Online Non-Convex Learning: Following the Perturbed Leader is Optimal

Arun Suggala
Carnegie Mellon University

Praneeth Netrapalli
Microsoft Research, India

Motivation

Many problems in ML, statistics involve non-convex non-concave games

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} F(\boldsymbol{x}, \boldsymbol{y})$$

Motivation

Many problems in ML, statistics involve non-convex non-concave games

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} F(\boldsymbol{x}, \boldsymbol{y})$$

- Generative Adversarial Networks

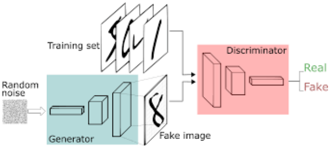

Motivation

Many problems in ML, statistics involve non-convex non-concave games

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} F(\boldsymbol{x}, \boldsymbol{y})$$

- Generative Adversarial Networks
- Robust optimization

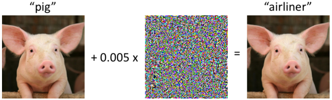

Motivation

Many problems in ML, statistics involve non-convex non-concave games

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} F(\boldsymbol{x}, \boldsymbol{y})$$

- Generative Adversarial Networks
- Robust optimization
- Minimax Estimators

Outline

- Introduction
- Background
- Main Result
- FTPL
- Optimistic FPTL
- Min-Max Games

Setup

- Time: $1, 2, \ldots T$
- At time t, learner predicts $\boldsymbol{x}_t \in \mathcal{X}$
- Adversary simultaneously reveals loss function f_t
- Goal: minimize cumulative loss $\sum_{t=1}^{T} f_t(\boldsymbol{x}_t)$

Setup

- At time t, predict $\boldsymbol{x}_t$ and observer loss function f_t
- Goal: minimize cumulative loss $\sum_{t=1}^{T} f_t(\boldsymbol{x}_t)$.
- Benchmark: $\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{t=1}^{T} f_t(\boldsymbol{x})$ - best fixed policy in hindsight.
- Regret: $\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \min_{\boldsymbol{x} \in \mathcal{X}} \sum_{t=1}^{T} f_t(\boldsymbol{x})$

Minimize Regret

History

- Online linear learning: dates back to [Brown and Von Neumann, 1950]
- Online convex learning: heavily studied since [Zinkevich, 2003]
- **Regret**

$$\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \min_{\boldsymbol{x}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\boldsymbol{x}) = O(\sqrt{T})$$

Online Non-Convex Learning

- Computationally intractable even if all $f_t(\cdot)$ are the same

What can we do?

1. Weaker notions of regret (such as stationarity in optimization) [Hazan et al., 2017]
2. Assume access to offline optimization oracles (only deal with learning) [Agarwal et al., 2018]

Main Result

Theorem

Suppose:

- $f_t(\cdot)$ *is Lipschitz continuous*
- $\boldsymbol{x}_t \in \mathcal{X}$ *with bounded diameter*
- *we have access to offline optimization oracle*

There exists a randomized algorithm such that

$$\mathbb{E}\left[\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \min_{\boldsymbol{x}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\boldsymbol{x})\right] = O(\sqrt{T}).$$

- Previous best: $O(T^{2/3})$ [Agarwal et al., 2018]

Algorithm: Follow The Perturbed Leader

Algorithm

- $\sigma_t \sim \mathrm{Unif}(0, \sqrt{T})$
- $\boldsymbol{x}_t \stackrel{def}{=} \arg\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) - \langle \sigma_t, \boldsymbol{x} \rangle$

- Studied by [Hannan, 1957, Kalai and Vempala, 2016] for linear losses
  - Regret $= O(\sqrt{T})$

Main Intuitions

Step 1

Reduction to oblivious adversary [Cesa-Bianchi and Lugosi, 2006]

- assume adversary fixes choices ahead of time
- suffices to work with a single random vector σ

Main Intuitions

Step 1

Reduction to oblivious adversary [Cesa-Bianchi and Lugosi, 2006]

- assume adversary fixes choices ahead of time
- suffices to work with a single random vector σ

Step 2

Be the perturbed leader lemma

- Recall, $\boldsymbol{x}_t \overset{def}{=} \arg\min_{\boldsymbol{x}\in\mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) - \langle \sigma, \boldsymbol{x} \rangle$
- $\mathbb{E}\left[\sum_{t=1}^{T} f_t(\boldsymbol{x}_{t+1}) - \min_{\boldsymbol{x}\in\mathcal{X}} f_t(\boldsymbol{x})\right] = O(\sqrt{T})$, since $\sigma \leq \sqrt{T}$

Main Intuitions (contd.)

Step 3

Stability

- Recall: $\boldsymbol{x}_t \stackrel{def}{=} \arg\min_{\boldsymbol{x}\in\mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) - \langle \sigma, \boldsymbol{x} \rangle$
- $\mathbb{E}\left[\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \sum_{t=1}^{T} f_t(\boldsymbol{x}_{t+1})\right] \leq L \sum_{t=1}^{T} \mathbb{E}\left[\|\boldsymbol{x}_t - \boldsymbol{x}_{t+1}\|\right]$

Stability Question

- Recall: $\boldsymbol{x}_t \stackrel{def}{=} \arg\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) - \langle \sigma, \boldsymbol{x} \rangle$
- How large can $\mathbb{E}\left[\|\boldsymbol{x}_t - \boldsymbol{x}_{t+1}\|\right]$ be?
- [Agarwal et al., 2018]: $\mathbb{E}\left[\|\boldsymbol{x}_t - \boldsymbol{x}_{t+1}\|\right] = O(T^{-1/3})$

Our Improvement

$$\mathbb{E}\left[\|\boldsymbol{x}_t - \boldsymbol{x}_{t+1}\|\right] = O(T^{-1/2})$$

Weak Monotonicity Property

- $\boldsymbol{x}_t(\sigma) \stackrel{def}{=} \arg\min_{\boldsymbol{x}\in\mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) - \langle \sigma, \boldsymbol{x} \rangle$

Weak Monotonicity Property

- $\boldsymbol{x}_{t,i}(\sigma + ce_i) \geq \boldsymbol{x}_{t,i}(\sigma)$ for all $\sigma, c \geq 0$

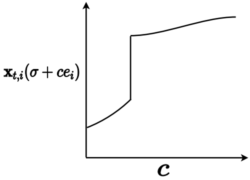

Strong Monotonicity Property

- $\boldsymbol{x}_t(\sigma) \stackrel{def}{=} \arg\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) - \langle \sigma, \boldsymbol{x} \rangle$

Strong Monotonicity Property (1D)

Forall $c \geq L$

- $\max\{\boldsymbol{x}_t(\sigma), \boldsymbol{x}_{t+1}(\sigma)\} \leq \min\{\boldsymbol{x}_t(\sigma + c), \boldsymbol{x}_{t+1}(\sigma + c)\}$

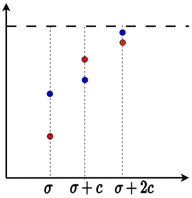

Strong Monotonicity Property

- Recall: $\boldsymbol{x}_t(\sigma) \stackrel{def}{=} \arg\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) - \langle \sigma, \boldsymbol{x} \rangle$

Strong Monotonicity Property (High Dim.)

- Suppose $\|\boldsymbol{x}_t(\sigma) - \boldsymbol{x}_{t+1}(\sigma)\|_1 \leq 10d \cdot |\boldsymbol{x}_{t,i}(\sigma) - \boldsymbol{x}_{t+1,i}(\sigma)|$
- Then for $\sigma' = \sigma + 100Lde_i$

$$\max\left(\boldsymbol{x}_{t,i}(\sigma), \boldsymbol{x}_{t+1,i}(\sigma)\right) \leq \min\left(\boldsymbol{x}_{t,i}(\sigma'), \boldsymbol{x}_{t+1,i}(\sigma')\right) + \frac{1}{10}|\boldsymbol{x}_{t,i}(\sigma) - \boldsymbol{x}_{t+1,i}(\sigma)|$$

Strong Monotonicity Property

- Recall: $\boldsymbol{x}_t(\sigma) \stackrel{def}{=} \arg\min_{\boldsymbol{x}\in\mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) - \langle \sigma, \boldsymbol{x} \rangle$

Strong Monotonicity Property (High Dim.)

- Suppose $\|\boldsymbol{x}_t(\sigma) - \boldsymbol{x}_{t+1}(\sigma)\|_1 \leq 10d \cdot |\boldsymbol{x}_{t,i}(\sigma) - \boldsymbol{x}_{t+1,i}(\sigma)|$
- Then for $\sigma' = \sigma + 100Lde_i$

$$\max\left(\boldsymbol{x}_{t,i}(\sigma), \boldsymbol{x}_{t+1,i}(\sigma)\right) \leq \min\left(\boldsymbol{x}_{t,i}(\sigma'), \boldsymbol{x}_{t+1,i}(\sigma')\right) + \frac{1}{10}|\boldsymbol{x}_{t,i}(\sigma) - \boldsymbol{x}_{t+1,i}(\sigma)|$$

- The two monotonicity properties give us $O(T^{-1/2})$ stability bound

Recap

- Follow the perturbed leader
- Be the leader lemma: playing $\boldsymbol{x}_{t+1}$ at time t is very good
- Stability: With perturbations, $\|\boldsymbol{x}_t - \boldsymbol{x}_{t+1}\|$ is very small
- Key technical results: Tight monotonicity lemmas

Upshot

FTPL with access to offline optimization oracle achieves $O(\sqrt{T})$ regret

Optimistic FTPL

- Better regret bounds when sequence of losses are predictable
- $g_t[f_1, \ldots f_{t-1}]$: guess for f_t at iteration t

Algorithm

- $\sigma_t \sim \mathrm{Unif}(0, \sqrt{T})$
- $\boldsymbol{x}_t \stackrel{def}{=} \arg\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{i=1}^{t-1} f_i(\boldsymbol{x}) + g_t(\boldsymbol{x}) - \langle \sigma_t, \boldsymbol{x} \rangle$

Optimistic FTPL

- **Regret**

$$\sum_{t=1}^{T} f_t(\boldsymbol{x}_t) - \min_{\boldsymbol{x}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\boldsymbol{x}) = O(L_t\sqrt{T})$$

L_t is Lipschitz constant of $(g_t - f_t)$

Non-Convex Non-Concave Games

- Two player zero-sum game

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} F(\boldsymbol{x}, \boldsymbol{y})$$

$F(\cdot, \boldsymbol{y})$ non-convex in $\boldsymbol{x}$, $F(\boldsymbol{x}, \cdot)$ non-concave in $\boldsymbol{y}$

Non-Convex Non-Concave Games

- Two player zero-sum game

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} F(\boldsymbol{x}, \boldsymbol{y})$$

  $F(\cdot, \boldsymbol{y})$ non-convex in $\boldsymbol{x}$, $F(\boldsymbol{x}, \cdot)$ non-concave in $\boldsymbol{y}$
- Both players use online learning algorithms against each other

Non-Convex Non-Concave Games

- Two player zero-sum game

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} F(\boldsymbol{x}, \boldsymbol{y})$$

$F(\cdot, \boldsymbol{y})$ non-convex in $\boldsymbol{x}$, $F(\boldsymbol{x}, \cdot)$ non-concave in $\boldsymbol{y}$

- Both players use online learning algorithms against each other

	Lipschitz + Smooth
FTPL	$O(T^{-1/2})$
Optimistic FTPL ($g_t = f_{t-1}$)	$O(T^{-3/4})$

Table: Rate of convergence to equilibirum

Conclusion

- FTPL achieves optimal $O(\sqrt{T})$ regret for online non-convex learning
- Optimistic FTPL can provide better rates if losses are non-adversarial

Conclusion

- FTPL achieves optimal $O(\sqrt{T})$ regret for online non-convex learning
- Optimistic FTPL can provide better rates if losses are non-adversarial

Questions?

References I

[Agarwal et al., 2018] Agarwal, N., Gonen, A., and Hazan, E. (2018).
Learning in non-convex games with an optimization oracle.
arXiv preprint arXiv:1810.07362.

[Brown and Von Neumann, 1950] Brown, G. W. and Von Neumann, J. (1950).
Solutions of games by differential equations.
Technical report, RAND CORP SANTA MONICA CA.

[Cesa-Bianchi and Lugosi, 2006] Cesa-Bianchi, N. and Lugosi, G. (2006).
Prediction, learning, and games.
Cambridge university press.

[Hannan, 1957] Hannan, J. (1957).
Approximation to bayes risk in repeated play.
Contributions to the Theory of Games, 3:97–139.

[Hazan et al., 2017] Hazan, E., Singh, K., and Zhang, C. (2017).
Efficient regret minimization in non-convex games.
arXiv preprint arXiv:1708.00075.

[Kalai and Vempala, 2016] Kalai, A. T. and Vempala, S. (2016).
Efficient algorithms for on-line optimization.
Journal of Computer and System Sciences, 71.

References II

[Zinkevich, 2003] Zinkevich, M. (2003).
Online convex programming and generalized infinitesimal gradient ascent.
In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 928–936.