Sequential Monte Carlo methods

Discussion Seminar 1

2021-08-18

Importance sampling proposals

In importance sampling, . . .

1. . . . why is it required that the support of the proposal covers the support of the target? Provide an intuitive explanation. What would happen otherwise?
2. . . . why is it important that the proposal has at least as “thick tails” as the target? Provide an intuitive explanation. What would happen otherwise?
3. . . . for a specific target function φ what would be the optimal proposal?

What is estimated?

Assume that we are running a **bootstrap particle filter** for some **state-space model**, with target function φ.

For each of the following expressions, what is being estimated? If any of the expressions estimate the same thing, which one you would prefer and why?

1. $\sum_{i=1}^{N} w_{t-1}^{i} \varphi(x_{t-1}^{i})$
2. $\frac{1}{N} \sum_{i=1}^{N} \varphi(x_{t-1}^{a_t^i})$
3. $\frac{1}{N} \sum_{i=1}^{N} \varphi(x_t^i)$
4. $\frac{1}{N} \sum_{i=1}^{N} \tilde{w}_t^i$
5. $\sum_{i=1}^{N} w_t^i \varphi(x_t^i)$

Deterministic transition

Consider a nonlinear state space model,

$$X_t = f(X_{t-1}, \theta),$$
$$Y_t = g(X_t, \theta) + E_t,$$

where the state transition is deterministic (there is no process noise). Assume that we use a **bootstrap particle filter** for this model.

1. What will the propagation step of the filter look like?
2. Do you see any potential issues with the filter? In particular, what will happen as t becomes large?
3. How does this relate to the stability of the particle filter discussed during lecture 5?

Deterministic observation

Consider a nonlinear state space model,

$$X_t = f(X_{t-1}, \theta) + V_t,$$
$$Y_t = g(X_t, \theta),$$

where the observations are deterministic (there is no measurement noise). Assume that we use a **bootstrap particle filter** for this model.

1. What will the weighting step of the filter look like?
2. Do you see any potential issues with the filter?
3. Assume instead that we have **very small** measurement noise.

$$Y_t = g(X_t, \theta) + E_t, \qquad \text{Var}(E_t) \approx 0.$$

   Answer the above questions again, can you think of any solutions to the problems?

Auxiliary variables

Consider a target distribution $\pi(x)$ of interest. We introduce an auxiliary variable u, so that the joint distribution for the pair (x, u) is $\pi(x, u)$.

What are the basic requirements that we put on $\pi(x, u)$ for this to be a **valid** and **useful** approach?

Likelihood estimate

For the bootstrap particle filter we have the likelihood estimate

$$\widehat{z} \approx \prod_{t=1}^{T} \left(\frac{1}{N} \sum_{i=1}^{N} \widetilde{w}_t^i \right).$$

1. Look at the Sequential Importance Sampler (end of lecture 3). What would the likelihood estimator be for this algorithm? (*hint: remember that the likelihood is the normalizing constant*)
2. Using this result, how would you estimate the likelihood if **adaptive resampling** is used? (Algorithm 3 lecture 8)