus static gains for the curve velocity correction gains.

Manual segmentations of the video sequences provide the ground truth. For quantitative comparisons of the results obtained with the different techniques, we used the L_2 and L_∞ norms on the group state. Local and global shape metrics are used to quantify deviations of the shapes from the truth. The number of misclassified pixels and the mean and maximum Laplace errors [14] are used for this purpose. Lastly, the number of frames tracked throughout the sequence is indicated.

Figures 5-8 depict samples of the segmentations obtained at given frames for each sequence and technique. Table I summarizes the performances of the different techniques under the comparison metrics. Strikeouts in the text indicate loss of track. Of note, in sequence 3 the measurement model fails due to shared statistics of the target and background segmentation models. The observer is able to gracefully handle the poor model.

VI. Conclusion

This paper proposed a method for estimating the locally optimal gain of a curve filtering strategy. The curve error is based upon a Laplace-Poisson PDE solution and provides a set of linear coordinate frames from which to perform curve operations and also compute curve variances. The curve filter was tested on static images and compared against a 1D Kalman filtered system to verify that convergence behavior is as expected. Additional validation was done by applying the filter to a set of image sequences. Future work seeks to extend the Kalman filter to the velocity states.

(a) Sequence 1

Metric / Algorithm	AC	Deformotion	Observer
Trackpt error (L_2/L_∞)	5.4 / 7.5	5.5 / 7.7	5.4 / 7.5
NMP (avg/max)	363 / 616	429 / 692	348 / 627
Mean Laplace (avg/max)	1.0 / 1.6	1.2 / 1.8	0.9 / 1.6
Max Laplace (avg/max)	4.0 / 7.9	4.0 / 8.9	3.6 / 6.8
# Frames tracked	40	40	40

(b) Sequence 2

Metric / Algorithm	AC	Deformotion	Shape	Observer
Trackpt error (L_2/L_∞)	~~2.2 / 6.6~~	~~2.2 / 9.6~~	~~7.6 / 18.5~~	1.8 / 6.2
NMP (avg/max)	~~78 / 202~~	~~72 / 172~~	~~87 / 160~~	63 / 111
Mean Laplace (avg/max)	~~1.0 / 3.7~~	~~0.9 / 3.1~~	~~1.2 / 2.6~~	0.7 / 1.3
Max Laplace (avg/max)	~~2.9 / 8.9~~	~~2.3 / 7.9~~	~~3.4 / 8.4~~	2.0 / 3.5
# Frames tracked	~~109~~	~~109~~	~~115~~	350

(c) Sequence 3

Metric / Algorithm	AC	Deformotion	Shape	Observer
Trackpt error (L_2/L_∞)	2.7 / 6.5	2.7 / 6.5	4.5 / 20.5	2.5 / 9.5
NMP (avg/max)	131 / 276	130 / 290	162 / 330	113 / 180
Mean Laplace (avg/max)	1.2 / 2.7	1.1 / 2.8	1.9 / 4.8	1.0 / 2.1
Max Laplace (avg/max)	3.6 / 7.3	2.3 / 7.9	7.1 / 16.5	3.6 / 8.3
# Frames tracked	150	150	150	150

(d) Sequence 4

Metric / Algorithm	AC	Deformotion	Shape	Observer
Trackpt error (L_2/L_∞)	4.4 / 9.4	4.3 / 9.3	~~3.3 / 9.7~~	2.3 / 5.5
NMP (avg/max)	52 / 149	49 / 125	~~98 / 195~~	48 / 95
Mean Laplace (avg/max)	1.2 / 14.9	1.1 / 18.4	~~2.7 / 8.4~~	0.8 / 2.4
Max Laplace (avg/max)	3.0 / 18.3	2.7 / 21.5	~~8.6 / 18.9~~	2.3 / 6.4
# Frames tracked	1014	1014	~~430~~	1014

TABLE I

PERFORMANCE TABLE USING QUANTITATIVE COMPARISON METRICS.

REFERENCES

[1] A. Blake and M. Isard. *Active Contours*. Springer Verlag, 1998.

[2] R. Brockett and A. Blake. Estimating the shape of a moving contour. In *Proceedings of IEEE Conference on Decision and Control*, pages 3247–3251, 1994.

[3] R.G. Brown and P.Y.C. Hwang. *Introduction to Random Signals and Applied Kalman Filtering*. John Wiley and Sons, Inc., New York, 1992.

[4] D. Cremers. Dynamical statistical shape priors for level set-based tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28:1262–1273, 2006.

[5] S. Dambreville, Y. Rathi, and Tannenbaum A. Tracking deformable objects with unscented Kalman filtering and geometric active contours. In *Proceedings of the American Control Conference*, pages 2856–2861, 2006.

[6] S. Haker, G. Sapiro, A. Tannenbaum, and D. Washburn. Missile tracking using knowledge-based adaptive thresholding: Tracking of high speed projectiles. In *Proceedings of the IEEE International Conference on Image Processing*, pages 786–780, 2001.

[7] B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

[8] J.D. Jackson, A.J. Yezzi, and S. Soatto. Tracking deformable moving objects under severe occlusions. In *Proceedings of IEEE Conference on Decision and Control*, pages 2990–2995, 2004.

[9] J. Malcolm, Y. Rathi, and A. Tannenbaum. Tracking through clutter using graph cuts. In *Proceedings of the British Machine Vision Conference*, page 116, 2007.

[10] 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*, 23(1):74–113, 2007.

[11] M. Niethammer, A. Tannenbaum, and S. Angenent. Dynamic active contours for visual tracking. *IEEE Transactions on Automatic Control*, 51(4):562–579, 2006.

[12] M. Niethammer, P.A. Vela, and A. Tannenbaum. Geometric observers for dynamically evolving curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(6):1093–1108, 2008.

[13] N. Papadakis and E. Mémin. A variational technique for time consistent tracking of curves and motion. *Journal of Mathematical Imaging and Vision*, 31(1):81–103, 2008.

[14] E. Pichon, D. Nain, and M. Niethammer. A Laplace equation approach for shape comparison. *SPIE Medical Imaging*, 2:24–132, 2006.

[15] Y. Rathi, N. Vaswani, and A. Tannenbaum. A generic framework for tracking using particle filter with dynamic shape prior. *IEEE Transactions on Image Processing*, 16(5):1370–1382, 2007.

[16] M. Rousson and R. Deriche. A variational framework for active and adaptative segmentation of vector valued images. In *Proceedings IEEE Workshop on Motion and Video Computing*, pages 56–61, 2002.

[17] J. Sethian. *Level Sets Methods and Fast Marching Methods*. Cambridge University Press, 1999.

[18] N. Vaswani, A. Yezzi, Y. Rathi, and A. Tannenbaum. Time-varying finite dimensional basis for tracking contour deformations. In *Proceedings of IEEE Conference on Decision and Control*, 2006.

[19] A.J. Yezzi and S. Soatto. Deformotion: Deforming motion, shape average and the joint registration and approximation of structures in images. *International Journal of Computer Vision*, 53(2/2):153–167, 2003.

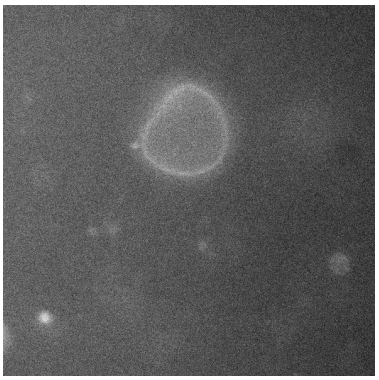

(a) Sequence 1

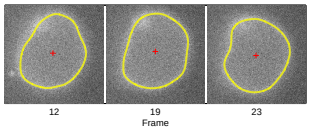

(b) Ground Truth

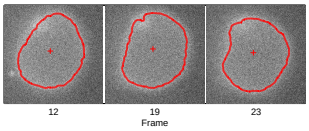

(c) Active Contour

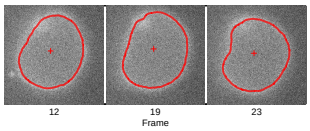

(d) Deformotion Filter

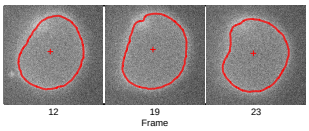

(e) Local Kalman Observer

Fig. 5. Snapshots of Sequence 1.

(a) Sequence 2

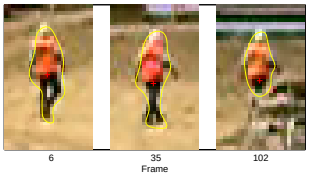

(b) Ground Truth

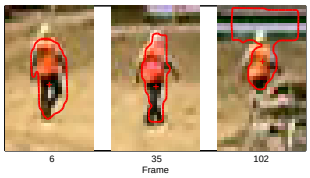

(c) ~~Active Contour~~

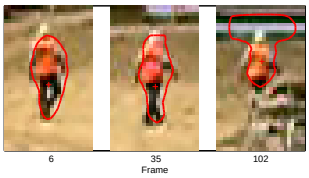

(d) ~~Deformotion Filter~~

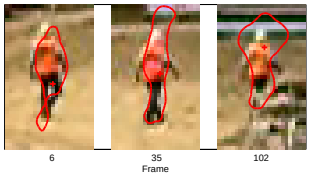

(e) ~~Shape-based~~

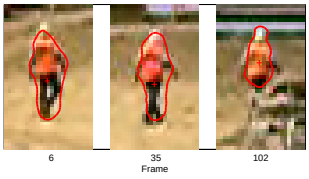

(f) Local Kalman Observer

Fig. 6. Snapshots of Sequence 2 (strikeouts indicate loss of track).

(a) Sequence 3

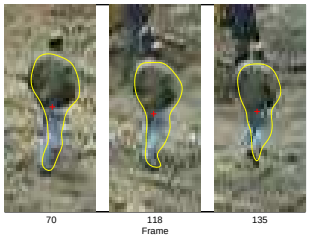

(b) Ground Truth

(c) Active Contour

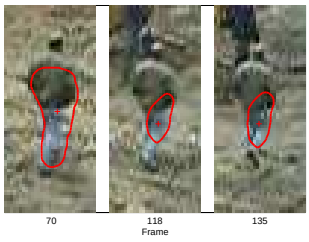

(d) Deformotion Filter

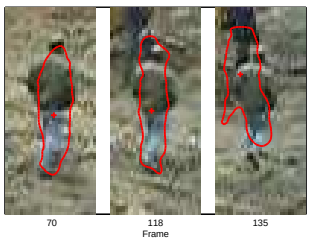

(e) Shape-based

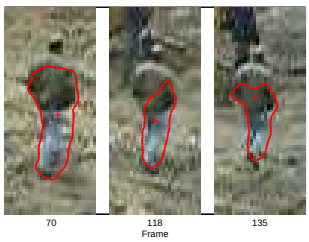

(f) Local Kalman Observer

Fig. 7. Snapshots of Sequence 3.

(a) Sequence 4

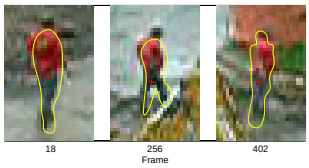

(b) Ground Truth

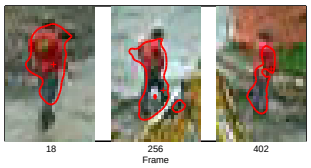

(c) Active Contour

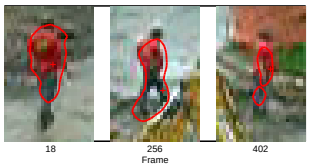

(d) Deformotion Filter

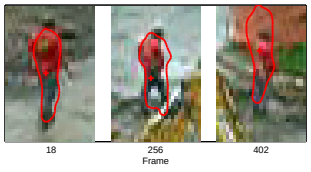

(e) ~~Shape-based~~

(f) Local Kalman Observer

Fig. 8. Snapshots of Sequence 4 (strikeouts indicate loss of track).