

A Local Extended 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 applies estimation theory to the problem of tracking deformable moving objects in an image sequence. We extend previous work to derive a sub-optimal second-order curve filtering strategy. The second-order model accounts naturally for the curve velocities, resulting in better curve predictions. The second-order curve dynamics are nonlinear, so an extended Kalman filtering approach is utilized to estimate the position and deformations of a curve as it evolves in the plane. Application to visual tracking is emphasized through experiments utilizing recorded imagery and providing objective comparisons to other tracking methods.

I. INTRODUCTION

This paper considers the problem of recursively tracking deformable moving objects in an image sequence. Due to image clutter, poor resolution, imaging noise, compression artifacts, or imperfect segmentation models, performing a series of static detections on individual frames extracted from the image sequence results in poor resilience to perturbations. Consequently, information gathered from past detections has been introduced in a variety of manners to increase robustness. Related literature includes methods exploiting temporal consistency of the target [11], methods processing volumetrically the image sequence with batch techniques [13], and methods introducing strong shape priors into the detection [2], [17].

It is common to view the contour tracking problem as an estimation problem [18] given uncertain measurements. The introduction of a Markovian assumption simplifies it to one of recursive estimation, for which many solutions involve the use of *estimators*. The references [1], [3], [4], [7], [9], [12], [19], [20], [21], [22] discuss various filter and estimation strategies for tracking of deformable objects. Given the complications arising from the infinite-dimensional nature of the space of smooth, closed planar curves, many approaches propose estimator designs requiring manual gain selection. Fixed high gains are typically used to conform the state estimates to the measurements when the image sequence is believed to have a high signal-to-noise ratio. Under severe perturbations, fixed low gains are adopted to moderate the injection of measurements into the state estimate. Particle filters do not require gain selection, however the probability distribution of the evolving state needs to be estimated and propagated at

each time-step. For visual tracking applications, this need to maintain a density estimate in the infinite-dimensional shape space introduces a substantial computational burden [15], [20]. Low dimensional parametrized approximations of the shape space are then often used to minimize the number of particles used [3]. In [8], we presented a local, optimal curve filtering strategy. The main contribution of the work consisted in developing a framework in which the optimal gain associated to the curve filtering process could be derived, given knowledge of quantitative uncertainty levels on the image sequence. However, the optimality was only valid for the curve *position*. Curve *velocity* estimation required selection of a constant gain.

Contribution: This paper extends the work in [8]. Here, we present a second-order extended Kalman curve filtering strategy. The nonlinearity of the second-order curve dynamics prevents the derivation of a linear Kalman curve filtering strategy. The second-order model accounts naturally for the curve velocities and couples the update of the curve's position and velocity, which results in better curve estimates. The curve filtering is then utilized in conjunction with standard filtering strategies on the (object) pose to estimate the position, shape, and local shape deformations of a deformable moving object in an image sequence. Optimality is examined in the sense of a quadratic Laplacian error [14] between the curve estimate and the true curve.

Outline: The paper is organized as follows. Section II presents the state space model and the family of transverse curve coordinates used to form the collection of coordinate frames within which curve operations are linear. Section III describes the local second-order curve filtering strategy. Section IV discusses the experimental setup and the corresponding results. The conclusion and further research directions are discussed in Section V.

II. STATE-SPACE REPRESENTATION

This section describes the state-space model used in the curve and pose estimation.

A. Pose-shape decomposition

Under *deformation* [23], a deformable moving object in a video sequence can be described by a group motion and

a shape deformation. Local shape deformations serve to describe the non-rigid curve motion. Naturally then, the state description will include a rigid group component and a shape component. The second-order shape description consists of the planar curve $\mathcal{C} : S^1 \rightarrow \mathbb{R}^2$, and its normal velocity $\beta : S^1 \rightarrow \mathbb{R}$, where S^1 denotes the circle. The circle will be described by a non-empty connected subset of \mathbb{R} that describes the curve parameter (e.g. $S^1 = [0, 1]$ for the arclength parameter s), with the end-points identified to be equal. In what follows, the curve will also be described implicitly by a signed-distance level-set function $\Psi : \mathbb{R}^2 \rightarrow \mathbb{R}$ where $\mathcal{C} = \Psi^{-1}(0)$. When this is the case, the normal velocity will be extended over the level-set domain to define $\hat{\beta} : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that $\hat{\beta} \circ \mathcal{C} = \beta$. For more details on this setup, see [10]. The group motion represents the pose and the pose velocities of the object in a reference coordinate frame; here it is denoted g and ξ , respectively, and will be the special Euclidean group $SE(2)$, or its subgroup $E(2)$, and its Lie algebra. Since the group variable and the Lie algebra live in a finite-dimensional space, standard filtering strategies [5], [6], [18] on finite-dimensional spaces can be applied to the pose and its velocity. This work focuses on the curve filtering process.


B. Curve representation and point notation

The curve \mathcal{C} is implicitly described as the zero level set of a higher-dimensional signed distance function $\Psi : D \rightarrow \mathbb{R}$, where $\mathcal{C} = \{x \in D \mid \Psi(x) = 0\}$ and $D \subset \mathbb{R}^2$ is a compact domain of the plane. The remainder of this subsection briefly describes a local, linear description for planar curve variation and curve uncertainty. It consists of a family of non-intersecting trajectories transverse to the curve \mathcal{C} . The construction of the family of transverse curves defines a coordinate system transverse to the curve.


Given two smooth curves \mathcal{C}_0 and \mathcal{C}_1 , a characteristic error vector field X_{char} is obtained by solving a series of Laplace and Poisson equations over the domain, with boundary conditions defined on the curves and domain boundary [8]. The error vector field computed this way exhibits some beneficial properties. In particular, its characteristic curves are non-intersecting, they approach the curves \mathcal{C}_0 and \mathcal{C}_1 from normal directions, and their speed never vanishes. For a particle $x_0 \in \mathcal{C}$, its traveling distance, d , at position x along the characteristic of X_{char} through x_0 is defined as the arc-length of the characteristic curve connecting x_0 and x . Measuring these traveling distances from a complete set of initial locations is obtained by solving

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

where τ is an artificial time parameter for the PDE equation and $d : \mathbb{R}^2 \times \mathbb{R}^+ \rightarrow \mathbb{R}$. Figure 1-(a) shows an example of the transverse curves and the distance map obtained.



(a) Distance map and transverse curves.



(b) Transverse coordinates.

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

For the curve \mathcal{C} local to \mathcal{C}_0 and \mathcal{C}_1 , the traveling distance map and its characteristics define a 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

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

where $sd(p; \mathcal{C})$ describes the signed traveling distance from a point x to its corresponding point on \mathcal{C} . The distance is negative when x lies interior to \mathcal{C} and positive exterior to \mathcal{C} . Figure 1-(b) depicts the coordinate system along one of the transverse curves (where \mathcal{C} is in red).

C. State-space

The goal of the filtering process is to arrive at estimates of the curve and its normal velocity, denoted by $\hat{\mathcal{C}}$ and $\hat{\beta}$, respectively. To explicitly exploit 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 $x(s)$, whose value is given by $x(s) = sd(\hat{\mathcal{C}}(s); \mathcal{C}(s))$, where $\mathcal{C}(s)$ is presumed to be the true curve state. Likewise, the normal velocity of the curve will be given by $v(s)$. This notation is called 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. For simplicity, let $\mathbf{x}(s) = [x(s), v(s)]^T$. The state $\mathbf{x}(\cdot)$ locally defines the second-order curve state.

III. SECOND-ORDER CURVE FILTERING

This section describes a second-order filtering method for the space of simple, closed planar curves. We discuss a

local, sub-optimal observer whose internal state is used to estimate the state of a moving object in an image sequence. The curve evolution equations are continuous, while the curve and covariance updates and corrections are discrete. In addition to the curve state, the filter state is composed of the curve covariance matrix $\mathbb{P} : S^1 \rightarrow \mathbb{R}^{2 \times 2}$, which is a measure of the curve uncertainty transverse to the curve and the normal curve velocity uncertainty. Due to the infinite-dimensional nature of curve geometry, some technical care is necessary to ensure that the evolving system preserves smoothness of the curve: we retain the filtering assumptions made in [8]. In particular, it is expected that all curves are aligned or have been registered prior to filtering. The curve estimate error is such that local curve metrics sufficiently quantify the error. Further, the measurement error covariance $\mathbb{R} : S^1 \rightarrow \mathbb{R}^{2 \times 2}$, and the initial curve uncertainty \mathbb{P} are presumed to vary smoothly along the curve. Since the curve filtering process is reduced to filterings of curve particles along the single-dimensional transverse curves, these smoothness and locality assumptions guarantee that the estimated state can be smoothly reconstructed.

A. Dynamical Prediction Model

The dynamical prediction model produces an estimate of the filter state at a future time from past estimates. Usually, it relies on prior knowledge about the state evolution resulting in phenomenological motion models. Here, the prediction model does not assume such priors. Rather a general purpose second-order model is used [12]. It is

$$\begin{aligned} \hat{C}_t &= \beta \mathcal{N}, \\ \hat{\beta}_t &= \left(\frac{1}{2} \hat{\beta}^2 + \frac{a}{\mu} \right) \kappa, \end{aligned} \quad (2)$$

where \mathcal{N} is the curve normal, a is a regularization constant, κ is the curvature and μ is the mass density constant for the curve. In the implicit signed-distance level-set representation, the equivalent curve dynamics are given by

$$\begin{aligned} \hat{\Psi}_t &= \hat{\beta} \cdot \left\| \nabla \hat{\Psi} \right\|, \\ \hat{\beta}_t &= \left(\frac{1}{2} \hat{\beta}^2 + \frac{a}{\mu} \right) \nabla \cdot \left(\frac{\nabla \hat{\Psi}}{\left\| \nabla \hat{\Psi} \right\|} \right), \end{aligned} \quad (3)$$

which relies on the facts that Ψ is a signed-distance function and that $\nabla \Psi \perp \nabla \hat{\beta}$ (since $\hat{\beta}$ is β extended along the normal characteristics of the curve).

In addition to the curve states, the curve covariances must be propagated. Given that the evolution equations are nonlinear in nature, a linear discrete approximation to the covariance update is needed. The approximation is derived directly from the linearization of Equation (2). Here, the first-order approximation is

$$\mathbb{P}(s, t + \Delta t) = \mathbb{F} \cdot \mathbb{P}(s, t) \cdot \mathbb{F}^T + \mathbb{Q} \cdot \Delta t^2,$$

where \mathbb{F} , as obtained from (2), is

$$\mathbb{F} = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 + \beta \kappa \Delta t \end{bmatrix} \quad \text{and} \quad \mathbb{Q} = \begin{bmatrix} 0 & 0 \\ 0 & \sigma \beta \end{bmatrix},$$

where \mathbb{Q} is the process noise covariance, $\mathbb{Q} : S^1 \rightarrow \mathbb{R}^{2 \times 2}$, due to uncertainty in the curve acceleration. The curve covariance is updated discretely with the curve. Furthermore, the update occurs for the extended covariance matrix $\bar{\mathbb{P}} : D \rightarrow \mathbb{R}^{2 \times 2}$, which is the covariance matrix \mathbb{P} extended to the entire level-set domain. After each time-step, the covariance is extended, much like the normal velocity (see [10] for details on evolving extended quantities).

The final predicted state will be denoted by \mathcal{C}^- , β^- , and \mathbb{P}^- . In the implicit representation with extended fiber elements, the prediction consists of Ψ^- , $\bar{\beta}^-$, and $\bar{\mathbb{P}}^-$.

B. Measurement Model

The curve position measurement \mathcal{C}^m for the current frame is obtained from any foreground/background segmentation algorithm applied to the current image, so long as the final result is a level-set function Ψ^m . To ensure proper alignment of the curves, either a localization procedure should be applied before segmentation or a registration method after segmentation.

Measurements for the curve velocities are obtained by computing the optical flow $(u, w)^T$ between two consecutive (aligned) images and projecting it onto the unit normals of Ψ^m :

$$\bar{\beta}^m = \begin{pmatrix} u \\ w \end{pmatrix} \cdot \frac{\nabla \Psi^m}{\left\| \nabla \Psi^m \right\|}.$$

Note that this measurement is defined over the entire image. The measured velocities, β^m , will be the normal velocities obtained by evaluation of $\hat{\beta}^m$ on the zero level-set, $\beta^m = \bar{\beta}^m \circ (\Psi^m)^{-1}(0)$.

Once the measurements are available, they need to be converted into the transverse coordinate system obtained from the prediction and the measurement curves. The procedure, as delineated in §II-B, involves the construction of the characteristic error vector field X_{char} arising from the curves $\hat{\Psi}^-$ and Ψ^m . The zero coordinate refers to the predicted curve \mathcal{C}^- , thereby yielding the local distance to the measurements x^m , and the normal velocity v^m .

C. Update Model

The role of the update model is to refine the prediction once state measurements become available. Here, for each particle along its corresponding single-dimensional transverse curve, an update is performed on the state prediction. Given the setup, along one of these transverse curves, the estimation problem faced is one of linear filtering for which an optimal solution is described by the Kalman filter [6], [18]. Following the previous work [8], define the curve error, the predicted curve error, and the measurement curve error by

$$\begin{aligned} \hat{e}(s) &= \hat{\mathbf{x}}(s) - \mathbf{x}(s), \\ \hat{e}^-(s) &= \hat{\mathbf{x}}^-(s) - \mathbf{x}(s), \quad \text{and} \\ \mathbf{e}^m(s) &= \mathbf{x}^m(s) - \mathbf{x}(s), \end{aligned}$$

and curve covariances

$$\begin{aligned} \mathbb{P}(s) &= E(\hat{\mathbf{e}}(s)\hat{\mathbf{e}}^T(s)) > 0, \\ \mathbb{P}^-(s) &= E(\mathbf{e}^m(s)(\mathbf{e}^m(s))^T) > 0, \end{aligned}$$

and $\mathbb{P}^m = \mathbb{R} > 0$ where \mathbb{R} is the measurement error covariance (smoothly varying along the curve). It is assumed that the measurements and predictions are independent, $\text{Cov}(\mathbf{x}^-, \mathbf{x}^m) = 0$.

The optimal correction gain associated to these errors and covariances under the update law, given in point representation,

$$\mathbf{x}^+(s) = \mathbf{x}^- + \mathbb{K}(\mathbf{x}^m - \mathbf{x}^-), \quad (4)$$

is $\mathbb{K} = \mathbb{P}^-(\mathbb{P}^- + \mathbb{R})^{-1}$. In the remainder of the paper, the gain matrix \mathbb{K} is decomposed as:

$$\mathbb{K} = \begin{bmatrix} K_{xx} & K_{xv} \\ K_{vx} & K_{vv} \end{bmatrix}.$$

The following discussion sketches the implementation of the update equations for the implicit representation of the curve and its fiber (velocity plus covariance).

1) *Curve position correction:* In point notation, the update law of the curve position is obtained from Equation (4),

$$\hat{\mathbf{x}}^+ = \hat{\mathbf{x}}^- + K_{xx} \cdot (\mathbf{x}^m - \hat{\mathbf{x}}^-) + K_{xv} \cdot (\mathbf{v}^m - \hat{\mathbf{v}}^-). \quad (5)$$

The curve defined by $\hat{\mathbf{x}}^+(s)$ becomes the posterior curve estimate.

2) *Curve velocity correction:* The velocity component is updated according to

$$\hat{\mathbf{v}}^+ = \hat{\mathbf{v}}^- + K_{vx} \cdot (\mathbf{x}^m - \hat{\mathbf{x}}^-) + K_{vv} \cdot (\mathbf{v}^m - \hat{\mathbf{v}}^-). \quad (6)$$

3) *Covariance update.:* The covariance update is

$$\mathbb{P}^+ = (\mathbb{1} - \mathbb{K}) \mathbb{P}^-. \quad (7)$$

The position and velocity updates are separated above because each must be dealt with separately. The curve update follows naturally from Equation (4), however the velocity and covariance need a special procedure. To perform the updates, predicted and measured curve velocities and covariances need to be transported to the updated curve location where they can be compared. Considering the velocity case only, transport is done by solving the advection equations:

$$\begin{cases} \phi_\tau + X_{char}^T \cdot \nabla \phi = 0 \\ \psi_\tau + X_{char}^T \cdot \nabla \psi = 0 \end{cases} \quad (8)$$

with $\phi(\cdot; \tau = 0) = \mathbf{v}^m$, $\psi(\cdot; \tau = 0) = \hat{\mathbf{v}}^-$, and τ an artificial time parameter. A similar procedure is applied to the covariance matrix. By extending the velocity and covariance along the characteristics and performing the update over the domain, the update is performed for the implicit representation. Such an approach facilitates reconstruction of the updated curve, its normal velocity, and its covariance after which the curve is re-initialized to be a signed-distance function while the normal velocity and covariance are extended along the normal characteristics.

IV. EXPERIMENTS AND RESULTS

The following experiments are described in order to validate the filter design and resulting performance. Since the second-order model is primarily motivated by the need to better estimate curve's deformations, we chose three test sequences presenting large and fast changes in shape. The first sequence of images comes from a construction database and the remaining two come from an aquarium database. In addition to the fast shape variations, poor resolution, clutter and non-stationary camera make the tracking task difficult. In all cases, the mass density was set to $\mu = 1$ and the curve regularization was set to $a = 1$.

As for the test algorithms, base measurements are obtained by applying an active contour tracking technique [16] to the test sequences. Subsequently, the deformation filtering technique [4] and our sub-optimal second-order filter are applied to the measurements. For comparison purposes, we manually segmented the construction sequence and 100 frames from each aquarium sequence to yield the ground truth. In the figures, the tracking techniques are labelled as Ground Truth, AC Measurements, Deformation Filter and AC Filter.

In order to provide objective comparison, the number of misclassified pixels (NMP) is used as a quantitative error metric. It represents a measure of fidelity to the true curve being estimated. Additionally, a smoothness index is employed. At each frame, it consists of a dissimilarity measure between the shapes at said frame and previous frame. Hence, this metric provides a measure of the temporal consistency of the target and indicates how fast the shape variations are.

In the first experiment, the aim is to track a construction worker. Active contour measurements spike when other workers presenting similar color distributions approach the target. Our second-order filter is able to attenuate these perturbations when they occur but also capture the correct shape when facing moderate to low segmentation noise. Comparatively, the deformation filter eliminates the perturbations but tends to oversmooth. Figure 2 depicts sample estimates obtained with the different tracking techniques as well as the estimation error and the smoothness indices.

The second experiment targets a fish in an aquarium. The movement and shape deformations of the target are dramatically fast in this test sequence. Moreover, the color distributions of the target and elements of the background overlap significantly. This results in deteriorating performance for the active contour measurements. Both the deformation filter and the local second-order filter drastically attenuate the perturbations and maintain temporal consistency of the target's shape (see Figure 3).

Figure 4 depicts sample estimates and the results corresponding to the third experiment. Here, a fish is quickly deforming with fast movements of the caudal fin. A small overlap between target and background color distributions results in acceptable active contour measurements with

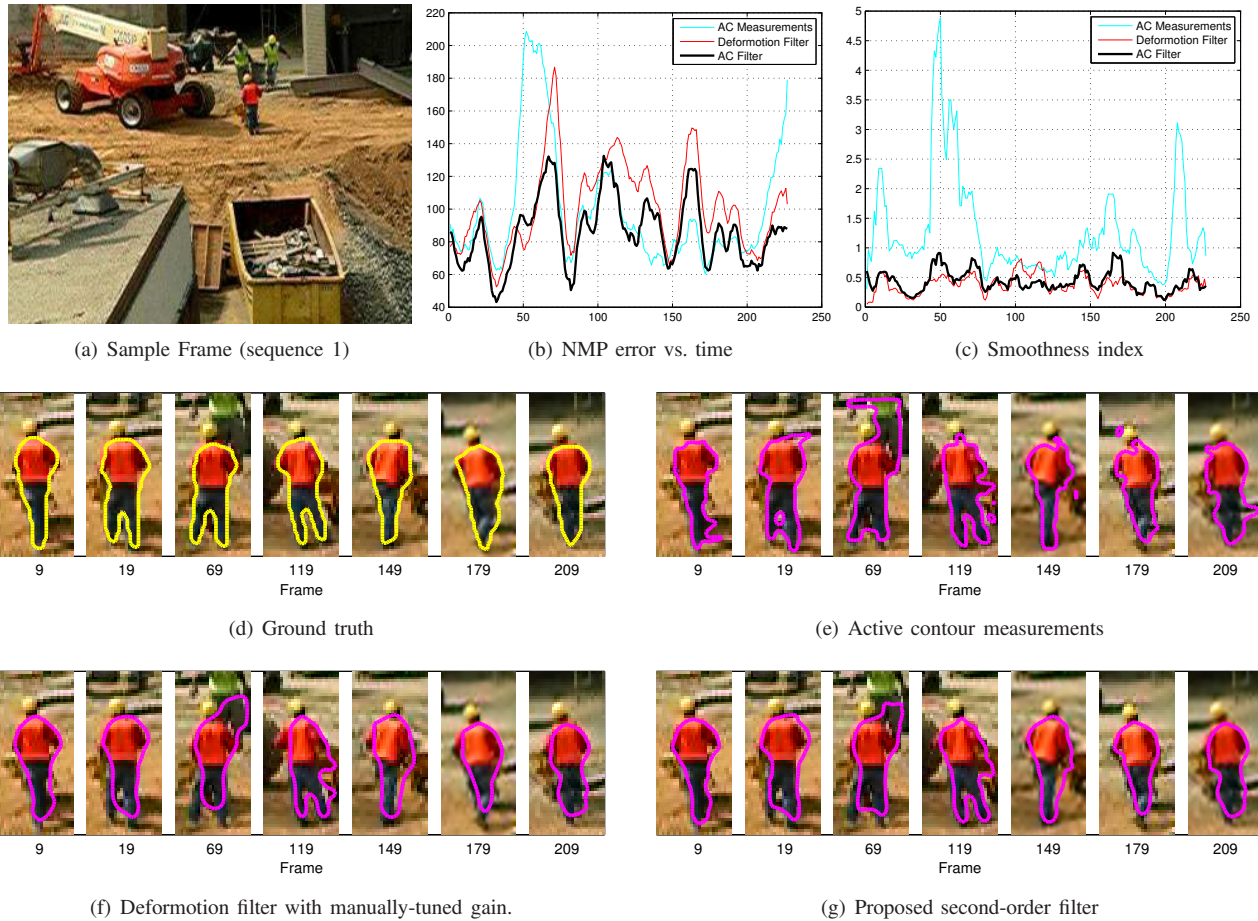


Fig. 2. Quantitative comparison of performance for Sequence 1. The top row depicts a sample frame, the number of misclassified pixels and a smoothness index capturing the temporal consistency of the target. Sample estimates are also displayed for each tracking technique.

small local perturbations and some jittering on a frame-to-frame basis. Due to the shape variations, application of the deformation filter smoothes out the shape perturbations but introduces a lag between estimates and the true shapes. The local second-order filter is able to eliminate those perturbations while adapting much faster to the shape variations thanks to the second-order model involving the curve velocities.

These experiments indicate that the sub-optimal second-order filter is capable of attenuating curve perturbations and maintaining a good fidelity to the true curve being estimated. While temporal consistency of the shape is obtained, it is not achieved with detrimental oversmoothing at the expense of curve fidelity.

V. CONCLUSION

This paper presented the design of a local sub-optimal second-order curve filtering strategy. Following up on previous work, we extend the filter’s state to include both the curve’s position and velocity. The resulting second-order model accounts more accurately for curve velocities, which results in better curve estimates. The filter design is validated by applying the technique to recorded im-

agery and comparing objectively its performance to other tracking techniques using a quantitative error metric. These experiments show that the technique presented here is well suited to estimate the position and deformations of highly deformable curves, especially in the face of noisy measurements as induced by image disturbances. The difficulty associated with the nonlinear dynamics indicate that second-order optimal filtering for dynamic curves may be best handled within the context of curve geodesics by improving upon [19].

a) Acknowledgement.: This work is supported by NSF ECCS#0622006.

REFERENCES

- [1] R. Brockett and A. Blake. Estimating the shape of a moving contour. In *proceedings of the IEEE Conference on Decision and Control*, pages 3247–3251, 1994.
- [2] D. Cremers, F. R. Schmidt, and F. Barthel. Shape priors in variational image segmentation: Convexity, lipschitz continuity and globally optimal solutions. Anchorage, Alaska, June 2008.
- [3] S. Dambreville, Y. Rathi, and A. Tannenbaum. Tracking deformable objects with unscented Kalman filtering and geometric active contours. In *proceedings of the IEEE American Control Conference*, pages 2856–2861, 2006.

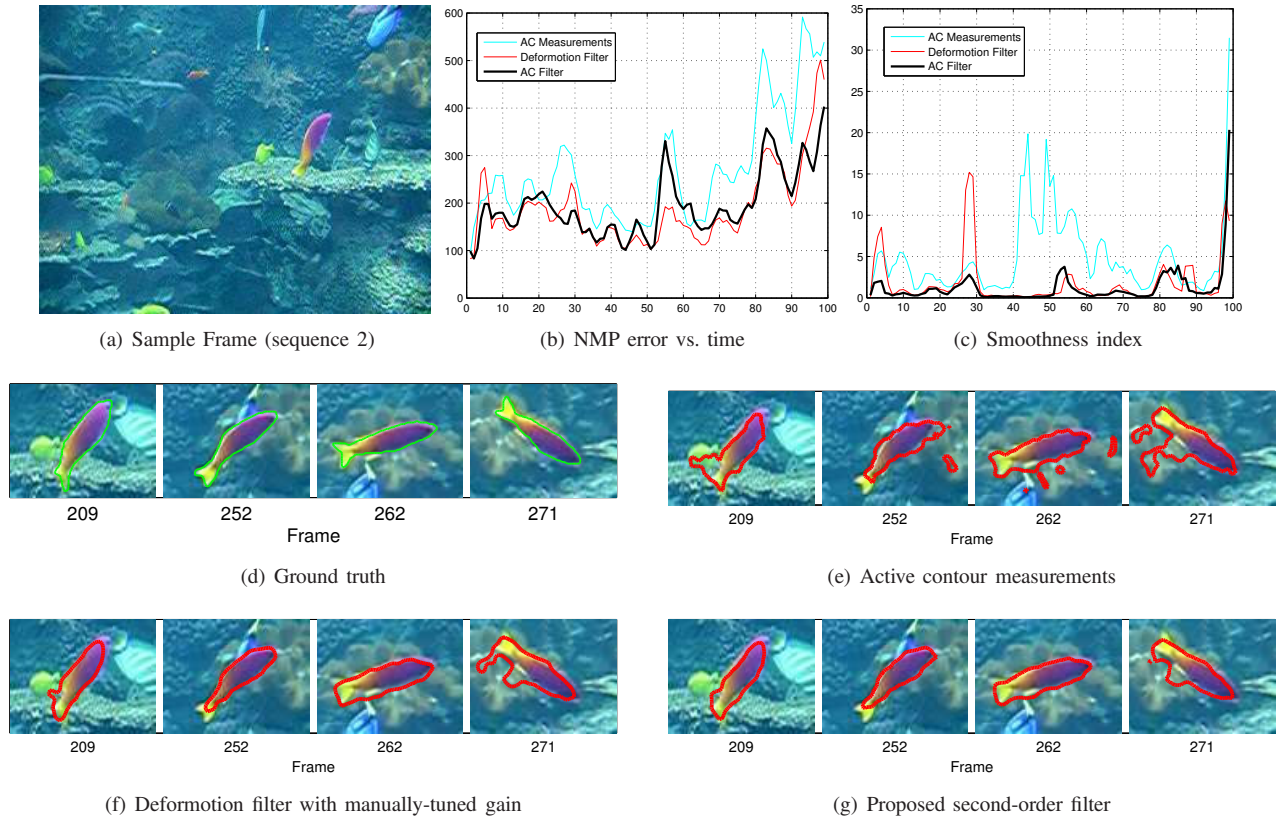



Fig. 3. Quantitative comparison of performance for Sequence 2. The top row depicts a sample frame, the number of misclassified pixels and a smoothness index capturing the temporal consistency of the target. Sample estimates are also displayed for each tracking technique.


- [4] J.D. Jackson, A.J. Yezzi, and S. Soatto. Tracking deformable moving objects under severe occlusions. In *proceedings of the IEEE Conference on Decision and Control*, pages 2990–2995, 2004.
- [5] S.J. Julier, J.K. Uhlmann, and Durrant-Whyte H.F. A new approach for filtering nonlinear systems. In *proceedings of the IEEE American Control Conference*, pages 1628–1632, 1997.
- [6] R.E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82(Series D):35–45, 1960.
- [7] I.J. Ndiour, J. Teizer, and P.A. Vela. A probabilistic contour observer for online visual tracking. *To appear in SIAM Journal on Imaging Sciences*, 2010.
- [8] I.J. Ndiour and P.A. Vela. Towards a local Kalman filter for visual tracking. In *proceedings of the IEEE Conference on Decision and Control*, pages 2420–2426, 2009.
- [9] I.J. Ndiour and P.A. Vela. Optimal estimation applied to visual contour tracking. In *proceedings of the IEEE American Control Conference*, pages 4564–4569, 2010.
- [10] M. Niethammer. *Dynamic Level Sets for Visual Tracking*. PhD thesis, Georgia Institute of Technology, 2004. <http://hdl.handle.net/1853/7606>.
- [11] M. Niethammer, A. Tannenbaum, and S. Angenent. Dynamic active contours for visual tracking. *IEEE Transactions on Automatic Control*, 53(2/2):153–167, 2003.
- [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. Variational optimal control technique for the tracking of deformable objects. In *proceedings of the IEEE International Conference on Computer Vision*, pages 1–7, 2007.
- [14] E. Pichon, D. Nain, and M. Niethammer. A Laplace equation approach for shape comparison. In *proceedings of the 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 adaptive segmentation of vector valued images. In *proceedings of the IEEE Workshop on Motion and Video Computing*, pages 56–61, 2002.
- [17] F. R. Schmidt and D. Cremers. A closed-form solution for image sequence segmentation with dynamical shape priors. In *Symposium of the German Association for Pattern Recognition (DAGM)*, Jena, Germany, September 2009.
- [18] D. Simon. *Optimal State Estimation: Kalman, H-infinity, and Nonlinear Approaches*. Wiley and Sons, 2006.
- [19] G. Sundaramoorthi, A. Mennucci, S. Soatto, and A.J. Yezzi. Tracking deforming objects by filtering and prediction in the space of curves. In *proceedings of the IEEE Conference on Decision and Control*, pages 2395–2401, 2009.
- [20] N. Vaswani, Y. Rathi, A.J. Yezzi, and A. Tannenbaum. Deform PF-MT: Particle filter with mode tracker for tracking non-affine contour deformations. *IEEE Transactions on Image Processing*, 19(4):841–857, 2010.
- [21] N. Vaswani, A. Yezzi, Y. Rathi, and A. Tannenbaum. Time-varying finite dimensional basis for tracking contour deformations. In *proceedings of the IEEE Conference on Decision and Control*, pages 1665–1672, 2006.
- [22] P.A. Vela and I.J. Ndiour. Estimation theory and tracking of deformable objects. *To appear in proceedings of the IEEE Multi-Conference on Systems and Control*, 2010.
- [23] A.J. Yezzi and S. Soatto. Deformation: 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) Sample Frame (sequence 3)




(b) Ground truth



(c) Active contour measurements



(d) Deformation filter with manually-tuned gain



(e) Proposed second-order filter

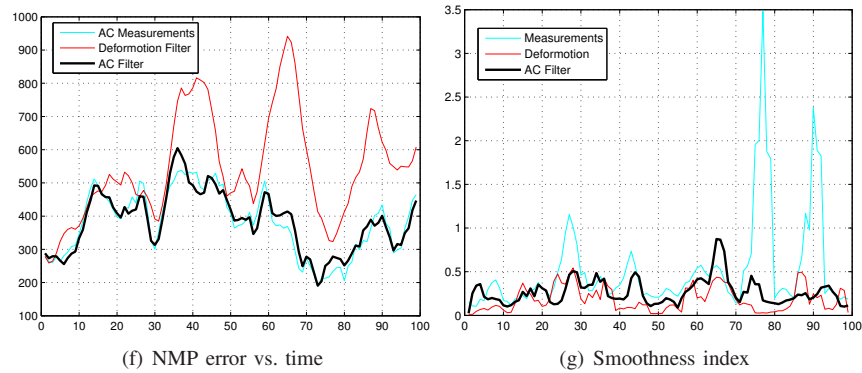


Fig. 4. Quantitative comparison of performance for Sequence 3. Sample estimates are displayed for each tracking technique, along with the number of misclassified pixels and a smoothness index capturing the temporal consistency of the target.