- Filtering → iteratively updating beliefs

  $bel(t) = P(x \mid z_{1:t}, u_{1:t})$

- Discrete distributions can be represented (and updated) as histograms

- What about continuous distributions

→ Particle filter

Represent $bel(x_t)$ by a set of <u>random state samples</u>

<u>Example</u>

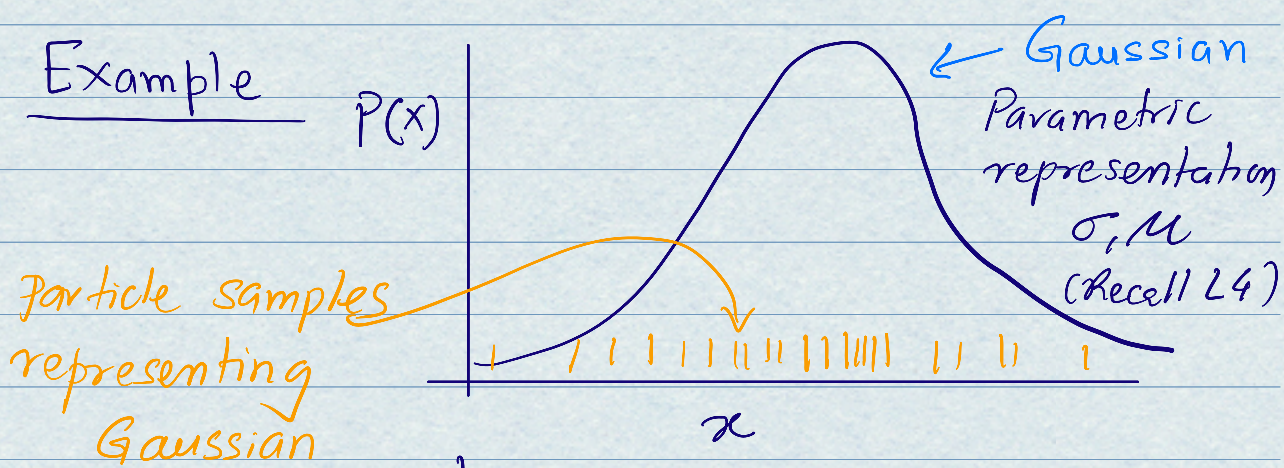

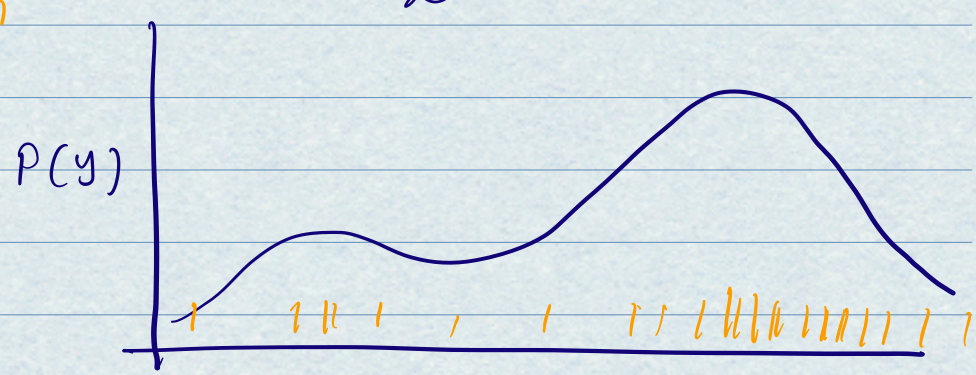

Advantages

Non parametric representation

① → Can represent a much broader set of distributions

② Can easily implement non-linear transformations of the distribution e.g. operations needed for the predict & correct stages of the filter

Def. The samples of $bel(x_t)$ distribution are called particles.

We denote M particles as

$$x_t^{[1]}, x_t^{[2]}, \ldots\ldots, x_t^{[M]} \in \mathbb{R}^n$$

Together $\mathcal{X}_t = \{x_t^{[1]}, \ldots\ldots, x_t^{[M]}\}$

Each $x_t^{[m]}$ is a concrete state at time t
$1 \leq m \leq M$

For example for the rear wheel vehicle model $x_t^{[m]} = \langle pos_x, pos_y, \theta \rangle$

$M \approx 100$

In representing $bel(x_t)$ with $\mathcal{X}_t$, ideally $x_t^{[m]}$ should be included in $\mathcal{X}_t$ with probability proportional to $bel(x_t) = P(x_t | \ldots)$

This will hold asymptotically as $M \to \infty$

Basic Particle Filtering algorithm

- Same structure as Bayes filter

  input bel$(x_{t-1}) \leftrightarrow X_{t-1}$

  control u_t measurement z_t

  output bel$(x_t) \leftrightarrow X_t$

  prediction from bel(x_{t-1}) using u_t

  Correction from $\overline{bel}(x_t)$ using z_t

State transition probability

$X_t = x_t^{[1]}, x_t^{[2]}, \dots x_t^{[M]}$ particles

Algorithm Particle_filter(X_{t-1}, u_t, z_t):

$\bar{X}_{t-1} = X_t = \emptyset$

for all m in [M] do:

importance sample $x_t^{[m]} \sim p(x_t | u_t, x_{t-1}^{[m]})$

factor $w_t^{[m]} = p\left(z_t \middle| x_t^{[m]}\right)$ → set of particles at this stage defines $\overline{bel}(x_t)$

$\overline{bel}$ $\bar{X}_t = \bar{X}_t + \langle x_t^{[m]}, w_t^{[m]} \rangle$

end for

for all m in [M] do:

draw i *with probability* $\propto w_t^{[i]}$

add $x_t^{[i]}$ *to* X_t

end for

Resampling or importance sampling "trick"

return X_t

Resampling step

"Survival of the fittest" based on measurement

We want to sample particles from $bel(x_t)$

But we only have samples from $\overline{bel}(x_t)$

More generally we want to sample from a distribution f

but we can only sample from another distribution g

How?

$$E_f[I(x \in A)] = \int f(x)\, I(x \in A)\, dx$$

$$= \int \frac{f(x)}{g(x)}^{*} \cdot g(x)\, I(x \in A)\, dx$$

Provided $f(x) > 0 \Rightarrow g(x) > 0$

$$= E_g[\omega(x)\, I(x \in A)]$$

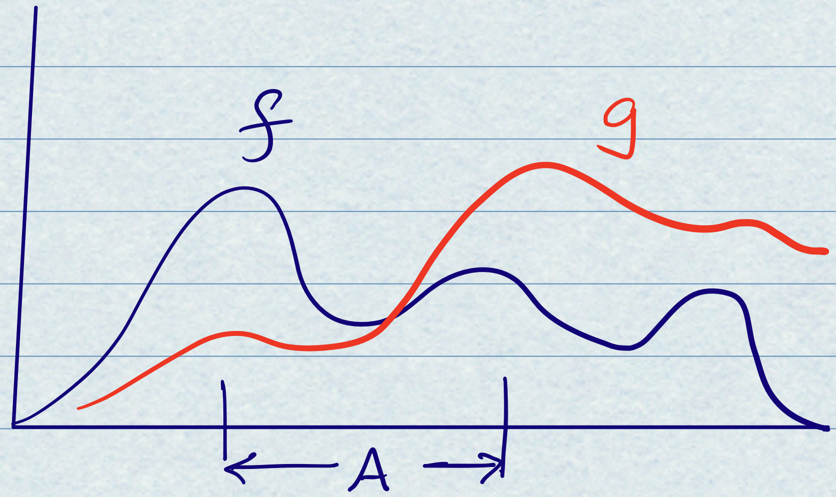

f
g
A

(5)

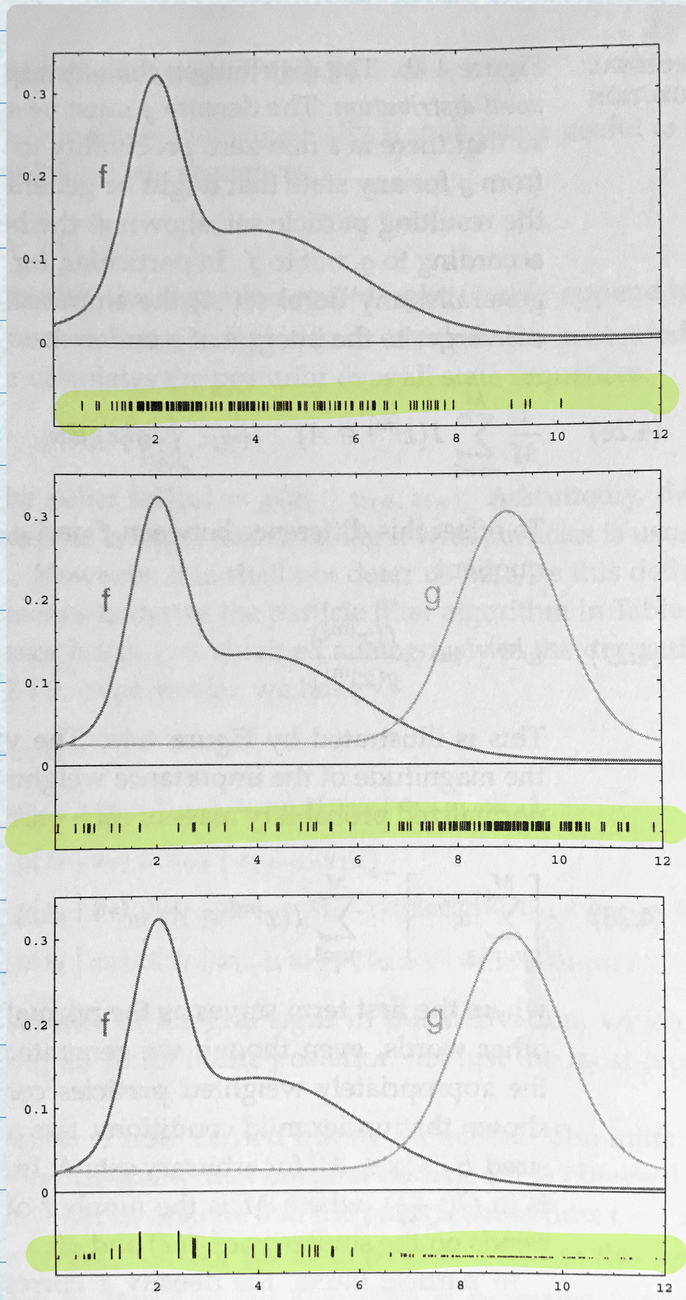

Samples from f

bel(x_t)

Samples from g

$\overline{bel}(x_t)$

Samples of f obtained by attaching a weight $\frac{f(x)}{g(x)}$ to each sample of g