Particle Filtering

KF

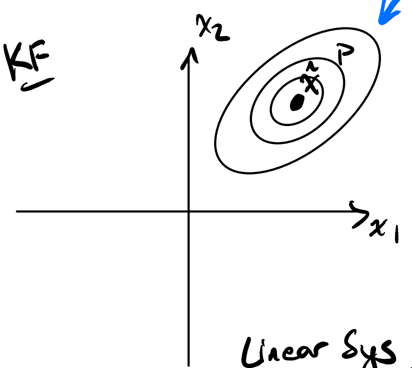

Linear Sys
Gaussian Noise

General PF

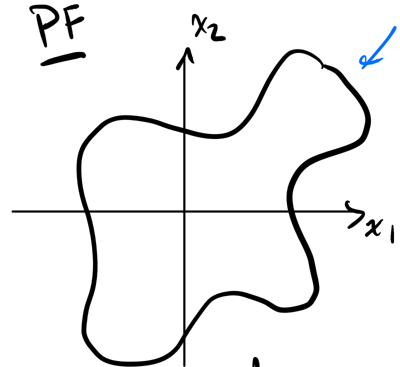

$p(x_k \mid \tilde{Y}_k)$ $\quad \hat{x}_k, P_k \leftarrow$

$e^{-\frac{1}{2}(x-\hat{x})P^{-1}(x-\hat{x})}$

$p(x_k \mid \tilde{Y}_k)$ $\quad \cancel{\hat{x}}, \cancel{P}$

$\tilde{Y}_k = (\tilde{y}_0, \tilde{y}_1, \cdots \tilde{y}_k)$

how do we frame the problem of iteratively estimating $p(x \mid \tilde{y})$ as the system evolves in time

when not linear
not Gaussian

Iterative Method $\longleftrightarrow$

Previous cond.
$p(x_k \mid \tilde{Y}_k)$

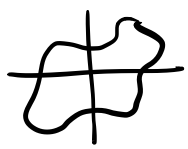

Dynamics
$p(x_{k+1} \mid x_k)$
- $x_{k+1} = f(x_k, u_k, \bar{w}_k)$
- $p(\bar{w}_k) \leftarrow$

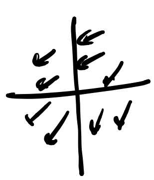

Measurement
$p(\tilde{y}_{k+1} \mid x_{k+1})$
- $\tilde{y}_{k+1} = h(x_{k+1}, v_{k+1})$
- $p(v_{k+1}) \leftarrow$

partial observation

$$\left.\right\} \; p(x_{k+1} \mid \tilde{Y}_{k+1})$$

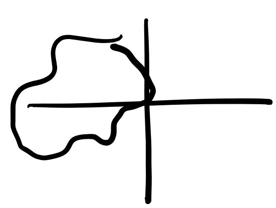

Goal! $p(x_{k+1} \mid \tilde{Y}_{k+1}) \longrightarrow$ 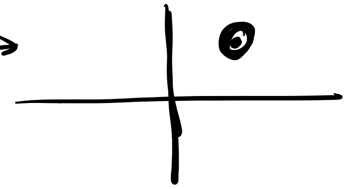

Very General Math Problem:

Bayesian Estimation, Recursive Estimation

one ex. particle filtering...

Relationship between dist..

want $p(x_{k+1} \mid \tilde{Y}_{k+1})$

apply Bayes rule to

newstate

old meas before k+1

$$p(\underbrace{x_{k+1}, \tilde{y}_{k+1}}_{\text{new meas}} \mid \tilde{Y}_k)$$

$$p(x_{k+1} \mid \tilde{y}_{k+1}, \tilde{Y}_k)\, p(\tilde{y}_{k+1} \mid \tilde{Y}_k) = p(\tilde{y}_{k+1} \mid x_{k+1}, \tilde{Y}_k)\, p(x_{k+1} \mid \tilde{Y}_k)$$

$$p(x_{k+1} \mid \tilde{y}_{k+1}, \tilde{Y}_k) = \frac{p(\tilde{y}_{k+1} \mid x_{k+1})\, p(x_{k+1} \mid \tilde{Y}_k)}{p(\tilde{y}_{k+1} \mid \tilde{Y}_k)}$$

Distribution Update:

$$\boxed{p(\underbrace{x_{k+1}}_{\text{new state}} \mid \underbrace{\tilde{Y}_{k+1}}_{\text{all meas}}) = \frac{\overbrace{p(\tilde{y}_{k+1} \mid x_{k+1})}^{\text{meas model}}\, p(x_{k+1} \mid \tilde{Y}_k)}{\int p(\tilde{y}_{k+1} \mid x_{k+1})\, p(x_{k+1} \mid \tilde{Y}_k)\, dx_{k+1}}}$$

$$p(\underline{x_{k+1}} \mid \underline{\tilde{Y}_k}) = \int \underbrace{p(x_{k+1} \mid x_k)}_{\text{dynamics}} \underbrace{p(x_k \mid \tilde{Y}_k)}_{\text{prev cond.}} \, \underline{dx_k}$$

Dynamics if no noise $p(x_{k+1} \mid x_k) = f(x_k, u_k)$

$$p(x_{k+1} \mid x_k) = \int \delta\left[x_{k+1} - f(x_k, u_k, w_k)\right] p(w_k) \, dw_k$$

$$p(\tilde{y}_{k+1} \mid x_{k+1}) = \int \delta\left(\tilde{y}_{k+1} - h_{k+1}(x_{k+1}, v_{k+1})\right) p(v_{k+1}) \, dv_{k+1}$$

Formulas are more complicated because we're tracking whole distributions as opposed to individual vectors...

Next:
- how to store these quantities in a memory efficient way...
- how to simplify different steps (when can we simplify)

1. HOW TO KEEP TRACK OF THESE DISTRIBUTIONS

$p(x)$

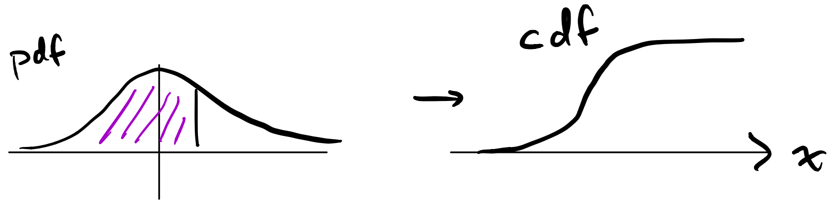

$$E_p[f(x)] = \int f(x)\, p(x)\, dx$$

Sampling

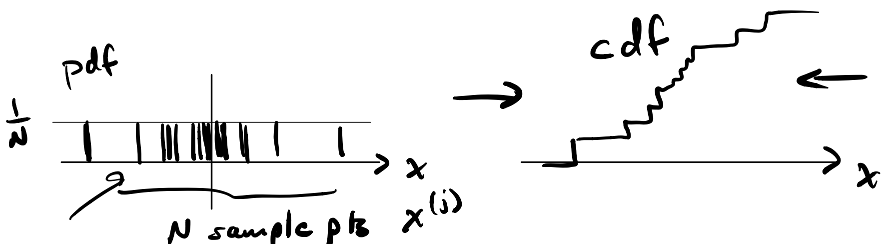

N sample pts $x^{(j)}$

if samples are drawn from $p(x)$

$$p(x) \approx \frac{1}{N} \sum_{j=1}^{N} \delta(x - x^{(j)})$$

$$E_p[f(x)] \approx \sum_{j=1}^{N} \frac{1}{N} f(x^{(j)})$$

ea sample has equal weight and is drawn from $p(x)$

Question: do we have to know $p(x)$ exactly in order to sample from it

Ex. $p(x)$

Another dist. $q(x)$ uniform

sample

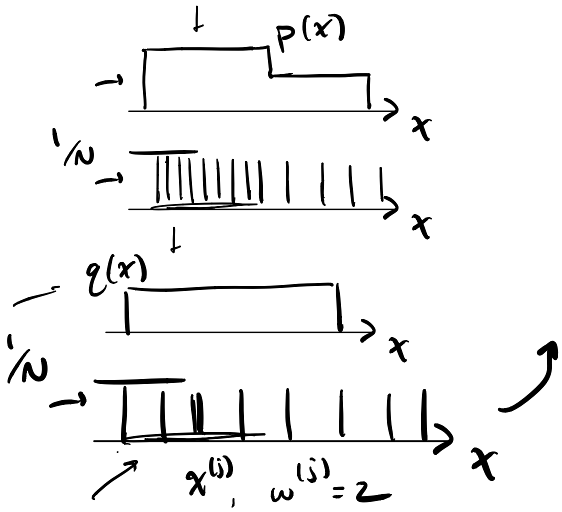

Can i use this to approx. $E_p[f(x)]$?

for ea. sample pt $x^{(j)}$, assign $w^{(j)}$

$$\underline{w^{(j)}} \underline{\propto} \frac{p(x^{(j)})}{q(x^{(j)})} \qquad \sum_{j=1}^{N} w^{(j)} = 1.$$

then

$$E_p[f(x)] \approx \sum_{j=1}^{N} \underline{w^{(j)}} f(x^{(j)})$$

← here weights are different

make $w^{(j)}$ larger for places where $p(x^{(j)}) > q(x^{(j)})$

if $w^{(j)} = 2$, think of this as counting $x^{(j)}$ twice.

we can sample from a diff distribution q → but we need the right weights.

Be careful if

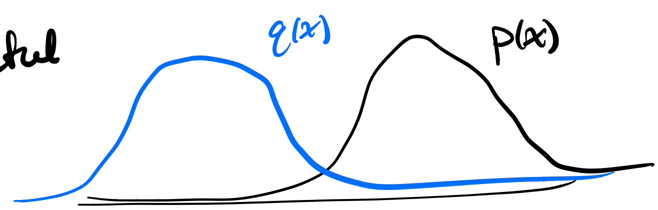

→ lead to a poor estimate

estimate is better if $q(x)$ similar to $p(x)$

Sequential Importance Sampling

$$\{(x_k^{(j)}, w_k^{(j)})\}_{j=1}^{N} \longrightarrow \{x_{k+1}^{(j)}, w_{k+1}^{(j)}\}_{j=1}^{N}$$

sample weights

how to update samples & weights...

$q(x)$: importance function

if $x_k^{(j)}$ is in a dense part of dist. → "give it a higher weight"

How to update the weights...

$$w_{k+1}^{(j)} \propto \frac{p(X_{k+1}^{(j)} \mid \tilde{Y}_{k+1})}{q(X_{k+1}^{(j)} \mid \tilde{Y}_{k+1})}$$

Time histories

$$\tilde{Y}_{k+1} = (\tilde{y}_0, \dots, \tilde{y}_{k+1})$$

$$X_{k+1}^{(j)} = (x_0^{(j)}, \dots, x_{k+1}^{(j)})$$

$$w_{k+1}^{(j)} \propto \frac{p(X_{k+1}^{(j)} \mid \tilde{Y}_k)\, p(\tilde{y}_{k+1} \mid X_{k+1}^{(j)}, \tilde{Y}_k)}{q(X_k^{(j)} \mid \tilde{Y}_k)\, q(x_{k+1}^{(j)} \mid X_k^{(j)} \tilde{Y}_{k+1})} \frac{1}{p(\tilde{y}_{k+1} \mid \tilde{Y}_k)}$$

$$\propto \frac{p(x_{k+1}^{(j)} \mid X_k^{(j)}, \tilde{Y}_k)\, p(X_k^{(j)} \mid \tilde{Y}_k)\, p(\tilde{y}_{k+1} \mid X_{k+1}^{(j)}, \tilde{Y}_k)}{q(X_k^{(j)} \mid \tilde{Y}_k)\, q(x_{k+1}^{(j)} \mid X_k^{(j)} \tilde{Y}_{k+1})} \frac{1}{p(\tilde{y}_{k+1} \mid \tilde{Y}_k)}$$

$$\propto \frac{1}{p(\tilde{y}_{k+1} \mid \tilde{Y}_k)} \frac{p(x_{k+1}^{(j)} \mid X_k^{(j)}, \tilde{Y}_k)\quad p(\tilde{y}_{k+1} \mid X_{k+1}^{(j)}, \tilde{Y}_k)}{q(x_{k+1}^{(j)} \mid X_k^{(j)} \tilde{Y}_{k+1})} \frac{p(X_k^{(j)} \mid \tilde{Y}_k)}{q(X_k^{(j)} \mid \tilde{Y}_k)}$$

doesn't depend on $x^{(j)}$

$$\propto \boxed{\frac{1}{p(\tilde{y}_{k+1} \mid \tilde{Y}_k)}} \frac{p(x_{k+1}^{(j)} \mid X_k^{(j)}, \tilde{Y}_k)\quad p(\tilde{y}_{k+1} \mid x_{k+1}^{(j)})}{q(x_{k+1}^{(j)} \mid X_k^{(j)} \tilde{Y}_{k+1})} \boxed{\frac{p(X_k^{(j)} \mid \tilde{Y}_k)}{q(X_k^{(j)} \mid \tilde{Y}_k)}}$$

if past history doesn't effect dynamics

$$p(x_{k+1}^{(j)} \mid X_k^{(j)}, \tilde{Y}_k) = p(x_{k+1}^{(j)} \mid x_k^{(j)})$$

$$q(x_{k+1}^{(j)} \mid X_k^{(j)}, \tilde{Y}_{k+1}) = q(x_{k+1}^{(j)} \mid x_k^{(j)}, \tilde{y}_{k+1})$$

$$w_{k+1}^{(j)} \propto \frac{p(x_{k+1}^{(j)} \mid x_k^{(j)})\, p(\tilde{y}_{k+1} \mid x_{k+1}^{(j)})}{q(x_{k+1}^{(j)} \mid x_k^{(j)}, \tilde{y}_{k+1})} w_k^{(j)}$$

BOOTSTRAP FILTER: ⟸ Brian Douglas video

$q(x_{k+1} \mid x_k^{(j)}, \tilde{y}_{k+1})$ dist get to pick

simple choice: $q(x_{k+1} \mid x_k^{(j)}, \tilde{y}_{k+1}) = p(x_{k+1} \mid x_k^{(j)})$

intuition

Sampling x_{k+1} without considering the new meas $\tilde{y}_{k+1}$

intuition

$$\boxed{w_{k+1}^{(j)} \propto p(\tilde{y}_{k+1} \mid x_{k+1}^{(j)})\, w_k^{(j)}}$$

if this is close to 1 $w_{k+1}^{(j)}$ gets bigger

prev weight

Summary:

track particles $x_k^{(j)}$

Prediction

$$x_{k+1}^{(j)} = f\left(x_k^{(j)}, u_k, \bar{w}_k^{(j)}\right) \qquad \bar{w}_k^{(j)} \sim p(\tilde{w}_k)$$

Update the weights

$$w_{k+1}^{(j)} = w_k^{(j)} \, \underline{p\left(\tilde{y}_{k+1} \mid x_{k+1}^{(j)}\right)}$$

in robot localization

Matching w/ lidar

$$w_{k+1}^{(j)} \leftarrow \frac{w_{k+1}^{(j)}}{\sum_{j=1}^{N} w_{k+1}^{(j)}}$$

Resample from the distribution ...