

Exercises set II

PhD course on Sequential Monte Carlo methods 2017

Department of Information Technology, Uppsala University

August 27, 2019

This document contains exercises to make you familiar with the content of the course. *The exercises in this document are not mandatory, and you do not need to hand in your solutions.* The mandatory assignment is found in a separate document named "Hand-in". We strongly recommend that you carefully work through these exercises before starting with the mandatory assignments.

II.1 Likelihood estimates for the stochastic volatility model


Consider again (cf. I.4) the stochastic volatility model

$$x_t | x_{t-1} \sim \mathcal{N}(x_t; \phi x_{t-1}, \sigma^2), \quad (1a)$$

$$y_t | x_t \sim \mathcal{N}(y_t; 0, \beta^2 \exp(x_t)), \quad (1b)$$

where the parameter vector is given by $\theta = \{\phi, \sigma, \beta\}$ and the data is found in `seOMXlogreturns2012to2014.csv`.

- (a) Let β be unknown, and assume the other parameters are $\phi = 0.98$ and $\sigma = 0.16$. Make a reasonably coarse grid for β between 0 to 2, and implement the bootstrap particle filter to estimate the likelihood for each of these values of β . Run the particle filter 10 times for every parameter combination, and present the result as a box plot similar to this:



For numerical reasons, it is usually better to consider the *log* likelihood, i.e., the logarithm of (10)

$$\log \hat{p}(y_{1:T}) = \log \prod_{t=1}^T \frac{1}{N} \sum_{i=1}^N \underbrace{p(y_t | x_t^i)}_{\tilde{w}_t^i} = \sum_{t=1}^T \left(\log \sum_{i=1}^N \tilde{w}_t^i - \log N \right). \quad (2)$$

It is, however, important to realize that $\mathbb{E}[\hat{p}(y_{1:T})] = p(y_{1:T})$ does *not* imply $\mathbb{E}[\log \hat{p}(y_{1:T})] = \log p(y_{1:T})$!

- (b) Study how N and T affects the variance in the log likelihood estimate.
- (c) Remove the resampling step from your particle filter algorithm, and study its effect on the variance of the estimator.

II.2 Fully adapted particle filter

(a) Motivate for each of these model why it is/is not possible to implement the fully adapted particle filter for it.

(i)

$$x_{t+1} = 0.4x_t + v_t, \quad v_t \sim \mathcal{N}(0, 1), \quad (3)$$

$$y_t = -0.5x_t + e_t \quad e_t \sim \mathcal{N}(0, 0.01). \quad (4)$$

(ii)

$$x_{t+1} = \cos(x_t)^2 + v_t, \quad v_t \sim \mathcal{N}(0, 1), \quad (5)$$

$$y_t = 2x_t + e_t \quad e_t \sim \mathcal{N}(0, 0.01). \quad (6)$$

(iii)

$$x_{t+1} = \cos(x_t + v_t)^2, \quad v_t \sim \mathcal{N}(0, 1), \quad (7)$$

$$y_t = 2x_t + e_t \quad e_t \sim \mathcal{U}([-2, 2]). \quad (8)$$

(b) Implement the fully adapted particle filter for model (ii), and make a simulation study to compare the variance in the estimates of $\mathbb{E}[X_t | y_{1:t}]$ to the estimates obtained by a bootstrap particle filter.

II.3 Likelihood estimator for the APF.

The particle filter likelihood estimator is given by

$$\hat{p}(y_{1:T}) = \prod_{t=1}^T \left\{ \frac{1}{N} \sum_i \tilde{w}_t^i \right\} \quad (9)$$

For the bootstrap particle filter, given in Algorithm 1, a sketchy derivation of this estimator can be done as:

$$p(y_{1:T}) = \prod_{t=1}^T p(y_t | y_{1:t-1}) = \prod_{t=1}^T \int p(y_t | x_t) p(x_t | y_{1:t-1}) dx_t \approx \prod_{t=1}^T \frac{1}{N} \sum_{i=1}^N \underbrace{p(y_t | x_t^i)}_{\tilde{w}_t^i} \quad (10)$$

where the particles x_t^i sampled at time t in the bootstrap particle filter (before weighting) can be viewed as approximately distributed according to the predictive distribution $p(x_t | y_{1:t-1})$.

However, the likelihood estimator (9) is valid for the general auxiliary particle filter, given in Algorithm 2, as well. Derive this estimator for the auxiliary particle filter, in a similar fashion as was done above for the bootstrap particle filter.

Hint: You need to take the auxiliary variables into account. That is, write the pdf $p(y_t | y_{1:t-1})$ as an integral over (x_t, a_t) (more precisely, an integral over x_t and sum over a_t). Then interpret this integral as an expected value with respect to the joint proposal used in the auxiliary particle filter.

N.B The expression (9) assumes that the weights are computed as stated in Algorithm 2, i.e. that the unnormalized weights at time t , \tilde{w}_t , are expressed in terms of the normalized weights at time $t - 1$, w_{t-1} .

Algorithm 1 Bootstrap particle filter (for $i = 1, \dots, N$)

- (a) **Initialization** ($t = 0$):
- i. Sample $x_0^i \sim p(x_0)$.
 - ii. Set initial weights: $w_0^i = 1/N$.
- (b) **for** $t = 1$ **to** T **do**
- i. **Resample**: sample ancestor indices $a_t^i \sim \mathcal{C}(\{w_{t-1}^j\}_{j=1}^N)$.
 - ii. **Propagate**: sample $x_t^i \sim p(x_t | x_{t-1}^{a_t^i})$.
 - iii. **Weight**: compute $\tilde{w}_t^i = p(y_t | x_t^i)$ and normalize $w_t^i = \tilde{w}_t^i / \sum_{j=1}^N \tilde{w}_t^j$.
-

Algorithm 2 Auxiliary particle filter (for $i = 1, \dots, N$)

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

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

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

II.4 Forgetting

Consider the bootstrap particle filter for the LGSS model with 0 variance in the transition model

$$X_t = 0.7X_{t-1} + V_t, \quad V_t \sim \mathcal{N}(0, 0), \quad (11a)$$

$$Y_t = 0.5X_t + E_t, \quad E_t \sim \mathcal{N}(0, 0.1). \quad (11b)$$

What happens to the errors in the particle filter (compared to the Kalman filter, the exact solution) along the time dimension? Specifically, run the particle filter, say, 100 times (using a fixed N) for the same data and compute the mean-squared-error of the test function $\varphi(x_t) = x_t$ with respect to the Kalman filter solution,

$$\frac{1}{100} \sum_{\ell=1}^{100} \left(\hat{I}_{t,N}^{\text{PF},\ell}(\varphi) - \mathbb{E}[X_t | y_{1:t}] \right)^2$$

for each time step $t = 1, 2, \dots$, where $\hat{I}_{t,N}^{\text{PF},\ell}(\varphi)$ is the estimate of $\mathbb{E}[X_t | y_{1:t}]$ obtained from the ℓ th run of the particle filter. Is the particle filter stable?