UPPSALA
UNIVERSITET

Sequential Monte Carlo methods

Lecture 6 – Auxiliary particle filters

Thomas Schön, Uppsala University

2021-08-17

Outline – Lecture 6

Aim: Show how we can improve the proposals for the particle filter by using auxiliary variables.

Outline:

1. Auxiliary variables
2. Ancestor indices as auxiliary variables
3. Improving the proposal distributions

Auxiliary variables

Target distribution: $\pi(x)$, difficult to sample from

Idea: Introduce another variable U with conditional distribution $\pi(u \mid x)$

The joint distribution $\pi(x, u) = \pi(u \mid x)\pi(x)$ admits $\pi(x)$ as a marginal by construction, i.e., $\int \pi(x, u)\mathrm{d}u = \pi(x)$.

Sampling from the joint $\pi(x, u)$ may be easier than directly sampling from the marginal $\pi(x)$!

The variable U is an **auxiliary variable**. It may have some "physical" interpretation (an unobserved measurement, unknown temperature, ...) but this is not necessary.

ex) Auxiliary variables

Importance sampling setting with **target** $\pi(x)$ and **proposal** $q(x)$. We **assume** $\widetilde{\pi}(x) \leq q(x)$.

To sample from $\pi(x)$, introduce an auxiliary variable

$$U \mid (X = x) \sim \mathcal{U}(0, \widetilde{\pi}(x)).$$

Joint target: $\pi(u, x) = \pi(u \mid x)\pi(x) = \mathcal{U}(u \mid 0, \widetilde{\pi}(x))\pi(x)$
Joint proposal: $q(u, x) = q(u \mid x)q(x) = \mathcal{U}(u \mid 0, q(x))q(x)$

The weights are,

$$\frac{\mathcal{U}(u \mid 0, \widetilde{\pi}(x))}{\mathcal{U}(u \mid 0, q(x))} \frac{\widetilde{\pi}(x)}{q(x)} = \mathbb{1}(u \leq \widetilde{\pi}(x)) \frac{q(x)}{\widetilde{\pi}(x)} \frac{\widetilde{\pi}(x)}{q(x)} = \mathbb{1}(u \leq \widetilde{\pi}(x))$$

In fact, conditionally on $\widetilde{w}^i = 1$, a sample x^i is an exact draw from $\pi(x)$ — referred to as **rejection sampling.**

Auxiliary particle filter

Sampling from the joint proposal

Sampling from the **joint proposal** $q(x_t, a_t \mid y_{1:t}) = \nu_{t-1}^{a_t} q(x_t \mid x_{t-1}^{a_t}, y_t)$:

1. Sample the auxiliary variable (**resampling**),

$$a_t^i \sim \mathcal{C}(\{\nu_{t-1}^j\}_{j=1}^N).$$

2. Sample x_t conditionally on the auxiliary variable (**propagation**),

$$x_t^i \sim q(x_t \mid x_{t-1}^{a_t^i}, y_t).$$

Repeat N times for $i = 1, \ldots, N$.

Algorithmically, sampling from the proposal is done exactly in the same way as before!

Computing the weights

The importance weights are given by the ratio between the **joint** target and **joint** proposal,

$$\widetilde{w}_t^i = \frac{w_{t-1}^{a_t^i}}{\nu_{t-1}^{a_t^i}} \frac{p(y_t \mid x_t^i) p(x_t^i \mid x_{t-1}^{a_t^i})}{q(x_t^i \mid x_{t-1}^{a_t^i}, y_t)}, \qquad i = 1, \ldots, N.$$

The weights can be computed in $O(N)$ computational time for quite arbitrary choices of $\{\nu_{t-1}^i\}_{i=1}^N$ and $q(\cdot)$.

Note that the resampling weights $\{\nu_{t-1}^i\}_{i=1}^N$

- can be different from the importance weights $\{w_{t-1}^i\}_{i=1}^N$,
- may depend on $\{x_{t-1}^i\}_{i=1}^N$ as well as on y_t.

Auxiliary particle filter

Algorithm 1 Auxiliary particle filter (for $i = 1, \ldots, N$)

1. **Initialization ($t = 0$):**
   (a) Sample $x_0^i \sim p(x_0)$.
   (b) Set initial weights: $w_0^i = 1/N$.
2. **for** $t = 1$ **to** T **do**
   (a) **Resample:** sample ancestor indices $a_t^i \sim \mathcal{C}(\{\nu_{t-1}^j\}_{j=1}^N)$.
   (b) **Propagate:** sample $x_t^i \sim q(x_t \mid x_{t-1}^{a_t^i}, y_t)$.
   (c) **Weight:** compute

$$\widetilde{w}_t^i = \frac{w_{t-1}^{a_t^i}}{\nu_{t-1}^{a_t^i}} \frac{p(y_t \mid x_t^i) p(x_t^i \mid x_{t-1}^{a_t^i})}{q(x_t^i \mid x_{t-1}^{a_t^i}, y_t)}$$

   and normalize $w_t^i = \widetilde{w}_t^i / \sum_{j=1}^N \widetilde{w}_t^j$.

Selecting the proposals

How do we select the proposals?

There is freedom in selecting the resampling weights $\{\nu_{t-1}^i\}_{i=1}^N$ and proposal $q(\cdot)$. **How are they chosen in practice?!**

With $\nu_{t-1}^i = w_{t-1}^i$ and $q(x_t \mid x_{t-1}, y_t) = p(x_t \mid x_{t-1})$ we recover exactly the bootstrap particle filter.

Is it possible to select the proposals so that $w_t^i \equiv \frac{1}{N}$?

An alternative motivation for the APF

Recall that the bootstrap PF makes use of the following unnormalized target distribution

$$\widetilde{\pi}(x_t) = p(y_t \mid x_t) \sum_{i=1}^{N} w_{t-1}^{i} p(x_t \mid x_{t-1}^{i})$$

and propose particles according to

$$q(x_t \mid x_t^{i}) = \sum_{i=1}^{N} w_{t-1}^{i} p(x_t \mid x_{t-1}^{i})$$

Problem: If the observation y_t is informative, the target and the proposal will be quite different, which results in unbalanced weights.

Reduce the problem by searching for a proposal distribution that is closer to the target, which leads to the introduction of an auxiliary variable A_t representing the mixture index.

A few concepts to summarize lecture 6

Auxiliary variable: a variable by which the target distribution is extended to improve efficiency or enable sampling from the target.

Ancestor index: auxiliary variable used in the particle filter, representing one of the components in the mixture target distribution.

Auxiliary particle filter: particle filter explicitly using the ancestor indices as auxiliary variables.