

Unconstrained Online Linear Learning in Hilbert Spaces: Minimax Algorithms and Normal Approximations

H. Brendan McMahan

Google, Seattle, WA

MCMAHAN@GOOGLE.COM

Francesco Orabona

Toyota Technological Institute at Chicago, Chicago, IL

FRANCESCO@ORABONA.COM

Abstract

We study algorithms for online linear optimization in Hilbert spaces, focusing on the case where the player is unconstrained. We develop a novel characterization of a large class of minimax algorithms, recovering, and even improving, several previous results as immediate corollaries. Moreover, using our tools, we develop an algorithm that provides a regret bound of $\mathcal{O}\left(U\sqrt{T\log(U\sqrt{T}\log^2 T + 1)}\right)$, where U is the L_2 norm of an arbitrary comparator and both T and U are unknown to the player. This bound is optimal up to $\sqrt{\log\log T}$ terms. When T is known, we derive an algorithm with an optimal regret bound (up to constant factors). For both the known and unknown T case, a Normal approximation to the conditional value of the game proves to be the key analysis tool.

Keywords: Online learning, minimax analysis, online convex optimization

1. Introduction

The online learning framework provides a scalable and flexible approach for modeling a wide range of prediction problems, including classification, regression, ranking, and portfolio management. Online algorithms work in rounds, where at each round a new instance is given and the algorithm makes a prediction. Then the environment reveals the label of the instance, and the learning algorithm updates its internal hypothesis. The aim of the learner is to minimize the cumulative loss it suffers due to its prediction error.

Research in this area has mainly focused on designing new prediction strategies and proving theoretical guarantees for them. However, recently, minimax analysis has been proposed as a general tool to design optimal prediction strategies (Rakhlin et al., 2012, 2013; McMahan and Abernethy, 2013). The problem is cast as a sequential multi-stage zero-sum game between the player (the learner) and an adversary (the environment), providing the *optimal* strategies for both. In some cases the value of the game can be calculated exactly in an efficient way (Abernethy et al., 2008a), in others upper bounds on the value of the game (often based on the sequential Rademacher complexity) are used to construct efficient algorithms with theoretical guarantees (Rakhlin et al., 2012).

While most of the work in this area has focused on the setting where the player is constrained to a bounded convex set (Abernethy et al., 2008a) (with the notable exception of McMahan and Abernethy (2013)), in this work we are interested in the general setting of unconstrained online learning with linear losses in Hilbert spaces. In Section 4, extending the work of McMahan and Abernethy (2013), we provide novel and general sufficient conditions to be able to compute the exact minimax strategy for both the player and the adversary, as well as the value of the game. In particular, we

show that under these conditions the optimal play of the adversary is *always orthogonal or always parallel to the sum of his previous plays*, while *the optimal play of the player is always parallel*. On the other hand, for some cases where the exact minimax strategy is hard to characterize, we introduce a new relaxation procedure based on a Normal approximation. In the particular application of interest, we show the relaxation is strong enough to yield an optimal regret bound, up to constant factors.

In Section 5, we use our new tools to recover and extend previous results on minimax strategies for linear online learning, including results for bounded domains. In fact, we show how to obtain a family of minimax strategies that smoothly interpolates between the minimax algorithm for a bounded feasible set and a minimax optimal algorithm in fact equivalent to unconstrained gradient descent. We emphasize that *all* the algorithms from this family are exactly minimax optimal,¹ in a sense we will make precise in the next section. Moreover, if you are allowed to play outside of the comparator set, we show that some members of this family have a non-vacuous regret bound for the unconstrained setting, while remaining optimal for the constrained one.

When studying unconstrained problems, a natural question is how small we can make the dependence of the regret bound on U , the L_2 norm of an arbitrary comparator point, while still maintaining a \sqrt{T} dependency on the time horizon. The best algorithm from the above family achieves $\text{Regret}(U) \leq \frac{1}{2}(U^2 + 1)\sqrt{T}$. Streeter and McMahan (2012) and Orabona (2013) show it is possible to reduce the dependence on U to $\mathcal{O}(U \log UT)$. In order to improve on this, in Section 6 we apply our techniques to analyze a strategy, based on a Normal potential function, that gives a regret bound of $\mathcal{O}\left(U\sqrt{T \log(U\sqrt{T} \log^2 T + 1)}\right)$ where U is the L_2 norm of a comparator, and both T and U are unknown. This bound is optimal up to $\sqrt{\log \log T}$ terms. Moreover, when T is known, we propose an algorithm based on a similar potential function that is optimal up to constant terms. This solves the open problem posed in those papers, matching the lower bound for this problem. Table 1 summarizes the regret bounds we prove, along with those for related algorithms.

Our analysis tools for both known- T and unknown horizon algorithms rest heavily on the relationship between the reward (negative loss) achieved by the algorithm, potential functions that provide a benchmark for the amount of reward the algorithm should have, the regret of the algorithm with respect to a post-hoc comparator u , and the conditional value of the game. These are familiar concepts from the literature, but we summarize these relationships and provide some modest generalizations in Section 3.

2. Notation and Problem Formulation

Let \mathcal{H} be a Hilbert space with inner product $\langle \cdot, \cdot \rangle$. The associated norm is denoted by $\|\cdot\|$, i.e. $\|x\| = \sqrt{\langle x, x \rangle}$. Given a closed and convex function f with domain $S \subseteq \mathcal{H}$, we will denote its Fenchel conjugate by $f^* : \mathcal{H} \rightarrow \mathbb{R}$ where $f^*(u) = \sup_{v \in S} (\langle v, u \rangle - f(v))$.

We consider a version of online linear optimization, a standard game for studying repeated decision making. On each of a sequence of rounds, a *player* chooses an action $w_t \in \mathcal{H}$, an *adversary* chooses a linear cost function $g_t \in \mathcal{G} \subseteq \mathcal{H}$, and the player suffers loss $\langle w_t, g_t \rangle$. For any sequence of

1. In this work, we use the term “minimax” to refer to the exact minimax solution to the zero sum game, as opposed to algorithms that only achieve the minimax optimal rate up to say constant factors.

Regret bounds for known- T algorithms			
(A)	Minimax Regret	\sqrt{T} for $u \in \mathcal{W}$, $\mathcal{O}(T)$ otherwise	Abernethy et al. (2008a)
(B)	OGD, fixed η	$\frac{1}{2}(1 + U^2)\sqrt{T}$	E.g., Shalev-Shwartz (2012)
(C)	pq -Algorithm	$(\frac{1}{p} + \frac{1}{q}U^q)\sqrt{T}$	Cor. 9, which also covers (A) and (B)
(D)	Reward Doubling	$\mathcal{O}(U\sqrt{T} \log(d(U+1)T))$	Streeter and McMahan (2012)
(E)	Normal Potential, $\epsilon = 1$	$\mathcal{O}\left(U\sqrt{T \log(UT+1)}\right)$	Theorem 11
(F)	Normal Potential, $\epsilon = \sqrt{T}$	$\mathcal{O}\left((U+1)\sqrt{T \log(U+1)}\right)$	Theorem 11
Regret bounds for adaptive algorithms for unknown T			
(G)	Adaptive FTRL/RDA	$(1 + \frac{1}{2}U^2)\sqrt{T}$	Shalev-Shwartz (2007); Xiao (2009)
(H)	Dim. Free Exp. Grad.	$\mathcal{O}(U\sqrt{T} \log(UT+1))$	Orabona (2013)
(I)	AdaptiveNormal	$\mathcal{O}^*\left(U\sqrt{T \log(UT+1)}\right)$	Theorem 12

Table 1: Here $U = \|u\|$ is the norm of a comparator, with U unknown to the algorithm. We let $\mathcal{W} = \{w : \|w\| \leq 1\}$; the adversary plays gradients with $\|g_t\| \leq 1$. (A) is minimax optimal for regret against points in \mathcal{W} , and always plays points from \mathcal{W} . The other algorithms are unconstrained. Even though (A) is minimax optimal for regret, other algorithms (e.g. (B)) offer strictly better bounds for arbitrary U . (C) corresponds to a family of minimax optimal algorithms where $\frac{1}{p} + \frac{1}{q} = 1$; $p = 2$ yields (B) and as $p \rightarrow 1$ the algorithm becomes (A); Corollary 9 covers (A) exactly. Only (D) has a dependence on d , the dimension of \mathcal{H} . The \mathcal{O}^* in (I) hides an additional $\log^2(T+1)$ term inside the log.

plays w_1, \dots, w_T and g_1, \dots, g_T , we define the *regret* against a comparator u in the standard way:

$$\text{Regret}(u) \equiv \sum_{t=1}^T \langle g_t, w_t - u \rangle.$$

This setting is general enough to cover the cases of online learning in, for example, \mathbb{R}^d , in the vector space of matrices, and in a RKHS. We also define the *reward* of the algorithm, which is the earnings (or negative losses) of the player throughout the game:

$$\text{Reward} \equiv \sum_{t=1}^T \langle -g_t, w_t \rangle.$$

We write $\theta_t \equiv -g_{1:t}$, where we use the compressed summation notation $g_{1:t} \equiv \sum_{s=1}^t g_s$.

The Minimax View It will be useful to consider a full game-theoretic characterization of the above interaction when the number of rounds T is known to both players. This approach that has received significant recent interest (Abernethy et al., 2008a, 2007; Abernethy and Warmuth, 2010; Abernethy et al., 2008b; Streeter and McMahan, 2012).

In the constrained setting, where the comparator vector $u \in \mathcal{W}$, we have that the *value of the game*, that is the regret when both the player and the adversary play optimally, is

$$\begin{aligned} V &\equiv \min_{w_1 \in \mathcal{H}} \max_{g_1 \in \mathcal{G}} \cdots \min_{w_T \in \mathcal{H}} \max_{g_T \in \mathcal{G}} \left(\sup_{u \in \mathcal{W}} \sum_{t=1}^T \langle w_t - u, g_t \rangle \right) \\ &= \min_{w_1 \in \mathcal{H}} \max_{g_1 \in \mathcal{G}} \cdots \min_{w_T \in \mathcal{H}} \max_{g_T \in \mathcal{G}} \left(\sum_{t=1}^T \langle w_t, g_t \rangle + \sup_{u \in \mathcal{W}} \langle u, \theta_T \rangle \right) \\ &= \min_{w_1 \in \mathcal{H}} \max_{g_1 \in \mathcal{G}} \cdots \min_{w_T \in \mathcal{H}} \max_{g_T \in \mathcal{G}} \left(\sum_{t=1}^T \langle w_t, g_t \rangle + B(\theta_T) \right), \end{aligned}$$

where

$$B(\theta) = \sup_{w \in \mathcal{W}} \langle w, \theta \rangle. \quad (1)$$

Following McMahan and Abernethy (2013), we generalize the game in terms of a generic convex *benchmark function* $B : \mathcal{H} \rightarrow \mathbb{R}$, instead of using the definition (1). This allows us to analyze the constrained and unconstrained setting in a unified way. Hence, the value of the game is the difference between the benchmark reward $B(\theta_T)$ and the actual reward achieved by the player (under optimal play by both parties). Intuitively, viewing the units of loss/reward as dollars, V is the amount of starting capital we need (equivalently, the amount we need to borrow) to ensure we end the game with $B(\theta_T)$ dollars. The motivation for defining the game in terms of an arbitrary B is made clear in the next section: It will allow us to derive Regret bounds in terms of the Fenchel conjugate of B .

We define inductively the *conditional value of the game* after g_1, \dots, g_t have been played by

$$V_t(\theta_t) = \min_{w \in \mathcal{H}} \max_{g \in \mathcal{G}} (\langle g, w \rangle + V_{t+1}(\theta_t - g)) \quad \text{with} \quad V_T(\theta_T) = B(\theta_T).$$

Thus, we can view the notation V for the value of the game as shorthand for $V_0(\mathbf{0})$. Under minimax play by both players, unrolling the previous equality, we have $\sum_{s=1}^t \langle g_s, w_s \rangle + V_t(-g_{1:t}) = V$, or for $t = T$,

$$\text{Reward} = \sum_{t=1}^T \langle -g_t, w_t \rangle = B(\theta_T) - V. \quad (2)$$

We also have that, given the conditional value of the game, a minimax-optimal strategy is

$$w_{t+1} = \arg \min_w \max_{g \in \mathcal{G}} \langle g, w \rangle + V_{t+1}(\theta_t - g). \quad (3)$$

McMahan and Abernethy (2013, Cor. 2) showed that in the unconstrained case, V_t is a smoothed version of B , where the smoothing comes from an expectation over future plays of the adversary. In this work, we show that in some cases (Theorem 4) we can find a closed form for V_t in terms of B , and in fact the solution to (3) will simply be the gradient of V_t , or equivalently, an FTRL algorithm with regularizer V_t^* . On the other hand, to derive our main results, we face a case (Theorem 6) where V_t is generally *not* expressible in closed form, and the resulting algorithm does *not* look like FTRL. We solve the first problem by using a Normal approximation to the adversary's future moves, and we solve the second by showing (3) can still be solved in closed form with respect to this approximation to V_t .

3. Potential Functions and the Duality of Reward and Regret

In the present section we will review some existing results in online learning theory as well as provide a number of mild generalizations for our purposes. *Potential functions* play a major role in the design and analysis of online learning algorithms (Cesa-Bianchi and Lugosi, 2006). We will use $q : \mathcal{H} \rightarrow \mathbb{R}$ to describe the potential, and the key assumptions are that q should depend solely on the cumulative gradients $g_{1:T}$ and that q is convex in this argument.² Since our aim is adaptive algorithms, we often look at a sequence of changing potential functions q_1, \dots, q_T , each of which takes as argument $-g_{1:t}$ and is convex. These functions have appeared with different interpretations in many papers, with different emphasis. They can be viewed as 1) the conjugate of an (implicit) time-varying regularizer in a Mirror Descent or Follow-the-Regularized-Leader (FTRL) algorithm (Cesa-Bianchi and Lugosi, 2006; Shalev-Shwartz, 2007; Rakhlin, 2009), 2) as proxy for the conditional value of the game in a minimax setting (Rakhlin et al., 2012), or 3) a potential giving a bound on the amount of reward we want the algorithm to have obtained at the end of round t (Streeter and McMahan, 2012; McMahan and Abernethy, 2013).

These views are of course closely connected, but can lead to somewhat different analysis techniques. Following the last view, suppose we interpret $q_t(\theta_t)$ as the desired reward at the end of round t , given the adversary has played $\theta_t = -g_{1:t}$ so far. Then, if we can bound our actual final reward in terms of $q_T(\theta_T)$, we also immediately get a regret bound stated in terms of the Fenchel conjugate q_T^* . Generalizing Streeter and McMahan (2012, Thm. 1), we have the following result (all omitted proofs can be found in the Appendix).

Theorem 1 *Let $\Psi : \mathcal{H} \rightarrow \mathbb{R}$ be a convex function. An algorithm for the player guarantees*

$$\text{Reward} \geq \Psi(-g_{1:T}) - \hat{\epsilon} \quad \text{for any } g_1, \dots, g_T \quad (4)$$

for a constant $\hat{\epsilon} \in \mathbb{R}$ if and only if it guarantees

$$\text{Regret}(u) \leq \Psi^*(u) + \hat{\epsilon} \quad \text{for all } u \in \mathcal{H}. \quad (5)$$

First we consider the minimax setting, where we define the game in terms of a convex benchmark B . Then, (2) gives us an immediate lower bound on the reward of the minimax strategy for the player (against any adversary), and so applying Theorem 1 with $\Psi = B$ gives

$$\forall u \in \mathcal{H}, \quad \text{Regret}(u) \leq B^*(u) + V. \quad (6)$$

The fundamental point, of which we will make much use, is this: even if one only cares about the traditional definition of regret, the study of the minimax game defined in terms of a general comparator benchmark B may be interesting, as the minimax algorithm for the player may then give novel bounds on regret. Note when B is defined as in (1), the theorem implies $\forall u \in \mathcal{W}$, $\text{Regret}(u) \leq V$. More generally, even for non-minimax algorithms, Theorem 1 states that understanding the reward (equivalently, loss) of an algorithm as a function of the sum of gradients chosen by the adversary is both necessary and sufficient for understanding the regret of the algorithm.

Now we consider the potential function view. The following general bound for any sequence of plays w_t against gradients g_t , for an arbitrary sequence of potential functions q_t , has been used

2. It is sometimes possible to generalize to potentials $q_t(g_1, \dots, g_t)$ that are functions of each gradient individually.

numerous times (see [Orabona \(2013, Lemma 1\)](#) and references therein). The claim is that

$$\text{Regret}(u) \leq q_T^*(u) + \sum_{t=1}^T (q_t(\theta_t) - q_{t-1}(\theta_{t-1}) + \langle w_t, g_t \rangle), \quad (7)$$

where we take $\theta_0 = \mathbf{0}$, and assume $q_0(\mathbf{0}) = 0$. In fact, this statement is essentially equivalent to the argument of (4) and (5). For intuition, we can view $q_t(\theta_t)$ as the amount of money we wish to have available at the end of round t . Suppose at the end of each round t , we borrow an additional sum ϵ_t as needed to ensure we actually have $q_t(\theta_t)$ on hand. Then, based on this invariant, the amount of reward we actually have after playing on round t is $q_{t-1}(\theta_{t-1}) + \langle w_t, -g_t \rangle$, the money we had at the beginning of the round, plus the reward we get for playing w_t . Thus, the additional amount we need to borrow at the end of round t in order to maintain the invariant is exactly

$$\epsilon_t(\theta_{t-1}, g_t) \equiv \underbrace{q_t(\theta_t)}_{\text{Reward desired}} - \underbrace{(q_{t-1}(\theta_{t-1}) + \langle w_t, -g_t \rangle)}_{\text{Reward achieved}}, \quad (8)$$

recalling $\theta_t = \theta_{t-1} - g_t$. Thus, if we can find bounds $\hat{\epsilon}_t$ such that for all t , θ_{t-1} , and $g \in \mathcal{G}$,

$$\hat{\epsilon}_t \geq \epsilon_t(\theta_{t-1}, g_t) \quad (9)$$

we can re-state (7) as exactly (5) with $\Psi = q_T$ and $\hat{\epsilon} = \hat{\epsilon}_{1:T}$. Further, solving (8) for the per-round reward $\langle w_t, -g_t \rangle$, summing from $t = 1$ to T and canceling telescoping terms gives exactly (4). Not surprisingly, both Theorem 1 and (7) can be proved in terms of the Fenchel-Young inequality.

When T is known, and the q_t are chosen carefully, it is possible to obtain $\hat{\epsilon}_t = 0$. On the other hand, when T is unknown to the players, typically we will need bounds $\hat{\epsilon}_t > 0$. For example, in both [Streeter and McMahan \(2012, Thm. 6\)](#) and [Orabona \(2013\)](#), the key is showing the sum of these $\hat{\epsilon}_t$ terms is always bounded by a constant. For completeness, we also state standard results where we interpret q_t^* as a regularizer.

The conjugate regularizer and Bregman divergences The updates of many algorithms are based on a time-varying version of the FTRL strategy,

$$w_{t+1} = \nabla q_t(\theta_t) = \arg \min_w \langle g_{1:t}, w \rangle + q_t^*(w), \quad (10)$$

where we view q_t^* as a time-varying regularizer (see [Orabona et al. \(2013\)](#) and references therein). Regret bounds can be easily obtained using (7) when the regularizers $q_t^*(w)$ are increasing with t , and they are strongly convex w.r.t. a norm $\|\cdot\|_*$, using the fact that the potential functions q_t will be strongly smooth. Then strong smoothness and particular choice of w_t implies

$$q_{t-1}(\theta_t) \leq q_{t-1}(\theta_{t-1}) - \langle w_t, g_t \rangle + \frac{1}{2} \|g_t\|^2, \quad (11)$$

which leads to the bound

$$\epsilon_t(\theta_t, g_t) = q_t(\theta_t) - q_{t-1}(\theta_{t-1}) + \langle w_t, g_t \rangle \leq q_t(\theta_t) - q_{t-1}(\theta_t) + \frac{1}{2} \|g_t\|^2 \leq \frac{1}{2} \|g_t\|^2,$$

where the last inequality follows from the fact that if $f(x) \leq g(x)$, then $f^*(y) \geq g^*(y)$ (immediate from the definition of the conjugate).

When the regularizer q^* is fixed, that is, $q_t = q$ for all t for some convex function q , we get the approach pioneered by [Grove et al. \(2001\)](#) and [Kivinen and Warmuth \(2001\)](#):

$$\epsilon_t(\theta_t, g_t) = q(\theta_t) - q(\theta_{t-1}) + \langle w_t, g_t \rangle = q(\theta_t) - (q(\theta_{t-1}) + \langle \nabla q(\theta_{t-1}), g_t \rangle) = D_q(\theta_t, \theta_{t-1}),$$

where D_q is the Bregman Divergence with respect to q , and we predict with $w_t = \nabla q(\theta_{t-1})$.

Admissible relaxations and potentials We extend the notion of relaxations of the conditional value of the game of [Rakhlin et al. \(2012\)](#) to the present setting. We say v_t with corresponding strategy w_t is a relaxation of V_t if

$$\forall \theta, \quad v_T(\theta) \geq B(\theta) \quad \text{and} \quad (12)$$

$$\forall t \in \{0, \dots, T-1\}, g \in \mathcal{G}, \theta \in \mathcal{H}, \quad v_t(\theta) + \hat{\epsilon}_{t+1} \geq \langle g, w_{t+1} \rangle + v_{t+1}(\theta - g), \quad (13)$$

for constants $\hat{\epsilon}_t \geq 0$. This definition matches Eq. (4) of [Rakhlin et al. \(2012\)](#) if we force all $\hat{\epsilon}_t = 0$, but if we allow some slack $\hat{\epsilon}_t$, (13) corresponds exactly to (8) and (9).

Note that (13) is invariant to adding a constant to all v_t . In particular, given an admissible v_t , we can define $q_t(\theta) = v_t(\theta) - v_0(\mathbf{0})$ so $q_t(\mathbf{0}) = 0$ and q satisfies (9) with the same $\hat{\epsilon}_t$ values for which v_t satisfies (13). Or we could define $q_0(\mathbf{0}) = 0$ and $q_t(\theta) = v_t(\theta)$ for $t \geq 1$, and take $\hat{\epsilon}_1 \leftarrow \hat{\epsilon}_1 + v_0(\mathbf{0})$ (or any other way of distributing the $v_0(\mathbf{0})$ into the $\hat{\epsilon}$). Generally, when T is known we will find working with admissible relaxations v_t to be most useful, while for unknown horizons T , potential functions with $q_0(\mathbf{0}) = 0$ will be more natural.

For our admissible relaxations, we have a result that closely mirrors [Theorem 1](#):

Corollary 2 *Let v_0, \dots, v_T be an admissible relaxation for a benchmark B . Then, for any sequence g_1, \dots, g_T , for any w_t chosen so (13) and (12) are satisfied, we have*

$$\text{Reward} \geq B(\theta_T) - v_0(0) - \hat{\epsilon}_{1:T} \quad \text{and} \quad \text{Regret}(u) \leq B^*(u) + v_0(0) + \hat{\epsilon}_{1:T}.$$

Proof For the first statement, re-arranging and summing (13) shows $\text{Reward}_t \geq v_t(\theta_t) - \hat{\epsilon}_{1:t} - v_0(0)$ and so final $\text{Reward} \geq B(\theta) - v_0(0) - \hat{\epsilon}_{1:T}$; the second result then follows from [Theorem 1](#). ■

The regret bound corresponds to (6); in particular, if we take v_t to be the conditional value of the game, then (12) and (13) hold with equality with all $\hat{\epsilon}_t = 0$. Note if we define B as in (1), the regret guarantee becomes $\forall u \in \mathcal{W}, \text{Regret}(u) \leq v_0(0) + \hat{\epsilon}_{1:T}$, analogous to ([Rakhlin et al., 2012, Prop. 1](#)) when $\hat{\epsilon}_{1:T} = 0$.

Deriving algorithms Consider an admissible relaxation v_t . Given the form of the regret bounds we have proved, a natural strategy is to choose w_{t+1} so as to minimize $\hat{\epsilon}_{t+1}$, that is,

$$w_{t+1} = \arg \min_w \max_{g \in \mathcal{G}} v_{t+1}(\theta_t - g) - v_t(\theta_t) + \langle g, w \rangle = \arg \min_w \max_{g \in \mathcal{G}} \langle g, w \rangle + v_{t+1}(\theta_t - g), \quad (14)$$

following [Rakhlin et al. \(2012, Eq. \(5\)\)](#), [Rakhlin et al. \(2013\)](#), and [Streeter and McMahan \(2012, Eq. \(8\)\)](#). We see that v_{t+1} is standing in for the conditional value of the game in (3). Since additive constants do not impact the argmin, we could also replace v_t with a potential q_t , say $q_t(\theta) = v_t(\theta) - v_0(0)$.

4. Minimax Analysis Approaches for Known-Horizon Games

In general, the problem of calculating the conditional value of a game $V_t(\theta)$ is hard. And even for a known potential, deriving an optimal solution via (14) is also in general a hard problem. When the player is unconstrained, we can simplify the computation of V_t and the derivation of optimal strategies. For example, following ideas from [McMahan and Abernethy \(2013\)](#),

$$\epsilon_t(\theta_t) = \max_{p \in \Delta(\mathcal{G}), \mathbb{E}_{g \sim p}[g] = 0} \mathbb{E}_{g \sim p} [q_{t+1}(\theta_t - g)] - q_t(\theta_t),$$

where $\Delta(\mathcal{G})$ is the set of probability distributions on \mathcal{G} . McMahan and Abernethy (2013) shows that in some cases is possible to easily calculate this maximum, in particular when $\mathcal{G} = [-G, G]^d$ and q_t decomposes on a per-coordinate spaces (that is, when the problem is essentially d independent, one-dimensional problems).

In this section we will state two quite general cases where we can obtain the exact value of the game, even though the problem does not decompose on a per coordinate basis. Note that in both cases the optimal strategy for w_{t+1} will be in the direction of θ_t .

We study the game when the horizon T is known, with a benchmark function of the form $B(\theta) = f(\|\theta\|)$ for an increasing convex function $f : [0, +\infty] \rightarrow \mathbb{R}$ (which ensures B is convex). Note this form for B is particularly natural given our desire to prove results that hold for general Hilbert spaces. We will then be able to derive regret bounds using Theorem 1, and the following technical lemma:

Lemma 3 *Let $B(\theta) = f(\|\theta\|)$ for $f : \mathbb{R} \rightarrow (-\infty, +\infty]$ even. Then, $B^*(u) = f^*(\|u\|)$.*

Recall that f is even if $f(x) = f(-x)$. Our key tool will be a careful study of the one-round version of this game. For this section, we let $h : \mathbb{R} \rightarrow \mathbb{R}$ be an even convex function that is increasing on $[0, \infty]$, $\mathcal{G} = \{g : \|g\| \leq G\}$, and d the dimension of \mathcal{H} . We consider the one-round game

$$H \equiv \min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + h(\|\theta - g\|), \quad (15)$$

where $\theta \in \mathcal{H}$ is a fixed parameter. For results regarding this game, we let $H(w, g) = \langle w, g \rangle + h(\|\theta - g\|)$, $w^* = \arg \min_w \max_{g \in \mathcal{G}} H(w, g)$, and $g^* = \arg \max_{g \in \mathcal{G}} H(w^*, g)$. Also, let $\hat{\theta} = \frac{\theta}{\|\theta\|}$ if $\|\theta\| \neq 0$, and $\mathbf{0}$ otherwise.

4.1. The case of the orthogonal adversary

Let $B(\theta) = f(\|\theta\|)$ for an increasing convex function $f : [0, \infty] \rightarrow \mathbb{R}$, and define $f_t(x) = f(\sqrt{x^2 + G^2(T-t)})$. Note that $f_t(\|\theta\|)$ can be viewed as a smoothed version of $B(\theta)$, since $\sqrt{\|\theta\|^2 + C}$ is a smoothed version of $\|\theta\|$ for a constant $C > 0$. Moreover, $f_0(\|\theta\|) = B(\theta)$.

Our first key result is the following:

Theorem 4 *Let the adversary play from $\mathcal{G} = \{g : \|g\| \leq G\}$ and assume all the f_t satisfy*

$$\min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + f_{t+1}(\|\theta - g\|) = f_{t+1}(\sqrt{\|\theta\|^2 + G^2}). \quad (16)$$

Then the value of the game is $f(G\sqrt{T})$, the conditional value is $V_t(\theta) = f_t(\|\theta\|) = f_{t+1}(\sqrt{\|\theta\|^2 + G^2})$, and the optimal strategy can be found using (14) on V_t .

Further, a sufficient condition for (16) is that $d > 1$, f is twice differentiable, and $f''(x) \leq f'(x)/x$, for all $x > 0$. In this case we also have that the minimax optimal strategy is

$$w_{t+1} = \nabla V_t(\theta_t) = \theta_t \frac{f'(\sqrt{\|\theta_t\|^2 + G^2(T-t)})}{\sqrt{\|\theta_t\|^2 + G^2(T-t)}}. \quad (17)$$

In this case, the minimax optimal strategy (20) is equivalent to the FTRL strategy in (10) with the time varying regularizer $V_t^*(w)$. The key lemma needed for the proof is the following:

Lemma 5 Consider the game of (15). Then, if $d > 1$, h is twice differentiable, and $h''(x) \leq \frac{h'(x)}{x}$ for $x > 0$, we have:

$$H = h\left(\sqrt{\|\theta\|^2 + G^2}\right) \quad \text{and} \quad w^* = \frac{\theta}{\sqrt{\|\theta\|^2 + G^2}} h'\left(\sqrt{\|\theta\|^2 + G^2}\right).$$

Any g^* such that $\langle \theta, g^* \rangle = 0$ and $\|g^*\| = G$ is a minimax play for the adversary.

We defer the proofs to the Appendix (of the proofs in the appendix, the proof of Lemmas 5 and 8 are perhaps the most important and instructive). Since the best response of the adversary is always to play a g^* orthogonal to θ , we call this the case of the orthogonal adversary.

4.2. The case of the parallel adversary, and Normal approximations

We analyze a second case where (15) has closed-form solution, and hence derive a class of games where we can cleanly state the value of the game and the minimax optimal strategy. The results of McMahan and Abernethy (2013) can be viewed as a special case of the results in this section.

First, we introduce some notation. We write $\tau \equiv T - t$ when T and t are clear from context. We write $r \sim \{-1, 1\}$ to indicate r is a Rademacher random variable, and $r_\tau \sim \{-1, 1\}^\tau$ to indicate r_τ is the sum of τ IID Rademacher random variables. Let $\sigma = \sqrt{\pi/2}$. We write ϕ for a random variable with distribution $N(0, \sigma^2)$, and similarly define $\phi_\tau \sim N(0, (T - t)\sigma^2)$. Then, define

$$f_t(x) = \mathbb{E}_{r_\tau \sim \{-1, 1\}^\tau} [f(|x + r_\tau G|)] \quad \text{and} \quad \hat{f}_t(x) = \mathbb{E}_{\phi_\tau \sim N(0, \tau\sigma^2)} [f(|x + \phi_\tau G|)], \quad (18)$$

and note $B(\theta) = f_T(\|\theta\|) = \hat{f}_T(\|\theta\|)$ since ϕ_0 and r_0 are always zero. These functions are exactly smoothed version of the function f used to define B . With these definitions, we can now state:

Theorem 6 Let $B(\theta) = f(\|\theta\|)$ for an increasing convex function $f : [0, \infty] \rightarrow \mathbb{R}$, and let the adversary play from $\mathcal{G} = \{g : \|g\| \leq G\}$. Assume f_t and \hat{f}_t as in (18) for all t . If all the f_t satisfy

$$\min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + f_{t+1}(\|\theta - g\|) = \mathbb{E}_{r \sim \{-1, 1\}} [f_{t+1}(\|\theta\| + rG)], \quad (19)$$

then $V_t(\theta) = f_t(\|\theta\|)$ is exactly the conditional value of the game, and (14) gives the minimax optimal strategy:

$$w_{t+1} = \hat{\theta} \frac{\hat{f}_{t+1}(\|\theta\| + G) - f_{t+1}(\|\theta\| - G)}{2G}. \quad (20)$$

Similarly, suppose the \hat{f}_t satisfy the equality (19) (with \hat{f}_t replacing f_t). Then $q_t(\theta) = \hat{f}_t(\|\theta\|)$ is an admissible relaxation of V_t , satisfying (13) with $\hat{\epsilon}_t = 0$, using w_{t+1} based on (14). Further, a sufficient condition for (19) is that $d = 1$, or $d > 1$, the f_t (or \hat{f}_t , respectively) are twice differentiable, and satisfy and $f_t''(x) \geq f_t'(x)/x$ for all $x > 0$.

Contrary to the case of the orthogonal adversary, the strategy in (20) cannot easily be interpreted as an FTRL algorithm. The proof is based on two lemmas. The first provides the key tool in supporting the Normal relaxation:

Lemma 7 Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a convex function and $\sigma^2 = \pi/2$. Then,

$$\mathbb{E}_{g \sim \{-1, 1\}} [f(g)] \leq \mathbb{E}_{\phi \sim N(0, \sigma^2)} [f(\phi)].$$

Proof First observe that $\mathbb{E}[(\phi - 1)\mathbf{1}\{\phi > 0\}] = 0$ and $\mathbb{E}[(\phi + 1)\mathbf{1}\{\phi < 0\}] = 0$ by our choice of σ . We will use two lower bounds on the function f , which follow from convexity:

$$f(x) \geq f(1) + f'(1)(x - 1) \quad \text{and} \quad f(x) \geq f(-1) + f'(-1)(x + 1).$$

Writing out the value of $\mathbb{E}[f(\phi)]$ explicitly we have

$$\begin{aligned} \mathbb{E}[f(\phi)] &= \mathbb{E}[f(\phi)\mathbf{1}\{\phi < 0\}] + \mathbb{E}[f(\phi)\mathbf{1}\{\phi > 0\}] \\ &\geq \mathbb{E}[(f(-1) + f'(-1)(\phi + 1))\mathbf{1}\{\phi < 0\}] + \mathbb{E}[(f(1) + f'(1)(\phi - 1))\mathbf{1}\{\phi > 0\}] \\ &= \frac{f(-1) + f(1)}{2} + f'(-1)\mathbb{E}[(\phi + 1)\mathbf{1}\{\phi < 0\}] + f'(1)\mathbb{E}[(\phi - 1)\mathbf{1}\{\phi > 0\}]. \end{aligned}$$

The latter two terms vanish, giving the stated inequality. \blacksquare

The second lemma is used to prove the sufficient condition by solving the one-round game; again, the proof is deferred to the Appendix. Note that functions of the form $h(x) = g(x^2)$, with g convex always satisfies the conditions of the following Lemma.

Lemma 8 *Consider the game of (15). Then, if $d = 1$, or if $d > 1$, h is twice differentiable, and $h''(x) > \frac{h'(x)}{x}$ for $x > 0$, then*

$$H = \frac{h(\|\theta\| + G) + h(\|\theta\| - G)}{2} \quad \text{and} \quad w^* = \hat{\theta} \frac{h(\|\theta\| + G) - h(\|\theta\| - G)}{2G}.$$

Any g^* that satisfies $|\langle \theta, g^* \rangle| = G\|\theta\|$ and $\|g^*\| = G$ is a minimax play for the adversary.

The adversary can always play $g^* = G \frac{\theta}{\|\theta\|}$ when $\theta \neq \mathbf{0}$, and so we describe this as the case of the parallel adversary. In fact, inductively this means that *all* the adversary's plays g_t can be on the same line, providing intuition for the fact that this lemma also applies in the 1-dimensional case.

Theorem 6 provides a recipe to produce suitable relaxations q_t which may, in certain cases, exhibit nice closed form solutions. The interpretation here is that a ‘‘Gaussian adversary’’ is stronger than one playing from the set $[-1, 1]$ which leads to IID Rademacher behavior, and this allows us to generate such potential functions via Gaussian smoothing. In this view, note that our choice of σ^2 gives $\mathbb{E}_\phi[\|\phi\|] = 1$.

5. A Power Family of Minimax Algorithms

We analyze a family of algorithms based on potentials $B(\theta) = f(\|\theta\|)$ where $f(x) = \frac{W}{p}|x|^p$ for parameters $W > 0$ and $p \in [1, 2]$, when the dimension is at least two. This is reminiscent of p -norm algorithms (Gentile, 2003), but the connection is superficial—the norm we use to measure θ is always the norm of our Hilbert space. Our main result is:

Corollary 9 *Let $d > 1$ and $W > 0$, and let f and B be defined as above. Define $f_t(x) = \frac{W}{p}(x^2 + (T - t)G)^{p/2}$. Then, $f_t(\|\theta\|)$ is the conditional value of the game, and the optimal strategy is as in Theorem 4. If $p \in (1, 2]$, letting $q \geq 2$ such that $1/p + 1/q = 1$, we have a bound*

$$\text{Regret}(u) \leq \frac{1}{W^{q-1}q} \|u\|^q + \frac{W}{p} (G\sqrt{T})^p \leq \left(\frac{1}{p} + \frac{1}{q} \|u\|^q\right) G\sqrt{T},$$

where the second inequality comes by taking $W = (G\sqrt{T})^{1-p}$. For all u , the bound $(\frac{1}{p} + \frac{1}{q}\|u\|^q)G\sqrt{T}$ is minimized by taking $p = 2$. For $p = 1$, we have

$$\forall u : \|u\| \leq W, \quad \text{Regret}(u) \leq WG\sqrt{T}.$$

Proof Let $f(x) = \frac{W}{p}|x|^p$ for $p \in [1, 2]$, Then, $f''(x) \leq f'(x)/x$, in fact basic calculations show $\frac{f'(x)/x}{f''(x)} = \frac{1}{p-1} \geq 1$ when $p \leq 2$. Hence, we can apply Theorem 4, proving the claim on the f_t . The regret bounds can then be derived from Corollary 2, which gives $\text{Regret}(u) \leq f^*(u) + f(G\sqrt{T})$, noting $f^*(u) = \frac{W}{q}|\frac{u}{W}|^q$ when $p > 1$. The fact that $p = 2$ is an optimal choice in the first bound follows from the fact that $\frac{d}{dp}(\frac{1}{p} + \frac{1}{q}\|u\|^q) \leq 0$ for $p \in (1, 2]$ with $q = \frac{p}{p-1}$. ■

The $p = 1$ case in fact exactly recaptures the result of Abernethy et al. (2008a) for linear functions, extending it also to spaces of dimension equal to two. The optimal update is $w_{t+1} = \nabla f_t(\|\theta_t\|) = W\theta_t/\sqrt{\|\theta_t\|^2 + G^2(T-t)}$. In addition to providing a regret bound for the comparator set $\mathcal{W} = \{u : \|u\| \leq W\}$, the algorithm will in fact only play points from this set.

For $p = q = 2$, writing $W = \eta$, we have

$$\text{Regret}(u) \leq \frac{1}{2\eta}\|u\|^2 + \frac{\eta}{2}G^2T,$$

for any u . In this case we see $W = \eta$ is behaving not like the radius of a comparator set, but rather as a learning rate. In fact, we have $w_{t+1} = \nabla V_t(\theta_t) = \eta\theta_t = -\eta g_{1:t}$, and so we see this minimax-optimal algorithm is in fact constant-step-size gradient descent. Taking $\eta = \frac{1}{G\sqrt{T}}$ yields $\frac{1}{2}(\|u\|^2 + 1)G\sqrt{T}$. This result complements McMahan and Abernethy (2013, Thm. 7), which covers the $d = 1$ case, or $d > 1$ when the adversary plays from $\mathcal{G} = [-1, 1]^d$.

Comparing the $p = 1$ and $p > 1$ algorithms reveals an interesting fact. For simplicity, take $G = 1$. Then, the $p = 1$ algorithm with $W = 1$ is exactly the minimax optimal algorithm for minimizing regret against comparators in the L_2 ball (for $d > 1$): the value of this game is \sqrt{T} and we can do no better (even by playing outside of the comparator set). However, picking $p > 1$ gives us algorithms that *will* play outside of the comparator set. While they cannot do better than \sqrt{T} , taking $G = 1$ and $\|u\| = 1$ shows that *all* algorithms in this family in fact achieve $\text{Regret}(u) \leq \sqrt{T}$ when $\|u\| \leq 1$, matching the exact minimax optimal value. Further, the algorithms with $p > 1$ provide much stronger guarantees, since they also give non-vacuous guarantees for $\|u\| > 1$, and tighter bounds when $\|u\| < 1$. This suggests that the $p = 2$ algorithm will be the most useful algorithm in practice, something that indeed has been observed empirically (given the prevalence of gradient descent in real applications). This result also clearly demonstrates the value of studying minimax-optimal algorithms for different choices of the benchmark B , as this can produce algorithms that are no worse and in some cases significantly better than minimax algorithms defined in terms of regret minimization directly (i.e., via (1)).

The key difference in these algorithms is not how they play against a minimax optimal adversary for the regret game, but how they play against non-worst-case adversaries. In fact, a simple induction based on Lemma 5 shows that any minimax-optimal adversary will play so that $\sqrt{\|\theta_t\|^2 + G^2(T-t)} = G\sqrt{T}$. Against such an adversary, the $p = 1$ algorithm is identical to the $p = 2$ algorithm with learning rate $\eta = \frac{1}{G\sqrt{T}}$. In fact, using the choice of W from Corollary 9, all of these algorithms play identically against a minimax adversary for the regret game.

6. Tight Bounds for Unconstrained Learning

In this section we analyze algorithms based on benchmarks and potentials of the form $\exp(\|\theta\|^2/t)$, and show they lead to a minimal dependence on $\|u\|$ in the corresponding regret bounds for a given upper bound on regret against the origin (equal to the loss of the algorithm).

First, we derive a lower bound for the known T game. Using Lemma 14 in the Appendix, we can show that the $B(\theta) = \exp(\|\theta\|^2/T)$ benchmark approximately corresponds to a regularizer of the form $\|u\| \sqrt{T \log(\sqrt{T}\|u\| + 1)}$; there is actually some technical challenge here, as the conjugate B^* cannot be computed in closed form—the given regularizer is an upper bound. This kind of regularizer is particularly interesting because it is related to parameter-free sub-gradient descent algorithms (Orabona, 2013); a similar potential function was used for a parameter-free algorithm by (Chaudhuri et al., 2009). The lower bound for this game was proven in Streeter and McMahan (2012) for 1-dimensional spaces, and Orabona (2013) extended it to Hilbert spaces and improved the leading constant. We report it here for completeness.

Theorem 10 *Fix a non-trivial Hilbert space \mathcal{H} and a specific online learning algorithm. If the algorithm guarantees a zero regret against the competitor with zero norm, then there exists a sequence of T cost vectors in \mathcal{H} , such that the regret against any other competitor is $\Omega(T)$. On the other hand, if the algorithm guarantees a regret at most of $\epsilon > 0$ against the competitor with zero norm, then, for any $0 < \eta < 1$, there exists a T_0 and a sequence of $T \geq T_0$ unitary norm vectors $g_t \in \mathcal{H}$, and a vector $u \in \mathcal{H}$ such that*

$$\text{Regret}(u) \geq (1 - \eta)\|u\| \sqrt{\frac{1}{\log 2}} \sqrt{T \log \frac{\eta\|u\|\sqrt{T}}{3\epsilon}} - 2.$$

6.1. Deriving a known- T algorithm with minimax rates via the Normal approximation

Consider the game with fixed known T , an adversary that plays from $\mathcal{G} = \{g \in \mathcal{H} \mid \|g\| \leq G\}$, and

$$B(\theta) = \epsilon \exp\left(\frac{\|\theta\|^2}{2aT}\right),$$

for constants $a > 1$ and $\epsilon > 0$. We will show that we are in the case of the parallel adversary, Section 4.2. Both computing the f_t based on Rademacher expectations and evaluating the sufficient condition for those f_t appear quite difficult, so we turn to the Normal approximation. We then have

$$\hat{f}_t(x) = \mathbb{E}_{\phi_\tau} \left[\epsilon \exp\left(\frac{(x + \phi_\tau G)^2}{2at}\right) \right] = \epsilon \left(1 - \frac{\pi G^2(T-t)}{2aT}\right)^{-\frac{1}{2}} \exp\left(\frac{x^2}{2aT - \pi G^2(T-t)}\right),$$

where we have computed the expectation in a closed form for the second equality. One can quickly verify that it satisfies the hypothesis of Theorem 6 for $a > G^2\pi/2$, hence $q_t(\theta) = \hat{f}_t(\|\theta\|)$ will be an admissible relaxation. Thus, by Corollary 2, we immediately have

$$\text{Regret}(u) \leq B^*(\theta_T) + \epsilon \left(1 - \frac{\pi G^2}{2a}\right)^{-\frac{1}{2}},$$

and so by Lemma 14 in the Appendix, we can state the following Theorem, that matches the lower bound up to a constant multiplicative factor.

Theorem 11 *Let $a > G^2\pi/2$, and $\mathcal{G} = \{g : \|g\| \leq G\}$. Denote by $\hat{\theta} = \frac{\theta}{\|\theta\|}$ if $\|\theta\| \neq 0$, and $\mathbf{0}$ otherwise. Fix the number of rounds T of the game, and consider the strategy*

$$w_{t+1} = \hat{\theta}_t \frac{\exp\left(\frac{(\|\theta_t\|+G)^2}{2aT-\pi G^2(T-t-1)}\right) - \exp\left(\frac{(\|\theta_t\|-G)^2}{2aT-\pi G^2(T-t-1)}\right)}{2G\sqrt{1-\frac{\pi G^2(T-t-1)}{2aT}}}.$$

Then, for any sequence of linear costs $\{g_t\}_{t=1}^T$, and any $u \in \mathcal{H}$, we have

$$\text{Regret}(u) \leq \|u\| \sqrt{2aT \log\left(\frac{\sqrt{aT}\|u\|}{\epsilon} + 1\right)} + \epsilon \left(\left(1 - \frac{\pi G^2}{2a}\right)^{-\frac{1}{2}} - 1 \right).$$

6.2. AdaptiveNormal: an adaptive algorithm for unknown T

Our techniques suggest the following recipe for developing adaptive algorithms: analyze the known T case, define a potential $q_t(\theta) \approx V_T(\theta)$, and then analyze the incrementally-optimal algorithm for this potential (14) via Theorem 1. We follow this recipe in the current section. Again consider the game where an adversary that plays from $\mathcal{G} = \{g \in \mathcal{H} \mid \|g\| \leq G\}$. Define the function f_t as

$$f_t(x) = \beta_t \exp\left(\frac{x}{2at}\right),$$

where $a > \frac{3\pi G^2}{4}$, and the β_t is a decreasing sequence that will be specified in the following. From this, we define the potential $q_t(\theta) = f_t(\|\theta\|^2)$. Suppose we play the incrementally-optimal algorithm of (14). Using Lemma 8 we can write the minimax value for the one-round game,

$$\begin{aligned} \epsilon_t(\theta_t) &= \mathbb{E}_{r \sim \{-1,1\}} [f_{t+1}((\|\theta_t\| + rG)^2)] - q_t(\theta_t) \\ &\leq \mathbb{E}_{\phi \sim N(0,\sigma^2)} [f_{t+1}((\|\theta_t\| + \phi G)^2)] - q_t(\theta_t). \end{aligned} \quad \text{Lemma 7.}$$

Using Lemma 17 in the Appendix and our hypothesis on a , we have that the RHS of this inequality is maximized for $\|\theta_t\| = 0$. Hence, using the inequality $\sqrt{a+b} \leq \sqrt{a} + \frac{b}{2\sqrt{a}}$, $\forall a, b > 0$, we get

$$\epsilon_t(\theta_t) \leq \beta_{t+1} \sqrt{1 + \frac{\pi G^2}{2a(t+1) - \pi G^2}} - \beta_t \leq \frac{\beta_t}{2} \frac{\pi G^2}{2a(t+1) - \pi G^2} \leq \frac{\pi G^2 \beta_t}{4at}.$$

Thus, choosing $\beta_t = \epsilon / \log^2(t+1)$, for example, is sufficient to prove that $\epsilon_{1:T}$ is bounded by $\epsilon \frac{\pi G^2}{a}$ (Baxley, 1992). Hence, again using Corollary 2 and Lemma 14 in the Appendix, we can state the following Theorem.

Theorem 12 *Let $a > 3G^2\pi/4$, and $\mathcal{G} = \{g : \|g\| \leq G\}$. Denote by $\hat{\theta} = \frac{\theta}{\|\theta\|}$ if $\|\theta\| \neq 0$, and $\mathbf{0}$ otherwise. Consider the strategy*

$$w_{t+1} = \hat{\theta}_t \left(\exp\left(\frac{(\|\theta_t\| + G)^2}{2a(t+1)}\right) - \exp\left(\frac{(\|\theta_t\| - G)^2}{2a(t+1)}\right) \right) (2G \log^2(t+2))^{-1}.$$

Then, for any sequence of linear costs $\{g_t\}_{t=1}^T$, and any $u \in \mathcal{H}$, we have

$$\text{Regret}(u) \leq \|u\| \sqrt{2aT \log\left(\frac{\sqrt{aT}\|u\| \log^2(T+1)}{\epsilon} + 1\right)} + \epsilon \left(\frac{\pi G^2}{a} - 1 \right).$$

Acknowledgments

We thank Jacob Abernethy for many useful conversations about this work.

References

- J. Abernethy and M.K. Warmuth. Repeated games against budgeted adversaries. *Advances in Neural Information Processing Systems*, 22, 2010.
- J. Abernethy, J. Langford, and M. K. Warmuth. Continuous experts and the Binning algorithm. In *Proceedings of the 19th Annual Conference on Learning Theory (COLT06)*, pages 544–558. Springer, June 2007.
- J. Abernethy, P. L. Bartlett, A. Rakhlin, and A. Tewari. Optimal strategies and minimax lower bounds for online convex games. In *COLT*, 2008a.
- J. Abernethy, M. K. Warmuth, and J. Yellin. Optimal strategies from random walks. In *Proceedings of the 21st Annual Conference on Learning Theory (COLT 08)*, pages 437–445, July 2008b.
- J. V. Baxley. Euler’s constant, Taylor’s formula, and slowly converging series. *Mathematics Magazine*, 65(5):302–313, 1992.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- K. Chaudhuri, Y. Freund, and D. Hsu. A parameter-free hedging algorithm. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 297–305. 2009.
- C. Gentile. The robustness of the p-norm algorithms. *Machine Learning*, 53(3):265–299, 2003.
- A. J. Grove, N. Littlestone, and D. Schuurmans. General convergence results for linear discriminant updates. *Machine Learning*, 43(3):173–210, 2001.
- J. Kivinen and M. K. Warmuth. Relative loss bounds for multidimensional regression problems. *Machine Learning*, 45(3):301–329, 2001.
- H. B. McMahan and J. Abernethy. Minimax optimal algorithms for unconstrained linear optimization. In *NIPS*, 2013.
- F. Orabona. Dimension-free exponentiated gradient. In *NIPS*, 2013.
- F. Orabona, K. Crammer, and N. Cesa-Bianchi. A generalized online mirror descent with applications to classification and regression, 2013. arXiv:1304.2994.
- A. Rakhlin. Lecture notes on online learning. Technical report, 2009.
- A. Rakhlin, O. Shamir, and K. Sridharan. Localization and adaptation in online learning. In *AISTATS*, 2013.
- S. Rakhlin, O. Shamir, and K. Sridharan. Relax and randomize: From value to algorithms. In *NIPS*, 2012.
- S. Shalev-Shwartz. Online learning: Theory, algorithms, and applications. Technical report, The Hebrew University, 2007. PhD thesis.
- S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 2012.
- M. Streeter and H. B. McMahan. No-regret algorithms for unconstrained online convex optimization. In *NIPS*, 2012.
- L. Xiao. Dual averaging method for regularized stochastic learning and online optimization. In *NIPS*, 2009.

Appendix A. Proofs

A.1. Proof of Theorem 1

Proof Suppose the algorithm provides the reward guarantee (4). First, note that for any comparator u , by definition we have

$$\text{Regret}(u) = -\text{Reward} - \langle g_{1:T}, u \rangle. \quad (21)$$

Then, applying the definitions of Reward, Regret, and the Fenchel conjugate, we have

$$\begin{aligned} \text{Regret}(u) &= \theta_T \cdot u - \text{Reward} && \text{By (21)} \\ &\leq \theta_T \cdot u - q_T(\theta_T) + \hat{\epsilon}_{1:T} && \text{By assumption (4)} \\ &\leq \max_{\theta} (\theta \cdot u - q_T(\theta) + \hat{\epsilon}_{1:T}) \\ &= q_T^*(u) + \hat{\epsilon}_{1:T}. \end{aligned}$$

For the other direction, assuming (5), we have for any comparator u ,

$$\begin{aligned} \text{Reward} &= \theta_T \cdot u - \text{Regret}(u) && \text{By (21)} \\ &= \max_v (\theta \cdot v - \text{Regret}(v)) \\ &\geq \max_v (\theta \cdot v - q_T^*(v) - \hat{\epsilon}_{1:T}) && \text{By assumption (5)} \\ &= q_T(\theta) - \hat{\epsilon}_{1:T}. \end{aligned}$$

Alternatively, one can prove this from the Fenchel-Young inequality. ■

A.2. Proof of Lemma 3

Proof We have $B^*(u) = \sup_{\theta} \langle u, \theta \rangle - f(\|\theta\|)$. If $\|u\| = 0$, the stated equality is correct, in fact

$$B^*(u) = \sup_{\theta} -f(\|\theta\|) = \sup_{\alpha \geq 0} -f(\alpha) = \sup_{\alpha \in \mathbb{R}} -f(\alpha) = f^*(0).$$

Hence we can assume $\|u\| \neq 0$, and by inspection we can take $\theta = \alpha u / \|u\|$, with $\alpha \geq 0$, and so

$$B^*(u) = \sup_{\alpha \geq 0} \alpha \|u\| - f(\alpha) = \sup_{\alpha \in \mathbb{R}} \alpha \|u\| - f(\alpha) = f^*(\|u\|).$$
■

A.3. Proof of Theorem 4

Proof First we show that if f satisfies the condition on the derivatives, the same conditions is satisfied by f_t , for all t . We have that all the f_t have the form $h(x) = f(\sqrt{x^2 + a})$, where $a \geq 0$. Hence we have to prove that $\frac{xh'(x)}{h'(x)} \leq 1$. We have that $h'(x) = \frac{xf'(\sqrt{x^2+a})}{\sqrt{x^2+a}}$, and $h''(x) = \frac{x^2 f''(\sqrt{x^2+a}) + \frac{a}{\sqrt{x^2+a}} f'(\sqrt{x^2+a})}{x^2+a}$, so

$$\frac{xh''(x)}{h'(x)} = \frac{x^2 f''(\sqrt{x^2+a}) \sqrt{x^2+a}}{f'(\sqrt{x^2+a})(x^2+a)} + \frac{a}{x^2+a} \leq \frac{x^2}{x^2+a} + \frac{a}{x^2+a} = 1,$$

where in the inequality we used the hypothesis on the derivatives of f .

We show V_t has the stated form by induction from T down to 0. The base case for $t = T$ is immediate. For the induction step, we have

$$\begin{aligned}
 V_t(\theta) &= \min_w \max_g \langle w, g \rangle + V_{t+1}(\theta - g) && \text{Defn.} \\
 &= \min_w \max_g \langle w, g \rangle + f_{t+1}(\|\theta - g\|) && \text{(IH)} \\
 &= f_{t+1}\left(\sqrt{\|\theta\|^2 + G^2}\right) && \text{Assumption (16)} \\
 &= f\left(\sqrt{\|\theta\|^2 + G^2(T-t)}\right).
 \end{aligned}$$

The sufficient condition for (16) follow immediately from Lemma 5. ■

A.4. Proof of Theorem 6

Proof First, we need to show the functions f_t and \hat{f}_t of (18) are even. Let r be a random variable draw from any symmetric distribution. Then, we have

$$f_t(x) = \mathbb{E}[f(|x + r|)] = \mathbb{E}[f(|-x - r|)] = \mathbb{E}[f(|-x + r|)] = f_t(-x),$$

where we have used the fact that $|\cdot|$ is even and the symmetry of r .

We show $f_t(\|\theta\|) = V_t(\theta)$ inductively from $t = T$ down to $t = 0$. The base case $T = t$ follows from the definition of f_t . Then, suppose the result holds for $t + 1$. We have

$$\begin{aligned}
 V_t(\theta) &= \min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + V_{t+1}(\theta - g) && \text{Defn.} \\
 &= \min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + f_{t+1}(\|\theta - g\|) && \text{IH} \\
 &= \mathbb{E}_{r \sim \{-1,1\}} [f_{t+1}(\|\theta\| + rG)] && \text{Lemma 8} \\
 &= \mathbb{E}_{r \sim \{-1,1\}} \left[\mathbb{E}_{r_{\tau-1} \sim \{-1,1\}^{\tau-1}} [f(\|\theta\| + rG + r_{\tau-1}G)] \right] \\
 &= f_t(\|\theta\|),
 \end{aligned}$$

where the last two lines follow from the definition of f_t and f_{t+1} . The case for \hat{f}_t is similar, using the hypothesis of the Theorem we have

$$\begin{aligned}
 \min_w \max_{g \in \mathcal{G}} \langle g, w \rangle + \hat{f}_{t+1}(\|\theta - g\|) &= E_{r \sim \{-1,1\}} [\hat{f}_{t+1}(\|\theta\| + rG)] \leq E_{r \sim \mathcal{N}(0, \sigma^2)} [\hat{f}_{t+1}(\|\theta\| + \phi G)] \\
 &= \hat{f}_t(\|\theta\|),
 \end{aligned}$$

where in the inequality we used Lemma 7, and in the second equality the definition of \hat{f}_t . Hence, $q_t(\theta) = \hat{f}_t(\|\theta\|)$ satisfy (13) with $\hat{\epsilon}_t = 0$. Finally, the sufficient conditions come immediately from Lemma 8. ■

A.5. Analysis of the one-round game: Proofs of Lemmas 5 and 8

In the process of proving these lemmas, we also show the following general lower bound:

Lemma 13 *Under the same definitions as in Lemma 5, if $d > 1$, we have*

$$H \geq h\left(\sqrt{\|\theta\|^2 + G^2}\right).$$

We now proceed with the proofs. The $d = 1$ case for Lemma 8 was proved in [McMahan and Abernethy \(2013\)](#).

Before proving the other results, we simplify a bit the formulation of the minimax problem. For the other results, the maximization wrt g of a convex function is always attained when $\|g\| = G$. Moreover, in the case of $\|\theta\| = 0$ the other results are true, in fact

$$\min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + h(\|\theta - g\|) = \min_w \max_{\|g\|=G} \langle w, g \rangle + h(\|g\|) = \min_w G\|w\| + h(G) = h(G).$$

Hence, without loss of generality, in the following we can write $w = \alpha \frac{\theta}{\|\theta\|} + \hat{w}$, where $\langle \hat{w}, \theta \rangle = 0$. It is easy to see that in all the cases the optimal choice of g turns out to be $g = \beta \frac{\theta}{\|\theta\|} + \gamma \hat{w}$, where $\gamma \geq 0$. With these settings, the minimax problem is equivalent to

$$\min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + h(\|\theta - g\|) = \min_{\alpha, \hat{w}} \max_{\beta^2 + \|\hat{w}\|^2 \gamma^2 = G^2} \alpha\beta + \gamma\|\hat{w}\|^2 + h(\sqrt{\|\theta\|^2 - 2\beta\|\theta\| + G^2}).$$

By inspection, the player can always choose $\hat{w} = \mathbf{0}$ so $\gamma\|\hat{w}\|^2 = 0$. Hence we have a simplified and equivalent form of our optimization problem

$$\min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + h(\|\theta - g\|) = \min_{\alpha} \max_{\beta^2 \leq G^2} \alpha\beta + h(\sqrt{\|\theta\|^2 - 2\beta\|\theta\| + G^2}). \quad (22)$$

For Lemma 13, it is enough to set $\beta = 0$ in (22).

For Lemma 5, we upper bound the minimum wrt to α with the specific choice of α . In particular, we set $\alpha = \frac{\|\theta\|}{\sqrt{\|\theta\|^2 + G^2}} h'(\sqrt{\|\theta\|^2 + G^2})$ in (22), and get

$$\min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + h(\|\theta - g\|) \leq \max_{\beta^2 \leq G^2} \frac{\beta\|\theta\| h'(\sqrt{\|\theta\|^2 + G^2})}{\sqrt{\|\theta\|^2 + G^2}} + h(\sqrt{\|\theta\|^2 - 2\beta\|\theta\| + G^2}).$$

The derivative of argument of the max wrt β is

$$\frac{\|\theta\| h'(\sqrt{\|\theta\|^2 + G^2})}{\sqrt{\|\theta\|^2 + G^2}} - \frac{\|\theta\| h'(\sqrt{\|\theta\|^2 - 2\beta\|\theta\| + G^2})}{\sqrt{\|\theta\|^2 - 2\beta\|\theta\| + G^2}}. \quad (23)$$

We have that if $\beta = 0$ the first derivative is 0. Using the hypothesis on the first and second derivative of h , we have that the second term in (23) increases in β . Hence $\beta = 0$ is the maximum. Comparing the obtained upper bound with the lower bound in Lemma 13, we get the stated equality.

For Lemma 8, the second derivative wrt β of the argument of the minimax problem in (22) is

$$-\|\theta\| \frac{-\|\theta\| h''(\sqrt{\|\theta\|^2 + G^2 - 2\beta\|\theta\|}) + h'(\sqrt{\|\theta\|^2 + G^2 - 2\beta\|\theta\|}) \frac{\|\theta\|}{\sqrt{\|\theta\|^2 + G^2 - 2\beta\|\theta\|}}}{\|\theta\| + G^2 - 2\beta\|\theta\|}$$

that is non negative, for our hypothesis on the derivatives of h . Hence, the argument of the minimax problem is convex wrt β , hence the maximum is achieved at the boundary of the domains, that is $\beta^2 = G^2$. So, we have

$$\min_w \max_{g \in \mathcal{G}} \langle w, g \rangle + h(\|\theta - g\|) = \max(-G\alpha + h(\|\theta\| + G), G\alpha + h(\|\theta\| - G)) .$$

The argmin of this quantity wrt to α is obtained when the the two terms in the max are equal, so we obtained the stated equality.

A.6. Lemma 14

Lemma 14 Define $f(\theta) = \beta \exp \frac{\|\theta\|^2}{2\alpha}$, for $\alpha, \beta > 0$. Then

$$f^*(w) \leq \|w\| \sqrt{2\alpha \log \left(\frac{\sqrt{\alpha}\|w\|}{\beta} + 1 \right)} - \beta .$$

Proof From the definition of Fenchel dual, we have

$$f^*(w) = \max_{\theta} \langle \theta, w \rangle - f(\theta) \leq \langle \theta^*, w \rangle - \beta .$$

where $\theta^* = \arg \max_{\theta} \langle \theta, w \rangle - f(\theta)$. We now use the fact that θ^* satisfies $w = \nabla f(\theta^*)$, that is

$$w = \theta^* \frac{\beta}{\alpha} \exp \left(\frac{\|\theta^*\|^2}{2\alpha} \right) ,$$

in other words we have that θ^* and w are in the same direction. Hence we can set $\theta^* = qw$, so that $f^*(w) \leq q\|w\|^2 - \beta$. We now need to look for $q > 0$, solving

$$\begin{aligned} \frac{q\beta}{\alpha} \exp \left(\frac{\|w\|^2 q^2}{2\alpha} \right) = 1 &\Leftrightarrow \frac{\|w\|^2 q^2}{2\alpha} + \log \frac{\beta q}{\alpha} = 0 \\ &\Leftrightarrow q = \sqrt{\frac{2\alpha}{\|w\|^2} \log \frac{\alpha}{q\beta}} = \sqrt{\frac{2\alpha}{\|w\|^2} \log \left(\frac{\sqrt{\alpha}\|w\|}{\beta \sqrt{2 \log \frac{\alpha}{q\beta}}} \right)} . \end{aligned}$$

Using the elementary inequality $\log x \leq \frac{m}{e} x^{\frac{1}{m}}$, $\forall m > 0$, we have

$$\begin{aligned} q^2 &= \frac{2\alpha}{\|w\|^2} \log \frac{\alpha}{q\beta} \leq \frac{2m\alpha}{e\|w\|^2} \left(\frac{\alpha}{q\beta} \right)^{\frac{1}{m}} \\ \Rightarrow q^{\frac{2m+1}{m}} &\leq \frac{2m\alpha^{\frac{m+1}{m}}}{e\|w\|^2 \beta^{\frac{1}{m}}} \Rightarrow q \leq \left(\frac{2m}{e\|w\|^2} \right)^{\frac{m}{2m+1}} \alpha^{\frac{m+1}{2m+1}} \beta^{-\frac{1}{2m+1}} \\ \Rightarrow \frac{\alpha}{\beta q} &\geq \left(\frac{e\|w\|^2}{2m} \right)^{\frac{m}{2m+1}} \alpha^{1-\frac{m+1}{2m+1}} \beta^{\frac{1}{2m+1}-1} = \frac{\alpha}{\beta q} \geq \left(\sqrt{\frac{e}{2m}} \frac{\|w\| \sqrt{\alpha}}{\beta} \right)^{\frac{2m}{2m+1}} . \end{aligned}$$

Hence we have

$$q \leq \sqrt{\frac{2\alpha}{\|w\|^2} \log \left(\frac{\sqrt{\alpha}\|w\|}{\beta \sqrt{\frac{4m}{2m+1} \log \sqrt{\frac{e}{2m}} \frac{\|w\| \sqrt{\alpha}}{\beta}}}} \right)} .$$

We set m such that $\sqrt{\frac{e}{2m}} \frac{\|w\|\sqrt{\alpha}}{\beta} = \sqrt{e}$, that is $\frac{1}{2} \left(\frac{\|w\|\sqrt{\alpha}}{\beta} \right)^2 = m$. Hence we have $\log \sqrt{\frac{e}{2m}} \frac{\|w\|\sqrt{\alpha}}{\beta} = \frac{1}{2}$ and $\frac{\|w\|\sqrt{\alpha}}{\sqrt{2m}} = 1$, and obtain

$$f^*(w) \leq \|w\| \sqrt{2\alpha \log \left(\sqrt{\left(\frac{\sqrt{\alpha}\|w\|}{\beta} \right)^2 + 1} \right)} - \beta \leq \|w\| \sqrt{2\alpha \log \left(\frac{\sqrt{\alpha}\|w\|}{\beta} + 1 \right)} - \beta.$$

■

A.7. Lemma 17

Lemma 15 *Let $f(x) = b \exp\left(\frac{x^2}{a}\right) - \exp\left(\frac{x^2}{c}\right)$. If $a \geq c > 0$, $b \geq 0$, and $bc \leq a$, then the function $f(x)$ is decreasing for $x \geq 0$.*

Proof The proof is immediate from the study of the first derivative. ■

Lemma 16 *Let $f(t) = \frac{a^{\frac{3}{2}} t \sqrt{t+1}}{(a(t+1)-b)^{\frac{3}{2}}}$, with $a \geq 3/2b > 0$. Then $f(t) \leq 1$ for any $t \geq 0$.*

Proof The sign of the first derivative of the function has the same sign of

$$(2a - 3b)(t + 1) + b,$$

hence from the hypothesis on a and b the function is strictly increasing. Moreover the asymptote for $t \rightarrow \infty$ is 1, hence we have the stated upper bound. ■

Lemma 17 *Let $f_t(x) = \beta_t \exp\left(\frac{x}{2at}\right)$, $\beta_{t+1} \leq \beta_t$, $\forall t$. If $a \geq \frac{3\pi G^2}{4}$, then*

$$\arg \max_x \mathbb{E}_{\phi \sim N(0, \sigma^2)} [f_{t+1}(x + \phi G)] - f_t(x) = 0,$$

where $\sigma^2 = \frac{\pi}{2}$.

Proof We have

$$\mathbb{E}_{\phi \sim N(0, \sigma^2)} [f_{t+1}(x + \phi G)] = \beta_{t+1} \sqrt{\frac{a(t+1)}{a(t+1) - \sigma^2 G^2}} \exp\left(\frac{x^2}{2[a(t+1) - \sigma^2 G^2]}\right),$$

so we have to study the max of

$$\beta_{t+1} \sqrt{\frac{a(t+1)}{a(t+1) - \sigma^2 G^2}} \exp\left(\frac{x^2}{2[a(t+1) - \sigma^2 G^2]}\right) - \beta_t \exp\left(\frac{x^2}{2at}\right).$$

The function is even, so we have a maximum in zero iff the function is decreasing for $x > 0$. Observe that, from Lemma 16, for any $t \geq 0$

$$\sqrt{\frac{a(t+1)}{a(t+1) - \sigma^2 G^2} \frac{at}{a(t+1) - \sigma^2 G^2}} \leq 1.$$

Hence, using Lemma 15, we obtain that the stated result. ■