MATS442 Stochastic simulation — Problems 6, 21.2.2020

The problems marked with **C** have a part which needs to be implemented in a computer programming language. Please return a single code file with the answers to these exercises by email to ⟨santeri.j.karppinen(·at·)jyu.fi⟩.

1. (2 points) Consider the particle filter (Algorithm 8.15). Recall that the resampling step (iii) may be implemented as follows (cf. Theorem 2.3):

(i) Draw $U_1, \ldots, U_n \overset{\text{i.i.d.}}{\sim} U(0,1)$,
(ii) Set $A_{t-1}^{(i)} := g(U_i; \bar{\omega}_{t-1}^{(1:n)})$,

where $g(u; \bar{\omega}^{(1:n)}) := \min\big\{k \in \{1{:}n\} \,:\, \sum_{j=1}^k \bar{\omega}^{(j)} \geq u\big\}$.

Consider the following two alternative resampling algorithms, described for general fixed probability distribution $\bar{\omega}^{(1:n)}$:

(a) *Stratified resampling*:
  (i) Draw $U_1, \ldots, U_n \overset{\text{i.i.d.}}{\sim} U(0,1)$ and set $\hat{U}_i := \frac{i-1+U_i}{n}$.
  (ii) Set $\hat{A}^{(i)} := g(\hat{U}_i); \bar{\omega}^{(1:n)}$.
(b) *Systematic resampling*:
  (i) Draw a *single* $U \sim U(0,1)$ and set $\tilde{U}_i := \frac{i-1+U}{n}$ for $i = 1{:}n$.
  (ii) Set $\tilde{A}^{(i)} := g(\tilde{U}_i; \bar{\omega}^{(1:n)})$.

Show that both stratified and systematic resampling satisfy the required unbiasedness condition

$$\mathbb{E}\bigg[\frac{1}{n}\sum_{i=1}^n \mathbf{1}\left(\hat{A}^{(i)} = j\right)\bigg] = \mathbb{E}\bigg[\frac{1}{n}\sum_{i=1}^n \mathbf{1}\left(\tilde{A}^{(i)} = j\right)\bigg] = \bar{\omega}^{(j)}, \qquad j = 1{:}n.$$

2.C (2 points) Consider the noisy AR(1) model in Example 8.2, with $\rho = 1$ and $\sigma_1^2 = \sigma_x^2 = \sigma_y^2 = 1$, suppose $T = 15$ and let

$$y_{1:T} := (0, -1, -2, -1, -4, -1, -2, 0, 2, -1, 1, 2, -2, 1, 3).$$

Implement the particle filter with the following proposal distributions:

$$M_t(x_t \mid x_{1:t-1}) = M_t(x_t \mid x_{t-1}) = N\bigg(x_t; \frac{\rho x_{t-1} + y_t}{2}, \frac{1}{2}\bigg)$$

with $x_0 := 0$ and $G_t(\,\cdot\,)$ chosen as in Remark 8.8.

Estimate $\mathbb{E}_p[X_T]$ and $\mathrm{Var}_p(X_T)$ with $n = 100$, $n = 1000$ and $n = 10000$ particles.

(Hint: You are welcome to (but need not) start with the example PF code in the lecture notes. In case you plan to reuse the code, please be aware that here the function G_t will depend also on x_{t-1}. The true values are approximately 1.9898 and 0.6180, respectively.)

3.C (2 points) Consider the setting in Problem 2, and assume still that $\sigma_1^2 = \sigma_x^2 = \sigma_y^2 = 1$, but that ρ is unknown, with the prior distribution $\mathrm{pr}(\rho) = \mathrm{Unif}(-10, 10)$.

Implement particle marginal Metropolis-Hastings targetting

$$p(\rho, x_{1:T}) \propto \mathrm{pr}(\rho) p^{(\rho)}(x_{1:T}, y_{1:T}),$$

In particular, use the Gaussian random-walk Metropolis algorithm with increment standard deviation 1/2 to propose changes in ρ, and the particle filter of Problem 2 in the pseudo-marginal algorithm with $n = 100$ particles. (Note that now $M_t = M_t^{(\rho)}$ and $G_t = G_t^{(\rho)}$!)

Estimate the mean and variance of ρ with respect to p, over (at least) 10000 iterations of the algorithm.

Hint: Because we are only interested in functions $f(\theta, x_{1:T}) = f(\theta)$, the variables $V_k^{(1:n)}$ and $\mathbf{X}_k^{(1:n)}$ need not be stored, only (Θ_k) and $\sum_{i=1}^n V_k^{(i)} \ldots$

4. Problem 6 from Demo 5 (if you did not do it already). Hint: Convince yourself that the transition probability K in a) can be written down as

$$\begin{aligned} [K]_{(x,z),(y,t)} &= \mathbb{P}\big((X_k, Z_k) = (y,t) \mid (X_{k-1}, Z_{k-1}) = (x,z)\big) \\ &= q(x,y)h(y,t)\alpha\big((x,z),(y,t)\big) + 1\big((y,t) = (x,z)\big)\rho(x,z), \end{aligned}$$

with

$$\begin{aligned} \alpha\big((x,z),(y,t)\big) &= \min\left\{1, \frac{t}{z}\frac{q(y,x)}{q(x,y)}\right\}, \quad \text{(if } zq(x,y) > 0 \text{ and 0 otherwise)}, \\ \rho(x,z) &= 1 - \sum_{y \in \mathbb{X}, t \in \mathbb{N}} q(x,y)h(y,t)\alpha\big((x,z),(y,t)\big). \end{aligned}$$

The solution for b) is a similar computation we have already done in previous exercises..