Optimization with Markovian Noise

Zero Order Upper bound

Chebykin Semen

Adviser - Aleksandr Beznosikov

May 17, 2024

Table of Contents

Table of Contents

Stochastic optimization problem

$$\min_{x \in \mathbb{R}^d} f(x) := \mathbb{E}_{\xi \sim D}[F(x, \xi)] \tag{1}$$

is well-studied [3].

We consider a generalization:

$$\min_{x \in \mathbb{R}^d} f(x) := \mathbb{E}_{Z \sim \pi}[F(x, Z)] \tag{2}$$

Where Z is a Markov Chain and π is it's invariant distribution.

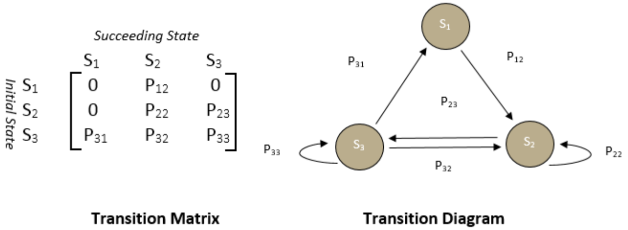

- Decentralized optimization (Token algorithms, Privacy)
- RL (Markov Decision Processes)

Table of Contents

Our research aims to expand the results of [2].

Following the article, we consider additional assumptions:

- (A1,A2) f is (μ, L)-smooth.
- (A3) Z is uniformly geometrically ergodic with mixing time τ.
- (A4) F is (σ, δ)-bounded:

$$\|\nabla F(x, Z) - \nabla f(x)\|^2 < \sigma^2 + \delta^2 \|\nabla f(x)\|^2$$

The main result of [2] is the following

Theorem. Under A1-A4 the problem can be solved (in terms of $\mathbb{E}\left[\|x^{(k)} - x^*\|^2\right] \leq \varepsilon$) in

$$\tilde{\mathcal{O}}\left(\tau\left[(1+\delta^2)\sqrt{\frac{L}{\mu}}\log\frac{1}{\varepsilon} + \frac{\sigma^2}{\mu^2\varepsilon}\right]\right) \quad \text{oracle calls}\,. \tag{3}$$

- Our goal is to adapt algorithm from [2] to zero-order oracle, meaning we can only access $F(x, Z)$ and not $\nabla F(x, Z)$.
- Moreover, we are working under one-point feedback, i.e. we can only access $F(\cdot, Z_1)$ with the same randomness once.
- Since we do not have access to ∇F, A4 transforms into: (A4') $F(x, Z) = f(x) + h(x, Z)$ with $h(x, Z)$ uniformly bounded:

$$|h(x, Z)|^2 < \sigma^2$$

Results

Tools

- $\nabla f \approx \underbrace{d\frac{f(x+te)-f(x-te)}{2t}e}_{*} \approx d\frac{F(x+te,Z_1)-F(x-te,Z_2)}{2t}e$
- But from [1]: $E_e[*] = \nabla \tilde{f}$ for $\tilde{f} = E_e f(x+te)$
- And $\tilde{f}$ is also (μ, L)-smooth and $|\tilde{f} - f| \leq Lt^2$
- Adapted from [2]:
  For any distribution ξ: $E_\xi[n^{-1}\sum_{i=1}^{n} h(x, Z_i)^2] \leq \frac{C_1\tau}{n}\sigma^2$

Under assumptions (A1-A4') the problem 2 can be solved (in terms of $\mu\|x^N - x^*\| + f(x^N) - f(x^*)$) in N iterations:

$$N = \mathcal{O}\left(d\sqrt{\frac{L}{\mu}}\log\frac{1}{\varepsilon} + d^2\frac{\tau L}{\mu\varepsilon^2}\sigma^2 + d^2\frac{L}{\mu}\right).$$

By setting $\tau = 1$ we can transform the problem into iid setting, in which the optimal results are of form $\mathcal{O}(\frac{d^2}{\mu\varepsilon^2})$

Table of Contents

Future work

- Lower bounds
- Better argument convergence
- Two-point feedback

Table of Contents

Bibliography I

[1] Arya Akhavan, Massimiliano Pontil, and Alexandre B. Tsybakov. "Exploiting Higher Order Smoothness in Derivative-free Optimization and Continuous Bandits". In: *CoRR* abs/2006.07862 (2020). arXiv: 2006.07862. URL: https://arxiv.org/abs/2006.07862.

[2] Aleksandr Beznosikov et al. *First Order Methods with Markovian Noise: from Acceleration to Variational Inequalities*. 2023. arXiv: 2305.15938 [math.OC].

[3] Boris Polyak. *Introduction to Optimization*. July 2020.