Stochastic Filtering (SS2016) Exercise Sheet 10

Lecture and Exercises: JProf. Dr. Philipp Harms
Due date: July 8, 2016

10.1. Monte Carlo simulation

Let X^i, $i \in \mathbb{N}$, be i.i.d. copies of a random variable X and let f be a bounded function. Then $\frac{1}{N}\sum_{i=1}^N f(X^i)$ is called Monte Carlo estimator for $\mathbb{E}[f(X)]$.

a) Show that the Monte Carlo estimator converges of order $\sqrt{N}$ with respect to the L^2 norm $\|\cdot\|_{L^2} = \sqrt{\mathbb{E}[(\cdot)^2]}$ as $N \to \infty$.

b) Is the estimator consistent (this means convergence in probability)? Is it strongly consistent (this means convergence almost surely)? Is it unbiased (this means that it has the correct mean)?

Remark: The same rate of convergence is retained by the particle filters below.

10.2. A naïve particle filter

Assume the following general setup: (X,Y) is a HMM with state kernel P, observation kernel $K = \lambda\phi$, and initial distribution $X_0 \sim \mu$; $\mathbb{P}$ is the law of (X,Y), and $\tilde{P}$ is the law of the reference HMM where λ is replaced by 1. Let X^i, $i \in \mathbb{N}$, be i.i.d. copies of the process X, called "particles." Then Exercise 10.1 suggests to use the estimators

$$\sigma_n(f) = \tilde{\mathbb{E}}\left[f(X_n) \prod_{k=0}^{n} \lambda(X_n, Y_n) \middle| \mathscr{F}_n(Y) \right] \approx \frac{1}{N}\sum_{i=1}^{N} f(X_n^i) \prod_{k=0}^{n} \lambda(X_k^i, Y_k),$$

$$\pi_n(f) = \frac{\sigma_n(f)}{\sigma_n(1)} \approx \frac{\sum_{i=1}^N f(X_n^i) \prod_{k=0}^n \lambda(X_k^i, Y_k)}{\sum_{i=1}^N \prod_{k=0}^n \lambda(X_k^i, Y_k)},$$

for any bounded measurable function f.

Note: the estimator of $\pi_n(f)$ can be written as $\sum_{i=1}^{N} w_n^i f(X_n^i)$ with weights

$$w_n^i = \frac{\prod_{k=0}^{n} \lambda(X_k^i, Y_k)}{\sum_{i=1}^{N} \prod_{k=0}^{n} \lambda(X_k^i, Y_k)}, \qquad i = 1, \dots, N.$$

Thus, the particle filter can be seen as a kind of important sampling, where the weights adjust the reference probability to the observations.

a) Are the estimators for $\sigma_n(f)$ and $\pi_n(f)$ (strongly) consistent? Are they unbiased?

b) The sequential importance sampling (SIS) algorithm is an implementation of the above particle filter where the weights are calculated recursively in n (c.f. [1, Algorithm 4.1]). Implement the SIS algorithm for the HMM of Exercise 4.2. Plot the mean and 95% confidence intervals of the particle filter and compare your results to the exact Kalman filter. What rate of convergence do you observe as you increase the number of particles?

c) Plot the weights of the particle filter to get some intuition for the following exercise.

10.3. Bootstrap particle filter

For medium to large n, the weights of the particle filter in Exercise 10.2 quickly become degenerate, meaning that all but one of the weights will be nearly zero. This renders the particle filter inefficient because the Monte Carlo approximation effectively consists of one sample rather than N samples.

This problem is addressed by the sequential importance sampling/resampling (SIS-R) algorithm [1, Algorithm 4.2], which is also called bootstrap particle filter. The SIS-R algorithm is a modification of the SIS algorithm where at each step n, the particles $X_n^1, \dots, X_n^N$ of the SIS algorithm are re-sampled from the empirical distribution $\sum_{i=1}^{N} \delta_{X_n^i}/N$ and the weights are reset to $1/N$.

a) Explain how to generate random draws from an empirical distribution $\sum_{i=1}^{N} \delta_{X_n^i}/N$.

b) Implement the SIS-R algorithm for the HMM of Exercise 4.2. Do you see improvements compared to Exercise 10.2?

References

[1] Ramon van Handel. *Hidden Markov Models*. Lecture Notes. Princeton University, 2008.