

Multiclass Online Learning and Uniform Convergence

Steve Hanneke

Purdue University

STEVE.HANNEKE@GMAIL.COM

Shay Moran

Technion and Google Research

SMORAN@TECHNION.AC.IL

Vinod Raman

University of Michigan

VKRAMAN@UMICH.EDU

Unique Subedi

University of Michigan

SUBEDI@UMICH.EDU

Ambuj Tewari

University of Michigan

TEWARIA@UMICH.EDU

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

We study multiclass classification in the agnostic adversarial online learning setting. As our main result, we prove that any multiclass concept class is agnostically learnable if and only if its Littlestone dimension is finite. This solves an open problem studied by Daniely, Sabato, Ben-David, and Shalev-Shwartz (2011,2015) who handled the case when the number of classes (or labels) is bounded. We also prove a separation between online learnability and online uniform convergence by exhibiting an easy-to-learn class whose sequential Rademacher complexity is unbounded.

Our learning algorithm uses the multiplicative weights algorithm, with a set of experts defined by executions of the Standard Optimal Algorithm on subsequences of size Littlestone dimension. We argue that the best expert has regret at most Littlestone dimension relative to the best concept in the class. This differs from the well-known covering technique of Ben-David, Pál, and Shalev-Shwartz (2009) for binary classification, where the best expert has regret zero.

Keywords: Online Learning, Multiclass Classification, Agnostic Learning, Regret Bound, Littlestone Dimension, Learnability

1. Introduction

Many important machine learning tasks involve a large prediction space; for instance, in language models the prediction space corresponds to the language size (number of words). Other examples include recommendation systems, image object recognition, protein folding prediction, and more.

Consequently, multiclass prediction problems have been studied extensively in the literature. [Natarajan and Tadepalli \(1988\)](#) and [Natarajan \(1989\)](#) initiated the study of multiclass prediction in the basic PAC setting. They characterized the concept classes that satisfy uniform convergence via a natural combinatorial parameter called the graph dimension. They further characterized PAC learnability in the case when the number of labels is bounded via another combinatorial parameter called the Natarajan dimension. Whether the Natarajan dimension characterizes learnability in general (i.e., even when the number of labels can be infinite) has remained open, until recently [Brukhim, Carmon, Dinur, Moran, and Yehudayoff \(2022\)](#) exhibited an unlearnable concept class with Natarajan dimension 1. [Brukhim et al. \(2022\)](#) showed that multiclass PAC learnability is in

fact captured by a different combinatorial parameter which they called the Daniely Shalev-Shwartz (DS) dimension, after [Daniely and Shalev-Shwartz \(2014\)](#) who defined it.

Remarkably, the nature of multiclass PAC learnability with a bounded label space is very different than the one when the number of labels is infinite. For instance, in the former PAC learnability and uniform convergence are equivalent¹. In contrast, already in the 80's, [Natarajan \(1988\)](#) demonstrated an easy-to-learn concept class over an infinite label space which does not satisfy uniform convergence. More recently, [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#); [Daniely and Shalev-Shwartz \(2014\)](#) studied variants of the ERM principle in multiclass learning, and even demonstrated a PAC learnable class which cannot be learned properly (i.e., by learners whose hypothesis is always an element of the concept class).

How about agnostic versus realizable PAC learning? Here, it turns out that the two are equivalent; that is, any class that is learnable in the realizable case is also learnable in the agnostic case. This was shown by [David, Moran, and Yehudayoff \(2016\)](#) using sample compression arguments.

Perhaps surprisingly, the corresponding questions in the setting of online multiclass classification are still open. [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#) initiated the study of online multiclass classification and show that, like in the binary case, the Littlestone dimension characterizes learnability in the realizable setting. They also studied the agnostic setting, and showed that when the number of labels is finite, then agnostic learnability is equivalent to realizable case learnability. They left open whether this equivalence extends to unbounded label space (or equivalently whether the dependence on the number of labels in the optimal regret can be removed). In this work we resolve this question by showing that agnostic- and realizable-case learnability remain equivalent, even when the number of labels is infinite.

How about uniform convergence versus learnability? In the past 15 years an online analogue of uniform convergence has emerged from the introduction of the sequential Rademacher complexity ([Rakhlin, Sridharan, and Tewari, 2010, 2015a](#)), and of the adversarial laws of large numbers framework ([Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev, 2021a](#)). This raises the question whether online uniform convergence and online learnability are equivalent. When the number of labels is finite, known results imply that indeed the two are equivalent. How about classes with an infinite number of labels? In this work we resolve this question by introducing a combinatorial parameter which we call the *sequential graph dimension* that characterizes online uniform convergence in the multiclass setting. Furthermore we identify an online learnable class with an unbounded sequential graph dimension, thus separating online uniform convergence from online learnability.

In both the PAC and online settings, our interest in studying unbounded label spaces has multiple motivations. One main interest is in establishing sharp enough guarantees to reflect the intuitive fact that the optimal performance should not inherently depend on the number of possible labels: that is, that the latter has no explicit significance to the optimal sample complexity (in PAC learning) or regret (in online learning), even when finite. As many modern learning problems have enormous label spaces (e.g., face recognition), this is quite relevant, and has recently been studied in the machine learning literature under the name “extreme classification”, where the number of possible labels is vast (possibly exceeding even the data set size). More abstractly, often in mathematics it is the case that infinities clarify concepts and phenomena (e.g., the notion of continuity, based on limits). Similarly, focusing on infinite label spaces is a natural way to abstract away an irrelevant detail, which helps to clarify what is the “correct” way to approach the problem.

1. This equivalence yields the fundamental *empirical risk minimization principle* in PAC learning.

1.1. Setting and Summary of Main Results

We begin with the basic setup. Let \mathcal{X} and \mathcal{Y} be arbitrary non-empty sets, called the *instance space* and *label space*, respectively. Let $\mathbb{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be an arbitrary set of functions, called the *concept class*. Our results will hold for any $(\mathcal{X}, \mathcal{Y}, \mathbb{H})$, and hence are fully general.

Denote by $\Pi(\mathcal{Y})$ the set of probability measures on \mathcal{Y} .² A *learning algorithm* is a mapping $\mathbb{A} : (\mathcal{X} \times \mathcal{Y})^* \times \mathcal{X} \rightarrow \Pi(\mathcal{Y})$. Intuitively, $\mathbb{A}(X_1, Y_1, \dots, X_{t-1}, Y_{t-1}, X_t)$ outputs a prediction \hat{y}_t of the label Y_t of the test point X_t , after observing the history $(X_1, Y_1), \dots, (X_{t-1}, Y_{t-1})$ and the test point X_t . The output space $\Pi(\mathcal{Y})$ generalizes this to allow for randomized algorithms: that is, the predicted label \hat{y}_t may be *randomized*. This is known to be necessary for agnostic online learning [Cesa-Bianchi and Lugosi \(2006\)](#). For simplicity, when \mathbb{A} and $X_1, Y_1, \dots, X_{t-1}, Y_{t-1}, X_t$ are clear from the context, we denote by $p_t \in \Pi(\mathcal{Y})$ the output $\mathbb{A}(X_1, Y_1, \dots, X_{t-1}, Y_{t-1}, X_t)$, and for any $y \in \mathcal{Y}$ we denote by $p_t(y)$ the probability of the singleton set $\{y\}$ under the probability measure p_t .

For any $T \in \mathbb{N}$, called the *time horizon*, the *regret* of a learning algorithm \mathbb{A} is defined as

$$\text{regret}(\mathbb{A}, T) = \sup_{(X_1, Y_1), \dots, (X_T, Y_T)} \left(\sum_{t=1}^T (1 - p_t(Y_t)) \right) - \left(\min_{h \in \mathbb{H}} \sum_{t=1}^T \mathbb{1}[h(X_t) \neq Y_t] \right).$$

As is well-known, this may be interpreted as the worst-case value of the difference between the expected number of mistakes the algorithm \mathbb{A} makes in its predictions \hat{y}_t (i.e., how many times t have $\hat{y}_t \neq Y_t$) and the number of mistakes made by the best function in the class \mathbb{H} (where “best” means the function making fewest mistakes on the sequence).

This is often also interpreted as a sequential game between the learner and an adversary. On each round t , the adversary first chooses a value for X_t , the learner observes this value and chooses a probability measure p_t ; the adversary observes this p_t and chooses a class label Y_t , and the learner suffers a loss $1 - p_t(Y_t)$, representing the probability that its randomized prediction \hat{y}_t is incorrect. The overall objective of the learner is to achieve low regret on this sequence of plays, that is, $\left(\sum_{t=1}^T (1 - p_t(Y_t)) \right) - \left(\min_{h \in \mathbb{H}} \sum_{t=1}^T \mathbb{1}[h(X_t) \neq Y_t] \right)$, whereas the adversary’s objective is to maximize this quantity. $\text{regret}(\mathbb{A}, T)$ represents the value of this objective when the adversary plays optimally against the algorithm \mathbb{A} .

A concept class \mathbb{H} is said to be *agnostically online learnable* if

$$\inf_{\mathbb{A}} \text{regret}(\mathbb{A}, T) = o(T).$$

The key quantity of interest in online learning (be it binary or multiclass) is the *Littlestone dimension*, $L(\mathbb{H})$ (see Section 2 for the definition), whose finiteness is known to be necessary and sufficient for online learnability in the *realizable* case (i.e., when the subtracted term in $\text{regret}(\mathbb{A}, T)$ is zero). [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#) effectively asked the following question regarding whether this fact extends to the agnostic setting:

Is any concept class \mathbb{H} agnostically online learnable if and only if $L(\mathbb{H}) < \infty$?

[Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#) proved the necessity direction (which, as they note, follows readily from arguments of [Littlestone, 1988](#)), that is, that any class

2. The associated σ -algebra is of little consequence here, except that singleton sets $\{y\}$ should be measurable.

with $L(\mathbb{H}) = \infty$ is not online learnable (indeed, even in the realizable case). [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2015\)](#) further establish a lower bound when $L(\mathbb{H}) < \infty$: for any \mathbb{A} and any $T \geq L(\mathbb{H})$,

$$\text{regret}(\mathbb{A}, T) = \Omega\left(\sqrt{L(\mathbb{H})T}\right). \quad (1)$$

However, for the sufficiency direction, their upper bound has a dependence on the number of classes, and hence only establishes sufficiency for a bounded number of classes. This is analogous to the well-known gap in PAC learnability, which was only recently resolved by [Brukhim, Carmon, Dinur, Moran, and Yehudayoff \(2022\)](#). As our main result, we prove the following theorem.

Theorem 1 (Main Result) *Any concept class \mathbb{H} is agnostically online learnable iff $L(\mathbb{H}) < \infty$. Moreover, for any $T \in \mathbb{N}$, there is an online learning algorithm \mathbb{A} satisfying*

$$\text{regret}(\mathbb{A}, T) = \tilde{O}\left(\sqrt{L(\mathbb{H})T}\right).$$

In light of the lower bound (1), this further establishes that the *optimal* achievable regret for $T \geq L(\mathbb{H})$ satisfies

$$\min_{\mathbb{A}} \text{regret}(\mathbb{A}, T) = \tilde{\Theta}\left(\sqrt{L(\mathbb{H})T}\right).$$

While the result of [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2015\)](#), which has a dependence on the number of classes, is based on an extension of the algorithm of [Ben-David, Pál, and Shalev-Shwartz \(2009\)](#) for binary classification, our result modifies this approach. We take inspiration from the PAC setting, where the only known proof that the agnostic sample complexity is characterized by the DS dimension proceeds by a reduction to the realizable case, wherein the algorithm first identifies a maximal subset of the data, and then applies a compression scheme for the realizable case to this subset ([David, Moran, and Yehudayoff, 2016](#)). In our case, this maximal realizable subset appears only in the *analysis*, but serves an important role. Specifically, our algorithm applies the well-known multiplicative weights experts algorithm, with a family of experts defined by predictions of all possible executions of the SOA (see Section 2.1) that are constrained to only *update* their predictor in at most $L(\mathbb{H})$ pre-specified time steps. The important property is that one of these experts corresponds precisely to executing the SOA on a maximal realizable subsequence of the data sequence, and hence makes at most $L(\mathbb{H})$ mistakes on this subsequence (see Section 2.1). This expert therefore has regret only $L(\mathbb{H})$ compared to the best function in \mathbb{H} , and a regret bound for the overall algorithm then follows from classical analysis of prediction with expert advice. We present the detailed proof in Section 3.

The \tilde{O} in Theorem 1 hides a factor $\sqrt{\log(T/L(\mathbb{H}))}$. An analogous factor has recently been removed from the best known regret bound for *binary* classification ([Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev, 2021a](#)), yielding a regret bound that is optimal up to numerical constants. It remains open whether an analogous refinement is possible in the multiclass setting. Specifically, we pose the following open problem.

Open Problem 1 *Is it true that, for any concept class \mathbb{H} , the optimal regret is $\Theta\left(\sqrt{L(\mathbb{H})T}\right)$?*

In addition to the above results for learnability, we also study the related question of adversarial uniform laws of large numbers; see Section 4 for definitions. It was shown in the PAC setting that multiclass learnability is *not* equivalent to a uniform law of large numbers ([Natarajan, 1988](#);

Daniely, Sabato, Ben-David, and Shalev-Shwartz, 2011, 2015), in contrast to binary classification where it has long been established that these are equivalent (Vapnik and Chervonenkis, 1974). Again in the binary classification setting, Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogeve (2021a) have established an equivalence between a type of adversarial uniform law of large numbers (AULLN), online learnability, and convergence of the sequential Rademacher complexity (the latter two were already known to be equivalent, Rakhlin, Sridharan, and Tewari, 2015a). In the present work, we find that the analogy between PAC and online learning settings holds true once again. That is, in the multiclass setting with unbounded label space, while the AULLN is again equivalent to convergence of the sequential Rademacher complexity, in this case these properties are *not* equivalent to online learnability. Indeed, we show (Theorem 7) that AULLN and convergence of sequential Rademacher complexity are characterized by finiteness of a different combinatorial parameter $d_{\text{SG}}(\mathbb{H})$ called the *sequential graph dimension* (Definition 6), and we exhibit an example in which $L(\mathbb{H}) = 1$ but $d_{\text{SG}}(\mathbb{H}) = \infty$. All of this remains perfectly analogous to known results for multiclass PAC learning. Carrying the parallel further, in the case of bounded label spaces, we relate these two parameters by $d_{\text{SG}}(\mathbb{H}) = O(L(\mathbb{H}) \log(|\mathcal{Y}|))$ (Theorem 8). Based on the above, we also state new bounds on the achievable regret (Theorem 9): namely, $\text{regret}(\mathbb{A}, T) = O\left(\sqrt{d_{\text{SG}}(\mathbb{H})T}\right)$, which in the finite $|\mathcal{Y}|$ case, further implies $\text{regret}(\mathbb{A}, T) = O\left(\sqrt{L(\mathbb{H})T \log(|\mathcal{Y}|)}\right)$.

2. Definitions

We begin with some basic useful notation. For any sequence z_1, z_2, \dots , for any $t \in \mathbb{N} \cup \{0\}$, we denote by $z_{\leq t} = (z_1, \dots, z_t)$ and $z_{< t} = (z_1, \dots, z_{t-1})$, interpreting $z_{\leq 0} = z_{< 1} = ()$, the empty sequence. For simplicity, for a sequence $(x_1, y_1), \dots, (x_n, y_n)$, we denote by $(x_{< t}, y_{< t}) = ((x_1, y_1), \dots, (x_{t-1}, y_{t-1}))$ and $(x_{\leq t}, y_{\leq t}) = ((x_1, y_1), \dots, (x_t, y_t))$. For any $h \in \mathbb{H}$, $n \in \mathbb{N}$, and $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n$, denote by $h(\mathbf{x}) = (h(x_1), \dots, h(x_n))$. For any $V \subseteq \mathbb{H}$, $n \in \mathbb{N}$, and any sequence $(\mathbf{x}, \mathbf{y}) = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathcal{X} \times \mathcal{Y})^n$, denote by $V_{(\mathbf{x}, \mathbf{y})} = \{h \in V : h(\mathbf{x}) = \mathbf{y}\}$, called a *version space* of V . We say a sequence $(\mathbf{x}, \mathbf{y}) \in (\mathcal{X} \times \mathcal{Y})^n$ is *realizable* by \mathbb{H} if $\mathbb{H}_{(\mathbf{x}, \mathbf{y})} \neq \emptyset$: that is, $\exists h \in \mathbb{H}$ with $h(\mathbf{x}) = \mathbf{y}$. For any value $\alpha \in \mathbb{R}$, denote by $\log(\alpha) = \ln(\max\{\alpha, e\})$.

The following is the primary definition of dimension central to this work: namely, the *Littlestone dimension* $L(\mathbb{H})$ of a concept class \mathbb{H} . The definition is due to Daniely, Sabato, Ben-David, and Shalev-Shwartz (2011, 2015), representing the natural generalization of the classic definition of Littlestone (1988) (for binary classification) to multiclass.

We first state a clear intuitive definition in terms of binary trees. Specifically, a Littlestone tree T is a rooted binary tree, where each internal node is labeled by an instance $x \in \mathcal{X}$, and each edge between that node and a child is labeled by (x, y) for a class label $y \in \mathcal{Y}$, with the restriction that if the node has two children, then the corresponding class labels y, y' must be distinct: i.e., if the edges are labeled $(x, y), (x, y')$, respectively, then $y \neq y'$. A finite-depth Littlestone tree T is *shattered* by \mathbb{H} if, for every leaf node of any depth d , the sequence (\mathbf{x}, \mathbf{y}) of labels of the edges along the path from the root to this leaf is realizable by \mathbb{H} . The *Littlestone dimension*, $L(\mathbb{H})$, is then the maximum depth of a *perfect* Littlestone tree shattered by \mathbb{H} , where the term *perfect* means that every node has 2 children and every leaf has equal depth. If there are arbitrarily large depths d for which there exist perfect Littlestone trees shattered by \mathbb{H} , then $L(\mathbb{H}) = \infty$.

We state this definition formally (giving notation to the labels of nodes and edges) as follows.

Definition 2 The Littlestone dimension of \mathbb{H} , denoted $L(\mathbb{H})$, is defined as the largest $n \in \mathbb{N} \cup \{0\}$ for which $\exists\{(x_{\mathbf{b}}, y_{(\mathbf{b},0)}, y_{(\mathbf{b},1)}) : \mathbf{b} \in \{0, 1\}^t, t \in \{0, \dots, n-1\}\} \subseteq \mathcal{X} \times \mathcal{Y}^2$ (interpreting $\{0, 1\}^0 = \{()\}$) with the property that $\forall b_1, \dots, b_n \in \{0, 1\}, \exists h \in \mathbb{H}$ with

$$h(x_{b_{<1}}, x_{b_{<2}}, \dots, x_{b_{<n}}) = (y_{b_{\leq 1}}, y_{b_{\leq 2}}, \dots, y_{b_{\leq n}}).$$

If no such largest n exists, define $L(\mathbb{H}) = \infty$. Also define $L(\emptyset) = -1$.

When $L(\mathbb{H}) < \infty$, one can show that $L(\mathbb{H})$ can equivalently be defined inductively as

$$\max_x \max_{y_0 \neq y_1} \min_{i \in \{0,1\}} L(\mathbb{H}_{(x, y_i)}) + 1.$$

To see the correspondence between Definition 2 and the definition in terms of shattered Littlestone trees, we take the nodes at depth $t \geq 0$ (counting the root as depth 0) to be labeled $x_{b_{\leq t}}$, with the edge connecting to its left child labeled $(x_{b_{\leq t}}, y_{(b_{\leq t}, 0)})$ and the edge connecting to its right child labeled $(x_{b_{\leq t}}, y_{(b_{\leq t}, 1)})$.

2.1. Background

It was shown by [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#) that \mathbb{H} is online learnable in the realizable case if and only if $L(\mathbb{H}) < \infty$. That is, there is an algorithm guaranteeing a bounded number of mistakes on all $(X_1, Y_1), \dots, (X_T, Y_T)$ realizable by \mathbb{H} . Specifically, they employ a natural multiclass generalization of Littlestone’s *Standard Optimal Algorithm* (SOA), defined as follows. For any $(X_1, Y_1), \dots, (X_{t-1}, Y_{t-1}), X_t$, define

$$\text{SOA}(X_{<t}, Y_{<t}, X_t) = \operatorname{argmax}_{y \in \mathcal{Y}} L(\mathbb{H}_{((X_{<t}, Y_{<t}), (X_t, y))}),$$

where ties are broken arbitrarily. In other words, it (deterministically) predicts a label \hat{y}_t which maximizes the Littlestone dimension of the version space constrained by that label. Following the classic argument of [Littlestone \(1988\)](#) for binary classification, [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#) showed that SOA makes at most $L(\mathbb{H})$ mistakes on any sequence $(X_1, Y_1), \dots, (X_T, Y_T)$ realizable by \mathbb{H} : that is,

$$\sum_{t=1}^T \mathbb{1}[\text{SOA}(X_{<t}, Y_{<t}, X_t) \neq Y_t] \leq L(\mathbb{H}).$$

The reason is clear from the inductive version of the definition of $L(\mathbb{H})$. When $L(\mathbb{H}) < \infty$, for any $V \subseteq \mathbb{H}$ and $x \in \mathcal{X}$, at most one label y can have $L(V_{(x,y)}) = L(V)$, and hence (since L is integer-valued) every other y' must have $L(V_{(x,y')}) \leq L(V) - 1$. Thus, for $V = \mathbb{H}_{(X_{<t}, Y_{<t})}$, by predicting the label \hat{y}_t with maximum $L(\mathbb{H}_{((X_{<t}, Y_{<t}), (X_t, \hat{y}_t))})$, we are guaranteed that if $\hat{y}_t \neq Y_t$, then Y_t cannot have $L(\mathbb{H}_{(X_{\leq t}, Y_{\leq t})}) = L(\mathbb{H}_{(X_{<t}, Y_{<t})})$. That is, every mistake guarantees that the version space $\mathbb{H}_{(X_{\leq t}, Y_{\leq t})}$ has a smaller Littlestone dimension (by at least 1) than $\mathbb{H}_{(X_{<t}, Y_{<t})}$. Since $(X_1, Y_1), \dots, (X_T, Y_T)$ is realizable by \mathbb{H} , we always have $\mathbb{H}_{(X_{\leq t}, Y_{\leq t})} \neq \emptyset$, and hence $L(\mathbb{H}_{(X_{\leq t}, Y_{\leq t})}) \geq 0$, so that we can have $\hat{y}_t \neq Y_t$ at most $L(\mathbb{H})$ times.

Again following [Littlestone \(1988\)](#), the work of [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#) also shows a lower bound, establishing that for any deterministic online learning algorithm, there exists a realizable sequence where it makes at least $L(\mathbb{H})$ mistakes, so that SOA

is optimal in this regard among all deterministic online learning algorithms. Moreover, even for any randomized learning algorithm, there always exists a sequence realizable by \mathbb{H} for which the *expected* number of mistakes is at least $L(\mathbb{H})/2$. Thus, there exists an algorithm guaranteeing a bounded number of mistakes on all realizable sequences if and only if $L(\mathbb{H}) < \infty$: that is, \mathbb{H} is online learnable in the realizable case iff $L(\mathbb{H}) < \infty$.

In the case of *agnostic* online learning for multiclass classification, [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2015\)](#) establish a lower bound (for $T \geq L(\mathbb{H})$), $\text{regret}(\mathbb{A}, T) = \Omega(\sqrt{L(\mathbb{H})T})$, holding for any algorithm \mathbb{A} . Moreover, in the case of $|\mathcal{Y}| < \infty$, they propose a learning algorithm \mathbb{A} guaranteeing $\text{regret}(\mathbb{A}, T) = O\left(\sqrt{L(\mathbb{H})T \log(T|\mathcal{Y}|)}\right)$. Thus, in the case of bounded label spaces, the Littlestone dimension characterizes agnostic learnability. They left open the question of whether this remains true for unbounded label spaces. Note that the dependence on $|\mathcal{Y}|$ in the above regret bound makes the bound vacuous for infinite \mathcal{Y} spaces. We resolve this question (positively) in the present work (Theorem 1), and refine the above finite label bound as well (Theorem 8).

3. Agnostic Online Learnability of Littlestone Classes

This section presents the main result and algorithm in detail. As one of the main components of the algorithm, we make use of the following classic result for learning from expert advice (e.g., [Vovk, 1990, 1992](#); [Littlestone and Warmuth, 1994](#)); see Theorem 2.2 of [Cesa-Bianchi and Lugosi \(2006\)](#).

Lemma 3 ([Cesa-Bianchi and Lugosi, 2006, Theorem 2.2](#)) *For any $N, T \in \mathbb{N}$ and an array of values $e_{i,t} \in \{0, 1\}$, with $i \in \{1, \dots, N\}$, $t \in \{1, \dots, T\}$, letting $\eta = \sqrt{(8/T) \ln(N)}$, for any $y_1, \dots, y_T \in \{0, 1\}$, letting $w_{i,1} = 1$ and $w_{i,t} = e^{-\eta \sum_{s<t} \mathbb{1}[e_{i,s} \neq y_s]}$ for each $t \in \{2, \dots, T\}$ and $i \in \{1, \dots, N\}$, letting $p_t = \sum_{i=1}^N w_{i,t} e_{i,t} / \sum_{i'=1}^N w_{i',t}$, it holds that*

$$\sum_{t=1}^T |p_t - y_t| - \min_{1 \leq i \leq N} \sum_{t=1}^T \mathbb{1}[e_{i,t} \neq y_t] \leq \sqrt{(T/2) \ln(N)}.$$

We are now ready to describe our agnostic online learning algorithm, denoted by \mathbb{A}_{AG} , specified based on a given time horizon $T \in \mathbb{N}$. As above, let $(X_1, Y_1), \dots, (X_T, Y_T)$ denote the data sequence. For any $J \subseteq \{1, \dots, T\}$, denote by $X_J = \{X_t : t \in J\}$ and $Y_J = \{Y_t : t \in J\}$, and for any $t \in \mathbb{N}$, denote by $J_{<t} = J \cap \{1, \dots, t-1\}$. We consider a set of *experts* defined as follows. Let $\mathcal{J} = \{J \subseteq \{1, \dots, T\} : |J| \leq L(\mathbb{H})\}$. For each $J \in \mathcal{J}$, for each $t \in \mathbb{N}$, define

$$g_t^J = \text{SOA}(X_{J_{<t}}, Y_{J_{<t}}, X_t).$$

That is, g_t^J is the prediction the SOA would make on X_t , given that its previous sequence of examples were $(X_{J_{<t}}, Y_{J_{<t}})$: namely, $\text{argmax}_y L(\mathbb{H}((X_{J_{<t}}, Y_{J_{<t}}), (X_t, y)))$. Based on the set of experts $\{g^J : J \in \mathcal{J}\}$, we apply the *multiplicative weights* algorithm. Explicitly, letting $\eta = \sqrt{(8/T) \ln(|\mathcal{J}|)}$, and defining $w_{J,1} = 1$ and $w_{J,t} = e^{-\eta \sum_{s<t} \mathbb{1}[g_s^J \neq Y_s]}$ for $t \in \{2, \dots, T\}$ and $J \in \mathcal{J}$, we define the prediction at time t as

$$p_t = \frac{\sum_{J \in \mathcal{J}} w_{J,t} g_t^J}{\sum_{J' \in \mathcal{J}} w_{J',t}}.$$

Note that the values g_t^J of the experts at time t only depend on $X_{<t}, Y_{<t}, X_t$, so that this is a valid prediction. We have the following result, representing the main theorem of this work. In particular, in conjunction with the necessity of $L(\mathbb{H}) < \infty$ for online learnability, established by [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2015\)](#), our [Theorem 1](#) immediately follows from this.

Theorem 4 *For any concept class \mathbb{H} and $T \geq 2L(\mathbb{H})$, the algorithm \mathbb{A}_{AG} defined above satisfies*

$$\text{regret}(\mathbb{A}_{AG}, T) = O\left(\sqrt{L(\mathbb{H})T \log\left(\frac{T}{L(\mathbb{H})}\right)}\right).$$

Proof Let $h^* \in \mathbb{H}$ satisfy $\sum_{t=1}^T \mathbb{1}[h^*(X_t) \neq Y_t] = \min_{h \in \mathbb{H}} \sum_{t=1}^T \mathbb{1}[h(X_t) \neq Y_t]$. Denote by $R^* = \{t \in \{1, \dots, T\} : h^*(X_t) = Y_t\}$, and note that the subsequence (X_{R^*}, Y_{R^*}) is realizable by \mathbb{H} . Define a sequence j_r inductively, as follows. Let $j_1 = \min\{t \in R^* : \text{SOA}(\emptyset, \emptyset, X_t) \neq Y_t\}$ if it exists. For $r \geq 2$, if j_{r-1} is defined, let

$$j_r = \min\{t \in R^* : t > j_{r-1} \text{ and } \text{SOA}(X_{j_{<r}}, Y_{j_{<r}}, X_t) \neq Y_t\}$$

if it exists. Finally, define

$$J^* = \{j_r : r \in \mathbb{N} \text{ and } j_r \text{ exists}\}.$$

In other words, J^* represents the sequence of mistakes SOA would make on the sequence R^* using conservative updates: that is, only adding an example (X_t, Y_t) to its history if it is a mistake point.

In particular, note that if $J^* = \emptyset$, then every $t \in R^*$ satisfies $g_t^{J^*} = \text{SOA}(X_{J^*_{<t}}, Y_{J^*_{<t}}, X_t) = \text{SOA}(\emptyset, \emptyset, X_t) = Y_t$. If $J^* \neq \emptyset$, then any $t \in R^*_{<j_1}$ has $g_t^{J^*} = \text{SOA}(X_{J^*_{<t}}, Y_{J^*_{<t}}, X_t) = \text{SOA}(\emptyset, \emptyset, X_t) = Y_t$; similarly, any $r \in \{1, \dots, |J^*| - 1\}$ and $t \in R^*_{<j_{r+1}} \setminus R^*_{\leq j_r}$ has $g_t^{J^*} = \text{SOA}(X_{J^*_{<t}}, Y_{J^*_{<t}}, X_t) = \text{SOA}(X_{j_{<r+1}}, Y_{j_{<r+1}}, X_t) = Y_t$; also, for $r = |J^*|$ (the largest r for which j_r is defined), any $t \in R^* \setminus R^*_{\leq j_r}$ has $g_t^{J^*} = \text{SOA}(X_{J^*_{<t}}, Y_{J^*_{<t}}, X_t) = \text{SOA}(X_{j_{<r+1}}, Y_{j_{<r+1}}, X_t) = Y_t$. In other words, $g_t^{J^*} = Y_t$ for every $t \in R^* \setminus J^*$. Therefore,

$$\sum_{t=1}^T \mathbb{1}[g_t^{J^*} \neq Y_t] \leq |J^*| + (T - |R^*|) = |J^*| + \min_{h \in \mathbb{H}} \sum_{t=1}^T \mathbb{1}[h(X_t) \neq Y_t].$$

Moreover, since (X_{R^*}, Y_{R^*}) is realizable by \mathbb{H} , and $J^* \subseteq R^*$, the subsequence (X_{J^*}, Y_{J^*}) is also realizable by \mathbb{H} . By definition, SOA makes a mistake on every time when run through the subsequence (X_{J^*}, Y_{J^*}) : that is, $\text{SOA}(X_{j_{<r}}, Y_{j_{<r}}, X_{j_r}) \neq Y_{j_r}$ for every $j_r \in J^*$. Thus, by the guaranteed mistake bound $L(\mathbb{H})$ for SOA on realizable sequences, we conclude that $|J^*| \leq L(\mathbb{H})$. In particular, this implies $J^* \in \mathcal{J}$. Altogether, by [Lemma 3](#), we have that

$$\begin{aligned} \sum_{t=1}^T |p_t - Y_t| &\leq \left(\sum_{t=1}^T \mathbb{1}[g_t^{J^*} \neq Y_t] \right) + \sqrt{(T/2) \ln(|\mathcal{J}|)} \\ &\leq \min_{h \in \mathbb{H}} \sum_{t=1}^T \mathbb{1}[h(X_t) \neq Y_t] + L(\mathbb{H}) + \sqrt{(T/2)L(\mathbb{H}) \ln\left(\frac{eT}{L(\mathbb{H})}\right)}, \end{aligned}$$

$$\implies \text{regret}(\mathbb{A}_{AG}, T) \leq L(\mathbb{H}) + \sqrt{(T/2)L(\mathbb{H}) \ln\left(\frac{eT}{L(\mathbb{H})}\right)} = O\left(\sqrt{L(\mathbb{H})T \log\left(\frac{T}{L(\mathbb{H})}\right)}\right). \blacksquare$$

4. Online Uniform Convergence Versus Online Learnability

In PAC learning for binary classification, for a given data distribution, a class \mathbb{H} satisfies the uniform law of large numbers (i.e., is P -Glivenko-Cantelli) if and only if the (normalized) Rademacher complexity converges to 0 in sample size. Moreover, the rate of uniform convergence is dominated by a converging distribution-free function of sample size if and only if the VC dimension is finite. By a chaining argument (Talagrand, 1994; van der Vaart and Wellner, 1996), the optimal form for this rate is $\Theta(\sqrt{\text{VC}(\mathbb{H})/n})$, where n is the sample size and $\text{VC}(\mathbb{H})$ is the VC dimension of \mathbb{H} .

Similarly, in adversarial online learning for binary classification, Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev (2021a) showed that a class \mathbb{H} satisfies an *adversarial* uniform law of large numbers if and only if the (normalized) *sequential* Rademacher complexity converges to 0 in the sequence length.³ The rate of adversarial uniform convergence, uniform over sequences, is then controlled by the Littlestone dimension. Again by a chaining argument (Alon et al., 2021a), the optimal form for this rate is $\Theta(\sqrt{L(\mathbb{H})T})$.⁴

Importantly, for both of these facts, the complexity measure controlling the rate of uniform convergence is the same as the complexity measure that determines learnability: for PAC learning, the VC dimension, and for online learning, the Littlestone dimension. In particular, together with separately established lower bounds for learning, these facts imply that the optimal rate of convergence of expected excess error in agnostic PAC learning is $\Theta(\sqrt{\text{VC}(\mathbb{H})/n})$ (Talagrand, 1994), whereas the optimal regret bound for agnostic online learning is $\Theta(\sqrt{L(\mathbb{H})T})$ (Alon et al., 2021a).

In the case of PAC learning for *multiclass* classification, it again holds that the classification losses satisfy the uniform law of large numbers if and only if the Rademacher complexity converges to 0 in sample size. However, in this case, there exists a distribution-free bound on the rate of uniform convergence if and only if the *graph dimension* is finite (Ben-David, Cesa-Bianchi, Haussler, and Long, 1995; Daniely, Sabato, Ben-David, and Shalev-Shwartz, 2011, 2015), and the optimal such bound is $\Theta(\sqrt{d_G(\mathbb{H})/n})$, where n is the sample size and $d_G(\mathbb{H})$ is the graph dimension. Notably, the graph dimension does *not* control PAC learnability of the class (Natarajan, 1988; Daniely et al., 2011, 2015); rather, a recent result of Brukhim et al. (2022) established that multiclass PAC learnability (including agnostic learnability) is controlled by a quantity they call the *DS dimension* (originally proposed by Daniely and Shalev-Shwartz, 2014, who proved it provides a lower bound), and there are simple examples where the DS dimension is finite while the graph dimension is infinite (necessarily with an infinite number of possible class labels). Thus, in the case of multiclass classification, we see that the uniform law of large numbers, and PAC learnability, are controlled by *different* parameters of the class, and there are PAC learnable classes which do not satisfy the uniform law of large numbers. Thus, PAC learning algorithms generally cannot rely on a uniform law of large numbers, in contrast to binary classification.

3. Rakhlin, Sridharan, and Tewari (2010, 2015a,b) also studied a notion of sequential uniform law of large numbers. Though formulated somewhat differently, that notion was also shown to be equivalent to convergence of sequential Rademacher complexity to 0, and hence is satisfied iff the notion of AULLN studied by Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev (2021a) is satisfied.

4. This result of Alon et al. (2021a) was in fact only shown under a further technical restriction on \mathbb{H} (rooted in the work of Rakhlin et al., 2010, 2015a), so that the rounds of the online learning game satisfy the *minimax theorem*. As the results in this section are based on this result of Alon et al. (2021a), we also suppose this condition is appropriately satisfied. It remains open whether this $\sqrt{L(\mathbb{H})T}$ regret for binary classification, and consequently our theorems below, remain valid without any restrictions on \mathbb{H} .

In this section, we note an analogous result holds for the *adversarial* uniform law of large numbers. We find that, while the adversarial uniform law of large numbers is again satisfied if and only if the sequential Rademacher complexity converges to 0 in sequence length, the optimal sequence-independent bound on the convergence depends on the *sequential graph dimension*. Thus, in light of our Theorem 1, establishing that agnostic online learnability is controlled by the *Littlestone dimension*, we again see that the parameter controlling the adversarial uniform law of large numbers differs from the parameter controlling online learnability (including agnostic learnability). Moreover, we provide a simple example where the sequential graph dimension is infinite while the Littlestone dimension is finite, thus showing that not all online learnable classes satisfy the AULLN.

We begin by recalling the following definitions of Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev (2021a). An adversarially uniform law of large numbers (AULLN) can be viewed as a sequential game between a *sampler* \mathcal{S} and an adversary. On each round t , the adversary chooses $(x_t, y_t) \in \mathcal{X} \times \mathcal{Y}$, and the sampler decides whether to include (x_t, y_t) in its “sample”: a subsequence K . The sampler may be randomized, and the adversary may observe and adapt to the sampler’s past decisions (though not its internal random bits).

Definition 5 (Alon et al., 2021a) *A concept class \mathbb{H} satisfies the adversarial uniform law of large numbers (AULLN) if, for any $\varepsilon, \delta \in (0, 1)$, there exists $k(\varepsilon, \delta) \in \mathbb{N}$ and a sampler \mathcal{S} such that, for any adversarially-produced sequence $(\bar{x}, \bar{y}) = \{(x_t, y_t)\}_{t=1}^T$ of any length T , the sample K selected by \mathcal{S} always satisfies $|K| \leq k(\varepsilon, \delta)$, and with probability at least $1 - \delta$, K forms an ε -approximation of (\bar{x}, \bar{y}) (with respect to \mathbb{H}), meaning⁵*

$$\sup_{h \in \mathbb{H}} \left| \frac{1}{|K|} \sum_{(x_t, y_t) \in K} \mathbb{1}[h(x_t) \neq y_t] - \frac{1}{T} \sum_{t=1}^T \mathbb{1}[h(x_t) \neq y_t] \right| \leq \varepsilon.$$

The AULLN was shown by Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev (2021a) to be intimately related to the *sequential Rademacher complexity* of the class of indicator functions, in this case, $(x, y) \mapsto \mathbb{1}[h(x) \neq y]$, $h \in \mathbb{H}$. Formally, we recall the definition of sequential Rademacher complexity from the work of Rakhlin, Sridharan, and Tewari (2010, 2015a). Let $\varepsilon = \{\varepsilon_t\}_{t \in \mathbb{N}}$ be i.i.d. Uniform($\{-1, 1\}$) random variables. Let $z = \{(x_t, y_t)\}_{t \in \mathbb{N}}$ be any sequence of functions $(x_t, y_t) : \{-1, 1\}^{t-1} \rightarrow \mathcal{X} \times \mathcal{Y}$, denoting by $(x_t(\varepsilon_{<t}), y_t(\varepsilon_{<t}))$ the value of this function on $\varepsilon_{<t}$. For any $T \in \mathbb{N}$, the sequential Rademacher complexity is defined as

$$\text{Rad}_T(\mathbb{H}) = \sup_z \mathbb{E} \left[\sup_{h \in \mathbb{H}} \frac{1}{T} \sum_{t=1}^T \varepsilon_t \mathbb{1}[h(x_t(\varepsilon_{<t})) \neq y_t(\varepsilon_{<t})] \right].$$

Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev (2021a) proved that, for binary classification, the best achievable bound on the sequential Rademacher complexity is the Littlestone dimension of \mathbb{H} . However, binary classification enjoys a special property that composition of \mathbb{H} with the 0-1 loss does not change the complexity. In contrast, in multiclass classification, the Littlestone dimension of \mathbb{H} composed with the 0-1 loss is a different quantity, which we term the *sequential graph dimension*:

5. To be clear, the sum over $(x_t, y_t) \in K$ treats duplicates as distinct: that is, there are $|K|$ terms in the sum.

Definition 6 *The sequential graph dimension of \mathbb{H} , denoted $d_{\text{SG}}(\mathbb{H})$, is defined as*

$$d_{\text{SG}}(\mathbb{H}) = L(\{(x, y) \mapsto \mathbb{1}[h(x) \neq y] : h \in \mathbb{H}\}).$$

Explicitly, $d_{\text{SG}}(\mathbb{H})$ is the largest $n \in \mathbb{N} \cup \{0\}$ s.t. $\exists\{(x_{\mathbf{b}}, y_{\mathbf{b}}) : \mathbf{b} \in \{0, 1\}^t, t \in \{0, \dots, n-1\}\} \subseteq \mathcal{X} \times \mathcal{Y}$ with the property that $\forall b_1, \dots, b_n \in \{0, 1\}, \exists h \in \mathbb{H}$ with

$$(\mathbb{1}[h(x_{b_{<1}}) \neq y_{b_{<1}}], \mathbb{1}[h(x_{b_{<2}}) \neq y_{b_{<2}}], \dots, \mathbb{1}[h(x_{b_{<n}}) \neq y_{b_{<n}}]) = (b_1, \dots, b_n).$$

If no such largest n exists, define $d_{\text{SG}}(\mathbb{H}) = \infty$. Also define $d_{\text{SG}}(\emptyset) = -1$.

Here we state an extension of the result of [Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev \(2021a\)](#) for AULLN to the multiclass setting. The result follows immediately by applying Theorems 2.2 and 2.3 of [Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev \(2021b\)](#) to the class of binary functions $\mathcal{G} = \{(x, y) \mapsto \mathbb{1}[h(x) \neq y] : h \in \mathbb{H}\}$, along with (for the final claim about $\text{Rad}_T(\mathbb{H})$) the analogous application to \mathcal{G} of Corollary 12 of [Rakhlin, Sridharan, and Tewari \(2015a\)](#) and the upper bound on sequential Rademacher complexity in the proof of Theorem 12.1 of [Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev \(2021b\)](#).

Theorem 7 *For any concept class \mathbb{H} , the following are equivalent:*

1. \mathbb{H} satisfies the adversarial uniform law of large numbers
2. $\text{Rad}_T(\mathbb{H}) \rightarrow 0$
3. $d_{\text{SG}}(\mathbb{H}) < \infty$.

Moreover, if \mathbb{H} satisfies AULLN, then Definition 5 is satisfied with $k(\varepsilon, \delta) = O\left(\frac{d_{\text{SG}}(\mathbb{H}) + \log(1/\delta)}{\varepsilon^2}\right)$, and the minimal achievable $k(\varepsilon, \delta)$ satisfies $k(\varepsilon, \delta) = \Omega\left(\frac{d_{\text{SG}}(\mathbb{H})}{\varepsilon^2}\right)$. Additionally, for $T \geq d_{\text{SG}}(\mathbb{H})$, it holds that $\text{Rad}_T(\mathbb{H}) = \Theta\left(\sqrt{\frac{d_{\text{SG}}(\mathbb{H})}{T}}\right)$.

While Theorem 7 indeed expresses a fairly tight relation between AULLN, seq. Rademacher complexity, and the seq. graph dimension, the corresponding result of [Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogev \(2021a\)](#) for binary classification additionally found an equivalence to online learnability and agnostic online learnability. Since our Theorem 1 establishes that agnostic online learnability is characterized by finiteness of the *Littlestone dimension* (and this is true for realizable online learning as well), rather than the sequential graph dimension, we see a separation in multiclass classification between AULLN and online learnability. To formally establish this separation, we present the following example, exhibiting a concept class with finite Littlestone dimension but infinite sequential graph dimension.

Example 1: The following construction is identical to a known example of [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#), originally constructed to show a class that is PAC learnable but has infinite (non-sequential) *graph* dimension. For completeness, we include the full details of the construction here. Let \mathcal{X} be a countable set, let $\mathcal{Y} = 2^{\mathcal{X}} \cup \{*\}$, and for each $A \subseteq \mathcal{X}$, define

$$h_A(x) = \begin{cases} A & \text{if } x \in A \\ * & \text{otherwise} \end{cases}.$$

Define $\mathbb{H} = \{h_A : A \subseteq \mathcal{X}\}$. For any $x \in \mathcal{X}$ and distinct $y_0, y_1 \in \mathcal{Y}$, it must be that one of y_0, y_1 is equal some $A \subseteq \mathcal{X}$. Since only one $h \in \mathbb{H}$ could possibly have $h(x) = A$ (namely, h_A), and even then only if $x \in A$, we would have $L(\mathbb{H}_{(x,A)}) \in \{-1, 0\}$. It follows that $L(\mathbb{H}) \leq 1$. It can be seen that $L(\mathbb{H}) = 1$ by choosing $y_0 = *$ and $y_1 = A$ for any A such that $x \in A$. On the other hand, for any $n \in \mathbb{N}$ and distinct $x_1, \dots, x_n \in \mathcal{X}$, for any $b_1, \dots, b_n \in \{0, 1\}$, letting $x_{b_{<t}} = x_t$ and $y_{b_{<t}} = *$ for each $t \leq n$, the function h_A , with $A = \{x_t : b_t = 1\}$, satisfies $\mathbb{1}[h_A(x_{b_{<t}}) \neq y_{b_{<t}}] = b_t$ for every $t \in \{1, \dots, n\}$. Thus, $d_{\text{SG}}(\mathbb{H}) \geq n$. Since this is true of every $n \in \mathbb{N}$, we see that $d_{\text{SG}}(\mathbb{H}) = \infty$. In particular, in light of Theorems 1 and 7, this class \mathbb{H} is agnostically online learnable (and realizable online learnable), but does not satisfy the adversarial uniform law of large numbers (nor satisfy $\text{Rad}_T(\mathbb{H}) \rightarrow 0$). This is precisely analogous to the findings of Daniely, Sabato, Ben-David, and Shalev-Shwartz (2011, 2015) that this class is PAC learnable but does not satisfy the (non-adversarial) uniform law of large numbers.

4.1. Finite Label Spaces

In regard to agnostic online learning in the case $|\mathcal{Y}| < \infty$, the best previous regret bound, due to Daniely, Sabato, Ben-David, and Shalev-Shwartz (2015), is $O(\sqrt{L(\mathbb{H})T \log(T|\mathcal{Y}|)})$. Based on Theorem 7, in this section, we improve this regret to $O(\sqrt{d_{\text{SG}}(\mathbb{H})T})$. Further, we prove in Theorem 8 that $d_{\text{SG}}(\mathbb{H}) = O(L(\mathbb{H}) \log(|\mathcal{Y}|))$, so that this additionally implies a guarantee $\text{regret}(\mathbb{A}, T) = O(\sqrt{L(\mathbb{H})T \log(|\mathcal{Y}|)})$. On the one hand, this improves over Theorem 4 by removing a factor $\sqrt{\log(T/L(\mathbb{H}))}$, but on the other hand, includes a factor $\sqrt{\log(|\mathcal{Y}|)}$ not present in Theorem 1. In light of the lower bound $\Omega(\sqrt{L(\mathbb{H})T})$ of Daniely, Sabato, Ben-David, and Shalev-Shwartz (2015) holding for any \mathbb{A} , we see that this $O(\sqrt{L(\mathbb{H})T \log(|\mathcal{Y}|)})$ regret guarantee is optimal up to the $\sqrt{\log(|\mathcal{Y}|)}$ factor. As stated in Open Problem 1, it remains an open problem to determine whether the optimal regret is always of the form $\Theta(\sqrt{L(\mathbb{H})T})$. Formally:

Theorem 8 *If $|\mathcal{Y}| < \infty$, for any concept class \mathbb{H} , $d_{\text{SG}}(\mathbb{H}) = O(L(\mathbb{H}) \log(|\mathcal{Y}|))$.*

Proof We follow a strategy of *adaptive experts*, rooted in the work of Ben-David, Pál, and Shalev-Shwartz (2009) (and similar to the extension thereof used by Daniely, Sabato, Ben-David, and Shalev-Shwartz, 2015 in their multiclass agnostic online learner). Let $n \in \mathbb{N}$ be any number with $n \leq d_{\text{SG}}(\mathbb{H})$. Let \mathcal{Q} denote the set of all (J, Y) such that $J \subseteq \{1, \dots, n\}$ with $|J| \leq L(\mathbb{H})$, and $Y = \{Y_j\}_{j \in J}$ is a sequence of values in \mathcal{Y} . For any $(J, Y) \in \mathcal{Q}$ and any $t \in \{1, \dots, n\}$ and $x_1, \dots, x_t \in \mathcal{X}$, define a value $y_t^{J,Y}(x_{\leq t})$ inductively, as

$$y_t^{J,Y}(x_{\leq t}) = \begin{cases} \text{SOA}(x_{<t}, y_{<t}^{J,Y}, x_t) & \text{if } j \notin J \\ Y_j & \text{if } j \in J \end{cases}$$

Consider any set $\{(x_b, y_b) : b \in \{0, 1\}^t, t \in \{0, \dots, n-1\}\} \subseteq \mathcal{X} \times \mathcal{Y}$ satisfying the property in Definition 6. Now we inductively construct a sequence $b_1, \dots, b_n \in \{0, 1\}$ as follows. For $t \in \{1, \dots, n\}$, suppose we have already defined b_1, \dots, b_{t-1} , and let

$$V_{<t} = \{(J, Y) \in \mathcal{Q} : \forall i \in \{1, \dots, t-1\}, \mathbb{1}[y_i^{J,Y}(x_{b_{<i}}) \neq y_{b_{<i}}] = b_i\}.$$

Define

$$b_t = \underset{b \in \{0,1\}}{\text{argmin}} |\{(J, Y) \in V_{<t} : \mathbb{1}[y_t^{J,Y}(x_{b_{<1}}, \dots, x_{b_{<t}}) \neq y_{b_{<t}}] = b\}|.$$

This completes the inductive definition of b_1, \dots, b_n . In particular, note that for every $t \in \{1, \dots, n\}$ satisfies $|V_t| \leq \frac{1}{2}|V_{t-1}|$, so that $|V_n| \leq 2^{-n}|\mathcal{Q}|$.

On the other hand, by definition of $(x_{\mathbf{b}}, y_{\mathbf{b}})$, there exists $h \in \mathbb{H}$ with $\mathbb{1}[h(x_{b_{<t}}) \neq y_{b_{<t}}] = b_t$ simultaneously for every $t \in \{1, \dots, n\}$. Let

$$J^* = \{t \in \{1, \dots, n\} : \text{SOA}(x_{b_{<t-1}}, h(x_{b_{<t-1}}), x_{b_{<t}}) \neq h(x_{b_{<t}})\},$$

interpreting $(x_{b_{<0}}, h(x_{b_{<0}})) = (\emptyset, \emptyset)$. Recall from Section 2.1 that [Daniely, Sabato, Ben-David, and Shalev-Shwartz \(2011, 2015\)](#) proved a mistake bound of $L(\mathbb{H})$ for SOA, and hence $|J^*| \leq L(\mathbb{H})$. For each $t \in J^*$, define $Y_t^* = h(x_{b_{<t}})$, and let $Y^* = \{Y_t^*\}_{t \in J^*}$. By definition of (J^*, Y^*) , we have

$$\forall t \in \{1, \dots, n\}, y_t^{J^*, Y^*}(x_{b_{<1}}, \dots, x_{b_{<t}}) = h(x_{b_{<t}}).$$

In particular, this implies $(J^*, Y^*) \in V_n$, so that $|V_n| \geq 1$. Altogether, we have

$$1 \leq |V_n| \leq 2^{-n}|\mathcal{Q}| = 2^{-n} \sum_{i=1}^{L(\mathbb{H})} \binom{n}{i} |\mathcal{Y}|^i \leq 2^{-n} \left(\frac{en}{L(\mathbb{H})} \right)^{L(\mathbb{H})} |\mathcal{Y}|^{L(\mathbb{H})}.$$

Multiplying the leftmost and rightmost expressions by 2^n and taking logarithms yields

$$n \leq L(\mathbb{H}) \log_2 \left(\frac{2n}{L(\mathbb{H})} \right) + L(\mathbb{H}) \log_2(|\mathcal{Y}|). \quad (2)$$

Solving for an upper bound on n reveals that $n \leq 2L(\mathbb{H}) \log_2(e|\mathcal{Y}|)$. Specifically, this claim holds trivially when $|\mathcal{Y}| = 1$, and for $|\mathcal{Y}| \geq 2$ it follows from (2) by Lemma 4.6 of [Vidyasagar \(2003\)](#). ■

Theorem 9 *If $|\mathcal{Y}| < \infty$, for any concept class \mathbb{H} with $L(\mathbb{H}) < \infty$, there exists an algorithm \mathbb{A} satisfying*

$$\text{regret}(\mathbb{A}, T) = O\left(\sqrt{d_{\text{SG}}(\mathbb{H})T}\right).$$

Moreover, this implies $\text{regret}(\mathbb{A}, T) = O\left(\sqrt{L(\mathbb{H})T \log(|\mathcal{Y}|)}\right)$.

Proof The proof follows identically the proof of Theorem 12.1 of [Alon, Ben-Eliezer, Dagan, Moran, Naor, and Yogeve \(2021b\)](#). Specifically, Theorem 7 of [Rakhlin, Sridharan, and Tewari \(2015a\)](#) provides that $\text{regret}(\mathbb{A}, T) \leq 2\text{Rad}_T(\mathbb{H})$. The theorem then follows directly from Theorems 7 and 8. ■

Acknowledgments

Shay Moran is a Robert J. Shillman Fellow; he acknowledges support by ISF grant 1225/20, by BSF grant 2018385, by an Azrieli Faculty Fellowship, by Israel PBC-VATAT, by the Technion Center for Machine Learning and Intelligent Systems (MLIS), and by the the European Union (ERC, GENERALIZATION, 101039692). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- N. Alon, O. Ben-Eliezer, Y. Dagan, S. Moran, M. Naor, and E. Yogev. Adversarial laws of large numbers and optimal regret in online classification. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 2021a.
- N. Alon, O. Ben-Eliezer, Y. Dagan, S. Moran, M. Naor, and E. Yogev. Adversarial laws of large numbers and optimal regret in online classification. *arXiv:2101.09054*, 2021b.
- S. Ben-David, N. Cesa-Bianchi, D. Haussler, and P. Long. Characterizations of learnability for classes of $\{0, \dots, n\}$ -valued functions. *Journal of Computer and System Sciences*, 50:74–86, 1995.
- S. Ben-David, D. Pál, and S. Shalev-Shwartz. Agnostic online learning. In *Proceedings of the 22nd Conference on Learning Theory*, 2009.
- N. Brukhim, D. Carmon, I. Dinur, S. Moran, and A. Yehudayoff. A characterization of multiclass learnability. In *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science*, 2022.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- A. Daniely and S. Shalev-Shwartz. Optimal learners for multiclass problems. In *Proceedings of the 27th Conference on Learning Theory*, 2014.
- A. Daniely, S. Sabato, S. Ben-David, and S. Shalev-Shwartz. Multiclass learnability and the ERM principle. In *Proceedings of the 24th Conference on Learning Theory*, 2011.
- A. Daniely, S. Sabato, S. Ben-David, and S. Shalev-Shwartz. Multiclass learnability and the ERM principle. *Journal of Machine Learning Research*, 16(12):2377–2404, 2015.
- O. David, S. Moran, and A. Yehudayoff. On statistical learning via the lens of compression. In *Advances in Neural Information Processing* 29, 2016.
- N. Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2:285–318, 1988.
- N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.
- B. K. Natarajan. On learning sets and functions. *Machine Learning*, 4:67–97, 1989.
- Balas K. Natarajan. Some results on learning. *Unpublished manuscript*, 1988.
- Balas K. Natarajan and Prasad Tadepalli. Two new frameworks for learning. In *ICML*, pages 402–415, 1988.
- A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Random averages, combinatorial parameters, and learnability. In *Advances in Neural Information Processing Systems* 23, 2010.

- A. Rakhlin, K. Sridharan, and A. Tewari. Online learning via sequential complexities. *Journal of Machine Learning Research*, 16(1):155–186, 2015a.
- A. Rakhlin, K. Sridharan, and A. Tewari. Sequential complexities and uniform martingale laws of large numbers. *Probability Theory and Related Fields*, 161:111–153, 2015b.
- M. Talagrand. Sharper bounds for gaussian and empirical processes. *The Annals of Probability*, 22: 28–76, 1994.
- A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes*. Springer, 1996.
- V. Vapnik and A. Chervonenkis. *Theory of Pattern Recognition*. Nauka, Moscow, 1974.
- M. Vidyasagar. *Learning and Generalization with Applications to Neural Networks*. Springer-Verlag, 2nd edition, 2003.
- V. Vovk. Aggregating strategies. In *Proceedings of the 3rd Annual Workshop on Computational Learning Theory*, 1990.
- V. Vovk. Universal forecasting algorithms. *Information and Computation*, 96(2):245–277, 1992.