Note Title 4/23/2006

Kalman Filters & Particle Filters.

Another way to deal with one-dimensional graphs.

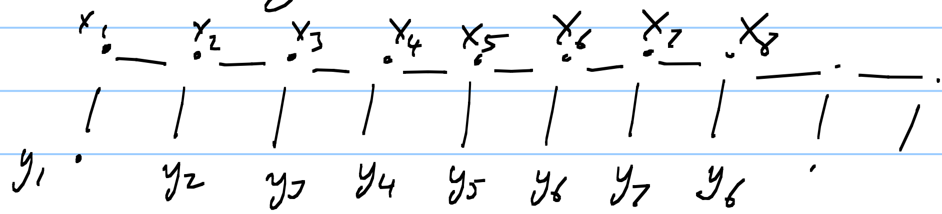

update model $p(x_{t+1} | x_t)$ $\qquad$ $p(y_t | x_t)$ observation model.

$\{x_t : t = 1, ..\}$ is the state of the system.
$\{y_t : t = 1, ..\}$ are the observations

Let $Y_t = \{y_t, y_{t-1}, y_{t-2}, \dots y_1\}$
all observations prior to time t.

Want to know
$p(x_t | Y_t)$
and update to $p(x_{t+1} | Y_{t+1})$.

$Y_{t+1} = y_{t+1}, Y_t$.

Two stages:
(i) prediction. $p(x_{t+1} | Y_t)$
(ii) Correction for new observation
$p(x_{t+1} | Y_{t+1})$.

Tracking Airplanes, Spacecraft,

Prediction

(1) $$p(x_{t+1} | Y_t) = \sum_{x_t} p(x_{t+1} | x_t)\, p(x_t | Y_t)$$

Correction.

(2) $$p(x_{t+1} | Y_{t+1}) = \frac{p(y_{t+1} | x_{t+1})\, p(x_{t+1} | Y_t)}{p(Y_{t+1})}$$

with $$p(Y_{t+1}) = \sum_{x_{t+1}} p(y_{t+1} | x_{t+1})\, p(x_{t+1} | Y_t)$$

normalization constant.

Problems:

It may be difficult to compute these two stages.

There is a very important special case - The Kalman Filter. Described in 1-D.

In this case:

$$p(y_t | x_t) = \frac{1}{\sqrt{2\pi}\,\sigma_m} e^{-(x_t - y_t)^2 / 2\sigma_m^2}$$

Gaussian Model

$$p(x_{t+1} | x_t) = \frac{1}{\sqrt{2\pi}\,\sigma_p} e^{-(x_{t+1} - x_t - \mu)^2 / 2\sigma_p^2}$$

Gaussian.

Note: for special case $\sigma_m = 0$
i.e. perfect measurements, then $\sigma_{t+1} = 0$
and $\mu_{t+1} = y_{t+1}$

Then the distributions are all Gaussian

$$p(x_t | Y_t) \sim N(\mu_t, \sigma_t)$$

$$p(x_{t+1} | Y_t) = \int dx_t \, p(x_{t+1} | x_t) \, p(x_t | Y_t)$$

$$= \int_{-\infty}^{\infty} dx_t \frac{1}{\sqrt{2\pi}\,\sigma_p} e^{-(x_{t+1}-x_t-\mu)^2/2\sigma_p^2} \frac{1}{\sqrt{2\pi}\,\sigma_t} e^{-\frac{(x_t-\mu_t)^2}{2\sigma_t^2}}$$

$$= \frac{1}{\sqrt{2\pi}\,(\sigma_p^2+\sigma_t^2)^{\frac{1}{2}}} e^{-(x_{t+1}-\mu-\mu_t)^2/2(\sigma_p^2+\sigma_t^2)}$$

$$p(x_{t+1} | Y_{t+1}) = \frac{p(y_{t+1} | x_{t+1}) \, p(x_{t+1} | Y_t)}{p(Y_t)}$$

$$p(x_{t+1} | Y_{t+1}) = N(\mu_{t+1}, \sigma_{t+1})$$

$$\mu_{t+1} = \underbrace{\mu + \mu_t}_{\text{Prediction}} - \frac{(\sigma_t^2+\sigma_p^2)}{\sigma_m^2 + (\sigma_t^2+\sigma_p^2)} \underbrace{\{(\mu+\mu_t) - y_{t+1}\}}_{\text{Correction.}}$$

$$\sigma_{t+1}^2 = \frac{\sigma_m^2(\sigma_t^2+\sigma_p^2)}{\sigma_m^2 + (\sigma_t^2+\sigma_p^2)}$$

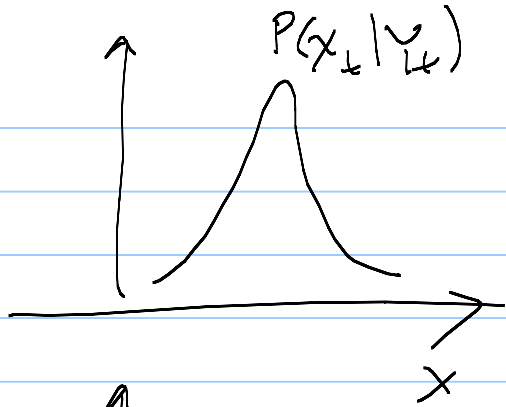

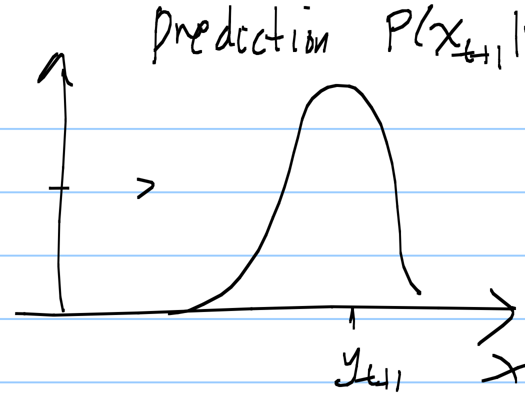

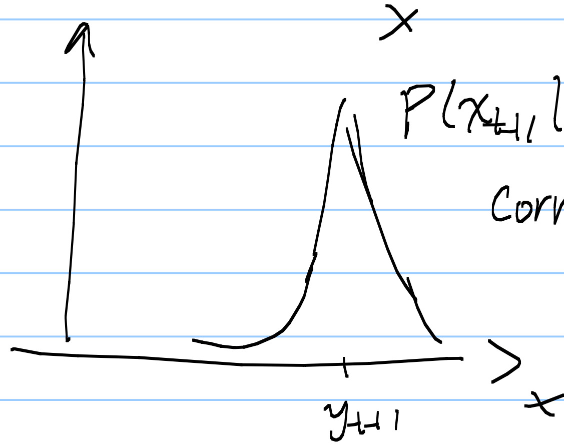

Correction due to new observation y_{t+1}.

Kalman Filter is very efficient, but requires the distributions $p(x_{t+1} | x_t)$ and $p(y_t | x_t)$ to be Gaussian.

Breakdown:

(1.) Tracking an object, and another object appears nearby. (next lecture).

(2.) x_t, y_t take finite set of values (e.g. Biology Applications)

Special Cases:

(1) Suppose $\sigma_m = 0$, i.e. the measurements are perfect.

Then it follows that:

(a) $\mu_{t+1} = (\mu+\mu_t) - (\mu+\mu_t) + y_{t+1} = y_{t+1}$

— current measurement

(b) $\sigma_{t+1} = 0$. Extreme case with perfect measurement

(2) Suppose $\sigma_p = 0$, i.e. we have perfect prediction.

Then:

(a) $$\mu_{t+1} = \frac{\sigma_t^2}{\sigma_m^2+\sigma_t^2} y_{t+1} + \frac{\sigma_m^2}{\sigma_m^2+\sigma_t^2}(\mu+\mu_t)$$ weighted average

(b) $$\sigma_{t+1}^2 = \frac{\sigma_m^2 \sigma_t^2}{\sigma_m^2+\sigma_t^2}$$

If we also have $\mu = 0$ (so x_t is constant)

then $$\mu_{t+1} = \frac{\sigma_t^2}{\sigma_m^2+\sigma_t^2} y_{t+1} + \frac{\sigma_m^2}{\sigma_m^2+\sigma_t^2} \mu_t$$

This is an incremental way to estimate the MAP for the

$$P(y_1, \dots, y_t \mid x) = \prod_{i=1}^{t} P(y_i \mid x)$$ Gaussian $N(0, \sigma_m)$

$P(x) \sim$ Gaussian $N(0, \sigma_1^2)$

Hence, reduces to MAP estimation for the static case.