

ITERATIVE WEAK LEARNABILITY AND MULTI-CLASS ADABOOST

IN-KOO CHO AND JONATHAN LIBGOBER

ABSTRACT. We construct an efficient recursive ensemble algorithm for the multi-class classification problem, inspired by SAMME (Zhu, Zou, Rosset, and Hastie (2009)). We strengthen the weak learnability condition in Zhu, Zou, Rosset, and Hastie (2009) by requiring that the weak learnability condition holds for any subset of labels with at least two elements. This condition is simpler to check than many proposed alternatives (e.g., Mukherjee and Schapire (2013)). As SAMME, our algorithm is reduced to the Adaptive Boosting algorithm (Schapire and Freund (2012)) if the number of labels is two, and can be motivated as a functional version of the steepest descending method to find an optimal solution. In contrast to SAMME, our algorithm’s final hypothesis converges to the correct label with probability 1. For any number of labels, the probability of misclassification vanishes exponentially as the training period increases. The sum of the training error and an additional term, that depends only on the sample size, bounds the generalization error of our algorithm as the Adaptive Boosting algorithm.

KEYWORDS. Adaptive Boosting Algorithm, SAMME, Iterative Weak Learnability, Efficient Ensemble Algorithm, Generalization Error

1. INTRODUCTION

While the Adaptive Boosting (AdaBoost) algorithm (Schapire and Freund (2012)) and its variants have been remarkably successful for binary classification problems, difficulties with generalizing to the multi-label case are well-known.¹ Early approaches to this problem proposed dividing the multi-class problem into multiple two-class problems (e.g., Schapire and Singer (1999)). These approaches either do not inherit all of the desirable properties of AdaBoost or lack its intuitive descriptions. The extended algorithms are expected to be efficient: the algorithm must produce a forecast, which converges to the correct label in probability, and the probability of misclassification vanishes at an exponential rate. However, the extended algorithm is often substantially different from the original formula of Schapire and Freund (2012) and is significantly more elaborate than the original formula. The aesthetic differences with AdaBoost are substantive. First, it is less straightforward to analyze the properties of these algorithms when they are more elaborate. Second, the conditions required to ensure they work are often significantly more stringent.

Zhu, Zou, Rosset, and Hastie (2009) developed a multi-class generalization of AdaBoost that maintains the same structure as the binary label case, called Stagewise Additive Modeling using a Multi-class Exponential loss function (SAMME). The algorithm has a

Date: January 26, 2021.

We are grateful for the hospitality and support from the University of Southern California. We thank Shaowei Ke, Tim Roughgarden, and Rob Schapire for their helpful comments. The first author gratefully acknowledges financial support from the National Science Foundation and the Korea Research Foundation.

¹For a comprehensive survey, see Friedman, Hastie, and Tibshirani (2000).

statistical foundation as a functional version of the recursive steepest descending method to calculate the minimum of an exponential loss function. SAMME is reduced to the original AdaBoost if the number of labels is two. Zhu, Zou, Rosset, and Hastie (2009) demonstrated through numerical simulations that if the set of weak hypothesis satisfies a version of the weak learnability condition, which extends the condition of Schapire and Freund (2012), SAMME converges to the correct labels.

However, SAMME fails to generate a forecast, which converges to the correct label in probability. Mukherjee and Schapire (2013) showed that the weak learnability condition of Zhu, Zou, Rosset, and Hastie (2009) is not adequate for constructing an efficient recursive ensemble algorithm. Mukherjee and Schapire (2013) instead imposed a condition called *Edge-over-Random* (EOR), which involves checking a richer set of conditions. Verifying EOR is not straightforward because we have to weigh the loss differently for every kind of misclassification.

This paper proposes an alternative approach. We show that a slight modification of SAMME, combined with a slightly stronger version of the weak learnability of Zhu, Zou, Rosset, and Hastie (2009), allows us to construct an efficient recursive ensemble algorithm. We strengthen the weak learnability condition by requiring that the same weak learnability continue to hold for any non-empty subset of labels with at least two elements. We call the stronger version of the weak learnability *iterative weak learnability*. One might wonder whether the stronger version of weak learnability requires an exceedingly elaborate class of weak hypotheses. On the contrary, the set of single threshold classifiers satisfies iterative weak learnability, which is the set of the simplest non-trivial classifiers.

We depart from the conventional view of a boosting algorithm as a process to produce a final hypothesis, which is increasingly accurate as the training period increases. Instead, we consider the boosting algorithm as a process to eliminate inaccurate labels recursively, leaving the correct label in each round of elimination. If the set of labels has two elements, the two views are two sides of the same coin. As an example, let us consider the original AdaBoost with two labels. One can view the algorithm as a process to identify a correct label for each observation. Instead, we view the same algorithm as a process to eliminate an incorrect label. If the number of labels is larger than two, the second view is much easier to generalize than the conventional view on the boosting algorithm.

Our algorithm can be motivated as an implementation of Elimination by Aspects (EBA) (Tversky (1972)). If we regard a label as an attribute, our exercise can be viewed as a procedure to identify the best possible attribute for the observation. The decision maker evaluates an attribute according to a linear function of different attributes after a sequence of observations and eliminates an attribute with the negative value of the linear function. Our algorithm selects the attribute which survives the iterative elimination process.²

We modify the SAMME algorithm to construct an ensemble algorithm to eliminate incorrect labels for each observation. By eliminating incorrect labels iteratively, the number of remaining labels becomes two, and then, our algorithm becomes the original AdaBoost algorithm to identify the correct label. We show that the probability of eliminating the correct label vanishes at an exponential rate as the training period increases, which implies

²We are grateful to Shaowei Ke for pointing out the possible link and the reference.

that our algorithm accurately eliminates incorrect labels efficiently, in contrast to the original proposal. One of the benefits of our approach is that we can analyze this algorithm following steps similar to the original binary label AdaBoost algorithm. Our algorithm produces a generalization error bound closely related to that of the original formulation.

Section 2 illustrates the problem. Section 3 constructs the algorithm and states the main result. Section 4 shows that the generalization error vanishes as the number of samples increases, proving that our algorithm is robust against the overfitting as the original AdaBoost is. Section 5 concludes the paper.

2. PRELIMINARIES

2.1. Basics. Given set X , let $|X|$ be the number of elements in X . Let A be the set of labels, and \mathcal{P} be the set of observations, which is endowed with a probability distribution. A classifier is

$$h : \mathcal{P} \rightarrow A$$

which labels each observation. We often call h a weak hypothesis (or simply, a hypothesis) in place of a classifier. Let \mathcal{H} be the set of all feasible classifiers.

Let

$$y : \mathcal{P} \rightarrow A$$

be the correct classifier. We say that h perfectly classifies y if

$$\mathbf{P}(h(p) = y(p)) = 1.$$

We assume that a decision maker is endowed with \mathcal{H} , but y is often significantly more complex than any element in \mathcal{H} . As a result, no $h \in \mathcal{H}$ can perfectly classify y . An important question is whether a decision maker who is endowed with primitive \mathcal{H} can “learn” y .

Definition 2.1. Let Γ be the set of all classifiers, and $\tilde{\Gamma} \subset \Gamma$ denote a subset of classifiers. A statistical procedure or algorithm is an onto function

$$\tau : \mathcal{D} \rightarrow \tilde{\Gamma},$$

where \mathcal{D} is a set of histories, \mathcal{T} is the set of feasible algorithms.

To incorporate learning, we need to allow a decision maker to observe data over time. We consider a dynamic interaction, where the algorithm can update its output in response to new data. Suppose that time is discrete: $t = 1, 2, \dots$ and D_t is the data available at the beginning of time t . $\tau(D_t)$ is the classifier which algorithm τ generates from data D_t . Let $\tau(D_t)(p)$ be the label assigned to observation p by $\tau(D_t)$.

Let $T \geq 1$ be the training period, when the decision maker updates $\tau(D_t)$ in response to D_t . The final hypothesis is the output $\tau(D_T)$ at the end of the training period.

We focus on *recursive ensemble algorithm* for its simplicity.

Definition 2.2. $\tau(D_K)$ is an ensemble of \mathcal{H} if $\exists h_1, \dots, h_K \in \mathcal{H}$ and $\alpha_1, \dots, \alpha_K \geq 0$ such that $\forall p \in \mathcal{P}$

$$\tau(D_K)(p) = \arg \max_a \sum_{k=1}^K \alpha_k \mathbf{1}[a = h_k(p)]$$

τ is a recursive algorithm, if $\tau(D_k)$ is determined as a function of $\tau(D_{k-1})$ and events in time k .

We require that an algorithm produce an accurate classifier as the final hypothesis and achieve the desired accuracy level with a reasonable amount of data.

Definition 2.3. Algorithm τ is consistent if

$$\lim_{K \rightarrow \infty} \mathbf{P}(\tau(D_K) \neq y(p)) = 0.$$

Algorithm τ is efficient if $\exists \rho > 0, K > 0$ such that

$$\mathbf{P}(\exists k \geq K, \tau(D_k)(p) \neq y(p)) \leq e^{-\rho K}.$$

2.2. Adaptive Boosting Algorithm. Let us illustrate the Adaptive Boosting (AdaBoost) algorithm by Schapire and Freund (2012). Instead of the original formula, we follow the framework of SAMME (Zhu, Zou, Rosset, and Hastie (2009)). Let us denote the algorithm as τ_A .

Initialize $d_1(p) = 1/|\mathcal{P}|$ as the uniform distribution over \mathcal{P} . For $k \geq 1$, D_k is the history at the beginning of time k . Iterate the following steps.

- (1) (calculation of forecasting error) Choose $h_k \in \mathcal{H}$.

$$\epsilon_k = \sum_p d_k(p) \mathbb{I}(h_k(p) \neq y(p))$$

If $\epsilon_k = 0$, then stop the iteration and set $\tau(D_k) = h_k$. If not, continue.

- (2) (calculating the weight)³

$$\alpha_k = \frac{(|A| - 1)^2}{|A|} \log \left(\frac{(|A| - 1)(1 - \epsilon_k)}{\epsilon_k} \right) \quad (2.1)$$

- (3) (preparing for the next round)

$$w_{k+1}(p) = \begin{cases} d_k(p) \exp(-(|A| - 1)\alpha_k) & \text{if } h_k(p) = y(p) \\ d_k(p) \exp(\alpha_k) & \text{if } h_k(p) \neq y(p). \end{cases}$$

$$d_{k+1}(p) = \frac{w_{k+1}(p)}{\sum_{p'} w_{k+1}(p')}.$$

- (4) Reset $k + 1$ as k and start the first step of selecting an optimal weak hypothesis.
(5) (final hypothesis)

$$\tau_A(D_t)(p) = \arg \max_{a \in A} \sum_{s=1}^t \alpha_s \mathbb{I}(h_s(p) = a) \quad \forall p \in \mathcal{P}$$

If $|A| = 2$, τ_A is the Adaptive Boosting (AdaBoost) algorithm of Schapire and Freund (2012).

The description of the algorithm is incomplete because if $\alpha_k < 0$, the algorithm is not well defined.

³The following formula is from equation (11) on page 353 of Zhu, Zou, Rosset, and Hastie (2009), which becomes the coefficient of AdaBoost if $|A| = 2$. This formula differs slightly from (c) on page 351.

Definition 2.4. Let A be the set of labels. $\frac{1}{|A|}$ random guess, or simply a random guess, is to predict each label with probability $\frac{1}{|A|}$.

τ_A is meaningful only if $h_k \in \mathcal{H}$ performs better than a random guess.

$$\sum_{p \in \mathcal{P}} (\mathbb{I}[h_k(p) = y(p)] - \mathbb{I}[h_k(p) \neq y(p)]) d(p) \geq \frac{2 - |A|}{|A|} + \rho, \quad (2.2)$$

which implies that $\alpha_k > 0$. Note that $\frac{2 - |A|}{|A|}$ is the score from the random guess.

The weak learnability implies the existence of such h_k . We state the condition according to Zhu, Zou, Rosset, and Hastie (2009).

Definition 2.5. A hypothesis class \mathcal{H} is weakly learnable if, for every probability distribution d over observations $p \in \mathcal{P}$ and labels $y(p)$, we have:

$$\min_{d \in \Delta(\mathcal{P})} \max_{h \in \mathcal{H}} \sum_{p \in \mathcal{P}} (\mathbb{I}[h(p) = y(p)] - \mathbb{I}[h(p) \neq y(p)]) d(p) \geq \frac{2 - |A|}{|A|} + \rho, \quad (2.3)$$

for some $\rho > 0$.

The weak learnability is sufficient for the efficiency of τ_A if $|A| = 2$.

Theorem 2.6. If \mathcal{H} is weakly learnable and $|A| = 2$, then τ_A is an efficient algorithm.

Proof. See Schapire and Freund (2012). \square

Mukherjee and Schapire (2013) pointed out that the notion of weak learnability of Zhu, Zou, Rosset, and Hastie (2009) is too weak to guarantee the convergence of the algorithm to the correct label if $|A| > 2$.

Let us examine an example in Mukherjee and Schapire (2013).

Example 2.7. $\mathcal{P} = \{a, b\}$, $A = \{1, 2, 3\}$, $y(a) = 1, y(b) = 2$. $\mathcal{H} = \{h_1, h_2\}$ where $h_1(a) = h_1(b) = 1$ and $h_2(a) = h_2(b) = 2$.

Since $|A| = 3$, the weak learnability is satisfied if the performance of a weak hypothesis is strictly better than

$$\frac{2 - 3}{3} = -\frac{1}{3}.$$

Let $d = (d(a), d(b))$ be a probability distribution over \mathcal{P} . The performance score of \mathcal{H} is

$$\max [1 - 2d(a), -1 + 2d(a)] \geq 0$$

where the equality holds if $d(a) = 0.5$. Thus, \mathcal{H} is weakly learnable. Any convex combination of h_1 and h_2 in \mathcal{H} assigns the same level to two observations. Since y assigns a different label on a different observation, SAMME fails to produce a consistent forecast, although \mathcal{H} satisfies the weak learnability.

Example 2.7 reveals two issues we must address to extend AdaBoost from binary to multiple labels. First, to construct a boosting algorithm, it is necessary to strengthen the notion of weak learnability. Second, the set of weak hypotheses \mathcal{H} cannot be too restrictive. The issue identified in the previous example is that the requirement of outperforming a random guess by ρ becomes less restrictive when the number of labels gets larger. Hence some distributions can satisfy 2.3 for many labels but can fail for a small number of labels. We propose to consider the minimal departure necessary to rule this out.

2.3. Iterative Weak Learnability. We strengthen the weak learnability by requiring that (2.3) holds for any subset of A with more than a single element. Let $A_i \subset A$ be a non-empty subset of A with $|A_i| \geq 2$. Define

$$\mathcal{H}_i = \{h : \mathcal{P} \rightarrow A_i\}$$

as the collection of all weak hypotheses with labels from A_i . Clearly, $\mathcal{H}_i \subset \mathcal{H}$.

Definition 2.8. \mathcal{H} is iteratively weakly learnable if $\forall A_i \subset A$ with $|A_i| \geq 2$, $\exists \rho_i > 0$ such that for every distribution d over observations $p \in \mathcal{P}$ and labels $y(p)$,

$$\min_{d \in \Delta(\mathcal{P})} \max_{h \in \mathcal{H}_i} \sum_{p \in \mathcal{P}} (\mathbb{I}[h(p) = y(p)] - \mathbb{I}[h(p) \neq y(p)]) d(p) \geq \frac{2 - |A_i|}{|A_i|} + \rho_i. \quad (2.4)$$

Note that $\frac{2 - |A_i|}{|A_i|}$ is a decreasing function of $|A_i|$. As $|A_i|$ decreases, the requirement for weak learnability of \mathcal{H}_i becomes more stringent, as the minimum performance criterion (the right hand side of (2.4) increases).

In Example 2.7, \mathcal{H} is weakly learnable, but not iteratively weakly learnable. To see this, note that if $A_i = \{1, 2\} \subset A = \{1, 2, 3\}$, then $\mathcal{H} = \mathcal{H}_i$.

$$\min_{0 \leq d(a) \leq 1} \max [1 - 2d(a), -1 + 2d(a)] = 0.$$

However, the weak learnability over A_i requires that the minimum performance be *strictly larger* than 0.

2.4. Sufficiently Rich \mathcal{H} . A fundamental question is whether we can find a class of simple hypothesis which is weakly learnable. With the set of the simplest classifiers, we can do better than the random guess over any subset of A with at least two elements.

Suppose that A and \mathcal{P} are subsets of a vector space or a finite dimensional Euclidean space. Consider h , which classifies \mathcal{P} according to the value of a linear function, thus called a linear classifier. Let \mathbf{H} be a hyperplane in \mathbb{R}^n : $\exists \lambda \in \mathbb{R}^n$ and $\omega \in \mathbb{R}$ such that

$$\mathbf{H} = \{p \in \mathbb{R}^n \mid \lambda p = \omega\}.$$

Define \mathbf{H}_+ as the close half space above \mathbf{H} :

$$\mathbf{H}_+ = \{p \in \mathbb{R}^n \mid \lambda p \geq \omega\}.$$

Definition 2.9. A single threshold (linear) classifier is a mapping

$$h : \mathcal{P} \rightarrow A$$

where $\exists a_+, a_- \in A$ such that

$$h(p) = \begin{cases} a^+ & \text{if } p \in \mathbf{H}_+ \\ a^- & \text{if } p \notin \mathbf{H}_+. \end{cases}$$

If h is a single threshold classifier, $|h(\mathcal{P})| = 2$, which may be strictly smaller than $|A|$. On the other hand, the set of single threshold classifiers is *closed under permutations*: if $\pi : A \rightarrow A$ is a bijection, then $\pi(h(p)) \in \mathcal{H}$ for all π . Being closed under permutations turns out to be the key property to obtain weak learnability.

Lemma 2.10. *Consider an environment where the image of $y(p)$ has $|A'|$ elements where $A' \subset A$. Any hypothesis class that is closed under permutations can do as well as $1/|A'|$ random guessing.*

Proof. See Lemma A.1 in Cho and Libgober (2020). □

We use this result to show that the set of single-threshold classifiers is iteratively weakly learnable.

Corollary 2.11. *The set of single threshold classifiers is iteratively weakly learnable.*

Proof. $\forall A_i \subset A$, the set of single threshold classifiers

$$\mathcal{H}_i = \{h : \mathcal{P} \rightarrow A_i\}$$

is closed under permutations, from which the conclusion follows. □

In Example 2.7, using single-threshold classifiers ensures that an extreme point of the convex hull of \mathcal{P} can always be labeled correctly, and the rest can at least do as well as a random guess. The analogous result is known to hold for the case of $|A| = 2$. Interestingly, the critical property in generalizing to the multi-class case is not changing the set of classifiers but instead allowing the labels to be sufficiently rich.

3. CONSTRUCTION OF ALGORITHM

The algorithm, which we call τ_S^K , iterates over two loops: within an epoch (inner loop) and across epochs (outer loop). K is the minimum length of an epoch, which we increase to improve the accuracy of the final hypothesis of the algorithm. The size of each epoch is a random variable $K_i \geq K$. The algorithm has I number of epochs, which is also a random variable. While the total number of periods to generate the final hypothesis is random variable $\sum_{i=1}^I K_i$, we show that as K becomes large, the total number of periods is bounded by $\min(|A|, |P|)K$ with a probability close to 1.

3.1. Preliminaries. The algorithm is run in discrete period $k = 1, 2, \dots$ within epoch $i = 1, 2, \dots$. Each epoch lasts at least K periods, but the terminal round of an epoch is random. At the start of an epoch, we set the period to $k = 1$ and re-start the procedure.

It is convenient to enumerate the elements in $A = \{1, \dots, |A|\}$ as the first $|A|$ positive integers, highlighting the fact that the weak learnability has little to do with the name of the labels but relies solely on the ordinal properties of A . Let $\pi : A \rightarrow A$ be a permutation over A :

$$\pi(j) = i_j \in \{1, \dots, |A|\} \quad \forall j.$$

Let $A_i = \{1, \dots, |A_i|\} \subset A$ be the collection of the first $|A_i|$ positive integers which is the set of feasible labels at the beginning of epoch i with $A_1 = A$. Let π_i be a permutation over A_i .

3.2. Within Epoch i . At the beginning of epoch i , $A_i \subset A$ and $y_i : \mathcal{P} \rightarrow A_i$ are given. Set $k = 1$. $d_1(p)$ as the uniform distribution over \mathcal{P} . The output of each stage consists three key components: an artificial probability distribution $d_k(p)$ over \mathcal{P} , a threshold rule h_k in step k , and a positive weight α_k .

Suppose that $d_k(p)$ is defined $\forall p \in \mathcal{P}$. Choose h_k by solving

$$\max_{h \in \mathcal{H}_i} \sum_{p \in \mathcal{P}} (\mathbb{I}[h(p) = y_i(p)] - \mathbb{I}[h(p) \neq y_i(p)]) d_k(p). \quad (3.5)$$

We call h_k an optimal classifier. Define

$$\epsilon_k = \mathbf{P}_{d_k} (h_k(p) \neq y_i(p)) \quad (3.6)$$

as the probability that the optimal classifier h_k at k round misclassifies p under d_k . If $\epsilon_k = 0$, then we stop the training and output h_k as the final hypothesis, which perfectly forecasts $y_i(p)$.

Suppose that $\epsilon_k > 0$. Define

$$\alpha_k = \frac{1}{2(|A_i| - 1)} \left[\log \frac{(|A_i| - 1)(1 - \epsilon_k)}{\epsilon_k} \right]. \quad (3.7)$$

Note that the only difference between (3.7) and (2.1) is the coefficient $\frac{1}{2(|A_i| - 1)}$. If $|A_i| = 2$, our algorithm is reduced to AdaBoost.

Define for each p , and each pair $(p, y(p))$,

$$d_{k+1}(p) = \frac{d_k(p) \exp \left(-\alpha_k [(|A_i| - 1) \mathbb{I}(h_k(p) = y_i(p)) - \mathbb{I}(h_k(p) \neq y_i(p))] \right)}{Z_k}$$

where

$$Z_k = \sum_{p \in \mathcal{P}} d_k(p) \exp \left(-\alpha_k [(|A_i| - 1) \mathbb{I}(h_k(p) = y_i(p)) - \mathbb{I}(h_k(p) \neq y_i(p))] \right).$$

Given d_{k+1} , we can recursively define h_{k+1} and ϵ_{k+1} .

Define $\forall a \in A_i = \{1, \dots, |A_i|\}$,

$$F_a^k(p) = \sum_{s=1}^k \alpha_s \mathbb{I}(h_s(p) = a)$$

and

$$\Psi_a^k(p) = (|A_i| - 1) F_a^k(p) - \sum_{a' \neq a} F_{a'}^k(p).$$

Note that

$$\begin{aligned}
\Psi_a^k(p) &= (|A_i| - 1)F_a^k(p) - \sum_{a' \neq a} F_{a'}^k(p) \\
&= \sum_{s=1}^k \alpha_s \left[(|A_i| - 1)\mathbb{I}(h_s(p) = a) - \sum_{a' \neq a} \mathbb{I}(h_s(p) = a') \right] \\
&= \sum_{s=1}^k \alpha_s \left[|A_i| \mathbb{I}(h_s(p) = a) - \left(\mathbb{I}(h_s(p) = a) + \sum_{a' \neq a} \mathbb{I}(h_s(p) = a') \right) \right] \\
&= \sum_{s=1}^k \alpha_s (|A_i| \mathbb{I}(h_s(p) = a) - 1). \tag{3.8}
\end{aligned}$$

Therefore, $\forall p$,

$$a \in \arg \max_{A_i} F_a^k(p) \tag{3.9}$$

if and only if

$$a \in \arg \max_{A_i} \Psi_a^k(p). \tag{3.10}$$

We use $\Psi_a^k(p)$ to predict $y_i(p)$, instead of $F_a^k(p)$. Since $\Psi_a^k(p)$ is a linear combination of $\mathbb{I}(h_s(p) = a)$, our algorithm is an ensemble algorithm. Since $\Psi_a^k(p)$ is calculated recursively within an epoch, and the change between two epochs depends only on the final outcome from the previous epoch, we regard our algorithm a recursive algorithm.

Except for the case of $|A_i| = 2$, (3.9) does not guarantee that $a = y_i(p)$ with probability close to 1, as $k \rightarrow \infty$. Mukherjee and Schapire (2013) argues that the failure is due to the fact that the weak learnability (2.3) is too weak, and proposed a stronger criterion of EOR (Edge over Random) criterion.

We take a different tack. We strengthen (2.3) to (2.4) by requiring the same condition for any subset of A with at least two elements. Instead of solving (3.10), we eliminate $a \in A_i$ which is not an element of $\arg \max_{A_i} \Psi_a^k(p)$. A critical step is to show that $a = y_i(p) \in A_i$ survives the elimination process with a probability close to 1. Since $|A| < \infty$, we can “solve” (3.10) in a finite number of iterations. Roughly speaking, (2.3) may be too weak to solve (3.9) directly, but is sufficiently strong to solve the equivalent problem (3.10) indirectly, if (2.3) holds for any subset with at least two elements.

Recall that K is the minimum number of periods in epoch i , which is a parameter of the algorithm. Define

$$K_i = \min \left\{ k \geq K \mid \forall p \in \mathcal{P}, \exists a \in A_i, \Psi_a^k(p) < 0 \right\}$$

as the first period in epoch i when each p has a label $a \in A_i$ where $\Psi_a^k(p) < 0$. Epoch i terminates at K_i period.

Before we start epoch $i + 1$, we need to update A_i and y_i . Roughly speaking, A_{i+1} is obtained by eliminating $a \in A_i$ with $\Psi_a^{K_i}(p) < 0$. Because the label with $\Psi_a^{K_i}(p) < 0$ can vary across different p , we need additional work. First, we need to eliminate the equal

number of labels from A_i for each p , and then, permute the remaining elements so that we can “homogenize” the set of labels for each p .

Define

$$B_i(p) = \{a \in A_i \mid \Psi_a^{K_i}(p) < 0\}$$

as the set of labels in A_i with a negative value of $\Psi_a^{K_i}(p)$, which are the candidates of elimination. Because $|B_i(p)|$ can vary across p , we cannot eliminate all elements in $B_i(p)$ from A_i , if we want to keep the number of surviving elements in A_i the same across p . Define

$$N_i = \min_p |B_i(p)| \geq 1$$

which is at least 1, by the definition of K_i . Recall that each element in $B_i(p)$ is a positive integer. Instead of eliminating all elements in $B_i(p)$, we eliminate N_i elements from $B_i(p)$ $\forall p \in \mathcal{P}$. Recall that

$$A_i = \{1, \dots, |A_i|\}.$$

We can write

$$B_i(p) = \{i(p)_1, \dots, i(p)_{|B_i(p)|}\}.$$

For convenience, we select the last N_i elements in $B_i(p)$ to construct $\underline{B}_i(p)$:

$$\underline{B}_i(p) = \left\{ i(p)_{|B_i(p)|-N_i+1}, \dots, i(p)_{|B_i(p)|} \right\}.$$

By the construction of $\underline{B}_i(p)$,

$$|A_i \setminus \underline{B}_i(p)| = |A_i \setminus \underline{B}_i(p')| \quad \forall p \neq p' \in \mathcal{P}. \quad (3.11)$$

The remaining step is to re-label each element in $A_i \setminus \underline{B}_i(p)$ so that each p has the same set of labels.

Define

$$A_{i+1} = \{1, \dots, |A_i \setminus \underline{B}_i(p)|\}$$

as the collection of the first $|A_i \setminus \underline{B}_i(p)|$ positive integers. Note

$$|A_{i+1}| = |A_i \setminus \underline{B}_i(p)| \quad \forall p \in \mathcal{P}$$

by (3.11).

We arrange the elements in $A_i \setminus \underline{B}_i(p)$ according to the index's order. For each $p \in \mathcal{P}$, define a permutation

$$\pi_i^p : A_{i+1} \rightarrow A_i \setminus \underline{B}_i(p)$$

where

$$\pi_i^p(1) = \min\{a \in A_i \setminus \underline{B}_i(p)\}$$

and let $\pi_i^p(1) = a_1$. Given $a_1, \dots, a_{j-1} \in A_i \setminus \underline{B}_i(p)$,

$$\pi_i^p(j) = \min\left\{a \in A_i \setminus [\underline{B}_i(p) \cup \{a_1, \dots, a_{j-1}\}]\right\}.$$

Since π_i^p is a bijection over A_{i+1} , $(\pi_i^p)^{-1}$ exists. The correct label in epoch i is

$$y_{i+1}(p) = (\pi_i^p)^{-1}(y_i(p)) \quad \forall p.$$

3.3. Iteration over epoch. Replace A_i and y_i by A_{i+1} and y_{i+1} . Start the entire process described in Section 3.2. The iteration over epoch stops at the first i where

$$|A_i \setminus \underline{B}_i(p)| = 1 \quad \forall p \quad (3.12)$$

Define

$$I = \min \{i \mid |A_i \setminus \underline{B}_i(p)| = 1 \quad \forall p\}.$$

as the number of epoch when the algorithm stops. Since epoch i last K_i periods, the total number of periods to train the algorithm is

$$t = \sum_{i=1}^I K_i$$

which is a random variable.

Let us call the constructed algorithm τ_S^K because the algorithm is parameterized by the minimum length of each epoch K . The element in $A_I \setminus \underline{B}_I(p)$ is the source of the final hypothesis of τ_S^K . Since we know $\pi_i^p \forall p \in \mathcal{P}$ and $\forall i \in \{1, \dots, I\}$, the final hypothesis of τ_S^K is

$$\tau_S^K(D_t)(p) = \pi_1^p \left(\dots \pi_I^p (A_I \setminus \underline{B}_I(p)) \right) \quad \forall p \in \mathcal{P}.$$

We are interested in $\mathbf{P} \left(\tau_S^K(D_t)(p) = y(p) \right)$.

Theorem 3.1. *Assume that $|A| < \infty$ and $|\mathcal{P}| < \infty$, and that \mathcal{H} is iteratively weakly learnable.*

- (1) $I \leq \min(|A|, |\mathcal{P}|)$
- (2) $\exists \rho > 0$ and \overline{K} such that $\forall K \geq \overline{K}$, with probability at least $1 - e^{-\rho K}$,
 - (a) $K_i = K \forall i \in \{1, \dots, I\}$
 - (b) $\mathbf{P} \left(\tau_S^K(D_t) = y(p) \right) \geq 1 - e^{-\rho K}$ where $t = \sum_{i=1}^I K_i$.

Proof. See Appendix A. □

4. GENERALIZATION ERROR

Our last result shows how our approach lends itself to a simple formulation of the generalization error. Let \mathcal{S} be the sample space generated by K samples and $\mathbf{P}_\mathcal{S}$ be the probability measure over the sample space. Theorem 3.1 proved that the training error vanishes at an exponential rate:

$$\mathbf{P}_\mathcal{S} \left(\tau_S^K(D_{KI})(p) \neq y(p) \right) \leq e^{-\rho K}$$

where K is the length of an epoch and I is the number of epochs. We know that

$$I \leq \min(|A|, |\mathcal{P}|).$$

We can show that the generalization error of τ_S^K vanishes as $K \rightarrow \infty$. Let $\mathbf{P}_\mathcal{D}$ be the population probability distribution.

Theorem 4.1. $\forall m \geq K$,

$$\mathbf{P}_\mathcal{D} \left(\tau_S^K(D_{KI})(p) \neq y(p) \right) \leq \mathbf{P}_\mathcal{S} \left(\tau_S^K(D_{KI})(p) \neq y(p) \right) + \mathcal{O} \left(\frac{1}{\sqrt{m}} \right).$$

Proof. See Appendix B. □

5. CONCLUSION

As τ_S^K is built on the framework of SAMME, the algorithm inherits the same statistical foundation as SAMME and is easier to motivate than other extensions of AdaBoost. Some extensions of AdaBoost (e.g., Mukherjee and Schapire (2013)) may have to “tune parameters” of the algorithm to achieve efficiency. The selection of a proper parameter of the algorithm requires information about \mathcal{H} . In contrast, the construction of τ_S^K requires no knowledge about \mathcal{H} , other than the fact that the set of weak hypothesis satisfies iterative weak learnability.

It is straightforward to check for hypothesis classes that are closed under permutations, making it easy to verify the iterative weak learnability. This allowed us to exhibit the algorithm straightforwardly, without needing to check that the algorithm can outperform random guessing for a rich set of possible cost functions.

APPENDIX A. PROOF OF THEOREM 3.1

A.1. **Uniform Bound for I .** We show

$$I \leq \min(|A|, |\mathcal{P}|).$$

Recall that $\forall d_k, h_k$ is maximizing

$$\sum_{p \in \mathcal{P}} (\mathbb{I}[h(p) = y_i(p)] - \mathbb{I}[h(p) \neq y_i(p)]) d_k(p).$$

If $\exists p \in \mathcal{P}$ such that $h(p) = a \in A \setminus y(\mathcal{P})$, then h cannot be an optimal weak hypothesis. Suppose that $a = a^+ \in A \setminus y(\mathcal{P})$. The decision maker can increase the performance by choosing a^+ from $y_i(\mathcal{P})$. Thus, $\forall k, h_k(\mathcal{P}) \subset y_i(\mathcal{P})$.

We have $\forall k, \forall a \notin y_i(\mathcal{P}) \subset A$,

$$\Psi_a^k(p) = \sum_{s=1}^k \alpha_s (|A| \mathbb{I}(h_s(p) = a) - 1) = - \sum_{s=1}^k \alpha_s < 0,$$

since $h_s(p) \neq a \forall a \in A \setminus y(\mathcal{P})$. Every $a \in A \setminus y(\mathcal{P})$ is eliminated at the end of the first epoch. Since $|y(\mathcal{P})| \leq |\mathcal{P}|$, the number of remaining elements after the first epoch is no more than $\min(|A|, |\mathcal{P}|)$.

A.2. **Within Epoch.** We show that if the minimum length K of an epoch is sufficiently large, then $K_i = K$ with a probability close to 1. That is, by the time when the algorithm reaches the scheduled end of i th epoch, $\exists a_i(p) \in A_i$ such that $\Psi_{a_i(p)}^K(p) < 0$. We calculate the lower bound of the probability. There exists $\rho_i > 0$ such that the probability is bounded from below by $1 - e^{-\rho_i K}$.

Recall that $\forall a \in A_i$,

$$F_a^k(p) = \sum_{s=1}^k \alpha_s \mathbb{I}(h_s(p) = a)$$

and

$$\begin{aligned} \Psi_a^k(p) &= (|A_i| - 1)F_a^k(p) - \sum_{a' \neq a} F_{a'}^k(p) \\ &= \sum_{s=1}^k \alpha_s (|A_i| \mathbb{I}(h_s(p) = a) - 1). \end{aligned}$$

Thus,

$$\begin{aligned} \sum_a \Psi_a^k(p) &= \sum_a \left(\sum_{s=1}^k \alpha_s [|A_i| \mathbb{I}(h_s(p) = a) - 1] \right) \\ &= \sum_{s=1}^k \alpha_s \sum_a [|A_i| \mathbb{I}(h_s(p) = a) - 1] \\ &= \sum_{s=1}^k \alpha_s \left[|A_i| \sum_a \mathbb{I}(h_s(p) = a) - \sum_a 1 \right] \\ &= \sum_{s=1}^k \alpha_s [|A_i| - |A_i|] = 0. \end{aligned}$$

Iterative weak learnability implies that $\forall d_k, \exists h_k \in \mathcal{H}_i$ so that

$$\sum_{p \in \mathcal{P}} (\mathbb{I}[h_k(p) = y_i(p)] - \mathbb{I}[h_k(p) \neq y_i(p)]) d(p) \geq \frac{2 - |A_i|}{|A_i|} + \gamma_k.$$

for some $\gamma_k > 0$. The weak learnability implies that $\inf_k \gamma_k > 0$.

Lemma A.1.

$$\mathbf{P} \left(\Psi_{y_i(p)}^k(p) \leq 0 \right) \leq e^{-\rho_i k}.$$

Proof. Note that

$$\begin{aligned}
\mathbf{P}\left(\Psi_{y_i(p)}^k(p) \leq 0\right) &= \mathbf{E}_{d_1} \mathbb{I}\left(\Psi_{y_i(p)}^k(p) \leq 0\right) \\
&= \mathbf{E}_{d_1} \mathbb{I}\left(-\Psi_{y_i(p)}^k(p) \geq 0\right) \\
&\leq \mathbf{E}_{d_1} \exp\left(-\Psi_{y_i(p)}^k(p)\right).
\end{aligned} \tag{A.13}$$

From the updating formula of d_t ,

$$\begin{aligned}
d_{k+1}(p) &= \frac{d_k(p) \exp\left[-\alpha_k(|A_i| - 1)\mathbb{I}(h_k(p) = y_i(p)) + \alpha_k\mathbb{I}(h_k(p) \neq y_i(p))\right]}{Z_k} \\
&= \frac{d_1(p) \exp\left[-\sum_{s=1}^k \alpha_s(|A_i| - 1)\mathbb{I}(h_s(p) = y_i(p)) + \alpha_s\mathbb{I}(h_s(p) \neq y_i(p))\right]}{\prod_{s=1}^k Z_s} \\
&= \frac{d_1(p) \exp\left[-(|A_i| - 1)F_{y_i(p)}^k(p) + \sum_{a \neq y_i(p)} F_a^k(p)\right]}{\prod_{s=1}^k Z_s}
\end{aligned} \tag{A.14}$$

$$= \frac{d_1(p) \exp\left[-\Psi_{y_i(p)}^k(p)\right]}{\prod_{s=1}^k Z_s}. \tag{A.15}$$

By summing over p , we have

$$\sum_p d_1(p) \exp\left[-\Psi_{y_i(p)}^k(p)\right] = \prod_{s=1}^k Z_s.$$

Note $\forall k \geq 1, \forall s \in \{1, \dots, k\}$,

$$\begin{aligned}
Z_s &= \sum_p d_s(p) \exp\left[-\alpha_s(|A_i| - 1)\mathbb{I}(h_s(p) = y_i(p)) + \alpha_s\mathbb{I}(h_s(p) \neq y_i(p))\right] \\
&= \sum_{h_s(p)=y_i(p)} d_s(p) e^{-(|A_i|-1)\alpha_s} + \sum_{h_s(p) \neq y_i(p)} d_s(p) e^{\alpha_s} \\
&= e^{-\alpha_s(|A_i|-1)}(1 - \epsilon_s) + e^{\alpha_s} \epsilon_s \\
&= (|A_i| - 1)^{-\frac{1}{2}} \epsilon_s^{\frac{1}{2}} (1 - \epsilon_s)^{\frac{1}{2}} + (|A_i| - 1)^{\frac{1}{2(|A_i|-1)}} (1 - \epsilon_s)^{\frac{1}{2(|A_i|-1)}} \epsilon_s^{1 - \frac{1}{2(|A_i|-1)}} \\
&\equiv \varphi(\epsilon_s).
\end{aligned}$$

Lemma A.2. $\varphi(\epsilon)$ is a strictly concave function of $\epsilon \in [0, 1]$.

$$\varphi\left(1 - \frac{1}{|A_i|}\right) = 1 \tag{A.16}$$

and

$$\varphi'\left(1 - \frac{1}{|A_i|}\right) = 0. \tag{A.17}$$

Proof. $\epsilon^{\frac{1}{2}}(1 - \epsilon)^{\frac{1}{2}}$ and $(1 - \epsilon)^{\frac{1}{2(|A_i|-1)}} \epsilon^{1 - \frac{1}{2(|A_i|-1)}}$ are strictly concave functions over ϵ . $\varphi(\epsilon)$ is the sum of two strictly concave functions. Thus, $\varphi(\epsilon)$ is strictly concave.

(A.16) follows from a simple calculation.

$$\begin{aligned}
&\varphi\left(1 - \frac{1}{|A_i|}\right) \\
&= (|A_i| - 1)^{-\frac{1}{2}} \left(\frac{|A_i| - 1}{|A_i|}\right)^{\frac{1}{2}} \left(\frac{1}{|A_i|}\right)^{\frac{1}{2}} + (|A_i| - 1)^{\frac{1}{2(|A_i|-1)}} \left(\frac{|A_i| - 1}{|A_i|}\right)^{1 - \frac{1}{2(|A_i|-1)}} \left(\frac{1}{|A_i|}\right)^{\frac{1}{2(|A_i|-1)}} \\
&= \frac{1}{|A_i|} + \left(1 - \frac{1}{|A_i|}\right) = 1.
\end{aligned}$$

The calculation of φ' is a little more involved.

$$\begin{aligned}
& \varphi' \left(1 - \frac{1}{|A_i|} \right) \\
&= (|A_i| - 1)^{-\frac{1}{2}} \left[\frac{1}{2} \left(\frac{1}{|A_i| - 1} \right)^{\frac{1}{2}} - \frac{1}{2} (|A_i| - 1)^{\frac{1}{2}} \right] \\
&\quad + (|A_i| - 1)^{\frac{1}{2(|A_i| - 1)}} \left[\left(1 - \frac{1}{2(|A_i| - 1)} \right) \left(\frac{1}{|A_i| - 1} \right)^{\frac{1}{2(|A_i| - 1)}} - \frac{1}{2(|A_i| - 1)} (|A_i| - 1)^{1 - \frac{1}{2(|A_i| - 1)}} \right] \\
&= \left[\frac{1}{2} \frac{|A_i|}{|A_i| - 1} - 1 \right] + \left[1 - \frac{1}{2} \frac{|A_i|}{|A_i| - 1} \right] \\
&= 0.
\end{aligned}$$

□

Since \mathcal{H}_i is weakly learnable, $\exists \gamma_k > 0$ such that

$$0 \leq \epsilon_k \leq 1 - \frac{1}{|A_i|} - \gamma_k < 1 - \frac{1}{|A_i|}.$$

Weak learnability implies that

$$\gamma_k > 0 \quad \text{and} \quad \gamma = \inf_k \gamma_k > 0.$$

Since φ is strictly concave, $\varphi'(\epsilon) > 0 \forall \epsilon < 1 - \frac{1}{|A_i|}$. Thus, $\exists \hat{\gamma} > 0$ such that $\forall \epsilon_k \leq 1 - \frac{1}{|A_i|} - \gamma$,

$$\varphi(\epsilon_k) \leq \varphi \left(1 - \frac{1}{|A_i|} - \gamma \right) = 1 - \hat{\gamma} < 1.$$

Thus,

$$Z_k = \varphi(\gamma_k) \leq \varphi(\gamma) \leq 1 - \hat{\gamma}$$

for some $\hat{\gamma} > 0$. Thus,

$$d_{k+1}(p) = d_1(p) \prod_{s=1}^k Z_s \leq \frac{1}{|\mathcal{P}|} (1 - \hat{\gamma})^k = \frac{1}{|\mathcal{P}|} e^{-\rho_i k}$$

where

$$\rho_i = -\log(1 - \hat{\gamma}) > 0. \tag{A.18}$$

□

By the definition, $\forall p, \forall k$,

$$\sum_{a \in A_i} \Psi_a^k(p) = 0. \tag{A.19}$$

By (A.19), $\forall p$, if $\Psi_{y(p)}^K(p) > 0$, then $\exists \underline{a}_i(p) \in A_i$ such that

$$\Psi_{\underline{a}_i(p)}^K(p) < 0.$$

Lemma A.1 implies

$$\underline{a}_1(p) \neq y_i(p)$$

with probably at least $1 - e^{-\rho_i K}$. Hence, $K_i = K$ with probability at least $1 - e^{-\rho_i K}$.

As Lemma A.1 indicates, if $a = y(p)$, then $\Psi_a^k(p) > 0$ almost surely as $K \rightarrow \infty$. If $|A| = 2$, then (A.19) implies that exactly one label is associated with a positive value of $\Psi_a^k(p)$. Thus, searching for a with $\Psi_a^k(p) > 0$ is the proper way to identify a correct label. However, if $|A| > 2$, then the positive sign of $\Psi_a^k(p)$ does not imply that $a = y(p)$. However, if $\Psi_a^k(p) < 0$, then Lemma A.1 implies that $a \neq y(p)$ almost surely, as $K \rightarrow \infty$. The idea of τ_S^K is to eliminate a with $\Psi_a^k(p) < 0$ in each epoch, until we have a single element left.

A.3. Across Epochs. We construct A_{i+1} by eliminating some elements from A_i and permuting the remaining elements according to π_i . With probability $1 - e^{-\rho_i K}$, the true label survives the elimination from A_i , and is included in A_{i+1} , permuted according to π_i .

Since we eliminate at least one element from A_i , the total number of epochs cannot be larger than $|A|$. Let I be the epoch when A_{I+1} is a singleton. With probability

$$\prod_{i=1}^I (1 - e^{-\rho_i K}),$$

the remaining label $y_I(p)$ is the real label which has been permuted by a series of π_1, \dots, π_I . We can recover the real label

$$y(p) = \pi_1 (\dots \pi_I(y_I(p))).$$

Define

$$\underline{\rho} = \frac{1}{2} \min \rho_i > 0.$$

Then, for a sufficiently large K ,

$$\prod_{i=1}^I (1 - e^{-\rho_i K}) \geq 1 - |A| e^{-K \min_i \rho_i} = 1 - e^{-K(\min_i \rho_i + \frac{1}{K} \log |A|)} \geq 1 - e^{-\underline{\rho} K}$$

which completes the proof.

APPENDIX B. PROOF OF THEOREM 4.1

The proof follows the same logic as the proof of Theorem 5.1 on page 98 of Schapire and Freund (2012). We need to change the notation to accommodate multiple classes ($|A| \geq 2$). For the reference, we sketch the proof by following the original proof of Theorem 5.1 in Schapire and Freund (2012), while pointing out the place where we need a change and referring the rest to Schapire and Freund (2012).

B.1. Preliminaries. Recall that $\forall i \leq I, \forall a \in A_i$

$$\Psi_a^k(p) = \sum_{s=1}^k \alpha_s [|A_i| \mathbb{I}(h_s(p) = a) - 1].$$

Define

$$g_a^s(p) = |A_i| \mathbb{I}(h_s(p) = a) - 1,$$

$$\beta_s = \frac{\alpha_s}{\sum_{s'=1}^k \alpha_{s'}}$$

and

$$\psi_a^k(p) = \sum_{s=1}^k \beta_s g_a^s(p).$$

We drop subscript a because we focus on $a = y(p)$ and write instead

$$g(p) = g_{y(p)}(p) \text{ and } \psi^k(p) = \psi_{y(p)}^k(p).$$

Recall that \mathcal{H} is the collection of feasible weak hypotheses, whose generic element is h . Define

$$\mathcal{G} = \{g \mid \exists h \in \mathcal{H}, g(p) = |A_i| \mathbb{I}(h(p) = y(p)) - 1\}$$

For a fixed n , define

$$\mathcal{A}_n = \left\{ f : p \mapsto \frac{1}{n} \sum_{j=1}^n g_j(p) \mid g_1, \dots, g_n \in \mathcal{G} \right\}.$$

Let $\tilde{g}_1, \dots, \tilde{g}_n$ be the collection of n randomly selected elements from \mathcal{G} and

$$\tilde{f} = \frac{1}{n} \sum_{j=1}^n \tilde{g}_j(p)$$

where

$$\mathbf{E}g_j = \sum_{s=1}^k \beta_s g_{y(p)}^s(p) = \psi_{y(p)}^k(p).$$

If

$$\tau_S^K(D_{KI})(p) \neq y(p),$$

then $y(p)$ is eliminated from A_i in i -th epoch $i < I$, which implies

$$\psi_{y(p)}^K(p) < 0.$$

Thus,

$$\begin{aligned} & \mathbf{P}_{\mathcal{D}} \left(\tau_S^K(D_{KI}) \neq y \right) \\ & \leq \prod_{i=1}^I \mathbf{P}_{\mathcal{D}} \left(\exists i < I, \psi_{y(p)}^K(p) < 0 \mid \psi_{y(p)}^K(p) \geq 0, j \leq i \right) \\ & = \prod_{i=1}^I \mathbf{P}_{\mathcal{D}} \left(\psi_{y(p)}^K(p) < 0 \text{ at the end of } i\text{-th epoch} \right). \end{aligned}$$

The last equality stems from the facts that the samples are independently drawn over periods, and that we re-initialize the algorithm at the beginning of each epoch.

We only show that in the first epoch

$$\mathbf{P}_{\mathcal{D}} \left(\psi_{y(p)}^K(p) < 0 \right) \leq \mathbf{P}_{\mathcal{S}} \left(\psi_{y(p)}^K(p) < 0 \right) + \mathcal{O} \left(\frac{1}{\sqrt{m}} \right)$$

where \mathcal{S} is the collection of all m samples. The proof for the remaining epochs follows the same logic. Since $I \leq \min(|A|, |\mathcal{P}|) \leq |A|$, we then have

$$\mathbf{P}_{\mathcal{D}} \left(\tau_S^K(D_{KI}) \neq y \right) \leq \mathbf{P}_{\mathcal{S}} \left(\tau_S^K(D_{KI}) \neq y \right) + |A| \mathcal{O} \left(\frac{1}{\sqrt{m}} \right).$$

B.2. Sketch of Proof.

Lemma B.1. (Re-statement of Lemma 5.2 in Schapire and Freund (2012)) $\forall \theta$,

$$\mathbf{P}_{\tilde{f}} \left(|\tilde{f}(p) - f(p)| \geq \frac{\theta}{2} \right) \leq 2e^{-\frac{\theta^2}{2|A|^2}} \equiv \beta_{n,\theta}. \quad (\text{B.20})$$

Proof. Note that the only change from the original statement in Schapire and Freund (2012) is $\beta_{n,\theta}$ which needs to be scaled to accommodate the case $|A| > 2$. Note that $g(p) \in \{|A| - 1, -1\}$. We need to scale g into a random variable in $[-1, 1]$ and modify $\beta_{n,\theta}$ accordingly. The remainder of the proof is identical with the original proof of Lemma 5.2 in Schapire and Freund (2012). \square

Lemma 5.3 of Schapire and Freund (2012) is not needed. The proof of Lemma 5.4 of Schapire and Freund (2012) applies to our case without modifying the parameter, from which the conclusion of Theorem 4.1 follows.

REFERENCES

- CHO, I.-K., AND J. LIBGOBER (2020): “Machine Learning for Strategic Inference,” Emory University and University of Southern California.
- FRIEDMAN, J., T. HASTIE, AND R. TIBSHIRANI (2000): “Additive Logistic Regression: A Statistical View of Boosting,” *The Annals of Statistics*, 28(2), 337–407.
- MUKHERJEE, I., AND R. E. SCHAPIRE (2013): “A Theory of Multiclass Boosting,” *Journal of Machine Learning Research*, 14, 437–497.
- SCHAPIRE, R. E., AND Y. FREUND (2012): *Boosting: Foundations and Algorithms*. MIT Press.
- SCHAPIRE, R. E., AND Y. SINGER (1999): “Improved Boosting Algorithms Using Confidence-Rate Predictions,” *Machine Learning*, 37, 297–336.
- TVERSKY, A. (1972): “Elimination by aspects: A theory of choice,” *Psychological Review*, 79(4), 281–299.
- ZHU, J., H. ZOU, S. ROSSET, AND T. HASTIE (2009): “Multi-class AdaBoost,” *Statistics and Its Interface*, 2, 349–360.

DEPARTMENT OF ECONOMICS, EMORY UNIVERSITY, ATLANTA, GA 30322 USA

Email address: `icho30@emory.edu`

URL: `https://sites.google.com/site/inkoocho`

DEPARTMENT OF ECONOMICS, UNIVERSITY OF SOUTHERN CALIFORNIA, LOS ANGELES, CA 90089
USA

Email address: `libgober@usc.edu`

URL: `http://www.jonlib.com/`