
Online Learning: Stochastic, Constrained, and Smoothed Adversaries

Alexander Rakhlin
Department of Statistics
University of Pennsylvania
rakhlin@wharton.upenn.edu

Karthik Sridharan
Toyota Technological Institute at Chicago
karthik@ttic.edu

Ambuj Tewari
Computer Science Department
University of Texas at Austin
ambuj@cs.utexas.edu

Abstract

Learning theory has largely focused on two main learning scenarios: the classical statistical setting where instances are drawn i.i.d. from a fixed distribution, and the adversarial scenario wherein, at every time step, an adversarially chosen instance is revealed to the player. It can be argued that in the real world neither of these assumptions is reasonable. We define the minimax value of a game where the adversary is restricted in his moves, capturing stochastic and non-stochastic assumptions on data. Building on the sequential symmetrization approach, we define a notion of distribution-dependent Rademacher complexity for the spectrum of problems ranging from i.i.d. to worst-case. The bounds let us immediately deduce variation-type bounds. We study a smoothed online learning scenario and show that exponentially small amount of noise can make function classes with infinite Littlestone dimension learnable.

1 Introduction

In the papers [1, 10, 11], an array of tools has been developed to study the minimax value of diverse sequential problems under the *worst-case* assumption on Nature. In [10], many analogues of the classical notions from statistical learning theory have been developed, and these have been extended in [11] for performance measures well beyond the additive regret. The process of *sequential symmetrization* emerged as a key technique for dealing with complicated nested minimax expressions. In the worst-case model, the developed tools give a unified treatment to such sequential problems as regret minimization, calibration of forecasters, Blackwell's approachability, Φ -regret, and more.

Learning theory has been so far focused predominantly on the i.i.d. and the worst-case learning scenarios. Much less is known about learnability in-between these two extremes. In the present paper, we make progress towards filling this gap by proposing a framework in which it is possible to variously restrict the behavior of Nature. By restricting Nature to play i.i.d. sequences, the results boil down to the classical notions of statistical learning in the supervised learning scenario. By not placing any restrictions on Nature, we recover the worst-case results of [10]. Between these two endpoints of the spectrum, particular assumptions on the adversary yield interesting bounds on the minimax value of the associated problem. Once again, the sequential symmetrization technique arises as the main tool for dealing with the minimax value, but the proofs require more care than in the i.i.d. or completely adversarial settings.

Adapting the game-theoretic language, we will think of the learner and the adversary as the two players of a zero-sum repeated game. Adversary’s moves will be associated with “data”, while the moves of the learner – with a function or a parameter. This point of view is not new: game-theoretic minimax analysis has been at the heart of statistical decision theory for more than half a century (see [3]). In fact, there is a well-developed theory of minimax estimation when restrictions are put on either the choice of the adversary or the allowed estimators by the player. *We are not aware of a similar theory for sequential problems with non-i.i.d. data.*

The main contribution of this paper is the development of tools for the analysis of online scenarios where the adversary’s moves are restricted in various ways. In addition to general theory, we consider several interesting scenarios which can be captured by our framework. All proofs are deferred to the appendix.

2 Value of the Game

Let \mathcal{F} be a closed subset of a complete separable metric space, denoting the set of moves of the learner. Suppose the adversary chooses from the set \mathcal{X} . Consider the `Online Learning Model`, defined as a T -round interaction between the learner and the adversary: On round $t = 1, \dots, T$, the learner chooses $f_t \in \mathcal{F}$, the adversary simultaneously picks $x_t \in \mathcal{X}$, and the learner suffers loss $f_t(x_t)$. The goal of the learner is to minimize regret, defined as $\sum_{t=1}^T f_t(x_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^T f(x_t)$. It is a standard fact that simultaneity of the choices can be formalized by the first player choosing a mixed strategy; the second player then picks an action based on this mixed strategy, but not on its realization. We therefore consider randomized learners who predict a distribution $q_t \in \mathcal{Q}$ on every round, where \mathcal{Q} is the set of probability distributions on \mathcal{F} , assumed to be weakly compact. The set of probability distributions on \mathcal{X} (mixed strategies of the adversary) is denoted by \mathcal{P} .

We would like to capture the fact that sequences (x_1, \dots, x_T) cannot be arbitrary. This is achieved by defining restrictions on the adversary, that is, subsets of “allowed” distributions for each round. These restrictions limit the scope of available mixed strategies for the adversary.

Definition 1. A restriction $\mathcal{P}_{1:T}$ on the adversary is a sequence $\mathcal{P}_1, \dots, \mathcal{P}_T$ of mappings $\mathcal{P}_t : \mathcal{X}^{t-1} \mapsto 2^{\mathcal{P}}$ such that $\mathcal{P}_t(x_{1:t-1})$ is a convex subset of \mathcal{P} for any $x_{1:t-1} \in \mathcal{X}^{t-1}$.

Note that the restrictions depend on the past moves of the adversary, but not on those of the player. We will write \mathcal{P}_t instead of $\mathcal{P}_t(x_{1:t-1})$ when $x_{1:t-1}$ is clearly defined. Using the notion of restrictions, we can give names to several types of adversaries that we will study in this paper.

- (1) A *worst-case adversary* is defined by vacuous restrictions $\mathcal{P}_t(x_{1:t-1}) = \mathcal{P}$. That is, any mixed strategy is available to the adversary, including any deterministic point distribution.
- (2) A *constrained adversary* is defined by $\mathcal{P}_t(x_{1:t-1})$ being the set of all distributions supported on the set $\{x \in \mathcal{X} : C_t(x_1, \dots, x_{t-1}, x) = 1\}$ for some deterministic binary-valued constraint C_t . The deterministic constraint can, for instance, ensure that the length of the path determined by the moves x_1, \dots, x_t stays below the allowed budget.
- (3) A *smoothed adversary* picks the worst-case sequence which gets corrupted by i.i.d. noise. Equivalently, we can view this as restrictions on the adversary who chooses the “center” (or a parameter) of the noise distribution.

Using techniques developed in this paper, we can also study the following adversaries (omitted due to lack of space):

- (4) A *hybrid adversary* in the supervised learning game picks the worst-case label y_t , but is forced to draw the x_t -variable from a fixed distribution [6].
- (5) An *i.i.d. adversary* is defined by a time-invariant restriction $\mathcal{P}_t(x_{1:t-1}) = \{p\}$ for every t and some $p \in \mathcal{P}$.

For the given restrictions $\mathcal{P}_{1:T}$, we define the value of the game as

$$\mathcal{V}_T(\mathcal{P}_{1:T}) \triangleq \inf_{q_1 \in \mathcal{Q}} \sup_{p_1 \in \mathcal{P}_1} \mathbb{E}_{f_1, x_1} \inf_{q_2 \in \mathcal{Q}} \sup_{p_2 \in \mathcal{P}_2} \mathbb{E}_{f_2, x_2} \dots \inf_{q_T \in \mathcal{Q}} \sup_{p_T \in \mathcal{P}_T} \mathbb{E}_{f_T, x_T} \left[\sum_{t=1}^T f_t(x_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^T f(x_t) \right] \quad (1)$$

where f_t has distribution q_t and x_t has distribution p_t . As in [10], the adversary is *adaptive*, that is, chooses p_t based on the history of moves $f_{1:t-1}$ and $x_{1:t-1}$. At this point, the only difference from

the setup of [10] is in the restrictions \mathcal{P}_t on the adversary. Because these restrictions might not allow point distributions, suprema over p_t 's in (1) cannot be equivalently written as the suprema over x_t 's.

A word about the notation. In [10], the value of the game is written as $\mathcal{V}_T(\mathcal{F})$, signifying that the main object of study is \mathcal{F} . In [11], it is written as $\mathcal{V}_T(\ell, \Phi_T)$ since the focus is on the complexity of the set of transformations Φ_T and the payoff mapping ℓ . In the present paper, the main focus is indeed on the restrictions on the adversary, justifying our choice $\mathcal{V}_T(\mathcal{P}_{1:T})$ for the notation.

The first step is to apply the minimax theorem. To this end, we verify the necessary conditions. Our assumption that \mathcal{F} is a closed subset of a complete separable metric space implies that \mathcal{Q} is tight and Prokhorov's theorem states that compactness of \mathcal{Q} under weak topology is equivalent to tightness [14]. Compactness under weak topology allows us to proceed as in [10]. Additionally, we require that the restriction sets are compact and convex.

Theorem 1. *Let \mathcal{F} and \mathcal{X} be the sets of moves for the two players, satisfying the necessary conditions for the minimax theorem to hold. Let $\mathcal{P}_{1:T}$ be the restrictions, and assume that for any $x_{1:t-1}$, $\mathcal{P}_t(x_{1:t-1})$ satisfies the necessary conditions for the minimax theorem to hold. Then*

$$\mathcal{V}_T(\mathcal{P}_{1:T}) = \sup_{p_1 \in \mathcal{P}_1} \mathbb{E}_{x_1 \sim p_1} \dots \sup_{p_T \in \mathcal{P}_T} \mathbb{E}_{x_T \sim p_T} \left[\sum_{t=1}^T \inf_{f_t \in \mathcal{F}} \mathbb{E}_{x_t \sim p_t} [f_t(x_t)] - \inf_{f \in \mathcal{F}} \sum_{t=1}^T f(x_t) \right]. \quad (2)$$

The nested sequence of suprema and expected values in Theorem 1 can be re-written succinctly as

$$\begin{aligned} \mathcal{V}_T(\mathcal{P}_{1:T}) &= \sup_{\mathbf{p} \in \mathfrak{P}} \mathbb{E}_{x_1 \sim p_1} \mathbb{E}_{x_2 \sim p_2(\cdot|x_1)} \dots \mathbb{E}_{x_T \sim p_T(\cdot|x_{1:T-1})} \left[\sum_{t=1}^T \inf_{f_t \in \mathcal{F}} \mathbb{E}_{x_t \sim p_t} [f_t(x_t)] - \inf_{f \in \mathcal{F}} \sum_{t=1}^T f(x_t) \right] \\ &= \sup_{\mathbf{p} \in \mathfrak{P}} \mathbb{E} \left[\sum_{t=1}^T \inf_{f_t \in \mathcal{F}} \mathbb{E}_{x_t \sim p_t} [f_t(x_t)] - \inf_{f \in \mathcal{F}} \sum_{t=1}^T f(x_t) \right] \end{aligned} \quad (3)$$

where the supremum is over all joint distributions \mathbf{p} over sequences, such that \mathbf{p} satisfies the restrictions as described below. Given a joint distribution \mathbf{p} on sequences $(x_1, \dots, x_T) \in \mathcal{X}^T$, we denote the associated conditional distributions by $p_t(\cdot|x_{1:t-1})$. We can think of the choice \mathbf{p} as a sequence of oblivious strategies $\{p_t : \mathcal{X}^{t-1} \mapsto \mathcal{P}_t\}_{t=1}^T$, mapping the prefix $x_{1:t-1}$ to a conditional distribution $p_t(\cdot|x_{1:t-1}) \in \mathcal{P}_t(x_{1:t-1})$. We will indeed call \mathbf{p} a ‘‘joint distribution’’ or an ‘‘oblivious strategy’’ interchangeably. We say that a joint distribution \mathbf{p} *satisfies restrictions* if for any t and any $x_{1:t-1} \in \mathcal{X}^{t-1}$, $p_t(\cdot|x_{1:t-1}) \in \mathcal{P}_t(x_{1:t-1})$. The set of all joint distributions satisfying the restrictions is denoted by \mathfrak{P} . We note that Theorem 1 cannot be deduced immediately from the analogous result in [10], as it is not clear how the restrictions on the adversary per each round come into play after applying the minimax theorem. Nevertheless, it is comforting that the restrictions directly translate into the set \mathfrak{P} of oblivious strategies satisfying the restrictions.

Before continuing with our goal of upper-bounding the value of the game, we state the following interesting facts.

Proposition 2. *There is an oblivious minimax optimal strategy for the adversary, and there is a corresponding minimax optimal strategy for the player that does not depend on its own moves.*

The latter statement of the proposition is folklore for worst-case learning, yet we have not seen a proof of it in the literature. The proposition holds for all online learning settings with legal restrictions $\mathcal{P}_{1:T}$, encompassing also the no-restrictions setting of worst-case online learning [10]. The result crucially relies on the fact that the objective is external regret.

3 Symmetrization and Random Averages

Theorem 1 is a useful representation of the value of the game. As the next step, we upper bound it with an expression which is easier to study. Such an expression is obtained by introducing Rademacher random variables. This process can be termed *sequential symmetrization* and has been exploited in [1, 10, 11]. The restrictions \mathcal{P}_t , however, make sequential symmetrization considerably more involved than in the papers cited above. The main difficulty arises from the fact that the set $\mathcal{P}_t(x_{1:t-1})$ depends on the sequence $x_{1:t-1}$, and symmetrization (that is, replacement of x_s with x'_s) has to be done with care as it affects this dependence. Roughly speaking, in the process of symmetrization, a tangent sequence x'_1, x'_2, \dots is introduced such that x_t and x'_t are independent and

identically distributed given “the past”. However, “the past” is itself an interleaving choice of the original sequence and the tangent sequence.

Define the “selector function” $\chi : \mathcal{X} \times \mathcal{X} \times \{\pm 1\} \mapsto \mathcal{X}$ by $\chi(x, x', \epsilon) = x'$ if $\epsilon = 1$ and $\chi(x, x', \epsilon) = x$ if $\epsilon = -1$. When x_t and x'_t are understood from the context, we will use the shorthand $\chi_t(\epsilon) := \chi(x_t, x'_t, \epsilon)$. In other words, χ_t selects between x_t and x'_t depending on the sign of ϵ . Throughout the paper, we deal with binary trees, which arise from symmetrization [10]. Given some set \mathcal{Z} , an \mathcal{Z} -valued tree of depth T is a sequence $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_T)$ of T mappings $\mathbf{z}_i : \{\pm 1\}^{i-1} \mapsto \mathcal{Z}$. The T -tuple $\epsilon = (\epsilon_1, \dots, \epsilon_T) \in \{\pm 1\}^T$ defines a path. For brevity, we write $\mathbf{z}_t(\epsilon)$ instead of $\mathbf{z}_t(\epsilon_{1:t-1})$.

Given a joint distribution \mathbf{p} , consider the “ $(\mathcal{X} \times \mathcal{X})^{T-1} \mapsto \mathcal{P}(\mathcal{X} \times \mathcal{X})$ ”-valued probability tree $\boldsymbol{\rho} = (\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_T)$ defined by

$$\boldsymbol{\rho}_t(\epsilon_{1:t-1})((x_1, x'_1), \dots, (x_{T-1}, x'_{T-1})) = (p_t(\cdot | \chi_1(\epsilon_1), \dots, \chi_{t-1}(\epsilon_{t-1})), p_t(\cdot | \chi_1(\epsilon_1), \dots, \chi_{t-1}(\epsilon_{t-1}))).$$

In other words, the values of the mappings $\boldsymbol{\rho}_t(\epsilon)$ are products of conditional distributions, where conditioning is done with respect to a sequence made from x_s and x'_s depending on the sign of ϵ_s . We note that the difficulty in intermixing the x and x' sequences *does not arise in i.i.d. or worst-case symmetrization*. However, in-between these extremes the notational complexity seems to be unavoidable if we are to employ symmetrization and obtain a version of Rademacher complexity.

As an example, consider the “left-most” path $\epsilon = -1$ in a binary tree of depth T , where $\mathbf{1} = (1, \dots, 1)$ is a T -dimensional vector of ones. Then all the selectors $\chi(x_t, x'_t, \epsilon_t)$ choose the sequence x_1, \dots, x_T . The probability tree $\boldsymbol{\rho}$ on the “left-most” path is, therefore, defined by the conditional distributions $p_t(\cdot | x_{1:t-1})$; on the path $\epsilon = \mathbf{1}$, the conditional distributions are $p_t(\cdot | x'_{1:t-1})$.

Slightly abusing the notation, we will write $\boldsymbol{\rho}_t(\epsilon)((x_1, x'_1), \dots, (x_{t-1}, x'_{t-1}))$ for the probability tree since $\boldsymbol{\rho}_t$ clearly depends only on the prefix up to time $t - 1$. Throughout the paper, it will be understood that the tree $\boldsymbol{\rho}$ is obtained from \mathbf{p} as described above. Since all the conditional distributions of \mathbf{p} satisfy the restrictions, so do the corresponding distributions of the probability tree $\boldsymbol{\rho}$. By saying that $\boldsymbol{\rho}$ satisfies restrictions we then mean that $\mathbf{p} \in \mathfrak{P}$.

Sampling of a pair of \mathcal{X} -valued trees from $\boldsymbol{\rho}$, written as $(\mathbf{x}, \mathbf{x}') \sim \boldsymbol{\rho}$, is defined as the following recursive process: for any $\epsilon \in \{\pm 1\}^T$, $(\mathbf{x}_1(\epsilon), \mathbf{x}'_1(\epsilon)) \sim \boldsymbol{\rho}_1(\epsilon)$ and

$$(\mathbf{x}_t(\epsilon), \mathbf{x}'_t(\epsilon)) \sim \boldsymbol{\rho}_t(\epsilon)((\mathbf{x}_1(\epsilon), \mathbf{x}'_1(\epsilon)), \dots, (\mathbf{x}_{t-1}(\epsilon), \mathbf{x}'_{t-1}(\epsilon))) \quad \text{for } 2 \leq t \leq T \quad (4)$$

To gain a better understanding of the sampling process, consider the first few levels of the tree. The roots $\mathbf{x}_1, \mathbf{x}'_1$ of the trees \mathbf{x}, \mathbf{x}' are sampled from p_1 , the conditional distribution for $t = 1$ given by \mathbf{p} . Next, say, $\epsilon_1 = +1$. Then the “right” children of \mathbf{x}_1 and \mathbf{x}'_1 are sampled via $\mathbf{x}_2(+1), \mathbf{x}'_2(+1) \sim p_2(\cdot | \mathbf{x}'_1)$ since $\chi_1(+1)$ selects \mathbf{x}'_1 . On the other hand, the “left” children $\mathbf{x}_2(-1), \mathbf{x}'_2(-1)$ are both distributed according to $p_2(\cdot | \mathbf{x}_1)$. Now, suppose $\epsilon_1 = +1$ and $\epsilon_2 = -1$. Then, $\mathbf{x}_3(+1, -1), \mathbf{x}'_3(+1, -1)$ are both sampled from $p_3(\cdot | \mathbf{x}'_1, \mathbf{x}_2(+1))$.

The proof of Theorem 3 reveals why such intricate conditional structure arises, and Proposition 5 below shows that this structure greatly simplifies for i.i.d. and worst-case situations. Nevertheless, the process described above allows us to define a *unified notion of Rademacher complexity* for the spectrum of assumptions between the two extremes.

Definition 2. The *distribution-dependent sequential Rademacher complexity* of a function class $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{X}}$ is defined as

$$\mathfrak{R}_T(\mathcal{F}, \mathbf{p}) \triangleq \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \boldsymbol{\rho}} \mathbb{E}_{\epsilon} \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t f(\mathbf{x}_t(\epsilon)) \right]$$

where $\epsilon = (\epsilon_1, \dots, \epsilon_T)$ is a sequence of i.i.d. Rademacher random variables and $\boldsymbol{\rho}$ is the probability tree associated with \mathbf{p} .

We now prove an upper bound on the value $\mathcal{V}_T(\mathcal{P}_{1:T})$ of the game in terms of this distribution-dependent sequential Rademacher complexity. The result cannot be deduced directly from [10], and it greatly increases the scope of problems whose learnability can now be studied in a unified manner.

Theorem 3. *The minimax value is bounded as*

$$\mathcal{V}_T(\mathcal{P}_{1:T}) \leq 2 \sup_{\mathbf{p} \in \mathfrak{P}} \mathfrak{R}_T(\mathcal{F}, \mathbf{p}). \quad (5)$$

More generally, for any measurable function M_t such that $M_t(\mathbf{p}, f, \mathbf{x}, \mathbf{x}', \epsilon) = M_t(\mathbf{p}, f, \mathbf{x}', \mathbf{x}, -\epsilon)$,

$$\mathcal{V}_T(\mathcal{P}_{1:T}) \leq 2 \sup_{\mathbf{p} \in \mathfrak{P}} \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \rho} \mathbb{E}_\epsilon \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t (f(\mathbf{x}_t(\epsilon)) - M_t(\mathbf{p}, f, \mathbf{x}, \mathbf{x}', \epsilon)) \right]$$

The following corollary provides a natural ‘‘centered’’ version of the distribution-dependent Rademacher complexity. That is, the complexity can be measured by relative shifts in the adversarial moves.

Corollary 4. *For the game with restrictions $\mathcal{P}_{1:T}$,*

$$\mathcal{V}_T(\mathcal{P}_{1:T}) \leq 2 \sup_{\mathbf{p} \in \mathfrak{P}} \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \rho} \mathbb{E}_\epsilon \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t \left(f(\mathbf{x}_t(\epsilon)) - \mathbb{E}_{t-1} f(\mathbf{x}_t(\epsilon)) \right) \right]$$

where \mathbb{E}_{t-1} denotes the conditional expectation of $\mathbf{x}_t(\epsilon)$.

Example 1. *Suppose \mathcal{F} is a unit ball in a Banach space and $f(x) = \langle f, x \rangle$. Then*

$$\mathcal{V}_T(\mathcal{P}_{1:T}) \leq 2 \sup_{\mathbf{p} \in \mathfrak{P}} \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \rho} \mathbb{E}_\epsilon \left\| \sum_{t=1}^T \epsilon_t \left(\mathbf{x}_t(\epsilon) - \mathbb{E}_{t-1} \mathbf{x}_t(\epsilon) \right) \right\|$$

Suppose the adversary plays a simple random walk (e.g., $p_t(x|x_1, \dots, x_{t-1}) = p_t(x|x_{t-1})$ is uniform on a unit sphere). For simplicity, suppose this is the only strategy allowed by the set \mathfrak{P} . Then $\mathbf{x}_t(\epsilon) - \mathbb{E}_{t-1} \mathbf{x}_t(\epsilon)$ are independent increments when conditioned on the history. Further, the increments do not depend on ϵ_t . Thus, $\mathcal{V}_T(\mathcal{P}_{1:T}) \leq 2\mathbb{E} \left\| \sum_{t=1}^T Y_t \right\|$ where $\{Y_t\}$ is the corresponding random walk.

We now show that the distribution-dependent sequential Rademacher complexity for *i.i.d.* data is precisely the classical Rademacher complexity, and further show that the distribution-dependent sequential Rademacher complexity is always upper bounded by the worst-case sequential Rademacher complexity defined in [10].

Proposition 5. *First, consider the *i.i.d.* restrictions $\mathcal{P}_t = \{p\}$ for all t , where p is some fixed distribution on \mathcal{X} , and let ρ be the process associated with the joint distribution $\mathbf{p} = p^T$. Then*

$$\mathfrak{R}_T(\mathcal{F}, \mathbf{p}) = \mathfrak{R}_T(\mathcal{F}, p), \quad \text{where} \quad \mathfrak{R}_T(\mathcal{F}, p) \triangleq \mathbb{E}_{x_1, \dots, x_T \sim p} \mathbb{E}_\epsilon \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t f(x_t) \right] \quad (6)$$

is the classical Rademacher complexity. Second, for any joint distribution \mathbf{p} ,

$$\mathfrak{R}_T(\mathcal{F}, \mathbf{p}) \leq \mathfrak{R}_T(\mathcal{F}), \quad \text{where} \quad \mathfrak{R}_T(\mathcal{F}) \triangleq \sup_{\mathbf{x}} \mathbb{E}_\epsilon \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t f(\mathbf{x}_t(\epsilon)) \right] \quad (7)$$

is the sequential Rademacher complexity defined in [10].

In the case of hybrid learning, adversary chooses a sequence of pairs (x_t, y_t) where the instance x_t 's are *i.i.d.* but the labels y_i 's are fully adversarial. The distribution-dependent Rademacher complexity in such a hybrid case can be upper bounded by a vary natural quantity: a random average where expectation is taken over x_t 's and a supremum over \mathcal{Y} -valued trees. So, the distribution dependent Rademacher complexity itself becomes a hybrid between the classical Rademacher complexity and the worst case sequential Rademacher complexity. For more details, see Lemma 17 in the Appendix as another example of an analysis of the distribution-dependent sequential Rademacher complexity.

Distribution-dependent sequential Rademacher complexity enjoys many of the nice properties satisfied by both classical and worst-case Rademacher complexities. As shown in [10], these properties are handy tools for proving upper bounds on the value in various examples. We have: (a) If $\mathcal{F} \subset \mathcal{G}$, then $\mathfrak{R}(\mathcal{F}, \mathbf{p}) \leq \mathfrak{R}(\mathcal{G}, \mathbf{p})$; (b) $\mathfrak{R}(\mathcal{F}, \mathbf{p}) = \mathfrak{R}(\text{conv}(\mathcal{F}), \mathbf{p})$; (c) $\mathfrak{R}(c\mathcal{F}, \mathbf{p}) = |c| \mathfrak{R}(\mathcal{F}, \mathbf{p})$ for all $c \in \mathbb{R}$; (d) For any h , $\mathfrak{R}(\mathcal{F} + h, \mathbf{p}) = \mathfrak{R}(\mathcal{F}, \mathbf{p})$ where $\mathcal{F} + h = \{f + h : f \in \mathcal{F}\}$.

In addition to the above properties, upper bounds on $\mathfrak{R}(\mathcal{F}, \mathbf{p})$ can be derived via sequential covering numbers defined in [10]. This notion of a cover captures the sequential complexity of a function class on a given \mathcal{X} -valued tree \mathbf{x} . One can then show an analogue of the Dudley integral bound, where the complexity is averaged with respect to the underlying process $(\mathbf{x}, \mathbf{x}') \sim \rho$.

4 Application: Constrained Adversaries

In this section, we consider adversaries who are deterministically constrained in the sequences of actions they can play. It is often useful to consider scenarios where the adversary is worst case, yet has some *budget* or *constraint* to satisfy while picking the actions. Examples of such scenarios include, for instance, games where the adversary is constrained to make moves that are close in some fashion to the previous move, linear games with bounded variance, and so on. Below we formulate such games quite generally through arbitrary constraints that the adversary has to satisfy on each round. We easily derive several results to illustrate the versatility of the developed framework.

For a T round game consider an adversary who is only allowed to play sequences x_1, \dots, x_T such that at round t the constraint $C_t(x_1, \dots, x_t) = 1$ is satisfied, where $C_t : \mathcal{X}^t \mapsto \{0, 1\}$ represents the constraint on the sequence played so far. The constrained adversary can be viewed as a stochastic adversary with restrictions on the conditional distribution at time t given by the set of all Borel distributions on the set $\mathcal{X}_t(x_{1:t-1}) \triangleq \{x \in \mathcal{X} : C_t(x_1, \dots, x_{t-1}, x) = 1\}$. Since this set includes all point distributions on each $x \in \mathcal{X}_t$, the sequential complexity simplifies in a way similar to worst-case adversaries. We write $\mathcal{V}_T(C_{1:T})$ for the value of the game with the given constraints. Now, assume that for any $x_{1:t-1}$, the set of all distributions on $\mathcal{X}_t(x_{1:t-1})$ is weakly compact in a way similar to compactness of \mathcal{P} . That is, $\mathcal{P}_t(x_{1:t-1})$ satisfy the necessary conditions for the minimax theorem to hold. We have the following corollaries of Theorems 1 and 3.

Corollary 6. *Let \mathcal{F} and \mathcal{X} be the sets of moves for the two players, satisfying the necessary conditions for the minimax theorem to hold. Let $\{C_t : \mathcal{X}^{t-1} \mapsto \{0, 1\}\}_{t=1}^T$ be the constraints. Then*

$$\mathcal{V}_T(C_{1:T}) = \sup_{\mathbf{p} \in \mathfrak{P}} \mathbb{E} \left[\sum_{t=1}^T \inf_{f_t \in \mathcal{F}} \mathbb{E}_{x_t \sim p_t} [f_t(x_t)] - \inf_{f \in \mathcal{F}} \sum_{t=1}^T f(x_t) \right] \quad (8)$$

where \mathbf{p} ranges over all distributions over sequences (x_1, \dots, x_T) such that $\forall t, C_t(x_{1:t-1}) = 1$.

Corollary 7. *Let the set \mathcal{T} be a set of pairs $(\mathbf{x}, \mathbf{x}')$ of \mathcal{X} -valued trees with the property that for any $\epsilon \in \{\pm 1\}^T$ and any $t \in [T]$, $C(\chi_1(\epsilon_1), \dots, \chi_{t-1}(\epsilon_{t-1}), \mathbf{x}_t(\epsilon)) = C(\chi_1(\epsilon_1), \dots, \chi_{t-1}(\epsilon_{t-1}), \mathbf{x}'_t(\epsilon)) = 1$. The minimax value is bounded as*

$$\mathcal{V}_T(C_{1:T}) \leq 2 \sup_{(\mathbf{x}, \mathbf{x}') \in \mathcal{T}} \mathfrak{R}_T(\mathcal{F}, \mathbf{p}).$$

More generally, for any measurable function M_t such that $M_t(f, \mathbf{x}, \mathbf{x}', \epsilon) = M_t(f, \mathbf{x}', \mathbf{x}, -\epsilon)$,

$$\mathcal{V}_T(C_{1:T}) \leq 2 \sup_{(\mathbf{x}, \mathbf{x}') \in \mathcal{T}} \mathbb{E}_\epsilon \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t (f(\mathbf{x}_t(\epsilon)) - M_t(f, \mathbf{x}, \mathbf{x}', \epsilon)) \right].$$

Armed with these results, we can recover and extend some known results on online learning against budgeted adversaries. The first result says that if the adversary is not allowed to move by more than σ_t away from its previous average of decisions, the player has a strategy to exploit this fact and obtain lower regret. For the ℓ_2 -norm, such “total variation” bounds have been achieved in [4] up to a $\log T$ factor. Our analysis seamlessly incorporates variance measured in arbitrary norms, not just ℓ_2 . We emphasize that such certificates of learnability are not possible with the analysis of [10].

Proposition 8 (Variance Bound). *Consider the online linear optimization setting with $\mathcal{F} = \{f : \Psi(f) \leq R^2\}$ for a λ -strongly function $\Psi : \mathcal{F} \mapsto \mathbb{R}_+$ on \mathcal{F} , and $\mathcal{X} = \{x : \|x\|_* \leq 1\}$. Let $f(x) = \langle f, x \rangle$ for any $f \in \mathcal{F}$ and $x \in \mathcal{X}$. Consider the sequence of constraints $\{C_t\}_{t=1}^T$ given by $C_t(x_1, \dots, x_{t-1}, x) = 1$ if $\|x - \frac{1}{t-1} \sum_{\tau=1}^{t-1} x_\tau\|_* \leq \sigma_t$ and 0 otherwise. Then*

$$\mathcal{V}_T(C_{1:T}) \leq 2\sqrt{2}R \sqrt{\lambda^{-1} \sum_{t=1}^T \sigma_t^2}$$

In particular, we obtain the following ℓ_2 variance bound. Consider the case when $\Psi : \mathcal{F} \mapsto \mathbb{R}_+$ is given by $\Psi(f) = \frac{1}{2} \|f\|^2$, $\mathcal{F} = \{f : \|f\|_2 \leq 1\}$ and $\mathcal{X} = \{x : \|x\|_2 \leq 1\}$. Consider the constrained game where the move x_t played by adversary at time t satisfies $\left\| x_t - \frac{1}{t-1} \sum_{\tau=1}^{t-1} x_\tau \right\|_2 \leq \sigma_t$. In this case we can conclude that $\mathcal{V}_T(C_{1:T}) \leq 2\sqrt{2} \sqrt{\sum_{t=1}^T \sigma_t^2}$. We can also derive a variance bound

over the simplex. Let $\Psi(f) = \sum_{i=1}^d f_i \log(df_i)$ is defined over the d -simplex \mathcal{F} , and $\mathcal{X} = \{x : \|x\|_\infty \leq 1\}$. Consider the constrained game where the move x_t played by adversary at time t satisfies $\max_{j \in [d]} \left| x_t[j] - \frac{1}{t-1} \sum_{\tau=1}^{t-1} x_\tau[j] \right| \leq \sigma_t$. For any $f \in \mathcal{F}$, $\Psi(f) \leq \log(d)$ and so we conclude that $\mathcal{V}_T(C_{1:T}) \leq 2\sqrt{2} \sqrt{\log(d) \sum_{t=1}^T \sigma_t^2}$.

The next Proposition gives a bound whenever the adversary is constrained to choose his decision from a small ball around the previous decision.

Proposition 9 (Slowly-Changing Decisions). *Consider the online linear optimization setting where adversary's move at any time is close to the move during the previous time step. Let $\mathcal{F} = \{f : \Psi(f) \leq R^2\}$ where $\Psi : \mathcal{F} \mapsto \mathbb{R}_+$ is a λ -strongly function on \mathcal{F} and $\mathcal{X} = \{x : \|x\|_* \leq B\}$. Let $f(x) = \langle f, x \rangle$ for any $f \in \mathcal{F}$ and $x \in \mathcal{X}$. Consider the sequence of constraints $\{C_t\}_{t=1}^T$ given by $C_t(x_1, \dots, x_{t-1}, x) = 1$ if $\|x - x_{t-1}\|_* \leq \delta$ and 0 otherwise. Then,*

$$\mathcal{V}_T(C_{1:T}) \leq 2R\delta \sqrt{2T/\lambda}.$$

In particular, consider the case of a Euclidean-norm restriction on the moves. Let $\Psi : \mathcal{F} \mapsto \mathbb{R}_+$ is given by $\Psi(f) = \frac{1}{2} \|f\|^2$, $\mathcal{F} = \{f : \|f\|_2 \leq 1\}$ and $\mathcal{X} = \{x : \|x\|_2 \leq 1\}$. Consider the constrained game where the move x_t played by adversary at time t satisfies $\|x_t - x_{t-1}\|_2 \leq \delta$. In this case we can conclude that $\mathcal{V}_T(C_{1:T}) \leq 2\delta\sqrt{2T}$. For the case of decision-making on the simplex, we obtain the following result. Let $\Psi(f) = \sum_{i=1}^d f_i \log(df_i)$ is defined over the d -simplex \mathcal{F} , and $\mathcal{X} = \{x : \|x\|_\infty \leq 1\}$. Consider the constrained game where the move x_t played by adversary at time t satisfies $\|x_t - x_{t-1}\|_\infty \leq \delta$. In this case note that for any $f \in \mathcal{F}$, $\Psi(f) \leq \log(d)$ and so we can conclude that $\mathcal{V}_T(C_{1:T}) \leq 2\delta\sqrt{2T \log(d)}$.

5 Application: Smoothed Adversaries

The development of *smoothed analysis* over the past decade is arguably one of the landmarks in the study of complexity of algorithms. In contrast to the overly optimistic *average complexity* and the overly pessimistic *worst-case complexity*, smoothed complexity can be seen as a more realistic measure of algorithm's performance. In their groundbreaking work, Spielman and Teng [13] showed that the smoothed running time complexity of the simplex method is polynomial. This result explains good performance of the method in practice despite its exponential-time worst-case complexity. In this section, we consider the effect of smoothing on *learnability*.

It is well-known that there is a gap between the i.i.d. and the worst-case scenarios. In fact, we do not need to go far for an example: a simple class of threshold functions on a unit interval is learnable in the i.i.d. supervised learning scenario, yet difficult in the online worst-case model [8, 2, 9]. This fact is reflected in the corresponding combinatorial dimensions: the Vapnik-Chervonenkis dimension is one, whereas the Littlestone dimension is infinite. The proof of the latter fact, however, reveals that the infinite number of mistakes on the part of the player is due to the infinite resolution of the carefully chosen adversarial sequence. We can argue that this infinite precision is an unreasonable assumption on the power of a real-world opponent. The idea of limiting the power of the malicious adversary through perturbing the sequence can be traced back to Posner and Kulkarni [9]. The authors considered on-line learning of functions of bounded variation, but in the so-called realizable setting (that is, when labels are given by some function in the given class).

We define the *smoothed online learning model* as the following T -round interaction between the learner and the adversary. On round t , the learner chooses $f_t \in \mathcal{F}$; the adversary simultaneously chooses $x_t \in \mathcal{X}$, which is then perturbed by some noise $s_t \sim \sigma$, yielding a value $\tilde{x}_t = \omega(x_t, s_t)$; and the player suffers $f_t(\tilde{x}_t)$. Regret is defined with respect to the perturbed sequence. Here $\omega : \mathcal{X} \times S \mapsto \mathcal{X}$ is some measurable mapping; for instance, additive disturbances can be written as $\tilde{x} = \omega(x, s) = x + s$. If ω keeps x_t unchanged, that is $\omega(x_t, s_t) = x_t$, the setting is precisely the standard online learning model. In the full information version, we assume that the choice \tilde{x}_t is revealed to the player at the end of round t . We now recognize that the setting is nothing but a particular way to restrict the adversary. That is, the choice $x_t \in \mathcal{X}$ defines a parameter of a mixed strategy from which a actual move $\omega(x_t, s_t)$ is drawn; for instance, for additive zero-mean Gaussian noise, x_t defines the center of the distribution from which $x_t + s_t$ is drawn. In other words, noise does not allow the adversary to play any desired mixed strategy.

The value of the smoothed online learning game (as defined in (1)) can be equivalently written as

$$\mathcal{V}_T = \inf_{q_1} \sup_{x_1} \mathbb{E}_{\substack{f_1 \sim q_1 \\ s_1 \sim \sigma}} \inf_{q_2} \sup_{x_2} \mathbb{E}_{\substack{f_2 \sim q_2 \\ s_2 \sim \sigma}} \cdots \inf_{q_T} \sup_{x_T} \mathbb{E}_{\substack{f_T \sim q_T \\ s_T \sim \sigma}} \left[\sum_{t=1}^T f_t(\omega(x_t, s_t)) - \inf_{f \in \mathcal{F}} \sum_{t=1}^T f(\omega(x_t, s_t)) \right]$$

where the infima are over $q_t \in \mathcal{Q}$ and the suprema are over $x_t \in \mathcal{X}$. Using sequential symmetrization, we deduce the following upper bound on the value of the smoothed online learning game.

Theorem 10. *The value of the smoothed online learning game is bounded above as*

$$\mathcal{V}_T \leq 2 \sup_{x_1 \in \mathcal{X}} \mathbb{E}_{s_1 \sim \sigma} \mathbb{E}_{\epsilon_1} \cdots \sup_{x_T \in \mathcal{X}} \mathbb{E}_{s_T \sim \sigma} \mathbb{E}_{\epsilon_T} \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t f(\omega(x_t, s_t)) \right]$$

We now demonstrate how Theorem 10 can be used to show learnability for smoothed learning of threshold functions. First, consider the supervised game with threshold functions on a unit interval (that is, non-homogenous hyperplanes). The moves of the adversary are pairs $x = (z, y)$ with $z \in [0, 1]$ and $y \in \{0, 1\}$, and the binary-valued function class \mathcal{F} is defined by

$$\mathcal{F} = \{f_\theta(z, y) = |y - \mathbf{1}\{z < \theta\}| : \theta \in [0, 1]\}, \quad (9)$$

that is, every function is associated with a threshold $\theta \in [0, 1]$. The class \mathcal{F} has infinite Littlestone's dimension and is not learnable in the worst-case online framework. Consider a smoothed scenario, with the z -variable of the adversarial move (z, y) perturbed by an additive uniform noise $\sigma = \text{Unif}[-\gamma/2, \gamma/2]$ for some $\gamma \geq 0$. That is, the actual move revealed to the player at time t is $(z_t + s_t, y_t)$, with $s_t \sim \sigma$. Any non-trivial upper bound on regret has to depend on particular noise assumptions, as $\gamma = 0$ corresponds to the case with infinite Littlestone dimension. For the uniform disturbance, the intuition tells us that noise implies a margin, and we should expect a $1/\gamma$ complexity parameter appearing in the bounds. The next lemma quantifies the intuition that additive noise limits precision of the adversary.

Lemma 11. *Let $\theta_1, \dots, \theta_N$ be obtained by discretizing the interval $[0, 1]$ into $N = T^a$ bins $[\theta_i, \theta_{i+1})$ of length T^{-a} , for some $a \geq 3$. Then, for any sequence $z_1, \dots, z_T \in [0, 1]$, with probability at least $1 - \frac{1}{\gamma T^{a-2}}$, no two elements of the sequence $z_1 + s_1, \dots, z_T + s_T$ belong to the same interval $[\theta_i, \theta_{i+1})$, where s_1, \dots, s_T are i.i.d. $\text{Unif}[-\gamma/2, \gamma/2]$.*

We now observe that, conditioned on the event in Lemma 11, the upper bound on the value in Theorem 10 is a supremum of N martingale difference sequences! We then arrive at:

Proposition 12. *For the problem of smoothed online learning of thresholds in 1-D, the value is*

$$\mathcal{V}_T \leq 2 + \sqrt{2T(4 \log T + \log(1/\gamma))}$$

What we found is somewhat surprising: for a problem which is not learnable in the online worst-case scenario, an exponentially small noise added to the moves of the adversary yields a learnable problem. This shows, at least in the given example, that the worst-case analysis and Littlestone's dimension are brittle notions which might be too restrictive in the real world, where some noise is unavoidable. It is comforting that small additive noise makes the problem learnable!

The proof for smoothed learning of half-spaces in higher dimension follows the same route as the one-dimensional exposition. For simplicity, assume the hyperplanes are homogenous and $\mathcal{Z} = S_{d-1} \subset \mathbb{R}^d$, $\mathcal{Y} = \{-1, 1\}$, $\mathcal{X} = \mathcal{Z} \times \mathcal{Y}$. Define $\mathcal{F} = \{f_\theta(z, y) = \mathbf{1}\{y \langle z, \theta \rangle > 0\} : \theta \in S_{d-1}\}$, and assume that the noise is distributed uniformly on a square patch with side-length γ on the surface of the sphere S_{d-1} . We can also consider other distributions, possibly with support on a d -dimensional ball instead.

Proposition 13. *For the problem of smoothed online learning of half-spaces,*

$$\mathcal{V}_T = O \left(\sqrt{dT \left(\log \left(\frac{1}{\gamma} \right) + \frac{3}{d-1} \log T \right)} + v_{d-2} \cdot \left(\frac{1}{\gamma} \right)^{\frac{3}{d-1}} \right)$$

where v_{d-2} is constant depending only on the dimension d .

We conclude that half spaces are online learnable in the smoothed model, since the upper bound of Proposition 13 guarantees existence of an algorithm which achieves this regret. In fact, for the two examples considered in this section, the Exponential Weights Algorithm on the discretization given by Lemma 11 is a (computationally infeasible) algorithm achieving the bound.

References

- [1] J. Abernethy, A. Agarwal, P. Bartlett, and A. Rakhlin. A stochastic view of optimal regret through minimax duality. In *COLT*, 2009.
- [2] S. Ben-David, D. Pal, and S. Shalev-Shwartz. Agnostic online learning. In *Proceedings of the 22th Annual Conference on Learning Theory*, 2009.
- [3] J.O. Berger. *Statistical decision theory and Bayesian analysis*. Springer, 1985.
- [4] E. Hazan and S. Kale. Better algorithms for benign bandits. In *SODA*, 2009.
- [5] S.M. Kakade, K. Sridharan, and A. Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. *NIPS*, 22, 2008.
- [6] A. Lazaric and R. Munos. Hybrid Stochastic-Adversarial On-line Learning. In *COLT*, 2009.
- [7] M. Ledoux and M. Talagrand. *Probability in Banach Spaces*. Springer-Verlag, New York, 1991.
- [8] N. Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 04 1988.
- [9] S. Posner and S. Kulkarni. On-line learning of functions of bounded variation under various sampling schemes. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 439–445. ACM, 1993.
- [10] A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Random averages, combinatorial parameters, and learnability. In *NIPS*, 2010. Full version available at arXiv:1006.1138.
- [11] A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Beyond regret. In *COLT*, 2011. Full version available at arXiv:1011.3168.
- [12] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability, stability and uniform convergence. *JMLR*, 11:2635–2670, Oct 2010.
- [13] D. A. Spielman and S. H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51(3):385–463, 2004.
- [14] A. W. Van Der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes : With Applications to Statistics*. Springer Series, March 1996.