Stochastic Gradient Techniques
(Period 2, 2021/2022)
Exam 21 December 2021 12:14-14:15 hr

Problem 1

Let $X(\theta)$ be a random variable with cumulative distribution function

$$F_\theta(x) = \mathbb{P}(X(\theta) \le x) = \begin{cases} 0 & \text{for } x < 0; \\ \sqrt{\theta x} & \text{for } 0 \le x < 1/\theta; \\ 1 & \text{for } x \ge 1/\theta. \end{cases} \tag{1}$$

The parameter is $\theta \in \Theta = [1, 6]$. A cost function $h\big(X(\theta), \theta\big)$ is given by

$$h\big(X(\theta), \theta\big) = 100X(\theta) + \theta^2,$$

and the objective function is

$$J(\theta) = \mathbb{E}\big[h\big(X(\theta), \theta\big)\big].$$

Figure 1 shows the graph of the objective function $J(\theta)$ on $[1, 6]$.

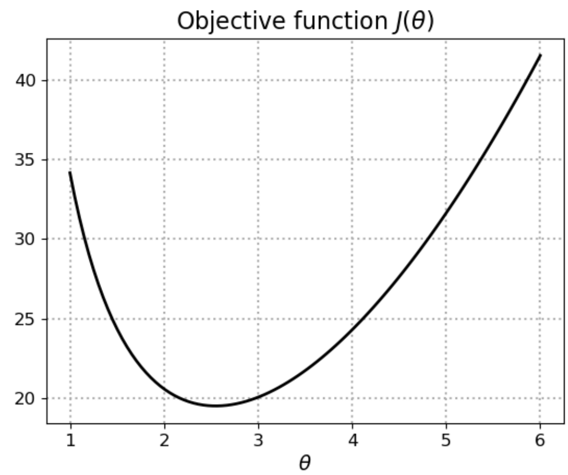

Figure 1: The objective function of Problem 1.

(a). Show how to generate samples from the distribution (1).

(b). Consider the optimization problem $\min_{\theta>0} J(\theta)$. Argue that a gradient descent algorithm will converge to the true minimizer. How do you adjust the gradient descent for the fact that $\theta > 0$?

(c). Compute the infinitesimal perturbation analysis (IPA) estimator for $dJ(\theta)/d\theta$. You may assume that interchange conditions apply.

(d). Give the complete simulation algorithm (in pseudocode) that would implement the stochastic approximation iteration starting at some $\theta_0 \in (1, 6)$ using a decreasing stepsize (starting at some $\epsilon_0 > 0$) and using the unbiased derivative estimator based on the IPA method. Include how to draw samples of the $X(\theta)$ distribution.

(e). Figure 2 shows the score function method (SFM) estimator of $dJ(\theta)/d\theta$. Clearly, it is not a good estimator. What is the reason that the SFM estimator is not good.

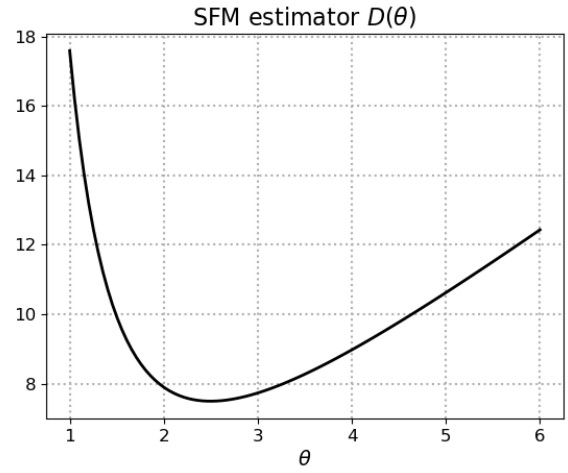

Figure 2: The score function estimator of Problem 1.

Solution:

(a). Apply the inverse transfomr method.

$$F_\theta(x) = u \quad \Leftrightarrow \quad \sqrt{\theta x} = u \quad \Leftrightarrow \quad x = u^2/\theta.$$

(b). The figure shows that $J(\theta)$ is continuous, differentiable, convex and unimodal on $[1, 6]$. Let θ^* be the true minimizer. When a proposal solution $\theta_n > \theta^*$, the next solution $\theta_{n+1} < \theta_n$ because $J(\theta)$ is increasing for $\theta > \theta^*$, and the gradient descent iteration is $\theta_{n+1} = \theta_n - \epsilon_n J'(\theta_n)$ with $\epsilon_n > 0$. Similarly, when $0 < \theta_n < \theta^*$, the next solution is $\theta_{n+1} > \theta_n$. Thus, when gain size ϵ_n is sufficiently small, the next solution is closer to θ^*.

When $\theta_n > \theta^*$, and the step $\epsilon_n J'(\theta_n) > \theta_n$, the next solution would be $\theta_{n+1} < 0$, which is not feasible. To cope with this, the negative θ_{n+1} is projected back into $(0, \infty)$, for instance, $\theta_{n+1} = 1$.

(c).

$$X(\theta) = U^2/\theta \quad \Rightarrow \quad X'(\theta) = -U^2/\theta^2 = -X(\theta)/\theta.$$

When it is allowed to interchange, the IPA estimator is

$$D^{\text{IPA}}(\theta) = \frac{d}{d\theta} h\big(X(\theta), \theta\big) = \frac{d}{d\theta}\big(100X(\theta) + \theta^2\big)$$
$$= 100X'(\theta) + 2\theta = -\frac{X(\theta)}{\theta} + 2\theta.$$

(d). We implement SA iteration $\theta_{n+1} = \theta_n + \epsilon_n Y_n$, with update $Y_n = -\frac{1}{k}\sum_{i=1}^{k} D_i^{\text{IPA}}(\theta_n)$, the average of k samples of the IPA estimator.

Algorithm 1 Stochastic Approximation Iteration

Require: θ_0 {*initial point*}
Require: ϵ_0 {*initial stepsize*}
Require: m {*number of iterations*}
Require: k {*size of mini batch*}
for $n = 0$ **to** $m-1$ **do**
  for $i = 1$ **to** k **do**
    Generate U, uniform $(0,1)$ $\{RNG\}$
    $X \leftarrow U^2/\theta_n$
    Compute $D_i = -100X/\theta_n + 2\theta_n$ {*i-th IPA sample*}
  end for
  $D \leftarrow \frac{1}{k}\sum_{i=1}^{k} D_i$ {*IPA estimator*}
  $\theta_{n+1} = \theta_n - \epsilon_0/(n+1) \times D$ {*SA iteration*}
end for
return θ_m

(e). The interchange conditions do not apply, thus the estimator is not unbiased. The condition that all densities should be defined on common support, is not satisfied here. The θ-distribution is defined on $[0,\theta]$. This complicated differentiating wrt θ.

In more detail:
The score function is derived as follows,

$$F_\theta(x) = \sqrt{\theta x} \quad \Rightarrow \quad f_\theta(x) = \frac{1}{2}\sqrt{\theta/x}$$

$$\Rightarrow \quad S(\theta; x) = \frac{\partial}{\partial\theta}\log f_\theta(x) = \frac{d}{d\theta}\log\sqrt{\theta} = \frac{1}{2\theta} \quad (0 < x < 1/\theta).$$

Suppose that you would compute

$$\frac{\partial}{\partial\theta}\int 100x f_\theta(x)\,dx = \int 100x \frac{\partial}{\partial\theta} f_\theta(x)\,dx = \int 100x S(\theta,x)\, f_\theta(x). \tag{2}$$

Then the SFM estimator for $J'(\theta)$ would be

$$D^{\mathrm{SFM}}(\theta) = 100X(\theta)S(\theta, X(\theta)) + 2\theta = 50X(\theta)/\theta + 2\theta.$$

This estimator is plotted in Figure 2. However, the interchange (2) is not valid, because the integral is actually

$$\int_0^\theta 100x f_\theta(x)\,dx.$$

Problem 2

Is $X(\theta)$ almost surely Lipschitz continuous in the following cases? And if so, is the Lipschitz modulus integrable? Just an answer is not sufficient. Provide an analysis where you proof your claims.

(a). $X(\theta) = U^\theta$, where U is the uniform (0,1) random variable, and $\theta \in (0,b)$ $(b > 0)$.

(b). $X(\theta) = \theta I\{E < \theta\}$, where E is the standard exponential random variable (i.e. with mean 1), and $\theta \in (a,b) \subset (0,\infty)$.

(c).

$$X(\theta) = \frac{\theta}{\sqrt{|NU - 1|}},$$

where U is the uniform (0,1) random variable, N is the Poisson random variable with mean 1, U and N are independent, and $\theta \in (0, 1)$.
Hint: consider first an arbitrary integer, e.g $N = 3$, and the function $X(\theta) = \theta/\sqrt{|3U - 1|}$.

Solution:

(a). Yes. For any $u \in (0, 1)$ sample of U, is $X(\theta) = u^\theta$ differentiable as function of θ on $[0, \infty)$, with $X'(\theta) = u^\theta \log u$. Note that u^θ is decreasing as function of θ on $[0, \infty)$, thus

$$\sup_{\theta \in (0,b)} |X'(\theta)| = |X'(0)| = |\log u| = -\log u = K(u) < \infty.$$

Thus $X(\theta)$ is Lipschitz with probability 1.

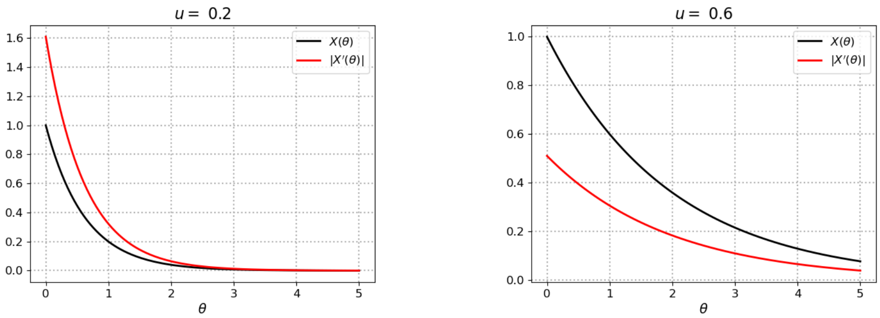

Figure 3: Two functions $X(\theta) = u^\theta$ on $[0, 5]$, and their derivatives.

The modulus is integrable:

$$\mathbb{E}[K] = \int_0^1 (-\log u)\, du \overset{u=e^{-x}}{=} \int_0^\infty x e^{-x}\, dx = 1.$$

(b). No. For any $y > 0$ sample of E, is

$$X(\theta) = \begin{cases} 0 & \text{for } 0 < \theta \leq y; \\ \theta & \text{for } \theta > y. \end{cases}$$

This function (of θ on (a, b)) is discontinuous if $y \in (a, b)$, namely in $\theta = y$, and then it is not Lipschitz.

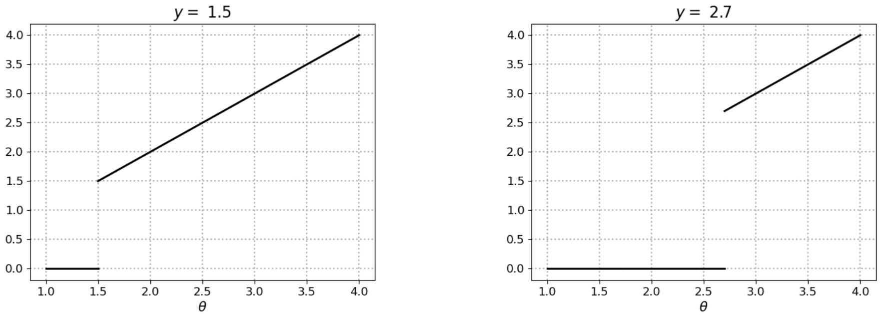

Figure 4: Two functions $X(\theta) = \theta I\{\theta > y\}$ on $[1, 4]$.

The probability that E samples from (a, b) is positive.

(c). Yes. First, the case $N = 3$, then for any sample $u \in (0, 1)$, $u \neq 1/3$, the function $X(\theta) = \theta/\sqrt{|3u - 1|}$ is linear as function of θ on $\mathbb{R}$. Hence, Lipschitz with finite modulus $K(u) = |X'(\theta)| = 1/\sqrt{|3u - 1|}$.

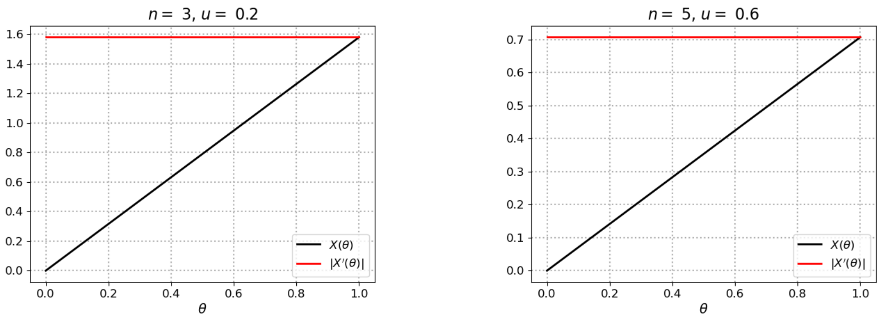

Figure 5: Two functions $X(\theta) = \theta/\sqrt{|nu - 1|}$ on $[0, 1]$, and their derivatives.

The modulus satisfies

$$\mathbb{E}[K] = \int_0^1 \frac{du}{\sqrt{|3u-1|}} = \int_0^{1/3} \frac{du}{\sqrt{1-3u}} + \int_{1/3}^1 \frac{du}{\sqrt{3u-1}}$$
$$= \Big[-\frac{2}{3}\sqrt{1-3u} \Big]_0^{1/3} + \Big[\frac{2}{3}\sqrt{3u-1} \Big]_{1/3}^1 = \frac{2}{3} + \frac{2}{3}\sqrt{2} < \infty.$$

The function $X(\theta) = \theta/\sqrt{|3u - 1|}$ is not defined, and thus not Lipschitz if $u = 1/3$, but, the probability that $U = 1/3$ is 0. Conclusion, if $N = 3$, the function $X(\theta)$ is Lipschitz with probability 1.

Now, take N random. For any $n = 0, 1, \ldots$ and $u \in (0, 1)$, such that $nu \neq 1$, the function $X(\theta) = \theta/\sqrt{|nu - 1|}$ is linear as function of θ on $\mathbb{R}$. Hence, Lipschitz with finite modulus $K(n,u) = 1/\sqrt{|nu-1|}$, which satisfies

$$\mathbb{E}[K] = \mathbb{E}\Big[\frac{1}{\sqrt{|NU-1|}}\Big] = \sum_{n=0}^{\infty} \mathbb{P}(N=n)\mathbb{E}\Big[\frac{1}{\sqrt{|nU-1|}}\Big],$$

with for $n \geq 1$,

$$\mathbb{E}\Big[\frac{1}{\sqrt{|nU-1|}}\Big] = \int_0^1 \frac{du}{\sqrt{|nu-1|}} = \int_0^{1/n} \frac{du}{\sqrt{1-nu}} + \int_{1/n}^1 \frac{du}{\sqrt{nu-1}}$$
$$= \Big[-\frac{2}{n}\sqrt{1-nu}\Big]_0^{1/n} + \Big[\frac{2}{n}\sqrt{nu-1}\Big]_{1/n}^1 = \frac{2}{n} + \frac{2}{n}\sqrt{n-1}.$$

Thus,

$$\mathbb{E}[K] = \mathbb{P}(N=0) + \sum_{n=1}^{\infty} \mathbb{P}(N=n)\Big(\frac{2}{n} + \frac{2}{n}\sqrt{n-1}\Big)$$
$$= \mathbb{P}(N=0) + 2e^{-1}\sum_{n=1}^{\infty} \frac{1}{n(n!)} + 2e^{-1}\sum_{n=1}^{\infty} \frac{\sqrt{n-1}}{n(n!)}$$
$$\leq \mathbb{P}(N=0) + 2e^{-1}\sum_{n=1}^{\infty} \frac{1}{n!} + 2e^{-1}\sum_{n=1}^{\infty} \frac{1}{n!} < \infty.$$

The function $X(\theta) = \theta/\sqrt{|nu-1|}$ is not defined, and thus not Lipschitz if $nu = 1$. But, the probability

$$\mathbb{P}(NU = 1) = \sum_{n=0}^{\infty} \mathbb{P}(N=n)\mathbb{P}(U = 1/n) = \sum_{n=0}^{\infty} \mathbb{P}(N=n) \times 0 = 0.$$

Conclusion, the function $X(\theta)$ is Lipschitz with probability 1.

Problem 3

Consider the following network of 6 nodes $(A, \ldots, F)$, and 8 links, numbered $1, \ldots 8$ as shown.

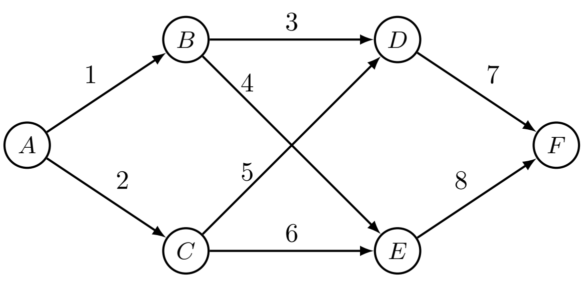

It takes a random time X_i to travel on link i. The travel times $X_1 = X_1(\theta)$ and $X_2 = X_2(\theta)$ depend on the same parameter, and are exponentially disributed with mean θ. The other travel times $X_3, \ldots, X_8$ are all exponentially disributed with mean 1. Moreover, $X_1, \ldots, X_8$ are independent.

There are 4 routes from A to F:

$$R_1 = A \to B \to D \to F;$$
$$R_2 = A \to B \to E \to F;$$
$$R_3 = A \to C \to D \to F;$$
$$R_4 = A \to C \to E \to F.$$

The route time on route j is the sum of the travel times on the links of route j: $T_j = T_j(\theta)\sum_{i\in R_j} X_i$. The maximal route time is

$$L(\theta) = L\big(X_1(\theta), X_2(\theta), X_3, \ldots, X_8\big) = \max_{j=1,2,3,4} T_j(\theta)$$

We are interested in the objective function

$$J(\theta) = \mathbb{E}\big[L(\theta)\big].$$

(a). Derive the measure valued differentiation (MVD) estimator $D^{\text{MVD}}(\theta)$ for $J'(\theta)$. You may assume that interchange applies, and you may use the MVD decomposition of the partial derivative of the exponential probability density function. Give an algorithm in pseudocode for computing the estimator, including sampling of the random variables.

(b). Provide a randomized MVD estimator $D^{\text{MVDrand}}(\theta)$ that is about twice as fast to compute as the $D^{\text{MVD}}(\theta)$ estimator of (a). Show that

$$\mathbb{E}\big[D^{\text{MVDrand}}(\theta)\big] = \mathbb{E}\big[D^{\text{MVD}}(\theta)\big].$$

(c). Now, assume that $X_1 = X_1(\theta_1)$ and $X_2 = X_2(\theta_2)$ are parameterized with different parameters θ_1 and θ_2, respectively. Futhermore, the objective function is

$$J(\theta_1,\theta_2) = \mathbb{E}\big[L\big(X_1(\theta_1), X_2(\theta_2), X_3, \ldots, X_8\big)\big] + \frac{10}{\theta_1\theta_2}.$$

Work out the details of the simultaneous perturbation stochastic approximation (SPSA) for the optimization problem $\min_{(\theta_1,\theta_2)} J(\theta_1,\theta_2)$. Give pseudocode of the associated algorithm.

Solution:

(a). The PDF of the $X_1(\theta)$ and $X_2(\theta)$ is

$$f_\theta(x) = \frac{1}{\theta} e^{-x/\theta} \quad (x > 0).$$

We use that the derivative is decomposed as

$$\frac{\partial}{\partial\theta} f_\theta(x) = \Big(\frac{x}{\theta^3} - \frac{1}{\theta^2}\Big)e^{-x/\theta} = \frac{1}{\theta}\Big(\underbrace{\frac{x}{\theta^2}e^{-x/\theta}}_{f_\theta^+ \sim \text{Gamma}} - \underbrace{\frac{1}{\theta}e^{-x/\theta}}_{f_\theta^- \sim \text{Exponential}}\Big).$$

The associated random variables are denoted by X^+, and X^-, respectively. The joint PDF of $X_1(\theta), X_2(\theta), X_3, \ldots, X_8$ is

$$f(x_1,\ldots,x_8) = f_\theta(x_1) f_\theta(x_2) \prod_{i=3}^{8} f(x_i),$$

with derivative (apply the product rule),

$$\frac{\partial}{\partial\theta} f(x_1,\ldots,x_8) = \Big(\frac{\partial}{\partial\theta} f_\theta(x_1)\Big) f_\theta(x_2) \prod_{i=3}^{8} f(x_i) + f_\theta(x_1)\Big(\frac{\partial}{\partial\theta} f_\theta(x_2)\Big) \prod_{i=3}^{8} f(x_i)$$

$$= \Big(\frac{1}{\theta}\big(f_\theta^+(x_1) - f_\theta^-(x_1)\big)\Big) f_\theta(x_2) \prod_{i=3}^{8} f(x_i)$$

$$+ f_\theta(x_1)\Big(\frac{1}{\theta}\big(f_\theta^+(x_2) - f_\theta^-(x_2)\big)\Big) \prod_{i=3}^{8} f(x_i)$$

Then,

$$J(\theta) = \int_{(0,\infty)^8} L(x_1,\ldots,x_8) f(x_1,\ldots,x_8)\, dx_1,\cdots dx_8$$

$$= \int_{(0,\infty)^8} L(x_1,\ldots,x_8) f_\theta(x_1) f_\theta(x_2) \prod_{i=3}^{8} f(x_i)\, dx_1,\cdots dx_8.$$

Because interchange is allowed,

$$J'(\theta) = \int_{(0,\infty)^8} L(x_1,\ldots,x_8) \frac{\partial}{\partial\theta} f(x_1,\ldots,x_8)\, dx_1,\cdots dx_8$$

$$= \int_{(0,\infty)^8} L(x_1,\ldots,x_8) \frac{1}{\theta}\Big($$

$$f_\theta^+(x_1) f_\theta(x_2) \prod_{i=3}^{8} f(x_i) - f_\theta^-(x_1) f_\theta(x_2) \prod_{i=3}^{8} f(x_i)$$

$$+ f_\theta(x_1) f_\theta^+(x_2) \prod_{i=3}^{8} f(x_i) - f_\theta(x_1) f_\theta^-(x_2) \prod_{i=3}^{8} f(x_i)\Big) dx_1,\cdots dx_8.$$

which we are going to model as $\mathbb{E}[D^{\mathrm{MVD}}(\theta)]$. Hence, the MVD estimator is

$$D^{\mathrm{MVD}}(\theta) = \frac{1}{\theta}\Big(L\big(X_1^+(\theta), X_2(\theta), X_3,\ldots,X_8\big) - L\big(X_1^-(\theta), X_2(\theta), X_3,\ldots,X_8\big)\Big)$$
$$+ \frac{1}{\theta}\Big(L\big(X_1(\theta), X_2^+(\theta), X_3,\ldots,X_8\big) - L\big(X_1(\theta), X_2^-(\theta), X_3,\ldots,X_8\big)\Big)$$

In this expression, the parameterized random variables $X_1(\theta)$ and $X_2(\theta)$ are decomposed in a $+$ and a $-$ version, as mentioned above. The $+$ version has an Erlang-2 distribution with scale θ, and the $-$ version has the original exponential distribution with scale θ. An Erlang-2 is the sum of two IID exponentially distributed random variables. In the algorithm below, $X_i(\theta)$ and $Y_i(\theta)$ $(i = 1, 2)$ are IID from the exponential distribution with mean θ. Also we use that exponential variates with mean μ are generated by $-\mu \log(1 - U)$ for $(0, 1)$-uniform U. The algorithm exploits that

$$L\big(X_1^-(\theta), X_2(\theta), X_3,\ldots,X_8\big) = L\big(X_1(\theta), X_2^-(\theta), X_3,\ldots,X_8\big)$$
$$= L\big(X_1(\theta), X_2(\theta), X_3,\ldots,X_8\big).$$

Algorithm 2 MVD Estimator

Require: $\theta > 0$ {*parameter of interest*}
for $i = 1$ **to** 2 **do**
  Generate U_i from random number generator
  $X_i(\theta) \leftarrow -\theta \log(1 - U_i)$ {*IID θ-exponentials*}
end for
for $i = 3$ **to** 8 **do**
  Generate U_i from random number generator
  $X_i(\theta) \leftarrow -\log(1 - U_i)$ {*IID standard exponentials*}
end for
for $j = 1$ **to** 4 **do**
  Compute T_j {*route travel times*}
end for
$L^{(-)} \leftarrow \max_j T_j$
for $i = 1$ **to** 2 **do**
  Generate U_i from random number generator
  $Y_i(\theta) \leftarrow -\theta \log(1 - U_i)$ {*θ-exponential*}
  $X_i(\theta) \leftarrow X_i(\theta) + Y_i(\theta)$ {*Erlang distribution*}
  for $j = 1$ **to** 4 **do**
    Compute T_j {*new route travel times*}
  end for
  $L_i^{(+)} \leftarrow \max T_j$
end for
return $D \leftarrow \sum_i \left(L_i^{(+)} - L^{(-)}\right)/\theta$

(b). Note that we need the decomposition of both X_1 and X_2. In the randomized version of MVD, we choose one of these two compositions at random. In detail, define the random variable $\sigma \in \{1, 2\}$ with uniform distribution $\mathbb{P}(\sigma = 1) = \mathbb{P}(\sigma = 2) = 0.5$. If $\sigma = i$, decompose variable X_i, $i = 1, 2$. Hence, the estimator is

$$D^{\text{MVDrand}}(\theta) = \frac{2}{\theta}\Big(L\big(X_\sigma^+(\theta), \ldots, X_8\big) - L\big(X_\sigma^-(\theta), \ldots, X_8\big)\Big).$$

This estimator needs 2 computations (or runs) of the cost function L, whereas the full MVD of part (a) needs 4 computations. Note that the constant $1/\theta$ is adapted for taking account the randomization. The randomized estimator is unbiased:

$$\begin{aligned}
&\mathbb{E}\big[D^{\text{MVDrand}}(\theta)\big] \\
&= \mathbb{P}(\sigma = 1)\frac{2}{\theta}\mathbb{E}\Big[L\big(X_1^+(\theta), X_2(\theta), X_3, \ldots, X_8\big) - L\big(X_1^-(\theta), X_2(\theta), X_3, \ldots, X_8\big)\Big] \\
&+ \mathbb{P}(\sigma = 2)\frac{2}{\theta}\mathbb{E}\Big[L\big(X_1(\theta), X_2^+(\theta), X_3, \ldots, X_8\big) - L\big(X_1(\theta), X_2^-(\theta), X_3, \ldots, X_8\big)\Big] \\
&= \frac{1}{\theta}\mathbb{E}\Big[L\big(X_1^+(\theta), X_2(\theta), X_3, \ldots, X_8\big) - L\big(X_1^-(\theta), X_2(\theta), X_3, \ldots, X_8\big)\Big] \\
&+ \frac{1}{\theta}\mathbb{E}\Big[L\big(X_1(\theta), X_2^+(\theta), X_3, \ldots, X_8\big) - L\big(X_1(\theta), X_2^-(\theta), X_3, \ldots, X_8\big)\Big] \\
&= \mathbb{E}\big[D^{\text{MVD}}(\theta)\big]
\end{aligned}$$

(c). The idea is to approximate the minimizer by the iteration

$$\theta_{n+1} = \theta_n + \epsilon_n Y_n, \quad n = 0, 1, \ldots,$$

where each approximation $\theta_n = (\theta_{n1}, \theta_{n2}) \in (0, \infty)^2$. The update Y_n is an approximation of the gradient $\nabla J(\theta_1, \theta_2)$. The exact gradient is

$$\begin{aligned} \nabla J(\theta_1, \theta_2) &= \nabla \mathbb{E}\big[L(\theta_1, \theta_2)\big] + \nabla \frac{10}{\theta_1 \theta_2} \\ &= \begin{pmatrix} \frac{\partial}{\partial \theta_1} \mathbb{E}\big[L(\theta_1, \theta_2)\big] \\ \frac{\partial}{\partial \theta_2} \mathbb{E}\big[L(\theta_1, \theta_2)\big] \end{pmatrix} - \begin{pmatrix} \frac{10}{\theta_1^2 \theta_2} \\ \frac{10}{\theta_1 \theta_2^2} \end{pmatrix} \end{aligned}$$

In iteration n of the SPSA algorithm, the current approximation of the minimizer is $\theta_n = (\theta_{n1}, \theta_{n2})$. Then, approximating the gradient $\nabla \mathbb{E}\big[L(\theta_{n1}, \theta_{n2})\big]$ is based on simultaneous randomly perturbing the two coordinates of θ_n. As follows, let Δ_{n1}, Δ_{n2} be i.i.d. uniformly distributed on $\{-1, 1\}$, and define the estimator $\widehat{\nabla} \mathbb{E}\big[L(\theta_n)\big] \in \mathbb{R}^2$ by its two coordinates

$$\begin{aligned} \Big(\widehat{\nabla} \mathbb{E}\big[L(\theta_n)\big]\Big)_i &= \frac{1}{2 \eta_n \Delta_{ni}} \\ &\quad \Big(L\big(X_1(\theta_{n1} + \eta_n \Delta_{n1}), X_2(\theta_{n2} + \eta_n \Delta_{n2}), X_3, \ldots, X_8\big) \\ &\quad - L\big(X_1(\theta_{n1} - \eta_n \Delta_{n1}), X_2(\theta_{n2} - \eta_n \Delta_{n2}), X_3, \ldots, X_8\big)\Big) \quad (i = 1, 2), \end{aligned}$$

where $\eta_n > 0$ is the perturbation size of θ_n in the n-th iteration. The update in the SPSA is

$$Y_n = -\widehat{\nabla} \mathbb{E}\big[L(\theta_n)\big] - \nabla \frac{10}{\theta_1 \theta_2}.$$

The algorithm below executes m iterations while decreasing the step size by $\epsilon_n = \epsilon_0/(n+1)$, and the perturbation size by $\eta_n = \eta_0/\sqrt[4]{n+1}$.

Algorithm 3 Simultaneous Perturbation Stochastic Approximation

Require: θ_0 {*initial point*}
Require: ϵ_0 {*initial stepsize*}
Require: η_0 {*initial perturbation*}
Require: m {*number of iterations*}
for $n = 0$ **to** $m - 1$ **do**
  $\epsilon_n \leftarrow \epsilon_0/(n+1)$
  $\eta_n \leftarrow \eta_0/\sqrt[4]{n+1}$
  for $i = 1$ **to** 2 **do**
    Generate $\Delta_i \in \{-1, 1\}$ at random
    Generate $X_i = X_i(\theta_{ni} + \eta_n \Delta_i)$ from exponential distribution
    Generate $Y_i = X_i(\theta_{ni} - \eta_n \Delta_i)$ from exponential distribution
  end for
  for $i = 3$ **to** 8 **do**
    Generate X_i from exponential distribution
  end for
  $L^+ \leftarrow L\big(X_1, X_2, X_3, \ldots, X_8\big)$
  $L^- \leftarrow L\big(Y_1, Y_2, X_3, \ldots, X_8\big)$
  $D_1 \leftarrow (L^+ - L^-)/(2\eta_n \Delta_1) - 10/(\theta_{n1}^2 \theta_{n2})$ {*coordinate 1*}
  $D_2 \leftarrow (L^+ - L^-)/(2\eta_n \Delta_2) - 10/(\theta_{n1} \theta_{n2}^2)$ {*coordinate 2*}
  $\theta_{n+1} = \theta_n - \epsilon_n \times D$ {*SPSA iteration*}
end for
return θ_m

Problem 4

Two single server stations are forming a closed loop. Jobs that complete their service at station 1 move to station 2 (as soon as there is a free place available), and jobs from station 2 move to station 1 (as soon as there is free place available). Next to the service place, station 1 has $B_1 = 5$ buffer places and station 2 has $B_2 = 4$ buffer places. There is a population of 8 jobs that circulate through the system.

(a). Give a DES description of this model. What is the event set, what is the physical state space, how is the event list defined, and what is the sate-transition mapping?

(b). Does this system satisfy the commuting condition?

Solution:

(a). DES description.

- Events $E = \{\beta_1, \beta_2 B\}$, where β_i is a service completion at station i.
- A state is a vector $s = (s_1, b_1, s_2, b_2))$, where s_i is the number of jobs at station i and $b_i = 0$ if station i not blocked and $b_i = 1$ if station i is blocked.
- Event list $L(s) = \{\beta_1\}$ if $s_2 = 0, s_1 > 0$, $L(s) = \{\beta_2\}$ if $s_1 = 0, s_2 > 0$, and $L(s) = \{\beta_1, \beta_2\}$ else.

- The state transition mapping is

$$\phi(\beta_1, s_1, 0, s_2, 0) = (s_1 - 1, 0, s_2 + 1, 0) \text{ if } s_2 \leq B_2,$$

and

$$\phi(\beta_1, s_1, 0, B_2 + 1, 0) = (s_1, 1, s_2 + 1, 0) \text{ if } s_2 > B_2,$$
$$\phi(\beta_2, s_1, 0, s_2, 0) = (s_1 + 1, 0, s_2 - 1, 0) \text{ if } s_1 \leq B_1,$$

and

$$\phi(\beta_2, B_1 + 1, 0, B_2 + 1, 0) = (B_1 + 1, 0, s_2, 1) \text{ if } s_1 > B_2,$$

Special cases go like

$$\phi(\beta_1, B_1 + 1, 0, s_2, 1) = (B_1 + 1, 0, s_2, 0),$$

and

$$\phi(\beta_2, s_1, 1, B_2 + 1, 0) = (s_1, 0, B_2 + 1, 0).$$

(b). To check (CC) for some state s, note that

$$\begin{aligned} \phi(\beta_1, \phi(\beta_2, (2, 0, B_2 + 1, 0)) &= \phi(\beta_1, (3, 0, B_2, 0)) \\ &= (2, 0, B_2 + 1, 0) \\ &= \phi(\beta_2, (2, 1, B_2 + 1, 0)) \\ &= \phi(\beta_2, \phi(\beta_1, (2, 0, B_2 + 1, 0)) \end{aligned}$$

Hence, (CC) holds. All other states follow the same way.