Université Paris 1 Panthéon-Sorbonne MMMEF
Stochastic Approximation Algorithm
Noufel Frikha

Practical Session 2 : Introduction to Stochastic Algorithms

This practical session provides an introduction to stochastic algorithms through two examples. These are optimisation procedures designed to estimate the zero θ^* of a function h of the form

$$h(\theta) = \mathbb{E}[H(\theta, U)],$$

where $H : \mathbb{R}^d \times \mathbb{R}^q \to \mathbb{R}^d$ and U is a random variable taking values in $\mathbb{R}^q$.

If the function h were computable at reasonable cost, one could implement the following deterministic algorithm :

$$\theta_{p+1} = \theta_p - \gamma_{p+1} h(\theta_p), \tag{1}$$

where $\theta_0 \in \mathbb{R}^d$ and $(\gamma_p)_{p\geq 1}$ is a positive sequence of step sizes satisfying

$$\sum_{p\geq 1} \gamma_p = +\infty \qquad \text{and} \qquad \sum_{p\geq 1} \gamma_p^2 < +\infty.$$

It is known that if h satisfies

$$\forall \theta \neq \theta^*, \quad \langle \theta - \theta^*, h(\theta) \rangle > 0, \tag{2}$$

and if, for some constant $C > 0$,

$$|h(\theta)|^2 \leq C(1 + |\theta - \theta^*|^2),$$

then the algorithm (1) satisfies

$$\theta_p \to \theta^* \qquad \text{as} \qquad p \to +\infty.$$

In many applications in mathematical finance, the function h cannot be computed explicitly. One therefore modifies (1) as follows :

$$\theta_{p+1} = \theta_p - \gamma_{p+1} H(\theta_p, U^{p+1}), \tag{3}$$

where $(U^p)_{p\geq 1}$ is an i.i.d. sequence with the same distribution as U. This defines a recursive algorithm in which, at each step p, the variable θ_p is updated using an independent simulation of the random variable U.

It can be shown (see the slides of Chapter 2) that if (2) holds and if H satisfies

$$\forall \theta \in \mathbb{R}^d, \quad \mathbb{E}[|H(\theta, U)|^2] \leq C(1 + |\theta - \theta^*|^2), \tag{4}$$

then the sequence $(\theta_p)_{p\geq 1}$ satisfies

$$\theta_p \xrightarrow{a.s.} \theta^*, \qquad \text{as} \qquad p \to +\infty.$$

1 Estimation of the Quantile in the Black–Scholes Model

We consider the Black–Scholes dynamics

$$S_T = S_0 \exp\left((\mu - \tfrac{\sigma^2}{2})T + \sigma W_T \right).$$

We fix $T = 1$, $\sigma = 40\%$, and $\mu = 10\%$. Our aim here is to estimate the quantile of order $\alpha \in (0, 1)$ of the distribution of S_T, that is, the unique θ^* solution of

$$\mathbb{P}(S_T \leq \theta) = \alpha. \tag{5}$$

A straightforward calculation shows that

$$\theta^* = S_0 \exp\left((\mu - \tfrac{\sigma^2}{2})T + \sigma\sqrt{T}\, \phi^{-1}(\alpha) \right), \tag{6}$$

where ϕ denotes the cumulative distribution function of the standard normal distribution $\mathcal{N}(0, 1)$.

1. Show that (5) can be reformulated as the search for the zero of a function h of the form $h(\theta) = \mathbb{E}[H(\theta, U)]$, where $H : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ and $U \sim \mathcal{N}(0, 1)$. Identify the function H.
2. Verify that assumptions (2) and (4) are satisfied.
3. Write a procedure implementing the stochastic algorithm (3).
4. Compare the value obtained after M iterations of the stochastic algorithm with the theoretical value (6). To this end, one can plot the trajectories of the stochastic approximation algorithm with respect to the number of iterates and compare it with the target value $\theta^\star$

2 Estimation of the Optimal Importance Sampling Parameter

In the introduction of the course, we saw that the price of a European vanilla option often boils down to estimating

$$\mathbb{E}[f(G)],$$

where $f : \mathbb{R}^d \to \mathbb{R}_+$ is the payoff function and $G \sim \mathcal{N}(0, I_d)$. This can be done by relying on the Monte Carlo method whose convergence rate crucially depends on the variance of $f(G)$.

The importance sampling method is a variance reduction method which in this context consists in minimising the strictly convex function $v \in \mathcal{C}^2(\mathbb{R}^d, \mathbb{R})$ defined by

$$v(\theta) = \mathbb{E}\left[f^2(G) e^{-\theta \cdot G - |\theta|^2/2}\right].$$

This function admits a unique minimiser θ^*, characterised as the unique solution of

$$\nabla v(\theta) = \mathbb{E}\left[f^2(G)(\theta - G) e^{-\theta \cdot G + |\theta|^2/2}\right] = 0.$$

In this section, we assume that f satisfies the following exponential growth assumption

$$|f(x)| \leq C e^{a|x|}.$$

1. Show that

$$\nabla v(\theta) = e^{|\theta|^2}\,\mathbb{E}\big[f^2(G-\theta)(2\theta - G)\big].$$

2. Define

$$H(\theta, G) = e^{-a(|\theta|+1)} f^2(G-\theta)(2\theta - G), \quad h(\theta) = \mathbb{E}[H(\theta, G)].$$

Show that $\{h = 0\} = \{\theta^*\}$ and that assumptions (2) and (4) hold provided

$$\mathbb{E}[e^{2a|G|}|G|^2] < +\infty.$$

3. Propose a procedure to estimate the parameter θ^* using a stochastic algorithm.
4. Consider a digital option with payoff $\mathbf{1}_{\{S_T > L\}}$, $L > 0$. The asset price follows the Black–Scholes dynamics

$$S_T = S_0 e^{(r-\frac{\sigma^2}{2})T + \sigma W_T}.$$

We choose $T = 1$, $S_0 = 100$, $\sigma = 0.2$, $r = 0.05$, $L = 140$. Fix $M = 10{,}000$ iterations of the stochastic algorithm. The exact price of the digital option is 0.05968. Implement the stochastic algorithm and demonstrate empirically that the new estimator achieves variance reduction.

3 Stochastic Mirror Descent for Expected Shortfall Minimisation

In this section, we focus on portfolio optimisation under the Expected Shortfall (ES) risk measure. Let $X \in \mathbb{R}^d$ denote the random vector of asset returns, and let $u \in \Delta_d$ be a portfolio weight vector on the simplex

$$\Delta_d = \left\{ u \in \mathbb{R}^d \ : \ u_i \geq 0, \ \sum_{i=1}^{d} u_i = 1 \right\}.$$

We consider the portfolio loss $L(u) = -\langle u, X\rangle$. For a fixed confidence level $\alpha \in (0,1)$, the Expected Shortfall at level α is defined by

$$\rho(L(u)) = \mathrm{ES}_\alpha(L(u)).$$

Using the Rockafellar-Uryasev representation, Expected Shortfall admits the variational formulation

$$\min_{u \in \Delta_d} \rho(-\langle u, X\rangle) = \min_{(u,\xi) \in \Delta_d \times \mathbb{R}} \mathbb{E}\left[\xi + \frac{1}{1-\alpha}\big(-\langle u, X\rangle - \xi\big)_+\right].$$

We therefore introduce the convex objective function

$$V(u, \xi) = \mathbb{E}\left[\xi + \frac{1}{1-\alpha}\big(-\langle u, X\rangle - \xi\big)_+\right].$$

1) Prove that the function V is convex on $\Delta_d \times \mathbb{R}$ and that for any $u \in \Delta_d$

$$\lim_{|\xi| \to \infty} V(u, \xi) = +\infty.$$

Stochastic Gradient Structure

Let $(X^p)_{p\geq 1}$ be an i.i.d. sequence with the same distribution as X. A stochastic estimator of the objective is given by

$$H(u,\xi,X) = \xi + \frac{1}{1-\alpha}\big(-\langle u, X\rangle - \xi\big)_+.$$

2) Prove that the subgradient of V is given by $\mathbb{E}[G(u,\xi,X)]$ with

$$G(u,\xi,X) = \begin{pmatrix} -\dfrac{1}{1-\alpha} X \mathbf{1}_{\{-\langle u,X\rangle - \xi > 0\}} \\ 1 - \dfrac{1}{1-\alpha}\mathbf{1}_{\{-\langle u,X\rangle - \xi > 0\}} \end{pmatrix}.$$

Stochastic Mirror Descent Algorithm

We minimise V over $\Delta_d \times \mathbb{R}$ using stochastic mirror descent.

Geometry on the simplex. For the u-component, we use the entropic mirror map

$$\Psi(u) = \sum_{i=1}^{d} u_i \log(u_i),$$

whose associated Bregman divergence is the Kullback–Leibler divergence. This ensures that the iterates remain in Δ_d.

Euclidean geometry for ξ. For the scalar variable ξ, we use a standard Euclidean update

$$\Phi(\xi) = \frac{1}{2}\xi^2.$$

3) Prove that the SMD algorithm $(u_p, \xi_p)_{p\geq 0}$ writes as follows

Update of u :

$$u_{p+1,i} = \frac{u_{p,i}\exp\Big(-\gamma_{p+1}\frac{-X_i^{p+1}}{1-\alpha}\mathbf{1}_{\{-\langle u_p, X^{p+1}\rangle - \xi_p > 0\}}\Big)}{\sum_{j=1}^{d} u_{p,j}\exp\Big(-\gamma_{p+1}\frac{-X_j^{p+1}}{1-\alpha}\mathbf{1}_{\{-\langle u_p, X^{p+1}\rangle - \xi_p > 0\}}\Big)}. \tag{7}$$

Update of ξ :

$$\xi_{p+1} = \xi_p - \gamma_{p+1}\left(1 - \frac{1}{1-\alpha}\mathbf{1}_{\{-\langle u_p, X^{p+1}\rangle - \xi_p > 0\}}\right). \tag{8}$$

where $(\gamma_p)_{p\geq 1}$ is the step sizes.

Example : Gaussian Log-Returns Model

We consider a d-dimensional Black–Scholes-type model in discrete time. Let $(Y^1, \ldots, Y^d)$ denote the vector of log-returns over one period, and assume

$$Y \sim \mathcal{N}(m, \Sigma),$$

where

$$m = (m_1, \ldots, m_d)^\top \in \mathbb{R}^d, \qquad \Sigma \in \mathbb{R}^{d \times d}$$

is a positive definite covariance matrix. The asset returns are defined by

$$X_i = e^{Y^i} - 1, \qquad i = 1, \ldots, d.$$

Explicit specification.

For numerical experiments, one may choose for instance :

$$d = 3, \qquad m = \begin{pmatrix} 0.05 \\ 0.07 \\ 0.09 \end{pmatrix},$$

and

$$\Sigma = \begin{pmatrix} 0.04 & 0.02 & 0.01 \\ 0.02 & 0.09 & 0.03 \\ 0.01 & 0.03 & 0.16 \end{pmatrix}.$$

This corresponds to annual volatilities

$$\sigma_1 = 20\%, \quad \sigma_2 = 30\%, \quad \sigma_3 = 40\%,$$

with moderate positive correlations.

4) Implement and illustrate the convergence of the SMD in this example. In particular, to observe the convergence, one can plot the trajectories of $u_{p,1}$, $u_{p,2}$, $u_{p,3}$ and ξ_p with respect to p.
5) What is the influence of the choice of the step size $(\gamma_p)_{p\geq 1}$.