Stochastic Optimization under Hidden Convexity

ETH zürich — EPFL

Ilyas Fatkhullin Niao He Yifan Hu

Problem Setting

Nonconvex Optimization:

$$\min_{x\in\mathcal{X}} F(x) := \mathbb{E}_{\xi\sim\mathcal{D}}[f(x,\xi)],$$

$\mathcal{X} \subset \mathbb{R}^d$ – convex.

Convex Reformulation:

$$\min_{u\in\mathcal{U}} H(u) := F(c^{-1}(u)),$$

$\mathcal{U} = c(\mathcal{X}) \subset \mathbb{R}^d$ – convex.

Unknown distribution $\mathcal{D}$ and transformation $c(\cdot)$.

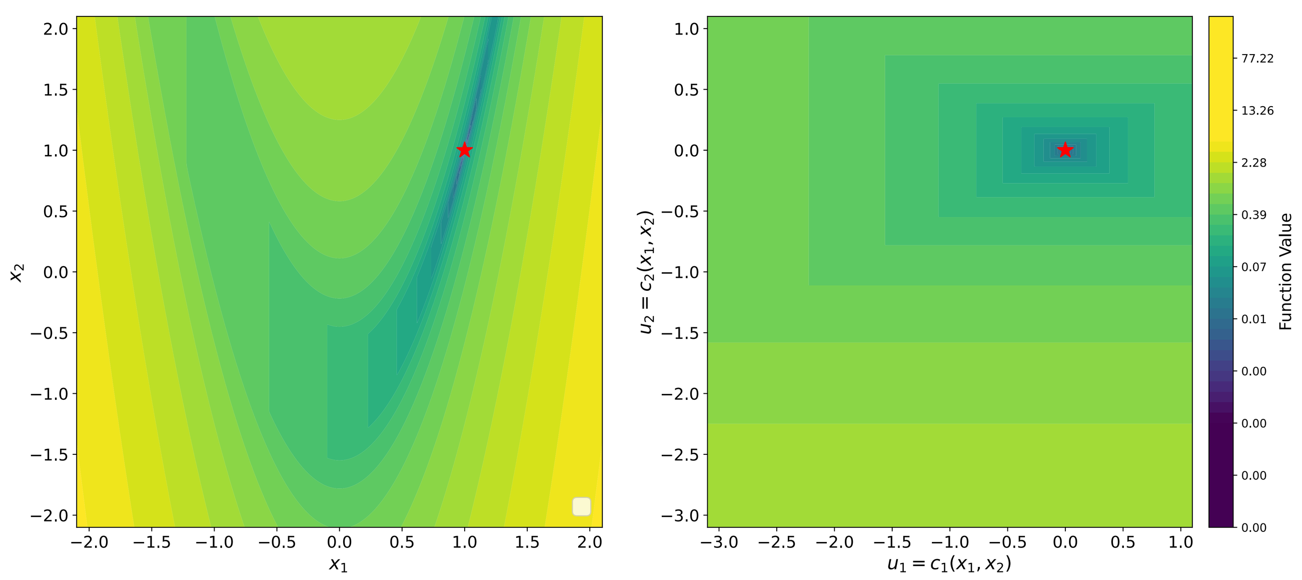

Example: $F(x_1,x_2) = \max\left\{\frac{1}{4}|x_1-1|, \frac{1}{2}|2x_1^2 - x_2 - 1|\right\}$.

Motivating Examples

Convex Reinforcement Learning [1]. MDP $\mathbb{M}(\mathcal{S},\mathcal{A},\mathcal{P},H,\rho,\gamma)$.
Parameter of a policy $\pi \in \Pi$, $\Pi \subset \mathbb{R}^{|\mathcal{S}|\times|\mathcal{A}|}$ is product of simplex sets.
State-action occupancy measure λ^π for $\pi \in \mathcal{X}$.

$$\lambda^\pi(s,a) := \sum_{h=0}^{+\infty} \gamma^h \mathbb{P}_{\rho,\pi}(s_h = s, a_h = a),$$

$H : \mathcal{U} \to \mathbb{R}$ is a general (convex) utility. The goal

$$\min_{\pi\in\mathcal{X}} F(\pi) := H(\lambda^\pi).$$

- Standard RL, $H(\lambda^\pi) = r^\top\lambda^\pi$ is linear in λ^π.
- Pure exploration, $H(\lambda^\pi)$ – negative entropy of λ^π.
- Imitation learning, $H(\lambda^\pi)$ is KL divergence.

Revenue Management and Inventory Control [2].

$$\min_{x\in[0,D]^d} F(x) := \mathbb{E}_\xi[f(x\wedge\xi)]$$
$$H(u) := \mathbb{E}_\xi[f(c^{-1}(u)\wedge\xi)].$$

Under transformation $u = c(x) = \mathbb{E}_\xi[x\wedge\xi]$, $H(u)$ is convex.

- Revenue management: $f(x) = r^\top x - \mathbb{E}_\eta\Gamma(x,\eta)$.
  Booking limit threshold v.s. Expected accepted reservations.
- Inventory with random capacity/supply: $f(x)$ newsvendor objectives.
  Ordering quantity v.s. Expected replenishment.

System Level Synthesis in Optimal Control [3].
Dynamics with finite horizon: $x(t+1) = A_t x(t) + B_t u(t) + w(t)$,

$$\min_{\mathbf{K}} F(\mathbf{K}) := \mathbb{E}\left[\mathbf{x}^\top\mathcal{Q}\mathbf{x} + \mathbf{u}^\top\mathcal{R}\mathbf{u}\right], \quad u(t) = \sum_{i=0}^{t} K(t,t-i)x(i).$$

Original control variable $\mathbf{K}$ v.s. New variable $\Phi := (\Phi_\mathbf{u}, \Phi_\mathbf{x})$.
Variable change: $\mathbf{K} = c^{-1}(\Phi) := \Phi_\mathbf{u}\Phi_\mathbf{x}^{-1}$; $\Phi_\mathbf{x}, \Phi_\mathbf{u}$ lower-block-triangular.

$$\min_{\Phi_\mathbf{x},\Phi_\mathbf{u}} H(\Phi_\mathbf{x},\Phi_\mathbf{u}), \quad \text{s.t. } \mathbf{M}\begin{bmatrix}\Phi_\mathbf{x}\\ \Phi_\mathbf{u}\end{bmatrix} = I, \quad H(\cdot) \text{ is quadratic in } \Phi_\mathbf{x}, \Phi_\mathbf{u}.$$

Transformation $c(\cdot)$ is Unknown

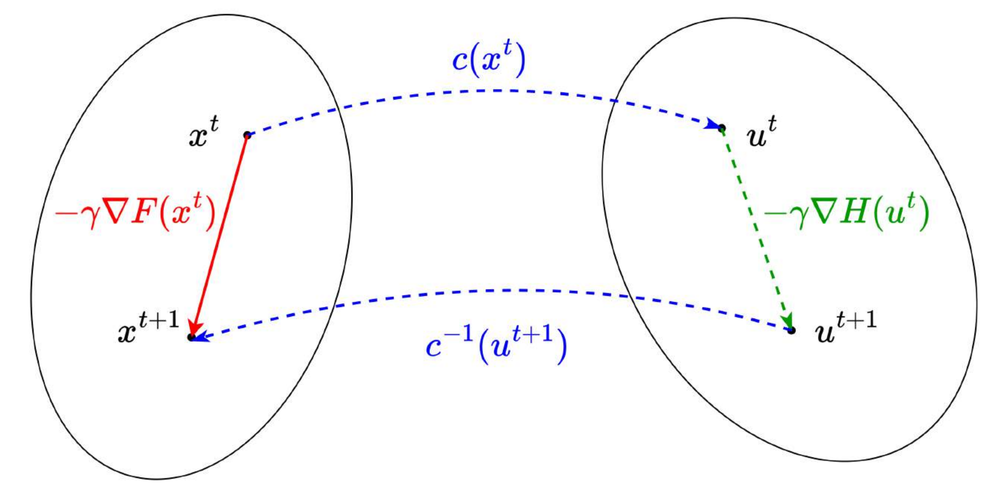

How does Projected (Sub)gradient Method behave?

- Simple to implement.
- Does not require transformation information.
- Run in an online fashion.

(Implicit) Hidden Convexity

C.1. $H : \mathcal{U} \to \mathbb{R}$ is convex and $\min_{u\in\mathcal{U}} H(u)$ admits a solution.

C.2. $c : \mathcal{X} \to \mathcal{U}$ is invertible. There exists $\mu_c > 0$ such that

$$\|c(x) - c(y)\| \geq \mu_c\|x-y\| \quad \text{for all } x,y\in\mathcal{X}.$$

Proposition 1. Let **C.1.** and **C.2.** hold. For any $\alpha \in [0,1]$, $x^* \in \mathcal{X}^*$ and $x \in \mathcal{X}$, define $x_\alpha := c^{-1}((1-\alpha)c(x) + \alpha c(x^*))$. Then

$$F(x_\alpha) \leq (1-\alpha)F(x) + \alpha F(x^*), \qquad \|x_\alpha - x\| \leq \frac{\alpha}{\mu_c}\|c(x) - c(x^*)\|.$$

Subgradient Method

Non-smooth setting:
A.1. $F(\cdot)$ is ℓ-weakly convex, i.e., $F(x) + \frac{\ell}{2}\|x-y\|^2$ is convex in x.
A.2. Stochastic sub-gradients with $\mathbb{E}[g(x,\xi)] \in \partial F(x)$ and

$$\mathbb{E}\left[\|g(x,\xi)\|^2\right] \leq G_F^2.$$

Remark. If $H(\cdot)$ is Lipschitz and $c(\cdot)$ is smooth, then **A.1.** holds.

$$\textbf{SM}: \quad x^{t+1} = \Pi_\mathcal{X}(x^t - \eta\, g(x^t,\xi^t)).$$

Analysis based on Moreau envelope [4]:

$$\Lambda_t^{\text{SM}} := \mathbb{E}\left[F_{1/\rho}(x^t) - F(x^*)\right],$$

$$F_{1/\rho}(x) := \min_{y\in\mathcal{X}}\left\{F(y) + \frac{\rho}{2}\|y-x\|^2\right\}.$$

Convergence of SM

Theorem 1. Let **C.1.**, **C.2.**, **A.1.**, **A.2.** hold, $\mathrm{diam}(\mathcal{U}) \leq D_\mathcal{U}$. Fix $\varepsilon > 0$, set $\eta = \frac{1}{2\ell}\min\left\{1, \frac{\mu_c^2\varepsilon^2}{D_\mathcal{U}^2 G_F^2}\right\}$. Then we have $\Lambda_T^{\text{SM}} \leq \varepsilon$ after

$$T = \widetilde{\mathcal{O}}\left(\frac{\ell D_\mathcal{U}^2}{\mu_c^2}\frac{1}{\varepsilon} + \frac{\ell D_\mathcal{U}^4 G_F^2}{\mu_c^4}\frac{1}{\varepsilon^3}\right).$$

- $F(\cdot)$ is non-smooth/non-convex, but **SM** converges in function value.
- Theorem 1 extends to smooth case under **A.1.'** and **A.2.'**.

Projected SGD with Momentum

Smooth setting:
A.1.' $F(\cdot)$ is L-smooth.
A.2.' Stochastic gradients with $\mathbb{E}[\nabla f(x,\xi)] = \nabla F(x)$:

$$\mathbb{E}\left[\|\nabla f(x,\xi) - \nabla F(x)\|^2\right] \leq \sigma^2.$$

$$\textbf{Proj-SGDM}: \quad \begin{aligned} x^{t+1} &= \Pi_\mathcal{X}(x^t - \eta\, g^t),\\ g^{t+1} &= (1-\beta)\, g^t + \beta\,\nabla f(x^{t+1},\xi^{t+1}).\end{aligned}$$

Analysis based on Lyapunov function [5]:

$$\Lambda_t^{HB} := \mathbb{E}\left[F(x^t) - F(x^*) + \frac{\eta}{\beta}\left\|g^t - \nabla F(x^t)\right\|^2\right].$$

Convergence of Proj-SGDM

Theorem 2. Let **C.1.**, **C.2.**, **A.1.'**, **A.2.'** hold and $\mathrm{diam}(\mathcal{U}) = D_\mathcal{U}$. Fix $\varepsilon > 0$, set $\eta = \frac{\beta}{4L}$, $\beta = \min\left\{1, \frac{\mu_c^2\varepsilon^2}{D_\mathcal{U}^2\sigma^2}\right\}$. Then we have $\Lambda_T^{\text{HB}} \leq \varepsilon$ after

$$T = \widetilde{\mathcal{O}}\left(\frac{L D_\mathcal{U}^2}{\mu_c^2}\frac{1}{\varepsilon} + \frac{L D_\mathcal{U}^4\sigma^2}{\mu_c^4}\frac{1}{\varepsilon^3}\right).$$

- Last iterate convergence.
- We have $F(x^t) \to F(x^*)$ and $g^t \to \nabla F(x^*)$ in expectation as $t \to \infty$.
- When $H(\cdot)$ is μ_H-strongly convex: $T = \widetilde{\mathcal{O}}\left(\frac{L}{\mu_c^2\mu_H} + \frac{L\sigma^2}{\mu_c^4\mu_H^2}\frac{1}{\varepsilon}\right)$.

References

[1] J. Zhang, C. Ni, C. Szepesvari, M. Wang. *On the convergence and sample efficiency of variance-reduced policy gradient method.* NeurIPS 2021.
[2] X. Chen, N. He, Y. Hu, Z. Ye. *Efficient Algorithms for Minimizing Compositions of Convex Functions and Random Functions and Its Applications in Network Revenue Management.* arXiv:2205.01774, 2022.
[3] J. Anderson, J. C. Doyle, S. H. Low, N. Matni. *System level synthesis.* Annual Reviews in Control 2019.
[4] D. Davis, D. Drusvyatskiy. *Stochastic subgradient method converges at the rate $O(k^{-1/4})$ on weakly convex functions.* arXiv:1802.02988, 2018.
[5] I. Fatkhullin, A. Tyurin, P. Richtárik. *Momentum provably improves error feedback!* NeurIPS 2023.