Lecture 13 Two Examples of Bayes-Kalman.

Note Title 4/27/2006

Recall $Y_t = \{y_1, \ldots, y_t\}$, state x_t

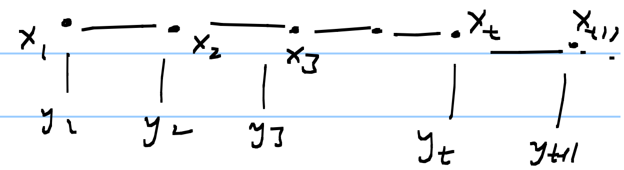

prediction

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

correction

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

$p(x_{t+1} | x_t)$

$p(y_t | x_t)$

Particle filters:

Approximate $p(x_t | Y_t)$ by samples $(x_t^1, \ldots, x_t^m)$ define recursively

update by sampling from $p(x_{t+1} | x_t^i)$ to get $\{x_{t+1}^1, \ldots, x_{t+1}^m\}$ represents $p(x_{t+1} | Y_t)$.

give each sample a weight $\omega_{t+1}^i \propto p(y_{t+1} | x_{t+1}^i)$

("$\propto$" means proportional to)

Resample from $\{x_{t+1}^1, \ldots, x_{t+1}^m\}$ with probability proportional to ω_{t+1}^i (with replacement) to get new samples

$\{x_{t+1}^{1*}, \ldots, x_{t+1}^{m*}\}$ represent $p(x_{t+1} | Y_{t+1})$

(2) **Example 1: Tracking. Gaussian Models.**

State variable $x_t = (x_{t,1}, x_{t-1,2})$

$x_{t,1}$ location $x_{t,2}$ velocity

Update: $x_{t,1} = x_{t-1,1} + x_{t-1,2} + \frac{1}{2}\omega_t$

$x_{t,2} = x_{t-1,2} + \omega_t$

ω_t i.i.d. normal $N(0, q^2)$.

Equivalently: $P(x_t | x_{t-1}) = p(x_{t:1} | x_{t-1:1}, x_{t-2:1}) P(x_{t:2} | x_{t-1:2})$.

with $p(x_{t,1} | x_{t-1,1}, x_{t-1,2}) = N(x_{t-1,1} + x_{t-1,2}, q^2/4)$

$P(x_{t,2} | x_{t-1,2}) = N(x_{t-1,2}, q^2)$

Observations $y_t = x_{t,1} + v_t \quad v_t \sim N(0, r^2)$

Note: we only observe the positions $x_{t,1}$. We cannot directly observe the velocity $x_{t:2}$.

$P(y_t | x_t) = N(x_{t,1}, r^2)$

For this example all the distributions are Gaussians. So Bayes-Kalman reduces to updating means and covariances.

Next we consider a harder problem that needs particle filters. The target object may be invisible and there may be distractor objects present.

(Page 3)

Example 2: The target object has a Gaussian prior $p(x_{t+1}|x_t) = N(x_t+\mu, \sigma_p^2)$ and the distribution of the "true" observation is Gaussian $N(x_t, r^2)$ if target is visible.

The Observation Model

observations occur in a window size Δ.

At time t, the number of observations in the window is a random variable m_t, which can be measured.

The observations are represented by a vector:

$\underline{y}_t = (y_{t,1}, \ldots, y_{t,m_t})$, where $y_{t,i}$ is position of i^{th} obs.

We assume that:

(1) the probability that target is visible is P_d;

(2) the prob dist of number of distractors is Poisson, with parameter $\lambda\Delta$.

(3) the position of a distractor is uniformly distributed in the window (i.e. prob density is $1/\Delta$)

(4) the position of target object, if present, is modeled by $N(x_t, r^2)$

Define an indicator variable:

$$I_k = \begin{cases} 0, & \text{if target is invisible} \\ k, & \text{if target is } k^{th} \text{ observation.} \end{cases}$$

$$p(\underline{y}_t | x_t, I_t = 0) = \Delta^{-m_t} \frac{(\lambda\Delta)^{m_t}}{m_t!} e^{-\lambda\Delta} = \frac{\lambda^{m_t}}{m_t!} e^{-\lambda\Delta}$$

$$p(\underline{y}_t | x_k, I_t = k) = \frac{\lambda^{m_t-1}}{(m_t-1)!} e^{-\lambda\Delta} \frac{1}{\sqrt{2\pi}\,r} \exp\left\{-\frac{(y_{t,k}-x_t)^2}{2r^2}\right\}$$

(Page 4) By assumption (1):

$$P(I_t=0) = 1-P_d, \quad P(I_t=k) = P_d/m_t.$$

<u>Hence</u> : $P(\underline{y}_t, I_t | X_t) = p(\underline{y}_t | I_t, X_t) P(I_t)$

(Note α is proportional)
$$\alpha \begin{cases} (1-P_d)\lambda, & \text{if } I_t = 0 \\ P_d \frac{1}{(2\pi)^{1/2} r} \exp\left\{-\frac{(y_{t,k}-x_t)^2}{2r^2}\right\}, & \text{if } I_k = k. \end{cases}$$

where the proportionality factor is $\frac{\lambda^{m_t-1}}{m_t!} e^{-\lambda\Delta}$

Now, I_t is unknown (it is a missing/latent/hidden variable) so we must sum it out to obtain the observation model:

$$p(\underline{y}_t | X_t) = \sum_{I_t} P(\underline{y}_t, I_t | X_t) \underset{\text{proportional to.}}{\alpha} (1-P_d)\lambda + \sum_{k=1}^{m_t} P_d \frac{1}{(2\pi)^{1/2} r} \exp\left\{-\frac{(y_{t,k}-x_t)^2}{2r^2}\right\}$$

The observation model is non-Gaussian so we cannot apply Kalman's update equations

Instead we can use particle filters. This only requires:

(i) that we can draw samples from $p(x_{t+1}|x_t)$
— easy because it is Gaussian

(ii) that we can evaluate $p(\underline{y}_t | X_t)$ up to a normalization constant. — which is also easy in this case.

<u>Note</u> that we can use particle filtering for this problem even if $p(x_{t+1}|x_t)$ is non-Gaussian, provided we can sample from it, and even if the likelihood term $p(\underline{y}_t | X_t)$ is more complicated.