A Penalty-Based Exterior-Interior Point Method for Stochastically Constrained Simulation Optimization

Xiangyu Yang

School of Management, Shandong University

June 2025

Co-authors

A Penalty-Based Exterior-Interior Point Method for Stochastically Constrained Simulation Optimization

Jiaqiao Hu
(Stony Brook University)

Jianghua Zhang
(Shandong University)

Outline

Outline

What we did

Proposed a new algorithm for solving the following stochastically constrained simulation optimization problem of the form

$$\begin{aligned} \min_{x\in\Theta} \quad & f(x) := \mathrm{E}[F(x,\xi_f)] \\ \text{s.t.} \quad & g^{[\ell]}(x) := \mathrm{E}\left[G^{[\ell]}\left(x,\xi_g^{[\ell]}\right)\right] \le 0, \quad \forall \ell = 1,\dots,q. \end{aligned}$$

Proved the almost sure convergence and expected convergence rate of the algorithm, and numerically validated its effectiveness.

Application example: NPC

Neyman–Pearson Classification (NPC)

$$\text{type I error} := \mathrm{E}[\mathrm{I}^+\{\hbar(a;x)\}|b=-1],$$
$$\text{type II error} := \mathrm{E}[\mathrm{I}^+\{-\hbar(a;x)\}|b=+1],$$

where $\hbar(a;x)$ is a classifier parameterized by x, trained using a dataset a, $b \in \{-1,+1\}$ is a label, and $\mathrm{I}^+\{\cdot\}$ is an indicator such that $\mathrm{I}^+\{z\} = 1$ if $z \geq 0$ and $\mathrm{I}^+\{z\} = 0$ if $z < 0$.

$$\begin{aligned} \min_x \quad & \text{type II error} \\ \text{s.t.} \quad & \text{type I error} \leq \imath, \end{aligned}$$

where $\imath > 0$ is a given threshold.

What's new

$$\begin{aligned} \min_{x\in\Theta} \quad & f(x) := \mathrm{E}[F(x,\xi_f)] \\ \text{s.t.} \quad & g^{[\ell]}(x) := \mathrm{E}\left[G^{[\ell]}\left(x,\xi_g^{[\ell]}\right)\right] \le 0, \quad \forall \ell = 1,\ldots,q. \end{aligned}$$

- Sample average approximation (SAA) vs.
  <u>Stochastic gradient descent (SGD)</u>
- Unbiased stochastic gradient with bounded variance vs.
  <u>Biased stochastic gradient with unbounded variance</u>
- Explicit feasibility check vs. <u>Implicit feasibility check</u>

Outline

New penalty functions

Define a family of parameterized penalty functions $\{r_\varepsilon(\cdot) : \varepsilon > 0\}$. We require them to meet some technical conditions.

Examples 2.1

- exponential penalty: $r_\varepsilon(y) = \exp(y/\varepsilon)$
- logarithmic (softplus) penalty: $r_\varepsilon(y) = \ln(1 + \exp(y/\varepsilon))$

Define a (deterministic) penalized problem:

$$\min_{x \in \Theta} h_\varepsilon(x) := f(x) + \sum_{\ell=1}^{q} r_\varepsilon\left(g^{[\ell]}(x)\right).$$

Denote by $x_\varepsilon^* := \arg\min_{x \in \Theta} h_\varepsilon(x)$ the penalized solution.

Convergence for the deterministic problem

Proposition 2.2

Given any positive sequence $\{\varepsilon_k\}$ such that $\varepsilon_k \to 0$. Applying the proposed penalty functions. Then there exists an integer N such that for all $k \geq N$, $\max_\ell g^{[\ell]}(x^*_{\varepsilon_k}) < 0$, and $\lim_{k\to\infty} f(x^*_{\varepsilon_k}) = f(x^*)$.

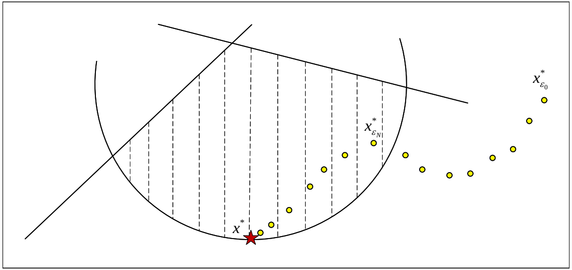

Figure 1: Illustration of the convergence behavior of the penalized problem

Multi-timescale stochastic approximation (SA)

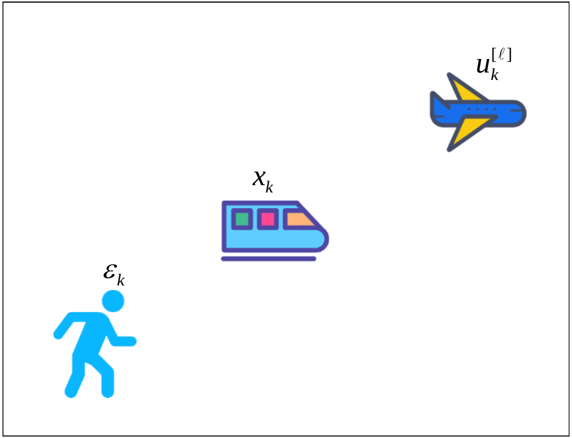

Figure 2: Multi-timescale SA algorithm design

Outline

Convergence

Theorem 3.1

Suppose $h_\varepsilon(\cdot)$ is strictly convex. Under appropriate conditions (where the stochastic gradients are asymptotically unbiased and their variances can be increasing), the sequence $\{f(x_k)\}$ almost surely converges to $f(x^*)$; i.e.,

$$f(x_k) \to f(x^*) \text{ as } k \to \infty \text{ w.p.1.}$$

Rate of convergence

Applying logarithmic penalty functions; objective function $f(\cdot)$ is strongly convex, constraint $g^{[\ell]}(\cdot)$ is convex.

- Step sizes: α_k, β_k, γ_k
- Decreasing rate of biases: b_k
- Increasing rate of variances: σ_k

Theorem 3.2

Under appropriate conditions,

$$\mathrm{E}[f(x_k) - f(x^*)]$$
$$= O(b_k) + O\left(\frac{\beta_k \sigma_k^2}{\alpha_k \varepsilon_k^3}\right) + O\left(\sqrt{\frac{\alpha_k}{\varepsilon_k}}\sigma_k\right) + O\left(\frac{\gamma_k}{\beta_k \varepsilon_k}\right) + O\left(\varepsilon_k \ln\left(\frac{1}{\varepsilon_k}\right)\right).$$

Rate of convergence (cont'd)

Use unbiased stochastic gradient estimators with bounded variances

- The optimal solution x^* is on the boundary of the feasible region: Rate $O(k^{-1/12})$
- The optimal solution x^* is in the interior of the feasible region: Rate $O(k^{-1/4})$
- ε-suboptimal solution: Rate $O(k^{-1/3})$

Numerical experiment

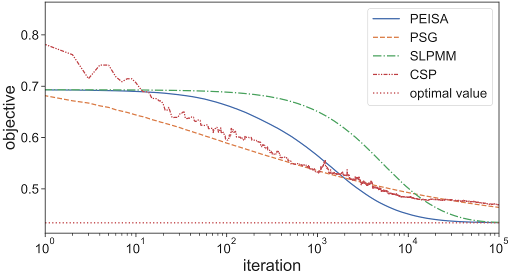

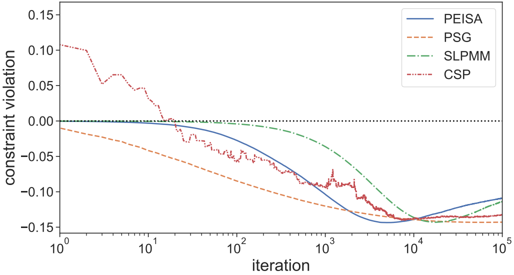

objective/iteration

constraint violation/iteration

Figure 3: Performance of comparison algorithms on the NPC problem, based on 30 independent runs

Outline

Takeaway

- A Novel penalty-based exterior-interior point method
- Multi-timescale SA coordinates the multi-layer loop problem
- Explore more scenarios of constrained optimization

Thank you for your time, and welcome your comments and suggestions!