Bridging Practical Needs and Theoretical Analysis of Machine Learning Algorithms

Shaocong Ma

November 1, 2023

University of Utah
s.ma@utah.edu

Shaocong Ma

Backgrounds

- Final year PhD student in EE (University of Utah).
- M.A. Degree in Statistics (UC Santa Barbara).
- B.S. Degree in Statistics (Sichuan University).

Research Fields

- *Optimization and Stochastic Approximation*: I re-investigate unreasonable theoretical assumptions based on practical needs.
- *Reinforcement Learning*: I design data-efficient algorithms that are substantiated with theoretical guarantees.

Overview

Optimization and
Stochastic Approximation

- Challenges from Real-World: Non-IID Data
- Challenges from Real-World: Non-Differentiability
- More Complicated Structure: Bilevel/Minimax Optimization

Reinforcement Learning

- Efficient Algorithms in Reinforcement Learning
- Robust Algorithms in Reinforcement Learning

Optimization and Stochastic Approximation

Challenges from Real-World: Non-IID Data

Example (Federated Learning)
The data points $x^{(i)}$ and $x^{(j)}$ in FL come from different devices.

$$x^{(i)} \sim X^{(i)}, x^{(j)} \sim X^{(j)}$$

Each device may have non-identical distribution.

Figure 1: Illustration of Federated Learning

Challenges from Real-World: Non-IID Data

Example (Reinforcement Learning)
The data point (s_t, a_t, r_t, s_{t+1}) in RL comes from a trajectory:

$$s_1, a_1, r_1, s_2, a_2, r_2, \ldots$$

Not ind. + Non-identical distribution.

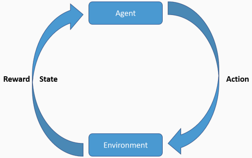

Figure 2: Agent-Environment Interaction

Characterization of Data Dependency: We use the mixing coefficient to measure the data dependency.

Definition

- $\{\xi_t\}_t$: a process with a stationary distribution μ.
- $\mathbb{P}(\xi_{t+k} \in \cdot | \mathcal{F}_t)$: the dist. of ξ_{t+k} cond. on $\mathcal{F}_t$.
- d_{TV}: the total variation distance.

The process $\{\xi_t\}_t$ is called ϕ-mixing if

$$\underbrace{\phi(k)}_{\text{mixing coef.}} := \sup_{t\in\mathbb{N}, A\in\mathcal{F}_t} 2d_{\mathrm{TV}}\Big(\mathbb{P}(\xi_{t+k} \in \cdot | A), \mu\Big) \to 0,$$

as $k \to \infty$.

Solutions to Non-IID Data

Example (Highly-Dependent Dataset)

- High-frequency trading;
- stock prices (B-S model with jumps);
- defaultable bonds (CIR Diffusion);
- digital currency (most noisy market).

Convergence of SGD over Dependent Data

	SGD	Mini-batch SGD
Weakly dependent; $\phi(k) \sim e^{-k}$	$\tilde{O}(\varepsilon^{-2})$	$O(\varepsilon^{-2})$
Medium dependent; $\phi(k) \sim \frac{1}{k^\theta}, \theta \geq 1.$	$O(\varepsilon^{-2-\frac{2}{\theta}})$	$\tilde{O}(\varepsilon^{-2})$
Highly dependent; $\phi(k) \sim \frac{1}{k^\theta}, 0 < \theta < 1.$	$O(\varepsilon^{-2-\frac{2}{\theta}})$	$O(\varepsilon^{-1-\frac{1}{\theta}})$

Table 1: Using a larger batch size is always helpful for dependent data[1].

[1] Ma, S., et al. "Data Sampling Affects the Complexity of Online SGD over Dependent Data." UAI 2022.

Future Works on Non-IID Data

- **Federated Learning on Non-IID Data**
  - Optimization theory without accessing the private data in each device.
- **Cybersecurity Problems**
  - Attacks generated from an adaptive adversarial attacker forms a Non-IID dataset.
- **Large Language Models/Multi-Modal Models**
  - Videos and speeches naturally form Non-IID Non-Stationary datasets.

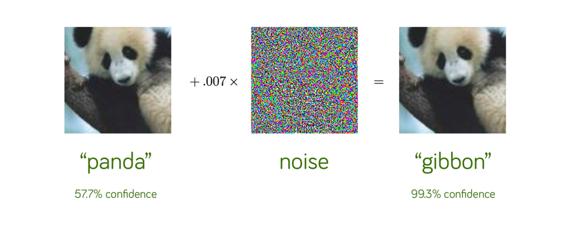

Figure 3: Adversarial Attacking

Figure 4: Gibbon

Challenges from Real-World: Non-Differentiability

Example (Air Flow Prediction)
Most of existing physical simulators (Matlab, SU2, etc.) are not differentiable or require additional efforts to implement differentiability.

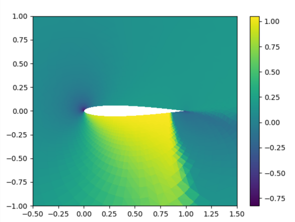

Figure 5: Simulation of Air Flow near an Airfoil

Solution: We propose the hybrid model[2]:

- optimize non-differentiable params using zeroth-order method;
- optimize neural network using standard optimizer.

[2]Ma, S., et al. "End-to-End Mesh Optimization of a Hybrid Deep Learning Black-Box PDE Solver." NeurIPS 2023 (ML4PS Workshop)

Challenges from Real-World: Non-Differentiability

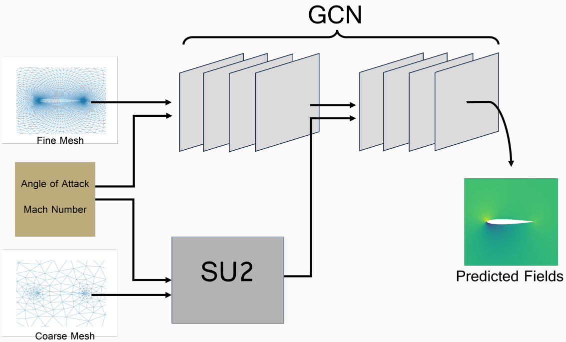

Figure 6: Physic-Informed GCN Model for Fluid Flow Prediction

We also focus on more complicated optimization problem:

Standard Opt. Problem $\min_{x \in X} f(x)$

Minimax Opt. Problem $\min_{x \in X} \max_{y \in Y} f(x, y)$

Applications: Adversarial Attacking[3], Reinforcement Learning[4], etc.

[3] Chen, Ziyi, Shaocong Ma, and Yi Zhou. "Accelerated Proximal Alternating Gradient-Descent-Ascent for Nonconvex Minimax Machine Learning." IEEE ISIT 2022.

[4] Chen, Ziyi, Shaocong Ma, and Yi Zhou. "Sample efficient stochastic policy extragradient algorithm for zero-sum markov game." ICLR 2021.

Reinforcement Learning

Efficient Algorithms in Reinforcement Learning

Key challenges in RL: Expensive Agent-Environment Interaction; e.g. Self-driven car.

Solutions: Propose efficient algorithms with optimal sample complexity.

- *Policy Evaluation*[5]: Quantify the performance of an agent.
- *Optimal Control*[6]: Learn the best strategy in a specific task.
- *Solve the Equilibrium of Stochastic Games*[7]: Learn the best strategy in a competitive environment.

---

[5] Ma, Shaocong, Yi Zhou, and Shaofeng Zou. "Variance-Reduced Off-Policy TDC learning: Non-Asymptotic Convergence Analysis." NeurIPS 2020.

[6] Ma, S., et al. "Greedy-GQ with Variance Reduction: Finite-Time Analysis and Improved Complexity." ICLR 2021.

[7] Chen, Ziyi, Shaocong Ma, and Yi Zhou. "Finding correlated equilibrium of constrained Markov game: A primal-dual approach." NeurIPS 2022.

Key challenges in RL: Environment is changing over time; e.g. Self-driven car.

Solutions: Let the agent consider the worst-case environment.

- *Solve the Robust Equilibrium of Stochastic Games*[8]: Quantify the performance of an agent.

[8] Ma, S., et al. "Decentralized Robust V-Learning for Solving Markov Games with Model Uncertainty." Submitted.

Thank You!

Any Questions?