
On Agnostic PAC Learning in the Small Error Regime

Julian Asilis
USC
asilis@usc.edu

Mikael Møller Høgsgaard
Aarhus University
hogsgaard@cs.au.dk

Grigoris Veleghas*
Google Research
gvelegkas@google.com

Abstract

Binary classification in the classic PAC model exhibits a curious phenomenon: Empirical Risk Minimization (ERM) learners are suboptimal in the realizable case yet optimal in the agnostic case. Roughly speaking, this owes itself to the fact that non-realizable distributions $\mathcal{D}$ are more difficult to learn than realizable distributions – even when one discounts a learner’s error by $\text{err}(h_{\mathcal{D}}^*)$, i.e., the error of the best hypothesis in $\mathcal{H}$. Thus, optimal agnostic learners are permitted to incur excess error on (easier-to-learn) distributions $\mathcal{D}$ for which $\tau = \text{err}(h_{\mathcal{D}}^*)$ is small. Recent work of Hanneke, Larsen, and Zhivotovskiy (FOCS ’24) addresses this shortcoming by including τ itself as a parameter in the agnostic error term. In this more fine-grained model, they demonstrate tightness of the error lower bound $\tau + \Omega(\sqrt{\tau(d+\log(1/\delta))/m} + d+\log(1/\delta)/m)$ in a regime where $\tau > d/m$, and leave open the question of whether there may be a higher lower bound when $\tau \approx d/m$, with d denoting $\text{VC}(\mathcal{H})$. In this work, we resolve this question by exhibiting a learner which achieves error $c \cdot \tau + O(\sqrt{\tau(d+\log(1/\delta))/m} + d+\log(1/\delta)/m)$ for a constant $c \leq 2.1$, matching the lower bound and demonstrating optimality when $\tau = O(d/m)$. Further, our learner is computationally efficient and is based upon careful aggregations of ERM classifiers, making progress on two other questions of Hanneke, Larsen, and Zhivotovskiy (FOCS ’24). We leave open the interesting question of whether our approach can be refined to lower the constant from 2.1 to 1, which would completely settle the complexity of agnostic learning.

1 Introduction

Binary classification stands as perhaps the most fundamental setting in supervised learning, and one of its best-understood. In Valiant’s celebrated Probably Approximately Correct (PAC) framework, learning is formalized using a domain $\mathcal{X}$ in which unlabeled data points reside, the label space $\mathcal{Y} = \{\pm 1\}$, and a probability distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$. The purpose of a learner $\mathcal{A}$ is to receive a training set $S = (x_i, y_i)_{i=1}^m$ of points drawn i.i.d. from $\mathcal{D}$ and then emit a hypothesis $\mathcal{A}(S) : \mathcal{X} \rightarrow \mathcal{Y}$ which is unlikely to mispredict the label of a new data point (x, y) drawn from $\mathcal{D}$.

At the center of a learning problem is a hypothesis class of functions $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with which the learner $\mathcal{A}$ must compete. More precisely, $\mathcal{A}$ is tasked with emitting a hypothesis f whose *error*, denoted $\mathcal{L}_{\mathcal{D}}(f) := \mathbb{P}_{(x,y) \sim \mathcal{D}}(f(x) \neq y)$, is nearly as small as that of the best hypothesis in $\mathcal{H}$. We assume for simplicity such a best hypothesis exists and denote it as $h_{\mathcal{D}}^* := \arg \min_{h \in \mathcal{H}} \mathcal{L}_{\mathcal{D}}(h)$ or when $\mathcal{D}$ is clear from context with h^* . Of course, the distribution $\mathcal{D}$ is unknown to $\mathcal{A}$, and $\mathcal{A}$ is judged on its performance across a broad family of allowable distributions $\mathbb{D}$. In *realizable* learning, $\mathbb{D} = \{\mathcal{D} : \inf_{h \in \mathcal{H}} \mathcal{L}_{\mathcal{D}}(h) = 0\}$, meaning $\mathcal{A}$ is promised that there always exists a *ground truth* hypothesis $h^* \in \mathcal{H}$ attaining zero error. In this case, for $\mathcal{A}$ to compete with the performance of $h_{\mathcal{D}}^*$ simply requires that it attains error $\leq \epsilon$ for all realizable distributions $\mathcal{D}$ (with high probability over the training set $S \sim \mathcal{D}^m$, and using the smallest amount of data m possible). In *agnostic* learning, $\mathbb{D}$ is defined as the set of all probability distributions over $\mathcal{X} \times \mathcal{Y}$, meaning any data-generating process

is allowed. Thus, $\mathcal{A}$ is not equipped with any information concerning $\mathcal{D} \in \mathbb{D}$ at the outset, but in turn it is only required to emit a hypothesis f with $\mathcal{L}_{\mathcal{D}}(f) \leq \mathcal{L}_{\mathcal{D}}(h_{\mathcal{D}}^*) + \epsilon$.

Perhaps the most fundamental questions in PAC learning ask: What is the minimum number of samples required to achieve error at most ϵ ? And which learners attain these rates? For realizable binary classification, the optimal sample complexity of learning remained a foremost open question for decades, and was finally resolved by breakthrough work of [Hanneke \(2016\)](#) which built upon that of [Simon \(2015\)](#). Notably, optimal binary classification requires the use of learners which emit hypotheses f outside of the underlying hypothesis class $\mathcal{H}$. Such learners are referred to as being *improper*. This has the effect of excluding Empirical Risk Minimization (ERM) from contention as an optimal learner for the realizable case. Recall that ERM learners proceed by selecting a hypothesis $h_S \in \mathcal{H}$ incurring the fewest errors on the training set S . (For realizable learning, note that h_S will then incur zero error on S .) ERM, then, is an example of a *proper* learner, which always emits a hypothesis in the underlying class $\mathcal{H}$. Agnostic learning, however, exhibits no such properness barrier for optimal learning. In fact, ERM algorithms are themselves optimal agnostic learners, despite their shortcomings on realizable distributions. This somewhat counter-intuitive behavior owes itself to the fact that agnostic learners are judged on a worst-case basis across all possible distributions $\mathcal{D}$. Simply put, non-realizable distributions are more difficult to learn than realizable distributions — even when one discounts the error term by $\mathcal{L}_{\mathcal{D}}(h_{\mathcal{D}}^*)$ — and thus ERM learners are permitted to incur some unnecessary error on realizable distributions, so long as they remain within the optimal rates induced by the more difficult (non-realizable) distributions.

In light of this behavior, it is natural to ask for a more refined perspective on agnostic learning, in which the error incurred beyond $\tau := \mathcal{L}_{\mathcal{D}}(h_{\mathcal{D}}^*)$ is itself studied as a function of τ . Note that this formalism addresses the previously-described issue by demanding that learners attain superior performance on distributions with smaller values of τ (such as realizable distributions, for which $\tau = 0$). Precisely this perspective was recently studied by [Hanneke et al. \(2024\)](#), who established that all proper learners, including ERM, are sub-optimal for $\tau \in [\Omega(\ln^{10}(m/d)(d + \ln(1/\delta))/m), o(1)]$, by establishing a new lower bound and an optimal learner in this regime. While settling the sample complexity for a wide range of τ , [Hanneke et al. \(2024\)](#) leave open several interesting questions.

Problem 1.1. *Let $d = \text{VC}(\mathcal{H})$ and $m = |S|$. What is the optimal sample complexity of learning in the regime where $\tau \approx d/m$?*

Problem 1.2. *Are learners based upon majority voting, such as bagging, optimal in the τ -based agnostic learning framework?*

Problem 1.3. *Can one design a computationally efficient learner which is optimal in the τ -based agnostic learning framework?*

The primary focus of our work is to resolve Open Problem 1.1, thereby extending the understanding of optimal error rates across a broader range of τ . As part of our efforts, we also make progress on Open Problems 1.2 and 1.3, as we now describe.

1.1 Overview of Main Results

We now present our primary result and a brief overview of our approach.

Theorem 1.4. *For any domain $\mathcal{X}$, hypothesis class $\mathcal{H}$ of VC dimension d , number of samples m , parameter $\delta \in (0, 1)$, there is an algorithm such that for any distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, 1\}$ it returns a classifier $h_S : \mathcal{X} \rightarrow \{-1, 1\}$ that, with probability at least $1 - \delta$, has error bounded by*

$$\min \left\{ 2.1 \cdot \tau + O \left(\sqrt{\frac{\tau(d + \ln(1/\delta))}{m}} + \frac{d + \ln(1/\delta)}{m} \right), \tau + O \left(\sqrt{\frac{\tau(d + \ln(1/\delta))}{m}} + \frac{\ln^5(m/d) \cdot (d + \ln(1/\delta))}{m} \right) \right\},$$

where τ is the error of the best hypothesis in $\mathcal{H}$.

Notably, Theorem 1.4 settles the sample complexity of agnostic learning in the regime $\tau = O(d/m)$, by exhibiting an optimal learner which attains existing lower bounds. Furthermore, it improves upon the learner of [Hanneke et al. \(2024\)](#) when $\tau = o(\ln^5(m/d) \cdot d/m)$. In light of Theorem 1.4, only the polylog range $\tau \in [\omega(1), o(\ln^{10}(m/d))] \cdot \frac{d + \ln(1/\delta)}{m}$ remains to have its optimal error rates characterized.

Let us briefly describe our approach. [Hanneke et al. \(2024\)](#), building upon [Devroye et al. \(1996\)](#), states that any learner, upon receiving m i.i.d. samples from $\mathcal{D}$, must produce with probability at least

δ a hypothesis incurring error $\tau + \Omega(\sqrt{\tau(d + \ln(1/\delta))/m} + (d + \ln(1/\delta))/m)$ for worst-case $\mathcal{D}$. Our first elementary observation is that, due to this lower bound, by designing a learner whose error is bounded by $c \cdot \tau + O(\sqrt{\tau(d + \ln(1/\delta))/m} + (d + \ln(1/\delta))/m)$ for some numerical constant $c \geq 1$, we can get a tight bound in the regime $\tau = O(d/m)$, thus resolving Open Problem 1.1. Furthermore, to make progress across the full regime $\tau \in [0, 1/2]$, it is crucial to obtain a constant c whose value is close to 1. Motivated by this observation, our primary result gives an algorithm that proceeds by taking majority votes over ERM’s trained on carefully crafted subsamples of the training set S , and which achieves error $2.1 \cdot \tau + O(\sqrt{\tau(d + \ln(1/\delta))/m} + (d + \ln(1/\delta))/m)$. Notably, this brings us within striking distance of the constant $c = 1$ — we leave open the intriguing question of whether our technique can be refined to lower c to 1, which would completely settle the complexity of agnostic learning.

Our approach is inspired by Hanneke (2016), but introduces several new algorithmic components and ideas in the analysis. More concretely, we first modify Hanneke’s sample splitting scheme, and then randomly select a small fraction of the resulting subsamples on which to run ERM, rather than running ERM on all such subsamples. This improves the ERM-oracle efficiency of the algorithm. In order to decrease the value of the constant multiplying τ , our main insight is to run *two* independent copies of the above classifier, as well as one ERM that is trained on elements coming from a certain “region of disagreement” of the previous two classifiers. For any test point x , if both of the voting classifiers agree on a label y and have “confidence” in their vote, we output y ; otherwise we output the prediction of the ERM. We hope that this idea of breaking ties between voting classifiers using ERM’s that are trained on their region of disagreement can find further applications. Our resulting algorithm employs a voting scheme at its center, making progress on Open Problem 1.2, and is computationally efficient with respect to ERM oracle calls, making considerable progress on Open Problem 1.3. Finally, we combine this algorithm with the learner of Hanneke et al. (2024) to obtain a “best-of-both-worlds” result. An overview of all the steps is presented in the start of Section 3.2.

1.2 Related Work

The PAC learning framework for statistical learning theory dates to the seminal work of Valiant (1984), with roots in prior work of Vapnik and Chervonenkis (Vapnik and Chervonenkis, 1964, 1974). In binary classification, finiteness of the VC dimension was first shown to characterize learnability by Blumer et al. (1989). Tight lower bounds on the sample complexity of learning VC classes in the realizable case were established by Ehrenfeucht et al. (1989) and matched by upper bounds of Hanneke (2016), building upon work of Simon (2015). Subsequent works have established different optimal PAC learners for the realizable setting (Aden-Ali et al., 2023, 2024; Høgsgaard, 2025; Larsen, 2023). For agnostic learning in the standard PAC framework, ERM is known to achieve sample complexity matching existing lower bounds (Anthony and Bartlett, 2009; Boucheron et al., 2005; Haussler, 1992). As described, we direct our attention to a more fine-grained view of agnostic learning, in which the error incurred by a learner above the best-in-class hypothesis $h_{\mathcal{D}}^*$ is itself studied as a function of $\tau = \mathcal{L}_{\mathcal{D}}(h_{\mathcal{D}}^*)$. Bounds employing τ in the error term are sometimes referred to as *first-order bounds* and have been previously analyzed in fields such as online learning (Maurer and Pontil, 2009; Wagenmaker et al., 2022). Hanneke et al. (2024) appear to be the first to consider τ -optimal-dependence for upper bounds in PAC learning; we adopt their perspective in this work.

2 Preliminaries

Notation For a natural number $m \in \mathbb{N}$, $[m]$ denotes the set $\{1, \dots, m\}$. Random variables are written in bold face (e.g., $\mathbf{x}$) and their realizations in non-bold type face (e.g., x). For a set Z , Z^* denotes the set of all finite sequences in Z , i.e., $Z^* = \bigcup_{i \in \mathbb{N}} Z^i$. For a sequence S of length m and indices $i \leq j \in [m]$, $S[i : j]$ denotes the smallest contiguous subsequence of S which includes both its i th and j th entries. Furthermore, we employ 1-indexing for sequences. For $S = (a, b, c)$, for instance, $S[1 : 2] = (a, b)$. The symbol $\sqcup$ is used to denote concatenation of sequences, as in $(a, b) \sqcup (c, d) = (a, b, c, d)$. When $S, S' \in Z^*$ are sequences in Z and each element $s \in S$ appears in S' no less frequently than in S , we write $S \sqsubseteq S'$. For a sequence S and a set A we denote by $S \sqcap A$ the longest subsequence of S that consists solely of elements of A . If S is a finite set, then $\mathbb{E}_{x \sim S}[f(x)]$ denotes the expected value of f over a uniformly random draw of $x \in S$.

Learning Theory Let us briefly recall the standard language of supervised learning. Unlabeled data points are drawn from a **domain** $\mathcal{X}$, which we permit to be arbitrary throughout the paper. We

study binary classification, in which data points are labeled by one of two labels in the **label set** $\mathcal{Y} = \{\pm 1\}$. A function $h : \mathcal{X} \rightarrow \mathcal{Y}$ is referred to as a **hypothesis** or **classifier**, and a collection of such functions $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ is a **hypothesis class**. Throughout the paper, we employ the 0-1 loss function $\ell_{0-1} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ defined by $\ell_{0-1}(y, y') = \mathbb{1}[y \neq y']$. A **training set** is a sequence of labeled data points $S = ((x_1, y_1), \dots, (x_m, y_m)) \in (\mathcal{X} \times \mathcal{Y})^*$. A **learner** is a function which receives training sets and emits hypotheses, e.g., $\mathcal{A} : (\mathcal{X} \times \mathcal{Y})^* \rightarrow \mathcal{Y}^{\mathcal{X}}$. The purpose of a learner is to emit a hypothesis h which attains low **error**, or **true error**, with respect to an unknown probability distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$. That is, $\mathcal{L}_{\mathcal{D}}(h) = \mathbb{E}_{(x,y) \sim \mathcal{D}} [\mathbb{1}[h(x) \neq y]]$. A natural proxy for the true error of h is its **empirical error** on a training set $S = (x_i, y_i)_{i \in [m]}$, denoted $\mathcal{L}_S(h) = \mathbb{E}_{(x,y) \sim S} [\mathbb{1}[h(x) \neq y]]$. If $\mathcal{A}$ is a learner for $\mathcal{H}$ with the property that $\mathcal{A}(S) \in \arg \min_{h \in \mathcal{H}} \mathcal{L}_S(h)$ for all training sets S , then $\mathcal{A}$ is said to be an **empirical risk minimization (ERM)** learner for $\mathcal{H}$. Throughout the paper we will use $\mathcal{A}$ to denote an arbitrary ERM learner.

3 Proof Sketch

We now provide a detailed explanation of our approach and a comprehensive sketch of the proof. We divide our discussion into two parts. In the first, we present a simple approach that achieves an error bound of $15\tau + O(\sqrt{\tau(d+\ln(1/\delta))/m} + (d+\ln(1/\delta))/m)$. Recall that this resolves the optimal sample complexity for the regime $\tau \approx d/m$. In the second part, we describe several modifications to the algorithm and new ideas in its analysis which drive the error down to $2.1\tau + O(\sqrt{\tau(d+\ln(1/\delta))/m} + (d+\ln(1/\delta))/m)$.

3.1 First Approach: Multiplicative Constant 15

A crucial component of our algorithm is a scheme $\mathcal{S}'$ for recursively splitting the input training sequence S into subsequences, which adapts Hanneke's recursive splitting algorithm (Hanneke, 2016) and is formalized in Algorithm 1 and depicted in Figure 1. The algorithm takes two training sequences as input: S , the active set, and T , the union of elements chosen in previous recursive calls.

Algorithm 1: Splitting algorithm $\mathcal{S}'$

Input: Training sequences $S, T \in (\mathcal{X} \times \mathcal{Y})^*$, where $|S| = 3^k$ for $k \in \mathbb{N}$.

Output: Family of training sequences.

if $k \geq 6$ **then**

 Partition S into S_1, S_2, S_3 , with S_i being the $(i-1)|S|/3 + 1$ to the $i|S|/3$ training examples of S . Set for each i

$$S_{i,\sqcup} = S_i[1 : 3^{k-4}], \quad S_{i,\cap} = S_i[3^{k-4} + 1 : 3^k],$$

return $[\mathcal{S}'(S_{1,\sqcup}; S_{1,\cap} \sqcup T), \mathcal{S}'(S_{2,\sqcup}; S_{2,\cap} \sqcup T), \mathcal{S}'(S_{3,\sqcup}; S_{3,\cap} \sqcup T)]$

else

return $S \sqcup T$




Figure 1: The splitting process of algorithm $\mathcal{S}'$. The active set S is split into three disjoint sets S_1, S_2, S_3 . Each of these is then split into a new active set (green) and a set of previously recursed-on samples (grey), which are passed down to subsequent recursive calls.

In comparison to Hanneke’s scheme, we have made two modifications which allow us to control the constant multiplying τ , as we will explain shortly. First, we create disjoint splits of the data, in contrast to the overlapping subsets employed by Hanneke. Second, we include more elements in the sets $S_{i,\sqcap}$, which are included in all subsequent training subsequences, than in $S_{i,\sqcup}$, the set on which we make a recursive call.

Let us introduce additional notation. We denote $s'_{\sqcap} := |S|/|S_{1,\sqcap}|$, for $S_{1,\sqcap}$ as defined in Algorithm 1. For $\mathcal{A}$ an ERM learner and $S'(S;T)$ the splitting scheme of Algorithm 1, we set $\mathcal{A}'(S;T) := (\mathcal{A}(S'))_{S' \in S'(S;T)}$, thought of as a multiset. For an example (x,y) , we define $\text{avg}_{\neq}(\mathcal{A}'(S;T))(x,y) = \sum_{h \in \mathcal{A}'(S;T)} \mathbb{1}\{h(x) \neq y\} / |\mathcal{A}'(S;T)|$, i.e., the fraction of incorrect hypotheses in $\mathcal{A}'(S;T)$ for (x,y) . For a natural number $t \in \mathbb{N}$, we let $\hat{\mathcal{A}}'_t(S;T)$, be the random multiset of t hypotheses drawn independently and uniformly at random from $\mathcal{A}'(S;T)$. In a slight overload of notation, we use $\mathbf{t}$ (i.e., t in bold face) to denote the randomness used to draw the t hypotheses from $\mathcal{A}'(S;T)$. Intuitively, one can think of $\hat{\mathcal{A}}'_t$ as a bagging algorithm where the subsampled training sequences are restricted to subsets of $S'(S;T)$. Similarly, we define $\text{avg}_{\neq}(\hat{\mathcal{A}}'_t(S;T))(x,y) = \sum_{h \in \hat{\mathcal{A}}'_t(S;T)} \mathbb{1}\{h(x) \neq y\} / |\hat{\mathcal{A}}'_t(S;T)|$. For a distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, 1\}$, training sequences $S, T \in (\mathcal{X} \times \{-1, 1\})^*$, and $\alpha \in [0, 1]$, we let $\mathcal{L}_{\mathcal{D}}^{\alpha}(\mathcal{A}'(S;T)) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(S;T))(\mathbf{x}, \mathbf{y}) \geq \alpha]$, i.e., the probability that at least an α -fraction of the hypotheses in $\mathcal{A}'(S;T)$ err on a new example drawn from $\mathcal{D}$. We will overload the notation $\mathcal{L}_{\mathcal{D}}$ when considering majorities and write $\mathcal{L}_{\mathcal{D}}(\mathcal{A}'(S;T)) = \mathcal{L}_{\mathcal{D}}^{0.5}(\mathcal{A}'(S;T))$ —the probability of an equal-weighted majority vote fails. Similarly, we define $\mathcal{L}_{\mathcal{D}}^{\alpha}(\hat{\mathcal{A}}'_t(S;T)) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\hat{\mathcal{A}}'_t(S;T))(\mathbf{x}, \mathbf{y}) \geq \alpha]$ and $\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}'_t(S;T)) = \mathcal{L}_{\mathcal{D}}^{0.5}(\hat{\mathcal{A}}'_t(S;T))$. Finally, we let $\hat{\mathcal{A}}'_t(S) = \hat{\mathcal{A}}'_t(S; \emptyset)$, and $\hat{\mathcal{A}}'_t(S)(x) = \text{sign}(\sum_{h \in \hat{\mathcal{A}}'_t(S;T)} h(x))$.

We now describe our first approach, which we break into three steps. The first step relates the error of $\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}'_t(\mathbf{S}))$ to $\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}, \emptyset))$, while the second and third steps bound the error of $\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}, \emptyset))$. The first step borrows ideas from Larsen (2023) and the last step from Hanneke (2016), but there are several technical bottlenecks in the analysis that do not appear in these works, as they consider the realizable setting, for which $\tau = 0$.

Relating the error of $\hat{\mathcal{A}}'_t(\mathbf{S})$ to $\mathcal{A}'(\mathbf{S}, \emptyset)$: We will demonstrate how to bound the error of the random classifier $\hat{\mathcal{A}}'_t(\mathbf{S})$ using the error of $\mathcal{A}'(\mathbf{S}, \emptyset)$. First, let $\mathbf{S} \sim \mathcal{D}^m$ and assume we have shown

$$\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}; \emptyset)) \leq 15\tau + O\left(\sqrt{\frac{\tau(d + \ln(1/\delta))}{m}} + \frac{d + \ln(1/\delta)}{m}\right) \quad (1)$$

with probability at least $1 - \delta/2$. Consider the event $E = \{(x, y) : \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 49/100\}$. By the law of total expectation, we bound the error of $\hat{\mathcal{A}}'_t(\mathbf{S})$ as $\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}'_t(\mathbf{S})) \leq \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}(E) + \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\mathbb{1}\{\hat{\mathcal{A}}'_t(\mathbf{S})(\mathbf{x}) \neq \mathbf{y}\} \mid \bar{E}]$. We see that the first term on the right-hand side is $\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}; \emptyset))$, which we have assumed for the moment can be bounded by Equation (1). We now argue that the second term can be bounded by $O((d + \ln(1/\delta))/m)$. Note that under the event $\bar{E} = \{(x, y) : \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}; \emptyset))(x, y) < 49/100\}$, strictly more than half of the hypotheses in $\mathcal{A}'(\mathbf{S}; \emptyset)$ — from which the hypotheses of $\hat{\mathcal{A}}'_t(\mathbf{S})$ are drawn — are correct. Using this, combined with Hoeffding’s inequality and switching the order of expectation (as $\mathbf{t}$ and $(\mathbf{x}, \mathbf{y})$ are independent), we get that $\mathbb{E}_{\mathbf{t}}[\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\mathbb{1}\{\hat{\mathcal{A}}'_t(\mathbf{S})(\mathbf{x}) \neq \mathbf{y}\} \mid \bar{E}]] \leq \exp(-\Theta(t))$. Setting $t = \Theta(\ln(m/(\delta(\ln(1/\delta) + d))))$ then gives $\mathbb{E}_{\mathbf{t}}[\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\mathbb{1}\{\hat{\mathcal{A}}'_t(\mathbf{S})(\mathbf{x}) \neq \mathbf{y}\} \mid \bar{E}]] = O((\delta(\ln(1/\delta) + d))/m)$. By an application of Markov’s inequality, this implies with probability at least $1 - \delta/2$ over the draws of hypotheses in $\hat{\mathcal{A}}'_t(\mathbf{S})$ that $\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\mathbb{1}\{\hat{\mathcal{A}}'_t(\mathbf{S})(\mathbf{x}) \neq \mathbf{y}\} \mid \bar{E}] = O((d + \ln(1/\delta))/m)$. This bounds the second term in the error decomposition of $\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}'_t(\mathbf{S}; \emptyset))$ and gives the claimed bound on $\hat{\mathcal{A}}'_t(\mathbf{S})$.

Bounding the error of $\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}, \emptyset))$: We now give the proof sketch of Equation (1). To this end, for a training sequence S' and hypothesis h we define $\sum_{\neq}(h, S') = \sum_{(x,y) \in S'} \mathbb{1}\{h(x) \neq y\}$. Assume for the moment that for $h^* \in \arg \min \mathcal{L}_{\mathcal{D}}(h)$ we have demonstrated that with probability at least $1 - \delta/4$ over $\mathbf{S}$, and for some numerical constants c_b, c_c ,

$$\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}; \emptyset)) \leq \max_{S' \in S'(\mathbf{S}, \emptyset)} \frac{14\sum_{\neq}(h^*, S')}{m/s'_{\sqcap}} + \sqrt{\frac{c_b(d + \ln(4/\delta)) \frac{14\sum_{\neq}(h^*, S')}{(m/s'_{\sqcap})}}{m}} + \frac{c_c(d + \ln(4/\delta))}{m}. \quad (2)$$

In order to exploit Equation (2), we first observe that for each $i \in \{1, 2, 3\}$, we have that

$$\max_{S' \in S'(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup \emptyset)} \frac{14\Sigma_{\neq}(h^*, S')}{m/s'_{\sqcap}} \leq \frac{14|\mathbf{S}_i| \Sigma_{\neq}(h^*, \mathbf{S}_i)}{|\mathbf{S}_i| m/s'_{\sqcap}} \leq 15 \mathcal{L}_{\mathbf{S}_i}(h^*), \quad (3)$$

as $S' \sqsubseteq \mathbf{S}_i$ and $|\mathbf{S}_i| s'_{\sqcap}/m = |\mathbf{S}_i|/|\mathbf{S}_{i,\sqcap}| = 1/(1 - \frac{1}{27})$. We briefly remark that the multiplication with $|\mathbf{S}_i|/|\mathbf{S}_{i,\sqcap}|$ is a source of a multiplicative factor on τ . However, it can be made arbitrarily close to 1 by making the split between $\mathbf{S}_{i,\sqcup}$ and $\mathbf{S}_{i,\sqcap}$ more imbalanced in the direction of $\mathbf{S}_{i,\sqcap}$, at the cost of larger constants c_b, c_c . Now, by an application of Bernstein's inequality on the hypothesis h^* (Lemma B.3) for each $i \in \{1, 2, 3\}$, we have that $15 \mathcal{L}_{\mathbf{S}_i}(h^*)$ is at most $15(\tau + O(\sqrt{\tau \ln(1/\delta)/m} + \ln(1/\delta)/m))$, with probability at least $1 - \delta/4$ over $\mathbf{S}_i$. Invoking the union bound and using that $\max_{S' \in S'(\mathbf{S}_i)} = \max_{i \in \{1, 2, 3\}} \max_{S' \in S'(\mathbf{S}_{i,\sqcup}, \mathbf{S}_{i,\sqcap} \sqcup \emptyset)}$, we then have that $\max_{S' \in S'(\mathbf{S}, \emptyset)} 14\Sigma_{\neq}(h^*, S')/(m/s'_{\sqcap}) = 15(\tau + O(\sqrt{\tau \ln(1/\delta)/m} + \ln(1/\delta)/m))$ with probability at least $1 - 3\delta/4$ over $\mathbf{S}$. Finally, union bounding with the event in Equation (2) yields that both events hold with probability at least $1 - \delta$ over $\mathbf{S}$, from which the error bound of Equation (1) follows by inserting the bound of $\max_{S' \in S'(\mathbf{S}, \emptyset)} 14\Sigma_{\neq}(h^*, S')/(m/s'_{\sqcap})$ into Equation (2).

Relating the error of $\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}, \emptyset))$ to the empirical error of h^* : Here we draw inspiration from the seminal ideas of Hanneke (2016) by analyzing the learner's loss recursively. However, certain aspects of our analysis will diverge from Hanneke's approach, which is tailored to the realizable setting. As in Hanneke (2016), we have the splitting scheme $S'(\cdot, \cdot)$ receive two arguments, the second of which can be thought of as the concatenation of all previous training sequences created by the recursive calls of Algorithm 1. As we are in the agnostic case, this second argument can be any training sequence $T \in (\mathcal{X} \times \{-1, 1\})^*$. We will first demonstrate that with probability at least $1 - \delta$ over $\mathbf{S}$,

$$\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}; T)) \leq \max_{S' \in S'(\mathbf{S}; T)} \frac{14\Sigma_{\neq}(h^*, S')}{m/s'_{\sqcap}} + \sqrt{\frac{c_b(d + \ln(1/\delta)) \frac{14\Sigma_{\neq}(h^*, S')}{(m/s'_{\sqcap})}}{m}} + \frac{c_c(d + \ln(1/\delta))}{m}. \quad (4)$$

Setting $T = \emptyset$ and rescaling δ then yields Equation (2).

The first step of our analysis is to relate the error of $\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}; T))$ to that of the previous calls $\mathcal{A}'(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup T)$. First note that for any (x, y) such that $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}; T))(x, y) \geq 49/100$, at least one of the calls $\mathcal{A}'(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup T)$ for $i \in \{1, 2, 3\}$ must have $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup T))(x, y) \geq 49/100$. Further, there must be at least $(\frac{49}{100} - \frac{1}{3}) \frac{3}{2} = \frac{47}{200}$ of the hypotheses in $\cup_{j \in \{1, 2, 3\} \setminus i} \mathcal{A}'(\mathbf{S}_{j,\sqcup}; \mathbf{S}_{j,\sqcap} \sqcup T)$ that also fail on (x, y) . Using these observations, we have that if we draw a random index $\mathbf{I} \in \{1, 2, 3\}$ and a random hypothesis $\hat{\mathbf{h}} \in \cup_{j \in \{1, 2, 3\} \setminus \mathbf{I}} \mathcal{A}'(\mathbf{S}_{j,\sqcup}; \mathbf{S}_{j,\sqcap} \sqcup T)$, then for (x, y) such that $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}; T))(x, y) \geq 49/100$ it holds that $\mathbb{P}_{\mathbf{I}, \hat{\mathbf{h}}}[\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{\mathbf{I},\sqcup}; \mathbf{S}_{\mathbf{I},\sqcap} \sqcup T))(x, y) \geq 49/100, \hat{\mathbf{h}}(x) \neq y] \geq \frac{1}{3} \frac{47}{200} \geq \frac{1}{13}$. Hence, we have that $13 \mathbb{E}_{\mathbf{I}, \hat{\mathbf{h}}}[\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{\mathbf{I},\sqcup}; \mathbf{S}_{\mathbf{I},\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq 49/100, \hat{\mathbf{h}}(\mathbf{x}) \neq \mathbf{y}]] \geq \mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}; T))$. Furthermore, by the definition of $\mathbf{I} \in \{1, 2, 3\}$ and $\hat{\mathbf{h}}$ being a random hypothesis from $\cup_{j \in \{1, 2, 3\} \setminus \mathbf{I}} \mathcal{A}'(\mathbf{S}_{j,\sqcup}; \mathbf{S}_{j,\sqcap} \sqcup T)$, we conclude that

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}; T)) &\leq 13 \mathbb{E}_{\mathbf{I}, \hat{\mathbf{h}}} \left[\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{\mathbf{I},\sqcup}; \mathbf{S}_{\mathbf{I},\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq 49/100, \hat{\mathbf{h}}(\mathbf{x}) \neq \mathbf{y} \right] \right] \\ &\leq 13 \max_{\substack{i \neq j \in \{1, 2, 3\}, \\ h \in \mathcal{A}'(\mathbf{S}_{j,\sqcup}; \mathbf{S}_{j,\sqcap} \sqcup T)}} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq 49/100, h(\mathbf{x}) \neq \mathbf{y} \right]. \end{aligned} \quad (5)$$

We will demonstrate Equation (4) by induction on $|\mathbf{S}| = 3^k$. By setting c_b and c_c sufficiently large, we can assume the claim holds for $k < 9$. By Equation (5) and a union bound, it suffices to bound each of the six terms for different combinations of i, j by Equation (4) with probability at least $1 - \delta/6$. Using symmetry, we can consider the case $j = 1$ and $i = 2$. Note that $13 \max_{h \in \mathcal{A}'(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq 49/100, h(\mathbf{x}) \neq \mathbf{y}]$ equals

$$\max_{\substack{h \in \mathcal{A}'(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)}} \left\{ 13 \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [h(\mathbf{x}) \neq \mathbf{y} \mid \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq 0.49] \mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T)) \right\}. \quad (6)$$

Supposing that $\mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))$ is upper bound by

$$\frac{1}{13} \left(\max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{14\Sigma_{\neq}(h^*, S')}{m/s'_{\sqcap}} + \sqrt{\frac{c_b(d + \ln(1/\delta)) \frac{14\Sigma_{\neq}(h^*, S')}{(m/s'_{\sqcap})}}{m}} + \frac{c_c(d + \ln(1/\delta))}{m} \right), \quad (7)$$

then Equation (6) is bounded as in Equation (4) and we are done. Using that $|\mathbf{S}_{2,\sqcup}| = m/3^4$, the inductive hypothesis gives that with probability at least $1 - \delta/16$ over $\mathbf{S}_2$ that $\mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))$ is at most

$$\max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{14\Sigma_{\neq}(h^*, S')}{m/(3^4 s'_{\sqcap})} + \sqrt{\frac{c_b(d + \ln(16/\delta)) \frac{14\Sigma_{\neq}(h^*, S')}{(m/s'_{\sqcap})}}{m/3^4}} + \frac{c_c(d + \ln(16/\delta))}{m/3^4}. \quad (8)$$

Where we have used that $\mathcal{S}(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T) \subseteq \mathcal{S}(\mathbf{S}; T)$. Thus, it suffices to consider the case in which $\mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))$ lies between the expressions in Equation (7) and Equation (8).

Let $A = \{(x, y) | \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))(x, y) \geq 49/100\}$ and $\mathbf{N} = |\mathbf{S}_{1,\sqcap} \cap A|$. Using a Chernoff bound, one has that $\mathbf{N} \geq 13m/(14s'_{\sqcap}) \mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))$ with probability at least $1 - \delta/16$. As $\mathbf{S}_{1,\sqcap} \cap A \sim \mathcal{D}(\cdot | \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq 49/100)$, uniform convergence over $\mathcal{H}$ (see Lemma B.6) yields that with probability $1 - \delta/16$ over $\mathbf{S}_{1,\sqcap} \cap A$, Equation (6) is bounded by

$$13 \max_{h \in \mathcal{A}'(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \left(\mathcal{L}_{\mathbf{S}_{1,\sqcap} \cap A}(h) + \sqrt{\frac{C(d + 2 \ln(48/\delta))}{\mathbf{N}}} \right) \mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T)). \quad (9)$$

Where $C > 0$ is a universal constant. We now bound each term in Equation (9), considered after multiplying by $\mathcal{L}_D^{0.49}$. Notably, this component of the analysis diverges considerably from that of Hanneke (2016) for the realizable case, in which one is assured that $\mathcal{L}_{\mathbf{S}_{1,\sqcap} \cap A}(h) = 0$. For the first, we use the definitions of $\mathbf{N}$ and $\mathcal{L}_{\mathbf{S}_{1,\sqcap} \cap A}$ (and abbreviate $h \in \mathcal{A}'(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)$ as $h \in \mathcal{A}'$) to observe that

$$13 \max_{h \in \mathcal{A}'} \mathcal{L}_{\mathbf{S}_{1,\sqcap} \cap A}(h) \mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T)) \leq 14 \max_{h \in \mathcal{A}'} \sum_{(x,y) \in \mathbf{S}_{1,\sqcap} \cap A} \frac{\mathbb{1}\{h(x) \neq y\}}{m/s'_{\sqcap}}. \quad (10)$$

Further, using that for any $S' \in \mathcal{S}'(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)$ we have $\mathbf{S}_{1,\sqcap} \cap A \subseteq S'$, and $h = \mathcal{A}(S')$ is a minimizer of $\sum_{(x,y) \in S'} \mathbb{1}\{h(x) \neq y\}$, combined with that $\mathcal{S}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T) \subseteq \mathcal{S}(\mathbf{S}; T)$, one can invoke Equation (10) to conclude that

$$13 \max_{h \in \mathcal{A}'(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \mathcal{L}_{\mathbf{S}_{1,\sqcap} \cap A}(h) \mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T)) \leq 14 \max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{\Sigma_{\neq}(h^*, S')}{m/s'_{\sqcap}}. \quad (11)$$

We note that the above step introduces a factor of $\frac{14}{13}$ on τ , which can be brought arbitrarily close to 1 by using a tighter Chernoff bound for the size of $\mathbf{N}$, at the cost of larger constants c_c and c_b in Equation (4). For the second term $\sqrt{C(d + 2 \ln(48/\delta))/\mathbf{N}} \mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))$ of Equation (9), we again use that $\mathbf{N} \geq (13m)/(14s'_{\sqcap}) \mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))$ to bound it from above by $\sqrt{(14C(d + 2 \ln(48/\delta)) \mathcal{L}_D^{0.49}(\mathcal{A}'(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T)))/(13m/s'_{\sqcap})}$. As we are considering the case in which $\mathcal{L}_D^{0.49}$ is bounded by the expression of Equation (8), one can show that the following upper bound holds for sufficiently large c_b and c_c :

$$\sqrt{\frac{c_b(d + \ln(1/\delta)) \frac{14\Sigma_{\neq}(h^*, S')}{(m/s'_{\sqcap})}}{m}} + \frac{c_c(d + \ln(1/\delta))}{m}.$$

Combining these bounds on each term of Equation (9) yields an upper bound of the form

$$\max_{S' \in \mathcal{S}'(\mathbf{S}; T)} \frac{14\Sigma_{\neq}(h^*, S')}{m/s'_{\sqcap}} + \sqrt{\frac{c_b(d + \ln(1/\delta)) \frac{14\Sigma_{\neq}(h^*, S')}{(m/s'_{\sqcap})}}{m}} + \frac{c_c(d + \ln(1/\delta))}{m}, \quad (12)$$

completing the inductive step and establishing Equation (4), as desired.

3.2 Improved Approach: Multiplicative Constant 2.1

Roughly speaking, we apply two modifications to the previous approach in order to achieve a considerably smaller constant factor on τ . First, we use a different splitting scheme (Algorithm 2) which recursively splits the dataset into 27 subsamples, rather than 3. Second, rather than taking a majority vote over ERM learners on a single instance of the splitting algorithm, we split the training set into three parts, run two independent instances of the voting classifier on the first two parts, and one instance of ERM on a subsample of the third part. The subsample we train the ERM on is carefully chosen and depends on a certain “region of disagreement” of the two voting classifiers. The final prediction at a given datapoint is as follows: If both voting classifiers agree on the predicted label with a certain notion of “margin,” then we output that label, otherwise we output the prediction of the tiebreaker ERM. Lastly, we combine our approach with that of Hanneke et al. (2024) to achieve a best-of-both-worlds bound. We summarize our approach in the following point form and Figure 2.

1. Split the training set S into three equally-sized parts S_1, S_2, S_3 .
2. Split S_1 into three equally-sized parts, $S_{1,1}, S_{1,2}, S_{1,3}$.
 - 2.1 Run the splitting scheme of Algorithm 2 on $(S_{1,1}, \emptyset)$ and $(S_{1,2}, \emptyset)$. Let S_1, S_2 be the resulting sets of training subsequences.
 - 2.2 Sample $t = O(\ln(m/(\delta(d + \ln(1/\delta)))))$ sequences from each of S_1 and S_2 uniformly at random. Let $\hat{S}_1, \hat{S}_2$ be the resulting collections of sequences.
 - 2.3 Train ERM learners on each sequence appearing in $\hat{S}_1, \hat{S}_2$. Denote by $\hat{A}_{t_1}, \hat{A}_{t_2}$ the resulting collections of classifiers produced by these ERMs, respectively.
 - 2.4 Define the set $S_3^\neq$ as follows: for any $(x, y) \in S_{1,3}$ if at least an $11/243$ -fraction of the classifiers in $\hat{A}_{t_1}$ do not predict the label y for x , or likewise for $\hat{A}_{t_2}$, then $(x, y) \in S_3^\neq$. Train an ERM on $S_3^\neq$, producing the classifier h_{tie} .
 - 2.5 Let $\tilde{A}_1$ be the classifier which acts as follows on any given x : if a $232/243$ -fraction of the classifiers of both $\hat{A}_{t_1}$ and $\hat{A}_{t_2}$ agree on a label y for x , then $\tilde{A}_1$ predicts y . Otherwise, it predicts $h_{\text{tie}}(x)$.
3. Use S_2 to train the algorithm of Hanneke et al. (2024). Let $\tilde{A}_2$ be the resulting classifier.
4. Output the classifier among $\tilde{A}_1$ and $\tilde{A}_2$ which attains superior performance on S_3 .




Figure 2: A flowchart of the final algorithm. The initial sample S is split into three parts. S_1 is used to construct our tie-breaking classifier $\tilde{A}_1$. S_2 is used to train the algorithm of Hanneke et al. (2024), yielding $\tilde{A}_2$. Finally, S_3 is used as a hold-out set to select the better of the two classifiers.

Remark 3.1. Intuitively, h_{tie} can be seen as “stabilizing” the predictions produced by the classifiers in the collections $\hat{A}_{t_1}$ and $\hat{A}_{t_2}$ when at least one collection is judged to have low confidence. However, strictly speaking, we cannot demonstrate that the learner $\tilde{A}_1$ described by Step (2.) is stable in the sense of Bousquet and Elisseeff (2002). In particular, the classifiers in each of $\hat{A}_{t_1}$ and $\hat{A}_{t_2}$ are trained on overlapping samples, meaning that many classifiers could change by altering even one example in $S_{1,1}$ or $S_{1,2}$, thereby altering $S_3^\neq$ as well.

We now derive the generalization error for the steps under points 2.1 – 2.5. The derivation is presented in a bottom-up fashion, starting by examining the failure modes of the approach in Section 3.1 (which leads to the bound with 15τ), and amending these to obtain the approach of 2.1 – 2.5 bounded by 2.1τ . To this end, let us recall the three sources leading to multiplicative factors on τ in Section 3.1. First is

the balance between the splits of S_i in Algorithm 1, i.e. $|S_i|/|S_{i,\square}| = 1/(1 - 1/27)$ in Equation (3). Recall that this constant can be driven down to 1 by considering even more imbalanced splits of the dataset, at the expense of larger constants multiplying the remaining terms of the bound. Second is the error arising from Equation (11), which controls the size of N via a Chernoff bound. This can similarly be driven arbitrarily close to 1. Hence, the primary multiplicative overhead results from the third source: the argument in Equation (5) relating the error of $\mathcal{A}'(\mathbf{S}, T)$ to that of its recursive calls.

The constant arises from relating $\mathcal{L}_{\mathcal{D}}^{0.49}(\mathcal{A}'(\mathbf{S}; T))$ to one of its previous iterates erring on a $49/100$ -fraction of its classifiers, and one hypothesis from the remaining iterates also erring. Recall that to get this bound we first observe that, with probability at least $1/3$ over a randomly drawn index $\mathbf{I} \sim \{1, 2, 3\}$, $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{\mathbf{I},\square}; \mathbf{S}_{\mathbf{I},\square} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq 0.49$ and a random hypothesis $\hat{\mathbf{h}} \in \sqcup_{j \in \{1, 2, 3\} \setminus \mathbf{I}} \mathcal{A}'(\mathbf{S}_{j,\square}; \mathbf{S}_{j,\square} \sqcup T)$, would, with probability at least $3/2(49/100 - 1/3)$, make an error. Thus, we get the inequality in Equation (5), which multiplies τ by ≈ 13 . To decrease this constant we first make the following observation: assume that we have two voting classifiers $\mathcal{A}'(\mathbf{S}_1; \emptyset)$, $\mathcal{A}'(\mathbf{S}_2; \emptyset)$ and an (x, y) such that $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(x, y) \geq 232/243$ and $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(x, y) \geq 232/243$. Using a similar analysis as in Equation (5), we have

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 232/243, \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 232/243] \quad (13)$$

$$\leq \frac{243}{232} \max_{h \in \mathcal{A}'(\mathbf{S}_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [h(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 232/243]. \quad (14)$$

Following a similar approach to that used to obtain Equation (1), we can bound Equation (14) by $1.1\tau + O(\sqrt{\tau(d + \ln(1/\delta))/m} + (d + \ln(1/\delta))/m)$. Thus, if we could bound the error using the event in Equation (13), we could obtain a better bound for the overall error of our algorithm.

To this end, consider a third independent training sequence $\mathbf{S}_3$ and let $\mathbf{S}_3^{\neq}$ be the training examples (x, y) in $\mathbf{S}_3$ such that $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(x, y) \geq 11/243$ or $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(x, y) \geq 11/243$. Let $h_{\text{tie}} = \mathcal{A}(\mathbf{S}_3^{\neq})$, i.e. h_{tie} is the output of ERM-learner trained on $\mathbf{S}_3^{\neq}$. We now introduce our tie-breaking idea. For a point x we let $\text{Tie}^{11/243}(\mathcal{A}'(\mathbf{S}_1; \emptyset), \mathcal{A}'(\mathbf{S}_2; \emptyset); h_{\text{tie}})(x) = y$, where $y \in \{\pm 1\}$ is the unique number for which $\sum_{h \in \mathcal{A}'(\mathbf{S}_1; \emptyset)} \mathbb{1}\{h(x) = y\} / |\mathcal{A}'(\mathbf{S}_1; \emptyset)| \geq 232/243$ and $\sum_{h \in \mathcal{A}'(\mathbf{S}_2; \emptyset)} \mathbb{1}\{h(x) = y\} / |\mathcal{A}'(\mathbf{S}_2; \emptyset)| \geq 232/243$. If such a number does not exist, we set $y = h_{\text{tie}}(x)$. Thus, this classifier errs on (x, y) if 1) $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(x, y) \geq 232/243$ and $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(x, y) \geq 232/243$, or 2) for all $y' \in \{\pm 1\}$, $\sum_{h \in \mathcal{A}'(\mathbf{S}_1; \emptyset)} \mathbb{1}\{h(x) = y'\} / |\mathcal{A}'(\mathbf{S}_1; \emptyset)| < 232/243$ or $\sum_{h \in \mathcal{A}'(\mathbf{S}_2; \emptyset)} \mathbb{1}\{h(x) = y'\} / |\mathcal{A}'(\mathbf{S}_2; \emptyset)| < 232/243$ and $h_{\text{tie}}(x) \neq y$. Using the definition of condition 2), we have that

$$\begin{aligned} & \mathcal{L}_{\mathcal{D}}(\text{Tie}^{11/243}(\mathcal{A}'(\mathbf{S}_1; \emptyset), \mathcal{A}'(\mathbf{S}_2; \emptyset); h_{\text{tie}})) \\ & \leq \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 232/243, \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 232/243] \\ & + \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ or } \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ and } h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}]. \end{aligned}$$

As observed below Equation (13), the first term can be bounded by $1.1\tau + O(\sqrt{\tau(d + \ln(1/\delta))/m} + (d + \ln(1/\delta))/m)$. Thus, if we could bound the second term by $\tau + O(\sqrt{\tau(d + \ln(1/\delta))/m} + (d + \ln(1/\delta))/m)$, we would be done. Proceeding as in Equation (6), let $\mathcal{D}_{\neq}$ be the conditional distribution given $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243$ or $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243$. Then,

$$\begin{aligned} & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ or } \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ and } h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}] = \\ & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_{\neq}} [h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}] \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ or } \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243]. \end{aligned}$$

We notice that h_{tie} is the output of an ERM learner trained on $\mathbf{S}_3^{\neq} \sim \mathcal{D}_{\neq}$, thus, by standard guarantees on ERM learners (Theorem C.2), it holds with high probability that

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_{\neq}} [h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}] \leq \inf_{h \in \mathcal{H}} \mathcal{L}_{\mathcal{D}_{\neq}}(h) + \sqrt{\frac{d + \ln(1/\delta)}{|\mathbf{S}_3^{\neq}|}} \leq \mathcal{L}_{\mathcal{D}_{\neq}}(h^*) + \sqrt{\frac{d + \ln(1/\delta)}{|\mathbf{S}_3^{\neq}|}},$$

where the second inequality follows from that $h^* \in \mathcal{H}$. Thus, we have that

$$\begin{aligned} & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ or } \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ and } h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}] \leq \\ & \tau + \sqrt{\frac{(d + \ln(1/\delta)) \mathbb{P}_{\mathbf{x}, \mathbf{y} \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ or } \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243]^2}{|\mathbf{S}_3^{\neq}|}} \end{aligned}$$

Using a similar argument as in Equation (9), a Chernoff bound yields that $|\mathbf{S}_3^\neq| = \Theta(m \mathbb{P}_{\mathbf{x}, \mathbf{y} \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_1; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ or } \text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_2; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243])$ with high probability. Thus, if the probability term were at most $c(\tau + (d + \ln(1/\delta))/m)$, the above argument would be complete. However, the analysis from Section 3.1 only bounds the probability term when it is $49/100$, instead of $11/243$, which is not sufficient. Thus, we will now remedy this using a different splitting algorithm that allows for showing the more fine-grained error is small.

An approach with more than 3 splits: We have seen that the above argument requires that $\mathbb{P}[\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}; \emptyset))(\mathbf{x}, \mathbf{y}) \geq 11/243] \leq c(\tau + (d + \ln(1/\delta))/m)$ with high probability. However, attempting to prove this by induction, as in Section 3.1, breaks down in the step at which we derived Equation (5). Recall that this is where we relate the condition $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}; T))(x, y) \geq 11/243$ to a previous recursive call also erring on an $11/243$ -fraction of its voters and one of the hypotheses in the remaining recursive calls erring on (x, y) . To see why this argument fails, now consider (x, y) such that $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}; T))(x, y) \geq 11/243$. Picking an index $\mathbf{I} \sim \{1, 2, 3\}$ with probability at least $1/3$ still returns $\text{avg}_{\neq}(\mathcal{A}'(\mathbf{S}_{\mathbf{I}, \sqcup}; T \sqcup \mathbf{S}_{\mathbf{I}, \sqcap}))(x, y) \geq 11/243$. However, when picking a random hypothesis $\hat{\mathbf{h}} \in \sqcup_{j \in \{1, 2, 3\} \setminus \mathbf{I}} \mathcal{A}'(\mathbf{S}_{1, j, \sqcup}; \mathbf{S}_{j, \sqcap} \sqcup T)$, the probability of $\hat{\mathbf{h}}$ erring can only be lower bounded by $3/2(11/243 - 1/3)$, which is negative! (all the errors might be in the recursive call $\mathbf{I}$). Thus, we cannot guarantee a lower bound on this probability when making only 3 recursive calls in Algorithm 1. However, if we make more recursive calls, e.g., 27 (chosen large enough to make the argument work), we get that with probability at least $1/27$ over $\mathbf{I} \sim \{1, \dots, 27\}$, it holds that $\mathcal{A}(\mathbf{S}_{\mathbf{I}, \sqcup}; \mathbf{S}_{\mathbf{I}, \sqcap} \sqcup T) \geq 11/243$ and with probability at least $27/26(11/243 - 1/27) \approx 0.009$ over $\hat{\mathbf{h}} \in \sqcup_{j \in \{1, \dots, 27\} \setminus \mathbf{I}} \mathcal{A}(\mathbf{S}_{1, j, \sqcup}; \mathbf{S}_{j, \sqcap} \sqcup T)$, we have that $\hat{\mathbf{h}}(x) \neq y$. By $27 \cdot (1/0.009) \leq 3160$, this gives

$$\mathcal{L}_{\mathcal{D}}^{\frac{11}{243}}(\mathcal{A}(\mathbf{S}; T)) \leq 3160 \max_{\substack{i, j \in \{1, \dots, 27\}, i \neq j \\ h' \in \mathcal{A}(\mathbf{S}_{j, \sqcup}; \mathbf{S}_{j, \sqcap} \sqcup T)}} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h'(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{i, \sqcup}; \mathbf{S}_{i, \sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right].$$

This is precisely the scheme we propose in Algorithm 2, with $\mathcal{A}(\mathbf{S}; T) = \{\mathcal{A}(S')\}_{S' \in \mathcal{S}(\mathbf{S}; T)}$. Now, defining $\hat{\mathcal{A}}_t(\mathbf{S})$ as t voters drawn from $\mathcal{A}(\mathbf{S}; \emptyset)$ with $t = \Theta(\ln(m/(\delta(d + \ln(1/\delta)))))$, and mimicking the analysis of Section 3.1 yields that $\mathcal{L}_{\mathcal{D}}^{\frac{11}{243}}(\hat{\mathcal{A}}_t) \leq c(\tau + (d + \ln(1/\delta))/m)$. Finally, roughly following the above arguments for $\mathcal{L}_{\mathcal{D}}(\text{Tie}^{11/243}(\mathcal{A}'_{t_1}(\mathbf{S}_1), \mathcal{A}'_{t_2}(\mathbf{S}_2); h_{\text{tie}}))$, but now with $\mathcal{L}_{\mathcal{D}}(\text{Tie}^{11/243}(\hat{\mathcal{A}}_{t_1}(\mathbf{S}_1), \hat{\mathcal{A}}_{t_2}(\mathbf{S}_2); h_{\text{tie}}))$, we obtain, for the latter, the claimed generalization error of $2.1\tau + O(\sqrt{\tau(d + \ln(1/\delta))/m} + (d + \ln(1/\delta))/m)$.

Algorithm 2: Splitting algorithm $\mathcal{S}$

Input: Training sequences $S, T \in (\mathcal{X} \times \mathcal{Y})^*$, where $|S| = 3^k$ for $k \in \mathbb{N}$.

Output: Family of training sequences.

if $k \geq 6$ **then**

 Partition S into $S_1, \dots, S_{27}$, with S_i being the $(i-1)|S|/27 + 1$ to the $i|S|/27$ training examples of S . Set for each i

$$S_{i, \sqcup} = S_i[1 : 3^{k-6}], \quad S_{i, \sqcap} = S_i[3^{k-6} + 1 : 3^{k-3}],$$

return $[\mathcal{S}(S_{1, \sqcup}; S_{1, \sqcap} \sqcup T), \dots, \mathcal{S}(S_{27, \sqcup}; S_{27, \sqcap} \sqcup T)]$

else

return $S \sqcup T$

With the above classifier in hand, we can now, on two new independent training sequences, obtain the classifier of Hanneke et al. (2024) using the first sequences, and then on the second sequences choose the best of our classifier and the classifier of Hanneke et al. (2024) as the final classifier.

4 Conclusion

We study the fundamental problem of agnostic PAC learning and provide improved sample complexity bounds parametrized by τ , the error of the best-in-class hypothesis. Our results resolve the question of Hanneke et al. (2024) asking for optimal error rates in the regime $\tau \approx d/m$, and make progress on their questions regarding optimal learners for the full range of τ and efficient learners based upon majority votes of ERM. The most interesting future direction is whether an improved analysis of our voting scheme or a modification of it can lead to optimal algorithms for the full range of τ .

Acknowledgments and Disclosure of Funding

While this work was carried out, Mikael Møller Høsgaard was supported by DFF Grant No. 9064-00068B, and the European Union (ERC, TUCLA, 101125203). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Julian Asilis was supported by the NSF CAREER Award CCF-223926 and the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1842487. Part of the work was done while Grigoris Velegkas was a PhD student at Yale University supported in part by the AI Institute for Learning-Enabled Optimization at Scale (TILOS).

References

- Ishaq Aden-Ali, Yeshwanth Cherapanamjeri, Abhishek Shetty, and Nikita Zhivotovskiy. Optimal pac bounds without uniform convergence. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1203–1223. IEEE, 2023. 3
- Ishaq Aden-Ali, Mikael Møller Høandgsaard, Kasper Green Larsen, and Nikita Zhivotovskiy. Majority-of-three: The simplest optimal learner? In *The Thirty Seventh Annual Conference on Learning Theory*, pages 22–45. PMLR, 2024. 3
- Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. cambridge university press, 2009. 3
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Learnability and the vapnik-chervonenkis dimension. *Journal of the ACM (JACM)*, 36(4):929–965, 1989. 3
- Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: A survey of some recent advances. *ESAIM: probability and statistics*, 9:323–375, 2005. 3
- Olivier Bousquet and André Elisseeff. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002. 8
- Olivier Bousquet, Steve Hanneke, Shay Moran, and Nikita Zhivotovskiy. Proper learning, helly number, and an optimal svm bound. In *Conference on Learning Theory*, pages 582–609. PMLR, 2020. 13, 15
- Luc Devroye, László Györfi, and Gábor Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, 1996. 2
- Andrzej Ehrenfeucht, David Haussler, Michael Kearns, and Leslie Valiant. A general lower bound on the number of examples needed for learning. *Information and Computation*, 82(3):247–261, 1989. 3
- Steve Hanneke. The optimal sample complexity of pac learning. *Journal of Machine Learning Research*, 17(38):1–15, 2016. 2, 3, 4, 5, 6, 7
- Steve Hanneke, Kasper Green Larsen, and Nikita Zhivotovskiy. Revisiting agnostic pac learning. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1968–1982. IEEE, 2024. 2, 3, 8, 10, 34
- David Haussler. Decision theoretic generalizations of the pac model for neural net and other learning applications. *Information and computation*, 100(1):78–150, 1992. 3
- Mikael Møller Høsgaard. Efficient optimal pac learning. *arXiv preprint arXiv:2502.03620*, 2025. 3
- Kasper Green Larsen. Bagging is an optimal PAC learner. In Gergely Neu and Lorenzo Rosasco, editors, *The Thirty Sixth Annual Conference on Learning Theory, COLT 2023, 12-15 July 2023, Bangalore, India*, volume 195 of *Proceedings of Machine Learning Research*, pages 450–468. PMLR, 2023. URL <https://proceedings.mlr.press/v195/larsen23a.html>. 3, 5
- Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample-variance penalization. In *Annual Conference Computational Learning Theory*, 2009. URL <https://api.semanticscholar.org/CorpusID:17090214>. 3

Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014. [17](#), [20](#), [26](#)

Hans U Simon. An almost optimal pac algorithm. In *Conference on Learning Theory*, pages 1552–1563. PMLR, 2015. [2](#), [3](#)

Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984. [3](#)

Vladimir Vapnik and A Ya Chervonenkis. A class of algorithms for pattern recognition learning. *Avtomat. i Telemekh*, 25(6):937–945, 1964. [3](#)

Vladimir Vapnik and Alexey Chervonenkis. Theory of pattern recognition, 1974. [3](#)

Andrew J Wagenmaker, Yifang Chen, Max Simchowitz, Simon Du, and Kevin Jamieson. First-order regret in reinforcement learning with linear function approximation: A robust estimation approach. In *International Conference on Machine Learning*, pages 22384–22429. PMLR, 2022. [3](#)

A Preliminaries for Proof

In this section we give the preliminaries for the proof. For the reader’s convenience, we restate Algorithm [2](#).

Algorithm 2: Splitting algorithm $\mathcal{S}$

Input: Training sequences $S, T \in (\mathcal{X} \times \mathcal{Y})^*$, where $|S| = 3^k$ for $k \in \mathbb{N}$.

Output: Family of training sequences.

if $k \geq 6$ **then**

 Partition S into $S_1, \dots, S_{27}$, with S_i being the $(i-1)|S|/27 + 1$ to the $i|S|/27$ training examples of S . Set for each i

$$S_{i,\sqcup} = S_i[1 : 3^{k-6}], \quad S_{i,\sqcap} = S_i[3^{k-6} + 1 : 3^k],$$

return $\{\mathcal{S}(S_{1,\sqcup}; S_{1,\sqcap} \sqcup T), \dots, \mathcal{S}(S_{27,\sqcup}; S_{27,\sqcap} \sqcup T)\}$

else

return $S \sqcup T$

We first observe that for an input training sequence $m = |S| = 3^k$, the above algorithm makes l recursive calls where $l \in \mathbb{N}$ satisfies $k - 6(l-1) \geq 6$ and $k - 6l < 6$, that is, l is the largest number such that $k/6 \geq l$. As l is a natural number, we get that $l = \lfloor k/6 \rfloor$. Furthermore, since $k = \log_3(m)$ we get that $l = \lfloor \log_3(m)/6 \rfloor$. For each of the l recursive calls 27 recursions are made. Thus, the total number of training sequences created in $\mathcal{S}$ is $27^l \leq 3^{3 \log_3(m)/(2 \cdot 6)} = m^{1/(4 \ln(3))} \geq m^{0.22}$. In what follows, we will use the quantity $s_{\sqcap}$, which we define as $\frac{|S|}{|S_{1,\sqcap}|}$ when running $\mathcal{S}(S; T)$ with $S; T \in (\mathcal{X} \times \mathcal{Y})^*$ such that $|S| = 3^k$ and $k \geq 6$. We notice that

$$s_{\sqcap} := \frac{|S|}{|S_{1,\sqcap}|} = \frac{3^k}{3^{k-3} - 3^{k-6}} = \frac{3^k}{3^{k-3}(1 - \frac{1}{3^3})} = \frac{27}{1 - \frac{1}{27}}. \quad (15)$$

This ratio $s_{\sqcap}$ will later in the proof show up as a constant, $|S_i|/(|S|s_{\sqcap}) = (|S_{i,\sqcup}| + |S_{i,\sqcap}|)/|S_{i,\sqcap}| = 1/(1 - 1/27)$, multiplied onto τ . Thus, from this relation we see that if the split of S_i is imbalanced so that $|S_{i,\sqcap}|$ is larger than $|S_{i,\sqcup}|$ the constant multiplied on to τ become smaller.

Furthermore, in what follows, for the set of training sequences generated by $\mathcal{S}(S; T)$ and a ERM-algorithm $\mathcal{A}$, we write $\mathcal{A}(S; T)$ for the set of classifiers the ERM-algorithm outputs when run on the training sequences in $\mathcal{S}(S; T)$, i.e., $\mathcal{A}(S; T) = \{\mathcal{A}(S')\}_{S' \in \mathcal{S}(S; T)}$, where this is understood as a multiset if the output of the ERM-algorithm is the same for different training sequences in $\mathcal{S}(S; T)$. Furthermore for an example (x, y) , we define $\text{avg}_{\neq}(\mathcal{A}(S; T))(x, y) = \sum_{h \in \mathcal{A}(S; T)} \mathbb{1}\{h(x) \neq y\}/|\mathcal{A}(S; T)|$, i.e., the average number of incorrect hypotheses in $\mathcal{A}(S; T)$. We notice that by the above comment about $\mathcal{S}(S; T)$ having size at least $m^{0.22}$, we have that $\mathcal{A}(S; T)$ contains just as many hypothesis, each of which is the output of an ERM run on a training sequence of $\mathcal{S}(S; T)$. Thus as

allotted to earlier our algorithm do not run on all the sub training sequences created by $\mathcal{S}(S; T)$, as it calls the $\mathcal{A}$ algorithm $O(\ln(m/(\delta(d + \ln(1/\delta)))))$ -times. Which leads us considering the following classifier.

For a natural number t , we let $\hat{\mathcal{A}}_t(S; T)$, be the random collection of t hypotheses drawn uniformly with replacement from $\mathcal{A}(S; T)$, with the draws being independent, where we see $\hat{\mathcal{A}}_t(S; T)$ as a multiset so allowing for repetitions. We remark here that we will overload notation and use $\mathbf{t}$ (so t in bold font) to denote the randomness used to draw the t hypotheses from $\mathcal{A}(S; T)$ in the following analysis of $\hat{\mathcal{A}}_t(S; T)$. Intuitively one can think of $\hat{\mathcal{A}}_t$ as a bagging algorithm where the subsampled training sequences are restricted to subsets of $\mathcal{S}(S; T)$ rather than sampling with replacement from the training examples of S and T . In what follows we will consider this algorithm parametrized by $t = O(\ln(m/(\delta(d + \ln(1/\delta)))))$ leading to a classifier with the same order of call to the ERM as stated in. Similarly to $\mathcal{A}(S; T)$ we also define $\text{avg}_{\neq}(\hat{\mathcal{A}}_t(S; T))(x, y) = \sum_{h \in \hat{\mathcal{A}}_t(S; T)} \mathbb{1}\{h(x) \neq y\} / |\hat{\mathcal{A}}_t(S; T)|$

Now, for a distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, 1\}$, training sequences $S; T \in (\mathcal{X} \times \{-1, 1\})^*$, and $\alpha \in [0, 1]$ we will use $\mathcal{L}_{\mathcal{D}}^{\alpha}(\mathcal{A}(S; T)) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}(S; T))(\mathbf{x}, \mathbf{y}) \geq \alpha]$, i.e., the probability of at least a α -fraction of the hypotheses in $\mathcal{A}(S; T)$ erroring on a new example drawn according to $\mathcal{D}$. As above we also define $\mathcal{L}_{\mathcal{D}}^{\alpha}(\hat{\mathcal{A}}_t(S; T)) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\hat{\mathcal{A}}_t(S; T))(\mathbf{x}, \mathbf{y}) \geq \alpha]$, for $\hat{\mathcal{A}}_t(S; T)$. In the following we will for the case where T is the empty training sequence $\emptyset$ us $\hat{\mathcal{A}}_t(S) = \hat{\mathcal{A}}_t(S; \emptyset)$.

A.1 Difficulties of the Proof

Let us take the opportunity to briefly describe the challenges associated with settling the sample complexity of agnostic learning in the regime $\tau = O(d/m)$. In particular, we will demonstrate that a straightforward multiplicative Chernoff argument can *not* yield a result as general as Theorem 1.4.

To begin, notice that the true errors of N hypotheses can indeed be confidently estimated to within error τ using $O(\log(N/\delta)/\tau)$ samples, due to multiplicative Chernoff. However, if $\tau = O(\log(N/\delta)/m)$, then the above yields a vacuous bound. An additive Chernoff bound would also in this case imply an additive $O(\log(N/\delta)/m)$ error term. In the paper, we consider classes of finite VC dimension d but arbitrary, possibly infinite cardinality N , thereby preventing us from employing this bound. Even if one were to bound N as $O((m/d)^d)$ using Sauer-Shelah, this would imply that multiplicative Chernoff is unhelpful when $\tau = O((d \log(m/d) + \log(1/\delta))/m)$. This is a wide regime containing, for instance, $\tau = O(d/m)$, a setting in which we demonstrate that our algorithm is optimal. Furthermore, an additive Chernoff bound would incur a suboptimal additive $O(d \log(m/d)/m)$ term.

However, let us assume that a Chernoff bound argument for infinite hypothesis classes could yield Theorem B.2 for all τ (e.g., using an ϵ -covering argument). In this case, one would have demonstrated that an empirical risk minimizer h' has error $2.1\tau + O(\sqrt{\frac{\tau(d + \ln(1/\delta))}{m}}) + \frac{d + \ln(1/\delta)}{m}$. Then, taking $\tau = 0$ (i.e., considering the realizable case), would demonstrate that ERM results in an error of $O(\frac{d + \ln(1/\delta)}{m})$. However, this is not generally true, owing to the worst-case lower bound on ERM learners of $O(\frac{d \ln(1/\epsilon) + \ln(1/\delta)}{m})$, due to (Bousquet et al., 2020, Theorem 11). This strongly suggests that a Chernoff bound argument using a union bound over each function in the hypothesis class (or a discretization of it) cannot yield a result as general as Theorem 1.1, which holds across all τ .

Finally, we remark that our learner for Theorem B.2 is agnostic to τ (i.e., is oblivious to the value of τ), whereas the previous Chernoff procedure requires knowledge of τ . This is a notable distinction, as precise knowledge of τ will typically be unavailable to the learner. In particular, ERM-based estimates of τ will be uninformative in the regime $\tau = o(d/m)$.

A Preliminaries for Proof

In this section we give the preliminaries for the proof. For the reader's convenience, we restate Algorithm 2.

Algorithm 2: Splitting algorithm $\mathcal{S}$

Input: Training sequences $S, T \in (\mathcal{X} \times \mathcal{Y})^*$, where $|S| = 3^k$ for $k \in \mathbb{N}$.

Output: Family of training sequences.

if $k \geq 6$ **then**

 Partition S into $S_1, \dots, S_{27}$, with S_i being the $(i-1)|S|/27 + 1$ to the $i|S|/27$ training examples of S . Set for each i

$$S_{i,\sqcup} = S_i[1 : 3^{k-6}], \quad S_{i,\sqcap} = S_i[3^{k-6} + 1 : 3^k],$$

return $[S(S_{1,\sqcup}; S_{1,\sqcap} \sqcup T), \dots, S(S_{27,\sqcup}; S_{27,\sqcap} \sqcup T)]$

else

return $S \sqcup T$

We first observe that for an input training sequence $m = |S| = 3^k$, the above algorithm makes l recursive calls where $l \in \mathbb{N}$ satisfies $k - 6(l-1) \geq 6$ and $k - 6l < 6$, that is, l is the largest number such that $k/6 \geq l$. As l is a natural number, we get that $l = \lfloor k/6 \rfloor$. Furthermore, since $k = \log_3(m)$ we get that $l = \lfloor \log_3(m)/6 \rfloor$. For each of the l recursive calls 27 recursions are made. Thus, the total number of training sequences created in $\mathcal{S}$ is $27^l \leq 3^{3 \log_3(m)/(2 \cdot 6)} = m^{1/(4 \ln(3))} \geq m^{0.22}$. In what follows, we will use the quantity $s_\sqcap$, which we define as $\frac{|S|}{|S_{1,\sqcap}|}$ when running $\mathcal{S}(S; T)$ with $S; T \in (\mathcal{X} \times \mathcal{Y})^*$ such that $|S| = 3^k$ and $k \geq 6$. We notice that

$$s_\sqcap := \frac{|S|}{|S_{1,\sqcap}|} = \frac{3^k}{3^{k-3} - 3^{k-6}} = \frac{3^k}{3^{k-3}(1 - \frac{1}{3^3})} = \frac{27}{1 - \frac{1}{27}}. \quad (16)$$

This ratio $s_\sqcap$ will later in the proof show up as a constant, $|S_i|/(|S|s_\sqcap) = (|S_{i,\sqcup}| + |S_{i,\sqcap}|)/|S_{i,\sqcap}| = 1/(1 - 1/27)$, multiplied onto τ . Thus, from this relation we see that if the split of S_i is imbalanced so that $|S_{i,\sqcap}|$ is larger than $|S_{i,\sqcup}|$ the constant multiplied on to τ become smaller.

Furthermore, in what follows, for the set of training sequences generated by $\mathcal{S}(S; T)$ and a ERM-algorithm $\mathcal{A}$, we write $\mathcal{A}(S; T)$ for the set of classifiers the ERM-algorithm outputs when run on the training sequences in $\mathcal{S}(S; T)$, i.e., $\mathcal{A}(S; T) = \{\mathcal{A}(S')\}_{S' \in \mathcal{S}(S; T)}$, where this is understood as a multiset if the output of the ERM-algorithm is the same for different training sequences in $\mathcal{S}(S; T)$. Furthermore for an example (x, y) , we define $\text{avg}_{\neq}(\mathcal{A}(S; T))(x, y) = \sum_{h \in \mathcal{A}(S; T)} \mathbb{1}\{h(x) \neq y\} / |\mathcal{A}(S; T)|$, i.e., the average number of incorrect hypotheses in $\mathcal{A}(S; T)$. We notice that by the above comment about $\mathcal{S}(S; T)$ having size at least $m^{0.22}$, we have that $\mathcal{A}(S; T)$ contains just as many hypothesis, each of which is the output of an ERM run on a training sequence of $\mathcal{S}(S; T)$. Thus as allotted to earlier our algorithm do not run on all the sub training sequences created by $\mathcal{S}(S; T)$, as it calls the $\mathcal{A}$ algorithm $O(\ln(m/(\delta(d + \ln(1/\delta)))))$ -times. Which leads us considering the following classifier.

For a natural number t , we let $\hat{\mathcal{A}}_t(S; T)$, be the random collection of t hypotheses drawn uniformly with replacement from $\mathcal{A}(S; T)$, with the draws being independent, where we see $\hat{\mathcal{A}}_t(S; T)$ as a multiset so allowing for repetitions. We remark here that we will overload notation and use $\mathbf{t}$ (so t in bold font) to denote the randomness used to draw the t hypotheses from $\mathcal{A}(S; T)$ in the following analysis of $\hat{\mathcal{A}}_t(S; T)$. Intuitively one can think of $\hat{\mathcal{A}}_t$ as a bagging algorithm where the subsampled training sequences are restricted to subsets of $\mathcal{S}(S; T)$ rather than sampling with replacement from the training examples of S and T . In what follows we will consider this algorithm parametrized by $t = O(\ln(m/(\delta(d + \ln(1/\delta)))))$ leading to a classifier with the same order of call to the ERM as stated in. Similarly to $\mathcal{A}(S; T)$ we also define $\text{avg}_{\neq}(\hat{\mathcal{A}}_t(S; T))(x, y) = \sum_{h \in \hat{\mathcal{A}}_t(S; T)} \mathbb{1}\{h(x) \neq y\} / |\hat{\mathcal{A}}_t(S; T)|$

Now, for a distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, 1\}$, training sequences $S; T \in (\mathcal{X} \times \{-1, 1\})^*$, and $\alpha \in [0, 1]$ we will use $\mathcal{L}_{\mathcal{D}}^\alpha(\mathcal{A}(S; T)) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\mathcal{A}(S; T))(\mathbf{x}, \mathbf{y}) \geq \alpha]$, i.e., the probability of at least a α -fraction of the hypotheses in $\mathcal{A}(S; T)$ erroring on a new example drawn according to $\mathcal{D}$.

As above we also define $\mathcal{L}_{\mathcal{D}}^\alpha(\hat{\mathcal{A}}_t(S; T)) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\text{avg}_{\neq}(\hat{\mathcal{A}}_t(S; T))(\mathbf{x}, \mathbf{y}) \geq \alpha]$, for $\hat{\mathcal{A}}_t(S; T)$.

In the following we will for the case where T is the empty training sequence $\emptyset$ us $\hat{\mathcal{A}}_t(S) = \hat{\mathcal{A}}_t(S; \emptyset)$.

A.1 Difficulties of the Proof

Let us take the opportunity to briefly describe the challenges associated with settling the sample complexity of agnostic learning in the regime $\tau = O(d/m)$. In particular, we will demonstrate that a straightforward multiplicative Chernoff argument can *not* yield a result as general as Theorem 1.4.

To begin, notice that the true errors of N hypotheses can indeed be confidently estimated to within error τ using $O(\log(N/\delta)/\tau)$ samples, due to multiplicative Chernoff. However, if $\tau = O(\log(N/\delta)/m)$, then the above yields a vacuous bound. An additive Chernoff bound would also in this case imply an additive $O(\log(N/\delta)/m)$ error term. In the paper, we consider classes of finite VC dimension d but arbitrary, possibly infinite cardinality N , thereby preventing us from employing this bound. Even if one were to bound N as $O((m/d)^d)$ using Sauer-Shelah, this would imply that multiplicative Chernoff is unhelpful when $\tau = O((d \log(m/d) + \log(1/\delta))/m)$. This is a wide regime containing, for instance, $\tau = O(d/m)$, a setting in which we demonstrate that our algorithm is optimal. Furthermore, an additive Chernoff bound would incur a suboptimal additive $O(d \log(m/d)/m)$ term.

However, let us assume that a Chernoff bound argument for infinite hypothesis classes could yield Theorem B.2 for all τ (e.g., using an ϵ -covering argument). In this case, one would have demonstrated that an empirical risk minimizer h' has error $2.1\tau + O(\sqrt{\frac{\tau(d + \ln(1/\delta))}{m}} + \frac{d + \ln(1/\delta)}{m})$. Then, taking $\tau = 0$ (i.e., considering the realizable case), would demonstrate that ERM results in an error of $O(\frac{d + \ln(1/\delta)}{m})$. However, this is not generally true, owing to the worst-case lower bound on ERM learners of $O(\frac{d \ln(1/\epsilon) + \ln(1/\delta)}{m})$, due to (Bousquet et al., 2020, Theorem 11). This strongly suggests that a Chernoff bound argument using a union bound over each function in the hypothesis class (or a discretization of it) cannot yield a result as general as Theorem 1.1, which holds across all τ .

Finally, we remark that our learner for Theorem B.2 is agnostic to τ (i.e., is oblivious to the value of τ), whereas the previous Chernoff procedure requires knowledge of τ . This is a notable distinction, as precise knowledge of τ will typically be unavailable to the learner. In particular, ERM-based estimates of τ will be uninformative in the regime $\tau = o(d/m)$.

B Analysis of $\hat{\mathcal{A}}_t$

As described in the proof sketch, we require a bound on $\mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S; \emptyset))$ in order to upper bound $\mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_t)$. Thus, we now present our error bound for Algorithm 2 when running $\mathcal{A}$ on each dataset generated on $\mathcal{S}(\mathbf{S}, \emptyset)$. (We assume that $|\mathbf{S}| = 3^k$ for $k \in \mathbb{N}$, at the cost of discarding a constant fraction of training points.)

Lemma B.1. *There exists a universal constant $c \geq 1$ such that: For any hypothesis class $\mathcal{H}$ of VC dimension d , distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$, failure parameter $0 < \delta < 1$, training sequence size $m = 3^k$ for $k \geq 6$, and training sequence $\mathbf{S} \sim \mathcal{D}^m$, with probability at least $1 - \delta$ over $\mathbf{S}$ one has that*

$$\mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(\mathbf{S}; \emptyset)) \leq c\tau + \frac{c(d + \ln(e/\delta))}{m}.$$

Let us defer the proof of Lemma B.1 for the moment, and proceed with presenting the main theorem of this section, assuming the claim of Lemma B.1.

Theorem B.2. *There exists a universal constant $c \geq 1$ such that: For any hypothesis class $\mathcal{H}$ of VC dimension d , distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$, failure parameter $0 < \delta < 1$, training sequence size $m = 3^k$ for $k \geq 6$, training sequence $\mathbf{S} \sim \mathcal{D}^m$, and sampling size $t \geq 4 \cdot 243^2 \ln(2m/(\delta(d + \ln(1/\delta))))$, we have with probability at least $1 - \delta$ over $\mathbf{S}$ and the randomness $\mathbf{t}$ used to draw $\hat{\mathcal{A}}_t(\mathbf{S})$ that:*

$$\mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_t(\mathbf{S})) \leq c\tau + \frac{c(d + \ln(e/\delta))}{m}.$$

Proof. Let $\hat{\mathcal{A}}_t(S; \emptyset) = \{\hat{\mathbf{h}}_1, \dots, \hat{\mathbf{h}}_t\}$, considered as a multiset, and recall that the $\hat{\mathbf{h}}_i$ are drawn uniformly at random from $\mathcal{A}(S; \emptyset) = \{\mathcal{A}(S')\}_{S' \in \mathcal{S}(S; \emptyset)}$, which is likewise treated as a multiset. Let

$E_{\mathbf{S}}$ denote the event

$$E_{\mathbf{S}} = \left\{ (x, y) \mid \sum_{h \in \mathcal{A}(\mathbf{S}; \emptyset)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(\mathbf{S}; \emptyset)|} \geq \frac{10}{243} \right\}$$

and

$$\bar{E}_{\mathbf{S}} = \left\{ (x, y) \mid \sum_{h \in \mathcal{A}(\mathbf{S}; \emptyset)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(\mathbf{S}; \emptyset)|} < \frac{10}{243} \right\}$$

its complement. Now fix a realization S of $\mathbf{S}$. Using the fact that $\mathbb{P}[A] = \mathbb{P}[A \cap B] + \mathbb{P}[A \cap \bar{B}]$, we have that

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_{\mathbf{t}}(S; \emptyset)) &= \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_i(\mathbf{x}) \neq \mathbf{y}\} / t \geq \frac{11}{243}, E_S \right] \\ &\quad + \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_i(\mathbf{x}) \neq \mathbf{y}\} / t \geq \frac{11}{243}, \bar{E}_S \right] \\ &\leq \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [E_S] + \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_i(\mathbf{x}) \neq \mathbf{y}\} / t \geq \frac{11}{243}, \bar{E}_S \right]. \end{aligned} \quad (17)$$

We now bound each term in the above, i.e., in Equation (17).

In pursuit of bounding the second term, consider a labeled example $(x, y) \in \bar{E}_S$. We may assume that $\bar{E}_S$ is non-empty, as otherwise the term is simply 0. Now, for any such labeled example $(x, y) \in \bar{E}_S$ we have that $\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_i(x) = y\}$ has expectation

$$\mu_{(x, y)} := \mathbb{E}_{\mathbf{t}} \left[\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_i(x) = y\} \right] = t \cdot \left(\sum_{h \in \mathcal{A}(\mathbf{S}; \emptyset)} \frac{\mathbb{1}\{h(x) = y\}}{|\mathcal{A}(\mathbf{S}; \emptyset)|} \right) \geq (1 - 10/243 \cdot t) \geq t/2,$$

where the final inequality follows from the fact that $(x, y) \in \bar{E}_S$. (Recall that we use the boldface symbol $\mathbf{t}$ to denote the randomness underlying the random variables $\hat{\mathbf{h}}_1, \dots, \hat{\mathbf{h}}_t$.) Now, since $\mathbb{1}\{\hat{\mathbf{h}}_i(x) \neq y\}$ is a collection of i.i.d. $\{0, 1\}$ -random variables, we have by the Chernoff inequality that

$$\begin{aligned} \mathbb{P}_{\mathbf{t}} \left[\sum_{i=1}^t \frac{\mathbb{1}\{\hat{\mathbf{h}}_i(x) \neq y\}}{t} \leq \left(1 - \frac{1}{243} \right) \mu_{(x, y)} \right] &\leq \exp \left(\frac{-\mu_{(x, y)}}{2 \cdot 243^2} \right) \\ &\leq \frac{\delta(d + \ln(1/\delta))}{2m}, \end{aligned}$$

where the final inequality follows from the fact that $\mu_{(x, y)} \geq t/2$ and

$$t \geq 4 \cdot 243^2 \ln \left(\frac{2m}{\delta(d + \ln(1/\delta))} \right).$$

The above implies that with probability at least $1 - \frac{\delta(d + \ln(1/\delta))}{2m}$, we have

$$\sum_{i=1}^t \frac{\mathbb{1}\{\hat{\mathbf{h}}_i(x) = y\}}{t} > \left(1 - \frac{1}{10} \right) \mu_{(x, y)} \geq \left(1 - \frac{1}{243} \right) \left(1 - \frac{10}{243} \right) = 1 - \frac{11}{243}.$$

This further implies that $\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_i(x) \neq y\} / t < 11/243$ with probability at most $\frac{\delta(d + \ln(1/\delta))}{2m}$. As we demonstrated this fact for any pair $(x, y) \in \bar{E}_S$, an application of Markov's inequality yields that

$$\begin{aligned} \mathbb{P}_{\mathbf{t}} \left[\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_i(\mathbf{x}) \neq \mathbf{y}\} / t \geq \frac{11}{243}, \bar{E}_S \right] \geq \frac{(d + \ln(1/\delta))}{m} \right] \\ \leq \frac{(d + \ln(1/\delta)) \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\mathbb{P}_{\mathbf{t}} \left[\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_i(\mathbf{x}) \neq \mathbf{y}\} / t \geq \frac{11}{243} \right] \mathbb{1}\{\bar{E}_S\} \right]}{m} \\ \leq \frac{\delta}{2}. \end{aligned} \quad (18)$$

Note that the first inequality follows from an application of Markov's inequality and the observation that $\bar{E}_S$ depends only upon $(\mathbf{x}, \mathbf{y})$ and $\mathbf{t}$, which are independent from one another, meaning we can swap the order of expectation. The final inequality follows from the bound on the probability of $\sum_{i=1}^t \mathbb{1}\{\hat{\mathbf{h}}_t(\mathbf{x}) \neq \mathbf{y}\}/t \geq \frac{11}{243}$ happening over $\mathbf{t}$, for $x, y \in \bar{E}_S$.

Thus, we conclude that with probability at least $1 - \delta/2$ over $\mathbf{t}$, the random draw of the hypothesis in $\hat{\mathcal{A}}_{\mathbf{t}}(S; \emptyset)$ is such that

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\sum_{i=1}^t \frac{\mathbb{1}\{\hat{\mathbf{h}}_t(\mathbf{x}) \neq \mathbf{y}\}}{t} \geq \frac{11}{243}, \bar{E}_S \right] \leq \frac{d + \ln(1/\delta)}{2m}.$$

Furthermore, as we showed this for any realization S of $\mathbf{S}$ (and $\mathbf{t}$ and $\mathbf{S}$ are independent), we conclude that, with probability at least $1 - \delta/2$ over both $\mathbf{t}$ and $\mathbf{S}$,

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\sum_{i=1}^t \frac{\mathbb{1}\{\hat{\mathbf{h}}_t(\mathbf{x}) \neq \mathbf{y}\}}{t} \geq \frac{11}{243}, \bar{E}_{\mathbf{S}} \right] \leq \frac{d + \ln(1/\delta)}{m}.$$

Furthermore, by Lemma B.1, we have that with probability at least $1 - \delta/2$ over $\mathbf{S}$,

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [E_{\mathbf{S}}] \leq c\tau + \frac{c(d + \ln(2e/\delta))}{m}. \quad (19)$$

Furthermore, as this event does not depend on the randomness $\mathbf{t}$ employed in drawing hypotheses for $\hat{\mathcal{A}}_{\mathbf{t}}(\mathbf{S})$, we conclude that the above also holds with probability at least $1 - \delta/2$ over both $\mathbf{S}$ and $\mathbf{t}$. Now, applying a union bound over the event in Equation (18) and Equation (19), combined with the bound on $\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{\mathbf{t}}(\mathbf{S}))$ in Equation (17), we get that, with probability at least $1 - \delta$ over $\mathbf{S}$ and $\mathbf{t}$, it holds that

$$\mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_{\mathbf{t}}(\mathbf{S})) \leq c\tau + \sqrt{\frac{c\tau(d + \ln(2e/\delta))}{m}} + \frac{(1+c)(d + \ln(2e/\delta))}{m}.$$

As c is permitted by any absolute constant, this concludes the proof. $\square$

We now proceed to give the proof of Lemma B.1. Doing so will require two additional results, the first of which relates the empirical error of a hypothesis h to its true error.

Lemma B.3 (Shalev-Shwartz and Ben-David (2014) Lemma B.10). *Let $\mathcal{D}$ be a distribution over $\mathcal{X} \times \{-1, 1\}$, $h \in \{-1, 1\}^{\mathcal{X}}$ be a hypothesis, $\delta \in (0, 1)$ a failure parameter, and $m \in \mathbb{N}$ a natural number. Then,*

$$\mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m} \left[\mathcal{L}_{\mathbf{S}}(h) \leq \mathcal{L}_{\mathcal{D}}(h) + \sqrt{\frac{2\mathcal{L}_{\mathcal{D}}(h) \ln(1/\delta)}{3m}} + \frac{2\ln(1/\delta)}{m} \right] \geq 1 - \delta,$$

and

$$\mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m} \left[\mathcal{L}_{\mathcal{D}}(h) \leq \mathcal{L}_{\mathbf{S}}(h) + \sqrt{\frac{2\mathcal{L}_{\mathbf{S}}(h) \ln(1/\delta)}{m}} + \frac{4\ln(1/\delta)}{m} \right] \geq 1 - \delta.$$

The second result we require is one which bounds the error of $\mathcal{A}(S; T)$ for arbitrary training sets T (i.e., not merely $T = \emptyset$, as we have previously considered).

Theorem B.4. *There exists, universal constant $c \geq 1$ such that: For any a hypothesis class $\mathcal{H}$ of VC dimension d , any distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, 1\}$, any failure parameter $\delta \in (0, 1)$, any training sequence size $m = 3^k$, any training sequence $T \in (\mathcal{X} \times \{-1, 1\})^*$, and a random training sequence $\mathbf{S} \sim \mathcal{D}^m$, it holds with probability at least $1 - \delta$ over $\mathbf{S}$ that:*

$$\mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(\mathbf{S}; T)) \leq \max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{c\Sigma_{\neq}(h^*, S')}{m/s_{\square}} + \frac{c(d + \ln(1/\delta))}{m},$$

where

$$\Sigma_{\neq}(h^*, S') = \sum_{(x, y) \in S'} \mathbb{1}\{h^*(x) \neq y\}.$$

Let us first give the proof of Lemma B.1 assuming Theorem B.4, and subsequently offer the proof of Theorem B.4.

Proof of Lemma B.1. First note that for each $i \in \{1, \dots, 27\}$,

$$\begin{aligned} \max_{S' \in \mathcal{S}(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap})} \frac{\Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}} &= \max_{S' \in \mathcal{S}(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap})} \sum_{(x,y) \in S'} \frac{\mathbb{1}\{h^*(x) \neq y\}}{(m/s_{\sqcap})} \\ &\leq \frac{|\mathbf{S}_{i,\sqcap} \sqcup \mathbf{S}_{i,\sqcup}|}{(m/s_{\sqcap})} \sum_{(x,y) \in \mathbf{S}_{i,\sqcap} \sqcup \mathbf{S}_{i,\sqcup}} \frac{\mathbb{1}\{h^*(x) \neq y\}}{|\mathbf{S}_{i,\sqcap} \sqcup \mathbf{S}_{i,\sqcup}|} \\ &= 2 \mathcal{L}_{\mathbf{S}_i}(h^*). \end{aligned}$$

The inequality follows from the fact that any $S' \in \mathcal{S}(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap})$ satisfies $S' \sqsubset \mathbf{S}_{i,\sqcap} \sqcup \mathbf{S}_{i,\sqcup}$ and the final equality uses the facts that $|\mathbf{S}_{i,\sqcap} \sqcup \mathbf{S}_{i,\sqcup}| = |\mathbf{S}_i| = m/3^3$ and $s_{\sqcap} = 3^3/(1 - 1/27)$. Thus, invoking Lemma B.3 over $\mathbf{S}_i$ for $i \in \{1, \dots, 27\}$ with failure parameter $\delta/28$, we have by a union bound that with probability at least $1 - 27\delta/28$ over $\mathbf{S}$, each $i \in \{1, \dots, 27\}$ satisfies

$$\max_{S' \in \mathcal{S}(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap})} \frac{\Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}} \leq 2 \left(\tau + \left(\sqrt{\tau \frac{6 \ln(28/\delta)}{3m}} + \frac{6 \ln(28/\delta)}{m} \right) \right) \leq 4 \left(\tau + \frac{6 \ln(28/\delta)}{m} \right) \quad (20)$$

where we have used that $\sqrt{ab} \leq a + b$. Furthermore, by Theorem B.4 we have that with probability at least $1 - \delta/28$ over $\mathbf{S}$,

$$\mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(\mathbf{S}; \emptyset)) \leq \max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{c \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}} + \frac{c(d + \ln(28/\delta))}{m}$$

Again invoking a union bound, we have with probability at least $1 - \delta$ over $\mathbf{S}$ that

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(\mathbf{S}; \emptyset)) &\leq \max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{c \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}} + \frac{c(d + \ln(28/\delta))}{m} \\ &\leq 4c \left(\tau + \frac{6 \ln(28/\delta)}{m} \right) + \frac{c(d + \ln(28/\delta))}{m} \\ &\leq 7c\tau + \frac{7c(d + \ln(28/\delta))}{m}, \end{aligned}$$

where the first inequality follows from Equation (20). This concludes the proof. $\square$

We now direct our attention to proving Theorem B.4. We will make use of another set of two lemmas, the first of which permits us to make a recursive argument over $\mathcal{A}$ -calls based on sub-training sequences created in Algorithm 2.

Lemma B.5. *Let $S, T \in (\mathcal{X} \times \mathcal{Y})^*$ with $|S| = 3^k$ for $k \geq 6$, and let $\mathcal{D}$ be a distribution over $\mathcal{X} \times \mathcal{Y}$. Then,*

$$\mathcal{L}_{\mathcal{D}}^{10/243}(S(S; T)) \leq 5687 \max_{\substack{i, j \in \{1, \dots, 27\} \\ i \neq j}} \max_{h' \in \mathcal{A}(S_{i,\sqcup}; S_{j,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h'(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(S_{i,\sqcup}; S_{i,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right].$$

Proof. Let (x, y) be an example such that

$$\sum_{h \in \mathcal{A}(S; T)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(S; T)|} \geq \frac{10}{243}.$$

As $k \geq 6$, Algorithm 2 calls itself when called with $(S; T)$. Furthermore, as each of the 27 calls produce an equal number of subtraining sequences, it must be the case that

$$\frac{10}{243} \leq \text{avg}_{\neq}(\mathcal{A}(S; T))(x, y) = \sum_{i=1}^{27} \frac{1}{27} \sum_{h \in \mathcal{A}(S_{i,\sqcup}; S_{i,\sqcap} \sqcup T)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(S_{i,\sqcup}; S_{i,\sqcap} \sqcup T)|}.$$

This in turn implies that there exists an $\hat{i} \in [27]$ satisfying the above inequality, i.e., such that

$$\frac{10}{243} \leq \text{avg}_{\neq}(\mathcal{A}(S_{\hat{i},\sqcup}; S_{\hat{i},\sqcap} \sqcup T))(x, y) = \sum_{h \in \mathcal{A}(S_{\hat{i},\sqcup}; S_{\hat{i},\sqcap} \sqcup T)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(S_{\hat{i},\sqcup}; S_{\hat{i},\sqcap} \sqcup T)|}.$$

We further observe that for any $i \in [27]$,

$$\begin{aligned} \frac{10}{243} &\leq \sum_{j=1}^{27} \frac{1}{27} \sum_{h \in \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)|} \\ &\leq \sum_{j \in \{1, \dots, 27\} \setminus i} \frac{1}{27} \sum_{h \in \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)|} + \frac{1}{27}. \end{aligned}$$

This implies, again for any arbitrary choice of $i \in [27]$, that

$$\frac{1}{243} \leq \sum_{j \in \{1, \dots, 27\} \setminus i} \frac{1}{27} \sum_{h \in \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)|}.$$

Simply multiplying both sides by $27/26$, we have that

$$\frac{1}{234} \leq \sum_{j \in \{1, \dots, 27\} \setminus i} \frac{1}{26} \sum_{h \in \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)} \frac{\mathbb{1}\{h(x) \neq y\}}{|\mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)|}. \quad (21)$$

Using the above, we can conclude that when (x, y) is such that $\text{avg}_{\neq}(\mathcal{A}(S; T))(x, y) \geq \frac{10}{243}$, then there exists an $i \in [27]$ with $\text{avg}_{\neq}(\mathcal{A}(S_{i,\sqcup}; S_{i,\sqcap} \sqcup T))(x, y) \geq \frac{10}{243}$. Then by Equation (21), at least a $1/234$ -fraction of hypotheses in $\bigsqcup_{j \in \{1, \dots, 27\} \setminus i} \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)$ err on (x, y) . Thus, if we let $\mathbf{I}$ be drawn uniformly at random from $\{1, \dots, 27\}$ and $\hat{\mathbf{h}}$ be drawn uniformly at random from $\bigsqcup_{j \in \{1, \dots, 27\} \setminus \mathbf{I}} \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)$, then by the law of total probability we have that

$$\begin{aligned} &\mathbb{P}_{\mathbf{I}, \hat{\mathbf{h}}} \left[\hat{\mathbf{h}}(x) \neq y, \text{avg}_{\neq}(\mathcal{A}(S_{\mathbf{I},\sqcup}; S_{\mathbf{I},\sqcap} \sqcup T))(x, y) \geq \frac{10}{243} \right] \\ &= \mathbb{P}_{\mathbf{I}, \hat{\mathbf{h}}} \left[\hat{\mathbf{h}}(x) \neq y \mid \text{avg}_{\neq}(\mathcal{A}(S_{\mathbf{I},\sqcup}; S_{\mathbf{I},\sqcap} \sqcup T))(x, y) \geq \frac{10}{243} \right] \mathbb{P}_{\hat{\mathbf{h}}} \left[\text{avg}_{\neq}(\mathcal{A}(S_{\mathbf{I},\sqcup}; S_{\mathbf{I},\sqcap} \sqcup T))(x, y) \geq \frac{10}{243} \right] \\ &\geq \frac{1}{234} \cdot \frac{10}{243} \\ &\geq \frac{1}{5687}. \end{aligned}$$

This implies in turn that

$$5687 \mathbb{P}_{\mathbf{I}, \hat{\mathbf{h}}} \left[\hat{\mathbf{h}}(x) \neq y, \text{avg}_{\neq}(\mathcal{A}(S_{\mathbf{I},\sqcup}; S_{\mathbf{I},\sqcap} \sqcup T))(x, y) \geq \frac{10}{243} \right] \geq \mathbb{1} \left\{ \text{avg}_{\neq}(\mathcal{A}(S; T))(x, y) \geq \frac{10}{243} \right\}. \quad (22)$$

Taking expectations with respect to $(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}$, we have

$$5687 \mathbb{E}_{\mathbf{I}, \hat{\mathbf{h}}} \left[\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\hat{\mathbf{h}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(S_{\mathbf{I},\sqcup}; S_{\mathbf{I},\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right] \right] \geq \mathcal{L}_{\mathcal{D}}^{\frac{10}{243}}(\mathcal{A}(S; T)).$$

As $\hat{\mathbf{h}} \in \bigsqcup_{j \in \{1, \dots, 27\} \setminus \mathbf{I}} \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)$, it follows that

$$\begin{aligned} &\mathbb{E}_{\mathbf{I}, \hat{\mathbf{h}}} \left[\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\hat{\mathbf{h}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(S_{\mathbf{I},\sqcup}; S_{\mathbf{I},\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right] \right] \\ &\leq \mathbb{E}_{\mathbf{I}} \left[\max_{h' \in \bigsqcup_{j \in \{1, \dots, 27\} \setminus \mathbf{I}} \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h'(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(S_{\mathbf{I},\sqcup}; S_{\mathbf{I},\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right] \right]. \end{aligned} \quad (23)$$

And as $\mathbf{I} \in \{1, \dots, 27\}$, then clearly

$$\begin{aligned} &\mathbb{E}_{\mathbf{I}} \left[\max_{h' \in \bigsqcup_{j \in \{1, \dots, 27\} \setminus \mathbf{I}} \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h'(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(S_{\mathbf{I},\sqcup}; S_{\mathbf{I},\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right] \right] \\ &\leq \max_{i \in \{1, \dots, 27\}} \max_{h' \in \bigsqcup_{j \in \{1, \dots, 27\} \setminus i} \mathcal{A}(S_{j,\sqcup}; S_{j,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h'(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(S_{i,\sqcup}; S_{i,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right]. \end{aligned} \quad (24)$$

By combining Equations (22) to (24), we conclude that

$$\begin{aligned}
& \mathcal{L}_{\mathcal{D}}^{\frac{10}{243}}(\mathcal{A}(S; T)) \\
& \leq 5687 \max_{i \in \{1, \dots, 27\}} \max_{h' \in \sqcup_{j \in \{1, \dots, 27\} \setminus i} \mathcal{A}(S_{j, \sqcup}; S_{j, \sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h'(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(S_{i, \sqcup}; S_{i, \sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right] \\
& = 5687 \max_{\substack{i, j \in \{1, \dots, 27\} \\ i \neq j}} \max_{h' \in \mathcal{A}(S_{j, \sqcup}; S_{j, \sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h'(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\mathcal{A}(S_{i, \sqcup}; S_{i, \sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right],
\end{aligned}$$

which completes the proof. $\square$

The second lemma employed in the proof of Theorem B.4 is the standard uniform convergence property for VC classes.

Lemma B.6 (Shalev-Shwartz and Ben-David (2014), Theorem 6.8). *There exists a universal constant $C > 1$ such that for any distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, 1\}$ and any hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ with finite VC-dimension d , it holds with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}^m$ that for all $h \in \mathcal{H}$:*

$$\mathcal{L}_{\mathcal{D}}(h) \leq \mathcal{L}_{\mathbf{S}}(h) + \sqrt{\frac{C(d + \ln(e/\delta))}{m}}.$$

We now present the proof of Theorem B.4, which concludes the section.

Proof of Theorem B.4. We induct on $k \geq 1$. In particular, we will demonstrate that for each $k \geq 1$ and $\mathbf{S} \sim \mathcal{D}^m$ with $m = 3^k$, and for any $\delta \in (0, 1)$, $T \in (\mathcal{X} \times \mathcal{Y})^*$, one has with probability at least $1 - \delta$ over $\mathbf{S}$ that

$$\mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(\mathbf{S}; T)) \leq \max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{12000 \cdot \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}} + \sqrt{c_b \frac{C(d + \ln(e/\delta)) \frac{12000 \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}}}{m}} + c_c \frac{C(d + \ln(e/\delta))}{m}, \quad (25)$$

where $s_{\sqcap} = \frac{|\mathbf{S}|}{|S_{i, \sqcap}|} = \frac{27}{1-1/27}$ is the previously defined constant, $C \geq 1$ is the constant from Lemma B.6, and c_b and c_c are the following constants:

$$\begin{aligned}
c_b &= (5687^2 \cdot 4 \cdot 3^6 \ln(24e) s_{\sqcap})^2, \\
c_c &= 3^{12} \ln(24e)^2 \sqrt{c_b} s_{\sqcap}.
\end{aligned}$$

Note that applying $\sqrt{ab} \leq a + b$ to Equation (25) would in fact suffice to complete the proof of Theorem B.4.

Thus it remains only to justify Equation (25). For any choice of $\delta \in (0, 1)$ and $T \in (\mathcal{X} \times \mathcal{Y})^*$, first observe that if $k \leq 12$, the claim follows immediately from the fact that the right hand side of Equation (25) is at least 1. (Owing to the fact that $c_c \geq 3^{12}$.)

We now proceed to the inductive step. For the sake of brevity, we will often suppress the distribution from which random variables are drawn when writing expectations and probabilities, e.g., $\mathbb{P}_{\mathbf{S}}$ rather than $\mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m}$. Now fix a choice of $T \in (\mathcal{X} \times \mathcal{Y})^*$, $\delta \in (0, 1)$, and $k > 12$. Let $a_{\mathbf{S}}$ equal the right-hand side of Equation (25), i.e.,

$$a_{\mathbf{S}} = \max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{12000 \cdot \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}} + \sqrt{c_b \frac{C(d + \ln(e/\delta)) \frac{12000 \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}}}{m}} + c_c \frac{C(d + \ln(e/\delta))}{m}. \quad (26)$$

Then invoking Lemma B.5 and a union bound, we have that

$$\mathbb{P}_{\mathbf{S}} \left[\mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(\mathbf{S}; T)) > a_{\mathbf{S}} \right] \quad (27)$$

$$\begin{aligned} &\leq \mathbb{P}_{\mathbf{S}} \left[5687 \max_{\substack{i,j \in \{1, \dots, 27\} \\ i \neq j}} \max_{h \in \mathcal{A}(\mathbf{S}_{j,\sqcup}; \mathbf{S}_{j,\sqcap} \sqcup T)} \mathbb{P}_{\mathbf{x}, \mathbf{y}} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right] \\ &\leq \sum_{\substack{i,j \in \{1, \dots, 27\} \\ i \neq j}} \mathbb{P}_{\mathbf{S}} \left[5687 \max_{h \in \mathcal{A}(\mathbf{S}_{j,\sqcup}; \mathbf{S}_{j,\sqcap} \sqcup T)} \mathbb{P}_{\mathbf{x}, \mathbf{y}} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right]. \end{aligned} \quad (28)$$

Thus it suffices to show that for $i \neq j \in [27]$,

$$\mathbb{P}_{\mathbf{S}} \left[5687 \max_{h \in \mathcal{A}(\mathbf{S}_{j,\sqcup}; \mathbf{S}_{j,\sqcap} \sqcup T)} \mathbb{P}_{\mathbf{x}, \mathbf{y}} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right] \leq \frac{\delta}{26 \cdot 27}, \quad (29)$$

as one can immediately apply this inequality with Equation (27). Then it remains to establish Equation (29). As the pairs $(\mathbf{S}_{1,\sqcup}, \mathbf{S}_{1,\sqcap}), \dots, (\mathbf{S}_{27,\sqcup}, \mathbf{S}_{27,\sqcap})$ are all i.i.d., it suffices to demonstrate the inequality for, say, $j = 1$ and $i = 2$. To this end, fix arbitrary realizations $(S_k)_{3 \leq k \leq 27}$ of the random variables $(S_k)_{3 \leq k \leq 27}$; we will demonstrate the claim for any such realization.

First note that if we happen to have realizations $S_{2,\sqcup}, S_{2,\sqcap}$ of $\mathbf{S}_{2,\sqcup}, \mathbf{S}_{2,\sqcap}$ such that

$$\mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)) = \mathbb{P}_{\mathbf{x}, \mathbf{y}} \left[\text{avg}_{\neq}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right] \leq c_c \frac{C(d + \ln(e/\delta))}{5687m} \leq \frac{a_{\mathbf{S}}}{5687},$$

then we are done by monotonicity of measures, as

$$5687 \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \mathbb{P}_{\mathbf{x}, \mathbf{y}} \left[\text{avg}_{\neq}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] \leq a_{\mathbf{S}}. \quad (30)$$

Furthermore, consider any realization $S_{2,\sqcap}$ of $\mathbf{S}_{2,\sqcap}$. We note that by $m = |\mathbf{S}| = 3^k$ for $k > 12$, and by Algorithm 2, it holds that $|\mathbf{S}_{2,\sqcup}| = 3^{k-6} = m/3^6$. Thus, we may invoke the inductive hypothesis with $\mathcal{A}(\mathbf{S}_{2,\sqcup}; S_{2,\sqcap} \sqcup T)$ and failure parameter $\delta/24$ in order to conclude that with probability at least $1 - \delta/24$ over $\mathbf{S}_{2,\sqcup}$,

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(\mathbf{S}_{2,\sqcup}; S_{2,\sqcap} \sqcup T)) &\leq \max_{S' \in \mathcal{S}(\mathbf{S}_{2,\sqcup}; S_{2,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/(3^6 s_{\sqcap})} \\ &\quad + \sqrt{c_b \frac{C(d + \ln(24e/\delta))}{m/3^6} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/(3^6 s_{\sqcap})}} + c_c \frac{C(d + \ln(24e/\delta))}{m/3^6}. \end{aligned} \quad (31)$$

Furthermore, for any $a, b, c, d > 0$, we have that

$$a + \sqrt{abc} + cd \leq a + a\sqrt{b} + c\sqrt{b} + cd = (1 + \sqrt{b}) \cdot a + (d + \sqrt{b}) \cdot c, \quad (32)$$

where the inequality follows from the fact that $\sqrt{abc} \leq \max(\sqrt{ba^2}, \sqrt{bc^2}) \leq a\sqrt{b} + c\sqrt{b}$. Now, combining Equation (31) and Equation (32) (with $b = c_b, d = c_c$), we obtain

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(\mathbf{S}_{2,\sqcup}; S_{2,\sqcap} \sqcup T)) &\leq (1 + \sqrt{c_b}) \max_{S' \in \mathcal{S}(\mathbf{S}_{2,\sqcup}; S_{2,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/(3^6 s_{\sqcap})} \\ &\quad + (c_c + \sqrt{c_b}) \frac{C(d + \ln(24e/\delta))}{m/3^6} \\ &\leq 2\sqrt{c_b} \max_{S' \in \mathcal{S}(\mathbf{S}_{2,\sqcup}; S_{2,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/(3^6 s_{\sqcap})} + 2c_c \frac{C(d + \ln(24e/\delta))}{m/3^6}. \end{aligned} \quad (33)$$

Note that the second inequality makes use of the fact that $c_b \geq 1$ and $c_b \leq c_c$. We thus conclude that for any realization $S_{2,\cap}$ of $\mathbf{S}_{2,\cap}$, the above inequality holds with probability at least $1 - \delta/24$ over $\mathbf{S}_{2,\cup}$. Further, as $\mathbf{S}_{2,\cap}$ and $\mathbf{S}_{2,\cup}$ are independent, the inequality also holds with probability at least $1 - \delta/24$ over $\mathbf{S}_{2,\cap}, \mathbf{S}_{2,\cup}$.

We now let

$$a_{\mathbf{S}_2} = 2\sqrt{c_b} \max_{S' \in \mathcal{S}(\mathbf{S}_{2,\cup}; \mathbf{S}_{2,\cap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/(3^6 s_{\cap})} + 2c_c \frac{C(d + \ln(24e/\delta))}{(m/3^6)}$$

and consider the following three events over $\mathbf{S}_2 = (\mathbf{S}_{2,\cap}, \mathbf{S}_{2,\cup})$:

$$\begin{aligned} E_1 &= \left\{ c_c \frac{C(d + \ln(e/\delta))}{5687m} < \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T)) \leq a_{\mathbf{S}_2} \right\}, \\ E_2 &= \left\{ c_c \frac{C(d + \ln(e/\delta))}{5687m} \geq \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T)) \right\}, \\ E_3 &= \left\{ \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T)) > a_{\mathbf{S}_2} \right\}. \end{aligned}$$

By Equation (30), we have that for $\mathbf{S}_{2,\cap}, \mathbf{S}_{2,\cup} \in E_2$, the bound in Equation (25) holds. Furthermore, from the comment below Equation (33), we have that $\mathbf{S}_{2,\cap}, \mathbf{S}_{2,\cup} \in E_3$ happens with probability at most $\delta/24$ over $\mathbf{S}_{2,\cap}, \mathbf{S}_{2,\cup}$. For brevity, let $a_{\mathbf{S}}$ denote the right-hand side of Equation (25). Then, using the law of total probability along with independence of $\mathbf{S}_1$ and $\mathbf{S}_2$, we can conclude that

$$\begin{aligned} & \mathbb{P}_{\mathbf{S}_1, \mathbf{S}_2} \left[\frac{5687}{h \in \mathcal{A}(\mathbf{S}_{1,\cup}; \mathbf{S}_{1,\cap} \sqcup T)} \max_{\mathbf{x}, \mathbf{y}} \mathbb{P} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{2,\cup}; \mathbf{S}_{2,\cap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right] \quad (34) \\ & \leq \mathbb{E}_{\mathbf{S}_2} \left[\mathbb{P}_{\mathbf{S}_1} \left[\frac{5687}{h \in \mathcal{A}(\mathbf{S}_{1,\cup}; \mathbf{S}_{1,\cap} \sqcup T)} \max_{\mathbf{x}, \mathbf{y}} \mathbb{P} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{2,\cup}; \mathbf{S}_{2,\cap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right] \middle| E_1 \right] \mathbb{P}[E_1] \\ & + \mathbb{E}_{\mathbf{S}_2} \left[\mathbb{P}_{\mathbf{S}_1} \left[\frac{5687}{h \in \mathcal{A}(\mathbf{S}_{1,\cup}; \mathbf{S}_{1,\cap} \sqcup T)} \max_{\mathbf{x}, \mathbf{y}} \mathbb{P} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{2,\cup}; \mathbf{S}_{2,\cap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right] \middle| E_2 \right] \mathbb{P}[E_2] \\ & + \mathbb{E}_{\mathbf{S}_2} \left[\mathbb{P}_{\mathbf{S}_1} \left[\frac{5687}{h \in \mathcal{A}(\mathbf{S}_{1,\cup}; \mathbf{S}_{1,\cap} \sqcup T)} \max_{\mathbf{x}, \mathbf{y}} \mathbb{P} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{2,\cup}; \mathbf{S}_{2,\cap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right] \middle| E_3 \right] \mathbb{P}[E_3] \\ & \leq \mathbb{E}_{\mathbf{S}_2} \left[\mathbb{P}_{\mathbf{S}_1} \left[\frac{5687}{h \in \mathcal{A}(\mathbf{S}_{1,\cup}; \mathbf{S}_{1,\cap} \sqcup T)} \max_{\mathbf{x}, \mathbf{y}} \mathbb{P} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{2,\cup}; \mathbf{S}_{2,\cap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right] \middle| E_1 \right] + 0 + \delta/24. \end{aligned}$$

Note that the second inequality follows from Equation (30), Equation (33) and $\mathbb{P}[E_1] \leq 1$. Thus, if we can bound the first term of the final line by $2\delta/24$, it will follow that Equation (29) holds with probability at least $1 - \delta/(26 \cdot 27)$, as claimed.

To this end, consider a realization S_2 of $\mathbf{S}_2 \in E_1$. For such an S_2 , we have that

$$\mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T)) = \mathbb{P}_{\mathbf{x}, \mathbf{y}} \left[\text{avg}_{\neq}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right] > 0.$$

Then, again invoking the law of total probability, we have that

$$\begin{aligned} & \frac{5687}{h \in \mathcal{A}(\mathbf{S}_{1,\cup}; \mathbf{S}_{1,\cap} \sqcup T)} \max_{\mathbf{x}, \mathbf{y}} \mathbb{P} \left[\text{avg}_{\neq}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] \\ & = \frac{5687}{h \in \mathcal{A}(\mathbf{S}_{1,\cup}; \mathbf{S}_{1,\cap} \sqcup T)} \max_{\mathbf{x}, \mathbf{y}} \mathbb{P} \left[h(\mathbf{x}) \neq \mathbf{y} \mid \text{avg}_{\neq}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243} \right] \\ & \quad \times \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T)). \quad (35) \end{aligned}$$

Now let $A = \{(x, y) \in (\mathcal{X} \times \mathcal{Y}) \mid \text{avg}_{\neq}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T))(x, y) \geq \frac{10}{243}\}$ and $\mathbf{N}_1 = |\mathbf{S}_{1,\cap} \cap A|$. As $S_2 \in E_1$, we have that

$$\mathbb{P}_{\mathbf{x}, \mathbf{y}}[A] = \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\cup}; S_{2,\cap} \sqcup T)) \geq \frac{c_c C(d + \ln(e/\delta))}{5687 \cdot m}.$$

Then, owing to the fact that $\mathbf{S}_{1,\sqcap} \sim \mathcal{D}^{(m/s_\sqcap)}$ — note that $m/s_\sqcap = |\mathbf{S}|/(|\mathbf{S}|/|\mathbf{S}_\sqcap|) = |\mathbf{S}_\sqcap|$ — this implies that

$$\mathbb{E}_{\mathbf{S}_{1,\sqcap}} [\mathbf{N}_1] = \mathbb{P}_{\mathbf{x}, \mathbf{y} \sim \mathcal{D}} [A] \cdot (m/s_\sqcap) = \frac{c_c C(d + \ln(e/\delta))}{5687 \cdot s_\sqcap}.$$

Thus, by a multiplicative Chernoff bound, we have

$$\begin{aligned} \mathbb{P}_{\mathbf{S}_{1,\sqcap} \sim \mathcal{D}^{(m/s_\sqcap)}} \left[\mathbf{N}_1 > \mathbb{E}_{\mathbf{S}_{1,\sqcap} \sim \mathcal{D}^{(m/s_\sqcap)}} [\mathbf{N}_1] / 2 \right] &\geq 1 - \exp \left(- \frac{c_c C(d + \ln(e/\delta))}{8 \cdot 5687 s_\sqcap} \right) \\ &\geq 1 - (\delta/e)^{19}. \end{aligned} \quad (36)$$

Note that the second inequality uses the fact that $c_c = 3^{12} \ln(24e)^2 \sqrt{c_b s_\sqcap}$. Now let N_1 be any realization of $\mathbf{N}_1$ such that

$$N_1 > \frac{1}{2} \cdot \mathbb{E}_{\mathbf{S}_{1,\sqcap} \sim \mathcal{D}^{(m/s_\sqcap)}} [\mathbf{N}_1] = \frac{1}{2} \cdot \mathcal{L}_{\mathcal{D}}^{10/243} (\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(m/s_\sqcap).$$

Notice that $\mathbf{S}_{1,\sqcap} \sqcap A \sim \mathcal{D}^{N_1}(\cdot | \text{avg}_{\neq}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243})$.

Now, combining Equation (35) and Lemma B.6, we have that with probability at least $1 - \delta/24$ over $\mathbf{S}_{1,\sqcap} \sqcap A$,

$$\begin{aligned} &5687 \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \mathbb{P}_{\mathbf{x}, \mathbf{y}} \left[\text{avg}_{\neq}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] \\ &\leq 5687 \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \left(\mathcal{L}_{\mathbf{S}_{1,\sqcap} \sqcap A}(h) + \sqrt{C(d + \ln(24e/\delta)) / N_1} \right) \mathcal{L}_{\mathcal{D}}^{10/243} (\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)) \end{aligned} \quad (37)$$

We now bound each of the two terms on the right hand side of Equation (37), considered after multiplying out the term associated with $\mathcal{L}_{\mathcal{D}}^{10/243}$. Beginning with the first term, and recalling that $N_1 = |\mathbf{S}_{1,\sqcap} \sqcap A|$, we have

$$\begin{aligned} &\mathcal{L}_{\mathbf{S}_{1,\sqcap} \sqcap A}(h) \cdot \mathcal{L}_{\mathcal{D}}^{10/243} (\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)) \\ &= \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \sum_{(x,y) \in \mathbf{S}_{1,\sqcap} \sqcap A} \frac{\mathbb{1}\{h(x) \neq y\}}{N_1} \mathcal{L}_{\mathcal{D}}^{10/243} (\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)) \\ &\leq \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \sum_{(x,y) \in \mathbf{S}_{1,\sqcap} \sqcap A} \frac{2\mathbb{1}\{h(x) \neq y\}}{m/s_\sqcap} \\ &\leq \max_{S' \in \mathcal{S}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \sum_{(x,y) \in \mathbf{S}_{1,\sqcap} \sqcap A} \frac{2\mathbb{1}\{\mathcal{A}(S')(x) \neq y\}}{m/s_\sqcap} \\ &\leq \max_{S' \in \mathcal{S}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \sum_{(x,y) \in S'} \frac{2\mathbb{1}\{\mathcal{A}(S')(x) \neq y\}}{m/s_\sqcap} \\ &\leq \max_{S' \in \mathcal{S}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \frac{2\Sigma_{\neq}(h^*, S')}{m/s_\sqcap}. \end{aligned} \quad (38)$$

Note that the first inequality uses the fact that $N_1 \geq \mathcal{L}_{\mathcal{D}}^{10/243} (\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(m/s_\sqcap)/2$ and the second inequality uses that $h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)$, meaning there exists an $S' \in \mathcal{S}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)$ such that $h = \mathcal{A}(S')$. The third inequality follows from the fact that $\mathbf{S}_{1,\sqcap} \sqcap A \subset \tilde{S}$ for any $\tilde{S} \in \mathcal{S}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)$ (and especially for S') and the final inequality from both the ERM-property of $\mathcal{A}$ on S' and the definition of $\Sigma_{\neq}(h^*, S')$.

We now bound the second term of Equation (37). In what follows, let $\beta = C(d + \ln(24e/\delta))$. We will in the first inequality use that $N_1 \geq \frac{1}{2} \cdot \mathcal{L}_D^{10/243}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(m/s_\sqcap)$:

$$\begin{aligned}
& \mathcal{L}_D^{10/243}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)) \sqrt{C(d + \ln(24e/\delta))} / N_1 \\
& \leq \sqrt{\frac{2\beta \mathcal{L}_D^{10/243}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))}{m/s_\sqcap}} \\
& \leq \sqrt{\frac{4\beta \left(\sqrt{c_b} \max_{S' \in \mathcal{S}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/3^6 s_\sqcap} + c_c \frac{C(d + \ln(24e/\delta))}{m/3^6} \right)}{m/s_\sqcap}} \\
& \quad \text{(by definition of } E_1 \text{ and Equation (33))} \\
& \leq \sqrt{\frac{4\beta \sqrt{c_b} \max_{S' \in \mathcal{S}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/(3^6 s_\sqcap)}}{m/s_\sqcap}} + \sqrt{\frac{4\beta^2 c_c \frac{1}{m/3^6}}{m/s_\sqcap}} \\
& \quad \text{(by } \sqrt{a+b} \leq \sqrt{a} + \sqrt{b} \text{ and definition of } \beta) \\
& \leq \sqrt{4 \cdot 3^6 \sqrt{c_b} \ln(24e) s_\sqcap} \sqrt{\frac{C(d + \ln(e/\delta)) \max_{S' \in \mathcal{S}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/s_\sqcap}}{m}} \\
& \quad + \sqrt{4 \cdot 3^6 c_c (\ln(24e))^2 s_\sqcap \frac{C(d + \ln(e/\delta))}{m}},
\end{aligned} \tag{39}$$

where the last inequality follows from $\beta = C(d + \ln(24e/\delta)) \leq \ln(24e)C(d + \ln(e/\delta))$ and rearrangement. We now bound each of the constant terms under the square roots. Beginning with the first term, we have

$$\sqrt{4 \cdot 3^6 \sqrt{c_b} \ln(24e) s_\sqcap} = \sqrt{5687^2 \cdot 4 \cdot 3^6 \sqrt{c_b} \ln(24e) s_\sqcap} / 5687 \leq \sqrt{c_b} / 5687, \tag{40}$$

where the inequality follows from $c_b = (5687^2 \cdot 4 \cdot 3^6 \ln(24e) s_\sqcap)^2$. For the second term, we have that

$$\sqrt{4 \cdot 3^6 c_c (\ln(24e))^2 s_\sqcap} = \sqrt{5687^2 \cdot 4 \cdot 3^6 c_c (\ln(24e))^2 s_\sqcap} / 5687 \leq c_c, \tag{41}$$

where the inequality follows by $c_c = 3^{12} \ln(24e)^2 \sqrt{c_b} s_\sqcap$ and $c_b = (5687^2 \cdot 4 \cdot 3^6 \ln(24e) s_\sqcap)^2$. Then we conclude from Equation (39) that

$$\begin{aligned}
& \sqrt{\frac{C(d + \ln(24e/\delta))}{N_1}} \mathcal{L}_D^{10/243}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)) \\
& \leq \frac{1}{5687} \left(\sqrt{c_b \frac{C(d + \ln(e/\delta)) \max_{S' \in \mathcal{S}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/s_\sqcap}}{m}} + c_c \frac{C(d + \ln(e/\delta))}{m} \right),
\end{aligned} \tag{42}$$

Thus, by applying Equation (38) and Equation (42) to Equation (37), we obtain that for any realization N_1 of $\mathbf{N}_1 \geq \frac{1}{2} \cdot \mathcal{L}_D^{10/243}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(m/s_\sqcap)$, it holds with probability at least $1 - \delta/24$ over

$\mathbf{S}_{1,\sqcap} \sqcap A \sim \mathcal{D}^{N_1}(\cdot | \mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)(x) \neq y)^2$ that

$$5687 \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y})} \left[\text{avg}_{\neq}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] \quad (43)$$

$$\begin{aligned} &\leq \max_{S' \in \mathcal{S}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}} + \sqrt{c_b \frac{C(d + \ln(e/\delta)) \max_{S' \in \mathcal{S}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}}}{m}} \\ &\quad + c_c \frac{C(d + \ln(e/\delta))}{m} \\ &\leq \max_{S' \in \mathcal{S}(\mathbf{S}; T)} \frac{12000 \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}} + \sqrt{c_b \frac{C(d + \ln(e/\delta)) \frac{12000 \Sigma_{\neq}(h^*, S')}{m/s_{\sqcap}}}{m}} + c_c \frac{C(d + \ln(e/\delta))}{m} \\ &= a_{\mathbf{S}}. \end{aligned} \quad (44)$$

Note that the second inequality follows from the fact that $S' \in \mathcal{S}(\mathbf{S}_{i,\sqcup}; \mathbf{S}_{i,\sqcap} \sqcup T)$ for $i = 1, 2$, meaning $S' \in \mathcal{S}(\mathbf{S}; T)$. The equality follows simply from the definition of $a_{\mathbf{S}}$ in Equation (26).

Now, combining the above observations, we can conclude that for any realization $S_2 \in E_1$ of $\mathbf{S}_2$,

$$\begin{aligned} &\mathbb{P}_{\mathbf{S}_1} \left[5687 \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y})} \left[\text{avg}_{\neq}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] \leq a_{\mathbf{S}} \right] \\ &\geq \mathbb{P}_{\mathbf{S}_1} \left[5687 \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y})} \left[\text{avg}_{\neq}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] \leq a_{\mathbf{S}} \right. \\ &\quad \left. \left| \mathbf{N}_1 \geq \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(m/s_{\sqcap})/2 \right] \times \mathbb{P}_{\mathbf{S}_1} \left[\mathbf{N}_1 \geq \mathcal{L}_{\mathcal{D}}^{10/243}(\mathcal{A}(S_{2,\sqcup}; S_{2,\sqcap} \sqcup T))(m/s_{\sqcap})/2 \right] \right] \\ &\geq \left(1 - \frac{\delta}{24} \right) \left(1 - \left(\frac{\delta}{e} \right)^{19} \right) \\ &\geq 1 - 2 \frac{\delta}{2^{12}}. \end{aligned}$$

(Note that the first term on the right side of the first inequality is distributed across two lines of text.) In particular, the first inequality follows from the law of total expectation, and the second from Equation (36) (see the comment below the equations) and Equation (43). Thus, the above implies that the term in Equation (34) which conditions upon $\mathbf{S}_2 \in E_1$ is bounded by $\delta/2^{11}$. Altogether, we can conclude that

$$\mathbb{P}_{\mathbf{S}_1, \mathbf{S}_2} \left[5687 \max_{h \in \mathcal{A}(\mathbf{S}_{1,\sqcup}; \mathbf{S}_{1,\sqcap} \sqcup T)} \mathbb{P}_{(\mathbf{x}, \mathbf{y})} \left[\text{avg}_{\neq}(\mathcal{A}(\mathbf{S}_{2,\sqcup}; \mathbf{S}_{2,\sqcap} \sqcup T))(\mathbf{x}, \mathbf{y}) \geq \frac{10}{243}, h(\mathbf{x}) \neq \mathbf{y} \right] > a_{\mathbf{S}} \right] \leq \frac{\delta}{2^{10}}.$$

As $\delta/2^{10} \leq \delta/2^{6 \cdot 27}$, we arrive at Equation (29), which as previously argued concludes the proof due to the fact that $(\mathbf{S}_{1,\sqcup}, \mathbf{S}_{1,\sqcap}), \dots, (\mathbf{S}_{27,\sqcup}, \mathbf{S}_{27,\sqcap})$ are i.i.d. $\square$

C Augmentation of $\hat{\mathcal{A}}_t$ Through Tie-breaking

Let us assume we are given a training sequence $\mathbf{S}$ of size $m = 3^k$ for $k \geq 1$. We then take $\mathbf{S}$ and split it into three disjoint, equal-sized training sequences $\mathbf{S}_1$, $\mathbf{S}_2$, and $\mathbf{S}_3$. We denote the sizes of $\mathbf{S}_1$, $\mathbf{S}_2$ and $\mathbf{S}_3$ as $m' = m/3$. On $\mathbf{S}_1$ and $\mathbf{S}_2$, we train $\hat{\mathcal{A}}_{t_1}(S_1)$ and $\hat{\mathcal{A}}_{t_2}(S_2)$, where we recall that t_1 and t_2 denote the randomness used to draw the hypothesis in $\hat{\mathcal{A}}_{t_1}(S_1)$ and $\hat{\mathcal{A}}_{t_2}(S_2)$ from, respectively, $\mathcal{A}(\mathbf{S}_1; \emptyset)$ and $\mathcal{A}(\mathbf{S}_2; \emptyset)$.

We now evaluate $\hat{\mathcal{A}}_{t_1}(\mathbf{S}_1)$ and $\hat{\mathcal{A}}_{t_2}(\mathbf{S}_2)$ on $\mathbf{S}_3$ and consider all the examples $(x, y) \in \mathbf{S}_3$, where $\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(x, y) \geq 11/243$ or $\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(x, y) \geq 11/243$. Denote the set of all such examples as $\mathbf{S}_3^{\neq}$. We now run the ERM-algorithm $\mathcal{A}$ on $\mathbf{S}_3^{\neq}$ to obtain $h_{\text{tie}} = \mathcal{A}(\mathbf{S}_3^{\neq})$.

²I.e., the restriction of $\mathcal{D}$ to those (x, y) pairs satisfying the given condition.

For a point x , let $\text{Tie}^{11/243}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1), \hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2); h_{\text{tie}})(x)$ be equal to the label y if both

$$\sum_{h \in \hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)} \mathbb{1}\{h(x) = y\} / |\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)| \geq 232/243$$

and

$$\sum_{h \in \hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)} \mathbb{1}\{h(x) = y\} / |\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)| \geq 232/243.$$

Otherwise, we set it to $h_{\text{tie}}(x)$. In other words, if both $\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)$ and $\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)$ have at least $232/243$ of their hypotheses agreeing on the same label y , we output that label; otherwise, we output the label of $h_{\text{tie}}(x)$.

Notice that if there were a true label y and point x , such that we ended up outputting the answer of $h_{\text{tie}}(x)$, then at least one of $\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)$ and $\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)$ has more than $11/243$ incorrect answers, not equal to y on x , which we know by Lemma B.1 is unlikely. Furthermore, in the former case and the tie erring on (x, y) , then both $\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)$ and $\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)$ err with at least a $232/243$ fraction of their hypotheses, which again is unlikely by Lemma B.1. Thus, both cases of possible error are unlikely, which we exploit in order to demonstrate the following theorem.

Theorem C.1. *There exists a universal constant $c \geq 1$ such that for any hypothesis class $\mathcal{H}$ of VC dimension d , distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$, failure parameter $0 < \delta < 1$, training sequence size $m = 3^k$ for $k \geq 5$, training sequence $\mathbf{S} \sim \mathcal{D}^m$, and sampling size $t_1, t_2 \geq 4 \cdot 243^2 \ln(2m/(\delta(d + \ln(86/\delta))))$, we have, with probability at least $1 - \delta$ over $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3 \sim \mathcal{D}^{m/3}$ and the randomness $\mathbf{t}_1, \mathbf{t}_2$ used to draw $\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1; \emptyset)$ and $\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2; \emptyset)$, that:*

$$\mathcal{L}_{\mathcal{D}}\left(\text{Tie}^{11/243}\left(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1), \hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2), h_{\text{tie}}\right)\right) \leq 2.0888\tau + \sqrt{\frac{c\tau(d + \ln(e/\delta))}{m}} + \frac{c(d + \ln(e/\delta))}{m}.$$

In the proof of Theorem C.1 we will need the following ERM-theorem. Recall that we take $\mathcal{A}$ to be an ERM-algorithm, meaning $\mathcal{A}$ is proper (i.e., it always emits hypotheses in $\mathcal{H}$), and for any training sequence S that $\mathcal{L}_S(\mathcal{A}(S)) = \min_{h \in \mathcal{H}} \mathcal{L}_S(h)$.

Theorem C.2. *[Shalev-Shwartz and Ben-David (2014) Theorem 6.8] There exists a universal constant $C' > 1$ such that for any distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, 1\}$, any hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ with VC dimension d , and any ERM-algorithm $\mathcal{A}$, it holds with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}^m$ that for all $h \in \mathcal{H}$:*

$$\mathcal{L}_{\mathcal{D}}(\mathcal{A}(\mathbf{S})) \leq \inf_{h \in \mathcal{H}} \mathcal{L}_{\mathcal{D}}(h) + \sqrt{\frac{C'(d + \ln(e/\delta))}{m}}.$$

Proof of Theorem C.1. First note that by the definition of $\text{Tie}^{11/243}$, for $\text{Tie}^{11/243}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1), \hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2), h_{\text{tie}})$ to err on a fixed example (x, y) , it must be the case that either there exists $y' \neq y \in \{-1, 1\}$ such that

$$\begin{aligned} & \sum_{h \in \hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)} \mathbb{1}\{h(x) = y'\} / |\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)| \geq 232/243 \text{ and } \sum_{h \in \hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)} \mathbb{1}\{h(x) = y'\} / |\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)| \geq 232/243 \\ \Leftrightarrow & \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(x, y) \geq \frac{232}{243}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(x, y) \geq \frac{232}{243} \end{aligned}$$

or the case that for all $y' \in \{-1, 1\}$,

$$\begin{aligned} & \left(\sum_{h \in \hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)} \mathbb{1}\{h(x) = y'\} < 232/243 \text{ or } \sum_{h \in \hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)} \mathbb{1}\{h(x) = y'\} < 232/243 \right) \text{ and } h_{\text{tie}}(x) \neq y \\ \Rightarrow & \left(\text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(x, y) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(x, y) \geq \frac{11}{243} \right) \text{ and } h_{\text{tie}}(x) \neq y \end{aligned}$$

Thus, we have that

$$\begin{aligned} & \mathcal{L}_{\mathcal{D}} \left(\text{Tie}^{11/243} \left(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1), \hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2), h_{\text{tie}} \right) (\mathbf{x}) \neq \mathbf{y} \right) \\ & \leq \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243} \right] \\ & \quad + \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}} \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right]. \end{aligned} \quad (45)$$

We now bound each of these terms separately. The first term we will soon bound by

$$\begin{aligned} & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243} \right] \\ & \leq 1.0888\tau + \sqrt{\frac{16cC\tau(d + 2\ln(3c_u/\delta))}{m'}} + \frac{6cC(d + \ln(ec_u/\delta))}{m'} \end{aligned} \quad (46)$$

with probability at least $1 - 82\delta/c_u$, over $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$. Likewise, the second term we will soon bound by

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}, \hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)(\mathbf{x}) \neq \hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)(\mathbf{x}) \right] \leq \tau + \sqrt{\frac{4\tau cC'(d + \ln(c_u/\delta))}{m'}} + \frac{5cC' \ln(c_u e/\delta)}{m} \quad (47)$$

with probability $1 - 4\delta/c_u$ at least over $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3, \mathbf{t}_1$ and $\mathbf{t}_2$, where $c, C, C' \geq 1$ are universal constants and $c_u = 86$. Applying a union bound over the above two events establishes the claim of the theorem.

Let us begin by pursuing Equation (46). To this end, consider any realizations of S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$. We note that for an example (x, y) with $\text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(x, y) \geq 232/243$ and $\text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(x, y) \geq 232/243$, it must be the case that both $\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)$ and $\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2)$ have at least a $232/243$ -fraction of hypotheses which err at (x, y) . Now let $\hat{\mathbf{h}}$ be a random hypothesis drawn from $\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)$. Then with probability at least $232/243$, $\hat{\mathbf{h}}(x) \neq y$. Thus, for any such example (x, y) , we conclude that

$$\mathbb{P}_{\hat{\mathbf{h}}} \left[\hat{\mathbf{h}}(x) \neq y, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(x, y) \geq \frac{232}{243} \right] \geq \frac{232}{243} \mathbb{1} \left\{ \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(x, y) \geq \frac{232}{243}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(x, y) \geq \frac{232}{243} \right\}.$$

Multiplying both sides of the above equation by $243/232$ and taking expectation with respect to $(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}$ on both sides, we obtain

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243} \right] \quad (48)$$

$$\leq \frac{243}{232} \mathbb{E}_{\hat{\mathbf{h}}} \left[\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\hat{\mathbf{h}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243} \right] \right]. \quad (49)$$

Now by $\hat{\mathbf{h}}$ being drawn from $\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1)$, which in turn is drawn from $\mathcal{A}(S_1; \emptyset)$, we conclude that $\hat{\mathbf{h}}$ is contained in $\mathcal{A}(S_1; \emptyset)$. Thus,

$$\mathbb{E}_{\hat{\mathbf{h}}} \left[\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\hat{\mathbf{h}}(x) \neq y, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq 232/243 \right] \right]$$

can be upper bounded by $\max_{h \in \mathcal{A}(S_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(x) \neq y, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq 232/243 \right]$.

Using this observation and substituting it into Equation (48), we obtain that

$$\begin{aligned} & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243} \right] \\ & \leq \frac{243}{232} \max_{h \in \mathcal{A}(S_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{232}{243} \right] \\ & \leq \frac{243}{232} \max_{h \in \mathcal{A}(S_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right]. \end{aligned}$$

As we demonstrated the above inequality for any realizations of S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$, the inequality also holds for the random variables. We now demonstrate that the right-hand side of the above expression can be bounded by

$$\begin{aligned} & \frac{243}{232} \max_{h \in \mathcal{A}(S_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ & \leq 1.0888\tau + \sqrt{\frac{16cC\tau(d + 2\ln(3c_u/\delta))}{m'}} + \frac{6cC(d + \ln(ec_u/\delta))}{m'} \end{aligned} \quad (50)$$

with probability at least $1 - 82\delta/c_u$ over $\mathbf{S}_1, \mathbf{S}_2$, and $\mathbf{t}_2$. We denote the above event over $\mathbf{S}_1, \mathbf{S}_2$, and $\mathbf{t}_2$ as E_G . In pursuit of Equation (50), we now consider the following events over $\mathbf{S}_2$ and $\mathbf{t}_2$:

$$\begin{aligned} E_1 &= \left\{ \mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_{t_2}(S_2)) \leq \frac{c \ln(c_u e/\delta)}{m'} \right\} \\ E_2 &= \left\{ \frac{c \ln(c_u e/\delta)}{m'} < \mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_{t_2}(S_2)) \leq c\tau + \frac{c(d + \ln(c_u e/\delta))}{m'} \right\} \\ E_3 &= \left\{ \mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_{t_2}(S_2)) > c\tau + \frac{c(d + \ln(c_u e/\delta))}{m'} \right\}, \end{aligned}$$

where c is at least the constant of Theorem B.2 and also greater than $c \geq 2 \cdot 10^6 s_{\square}$. We first notice that if S_2 and t_2 are realizations of $\mathbf{S}_2$ and $\mathbf{t}_2$ in E_1 , then by monotonicity of measures, we have that

$$\frac{243}{232} \max_{h \in \mathcal{A}(S_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \leq \frac{2c \ln(c_u e/\delta)}{m'}, \quad (51)$$

which would imply the event E_G in Equation (50).

We now consider realizations S_2 and t_2 of $\mathbf{S}_2$ and $\mathbf{t}_2$ in E_2 . Let $\mathbf{S}_{1,i,\square}$ denote $(\mathbf{S}_1)_{i,\square}$ and $\mathbf{S}_{1,i,\sqcup} = (\mathbf{S}_1)_{i,\sqcup}$ for $i \in \{1, \dots, 27\}$. Using this notation, we have that

$$\begin{aligned} & \frac{243}{232} \max_{h \in \mathcal{A}(S_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ &= \frac{243}{232} \max_{i \in \{1, \dots, 27\}} \max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\sqcup}; \mathbf{S}_{1,i,\square})} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right]. \end{aligned}$$

We now bound the max term associated to each $i \in [27]$. To this end, fix such an i and let $A = \{(x, y) \mid \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(x, y) \geq 11/243\}$. Also, let $\mathbf{N}_i$ denote the number of examples in $\mathbf{S}_{1,i,\square}$ landing in A , i.e., $\mathbf{N}_i = |\mathbf{S}_{1,i,\square} \cap A|$. Now by S_2 and t_2 being realizations of $\mathbf{S}_2$ and $\mathbf{t}_2$ in E_2 , we have that $\mathbb{P}[A] = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[A] = \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \geq \frac{c \ln(c_u e/\delta)}{m'}$. Thus, as $\mathbf{S}_{1,i,\square} \sim \mathcal{D}^{m'/s_{\square}}$, we have that $\mathbb{E}[\mathbf{N}_i] = \mathbb{P}[A]m'/s_{\square}$. Furthermore, this implies by Chernoff that with probability at least $1 - \delta/c_u$ over $\mathbf{S}_{1,i,\square}$,

$$\mathbb{P}_{\mathbf{N}_i}[\mathbf{N}_i > (1 - 1/1000) \mathbb{P}[A]m'/s_{\square}] \geq 1 - \exp(-\mathbb{P}[A]m'/(2 \cdot 1000^2 s_{\square})) \geq 1 - \delta/c_u$$

where the last inequality follows by $\mathbb{P}[A] \geq \frac{c \ln(c_u e/\delta)}{m'}$ and $c \geq 2 \cdot 10^6 s_{\square}$. Let now $\mathcal{D}(\cdot \mid A)$ be the conditional distribution of A , i.e., for an event E over $\mathcal{X} \times \mathcal{Y}$, we have that $\mathcal{D}(E \mid A) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[(\mathbf{x}, \mathbf{y}) \in E \cap A] / \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[(\mathbf{x}, \mathbf{y}) \in A]$. Since $\mathbf{S}_{1,i,\square} \sim \mathcal{D}$, it follows that $\mathbf{S}_{1,i,\square} \cap A \sim \mathcal{D}(\cdot \mid A)^{\mathbf{N}_i}$. Consider now a realization N_i of $\mathbf{N}_i$ with

$$N_i \geq (1 - 1/1000) \mathbb{P}[A]m'/s_{\square} = (1 - 1/1000) \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2))m'/s_{\square}.$$

Then by the law of total probability and definition of $\mathcal{D}(\cdot \mid A)$, we have that

$$\begin{aligned} & \max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\sqcup}; \mathbf{S}_{1,i,\square})} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ &= \max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\sqcup}; \mathbf{S}_{1,i,\square})} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(\mathbf{x}) \neq \mathbf{y} \mid \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ &= \max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\sqcup}; \mathbf{S}_{1,i,\square})} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}(\cdot \mid A)} [h(\mathbf{x}) \neq \mathbf{y}] \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)), \end{aligned}$$

where the last equality follows by the definition of $\mathcal{D}(\cdot|A)$. Furthermore, by Lemma B.6, we have with probability at least $1 - \delta/c_u$ over $\mathbf{S}_{1,i,\square} \cap A \sim \mathcal{D}(\cdot|A)^{N_i}$ that

$$\begin{aligned} & \max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}(\cdot|A)} [h(\mathbf{x}) \neq \mathbf{y}] \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \\ & \leq \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \left(\max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \mathcal{L}_{\mathbf{S}_{1,i,\square} \cap A}(h) + \sqrt{\frac{C(d + \ln(ec_u/\delta))}{N_i}} \right), \end{aligned} \quad (52)$$

where $C \geq 1$ is the universal constant of Lemma B.6. We now bound each term, starting with the first. Now, $\max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \mathcal{L}_{\mathbf{S}_{1,i,\square} \cap A}(h)$ is equal to $\max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \Sigma_{\neq}(h, \mathbf{S}_{1,i,\square} \cap A)/N_i$. And as any $h \in \mathcal{A}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})$ is equal to $h = \mathcal{A}(S')$ for some $S' \in \mathcal{S}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})$, we get that $\max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \Sigma_{\neq}(h, \mathbf{S}_{1,i,\square} \cap A)$ is equal to $\max_{S' \in \mathcal{S}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \Sigma_{\neq}(\mathcal{A}(S'), \mathbf{S}_{1,i,\square} \cap A)$. Furthermore, since any $S' \in \mathcal{S}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})$ contains the training sequence $\mathbf{S}_{1,i,\square} \cap A$, we get that $\max_{S' \in \mathcal{S}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \Sigma_{\neq}(\mathcal{A}(S'), \mathbf{S}_{1,i,\square} \cap A) \leq \max_{S' \in \mathcal{S}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \Sigma_{\neq}(\mathcal{A}(S'), S')$.

Now we have that $\mathcal{A}(S')$ is a ERM-algorithm run on S' , thus $\mathcal{A}(S') \in \mathcal{H}$ and any other hypothesis in $\mathcal{H}$ has a larger empirical error on S' than $\mathcal{A}(S')$, including h^* . We thus have that $\max_{S' \in \mathcal{S}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \Sigma_{\neq}(\mathcal{A}(S'), S') \leq \max_{S' \in \mathcal{S}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \Sigma_{\neq}(h^*, S')$, which, since S' is contained in $\mathbf{S}_{1,i} = (\mathbf{S}_1)_i$, can further be upper bounded by $\Sigma_{\neq}(h^*, \mathbf{S}_{1,i})$. We can thus conclude that

$$\max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \mathcal{L}_{\mathbf{S}_{1,i,\square} \cap A}(h) \leq \Sigma_{\neq}(h^*, \mathbf{S}_{1,i})/N_i = m' \mathcal{L}_{\mathbf{S}_{1,i}}(h^*)/(3^3 N_i).$$

Note we have used that $|\mathbf{S}_{1,i}| = m'/3^3$. Now, using the fact that we have crystallized an event N_i with $N_i \geq (1 - 1/1000) \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2))m'/s_{\square}$, we have that the first term of Equation (52) (after factoring out the multiplication) can be bounded as follows:

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\square}; \mathbf{S}_{1,i,\square})} \mathcal{L}_{\mathbf{S}_{1,i,\square} \cap A}(h) & \leq \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \frac{m'}{3N_i} \mathcal{L}_{\mathbf{S}_{1,i}}(h^*) \\ & \leq \frac{1000s_{\square}}{3^3 999} \mathcal{L}_{\mathbf{S}_{1,i}}(h^*) \\ & \leq \frac{1000 \cdot 27}{999 \cdot 26} \mathcal{L}_{\mathbf{S}_{1,i}}(h^*) \end{aligned} \quad (53)$$

where we, in the last inequality, have used that $s_{\square} = \frac{3^3}{1 - \frac{1}{27}}$. This concludes the bound on the first term in Equation (52).

Now, using that $N_i \geq (1 - 1/1000) \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2))m'/s_{\square}$, we get that the second term in Equation (52) can be bound as follows:

$$\begin{aligned} \sqrt{\frac{C(d + \ln(ec_u/\delta))}{N_i}} \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) & \leq \sqrt{\frac{1000C(d + \ln(ec_u/\delta)) \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2))s_{\square}}{999m'}} \\ & \leq \sqrt{\frac{3^3 \cdot 27 \cdot 1000C(d + \ln(ec_u/\delta)) \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2))}{26 \cdot 999m'}} \\ & \leq \sqrt{\frac{29C(d + \ln(ec_u/\delta)) \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2))}{m'}}. \end{aligned}$$

In the second inequality we have used that $s_{\square} = \frac{3^3}{1 - \frac{1}{27}}$. We now use the assumption that we had realizations S_2 and t_2 of $\mathbf{S}_2$ and $\mathbf{t}_2$ in E_2 , i.e., such that $\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \leq c\tau + \frac{c(d + \ln(ec_u/\delta))}{m'}$. This allows us to conclude that

$$\sqrt{\frac{4C(d + \ln(ec_u/\delta)) \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2))}{m'}} \leq \sqrt{\frac{4C(d + \ln(ec_u/\delta)) \left(c\tau + \frac{c(d + \ln(ec_u/\delta))}{m'} \right)}{m'}} \quad (54)$$

$$\leq \sqrt{\frac{4cC\tau(d + \ln(ec_u/\delta))}{m'}} + \frac{4cC(d + \ln(ec_u/\delta))}{m'}, \quad (55)$$

where we have used the inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b \geq 0$ in the last step.

Then, to summarize, we have seen that the event $\mathbf{N}_i \geq (1 - 1/1000) \mathbb{P}[A]m'/s_{\square} = (1 - 1/1000) \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2))m'/s_{\square}$ occurs with probability at least $1 - \delta/c_u$, and that conditioned on this event, Equation (52) holds with probability at least $1 - \delta/c_u$. Consequently, we have that with probability at least $1 - 2\delta/c_u$ over $\mathbf{S}_{1,i}$ each of Equation (52), Equation (53) and Equation (54) hold. Then, with probability at least $1 - 2\delta/c_u$ over $\mathbf{S}_{1,i}$ we have that

$$\begin{aligned} & \frac{243}{232} \max_{h \in \mathcal{A}(\mathbf{S}_{1,i,\cup}; \mathbf{S}_{1,i,\cap})} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(x) \neq y, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ & \leq \frac{243}{232} \frac{1000 \cdot 27}{999 \cdot 26} \mathcal{L}_{\mathbf{S}_{1,i}}(h^*) + \sqrt{\frac{4cC\tau(d + \ln(ec_u/\delta))}{m'}} + \frac{4cC(d + \ln(ec_u/\delta))}{m'} \\ & \leq 1.0888 \mathcal{L}_{\mathbf{S}_{1,i}}(h^*) + \sqrt{\frac{4cC\tau(d + \ln(ec_u/\delta))}{m'}} + \frac{4cC(d + \ln(ec_u/\delta))}{m'}. \end{aligned}$$

Now, invoking a union bound over $i \in \{1, \dots, 27\}$, we have that with probability at least $1 - 54\delta/c_u$ over $\mathbf{S}_1$, it holds that

$$\begin{aligned} & \frac{243}{232} \max_{h \in \mathcal{A}(\mathbf{S}_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(x) \neq y, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ & \leq 1.0888 \max_{i \in \{1, \dots, 27\}} \mathcal{L}_{\mathbf{S}_{1,i}}(h^*) + \sqrt{\frac{4cC\tau(d + \ln(ec_u/\delta))}{m'}} + \frac{4cC(d + \ln(ec_u/\delta))}{m'}. \quad (56) \end{aligned}$$

Furthermore, by Lemma B.3 and another union bound over $\mathbf{S}_{1,1}, \dots, \mathbf{S}_{1,27}$, we have that with probability at least $1 - 27\delta/c_u$ over $\mathbf{S}_1 \sim \mathcal{D}^{m'}$ it holds that

$$\max_{i \in \{1, \dots, 27\}} \mathcal{L}_{\mathbf{S}_{1,i}}(h^*) \leq \tau + \sqrt{\frac{2\tau \ln(c_u/\delta)}{3m'}} + \frac{2 \ln(c_u/\delta)}{m'}. \quad (57)$$

Thus, by applying the union bound over the events in Equation (56) and Equation (57), we get that with probability at least $1 - 81\delta/c_u$ over $\mathbf{S}_1$, it holds that

$$\begin{aligned} & \frac{243}{232} \max_{h \in \mathcal{A}(\mathbf{S}_1; \emptyset)} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h(x) \neq y, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ & \leq 1.0888 \left(\tau + \sqrt{\frac{2\tau \ln(c_u/\delta)}{3m'}} + \frac{2 \ln(c_u/\delta)}{m'} \right) + \sqrt{\frac{4cC\tau(d + \ln(ec_u/\delta))}{m'}} + \frac{4cC(d + \ln(ec_u/\delta))}{m'} \\ & \leq 1.0888\tau + \sqrt{\frac{16cC\tau(d + 2 \ln(3c_u/\delta))}{m'}} + \frac{6cC(d + \ln(ec_u/\delta))}{m'}. \quad (58) \end{aligned}$$

Note that this suffices to give the event E_G in Equation (50). We remark that we demonstrated the above for any realizations S_2 and t_2 of $\mathbf{S}_2$ and $\mathbf{t}_2$ in E_2 .

We now notice that by Theorem B.2, the fact that $t_1 \geq 4 \cdot 243^2 \ln(2m/(\delta(d + \ln(86/\delta))))$, and the choice of $c_u = 86$, we have that $\mathbb{P}_{\mathbf{S}_2, \mathbf{t}_2}[E_3] \leq \delta/c_u$. Combining this with the conclusion below Equation (51) and Equation (58), and with the fact that $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_2$ are independent, we have that

$$\begin{aligned} \mathbb{P}_{\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_2}[E_G] &= \mathbb{E}_{\mathbf{S}_2, \mathbf{t}_2} \left[\mathbb{P}_{\mathbf{S}_1}[E_G] \mathbb{1}\{E_1\} \right] + \mathbb{E}_{\mathbf{S}_2, \mathbf{t}_2} \left[\mathbb{P}_{\mathbf{S}_1}[E_G] \mathbb{1}\{E_2\} \right] + \mathbb{E}_{\mathbf{S}_2, \mathbf{t}_2} \left[\mathbb{P}_{\mathbf{S}_1}[E_G] \mathbb{1}\{E_3\} \right] \\ &\geq \mathbb{P}_{\mathbf{S}_2, \mathbf{t}_2}[E_1] + (1 - 81\delta/c_u) \mathbb{P}_{\mathbf{S}_2, \mathbf{t}_2}[E_2] \\ &\geq (1 - 81\delta/c_u)(1 - \mathbb{P}_{\mathbf{S}_2, \mathbf{t}_2}[E_3]) \\ &\geq 1 - 82\delta/c_u. \end{aligned}$$

Note that first equality follows from E_1, E_2, E_3 partitioning the outcomes of $\mathbf{S}_2$ and $\mathbf{t}_2$ and the first inequality follows from the conclusions below Equation (51) and Equation (58), which state that E_G holds with probability 1 on E_1 and with probability at least $1 - 81\delta/c_u$ on E_2 . The second

inequality again follows from E_1, E_2, E_3 partitioning the outcomes of $\mathbf{S}_2$ and $\mathbf{t}_2$ and the bound $\mathbb{P}_{\mathbf{S}_2, \mathbf{t}_2}[E_3] \leq \delta/c_u$, which shows Equation (46).

We now proceed to show Equation (47), i.e., that with probability at least $1 - 4\delta/c_u$ over $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3, \mathbf{t}_1, \mathbf{t}_2$, it holds that

$$\begin{aligned} & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}} \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ & \leq \tau + \sqrt{\frac{4\tau c C'(d + \ln(c_u/\delta))}{m'}} + \frac{5cC' \ln(c_u e/\delta)}{m}. \end{aligned} \quad (59)$$

We denote this event E_F . Towards proving the claim, we consider the following event over $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$,

$$E_4 = \left\{ \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_1}(\mathbf{S}_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{\mathbf{t}_2}(\mathbf{S}_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \leq \frac{c \ln(c_u e/\delta)}{m'} \right\}.$$

As previously mentioned, we take c to be at least the constant of Theorem B.2 and to also satisfy $c \geq 2 \cdot 10^6 s_{\square}$.

Now, if we have realizations S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$ in E_4 , then by monotonicity of measures, we obtain that

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}} \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \leq \frac{c \ln(c_u e/\delta)}{m'},$$

which would imply the event E_F in Equation (59). Now consider a realization of S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$ on the complement $\bar{E}_4$ of E_4 . We then have that

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq 11/243 \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq 11/243 \right] \geq \frac{c \ln(c_u e/\delta)}{m'}.$$

We recall that $\mathbf{S}_3^{\neq}$, are the examples in $(x, y) \in \mathbf{S}_3$ for which

$$\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(x, y) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(x, y) \geq \frac{11}{243}.$$

In the following, let $\mathcal{D}(\cdot | \neq)$ be the conditional distribution of $\mathcal{D}$ given that $\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(x, y) \geq \frac{11}{243}$ or $\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(x, y) \geq \frac{11}{243}$. That is, for an event B over $(\mathcal{X} \times \mathcal{Y})$, we have

$$\mathcal{D}(B | \neq) = \frac{\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[(\mathbf{x}, \mathbf{y}) \in B, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right]}{\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right]}$$

Thus we have that $\mathbf{S}_3^{\neq} \sim \mathcal{D}(B | \neq)$. We now notice by $\mathbf{S}_3 \sim \mathcal{D}$ that

$$\mathbb{E}_{\mathbf{S}_3} [|\mathbf{S}_3^{\neq}|] = m' \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right].$$

Using a Chernoff bound, this implies

$$\begin{aligned} & \mathbb{P}_{\mathbf{S}_3} \left[|\mathbf{S}_3^{\neq}| \geq m' \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] / 2 \right] \\ & \geq 1 - \exp \left(-m' \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] / 8 \right) \\ & \geq 1 - \delta/c_u. \end{aligned}$$

Note that we are also using the facts that

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \geq \frac{c \ln(c_u e/\delta)}{m'}$$

and that $c \geq 2 \cdot 1000^2 s_{\square}$. Now consider an outcome of $N_{\neq} = |\mathbf{S}_3^{\neq}|$ where

$$N_{\neq} \geq \frac{m'}{2} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right].$$

Now, by Theorem C.2 we have that since $h_{\text{tie}} = \mathcal{A}(\mathbf{S}_3^\#)$ and $\mathbf{S}_3^\# \sim \mathcal{D}(\cdot | \neq)$, then with probability at least $1 - \delta/c_u$ over $\mathbf{S}_3^\#$,

$$\begin{aligned} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}(\cdot | \neq)} [h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}] &\leq \inf_{h \in \mathcal{H}} \mathcal{L}_{\mathcal{D}(\cdot | \neq)}(h) + \sqrt{\frac{C'(d + \ln(c_u/\delta))}{N_\#}} \\ &\leq \mathcal{L}_{\mathcal{D}(\cdot | \neq)}(h^*) + \sqrt{\frac{C'(d + \ln(c_u/\delta))}{N_\#}}. \end{aligned}$$

Note that the first inequality uses Theorem C.2 (and $C' > 1$ is the universal constant of Theorem C.2), and the second inequality uses that $h^* \in \mathcal{H}$ so it has error greater than the infimum. Now, using the law of total expectation, we have that

$$\begin{aligned} &\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ &= \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}(\cdot | \neq)} [h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}] \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ &\leq \left(\mathcal{L}_{\mathcal{D}(\cdot | \neq)}(h^*) + \sqrt{\frac{C'(d + \ln(c_u/\delta))}{N_\#}} \right) \\ &\quad \times \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \end{aligned}$$

We now bound each term in the above. First,

$$\begin{aligned} &\mathcal{L}_{\mathcal{D}(\cdot | \neq)}(h^*) \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) > 11/243 \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) > 11/243 \right] \\ &= \mathbb{P}_{\mathbf{x}, \mathbf{y} \sim \mathcal{D}} \left[h^*(\mathbf{x}) \neq \mathbf{y}, \hat{\mathcal{A}}_{t_1}(S_1)(\mathbf{x}) \neq \hat{\mathcal{A}}_{t_2}(S_2)(\mathbf{x}) \right] \end{aligned}$$

which is less than τ . Furthermore, for the second term, we have by

$$N_\# \geq \frac{m'}{2} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) > 11/243 \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) > 11/243 \right]$$

that

$$\begin{aligned} &\sqrt{\frac{C'(d + \ln(c_u/\delta))}{N_\#}} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ &\leq \sqrt{\frac{2C'(d + \ln(c_u/\delta)) \left(\mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_{t_1}(S_1)) + \mathcal{L}_{\mathcal{D}}^{11/243}(\hat{\mathcal{A}}_{t_2}(S_2)) \right)}{m'}} \end{aligned}$$

where the first inequality follows from plugging in that

$$N_\# \geq \frac{m'}{2} \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) > 11/243 \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) > 11/243 \right]$$

in the denominator, and the second inequality follows from a union bound over the event

$$\text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) > 11/243 \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) > 11/243.$$

Thus, we have shown that with probability at least $1 - 2\delta/c_u$ over $\mathbf{S}_3$, it holds that

$$\begin{aligned} &\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ &\leq \tau + \sqrt{\frac{2C'(d + \ln(c_u/\delta)) \left(\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_1}(S_1)) + \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \right)}{m'}} \end{aligned}$$

which, in particular, also implies that with probability at least $1 - 2\delta/c_u$ over $\mathbf{S}_3$

$$\begin{aligned} & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ & \leq \tau + \sqrt{\frac{2C'(d + \ln(c_u/\delta)) \left(\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_1}(S_1)) + \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \right)}{m'}} + \frac{c \ln(c_u e/\delta)}{m}. \end{aligned}$$

Thus, since the above also upper bounds

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right]$$

for any realization S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$, and $\mathbf{t}_2$, on E_4 , (with probability 1), we conclude that for any realization S_1, S_2, t_1 , and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$, and $\mathbf{t}_2$, it holds with probability at least $1 - 2\delta/c_u$ over $\mathbf{S}_3$ that

$$\begin{aligned} & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ & \leq \tau + \sqrt{\frac{2C'(d + \ln(c_u/\delta)) \left(\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_1}(S_1)) + \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \right)}{m'}} + \frac{c \ln(c_u e/\delta)}{m}. \end{aligned}$$

Let this event be denoted E_5 . Now, let E_6 be the event that

$$E_6 = \left\{ \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_1}(S_1)) + \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \leq 2 \left(c\tau + \frac{c(d + \ln(c_u e/\delta))}{m'} \right) \right\}$$

which, by Theorem B.2, the fact that $t_1, t_2 \geq 4 \cdot 243^2 \ln(2m/(\delta(d + \ln(86/\delta))))$, and a union bound, holds with probability at least $1 - 2\delta/c_u$. (We let $c > 1$ be a universal constant at least as large as the universal constant of Theorem B.2, and with $c \geq 2 \cdot 10^6 s_{\square}$.) Now, notice that for realizations S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$, on E_6 , it holds with probability at least $1 - 2\delta/c_u$ over $\mathbf{S}_3$ that

$$\begin{aligned} & \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[h_{\text{tie}}(\mathbf{x}) \neq \mathbf{y}, \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_1}(S_1))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \text{ or } \text{avg}_{\neq}(\hat{\mathcal{A}}_{t_2}(S_2))(\mathbf{x}, \mathbf{y}) \geq \frac{11}{243} \right] \\ & \leq \tau + \sqrt{\frac{2C'(d + \ln(c_u/\delta)) \left(\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_1}(S_1)) + \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{t_2}(S_2)) \right)}{m'}} + \frac{c \ln(c_u e/\delta)}{m} \\ & \leq \tau + \sqrt{\frac{4C'(d + \ln(c_u/\delta)) \left(c\tau + \frac{c(d + \ln(c_u e/\delta))}{m'} \right)}{m'}} + \frac{c \ln(c_u e/\delta)}{m} \\ & \leq \tau + \sqrt{\frac{4\tau c C'(d + \ln(c_u/\delta))}{m'}} + \frac{5c C' \ln(c_u e/\delta)}{m}. \end{aligned}$$

where the first inequality follows by E_5 , the second inequality by E_6 , and the third inequality by $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$. We notice that the above event is E_F in Equation (59). Thus, by the above holding with probability at least $1 - 2\delta/c_u$ for any outcome of S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$ on E_6 , and by $\mathbf{S}_1, \mathbf{S}_2$ and $\mathbf{S}_3$ being independent we conclude that

$$\begin{aligned} \mathbb{P}_{\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3, \mathbf{t}_1, \mathbf{t}_2} [E_F] & \geq \mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3, \mathbf{t}_1, \mathbf{t}_2} \left[\mathbb{P}_{\mathbf{S}_3} [E_F, E_5] \mathbb{1}\{E_6\} \right] \\ & = \mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3, \mathbf{t}_1, \mathbf{t}_2} \left[\mathbb{P}_{\mathbf{S}_3} [E_5] \mathbb{1}\{E_6\} \right] \\ & \geq \mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3, \mathbf{t}_1, \mathbf{t}_2} [(1 - 2\delta/c_u) \mathbb{1}\{E_6\}] \\ & \geq (1 - 2\delta/c_u)^2 \\ & \geq 1 - 4\delta/c_u. \end{aligned}$$

Note that the equality follows from our previous reasoning, i.e., that for outcomes of S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$ on E_6 , and outcomes S_3 of $\mathbf{S}_3$ on E_5 , E_F holds. The second inequality follows from the fact that E_5 holds with probability at least $1 - 2\delta/c_u$ over $\mathbf{S}_3$ for any realizations S_1, S_2, t_1 and t_2 of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$. The third inequality follows from the fact that E_6 holds with probability at least $1 - 2\delta/c_u$ for $\mathbf{S}_1, \mathbf{S}_2, \mathbf{t}_1$ and $\mathbf{t}_2$, which concludes the proof of Equation (47), as desired. $\square$

D Best-of-both-worlds Learner

In this section we demonstrate that splitting a training sample $\mathbf{S} \sim \mathcal{D}^m$ into $\{\mathbf{S}_i \sim \mathcal{D}^{m/3}\}_{i \in [3]}$ followed by running the algorithm of [Hanneke et al. \(2024\)](#) on $\mathbf{S}_1$, running the algorithm of Theorem B.2 on $\mathbf{S}_2$, and selecting the one $h_{\min}$ with smallest empirical error on $\mathbf{S}_3$ gives the following error bound:

$$\mathcal{L}_{\mathcal{D}}(h_{\min}) = \min \left(2.1 \cdot \tau + O \left(\sqrt{\frac{\tau(d + \ln(1/\delta))}{m}} + \frac{(d + \ln(1/\delta))}{m} \right), \right. \\ \left. \tau + O \left(\sqrt{\frac{\tau(d + \ln(1/\delta))}{m}} + \frac{\ln^5(m/d)(d + \ln(1/\delta))}{m} \right) \right).$$

In pursuit of the above, recall the error bound of ([Hanneke et al., 2024](#), Theorem 3), which establishes the existence of a learner $\tilde{\mathcal{A}}$ which with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}^m$ incurs error at most

$$\mathcal{L}_{\mathcal{D}}(\tilde{\mathcal{A}}(\mathbf{S})) \leq \tau + \sqrt{\frac{c'\tau(d + \ln(1/\delta))}{m}} + \frac{c' \ln^5(m/d)(d + \ln(1/\delta))}{m},$$

for a universal constant $c' \geq 1$. We will in the following use $\hat{\mathcal{A}}_t$ to denote the algorithm of Theorem C.1. In what follows let $m' = m/3$.

By invoking the previous bounds on $\hat{\mathcal{A}}_t$ and $\tilde{\mathcal{A}}$ with $\delta = \delta/4$, and further employing a union bound, we have that with probability at least $1 - \delta/2$ over $\mathbf{S}_1$ and $\mathbf{S}_2$, both $\tilde{\mathcal{A}}$ and $\hat{\mathcal{A}}_t$ will emit hypotheses satisfying:

$$\mathcal{L}_{\mathcal{D}}(\tilde{\mathcal{A}}(\mathbf{S}_1)) \leq \tau + \sqrt{\frac{c'\tau(d + \ln(4/\delta))}{m'}} + \frac{c' \ln^5(m'/d)(d + \ln(4/\delta))}{m'} \quad (60)$$

$$\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_t(\mathbf{S}_2; \emptyset)) \leq 2.1 \cdot \tau + \sqrt{\frac{c\tau(d + \ln(4/\delta))}{m'}} + \frac{c(d + \ln(4/\delta))}{m'}, \quad (61)$$

where c is the universal constant of Theorem C.1. Let G denote the event that $\mathbf{S}$ satisfies the previous condition, which, as we have noted, has probability at least $1 - \delta/2$.

Now consider realizations S_1 and S_2 of $\mathbf{S}_1$ and $\mathbf{S}_2$ with $(\mathbf{S}_1, \mathbf{S}_2) \in G$. Further, let $h_{\min} = \arg \min_{h' \in \{\tilde{\mathcal{A}}(S_1), \hat{\mathcal{A}}_t(S_2; \emptyset)\}} (\mathcal{L}_{\mathbf{S}_3}(h'))$. We now invoke Lemma B.3 on $\mathbf{S}_3$ with the classifiers $\tilde{\mathcal{A}}(S_1)$ and $\hat{\mathcal{A}}_t(S_2; \emptyset)$ (and failure probability $\delta/8$) along with a union bound to see that with probability at least $1 - \delta/2$ over $\mathbf{S}_3$, we have that both choices of $h \in \{\tilde{\mathcal{A}}(S_1), \hat{\mathcal{A}}_t(S_2; \emptyset)\}$ satisfy

$$\mathcal{L}_{\mathcal{D}}(h) \leq \mathcal{L}_{\mathbf{S}_3}(h) + \sqrt{\frac{2\mathcal{L}_{\mathbf{S}_3}(h) \ln(8/\delta)}{m'}} + \frac{4 \ln(8/\delta)}{m'}, \quad (62)$$

and

$$\mathcal{L}_{\mathbf{S}_3}(h) \leq \mathcal{L}_{\mathcal{D}}(h) + \sqrt{\frac{2\mathcal{L}_{\mathcal{D}}(h) \ln(8/\delta)}{3m'}} + \frac{2 \ln(8/\delta)}{m'} \leq 2\mathcal{L}_{\mathcal{D}}(h') + \frac{4 \ln(8/\delta)}{m'}. \quad (63)$$

Note that the final inequality follows from the fact that $\sqrt{ab} \leq a + b$ for $a, b > 0$. Then, using the previous inequalities along with the definition of $h_{\min}$, we have that:

$$\mathcal{L}_{\mathcal{D}}(h_{\min}) \leq \mathcal{L}_{\mathbf{S}_3}(h_{\min}) + \sqrt{\frac{2 \mathcal{L}_{\mathbf{S}_3}(h_{\min}) \ln(8/\delta)}{m'}} + \frac{4 \ln(8/\delta)}{m'} \quad (64)$$

$$\leq \min_{h \in \{\tilde{\mathcal{A}}(S_1), \tilde{\mathcal{A}}(S_2; \emptyset)\}} \left(\mathcal{L}_{\mathbf{S}_3}(h) + \sqrt{\frac{2 \mathcal{L}_{\mathbf{S}_3}(h) \ln(8/\delta)}{m'}} + \frac{4 \ln(8/\delta)}{m'} \right) \quad (65)$$

$$\leq \min_{h \in \{\tilde{\mathcal{A}}(S_1), \tilde{\mathcal{A}}(S_2; \emptyset)\}} \left(\mathcal{L}_{\mathcal{D}}(h) + \sqrt{\frac{2 \mathcal{L}_{\mathcal{D}}(h) \ln(8/\delta)}{m'}} + \frac{6 \ln(8/\delta)}{m'} \right. \\ \left. + \sqrt{\frac{2(2 \mathcal{L}_{\mathcal{D}}(h) + 4 \ln(8/\delta)/m') \ln(8/\delta)}{m'}} \right) \quad (66)$$

$$\leq \min_{h \in \{\tilde{\mathcal{A}}(S_1), \tilde{\mathcal{A}}(S_2; \emptyset)\}} \left(\mathcal{L}_{\mathcal{D}}(h) + 3 \sqrt{\frac{2 \mathcal{L}_{\mathcal{D}}(h) \ln(8/\delta)}{m'}} + \frac{9 \ln(8/\delta)}{m'} \right). \quad (67)$$

In particular, the first inequality follows from Equation (62) applied to $h_{\min}$, the second from the definition of $h_{\min}$, the third from Equation (63), and the fourth from the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b > 0$. Now using the equations in Equation (60) we have that

$$\mathcal{L}_{\mathcal{D}}(\tilde{\mathcal{A}}(\mathbf{S}_1)) \leq \tau + \sqrt{\frac{c' \tau (d + \ln(4/\delta))}{m'}} + \frac{c' \ln^5(m'/d) (d + \ln(4/\delta))}{m'} \leq 2\tau + \frac{2c' \ln^5(m'/d) (d + \ln(4/\delta))}{m'} \quad (68)$$

since $\sqrt{ab} \leq a + b$ for $a, b > 0$. (We are further assuming that $m = \Omega(d)$, a standard condition.) Then, using Equation (68) and the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b > 0$, we have

$$\sqrt{\frac{2 \mathcal{L}_{\mathcal{D}}(\tilde{\mathcal{A}}(\mathbf{S}_1)) \ln(8/\delta)}{m'}} \leq \sqrt{\frac{2 \left(2\tau + \frac{2c' \ln^5(m'/d) (d + \ln(4/\delta))}{m'} \right) \ln(8/\delta)}{m'}} \\ \leq \sqrt{\frac{4\tau \ln(8/\delta)}{m'}} + \frac{2c' \ln^5(m'/d) (d + \ln(8/\delta))}{m'}.$$

Using the previous inequality along with Equation (60), we have that

$$\mathcal{L}_{\mathcal{D}}(\tilde{\mathcal{A}}(\mathbf{S}_1)) + 3 \sqrt{\frac{2 \mathcal{L}_{\mathcal{D}}(\tilde{\mathcal{A}}(\mathbf{S}_1)) \ln(8/\delta)}{m'}} + \frac{9 \ln(8/\delta)}{m'} \\ \leq \tau + \sqrt{\frac{c' \tau (d + \ln(4/\delta))}{m'}} + \frac{c' \ln^5(m'/d) (d + \ln(4/\delta))}{m'} + 3 \sqrt{\frac{4\tau \ln(8/\delta)}{m'}} \\ + \frac{6c' \ln^5(m'/d) (d + \ln(8/\delta))}{m'} + \frac{9 \ln(8/\delta)}{m'} \\ \leq \tau + 7 \sqrt{\frac{c' \tau (d + \ln(8/\delta))}{m'}} + \frac{16c' \ln^5(m'/d) (d + \ln(8/\delta))}{m'}.$$

Similarly, using Equation (61) and $\sqrt{ab} \leq a + b$ we have that

$$\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{\mathbf{t}}(\mathbf{S}_2; \emptyset)) \leq 3.1\tau + \frac{2c(d + \ln(4/\delta))}{m'}.$$

Again using the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b > 0$, we have that

$$\sqrt{\frac{2 \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_{\mathbf{t}}(\mathbf{S}_2; \emptyset)) \ln(8/\delta)}{m'}} \leq \sqrt{\frac{6.2\tau \ln(8/\delta)}{m'}} + \frac{2c(d + \ln(8/\delta))}{m'}.$$

Applying Equation (61) then yields

$$\begin{aligned}
& \mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_t(\mathbf{S}_2; \emptyset)) + 3\sqrt{\frac{2\mathcal{L}_{\mathcal{D}}(\hat{\mathcal{A}}_t(\mathbf{S}_2; \emptyset)) \ln(8/\delta)}{m'}} + \frac{9\ln(8/\delta)}{m'} \\
& \leq 2.1\tau + \sqrt{\frac{c\tau(d + \ln(4/\delta))}{m'}} + \frac{c(d + \ln(4/\delta))}{m'} + 3\sqrt{\frac{6.2\tau \ln(8/\delta)}{m'}} + \frac{6c(d + \ln(8/\delta))}{m'} + \frac{9\ln(8/\delta)}{m'} \\
& \leq 2.1\tau + 9\sqrt{\frac{c\tau(d + \ln(8/\delta))}{m'}} + \frac{16c(d + \ln(8/\delta))}{m'}.
\end{aligned}$$

Now plugging in the above expressions into Equation (64) which is exactly the minimum of $\tilde{\mathcal{A}}(\mathbf{S}_1)$ and $\hat{\mathcal{A}}_t(\mathbf{S}_2; \emptyset)$ we arrive at

$$\begin{aligned}
\mathcal{L}_{\mathcal{D}}(h_{\min}) = \min & \left(2.1\tau + 9\sqrt{\frac{c\tau(d + \ln(1/\delta))}{m'}} + \frac{16c(d + \ln(1/\delta))}{m'}, \right. \\
& \left. \tau + 7\sqrt{\frac{c'\tau(d + \ln(1/\delta))}{m'}} + \frac{16c'\ln^5(m'/d)(d + \ln(1/\delta))}{m'} \right) \quad (69)
\end{aligned}$$

with probability at least $1 - \delta/2$ over $\mathbf{S}_3$ for any realizations S_1 and S_2 of $\mathbf{S}_1$ and $\mathbf{S}_2$ in G , now let T denote the event of Equation (69), then we have by independence of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3$ that

$$\mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m}[T] \geq \mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2 \sim \mathcal{D}^{m/3}} \left[\mathbb{P}_{\mathbf{S}_3 \sim \mathcal{D}^{m/3}} [\mathbb{1}\{T\}] \mathbb{1}\{G\} \right] \geq \mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2 \sim \mathcal{D}^{m/3}} [(1 - \delta/2) \mathbb{1}\{G\}]$$

where the first inequality follows by independence of $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3$ and G only depending upon $\mathbf{S}_1, \mathbf{S}_2$ and the second inequality by Equation (69) holding (the event of T) with probability at least $1 - \delta/2$ over $\mathbf{S}_3$ for any $\mathbf{S}_1, \mathbf{S}_2$ in G . Lastly, note that

$$\mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2 \sim \mathcal{D}^{m/3}} [(1 - \delta/2) \mathbb{1}\{G\}] \geq (1 - \delta/2)(1 - \delta/2) \geq 1 - \delta,$$

due to G having probability at least $1 - \delta/2$ over $\mathbf{S}_1, \mathbf{S}_2$ by Equation (60), which concludes the proof.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The error bound of theorem 1.4, presented in the introduction, is the main contribution of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discuss how there is still a small gap in the known lower bounds and upper bounds after our new result.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Each statement in the paper is proven.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: No experiments are included in the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - 4.1 If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - 4.2 If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - 4.3 If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - 4.4 We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: No experiments are included in the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: No experiments are included in the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: No experiments are included in the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: No experiments are included in the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have no concerns about our work violating the NeurIPS code of ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The work is theoretical and the results hold under a mathematical model.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not release any model or datasets.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: We do not have any code or dataset.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release any code or dataset.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not have any experiments.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We have not experiments, so it does not apply.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: We do not employ LLMs for any important or original components of the paper.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.