Particle Filtering Application.

Note Title — 4/27/2006

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

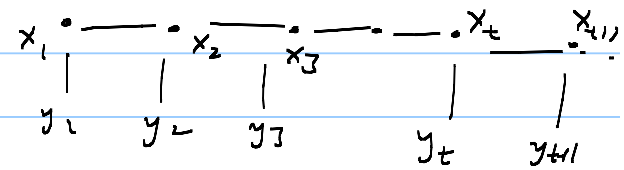

prediction

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

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:** Tracking

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

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

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

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

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

Models are Gaussian.
So updates can be expressed as Kalman filters.

E.G. only update the means & covariances

(*) the distribution of the "true" observation is Gaussian $N(x_t, r^2)$. But we don't know which observation is true.

The prediction model is Gaussian. $p(x_{t+1} | x_t) = N(x_t + \mu, \sigma_p^2)$

(3) Harder Problem:

Suppose the target object can be invisible and there may be distractor objects present. Impossible to use standard Kalman.

Notation: The observations occur in a window size Δ.

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

The observations are represented by a vector $\underline{y}_t = (y_{t,1}, \dots y_{t,m_t})$. Each $y_{t,i}$ is an observation.

Model Assumptions.

The probability that target is visible is P_d.

The distribution of the number of distractors is Poisson, with parameter $\lambda\Delta$.

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

Define an Indicator Variable $I_t = \begin{cases} 0, & \text{if target is invisible} \\ k, & \text{if } k^{th} \text{ object is the target} \end{cases}$ -

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

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

(4)

By assumption

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

So

$$p(\underline{y}_t, I_t \mid x_t) \propto \begin{cases} (1-P_d)\lambda & \text{if } I_t=0 \\ P_d \frac{1}{(2\pi)^{1/2} r} \exp\left\{-\frac{(y_{t,k}-x_t)^2}{2r^2}\right\} & \text{if } I_t=k \end{cases}$$

(we drop the common factor $\frac{\lambda^{m_t}}{(m_t-1)!} e^{-\lambda\Delta}$).

Now I_t is not known, so we have to sum it out to obtain

$$p(\underline{y}_t \mid x_t) = \sum_{I_t} p(\underline{y}_t, I_t \mid x_t)$$

$$\propto (1-P_d)\lambda + \sum_{k=1}^{m_t} P_d \frac{1}{(2\pi)^{1/2} r} \exp\left\{-\frac{(y_{t,k}-x_t)^2}{2r^2}\right\}$$

Now we know $p(x_{t+1} \mid x_t)$ & $p(\underline{y}_t \mid x_t)$ so we can apply particle/bootstrap filtering.

This requires the ability to draw samples from $p(x_{t+1} \mid x_t)$ — easy, because this is a Gaussian

And to evaluate the terms proportional to $p_t(\underline{y}_t \mid x_t)$ — straightforward, as above (see lecture notes 8).

Particle Filters will work, but there is a better way.