

Online Linear Optimization via Smoothing

Jacob Abernethy

JABERNET@UMICH.EDU

Chansoo Lee

CHANSOOL@UMICH.EDU

Computer Science and Engineering Division, University of Michigan, Ann Arbor

Abhinav Sinha

ABSI@UMICH.EDU

Electrical and Computer Engineering Division, University of Michigan, Ann Arbor

Ambuj Tewari

TEWARIA@UMICH.EDU

Department of Statistics, University of Michigan, Ann Arbor

Abstract

We present a new optimization-theoretic approach to analyzing Follow-the-Leader style algorithms, particularly in the setting where perturbations are used as a tool for regularization. We show that adding a strongly convex penalty function to the decision rule and adding stochastic perturbations to data correspond to deterministic and stochastic smoothing operations, respectively. We establish an equivalence between “Follow the Regularized Leader” and “Follow the Perturbed Leader” up to the smoothness properties. This intuition leads to a new generic analysis framework that recovers and improves the previous known regret bounds of the class of algorithms commonly known as Follow the Perturbed Leader.

1. Introduction

In this paper, we study *online learning* (other names include adversarial learning or no-regret learning) in which the learner iteratively plays actions based on the data received up to the previous iteration. The data sequence is chosen by an adversary and the learner’s goal is to minimize the worst-case regret. The key to developing optimal algorithms is regularization, interpreted as *hedging* against an adversarial future input and avoiding *overfitting* to the observed data. In this paper, we focus on regularization techniques for online linear optimization problems where the learner’s action is evaluated on a linear reward function.

Follow the Regularized Leader (FTRL) is an algorithm that uses explicit regularization *via penalty function*, which directly changes the optimization objective. At every iteration, FTRL selects an action by optimizing $\arg \max_w f(w, \Theta) - \mathcal{R}(w)$ where f is the true objective, Θ is the observed data, and \mathcal{R} is a strongly convex penalty function such as the well-known ℓ_2 -regularizer $\|\cdot\|_2$. The regret analysis of FTRL reduces to the analysis of the second-order behavior of the penalty function (Shalev-Shwartz, 2012), which is well-studied due to the powerful convex analysis tools. In fact, regularization via penalty methods for online learning in general are very well understood. Srebro et al. (2011) proved that Mirror Descent, a regularization via penalty method, achieves a nearly optimal regret guarantee for a general class of online learning problems, and McMahan (2011) showed that FTRL is equivalent to Mirror Descent under some assumptions.

Follow the Perturbed Leader (FTPL), on the other hand, uses implicit regularization *via perturbations*. At every iteration, FTPL selects an action by optimizing $\arg \max_w f(w, \Theta + u)$ where

Θ is the observed data and u is some random noise vector, often referred to as a “perturbation” of the input. Unfortunately, the analysis of FTPL lacks a generic framework and relies substantially on clever algebra tricks and heavy probabilistic analysis (Kalai and Vempala, 2005; Devroye et al., 2013; van Erven et al., 2014). Convex analysis techniques, which led to our current thorough understanding of FTRL, have not been applied to FTPL, partly because the decision rule of FTPL does not explicitly contain a convex function.

In this paper, we present a new analysis framework that makes it possible to analyze FTPL in the same way that FTRL has been analyzed, particularly with regards to *second-order properties* of convex functions. We show that both FTPL and FTRL naturally arise as *smoothing operations* of a non-smooth potential function and the regret analysis boils down to controlling the *smoothing parameters* as defined in Section 3. This new unified analysis framework not only recovers the known optimal regret bounds, but also gives a new type of generic regret bounds.

Prior to our work, Rakhlin et al. (2012) showed that both FTPL and FTRL naturally arise as *admissible relaxations* of the minimax value of the game between the learner and adversary. In short, adding a random perturbation and adding a regularization penalty function are both optimal ways to simulate the worst-case future input sequence. We establish a stronger connection between FTRL and FTPL; both algorithms are derived from smoothing operations and they are *equivalent* up to the smoothing parameters. This equivalence is in fact a very strong result, considering the fact that Harsanyi (1973) showed that there is no general bijection between FTPL and FTRL.

This paper also aligns itself with the previous work that studied the connection between explicit regularization *via penalty* and implicit regularization *via perturbations*. Bishop (1995) showed that adding Gaussian noise to features of the training examples is equivalent to Tikhonov regularization, and more recently Wager et al. (2013) showed that for online learning, dropout training (Hinton et al., 2012) is similar to AdagGrad (Duchi et al., 2010) in that both methods scale features by the Fisher information. These results are derived from Taylor approximations, but our FTPL-FTRL connection is derived from the convex conjugate duality.

An interesting feature of our analysis framework is that we can directly apply existing techniques from the optimization literature, and conversely, our new findings in online linear optimization may apply to optimization theory. In Section 4.3, a straightforward application of the results on Gaussian smoothing by Nesterov (2011) and Duchi et al. (2012) gives a generic regret bound for an arbitrary online linear optimization problem. In Section 4.1 and 4.2, we improve this bound for the special cases that correspond to canonical online linear optimization problems, and these results may be of interest to the optimization community.

2. Preliminaries

2.1. Convex Analysis

Let f be a differentiable, closed, and proper convex function whose domain is $\text{dom} f \subseteq \mathbb{R}^N$. We say that f is L -Lipschitz with respect to a norm $\|\cdot\|$ when f satisfies $|f(x) - f(y)| \leq L\|x - y\|$ for all $x, y \in \text{dom}(f)$.

The *Bregman divergence*, denoted $D_f(y, x)$, is the gap between $f(y)$ and the linear approximation of $f(y)$ around x . Formally, $D_f(y, x) = f(y) - f(x) - \langle \nabla f(x), y - x \rangle$. We say that f is β -strongly convex with respect to a norm $\|\cdot\|$ if we have $D_f(y, x) \geq \frac{\beta}{2}\|y - x\|^2$ for all $x, y \in \text{dom} f$. Similarly, f is said to be β -strongly smooth with respect to a norm $\|\cdot\|$ if we have

$D_f(y, x) \leq \frac{\beta}{2} \|y - x\|^2$ for all $x, y \in \text{dom} f$. The Bregman divergence measures how fast the gradient changes, or equivalently, how large the second derivative is. In fact, we can bound the Bregman divergence by analyzing the local behavior of Hessian, as the following adaptation of [Abernethy et al. \(2013, Lemma 4.6\)](#) shows.

Lemma 1 *Let f be a twice-differentiable convex function with $\text{dom} f \subseteq \mathbb{R}^N$. Let $x \in \text{dom} f$, such that $v^T \nabla^2 f(x + \alpha v)v \in [a, b]$ ($a \leq b$) for all $\alpha \in [0, 1]$. Then, $a\|v\|^2/2 \leq D_f(x+v, x) \leq b\|v\|^2/2$.*

The *Fenchel conjugate* of f is $f^*(\theta) = \sup_{w \in \text{dom}(f)} \{\langle w, \theta \rangle - f(w)\}$, and it is a dual mapping that satisfies $f = (f^*)^*$ and $\nabla f^* \in \text{dom}(f)$. By the strong convexity-strong smoothness duality, f is β -strongly smooth with respect to a norm $\|\cdot\|$ if and only if f^* is $\frac{1}{\beta}$ -strongly smooth with respect to the dual norm $\|\cdot\|_*$. For more details and proofs, readers are referred to an excellent survey by [Shalev-Shwartz \(2012\)](#).

2.2. Online Linear Optimization

Let \mathcal{X} and \mathcal{Y} be convex and closed subsets of \mathbb{R}^N . The online linear optimization is defined to be the following iterative process:

On round $t = 1, \dots, T$,

- the learner plays $w_t \in \mathcal{X}$.
- the adversary reveals $\theta_t \in \mathcal{Y}$.
- the learner receives a reward¹ $\langle w_t, \theta_t \rangle$.

We say \mathcal{X} is the *decision set* and \mathcal{Y} is the *reward set*. Let $\Theta_t = \sum_{s=1}^t \theta_s$ be the cumulative reward. The learner's goal is to minimize the (external) regret, defined as:

$$\text{Regret} = \underbrace{\max_{w \in \mathcal{X}} \langle w, \Theta_T \rangle}_{\text{baseline potential}} - \sum_{t=1}^T \langle w_t, \theta_t \rangle. \quad (1)$$

The *baseline potential function* $\Phi(\Theta) := \max_{w \in \mathcal{X}} \langle w, \Theta \rangle$ is the comparator term against which we define the regret, and it coincides with the *support function* of \mathcal{X} . For a bounded compact set \mathcal{X} , the support function of \mathcal{X} is sublinear² and Lipschitz continuous with respect to any norm $\|\cdot\|$ with the Lipschitz constant $\sup_{x \in \mathcal{X}} \|x\|$. For more details and proofs, readers are referred to [Rockafellar \(1997, Section 13\)](#) or [Molchanov \(2005, Appendix F\)](#).

3. Online Linear Optimization Algorithms via Smoothing

3.1. Gradient-Based Prediction Algorithm

Follow-the-Leader style algorithms solve an optimization objective every round and play an action of the form $w_t = \arg \max_{w \in \mathcal{X}} f(w, \Theta_{t-1})$ given a fixed Θ_{t-1} . For example, Follow the Regularized Leader maximizes $f(w, \Theta) = \langle w, \Theta \rangle - \mathcal{R}(w)$ where \mathcal{R} is a strongly convex regularizer, and Follow the Perturbed Leader maximizes $f = \langle w, \Theta + u \rangle$ where u is a random noise. A surprising

1. Our somewhat less conventional choice of maximizing the reward instead of minimizing the loss was made so that we directly analyze the convex function $\max(\cdot)$ without cumbersome sign changes.
 2. A function f is sublinear if it is positive homogeneous (i.e., $f(ax) = af(x)$ for all $a > 0$) and subadditive (i.e., $f(x) + f(y) \geq f(x + y)$).

fact about these algorithms is that there are many scenarios in which the action w_t is exactly *the gradient* of some scalar potential function Φ_t evaluated at Θ_{t-1} . This perspective gives rise to what we call the Gradient-based Prediction Algorithm (GBPA), presented below. Note that [Cesa-Bianchi and Lugosi \(2006, Theorem 11.6\)](#) presented a similar algorithm, but our formulation eliminates all dual mappings.

Algorithm 1: Gradient-Based Prediction Algorithm (GBPA)

Input: $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^N$

Initialize $\Theta_0 = 0$

for $t = 1$ **to** T **do**

The learner chooses differentiable $\Phi_t : \mathbb{R}^N \rightarrow \mathbb{R}$ whose gradient satisfies $\text{Image}(\nabla\Phi_t) \subseteq \mathcal{X}$

The learner plays $w_t = \nabla\Phi_t(\Theta_{t-1})$

The adversary reveals $\theta_t \in \mathcal{Y}$ and the learner gets a reward of $\langle w_t, \theta_t \rangle$

Update $\Theta_t = \Theta_{t-1} + \theta_t$

end

Lemma 2 (GBPA Regret) *Let Φ be the baseline potential function for an online linear optimization problem. The regret of the GBPA can be written as:*

$$\text{Regret} = \underbrace{\Phi(\Theta_T) - \Phi_T(\Theta_T)}_{\text{underestimation penalty}} + \sum_{t=1}^T \left(\underbrace{(\Phi_t(\Theta_{t-1}) - \Phi_{t-1}(\Theta_{t-1}))}_{\text{overestimation penalty}} + \underbrace{D_{\Phi_t}(\Theta_t, \Theta_{t-1})}_{\text{divergence penalty}} \right), \quad (2)$$

where $\Phi_0 \equiv \Phi$.

Proof See Appendix A.1. ■

In the existing FTPL analysis, the counterpart of the divergence penalty is $\langle w_{t+1} - w_t, \theta_t \rangle$, which is controlled by analyzing the probability that the noise would cause the two random variables w_{t+1} and w_t to differ. In our framework, w_t is the gradient of a function Φ_t of Θ , which means that if Φ_t is twice-differentiable, we can take the derivative of w_t with respect to Θ . This derivative is the Hessian matrix of Φ_t , which essentially controls $\langle w_t - w_{t-1} \rangle$ with the help of Lemma 1. Since we focus on the curvature property of *functions* as opposed to random vectors, our FTPL analysis involves less probabilistic analysis than [Devroye et al. \(2013\)](#) or [van Erven et al. \(2014\)](#) does.

We point out a couple of important facts about Lemma 2:

- (a) If $\Phi_1 \equiv \dots \equiv \Phi_T$, then the overestimation penalty sums up to $\Phi_1(0) - \Phi(0) = \Phi_T(0) - \Phi(0)$.
- (b) If Φ_t is β -strongly smooth with respect to $\|\cdot\|$, the divergence penalty at t is at most $\frac{\beta}{2}\|\theta_t\|^2$.

3.2. Smoothability of the Baseline Potential

Equation 2 shows that the regret of the GBPA can be broken into two parts. One source of regret is the Bregman divergence of Φ_t ; since θ_t is not known until playing w_t , the GBPA always ascends along the gradient that is one step behind. The adversary can exploit this and play θ_t to induce a large *gap* between $\Phi_t(\mathbf{x}_t)$ and the linear approximation of $\Phi_t(\Theta_t)$ around Θ_{t-1} . Of course, the learner can reduce this gap by choosing a *smooth* Φ_t whose gradient changes slowly. The learner,

however, cannot achieve low regret by choosing an arbitrarily smooth Φ_t , because the other source of regret is the difference between Φ_t and Φ . In short, the GBPA achieves low regret if the potential function Φ_t gives a favorable tradeoff between the two sources of regret. This tradeoff is captured by the following definition of *smoothability*.

Definition 3 (*Beck and Teboulle, 2012, Definition 2.1*) Let Φ be a closed proper convex function. A collection of functions $\{\hat{\Phi}_\eta : \eta \in \mathbb{R}\}$ is said to be an η -smoothing of a smoothable function Φ with smoothing parameters $(\alpha, \beta, \|\cdot\|)$, if for every $\eta > 0$

(i) There exists α_1 (underestimation bound) and α_2 (overestimation bound) such that

$$\sup_{\Theta \in \text{dom}(\Phi)} \Phi(\Theta) - \hat{\Phi}_\eta(\Theta) \leq \alpha_1 \eta \quad \text{and} \quad \sup_{\Theta \in \text{dom}(\Phi)} \hat{\Phi}_\eta(\Theta) - \Phi(\Theta) \leq \alpha_2 \eta$$

with $\alpha_1 + \alpha_2 = \alpha$.

(ii) $\hat{\Phi}_\eta$ is $\frac{\beta}{\eta}$ -strongly smooth with respect to $\|\cdot\|$.

We say α is the deviation parameter, and β is the smoothness parameter.

A straightforward application of Lemma 2 gives the following statement:

Corollary 4 Let Φ be the baseline potential for an online linear optimization problem. Suppose $\{\hat{\Phi}_\eta\}$ is an η -smoothing of Φ with parameters $(\alpha, \beta, \|\cdot\|)$. Then, the GBPA run with $\Phi_1 \equiv \dots \equiv \Phi_T \equiv \hat{\Phi}_\eta$ has regret at most

$$\text{Regret} \leq \alpha \eta + \frac{\beta}{2\eta} \sum_{t=1}^T \|\theta_t\|^2$$

In online linear optimization, we often consider the settings where the marginal reward vectors $\theta_1, \dots, \theta_t$ are constrained by a norm, i.e., $\|\theta_t\| \leq r$ for all t . In such settings, the regret grows in $O(\sqrt{r\alpha\beta T})$ for the optimal choice of α . The product $\alpha\beta$, therefore, is at the core of the GBPA regret analysis.

3.3. Algorithms

Follow the Leader (FTL) Consider the GBPA run with a fixed potential function $\Phi_t \equiv \Phi$ for $t = 1, \dots, T$, i.e., the learner chooses the baseline potential function every iteration. At iteration t , this algorithm plays $\nabla \Phi_t(\Theta_{t-1}) = \arg \max_w \langle w, \Theta_{t-1} \rangle$, which is equivalent to FTL (Cesa-Bianchi and Lugosi, 2006, Section 3.2). FTL suffers zero regret from the over- or underestimation penalty, but the divergence penalty grows linearly in T in the worst case, resulting in an $\Omega(T)$ regret.

Follow the Regularized Leader (FTRL) Consider the GBPA run with a *regularized potential*:

$$\forall t, \Phi_t(\Theta) = \mathcal{R}^*(\Theta) = \max_{w \in \mathcal{X}} \{\langle w, \Theta \rangle - \mathcal{R}(w)\} \quad (3)$$

where $\mathcal{R} : \mathcal{X} \rightarrow \mathbb{R}$ is a β -strongly convex function. At time t , this algorithm plays $\nabla \Phi_t(\Theta_{t-1}) = \arg \max_w \{\langle w, \Theta_{t-1} \rangle - \mathcal{R}(w)\}$, which is equivalent to FTRL. By the strong convexity-strong smoothness duality, Φ_t is $\frac{1}{\beta}$ -strongly smooth with respect to the dual norm $\|\cdot\|_*$. In Section 5, we give an alternative interpretation of FTRL as a deterministic smoothing technique called inf-conv smoothing.

Follow the Perturbed Leader (FTPL) Consider the GBPA run with a *stochastically smoothed potential*:

$$\forall t, \Phi_t(\Theta) = \tilde{\Phi}(\Theta; \eta, \mathcal{D}) \stackrel{\text{def}}{=} \mathbb{E}_{u \sim \mathcal{D}}[\Phi(\Theta + \eta u)] = \mathbb{E}_{u \sim \mathcal{D}} \left[\max_{w \in \mathcal{X}} \{ \langle w, \Theta + \eta u \rangle \} \right] \quad (4)$$

where \mathcal{D} is a smoothing distribution with the support \mathbb{R}^N and $\eta > 0$ is a *scaling parameter*. This technique of *stochastic smoothing* has been well-studied in the optimization literature for gradient-free optimization algorithms (Glasserman, 1991; Yousean et al., 2010) and accelerated gradient methods for non-smooth optimizations (Duchi et al., 2012). If the max expression inside the expectation has a unique maximizer with probability one, we can swap the expectation and gradient (Bertsekas, 1973, Proposition 2.2) to obtain

$$\nabla \Phi_t(\Theta_{t-1}) = \mathbb{E}_{u \sim \mathcal{D}} \left[\arg \max_{w \in \mathcal{X}} \{ \langle w, \Theta_{t-1} + \eta u \rangle \} \right]. \quad (5)$$

Each arg max expression is equivalent to the decision rule of FTPL (Hannan, 1957; Kalai and Vempala, 2005); the GBPA on a stochastically smoothed potential can thus be seen as playing the *expected action* of FTPL. Since the learner gets a linear reward in online linear optimization, the regret of the GBPA on a stochastically smoothed potential is equal to the *expected regret* of FTPL.

FTPL-FTRL Duality Our potential-based formulation of FTRL and FTPL reveals that a strongly convex regularizer defines a smooth potential function via duality, while adding perturbations is a *direct* smoothing operation on the baseline potential function. By the strong convexity-strong smoothness duality, if the stochastically smoothed potential function is $(1/\beta)$ -strongly smooth with respect to $\| \cdot \|_*$, then its Fenchel conjugate implicitly defines a regularizer that is β -strongly convex with respect to $\| \cdot \|$.

This connection via duality is a *bijection* in the special case where the decision set is one-dimensional. Previously it had been observed³ that the Hedge Algorithm (Freund and Schapire, 1997), which can be cast as FTRL with an entropic regularization $\mathcal{R}(w) = \sum_i w_i \log w_i$, is equivalent to FTPL with Gumbel-distributed noise. Hofbauer and Sandholm (2002, Section 2) gave a generalization of this fact to a much larger class of perturbations, although they focused on repeated game playing where the learner’s decision set \mathcal{X} is the probability simplex. The inverse mapping from FTPL to FTRL, however, does not appear to have been previously published.

Theorem 5 Consider the one-dimensional online linear optimization problem with $\mathcal{X} = \mathcal{Y} = [0, 1]$. Let $\mathcal{R} : \mathcal{X} \rightarrow \mathbb{R}$ be a strongly convex regularizer. Its Fenchel conjugate \mathcal{R}^* defines a valid CDF of a continuous distribution \mathcal{D} such that Equation 3 and Equation 4 are equal. Conversely, let $F_{\mathcal{D}}$ be a CDF of a continuous distribution \mathcal{D} with a finite expectation. If we define \mathcal{R} to be such that $\mathcal{R}(w) - \mathcal{R}(0) = - \int_0^w F_{\mathcal{D}}^{-1}(1 - z) dz$, then Equation 3 and Equation 4 are equal.

Proof In Appendix B.1. ■

3. Adam Kalai first described this result in personal communication and Warmuth (2009) expanded it into a short note available online. However, the result appears to be folklore in the area of probabilistic choice models, and it is mentioned briefly in Hofbauer and Sandholm (2002).

4. Online Linear Optimization via Gaussian Smoothing

Gaussian smoothing is a standard technique for smoothing a function. In computer vision applications, for example, image pixels are viewed as a function of the (x, y) -coordinates, and Gaussian smoothing is used to blur noises in the image. We first present basic results on Gaussian smoothing from the optimization literature.

Definition 6 (Gaussian smoothing) *Let $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}$ be a function. Then, we define its Gaussian smoothing, with a scaling parameter $\eta > 0$ and a covariance matrix Σ , as*

$$\tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, \Sigma)) = \mathbb{E}_{u \sim \mathcal{N}(0, \Sigma)} \Phi(\Theta + \eta u) = (2\pi)^{-\frac{N}{2}} \det(\Sigma)^{-\frac{1}{2}} \int_{\mathbb{R}^N} \Phi(\Theta + \eta u) e^{-\frac{1}{2}u^T \Sigma^{-1} u} du$$

In this section, when the smoothing parameters are clear from the context, we use a shorthand notation $\tilde{\Phi}$. An extremely useful property of Gaussian smoothing is that $\tilde{\Phi}$ is always twice-differentiable, even when Φ is not. The trick is to introduce a new variable $\tilde{\Theta} = \Theta + \eta u$. After substitutions, the variable Θ only appears in the exponent, which can be safely differentiated.

Lemma 7 (*Nesterov 2011, Lemma 2, and Bhatnagar 2007, Section 3*) *Let $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}$ be a function. For any positive η , $\tilde{\Phi}(\cdot; \eta, \mathcal{N}(0, \Sigma))$ is twice-differentiable and*

$$\nabla \tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, \Sigma)) = \frac{1}{\eta} \mathbb{E}_u [\Phi(\Theta + \eta u) \Sigma^{-1} u] \quad (6)$$

$$\nabla^2 \tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, \Sigma)) = \frac{1}{\eta^2} \mathbb{E}_u \left[\Phi(\Theta + \eta u) \left((\Sigma^{-1} u)(\Sigma^{-1} u)^T - \Sigma^{-1} \right) \right] \quad (7)$$

If $\Phi(\Theta + \eta u)$ is differentiable almost everywhere, then we can directly differentiate Equation 6 by swapping the expectation and gradient (Bertsekas, 1973, Proposition 2.2) and obtain an alternative expression for Hessian:

$$\nabla^2 \tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, \Sigma)) = \frac{1}{\eta} \mathbb{E}_u [\nabla \Phi(\Theta + \eta u) (\Sigma^{-1} u)^T]. \quad (8)$$

4.1. Experts Setting (ℓ_1 - ℓ_∞ case)

The experts setting is where $\mathcal{X} = \Delta^N \stackrel{\text{def}}{=} \{w \in \mathbb{R}^N : \sum_i w_i = 1, w_i \geq 0 \forall i\}$, and $\mathcal{Y} = \{\theta \in \mathbb{R}^N : \|\theta\|_\infty \leq 1\}$. The baseline potential function is $\Phi(\Theta) = \max_{w \in \mathcal{X}} \langle w, \Theta \rangle = \Theta_{i^*(\Theta)}$, where we define $i^*(z) := \min\{i : i \in \arg \max_j z_j\}$.

Our regret bound in Theorem 8 is data-dependent, and it is stronger than the previously known $O(\sqrt{T \log N})$ regret bounds of the algorithms that use similar perturbations. In the game theoretic analysis of Gaussian perturbations by Rakhlin et al. (2012), the algorithm uses the scaling parameter $\eta_t = \sqrt{T - t}$, which requires the knowledge of T and does not adapt to data. Devroye et al. (2013) proposed the Prediction by Random Walk (PRW) algorithm, which flips a fair coin every round and decides whether to add 1 to each coordinate. Due to the discrete nature of the algorithm, the analysis must assume the worst case where $\|\theta_t\|_* = 1$ for all t .

Theorem 8 *Let Φ be the baseline potential for the experts setting. The GBPA run with the Gaussian smoothing of Φ , i.e., $\Phi_t(\cdot) = \tilde{\Phi}(\cdot; \eta_t, \mathcal{N}(0, I))$ for all t has regret at most*

$$\text{Regret} \leq \sqrt{2 \log N} \left(\eta_T + \sum_{t=1}^T \frac{1}{\eta_t} \|\theta_t\|_\infty^2 \right). \quad (9)$$

If the algorithm selects $\eta_t = \sqrt{\sum_{s=1}^T \|\theta_s\|_\infty^2}$ for all t (with the help of hindsight), we have

$$\text{Regret} \leq 2\sqrt{2 \sum_{t=1}^T \|\theta_t\|_\infty^2 \log N}.$$

If the algorithm selects η_t adaptively according to $\eta_t = \sqrt{2(1 + \sum_{s=1}^{t-1} \|\theta_s\|_\infty^2)}$, we have

$$\text{Regret} \leq 4\sqrt{(1 + \sum_{t=1}^T \|\theta_t\|_\infty^2) \log N}.$$

Proof In order to apply Lemma 2, we need to upper bound (i) the overestimation and underestimation penalty, and (ii) the Bregman divergence. To bound (i), first note that due to convexity of Φ , the smoothed potential $\tilde{\Phi}$ is also convex and upper bounds the baseline potential. Hence, the underestimation penalty is at most 0, and when η_t is fixed for all t , it is straightforward to bound the overestimation penalty:

$$\Phi_T(0) - \Phi(0) \leq \mathbb{E}_{u \sim \mathcal{N}(0, I)}[\Phi(\eta_T u)] \leq \eta_T \sqrt{2 \log N}. \quad (10)$$

The first inequality is the triangle inequality. The second inequality is a well-known result and we included the proof in Appendix C.1 for completeness. For the adaptive η_t , we apply Lemma 10, which we prove at the end of this section, to get the same bound.

It now remains to bound the Bregman divergence. This is achieved in Lemma 9 where we upper bound $\sum_{i,j} |\nabla_{ij}^2 \Phi|$, which is an upper bound on $\max_{\theta: \|\theta\|_\infty=1} \theta^T (\nabla^2 \Phi) \theta$. The final step is to apply Lemma 1. ■

The proof of Theorem 8 shows that for the experts setting, the Gaussian smoothing is an η -smoothing with parameters $(O(\sqrt{\log N}), O(\sqrt{\log N}), \|\cdot\|)$. This is in contrast to the Hedge Algorithm (Freund and Schapire, 1997), which is an η -smoothing with parameters $(\log N, 1, \|\cdot\|)$ (See Section 5 for details). Interestingly, the two algorithms obtain the same optimal regret (up to constant factors) although they have different smoothing parameters.

Lemma 9 *Let Φ be the baseline potential for the experts setting. Let the Hessian matrix of the Gaussian-smoothed baseline potential be denoted H , i.e., $H = \nabla^2 \tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, I))$. Then,*

$$\sum_{i,j} |H_{ij}| \leq \frac{2\sqrt{2 \log N}}{\eta}.$$

Proof With probability one, $\Phi(\Theta + \eta u)$ is differentiable and from Lemma 7, we can write

$$H = \frac{1}{\eta} \mathbb{E}[\nabla \Phi(\Theta + \eta u) u^T] = \mathbb{E}[e_{i^*(\eta u + \Theta)} u^T],$$

where $e_i \in \mathbb{R}^N$ is the i -th standard basis vector.

First, we note that all off-diagonals of H are negative and all diagonal entries in H are positive. This is because the Hessian matrix is the covariance matrix between the probability that i -th coordinate is the maximum and the extra random Gaussian noise added to the j -th coordinate; for any positive number α , $u_j = \alpha$ and $u_j = -\alpha$ have the same probability, but the indicator for $i = i^*$ has a higher probability to be 1 when u_i is positive (hence $H_{ii} > 0$) and u_j is negative for $i \neq j$ (hence $H_{ij} < 0$ for $i \neq j$).

Second, the entries of H sum up to 0, as

$$\sum_{i,j} H_{ij} = \frac{1}{\eta} \mathbb{E} \left[\sum_j u_j \sum_i \mathbf{1}\{i = i^*(\Theta + u)\} \right] = \frac{1}{\eta} \mathbb{E} \left[\sum_j u_j \right] = 0.$$

Combining the two observations, we have

$$\sum_{i,j} |H_{ij}| = \sum_{i,j:H_{ij}>0} H_{ij} + \sum_{i,j:H_{ij}<0} -H_{ij} = 2 \sum_{i,j:H_{ij}>0} H_{ij} = 2\text{Tr}(H)$$

Finally, the trace is bounded as follows:

$$\begin{aligned} \text{Tr}(H) &= \frac{1}{\eta} \mathbb{E} \left[\sum_i u_i \mathbf{1}\{i = i^*(\Theta + u)\} \right] \leq \frac{1}{\eta} \mathbb{E} \left[(\max_k u_k) \sum_i \mathbf{1}\{i = i^*(\Theta + u)\} \right] \\ &= \frac{1}{\eta} \mathbb{E} [\max_k u_k] \leq \frac{1}{\eta} \sqrt{2 \log N}, \end{aligned}$$

where the final inequality is shown in Appendix C.1. Multiplying both sides by 2 completes the proof. \blacksquare

Time-Varying Scaling Parameters When the scaling parameter η_t changes every iteration, the overestimation penalty becomes a sum of T terms. The following lemma shows that using the sublinearity of the baseline potential, we can collapse them into one.

Lemma 10 *Let $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}$ be a sublinear function, and \mathcal{D} be a continuous distribution with the support \mathbb{R}^N . Let $\Phi_t(\Theta) = \check{\Phi}(\Theta; \eta_t, \mathcal{D})$ for $t = 0, \dots, T$ and choose η_t to be a non-decreasing sequence of non-negative numbers ($\eta_0 = 0$ so that $\Phi_0 = \Phi$). Then, the overestimation penalty in Equation 2 has the following upper bound:*

$$\sum_{t=1}^T \Phi_t(\Theta_{t-1}) - \Phi_{t-1}(\Theta_{t-1}) \leq \eta_T \mathbb{E}_{u \sim \mathcal{D}} [\Phi(u)].$$

Proof See Appendix C.2 \blacksquare

4.2. Online Linear Optimization over Euclidean Balls (ℓ_2 - ℓ_2 case)

The Euclidean balls setting is where $\mathcal{X} = \mathcal{Y} = \{x \in \mathbb{R}^N : \|x\|_2 \leq 1\}$. The baseline potential function is $\Phi(\Theta) = \max_{w \in \mathcal{X}} \langle w, \Theta \rangle = \|\Theta\|_2$.

Theorem 11 *Let Φ be the baseline potential for the Euclidean balls setting. The GBPA run with $\Phi_t(\cdot) = \tilde{\Phi}(\cdot; \eta, \mathcal{N}(0, I))$ for all t has regret at most*

$$\text{Regret} \leq \eta_T \sqrt{N} + \frac{1}{2\sqrt{N}} \sum_{t=1}^T \frac{1}{\eta_t} \|\theta_t\|_2^2. \quad (11)$$

If the algorithm selects $\eta_t = \sqrt{\sum_{s=1}^T \|\theta_s\|_2^2 / (2N)}$ for all t (with the help of hindsight), we have

$$\text{Regret} \leq \sqrt{2 \sum_{t=1}^T \|\theta_t\|_2^2}.$$

If the algorithm selects η_t adaptively according to $\eta_t = \sqrt{(1 + \sum_{s=1}^{t-1} \|\theta_s\|_2^2) / N}$, we have

$$\text{Regret} \leq 2\sqrt{1 + \sum_{t=1}^T \|\theta_t\|_2^2}$$

Proof The proof is mostly similar to that of Theorem 8. In order to apply Lemma 2, we need to upper bound (i) the overestimation and underestimation penalty, and (ii) the Bregman divergence.

The Gaussian smoothing always overestimates a convex function, so it suffices to bound the overestimation penalty. Furthermore, it suffices to consider the fixed η_t case due to Lemma 1. The overestimation penalty can be upper-bounded as follows:

$$\begin{aligned} \Phi_T(0) - \Phi(0) &= \mathbb{E}_{u \sim \mathcal{N}(0, I)} \|\Theta + \eta_T u\|_2 - \|\Theta\|_2 \leq \eta_T \mathbb{E}_{u \sim \mathcal{N}(0, I)} \|u\|_2 \\ &\leq \eta_T \sqrt{\mathbb{E}_{u \sim \mathcal{N}(0, I)} \|u\|_2^2} = \eta_T \sqrt{N} \end{aligned}$$

The first inequality is from the triangle inequality, and the second inequality is from the concavity of the square root.

For the divergence penalty, note that the upper bound on $\max_{v: \|\theta\|_2=1} \theta^T (\nabla^2 \tilde{\Phi}) \theta$ is exactly the maximum eigenvalue of the Hessian, which we bound in Lemma 12. The final step is to apply Lemma 1. \blacksquare

Lemma 12 *Let Φ be the baseline potential for the Euclidean balls setting. Then, for all $\Theta \in \mathbb{R}^N$ and $\eta > 0$, the Hessian matrix of the Gaussian smoothed potential satisfies*

$$\nabla^2 \tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, I)) \preceq \frac{1}{\eta\sqrt{N}} I.$$

Proof The Hessian of the Euclidean norm $\nabla^2 \Phi(\Theta) = \|\Theta\|_2^{-1} I - \|\Theta\|_2^{-3} \Theta \Theta^T$ diverges near $\Theta = 0$. Expectedly, the maximum curvature is at origin even after Gaussian smoothing (See Appendix C.3). So, it suffices to prove

$$\nabla^2 \Phi(0) = \mathbb{E}_{u \sim \mathcal{N}(0, I)} [\|u\|_2 (uu^T - I)] \preceq \sqrt{\frac{1}{N}} I,$$

where the Hessian expression is from Equation 8.

By symmetry, all off-diagonal elements of the Hessian are 0. Let $Y = \|u\|_2^2$, which is Chi-squared with N degrees of freedom. So,

$$\text{Tr}(\mathbb{E}[\|u\|_2 (uu^T - I)]) = \mathbb{E}[\text{Tr}(\|u\|_2 (uu^T - I))] = \mathbb{E}[\|u\|_2^3 - N\|u\|_2] = \mathbb{E}[Y^{\frac{3}{2}}] - N\mathbb{E}[Y^{\frac{1}{2}}]$$

Using the Chi-Squared moment formula (Harvey, 1965, p. 20), the above becomes:

$$\frac{2^{\frac{3}{2}}\Gamma(\frac{3}{2} + \frac{N}{2})}{\Gamma(\frac{N}{2})} - \frac{N2^{\frac{1}{2}}\Gamma(\frac{1}{2} + \frac{N}{2})}{\Gamma(\frac{N}{2})} = \frac{\sqrt{2}\Gamma(\frac{1}{2} + \frac{N}{2})}{\Gamma(\frac{N}{2})}. \quad (12)$$

From the log-convexity of the Gamma function,

$$\log \Gamma\left(\frac{1}{2} + \frac{N}{2}\right) \leq \frac{1}{2} \left(\log \Gamma\left(\frac{N}{2}\right) + \log \Gamma\left(\frac{N}{2} + 1\right) \right) = \log \Gamma\left(\frac{N}{2}\right) \sqrt{\frac{N}{2}}.$$

Exponentiating both sides, we obtain

$$\Gamma\left(\frac{1}{2} + \frac{N}{2}\right) \leq \Gamma\left(\frac{N}{2}\right) \sqrt{\frac{N}{2}},$$

which we apply to Equation 12 and get $\text{Tr}(\nabla^2\Phi(0)) \leq \sqrt{N}$. To complete the proof, note that by symmetry, each entry must have the same expected value, and hence it is bounded by $\sqrt{1/N}$. ■

4.3. General Bound

In this section, we will use a generic property of Gaussian smoothing to derive a regret bound that holds for any arbitrary online linear optimization problem.

Lemma 13 (Duchi et al., 2012, Lemma E.2) *Let Φ be a real-valued convex function on a closed domain which is a subset of \mathbb{R}^N . Suppose Φ is L -Lipschitz with respect to $\|\cdot\|_2$, and let $\hat{\Phi}_\eta$ be the Gaussian smoothing of Φ with the scaling parameter η and identity covariance. Then, $\{\hat{\Phi}_\eta\}$ is an η -smoothing of Φ with parameters $(L\sqrt{N}, L, \|\cdot\|_2)$.*

Consider an instance of online linear optimization with decision set \mathcal{X} and reward set \mathcal{Y} . The baseline potential function Φ is $\|\mathcal{X}\|_2$ -Lipschitz with respect to $\|\cdot\|_2$, where $\|\mathcal{X}\|_2 = \sup_{x \in \mathcal{X}} \|x\|_2$. From Lemma 13 and Corollary 4, it follows that

$$\text{Regret} \leq \eta\sqrt{N}\|\mathcal{X}\|_2 + \frac{\|\mathcal{X}\|_2}{2} \sum_{t=1}^T \|\theta_t\|_2^2,$$

which is $O(N^{\frac{1}{4}}\|\mathcal{X}\|_2\|\mathcal{Y}\|_2\sqrt{T})$ after tuning η . This regret bound, however, often gives a suboptimal dependence on the dimension N . For example, it gives $O(N^{\frac{3}{4}}T^{\frac{1}{2}})$ regret bound for the experts setting where $\|\mathcal{X}\|_2 = 1$ and $\|\mathcal{Y}\|_2 = \sqrt{N}$, and $O(\sqrt{NT})$ regret bound for the Euclidean balls setting where $\|\mathcal{X}\|_2 = \|\mathcal{Y}\|_2 = 1$.

4.4. Online Convex Optimization

In online convex optimization, the learner receives a sequence of convex functions f_t whose domain is \mathcal{X} and its subgradients are in the set \mathcal{Y} (Zinkevich, 2003). After the learner plays $w_t \in \mathcal{X}$, the reward function f_t is revealed. The learner gains $f_t(w_t)$ and observes $\nabla f_t(w_t)$, a subgradient of f_t at w_t .

A simple linearization argument shows that our regret bounds for online linear optimization generalize to online convex optimization. Let w^* be the optimal fixed point in hindsight. The true regret is upper bounded by the linearized regret, as $f_t(w^*) - f_t(w_t) \leq \langle w^* - w_t, \nabla f_t(w_t) \rangle$ for any subgradient $\nabla f_t(\cdot)$, and our analysis bounds the linearized regret. Unlike in the online linear optimization settings, however, the regret bound is valid only for the GBPA with smoothed potentials, which plays the *expected action* of FTPL.

5. Online Linear Optimization via Inf-conv Smoothing

Beck and Teboulle (2012) proposed inf-conv smoothing, which is an infimal convolution with a strongly smooth function. In this section, we will show that FTRL is equivalent to the GBPA run with the inf-conv smoothing of the baseline potential function.

Let $(\mathcal{X}, \|\cdot\|)$ be a normed vector space, and $(\mathcal{X}^*, \|\cdot\|_*)$ be its dual. Let $\Phi : \mathcal{X}^* \rightarrow \mathbb{R}$ be a closed proper convex function, and let \mathcal{S} be a β -strongly smooth function on \mathcal{X}^* with respect to $\|\cdot\|_*$. Then, the *inf-conv smoothing* of Φ with \mathcal{S} is defined as:

$$\Phi^{\text{ic}}(\Theta; \eta, \mathcal{S}) \stackrel{\text{def}}{=} \inf_{\Theta^* \in \mathcal{X}^*} \left\{ \Phi(\Theta^*) + \eta \mathcal{S} \left(\frac{\Theta - \Theta^*}{\eta} \right) \right\} = \max_{w \in \mathcal{X}} \left\{ \langle w, \Theta \rangle - \Phi^*(w) - \eta \mathcal{S}^*(w) \right\}. \quad (13)$$

The first expression with infimum is precisely the infimal convolution of $\Phi(\cdot)$ and $\eta \mathcal{S}(\frac{\cdot}{\eta})$, and the second expression with supremum is an equivalent dual formulation. The inf-conv smoothing $\Phi^{\text{ic}}(\Theta; \eta, \mathcal{S})$ is finite, and it is an η -smoothing of Φ (Definition 3) with smoothing parameters

$$\left(\max_{\Theta \in \mathcal{X}^*} \max_{w \in \partial \Phi(\Theta)} \mathcal{S}^*(w), \beta, \|\cdot\| \right). \quad (14)$$

where $\partial \Phi(\Theta)$ is a set of subgradients of Φ at Θ .

Connection to FTRL Consider an online linear optimization problem with decision set $\mathcal{X} \subseteq \mathbb{R}^N$. Then, the dual space \mathcal{X}^* is simply \mathbb{R}^N . Let \mathcal{R} be a β -strongly convex function on \mathcal{X} with respect to a norm $\|\cdot\|$. By the strong convexity-strong smoothness duality, \mathcal{R}^* is $\frac{1}{\beta}$ -strongly smooth. Consider the inf-conv smoothing of the baseline potential function Φ with \mathcal{R}^* , denoted $\Phi^{\text{ic}}(\Theta; \eta, \mathcal{R}^*)$. We will that the GBPA run with $\Phi^{\text{ic}}(\Theta; \eta, \mathcal{R}^*)$ is equivalent to FTRL with \mathcal{R} as the regularizer.

First, note that the baseline potential is the convex conjugate of the null regularizer, i.e., $\Phi^*(w) = 0$ for all $w \in \mathcal{X}$. The dual formulation of inf-conv smoothing (Equation 13) thus becomes

$$\Phi^{\text{ic}}(\Theta; \eta, \mathcal{S}) = \max_{w \in \mathcal{X}} \left\{ \langle w, \Theta \rangle - \eta \mathcal{R}(w) \right\},$$

which is identical to Equation 3 except that the above expression has an extra parameter η that controls the degree of smoothing. To simplify the deviation parameter in Equation 14, note that the subgradients of Φ always lie in \mathcal{X} because of duality. Hence, the two supremum expressions collapse to one supremum: $\max_{w \in \mathcal{X}} \mathcal{S}^*(w)$. Plugging the smoothing parameters into Corollary 4 gives the well-known FTRL regret bound as in Theorem 2.11 or 2.21 of Shalev-Shwartz (2012):

$$\text{Regret} \leq \eta \max_{w \in \mathcal{X}} \mathcal{S}^*(w) + \frac{\beta}{2\eta} \sum_{t=1}^T \|\theta_t\|^2.$$

Acknowledgments

CL and AT gratefully acknowledge the support of NSF under grant IIS-1319810. We thank the anonymous reviewers for their helpful suggestions. We would also like to thank Andre Wibisono for several very useful discussions and his help improving the manuscript. Finally we thank Elad Hazan for early support in developing the ideas herein.

References

- Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(2):12, 2013.
- Amir Beck and Marc Teboulle. Smoothing and first order methods: A unified framework. *SIAM Journal on Optimization*, 22(2):557–580, 2012.
- Dimitri P. Bertsekas. Stochastic optimization problems with nondifferentiable cost functionals. *Journal of Optimization Theory and Applications*, 12(2):218–231, 1973. ISSN 0022-3239. doi: 10.1007/BF00934819.
- Shalabh Bhatnagar. Adaptive newton-based multivariate smoothed functional algorithms for simulation optimization. *ACM Transactions on Modeling and Computer Simulation*, 2007.
- Chris M. Bishop. Training with noise is equivalent to tikhonov regularization. *Neural Computation*, 7(1):108–116, January 1995. ISSN 0899-7667.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006. ISBN 978-0-521-84108-5.
- Luc Devroye, Gábor Lugosi, and Gergely Neu. Prediction by random-walk perturbation. In *Proceedings of Conference on Learning Theory (COLT)*, 2013.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. In *Proceedings of Conference on Learning Theory (COLT)*, 2010.
- John Duchi, Peter L. Bartlett, and Martin J. Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012. doi: 10.1137/110831659.
- Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119 – 139, 1997. ISSN 0022-0000. doi: <http://dx.doi.org/10.1006/jcss.1997.1504>.
- Paul Glasserman. *Gradient Estimation Via Perturbation Analysis*. Kluwer international series in engineering and computer science: Discrete event dynamic systems. Springer, 1991. ISBN 9780792390954.
- James Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.
- John C. Harsanyi. Oddness of the number of equilibrium points: a new proof. *International Journal of Game Theory*, 2(1):235–250, 1973.
- James R. Harvey. Fractional moments of a quadratic form in noncentral normal random variables, April 1965.
- Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *ArXiv preprint*, arXiv:1207.0580, 2012.

- Josef Hofbauer and William H. Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294, 2002.
- Adam T. Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- H. Brendan McMahan. Follow-the-regularized-leader and mirror descent: Equivalence theorems and l_1 regularization. In *AISTATS*, pages 525–533, 2011.
- Ilya S. Molchanov. *Theory of random sets*. Probability and its applications. Springer, New York, 2005. ISBN 1-85233-892-X.
- Yurii Nesterov. Random gradient-free minimization of convex functions. *ECORE Discussion Paper*, 2011.
- Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Relax and randomize : From value to algorithms. In *Proceedings of Neural Information Processing Systems (NIPS)*, 2012.
- R.T. Rockafellar. *Convex Analysis*. Convex Analysis. Princeton University Press, 1997. ISBN 9780691015866.
- Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, February 2012. ISSN 1935-8237.
- Nati Srebro, Karthik Sridharan, and Ambuj Tewari. On the universality of online mirror descent. In *Proceedings of Neural Information Processing Systems (NIPS)*, pages 2645–2653, 2011.
- Tim van Erven, Wojciech Kotlowski, and Manfred K. Warmuth. Follow the leader with dropout perturbations. In *Proceedings of Conference on Learning Theory (COLT)*, 2014.
- Stefan Wager, Sida Wang, and Percy Liang. Dropout training as adaptive regularization. In *Proceedings of Neural Information Processing Systems (NIPS)*, 2013.
- Manfred Warmuth. A perturbation that makes “Follow the leader” equivalent to “Randomized weighted majority”. <http://classes.soe.ucsc.edu/cms290c/Spring09/lect/10/wmkalai-rewrite.pdf>, 2009.
- Farzad Yousean, Angelia Nedić, and Uday V. Shanbhag. Convex nondifferentiable stochastic optimization: A local randomized smoothing technique. In *Proceedings of American Control Conference (ACC)*, 2010, pages 4875–4880, June 2010.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *International Conference on Machine Learning (ICML)*, 2003.

Appendix A. Gradient-Based Prediction Algorithm

A.1. Proof of Lemma 2

Proof We note that since $\Phi_0(0) = 0$,

$$\begin{aligned}\Phi_T(\Theta_T) &= \sum_{t=1}^T \Phi_t(\Theta_t) - \Phi_{t-1}(\Theta_{t-1}) \\ &= \sum_{t=1}^T \left((\Phi_t(\Theta_t) - \Phi_t(\Theta_{t-1})) + (\Phi_t(\Theta_{t-1}) - \Phi_{t-1}(\Theta_{t-1})) \right)\end{aligned}$$

The first difference can be rewritten as:

$$\Phi_t(\Theta_t) - \Phi_t(\Theta_{t-1}) = \langle \nabla \Phi_t(\Theta_{t-1}), \Theta_t \rangle + D_{\Phi_t}(\Theta_t, \Theta_{t-1})$$

By combining the above two,

$$\begin{aligned}\text{Regret} &= \Phi(\Theta_T) - \sum_{t=1}^T \langle \nabla \Phi_t(\Theta_{t-1}), \Theta_t \rangle \\ &= \Phi(\Theta_T) - \Phi_T(\Theta_T) + \sum_{t=1}^T D_{\Phi_t}(\Theta_t, \Theta_{t-1}) + \Phi_t(\Theta_{t-1}) - \Phi_{t-1}(\Theta_{t-1})\end{aligned}$$

which completes the proof. ■

Appendix B. FTPL-FTRL Duality

B.1. Proof of Theorem 5

Proof Consider a one-dimensional online linear optimization prediction problem where the player chooses an action w_t from $\mathcal{X} = [0, 1]$ and the adversary chooses a reward θ_t from $\mathcal{Y} = [0, 1]$. This can be interpreted as a two-expert setting; the player's action $w_t \in \mathcal{X}$ is the probability of following the first expert and θ_t is the net excess reward of the first expert over the second. The baseline potential for this setting is $\Phi(\Theta) = \max_{w \in [0, 1]} w\Theta$.

Let us consider an instance of FTPL with a continuous distribution \mathcal{D} whose cumulative density function (cdf) is $F_{\mathcal{D}}$. Let $\tilde{\Phi}$ be the smoothed potential function (Equation 4) with distribution \mathcal{D} . Its derivative is

$$\tilde{\Phi}'(\Theta) = \mathbb{E}[\arg \max_{w \in \mathcal{K}} w(\Theta + u)] = \mathbb{P}[u > -\Theta] \quad (15)$$

because the maximizer is unique with probability 1. Notice, crucially, that the derivative $\tilde{\Phi}'(\Theta)$ is exactly the expected solution of our FTPL instance. Moreover, by differentiating it again, we see that the second derivative of $\tilde{\Phi}$ at Θ is exactly the pdf of \mathcal{D} evaluated at $(-\Theta)$.

We can now precisely define the mapping from FTPL to FTRL. Our goal is to find a convex regularization function \mathcal{R} such that $\mathbb{P}(u > -\Theta) = \arg \max_{w \in \mathcal{X}} (w\Theta - \mathcal{R}(w))$. Since this is a

one-dimensional convex optimization problem, we can differentiate for the solution. The characterization of \mathcal{R} is:

$$\mathcal{R}(w) - \mathcal{R}(0) = - \int_0^w F_{\mathcal{D}}^{-1}(1-z) dz. \quad (16)$$

Note that the cdf $F_{\mathcal{D}}(\cdot)$ is indeed invertible since it is a strictly increasing function.

The inverse mapping is just as straightforward. Given a regularization function \mathcal{R} well-defined over $[0, 1]$, we can always construct its Fenchel conjugate $\mathcal{R}^*(\Theta) = \sup_{w \in \mathcal{X}} \langle w, \Theta \rangle - \mathcal{R}(w)$. The derivative of \mathcal{R}^* is an increasing convex function, whose infimum is 0 at $\Theta = -\infty$ and supremum is 1 at $\Theta = +\infty$. Hence, \mathcal{R}^* defines a cdf, and an easy calculation shows that this perturbation distribution exactly reproduces FTRL corresponding to \mathcal{R} . \blacksquare

Appendix C. Gaussian smoothing

C.1. Proof of Equation 10

Let X_1, \dots, X_N be independent standard Gaussian random variables, and let $Z = \max_{i=1, \dots, N} X_i$. For any real number a , we have

$$\exp(a\mathbb{E}[Z]) \leq \mathbb{E} \exp(aZ) = \mathbb{E} \max_{i=1, \dots, N} \exp(aX_i) \leq \sum_{i=1}^N \mathbb{E}[\exp(aX_i)] = N \exp(a^2/2).$$

The first inequality is from the convexity of the exponential function, and the last equality is by the definition of the moment generating function of Gaussian random variables. Taking the natural logarithm of both sides and dividing by a gives

$$\mathbb{E}[Z] \leq \frac{\log N}{a} + \frac{a}{2}.$$

In particular, by choosing $a = \sqrt{2 \log N}$, we have $\mathbb{E}[Z] \leq \sqrt{2 \log N}$.

C.2. Proof of Lemma 10

Proof By the subadditivity (triangle inequality) of Φ ,

$$\tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, I)) - \tilde{\Phi}(\Theta; \eta', \mathcal{N}(0, I)) = \mathbb{E}_{u \sim \mathcal{N}(0, I)} [\Phi(\Theta + \eta u) - \Phi(\Theta + \eta' u)] \quad (17)$$

$$\leq \mathbb{E}_{u \sim \mathcal{N}(0, I)} [\Phi((\eta - \eta')u)] \quad (18)$$

and the statement follows from the positive homogeneity of Φ . \blacksquare

C.3. Proof that the origin is the worst case (Lemma 12)

Proof Let $\Phi(\Theta) = \|\Theta\|_2$ and η be a positive number. By continuity of eigenvectors, it suffices to show that the maximum eigenvalue of the Hessian matrix of the Gaussian smoothed potential $\tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, I))$ is decreasing in $\|\Theta\|$ for $\|\Theta\| > 0$.

By Lemma 7, the gradient can be written as follows:

$$\nabla \tilde{\Phi}(\Theta; \eta, \mathcal{N}(0, I)) = \frac{1}{\eta} \mathbb{E}_{u \sim \mathcal{N}(0, I)} [u \|\Theta + \eta u\|] \quad (19)$$

Let u_i be the i -th coordinate of the vector u . Since the standard normal distribution is spherically symmetric, we can rotate the random variable u such that its first coordinate u_1 is along the direction of Θ . After rotation, the gradient can be written as

$$\frac{1}{\eta} \mathbb{E}_{u \sim \mathcal{N}(0, I)} \left[u \sqrt{(\|\Theta\| + \eta u_1)^2 + \sum_{k=2}^N \eta^2 u_k^2} \right]$$

which is clearly independent of the coordinates of Θ . The pdf of standard Gaussian distribution has the same value at (u_1, u_2, \dots, u_n) and its sign-flipped pair $(u_1, -u_2, \dots, -u_n)$. Hence, in expectation, the two vectors cancel out every coordinate but the first, which is along the direction of Θ . Therefore, there exists a function α such that $\mathbb{E}_{u \sim \mathcal{N}(0, I)} [u \|\Theta + \eta u\|] = \alpha(\|\Theta\|)\Theta$.

Now, we will show that α is decreasing in $\|\Theta\|$. Due to symmetry, it suffices to consider $\Theta = t e_1$ for $t \in \mathbb{R}^+$, without loss of generality. For any $t > 0$,

$$\begin{aligned} \alpha(t) &= \mathbb{E}[u_1 \sqrt{(t + \eta u_1)^2 + u_{\text{rest}}^2}] / t \\ &= \mathbb{E}_{u_{\text{rest}}} [\mathbb{E}_{u_1} [u_1 \sqrt{(t + \eta u_1)^2 + b^2} | u_{\text{rest}} = b]] / t \\ &= \mathbb{E}_{u_{\text{rest}}} [\mathbb{E}_{a=\eta|u_1|} [a (\sqrt{(t+a)^2 + b^2} - \sqrt{(t-a)^2 + B}) | u_{\text{rest}} = b]] / t \end{aligned}$$

Let $g(t) = (\sqrt{(t+a)^2 + B} - \sqrt{(t-a)^2 + B}) / t$. Take the first derivative with respect to t , and we have:

$$\begin{aligned} g'(t) &= \frac{1}{t^2} \left(\sqrt{(t-a)^2 + b^2} - \frac{t(t-a)}{\sqrt{(t+a)^2 + b^2}} - \sqrt{(t+a)^2 + b^2} + \frac{t(t-a)}{\sqrt{(t+a)^2 + b^2}} \right) \\ &= \frac{1}{t^2} \left(\frac{a^2 + b^2 - at}{\sqrt{(t-a)^2 + b^2}} - \frac{a^2 + b^2 + at}{\sqrt{(t+a)^2 + b^2}} \right) \end{aligned}$$

$$\left((a^2 + b^2) - at \right)^2 \left((t+a)^2 + b^2 \right) - \left((a^2 + b^2) + at \right)^2 \left((t-a)^2 + b^2 \right) = -4ab^2t^3 < 0$$

because t, η, u', B are all positive. So, $g(t) < 0$, which proves that α is decreasing in Θ .

The final step is to write the gradient as $\nabla(\tilde{\Phi}; \eta, \mathcal{N}(0, I))(\Theta) = \alpha(\|\Theta\|)\Theta$ and differentiate it:

$$\nabla^2 f_\eta(\Theta) = \frac{\alpha'(\|\Theta\|)}{\|\Theta\|} \Theta \Theta^T + \alpha(\|\Theta\|) I$$

The Hessian has two distinct eigenvalues $\alpha(\|\Theta\|)$ and $\alpha(\|\Theta\|) + \alpha'(\|\Theta\|)\|\Theta\|$, which correspond to the eigenspace orthogonal to Θ and parallel to Θ , respectively. Since α' is negative, α is always the maximum eigenvalue and it decreases in $\|\Theta\|$. \blacksquare