6.801/866

Tracking with Non-linear Dynamic Models

T. Darrell

Tracking with Non-linear Dynamic Models

- Distribution propagation
- Problems with non-linearities
- Sampling densities
- Particle filtering
- Tracking people

Distribution propogation

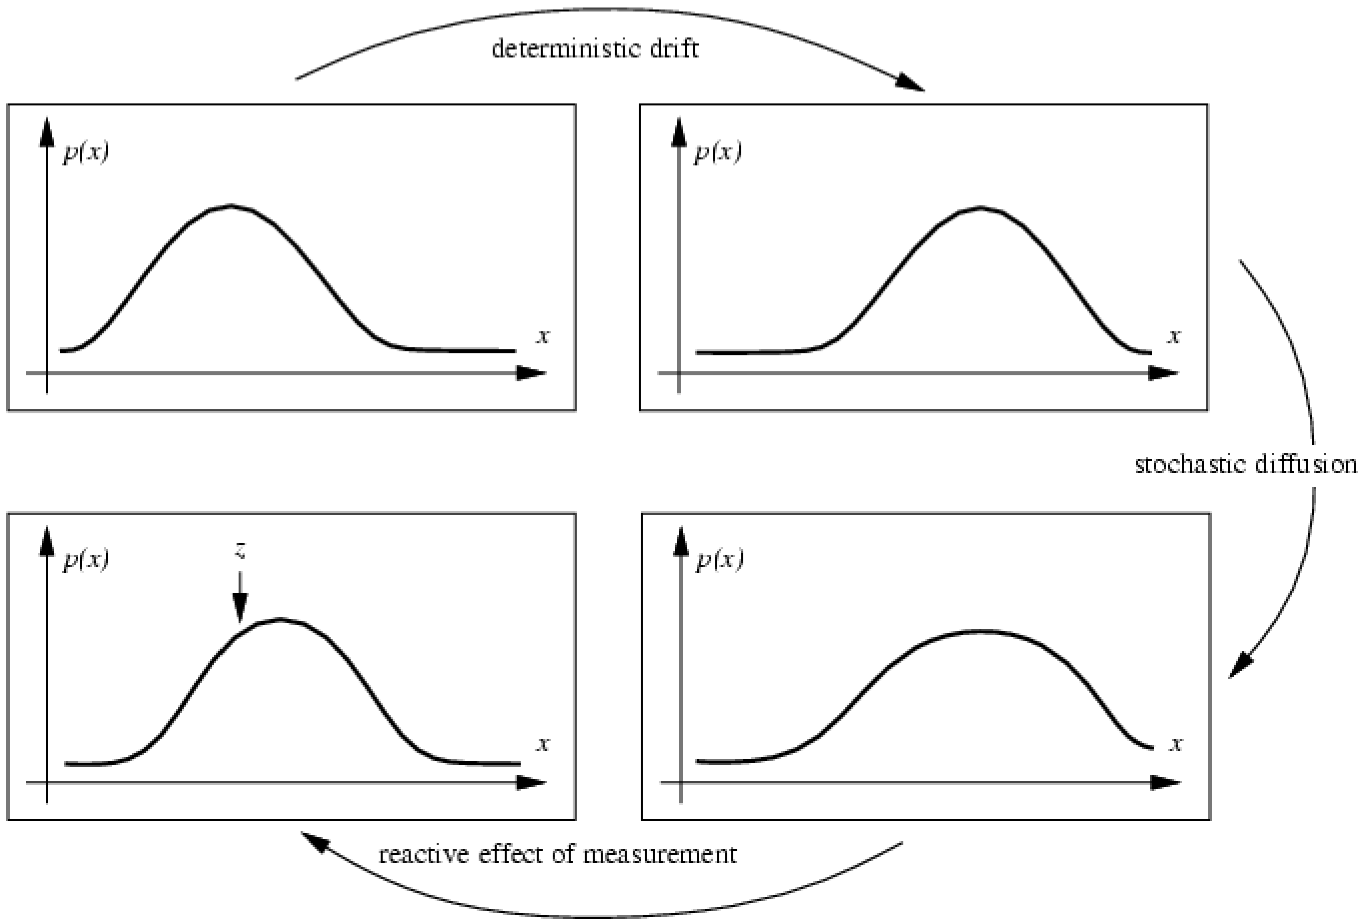

[Isard 1998]

Distribution propogation

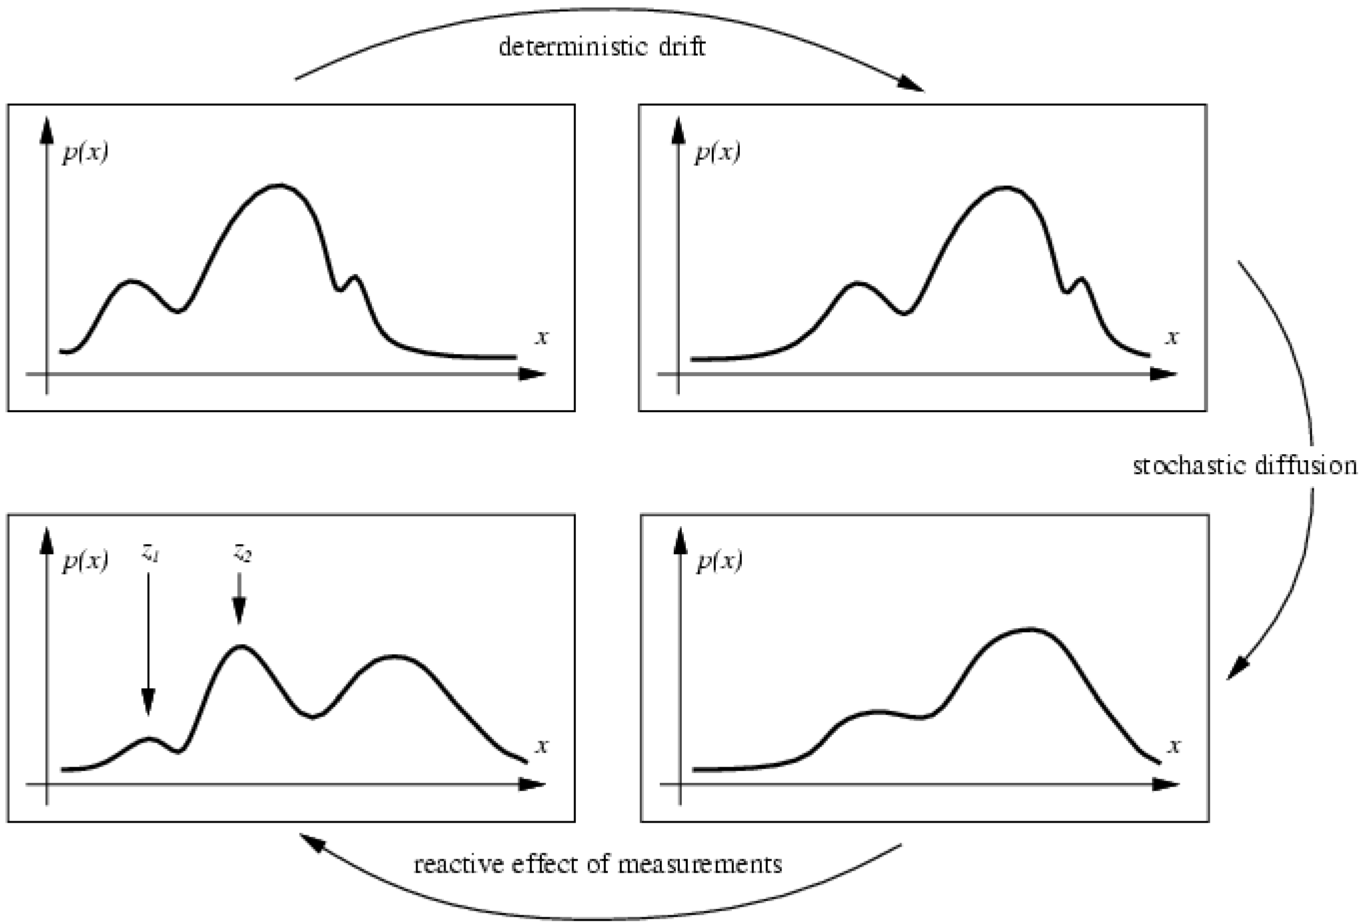

[Isard 1998]

A little nonlinearity becomes very non-Gaussian

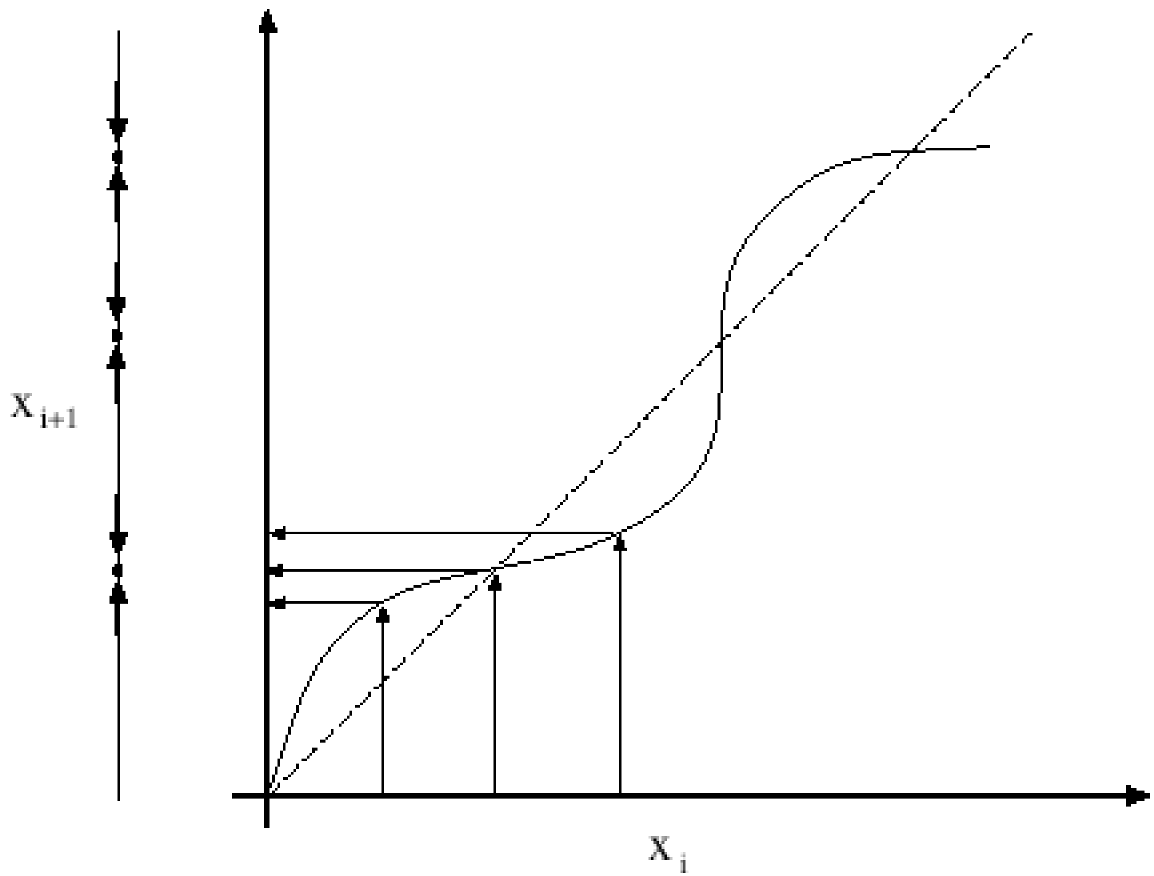

$$x_{i+1} = x_i + 0.1 \sin x_i$$

time evolution of the state of a set of 100 points

$$x_{i+1} = x_i + 0.1 \sin x_i$$

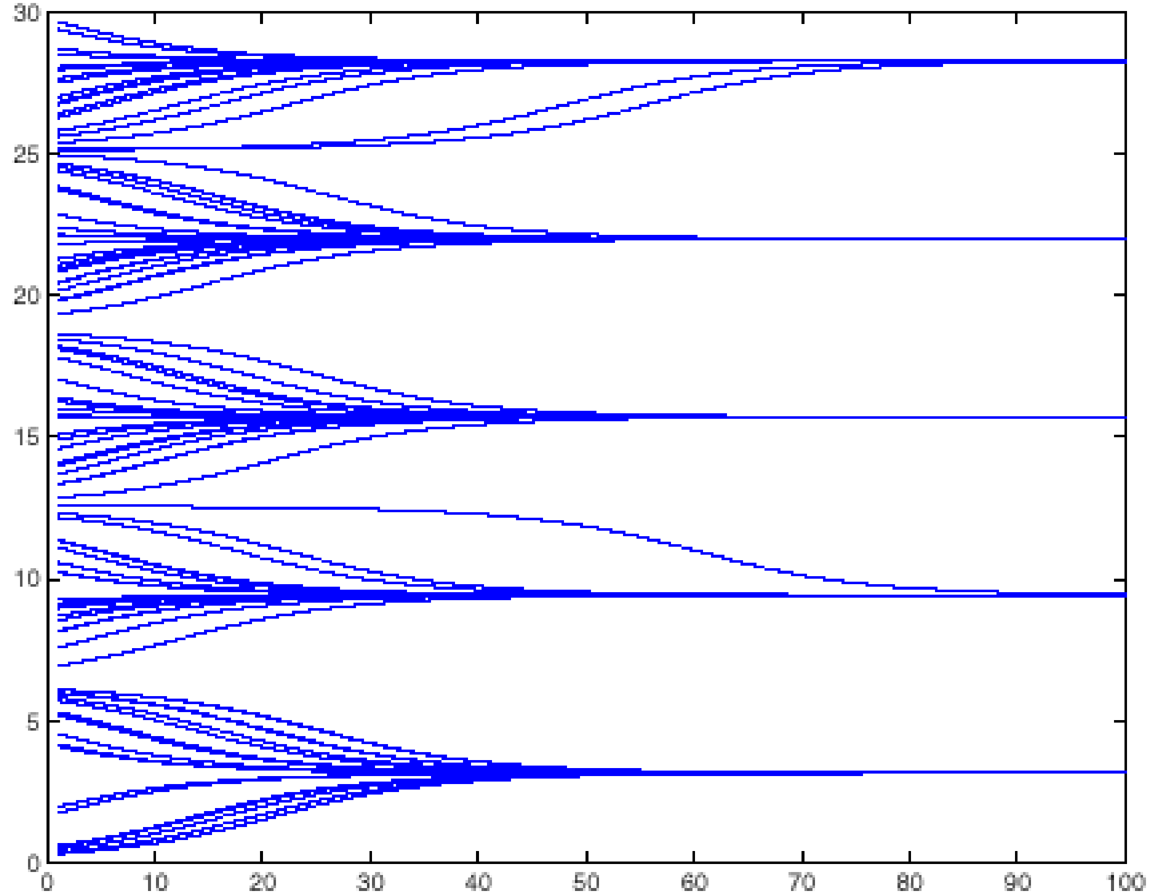

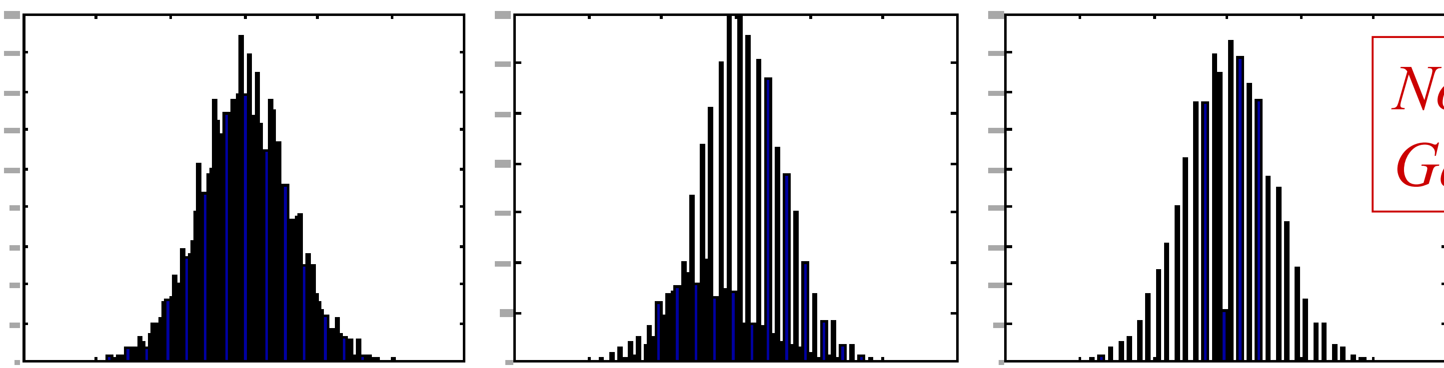

$P(x_0)$ $P(x_{20})$ $P(x_{70})$

Not Gaussian!

EKF

Linearize system at each time point to form an Extended Kalman Filter (EKF)

- Compute Jacobian matrix

$$\mathcal{J}(\boldsymbol{g}; \boldsymbol{x}_j)$$

whose (l,m)'th value is $\frac{\partial f_l}{\partial x_m}$ evaluated at $\boldsymbol{x}_j$

- use this for forward model at each step in KF

Useful in many engineering applications, but not as successful in computer vision….

Representing non-linear Distributions

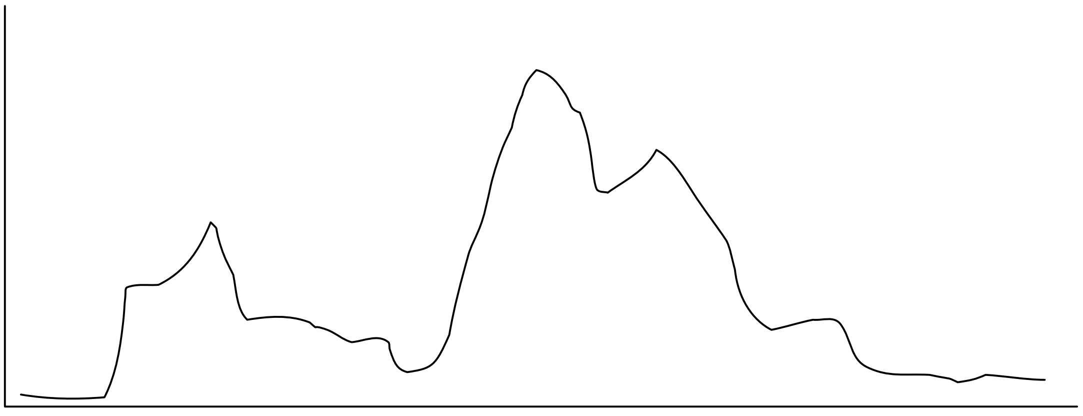

Representing non-linear Distributions

Unimodal parametric models fail to capture real-world densities…

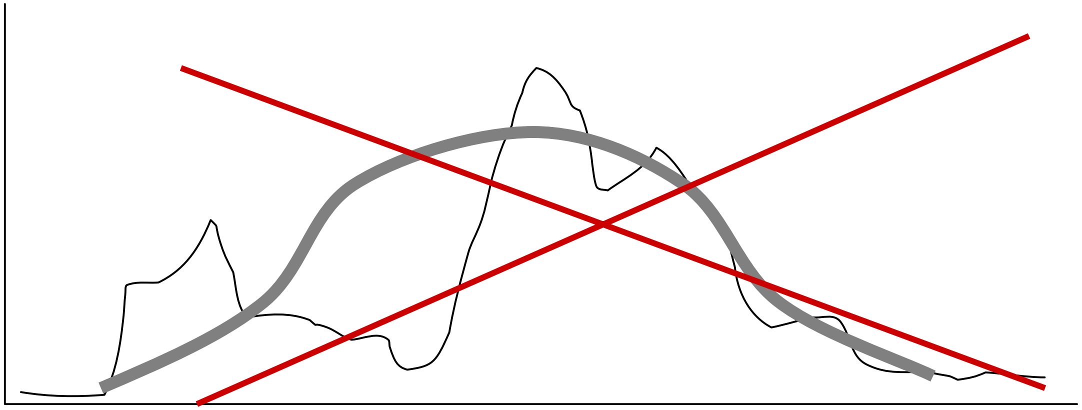

Representing non-linear Distributions

Mixture models are appealing, but very hard to propagate analytically!

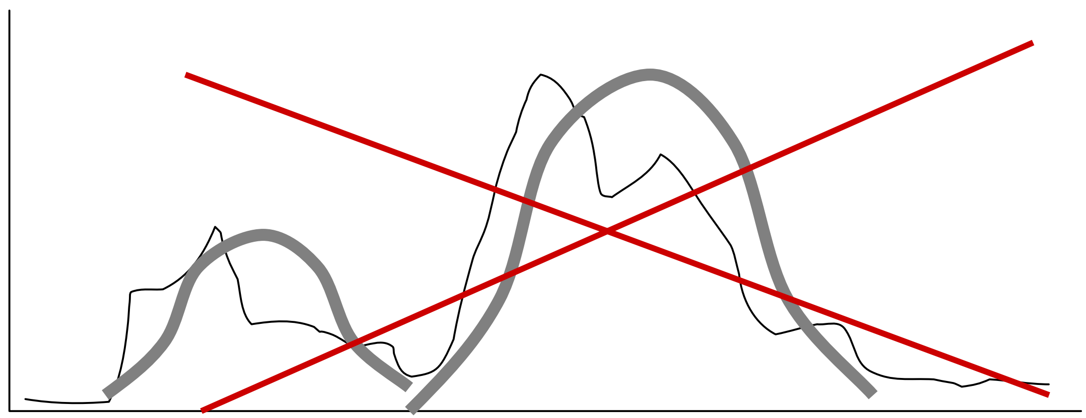

[but see Cham and Rehg’s MHT approach]

Representing Distributions using Weighted Samples

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

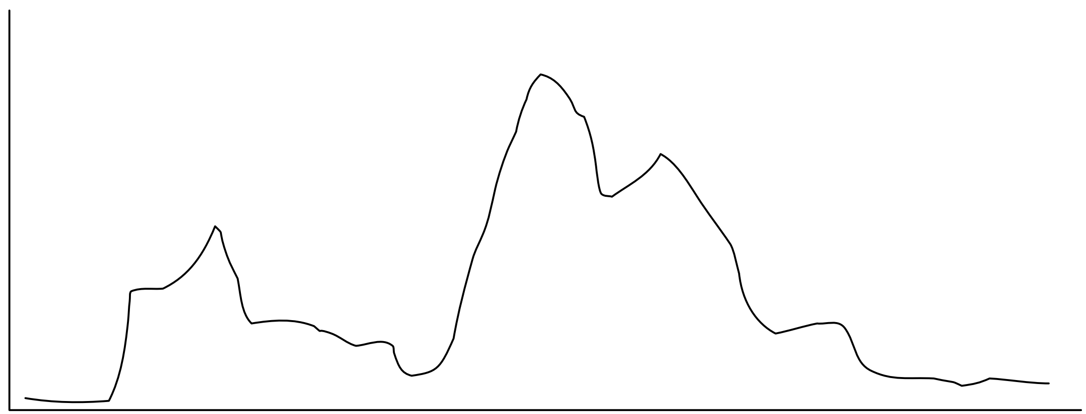

Representing Distributions using Weighted Samples

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

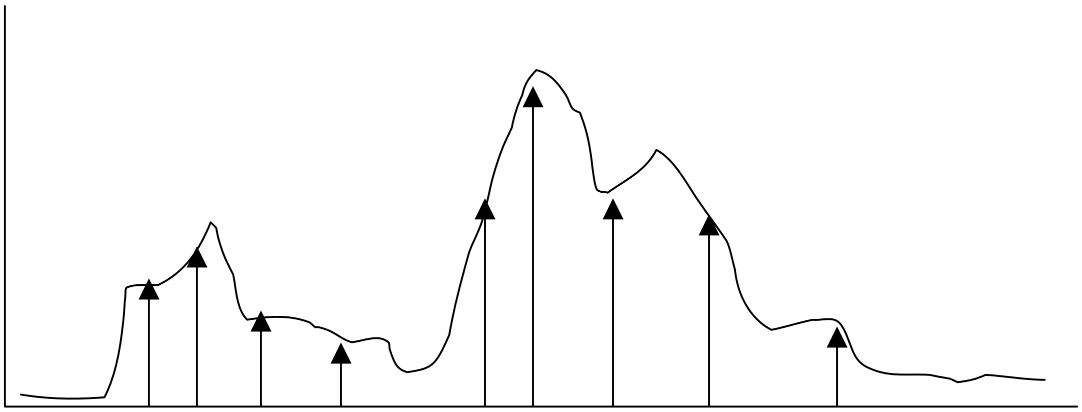

1. a set of sample locations
   from what distribution?
2. weights associated with those locations
   how does that distribution effect the weights?

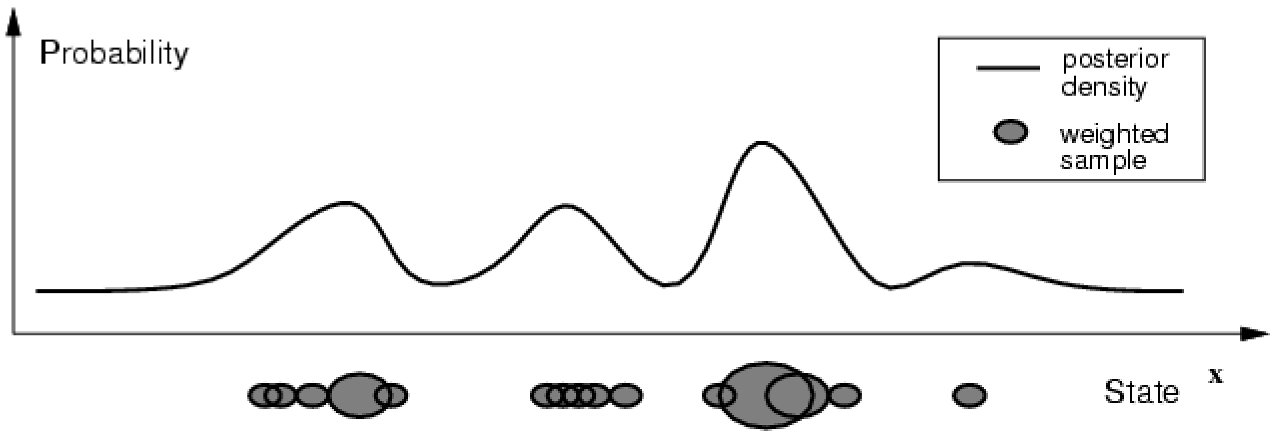

[Isard 1998]

Expectation over sampled density

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

then

$$\begin{aligned}
\mathrm{E}_{p_f}[g] &= \int g(\boldsymbol{U})p_f(\boldsymbol{U})d\boldsymbol{U} \\
&= \frac{\int g(\boldsymbol{U})f(\boldsymbol{U})d\boldsymbol{U}}{\int f(\boldsymbol{U})d\boldsymbol{U}} \\
&= \mathrm{E}\left[\frac{\sum_i g(\boldsymbol{u}_i)w_i}{\sum_i w_i}\right] \\
&\approx \frac{\sum_i g(\boldsymbol{u}_i)w_i}{\sum_i w_i}
\end{aligned}$$

$$w^i = f(\boldsymbol{u}^i)/s(\boldsymbol{u}^i)$$

$$\begin{aligned}
\mathrm{E}\left[\frac{1}{N}\sum_i w^i\right] &= \int 1\frac{f(\boldsymbol{U})}{s(\boldsymbol{U})}s(\boldsymbol{U})d\boldsymbol{U} \\
&= \int f(\boldsymbol{U})d\boldsymbol{U}
\end{aligned}$$

Expectation over sampled density

We have a representation of a probability distribution

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

by a set of weighted samples

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

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

$$\int g(\boldsymbol{U}) p_f(\boldsymbol{U}) d\boldsymbol{U} \approx \frac{\sum_{i=1}^N g(\boldsymbol{u}^i) w^i}{\sum_{i=1}^N w^i}$$

Sampled representation of a probability distribution

Represent a probability distribution

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

by a set of N weighted samples

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

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

Marginalizing a sampled density

If we have a sampled representation of a joint density

$$\left\{((\boldsymbol{m}^i, \boldsymbol{n}^i), w^i)\right\}$$

and we wish to marginalize over one variable:

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

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

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

Marginalizing a sampled density

Assume we have a sampled representation of a distribution

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

given by

$$\left\{((\boldsymbol{m}^i, \boldsymbol{n}^i), w^i)\right\}$$

Then

$$\left\{(\boldsymbol{m}^i, w^i)\right\}$$

is a representation of the marginal,

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

Sampled Bayes

Transforming a Sampled Representation of a Prior into a Sampled Representation of a Posterior:

$$\int g(\boldsymbol{U}) p(\boldsymbol{U}|\boldsymbol{V} = v_0) d\boldsymbol{U} = \frac{1}{K} \int g(\boldsymbol{U}) p(\boldsymbol{V} = \boldsymbol{v}_0 | \boldsymbol{U}) p(\boldsymbol{U}) d\boldsymbol{U}$$

$$\approx \frac{1}{K} \frac{\sum_{i=1}^{N} g(\boldsymbol{u}^i) p(\boldsymbol{V} = \boldsymbol{v}_0 | \boldsymbol{u}^i) w^i}{\sum_{i=1}^{N} w^i}$$

$$\approx \frac{\sum_{i=1}^{N} g(\boldsymbol{u}^i) p(\boldsymbol{V} = \boldsymbol{v}_0 | \boldsymbol{u}^i) w^i}{\sum_{i=1}^{N} p(\boldsymbol{V} = \boldsymbol{v}_0 | \boldsymbol{u}^i) w^i}$$

Sampled Bayes

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

$$\left\{(\boldsymbol{u}^i, w^i)\right\}$$

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

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

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

$$\left\{(\boldsymbol{u}^i, w'^i)\right\}$$

where

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

Sampled Prediction

$$P(\boldsymbol{x}_i|\boldsymbol{y}_0,\ldots,\boldsymbol{y}_{i-1}) \;=\; ?$$

$$p(\boldsymbol{X}_i,\boldsymbol{X}_{i-1}|\boldsymbol{y}_0,\ldots,\boldsymbol{y}_{i-1}) = p(\boldsymbol{X}_i|\boldsymbol{X}_{i-1})p(\boldsymbol{X}_{i-1}|\boldsymbol{y}_0,\ldots,\boldsymbol{y}_{i-1})$$

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

$$\left\{(\boldsymbol{u}_{i-1}^k, w_{i-1}^k)\right\}$$

$$\rightarrow \quad \boldsymbol{x}_i = \boldsymbol{f}(\boldsymbol{x}_{i-1}) + \xi_i \quad \rightarrow$$

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

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

Drop elements to get $P(\boldsymbol{x}_i|\boldsymbol{y}_0,\ldots,\boldsymbol{y}_{i-1}) \; \sim=$

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

Sampled Correction

Prior → posterior

Reweight with

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

yielding

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

Naïve PF Tracking

- Start with samples from something simple (Gaussian)
- Repeat
  - Correct

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

  - Predict

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

Doesn't work that well

Sample impoverishment

Test with linear case:

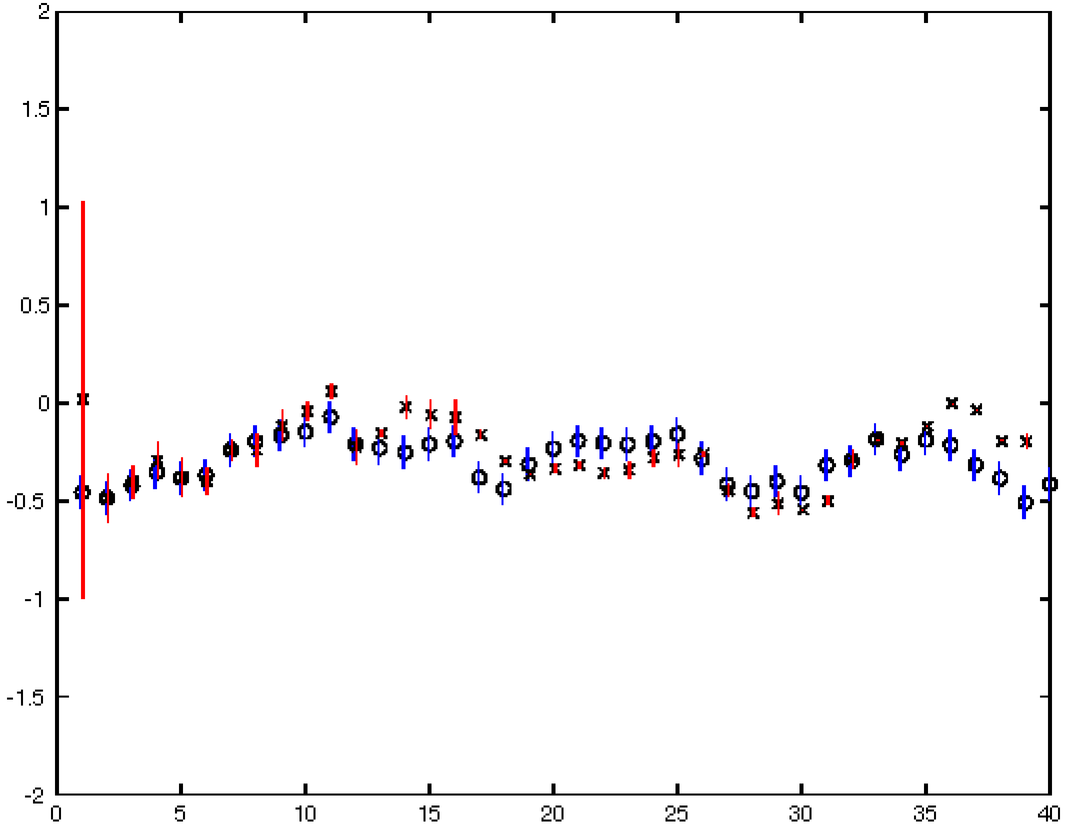

kf: x
pf: o

Sample impoverishment

10 of the 100 particles:

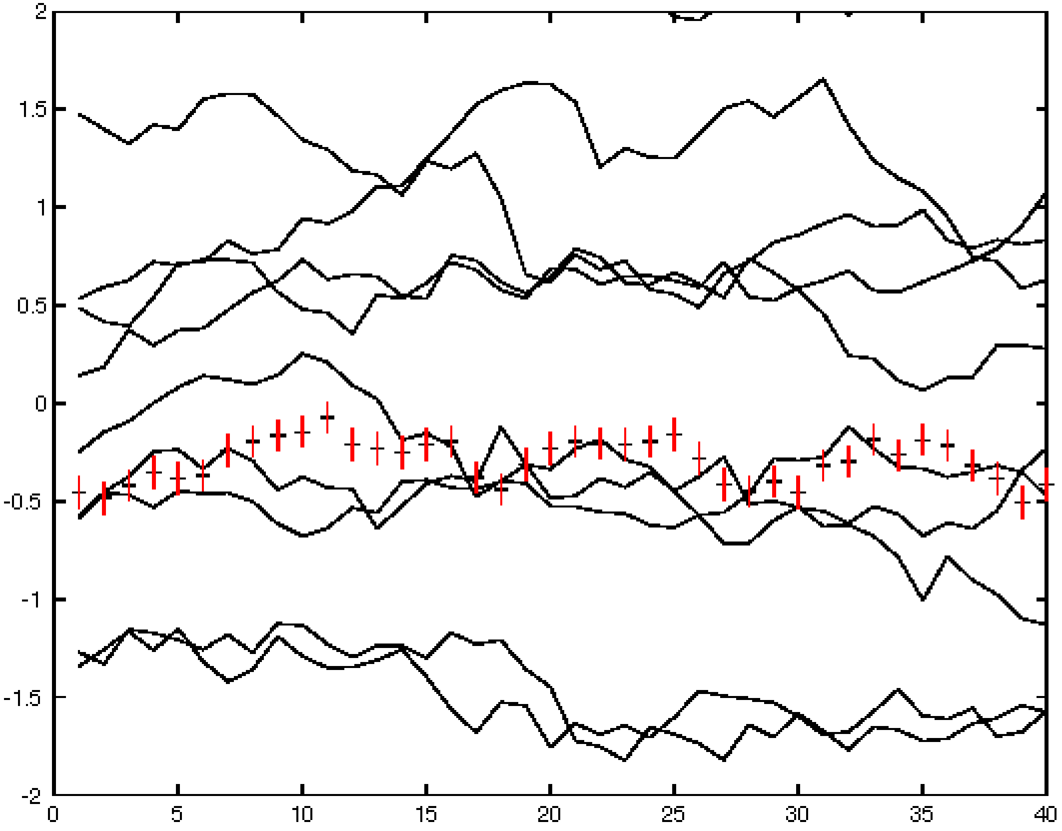

Resample the prior

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

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

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

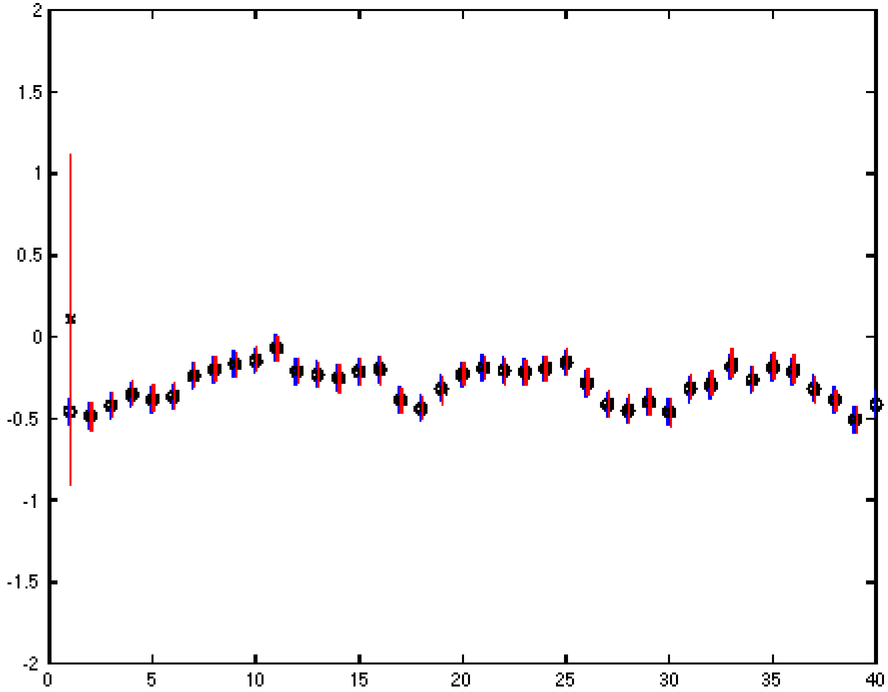

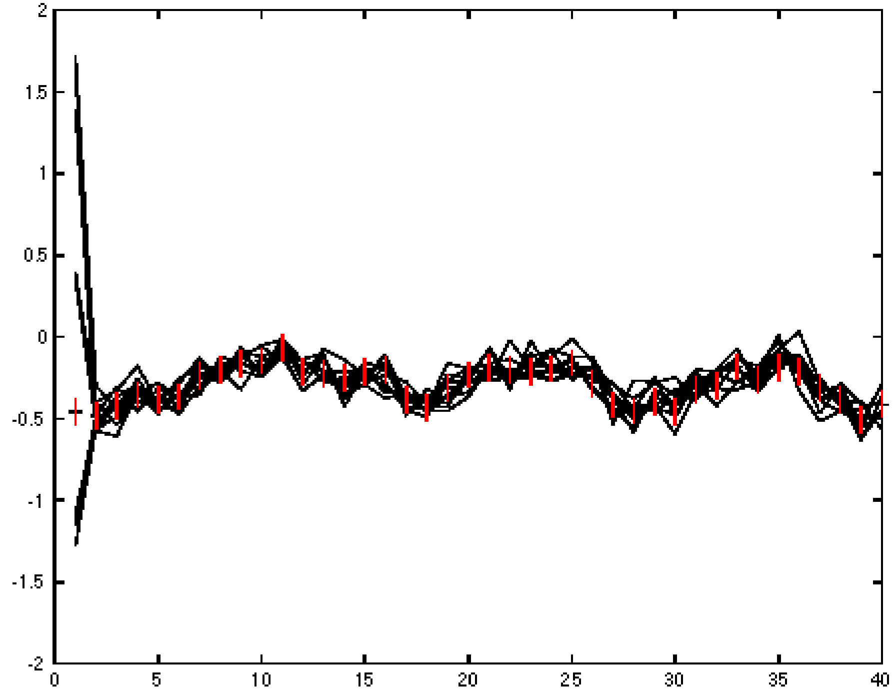

A practical particle filter with resampling

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 $1/N$.

A variant

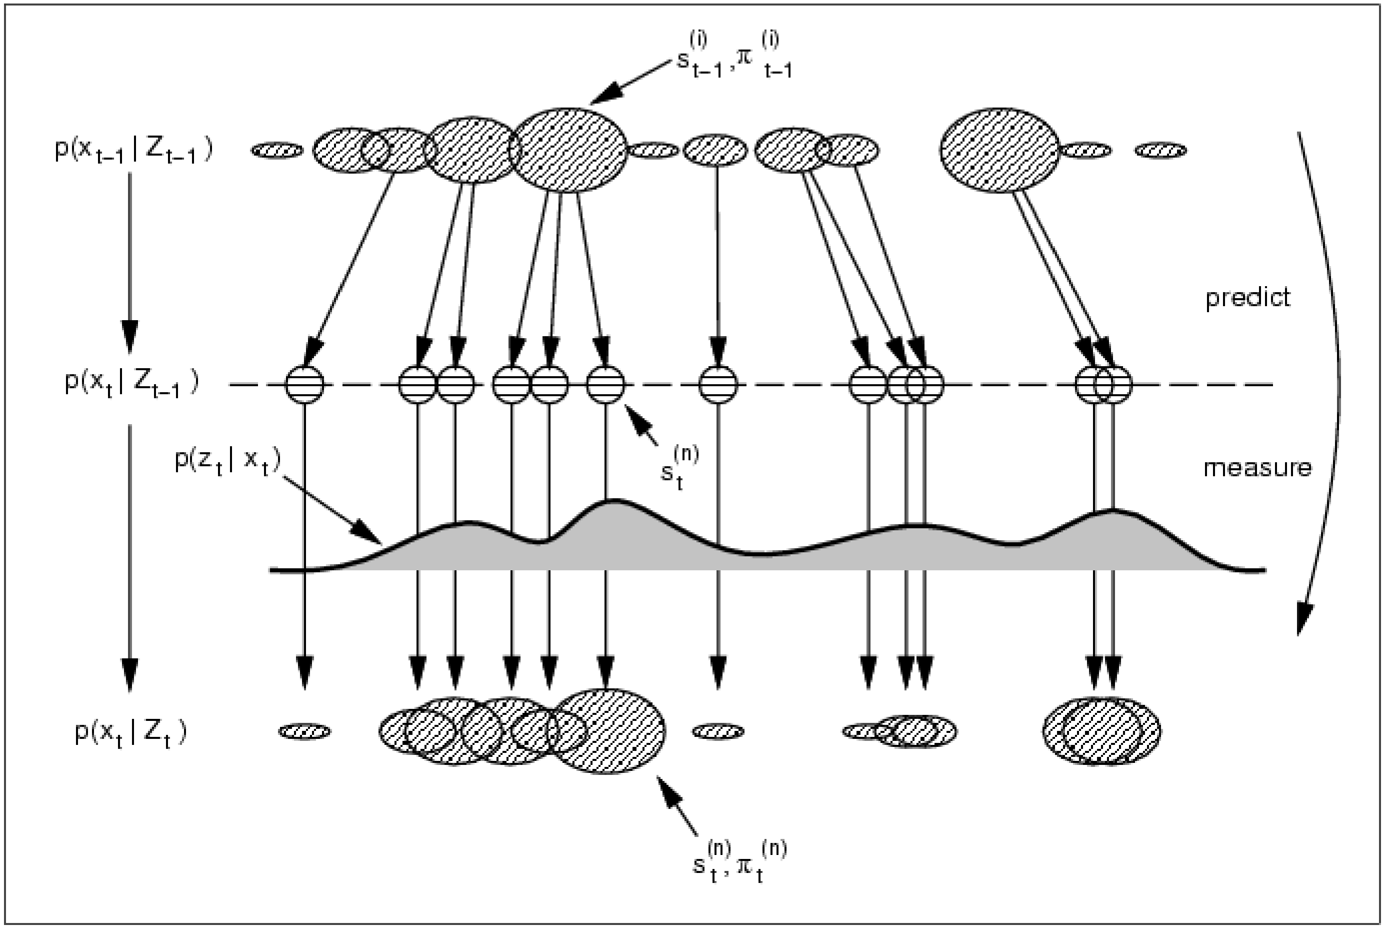

[Isard 1998]

A variant (animation)

[Isard 1998]

Applications

Tracking

- hands
- bodies
- leaves

Contour tracking

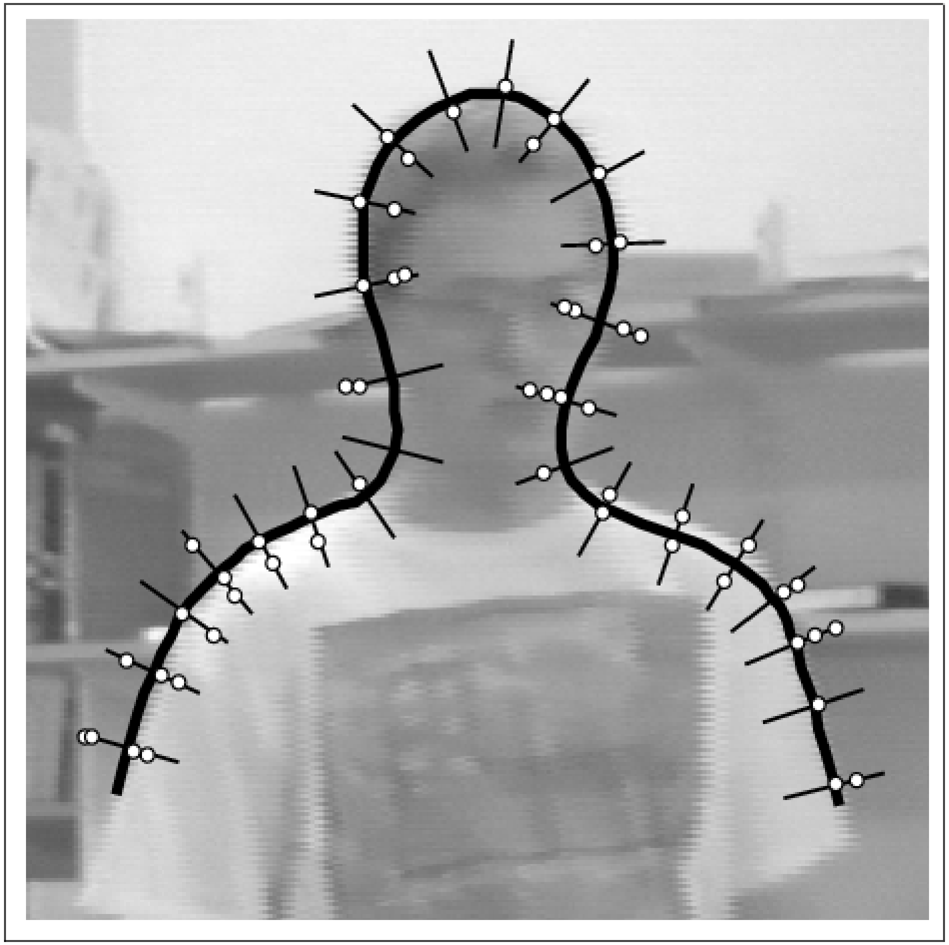

[Isard 1998]

Head tracking

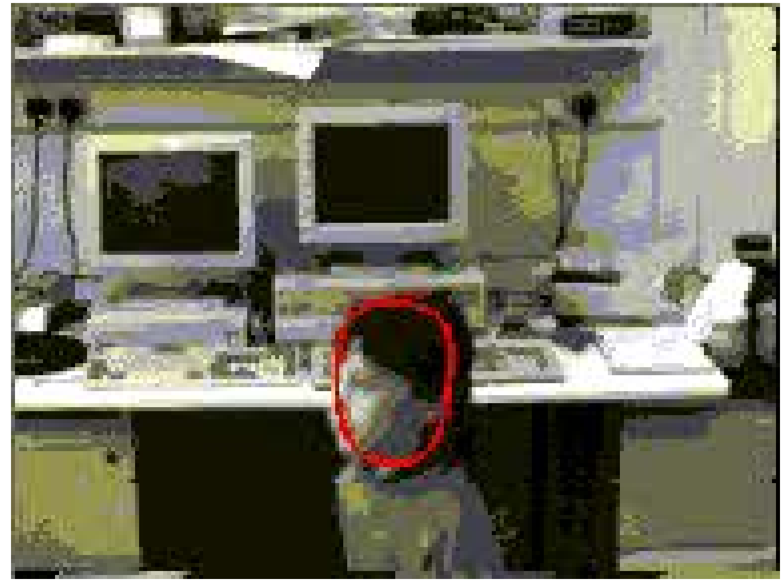

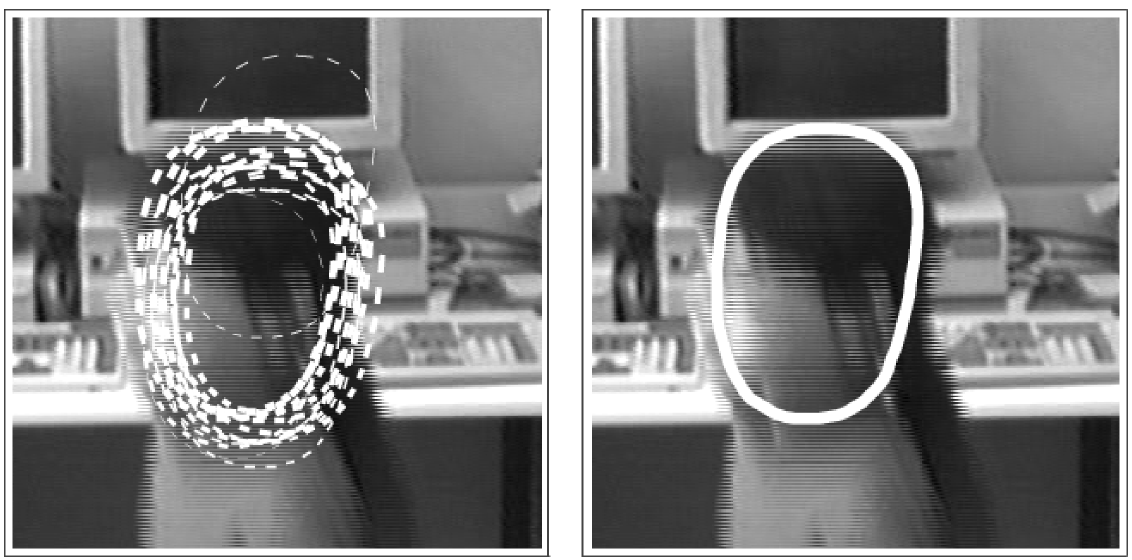

(a) (b)

[Isard 1998]

Leaf tracking

[Isard 1998]

Hand tracking

[Isard 1998]

Mixed state tracking

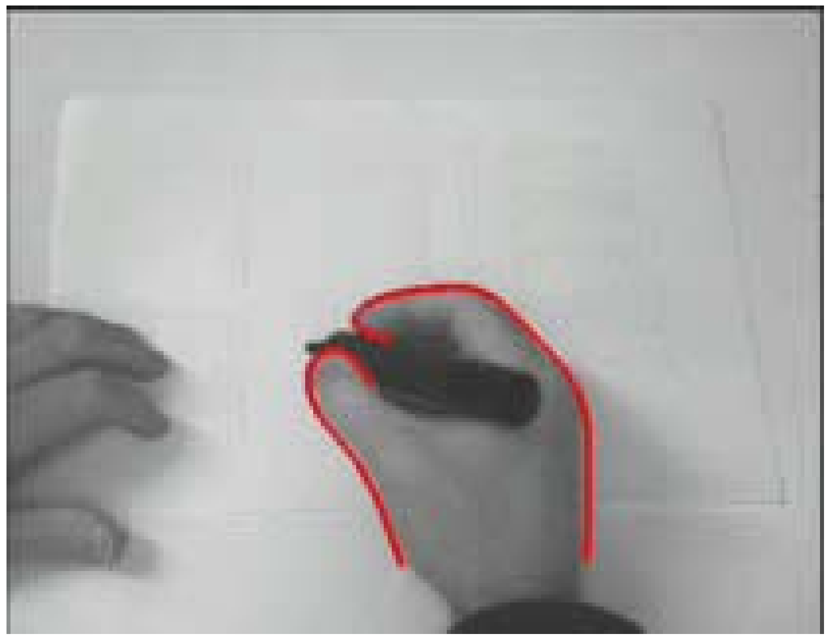

[Isard 1998]

A drawing interface

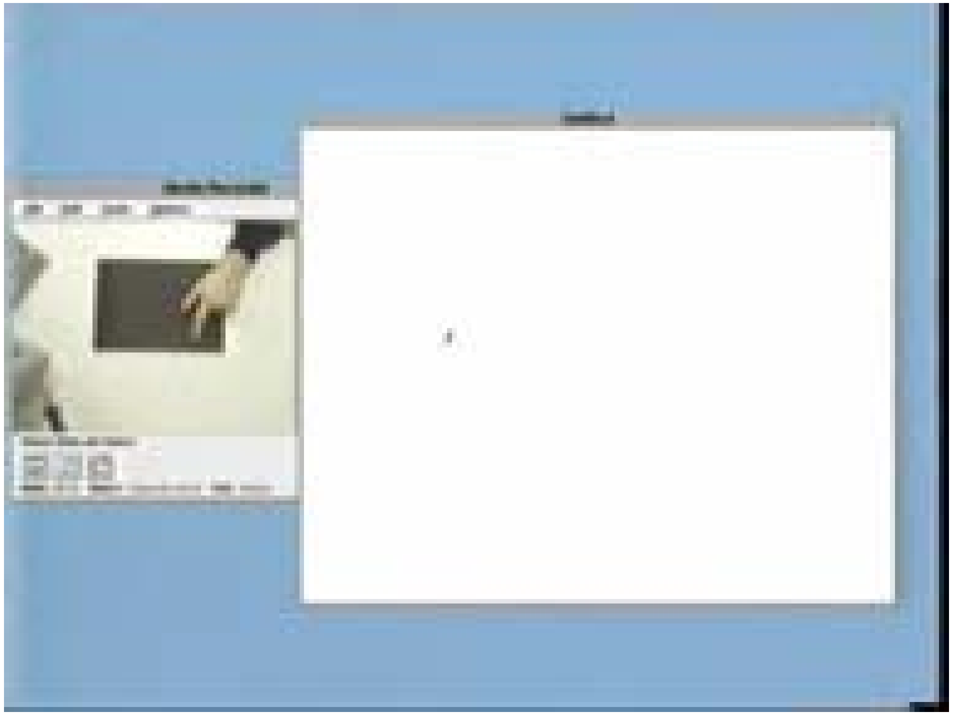

[Isard 1998]

Articulated tracking

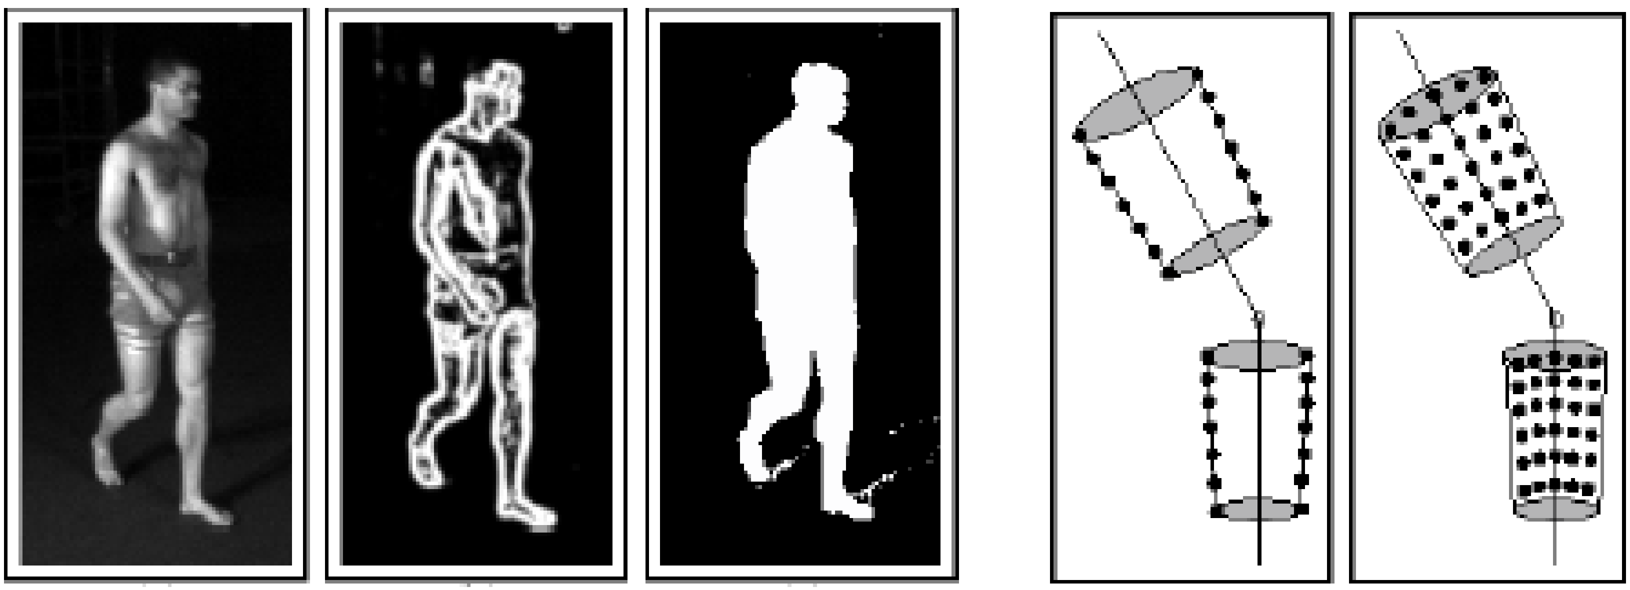

Interesting Extensions

- Multiple people / objects
  - state has to model multi-body configuration
  - resampling is tricky!

- Multiple modalities
  - fuse observation likelihoods
  - audio/visual localization and source separation

Tracking with Non-linear Dynamic Models

- Distribution propagation
- Problems with non-linearities
- Sampling densities
- Particle filtering
- Tracking people

[Figures from F&P except as noted]