Dynamic Curve Estimation for Visual Tracking

Ibrahima Jacques Ndiour

School of Electrical and Computer Engineering,
Thesis Defense

July 12, 2010

Preamble

- Problem Statement: We consider the problem of accurate online contour-based object tracking in the face of uncertainty caused by imaging noise and approximate/imperfect segmentation models.

- Thesis: The construction of estimators for temporally evolving curves will lead to improved performance for visual tracking systems.

Georgia Tech College of Engineering

Prior Related Work

1. N. Papadakis and E. Mémin. **Variational optimal control technique for the tracking of deformable objects.** ICCV 2007.
2. J.D. Jackson, A.J. Yezzi, and S. Soatto. **Tracking deformable moving objects under severe occlusions.** CDC 2004.
3. D. Cremers. **Dynamical statistical shape priors for level set-based tracking**. TPAMI 2006.
4. S. Dambreville, Y. Rathi and A. Tannenbaum. **Tracking deformable objects with unscented Kalman filtering and geometric active contours.** ACC 2006.
5. M. Niethammer, P.A. Vela, and A. Tannenbaum. **Geometric observers for dynamically evolving curves**. TPAMI 2008.
6. G. Sundaramoorthi, A. Mennucci, S. Soatto, and A.J. Yezzi. **Tracking deforming objects by filtering and prediction in the space of curves.** CDC 2009.
7. N. Vaswani, Y. Rathi, A.J. Yezzi and A. Tannenbaum. **Deform PF-MT: particle filter with mode tracker for tracking non-affine contour deformations.** TIP 2009.

Georgia Tech College of Engineering

Our approach to Visual Tracking

We propose the following observer structure:

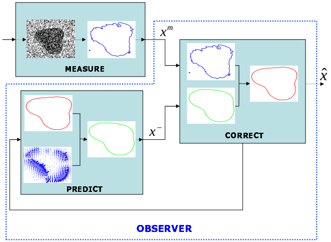

Such observer design presents several benefits:

- may be agnostic to measurement strategies.
- analyzes shape in a non-parametrized setting.
- provides modularity of the algorithmic blocks.
- increases robustness.

State-Space Representations

The state of the target is described by a pose state which represents the ensemble movement and a shape state which represents the local deformations.

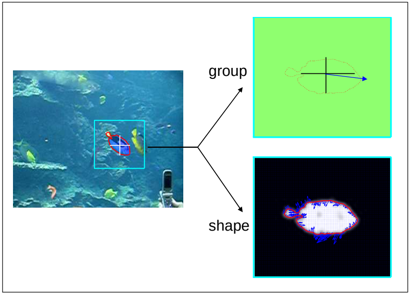

Figure: Group/Shape Decomposition.

Correction Models

- Linear finite-dimensional systems are updated according to:

$$\begin{bmatrix} x \\ \dot{x} \end{bmatrix}^{+} = \begin{bmatrix} x \\ \dot{x} \end{bmatrix}^{-} + \begin{pmatrix} K_{11} & K_{12} \\ K_{21} & K_{22} \end{pmatrix} \cdot \begin{bmatrix} err(x_m, x^-) \\ err(\dot{x}_m, \dot{x}^-) \end{bmatrix} \quad (1)$$

- The group and its velocity are then updated using equation (1).

- The shape variable being infinite-dimensional, standard update schemes are inappropriate.

Georgia Tech College of Engineering

Implicit Shape Representations

Probabilistic Contour Observer

- Propose: observer for online visual contour tracking.
- Contributions:
  1. formulates tracking as decoupled observer designs on group and shape.
  2. incorporates dynamical prediction models for shape space.
  3. defines a novel update method suited to shape descriptor.
  4. quantitatively validates the system's performance.

Dynamical Prediction Models

The prediction model uses the state estimate from the previous frame to produce an estimate at the current frame.

Constant group velocity	$\begin{cases} \dot{g} = \xi, \quad \dot{\xi} = 0 \\ \dot{P} = 0 \end{cases}$
Constant velocity	$\begin{cases} \dot{g} = \xi, & \dot{\xi} = 0 \\ \dot{P} + \nabla P \cdot \Theta = 0, & \dot{\Theta} = 0 \end{cases}$
Advected velocity	$\begin{cases} \dot{g} = \xi, & \dot{\xi} = 0 \\ \dot{P} + \nabla P \cdot \Theta = 0, & \dot{\Theta} + \nabla \Theta \cdot \Theta = 0 \end{cases}$

Measurement Model

Measurement involves the determination of the four substates $(g_m, \xi_m, P_m, \Theta_m)$:

- pose g_m is obtained through target localization.
- shape P_m is given by segmentation.
- shape velocity Θ_m can be measured by computing the optical flow between two subsequent aligned images.
- group velocity ξ_m is not available for measurement.

Correction Models

- Nonlocal shape correction can be performed by defining an error vector field $X_{err}(P^-, P_m)$, whose flow results in a homotopy: $P^+ = \Phi^{K_{11}}_{X_{err}}(P^-)$.

- Correction on the shape can also be achieved through a weighted geometric averaging procedure:

$$P^+(r) = \left(P^-(r)\right)^{1-K_{11}(r)} \cdot \left(P^m(r)\right)^{K_{11}(r)}.$$

- Example of shape interpolation.

Experimental setup

- Conduct experiments with real imagery to assess the observer's performance.
- Compare to other visual tracking techniques.
- Manually segment and determine track point for ground truth.
- Error metrics:
  1. pose: L_2 error.
  2. shape: number of misclassified pixels (NMP), Hausdorff and Sobolev distances.

Georgia Tech College of Engineering

Tracking Experiments: Sample Video 1

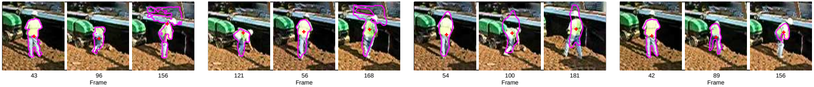

(a) Bayesian (b) Deformotion (c) Shape (d) Observer

Figure: Sample estimates.

Metric \ Algorithm	Bayesian	AC	Deformotion	Shape	Observer
Trackpt error (L_2/L_∞)	16.6 / 24.4	~~11.5 / 52.3~~	7.9 / 16.0	**5.4** / **12.3**	8.0 / 15.5
NMP (med/max)	253 / 1420	~~288 / 1328~~	202 / 755	299 / 536	**171** / **508**
Hausdorff (med/max)	10.2 / 35.0	~~30.0 / ∞~~	7.8 / 26.2	10.9 / **25.8**	**7.7** / 27.4
Sobolev (med/max)	8.2 / 70.6	~~100.0 / ∞~~	**5.8** / **35.3**	11.7 / 38.1	6.5 / 81.8
# Frames tracked	200	~~150~~	200	200	200

Table: Error metrics.

Tracking Experiments: Sample Video 2

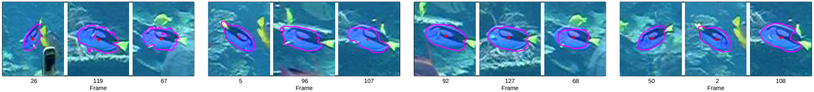

(a) Bayesian (b) AC (c) Deformotion (d) Observer

Figure: Sample estimates.

Metric \ Algorithm	Bayesian	AC	Deformotion	Shape	Observer
Trackpt error (L_2/L_∞)	8.6 / 13.2	2.8 / 7.0	**2.6** / 12.3	5.6 / 15.8	2.7 / **5.8**
NMP (med/max)	251 / 969	**244** / 549	248 / 769	575 / 833	279 / **478**
Hausdorff (med/max)	**10.9** / **18.4**	11.1 / 19.2	12.3 / 19.7	12.0 / 22.5	14.6 / 20.7
Sobolev (med/max)	**8.2** / 52.9	12.9 / 95.8	11.9 / 46.7	13.2 / 43.9	12.9 / **26.9**
# Frames tracked	477	**478**	477	475	**478**

Table: Error metrics.

Concluding Remarks

- Proposed: a recursive dynamic filter for tracking.
- Results:
  1. Achieves temporal consistency
  2. Equal to or more effective than other online, recursive methods.
  3. Does not require training.
  4. Low computational cost.
- Remaining: Optimal gain strategy for the filtering.

Georgia Tech College of Engineering

Optimal Contour Estimation

- Propose: derivation of an optimal contour estimator.
- Contributions:
  1. simplifies infinite-dimensional filtering problem into series of point-wise estimation tasks.
  2. derives an optimal estimator.
  3. quantitatively validates the estimator's performance.

Measurement Strategy: Bayesian Segmentation

- Image is composed of several classes.
- Each class is described by a likelihood distribution.
- Each pixel has an *a priori* probability of being assigned to a given class.
- Classifier selects the most likely class for a pixel through a *maximum a posteriori* approach using Bayes' rule:

$$\Pr(c_i = c | v_i = v) = \frac{\Pr(v_i = v | c_i = c) \Pr(c_i = c)}{\sum_{\gamma} \Pr(v_i = v | c_i = \gamma) \Pr(c_i = \gamma)}.$$

Georgia Tech College of Engineering

Measurement Strategy: Bayesian Segmentation

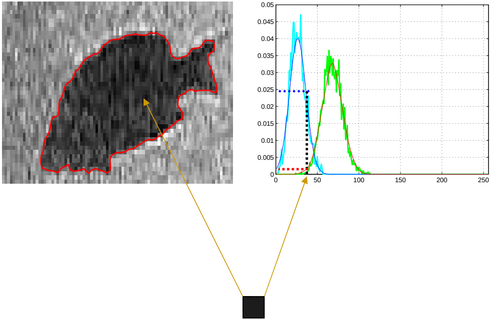

Figure: Likelihood Generation for Bayesian Segmentation

Multiplicative state uncertainty

- Consider an image I taking values in $\mathbb{R}$.
- Pixel intensity corrupted by additive noise ν following $\mathcal{N}(0, \sigma_\nu^2)$.
- Generation of likelihood $\zeta(x, y)$ at pixel $I(x, y)$ yields:

$$\zeta(x, y) = \sqrt{c} \cdot e^{-\frac{1}{2}\left(\frac{I(x,y)+\nu-\mu_F}{\sigma_F}\right)^2}.$$

Multiplicative state uncertainty

- When expanded, the measured classification likelihood is:

$$\zeta(x,y) = \underbrace{\sqrt{c} \cdot e^{-\frac{1}{2}\left(\frac{I(x,y)-\mu_F}{\sigma_F}\right)^2}}_{\rho(x,y)} \cdot \underbrace{e^{-\frac{1}{2}\left(\frac{\nu}{\sigma_F}\right)^2} \cdot e^{-\left(\frac{\nu\,(I(x,y)-\mu_F)}{\sigma_F^2}\right)}}_{\eta(x,y)}$$

- $\rho(x,y)$ is the true classification likelihood.
- $\eta(x,y)$ is the classification measurement noise.
- Additive image noise $\Longrightarrow$ Multiplicative state uncertainty
- Use of an appropriate estimator can resolve multiplicative uncertainty.

Georgia Tech | College of Engineering

The geometric averaging update model

- In the log-space associated to the densities:

$$\log(\zeta) = \log(\rho) + \log(\eta).$$

- Applying a constant gain, linear filtering strategy to filter the noise leads to:

$$\log(\hat{\rho}^{+}) = \log(\hat{\rho}^{-}) + K \left[\log(\zeta) - \log(\hat{\rho}^{-})\right].$$

- Rearranging the terms gives

$$\log(\hat{\rho}^{+}) = (1-K)\log(\hat{\rho}^{-}) + K\log(\zeta).$$

- Finally, we return to the densities by exponentiating:

$$\hat{\rho}^{+} = \left(\hat{\rho}^{-}\right)^{1-K} (\zeta)^{K}.$$

The prediction model

- The prediction step can be chosen to be static (propagation of the previous state estimate) or dynamical, given prior knowledge on the state evolution:

$$\hat{\rho}_t^- = f(\hat{\rho}_{t-1}^+),$$

where f represents the state transition function.

- When no sufficient prior information of the state evolution is known, the generic static prediction model can be used ($f = id$).

Georgia Tech College of Engineering

Noise statistics estimation

- Correlation between the likelihood and the measurement noise:

$$S = E\left(\log\left(\rho\right) \cdot \log\left(\eta\right)\right) = 0.$$

- Since $\frac{I(r)-\mu_F}{\sigma_F}$ and $\nu(r)$ follow normal distributions $\mathcal{N}(0,1)$ and $\mathcal{N}(0,\sigma_\nu^2)$ respectively, the measurement error covariance is derived:

$$\begin{aligned} R &= E\left(\left[\log(\eta)\right]^2\right) \\ &= \frac{1}{2}\left(\frac{\sigma_\nu}{\sigma_F}\right)^4 + \left(\frac{\sigma_\nu}{\sigma_F}\right)^2. \end{aligned}$$

Georgia Tech College of Engineering

Error Covariance Update

- Error variance $\hat{P}_t^+$ defined as $E\left(\left[\log(\rho_t) - \log(\hat{\rho}_t^+)\right]^2\right)$.

- It is a measure of the accuracy of the estimate $\hat{\rho}_t^+$ and needs to be estimated at the prediction step and updated at the correction step:

$$
\begin{aligned}
\hat{P}_t^+ &= E\left(\left[\log(\rho_t) - \log(\hat{\rho}_t^+)\right]^2\right) \\
&= E\left(\left[\log(\rho_t) - (1 - K_t)\log(\hat{\rho}_t{}^-) - K_t \log(\zeta_t)\right]^2\right) \\
&= E\left(\left[(1 - K_t)\left[\log(\rho_t) - \log(\hat{\rho}_t^-)\right] - K_t \log(\eta_t)\right]^2\right) \\
&= (1 - K_t)^2\ \hat{P}_t^- + K_t^2\, R.
\end{aligned}
$$

Georgia Tech College of Engineering

Error Covariance Prediction

- Assume a static prediction model with multiplicative process noise τ, i.e. $(\hat{\rho}_t^- = \hat{\rho}_{t-1}^+ \cdot \tau_t)$.
- Further assume the process noise to be independent from the observation noise.
- The predicted error variance $\hat{P}_t^-$ is then given by:

$$\begin{aligned}\hat{P}_t^- &= E\left(\left[\log(\rho_t) - \log(\hat{\rho}_t^-)\right]^2\right) \\ &= E\left(\left[\log(\rho_{t-1}) - \log(\hat{\rho}_{t-1}^+)\right]^2\right) + E\left(\left[\log(\tau_t)\right]^2\right) \\ &= \hat{P}_{t-1}^+ + Q,\end{aligned}$$

where $Q = E\left(\left[\log(\tau)\right]^2\right)$ represents the process error variance.

Optimal state estimation

- In the log-space associated to the densities, the problem considered is one of point-wise linear filtering for a system facing additive noise.
- Thus, the optimal selection of the gain K_t is given by the Kalman gain: $K_t = \hat{P}_t^-(\hat{P}_t^- + R)^{-1}$.

Table: Filtering equations for the visual tracking system

Prediction	$\begin{cases} \hat{\rho}_t^- = \hat{\rho}_{t-1}^+ \\ \hat{P}_t^- = \hat{P}_{t-1}^+ + Q \end{cases}$
Update	$\begin{cases} K_t = \hat{P}_t^-(\hat{P}_t^- + R)^{-1} \\ \hat{\rho}_t^+ = \left(\hat{\rho}_t^-\right)^{1-K_t} \cdot (\zeta_t)^{K_t} \\ \hat{P}_t^+ = (1-K_t)^2\ \hat{P}_t^- + K_t^2\ R \end{cases}$

Georgia Tech College of Engineering

Algorithm and implementation

The optimal estimation algorithm can be summarized as follows:

- Estimate the additive imaging noise offline.
- For every pixel, run two estimators to filter the foreground and background likelihoods ($\hat{\rho}_F(r)$ and $\hat{\rho}_B(r)$):
  1. obtain predictions with corresponding equations in Table 3.
  2. obtain measurement by performing Bayesian segmentation.
  3. obtain updates with corresponding equations in Table 3.
- The estimated classification probability field is obtained by normalization: $\frac{\hat{\rho}_F}{\hat{\rho}_F+\hat{\rho}_B}$.

Georgia Tech College of Engineering

Extension to vector-valued images

- Similarly to the scalar case, the measured likelihood can be expressed as:

$$\zeta(r) = \overbrace{\sqrt{\Delta} \cdot e^{-\frac{1}{2}(I-\mu_F)^T \Sigma_F^{-1} (I-\mu_F)}}^{\rho} \cdot \underbrace{e^{-\frac{1}{2}(I-\mu_F)^T \Sigma_F^{-1} \nu} \cdot e^{-\frac{1}{2}\nu^T \Sigma_F^{-1} (I-\mu_F)} \cdot e^{-\frac{1}{2}\nu^T \Sigma_F^{-1} \nu}}_{\eta}.$$

- The estimator retains the same structure, with the measurement error variance now given in the multivariate case by:

$$R = \frac{1}{2} Tr \left[\left(\Sigma_\nu \Sigma_F^{-1} \right)^2 \right] + Tr \left[\Sigma_\nu \cdot \Sigma_F^{-1} \right].$$

Georgia Tech College of Engineering

Modeling complex appearance models

The work can be extended to handle complex target and background appearance models by representing them with Gaussian mixture models.

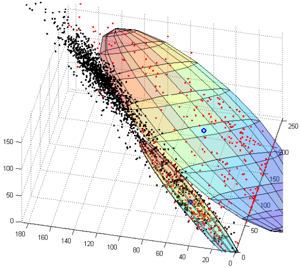

Distributed Filtering for Spatial Consistency

- Move beyond point-wise filtering.
- Assume that a given pixel and its m closest neighbors capture the same visual phenemonon, only from different but close viewpoints.

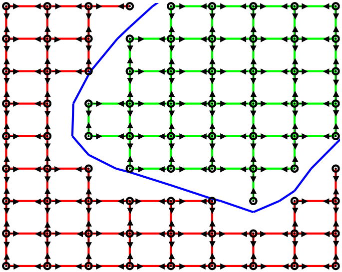

Figure: Network topology for distributed filtering (4-connectivity).

Distributed Filtering for Spatial Consistency

- Filtering equations:

Table: Filtering equations using the information form

Prediction	$\begin{cases} \hat{\rho}_t^- = \hat{\rho}_{t-1}^+ \\ \hat{P}_t^- = \hat{P}_{t-1}^+ + Q \end{cases}$
Update	$\begin{cases} \left(\tilde{P}_t^+\right)^{-1} = R^{-1} + \left(\hat{P}_t^-\right)^{-1} \\ K_t = \tilde{P}_t^+ \cdot R^{-1} \\ \tilde{\rho}_t^+ = \left(\hat{\rho}_t^-\right)^{1-K_t} \cdot (\zeta_t)^{K_t} \end{cases}$

- Assimilation equations:

$$\begin{cases} \left(\hat{P}_{t,i}^+\right)^{-1} = \left(\hat{P}_{t,i}^-\right)^{-1} + \sum_{j=1}^{m} \left[\left(\tilde{P}_{t,j}^-\right)^{-1} - \left(\hat{P}_{t,j}^-\right)^{-1}\right] \\ \hat{\rho}_{t,i}^+ = \left(\hat{\rho}_{t,i}^-\right)^{\hat{P}_{t,i}^+ \cdot \hat{P}_{t,i}^-} \cdot \prod_{j=1}^{m} \frac{\left(\tilde{\rho}_{t,j}\right)^{\hat{P}_{t,i}^+ \cdot \tilde{P}_{t,j}}}{\left(\hat{\rho}_{t,j}^-\right)^{\hat{P}_{t,i}^+ \cdot \hat{P}_{t,j}^-}} \end{cases}$$

Experimental setup

- Conduct experiments with real imagery to assess performance.
- Compare performance against fixed-gain filtering strategies and other visual tracking techniques.
- Ground truth through manual segmentations.
- Error metric given by the NMP.

Georgia Tech College of Engineering

Optimality Experiments: Grayscale

(a) Original

(b) $\sigma_\nu = 25$

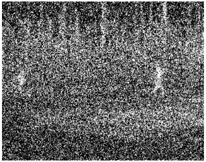

(c) $\sigma_\nu = 100$

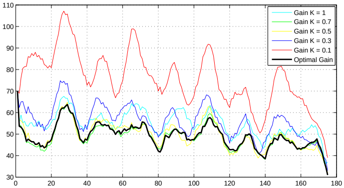

(d) NMP vs time ($\sigma_\nu = 25$)

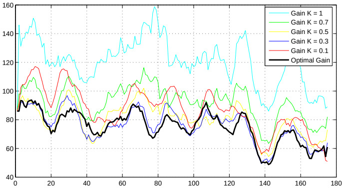

(e) NMP vs time ($\sigma_\nu = 100$)

Optimality Experiments: Color

(f) Original

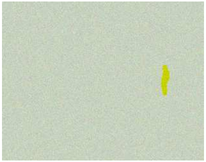

(g) $\Sigma_\nu = 10 \cdot \mathbb{1}$

(h) $\Sigma_\nu = 100 \cdot \mathbb{1}$

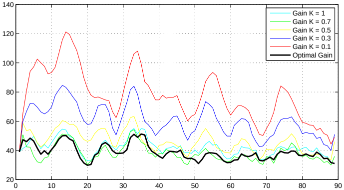

(i) NMP vs time ($\Sigma_\nu = 100 \cdot \mathbb{1}$)

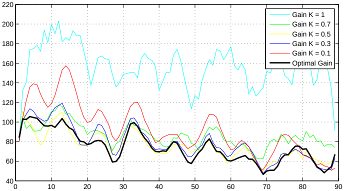

(j) NMP vs time ($\Sigma_\nu = 200 \cdot \mathbb{1}$)

Georgia Tech College of Engineering

Comparative Performance

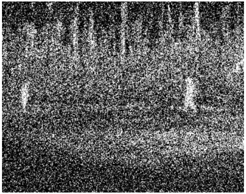

(a) Sample Frame

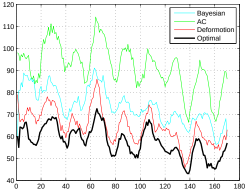

(b) NMP error metric

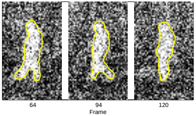

(c) Truth

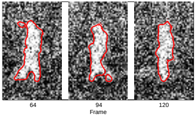

(d) Bayesian

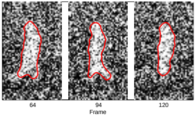

(e) Deformotion

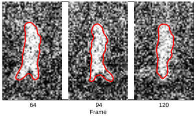

(f) Optimal

Figure: Noisy Synthetic Sequence.

Georgia Tech College of Engineering

Concluding Remarks

- Proposed: an optimal contour estimator.
- Results:
  1. Formally tied optimal gain to **measurable** uncertainty on image data.
  2. Does not require manual gain tuning.
  3. Able to handle severe noise perturbations.
  4. Compares favorably with other tracking methods.

Local Optimal Filters

- Propose: local optimal filters for closed curve filtering.
- Contributions:
    1. introduces a local, linear description for planar curve variation and curve uncertainty.
    2. derives mechanisms for estimating the optimal filtering gain, given quantitative uncertainty levels.
    3. quantitatively validates the filter's performance.
- Shape representation: signed distance function.

Georgia Tech College of Engineering

Transverse curve coordinates

- *Correspondence trajectories:*

  Solve the Laplace equation $\Delta u = 0$ (with boundary conditions) to obtain a harmonic field. The corresponding characteristic vector field is given by $\frac{\nabla u}{||\nabla u||}$.

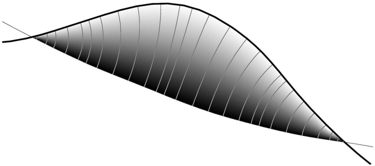

Figure: Characteristic Error vector Field.

Georgia Tech College of Engineering

Transverse curve coordinates

- *Curve coordinate system:*

  Following the distance characteristics starting at a curve point defines the local transverse coordinate system.

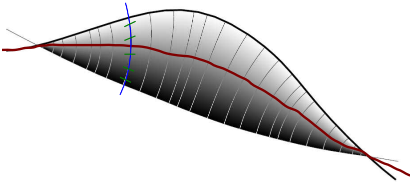

FIGURE: Transverse Coordinates.

First-Order Curve Filtering

- Consider estimates $\hat{\mathcal{C}}$, $\hat{\mathcal{C}}^-$ of the true curve $\mathcal{C}$ and a measurement $\mathcal{C}^m$.
- Assume that measurements are independent from predictions, i.e. $\mathrm{Cov}(\mathcal{C}^{\mathrm{m}}, \hat{\mathcal{C}}^-) = 0$.
- In point notation, the curve errors of the estimates are:

$$\begin{cases} \hat{e}^-(s) & = \hat{x}^-(s) - x(s) \\ \hat{e}(s) & = \hat{x}(s) - x(s) \end{cases}$$

- The variances associated with the errors are:

$$\begin{cases} P^-(s) & = E\left(\left[\hat{x}^-(s) - x(s)\right]^2\right) > 0 \\ P(s) & = E\left(\left[\hat{x}(s) - x(s)\right]^2\right) > 0 \end{cases}$$

First-Order Curve Filtering

- Assume that $R(s)$ varies smoothly along the curve.
- The optimal selection of the gain K minimizes the error covariance $P(s)$ under the update model:

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

- Given the setup, the optimal choice of K is given by the Kalman gain:

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

- The associated error variance is:

$$P^{+} = P^{-}\left(1 - K\right)$$

Georgia Tech College of Engineering

Second-Order Curve Filtering

- First-order filtering strategy: absence of a dynamical prediction model.
- Account for shape dynamics by considering both the curve's position and normal velocity: $\mathbf{x}(s) = [x(s), v(s)]$.
- Filter state also includes the curve covariance matrix $\mathbb{P} : S^1 \to \mathbb{R}^{2x2}$.
- Second-order curve dynamics may be nonlinear.

Dynamical Prediction Models

- Constant velocity model:

$$\begin{cases} \hat{\mathcal{C}}_t = \beta \mathcal{N} \\ \hat{\beta}_t = 0 \end{cases} \iff \begin{cases} \hat{\Psi}_t = \hat{\hat{\beta}} \cdot \left\| \nabla \hat{\Psi} \right\| \\ \hat{\hat{\beta}}_t = 0 \end{cases}$$

- General purpose second-order model:

$$\begin{cases} \hat{\mathcal{C}}_t = \beta \mathcal{N} \\ \hat{\beta}_t = \left(\frac{1}{2} \hat{\beta}^2 + \frac{a}{\mu} \right) \kappa \end{cases} \iff \begin{cases} \hat{\Psi}_t = \hat{\hat{\beta}} \cdot \left\| \nabla \hat{\Psi} \right\| \\ \hat{\hat{\beta}}_t = \left(\frac{1}{2} \hat{\hat{\beta}}^2 + \frac{a}{\mu} \right) \nabla \cdot \left(\frac{\nabla \hat{\Psi}}{\left\| \nabla \hat{\Psi} \right\|} \right) \end{cases}$$

Georgia Tech College of Engineering

Dynamical Prediction Models

Using a first-order linear discrete approximation leads to the covariance update:

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

- Constant velocity model:

$$\mathbb{F} = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix}.$$

- General purpose second-order model:

$$\mathbb{F} \approx \begin{bmatrix} 1 & \Delta t \\ 0 & 1 + \beta\kappa\, \Delta t \end{bmatrix}.$$

Georgia Tech College of Engineering

Update Model

- The optimal correction gain under the update law,

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

is $\mathbb{K} = \mathbb{P}^{-}\left(\mathbb{P}^{-} + \mathbb{R}\right)^{-1}$.

- Decomposing the gain matrix $\mathbb{K}$ as:

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

leads to the position update:

$$\hat{x}^{+} = \hat{x}^{-} + K_{xx} \cdot \left(x^{m} - \hat{x}^{-}\right) + K_{xv} \cdot \left(v^{m} - \hat{v}^{-}\right).$$

Georgia Tech College of Engineering

Update Model

- The velocity is updated according to:

$$\hat{v}^{+} = \hat{v}^{-} + K_{vx} \cdot \left(x^{m} - \hat{x}^{-}\right) + K_{vv} \cdot \left(v^{m} - \hat{v}^{-}\right).$$

- The covariance update is:

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

- Prior to covariance update, predicted and measured covariances need to be transported to the updated curve location where they can be compared.

Static Filtering Experimentation

- First-order filter was applied to a noisy static sequence.
- Then used fixed-gain filtering strategies. With 0.05 gain increments, a gain sweep (from 0.05 to 0.95) verified gain optimality.

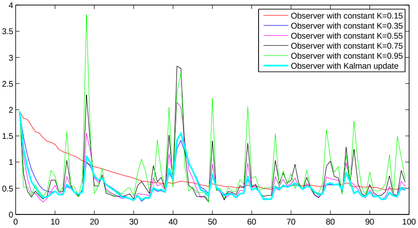

Georgia Tech College of Engineering

Comparison to an actual 1D system

Evolution of the error for a true 1D system and for a simulated static tracking scenario.

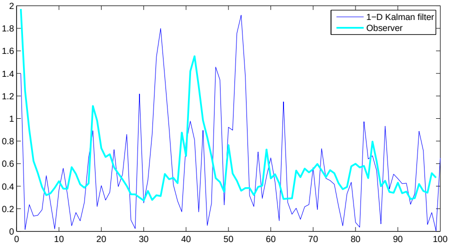

Figure: Error comparison against a true 1D system.

Georgia Tech College of Engineering

Tracking with the first-order filter

Figure: Snapshot

Metric \ Algorithm	AC	Deformotion	Shape	Filter
Group 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**

Figure: Comparative error metrics.

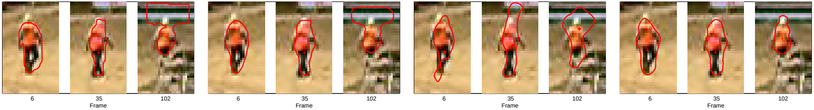

(a) ~~AC~~ (b) ~~Deformotion~~ (c) ~~Shape~~ (d) Local Filter

Figure: Sample estimates.

Georgia Tech College of Engineering

Tracking with the second-order filter

(a) Snapshot

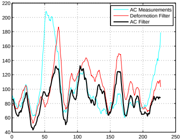

(b) NMP

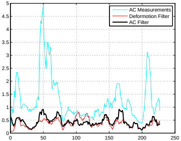

(c) Smoothness

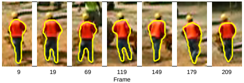

(d) Ground truth

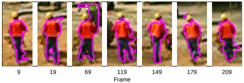

(e) Active contour

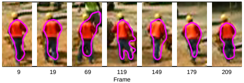

(f) Deformotion filter

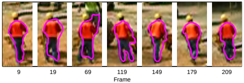

(g) Local Filter

Georgia Tech College of Engineering

Tracking with the second-order filter

(a) Snapshot

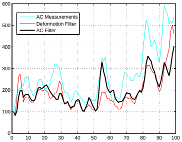

(b) NMP

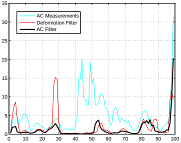

(c) Smoothness

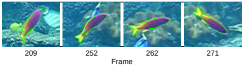

(d) Ground truth

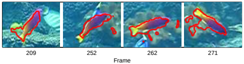

(e) Active contour

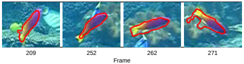

(f) Deformotion filter

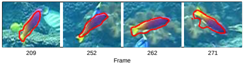

(g) Local Filter

Georgia Tech College of Engineering

Concluding Remarks

- Proposed: locally optimal curve filters.
- Results:
  1. Provides a set of linear coordinate frames from which to perform curve operations.
  2. Incorporates dynamical models to deal with both rigid-body and elastic objects.
  3. Validates design with visual tracking experiments.
  4. Able to estimate curve deformations in presence of image disturbances and imperfect segmentation models.

Georgia Tech College of Engineering

Geometric Avering for Statistical Methods

- Probability fields or confidence maps often used in computer vision, machine learning, and signal processing.
- Even when not naturally defined, it is relatively simple to generate one from existing similarity/distance maps.
- Performance is intimately linked to the SNR of the map.
- Application of a filtering procedure should then improve the overall performance.

Georgia Tech College of Engineering

Geometric Averaging for Ensemble Tracking

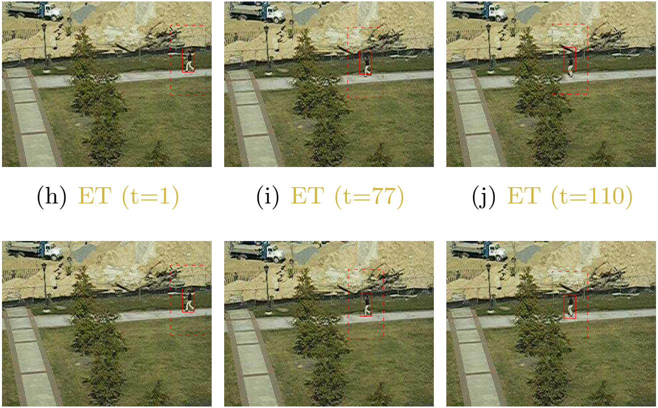

(h) ET (t=1) (i) ET (t=77) (j) ET (t=110)

(k) Filtered ET (l) Filtered ET (m) Filtered ET

Conclusion

- Explored filtering schemes for dynamic curve estimation.
- Developed optimal curve estimation strategies for different state-space representations:
    1. probabilistic shape descriptor.
    2. level set descriptor.
- Validated objectively the work using:
    1. recorded imagery.
    2. ground truth.
    3. relevant error metrics.
- Provided an effective class of solutions to the visual tracking problem.

Georgia Tech College of Engineering

List of related publications

- Refereed Journal Publications:
  1. Ibrahima J. Ndiour, Jochen Teizer, and Patricio A. Vela. **A Probabilistic Contour Observer for Online Visual Tracking**. To appear in SIAM Journal on Imaging Sciences, 2010.

- Selected Conference Publications:
  1. Ibrahima J. Ndiour and Patricio A. Vela. **A Local Extended Kalman Filter for Visual Tracking**. CDC 2010.
  2. Ibrahima J. Ndiour and Patricio A. Vela. **Optimal Estimation Applied to Visual Contour Tracking**. ACC 2010.
  3. Ibrahima J. Ndiour, Omar Arif, Jochen Teizer, and Patricio A. Vela. **A Probabilistic Observer for Visual Tracking**. ACC 2010.
  4. Patricio A. Vela and Ibrahima J. Ndiour. **Estimation Theory and Tracking of Deformable Objects**. MSC 2010.
  5. Ibrahima J. Ndiour and Patricio A. Vela. **Towards a Local Kalman Filter for Visual Tracking**. CDC 2009.
  6. Ibrahima J. Ndiour, Omar Arif, Jochen Teizer, and Patricio A. Vela. **A Probabilistic Shape Filter for Contour Tracking**. ICIP 2009.
  7. Ibrahima J. Ndiour and Patricio A. Vela. **Noise Estimation and Adaptive Filtering During Visual Tracking**. ICIP 2009.

Potential research directions

- Robustify the estimators.
- Investigate methods to accurately model the uncertainty.
- Study the balance between shape constraints and filtering schemes.
- Extend the ideas presented here beyond contour-based tracking methods to other statistical methods.

Thank you for your attention.

Questions?

Georgia Tech College of Engineering