

# C280, Computer Vision

Prof. Trevor Darrell

[trevor@eecs.berkeley.edu](mailto:trevor@eecs.berkeley.edu)

Lecture 19: Tracking

# Tracking scenarios

- Follow a point
- Follow a template
- Follow a changing template
- Follow all the elements of a moving person, fit a model to it.

# Things to consider in tracking

What are the dynamics of the thing being tracked?

How is it observed?

# Three main issues in tracking

- **Prediction:** we have seen  $\mathbf{y}_0, \dots, \mathbf{y}_{i-1}$  — what state does this set of measurements predict for the  $i$ 'th frame? to solve this problem, we need to obtain a representation of  $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_{i-1} = \mathbf{y}_{i-1})$ .
- **Data association:** Some of the measurements obtained from the  $i$ -th frame may tell us about the object's state. Typically, we use  $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_{i-1} = \mathbf{y}_{i-1})$  to identify these measurements.
- **Correction:** now that we have  $\mathbf{y}_i$  — the relevant measurements — we need to compute a representation of  $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_i = \mathbf{y}_i)$ .

# Simplifying Assumptions

- **Only the immediate past matters:** formally, we require

$$P(\mathbf{X}_i | \mathbf{X}_1, \dots, \mathbf{X}_{i-1}) = P(\mathbf{X}_i | \mathbf{X}_{i-1})$$


This assumption hugely simplifies the design of algorithms, as we shall see; furthermore, it isn't terribly restrictive if we're clever about interpreting  $\mathbf{X}_i$  as we shall show in the next section.

- **Measurements depend only on the current state:** we assume that  $\mathbf{Y}_i$  is conditionally independent of all other measurements given  $\mathbf{X}_i$ . This means that

$$P(\mathbf{Y}_i, \mathbf{Y}_j, \dots, \mathbf{Y}_k | \mathbf{X}_i) = P(\mathbf{Y}_i | \mathbf{X}_i) P(\mathbf{Y}_j, \dots, \mathbf{Y}_k | \mathbf{X}_i)$$

Again, this isn't a particularly restrictive or controversial assumption, but it yields important simplifications.

# Kalman filter graphical model and corresponding factorized joint probability



$$P(x_1, x_2, x_3, y_1, y_2, y_3) =$$

$$P(x_1)P(y_1 | x_1)P(x_2 | x_1)P(y_2 | x_2)P(x_3 | x_2)P(y_3 | x_3)$$

# Tracking as induction

- Make a measurement starting in the 0<sup>th</sup> frame
- Then: assume you have an estimate at the  $i$ th frame, after the measurement step.
- Show that you can do prediction for the  $i+1$ th frame, and measurement for the  $i+1$ th frame.

# Base case

Firstly, we assume that we have  $P(\mathbf{X}_0)$

$$P(\mathbf{X}_0 | \mathbf{Y}_0 = \mathbf{y}_0) = \frac{P(\mathbf{y}_0 | \mathbf{X}_0) P(\mathbf{X}_0)}{P(\mathbf{y}_0)}$$

$$\propto P(\mathbf{y}_0 | \mathbf{X}_0) P(\mathbf{X}_0)$$



# Prediction step

## Prediction

Prediction involves representing

$$P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1})$$

given

$$P(\mathbf{X}_{i-1} | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}).$$

Our independence assumptions make it possible to write

$$\begin{aligned} P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) &= \int P(\mathbf{X}_i, \mathbf{X}_{i-1} | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) d\mathbf{X}_{i-1} \\ &= \int P(\mathbf{X}_i | \mathbf{X}_{i-1}, \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) P(\mathbf{X}_{i-1} | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) d\mathbf{X}_{i-1} \\ &= \int P(\mathbf{X}_i | \mathbf{X}_{i-1}) P(\mathbf{X}_{i-1} | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) d\mathbf{X}_{i-1} \end{aligned}$$

# Update step

## Correction

Correction involves obtaining a representation of

$$P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_i)$$

given

$$P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1})$$

Our independence assumptions make it possible to write

$$\begin{aligned} P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_i) &= \frac{P(\mathbf{X}_i, \mathbf{y}_0, \dots, \mathbf{y}_i)}{P(\mathbf{y}_0, \dots, \mathbf{y}_i)} \\ &= \frac{P(\mathbf{y}_i | \mathbf{X}_i, \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) P(\mathbf{y}_0, \dots, \mathbf{y}_{i-1})}{P(\mathbf{y}_0, \dots, \mathbf{y}_i)} \\ &= P(\mathbf{y}_i | \mathbf{X}_i) P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) \frac{P(\mathbf{y}_0, \dots, \mathbf{y}_{i-1})}{P(\mathbf{y}_0, \dots, \mathbf{y}_i)} \\ &= \frac{P(\mathbf{y}_i | \mathbf{X}_i) P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1})}{\int P(\mathbf{y}_i | \mathbf{X}_i) P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) d\mathbf{X}_i} \end{aligned}$$

# The Kalman Filter

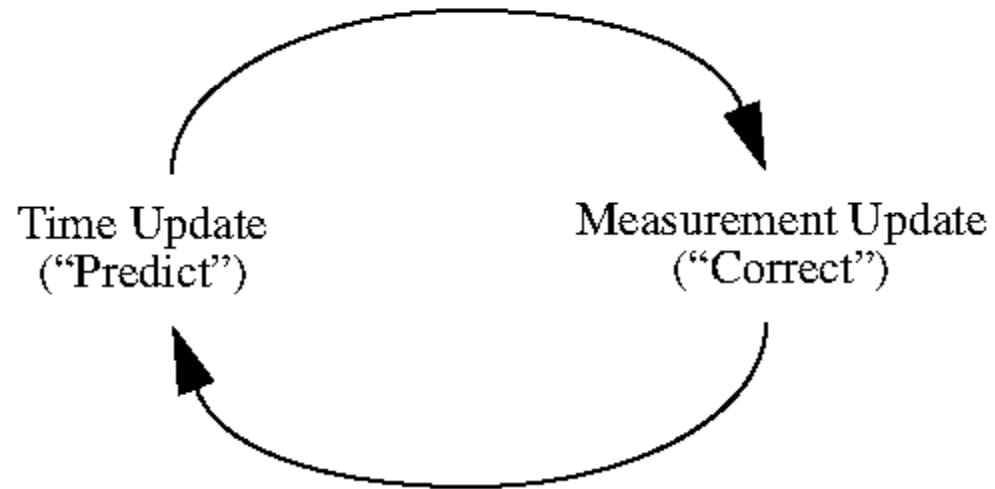
- Key ideas:
  - Linear models interact uniquely well with Gaussian noise - make the prior Gaussian, everything else Gaussian and the calculations are easy
  - Gaussians are really easy to represent --- once you know the mean and covariance, you're done

# Recall the three main issues in tracking

- **Prediction:** we have seen  $\mathbf{y}_0, \dots, \mathbf{y}_{i-1}$  — what state does this set of measurements predict for the  $i$ 'th frame? to solve this problem, we need to obtain a representation of  $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_{i-1} = \mathbf{y}_{i-1})$ .
- **Data association:** Some of the measurements obtained from the  $i$ -th frame may tell us about the object's state. Typically, we use  $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_{i-1} = \mathbf{y}_{i-1})$  to identify these measurements.
- **Correction:** now that we have  $\mathbf{y}_i$  — the relevant measurements — we need to compute a representation of  $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_i = \mathbf{y}_i)$ .

*(Ignore data association for now)*

# The Kalman Filter



# The Kalman Filter in 1D

- Dynamic Model

$$x_i \sim N(d_i x_{i-1}, \sigma_{d_i}^2)$$

$$y_i \sim N(m_i x_i, \sigma_{m_i}^2)$$

- Notation


mean of  $P(X_i | y_0, \dots, y_{i-1})$  as  $\bar{X}_i^-$  ← Predicted mean

mean of  $P(X_i | y_0, \dots, y_i)$  as  $\bar{X}_i^+$  ← Corrected mean

the standard deviation of  $P(X_i | y_0, \dots, y_{i-1})$  as  $\sigma_i^-$

of  $P(X_i | y_0, \dots, y_i)$  as  $\sigma_i^+$ .

# The Kalman Filter



# Prediction for 1D Kalman filter

- The new state is obtained by
  - multiplying old state by known constant
  - adding zero-mean noise
- Therefore, predicted mean for new state is
  - constant times mean for old state
- Old variance is normal random variable
  - variance is multiplied by square of constant
  - and variance of noise is added.

$$x_i \sim N(d_i x_{i-1}, \sigma_{d_i}^2)$$

$$\bar{X}_i^- = d_i \bar{X}_{i-1}^+$$

$$(\sigma_i^-)^2 = \sigma_{d_i}^2 + (d_i \sigma_{i-1}^+)^2$$



Dynamic Model:

$$x_i \sim N(d_i x_{i-1}, \sigma_{d_i})$$

$$y_i \sim N(m_i x_i, \sigma_{m_i})$$


Start Assumptions:  $\bar{x}_0^-$  and  $\sigma_0^-$  are known

Update Equations: Prediction

$$\bar{x}_i^- = d_i \bar{x}_{i-1}^+$$

$$\sigma_i^- = \sqrt{\sigma_{d_i}^2 + (d_i \sigma_{i-1}^+)^2}$$

# The Kalman Filter



# Measurement update for 1D Kalman filter

$$\mathbf{x}_i^+ = \left( \frac{\bar{\mathbf{x}}_i^- \sigma_{m_i}^2 + m_i y_i (\sigma_i^-)^2}{\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2} \right)$$

$$\sigma_i^+ = \sqrt{\left( \frac{\sigma_{m_i}^2 (\sigma_i^-)^2}{(\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2)} \right)}$$

Notice:

- if measurement noise is small,  
we rely mainly on the measurement,
- if it's large, mainly on the prediction
- $\sigma$  does not depend on  $y$

Dynamic Model:

$$x_i \sim N(d_i x_{i-1}, \sigma_{d_i})$$

$$y_i \sim N(m_i x_i, \sigma_{m_i})$$

Start Assumptions:  $\bar{x}_0^-$  and  $\sigma_0^-$  are known

Update Equations: Prediction


$$\bar{x}_i^- = d_i \bar{x}_{i-1}^+$$

$$\sigma_i^- = \sqrt{\sigma_{d_i}^2 + (d_i \sigma_{i-1}^+)^2}$$

Update Equations: Correction

$$x_i^+ = \left( \frac{\bar{x}_i^- \sigma_{m_i}^2 + m_i y_i (\sigma_i^-)^2}{\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2} \right)$$

$$\sigma_i^+ = \sqrt{\left( \frac{\sigma_{m_i}^2 (\sigma_i^-)^2}{(\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2)} \right)}$$



# Kalman filter for computing an on-line average

- What Kalman filter parameters and initial conditions should we pick so that the optimal estimate for  $x$  at each iteration is just the average of all the observations seen so far?

## Kalman filter model

$$x_i \sim N(d_i x_{i-1}, \sigma_{d_i})$$

$$y_i \sim N(m_i x_i, \sigma_{m_i})$$

$$d_i = 1, m_i = 1, \sigma_{d_i} = 0, \sigma_{m_i} = 1$$

### Initial conditions

$$\bar{x}_0^- = 0 \quad \sigma_0^- = \infty$$

$\bar{x}_0^-$  and  $\sigma_0^-$  are known  
prediction

Iteration	0	1	2
$\bar{x}_i^- = d_i \bar{x}_{i-1}^+$	0	$y_0$	$\frac{y_0 + y_1}{2}$
$\sigma_i^- = \sqrt{\sigma_{d_i}^2 + (d_i \sigma_{i-1}^+)^2}$	$\bar{x}_i^-$	$y_0$	$\frac{y_0 + y_1}{2}$
Correction	$\bar{x}_i^+$	$y_0$	$\frac{y_0 + y_1 + y_2}{3}$
$x_i^+ = \left( \frac{\bar{x}_i^- \sigma_{m_i}^2 + m_i y_i (\sigma_i^-)^2}{\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2} \right)$	$\sigma_i^-$	$\infty$	$\frac{1}{\sqrt{2}}$
$\sigma_i^+ = \sqrt{\left( \frac{\sigma_{m_i}^2 (\sigma_i^-)^2}{(\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2)} \right)}$	$\sigma_i^+$	1	$\frac{1}{\sqrt{3}}$

What happens if the  $x$  dynamics are given a non-zero variance?

## Kalman filter model

$$x_i \sim N(d_i x_{i-1}, \sigma_{d_i})$$

$$y_i \sim N(m_i x_i, \sigma_{m_i})$$

$$d_i = 1, m_i = 1, \sigma_{d_i} = 1, \sigma_{m_i} = 1$$

### Initial conditions

$$\bar{x}_0^- = 0 \quad \sigma_0^- = \infty$$

$\bar{x}_0^-$  and  $\sigma_0^-$  are known  
prediction

	Iteration	0	1	2
$\bar{x}_i^- = d_i \bar{x}_{i-1}^+$				
$\sigma_i^- = \sqrt{\sigma_{d_i}^2 + (d_i \sigma_{i-1}^+)^2}$	$\bar{x}_i^-$	0	$y_0$	$\frac{y_0 + 2y_1}{3}$
Correction	$\bar{x}_i^+$	$y_0$	$\frac{y_0 + 2y_1}{3}$	$\frac{y_0 + 2y_1 + 5y_2}{8}$
$x_i^+ = \left( \frac{\bar{x}_i^- \sigma_{m_i}^2 + m_i y_i (\sigma_i^-)^2}{\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2} \right)$	$\sigma_i^-$	$\infty$	$\sqrt{2}$	$\sqrt{\frac{5}{3}}$
$\sigma_i^+ = \sqrt{\left( \frac{\sigma_{m_i}^2 (\sigma_i^-)^2}{(\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2)} \right)}$	$\sigma_i^+$	1	$\sqrt{\frac{2}{3}}$	$\sqrt{\frac{5}{8}}$



# Linear dynamic models

- A linear dynamic model has the form

$$\mathbf{x}_i = N(\mathbf{D}_{i-1}\mathbf{x}_{i-1}; \Sigma_{d_i})$$

$$\mathbf{y}_i = N(\mathbf{M}_i\mathbf{x}_i; \Sigma_{m_i})$$

- This is much, much more general than it looks, and extremely powerful



# Examples of linear state space models

$$\mathbf{x}_i = N(\mathbf{D}_{i-1}\mathbf{x}_{i-1}; \Sigma_{d_i})$$

$$\mathbf{y}_i = N(\mathbf{M}_i\mathbf{x}_i; \Sigma_{m_i})$$

- Drifting points
  - assume that the new position of the point is the old one, plus noise

**D = Identity**



# Constant velocity

$$\mathbf{x}_i = N(\mathbf{D}_{i-1}\mathbf{x}_{i-1}; \Sigma_{d_i})$$

$$\mathbf{y}_i = N(\mathbf{M}_i\mathbf{x}_i; \Sigma_{m_i})$$

- We have

$$u_i = u_{i-1} + \Delta t v_{i-1} + \varepsilon_i$$

$$v_i = v_{i-1} + \zeta_i$$



– (the Greek letters denote noise terms)

- Stack (u, v) into a single state vector


$$\begin{pmatrix} u \\ v \end{pmatrix}_i = \begin{pmatrix} 1 & \Delta t \\ 0 & 1 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix}_{i-1} + \text{noise}$$

– which is the form we had above

$\mathbf{x}_i$                        $\mathbf{D}_{i-1}$                        $\mathbf{x}_{i-1}$



Constant  
Velocity  
Model



# Constant acceleration

$$\mathbf{x}_i = N(\mathbf{D}_{i-1}\mathbf{x}_{i-1}; \Sigma_{d_i})$$

$$\mathbf{y}_i = N(\mathbf{M}_i\mathbf{x}_i; \Sigma_{m_i})$$

- We have

$$u_i = u_{i-1} + \Delta t v_{i-1} + \varepsilon_i$$

$$v_i = v_{i-1} + \Delta t a_{i-1} + \zeta_i$$

$$a_i = a_{i-1} + \xi_i$$

– (the Greek letters denote noise terms)

- Stack (u, v) into a single state vector



$$\begin{pmatrix} u \\ v \\ a \end{pmatrix}_i = \begin{pmatrix} 1 & \Delta t & 0 \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} u \\ v \\ a \end{pmatrix}_{i-1} + \text{noise}$$

– which is the form we had above

$\mathbf{x}_i$

$\mathbf{D}_{i-1}$

$\mathbf{x}_{i-1}$



Constant  
Acceleration  
Model

## Periodic motion

$$\mathbf{x}_i = N(\mathbf{D}_{i-1}\mathbf{x}_{i-1}; \Sigma_{d_i})$$

$$\mathbf{y}_i = N(\mathbf{M}_i\mathbf{x}_i; \Sigma_{m_i})$$

Assume we have a point, moving on a line with a periodic movement defined with a differential eq:

$$\frac{d^2p}{dt^2} = -p$$

can be defined as

$$\frac{du}{dt} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} u = Su$$

with state defined as stacked position and velocity  $u=(p, v)$

# Periodic motion

$$\mathbf{x}_i = N(\mathbf{D}_{i-1}\mathbf{x}_{i-1}; \Sigma_{d_i})$$

$$\mathbf{y}_i = N(\mathbf{M}_i\mathbf{x}_i; \Sigma_{m_i})$$

$$\frac{du}{dt} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} u = Su$$

Take discrete approximation....(e.g., forward Euler integration with  $\Delta t$  stepsize.)

$$\begin{aligned} \mathbf{u}_i &= \mathbf{u}_{i-1} + \Delta t \frac{d\mathbf{u}}{dt} \\ &= \mathbf{u}_{i-1} + \Delta t S \mathbf{u}_{i-1} \\ &= \begin{pmatrix} 1 & \Delta t \\ -\Delta t & 1 \end{pmatrix} \mathbf{u}_{i-1} \end{aligned}$$

$\uparrow$   $\mathbf{x}_i$                        $\uparrow$   $\mathbf{D}_{i-1}$                        $\uparrow$   $\mathbf{x}_{i-1}$




# n-D

Generalization to n-D is straightforward but more complex.


# n-D

Generalization to n-D is straightforward but more complex.



# n-D Prediction

Generalization to n-D is straightforward but more complex.



Prediction:

- Multiply estimate at prior time with forward model:


$$\bar{\mathbf{x}}_i^- = \mathcal{D}_i \bar{\mathbf{x}}_{i-1}^+$$

- Propagate covariance through model and add new noise:

$$\Sigma_i^- = \Sigma_{d_i} + \mathcal{D}_i \sigma_{i-1}^+ \mathcal{D}_i$$

# n-D Correction

Generalization to n-D is straightforward but more complex.



Correction:

- Update *a priori* estimate with measurement to form *a posteriori*

# n-D correction

Find linear filter on innovations

$$\bar{\mathbf{x}}_i^+ = \bar{\mathbf{x}}_i^- + \mathcal{K}_i [\mathbf{y}_i - \mathcal{M}_i \bar{\mathbf{x}}_i^-]$$

which minimizes *a posteriori* error covariance:

$$E \left[ \left( x - \bar{x}^+ \right)^T \left( x - \bar{x}^+ \right) \right]$$

$\mathbf{K}$  is the *Kalman Gain* matrix. A solution is

$$\mathcal{K}_i = \Sigma_i^- \mathcal{M}_i^T [\mathcal{M}_i \Sigma_i^- \mathcal{M}_i^T + \Sigma_{m_i}]^{-1}$$

# Kalman Gain Matrix

$$\bar{\mathbf{x}}_i^+ = \bar{\mathbf{x}}_i^- + \mathcal{K}_i [\mathbf{y}_i - \mathcal{M}_i \bar{\mathbf{x}}_i^-]$$

$$\mathcal{K}_i = \Sigma_i^- \mathcal{M}_i^T [\mathcal{M}_i \Sigma_i^- \mathcal{M}_i^T + \Sigma_{m_i}]^{-1}$$

As measurement becomes more reliable,  $\mathbf{K}$  weights residual more heavily,

$$\lim_{\Sigma_m \rightarrow 0} \mathbf{K}_i = \mathbf{M}^{-1}$$

As prior covariance approaches 0, measurements are ignored:

$$\lim_{\Sigma_i^- \rightarrow 0} \mathbf{K}_i = 0$$

Dynamic Model:

$$\mathbf{x}_i \sim N(\mathcal{D}_i \mathbf{x}_{i-1}, \Sigma_{d_i})$$

$$\mathbf{y}_i \sim N(\mathcal{M}_i \mathbf{x}_i, \Sigma_{m_i})$$

Start Assumptions:  $\bar{\mathbf{x}}_0^-$  and  $\Sigma_0^-$  are known

Update Equations: Prediction

$$\bar{\mathbf{x}}_i^- = \mathcal{D}_i \bar{\mathbf{x}}_{i-1}^+$$



$$\Sigma_i^- = \Sigma_{d_i} + \mathcal{D}_i \Sigma_{i-1}^+ \mathcal{D}_i$$

Update Equations: Correction

$$\mathcal{K}_i = \Sigma_i^- \mathcal{M}_i^T [\mathcal{M}_i \Sigma_i^- \mathcal{M}_i^T + \Sigma_{m_i}]^{-1}$$


$$\bar{\mathbf{x}}_i^+ = \bar{\mathbf{x}}_i^- + \mathcal{K}_i [\mathbf{y}_i - \mathcal{M}_i \bar{\mathbf{x}}_i^-]$$

$$\Sigma_i^+ = [Id - \mathcal{K}_i \mathcal{M}_i] \Sigma_i^-$$




## Constant Velocity Model






This is figure 17.3 of Forsyth and Ponce. The notation is a bit involved, but is logical. We plot the true state as open circles, measurements as x's, predicted means as \*'s with three standard deviation bars, corrected means as +'s with three standard deviation bars.



The \*-s give  $\bar{x}_i^-$ , +-s give  $\bar{x}_i^+$ , vertical bars are 3 standard deviation bars




The o-s give state, x-s measurement.

The \*-s give  $\bar{x}_i^-$ , +s give  $\bar{x}_i^+$ , vertical bars are 3 standard deviation bars

# Smoothing

- Idea
  - We don't have the best estimate of state - what about the future?
  - Run two filters, one moving forward, the other backward in time.
  - Now combine state estimates
    - The crucial point here is that we can obtain a smoothed estimate by viewing the backward filter's prediction as yet another measurement for the forward filter


# Forward estimates.



The o-s give state, x-s measurement.

The \*-s give  $\bar{x}_i^-$ , +-s give  $\bar{x}_i^+$ , vertical bars are 3 standard deviation bars


# Backward estimates.



The o-s give state, x-s measurement. time

The \*-s give  $\bar{x}_i^-$ , +-s give  $\bar{x}_i^+$ , vertical bars are 3 standard deviation bars

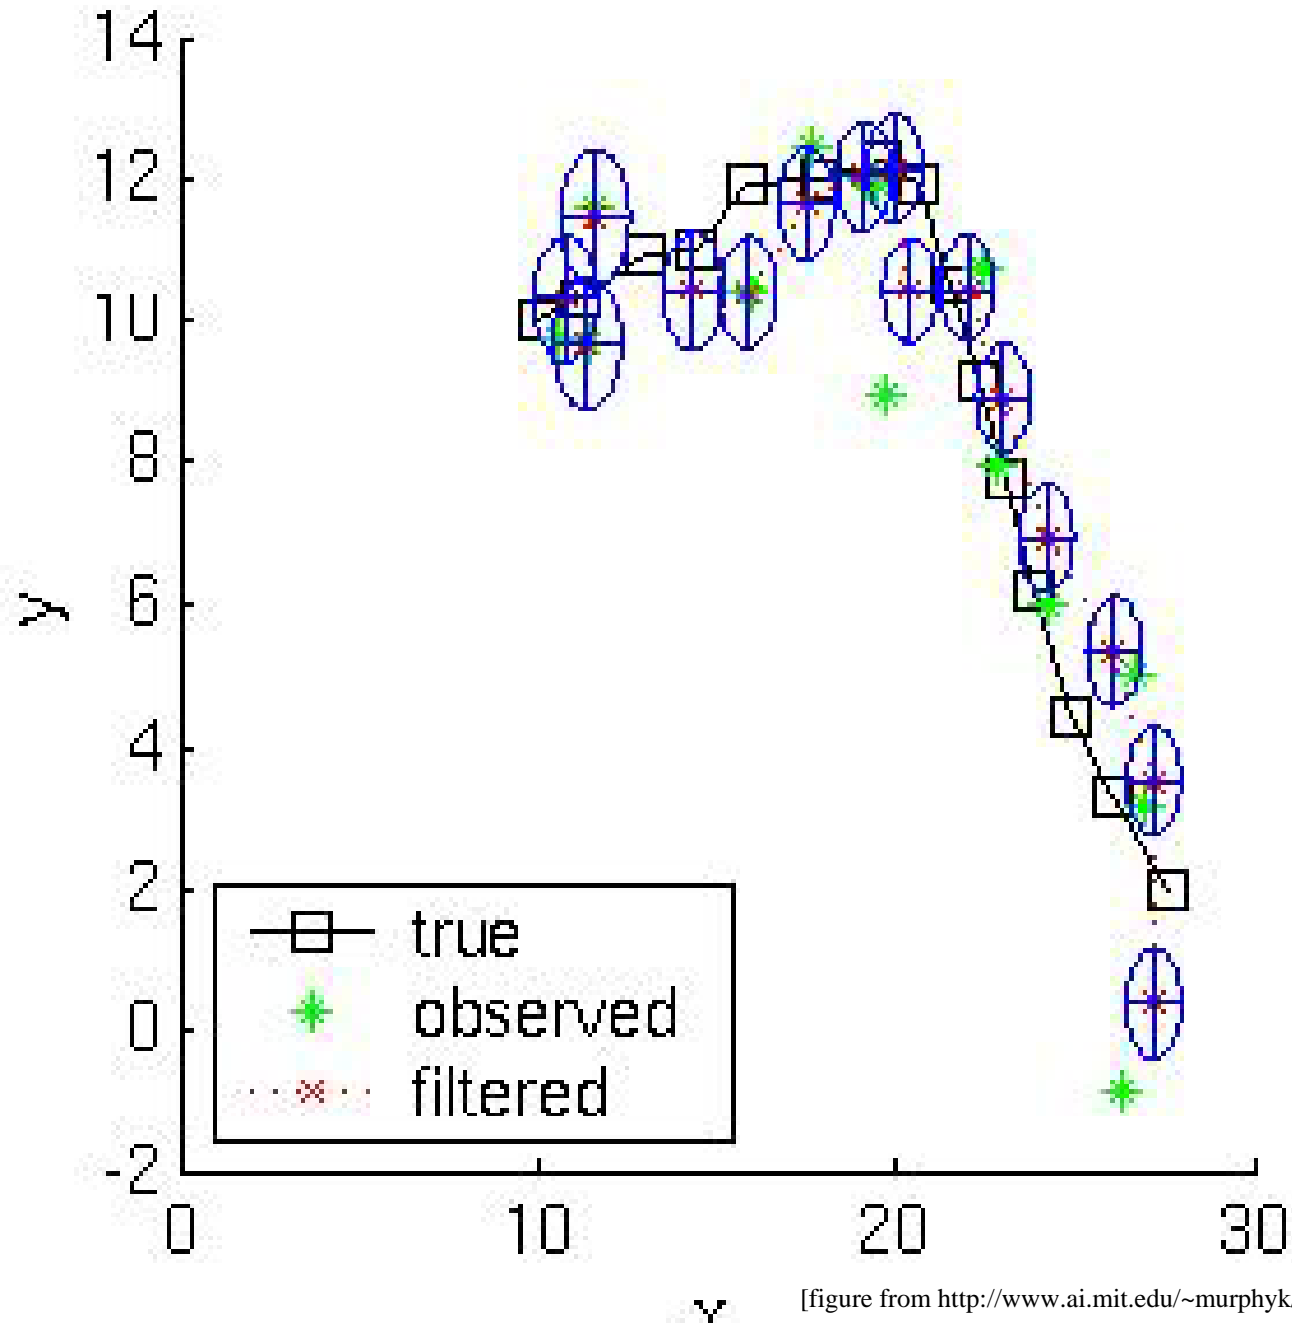
## Combined forward-backward estimates.





The o-s give state, x-s measurement. [time](#)

The \*-s give  $\bar{x}_i^-$ , +-s give  $\bar{x}_i^+$ , vertical bars are 3 standard deviation bars

## 2-D constant velocity example from Kevin Murphy's Matlab toolbox







## 2-D constant velocity example from Kevin Murphy's Matlab toolbox

- MSE of filtered estimate is 4.9; of smoothed estimate. 3.2.
- Not only is the smoothed estimate better, but we know that it is better, as illustrated by the smaller uncertainty ellipses
- Note how the smoothed ellipses are larger at the ends, because these points have seen less data.
- Also, note how rapidly the filtered ellipses reach their steady-state (“Ricatti”) values.

# Resources

- Kalman filter homepage

<http://www.cs.unc.edu/~welch/kalman/>

(kalman filter demo applet)

- Kevin Murphy's Matlab toolbox:

<http://www.ai.mit.edu/~murphyk/Software/Kalman/kalman.html>

# Embellishments for tracking

- Richer models of  $P(\mathbf{x}_n|\mathbf{x}_{n-1})$
- Richer models of  $P(\mathbf{y}_n|\mathbf{x}_n)$
- Richer model of probability distributions

# Abrupt changes

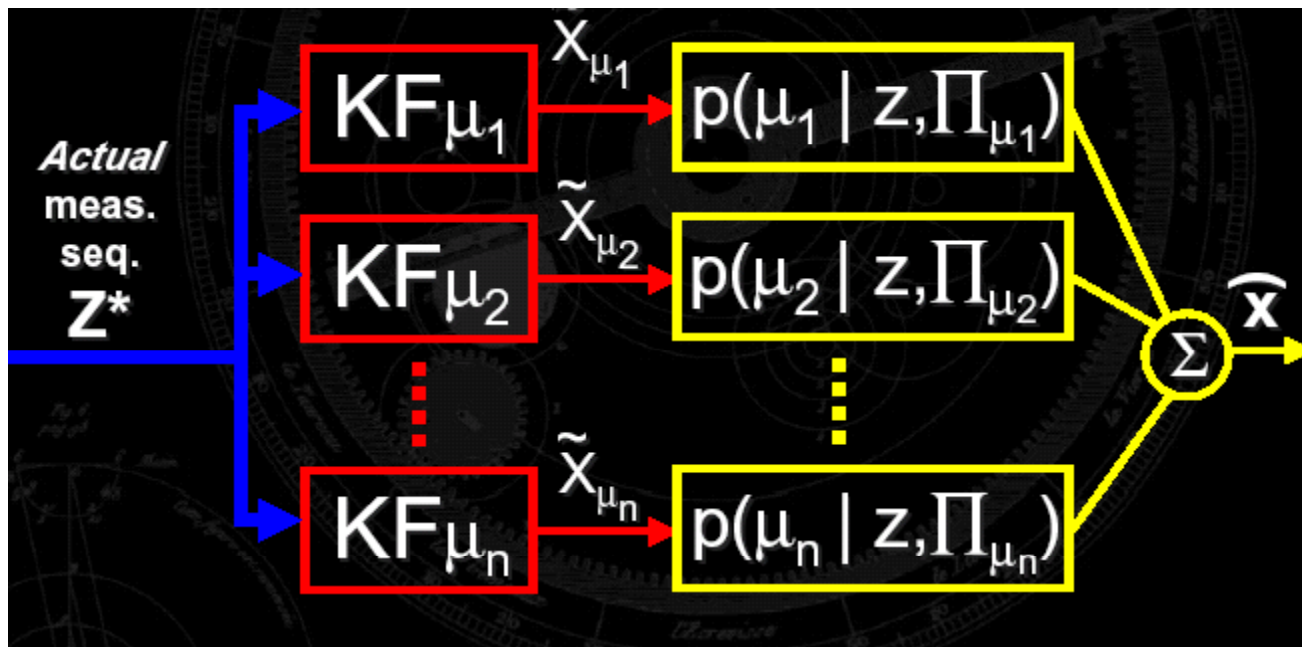
What if environment is sometimes unpredictable?

Do people move with constant velocity?

Test several models of assumed dynamics, use the best.

# Multiple model filters


Test several models of assumed dynamics



[figure from Welsh and Bishop 2001]


# MM estimate

Two models: Position (P), Position+Velocity (PV)




[figure from Welsh and Bishop 2001]

# P likelihood



[figure from Welsh and Bishop 2001]


# No lag



[figure from Welsh and Bishop 2001]



# Smooth when still



[figure from Welsh and Bishop 2001]

# Embellishments for tracking

- Richer models of  $P(\mathbf{y}_n|\mathbf{x}_n)$
- Richer model of probability distributions

# Jepson, Fleet, and El-Maraghi tracker

*IEEE Conference on Computer Vision and Pattern Recognition, Kauai, 2001, Vol. I, pp. 415–422*

## **Robust Online Appearance Models for Visual Tracking**

**Allan D. Jepson\* David J. Fleet† Thomas F. El-Maraghi\*†**


\* Department of Computer Science, University of Toronto, Toronto, M5S 1A4

† Xerox Palo Alto Research Center, 3333 Coyote Hill Rd, Palo Alto, CA 94304

# Wandering, Stable, and Lost appearance model

- Introduce 3 competing models to explain the appearance of the tracked region:
  - A stable model—Gaussian with some mean and covariance.
  - A 2-frame motion tracker appearance model, to rebuild the stable model when it gets lost
  - An outlier model—uniform probability over all appearances.
- Use an on-line EM algorithm to fit the (changing) model parameters to the recent appearance data.

# Jepson, Fleet, and El-Maraghi tracker for toy 1-pixel image model



Red line: observations

Blue line: true appearance state

Black line: mean of the stable process

Mixing probabilities for Stable (black), Wandering (red) and Lost (green) processes




Figure 2. Estimation using on-line EM. (top) The original data (thin red) with true state (dashed blue) and the estimated mean of the stable process (thick black). The noise is a mixture of Gaussian and uniform densities, with mixing probabilities (0.9, 0.1), except for 15 frames at 300 which are pure outliers. (bottom) Mixing probabilities for  $\mathcal{S}$  (black),  $\mathcal{W}$  (dashed red), and the  $\mathcal{L}$  (light green).

# Non-toy image representation

Phase of a steerable quadrature pair (G2, H2).

Steered to 4 different orientations, at 2 scales.

# The motion tracker

- Motion prior,  $P(x_n | x_{n-1})$ , prefers slow velocities and small accelerations.
- The WSL appearance model gives a likelihood for each possible new position, orientation, and scale of the tracked region.
- They combine that with the motion prior to find the most probable position, orientation, and scale of the tracked region in the next frame.
- Gives state-of-the-art tracking results.



# Jepson, Fleet, and El-Maraghi tracker



Figure 4. The adaptation of the model during tracking. (top) The target region in selected frames 200, 300, 480. (bottom) The stable component's mixing probability (left) and mean (right) for the selected frames.



# Jepson, Fleet, and El-Maraghi tracker

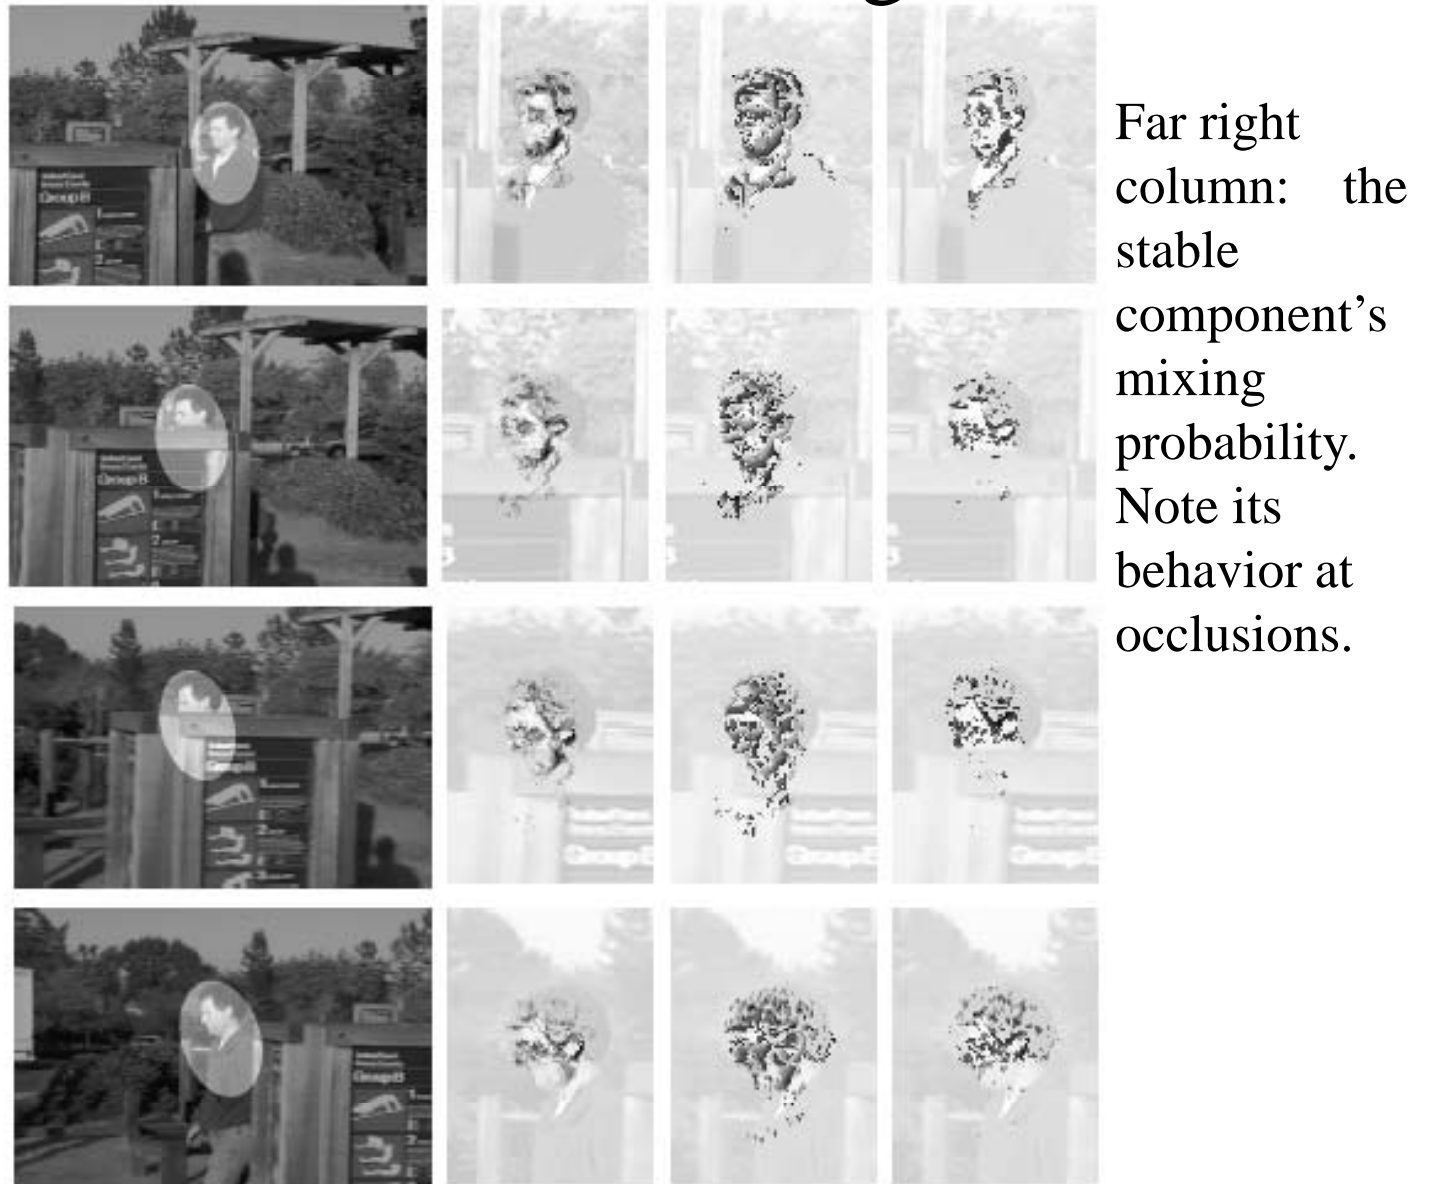


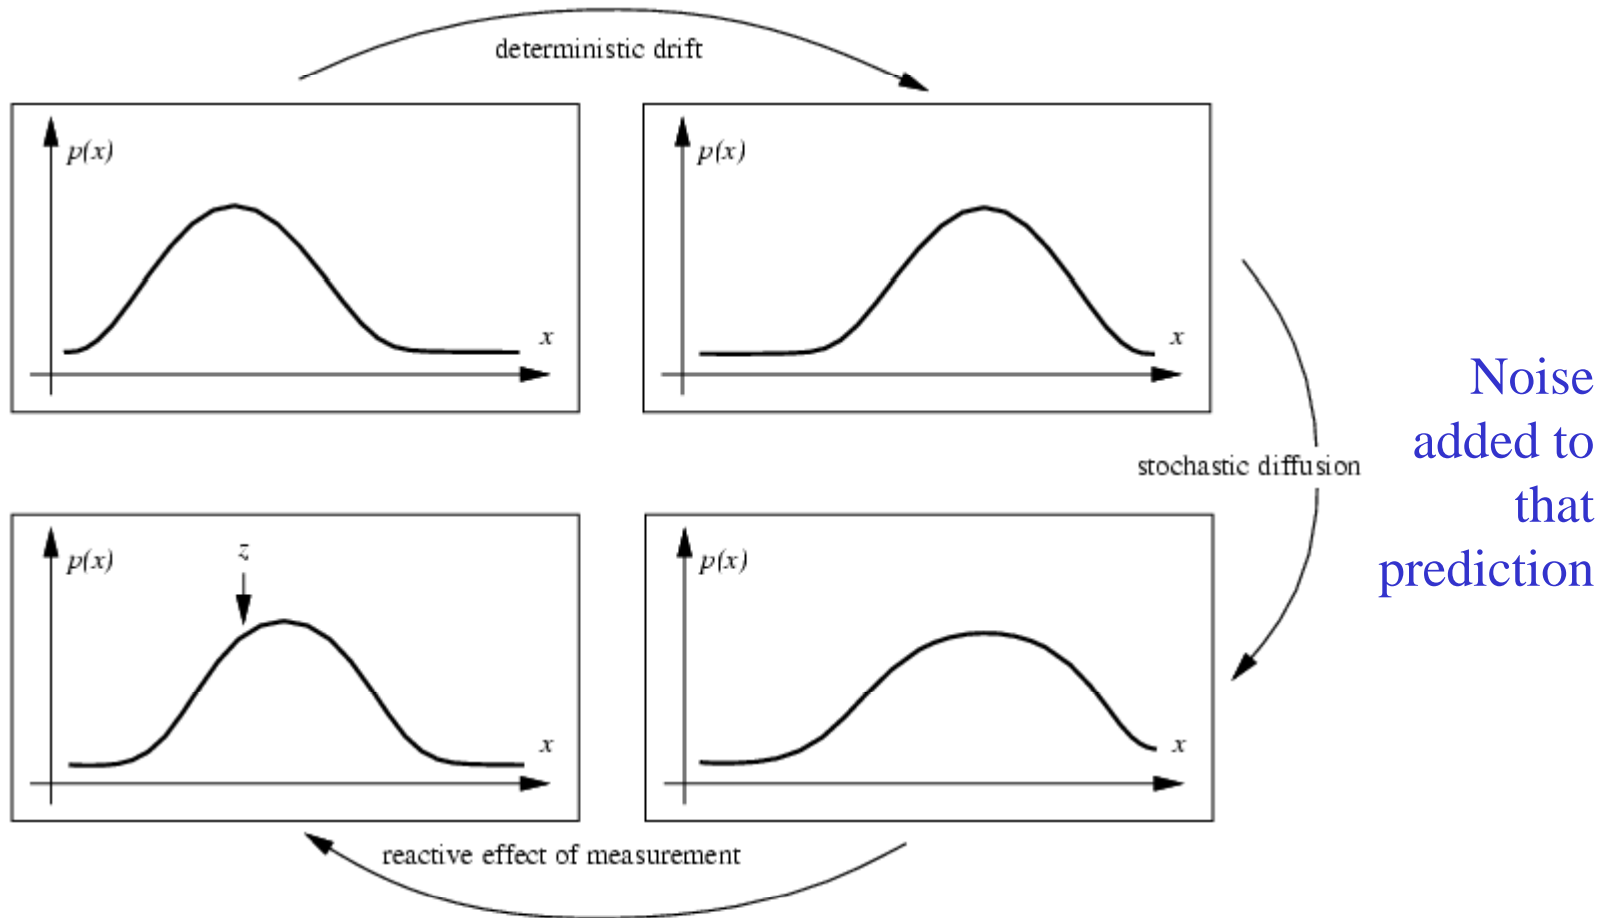
Figure 3. Each row shows, from left to right, the tracking region, the stable component's mixing probability  $m_s(x, t)$ , mean  $\mu_s(x, t)$ , and ownership probability  $o_s(x, t)$ . The rows correspond to frames 244, 259, 274, and 289, top to bottom. Note the model persistence and the drop in data ownership within the occluded region.

# Embellishments for tracking

- Richer model of  $P(y_n|x_n)$
- Richer model of probability distributions
  - Particle filter models, applied to tracking humans

# (KF) Distribution propagation


prediction from previous time frame




Noise added to that prediction

Make new measurement at next time frame

# Distribution propagation




# Representing non-linear Distributions




# Representing non-linear Distributions

Unimodal parametric models fail to capture real-world densities...



# Discretize by evenly sampling over the entire state space


Tractable for 1-d problems like stereo, but not for high-dimensional problems.






# Representing Distributions using Weighted Samples

Rather than a parametric form, use a set of samples  
to represent a density:



# Representing Distributions using Weighted Samples

Rather than a parametric form, use a set of samples to represent a density:




$$\{(u^i, w^i)\}$$

Sample positions

Probability mass at each sample

This gives us two knobs to adjust when representing a probability density by samples: the locations of the samples, and the probability weight on each sample.

# Representing distributions using weighted samples, another picture



# Sampled representation of a probability distribution

Represent a probability distribution

$$p_f(\mathbf{X}) = \frac{f(\mathbf{X})}{\int f(\mathbf{U})d\mathbf{U}}$$

by a set of  $N$  weighted samples


$$\{(\mathbf{u}^i, w^i)\}$$

where  $\mathbf{u}^i \sim s(\mathbf{u})$  and  $w^i = f(\mathbf{u}^i)/s(\mathbf{u}^i)$ .

You can also think of this as a sum of dirac delta functions, each of weight  $w$ :

$$p_f(x) = \sum_i w^i \delta(x - u^i)$$

# Tracking, in particle filter representation



$$p_f(x) = \sum_i w^i \delta(x - u^i)$$

$$P(x_n | y_1 \dots y_n) = k P(y_n | x_n) \int dx_{n-1} P(x_n | x_{n-1}) P(x_{n-1} | y_1 \dots y_{n-1})$$



# Particle filter

- Let's apply this sampled probability density machinery to generalize the Kalman filtering framework.
- More general probability density representation than uni-modal Gaussian.
- Allows for general state dynamics,  $f(x) + \text{noise}$

# Sampled Prediction

$$P(\mathbf{x}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) = ?$$

$$p(\mathbf{X}_{i-1} | \mathbf{y}_0, \dots, \mathbf{y}_{i-1})$$

$$\{(\mathbf{u}_{i-1}^k, w_{i-1}^k)\}$$

$$\longrightarrow \mathbf{x}_i = \mathbf{f}(\mathbf{x}_{i-1}) + \xi_i \longrightarrow$$

$$\{((f(\mathbf{u}_{i-1}^k) + \xi_i^l, \mathbf{u}_{i-1}^k), w_{i-1}^k)\}$$

$$p(\mathbf{X}_i, \mathbf{X}_{i-1} | \mathbf{y}_0, \dots, \mathbf{y}_{i-1})$$

Drop elements to marginalize to get

$$P(\mathbf{x}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) \approx$$

$$\{(f(\mathbf{u}_{i-1}^k) + \xi_i^l, w_{i-1}^k)\}$$

# Sampled Correction (Bayes rule)

Prior  $\rightarrow$  posterior

Reweight every sample with the likelihood of the observations, given that sample:

$$p(\mathbf{Y}_i = \mathbf{y}_i | \mathbf{X}_i = \mathbf{s}_i^{k,-}) w_i^{k,-}$$

yielding a set of samples describing the probability distribution after the correction (update) step:

$$\left\{ (\mathbf{s}_i^{k,-}, p(\mathbf{Y}_i = \mathbf{y}_i | \mathbf{X}_i = \mathbf{s}_i^{k,-}) w_i^{k,-}) \right\}$$



# Naïve PF Tracking

- Start with samples from something simple (Gaussian)

- Repeat

– Correct

Take each particle from the prediction step and modify the old weight by multiplying by the new likelihood

$$\left\{ (\mathbf{s}_i^{k,-}, p(\mathbf{Y}_i = \mathbf{y}_i | \mathbf{X}_i = \mathbf{s}_i^{k,-}) w_i^{k,-}) \right\}$$

– Predict


$$\left\{ (f(\mathbf{s}_{i-1}^k) + \xi_i^l, w_{i-1}^k) \right\}$$

Run every particle through the dynamics function and add noise.

But doesn't work that well because of sample impoverishment...

# Sample impoverishment

10 of the 100 particles, along with the true Kalman filter track, with variance:



# Resample the prior



In a sampled density representation, the frequency of samples can be traded off against weight:

$$(\mathbf{s}_k, w_k) \longrightarrow \begin{array}{c} (\mathbf{s}_k, 1) \\ (\mathbf{s}_k, 1) \\ (\mathbf{s}_k, 1) \\ \vdots \end{array} N_k \text{ copies} \quad \text{s.t.} \quad \frac{N_k}{\sum_k N_k} = w_k$$

These new samples are a representation of the same density.

I.e., make  $N$  draws with replacement from the original set of samples, using the weights as the probability of drawing a sample.

# Resampling concentrates samples



## A practical particle filter with resampling

**Initialization:** Represent  $P(X_0)$  by a set of  $N$  samples

$$\left\{ (s_0^{k,-}, w_0^{k,-}) \right\}$$

where

$$s_0^{k,-} \sim P_s(S) \text{ and } w_0^{k,-} = P(s_0^{k,-})/P_s(S = s_0^{k,-})$$

Ideally,  $P(X_0)$  has a simple form and  $s_0^{k,-} \sim P(X_0)$  and  $w_0^{k,-} = 1$ .

**Prediction:** Represent  $P(X_i|y_0, y_{i-1})$  by

$$\left\{ (s_i^{k,-}, w_i^{k,-}) \right\}$$

where

$$s_i^{k,-} = f(s_{i-1}^{k,+}) + \xi_i^k \text{ and } w_i^{k,-} = w_{i-1}^{k,+} \text{ and } \xi_i^k \sim N(0, \Sigma_{d_i})$$

**Correction:** Represent  $P(X_i|y_0, y_i)$  by


$$\left\{ (s_i^{k,+}, w_i^{k,+}) \right\}$$

where

$$s_i^{k,+} = s_i^{k,-} \text{ and } w_i^{k,+} = P(Y_i = y_i | X_i = s_i^{k,-}) w_i^{k,-}$$

**Resampling:** Normalise the weights so that  $\sum_i w_i^{k,+} = 1$  and compute the variance of the normalised weights. If this variance exceeds some threshold, then construct a new set of samples by drawing, with replacement,  $N$  samples from the old set, using the weights as the probability that a sample will be drawn. The weight of each sample is now  $1/N$ .

# Pictorial view



Contour tracking by stochastic propagation of conditional density - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://research.microsoft.com/users/misard/abstracts/eccv96.isard.html

Getting Started Latest Headlines

## Contour tracking by stochastic propagation of conditional density

**Michael Isard and Andrew Blake**  
**Proc. European Conference on Computer Vision, vol. 1, pp. 343--356, Cambridge UK, (1996).**

### Abstract

The problem of tracking curves in dense visual clutter is a challenging one. Trackers based on Kalman filters are of limited use; because they are based on Gaussian densities which are unimodal, they cannot represent simultaneous alternative hypotheses. Extensions to the Kalman filter to handle multiple data associations work satisfactorily in the simple case of point targets, but do not extend naturally to continuous curves. A new, stochastic algorithm is proposed here, the **Condensation** algorithm --- Conditional Density Propagation over time. It uses 'factored sampling', a method previously applied to interpretation of static images, in which the distribution of possible interpretations is represented by a randomly generated set of representatives. The **Condensation** algorithm combines factored sampling with learned dynamical models to propagate an entire probability distribution for object position and shape, over time. The result is highly robust tracking of agile motion in clutter, markedly superior to what has previously been attainable from Kalman filtering. Notwithstanding the use of stochastic methods, the algorithm runs in near real-time.

Click here for a [compressed postscript](#) version

### Back to

[Michael Isard's home page](#)

# The Condensation Algorithm

## Background



Tracking objects through highly cluttered scenes is difficult. We believe that for tracking to be robust when following agile moving objects, in the presence of dense background clutter, probabilistic algorithms are essential. Previous algorithms, for example the Kalman filter, have been limited in the range of probability distributions they represent. We have developed a new algorithm, the **Condensation** algorithm (**Conditional Density Propagation**) which allows quite general representations of probability. Experimental results show that this increased generality does indeed lead to a marked improvement in tracking performance. In addition to permitting high-quality tracking in clutter, the simplicity of the **Condensation** algorithm also allows the use of non-linear motion models more complex than those commonly used in Kalman filters. We have implemented a mixed discrete/continuous tracker in the **Condensation** framework which switches between multiple continuous Auto-Regressive Process motion models according to a discrete transition matrix. Also, by using the statistical technique of *importance sampling* it is possible to build a **Condensation** tracker which runs in real time, and we have implemented a real-time hand-tracker on a low-end SGI workstation. My [D.Phil. thesis](#) gives a thorough description of the algorithm and some applications.

## Sample Code

[Download](#) source code of a simple implementation of the **Condensation** algorithm.

## Results

Here is an MPEG (2.3Mb) showing the **Condensation** algorithm tracking a [leaf blowing in the wind](#), against a background of





# Animation of condensation algorithm



# Applications


## Tracking

- hands
- bodies
- Leaves


What might we expect?

Reliable, robust, slow

# Contour tracking




# Head tracking



(a)

Picture of the states represented by  
the top weighted particles




(b)

The mean state

# Leaf tracking



# Hand tracking



# Desired operations with a probability density

- Expected value
- Marginalization
- Bayes rule

# Computing an expectation using sampled representation

$$E_{p_f} [g] = \int g(\mathbf{U}) p_f(\mathbf{U}) d\mathbf{U}$$

using

$$p_f(x) = \sum_i w^i \delta(x - u^i)$$

$$\approx \frac{\sum_i g(u_i) w_i}{\sum_i w_i}$$



# Marginalizing a sampled density

If we have a sampled representation of a joint density

$$\{((\mathbf{m}^i, \mathbf{n}^i), w^i)\}$$

and we wish to marginalize over one variable:

$$p_f(\mathbf{M}) = \int p_f(\mathbf{M}, \mathbf{N}) d\mathbf{N}$$

we can simply ignore the corresponding components of the samples (!):

$$\begin{aligned} \int g(\mathbf{M}) p_f(\mathbf{M}) d\mathbf{M} &= \int g(\mathbf{M}) \int p_f(\mathbf{M}, \mathbf{N}) d\mathbf{N} d\mathbf{M} \\ &= \int \int g(\mathbf{M}) p_f(\mathbf{M}, \mathbf{N}) d\mathbf{N} d\mathbf{M} \\ &\approx \frac{\sum_{i=1}^N g(\mathbf{m}^i) w^i}{\sum_{i=1}^N w^i} \end{aligned}$$

# Marginalizing a sampled density

Assume we have a sampled representation of a distribution

$$p_f(\mathbf{M}, \mathbf{N})$$

given by

$$\{((\mathbf{m}^i, \mathbf{n}^i), w^i)\}$$

Then

$$\{(\mathbf{m}^i, w^i)\}$$

is a representation of the marginal,

$$\int p_f(\mathbf{M}, \mathbf{N}) d\mathbf{N}$$

# Sampled Bayes rule

Assume we have a representation of  $p(\mathbf{U})$  as

$$\{(\mathbf{u}^i, w^i)\}$$

Assume we have an observation  $\mathbf{V} = \mathbf{v}_0$ ,

and a likelihood model  $p(\mathbf{V}|\mathbf{U})$ .

The posterior,  $p(\mathbf{U}|\mathbf{V} = \mathbf{v}_0)$  is represented by

$$k p(\mathbf{V}=\mathbf{v}_0|\mathbf{U})p(\mathbf{U}) \{(\mathbf{u}^i, w'^i)\}$$

where

$$w'^i = p(\mathbf{V} = \mathbf{v}_0|\mathbf{u}^i)w^i$$