Tutorial On Sequential Sampling Methods

Javier R. Movellan

Let $H = (H_0, H_1, \cdots)$ a stochastic process representing some (hidden) state dynamics, $O = (O_0, O_1, \cdots)$ represent some observable dynamics. Let $\bar{H}_t = (H_0, \cdots, H_t)$, $\bar{O}_t = (O_0, \cdots, O_t)$ for $t = 1, 2, \cdots$. For a fixed sequence $x = (x_1, x_2, \cdots)$ we let $\bar{x}_t = (x_1, \cdots, x_t)$. To simplify the presentation we identify probability density functions by their arguments. For example the notation $p(o_t|h_t)$ stands for $p_{O_t|H_t}(o_t|h_t)$. Moreover we gloss over differences between continuous and discrete random variables by accepting' delta functions as proper probability density functions. The joint process (H, O) is assumed to have the following Markovian properties:

- System dynamics:
$$p(h_t \mid \bar{h}_{t-1}, \bar{o}_{t-1}) = p(h_t \mid h_{t-1}) \text{ for all } \bar{h}_t \in \mathbb{R}^t, \bar{o}_{t-1} \in \mathbb{R}^{t-1}, \tag{1}$$
- Observation dynamics:
$$p(o_t \mid \bar{h}_t, \bar{o}_{t-1}) = p(o_t \mid h_t) \text{ for all } \bar{h}_t \in \mathbb{R}^t, \bar{o}_t \in \mathbb{R}^t. \tag{2}$$

0.1 Forward Recursion Equation

Suppose we are given an observation sequence $\bar{o} = (o_1, o_2, \cdots)$. Our goal is to get an estimate of $p(h_t \mid \bar{o}_t)$ for $t = 0, 1, \cdots$. This would allow us to make inferences about the hidden process based on the observed sequence. First suppose that we know $p(h_{t-1} \mid \bar{o}_{t-1})$, the following recursion equation allows us to get $p(h_t \mid \bar{o}_t)$

$$p(h_t \mid \bar{o}_t) = \frac{p(\bar{o}_{t-1})}{p(\bar{o}_t)} p(o_t \mid h_t) \int dh_{t-1}\, p(h_{t-1} \mid \bar{o}_{t-1}) p(h_{t-1} \mid h_t). \tag{3}$$

Proof:

$$p(h_t \mid \bar{o}_t) = \frac{p(h_t, o_t, \bar{o}_{t-1})}{p(\bar{o}_t)} = \frac{p(\bar{o}_{t-1})}{p(\bar{o}_t)} p(h_t, o_t \mid \bar{o}_{t-1}) \tag{4}$$

$$= \frac{p(\bar{o}_{t-1})}{p(\bar{o}_t)} \int dh_{t-1}\, p(h_t, h_{t-1}, o_t \mid \bar{o}_{t-1}) \tag{5}$$

$$= \frac{p(\bar{o}_{t-1})}{p(\bar{o}_t)} \int dh_{t-1}\, p(h_{t-1} \mid \bar{o}_{t-1}) p(h_t \mid \bar{o}_{t-1}, h_{t-1}) p(o_t \mid \bar{o}_{t-1}, h_{t-1}, h_t) \tag{6}$$

$$= \frac{p(\bar{o}_{t-1})}{p(\bar{o}_t)} \int dh_{t-1}\, p(h_{t-1} \mid \bar{o}_{t-1}) p(h_t \mid h_{t-1}) p(o_t \mid h_t). \tag{7}$$

where in the last step we used the Markovian properties of the process.

□

0.2 Sequential Sampling

We will now use the forward recursion equation to devise a sequential Monte-Carlo sampling scheme that will give us estimates of $p(h_t \mid o_t)$ for all t. We represent probability estimates using hats ($\hat{}$).

Initialization: We get an estimate of $p(h_0 \mid o_0)$ by obtaining n i.i.d. random samples $h_0^{(1)}, \cdots, h_0^{(n)}$ from p_{H_0} and defining

$$\hat{p}(h_0 \mid o_0) = \frac{\sum_{i=1}^n \delta(h_0 - h_0^{(i)})\, p(o_0 \mid h_0^{(i)})}{\sum_{j=1}^n p(o_0 \mid h_0^{(j)})} \text{ for all } h_0 \in \mathbb{R}. \tag{8}$$

Note we are modeling the probability density function $p_{H_0|O_0}$ as a sum of delta functions (spikes) centered at the n i.i.d. samples. Each spike has strength proportional to the posterior probability of the observation given the sampled hidden state.

Recursion: Assuming we have $\hat{p}_{H_{t-1}|\bar{O}_{t-1}}$ we can get an estimate of $\hat{p}_{H_t|\bar{O}_t}$ using the forward recursion equation:

- Get n i.i.d. samples $\tilde{h}_{t-1}^{(1)}, \cdots, \tilde{h}_{t-1}^{(n)}$ from $\hat{p}_{H_{t-1}|\bar{O}_{t-1}}$.
- For each $\tilde{h}_{t-1}^{(i)}$ get a sample $h_t^{(i)}$ from $p_{H_t|H_{t-1}}(\cdot \mid \tilde{h}_{t-1}^{(i)})$. This results in n samples $h_t^{(1)}, \cdots, h_t^{(n)}$ from $\hat{p}_{H_t|\bar{O}_{t-1}}$.
- The estimate of $p_{H_t|\bar{O}_t}$ is defined as follows

$$\hat{p}(h_t \mid \bar{o}_t) = \frac{\sum_{i=1}^n \delta(h_t - h_t^{(i)})\, p(o_t \mid h_t^{(i)})}{\sum_{j=1}^n p(o_t \mid h_t^{(j)})} \text{ for all } h_t \in \mathbb{R}. \tag{9}$$

Notes: The sampling scheme requires we weight delta functions centered at $h_t^{(i)}$ by the the value of $p(o_t \mid h_t^{(i)})$. In practice we just need a number proportional to that value. Let $w(h_t, o_t) = k(o_t)p(o_t \mid h_t)$, where the proportionality constant $k(o_t)$ is independent of h_t. Then (9) can be modified as follows:

$$\hat{p}(h_t \mid \bar{o}_t) = \frac{\sum_i \delta(h_t - h_t^{(i)})\, w(h_t^{(i)}, o_t)}{\sum_{j=1}^n w(h_t^{(j)}, o_t)} \text{ for all } h_t \in \mathbb{R}. \tag{10}$$

This is of interest since in some cases it is easier to obtain a model of $p(h_t \mid o_t)$ than a model of $p(o_t \mid h_t)$. For example, neural networks can be trained to provide estimates of $p(h_t \mid o_t)$, i.e., for a given input o_t to the neural network the output can be interpreted as an estimate of the posterior probability of the state given the observation o_t. Using Bayes rule we have

$$p(o_t \mid h_t) = k(o_t)w(h_t, o_t), \tag{11}$$

where

$$k(o_t) = p(o_t), \tag{12}$$

$$w(h_t, o_t) = p(h_t \mid o_t)/p(h_t). \tag{13}$$

Here $p(h_t \mid o_t)$ is provided by the neural network and $p(h_t)$ can be interpreted as a model of the prior probability of the states.

0.3 Importance Sampling

In the previous sampling scheme the samples $h_t^{(1)}, \cdots, h_t^{(n)}$ are taken from $\hat{p}_{H_t|\bar{O}_{t-1}}(\cdot \mid \bar{o}_{t-1})$. To increase the efficiency of our estimates we may want to sample from another distribution $g_t(\cdot)$ and compensate by multiplying each sample by $\hat{p}_{H_t|\bar{O}_{t-1}}(\cdot \mid \bar{o}_{t-1})/g_t(\cdot)$. In particular let

$$\hat{p}(h_{t-1} \mid \bar{o}_{t-1}) = \sum_{i=1}^n \delta(h_{t-1} - h_{t-1}^{(i)})\, w_{t-1}(h_{t-1}^{(i)}, o_{t-1}). \tag{14}$$

Then

$$\hat{p}(h_t \mid \bar{o}_{t-1}) = \int dh_{t-1}\, \hat{p}(h_t \mid \bar{o}_{t-1})p(h_t \mid h_{t-1}) = \sum_{i=1}^n w(h_{t-1}^{(i)}, o_{t-1})p(h_t \mid h_{t-1}^{(i)}). \tag{15}$$

Now we sample $h_t^{(1)}, \cdots, h_t^{(n)}$ from $g_t(\cdot)$ to get

$$\hat{p}(h_t \mid \bar{o}_t) = \sum_{i=1}^n \delta(h_t - h_t^{(i)})w_t(h_t^{(i)}, o_t) \tag{16}$$

where

$$w_t(h_t, o_t) = p(o_t \mid h_t)\hat{p}(h_t \mid \bar{o}_{t-1})/g_t(h_t). \tag{17}$$

1 History

- The first version of this document was written by Javier R. Movellan in 1996 and used in one of the courses he taught at the Cognitive Science Department at UCSD.
- The document was made open source under the GNU Free Documentation License Version 1.2 on October 9 2003, as part of the Kolmogorov project.