A FTPL Approach to Online Non-convex Optimization

CSIE 5002 – Prediction, Learning, and Games

Cheng-Kang Chou, Yu-Chieh Kuo, Chao-Hsun Yang

Convexity

Why is convexity important?

Convexity:

- Ensures that the local minimum is the global minimum.
- Optimality condition helps capture the minimizer.

Definition (Non-convex)

We say a problem is convex if both the objective function f and the constraint set $\mathcal{X}$ are convex; otherwise, we say it is a non-convex problem.

From Convex to Non-convex

Convex:

- Convexity ensures the global minimum.

Non-convex:

- There exists potentially many local minima.
- Finding the global minima among all the local minima is hard (at least NP-Hard). [1]
- Saddle points present.

[1] See an overview by Danilova et al. (2021).

Challenge in Online Non-convex Optimization

- Suggala and Netrapalli (2020) prove that no deterministic algorithm can achieve $o(1)$ regret in the setting of online non-convex learning.
- The broadly developed methods such as Follow the Regularized Leader (FTRL) and Online Mirror Descent (OMD) cannot perform well in the non-convex setting.

A Simple Neural Network (NN) Example

Algorithm (Model Training in Classification)

Given training data $(x, y) \in (\mathcal{X}, \mathcal{Y})$, the activation function, and the loss function, LEARNER announces a weight $w_t \in \mathcal{W}$ corresponding to model structure at t-th epoch and receives a cost $f : (\mathcal{X}, \mathcal{Y}) \times \mathcal{W} \to \mathbb{R}$.

- Imagine an example where the activation function is `Sigmoid`, which is non-convex to $w \in \mathcal{W}$.

Non-convex Applications

Neural Networks:

- Szegedy et al. (2013) specify loss functions for neural networks are non-convex in general.
- Applications include adversarial training (Generative Adversarial Networks, GAN), efficient computation of an equilibrium in non-convex games, etc.

Optimization

Reduce Regret Minimization to Optimization

Online-to-offline reduction:

- Kalai and Vempala (2005) introduce Follow the Perturbed Leader (FTPL) to the regret minimization problem.
- The action set and loss functions for LEARNER is not assumed to be convex.
- The minimizer is assumed to be computed efficiently.

Some heuristic methods can find approximate global optima reasonably fast. See Drori and Shamir (2020).

Optimization Oracle

Definition (Value Oracle)

A value oracle is a procedure $\texttt{Val} : \mathcal{X} \times \mathcal{Y} \to [0,1]$ that for any action pair $x \in \mathcal{X}$, $y \in \mathcal{Y}$, return the loss value $\ell(x, y)$ in time $\mathcal{O}(1)$.

Definition (Optimization Oracle)

An optimization oracle is a procedure $\texttt{Opt}$ that receives a distribution $q \in \Delta(\mathcal{Y})$ and returns the best performing action w.r.t. q in $\mathcal{O}(1)$ on any input, i.e.,

$$\forall q \in \Delta(\mathcal{Y}),\ \texttt{Opt}(q) \in \arg\min_{x \in \mathcal{X}} \mathbb{E}_{y \sim q}[\ell(x, y)].$$

Efficient Computation Assumption

Online with offline:

- Littlestone (1989) shows that efficient online computation implies efficient offline approximation.
- What Kalai and Vempala (2005) and we discuss here is how to use efficient offline optimization algorithms for the online problem.

FTPL

Preliminaries

Setting and assumptions:

- $\mathcal{X} \subseteq \mathbb{R}^d$ is the action set for LEARNER which is bounded with ℓ_∞ diameter of D, i.e., $D = \sup_{x,y \in \mathcal{X}} \|x - y\|_\infty$.
- The sequence of loss function f_t chosen by REALITY are L-Lipschitz w.r.t. ℓ_1 norm, i.e., $|f_t(x) - f_t(y)| \leq L\|x - y\|_1$ for all $x, y \in \mathcal{X}$.
- LEARNER aims to choose a sequence of actions $\{x_t\}_{t=1}^T$ to minimize the regret

$$R_T(x) = \frac{1}{T}\sum_{t=1}^{T} f_t(x_t) - \frac{1}{T}\inf_{x \in \mathcal{X}} \sum_{t=1}^{T} f_t(x).$$

Recall of Follow the Leader (FTL)

Algorithm (Follow the Leader, FTL)

LEARNER plays a strategy $x_t \in \mathcal{X}$ at t-th round with minimal loss over past rounds, i.e.,

$$x_t = \arg\min_{x \in \mathcal{X}} \sum_{i=1}^{t-1} f_i(x).$$

Follow the Perturbed Leader (FTPL)

Randomness:

- Following FTL, we perturb the cumulative loss by adding a random perturbation at each round.

Algorithm 1 (FTPL, Agarwal et al. (2019))

Given a parameter $\eta > 0$, LEARNER draws an i.i.d. random vector $\sigma_t \sim (\exp(\eta))^d$ and plays a strategy

$$x_t = \arg\min_{x \in \mathcal{X}} \left\{ \sum_{i=1}^{t-1} f_i(x) - \sigma_t^\top x \right\}$$

at t-th round.

Analysis

Complexity:

- Algorithm 1 achieves the regret bound $\mathbb{E}[R_T(x)] \leq \mathcal{O}\left(T^{-1/3}\right)$.
- Cesa-Bianchi and Lugosi (2006) show that any algorithm that is guaranteed to work against an oblivious REALITY also works for a non-oblivious REALITY. That is, it suffices to work with a single random vector σ.

Strategy as an Oracle

LEARNER's strategy can be regarded as an offline optimization oracle.

Definition (Offline optimization oracle)

An offline optimization oracle is a procedure `Off-Opt` that receives a sequence of loss functions $f_1, \cdots, f_k$ and a d-dimensional vector σ and returns the best performing action

$$x^\star = \arg\min_{x \in \mathcal{X}} \left\{ \sum_{i=1}^{k} f_i(x) - \sigma^\top x \right\}.$$

Approximate Optimization Oracle

Definition ((α, β)-approximate optimization oracle)

An optimization oracle is called (α, β)-approximate optimization oracle (`Approx-Opt`) if it returns $x^\star \in \mathcal{X}$ such that

$$f(x^\star) - \sigma^\top x^\star \leq \inf_{x \in \mathcal{X}} \{f(x) - \sigma^\top x\} + (\alpha + \beta \|\sigma\|_1).$$

We denote it with $\texttt{Approx-Opt}_{\alpha,\beta}(f - \langle \sigma, \cdot \rangle)$.

FTPL with Approximate Optimization Oracle

Algorithm 2 (FTPL, Suggala and Netrapalli (2020))

Given a parameter $\eta > 0$, LEARNER draws an i.i.d. random vector $\sigma_t \sim (\exp(\eta))^d$ and plays a strategy

$$x_t = \texttt{Approx-Opt}_{\alpha,\beta}\left(\sum_{i=1}^{t-1} f_i - \langle \sigma_t, \cdot \rangle\right)$$

at t-th round.

Performance Analysis

- Suggala and Netrapalli (2020) give the regret bound

$$\mathbb{E}[R_T(x)] \leq \mathcal{O}\left(\eta d^2 D L^2 + \frac{d(\beta T + D)}{\eta T} + \alpha + \beta d L\right).$$

- By setting appropriate $\eta = \left(\sqrt{dT} - L\right)^{-1}$ and when $\alpha = \mathcal{O}\left(T^{-1/2}\right), \beta = \mathcal{O}(T^{-1})$, FTPL achieves the $O\left(T^{-1/2}\right)$ regret bound.

OFTPL

Online Learning with Predictable Sequences

Optimistic case:

- The general online learning setting assumes that the loss functions are chosen in an adversarial manner by REALITY.
- However, the loss functions might have some patterns and could be predictable.
- LEARNER can have prior knowledge about the loss functions (see Rakhlin and Sridharan (2013)).

OFTPL with Approximate Optimization Oracle

Algorithm 3 (OFTPL, Suggala and Netrapalli (2020))

Let g_t be LEARNER's guess of f_t depending on $f_1, \cdots, f_{t-1}$. Given a parameter $\eta > 0$, LEARNER draws an i.i.d. random vector $\sigma_t \sim (\exp(\eta))^d$ and plays a strategy

$$x_t = \texttt{Approx-Opt}_{\alpha,\beta}\left(\sum_{i=1}^{t-1} f_i + g_t - \langle \sigma_t, \cdot \rangle\right)$$

at t-th round.

Rakhlin and Sridharan (2013) give a thorough discussion on the choices of g_t.

Performance Analysis

- Suggala and Netrapalli (2020) give the regret bound

$$\mathbb{E}[R_T(x)] \leq \mathcal{O}\left(\eta d^2 D \sum_{t=1}^{T} \frac{L_t^2}{T} + \frac{d(\beta T + D)}{\eta T} + \alpha + \beta d \sum_{t=1}^{T} \frac{L_t}{T}\right),$$

  where they assume $g_t - f_t$ is L_t-Lipschitz w.r.t. ℓ_1 norm.
- The regret bound is expected to be smaller when the guess g_t is close to f_t.

Summary

Takeaway:

- Non-convex problems are common and hard to solve, but a slightly modified FTL can achieve a sublinear regret bound.
- The Exponential Weight Algorithm is another approach to solve non-convex problems.

Comment:

- Can other distributions improve the regret bound?

Reference

Naman Agarwal, Alon Gonen, and Elad Hazan: **Learning in non-convex games with an optimization oracle**, COLT, 2019.

Arun Sai Suggala and Praneeth Netrapalli: **Online Non-Convex Learning: Following the Perturbed Leader is Optimal**, ALT, 2020.

Alexander Rakhlin and Karthik Sridharan: **Online learning with predictable sequences**, JMLR, 2012.

Marina Danilova, Pavel Dvurechensky, Alexander Gasnikov, Eduard Gorbunov, Sergey Guminov, Dmitry Kamzolov, and Innokentiy Shibaev: **Recent Theoretical Advances in Non-Convex Optimization**, arxiv, 2021.

Reference

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus: **Intriguing properties of neural networks**, ICLR, 2013.

Adam Kalai and Santosh Vempala: **Efficient algorithms for on-line decision problems**, J. of Comput. and Sys. Sci., 2005.

Nick Littlestone: **From on-line to batch learning**, COLT, 1989.

Nicolo Cesa-Bianchi and Gabor Lugosi: **Prediction, learning, and games**, Cambridge university press, 2006.

Yoel Drori and Ohad Shamir: **The complexity of finding stationary points with stochastic gradient descent**, ICML, 2020.