

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

Arun Sai Suggala

*Machine Learning Department
Carnegie Mellon University*

ASUGGALA@CS.CMU.EDU

Praneeth Netrapalli

Microsoft Research, India

PRANEETH@MICROSOFT.COM

Editors: Aryeh Kontorovich and Gergely Neu

Abstract

We study the problem of online learning with non-convex losses, where the learner has access to an offline optimization oracle. We show that the classical Follow the Perturbed Leader (FTPL) algorithm achieves optimal regret rate of $O(T^{-1/2})$ in this setting. This improves upon the previous best-known regret rate of $O(T^{-1/3})$ for FTPL. We further show that an optimistic variant of FTPL achieves better regret bounds when the sequence of losses encountered by the learner is “predictable”.

Keywords: Online Learning, Non-Convex Losses, Perturbation

1. Introduction

In this work, we study the problem of online learning with non-convex losses, where, in each iteration, the learner chooses an action and observes a loss which could potentially be non-convex. The goal of the learner is to choose a sequence of actions which minimize the cumulative loss suffered over the course of learning. The paradigm of online learning has been studied in a number of fields, including game theory, machine learning, statistics and has several practical applications. In recent years a number of efficient algorithms have been developed for online learning. Convexity of the loss functions has played a central role in the development of many of these techniques. In this work, we consider a more general setting, where the sequence of loss functions encountered by the learner could be non-convex. Such a setting has numerous applications in machine learning, especially in adversarial training (Szegedy et al., 2013), robust optimization and training of Generative Adversarial Networks (GANs) (Goodfellow et al., 2014).

As mentioned above, most of the existing works on online optimization have focused on convex loss functions (Hazan, 2016). A number of computationally efficient approaches have been proposed for regret minimization in this setting. However, when the losses are non-convex, minimizing the regret is computationally hard. Recent works on learning with non-convex losses get over this computational barrier by either working with a restricted class of loss functions such as approximately convex losses (Gao et al., 2018) or by optimizing a computationally tractable notion of regret (Hazan et al., 2017). Consequently, the techniques studied in these papers do not guarantee vanishing regret for general non-convex losses. Another class of approaches consider general non-convex losses, but assume access to a sampling oracle (Maillard and Munos, 2010; Krichene et al., 2015) or an offline optimization oracle (Agarwal et al., 2019). Of these, assuming access to an

offline optimization oracle is reasonable, given that in practice, simple heuristics such as stochastic gradient descent seem to be able to find approximate global optima reasonably fast even for complicated tasks such as training deep neural networks.

In a recent work Agarwal et al. (2019) take this later approach, where they assume access to an offline optimization oracle, and show that the classical Follow the Perturbed Leader (FTPL) algorithm achieves $O(T^{-1/3})$ regret for general non-convex losses which are Lipschitz continuous. In this work, we improve upon this result and show that FTPL in fact achieves optimal $O(T^{-1/2})$ regret.

2. Problem Setup and Main Results

Let $\mathcal{X} \subseteq \mathbb{R}^d$ denote the set of all possible moves of the learner. In the online learning framework, on each round t , the learner makes a prediction $\mathbf{x}_t \in \mathcal{X}$ and the nature/adversary simultaneously chooses a loss function $f_t: \mathcal{X} \rightarrow \mathbb{R}$ and observe each others actions. The goal of the learner is to choose a sequence of actions $\{\mathbf{x}_t\}_{t=1}^T$ such that the following notion of regret is small

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

In this work we assume that \mathcal{X} is bounded and has ℓ_∞ diameter of D , which is defined as

$$D = \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|_\infty.$$

Moreover, we assume that the sequence of loss functions f_t chosen by the adversary are L -Lipschitz with respect to ℓ_1 norm, that is, for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, $|f_t(\mathbf{x}) - f_t(\mathbf{y})| \leq L\|\mathbf{x} - \mathbf{y}\|_1$.

Approximate Optimization Oracle. Our results rely on an offline optimization oracle which takes as input a function $f: \mathcal{X} \rightarrow \mathbb{R}$ and a d -dimensional vector σ and returns an approximate minimizer of $\mathbf{x} \mapsto f(\mathbf{x}) - \langle \sigma, \mathbf{x} \rangle$. An optimization oracle is called “ (α, β) -approximate optimization oracle” if it returns $\mathbf{x}^* \in \mathcal{X}$ such that

$$f(\mathbf{x}^*) - \langle \sigma, \mathbf{x}^* \rangle \leq \inf_{\mathbf{x} \in \mathcal{X}} [f(\mathbf{x}) - \langle \sigma, \mathbf{x} \rangle] + (\alpha + \beta\|\sigma\|_1),$$

We denote such an optimization oracle with $\mathcal{O}_{\alpha, \beta}(f - \langle \sigma, \cdot \rangle)$.

FTPL. Given access to an (α, β) -approximate offline optimization oracle, we study the FTPL algorithm which is described by the following prediction rule (see Algorithm 1).

$$\mathbf{x}_t = \mathcal{O}_{\alpha, \beta} \left(\sum_{i=1}^{t-1} f_i - \langle \sigma_t, \cdot \rangle \right), \tag{1}$$

where $\sigma_t \in \mathbb{R}^d$ is a random perturbation such that $\sigma_{t,j}$, the j^{th} coordiante of σ_t , is sampled from $\text{Exp}(\eta)$, the exponential distribution with parameter η ¹. We note that one can also generate the perturbations from other probability distributions such as uniform distribution and achieve similar regret bounds as presented in this paper.

1. Recall, Z is an exponential random variable with parameter η if $P(Z \geq s) = \exp(-\eta s)$

Algorithm 1 Follow the Perturbed Leader (FTPL)

- 1: **Input:** Parameter of exponential distribution η , approximate optimization oracle $\mathcal{O}_{\alpha,\beta}$
- 2: **for** $t = 1 \dots T$ **do**
- 3: Generate random vector σ_t such that $\{\sigma_{t,j}\}_{j=1}^d \stackrel{i.i.d.}{\sim} \text{Exp}(\eta)$
- 4: Predict \mathbf{x}_t as

$$\mathbf{x}_t = \mathcal{O}_{\alpha,\beta} \left(\sum_{i=1}^{t-1} f_i - \langle \sigma_t, \cdot \rangle \right).$$

- 5: Observe loss function f_t
 - 6: **end for**
-

Optimistic FTPL (OFTPL). In the general online learning setting considered above, we assumed that the loss functions could possibly be chosen in an adversarial manner by nature. However, in certain applications, the loss functions may not be adversarial. Instead, they might have some patterns and could be predictable. In such cases, [Rakhlin and Sridharan \(2012\)](#) present algorithms for online linear optimization which can exploit the predictability of losses to obtain better regret bounds. We show that the techniques of [Rakhlin and Sridharan \(2012\)](#) can be extended to the online non-convex optimization setting considered in this work.

Let $g_t[f_1 \dots f_{t-1}]$ be our guess of the loss f_t at the beginning of round t , with $g_1 = 0$. To simplify the notation, in the sequel, we suppress the dependence of g_t on $\{f_i\}_{i=1}^{t-1}$. Some potential choices for g_t that could be of interest are f_{t-1} , $\frac{1}{t-1} \sum_{i=1}^{t-1} f_i$. For a thorough discussion on the choices of g_t and concrete examples where predictable loss functions arise, we refer the reader to [Rakhlin and Sridharan \(2012, 2013\)](#). Given g_t , we predict \mathbf{x}_t in OFTPL as

$$\mathbf{x}_t = \mathcal{O}_{\alpha,\beta} \left(\sum_{i=1}^{t-1} f_i + g_t - \langle \sigma_t, \cdot \rangle \right) \tag{2}$$

When our guess g_t is close to f_t we expect OFTPL to have a smaller regret. In [Theorem 2](#) we show that the regret of OFTPL depends only on $(f_t - g_t)$.

2.1. Main Results

We present our main results for an oblivious adversary who fixes the sequence of losses $\{f_t\}_{t=1}^T$ ahead of the game. Following [Cesa-Bianchi and Lugosi \(2006\)](#), one can show that any algorithm that is guaranteed to work against an oblivious adversary also works for a non-oblivious adversary, whose actions are allowed to depend on the past predictions of the algorithm. For the sake of completeness, we present a proof of this reduction from non-oblivious to oblivious adversary model in [Appendix B](#).

Theorem 1 (Non-Convex FTPL) *Let D be the ℓ_∞ diameter of \mathcal{X} . Suppose the losses encountered by the learner are L -Lipschitz w.r.t ℓ_1 norm. Moreover, suppose the optimization oracle used by [Algorithm 1](#) is a “ (α, β) -approximate” optimization oracle. For any fixed η , the predictions of [Algorithm 1](#) satisfy the following regret bound*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T f_t(\mathbf{x}_t) - \frac{1}{T} \inf_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) \right] \leq O \left(\eta d^2 D L^2 + \frac{d(\beta T + D)}{\eta T} + \alpha + \beta d L \right).$$

Theorem 2 (Non-Convex OFTPL) *Let D be the ℓ_∞ diameter of \mathcal{X} . Suppose our guess g_t is such that $(g_t - f_t)$ is L_t -Lipschitz w.r.t ℓ_1 norm, for all $t \in [T]$. For any fixed η , OFTPL with access to a “ (α, β) -approximate” optimization oracle satisfies the following regret bound*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T f_t(\mathbf{x}_t) - \frac{1}{T} \inf_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) \right] \leq O \left(\eta d^2 D \sum_{t=1}^T \frac{L_t^2}{T} + \frac{d(\beta T + D)}{\eta T} + \alpha + \beta d \sum_{t=1}^T \frac{L_t}{T} \right).$$

The above result shows that for appropriate choice of η , FTPL achieves $O(d^{\frac{3}{2}}T^{-\frac{1}{2}} + \alpha + \beta d^{\frac{3}{2}}T^{\frac{1}{2}})$ regret. This also shows that when $\alpha = O(T^{-\frac{1}{2}}), \beta = O(T^{-1})$, FTPL achieves the optimal $O(T^{-\frac{1}{2}})$ regret. This improves upon the $O(T^{-\frac{1}{3}})$ regret bound obtained by Agarwal et al. (2019). We note that the above results can be generalized to infinite-dimensional spaces such as ℓ^1 space of sequences. To do this we assume that the domain \mathcal{X} is bounded and can be enclosed in a hyper-rectangle with edge length D_i along the i^{th} standard basis vector. Through a more careful analysis we can obtain regret bounds that depend on the *effective dimension* of \mathcal{X} , which is defined as $\frac{\sum_{i=1}^d D_i}{\max_i D_i}$, instead of d .

Before we conclude the section we point out that as an immediate consequence of the above regret bounds, we obtain algorithms for approximating the mixed strategy Nash equilibria of general non-convex non-concave saddle point problems of the form $\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{y} \in \mathcal{Y}} F(\mathbf{x}, \mathbf{y})$. This follows from the observation that saddle point problems can be solved by playing two online optimization algorithms against each other (Cesa-Bianchi and Lugosi, 2006; Hazan, 2016).

3. Background

In this section we briefly review the relevant literature on online learning in both convex and non-convex settings.

Online Convex Optimization. When the domain \mathcal{X} and the loss functions f_t encountered by the learner are convex, a number of efficient algorithms for regret minimization have been studied. Most of these algorithms fall into three broad categories, namely Follow the Regularized Leader (FTRL), Online Mirror Descent (OMD) (Hazan, 2016) and Follow the Perturbed Leader (FTPL) (Kalai and Vempala, 2016). FTRL algorithms make a prediction in each iteration by minimizing $\arg\min_{\mathbf{x}} \sum_{i=1}^{t-1} f_i(\mathbf{x}) + R(\mathbf{x})$, where R is a strongly convex regularizer. The regularization R plays a crucial role in the performance of the algorithm and helps avoid overfitting to the observed loss functions. Similar to FTRL, OMD also relies on explicit regularization to guarantee vanishing regret. In fact, under certain settings, both OMD and FTRL algorithms are known to be equivalent (McMahan, 2011). For a broad class of online convex optimization problems, FTRL and OMD are known to achieve optimal regret guarantees.

FTPL algorithms rely on random perturbation of loss functions to guarantee vanishing regret. This random perturbation can be viewed as having a similar role as the explicit regularization used in FTRL and OMD. In a recent work Abernethy et al. (2016) use duality to connect FTPL and FTRL. They show that every instance of FTPL is also an instance of FTRL.

Online Non-Convex Optimization. A natural question that arises in the context of online non-convex learning is whether there exist counterparts of FTRL and OMD which achieve vanishing regret. Unfortunately, the answer is no. As we show in the following Proposition, there exists no deterministic algorithm that can achieve vanishing regret when the losses are non-convex.

Proposition 3 *No deterministic algorithm can achieve $o(1)$ regret in the setting of online non-convex learning.*

The above Proposition shows that only randomized algorithms can achieve vanishing regret. Recent works of [Maillard and Munos \(2010\)](#); [Krichene et al. \(2015\)](#) consider the natural extension of Exponential Weight Algorithm to continuous domains and show that the resulting algorithm has vanishing regret in the setting of online non-convex learning. The algorithms studied in these works rely on an offline sampling oracle which can generate samples from any given probability distribution. In another line of work, [Agarwal et al. \(2019\)](#) study the classical FTPL algorithm with access to a certain offline optimization oracle and show that it achieves $O(T^{-1/3})$ regret. As an immediate consequence of this result, the authors show that both online adversarial learning model and statistical learning model are computationally equivalent.

4. Non-Convex FTPL

In this section, we present a proof of Theorem 1. Since we are in the oblivious adversary setting, it suffices to work with a single random vector σ , instead of generating a new random vector in each iteration. The first step in the proof involves relating the expected regret to the stability of prediction, which is a standard step in the analysis of many online learning algorithms.

Lemma 4 *The regret of Algorithm 1 can be upper bounded as*

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T f_t(\mathbf{x}_t) - \inf_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) \right] \leq \frac{L}{T} \sum_{t=1}^T \underbrace{\mathbb{E} [\|\mathbf{x}_t - \mathbf{x}_{t+1}\|_1]}_{\text{Stability}} + \frac{d(\beta T + D)}{\eta T} + \alpha. \quad (3)$$

In the rest of the proof we focus on bounding the stability term $\mathbb{E} [\|\mathbf{x}_t - \mathbf{x}_{t+1}\|_1]$. The randomness used in the algorithm is crucial for bounding its stability. The more randomness we add, the more stable the algorithm is. However, there is a price we pay for adding randomness. It causes the algorithm to make poor predictions, which leads to worse regret. This is evident in the second term in the upper bound in Equation (3), which increases as η decreases.

We first provide an brief sketch of the proof in the 1-dimensional case. Similar to the proof of [Agarwal et al. \(2019\)](#), our proof relies on showing certain monotonicity properties of the predictions of the algorithm. Letting $\mathbf{x}_t(\sigma)$ be the prediction in the t^{th} iteration of FTPL with random perturbation σ , we show that the predictions are monotonic functions of σ

$$\forall t, c > 0, \quad \mathbf{x}_t(\sigma + c) \geq \mathbf{x}_t(\sigma).$$

Moreover, we show that

$$\forall c > L, \quad \min \{\mathbf{x}_t(\sigma + c), \mathbf{x}_{t+1}(\sigma + c)\} \geq \max \{\mathbf{x}_t(\sigma), \mathbf{x}_{t+1}(\sigma)\}.$$

Since the domain is bounded, these two properties imply that the functions $\mathbf{x}_t(\sigma)$, $\mathbf{x}_{t+1}(\sigma)$ should be close to each other for sufficiently large values of σ (see Figure 1 for an illustration). The closeness of these two functions immediately implies the stability of the algorithm. In what follows, we formalize this argument and extend it to the high-dimensional case.




Figure 1: Illustration of monotonicity properties of the predictions of FTPL on a 1-dimensional example with $D = 10, L = 2$.

Lemma 5 (Monotonicity 1) *Let $\mathbf{x}_t(\sigma)$ be the prediction of FTPL in iteration t , with random perturbation σ . Let \mathbf{e}_i denote the i^{th} standard basis vector and $\mathbf{x}_{t,i}$ denote the i^{th} coordinate of \mathbf{x}_t . Then the following monotonicity property holds for any $c > 0$*

$$\mathbf{x}_{t,i}(\sigma + c\mathbf{e}_i) \geq \mathbf{x}_{t,i}(\sigma) - \frac{2(\alpha + \beta\|\sigma\|_1)}{c} - \beta.$$

Proof Let $f_{1:t}(\mathbf{x}) = \sum_{i=1}^t f_i(\mathbf{x})$ and $\sigma' = \sigma + c\mathbf{e}_i$. Moreover, let $\gamma(\sigma) = \alpha + \beta\|\sigma\|_1$ be the approximation error of the offline optimization oracle. From the approximate optimality of $\mathbf{x}_t(\sigma)$ we have

$$\begin{aligned} & f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma, \mathbf{x}_t(\sigma) \rangle \\ & \leq f_{1:t-1}(\mathbf{x}_t(\sigma')) - \langle \sigma, \mathbf{x}_t(\sigma') \rangle + \gamma(\sigma) \\ & = f_{1:t-1}(\mathbf{x}_t(\sigma')) - \langle \sigma', \mathbf{x}_t(\sigma') \rangle + c\mathbf{x}_{t,i}(\sigma') + \gamma(\sigma) \\ & \stackrel{(a)}{\leq} f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma', \mathbf{x}_t(\sigma) \rangle + c\mathbf{x}_{t,i}(\sigma') + \gamma(\sigma) + \gamma(\sigma') \\ & = f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma, \mathbf{x}_t(\sigma) \rangle + c(\mathbf{x}_{t,i}(\sigma') - \mathbf{x}_{t,i}(\sigma)) + \gamma(\sigma) + \gamma(\sigma'), \end{aligned}$$

where (a) follows from the approximate optimality of $\mathbf{x}_t(\sigma')$. Combining the first and last terms in the above expression, we get $\mathbf{x}_{t,i}(\sigma') \geq \mathbf{x}_{t,i}(\sigma) - \frac{2\gamma(\sigma)}{c} - \beta$. \blacksquare

Lemma 6 (Monotonicity 2) *Let $\mathbf{x}_t(\sigma)$ be the prediction of FTPL in iteration t , with random perturbation σ . Let \mathbf{e}_i denote the i^{th} standard basis vector and $\mathbf{x}_{t,i}$ denote the i^{th} coordinate of \mathbf{x}_t . Suppose $\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1 \leq 10d \cdot |\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)|$. For $\sigma' = \sigma + 100Ld\mathbf{e}_i$, we have*

$$\begin{aligned} \min(\mathbf{x}_{t,i}(\sigma'), \mathbf{x}_{t+1,i}(\sigma')) & \geq \max(\mathbf{x}_{t,i}(\sigma), \mathbf{x}_{t+1,i}(\sigma)) - \frac{1}{10}|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| \\ & \quad - \frac{3(\alpha + \beta\|\sigma\|_1)}{100Ld} - \beta. \end{aligned}$$

Proof Let $f_{1:t}(\mathbf{x}) = \sum_{i=1}^t f_i(\mathbf{x})$ and let $\gamma(\sigma) = \alpha + \beta\|\sigma\|_1$ be the approximation error of the offline optimization oracle. From the approximate optimality of $\mathbf{x}_t(\sigma)$, we have

$$\begin{aligned} & f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma, \mathbf{x}_t(\sigma) \rangle + f_t(\mathbf{x}_t(\sigma)) \\ & \leq f_{1:t-1}(\mathbf{x}_{t+1}(\sigma)) - \langle \sigma, \mathbf{x}_{t+1}(\sigma) \rangle + f_t(\mathbf{x}_t(\sigma)) + \gamma(\sigma) \\ & \stackrel{(a)}{\leq} f_{1:t-1}(\mathbf{x}_{t+1}(\sigma)) - \langle \sigma, \mathbf{x}_{t+1}(\sigma) \rangle + f_t(\mathbf{x}_{t+1}(\sigma)) + L\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1 + \gamma(\sigma) \\ & \stackrel{(b)}{\leq} f_{1:t-1}(\mathbf{x}_{t+1}(\sigma)) - \langle \sigma, \mathbf{x}_{t+1}(\sigma) \rangle + f_t(\mathbf{x}_{t+1}(\sigma)) + 10Ld|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| + \gamma(\sigma), \end{aligned}$$

where (a) follows from the Lipschitz property of $f_t(\cdot)$ and (b) follows from our assumption on $\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1$. Next, from the optimality of $\mathbf{x}_{t+1}(\sigma')$, we have

$$\begin{aligned} & f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma, \mathbf{x}_t(\sigma) \rangle + f_t(\mathbf{x}_t(\sigma)) \\ & = f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma', \mathbf{x}_t(\sigma) \rangle + f_t(\mathbf{x}_t(\sigma)) + \langle 100Ld\mathbf{e}_i, \mathbf{x}_t(\sigma) \rangle \\ & \geq f_{1:t-1}(\mathbf{x}_{t+1}(\sigma')) - \langle \sigma', \mathbf{x}_{t+1}(\sigma') \rangle + f_t(\mathbf{x}_{t+1}(\sigma')) + 100Ld\mathbf{x}_{t,i}(\sigma) - \gamma(\sigma') \\ & = f_{1:t-1}(\mathbf{x}_{t+1}(\sigma')) - \langle \sigma, \mathbf{x}_{t+1}(\sigma') \rangle + f_t(\mathbf{x}_{t+1}(\sigma')) + 100Ld(\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma')) - \gamma(\sigma') \\ & \geq f_{1:t-1}(\mathbf{x}_{t+1}(\sigma)) - \langle \sigma, \mathbf{x}_{t+1}(\sigma) \rangle + f_t(\mathbf{x}_{t+1}(\sigma)) + 100Ld(\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma')) - \gamma(\sigma') - \gamma(\sigma), \end{aligned}$$

where the last inequality follows from the optimality of $\mathbf{x}_{t+1}(\sigma)$. Combining the above two equations, we get

$$\mathbf{x}_{t+1,i}(\sigma') - \mathbf{x}_{t,i}(\sigma) \geq -\frac{1}{10}|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| - \frac{3\gamma(\sigma)}{100Ld} - \beta.$$

A similar argument shows that

$$\mathbf{x}_{t,i}(\sigma') - \mathbf{x}_{t+1,i}(\sigma) \geq -\frac{1}{10}|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| - \frac{3\gamma(\sigma)}{100Ld} - \beta.$$

Finally, from the monotonicity property in Lemma 5 we know that

$$\mathbf{x}_{t+1,i}(\sigma') - \mathbf{x}_{t+1,i}(\sigma) \geq -\frac{3\gamma(\sigma)}{100Ld} - \beta, \quad \mathbf{x}_{t,i}(\sigma') - \mathbf{x}_{t,i}(\sigma) \geq -\frac{3\gamma(\sigma)}{100Ld} - \beta.$$

Combining the above four inequalities gives us the required result. \blacksquare

Proof of Theorem 1. We now proceed to the proof of Theorem 1. We use the same notation as in Lemmas 5, 6. First note that $\mathbb{E}[\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1]$ can be written as

$$\mathbb{E}[\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1] = \sum_{i=1}^d \mathbb{E}[|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)|]. \quad (4)$$

To bound $\mathbb{E}[\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1]$ we derive an upper bound for $\mathbb{E}[|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)|]$, $\forall i \in [d]$. For any $i \in [d]$, define $\mathbb{E}_{-i}[|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)|]$ as

$$\mathbb{E}_{-i}[|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)|] := \mathbb{E}\left[|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| \middle| \{\sigma_j\}_{j \neq i}\right],$$

where σ_j is the j^{th} coordinate of σ . Let $\mathbf{x}_{\max,i}(\sigma) = \max(\mathbf{x}_{t,i}(\sigma), \mathbf{x}_{t+1,i}(\sigma))$ and $\mathbf{x}_{\min,i}(\sigma) = \min(\mathbf{x}_{t,i}(\sigma), \mathbf{x}_{t+1,i}(\sigma))$. Then $\mathbb{E}_{-i}[\|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)\|] = \mathbb{E}_{-i}[\mathbf{x}_{\max,i}(\sigma)] - \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma)]$. Define event \mathcal{E} as

$$\mathcal{E} = \{\sigma : \|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1 \leq 10d \cdot |\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)|\}.$$

Consider the following

$$\begin{aligned} \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma)] &= \mathbb{P}(\sigma_i < 100Ld) \underbrace{\mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma) | \sigma_i < 100Ld]}_{T_1} \\ &\quad + \underbrace{\mathbb{P}(\sigma_i \geq 100Ld) \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma) | \sigma_i \geq 100Ld]}_{T_2}. \end{aligned} \tag{5}$$

We now try to lower bound T_1, T_2 in the above equation. Since the domain of i^{th} coordinate lies within some interval of length D and since T_1 and $\mathbb{E}_{-i}[\mathbf{x}_{\max,i}(\sigma)]$ are points in this interval, their difference is bounded by D . So T_1 is lower bounded by $\mathbb{E}_{-i}[\mathbf{x}_{\max,i}(\sigma)] - D$. We next rewrite T_2 as follows.

$$\begin{aligned} T_2 &= \mathbb{P}(\sigma_i \geq 100Ld) \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma) | \sigma_i \geq 100Ld] = \int_{\sigma_i=100Ld}^{\infty} \mathbf{x}_{\min,i}(\sigma) P(\sigma_i) d\sigma_i \\ &= \int_{\sigma_i=100Ld}^{\infty} \mathbf{x}_{\min,i}(\sigma) \eta e^{-\eta\sigma_i} d\sigma_i \end{aligned}$$

We now do a *change of variables* in the above integration. Let $\sigma_i = \sigma'_i + 100Ld$ and $\sigma' = [\sigma_1, \dots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \dots]$ be the vector obtained by replacing the i^{th} coordinate of σ with σ'_i . Rewriting the RHS in terms of σ'_i and σ' , we get

$$\begin{aligned} \int_{\sigma_i=100Ld}^{\infty} \mathbf{x}_{\min,i}(\sigma) \eta e^{-\eta\sigma_i} d\sigma_i &= \int_{\sigma'_i=0}^{\infty} \mathbf{x}_{\min,i}(\sigma' + 100Ld\mathbf{e}_i) \eta e^{-\eta(\sigma'_i + 100Ld)} d\sigma'_i \\ &= e^{-100\eta Ld} \int_{\sigma'_i=0}^{\infty} \mathbf{x}_{\min,i}(\sigma' + 100Ld\mathbf{e}_i) \eta e^{-\eta\sigma'_i} d\sigma'_i \\ &= e^{-100\eta Ld} \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma' + 100Ld\mathbf{e}_i)]. \end{aligned}$$

This shows that $T_2 = e^{-100\eta Ld} \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma + 100Ld\mathbf{e}_i)]$. Substituting the lower bounds for T_1, T_2 in Equation (5), we get

$$\begin{aligned} \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma)] &\geq (1 - \exp(-100\eta Ld)) (\mathbb{E}_{-i}[\mathbf{x}_{\max,i}(\sigma)] - D) \\ &\quad + \exp(-100\eta Ld) \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma + 100Ld\mathbf{e}_i)], \end{aligned}$$

We can further lower bound $\mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma)]$ as follows

$$\begin{aligned} \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma)] &\geq (1 - \exp(-100\eta Ld)) (\mathbb{E}_{-i}[\mathbf{x}_{\max,i}(\sigma)] - D) \\ &\quad + \exp(-100\eta Ld) \mathbb{P}_{-i}(\mathcal{E}) \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma + 100Ld\mathbf{e}_i) | \mathcal{E}] \\ &\quad + \exp(-100\eta Ld) \mathbb{P}_{-i}(\mathcal{E}^c) \mathbb{E}_{-i}[\mathbf{x}_{\min,i}(\sigma + 100Ld\mathbf{e}_i) | \mathcal{E}^c], \end{aligned}$$

where $\mathbb{P}_{-i}(\mathcal{E})$ is defined as $\mathbb{P}_{-i}(\mathcal{E}) := \mathbb{P}(\mathcal{E} \mid \{\sigma_j\}_{j \neq i})$. We now use the monotonicity properties proved in Lemmas 5, 6 to further lower bound $\mathbb{E}_{-i}[\mathbf{x}_{min,i}(\sigma)]$. Let $\gamma(\sigma) = \alpha + \beta\|\sigma\|_1$ be the approximation error of the offline optimization oracle. Then

$$\begin{aligned}
 \mathbb{E}_{-i}[\mathbf{x}_{min,i}(\sigma)] &\geq (1 - \exp(-100\eta Ld)) (\mathbb{E}_{-i}[\mathbf{x}_{max,i}(\sigma)] - D) \\
 &\quad + \exp(-100\eta Ld) \mathbb{P}_{-i}(\mathcal{E}) \mathbb{E}_{-i} \left[\mathbf{x}_{max,i}(\sigma) - \frac{1}{10} |\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| - \frac{3\gamma(\sigma)}{100Ld} - \beta \mid \mathcal{E} \right] \\
 &\quad + \exp(-100\eta Ld) \mathbb{P}_{-i}(\mathcal{E}^c) \mathbb{E}_{-i} \left[\mathbf{x}_{min,i}(\sigma) - \frac{2\gamma(\sigma)}{100Ld} - \beta \mid \mathcal{E}^c \right] \\
 &\geq (1 - \exp(-100\eta Ld)) (\mathbb{E}_{-i}[\mathbf{x}_{max,i}(\sigma)] - D) \\
 &\quad + \exp(-100\eta Ld) \mathbb{P}_{-i}(\mathcal{E}) \mathbb{E}_{-i} \left[\mathbf{x}_{max,i}(\sigma) - \frac{1}{10} |\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| - \frac{3\gamma(\sigma)}{100Ld} - \beta \mid \mathcal{E} \right] \\
 &\quad + \exp(-100\eta Ld) \mathbb{P}_{-i}(\mathcal{E}^c) \mathbb{E}_{-i} \left[\mathbf{x}_{max,i}(\sigma) - \frac{1}{10d} \|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1 - \frac{2\gamma(\sigma)}{100Ld} - \beta \mid \mathcal{E}^c \right],
 \end{aligned}$$

where the first inequality follows from Lemmas 5, 6, the second inequality follows from the definition of \mathcal{E}^c . Rearranging the terms in the RHS and using $\mathbb{P}_{-i}(\mathcal{E}) \leq 1$ gives us

$$\begin{aligned}
 \mathbb{E}_{-i}[\mathbf{x}_{min,i}(\sigma)] &\geq (1 - \exp(-100\eta Ld)) (\mathbb{E}_{-i}[\mathbf{x}_{max,i}(\sigma)] - D) \\
 &\quad + \exp(-100\eta Ld) \mathbb{E}_{-i} \left[\mathbf{x}_{max,i}(\sigma) - \frac{3\gamma(\sigma)}{100Ld} - \beta \right] \\
 &\quad - \exp(-100\eta Ld) \mathbb{E}_{-i} \left[\frac{1}{10} |\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| + \frac{1}{10d} \|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1 \right] \\
 &\geq \mathbb{E}_{-i}[\mathbf{x}_{max,i}(\sigma)] - 100\eta LdD - \frac{3\gamma(\sigma)}{100Ld} - \beta \\
 &\quad - \mathbb{E}_{-i} \left[\frac{1}{10} |\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)| + \frac{1}{10d} \|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1 \right],
 \end{aligned}$$

where the last inequality uses the fact that $\exp(x) \geq 1 + x$. Rearranging the terms in the last inequality gives us

$$\begin{aligned}
 \mathbb{E}_{-i} [|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)|] &\leq \frac{1}{9d} \mathbb{E}_{-i} [\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1] \\
 &\quad + \frac{1000}{9} \eta LdD + \frac{\mathbb{E}_{-i}[\gamma(\sigma)]}{30Ld} + \frac{10}{9} \beta.
 \end{aligned}$$

Since the above bound holds for any $\{\sigma_j\}_{j \neq i}$, we get the following bound on the unconditioned expectation

$$\begin{aligned}
 \mathbb{E} [|\mathbf{x}_{t,i}(\sigma) - \mathbf{x}_{t+1,i}(\sigma)|] &\leq \frac{1}{9d} \mathbb{E} [\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1] \\
 &\quad + \frac{1000}{9} \eta LdD + \frac{\mathbb{E}[\gamma(\sigma)]}{30Ld} + \frac{10}{9} \beta.
 \end{aligned}$$

Plugging this in Equation (4) gives us the following bound on stability of predictions of FTPL

$$\mathbb{E} [\|\mathbf{x}_t(\sigma) - \mathbf{x}_{t+1}(\sigma)\|_1] \leq 125\eta Ld^2D + \frac{\beta d}{20\eta L} + 2\beta d + \frac{\alpha}{20L}.$$

Plugging the above bound in Equation (3) gives us the required bound on regret.

5. Non-Convex OFTPL

In this section, we present a proof of Theorem 2. Since we are in the oblivious adversary model, similar to the proof of Theorem 1, we work with a single random vector σ over the entire algorithm. We first relate the expected regret of OFTPL to the stability of its prediction. Unlike Lemma 4, the upper bound we obtain for OFTPL depends on the Lipschitz constant of $(f_t - g_t)$.

Lemma 7 *Let $\bar{\mathbf{x}}_t$ be any minimizer of $\sum_{i=1}^{t-1} f_i(\mathbf{x}) - \langle \sigma, \mathbf{x} \rangle$. The regret of OFTPL can be upper bounded as*

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T f_t(\mathbf{x}_t) - \inf_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) \right] \leq \sum_{t=1}^T \frac{L_t}{T} \mathbb{E} [\|\mathbf{x}_t - \bar{\mathbf{x}}_{t+1}\|_1] + \frac{d(\beta T + D)}{\eta T} + \alpha. \quad (6)$$

The rest of the proof of Theorem 2 involves bounding $\mathbb{E} [\|\mathbf{x}_t - \bar{\mathbf{x}}_{t+1}\|_1]$ and uses identical arguments as in the proof of Theorem 1 (see Appendix E).

6. Conclusion

In this work, we considered the problem of online learning with non-convex losses and showed that the classical FTPL algorithm with access to an offline optimization oracle achieves optimal regret rate of $O(T^{-1/2})$. We further showed that an optimistic variant of FTPL can achieve better regret bounds when the sequence of losses are predictable.

The problem of online non-convex learning has several important applications in machine learning. We believe the algorithms studied in this work can lead to improved training procedures for adversarial training and training of Generative Adversarial Networks, which currently rely on algorithms from online convex learning to solve the non-convex non-concave saddle point problems in their training objectives.

Appendix A. Proof of Proposition 3

For any deterministic algorithm, we show that there exists a sequence of loss functions over which the algorithm has $\Omega(T)$ regret. We work in the 1-dimensional setting and assume that the domain \mathcal{X} is equal to $[-D, D]$. Suppose the adversary chooses the loss functions from the following class of 1-Lipschitz functions $\mathcal{F} = \{g_a(\mathbf{x}) : a \in [-D, D]\}$, where g_a is given by

$$g_a(\mathbf{x}) = \max \left\{ 0, \frac{D}{2} - |\mathbf{x} - a| \right\}.$$

We now describe our construction of the sequence of losses that cause the deterministic algorithm to fail. Let $f_{<t} = \{f_1, \dots, f_{t-1}\}$ be the sequence of loss functions chosen until iteration $t - 1$. Let \mathbf{x}_t be the prediction of the deterministic learner at iteration t . Then we choose the loss at iteration t as $f_t(\mathbf{x}) = g_{\mathbf{x}_t}(\mathbf{x})$. It is easy to see that, after T iterations, the loss suffered by the learner is equal to $\frac{DT}{2}$. Whereas, the loss of the best action in hindsight can be upper bounded as

$$\inf_{\mathbf{x} \in [-D, D]} \sum_{t=1}^T f_t(\mathbf{x}) \leq \frac{DT}{4}.$$

This shows that the regret of any deterministic algorithm is $\Omega(1)$.

Appendix B. Non-oblivious to Oblivious Adversary Model

In the oblivious adversary model, the actions $\{f_t\}_{t=1}^T$ of the adversary are assumed to be independent of the predictions $\{\mathbf{x}_t\}_{t=1}^T$ of the FTPL/OFTPL algorithm. In this model, we assume that the sequence of losses $\{f_t\}_{t=1}^T$ is fixed ahead of time. Whereas in the non-oblivious adversary model, the actions of the adversary are allowed to depend on the past predictions of the algorithm, *i.e.*, each f_t is given by $f_t := F_t[\mathbf{x}_{<t}]$ for some function $F_t : \mathcal{X}^{t-1} \rightarrow \mathcal{F}$, where \mathcal{F} is the set of all possible actions of the adversary and $\mathbf{x}_{<t}$ is a shorthand for $\{\mathbf{x}_1 \dots \mathbf{x}_{t-1}\}$ and F_1 is a constant function. Note that the functions $F_1 \dots F_T$ uniquely determine a non-oblivious adversary.

Let P_t be the conditional distribution of the prediction \mathbf{x}_t of the FTPL/OFTPL algorithm, conditioned on the past predictions $\mathbf{x}_{<t}$. Note that when the adversary is oblivious, P_t is independent of $\mathbf{x}_{<t}$. Moreover, in both oblivious and non-oblivious models, P_t is fully determined by the past actions $f_{<t}$ of the adversary. Let $f_t(P_t)$ denote the expected loss $\mathbb{E}_{\mathbf{x} \sim P_t} [f_t(\mathbf{x}) | \mathbf{x}_{<t}]$.

The following Theorem shows that any algorithm which is guaranteed to work against an oblivious adversary also works against a non-oblivious adversary. This is an adaptation of Lemma 4.1 of [Cesa-Bianchi and Lugosi \(2006\)](#) to the setting studied in this paper.

Theorem 8 *Let B be a positive constant. Suppose the FTPL, OFTPL algorithms satisfy the following regret bound against an oblivious adversary*

$$\mathbb{E} \left[\sum_{t=1}^T f_t(\mathbf{x}_t) - \inf_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) \right] \leq B, \quad \forall f_1 \dots f_T \in \mathcal{F}. \quad (7)$$

Then these algorithms satisfy the following regret bound against a non-oblivious adversary

$$\sum_{t=1}^T f_t(P_t) - \inf_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) \leq B.$$

Proof Consider the non-oblivious adversary model. For any $\mathbf{x} \in \mathcal{X}$ we have

$$\begin{aligned} & \sum_{t=1}^T f_t(P_t) - \sum_{t=1}^T f_t(\mathbf{x}) \\ &= \sum_{t=1}^T F_t[\mathbf{x}_{<t}](P_t) - \sum_{t=1}^T F_t[\mathbf{x}_{<t}](\mathbf{x}) \\ &\stackrel{(a)}{\leq} \sup_{F_1, \dots, F_T} \left(\sum_{t=1}^T F_t[\mathbf{x}_{<t}](P_t) - \sum_{t=1}^T F_t[\mathbf{x}_{<t}](\mathbf{x}) \right) \\ &\stackrel{(b)}{=} \sup_{g_1 \in \mathcal{F}} \left(g_1(P_1) - g_1(\mathbf{x}) + \sup_{g_2 \in \mathcal{F}} \left(g_2(P_2) - g_2(\mathbf{x}) + \sup_{g_3 \in \mathcal{F}} \left(\dots + \sup_{g_T \in \mathcal{F}} g_T(P_T) - g_T(\mathbf{x}) \right) \right) \right), \end{aligned}$$

where the supremum in (a) is over all possible non-oblivious adversaries. To see why (b) holds, consider $T = 2$. Then

$$\begin{aligned} & \sup_{F_1, F_2} (F_1[\mathbf{x}_{<1}](P_1) - F_1[\mathbf{x}_{<1}](\mathbf{x}) + F_2[\mathbf{x}_{<2}](P_2) - F_2[\mathbf{x}_{<2}](\mathbf{x})) \\ &= \sup_{g_1 \in \mathcal{F}, F_2} (g_1(P_1) - g_1(\mathbf{x}) + F_2[\mathbf{x}_{<2}](P_2) - F_2[\mathbf{x}_{<2}](\mathbf{x})) \\ &= \sup_{g_1} \left(g_1(P_1) - g_1(\mathbf{x}) + \sup_{g_2 \in \mathcal{F}} g_2(P_2) - g_2(\mathbf{x}) \right). \end{aligned}$$

This shows that a good strategy for the adversary is to set $F_2[\mathbf{x}_{<2}]$ to be a maximizer of $g_2(P_2) - g_2(\mathbf{x})$. Using a similar argument we can show that (b) holds for $T > 2$.

Next, we show that

$$\begin{aligned} & \sup_{g_1 \in \mathcal{F}} \left(g_1(P_1) - g_1(\mathbf{x}) + \sup_{g_2 \in \mathcal{F}} \left(g_2(P_2) - g_2(\mathbf{x}) + \sup_{g_3 \in \mathcal{F}} \left(\cdots + \sup_{g_T \in \mathcal{F}} g_T(P_T) - g_T(\mathbf{x}) \right) \right) \right) \\ &= \sup_{g_1 \dots g_T \in \mathcal{F}} \left(\sum_{t=1}^T g_t(P_t) - g_t(\mathbf{x}) \right). \end{aligned}$$

Moreover, we show that the maximizers of the RHS objective are independent of the predictions $\{\mathbf{x}_t\}_{t=1}^T$ of the algorithm. This would then imply that the RHS is exactly equal to the regret of the algorithm under the oblivious adversary model, which is upper bounded by B . To see why the above statements are true, again consider the case of $T = 2$. First note that $g_1(P_1) - g_1(\mathbf{x})$ is independent of g_2 . So $g_1(P_1) - g_1(\mathbf{x})$ can be pushed inside the inner supremum. So we have

$$\begin{aligned} & \sup_{g_1 \in \mathcal{F}} \left(g_1(P_1) - g_1(\mathbf{x}) + \sup_{g_2 \in \mathcal{F}} (g_2(P_2) - g_2(\mathbf{x})) \right) \\ &= \sup_{g_1, g_2 \in \mathcal{F}} (g_1(P_1) - g_1(\mathbf{x}) + g_2(P_2) - g_2(\mathbf{x})) \end{aligned}$$

To see why the maximizers of the RHS are independent of $\mathbf{x}_1, \mathbf{x}_2$, note that P_1 is independent of $\mathbf{x}_1, \mathbf{x}_2$. Moreover, P_2 is fully determined by g_1 . So the objective is independent of $\mathbf{x}_1, \mathbf{x}_2$. This shows that the maximizers are independent of $\mathbf{x}_1, \mathbf{x}_2$. Using a similar argument we can show that the above claim holds for $T > 2$. Finally, from the regret bound against an oblivious adversary in Equation (7), we have

$$\sup_{g_1 \dots g_T \in \mathcal{F}} \left(\sum_{t=1}^T g_t(P_t) - g_t(\mathbf{x}) \right) = \sup_{g_1 \dots g_T \in \mathcal{F}} \mathbb{E} \left[\sum_{t=1}^T g_t(\mathbf{x}_t) - \sum_{t=1}^T g_t(\mathbf{x}) \right] \leq B.$$

This shows that for any $\mathbf{x} \in \mathcal{X}$, $\sum_{t=1}^T f_t(P_t) - \sum_{t=1}^T f_t(\mathbf{x}) \leq B$. ■

Appendix C. Proof of Lemma 4

Let $\gamma(\sigma) = \alpha + \beta\|\sigma\|_1$. For any $\mathbf{x}^* \in \mathcal{X}$ we have

$$\begin{aligned} & \sum_{t=1}^T [f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*)] \\ &= \sum_{t=1}^T [f_t(\mathbf{x}_t) - f_t(\mathbf{x}_{t+1})] + \sum_{t=1}^T [f_t(\mathbf{x}_{t+1}) - f_t(\mathbf{x}^*)] \\ &\leq \sum_{t=1}^T L\|\mathbf{x}_t - \mathbf{x}_{t+1}\|_1 + \sum_{t=1}^T [f_t(\mathbf{x}_{t+1}) - f_t(\mathbf{x}^*)]. \end{aligned}$$

We now use induction to show that $\sum_{t=1}^T [f_t(\mathbf{x}_{t+1}) - f_t(\mathbf{x}^*)] \leq \gamma(\sigma)T + \langle \sigma, \mathbf{x}_2 - \mathbf{x}^* \rangle$.

Base Case ($T = 1$). Since \mathbf{x}_2 is an approximate minimizer of $f_1(\mathbf{x}) - \langle \sigma, \mathbf{x} \rangle$, we have

$$f_1(\mathbf{x}_2) - \langle \sigma, \mathbf{x}_2 \rangle \leq \min_{\mathbf{x} \in \mathcal{X}} f_1(\mathbf{x}) - \langle \sigma, \mathbf{x} \rangle + \gamma(\sigma) \leq f_1(\mathbf{x}^*) - \langle \sigma, \mathbf{x}^* \rangle + \gamma(\sigma),$$

where the last inequality holds for any $\mathbf{x}^* \in \mathcal{X}$. This shows that $f_1(\mathbf{x}_2) - f_1(\mathbf{x}^*) \leq \gamma(\sigma) + \langle \sigma, \mathbf{x}_2 - \mathbf{x}^* \rangle$.

Induction Step. Suppose the claim holds for all $T \leq T_0 - 1$. We now show that it also holds for T_0 .

$$\begin{aligned} & \sum_{t=1}^{T_0} f_t(\mathbf{x}_{t+1}) \\ &\stackrel{(a)}{\leq} \left[\sum_{t=1}^{T_0-1} f_t(\mathbf{x}_{T_0+1}) + \langle \sigma, \mathbf{x}_2 - \mathbf{x}_{T_0+1} \rangle + \gamma(\sigma)(T_0 - 1) \right] + f_{T_0}(\mathbf{x}_{T_0+1}) \\ &= \left[\sum_{t=1}^{T_0} f_t(\mathbf{x}_{T_0+1}) - \langle \sigma, \mathbf{x}_{T_0+1} \rangle \right] + \langle \sigma, \mathbf{x}_2 \rangle + \gamma(\sigma)(T_0 - 1) \\ &\stackrel{(b)}{\leq} \sum_{t=1}^{T_0} f_t(\mathbf{x}^*) + \langle \sigma, \mathbf{x}_2 - \mathbf{x}^* \rangle + \gamma(\sigma)T_0, \quad \forall \mathbf{x}^* \in \mathcal{X}, \end{aligned}$$

where (a) follows since the claim holds for any $T \leq T_0 - 1$, and (b) follows from the approximate optimality of \mathbf{x}_{T_0+1} .

Using this result, we get the following upper bound on the expected regret of FTPL

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T f_t(\mathbf{x}_t) - \inf_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) \right] &\leq L \sum_{t=1}^T \mathbb{E} [\|\mathbf{x}_t - \mathbf{x}_{t+1}\|_1] + \mathbb{E} [\gamma(\sigma)T + \langle \sigma, \mathbf{x}_2 - \mathbf{x}^* \rangle] \\ &\leq L \sum_{t=1}^T \mathbb{E} [\|\mathbf{x}_t - \mathbf{x}_{t+1}\|_1] + (\beta T + D) \left(\sum_{i=1}^d \mathbb{E} [\sigma_i] \right) + \alpha T \end{aligned}$$

The proof of the Lemma now follows from the following property of exponential distribution

$$\mathbb{E} [\sigma_i] = \frac{1}{\eta_i}.$$

Appendix D. Proof of Lemma 7

The proof uses similar arguments as in the proof of [Rakhlin and Sridharan \(2012\)](#) for Optimistic FTRL. Let $\Delta_t(\mathbf{x}) = f_t(\mathbf{x}) - g_t(\mathbf{x})$ and $\gamma(\sigma) = \alpha + \beta\|\sigma\|_1$. For any $\mathbf{x}^* \in \mathcal{X}$ we have

$$\begin{aligned} & \sum_{t=1}^T [f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*)] \\ &= \sum_{t=1}^T [\Delta_t(\mathbf{x}_t) - \Delta_t(\bar{\mathbf{x}}_{t+1})] + \sum_{t=1}^T [g_t(\mathbf{x}_t) - g_t(\bar{\mathbf{x}}_{t+1})] + \sum_{t=1}^T [f_t(\bar{\mathbf{x}}_{t+1}) - f_t(\mathbf{x}^*)] \\ &\leq \sum_{t=1}^T L_t \|\mathbf{x}_t - \bar{\mathbf{x}}_{t+1}\|_1 + \sum_{t=1}^T [g_t(\mathbf{x}_t) - g_t(\bar{\mathbf{x}}_{t+1})] + \sum_{t=1}^T [f_t(\bar{\mathbf{x}}_{t+1}) - f_t(\mathbf{x}^*)]. \end{aligned}$$

We use induction to show that the following holds for any $T, \mathbf{x}^* \in \mathcal{X}$

$$\sum_{t=1}^T [g_t(\mathbf{x}_t) - g_t(\bar{\mathbf{x}}_{t+1})] + \sum_{t=1}^T [f_t(\bar{\mathbf{x}}_{t+1}) - f_t(\mathbf{x}^*)] \leq \langle \sigma, \bar{\mathbf{x}}_2 - \mathbf{x}^* \rangle + \gamma(\sigma)(T - 1).$$

Base Case ($T = 1$). First note that $g_1 = 0$. Since $\bar{\mathbf{x}}_2$ is a minimizer of $f_1(\mathbf{x}) - \langle \sigma, \mathbf{x} \rangle$, we have

$$f_1(\bar{\mathbf{x}}_2) - \langle \sigma, \bar{\mathbf{x}}_2 \rangle \leq f_1(\mathbf{x}^*) - \langle \sigma, \mathbf{x}^* \rangle, \quad \forall \mathbf{x}^* \in \mathcal{X}.$$

This shows that $f_1(\bar{\mathbf{x}}_2) - f_1(\mathbf{x}^*) \leq \langle \sigma, \bar{\mathbf{x}}_2 - \mathbf{x}^* \rangle$.

Induction Step. Suppose the claim holds for all $T \leq T_0 - 1$. We now show that it also holds for T_0 . Consider the following series of inequalities

$$\begin{aligned} & \sum_{t=1}^{T_0} [g_t(\mathbf{x}_t) - g_t(\bar{\mathbf{x}}_{t+1})] + \sum_{t=1}^{T_0} f_t(\bar{\mathbf{x}}_{t+1}) \\ &\stackrel{(a)}{\leq} \left[\sum_{t=1}^{T_0-1} f_t(\mathbf{x}_{T_0}) + \langle \sigma, \bar{\mathbf{x}}_2 - \mathbf{x}_{T_0} \rangle + \gamma(\sigma)(T_0 - 2) \right] + [g_{T_0}(\mathbf{x}_{T_0}) - g_{T_0}(\bar{\mathbf{x}}_{T_0+1}) + f_{T_0}(\bar{\mathbf{x}}_{T_0+1})] \\ &= \left[\sum_{t=1}^{T_0-1} f_t(\mathbf{x}_{T_0}) + g_{T_0}(\mathbf{x}_{T_0}) - \langle \sigma, \mathbf{x}_{T_0} \rangle \right] + [\langle \sigma, \bar{\mathbf{x}}_2 \rangle - g_{T_0}(\bar{\mathbf{x}}_{T_0+1}) + f_{T_0}(\bar{\mathbf{x}}_{T_0+1})] + \gamma(\sigma)(T_0 - 2) \\ &\stackrel{(b)}{\leq} \left[\sum_{t=1}^{T_0-1} f_t(\bar{\mathbf{x}}_{T_0+1}) + g_{T_0}(\bar{\mathbf{x}}_{T_0+1}) - \langle \sigma, \bar{\mathbf{x}}_{T_0+1} \rangle \right] + [\langle \sigma, \bar{\mathbf{x}}_2 \rangle - g_{T_0}(\bar{\mathbf{x}}_{T_0+1}) + f_{T_0}(\bar{\mathbf{x}}_{T_0+1})] + \gamma(\sigma)(T_0 - 1) \\ &= \left[\sum_{t=1}^{T_0} f_t(\bar{\mathbf{x}}_{T_0+1}) - \langle \sigma, \bar{\mathbf{x}}_{T_0+1} \rangle \right] + \langle \sigma, \bar{\mathbf{x}}_2 \rangle + \gamma(\sigma)(T_0 - 1) \\ &\stackrel{(c)}{\leq} \sum_{t=1}^{T_0} f_t(\mathbf{x}^*) + \langle \sigma, \bar{\mathbf{x}}_2 - \mathbf{x}^* \rangle + \gamma(\sigma)(T_0 - 1), \end{aligned}$$

where (a) follows since the claim holds for any $T \leq T_0 - 1$, (b) follows from the approximate optimality of \mathbf{x}_{T_0} and (c) follows from the optimality of $\bar{\mathbf{x}}_{T_0+1}$.

This gives the following upper bound on the regret of OFTPL

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \inf_{\mathbf{x} \in \mathcal{X}} \sum_{i=1}^T f_t(\mathbf{x}) \leq \sum_{t=1}^T L_t \|\mathbf{x}_t - \bar{\mathbf{x}}_{t+1}\|_1 + \frac{d(\beta T + D)}{\eta} + \alpha T.$$

Appendix E. Proof of Theorem 2

Lemma 9 *Let $\mathbf{x}_t(\sigma)$ be the prediction of OFTPL in iteration t , with random perturbation σ . Then the following monotonicity property holds for any $c > 0$*

$$\mathbf{x}_{t,i}(\sigma + c\mathbf{e}_i) \geq \mathbf{x}_{t,i}(\sigma) - \frac{2(\alpha + \beta\|\sigma\|_1)}{c} - \beta.$$

Proof Let $f_{1:t}(\mathbf{x}) = \sum_{i=1}^t f_i(\mathbf{x})$ and $\sigma' = \sigma + c\mathbf{e}_i$. Moreover, let $\gamma(\sigma) = \alpha + \beta\|\sigma\|_1$ be the approximation error of the offline optimization oracle. From the approximate optimality of $\mathbf{x}_t(\sigma)$ we have

$$\begin{aligned} & f_{1:t-1}(\mathbf{x}_t(\sigma)) + g_t(\mathbf{x}_t(\sigma)) - \langle \sigma, \mathbf{x}_t(\sigma) \rangle \\ & \leq f_{1:t-1}(\mathbf{x}_t(\sigma')) + g_t(\mathbf{x}_t(\sigma')) - \langle \sigma, \mathbf{x}_t(\sigma') \rangle + \gamma(\sigma) \\ & = f_{1:t-1}(\mathbf{x}_t(\sigma')) + g_t(\mathbf{x}_t(\sigma')) - \langle \sigma', \mathbf{x}_t(\sigma') \rangle + c\mathbf{x}_{t,i}(\sigma') + \gamma(\sigma) \\ & \stackrel{(a)}{\leq} f_{1:t-1}(\mathbf{x}_t(\sigma)) + g_t(\mathbf{x}_t(\sigma)) - \langle \sigma', \mathbf{x}_t(\sigma) \rangle + c\mathbf{x}_{t,i}(\sigma') + \gamma(\sigma) + \gamma(\sigma') \\ & = f_{1:t-1}(\mathbf{x}_t(\sigma)) + g_t(\mathbf{x}_t(\sigma)) - \langle \sigma, \mathbf{x}_t(\sigma) \rangle + c(\mathbf{x}_{t,i}(\sigma') - \mathbf{x}_{t,i}(\sigma)) + \gamma(\sigma) + \gamma(\sigma'), \end{aligned}$$

where (a) follows from the approximate optimality of $\mathbf{x}_t(\sigma')$. Combining the first and last terms in the above expression, we get $\mathbf{x}_{t,i}(\sigma') \geq \mathbf{x}_{t,i}(\sigma) - \frac{2\gamma(\sigma)}{c} - \beta$. \blacksquare

We note that a similar argument can be used to show that $\bar{\mathbf{x}}_{t,i}(\alpha + c\mathbf{e}_i) \geq \bar{\mathbf{x}}_{t,i}(\alpha)$.

Lemma 10 *Suppose $\|\mathbf{x}_t(\sigma) - \bar{\mathbf{x}}_{t+1}(\sigma)\|_1 \leq 10d \cdot |\mathbf{x}_{t,i}(\sigma) - \bar{\mathbf{x}}_{t+1,i}(\sigma)|$. For $\sigma' = \sigma + 100L_t d\mathbf{e}_i$, we have*

$$\begin{aligned} \min(\mathbf{x}_{t,i}(\sigma'), \bar{\mathbf{x}}_{t+1,i}(\sigma')) & \geq \max(\mathbf{x}_{t,i}(\sigma), \bar{\mathbf{x}}_{t+1,i}(\sigma)) - \frac{1}{10} |\mathbf{x}_{t,i}(\sigma) - \bar{\mathbf{x}}_{t+1,i}(\sigma)| \\ & \quad - \frac{3(\alpha + \beta\|\sigma\|_1)}{100L_t d} - \beta. \end{aligned}$$

Proof Let $f_{1:t}(\mathbf{x}) = \sum_{i=1}^t f_i(\mathbf{x})$ and let $\gamma(\sigma) = \alpha + \beta\|\sigma\|_1$ be the approximation error of the offline optimization oracle. From the approximate optimality of $\mathbf{x}_t(\sigma)$, we have

$$\begin{aligned} & f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma, \mathbf{x}_t(\sigma) \rangle + f_t(\mathbf{x}_t(\sigma)) \\ & \leq f_{1:t-1}(\bar{\mathbf{x}}_{t+1}(\sigma)) + g_t(\bar{\mathbf{x}}_{t+1}(\sigma)) - \langle \sigma, \bar{\mathbf{x}}_{t+1}(\sigma) \rangle \\ & \quad + f_t(\mathbf{x}_t(\sigma)) - g_t(\mathbf{x}_t(\sigma)) + \gamma(\sigma) \\ & \stackrel{(a)}{\leq} f_{1:t-1}(\bar{\mathbf{x}}_{t+1}(\sigma)) + g_t(\bar{\mathbf{x}}_{t+1}(\sigma)) - \langle \sigma, \bar{\mathbf{x}}_{t+1}(\sigma) \rangle \\ & \quad + f_t(\bar{\mathbf{x}}_{t+1}(\sigma)) - g_t(\bar{\mathbf{x}}_{t+1}(\sigma)) + L_t \|\mathbf{x}_t(\sigma) - \bar{\mathbf{x}}_{t+1}(\sigma)\|_1 + \gamma(\sigma) \\ & \stackrel{(b)}{\leq} f_{1:t-1}(\bar{\mathbf{x}}_{t+1}(\sigma)) + g_t(\bar{\mathbf{x}}_{t+1}(\sigma)) - \langle \sigma, \bar{\mathbf{x}}_{t+1}(\sigma) \rangle \\ & \quad + f_t(\bar{\mathbf{x}}_{t+1}(\sigma)) - g_t(\bar{\mathbf{x}}_{t+1}(\sigma)) + 10L_t d |\mathbf{x}_{t,i}(\sigma) - \bar{\mathbf{x}}_{t+1,i}(\sigma)| + \gamma(\sigma), \end{aligned}$$

where (a) follows from the Lipschitz property of $f_t(\cdot)$ and (b) follows from our assumption on $\|\mathbf{x}_t(\sigma) - \bar{\mathbf{x}}_{t+1}(\sigma)\|_1$. Next, from the optimality of $\bar{\mathbf{x}}_{t+1}(\sigma')$, we have

$$\begin{aligned} & f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma, \mathbf{x}_t(\sigma) \rangle + f_t(\mathbf{x}_t(\sigma)) \\ &= f_{1:t-1}(\mathbf{x}_t(\sigma)) - \langle \sigma', \mathbf{x}_t(\sigma) \rangle + f_t(\mathbf{x}_t(\sigma)) + \langle 100L_t d \mathbf{e}_i, \mathbf{x}_t(\sigma) \rangle \\ &\geq f_{1:t-1}(\bar{\mathbf{x}}_{t+1}(\sigma')) - \langle \sigma', \bar{\mathbf{x}}_{t+1}(\sigma') \rangle + f_t(\bar{\mathbf{x}}_{t+1}(\sigma')) + 100L_t d \mathbf{x}_{t,i}(\sigma) \\ &= f_{1:t-1}(\bar{\mathbf{x}}_{t+1}(\sigma')) - \langle \sigma, \bar{\mathbf{x}}_{t+1}(\sigma') \rangle + f_t(\bar{\mathbf{x}}_{t+1}(\sigma')) + 100L_t d(\mathbf{x}_{t,i}(\sigma) - \bar{\mathbf{x}}_{t+1,i}(\sigma')) \\ &\geq f_{1:t-1}(\bar{\mathbf{x}}_{t+1}(\sigma)) - \langle \sigma, \bar{\mathbf{x}}_{t+1}(\sigma) \rangle + f_t(\bar{\mathbf{x}}_{t+1}(\sigma)) + 100L_t d(\mathbf{x}_{t,i}(\sigma) - \bar{\mathbf{x}}_{t+1,i}(\sigma')), \end{aligned}$$

where the last inequality follows from the optimality of $\bar{\mathbf{x}}_{t+1}(\sigma)$. Combining the above two equations, we get

$$\bar{\mathbf{x}}_{t+1,i}(\sigma') - \mathbf{x}_{t,i}(\sigma) \geq -\frac{1}{10} |\mathbf{x}_{t,i}(\sigma) - \bar{\mathbf{x}}_{t+1,i}(\sigma)| - \frac{\gamma(\sigma)}{100L_t d}.$$

A similar argument shows that

$$\mathbf{x}_{t,i}(\sigma') - \bar{\mathbf{x}}_{t+1,i}(\sigma) \geq -\frac{1}{10} |\mathbf{x}_{t,i}(\sigma) - \bar{\mathbf{x}}_{t+1,i}(\sigma)| - \frac{\gamma(\sigma)}{100L_t d}.$$

Finally, from the monotonicity property in Lemma 5 we know that

$$\bar{\mathbf{x}}_{t+1,i}(\sigma') - \bar{\mathbf{x}}_{t+1,i}(\sigma) \geq 0, \quad \mathbf{x}_{t,i}(\sigma') - \mathbf{x}_{t,i}(\sigma) \geq -\frac{3\gamma(\sigma)}{100L_t d} - \beta.$$

Combining the above four inequalities gives us the required result. ■

The rest of the proof relies on the monotonicity properties showed in the above two Lemmas to bound $\mathbb{E}[\|\mathbf{x}_t - \bar{\mathbf{x}}_{t+1}\|_1]$ and uses identical arguments as in the proof of Theorem 1.

References

- Jacob Abernethy, Chansoo Lee, and Ambuj Tewari. Perturbation techniques in online learning and optimization. *Perturbations, Optimization, and Statistics*, page 233, 2016.
- Naman Agarwal, Alon Gonen, and Elad Hazan. Learning in non-convex games with an optimization oracle. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 18–29, Phoenix, USA, 25–28 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v99/agarwal19a.html>.
- Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Xiand Gao, Xiaobo Li, and Shuzhong Zhang. Online learning with non-convex losses and non-stationary regret. In *International Conference on Artificial Intelligence and Statistics*, pages 235–243, 2018.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Elad Hazan, Karan Singh, and Cyril Zhang. Efficient regret minimization in non-convex games. *arXiv preprint arXiv:1708.00075*, 2017.
- Adam Kalai and Santosh Vempala. Efficient algorithms for on-line optimization. *Journal of Computer and System Sciences*, 71, 2016.
- Walid Krichene, Maximilian Balandat, Claire Tomlin, and Alexandre Bayen. The hedge algorithm on a continuum. In *International Conference on Machine Learning*, pages 824–832, 2015.
- Odalric-Ambrym Maillard and Rémi Munos. Online learning in adversarial lipschitz environments. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 305–320. Springer, 2010.
- Brendan McMahan. Follow-the-regularized-leader and mirror descent: Equivalence theorems and ℓ_1 regularization. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 525–533, 2011.
- Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. *arXiv preprint arXiv:1208.3728*, 2012.
- Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems*, pages 3066–3074, 2013.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.