
Online Agnostic Multiclass Boosting

Vinod Raman

Department of Statistics
University of Michigan
Ann Arbor, MI 48104
vkraman@umich.edu

Ambuj Tewari

Department of Statistics
University of Michigan
Ann Arbor, MI 48018
tewaria@umich.edu

Abstract

Boosting is a fundamental approach in machine learning that enjoys both strong theoretical and practical guarantees. At a high-level, boosting algorithms cleverly aggregate weak learners to generate predictions with arbitrarily high accuracy. In this way, boosting algorithms convert weak learners into strong ones. Recently, Brukhim et al. [6] extended boosting to the online agnostic binary classification setting. A key ingredient in their approach is a clean and simple reduction to online convex optimization, one that efficiently converts an arbitrary online convex optimizer to an agnostic online booster. In this work, we extend this reduction to multiclass problems and give the first boosting algorithm for online agnostic multiclass classification. Our reduction also enables the construction of algorithms for statistical agnostic, online realizable, and statistical realizable multiclass boosting.

1 Introduction

Boosting is a fundamental technique in machine learning that cleverly aggregates the predictions of weak learners to produce a strong learner. Originally studied in the batch (realizable) PAC learning setting for binary classification, boosting has now been extended to a wide variety of settings, including multiclass classification, online learning, and agnostic learning. Recently, Brukhim et al. [6] extended boosting to online agnostic binary classification, marking the completion of boosting algorithms for all four regimes of statistical/online and agnostic/realizable binary classification. However, less can be said about multiclass classification, where boosting algorithms are only well studied under the assumption of realizability. Designing online agnostic multiclass boosting algorithms is important for several reasons. First, realizability is a very strong assumption in real-life: it requires the existence of an expert that *perfectly* labels the data. Second, the vast majority of classification tasks require more than two labels, a prominent example being image classification. Lastly, modern machine learning tasks often require sequential processing of data, making the design and development of online algorithms increasingly relevant. In this work we fill this gap in literature by studying online agnostic boosting for multiclass problems.

1.1 Main Results

We give the first weak learning conditions and algorithms for online agnostic multiclass boosting. A key idea of our algorithms is a reduction from boosting to online convex optimization (OCO), an idea borrowed from Brukhim et al. [6]. As a consequence of this reduction, we are also able to give algorithms for the three other settings of statistical agnostic, online realizable, and statistical realizable multiclass boosting. Finally, we give empirical results showcasing that our OCO-based boosting algorithms are fast and competitive with existing state-of-the-art multiclass boosting algorithms.

1.2 Related Works

Boosting was first studied for binary classification under the realizable PAC learning setting [33, 16–18, 32] and then later extended to the agnostic PAC learning setting [15, 21, 27, 26, 25, 28, 2]. The success of boosting for binary classification led to significant interest in designing boosting algorithms for multiclass problems. As a result, several multiclass boosting algorithms were proposed for the realizable batch setting [19, 14, 22, 20], culminating in the work by Mukherjee and Schapire [30], who unified the previous approaches under a cost matrix framework. Beyond the batch setting, online boosting algorithms for binary [31, 9, 4] and multiclass classification [10, 24] have been designed assuming realizability (mistake-bound). More recently, Brukhim et al. [6] give the first online *agnostic* (regret-bound) boosting algorithm for binary classification, marking complete all four regimes of statistical/online and agnostic/realizable boosting for *binary classification*.

Moving to agnostic *multiclass* boosting, Brukhim et al. [7] study the resources required for boosting in the statistical/batch setting as the number of labels k grows. However, they consider an alternative model of boosting where the weak learner is a strong agnostic PAC learner for a simple "easy-to-learn" base hypothesis class, and the goal is to learn target concepts outside the base class. Specifically, they assume the target concept can be represented by weighted plurality votes over the base class. In this way, the weakness of a weak learner is manifested in the base hypothesis class. Instead, in our work, we consider the standard boosting model of fixing the base hypothesis class, and ask whether a weak learner's performance can be improved relative to the best fixed hypothesis in that class.

Beyond classification, several other works have studied online agnostic boosting for real-valued loss functions under both full-information and bandit feedback settings [3, 1, 5]. In particular, the work by Brukhim and Hazan [5] reduce online boosting for regression tasks under bandit feedback to online linear optimization. A key difference between these works and ours is in the weak learning assumption: these works consider a weak learner that is a strong learner for a small base class of regression functions. The goal of boosting then is to produce a strong online learner for a larger class which contains linear spans of the base class. This is in contrast to this work, where again, we fix the base hypothesis class, and boost the regret bound.

2 Preliminaries and Notation

We first describe the basic setup for online agnostic multiclass boosting. There are k possible labels $\mathcal{B}_k := \{1, \dots, k\}$ and k is known to the booster and the weak learners. The booster maintains N copies of a weak learner, \mathcal{W} , which themselves are (randomized) online learning algorithms that sequentially process examples from instance space \mathcal{X} and output predictions in \mathcal{B}_k , which we denote as the set of basis vectors of length k . At each iteration $t = 1, \dots, T$, an adversary picks a labeled example $(x_t, y_t) \in \mathcal{X} \times \mathcal{B}_k$ and reveals x_t to the booster. Once the booster observes the unlabeled data x_t , it gathers the weak learners' predictions and makes a final (possibly randomized) prediction $\hat{y}_t \in \mathcal{B}_k$. After observing the booster's final decision, the adversary reveals the true label y_t , and the booster suffers the loss $\mathbb{1}\{\hat{y}_t \neq y_t\}$ or equivalently, *gains* the reward $2\mathbb{1}\{\hat{y}_t = y_t\} - 1$. Finally, the booster, after observing the true label y_t , updates each weak learner. Note that the loss/gain of the booster in round t can be written as $1 - \hat{y}_t \cdot y_t$ and $2\hat{y}_t \cdot y_t - 1$ respectively. We also let Δ_k represent the $(k - 1)$ -dimensional probability simplex and $\Delta_{\frac{k}{\gamma}}$ represent the $\frac{1}{\gamma}$ -scaled $(k - 1)$ -dimensional probability simplex for $\gamma \in (0, 1)$. Finally, we let $\mathbb{1}_k$ denote the k -dimensional ones vector.

Evaluation. Unlike the realizable setting, in the agnostic setting, we place no restrictions on how the stream of examples x_1, \dots, x_T are labelled. Thus, for a fixed hypothesis class $\mathcal{H} \subseteq \mathcal{B}_k^{\mathcal{X}}$, the goal of the booster is to output predictions \hat{y}_t such that the expected *regret*,

$$\mathbb{E} \left[\max_{h \in \mathcal{H}} \sum_{t=1}^T (2h(x_t) \cdot y_t - 1) - \sum_{t=1}^T (2\hat{y}_t \cdot y_t - 1) \right],$$

is minimized, where the expectation is over the randomness of the booster and that of the possibly adaptive adversary. Note, this is in contrast to the realizable setting where the stream is labelled by a $h \in \mathcal{H}$ and we wish the booster to minimize the (expected) number of *mistakes* (mistake-bound).

Agnostic Boosting. A key technique in agnostic boosting, first appearing in the work by Kanade and Kalai [27], is to update weak learners by feeding randomly *relabelled* examples. This is in contrast to the realizable setting where we typically update weak learners by passing *reweighted* examples.

Accordingly, in order to design a good boosting algorithm, we need to design the booster’s strategy for random relabelling while also quantifying the weak learner’s ability to maximize cumulative gain, even under relabelled data. The first task will be resolved by allowing the booster to use an Online Convex Optimization (OCO) oracle. In this way, we reduce boosting to OCO, an idea borrowed from Brukhim et al. [6]. For the second task, we give different possible weak learning conditions for the *same* algorithm, all of which capture the ability of a weak learner to maximize cumulative gain with respect to the best fixed competitor in hindsight.

Online Convex Optimization. Our booster will use an OCO oracle to update its weak learners. The OCO setting is a sequential game between an online player and adversary over N rounds (see [23] for an in-depth introduction). In each round, the player plays a point x_i in a compact convex set $\mathcal{K} \subset \mathbb{R}^d$, the adversary reveals a loss function f_i chosen from a family of bounded convex functions over \mathcal{K} , and the player suffers the loss $f_i(x_i)$. The goal of player is to also minimize *regret*, defined as

$$R(N) = \sum_{i=1}^N f_i(x_i) - \min_{x \in \mathcal{K}} \sum_{i=1}^N f_i(x).$$

We will denote an algorithm in this setting as $\text{OCO}(\mathcal{K}, N)$. If \mathcal{A} is a $\text{OCO}(\mathcal{K}, N)$, then we will denote its regret by $R_{\mathcal{A}}(N)$. For many OCO algorithms, like Online Gradient Descent (OGD), the regret $R_{\mathcal{A}}(N)$ is a sub-linear function of the time-horizon N .

3 Online Agnostic Boosting

In this section, we present an online agnostic boosting algorithm and analyze its performance. We begin in Subsection 3.1 by formally describing our algorithm which reduces boosting to OCO. Then, in Subsection 3.2, we state a weak learning condition and subsequently prove a regret bound for our proposed algorithm.

3.1 Algorithm

Pseudocode for our online agnostic boosting algorithm is provided in Algorithm 1. The booster maintains oracle access to N copies of a weak learner $\mathcal{W}_1, \dots, \mathcal{W}_N$ as well as a $\text{OCO}(\Delta_k, N)$ algorithm \mathcal{A} . Each weak learner is characterized by some advantage parameter $\gamma \in (0, 1]$ (directly proportional to its strength), and we will be precise about what exactly γ quantifies in Subsection 3.2. In round t , the booster uses the weak learners to make a prediction \hat{y}_t , observes the true label y_t , and finally simulates a game with the OCO algorithm \mathcal{A} to update each weak learner \mathcal{W}_i . Specifically, the booster uses the outputs of \mathcal{A} to feed relabelled examples (x_t, y_t^i) to \mathcal{W}_i . A critical component of our algorithm is the L_2 projection operator onto the probability simplex Δ_k , which we denote by Π . This is used by the booster in line 4 to make randomized predictions \hat{y}_t .

We now highlight some desirable properties of Algorithm 1. First, it is easy to implement and efficient assuming access to an efficient weak learner. In particular, for each round t , if the running time of the weak learner is Q , then the running time of our booster is $O(NQ + Nk \log(k))$. This is in contrast to the OnlineMBBM algorithm proposed by Jung et al. [24] for the realizable setting which is not efficient even assuming access to an efficient weak learner. Second, when $k = 2$, our algorithm reduces down to the online agnostic boosting algorithm proposed by Brukhim et al. [6] in the binary case. Indeed, one can verify that when $k = 2$, the L_2 projection operator reduces to the same projection used by Brukhim et al. [6] and the distribution over labels p_t^i induced by the outputs of OCO algorithm are equal for each weak learner in every round t .

While the framework in Algorithm 1 is inspired from Brukhim et al. [6] in the binary setting, several new complications arise in the multiclass setting. Most notably, when $k > 2$, we must figure out how the booster should make predictions, what loss functions the booster should construct and pass to \mathcal{A} , and lastly, how the booster should use the output of \mathcal{A} to update each weak learner. The interplay behind these three algorithmic pieces is delicate and they have been carefully designed in Algorithm 1 to enable the analysis in Subsection 3.2. Below, we provide some intuition behind these algorithmic decisions.

Randomized Prediction. At the start of each round, the booster averages the weak learners votes, scales the average by the parameter γ , projects the scaled vector back into the simplex, and finally samples a random label. When the L_2 projection operator is selected, this approach for randomized

Algorithm 1: Online Agnostic Multiclass Boosting via OCO

Input: Weak Learners $\mathcal{W}_1 \dots \mathcal{W}_N$, OCO(Δ_k , N) algorithm \mathcal{A} , Advantage parameter γ

```

1 for  $t = 1, \dots, T$  do
2   Receive example  $x_t$ 
3   Accumulate weak predictions  $h_t = \sum_{i=1}^N \mathcal{W}_i(x_t)$ 
4   Set  $\mathcal{D}_t = \prod(\frac{h_t}{\gamma N})$ 
5   Predict  $\hat{y}_t \sim \mathcal{D}_t$ 
6   Receive true label  $y_t$ 
7   for  $i = 1, \dots, N$  do
8     If  $i > 1$ , obtain  $p_t^i = \mathcal{A}(l_t^1, \dots, l_t^{i-1})$ . Else, initialize  $p_t^1 = \frac{1}{k}$ .
9     Reveal loss function:  $l_t^i(p) = p \cdot \left( \frac{2\mathcal{W}_i(x_t) - \mathbb{1}_k}{\gamma} - (2y_t - \mathbb{1}_k) \right)$ 
10    Sample random label  $y_t^i \sim p_t^i$ , and pass  $(x_t, y_t^i)$  to  $\mathcal{W}_i$ 
11  Reset  $\mathcal{A}$ 

```




Figure 1: The red, blue, and black lines correspond to $\frac{\Delta_2}{\gamma}$ for $\gamma = 0.50, 0.75$, and, 1.0 respectively. p_1 and p_2 are γ -scaled votes and p_1^* and p_2^* are their corresponding L_2 projection onto Δ_2 . \tilde{p}_1 and \tilde{p}_2 denote the orthogonal projections onto the plane containing Δ_2 .

prediction achieves a *polarization* effect: as γ gets smaller, the projection concentrates mass on a fewer number of labels, specifically those labels that have achieved the majority of the votes from the weak learners. When γ nears 0, the L_2 projection eventually places all mass on the label with the most votes.

One might think that the L_2 projection is not a natural projection operator for Δ_k (as KL or L_1 projections might seem better suited to the geometry of the probability simplex). However, in the case where we project from $\Delta_{\frac{k}{\gamma}}$ to Δ_k , we find that it is a natural choice from a geometric perspective. Figure 1 provides a visualization of the L_2 projection operator from $\Delta_{\frac{k}{\gamma}}$ to Δ_k for $k = 2$. Because the spaces $\Delta_{\frac{k}{\gamma}}$ and Δ_k are parallel, given any point $p \in \Delta_{\frac{k}{\gamma}}$, one can think of its L_2 projection in the following procedural manner:

1. Compute the *orthogonal* projection, \tilde{p} , onto the plane containing Δ_k .
2. If $\tilde{p} \in \Delta_k$, output $p^* = \tilde{p}$.
3. Else, output $p^* = \arg \min_{p \in \Delta_k} \|p - \tilde{p}\|_2$.

From this procedural perspective, Figure 1 lends a geometric intuition behind the polarization effect of the L_2 projection. As γ shrinks, the number of points in $\Delta_{\frac{k}{\gamma}}$ that lie orthogonally above Δ_k shrinks. Thus, the orthogonal projection of the vast majority of points in $\Delta_{\frac{k}{\gamma}}$ will lie outside Δ_k ,

leading to a greater number of sparse projections that lie on the boundary of Δ_k . From the booster’s viewpoint, this property of the L_2 projection is desirable. If the weak learners are very weak (small γ) but somehow concentrate votes on a few labels, then it may be likely that the true response is amongst these few labels. In this sense, γ controls how many weak learners need to agree on a particular label to convince the booster to deterministically predict that label.

Updating Weak Learners. Once the true label y_t is revealed in each round, the booster must update each weak learner. As mentioned in Section 2, one strategy for updating weak learners in the agnostic setting is via *random relabelling*. Indeed, in line 10, the booster passes back to the weak learner the example x_t with a random label $y_t^i \sim p_t^i$. Together, the specified loss function in line 9 and random relabelling strategy in line 10 achieve the following effect: if more weak learners make mistakes, the distribution over labels output by the OCO algorithm \mathcal{A} in line 8 concentrates on the true label y_t , increasing the likelihood that y_t is passed to subsequent weak learners. This is desirable as the outputs of the OCO guide each weak learner to correct for the mistakes of preceding learners.

3.2 Regret Analysis

Before we give the regret bound of Algorithm 1, we need to specify the capacity to which the weak learners can make predictions, even under potentially relabelled data. Unfortunately, in the agnostic setting, there is no canonical weak learning condition for multiclass problems. In this paper, we give *several* possible definitions of an Agnostic Weak Online Learner (AWOL) for the multiclass setting, all of which enable a regret analysis of Algorithm 1. We emphasize that we can derive different regret bounds for Algorithm 1 based on what condition we assume our weak learners to satisfy. For this section, we present a weak learning condition based loosely on a *one-vs-one* approach to multiclass classification. In Appendix C, we provide alternative weak learning conditions.

Define the gain function for input $z \in \mathcal{B}_k$,

$$\sigma_{y,\ell}(z) = \mathbb{1}\{z = y\} - \mathbb{1}\{z = \ell\} = z \cdot (y - \ell).$$

For some $y \in \mathcal{B}_k$ and $\ell \in \mathcal{B}_k \setminus \{y\}$, $\sigma_{y,\ell}(\cdot)$ can be thought of as the binary classification task between labels y and ℓ . Definition 1 requires that for any such sequence of binary classification tasks, an online agnostic weak learner must be able to eventually distinguish between every pair of labels to some non-trivial, but far from optimal, degree.

Definition 1 (Agnostic Weak Online Learning). *Let $\mathcal{H} \subseteq \mathcal{B}_k^{\mathcal{X}}$ be a class of experts and let $0 < \gamma \leq 1$ denote the “advantage”. An online learning algorithm \mathcal{W} is a (γ, T) -agnostic weak online learner (AWOL) for \mathcal{H} if for any adaptively chosen sequence of tuples $(x_t, y_t, \ell_t) \in \mathcal{X} \times \mathcal{B}_k \times \mathcal{B}_k$ where $\ell_t \neq y_t$, the algorithm outputs $\mathcal{W}(x_t) \in \mathcal{B}_k$ at every iteration $t \in [T]$ such that,*

$$\gamma \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t=1}^T \sigma_{y_t, \ell_t}(h(x_t)) \right] - \mathbb{E} \left[\sum_{t=1}^T \sigma_{y_t, \ell_t}(\mathcal{W}(x_t)) \right] \leq R_{\mathcal{W}}(T, k),$$

where the expectation is taken w.r.t. the randomness of the weak learner \mathcal{W} and that of the possibly adaptive adversary, and $R_{\mathcal{W}} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}_+$ is the additive regret: a non-decreasing, sub-linear function of T .

The precise dependence of $R_{\mathcal{W}}(T, k)$ on k is explored in more detail in Appendix D where we explicitly construct learners satisfying AWOL. We make few remarks about Definition 1. First, the strength of the weak learner varies directly with γ . Second, for $k = 2$, Definition 1 reduces to the weak learning condition by Brukhim et al. [6] in the binary setting. Finally, we emphasize that Definition 1 holds under an *adaptive* adversary, one that can choose (y_t, ℓ_t) based on $\{\mathcal{W}(x_i)\}_{i=1}^{t-1}$ and its own internal random bits. Importantly, we also allow the adversary to pick ℓ_t even after it has observed $\mathcal{W}(x_t)$ as this only controls how much loss the weak learner suffers when it is incorrect.

Under the assumption that weak learners satisfy Definition 1, Theorem 1 bounds the regret of Algorithm 1 under an *oblivious* adversary. Using a standard reduction, our results can then be generalized to an *adaptive* adversary (see Chapter 4 in [8]). In particular, a key requirement allowing an oblivious regret bound to generalize to an adaptive regret bound is that the learner’s predictions on round t should not depend on any of its past predictions from previous rounds. This is indeed true for our Booster.

Theorem 1 (Regret Bound). Assuming weak learners satisfy Definition 1, the expected regret bound of Algorithm 1 is,

$$\frac{1}{T} \mathbb{E} \left[\max_{h \in \mathcal{H}} \sum_{t=1}^T (2h(x_t) \cdot y_t - 1) - \sum_{t=1}^T (2\hat{y}_t \cdot y_t - 1) \right] \leq \frac{R_{\mathcal{W}}(T, k)}{\gamma T} + \frac{R_{\mathcal{A}}(N)}{N},$$

where the expectation is over the randomness of the algorithm and weak learners, and $R_{\mathcal{W}}(T, k)$, $R_{\mathcal{A}}(N)$ are the regret terms of the AWOL, OCO algorithms respectively.

If one picks \mathcal{A} to be *Online Gradient Descent* (OGD), then $R_{\mathcal{A}}(N) = O(GD\sqrt{N})$, where D is the diameter of Δ_k and G is the upper-bound on $\|\nabla_p \ell_t^i(p)\|$. In our setting, $D = \sqrt{2}$ and $G = O(\frac{1}{\gamma})$ (using Lemma 19) and hence the average regret further simplifies to:

$$\frac{1}{T} \mathbb{E} \left[\max_{h \in \mathcal{H}} \sum_{t=1}^T (2h(x_t) \cdot y_t - 1) - \sum_{t=1}^T (2\hat{y}_t \cdot y_t - 1) \right] \leq \frac{R_{\mathcal{W}}(T, k)}{\gamma T} + O\left(\frac{1}{\gamma\sqrt{N}}\right).$$

To exemplify the role of N , consider a scenario where $R_{\mathcal{W}}(T, k) = O(k\sqrt{T})$ (see Appendix D for examples). By setting $N = \frac{T}{\gamma^2}$ the overall regret of the booster becomes $O(\frac{k\sqrt{T}}{\gamma})$. In the next subsection, we give the proof of Theorem 1.

3.2.1 Proof of Theorem 1

We follow a similar procedure to Bruchim et al. [6] by lower and upper bounding the expected sum of losses passed to \mathcal{A} in terms of the regret of the weak learner and the regret of \mathcal{A} respectively. These bounds rely on several lemmas that have been abstracted out and provided in Appendix E. As mentioned previously, we also assume an oblivious adversary.

Let $(x_1, y_1), \dots, (x_T, y_T)$ be any sequence of example-label pairs. We start by giving a lower bound on the expected sum of losses passed to \mathcal{A} using Definition 1. Define $h^* = \arg \max_{h \in \mathcal{H}} \sum_{t=1}^T (2h(x_t) \cdot y_t - 1)$ as the optimal competitor in hindsight and let

$$\ell_t^i = \begin{cases} \mathcal{W}_i(x_t), & \text{if } \mathcal{W}_i(x_t) \neq y_t^i \\ \ell \in \mathcal{B}_k \setminus \{y_t^i\}, & \text{otherwise} \end{cases}.$$

Note that the precise choice of ℓ_t^i in the second case is not important. The proof below only uses the fact that $\ell_t^i \in \mathcal{B}_k \setminus \{y_t^i\}$ in that case. We then have,

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T \ell_t^i(p_t^i) \right] &= \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T p_t^i \cdot \left(\frac{2\mathcal{W}_i(x_t) - \mathbb{1}_k}{\gamma} - (2y_t - \mathbb{1}_k) \right) \right] \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2\mathcal{W}_i(x_t) - \mathbb{1}_k)] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2\mathcal{W}_i(x_t) \cdot y_t^i - 1] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \quad (\text{Lemma 15}) \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [\sigma_{y_t^i, \ell_t^i}(\mathcal{W}_i(x_t))] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)]. \end{aligned}$$

Using the weak learning condition in Definition 1,

$$\begin{aligned}
\frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} \left[\sigma_{y_t^i, \ell_t^i}(\mathcal{W}_i(x_t)) \right] &\geq \sum_{i=1}^N \max_{h \in \mathcal{H}} \sum_{t=1}^T \mathbb{E} \left[\sigma_{y_t^i, \ell_t^i}(h(x_t)) \right] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} && \text{(Definition 1)} \\
&\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} \left[\sigma_{y_t^i, \ell_t^i}(h^*(x_t)) \right] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \\
&\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2h^*(x_t) \cdot y_t^i - 1] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \\
&= \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2h^*(x_t) - \mathbb{1}_k)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma}. && \text{(Lemma 15)}
\end{aligned}$$

Putting things together, we find,

$$\begin{aligned}
\mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T l_t^i(p_t^i) \right] &\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2h^*(x_t) - 2y_t)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \\
&\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2(h^*(x_t) \cdot y_t - 1)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} && \text{(Lemma 16)} \\
&= N \sum_{t=1}^T 2(h^*(x_t) \cdot y_t - 1) - \frac{NR_{\mathcal{W}}(T, k)}{\gamma}.
\end{aligned}$$

Now, we compute an upper bound. For any $t \in [T]$ and arbitrary $p^* \in \Delta_k$:

$$\begin{aligned}
\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N l_t^i(p_t^i) \right] &\leq \frac{1}{N} \mathbb{E} \left[\min_{p \in \mathcal{K}} \sum_{i=1}^N l_t^i(p) \right] + \frac{R_{\mathcal{A}}(N)}{N} && \text{(OCO Regret)} \\
&\leq \frac{1}{N} \mathbb{E} \left[\sum_{i=1}^N l_t^i(p^*) \right] + \frac{R_{\mathcal{A}}(N)}{N} \\
&= \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[p^* \cdot \left(\frac{2\mathcal{W}_i(x_t) - \mathbb{1}_k}{\gamma} - (2y_t - \mathbb{1}_k) \right) \right] + \frac{R_{\mathcal{A}}(N)}{N} \\
&= \frac{1}{\gamma N} \sum_{i=1}^N (p^* \cdot \mathbb{E} [2\mathcal{W}_i(x_t) - \mathbb{1}_k]) - p^* \cdot (2y_t - \mathbb{1}_k) + \frac{R_{\mathcal{A}}(N)}{N} \\
&= p^* \cdot \left(\frac{1}{\gamma N} \sum_{i=1}^N \mathbb{E} [2\mathcal{W}_i(x_t) - \mathbb{1}_k] \right) - p^* \cdot (2y_t - \mathbb{1}_k) + \frac{R_{\mathcal{A}}(N)}{N} \\
&= \mathbb{E} \left[p^* \cdot \left(\frac{\frac{2}{N} \sum_{i=1}^N \mathcal{W}_i(x_t) - \mathbb{1}_k}{\gamma} - (2y_t - \mathbb{1}_k) \right) \right] + \frac{R_{\mathcal{A}}(N)}{N} \\
&\leq \mathbb{E} \left[2 \left(\prod \left(\frac{1}{\gamma N} \sum_{i=1}^N \mathcal{W}_i(x_t) \right) \cdot y_t - 1 \right) \right] + \frac{R_{\mathcal{A}}(N)}{N} && \text{(Lemma 17)} \\
&= 2 (\mathbb{E} [\hat{y}_t] \cdot y_t - 1) + \frac{R_{\mathcal{A}}(N)}{N}. && \text{(Law of total expectation)}
\end{aligned}$$

Summing over T ,

$$\mathbb{E} \left[\frac{1}{N} \sum_{t=1}^T \sum_{i=1}^N l_t^i(p_t^i) \right] \leq \sum_{t=1}^T 2 (\mathbb{E} [\hat{y}_t] \cdot y_t - 1) + \frac{TR_{\mathcal{A}}(N)}{N}.$$

Combining lower and upper bounds for $\mathbb{E} \left[\frac{1}{NT} \sum_{i=1}^N \ell_t^i(p_t^i) \right]$, we get

$$\frac{1}{T} \mathbb{E} \left[\max_{h \in \mathcal{H}} \sum_{t=1}^T (2h(x_t) \cdot y_t - 1) - \sum_{t=1}^T (2\hat{y}_t \cdot y_t - 1) \right] \leq \frac{R_{\mathcal{W}}(T, k)}{\gamma T} + \frac{R_{\mathcal{A}}(N)}{N},$$

which completes the proof. Since $2h(x_t) \cdot y_t - 1 = 1 - 2\mathbb{1}_k\{h(x_t) \neq y_t\}$, we can also write the average regret in terms of the number of *mistakes*,

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}\{\hat{y}_t \neq y_t\} - \min_{h \in \mathcal{H}} \sum_{t=1}^T \mathbb{1}\{h(x_t) \neq y_t\} \right] \leq \frac{R_{\mathcal{W}}(T, k)}{2\gamma T} + \frac{R_{\mathcal{A}}(N)}{2N}.$$

4 Beyond Online Agnostic Boosting

In this section, we present results of extending our reduction to the three other boosting settings, namely statistical agnostic, online realizable, and statistical realizable learning. The purpose of this section is to showcase the generality of the OCO-based boosting framework and not to achieve state-of-the-art bounds for these settings. Boosting algorithms and all associated proofs can be found in Appendix A and B. Throughout this section we will let \mathcal{W}_S denote the hypothesis output by a weak learner trained on a sample S .

4.1 Statistical Agnostic Boosting

In the statistical setting, our objective of interest is the *correlation*, which we define below. Let \mathcal{D} be a distribution over $\mathcal{X} \times \mathcal{B}_k$ and let $h : \mathcal{X} \rightarrow \mathcal{B}_k$ be an hypothesis. Define the *multiclass* correlation of h with respect to \mathcal{D} as

$$\text{cor}_{\mathcal{D}}(h) = \mathbb{E}_{(x,y) \sim \mathcal{D}} [h(x) \cdot (2y - \mathbb{1}_k)].$$

Like the online agnostic setting, the boosting algorithm for this setting (provided in Appendix A) can be analyzed under several candidate weak learning conditions. To showcase the dependence on k , we provide a weak learning condition based loosely on a *one-vs-all* approach to multiclass classification.

Definition 2 (Empirical Agnostic Weak Learning). *Let $\mathcal{H} \subseteq \mathcal{B}_k^{\mathcal{X}}$ be a hypothesis class and let $0 < \gamma \leq 1$ denote the “advantage”. Let $\mathbf{x} = (x_1, \dots, x_m) \in \mathcal{X}$ denote an unlabeled sample. A learning algorithm \mathcal{W} is a $(\gamma, \epsilon_0, m_0)$ -agnostic weak learner (AWL) for \mathcal{H} with respect to \mathbf{x} if for any labels $\mathbf{y} = (y_1, \dots, y_m) \in \mathcal{B}_k$, and every reference label $\ell \in \mathcal{B}_k$,*

$$\mathbb{E}_{S'} \left[\sum_{i: y_i = \ell} \mathcal{W}_{S'}(x_i) \cdot (2y_i - \mathbb{1}_k) \right] \geq \gamma \max_{h \in \mathcal{H}} \sum_{i: y_i = \ell} h(x_i) \cdot (2y_i - \mathbb{1}_k) - m\epsilon_0,$$

where S' is an independent sample of size m_0 drawn from the distribution which uniformly assigns to each example (x_i, y_i) probability $1/m$.

Under the assumption that our weak learners satisfy Definition 2, Corollary 2 bounds the expected correlation of a statistical agnostic boosting algorithm.

Corollary 2 (Empirical Agnostic Correlation Bound). *There exists a boosting algorithm whose output after T rounds, denoted by \bar{h} , satisfies:*

$$\mathbb{E} [\text{cor}_S(\bar{h})] \geq \max_{h \in \mathcal{H}} \mathbb{E} [\text{cor}_S(h)] - \frac{R_{\mathcal{A}}(T)}{T} - \frac{k\epsilon_0}{\gamma},$$

where S is the distribution which uniformly assigns to each example (x_i, y_i) probability $1/m$.

Letting \mathcal{A} be OGD, and setting $T = O(\frac{1}{\gamma^2 \epsilon^2})$ for any $\epsilon > 0$ gives an error rate of $\frac{k\epsilon_0}{\gamma} + \epsilon$. Although, to our best knowledge, there are no existing multiclass boosting algorithms in the agnostic setting, we point out that there is evidence to suggest that a sub-optimal dependence of the error rate on k might be unavoidable (see [7]).

4.2 Online Realizable Boosting

In the online realizable setting, we are guaranteed that the stream of examples $(x_1, y_1), \dots, (x_T, y_T)$ is perfectly labelled by some hypothesis $h \in \mathcal{H}$. Definition 3 then gives an appropriate online realizable weak learning condition.

Definition 3 (Realizable Weak Online Learning). *Let $\mathcal{H} \subseteq \mathcal{B}_k^{\mathcal{X}}$ be a class of experts and let $0 < \gamma \leq 1$ denote the “advantage”. An online learning algorithm \mathcal{W} is a (γ, T) -realizable weak online learner (**RWOL**) for \mathcal{H} if for any sequence $(x_1, y_1), \dots, (x_T, y_T) \in \mathcal{X} \times \mathcal{B}_k$ that is realizable by \mathcal{H} , at every iteration $t \in [T]$, the algorithm outputs $\mathcal{W}(x_t) \in \mathcal{B}_k$ such that,*

$$\mathbb{E} \left[\sum_{t=1}^T (2\mathcal{W}(x_t) \cdot y_t - 1) \right] \geq \gamma T - R_{\mathcal{W}}(T, k),$$

where the expectation is taken w.r.t. the randomness of the weak learner \mathcal{W} and $R_{\mathcal{W}} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}_+$ is the additive regret: a non-decreasing, sub-linear function of T .

Under Definition 3, Corollary 3 bounds the expected gain of an online boosting algorithm.

Corollary 3 (Mistake Bound). *There exists a boosting algorithm whose outputs $\hat{y}_1, \dots, \hat{y}_T$ satisfy:*

$$\frac{1}{T} \sum_{t=1}^T (2\mathbb{E}[\hat{y}_t] \cdot y_t - 1) \geq 1 - \frac{R_{\mathcal{A}}(N)}{N} - \frac{\tilde{R}_{\mathcal{W}}(T, k)}{\gamma T},$$

where $\tilde{R}_{\mathcal{W}}(T, k) = 2R_{\mathcal{W}}(T, k) + \tilde{O}(\sqrt{T})$.

If we consider a scenario where $\tilde{R}_{\mathcal{W}}(T, k) = O(k\sqrt{T})$, taking \mathcal{A} to be OGD, N to be $O(\frac{1}{\gamma^2\epsilon^2})$ and T to be $O(\frac{1}{\gamma^2\epsilon^2})$, gives a mistake-bound at most ϵkT . We point out that we cannot readily compare this bound with the existing bounds for online realizable boosting by Jung et al. [24] because the dependence of $R_{\mathcal{W}}(T, k)$ on k depends on the weak learner of choice.

4.3 Statistical Realizable Boosting

In the statistical realizable setting, our metric of interest is again the multiclass correlation. However, under the realizability assumption, $\max_{h^* \in \mathcal{H}} \text{cor}_D(h^*) = 1$. Accordingly, Definition 4 gives an appropriate realizable weak learning condition.

Definition 4 (Empirical Realizable Weak Learning). *Let $\mathcal{H} \subseteq \mathcal{B}_k^{\mathcal{X}}$ be a hypothesis class and let $0 < \gamma \leq 1$ denote the “advantage”. Let $S = \{(x_1, y_1), \dots, (x_m, y_m)\} \in \mathcal{X} \times \mathcal{B}_k$ be a sample. A learning algorithm \mathcal{W} is a (γ, m_0) -realizable weak learner (**RWL**) for \mathcal{H} with respect to S if for any distribution $\mathbf{p} = (p_1, \dots, p_m)$ over the examples,*

$$\mathbb{E}_{S'} [\text{cor}_{\mathbf{p}}(\mathcal{W}_{S'})] \geq \gamma,$$

where S' is an independent sample of size m_0 drawn from \mathbf{p} .

Under Definition 4, Corollary 4 bounds the expected correlation of a statistical realizable boosting algorithm.

Corollary 4 (Empirical Correlation Bound). *There exists a boosting algorithm whose output after T rounds, denoted \bar{h} , satisfies,*

$$\mathbb{E} [\text{cor}_S(\bar{h})] \geq 1 - \frac{R_{\mathcal{A}}(T)}{T},$$

where S is the distribution which uniformly assigns to each example (x_i, y_i) probability $1/m$.

Note that the cost of weak learning manifests in the regret term of the OCO algorithm. Precisely, if one picks OGD to be the OCO algorithm, then the bound on correlation can be expressed as $\mathbb{E} [\text{cor}_S(\bar{h})] \geq 1 - O\left(\frac{1}{\gamma\sqrt{T}}\right)$, which exhibits a decreasing lower bound as γ shrinks. Setting $T = O(\frac{1}{\gamma^2\epsilon^2})$ for any $\epsilon > 0$ ensures that at most ϵ error is obtained. It is difficult to compare our error rates to that achieved by existing multiclass boosting algorithm in the realizable setting because our weak learning condition is different and potentially stronger. Nevertheless, our bounds are sub-optimal to those in Mukherjee and Schapire [30] which only require setting $T = O(\frac{1}{\gamma^2} \log(\frac{k}{\epsilon}))$ to achieve ϵ error.

Table 1: Average accuracy and runtime of algorithms on 7 UCI datasets.

Dataset	k	Accuracy(%)				Runtime(s)			
		Agn	Opt	Ada	OCOR	Agn	Opt	Ada	OCOR
Balance	3	84.4	79.5	75.2	78.8	11.5	57.6	7.1	4.5
Cars	4	80.6	69.9	75.8	69.6	42.5	104.1	20.4	12.1
Landsat	6	67.0	80.8	56.9	79.6	1255.6	1814.3	81.8	440.4
Segment	7	75.0	78.9	68.6	79.3	647.2	2448.3	51.8	154.7
Mice	8	86.0	77.6	71.3	79.6	1025.8	1811.6	90.1	258.2
Yeast	10	42.3	39.8	41.7	47.6	348.3	1468.8	25.2	56.6
Abalone	28	22.1	24.9	19.2	22.0	4036.7	9775.3	108.0	248.2

5 Experiments

We performed experiments with Algorithm 1 on seven UCI datasets [13]. Since, to our knowledge, there are no known online *agnostic* boosting algorithms for multiclass problems, we benchmark performance against three online *realizable* multiclass boosting algorithms: the state-of-the-art optimal (OnlineMBBM) and adaptive (Adaboost.OLM) online boosting algorithms by Jung et al. [24], and the OCO-based online boosting algorithm alluded to in Section 4 (see Appendix B). For weak learners, we used the implementation of the VeryFastDecisionTree from the River package [29] and restricted the maximum depth of the tree to 1. We used Projected OGD [35] for the OCO algorithm and set the number of weak learners, N , to 100 for each boosting algorithm.

Our experiments using Algorithm 1 are performed using *fractional relabeling*, a technique borrowed from Kanade and Kalai [27]. That is, instead of passing just a single example-label pair (x_t, y_t^i) where $y_t^i \sim p_t^i$, we pass all k possible *weighted* example-label tuples $\{(x_t, \ell, p_t^i[\ell])\}_{\ell=1}^k$ to weak learner i in round t . Experiments with random relabeling showed that random relabeling runs faster but performs worse than fractional relabeling.

Table 1 summarizes the average accuracy and runtime over five independent shuffles of each dataset for each boosting algorithm. "Agn" refers to Algorithm 1, "Opt" to OnlineMBBM, "Ada" to Adaboost.OLM, and "OCOR" to the OCO-based online realizable algorithm. Standard errors for accuracies are provided in Appendix F. All algorithms **except** AdaBoost.OLM are parameterized by an advantage parameter $\gamma \in (0, 1)$. Thus, γ was tuned separately for each respective cell of Table 1. See Appendix F for more experimental details. All code is available at <https://github.com/vinodkraman/OnlineAgnosticMulticlassBoosting>.

Despite having the fastest overall runtime, Adaboost.OLM had the poorest performance. Compared to OnlineMBBM, our OCO-based boosting algorithms achieve comparable performance at a fraction of the runtime. Specifically, for three of our datasets, our online agnostic boosting algorithm achieves the highest accuracy. For the remaining datasets, our OCO-based realizable boosting algorithm achieves comparable accuracy to OnlineMBBM with shorter runtimes.

6 Discussion

We give the first weak learning conditions and algorithm for online agnostic multiclass boosting. Our algorithm relies on a clean and simple reduction from boosting to online convex optimization. This fruitful connection allows us to go beyond the online agnostic setting and design multiclass boosting algorithms for all four regimes of statistical/online and agnostic/realizable learning. As future work, we leave it open to identify the correct weak learning condition, construct adaptive versions of our boosting algorithms, improve regret upper bounds/prove lower bounds, design agnostic boosting algorithms under bandit feedback, and study the impact of the choice of weak learner and OCO algorithm on the empirical performance of our OCO-based boosting algorithms.

Acknowledgements. AT acknowledges the support of NSF via grant IIS-2007055.

References

- [1] N. Agarwal, N. Brukhim, E. Hazan, and Z. Lu. Boosting for control of dynamical systems. In *International Conference on Machine Learning*, pages 96–103. PMLR, 2020.
- [2] S. Ben-David, P. M. Long, and Y. Mansour. Agnostic boosting. In *International Conference on Computational Learning Theory*, pages 507–516. Springer, 2001.
- [3] A. Beygelzimer, E. Hazan, S. Kale, and H. Luo. Online gradient boosting. *Advances in neural information processing systems*, 28, 2015.
- [4] A. Beygelzimer, S. Kale, and H. Luo. Optimal and adaptive algorithms for online boosting. In *International Conference on Machine Learning*, pages 2323–2331. PMLR, 2015.
- [5] N. Brukhim and E. Hazan. Online boosting with bandit feedback. In *Algorithmic Learning Theory*, pages 397–420. PMLR, 2021.
- [6] N. Brukhim, X. Chen, E. Hazan, and S. Moran. Online agnostic boosting via regret minimization. *Advances in Neural Information Processing Systems*, 33:644–654, 2020.
- [7] N. Brukhim, E. Hazan, S. Moran, I. Mukherjee, and R. E. Schapire. Multiclass boosting and the cost of weak learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [8] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, USA, 2006. ISBN 0521841089.
- [9] S.-T. Chen, H.-T. Lin, and C.-J. Lu. An online boosting algorithm with theoretical justifications. *arXiv preprint arXiv:1206.6422*, 2012.
- [10] S.-T. Chen, H.-T. Lin, and C.-J. Lu. Boosting with online binary learners for the multiclass bandit problem. In *International Conference on Machine Learning*, pages 342–350. PMLR, 2014.
- [11] Y. Chen and X. Ye. Projection onto a simplex. *arXiv preprint arXiv:1101.6081*, 2011.
- [12] A. Daniely, S. Sabato, S. Ben-David, and S. Shalev-Shwartz. Multiclass learnability and the erm principle. *J. Mach. Learn. Res.*, 16(1):2377–2404, 2015.
- [13] C. L. DuBois. UCI network data repository, 2008. URL <http://networkdata.ics.uci.edu>.
- [14] G. Eibl and K.-P. Pfeiffer. Multiclass boosting for weak classifiers. *Journal of Machine Learning Research*, 6(2), 2005.
- [15] V. Feldman. Distribution-specific agnostic boosting. *arXiv preprint arXiv:0909.2927*, 2009.
- [16] Y. Freund. Boosting a weak learning algorithm by majority. *Information and computation*, 121(2):256–285, 1995.
- [17] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- [18] Y. Freund and R. E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.
- [19] Y. Freund, R. E. Schapire, et al. Experiments with a new boosting algorithm. In *icml*, volume 96, pages 148–156. Citeseer, 1996.
- [20] Y. Freund, R. Schapire, and N. Abe. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612, 1999.
- [21] D. Gavinsky. Optimally-smooth adaptive boosting and application to agnostic learning. *Journal of Machine Learning Research*, 4(May):101–117, 2003.
- [22] T. Hastie, S. Rosset, J. Zhu, and H. Zou. Multi-class adaboost. *Statistics and its Interface*, 2(3): 349–360, 2009.

- [23] E. Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [24] Y. H. Jung, J. Goetz, and A. Tewari. Online multiclass boosting. *Advances in neural information processing systems*, 30, 2017.
- [25] A. T. Kalai and R. A. Servedio. Boosting in the presence of noise. *Journal of Computer and System Sciences*, 71(3):266–290, 2005.
- [26] A. T. Kalai, Y. Mansour, and E. Verbin. On agnostic boosting and parity learning. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 629–638, 2008.
- [27] V. Kanade and A. Kalai. Potential-based agnostic boosting. *Advances in neural information processing systems*, 22, 2009.
- [28] Y. Mansour and D. McAllester. Boosting using branching programs. *Journal of Computer and System Sciences*, 64(1):103–112, 2002.
- [29] J. Montiel, M. Halford, S. M. Mastelini, G. Bolmier, R. Sourty, R. Vaysse, A. Zouitine, H. M. Gomes, J. Read, T. Abdessalem, and A. Bifet. River: machine learning for streaming data in python, 2020.
- [30] I. Mukherjee and R. E. Schapire. A theory of multiclass boosting. *Journal of Machine Learning Research*, 2013.
- [31] N. C. Oza and S. J. Russell. Online bagging and boosting. In *International Workshop on Artificial Intelligence and Statistics*, pages 229–236. PMLR, 2001.
- [32] R. E. Schapire. The strength of weak learnability. *Machine learning*, 5(2):197–227, 1990.
- [33] R. E. Schapire and Y. Freund. Boosting: Foundations and algorithms. *Kybernetes*, 2013.
- [34] M. Sion. On general minimax theorems. *Pacific Journal of mathematics*, 8(1):171–176, 1958.
- [35] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 928–936, 2003.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) See Section 6 where we identify future work.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[No\]](#) Due to the theoretical/general nature of the work, we do not foresee any major/direct ethical or societal consequences of the research we presented.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#) See Section 3.2, Section 4, Appendix A, Appendix B, and Appendix C
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#) See Section 3.2.1, Appendix A, Appendix B, Appendix D, and Appendix E
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) In order to preserve anonymity, we will include a URL to a github repository if accepted

- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 5 and Appendix F
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Appendix F
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix F
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 5
 - (b) Did you mention the license of the assets? [Yes] See Appendix F
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] None of the datasets we are using contain information about people.
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Statistical Agnostic Boosting

We fill in the gaps of Section 4 for the statistical agnostic setting by providing a boosting algorithm and its theoretical analysis.

A.1 Algorithm

We now describe a statistical agnostic boosting algorithm whose pseudocode is provided below. The booster is given as input a sample $S = (x_1, y_1), \dots, (x_m, y_m) \in \mathcal{X} \times \mathcal{B}_k$. The booster has black-box oracle access to T copies of a $(\gamma, \epsilon_0, m_0)$ -AWL algorithm, $\mathcal{W}_1, \dots, \mathcal{W}_T$, each satisfying Definition 2, and m copies of an $\text{OCO}(\Delta_k, T)$ algorithm, $\mathcal{A}^1, \dots, \mathcal{A}^m$. Importantly, note that in line 8, we denote by \prod as a randomized *prediction* operator that given a γ -scaled distribution over k labels, first computes the L_2 projection onto Δ_k , and then randomly samples a label from the resulting distribution.

Algorithm 2: Statistical Agnostic Multiclass Boosting via OCO

Input: $(\gamma, \epsilon_0, m_0)$ -AWLs $\mathcal{W}_1 \dots \mathcal{W}_T$, $\text{OCO}(\Delta_k, T)$ algorithms $\mathcal{A}^1 \dots \mathcal{A}^m$, and Sample $S = (x_1, y_1) \dots (x_m, y_m) \in \mathcal{X} \times \mathcal{B}_k$

- 1 **for** $t = 1, \dots, T$ **do**
- 2 If $t = 1$, set $P_1[i] = \frac{1}{k} \forall i \in [m]$
- 3 Pass m_0 examples to \mathcal{W}_t drawn from the following distribution:
- 4 Draw x_i w.p. $\frac{1}{m}$ and relabel by drawing $\tilde{y}_i \sim P_t[i]$
- 5 Let h_t be the weak hypothesis returned by \mathcal{W}_t
- 6 Reveal loss function: $l_t^i(p) = (2p - \mathbb{1}_k) \cdot \left(\frac{h_t(x_i)}{\gamma} - y_i \right) \forall i \in [m]$
- 7 Set $P_{t+1}[i] = \mathcal{A}^i(l_1^i, \dots, l_t^i) \forall i \in [m]$
- 8 **return** $\bar{h}(x) = \prod \left(\frac{1}{\gamma T} \sum_{t=1}^T h_t(x) \right)$

Once again, the high-level framework is inspired by the work of Brukhim et al. [6], but like in the online setting, several pieces need to be redesigned when $k > 2$. One key difference between our algorithm and the corresponding algorithm by Brukhim et al. [6] is the use of multiple OCO algorithms. The algorithmic choices in Algorithm 2 are similar to those made in the online setting. Thus, the intuition provided in Section 3.1 also follows here.

Under the assumption that the weak learners satisfy Definition 2, Theorem 5 bounds the *correlation* of Algorithm 2.

Theorem 5. *The output of Algorithm 2, which is denoted \bar{h} , satisfies,*

$$\mathbb{E} [\text{cor}_S(\bar{h})] \geq \max_{h \in \mathcal{H}} \text{cor}_S(h) - \frac{R_{\mathcal{A}}(T)}{T} - \frac{k\epsilon_0}{\gamma},$$

where S is the distribution which uniformly assigns to each example (x_i, y_i) probability $1/m$.

The proof of Theorem 5 is split over the next two subsections.

A.2 Improper game playing

An important step towards proving Theorem 5 is framing statistical agnostic multiclass boosting as an improper zero-sum game. A similar idea was used by Brukhim et al. [6] for binary classification. We take a detour and elaborate on this connection here.

Consider an improper zero-sum game. In this game, there are two players A and B, and a payoff function g that decomposes into the sum of m smaller independent convex-concave payoff functions f^1, \dots, f^m each of which depends on the players' strategies. The goal for Player A is to minimize g , while the goal for Player B is to maximize g . Let \mathcal{K}_A and \mathcal{K}_B be the convex, compact decisions sets of players A and B respectively. In addition, let \mathcal{K}_C be a convex, compact set, and let \mathcal{K}_A be a matrix consisting of m row vectors from \mathcal{K}_C . Because convexity and concavity are preserved under

(non-negative weighted) summation, g is also convex-concave. By Sion's minimax theorem [34], the value of the game is well-defined, which we denote as

$$\lambda^* = \min_{P \in \mathcal{K}_A} \max_{q \in \mathcal{K}_B} g(P, q) = \max_{q \in \mathcal{K}_B} \min_{P \in \mathcal{K}_A} g(P, q).$$

Let \mathcal{K}'_B be a convex compact set such that $\mathcal{K}_B \subseteq \mathcal{K}'_B$. Strategies in \mathcal{K}_B are proper, while those in \mathcal{K}'_B are improper. We allow g to be defined over \mathcal{K}'_B . In addition, we make the following three assumptions:

1. Player B has access to a randomized approximate optimization oracle \mathcal{W} . Given any $P \in \mathcal{K}_A$, \mathcal{W} outputs a response $q \in \mathcal{K}'_B$ such that $\mathbb{E}[g(P, q)] \geq \max_{q^* \in \mathcal{K}_B} g(P, q^*) - \epsilon$, where the expectation is over randomness of \mathcal{W}
2. Player B is allowed to play strategies in \mathcal{K}'_B
3. Player A has access to m copies of possibly (independently) randomized OCO(\mathcal{K}_C, T) algorithms $\mathcal{A}_1, \dots, \mathcal{A}_m$, all with regret $R_{\mathcal{A}}(T)$

Algorithm 3: Improper Game Playing

```

1 for  $t = 1, \dots, T$  do
2   Player A plays  $P_t$ 
3   Player B plays  $q_t \in \mathcal{K}'_B$ , where  $q_t = \mathcal{W}(P_t)$ 
4   Player A and B lose/gain payoff  $g(P_t, q_t) = \sum_{i=1}^m f^i(P_t[i], q_t)$ 
5   Define loss:  $\ell_t^i(p) = f^i(p, q_t) \forall i \in [m]$ 
6   Player A updates  $P_{t+1}[i] = \mathcal{A}^i(\ell_1^i, \dots, \ell_t^i) \forall i \in [m]$ 

```

Proposition 6. *If players A and B play according to Algorithm 3, then player B's average strategy $\bar{q} = \frac{1}{T} \sum_{t=1}^T q_t$, $\bar{q} \in \mathcal{K}'_B$, satisfies for any $P^* \in \mathcal{K}_A$,*

$$\lambda^* \leq \mathbb{E}[g(P^*, \bar{q})] + \frac{mR_{\mathcal{A}}(T)}{T} + \epsilon,$$

where the expectation is over the randomness of \mathcal{W} .

Proof. Since the game is well-defined over \mathcal{K}_A and \mathcal{K}_B , there exists a max-min strategy $q^* \in \mathcal{K}_B$ for player B such that for all $P \in \mathcal{K}_A$, $g(P, q^*) \geq \lambda^*$. Let $\bar{P} = \frac{1}{T} \sum_{t=1}^T P_t$. Then,

$$\begin{aligned} \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T g(P_t, q_t) \right] &\geq \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \max_{q \in \mathcal{K}_B} g(P_t, q) \right] - \epsilon \\ &\geq \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T g(P_t, q^*) \right] - \epsilon \\ &\geq \mathbb{E} [d(\bar{P}, q^*)] - \epsilon \\ &\geq \lambda^* - \epsilon. \end{aligned}$$

Now let $\bar{q} = \frac{1}{T} \sum_{t=1}^T q_t$. Note that $\bar{q} \in \mathcal{K}'_B$. Then we write:

$$\begin{aligned}
\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T g(P_t, q_t) \right] &= \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \sum_{i=1}^m f^i(P_t[i], q_t) \right] \\
&= \sum_{i=1}^m \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T f^i(P_t[i], q_t) \right] \\
&\leq \sum_{i=1}^m \left(\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T f^i(P^*[i], q_t) \right] + \frac{R_{\mathcal{A}}(T)}{T} \right) \\
&= \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \sum_{i=1}^m f^i(P^*[i], q_t) \right] + \frac{mR_{\mathcal{A}}(T)}{T} \\
&= \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T g(P^*, q_t) \right] + \frac{mR_{\mathcal{A}}(T)}{T} \\
&\leq \mathbb{E} [g(P^*, \bar{q})] + \frac{mR_{\mathcal{A}}(T)}{T}.
\end{aligned}$$

Here, P^* is any arbitrary matrix in \mathcal{K}_A . Combining the lower and upper bounds, we get:

$$\lambda^* \leq \mathbb{E} [g(P^*, \bar{q})] + \frac{mR_{\mathcal{A}}(T)}{T} + \epsilon,$$

which completes the proof. □

A.3 Proof of Theorem 5

Now we are ready to prove Theorem 5. The proof strategy will be as follows. We will first show how the proposed agnostic boosting algorithm is an instance of the improper game playing setup described above. Then, we will show that a weak learner satisfying Definition 2 corresponds to the randomized approximate optimization oracle. Finally, we will explicitly compute the value of the game, and derive the lower bound on correlation.

Under Definition 2, we can carefully argue that the agnostic weak learner \mathcal{W} induces an approximate optimization oracle for Player B. Specifically, we can show the following lemma.

Lemma 7. *For any $P \in \mathcal{K}_A$, the output $q' = \frac{\mathcal{W}(P)}{\gamma} \in \mathcal{K}'_B$ satisfies,*

$$\mathbb{E} [g(P, q')] \geq \max_{q \in \mathcal{K}_B} g(P, q) - \frac{mk\epsilon_0}{\gamma}.$$

Proof. Note that in line 3 of Algorithm 2, we pass re-labelled examples back to the weak learner. An alternative and equivalent approach to what is presented is to first relabel each example using P and then to uniformly sample and pass m_0 examples to \mathcal{W} . Let \tilde{y}_i correspond to the relabeled class of example i . Recall that $\tilde{y}_i \sim P[i]$. Let R be the distribution over (x_i, \tilde{y}_i) after relabelling and define $h^* = \max_{h \in \mathcal{H}} \text{cor}_R(h)$. Then, by the weak learning assumption,

$$\begin{aligned}
\mathbb{E}_{S'|\tilde{y}} \left[\sum_{i=1}^m \mathcal{W}_{S'}(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) \right] &= \mathbb{E}_{S'|\tilde{y}} \left[\sum_{\ell \in \mathcal{B}_k} \sum_{i:\tilde{y}_i=\ell} \mathcal{W}_{S'}(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) \right] \\
&\geq \sum_{\ell \in \mathcal{B}_k} \left(\gamma \max_{h \in \mathcal{H}} \sum_{i:\tilde{y}_i=\ell} h(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) - m\epsilon_0 \right) \\
&\geq \gamma \sum_{\ell \in \mathcal{B}_k} \sum_{i:\tilde{y}_i=\ell} h^*(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) - mk\epsilon_0 \\
&= \gamma \sum_{i=1}^m h^*(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) - mk\epsilon_0.
\end{aligned}$$

Taking the expectation of both sides,

$$\begin{aligned}
\mathbb{E}_{\tilde{y}} \left[\mathbb{E}_{S'|\tilde{y}} \left[\sum_{i=1}^m \mathcal{W}_{S'}(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) \right] \right] &= \mathbb{E} \left[\sum_{i=1}^m \mathcal{W}_{S'}(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) \right] \\
&= \mathbb{E}_{S'} \left[\mathbb{E}_{\tilde{y}|S'} \left[\sum_{i=1}^m \mathcal{W}_{S'}(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) \right] \right] \\
&\geq \mathbb{E}_{\tilde{y}} \left[\gamma \sum_{i=1}^m h^*(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) \right] - mk\epsilon_0.
\end{aligned}$$

For any $h \in \mathcal{H}$,

$$\begin{aligned}
\mathbb{E}_{\tilde{y}} \left[\frac{1}{m} \sum_{i=1}^m h(x_i) \cdot (2\tilde{y}_i - \mathbb{1}_k) \right] &= \mathbb{E}_{(x_i, y_i) \sim R} [h(x_i) \cdot (2y_i - \mathbb{1}_k)] \\
&= \mathbb{E}_{x_i} \left[\mathbb{E}_{y_i|x_i} [h(x_i) \cdot (2y_i - \mathbb{1}_k)] \right] \\
&= \mathbb{E}_{x_i} [h(x_i) \cdot (2\mathbb{E}[y_i|x_i] - \mathbb{1}_k)] \\
&= \mathbb{E}_{x_i} [h(x_i) \cdot (2P[i] - \mathbb{1}_k)] \\
&= \frac{1}{m} \sum_{i=1}^m h(x_i) \cdot (2P[i] - \mathbb{1}_k) \\
&= \mathbb{E}_{(x_i, y_i) \sim D} [h(x_i) \cdot (2P[i] - \mathbb{1}_k)].
\end{aligned}$$

Note the last equality follows from the fact that both R and D have the same marginal distribution over unlabeled examples x_i 's. Combining the results, multiplying by m , and dividing by γ gives,

$$\mathbb{E}_{S'} \left[\sum_{i=1}^m \frac{\mathcal{W}_{S'}(x_i)}{\gamma} \cdot (2P[i] - \mathbb{1}_k) \right] \geq \sum_{i=1}^m h^*(x_i) \cdot (2P[i] - \mathbb{1}_k) - \frac{mk\epsilon_0}{\gamma}.$$

Finally, note that $q'(x_i) = \frac{\mathcal{W}_{S'}(x_i)}{\gamma} \in \mathcal{K}'_B$. Then, recalling our definition of g ,

$$\begin{aligned}
\mathbb{E}[g(P, q')] &= \mathbb{E} \left[\sum_{i=1}^m (2P[i] - \mathbb{1}_k) \cdot (q'(x_i) - y_i) \right] \\
&= \mathbb{E} \left[\sum_{i=1}^m (2P[i] - \mathbb{1}_k) \cdot q'(x_i) \right] - \sum_{i=1}^m (2P[i] - \mathbb{1}_k) \cdot y_i \\
&\geq \sum_{i=1}^m h^*(x_i) \cdot (2P[i] - \mathbb{1}_k) - \frac{mk\epsilon_0}{\gamma} - \sum_{i=1}^m (2P[i] - \mathbb{1}_k) \cdot y_i \\
&= \sum_{i=1}^m (h^*(x_i) - y_i) \cdot (2P[i] - \mathbb{1}_k) - \frac{mk\epsilon_0}{\gamma} \\
&= \max_{q \in \mathcal{K}_B} g(P, q) - \frac{mk\epsilon_0}{\gamma}.
\end{aligned}$$

which completes the proof. □

Now, we return to proving Theorem 5. We start by explicitly computing the value of the above game. One can show that the dominant strategy for player A is to return a matrix where row i is the vector y_i . That is, $P = [y_1; y_2; \dots; y_m]$. Because g decomposes into the sum of smaller independent losses f , it is helpful to instead focus our attention on

$$f^i(P[i], q) = (2P[i] - \mathbb{1}_k) \cdot (q(x_i) - y_i).$$

Above, Player A has control over the vector $P[i]$ and wants to minimize f^i . We can show that regardless of what $q(x_i) \in \mathcal{K}_B$ is, Player A minimizes f by playing $P[i] = y_i$. Note, only the case where $q(x_i) \neq y_i$ is important. Under this scenario, Player A needs to maximize the value at the y_i th index and minimize the value at the $q(x_i)$ th index. This is precisely accomplished by setting $P[i] = y_i$. Since the smaller loss functions f^i are independent of one another, it follows that Player A minimizes g by playing matrix $P = [y_1; y_2; \dots; y_m]$. Under this fixed dominant strategy for A, the value of the game λ^* can be computed as

$$\begin{aligned}
\lambda^* &= \max_{q \in \mathcal{K}_B} \min_{P \in \mathcal{K}_A} g(P, q) \\
&= \max_{q \in \mathcal{K}_B} \sum_{i=1}^m (2y_i - \mathbb{1}_k) \cdot (q(x_i) - y_i) \\
&= \max_{q \in \mathcal{K}_B} \sum_{i=1}^m 2(q(x_i) \cdot y_i - 1) \\
&= m \cdot \text{cor}_S(h^*) - m.
\end{aligned}$$

Then, for any P^* using Proposition 6,

$$\begin{aligned}
m \cdot \max_{h \in \mathcal{H}} \text{cor}_S(h) - m &\leq \mathbb{E} [g(P^*, \bar{q})] + \frac{mR_{\mathcal{A}}(T)}{T} + \frac{mk\epsilon_0}{\gamma} \\
&= \mathbb{E} \left[\sum_{i=1}^m (2P^*[i] - \mathbb{1}_k) \cdot (\bar{q}(x_i) - y_i) \right] + \frac{mR_{\mathcal{A}}(T)}{T} + \frac{mk\epsilon_0}{\gamma} \\
&= \mathbb{E} \left[\sum_{i=1}^m (2P^*[i] - \mathbb{1}_k) \cdot \left(\frac{1}{\gamma T} \sum_{t=1}^T h_t(x_i) - y_i \right) \right] + \frac{mR_{\mathcal{A}}(T)}{T} + \frac{mk\epsilon_0}{\gamma} \\
&= \mathbb{E} \left[\sum_{i=1}^m P^*[i] \cdot \left(\frac{\frac{2}{T} \sum_{t=1}^T h_t(x_i) - \mathbb{1}_k}{\gamma} - (2y_i - \mathbb{1}_k) \right) \right] \\
&= \frac{mR_{\mathcal{A}}(T)}{T} + \frac{mk\epsilon_0}{\gamma}. \quad (\text{Lemma 19})
\end{aligned}$$

According to Lemma 17, there exists a $P^*[i]$ and therefore a P^* , such that

$$\begin{aligned}
m \cdot \max_{h \in \mathcal{H}} \text{cor}_S(h) - m &\leq \mathbb{E} \left[\sum_{i=1}^m 2 \left(\prod \left(\frac{\sum_{i=1}^T h_t(x_i)}{\gamma T} \right) \cdot y_i - 1 \right) \right] + \frac{mR_{\mathcal{A}}(T)}{T} + \frac{mk\epsilon_0}{\gamma} \\
&= \mathbb{E} \left[\sum_{i=1}^m 2\bar{h}(x_t) \cdot y_i - 1 \right] - m + \frac{mR_{\mathcal{A}}(T)}{T} + \frac{mk\epsilon_0}{\gamma}.
\end{aligned}$$

Subtracting and then dividing both sides by m ,

$$\begin{aligned}
\max_{h \in \mathcal{H}} \text{cor}_S(h) &\leq \mathbb{E} \left[\frac{1}{m} \sum_{i=1}^m 2\bar{h}(x_t) \cdot y_i - 1 \right] + \frac{R_{\mathcal{A}}(T)}{T} + \frac{k\epsilon_0}{\gamma} \\
&= \mathbb{E} [\text{cor}_S(\bar{h})] + \frac{R_{\mathcal{A}}(T)}{T} + \frac{k\epsilon_0}{\gamma}.
\end{aligned}$$

Rearranging, we have shown that

$$\mathbb{E} [\text{cor}_S(\bar{h})] \geq \max_{h \in \mathcal{H}} \text{cor}_S(h) - \frac{R_{\mathcal{A}}(T)}{T} - \frac{k\epsilon_0}{\gamma}$$

which completes the proof. If one further lets $R_{\mathcal{A}}(T)$ to be OGD, as in the online setting, then observe that we get

$$\mathbb{E} [\text{cor}_S(\bar{h})] \geq \max_{h \in \mathcal{H}} \text{cor}_S(h) - \frac{R_{\mathcal{A}}(T)}{T} - \frac{k\epsilon_0}{\gamma}$$

B Realizable Multiclass Boosting

In this section, we fill in the gaps of Section 4 for the online and statistical realizable settings by providing boosting algorithms and their theoretical analysis. Namely, we show explicitly how the OCO framework can also be used to construct multiclass boosting algorithms in the realizable setting. Unlike in the agnostic setting, our realizable boosting algorithms will update weak learners by *reweighting* examples.

B.1 Online Setting

We now describe an online realizable boosting algorithm whose pseudocode is provided below. The booster has black-box oracle access to N copies of a (γ, T) -RWOL algorithm $\mathcal{W}_1, \dots, \mathcal{W}_N$, each satisfying Definition 3, and an OCO($[0, 1], N$) algorithm \mathcal{A} . Importantly, note that in line 4, we denote by Π the L_2 projection onto Δ_k .

Algorithm 4: Online Realizable Multiclass Boosting via OCO

Input: (γ, T) -RWOL $\mathcal{W}_1 \dots \mathcal{W}_N$, OCO($[0, 1], N$) algorithm \mathcal{A}

```

1 for  $t = 1, \dots, T$  do
2   Receive example  $x_t$ 
3   Accumulate weak predictions  $h_t = \sum_{i=1}^N \mathcal{W}_i(x_t)$ 
4   Set  $\mathcal{D}_t = \Pi(\frac{h_t}{\gamma N})$ 
5   Predict  $\hat{y}_t \sim \mathcal{D}_t$ 
6   Receive true label  $y_t$ 
7   for  $i = 1, \dots, N$  do
8     If  $i > 1$ , obtain  $p_t^i = \mathcal{A}(l_t^1, \dots, l_t^{i-1})$ . Else, initialize  $p_t^1 = 0.5$ .
9     Reveal loss function:  $l_t^i(p) = p(\frac{2\mathcal{W}_i(x_t) \cdot y_t - 1}{\gamma} - 1)$ 
10    Pass  $(x_t, y_t)$  to  $\mathcal{W}_i$  w.p.  $p_t^i$ 
11  Reset  $\mathcal{A}$ 

```

Theorem 8. *The regret bound of Algorithm 4 satisfies:*

$$\frac{1}{T} \sum_{t=1}^T (2\mathbb{E}[\hat{y}_t] \cdot y_t - 1) \geq 1 - \frac{R_{\mathcal{A}}(N)}{N} - \frac{\tilde{R}_{\mathcal{W}}(T, k)}{\gamma T},$$

where $\tilde{R}_{\mathcal{W}}(T, k) = 2R_{\mathcal{W}}(T, k) + \tilde{O}(\sqrt{T})$.

Proof. As usual, the approach is to compute a lower and upper bound on the sum of the expected losses passed to the OCO oracle. Starting with the upper bound,

$$\begin{aligned}
\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N l_t^i(p_t^i) \right] &\leq \frac{1}{N} \mathbb{E} \left[\min_{p \in \mathcal{K}} \sum_{i=1}^N l_t^i(p) \right] + \frac{R_{\mathcal{A}}(N)}{N} && \text{(OCO Regret)} \\
&\leq \frac{1}{N} \mathbb{E} \left[\sum_{i=1}^N l_t^i(p^*) \right] + \frac{R_{\mathcal{A}}(N)}{N} \\
&= \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[p^* \left(\frac{2\mathcal{W}_i(x_t) \cdot y_t - 1}{\gamma} - 1 \right) \right] + \frac{R_{\mathcal{A}}(N)}{N} \\
&= \mathbb{E} \left[p^* \left(\frac{2 \left(\frac{1}{N} \sum_{i=1}^N \mathcal{W}_i(x_t) \right) \cdot y_t - 1}{\gamma} - 1 \right) \right] + \frac{R_{\mathcal{A}}(N)}{N} \\
&\leq \mathbb{E} \left[2 \left(\prod \left(\frac{1}{\gamma N} \sum_{i=1}^N \mathcal{W}_i(x_t) \right) \cdot y_t - 1 \right) \right] + \frac{R_{\mathcal{A}}(N)}{N} && \text{(Lemma 18)} \\
&= 2(\mathbb{E}[\hat{y}_t] \cdot y_t - 1) + \frac{R_{\mathcal{A}}(N)}{N}.
\end{aligned}$$

Finally, summing over $t \in [T]$,

$$\mathbb{E} \left[\frac{1}{N} \sum_{t=1}^T \sum_{i=1}^N l_t^i(p_t^i) \right] \leq \sum_{t=1}^T 2 (\mathbb{E} [\hat{y}_t] \cdot y_t - 1) + \frac{TR_{\mathcal{A}}(N)}{N}.$$

For the lower bound, we first need the following important Lemma adapted from [4] to the multiclass setting:

Lemma 9. *For any weak learner (γ, T) -RWOL \mathcal{W} , there exists $\tilde{R}_{\mathcal{W}}(T, k) = \tilde{\mathcal{O}}(\sqrt{\sum_t p_t}) + 2R_{\mathcal{W}}(T, k)$ such that for any sequence $p_1, \dots, p_T \in [0, 1]$,*

$$\sum_{t=1}^T p_t (2\mathcal{W}(x_t) \cdot y_t - 1) \geq \gamma \sum_{t=1}^T p_t - \tilde{R}_{\mathcal{W}}(T, k).$$

The proof of Lemma 9 follows exactly from Lemma 14 in Brukhim et al. [6] and Lemma 1 in Beygelzimer et al. [4] and so we omit it here. Using Lemma 9,

$$\begin{aligned} \frac{1}{\gamma} \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T p_t^i (2\mathcal{W}(x_t) \cdot y_t - 1) \right] &\geq \mathbb{E} \left[\frac{1}{\gamma} \sum_{i=1}^N \left(\gamma \sum_{t=1}^T p_t^i - \tilde{R}_{\mathcal{W}}(T, k) \right) \right] \\ &= \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i] - \frac{N}{\gamma} \tilde{R}_{\mathcal{W}}(T, k). \end{aligned}$$

Then,

$$\mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^N l_t^i(p_t^i) \right] = \frac{1}{\gamma} \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T p_t^i (2\mathcal{W}(x_t) \cdot y_t - 1) \right] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i] \geq -\frac{N}{\gamma} \tilde{R}_{\mathcal{W}}(T, k).$$

By combining upper and lower bounds for $\mathbb{E} \left[\frac{1}{NT} \sum_t \sum_i l_t^i(p_t^i) \right]$, we get

$$\frac{1}{T} \sum_{t=1}^T (2\mathbb{E} [\hat{y}_t] \cdot y_t - 1) \geq 1 - \frac{R_{\mathcal{A}}(N)}{N} - \frac{\tilde{R}_{\mathcal{W}}(T, k)}{\gamma T}.$$

which completes the proof. \square

B.2 Statistical Setting

We now describe a statistical realizable boosting algorithm whose pseudocode is provided below. The booster is given as input a sample $S = (x_1, y_1), \dots, (x_m, y_m) \in \mathcal{X} \times \mathcal{B}_k$. The booster has black-box oracle access to T copies of a (γ, m_0) -RWL algorithm $\mathcal{W}_1, \dots, \mathcal{W}_T$, each satisfying Definition 4, and m copies of an OCO($[0, 1], T$) algorithm, $\mathcal{A}^1, \dots, \mathcal{A}^m$. Importantly, note that in line 8, we denote by Π as a randomized *prediction* operator that given a γ -scaled distribution over k labels, first computes the L_2 projection onto Δ_k , and then randomly samples a label from the resulting distribution.

Algorithm 5: Statistical Realizable Boosting via OCO

Input: (γ, m_0) -RWLs $\mathcal{W}_1 \dots \mathcal{W}_T$, OCO $([0, 1], T)$ algorithms $\mathcal{A}^1 \dots \mathcal{A}^m$, and Sample $S = (x_1, y_1) \dots (x_m, y_m) \in \mathcal{X} \times \mathcal{B}_k$

```

1 for  $t = 1, \dots, T$  do
2   If  $t = 1$ , set  $P_1[i] = 1/m \forall i \in [m]$ 
3   Pass  $m_0$  examples to  $\mathcal{W}_t$  drawn from the following distribution:
4   Draw  $(x_i, y_i)$  w.p.  $\propto P_t[i]$ 
5   Let  $h_t$  be the weak hypothesis returned by  $\mathcal{W}_t$ 
6   Reveal loss function:  $l_t^i(p) = p(\frac{h_t(x_i)}{\gamma}) \cdot (2y_i - \mathbb{1}_k) - 1 \forall i \in [m]$ 
7   Set  $P_{t+1}[i] = \mathcal{A}^i(l_1^i, \dots, l_t^i) \forall i \in [m]$ 
8 return  $\bar{h}(x) = \prod \left( \frac{1}{\gamma^T} \sum_{t=1}^T h_t(x) \right)$ 

```

Theorem 10. The output of Algorithm 5, which is denoted \bar{h} , satisfies,

$$\mathbb{E} [\text{cor}_S(\bar{h})] \geq 1 - \frac{R_{\mathcal{A}}(T)}{T},$$

where S is the distribution which uniformly assigns to each example (x_i, y_i) probability $1/m$.

Proof. We will use a very similar proof strategy as in the agnostic setting by reducing to improper game playing. Let h^* be a hypothesis consistent with the input sample (i.e. $h^*(x_i) = y_i$ for $i \leq m$) and let $\mathcal{H}' = \mathcal{H} \cup h^*$. We begin by establishing the reduction to improper game playing. Take $\mathcal{K}_A = [0, 1]^m$, $\mathcal{K}_B = \Delta_{\mathcal{H}}$, $\mathcal{K}'_B = \frac{1}{\gamma} \mathcal{K}_B$. Define payoff functions:

$$f^i(p, q) = p(q(x_i)) \cdot (2y_i - \mathbb{1}_k) - 1$$

$$\begin{aligned} g(P, q) &= \sum_{i=1}^m f^i(P[i], q) \\ &= \sum_{i=1}^m P[i](q(x_i) \cdot (2y_i - \mathbb{1}_k) - 1). \end{aligned}$$

We will now show that the agnostic weak learner \mathcal{W} induces an approximate optimization oracle for Player B. Specifically, we will show the following lemma.

Lemma 11. For any $P \in \mathcal{K}_A$, the output $q' = \frac{\mathcal{W}(P)}{\gamma} \in \mathcal{K}'_B$ satisfies

$$\mathbb{E} [g(P, q')] \geq 0.$$

That is, the weak learner corresponds to an approximate oracle with 0 error.

Proof. Using the definition of correlation and the weak learning condition.

$$\mathbb{E} [\text{cor}_P(\mathcal{W}_{S'})] = \mathbb{E} \left[\sum_{i=1}^m \frac{P[i]}{\sum_{i=1}^m P[i]} \mathcal{W}_{S'}(x_i) \cdot (2y_i - \mathbb{1}_k) \right] \geq \gamma$$

and therefore,

$$\mathbb{E} \left[\sum_{i=1}^m P[i] \frac{\mathcal{W}_{S'}(x_i)}{\gamma} \cdot (2y_i - \mathbb{1}_k) \right] \geq \sum_{i=1}^m P[i].$$

Now, take $q'(x_i) = \frac{\mathcal{W}_{S'}(x_i)}{\gamma} \in \mathcal{K}'_B$. Then, recalling our definition of g ,

$$\begin{aligned}
\mathbb{E}[g(P, q')] &= \mathbb{E} \left[\sum_{i=1}^m P[i] (q'(x_i) \cdot (2y_i - \mathbb{1}_k) - 1) \right] \\
&= \mathbb{E} \left[\sum_{i=1}^m P[i] q'(x_i) \cdot (2y_i - \mathbb{1}_k) \right] - \sum_{i=1}^m P[i] \\
&\geq 0 = \max_{q \in \mathcal{K}_B} g(P, q).
\end{aligned}$$

□

Now, we will complete the last part of the overall proof. We start by explicitly computing the value of the above game. One can show that the dominant strategy for player B is to return h^* . Indeed, since h^* is consistent, we can show that $g(P, h^*) = 0$ for any P and thus $\lambda^* = 0$. Then, for any P^* using Proposition 6,

$$\begin{aligned}
0 &\leq \mathbb{E}[g(P^*, \bar{q})] + \frac{mR_{\mathcal{A}}(T)}{T} \\
&= \mathbb{E} \left[\sum_{i=1}^m P^*[i] (\bar{q} \cdot (2y_i - \mathbb{1}_k) - 1) \right] + \frac{mR_{\mathcal{A}}(T)}{T} \\
&= \mathbb{E} \left[\sum_{i=1}^m P^*[i] \left(\left(\frac{1}{\gamma T} \sum_{t=1}^T h_t(x_i) \right) \cdot (2y_i - \mathbb{1}_k) - 1 \right) \right] + \frac{mR_{\mathcal{A}}(T)}{T} \\
&= \mathbb{E} \left[\sum_{i=1}^m P^*[i] \left(\frac{\left(\frac{2}{T} \sum_{t=1}^T h_t(x_i) \right) \cdot y_i - 1}{\gamma} \right) \right] + \frac{mR_{\mathcal{A}}(T)}{T}.
\end{aligned}$$

According to Lemma 18, there exists a $P^*[i]$ and therefore a P^* , such that

$$\begin{aligned}
0 &\leq \mathbb{E} \left[\sum_{i=1}^m 2 \left(\prod \left(\frac{\sum_{i=1}^T h_t(x_i)}{\gamma T} \right) \cdot y_i - 1 \right) \right] + \frac{mR_{\mathcal{A}}(T)}{T} \\
&= \mathbb{E} \left[\sum_{i=1}^m 2\bar{h}(x_t) \cdot y_i - 1 \right] - m + \frac{mR_{\mathcal{A}}(T)}{T}.
\end{aligned}$$

Dividing both sides by m ,

$$\begin{aligned}
0 &\leq \mathbb{E} \left[\frac{1}{m} \sum_{i=1}^m 2\bar{h}(x_t) \cdot y_i - 1 \right] - 1 + \frac{R_{\mathcal{A}}(T)}{T} \\
&= \mathbb{E} [\text{cor}_S(\bar{h})] - 1 + \frac{R_{\mathcal{A}}(T)}{T}.
\end{aligned}$$

Thus, rearranging, we have shown that

$$\mathbb{E} [\text{cor}_S(\bar{h})] \geq 1 - \frac{R_{\mathcal{A}}(T)}{T}.$$

□

C Alternate Weak Learning Conditions

Definition 1 roughly requires that the online weak learner be able to distinguish between every pair of classes to some non-trivial, but far from optimal, degree. In this section, we propose other possible online agnostic weak learning conditions *without* changing Algorithm 1.

We start by first considering a setting where there always exists a hypothesis $h \in \mathcal{H}$ that, in expectation, makes at most $T/2$ mistakes (i.e. gets at least $1/2$ correct). Here, the weak learning condition below suffices as $\max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t=1}^T 2h(x_t) \cdot y_t - 1 \right]$ is guaranteed to be positive.

Definition 5. Let $\mathcal{H} \subseteq \mathcal{B}_k^{\mathcal{X}}$ be a class of experts and let $0 < \gamma \leq 1$ denote the “advantage”. An online learning algorithm \mathcal{W} is a (γ, T) -agnostic weak online learner for \mathcal{H} if for any adaptively chosen sequence of tuples $(x_t, y_t) \in \mathcal{X} \times \mathcal{B}_k$, at every iteration $t \in [T]$, the algorithm outputs $\mathcal{W}(x_t) \in \mathcal{B}_k$ such that,

$$\gamma \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t=1}^T 2h(x_t) \cdot y_t - 1 \right] - \mathbb{E} \left[\sum_{t=1}^T 2\mathcal{W}(x_t) \cdot y_t - 1 \right] \leq R_{\mathcal{W}}(T, k),$$

where the expectation is taken w.r.t. the randomness of the weak learner \mathcal{W} and that of the possibly adaptive adversary, $R_{\mathcal{W}} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}_+$ is the additive regret: a non-decreasing, sub-linear function of T .

For weak learners satisfying Definition 5, one can show, using the same proof strategy in Section 3.2, that Algorithm 1 achieves the average regret bound below.

Proposition 12. Assuming weak learners satisfy Definition 5, the expected regret bound of Algorithm 1 is

$$\frac{1}{T} \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t=1}^T 2h(x_t) \cdot y_t - 1 \right] - \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T 2\hat{y}_t \cdot y_t - 1 \right] \leq \frac{R_{\mathcal{W}}(T, k)}{\gamma T} + \frac{R_{\mathcal{A}}(N)}{N}.$$

For completeness sake, we include a partial proof below.

Proof. The proof follows the same strategy as that used to prove Theorem 1. That is, we will upper and lower bound the expected sum of loss passed to the OCO algorithm. Fortunately, the proof of the upper bound remains exactly the same as that in the proof of Theorem 1. Thus, we will only derive a lower bound on the sum of losses passed to the OCO algorithm. Define $h^* = \arg \max_{h \in \mathcal{H}} \sum_{t=1}^T (2h(x_t) \cdot y_t - 1)$ as the optimal competitor in hindsight. Then,

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T l_t^i(p_t^i) \right] &= \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T p_t^i \cdot \left(\frac{2\mathcal{W}_i(x_t) - \mathbb{1}_k}{\gamma} - (2y_t - \mathbb{1}_k) \right) \right] \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2\mathcal{W}_i(x_t) - \mathbb{1}_k)] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2\mathcal{W}_i(x_t) \cdot y_t^i - 1] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \quad (\text{Lemma 15}) \end{aligned}$$

Using the weak learning condition in Definition 5,

$$\begin{aligned}
\frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2\mathcal{W}_i(x_t) \cdot y_t^i - 1] &\geq \sum_{i=1}^N \max_{h \in \mathcal{H}} \sum_{t=1}^T \mathbb{E} [2h(x_t) \cdot y_t^i - 1] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \quad (\text{Definition 5}) \\
&\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2h^*(x_t) \cdot y_t^i - 1] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \\
&= \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2h^*(x_t) - \mathbb{1}_k)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma}. \quad (\text{Lemma 15})
\end{aligned}$$

Putting things together, we find,

$$\begin{aligned}
\mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T l_t^i(p_t^i) \right] &\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2h^*(x_t) - 2y_t)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \\
&\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2(h^*(x_t) \cdot y_t - 1)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \quad (\text{Lemma 16}) \\
&= N \sum_{t=1}^T 2(h^*(x_t) \cdot y_t - 1) - \frac{NR_{\mathcal{W}}(T, k)}{\gamma}.
\end{aligned}$$

Note that this is the same lower bound as in the proof of Theorem 1 and therefore the same regret bound as in Theorem 1 holds. \square

The assumption on the hypothesis class that there exists a hypothesis $h \in \mathcal{H}$ that, in expectation, makes at most $T/2$ mistakes is quite strong, especially if we allow a random adaptive adversary. Note that Definition 5 cannot be used when this assumption on \mathcal{H} does not hold - $\mathbb{E} \left[\sum_{t=1}^T 2h(x_t) \cdot y_t - 1 \right]$ can be potentially negative for every $h \in \mathcal{H}$, even for the randomly guessing competitor. One way to relax this assumption on \mathcal{H} is via Definition 1 presented in the main text. Another way, is by taking a *one-vs-all* perspective to multiclass classification, which we present below.

Definition 6. Let $\mathcal{H} \subseteq \mathcal{B}_k^{\mathcal{X}}$ be a class of experts and let $0 < \gamma \leq 1$ denote the “advantage”. An online learning algorithm \mathcal{W} is a (γ, T) -agnostic weak online learner for \mathcal{H} if for any adaptively chosen sequence of tuples $(x_t, y_t) \in \mathcal{X} \times \mathcal{B}_k$, at every iteration $t \in [T]$, for every label $\ell \in \mathcal{B}_k$, the algorithm outputs $\mathcal{W}(x_t) \in \mathcal{B}_k$ such that,

$$\gamma \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t: y_t = \ell} 2h(x_t) \cdot y_t - 1 \right] - \mathbb{E} \left[\sum_{t: y_t = \ell} 2\mathcal{W}(x_t) \cdot y_t - 1 \right] \leq R_{\mathcal{W}}(T, k),$$

where the expectation is taken w.r.t. the randomness of the weak learner \mathcal{W} and that of the possibly adaptive adversary, $R_{\mathcal{W}} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}_+$ is the additive regret: a non-decreasing, sub-linear function of T .

Intuitively, Definition 6 roughly requires that a weak learner be able to learn each class to a non-trivial (but far from optimal) degree of accuracy as it processes the stream of examples. In order to ensure that $\gamma \in (0, 1)$, we also need that our hypothesis class \mathcal{H} contains k hypothesis each of which predicts a reference class $\ell \in \mathcal{B}_k$ with probability a $1/2$. This assumption on the hypothesis class is mild: if our class \mathcal{H} does not contain these randomly guessing binary hypotheses, we can add them without substantially increasing the complexity of the class \mathcal{H} . That said, this is still a minor drawback compared to Definition 1, where no additional assumptions on \mathcal{H} are needed.

Under Definition 6, we can show that the following regret bound for Algorithm 1. Note that this regret bound is worse compared to the one derived assuming the weak learners satisfy Definition 1. Indeed, compared to Theorem 1, there is an extra factor of k in the term containing $R_{\mathcal{W}}(T, k)$ in the bound below.

Proposition 13. *Assuming weak learners satisfy Definition 6, the expected regret bound of Algorithm 1 is*

$$\frac{1}{T} \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t=1}^T 2h(x_t) \cdot y_t - 1 \right] - \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T 2\hat{y}_t \cdot y_t - 1 \right] \leq \frac{kR_{\mathcal{W}}(T, k)}{\gamma T} + \frac{R_{\mathcal{A}}(N)}{N}.$$

Proof. The proof follows the same strategy as that used to prove Theorem 1. Like above, the proof of the upper bound remains exactly the same as that in the proof of Theorem 1. Thus, we will only derive a lower bound on the sum of losses passed to the OCO algorithm. Let $h^* = \arg \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t=1}^T 2h(x_t) \cdot y_t - 1 \right]$ and $h_{i\ell}^* = \arg \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t: y_t^i = \ell}^T 2h(x_t) \cdot y_t^i - 1 \right]$. Then,

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T l_t^i(p_t^i) \right] &= \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T p_t^i \cdot \left(\frac{2\mathcal{W}_i(x_t) - \mathbb{1}_k}{\gamma} - (2y_t - \mathbb{1}_k) \right) \right] \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot 2(\mathcal{W}_i(x_t) - \mathbb{1}_k)] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2\mathcal{W}_i(x_t) \cdot y_t^i - 1] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \quad (\text{Lemma 15}) \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{\ell \in \mathcal{B}_k} \sum_{t: y_t = \ell} \mathbb{E} [2\mathcal{W}_i(x_t) \cdot y_t^i - 1] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \end{aligned}$$

Using the weak learning condition in Definition 6,

$$\begin{aligned} \frac{1}{\gamma} \sum_{i=1}^N \sum_{\ell \in \mathcal{B}_k} \sum_{t: y_t = \ell} \mathbb{E} [2\mathcal{W}_i(x_t) \cdot y_t^i - 1] &\geq \sum_{i=1}^N \sum_{\ell \in \mathcal{B}_k} \sum_{t: y_t = \ell} \mathbb{E} [2h_{i\ell}^*(x_t) \cdot y_t^i - 1] - \frac{kNR_{\mathcal{W}}(T)}{\gamma} \quad (\text{Definition 5}) \\ &\geq \sum_{i=1}^N \sum_{\ell \in \mathcal{B}_k} \sum_{t: y_t = \ell} \mathbb{E} [2h^*(x_t) \cdot y_t^i - 1] - \frac{kNR_{\mathcal{W}}(T)}{\gamma} \\ &= \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2h^*(x_t) \cdot y_t^i - 1] - \frac{kNR_{\mathcal{W}}(T)}{\gamma} \\ &= \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2h^*(x_t) - \mathbb{1}_k)] - \frac{kNR_{\mathcal{W}}(T)}{\gamma} \quad (\text{Lemma 15}) \end{aligned}$$

Putting things together, we find,

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T l_t^i(p_t^i) \right] &\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2h^*(x_t) - 2y_t)] - \frac{kNR_{\mathcal{W}}(T, k)}{\gamma} \\ &\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2(h^*(x_t) \cdot y_t - 1)] - \frac{kNR_{\mathcal{W}}(T, k)}{\gamma} \quad (\text{Lemma 16}) \\ &= N \sum_{t=1}^T 2(h^*(x_t) \cdot y_t - 1) - \frac{kNR_{\mathcal{W}}(T, k)}{\gamma}. \end{aligned}$$

Combining this lower bound with the upper bound on the expected sum of losses in the proof of Theorem 1 completes this proof. \square

We end this section by presenting one more online agnostic weak learning condition that places an *asymmetric* gain function on the best competitor and the weak learner. One of the benefits of this condition is that as k increases, it is more explicit how exactly the weak learning assumption becomes stronger.

Definition 7. Let $\mathcal{H} \subseteq \mathcal{B}_k^{\mathcal{X}}$ be a class of experts and let $0 < \gamma \leq 1$ denote the “advantage”. An online learning algorithm \mathcal{W} is a (γ, T) -agnostic weak online learner for \mathcal{H} if for any adaptively chosen sequence of tuples $(x_t, y_t) \in \mathcal{X} \times \mathcal{B}_k$, at every iteration $t \in [T]$, the algorithm outputs $\mathcal{W}(x_t) \in \mathcal{B}_k$ such that,

$$\gamma \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t=1}^T \frac{kh(x_t) \cdot y_t - 1}{k-1} \right] - \mathbb{E} \left[\sum_{t=1}^T 2\mathcal{W}(x_t) \cdot y_t - 1 \right] \leq R_{\mathcal{W}}(T, k),$$

where the expectation is taken w.r.t. the randomness of the weak learner \mathcal{W} and that of the possibly adaptive adversary, $R_{\mathcal{W}} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}_+$ is the additive regret: a non-decreasing, sub-linear function of T .

Along the same lines as above, under Definition 7, one can show that Algorithm 1 achieves the following average regret bound.

Proposition 14. Assuming weak learners satisfy Definition 7, the expected regret bound of Algorithm 1 is

$$\frac{1}{T} \max_{h \in \mathcal{H}} \mathbb{E} \left[\sum_{t=1}^T 2h(x_t) \cdot y_t - 1 \right] - \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T 2\hat{y}_t \cdot y_t - 1 \right] \leq \frac{R_{\mathcal{W}}(T, k)}{\gamma T} + \frac{R_A(N)}{N}.$$

Proof. Again, we only need to show a lower bound on the expected sum of losses passed to the OCO algorithm. Define $h^* = \arg \max_{h \in \mathcal{H}} \sum_{t=1}^T (2h(x_t) \cdot y_t - 1)$ as the optimal competitor in hindsight. Then,

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T l_t^i(p_t^i) \right] &= \mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T p_t^i \cdot \left(\frac{2\mathcal{W}_i(x_t) - \mathbb{1}_k}{\gamma} - (2y_t - \mathbb{1}_k) \right) \right] \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2\mathcal{W}_i(x_t) - \mathbb{1}_k)] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \\ &= \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2\mathcal{W}_i(x_t) \cdot y_t^i - 1] - \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2y_t - \mathbb{1}_k)] \quad (\text{Lemma 15}) \end{aligned}$$

Using the weak learning condition in Definition 7,

$$\begin{aligned} \frac{1}{\gamma} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2\mathcal{W}_i(x_t) \cdot y_t^i - 1] &\geq \sum_{i=1}^N \max_{h \in \mathcal{H}} \sum_{t=1}^T \mathbb{E} \left[\frac{kh(x_t) \cdot y_t^i - 1}{k-1} \right] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \quad (\text{Definition 7}) \\ &\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2h^*(x_t) \cdot y_t^i - 1] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \\ &= \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2h^*(x_t) - \mathbb{1}_k)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma}. \quad (\text{Lemma 15}) \end{aligned}$$

Putting things together, we find,

$$\begin{aligned}
\mathbb{E} \left[\sum_{i=1}^N \sum_{t=1}^T l_t^i(p_t^i) \right] &\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [p_t^i \cdot (2h^*(x_t) - 2y_t)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} \\
&\geq \sum_{i=1}^N \sum_{t=1}^T \mathbb{E} [2(h^*(x_t) \cdot y_t - 1)] - \frac{NR_{\mathcal{W}}(T, k)}{\gamma} && \text{(Lemma 16)} \\
&= N \sum_{t=1}^T 2(h^*(x_t) \cdot y_t - 1) - \frac{NR_{\mathcal{W}}(T, k)}{\gamma}.
\end{aligned}$$

Note that this is the same lower bound as in the proof of Theorem 1 and therefore the same regret bound as in Theorem 1 holds. \square

D Dependence of $R_{\mathcal{W}}(T, k)$ on k

We show concretely how the weak learner's regret, $R_{\mathcal{W}}(T, k)$, can depend on the number of classes k . Recall, that the weak learning condition in Definition 1 is written in terms of approximately *maximizing* a sequence of *gain* functions:

$$\sigma_{y,\ell}(z) = \begin{cases} -1, & \text{if } z = \ell \\ 1, & \text{if } z = y \\ 0, & \text{otherwise} \end{cases}$$

with an advantage parameter $0 < \gamma < 1$. By taking an affine transformation of $\sigma_{y,\ell}(z)$, Definition 1 can be equivalently expressed as approximately *minimizing* a sequence of bounded, non-negative *loss* functions

$$L_{y,\ell}(z) = \frac{1 - \sigma_{y,\ell}(z)}{2} = \begin{cases} 1, & \text{if } z = \ell \\ 0, & \text{if } z = y \\ 1/2, & \text{otherwise} \end{cases}$$

with an advantage parameter $\gamma > 1$. By redefining weak learning in terms of bounded, non-negative loss functions, we can tap into the rich literature of Prediction with Expert Advice to construct online agnostic weak learners. Specifically, we will use the celebrated (Randomized) Exponential Weights Algorithm (EWA). A nice fact about the EWA is that for 0-1 losses $\tilde{L}_{y_t}(z_t) = \mathbb{1}\{z_t \neq y_t\}$, and any *finite* set of experts \mathcal{H} , it enjoys the regret bound

$$\mathbb{E} \left[\sum_{t=1}^T \tilde{L}_{y_t}(\hat{y}_t) \right] \leq \frac{\eta}{1 - e^{-\eta}} \mathbb{E} \left[\min_{h \in \mathcal{H}} \sum_{t=1}^T \tilde{L}_{y_t}(h(x_t)) \right] + \frac{\ln(|\mathcal{H}|)}{1 - e^{-\eta}},$$

where \hat{y}_t denotes the prediction of the EWA in the t 'th iteration and $\eta > 0$ is a tuneable learning rate (see the the exponentially-weighted forecaster from Chapter 2 of [8]). We will be using this regret bound extensively in the next two subsections.

D.1 Finite Hypothesis Classes

In this section, we will consider two types of finite hypothesis classes: the set of discretized weight matrices and the set of discretized multiclass decision trees of depth 1.

Beginning with weight matrices, consider the multiclass classification setup with example label pairs $(x, y) \in \mathbb{R}^d \times \mathcal{B}_k$. Let \mathcal{H} denote a finite hypothesis class parameterized by discretized weight matrices $W_h \in \mathbb{R}^{k \times d}$, such that $|\mathcal{H}| = \left(\frac{1}{\delta}\right)^{kd}$, where $\delta \in (0, 1)$ is the level of discretization. For an input example $x \in \mathbb{R}^d$, a hypothesis $h \in \mathcal{H}$ makes its prediction as $\hat{y}_t = \arg \max_k W_h x$. Take \mathcal{H} to

be a set of experts. We will now show that an instance of the EWA, using \mathcal{H} and learning rate η , run over 0-1 losses corresponds to an agnostic weak online learner with advantage parameter $\gamma = \frac{2\eta}{1-e^{-\eta}}$ and regret $R_{\mathcal{W}}(T, k) = O(kd \log(\frac{1}{\delta}))$. Taking the discretization parameter $\delta = \frac{1}{T}$ ensures that the regret remains sublinear in T .

First, recall that for 0-1 losses $\tilde{L}_{y_t}(z_t) = \mathbb{1}\{z_t \neq y_t\}$, the EWA algorithm \mathcal{W} using hypothesis class \mathcal{H} and learning rate $\eta > 0$ guarantees the regret bound

$$\mathbb{E} \left[\sum_{t=1}^T \tilde{L}_{y_t}(\mathcal{W}(x_t)) \right] \leq \frac{\eta}{1-e^{-\eta}} \mathbb{E} \left[\min_{h \in \mathcal{H}} \sum_{t=1}^T \tilde{L}_{y_t}(h(x_t)) \right] + O(\log |\mathcal{H}|).$$

Next, for any values of ℓ_1, \dots, ℓ_T where $\ell_t \neq y_t$, observe that

$$\mathbb{E} \left[\sum_{t=1}^T L_{y_t, \ell_t}(\mathcal{W}(x_t)) \right] \leq \mathbb{E} \left[\sum_{t=1}^T \tilde{L}_{y_t}(\mathcal{W}(x_t)) \right]$$

and

$$\frac{1}{2} \mathbb{E} \left[\min_{h \in \mathcal{H}} \sum_{t=1}^T \tilde{L}_{y_t}(h(x_t)) \right] \leq \mathbb{E} \left[\min_{h \in \mathcal{H}} \sum_{t=1}^T L_{y_t, \ell_t}(h(x_t)) \right].$$

Putting things together, we get

$$\mathbb{E} \left[\sum_{t=1}^T L_{y_t, \ell_t}(\mathcal{W}(x_t)) \right] \leq \frac{2\eta}{1-e^{-\eta}} \mathbb{E} \left[\min_{h \in \mathcal{H}} \sum_{t=1}^T L_{y_t, \ell_t}(h(x_t)) \right] + O(\log |\mathcal{H}|).$$

Finally, observing that $|\mathcal{H}| = (\frac{1}{\delta})^{kd}$ completes the proof. As an example, if one takes $\eta = 1$, then the EWA algorithm corresponds to an agnostic weak online learner with advantage parameter $\gamma = \frac{2}{1-e^{-1}} \approx 3.16$ and regret $R_{\mathcal{W}}(T, k) = O(kd \log T)$.

A similar procedure can be performed when considering the same multiclass classification setup when the hypothesis class \mathcal{H} is the set of depth 1 multiclass decision trees. If one further restricts to a k -wise split on a feature $j \in [d]$, then the class of depth 1 decision trees can be abstractly represented by the function:

$$f(x) = \begin{cases} y_1, & \text{if } -\infty < x_j \leq \tau_1 \\ y_2, & \text{if } \tau_1 < x_j \leq \tau_2 \\ \dots & \\ y_k, & \text{if } \tau_{k-1} < x_j < \infty \end{cases}$$

for input $x \in \mathbb{R}^d$, thresholds $\tau_1, \dots, \tau_{k-1}$, and labels y_1, \dots, y_k . If one discretizes the thresholds using the parameter $\delta \in (0, 1)$, then $|\mathcal{H}| = O(d \cdot k^k \cdot (\frac{1}{\delta})^k)$. Therefore, EWA using this hypothesis class and 0-1 losses corresponds to an agnostic weak online learner with regret $R_{\mathcal{W}}(T, k) = O(k \log(kT) + \log d)$.

D.2 Infinite Hypothesis Classes

Our construction of weak learners for the two learning settings above crucially relied on the fact that the hypothesis class was finite. Below, we discuss how to construct agnostic weak online learners for infinite hypothesis classes. The key insight from the constructions above is that any agnostic weak online learner with advantage $\gamma > 1$ for the standard 0-1 loss can be converted into weak learner that satisfies an equivalent version of Definition 1 with the non-negative loss functions $L_{y, \ell}(z)$ with twice the advantage, namely 2γ . For finite hypothesis classes, we constructed weak learners for the 0-1 loss by running (randomized) EWA over \mathcal{H} with a fixed learning rate η . Similarly, for infinite hypothesis classes, we can construct a weak learner for 0-1 losses by setting the learning rate η to be a fixed

positive constant in the optimal multiclass agnostic online learner proposed in Daniely et al. [12]. Concretely, the multiclass agnostic online learner proposed in Section 5 of Daniely et al. [12] runs the (randomized) EWA over a carefully constructed finite set of experts of size at most $(Tk)^{LDim(\mathcal{H})}$, where $LDim(\mathcal{H})$ is the Multiclass Littlestone Dimension of \mathcal{H} . Following an identical analysis as in Daniely et al. [12], one can now show that if \mathcal{W} is a EWA algorithm, then for 0-1 losses over this carefully constructed set of experts, \mathcal{W} achieves regret

$$\mathbb{E} \left[\sum_{t=1}^T \tilde{L}_{y_t}(\mathcal{W}(x_t)) \right] \leq \frac{\eta}{1 - e^{-\eta}} \mathbb{E} \left[\min_{h \in \mathcal{H}} \sum_{t=1}^T \tilde{L}_{y_t}(h(x_t)) \right] + O(LDim(\mathcal{H}) \log(Tk)),$$

where \mathcal{H} is the hypothesis class of interest and $\eta > 0$. Picking $\eta = 1$, we get that the EWA \mathcal{W} is an agnostic online weak learner for 0-1 losses with advantage parameter ≈ 1.58 . Thus, using the key insight mentioned above, \mathcal{W} is also a agnostic online weak learner for the non-negative loss functions $L_{y,l}(z)$ with advantage parameter ≈ 3.16 . This suffices to show that \mathcal{W} is also an agnostic online weak learner with respect to Definition 1.

E Important Lemmas

We give the important lemmas (and their proofs) used in the main text and Appendix A and B.

Lemma 15. For any t and i , $\mathbb{E} [y_t^i \cdot \mathcal{W}_i(x_t)] = \mathbb{E} [p_t^i \cdot \mathcal{W}_i(x_t)]$.

Proof. First note that $\mathbb{E} [y_t^i | p_t^i, y_t] = p_t^i$. Next note that $\mathcal{W}_i(x_t)$ and y_t^i are conditionally independent given p_t^i and y_t . Then,

$$\begin{aligned} \mathbb{E} [y_t^i \cdot \mathcal{W}_i(x_t)] &= \mathbb{E}_{p_t^i, y_t} \left[\mathbb{E} [y_t^i \cdot \mathcal{W}_i(x_t) | p_t^i, y_t] \right] && \text{(law of total expectation)} \\ &= \mathbb{E}_{p_t^i, y_t} \left[\mathbb{E} [y_t^i | p_t^i, y_t] \cdot \mathbb{E} [\mathcal{W}_i(x_t) | p_t^i, y_t] \right] \\ &= \mathbb{E}_{p_t^i, y_t} \left[p_t^i \cdot \mathbb{E} [\mathcal{W}_i(x_t) | p_t^i, y_t] \right] \\ &= \mathbb{E}_{p_t^i, y_t} \left[\mathbb{E} [p_t^i \cdot \mathcal{W}_i(x_t) | p_t^i, y_t] \right] \\ &= \mathbb{E} [p_t^i \cdot \mathcal{W}_i(x_t)]. \end{aligned}$$

□

Lemma 16. For any pair $h, y \in \mathcal{B}_k$ and $p \in \Delta_k$,

$$p \cdot (2h - 2y) \geq 2(h \cdot y - 1)$$

Proof. Consider the case when $h = y$. Then the left and right side of the above inequality hold with equality at the value of 0 for any vector p in the simplex. Now, consider the case when $h \neq y$. Then, observe that setting $p = y$ achieves the minimum of $p \cdot (2h - 2y)$ at -2 which matches the right-hand side. □

Lemma 17. For every $0 < \gamma \leq 1$, $h \in \Delta_k$, and $y \in \mathcal{B}_k$, there exists $p \in \Delta_k$ such that:

$$p \cdot \left(\frac{2h - \mathbb{1}_k}{\gamma} - (2y - \mathbb{1}_k) \right) \leq 2 \left(\prod \left(\frac{h}{\gamma} \right) \cdot y - 1 \right),$$

where $\prod : \mathbb{R}^k \rightarrow \Delta_k$ is the L_2 projection operator onto the simplex.

Proof. We first review an important property of the L_2 projection onto the simplex. It is known that the solution to the L_2 projection $\prod(\cdot)$ onto the simplex is just a threshold operator (see Theorem 2.2 in [11]): for any input vector g , subtract a constant μ from each element. If the result is negative, replace it by zero. More formally, $\prod(g) = [g - \mu \mathbb{1}_k]_+$. Note that μ must satisfy the piecewise linear equation $\sum_{i=1}^k \max(0, g_i - \mu) = 1$. Let g' denote the sorted version of g in decreasing order. Then, define K as the largest integer within $\{1, 2, \dots, k\}$ such that $g'_K - \frac{\sum_{j=1}^K g'_j - 1}{K} > 0$. One can verify that $\mu = \frac{\sum_{j=1}^K g'_j - 1}{K}$ is the unique solution to the piece-wise linear equation above. We will use this property of L_2 projection extensively in our proof, which we formally begin below.

Define $\tilde{h} = \prod\left(\frac{h}{\gamma}\right)$. Let t be the number of zero entries in the projection \tilde{h} . Furthermore, define $h_y = h \cdot y$ and $\tilde{h}_y = \tilde{h} \cdot y$ as the y 'th index of h and $\prod\left(\frac{h}{\gamma}\right)$ respectively. We will show that when $t \leq k - 2$, setting $p = y$ achieves the desired bound and when $t = k - 1$, there exists some $p \in \Delta_k$ that achieves the desired bound.

We now begin with the case where $t \leq k - 2$. Our goal will be to first show the following lower bound on \tilde{h}_y ,

$$\tilde{h}_y \geq \frac{(k-t)h_y - 1}{\gamma(k-t)} + \frac{1}{k-t}.$$

Consider the subcase where $\tilde{h}_y = 0$. Then, by the properties of the projection operator above it must have been the case that $\frac{h_y}{\gamma} \leq \mu$, where using the definition above $\mu = \frac{\frac{1}{\gamma} \sum_{j=1}^{k-t} h'_j - 1}{k-t}$. Therefore, $h_y \leq \frac{\sum_{j=1}^{k-t} h'_j - \gamma}{k-t} \leq \frac{1-\gamma}{k-t}$. Substituting in, we find

$$\frac{(k-t)h_y - 1}{\gamma(k-t)} + \frac{1}{k-t} \leq 0 = \tilde{h}_y.$$

Now, consider the subcase where $\tilde{h}_y > 0$. Then, $\tilde{h}_y = \frac{h_y}{\gamma} - \mu$, where again

$$\mu = \frac{\frac{1}{\gamma} \sum_{j=1}^{k-t} h'_j - 1}{k-t} \leq \frac{1-\gamma}{\gamma(k-t)}.$$

Substituting in completes proving the lowerbound

$$\tilde{h}_y = \frac{h_y}{\gamma} - \mu \geq \frac{h_y}{\gamma} - \frac{1-\gamma}{\gamma(k-t)} = \frac{(k-t)h_y - 1}{\gamma(k-t)} + \frac{1}{k-t}.$$

Now that we have shown,

$$\tilde{h}_y = \frac{h_y}{\gamma} - \mu \geq \frac{(k-t)h_y - 1}{\gamma(k-t)} + \frac{1}{k-t},$$

we finally show this implies the inequality

$$\left(\frac{2h_y - 1}{\gamma} - 1\right) \leq 2\tilde{h}_y - 2,$$

which follows from plugging $p = y$ into the lemma. We start with

$$\begin{aligned}
2\tilde{h}_y - \frac{2h_y - 1}{\gamma} - 1 &\geq \frac{2(k-t)h_y - 2}{\gamma(k-t)} + \frac{2}{k-t} - \frac{2h_y - 1}{\gamma} - 1 \\
&= -\frac{2}{\gamma(k-t)} + \frac{1}{\gamma} + \frac{2}{k-t} - 1 \\
&= \left(\frac{1}{\gamma} - 1\right) - \frac{2}{k-t} \left(\frac{1}{\gamma} - 1\right) \\
&= \left(\frac{1}{\gamma} - 1\right) \left(1 - \frac{2}{k-t}\right) \geq 0,
\end{aligned}$$

where the last inequality follows from the assumption that $t \leq k - 2$. Thus, for the case where $t \leq k - 2$, we have shown the lemma holds by setting $p = y$.

Now we will consider the case where $t = k - 1$. Again, we will also consider the subcases where $\tilde{h}_y = 0$ and $\tilde{h}_y = 1$. We start with the subcase where $\tilde{h}_y = 0$. Here, we will show that picking $p = y$ is the right choice. Namely, we will show that the following inequality holds

$$\left(\frac{2h_y - 1}{\gamma} - 1\right) \leq 2\tilde{h}_y - 2.$$

Under the assumption that $\tilde{h}_y = 0$, the right hand side collapses to -2 . Thus, we need to show that $\frac{2h_y - 1}{\gamma} \leq -1$. Recall that if $\tilde{h}_y = 0$, then by projection properties, $\frac{h_y}{\gamma} \leq \mu$. Again, by definition,

$$\mu = \frac{\frac{1}{\gamma} \sum_{j=1}^{k-t} h'_j - 1}{k-t} = \max_j \frac{h_j}{\gamma} - 1 \leq \frac{1 - h_y}{\gamma} - 1.$$

Therefore we find that, $\frac{h_y}{\gamma} \leq \frac{1 - h_y}{\gamma} - 1$, from which it is easy to see that $\frac{2h_y - 1}{\gamma} \leq -1$, which completes this subcase.

We now move to the subcase where $\tilde{h}_y = 1$. Here, we will show that there exists a $p \in \mathcal{B}_k \setminus \{y\}$ that satisfies the required bound in the lemma. When $\tilde{h}_y = 1$, the right hand side collapses to 0. Thus, we need to show the existence of $p \in \mathcal{B}_k \setminus \{y\}$ s.t.

$$\frac{2h_p - 1}{\gamma} + 1 \leq 0,$$

where h_p denotes $h \cdot p$, the value of h at the p 'th index. We shall prove the bound above via contradiction. Assume that indeed for all $p \in \mathcal{B}_k \setminus \{y\}$,

$$\frac{2h_p - 1}{\gamma} + 1 > 0.$$

This implies that $h_p > \frac{1-\gamma}{2}$ for all $p \in \mathcal{B}_k \setminus \{y\}$. Observe that when $k \geq 3$, the probability mass over all labels other than y is *strictly* bounded below by $1 - \gamma$ and so it must be that $h_y < \gamma$. Recall that if $\tilde{h}_y = 1$, then by projection properties,

$$\mu = \frac{\frac{1}{\gamma} \sum_{j=1}^{k-t} h'_j - 1}{k-t} = \frac{h_y}{\gamma} - 1 < 0,$$

where the last inequality follows from the fact that h_y must have mass strictly less than γ . However, if $\mu < 0$, then since the solution to the projection operation is $\left[\frac{h}{\gamma} - \mu \mathbb{1}_k\right]_+$ and the entries of h are non-negative, all entries of $\frac{h}{\gamma} - \mu \mathbb{1}_k$ are strictly positive. This contradicts our original assumption that $t = k - 1$ (implying that all but one index are negative) which completes the proof for $k \geq 3$.

Now, when $k = 2$, $\mu = \frac{h_y}{\gamma} - 1$, and by projection properties,

$$\frac{h_p}{\gamma} - \mu = \frac{h_p - h_y}{\gamma} + 1 \leq 0.$$

Noting that $h_y = 1 - h_p$, completes the proof since it implies that for $p \neq y$, $\frac{2h_p - 1}{\gamma} + 1 \leq 0$. \square

Lemma 18. *For every $0 < \gamma \leq 1$, $h \in \Delta_k$, and $y \in \mathcal{B}_k$, there exists $p \in [0, 1]$ such that:*

$$p \left(\frac{2h \cdot y - 1}{\gamma} - 1 \right) \leq 2 \left(\prod \left(\frac{h}{\gamma} \right) \cdot y - 1 \right),$$

where $\prod : \mathbb{R}^k \rightarrow \Delta_k$ is the L_2 projection operator onto the simplex.

Proof. Observe that we can consider the same four cases as in the proof for Lemma 17. Indeed, for three out of the four cases, the optimal choice for Lemma 17 was selecting $p = y$. Under these three cases, we know that by picking $p = y$ the following inequality holds by substituting $p = y$ into Lemma 17:

$$\frac{2h \cdot y - 1}{\gamma} - 1 \leq 2 \left(\prod \left(\frac{h}{\gamma} \right) \cdot y - 1 \right).$$

Thus, for this lemma, for the same three cases of Lemma 17, it suffices to pick $p = 1$. In the last case of Lemma 17, $\tilde{h}_y = 1$, that is, all the mass after the projection falls on the label y . In this case, the right hand side of our inequality collapses to 0. Thus it suffices to pick $p = 0$ for this lemma to hold. \square

Lemma 19. *For $0 < \gamma \leq 1$, $h \in \Delta_k$, $y \in \mathcal{B}_k$, and all $p \in \Delta_k$:*

$$p \cdot \left(\frac{2h - \mathbb{1}_k}{\gamma} - (2y - \mathbb{1}_k) \right) = (2p - \mathbb{1}_k) \cdot \left(\frac{h}{\gamma} - y \right)$$

Proof.

$$\begin{aligned} p \cdot \left(\frac{2h - \mathbb{1}_k}{\gamma} - (2y - \mathbb{1}_k) \right) &= p \cdot \left(\frac{2h - \mathbb{1}_k}{\gamma} \right) - p \cdot (2y - \mathbb{1}_k) \\ &= 2p \cdot \frac{h}{\gamma} - \frac{1}{\gamma} - 2p \cdot y + 1 \\ &= 2p \cdot \left(\frac{h}{\gamma} - y \right) - \frac{1}{\gamma} + 1 \\ &= 2p \cdot \left(\frac{h}{\gamma} - y \right) - \mathbb{1}_k \cdot \left(\frac{h}{\gamma} - y \right) \\ &= (2p - \mathbb{1}_k) \cdot \left(\frac{h}{\gamma} - y \right). \end{aligned}$$

\square

F Experimental Details

Each cell in Table 1 is the average over the "best accuracies" of five independent shuffles of the dataset. For each shuffle, the "best accuracy" was computed as the maximum accuracy over five candidate γ values. For both OCO-based boosting algorithms, γ was tuned across $[0.10, 0.30, 0.50, 0.70, 1]$, and for OnlineMBBM, γ was tuned across $[0.001, 0.01, 0.05, 0.1, 0.3]$. The range of γ values used to tune OnlineMBBM is consistent with those used by Jung et al. [24]. In addition, in our own experiments, we found that for γ values larger than 0.30, OnlineMBBM performed significantly worse. Note that

Table 2: Dataset summaries.

Dataset	#datapoints	#covariates	#classes
Balance	625	4	3
Cars	1728	6	4
LandSat	6435	36	6
Segmentation	2310	19	7
Mice	1080	82	8
Yeast	1483	8	10
Abalone	4177	8	28

Table 3: Standard errors of accuracies.

Dataset	Standard Errors of Accuracies			
	Agn	Opt	Ada	OCOR
Balance	0.68	1.06	1.27	0.84
Cars	0.46	0.02	0.55	0.32
LandSat	0.89	0.36	1.23	0.21
Segmentation	0.70	0.48	1.56	0.54
Mice	0.58	0.60	2.53	0.26
Yeast	0.43	1.03	1.09	1.53
Abalone	0.53	0.42	0.28	0.31

AdaBoost.OLM does not have a tuning parameter γ , which is one of its advantages despite its subpar performance.

For both OCO-based boosting algorithms, we used Online Gradient Descent with learning rate $\eta = \frac{\gamma}{\sqrt{N}}$ as the OCO algorithm. This learning rate is optimal up to constant factors [35]. All computations were carried out on a Nehalem architecture 10-core 2.27 GHz Intel Xeon E7-4860 processors with 25 GB RAM per core. The total amount of computing time was around 500 hours. The River package [29], used to implement the VeryFastDecisionTree, is licensed under BSD 3-Clause.

Table 2 provides more information about each of the datasets used. For the Mice dataset, entries with missing data were replaced with the average value of their respective column.

Table 3 provides the standard error for each accuracy in Table 1.