

Improved Margin Generalization Bounds for Voting Classifiers

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

Kasper Green Larsen
Aarhus University
larsen@cs.au.dk

Abstract

In this paper we establish a new margin-based generalization bound for voting classifiers, refining existing results and yielding tighter generalization guarantees for widely used boosting algorithms such as AdaBoost [Freund and Schapire, 1997]. Furthermore, the new margin-based generalization bound enables the derivation of an optimal weak-to-strong learner: a Majority-of-3 large-margin classifiers with an expected error matching the theoretical lower bound. This result provides a more natural alternative to the Majority-of-5 algorithm by [Høgsgaard et al., 2024], and matches the Majority-of-3 result by [Aden-Ali et al., 2024] for the realizable prediction model.

1 Introduction

Creating ensembles of classifiers is a classic technique in machine learning, and thus studying generalization of such ensembles is natural. An example of the success of ensembles of classifiers is boosting algorithms, which are one of the pillars of classic machine learning, often significantly improving the accuracy of a base learning algorithm by creating an ensemble of multiple base classifiers/hypotheses. More formally, consider binary classification over an input domain \mathcal{X} and let \mathcal{A} be a base learning algorithm that on a (possibly weighted) training sequence $S \in (\mathcal{X} \times \{-1, 1\})^*$, returns a hypothesis $h_S \in \{-1, 1\}^{\mathcal{X}}$ better than guessing on the weighted training sequence. Boosting algorithms, like AdaBoost [Freund and Schapire, 1997] then iteratively invokes \mathcal{A} on reweighed version of the training sequence S to produce hypotheses h_1, \dots, h_t . In each step, the training sequence S is weighed to put more emphasis on data points that h_1, \dots, h_{t-1} misclassified, forcing h_t to focus on these points. The obtained hypotheses are finally combined via a weighted majority vote to obtain the *voting classifier* $f(x) = \text{sign}(\sum_i \alpha_i h_i(x))$ with $\alpha_i > 0$ for all i .

Much research has gone into understanding the impressive performance of boosting algorithms and voting classifiers, with one particularly influential line of work focusing on so-called *margins* [Schapire et al., 1998]. Given a voting classifiers $f(x) = \text{sign}(\sum_i \alpha_i h_i(x))$, the margin of f on a data point $(x, y) \in \mathcal{X} \times \{-1, 1\}$ is then defined as

$$\text{margin}(f, (x, y)) := \frac{y \sum_i \alpha_i h_i(x)}{\sum_i \alpha_i}.$$

Observe that the margin is a number in $[-1, 1]$. The margin is 1 if all hypotheses h_i in f agree and are correct, it is -1 if they agree and are incorrect, and it is 0 if there is a 50-50 weighted split in their predictions for the label of x . The margins can thus be thought of as a (signed) certainty in the prediction made by f .

Early boosting experiments [Schapire et al., 1998] showed that the accuracy of voting classifiers trained with AdaBoost often improves even when adding more hypotheses h_t to f after the point where f perfectly classifies the training data S . As adding more hypotheses to f results in a more complicated model, this behavior was surprising. Work by [Schapire et al., 1998] attribute these improvements in accuracy to improved margins on the training data. In more detail, they proved generalization bounds stating that large margins imply good generalization performance. To give a flavor of these bounds, assume for now that the hypotheses h_i used for constructing a voting classifier, all belong to a finite hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$. If we use $\mathcal{L}_S^\gamma(f)$ to denote the fraction of samples (x, y) in S for which f has margin no more than γ , then [Schapire et al., 1998] showed that for any data distribution \mathcal{D} over $\mathcal{X} \times \{-1, 1\}$, any $0 < \delta < 1$ and $0 < \gamma \leq 1$, it holds with probability at least $1 - \delta$ over a training sequence $\mathbf{S} \sim \mathcal{D}^n$ that every voting classifier f has

$$\mathcal{L}_{\mathcal{D}}(f) := \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[f(\mathbf{x}) \neq \mathbf{y}] = \mathcal{L}_{\mathbf{S}}^\gamma(f) + O\left(\sqrt{\frac{\ln(|\mathcal{H}|) \ln m}{\gamma^2 m} + \frac{\ln(1/\delta)}{m}}\right). \quad (1)$$

As can be seen from this bound, large margins $\gamma > 0$ improve generalization. For the case where all samples have margin at least γ , i.e. $\mathcal{L}_S^\gamma(f) = 0$, [Breiman, 1999] improved this to

$$\mathcal{L}_{\mathcal{D}}(f) = O\left(\frac{\ln(|\mathcal{H}|) \ln m}{\gamma^2 m} + \frac{\ln(1/\delta)}{m}\right). \quad (2)$$

The current state-of-the-art margin generalization bounds nicely interpolates between the two bounds above. Concretely, [Gao and Zhou, 2013] proved the following generalization of Eq. (1) and Eq. (2) often referred to as the *k'th margin bound* (for simplicity, we hide the dependency on δ):

$$\mathcal{L}_{\mathcal{D}}(f) = \mathcal{L}_{\mathbf{S}}^\gamma(f) + O\left(\sqrt{\frac{\mathcal{L}_{\mathbf{S}}^\gamma(f) \ln(|\mathcal{H}|) \ln m}{\gamma^2 m} + \frac{\ln(|\mathcal{H}|) \ln m}{\gamma^2 m}}\right). \quad (3)$$

Lower bounds show that this generalization bound is nearly tight. In particular, the work [Grønlund et al., 2020] showed that for any cardinality N , and parameters $1/m < \tau, \gamma < c$ for a sufficiently small constant $c > 0$, there is a data distribution \mathcal{D} and finite hypothesis class \mathcal{H} with $|\mathcal{H}| = N$, such that with constant probability over $\mathbf{S} \sim \mathcal{D}^m$, there is a voting classifier f over \mathcal{H} with $\mathcal{L}_{\mathbf{S}}^\gamma(f) \leq \tau$ and

$$\mathcal{L}_{\mathcal{D}}(f) = \tau + \Omega\left(\sqrt{\frac{\tau \ln(N) \ln(1/\tau)}{\gamma^2 m} + \frac{\ln(N) \ln m}{\gamma^2 m}}\right). \quad (4)$$

This matches the upper bound in Eq. (3) up to the gap between $\sqrt{\ln(1/\tau)} \approx \sqrt{\ln(1/\mathcal{L}_{\mathbf{S}}^\gamma(f))}$ and $\ln m$, improving by a $\sqrt{\ln(1/\tau)}$ factor over a previous lower bound by [Grønlund et al., 2019].

Note that we have simplified the lower bound slightly, as the true statement would have $\ln m$ replaced by $\ln(\gamma^2 m / \ln N)$. A similar substitution of $\ln m$ by $\ln(\gamma^2 m / \ln(|\mathcal{H}|))$ in the upper bound Eq. (3) is also possible.

If we instead turn to the more general case of voting classifiers over a possibly infinite hypothesis class \mathcal{H} of VC-dimension d , the current state of affairs is less satisfying. Also in the work by [Gao and Zhou, 2013] introducing the k 'th margin bound, they show that for any data distribution \mathcal{D} and hypothesis class \mathcal{H} of VC-dimension d , it holds with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}^m$ that every voting classifier f over \mathcal{H} satisfies

$$\mathcal{L}_{\mathcal{D}}(f) = \mathcal{L}_{\mathbf{S}}^{\gamma}(f) + O\left(\sqrt{\mathcal{L}_{\mathbf{S}}^{\gamma}(f) \left(\frac{d \ln(m/d) \ln m}{\gamma^2 m} + \frac{\ln(1/\delta)}{m}\right)} + \frac{d \ln(m/d) \ln m}{\gamma^2 m} + \frac{\ln(1/\delta)}{m}\right). \quad (5)$$

The only lower bound for finite VC-dimension is Eq. (4) with $\ln(N)$ replaced by d . The gap here is thus a logarithmic factor and the generalization bound in Eq. (5) has not seen any improvements since.

New Margin Generalization Bounds. Our first main contribution is improved and almost optimal generalization bounds for voting classifiers with large margins. Concretely, we show the following theorem

Theorem 1. [Informal statement of Theorem 10] For any hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ of VC-dimension d , distribution \mathcal{D} over $\mathcal{X} \times \{-1, 1\}$, failure parameter $0 < \delta < 1$ and any constant $0 < \varepsilon < 1$, it holds with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}^m$ that for any margin $0 < \gamma \leq 1$ and any voting classifier f over \mathcal{H} , we have

$$\mathcal{L}_{\mathcal{D}}(f) = \mathcal{L}_{\mathbf{S}}^{\gamma}(f) + O\left(\sqrt{\mathcal{L}_{\mathbf{S}}^{\gamma}(f) \left(\frac{d \Gamma\left(\frac{\gamma^2 m}{d}\right)}{\gamma^2 m} + \frac{\ln\left(\frac{1}{\delta}\right)}{m}\right)} + \frac{d \Gamma\left(\frac{\gamma^2 m}{d}\right)}{\gamma^2 m} + \frac{\ln\left(\frac{1}{\delta}\right)}{m}\right),$$

where $\Gamma(x) = \ln(x) \ln^2(\ln x)$.

Our theorem improves over Eq. (5) by nearly a logarithmic factor and the gap between our upper bound and the lower bound in Eq. (4) is essentially $\ln(\ln(\gamma^2 m/d))$ times the ratio between $\sqrt{\ln(1/\mathcal{L}_{\mathbf{S}}^{\gamma}(f))}$ and $\sqrt{\ln(\gamma^2 m/d)}$. Furthermore, all logarithmic factors are now $\ln(\gamma^2 m/d)$ instead of $\ln(m/d)$ and $\ln m$. While the improvement inside the \ln 's might seem minor, this has crucial implications for the development of a new boosting algorithm explained later. Furthermore, and unlike in the case of finite \mathcal{H} (see discussion after Eq. (4)), there does not seem to be a way of tweaking the proof of the previous bound in Eq. (5) to improve the factors $\ln(m/d)$ and $\ln m$ to $\ln(\gamma^2 m/d)$.

New Boosting Results. One of the prime motivations for studying generalization bounds for large margin voting classifiers, is their application to boosting algorithms. When studying boosting theoretically, we typically use the framework of *weak to strong* learning by [Kearns and Valiant, 1988, Kearns and Valiant, 1994]. Let $t \in \{-1, 1\}^{\mathcal{X}}$ be an unknown target concept assigning labels $t(x)$ to samples $x \in \mathcal{X}$. For a distribution \mathcal{D} over

\mathcal{X} , let \mathcal{D}_t be the distribution over $\mathcal{X} \times \{-1, 1\}$ obtained by drawing a sample $\mathbf{x} \sim \mathcal{D}$ and returning the pair $(\mathbf{x}, t(\mathbf{x}))$.

A γ -weak learner \mathcal{W} , is a learning algorithm that for any distribution \mathcal{D} over \mathcal{X} , when given m_0 i.i.d. samples $(\mathbf{x}_i, t(\mathbf{x}_i)) \sim \mathcal{D}_t$, \mathcal{W} produces with probability at least $1 - \delta_0$ a hypothesis h with $\mathcal{L}_{\mathcal{D}_t}(h) = \mathbb{P}_{(\mathbf{x}, t(\mathbf{x})) \sim \mathcal{D}_t}[h(\mathbf{x}) \neq t(\mathbf{x})] \leq 1/2 - \gamma$. Here m_0 and δ_0 are constants. A strong learner in contrast, is a learning algorithm that for any distribution \mathcal{D} over \mathcal{X} , when given $m(\varepsilon, \delta)$ i.i.d. samples from \mathcal{D}_t , it produces with probability at least $1 - \delta$ a hypothesis with $\mathcal{L}_{\mathcal{D}_t}(h) \leq \varepsilon$. A strong learner thus obtains arbitrarily high accuracy when given enough training data.

AdaBoost [Freund and Schapire, 1997] is the most famous algorithm for constructing a weak learner from a strong learner. Concretely, it can be shown that if AdaBoost is run for $O(\gamma^{-2} \ln m)$ iterations, then it produces a voting classifier f with margin $\Omega(\gamma)$ on all samples in a given training sequence S with $|S| = m$ [Schapire and Freund, 2012] [Theorem 5.8]. If the weak learner/base learning algorithm always returns hypotheses from a hypothesis class \mathcal{H} of VC-dimension d , this allows us to use our new generalization bound in Theorem 1 to conclude

Corollary 2. *For any γ -weak learner \mathcal{W} using a hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ of VC-dimension d , distribution \mathcal{D} over \mathcal{X} , target concept t , failure parameter $0 < \delta < 1$ and any constant $0 < \varepsilon < 1$, it holds with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}_t^m$, that the voting classifier f produced by AdaBoost on \mathbf{S} with weak learner \mathcal{W} has*

$$\mathcal{L}_{\mathcal{D}_t}(f) = O\left(\frac{d\Gamma(\gamma^2 m/d)}{\gamma^2 m} + \frac{\ln(1/\delta)}{m}\right),$$

where $\Gamma(x) = \ln(x) \ln^2(\ln x)$.

The previous best known generalization bound for AdaBoost followed from the margin generalization bound Eq. (5) and was $\mathcal{L}_{\mathcal{D}_t}(f) = O\left(\frac{d \ln(m/d) \ln m}{\gamma^2 m} + \frac{\ln(1/\delta)}{m}\right)$. Moreover, our new bound is tight up to a $\ln^2(\ln(\gamma^2 m/d))$ factor as demonstrated by a lower bound of [Høgsgaard et al., 2023] showing that for $c^{-1} \sqrt{d/m} < \gamma < c$ for sufficiently small constant $c > 0$ and VC-dimension $d = \Omega(\ln(1/\gamma))$, there is a data distribution \mathcal{D} , a weak learner \mathcal{W} using a hypothesis class of VC-dimension d , and a concept t , such that AdaBoost run with \mathcal{W} has $\mathcal{L}_{\mathcal{D}_t}(f) = \Omega\left(\frac{d \ln(\gamma^2 m/d)}{\gamma^2 m}\right)$, with constant probability over a training sequence $\mathbf{S} \sim \mathcal{D}_t^m$.

In addition to improving our understanding of AdaBoost, our new generalization bound for voting classifiers also allows us to design an improved weak to strong learner with an optimal in-expectation error. In the work [Larsen and Ritzert, 2022], it was shown that the optimal generalization error of any weak to strong learning algorithm with access to a γ -weak learner using a hypothesis class of VC-dimension d is

$$\mathcal{L}_{\mathcal{D}_t}(f) = \Theta\left(\frac{d}{\gamma^2 m} + \frac{\ln(1/\delta)}{m}\right). \quad (6)$$

In light of the lower bound above for AdaBoost, this implies that AdaBoost is not an optimal weak to strong learner. However, several optimal weak to strong learning algorithms have been developed. In [Larsen and Ritzert, 2022], the authors gave the first

such algorithm. This algorithm uses the sub-sampling idea of [Hanneke, 2016] from optimal realizable PAC learning and runs AdaBoost on $m^{\lg_4 3}$ many sub-samples $S_i \subset S$ of the training data. It combines the produced voting classifiers by taking a majority vote among their predictions, i.e. a majority-of-majorities. Since each S_i has $|S_i| = \Omega(m)$, this slows down AdaBoost by a factor $m^{\lg_4 3} \approx m^{0.79}$. Later work by [Larsen, 2023] showed that running Bagging [Breiman, 1996] to draw $O(\ln(m/\delta))$ many random sub-samples S_i from S and running AdaBoost on each also results in an optimal generalization error matching Eq. (6), thus reducing the computational overhead to a logarithmic factor. Finally, a recent work by [Høgsgaard et al., 2024] built on a Majority-of-3 result in realizable PAC learning [Aden-Ali et al., 2024] to show that simply partitioning a training sequence into 5 disjoint pieces of $m/5$ samples each, and outputting a majority vote among voting classifiers trained with AdaBoost on each sub-sample, results in an optimal in-expectation error of $\mathbb{E}_{\mathbf{S}}[\mathcal{L}_{\mathcal{D}_t}(f)] = O(d/(\gamma^2 m))$.

Our new generalization bound in Theorem 1 allows us to improve the Majority-of-5 algorithm to a more natural Majority-of-3

Corollary 3. [Follows from Theorem 11] For any γ -weak learner \mathcal{W} using a hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ of VC-dimension d , distribution \mathcal{D} over \mathcal{X} , and concept t , it holds that the voting classifiers $f_{\mathbf{S}_1}, f_{\mathbf{S}_2}, f_{\mathbf{S}_3}$ produced by AdaBoost on i.i.d. training sequences $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3 \sim \mathcal{D}_t^m$ with weak learner \mathcal{W} satisfy

$$\mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3 \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t}(\text{Maj}(f_{\mathbf{S}_1}, f_{\mathbf{S}_2}, f_{\mathbf{S}_3}))] = O\left(\frac{d}{\gamma^2 m}\right).$$

It is in this result that it is critical that our generalization bound in Theorem 1 has $\ln(\gamma^2 m/d)$ factors rather than $\ln(m/d)$ or $\ln(m)$ factors. We elaborate on this in Section 2.2.

2 Proof Overviews and Notation

In this section, we first describe the ideas going into our improved generalization bound for voting classifiers, stated in Theorem 1. We then proceed to give an overview of our proof that Majority-of-3 AdaBoosts gives an optimal in-expectation error for weak to strong learning as in Corollary 3. Along the way, we introduce notation that we will use in our proofs.

2.1 New Margin Generalization Bounds

Recall that our goal is to establish Theorem 1, showing that with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}^m$, it holds for all $\gamma \in [0, 1]$ and voting classifiers f over a hypothesis class \mathcal{H} of VC-dimension d that:

$$\mathcal{L}_{\mathcal{D}}(f) = \mathcal{L}_{\mathbf{S}}^{\gamma}(f) + O\left(\sqrt{\mathcal{L}_{\mathbf{S}}^{\gamma}(f) \left(\frac{d\Gamma\left(\frac{\gamma^2 m}{d}\right)}{\gamma^2 m} + \frac{\ln\left(\frac{1}{\delta}\right)}{m}\right)} + \frac{d\Gamma\left(\frac{\gamma^2 m}{d}\right)}{\gamma^2 m} + \frac{\ln\left(\frac{1}{\delta}\right)}{m}\right), \quad (7)$$

where $\Gamma(x) = \ln(x) \ln^2(\ln x)$.

Let us first introduce a more convenient way of representing voting classifiers. Recall that voting classifiers f are of the form $f(x) = \text{sign}(\sum_i \alpha_i h_i(x))$ with all $\alpha_i > 0$. Furthermore, the margin on a training example (x, y) is defined as $y \sum_i \alpha_i h_i(x) / \sum_j \alpha_j$. To avoid the tedious normalization by $\sum_j \alpha_j$, we henceforth assume all voting classifiers have $\sum_i \alpha_i = 1$. In addition, we will drop the $\text{sign}(\cdot)$ and instead write $f(x) = \sum_i \alpha_i h_i(x)$. In this way, we have that $\text{sign}(f(x)) \neq y$ if and only if $yf(x) \leq 0$ (we define $\text{sign}(0) = 0$). Furthermore, the margin is no more than γ if and only if $yf(x) \leq \gamma$. We hence define $\Delta(\mathcal{H})$ as the set of all convex combinations $\sum_i \alpha_i h_i$ for $h_i \in \mathcal{H}$ (i.e. $\sum_i \alpha_i = 1$ and $\alpha_i > 0$) and refer to $\Delta(\mathcal{H})$ as the set of voting classifiers over \mathcal{H} . With this notation, we have $\mathcal{L}_{\mathcal{D}}^{\gamma}(f) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\mathbf{y}f(\mathbf{x}) \leq \gamma]$, $\mathcal{L}_{\mathcal{S}}^{\gamma}(f)$ is the fraction of training examples $(x, y) \in S$ with $yf(x) \leq \gamma$. For $\gamma = 0$ we will use $\mathcal{L}_{\mathcal{D}}(f) = \mathcal{L}_{\mathcal{D}}^0(f) = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\mathbf{y}f(\mathbf{x}) \leq 0] = \mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}}[\text{sign}(f(\mathbf{x})) \neq \mathbf{y}]$.

Partitioning into Intervals. To establish Eq. (7), we first simplify the task by partitioning the range of γ and $\mathcal{L}_{\mathcal{S}}^{\gamma}(f)$ into small intervals $[\gamma_0^i, \gamma_1^i]$ and $[\tau_0^{i,j}, \tau_1^{i,j}]$, respectively. For each interval $[\gamma_0, \gamma_1] = [\gamma_0^i, \gamma_1^i]$ and $[\tau_0, \tau_1] = [\tau_0^{i,j}, \tau_1^{i,j}]$, we show that for any δ , with probability at least $1 - \delta$ we have for every $f \in \Delta(\mathcal{H})$ and $\gamma \in [\gamma_0, \gamma_1]$ that either: $\mathcal{L}_{\mathcal{S}}^{\gamma}(f) \notin [\tau_0, \tau_1]$ or

$$\mathcal{L}_{\mathcal{D}}(f) \leq \tau_1 + O \left(\sqrt{\tau_1 \left(\frac{d\Gamma\left(\frac{m\gamma_0^2}{d}\right)}{\gamma_0^2 m} + \frac{\ln\left(\frac{1}{\delta}\right)}{m}\right)} + \left(\frac{d\Gamma\left(\frac{m\gamma_0^2}{d}\right)}{\gamma_0^2 m} + \frac{\ln\left(\frac{1}{\delta}\right)}{m}\right) \right). \quad (8)$$

We can then union bound over all intervals, choosing appropriate values $\delta_{i,j} < \delta$, to conclude that with probability $1 - \delta$, the guarantee Eq. (8) holds simultaneously for all intervals. If we choose the length of the intervals small enough, all $\gamma \in [\gamma_0, \gamma_1]$ and $\mathcal{L}_{\mathcal{S}}^{\gamma}(f) \in [\tau_0, \tau_1]$ are sufficiently close that we may substitute all occurrences of γ_0 and τ_1 in Eq. (8) by γ and $\mathcal{L}_{\mathcal{S}}^{\gamma}(f)$. Such a partitioning is standard in proofs of margin bounds, although we have to be a little careful in defining the intervals small enough. Having done so, this recovers Eq. (7).

Ghost Set. Thus we focus on showing the claim in Eq. (8). As in many previous proofs of generalization bounds, we first seek to discretize the infinite hypothesis class $\Delta(\mathcal{H})$ and then apply a union bound over a finite set of events/hypotheses. In our proof, we will construct a $(\gamma_0/2)$ ℓ_{∞} -covering N of $\Delta(\mathcal{H})$. Ideally, such a covering would contain for every $f \in \Delta(\mathcal{H})$, a function $f' \in N$ such that $|f(x) - f'(x)| \leq \gamma_0/2$ all $x \in \mathcal{X}$. Unfortunately, there might not be a finite such N when requiring $|f(x) - f'(x)| \leq \gamma_0/2$ for all x in the full input domain \mathcal{X} . We thus start by introducing a *ghost set* $\mathbf{S}' \sim \mathcal{D}^m$ to replace all references to $\mathcal{L}_{\mathcal{D}}(f)$ (which depends on the full domain \mathcal{X}) by $\mathcal{L}_{\mathbf{S}'}(f)$ (which depends only on \mathbf{S}'). Using standard arguments relating $\mathcal{L}_{\mathcal{D}}(f)$ to $\mathcal{L}_{\mathbf{S}'}(f)$, we show that Eq. (8) follows if we can show that with probability $1 - \delta$ over the pair $(\mathbf{S}, \mathbf{S}')$, it holds for every $f \in \Delta(\mathcal{H})$ and $\gamma \in [\gamma_0, \gamma_1]$ that either $\mathcal{L}_{\mathcal{S}}^{\gamma}(f) \notin [\tau_0, \tau_1]$ or

$$\mathcal{L}_{\mathbf{S}'}(f) \leq \tau_1 + O \left(\sqrt{\tau_1 \left(\frac{d\Gamma\left(\frac{m\gamma_0^2}{d}\right)}{\gamma_0^2 m} + \frac{\ln\left(\frac{1}{\delta}\right)}{m}\right)} + \left(\frac{d\Gamma\left(\frac{m\gamma_0^2}{d}\right)}{\gamma_0^2 m} + \frac{\ln\left(\frac{1}{\delta}\right)}{m}\right) \right). \quad (9)$$

Observe that we substituted $\mathcal{L}_{\mathcal{D}}(f)$ by $\mathcal{L}_{\mathbf{S}'}(f)$ compared to Eq. (8).

Covering. To establish Eq. (9), consider first a fixed $f \in \Delta(\mathcal{H})$. Since \mathbf{S} and \mathbf{S}' are i.i.d. samples from \mathcal{D}^m , we have that $\mathcal{L}_{\mathbf{S}}^\gamma(f)$ and $\mathcal{L}_{\mathbf{S}'}(f)$ are strongly concentrated around their means $\mathcal{L}_{\mathcal{D}}^\gamma(f)$ and $\mathcal{L}_{\mathcal{D}}(f)$. Moreover, since $\mathcal{L}_{\mathcal{D}}^\gamma(f) \geq \mathcal{L}_{\mathcal{D}}(f)$, it is highly unlikely that $\mathcal{L}_{\mathbf{S}'}(f)$ is significantly larger than $\mathcal{L}_{\mathbf{S}}(f)$. This is precisely what we need to establish Eq. (9). Concretely, we need to show that there is no f with $\mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0, \tau_1]$ such that $\mathcal{L}_{\mathbf{S}'}(f)$ is much larger than τ_1 . Since this event is unlikely for a fixed f , we now introduce a discretization of $\Delta(\mathcal{H})$ that would preserve any large gap between $\mathcal{L}_{\mathbf{S}}^\gamma(f)$ and $\mathcal{L}_{\mathbf{S}'}(f)$.

To this end, we discretize $\Delta(\mathcal{H})$ on $\mathbf{X} = \mathbf{S} \cup \mathbf{S}'$ via a $\gamma_0/2$ ℓ_∞ -covering N , i.e., for any $f \in \Delta(\mathcal{H})$, there is an $f' \in N$ with $|f(x) - f'(x)| \leq \gamma_0/2$ for all $x \in \mathbf{X}$. Now observe that whenever $yf(x) > \gamma_0$, we also have $yf'(x) > \gamma_0/2$. Thus, for any $\gamma \in [\gamma_0, \gamma_1]$, we have $\mathcal{L}_{\mathbf{S}}^\gamma(f) \geq \mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f')$. Similarly, we have for $yf(x) \leq 0$ that $yf'(x) \leq \gamma_0/2$, and thus $\mathcal{L}_{\mathbf{S}'}(f) \leq \mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f')$. Therefore, a $\gamma_0/2$ ℓ_∞ -covering N preserves the imbalance between $\mathcal{L}_{\mathbf{S}}^\gamma(f)$ and $\mathcal{L}_{\mathbf{S}'}(f)$ via $\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f')$ and $\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f')$.

To construct a $\gamma_0/2$ ℓ_∞ -covering N of \mathbf{X} and union bound over it, we need the point set \mathbf{X} to be fixed - however we still want to be able to show that an imbalance between $\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f')$ and $\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f')$ for some $f' \in N$ is highly unlikely. As in previous works, we employ the following way of viewing the sampling of \mathbf{S} and \mathbf{S}' . First, we draw $\mathbf{X} \sim \mathcal{D}^{2m}$, consisting of $2m$ i.i.d. training examples from \mathcal{D} , and then let \mathbf{S} be m points drawn without replacement from \mathbf{X} , and \mathbf{S}' be the remaining m points of \mathbf{X} , i.e., $\mathbf{S}' = \mathbf{X} \setminus \mathbf{S}$. Taking this viewpoint of drawing \mathbf{S} and \mathbf{S}' allows us to fix the realization X of the points in \mathbf{X} , while still having which training examples ending up in \mathbf{S} and \mathbf{S}' being random. This still allows us to argue that an imbalance between $\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f')$ and $\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f')$ for some $f' \in N$ is unlikely.

Thus, we now consider an arbitrary but fixed realization X of \mathbf{X} , and let N be a $\gamma_0/2$ ℓ_∞ -covering of X . By the above arguments above, if we can show for any $0 < \delta < 1$ and an arbitrary $f \in N$, it holds with probability at least $1 - \delta$ over the random partitioning of X into \mathbf{S}, \mathbf{S}' that either $\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f) > \tau_1$ or

$$\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f) = \tau_1 + O\left(\sqrt{\tau_1 \frac{\ln(\frac{1}{\delta})}{m}} + \frac{\ln(\frac{1}{\delta})}{m}\right), \quad (10)$$

then we can union bound over all $f \in N$, with δ rescaled to $\delta/|N|$, to conclude that with probability $1 - \delta$ it holds for all $f \in N$ that either $\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f) > \tau_1$ or

$$\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f) = \tau_1 + O\left(\sqrt{\tau_1 \frac{\ln(\frac{|N|}{\delta})}{m}} + \frac{\ln(\frac{|N|}{\delta})}{m}\right).$$

Giving an appropriate upper bound on $|N|$ will then imply Eq. (9)

Now, to argue Eq. (10) for a fixed $f \in N$, we want to show that the event $\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f) \leq \tau_1$ and $\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f) = \tau_1 + \Omega(\sqrt{\tau_1 \ln(1/\delta)/m} + \ln(1/\delta)/m)$ happens with probability at most δ . Let μ denote the fraction of mistakes f makes on X and observe that $\mu = (\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f) + \mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f))/2$. We notice that μ has to be at least $\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f)/2 = (\tau_1 + \Omega(\sqrt{\tau_1 \ln(1/\delta)/m} + \ln(1/\delta)/m))/2$.

$\ln(1/\delta)/m)/2$ for the event to occur. Since μ is $\Omega(\ln(1/\delta)/m)$ and $\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f)$ has expectation equal to μ , it follows by an invocation of a Chernoff bound (without replacement) that with probability at least $1 - \delta$ over \mathbf{S} (drawn from X) that

$$\begin{aligned} \mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f) &\geq \left(1 - \sqrt{\frac{2 \ln(1/\delta)}{\mu m}}\right) \mu \\ &= (\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f) + \mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f))/2 - \sqrt{\frac{2 \ln(1/\delta)(\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f) + \mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f))/2}{m}}, \end{aligned}$$

where doing some rearrangements implies the following inequality

$$\frac{\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f)}{2} + \sqrt{\frac{\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f) \ln(1/\delta)}{m}} \geq \frac{\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f)}{2} - \sqrt{\frac{\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f) \ln(1/\delta)}{m}}.$$

We notice that the above inequality is implying that $\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f)$ cannot be too large compared to $\mathcal{L}_{\mathbf{S}}^{\gamma_0/2}(f)$. Specifically the inequality implies that for $\mathcal{L}_{\mathbf{S}}(f) \leq \tau_1$, we must have $\mathcal{L}_{\mathbf{S}'}^{\gamma_0/2}(f) = \tau_1 + O(\sqrt{(\tau_1 \ln(1/\delta))/m} + \ln(1/\delta)/m)$ as desired. Let us finally remark that applying Hoeffding's inequality would be insufficient to obtain our bounds in that we crucially exploit that Chernoff (or Bernstein's) gives bounds relative to the mean μ .

Clipping. While the above argument gives the correct type of bound $\tau_1 + \sqrt{\tau_1 \ln(|N|/\delta)/m} + \ln(|N|/\delta)/m$, the size of the above suggested cover N turns out to be too large. The intuitive reason is that the functions $f \in \Delta(\mathcal{H})$ take values in the range $[-1, 1]$, whereas we only really care about the values being larger than γ or smaller than $-\gamma$ in the losses $\mathcal{L}_{\mathbf{S}}^{\gamma}(f)$ and $\mathcal{L}_{\mathbf{S}'}^{\gamma}(f)$. Constructing a cover for the full range $[-1, 1]$ thus leads to a larger cover size than necessary and hence is too costly for a union bound. Our idea to remedy this, is to *clip* the voting classifiers in $\Delta(\mathcal{H})$. For this let $\gamma > 0$, and f be a function from \mathcal{X} into $[-1, 1]$, we then define $f_{[\gamma]}$ as the function from $\mathcal{X} \rightarrow [-1, 1]$ given by

$$f_{[\gamma]}(x) = \begin{cases} \gamma & \text{if } f(x) > \gamma \\ -\gamma & \text{if } f(x) < -\gamma, \\ f(x) & \text{else} \end{cases} \quad (11)$$

and $\Delta(\mathcal{H})_{[\gamma]} = \{f_{[\gamma]} : f \in \Delta(\mathcal{H})\}$, i.e. the functions in $\Delta(\mathcal{H})$ capped to respectively $-\gamma$ and γ if it goes below or above $-\gamma$ or γ . We will show that $\Delta(\mathcal{H})_{[\gamma]}$ has a small $\gamma_0/2$ -cover \mathcal{N} of cardinality just

$$\mathcal{N}_{\infty}(X, \Delta(\mathcal{H})_{[\gamma]}, \gamma_0/2) = \exp\left(O\left(\frac{d}{\gamma_0^2} \Gamma\left(\frac{m\gamma_0^2}{d}\right)\right)\right), \quad (12)$$

where the notation $\mathcal{N}_{\infty}(\cdot, \cdot, \cdot)$ for a point set $X \subseteq \mathcal{X}$ function class $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{X}}$ and precision parameter α means the smallest size, $\mathcal{N}_{\infty}(X, \mathcal{F}, \alpha)$, of an α ℓ_{∞} -covering of \mathcal{F} , N , on X .

Relating Covering Number and Fat Shattering. We finally need to bound the covering number as in Eq. (12), where a key part of the argument is that we use $\Delta(\mathcal{H})_{[\gamma_1]}$ instead of $\Delta(\mathcal{H})$. With the goal of establishing this bound on the cover size, we use a result by [Rudelson and Vershynin, 2006] relating the covering number to fat shattering. Let us first recall the definition of fat shattering.

For a point set $\{x_1, \dots, x_d\}$ of size d and a level parameter $\beta > 0$, we say that a function class \mathcal{F} β -shatters $\{x_1, \dots, x_d\}$ if there exists levels r_1, \dots, r_d such that for any $b \in \{-1, 1\}^d$, we have that there exists $f \in \mathcal{F}$ such that

$$\begin{aligned} f(x_i) &\leq r_i - \beta & \text{if } b_i = -1 \\ f(x_i) &\geq r_i + \beta & \text{if } b_i = 1, \end{aligned}$$

that is, the function class is rich enough to be β above or below the levels r_1, \dots, r_d on the point set $\{x_1, \dots, x_d\}$. For a function class \mathcal{F} and level $\beta > 0$ we define $\text{fat}_\beta(\mathcal{F})$ as the largest number d , such that there exists a point set x_1, \dots, x_d of size d , which is β -shattered by \mathcal{F} .

With the definition of $\text{fat}_\beta(\mathcal{F})$ in place, [Rudelson and Vershynin, 2006] [Theorem 4.4] says that for any $0 < \alpha < 1/2$, any $0 < \varepsilon < 1$, any function class \mathcal{F} with $\text{fat}_{c\alpha\varepsilon}$ -dimension $d_{c\alpha\varepsilon}$, for a constant $c > 0$, and any point set $X \subseteq \mathcal{X}$, with $|X| = m$, such that $\sum_{x \in X} |f(x)|/m \leq 1$ for any $f \in \mathcal{F}$, it holds that $\ln(\mathcal{N}_\infty(X, \mathcal{F}, \alpha)) = O(d_{c\alpha\varepsilon} \ln\left(\frac{m}{d_{c\alpha\varepsilon}\alpha}\right) \ln^\varepsilon\left(\frac{m}{d_{c\alpha\varepsilon}}\right))$. The above bound looks quite different from Eq. (12), but we will later choose the variable ε in an appropriate way and recover Eq. (12).

A first naive approach would be to invoke the result of [Rudelson and Vershynin, 2006] directly on $\Delta(\mathcal{H})_{[\gamma_1]}$ (or even the unclipped $\Delta(\mathcal{H})$), to conclude that

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma_1]}, \gamma_0/2)) = O\left(d_{c\gamma_0\varepsilon/2} \ln\left(\frac{m}{d_{c\gamma_0\varepsilon/2}\gamma_0}\right) \ln^\varepsilon\left(\left(\frac{m}{d_{c\gamma_0\varepsilon/2}}\right)\right)\right), \quad (13)$$

We will later show that $d_{c\gamma_0\varepsilon/2} = O(d/(\gamma_0\varepsilon)^{-2})$ for $\Delta(\mathcal{H})_{[\gamma_1]}$, with d being the VC-dimension of \mathcal{H} . Inserting this in the above and considering ε as a constant we see this fails to recover our claimed covering number in Eq. (12). In particular, if considering ε as a constant, the $\ln(m/(d_{c\gamma_0\varepsilon/2}\gamma_0))$ factor would become $\ln(\gamma_0 m/d)$ rather than the claimed $\ln(\gamma_0^2 m/d)$. Again, this difference turns out to be crucial for our Majority-of-3 algorithm as we will argue later.

To remedy this, we exploit the clipping. We observe that for $f \in \Delta(\mathcal{H})_{[\gamma_1]}$, we have that $\sum_{x \in X} |f(x)|/m \leq \gamma_1$. We may thus invoke the result of [Rudelson and Vershynin, 2006] on the scaled function class $\gamma_1^{-1} \cdot \Delta(\mathcal{H})_{[\gamma_1]} = \{f \mid \exists f' \in \Delta(\mathcal{H}), f = \gamma_1^{-1} f'\}$, i.e. the functions in $\Delta(\mathcal{H})_{[\gamma_1]}$ scaled by γ_1^{-1} to get that

$$\ln(\mathcal{N}_\infty(X, \gamma_1^{-1} \cdot \Delta(\mathcal{H})_{[\gamma_1]}, \gamma_0/(2\gamma_1))) = O\left(d_{c\gamma_0\varepsilon/(2\gamma_1)} \ln^{1+\varepsilon}\left(\frac{m}{d_{c\gamma_0\varepsilon/(2\gamma_1)}}\right)\right),$$

where $d_{c\gamma_0\varepsilon/(2\gamma_1)}$ is the $\text{fat}_{c\gamma_0\varepsilon/(2\gamma_1)}$ -dimension of $\gamma_1^{-1} \cdot \Delta(\mathcal{H})_{[\gamma_1]}$. Picking a minimal $\gamma_0/(2\gamma_1)$ covering N for $\gamma_1^{-1} \cdot \Delta(\mathcal{H})_{[\gamma_1]}$ and downscaling all functions in N by γ_1 results in $\gamma_1 N$ being a $\gamma_0/2$ covering for $\Delta(\mathcal{H})_{[\gamma_1]}$ as required. We have thus exploited the clipping to show that

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma_1]}, \gamma_0/2)) = O\left(d_{c\gamma_0\varepsilon/(2\gamma_1)} \ln^{1+\varepsilon}\left(\frac{m}{d_{c\gamma_0\varepsilon/(2\gamma_1)}}\right)\right),$$

All that remains is thus to bound $d_{c\gamma_0\varepsilon/(2\gamma_1)}$, the $\text{fat}_{c\gamma_0\varepsilon/(2\gamma_1)}$ -dimension of $\gamma_1^{-1}\Delta(\mathcal{H})_{\lceil\gamma_1\rceil}$ and then set ε appropriate. Now the $\text{fat}_{c\gamma_0\varepsilon/(2\gamma_1)}$ -dimension of $\gamma_1^{-1}\cdot\Delta(\mathcal{H})_{\lceil\gamma_1\rceil}$ is, due to the scale invariance of fat-dimension, the same as the $\text{fat}_{c\gamma_0\varepsilon/2}$ -dimension of $\Delta(\mathcal{H})_{\lceil\gamma_1\rceil}$. We have thus improved [Eq. \(13\)](#) to

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{\lceil\gamma_1\rceil}, \gamma_0/2)) = O\left(d_{c\gamma_0\varepsilon/2} \ln^{1+\varepsilon}\left(\frac{m}{d_{c\gamma_0\varepsilon/2}}\right)\right),$$

Inserting the claimed bound of $d_{c\gamma_0\varepsilon/2} = O(d(\gamma_0\varepsilon)^{-2})$ and setting $\varepsilon = 1/\ln(\ln(m\gamma_0^2/d))$ gives

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{\lceil\gamma_1\rceil}, \gamma_0/2)) = O\left(\frac{d}{\gamma_0^2\varepsilon^2} \ln^{1+\varepsilon}\left(\frac{m\gamma_0^2\varepsilon^2}{d}\right)\right) = O\left(\frac{d}{\gamma_0^2} \ln^2\left(\ln\left(\frac{m\gamma_0^2}{d}\right)\right) \ln\left(\frac{m\gamma_0^2}{d}\right)\right),$$

where in the last inequality have used that $\exp(\varepsilon \ln(\ln(m\gamma_0^2/d))) = O(1)$. Since $\Gamma(x) = \ln^2(\ln(x)) \ln(x)$ the above establishes [Eq. \(12\)](#) and completes our bound on the covering number and. All that remains is thus to argue that $d_{c\gamma_0\varepsilon/2} = O(d(\gamma_0\varepsilon)^{-2})$.

Bounding Fat Shattering Dimension. To bound $d_{c\gamma_0\varepsilon/2}$, we use an argument similar to the proof of [\[Larsen and Ritzert, 2022\]](#) [Lemma 9]. Assume $\Delta(\mathcal{H})_{\lceil\gamma_1\rceil}$ $c\gamma_0\varepsilon/2$ -shatters a set of n points x_1, \dots, x_n , with witness $r_1, \dots, r_n \in [-\gamma_1, \gamma_1]$. By definition of shattering, we then have that for any $b \in \{-1, 1\}^n$, there exists $f \in \Delta(\mathcal{H})_{\lceil\gamma_1\rceil}$ such that $b_i(f(x_i) - r_i) \geq c\gamma_0\varepsilon/2$ for all $i = 1, \dots, n$.

We next observe that since $f \in \Delta(\mathcal{H})_{\lceil\gamma_1\rceil}$ is equal to $\min(\max(f', -\gamma_1), \gamma_1)$ for $f' \in \Delta(\mathcal{H})$, we have that f' also satisfies $b_i(f'(x_i) - r_i) \geq c\gamma_0\varepsilon/2$, implying that $\Delta(\mathcal{H})$ also $c\gamma_0\varepsilon/2$ -shatters x_1, \dots, x_n with the same witness. We can now upper and lower bound the Rademacher complexity of $\Delta(\mathcal{H})$ as follows

$$c\gamma_0\varepsilon/2 \leq \mathbb{E}_{\sigma \sim \{-1, 1\}^n} \left[\sup_{f \in \Delta(\mathcal{H})} \sum_{i=1}^n \sigma_i (f(x_i) - r_i) / n \right] = \mathbb{E}_{\sigma \sim \{-1, 1\}^n} \left[\sup_{f \in \mathcal{H}} \sum_{i=1}^n \sigma_i f(x_i) \right] \leq c' \cdot \sqrt{\frac{d}{n}}, \quad (14)$$

for a constant $c' > 0$. The first inequality holds because for any $\sigma \in \{-1, 1\}^n$, by definition of $c\gamma_0\varepsilon/2$ -shattering, there is an $f \in \Delta(\mathcal{H})$ with $\sigma_i(f(x_i) - r_i) \geq c\gamma_0\varepsilon/2$ for all i . The equality holds because $-\sigma_i r_i/n$ is independent of f , and thus can be moved outside the sup, and we have $\mathbb{E}[\sigma_i r_i/n] = \mathbb{E}[\sigma_i] r_i/n = 0$. Furthermore, observe that in the equality, we also replace $\sup_{f \in \Delta(\mathcal{H})}$ by $\sup_{f \in \mathcal{H}}$. This is true since for any convex combination $f \in \Delta(\mathcal{H})$ with $f = \sum_j \alpha_j h_j$ we have $\sum_i \sigma_i f(x_i) = \sum_j \alpha_j \sum_i \sigma_i h_j(x_i) \leq \sup_{h \in \mathcal{H}} \sum_i \sigma_i h(x_i)$ implying $\sup_{f \in \Delta(\mathcal{H})} \sum_i \sigma_i f(x_i) \leq \sup_{h \in \mathcal{H}} \sum_i \sigma_i h(x_i)$. Furthermore, since $\mathcal{H} \subseteq \Delta(\mathcal{H})$ the opposite inequality also holds so we have an equality. The last inequality is by classic bounds on the Rademacher complexity of classes with bounded VC-dimension, due to a bound by [\[Dudley, 1978\]](#) [see e.g. [\[Hajek and Raginsky, 2021\]](#), Theorem 7.2]. By rearrangement of [Eq. \(14\)](#) we conclude that $n = O(d(\gamma_0\varepsilon)^{-2})$ as claimed, which concludes the proof sketch.

2.2 Majority-of-3

We finally describe the main ideas in our proof that Majority-of-3 AdaBoosts achieves an optimal in-expectation error of $O(d/(\gamma^2 m))$ as stated in [Corollary 3](#). We will also explain

why it is crucial for this result, that the logarithmic factors in our margin generalization bound [Theorem 1](#) are $\ln(\gamma^2 m/d)$ and not $\ln(m/d)$, $\ln(m)$ or $\ln(\gamma m/d)$. Let \mathcal{D} be the unknown data distribution over \mathcal{X} and let $t \in \{-1, 1\}^{\mathcal{X}}$ be the unknown target concept. Recall that if AdaBoost is run for $\Omega(\gamma^{-2} \ln m)$ iterations with a γ -weak learner \mathcal{W} , then it produces a voting classifier with margins $\Omega(\gamma)$ on all examples in the training sequence $\mathbf{S} \sim \mathcal{D}_t^m$. We now use an analysis idea by [\[Aden-Ali et al., 2024\]](#), used to show that the Majority-of-3 Empirical Risk Minimizers has an optimal in-expectation error for realizable PAC learning.

Consider partitioning a training sequence $\mathbf{S} \sim \mathcal{D}^{(2k-1)m}$ into $2k-1$ equal sized training sequences $\mathbf{S}_1, \dots, \mathbf{S}_{2k-1}$ (with $k=2$ for Majority-of-3 and $k=3$ for Majority-of-5) of m training examples each (rescaling m by $2k-1$ recovers the guarantees for a training sequence of size m). If we run AdaBoost on each to obtain voting classifiers $f_{\mathbf{S}_1}, \dots, f_{\mathbf{S}_{2k-1}}$, then each $f_{\mathbf{S}_i}$ has margins $\Omega(\gamma)$ on all of \mathbf{S}_i . For concreteness, let us say all margins are at least $\gamma/2$. Furthermore, for any point $x \in \mathcal{X}$, we have that if the majority vote $\text{Maj}(f_{\mathbf{S}_1}, \dots, f_{\mathbf{S}_{2k-1}})$ errs on x , where we define the majority vote as $\text{Maj}(f_{\mathbf{S}_1}, \dots, f_{\mathbf{S}_{2k-1}}) = \text{sign}(\sum_{i=1}^{2k-1} \text{sign}(f_i))$, then at least k of the voting classifiers err on x . Let us fix an $x \in \mathcal{X}$ and denote by p_x the probability that $f_{\mathbf{S}_i}$ errs on x , where the probability is over the random choice of \mathbf{S}_i . Observe that since the training sequences \mathbf{S}_i are i.i.d., this is the same probability for each \mathbf{S}_i . Moreover, by independence of the \mathbf{S}_i 's, we have that the probability that a fixed set of k of the voting classifiers all err on x is precisely p_x^k . A union bound over all $\binom{2k-1}{k} \leq 2^{2k}$ choices of k voting classifiers implies that

$$\mathbb{P}_{\mathbf{S}}[\text{Maj}(f_{\mathbf{S}_1}, \dots, f_{\mathbf{S}_{2k-1}})(x) \neq t(x)] \leq 2^{2k} p_x^k.$$

By swapping the order of expectation, we can bound the expected error of $\text{Maj}(f_{\mathbf{S}_1}, \dots, f_{\mathbf{S}_{2k-1}})$ as follows

$$\begin{aligned} \mathbb{E}_{\mathbf{S}}[\mathcal{L}_{\mathcal{D}_t}(\text{Maj}(f_{\mathbf{S}_1}, \dots, f_{\mathbf{S}_{2k-1}}))] &= \mathbb{E}_{(\mathbf{x}, t(\mathbf{x})) \sim \mathcal{D}_t} \left[\mathbb{P}_{\mathbf{S}}[\text{Maj}(f_{\mathbf{S}_1}, \dots, f_{\mathbf{S}_{2k-1}}) \neq t(\mathbf{x})] \right] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [2^{2k} p_{\mathbf{x}}^k]. \end{aligned}$$

Using the approach in [\[Aden-Ali et al., 2024\]](#), we now partition the input domain \mathcal{X} into regions R_i , such that $R_i = \{x \in \mathcal{X} : p_x \in (2^{-i-1}, 2^{-i}]\}$ for $i = 0, \dots, \infty$. Letting $\mathbb{P}[R_i]$ denote $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}}[\mathbf{x} \in R_i]$ and using the notation $\mathbb{E}_{\mathbf{x} \sim \mathcal{D}}[\cdot | R_i]$ to denote the conditional expectation, when conditioning on $\mathbf{x} \in R_i$, we now have that

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [2^{2k} p_{\mathbf{x}}^k] &= \sum_{i=0}^{\infty} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [2^{2k} p_{\mathbf{x}}^k | R_i] \mathbb{P}[R_i] \\ &\leq 2^{2k} \cdot \sum_{i=0}^{\infty} 2^{-ik} \mathbb{P}[R_i]. \end{aligned} \tag{15}$$

The goal is thus to bound $\mathbb{P}[R_i]$. For this, the key is to exploit that $p_x > 2^{-i-1}$ for $x \in R_i$. Let $m_i = \mathbb{P}[R_i]m$ denote the expected number of samples from R_i in a training sequence $\mathbf{S}_j \sim \mathcal{D}_t^m$. Since AdaBoost produces a voting classifier with margins $\Omega(\gamma)$ on all points in its training data set, it in particular has margins $\Omega(\gamma)$ on all data points in $\mathbf{S}_j \cap R_i$. We note that these samples are i.i.d. from the conditional distribution of a $\mathbf{x} \sim \mathcal{D}$, conditioned

on $\mathbf{x} \in R_i$. Let us denote this conditional distribution by $\mathcal{D}_t | R_i$. We can now invoke our new margin generalization bound in [Theorem 1](#) to conclude that

$$\mathbb{E}_{\mathbf{S}_i}[\mathcal{L}_{\mathcal{D}_t|R_i}(f_{\mathbf{S}_i})] = O\left(\frac{d\Gamma(\gamma^2 m_i/d)}{\gamma^2 m_i}\right), \quad (16)$$

with $\Gamma(x) = \ln(x) \ln^2(\ln x)$. Note that [Theorem 1](#) actually gives a high probability guarantee, which in particular implies the above guarantee on the expected error. The crucial point is that we invoke [Theorem 1](#) with the conditional distribution $\mathcal{D}_t | R_i$ and $\mathcal{L}_{\mathbf{S}_i}^{\gamma/2}(f_{\mathbf{S}_i}) = 0$ since we have all margins at least $\gamma/2$, and we have m_i samples from this distribution (in expectation). On the other hand, we have by definition of R_i that $\mathbb{E}_{\mathbf{S}_i}[\mathcal{L}_{\mathcal{D}_t|R_i}(f_{\mathbf{S}_i})] > 2^{-i-1}$. Writing $x = \gamma^2 m_i/d$ for short, we thus conclude that

$$\frac{\Gamma(x)}{x} = \Omega(2^{-i}) \Rightarrow x = O(i \ln^2(i) 2^i) \Rightarrow m_i = O\left(\frac{di \ln^2(i) 2^i}{\gamma^2}\right) \Rightarrow \mathbb{P}[R_i] = O\left(\frac{di \ln^2(i) 2^i}{\gamma^2 m}\right)$$

Inserting this in [Eq. \(15\)](#) bounds the expected error by (for constant k):

$$\mathbb{E}_{\mathbf{S}}[\mathcal{L}_{\mathcal{D}_t}(\text{Maj}(f_{\mathbf{S}_1}, \dots, f_{\mathbf{S}_{2k-1}}))] = O\left(\sum_{i=0}^{\infty} \frac{2^{-ik} i \ln^2(i) 2^i d}{\gamma^2 m}\right). \quad (17)$$

Inserting $k = 2$ (corresponding to Majority-of-3) gives the desired $O(d/(\gamma^2 m))$ as the $2^{-ik} = 2^{-2i}$ decreases fast enough to cancel the $i \ln^2(i) 2^i$ factors.

Failure of Previous Bounds. Let us now discuss why the $\Gamma(\gamma^2 m/d)$ factor is crucial compared to $\ln(m)$, $\ln(\gamma m/d)$ and $\ln(m/d)$ factors in the above analysis. Consider again [Eq. \(16\)](#) and assume for simplicity that the generalization bound had instead given us

$$\mathbb{E}_{\mathbf{S}_i}[\mathcal{L}_{\mathcal{D}_t|R_i}(f_{\mathbf{S}_i})] = O\left(\frac{d \ln(\gamma m_i/d)}{\gamma^2 m_i}\right), \quad (18)$$

i.e. a slightly suboptimal dependency on γ inside the $\ln(\cdot) \ln^2(\ln(\cdot))$. We would then get the inequality

$$\frac{d \ln(\gamma m_i/d)}{\gamma^2 m_i} = \Omega(2^{-i}).$$

Now again letting $x = \gamma^2 m_i/d$ then the above can be shown to imply

$$\frac{\ln(x/\gamma)}{x} = \Omega(2^{-i}) \Rightarrow x = O(\ln(2^i/\gamma) 2^i) \Rightarrow \mathbb{P}[R_i] = O\left(\frac{d \ln(2^i/\gamma) 2^i}{\gamma^2 m}\right).$$

Now plugging this bound $\mathbb{P}[R_i]$ into [Eq. \(15\)](#) (with $k = 2$), we get the following error bounds for $\mathbb{E}_{\mathbf{S}}[\mathcal{L}_{\mathcal{D}_t}(\text{Maj}(f_{\mathbf{S}_1}, f_{\mathbf{S}_2}, f_{\mathbf{S}_3}))]$ of

$$O\left(\sum_{i=0}^{\infty} \frac{2^{-2i} d \ln(2^i/\gamma) 2^i}{\gamma^2 m}\right) = O\left(\frac{d \ln(1/\gamma)}{\gamma^2 m}\right).$$

The obtained error bound of $\mathbb{E}_{\mathbf{S}}[\mathcal{L}_{\mathcal{D}_t}(\text{Maj}(f_{\mathbf{S}_1}, f_{\mathbf{S}_2}, f_{\mathbf{S}_3}))]$ thus has a superfluous $\ln(1/\gamma)$ factor.

These shortcomings of previous margins bounds used in the above analysis of the expected error $\mathbb{E}_{\mathbf{S}}[\mathcal{L}_{\mathcal{D}_t}(\text{Maj}(f_{\mathbf{S}_1}, f_{\mathbf{S}_2}, f_{\mathbf{S}_3}))]$, adding a superfluous $\ln(1/\gamma)$ -factor, is precisely the reason why previous work needed a Majority-of-5. Being unable to use the margin generalization bounds with sub-optimal $\ln(\cdot)$ factors, the work [Høgsgaard et al., 2024] instead relied on the much weaker guarantee that any voting classifier f with margins γ has $\mathcal{L}_{\mathcal{D}_t}(f) = O(\sqrt{d/(\gamma^2 m)})$, where this bound can be obtained by following the steps of [Schapire and Freund, 2012] [page 107-111] and using the stronger bound on the Rademacher complexity for a function class with VC-dimension d of $O(\sqrt{d/m})$ due to [Dudley, 1978] [See e.g. Theorem 7.2 [Hajek and Raginsky, 2021]], instead of the weaker $O(\sqrt{d \ln(m/d)/m})$ used in [Schapire and Freund, 2012]. This bound has the right behaviour for $m \approx d/\gamma^2$ unlike the bounds with sub-optimal logarithmic factors and results in the guarantee $\mathbb{P}[R_i] = O(2^{2i})$ instead of $O(i \ln^2(i) 2^i)$. This needs $k = 3$ for 2^{-ik} to dominate 2^{2i} in $\sum_{i=0}^{\infty} 2^{-ik} 2^{2i}$ from Eq. (17), whereas it suffices for us with $k = 2$ since we only need to bound $\sum_{i=0}^{\infty} 2^{-ik} i \ln^2(i) 2^i$, leading to the more natural Majority-of-3 instead of Majority-of-5.

Organization of Paper. The following sections are organized as follows. In Section 3 we prove the margin generalization bound in Eq. (8). In Section 4 we prove the bound on the covering number of the clipped function class $\Delta(\mathcal{H})_{[\gamma]}$. In Section 5, we put together the results from Section 3 and Section 4 to prove the main margin generalization bound in Theorem 1. In Section 6 we prove the main result on Majority-of-3 in Corollary 3.

3 Upper bound

In this section we prove Lemma 4, which is used in a union bound over suitable γ_0, γ_1 and τ_0 and τ_1 , giving the margin bound in Theorem 1 for all $f \in \Delta(\mathcal{H})$ and margins $0 < \gamma \leq 1$.

Lemma 4. *For a hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$, margin thresholds $0 < \gamma_0 \leq \gamma_1 \leq 1$, error thresholds $0 \leq \tau_0 \leq \tau_1$, failure parameter $0 < \delta < 1$, we have that*

$$\begin{aligned} & \mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m} \left[\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^{\gamma}(f) \in [\tau_0, \tau_1], \mathcal{L}_{\mathcal{D}}(f) \geq \tau_1 + 64 \left(\sqrt{\frac{\tau_1 \cdot 2 \ln(\frac{\epsilon}{\delta})}{m}} + \frac{2 \ln(\frac{\epsilon}{\delta})}{m} \right) \right] \\ & \leq \delta \cdot \sup_{X \in \mathcal{X}^{2m}} \mathcal{N}_{\infty}(X, \Delta(\mathcal{H})_{[2\gamma_1]}, \frac{\gamma_0}{2}). \end{aligned}$$

Proof. We first notice that if $\tau_1 > 1$ then we are done by $\mathcal{L}_{\mathcal{D}}(f) \leq 1$, thus we assume for the remaining of the proof that $\tau_1 \leq 1$. For ease of notation, let $\beta = \left(\sqrt{\tau_1 \cdot 2 \ln(\frac{\epsilon}{\delta})/m} + 2 \ln(\frac{\epsilon}{\delta})/m \right)$ in the following. We start by showing that

$$\begin{aligned} & \mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m} \left[\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^{\gamma}(f) \in [\tau_0, \tau_1], \mathcal{L}_{\mathcal{D}}(f) \geq \tau + 64\beta \right] \tag{19} \\ & \leq 2 \mathbb{P}_{\mathbf{S}, \mathbf{S}' \sim \mathcal{D}^m} \left[\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^{\gamma}(f) \in [\tau_0, \tau_1], \mathcal{L}_{\mathbf{S}'}(f) \geq \tau_1 + 32\beta \right]. \end{aligned}$$

To this end, we first notice that if $f \in \Delta(\mathcal{H})$ is such that $\mathcal{L}_{\mathcal{D}}(f) > \frac{2\ln(\frac{\epsilon}{\delta})}{m}$, then by Chernoff, we have that

$$\mathbb{P}_{\mathbf{S}'} \left[\mathcal{L}_{\mathbf{S}'}(f) \leq \left(1 - \sqrt{\frac{2\ln(\frac{\epsilon}{\delta})}{m \mathcal{L}_{\mathcal{D}}(f)}} \right) \mathcal{L}_{\mathcal{D}}(f) \right] \leq \exp \left(-\frac{2\ln(\frac{\epsilon}{\delta})}{2} \right) \leq \frac{\delta}{e}.$$

This implies that with probability at least $1 - \frac{\delta}{e}$ over \mathbf{S}' , we have

$$\mathcal{L}_{\mathbf{S}'}(f) \geq \mathcal{L}_{\mathcal{D}}(f) - \sqrt{\frac{\mathcal{L}_{\mathcal{D}}(f) \cdot 2\ln(\frac{\epsilon}{\delta})}{m}}.$$

Now, for $a > 0$, $x - \sqrt{ax}$ is increasing for $x \geq a/4$, since it has derivative $1 - \frac{a}{2\sqrt{ax}}$. Thus, for $\mathcal{L}_{\mathcal{D}}(f) \geq \tau_1 + 64\beta = \tau_1 + 64 \left(\sqrt{\frac{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} + \frac{2\ln(\frac{\epsilon}{\delta})}{m} \right) \geq \frac{2\ln(\frac{\epsilon}{\delta})}{m}$, we have by the above with $a = \frac{2\ln(\epsilon/\delta)}{m}$, that with probability at least $1 - \frac{\delta}{e}$ over \mathbf{S}' , we have

$$\begin{aligned} \mathcal{L}_{\mathbf{S}'}(f) &\geq \mathcal{L}_{\mathcal{D}}(f) - \sqrt{\frac{\mathcal{L}_{\mathcal{D}}(f) \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} \\ &\geq \tau_1 + 64 \left(\sqrt{\frac{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} + \frac{2\ln(\frac{\epsilon}{\delta})}{m} \right) - \sqrt{\frac{\left(\tau_1 + 64 \left(\sqrt{\frac{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} + \frac{2\ln(\frac{\epsilon}{\delta})}{m} \right) \right) \cdot 2\ln(\frac{\epsilon}{\delta})}{m}}, \end{aligned} \quad (20)$$

where the square root term can be upper bounded using $a + \sqrt{ab} + b \leq 2(a + b)$ and $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b > 0$, as follows

$$\begin{aligned} &\sqrt{\frac{\left(\tau_1 + 64 \left(\sqrt{\frac{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} + \frac{2\ln(\frac{\epsilon}{\delta})}{m} \right) \right) \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} \leq \sqrt{\frac{64 \left(\tau_1 + \sqrt{\frac{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} + \frac{2\ln(\frac{\epsilon}{\delta})}{m} \right) \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} \\ &\leq \sqrt{\frac{2 \cdot 64 \left(\tau_1 + \frac{2\ln(\frac{\epsilon}{\delta})}{m} \right) \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} \leq \sqrt{2 \cdot 64} \left(\sqrt{\frac{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} + \frac{2\ln(\frac{\epsilon}{\delta})}{m} \right). \end{aligned}$$

Now using the above upper bound combined with [Eq. \(20\)](#) we conclude that

$$\begin{aligned} \mathcal{L}_{\mathbf{S}'}(f) &\geq \tau_1 + (64 - \sqrt{2 \cdot 64}) \sqrt{\frac{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} + (64 - \sqrt{2 \cdot 64}) \frac{2\ln(\frac{\epsilon}{\delta})}{m} \\ &\geq \tau_1 + 32 \left(\sqrt{\frac{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})}{m}} + \frac{2\ln(\frac{\epsilon}{\delta})}{m} \right) = \tau + 32\beta, \end{aligned} \quad (21)$$

where the last equality uses that $\beta = \left(\sqrt{\tau_1 \cdot 2\ln(\frac{\epsilon}{\delta})/m} + 2\ln(\frac{\epsilon}{\delta})/m \right)$. Thus, we conclude

by the above and the law of total probability that

$$\begin{aligned}
& \mathbb{P}_{\mathbf{S}, \mathbf{S}' \sim \mathcal{D}^m} [\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0, \tau_1], \mathcal{L}_{\mathbf{S}'}(f) \geq \tau_1 + 32\beta] \quad (22) \\
&= \mathbb{P}_{\mathbf{S}, \mathbf{S}' \sim \mathcal{D}^m} \left[\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0, \tau_1], \mathcal{L}_{\mathbf{S}'}(f) \geq \tau_1 + 32\beta \right. \\
&\quad \left. | \exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0, \tau_1], \mathcal{L}_{\mathcal{D}}(f) \geq \tau_1 + 64\beta \right] \\
&\quad \cdot \mathbb{P}_{\mathbf{S}} [\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0, \tau_1], \mathcal{L}_{\mathcal{D}}(f) \geq \tau_1 + 64\beta] \\
&\geq \left(1 - \frac{\delta}{e}\right) \mathbb{P}_{\mathbf{S}} [\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0, \tau_1], \mathcal{L}_{\mathcal{D}}(f) \geq \tau_1 + 64\beta]
\end{aligned}$$

where the second to last inequality follows by the law of total probability and the last inequality by Eq. (21), whereby we have shown Eq. (19).

We now bound the probability of the expression on the second line of Eq. (19) by

$$\sup_{X \in \mathcal{X}^{2m}} \mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[2\gamma_1]}, \frac{\gamma_0}{6}) \frac{\delta}{e},$$

which would conclude the proof.

Now consider any γ such that $\gamma_0 \leq \gamma \leq \gamma_1$. We now recall that for a function f , we defined $f_{[2\gamma_1]}$ as follows:

$$f_{[2\gamma_1]}(x) = \begin{cases} 2\gamma_1 & \text{if } f(x) > 2\gamma_1 \\ -2\gamma_1 & \text{if } f(x) < -2\gamma_1 \\ f(x) & \text{else .} \end{cases}$$

Thus, $f_{[2\gamma_1]}(x)$ has the same sign as $f(x)$ and is equal to $f(x)$, when $|f(x)| \leq 2\gamma_1$. Using the above observations we conclude by $y \in \{-1, 1\}$ and $\gamma \leq \gamma_1$ that $\mathbb{1}\{f(x)y \leq \gamma\} = \mathbb{1}\{f_{[2\gamma_1]}(x)y \leq \gamma\}$ and $\mathbb{1}\{f(x)y \leq 0\} = \mathbb{1}\{f_{[2\gamma_1]}(x)y \leq 0\}$. Thus, we get that

$$\begin{aligned}
& \mathbb{P}_{\mathbf{S}, \mathbf{S}' \sim \mathcal{D}^m} [\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \leq \tau_1, \mathcal{L}_{\mathbf{S}'}(f) \geq \tau_1 + 32\beta] \quad (23) \\
&= \mathbb{P}_{\mathbf{S}, \mathbf{S}' \sim \mathcal{D}^m} [\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H})_{[2\gamma_1]} : \mathcal{L}_{\mathbf{S}}^\gamma(f) \leq \tau_1, \mathcal{L}_{\mathbf{S}'}(f) \geq \tau_1 + 32\beta].
\end{aligned}$$

By \mathbf{S} and \mathbf{S}' being i.i.d., we may see \mathbf{S} and \mathbf{S}' as being drawn in the following way: First, we draw $\bar{\mathbf{S}} \sim \mathcal{D}^{2m}$ and then let \mathbf{S} be created by drawing m times without replacement from $\bar{\mathbf{S}}$ and $\mathbf{S}' = \bar{\mathbf{S}} \setminus \mathbf{S}$. We denote this way of drawing \mathbf{S} and \mathbf{S}' from $\bar{\mathbf{S}}$ as $\mathbf{S}, \mathbf{S}' \sim \bar{\mathbf{S}}$. We then have that,

$$\begin{aligned}
& \mathbb{P}_{\mathbf{S}, \mathbf{S}' \sim \mathcal{D}^m} [\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H})_{[2\gamma_1]} : \mathcal{L}_{\mathbf{S}}^\gamma(f) \leq \tau_1, \mathcal{L}_{\mathbf{S}'}(f) \geq \tau_1 + 32\beta] \quad (24) \\
&\leq \mathbb{P}_{\bar{\mathbf{S}} \sim \mathcal{D}^{2m}} \left[\mathbb{P}_{\mathbf{S}, \mathbf{S}' \sim \bar{\mathbf{S}}} [\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H})_{[2\gamma_1]} : \mathcal{L}_{\mathbf{S}}^\gamma(f) \leq \tau_1, \mathcal{L}_{\mathbf{S}'}(f) \geq \tau_1 + 32\beta] \right] \\
&\leq \sup_{Z \in (\mathcal{X} \times \{-1, 1\})^{2m}} \mathbb{P}_{\mathbf{S}, \mathbf{S}' \sim Z} [\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H})_{[2\gamma_1]} : \mathcal{L}_{\mathbf{S}}^\gamma(f) \leq \tau_1, \mathcal{L}_{\mathbf{S}'}(f) \geq \tau_1 + 32\beta].
\end{aligned}$$

Let now $Z = (X, Y) \in (\mathcal{X} \times \{-1, 1\})^{2m}$ and N be a $\frac{\gamma_0}{2}$ -cover for $\Delta(\mathcal{H})_{\lceil 2\gamma_1 \rceil}$ on X with respect to ∞ -norm of minimal size. That is, for any $f \in \Delta(\mathcal{H})_{\lceil 2\gamma_1 \rceil}$, we have

$$\min_{f' \in N} \max_{x \in X} |f(x) - f'(x)| \leq \frac{\gamma_0}{2}.$$

and any other cover satisfying the above has size at least $|N|$. We notice that this implies that if γ is such that $\gamma_0 \leq \gamma \leq \gamma_1$ and $f \in \Delta(\mathcal{H})_{\lceil 2\gamma_1 \rceil}$, then for the function f' closest to f in N , i.e. $f' = \operatorname{argmin}_{f' \in N} \max_{x \in X} |f(x) - f'(x)|$, with ties broken arbitrarily, it holds for any $(x, y) \in Z$ with $f(x)y \leq 0$ that $f'(x)y = f(x)y + (f'(x) - f(x))y \leq 0 + \frac{\gamma_0}{2} \leq \frac{\gamma_0}{2}$. Similarly for any $(x, y) \in Z$ such that $f(x)y > \gamma$, we have $f'(x)y = f(x)y + (f'(x) - f(x))y > \gamma - \frac{\gamma_0}{2} > \frac{\gamma_0}{2}$, where the last inequality follows from $\gamma \geq \gamma_0$. Thus, we conclude that for any γ such that $\gamma_0 \leq \gamma \leq \gamma_1$ and $f \in \Delta(\mathcal{H})_{\lceil 2\gamma_1 \rceil}$, there exists $f' \in N$ such that $\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f') \geq \mathcal{L}_{\mathbf{S}'}^\gamma(f)$ and that $\mathcal{L}_{\mathbf{S}}^\gamma(f) \geq \mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f')$. By the above and the union bound, we conclude that

$$\begin{aligned} & \mathbb{P}_{\mathbf{s}, \mathbf{s}' \sim Z} \left[\exists \gamma \in [\gamma_0, \gamma_1], \exists f \in \Delta(\mathcal{H})_{\lceil 2\gamma_1 \rceil} : \mathcal{L}_{\mathbf{S}}^\gamma(f) \leq \tau_1, \mathcal{L}_{\mathbf{S}'}^\gamma(f) \geq \tau_1 + 32\beta \right] \\ & \leq \mathbb{P}_{\mathbf{s}, \mathbf{s}' \sim Z} \left[\exists f \in N : \mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \leq \tau_1, \mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \geq \tau_1 + 32\beta \right] \\ & \leq \sum_{f \in N} \mathbb{P}_{\mathbf{s}, \mathbf{s}' \sim Z} \left[\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \leq \tau_1, \mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \geq \tau_1 + 32\beta \right]. \end{aligned} \quad (25)$$

We now show that for each $f \in N$, each term/probability in the sum is at most $\frac{\delta}{e}$, so the sum is at most $\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{\lceil 2\gamma_1 \rceil}, \frac{\gamma_0}{2}) \frac{\delta}{e}$. Which implies by Eq. (24), that Eq. (23) is bounded by $\sup_{X \in \mathcal{X}^{2m}} \mathcal{N}_\infty(X, \Delta(\mathcal{H})_{\lceil 2\gamma_1 \rceil}, \frac{\gamma_0}{2}) \frac{\delta}{e}$ as claimed below Eq. (22) and would conclude the proof. To the end of showing the above let $f \in N$ for now.

Recall that $\beta = \left(\sqrt{\frac{\tau_1 \cdot 2 \ln(\frac{e}{\delta})}{m}} + \frac{2 \ln(\frac{e}{\delta})}{m} \right)$. Let 2μ denote the fraction of points in Z where f has less than $\gamma_0/2$ -margin, i.e.,

$$2\mu = \mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) + \mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) = \left| \{(x, y) \in Z : f(x)y \leq \frac{\gamma_0}{2}\} \right| / m.$$

We notice that μ is the expectation of $\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f)$ and that for the probability in Eq. (25) to be nonzero, it must be the case that $\mu \geq 32 \frac{\ln(\frac{e}{\delta})}{m}$. Thus, we assume this is the case. We notice that if $\mu > \frac{2 \ln(\frac{e}{\delta})}{m}$, we have by Chernoff (which due to [Hoeffding, 1963][see section 6] also holds when sampling without replacement) that

$$\mathbb{P}_{\mathbf{s}, \mathbf{s}' \sim Z} \left[\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \leq \left(1 - \sqrt{\frac{2 \ln(\frac{e}{\delta})}{\mu m}} \right) \mu \right] \leq \exp \left(-\frac{2 \ln(\frac{e}{\delta}) \mu m}{2\mu m} \right) \leq \frac{\delta}{e}$$

Thus, with probability at least $1 - \frac{\delta}{e}$ over \mathbf{S}, \mathbf{S}' , we have by definition of μ that

$$\begin{aligned}
\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) &\geq \mu - \sqrt{\frac{\mu 2 \ln\left(\frac{e}{\delta}\right)}{m}} = \frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) + \mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f)}{2} - \sqrt{\frac{\left(\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) + \mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f)\right) \ln\left(\frac{e}{\delta}\right)}{m}} \\
&\geq \frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) + \mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f)}{2} - \sqrt{\frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{e}{\delta}\right)}{m}} - \sqrt{\frac{\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{e}{\delta}\right)}{m}} \quad (\text{by } \sqrt{a+b} \leq \sqrt{a} + \sqrt{b}) \\
&\Rightarrow \frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f)}{2} + \sqrt{\frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{e}{\delta}\right)}{m}} \geq \frac{\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f)}{2} - \sqrt{\frac{\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{e}{\delta}\right)}{m}}. \quad (\text{by rearrangement})
\end{aligned} \tag{26}$$

Recall that $\beta = \left(\sqrt{\frac{\tau_1 \cdot 2 \ln\left(\frac{e}{\delta}\right)}{m}} + \frac{2 \ln\left(\frac{e}{\delta}\right)}{m}\right)$. We now show that if $\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \geq \tau_1 + 32\beta = \tau_1 + 32\left(\sqrt{\frac{\tau_1 \cdot 2 \ln\left(\frac{e}{\delta}\right)}{m}} + \frac{2 \ln\left(\frac{e}{\delta}\right)}{m}\right)$ and $\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \leq \tau_1$ then we have

$$\frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f)}{2} + \sqrt{\frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{e}{\delta}\right)}{m}} < \frac{\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f)}{2} - \sqrt{\frac{\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{e}{\delta}\right)}{m}}. \tag{27}$$

This implies that the event $\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \geq \tau_1 + 32\beta$ and $\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \leq \tau_1$ is in the complement of the event $\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \geq \mu - \sqrt{\frac{\mu 2 \ln\left(\frac{e}{\delta}\right)}{m}}$, which we just argued happens with probability at least $1 - \frac{\delta}{e}$ over \mathbf{S}, \mathbf{S}' . Thus we now show that $\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \geq \tau_1 + 32\beta$ and $\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \leq \tau_1$ implies [Eq. \(27\)](#), which would establish that each term in [Eq. \(25\)](#) is no more than $\frac{\delta}{e}$ and thereby conclude the proof.

If $\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \leq \tau_1$, then by $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, we have that

$$\frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f)}{2} + \sqrt{\frac{\mathcal{L}_{\mathbf{S}}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{e}{\delta}\right)}{m}} \leq \frac{\tau_1}{2} + \sqrt{\frac{\tau_1 \ln\left(\frac{e}{\delta}\right)}{m}} \tag{28}$$

Now if $\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \geq \tau_1 + 32\beta \Rightarrow \tau_1 + 32\left(\sqrt{\frac{\tau_1 \cdot 2 \ln\left(\frac{e}{\delta}\right)}{m}} + \frac{2 \ln\left(\frac{e}{\delta}\right)}{m}\right)$ and using that we argued earlier that $x - \sqrt{ax}$ is increasing for $x \geq \frac{a}{4}$, we get that

$$\begin{aligned}
&\frac{\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f)}{2} - \sqrt{\frac{\mathcal{L}_{\mathbf{S}'}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{e}{\delta}\right)}{m}} \\
&\geq \frac{\tau_1}{2} + 16\left(\sqrt{\frac{\tau_1 \cdot 2 \ln\left(\frac{e}{\delta}\right)}{m}} + \frac{2 \ln\left(\frac{e}{\delta}\right)}{m}\right) - \sqrt{\frac{\left(\tau_1 + 32\left(\sqrt{\frac{\tau_1 \cdot 2 \ln\left(\frac{e}{\delta}\right)}{m}} + \frac{2 \ln\left(\frac{e}{\delta}\right)}{m}\right)\right) \ln\left(\frac{e}{\delta}\right)}{m}}.
\end{aligned} \tag{29}$$

Now upper bounding the second term using by $a + \sqrt{ab} + b \leq 2(a+b)$ and $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$,

$$\begin{aligned} & \sqrt{\frac{\left(\tau_1 + 32\left(\sqrt{\frac{\tau_1 \cdot 2 \ln(\frac{\epsilon}{\delta})}{m}} + \frac{2 \ln(\frac{\epsilon}{\delta})}{m}\right)\right) \ln\left(\frac{\epsilon}{\delta}\right)}{m}} \leq \sqrt{\frac{32\left(\tau_1 + \sqrt{\frac{\tau_1 \cdot 2 \ln(\frac{\epsilon}{\delta})}{m}} + \frac{2 \ln(\frac{\epsilon}{\delta})}{m}\right) \ln\left(\frac{\epsilon}{\delta}\right)}{m}} \\ & \leq \sqrt{\frac{64\left(\tau_1 + \frac{2 \ln(\frac{\epsilon}{\delta})}{m}\right) \ln\left(\frac{\epsilon}{\delta}\right)}{m}} \leq \sqrt{64} \sqrt{\frac{\tau_1 \ln\left(\frac{\epsilon}{\delta}\right)}{m}} + \sqrt{128} \frac{\ln\left(\frac{\epsilon}{\delta}\right)}{m}, \end{aligned}$$

we conclude from Eq. (29) that

$$\frac{\mathcal{L}_{\mathcal{S}'}^{\frac{\gamma_0}{2}}(f)}{2} - \sqrt{\frac{\mathcal{L}_{\mathcal{S}'}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{\epsilon}{\delta}\right)}{m}} \geq \frac{\tau_1}{2} + (16\sqrt{2} - \sqrt{64}) \sqrt{\frac{\tau_1 \cdot \ln\left(\frac{\epsilon}{\delta}\right)}{m}} + (32 - \sqrt{128}) \frac{\ln\left(\frac{\epsilon}{\delta}\right)}{m}. \quad (30)$$

Thus, combining Eq. (28) and Eq. (30) and using that $16\sqrt{2} - \sqrt{64} \geq 14$ and $32 - \sqrt{128} \geq 20$, we conclude that if $\mathcal{L}_{\mathcal{S}'}^{\frac{\gamma_0}{2}}(f) \geq \tau_1 + 32\beta = \tau_1 + 32\left(\sqrt{\frac{\tau_1 \cdot 2 \ln(\frac{\epsilon}{\delta})}{m}} + \frac{2 \ln(\frac{\epsilon}{\delta})}{m}\right)$ and $\mathcal{L}_{\mathcal{S}}^{\frac{\gamma_0}{2}}(f) \leq \tau_1$, then it implies that

$$\frac{\mathcal{L}_{\mathcal{S}'}^{\frac{\gamma_0}{2}}(f)}{2} - \sqrt{\frac{\mathcal{L}_{\mathcal{S}'}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{\epsilon}{\delta}\right)}{m}} > \frac{\mathcal{L}_{\mathcal{S}}^{\frac{\gamma_0}{2}}(f)}{2} + \sqrt{\frac{\mathcal{L}_{\mathcal{S}}^{\frac{\gamma_0}{2}}(f) \ln\left(\frac{\epsilon}{\delta}\right)}{m}},$$

which as argued earlier concludes the proof. \square

4 Bound on infinity cover

In this section, we prove Lemma 5, where from Corollary 6 follows and gives us the covering bound needed as described in the proof sketch. To this end recall that we for a function class \mathcal{F} and a point set $X = \{x_1, \dots, x_m\} \subset \mathcal{X}$ of size m and precision γ define $\mathcal{N}_\infty(X, \mathcal{F}, \gamma)$ to be the minimal size of a γ -cover in infinity norm of \mathcal{F} on X , i.e. the smallest size of a set of functions N satisfying $\min_{f' \in N} \max_{x \in X} |f(x) - f'(x)| \leq \gamma$. Furthermore, we will use $\text{Ln}(x) = \ln(\max(x, e))$ for the truncated natural logarithm. We now state Lemma 5 and Corollary 6, and show how Lemma 5 implies Corollary 6.

Lemma 5. *There exist universal constants $c' \geq c > 0$ and $c' \geq 1$ such that: For a point set $X = \{x_1, \dots, x_m\} \subseteq \mathcal{X}$ of size m , $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ a hypothesis class of VC-dimension d , $0 < \epsilon < 1$, $0 < \alpha < \frac{1}{2}$ and $0 < \gamma \leq 1$ such that $\gamma^2 \geq \frac{cd}{\alpha^2 \epsilon^2 m}$, we have*

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}, \alpha\gamma)) \leq \frac{c'd}{\alpha^2 \epsilon^2 \gamma^2} \text{Ln}^{1+\epsilon}\left(\frac{8\alpha m \epsilon^2 \gamma^2}{cd}\right). \quad (31)$$

Corollary 6. *There exist universal constants $c' \geq c > 0$ and $c' \geq 1$ such that: For a point set $X = \{x_1, \dots, x_m\} \subseteq \mathcal{X}$ of size m , $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ a hypothesis class of VC-dimension d , $0 < \alpha < \frac{1}{2}$ and $0 < \gamma \leq 1$ such that $\gamma^2 \geq \frac{4cd \text{Ln}^2(\text{Ln}(\frac{8\alpha m \gamma^2}{cd}))}{\alpha^2 m}$, we have*

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}, \alpha\gamma)) \leq \frac{c'd}{\alpha^2 \gamma^2} \text{Ln}^2(\text{Ln}(\frac{8\alpha m \gamma^2}{cd})) \text{Ln}\left(\frac{8\alpha m \gamma^2}{cd}\right) = \frac{c'd}{\alpha^2 \gamma^2} \Gamma\left(\frac{8\alpha m \gamma^2}{cd}\right) \quad (32)$$

where $\Gamma(x) = \text{Ln}^2(\text{Ln}(x)) \text{Ln}(x)$.

Proof of Corollary 6. Let c', c be the constants from Lemma 5, and furthermore consider $\gamma^2 \geq \frac{4cd \text{Ln}^2(\text{Ln}(\frac{8\alpha m \gamma^2}{cd}))}{\alpha^2 m}$. To the end of invoking Lemma 5, notice that for $\varepsilon = 1/(2 \text{Ln}(\text{Ln}(\frac{8\alpha m \gamma^2}{cd}))) \leq 1/2$, we have that $\frac{cd}{\alpha^2 \varepsilon^2 m} \leq \frac{4cd \text{Ln}^2(\text{Ln}(\frac{8\alpha m \gamma^2}{cd}))}{\alpha^2 m} \leq \gamma^2$, and we thus may invoke Lemma 5 with this ε . We notice that for this ε we have that

$$\text{Ln}^\varepsilon\left(\frac{8\alpha m \varepsilon^2 \gamma^2}{cd}\right) = \exp\left(\varepsilon \ln\left(\text{Ln}\left(\frac{8\alpha m \varepsilon^2 \gamma^2}{cd}\right)\right)\right) \leq e,$$

where we in the inequality used the value of ε and that $\text{Ln}\left(\frac{8\alpha m \varepsilon^2 \gamma^2}{cd}\right) \leq \text{Ln}\left(\frac{8\alpha m \gamma^2}{cd}\right)$. Thus we conclude that

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}, \alpha\gamma)) \leq \frac{c'd}{\alpha^2 \varepsilon^2 \gamma^2} \text{Ln}^{1+\varepsilon}\left(\frac{8\alpha m \varepsilon^2 \gamma^2}{cd}\right) \quad (33)$$

$$\leq \frac{4ec'd}{\alpha^2 \gamma^2} \text{Ln}^2\left(\text{Ln}\left(\frac{8\alpha m \gamma^2}{cd}\right)\right) \text{Ln}\left(\frac{8\alpha m \gamma^2}{cd}\right). \quad (34)$$

whereby Corollary 6 follows from redefining c' to be $4ec'$. \square

We now move on to prove Lemma 5. To prove Lemma 5, we need the following two lemmas. The first lemma bounds the infinity cover of $\Delta(\mathcal{H})_{[\gamma]}$ in terms of the fat shattering dimension of $\Delta(\mathcal{H})_{[\gamma]}$.

Lemma 7. *There exists universal constants $\check{C} \geq 1$ and $\check{c} > 0$ such that: For $X = \{x_1, \dots, x_m\} \subseteq \mathcal{X}$ a point set of size m , $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ a hypothesis class, $0 < \varepsilon < 1$, $0 < \alpha < \frac{1}{2}$ and $0 < \gamma \leq 1$ if $\text{fat}_{\check{c}\alpha\varepsilon\gamma}(\Delta(\mathcal{H})_{[\gamma]}) \leq m$, then we have, for $d = \text{fat}_{\check{c}\alpha\varepsilon\gamma}(\Delta(\mathcal{H})_{[\gamma]})$, that*

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}, \alpha\gamma)) \leq \check{C}d \ln^{1+\varepsilon}\left(\frac{2m}{d\alpha}\right). \quad (35)$$

We also postpone the proof of Lemma 7 and give it after the proof of Lemma 5. The next lemma gives a bound on the fat shattering dimension of $\Delta(\mathcal{H})_{[\gamma]}$, in terms of γ and the VC-dimension of the hypothesis class \mathcal{H} .

Lemma 8. *There exists a universal constant $C \geq 1$ such that: For $1 \geq \gamma > 0$, $\beta > 0$, and $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ a hypothesis class with VC dimension d , we have that $\text{fat}_\beta(\Delta(\mathcal{H})_{[\gamma]}) \leq \frac{Cd}{\beta^2}$*

We postpone the proof of Lemma 8 for now and give the proof of Lemma 5 by combining Lemma 7 and Lemma 8.

Proof of Lemma 5. In what follows let, $\check{C} \geq 1$ and $\check{c} > 0$ be the universal constants from Lemma 7, and $C \geq 1$ be the universal constant from Lemma 8. Furthermore, let c denote the universal constant $c = \frac{16C}{\check{c}^2}$ and $c' = \frac{16\check{C}C}{\check{c}^2}$ in Lemma 5. Since $\check{C} \geq 1$ we have $c' \geq c$.

By Lemma 8, we have that $d' = \text{fat}_{\check{c}\alpha\varepsilon\gamma}(\Delta(\mathcal{H})_{[\gamma]})$ is upper bounded by $\frac{4Cd}{\check{c}^2 \alpha^2 \varepsilon^2 \gamma^2}$. Thus, we have that $d' = \text{fat}_{\check{c}\alpha\varepsilon\gamma}(\Delta(\mathcal{H})_{[\gamma]}) \leq m$, by the assumption $\gamma^2 \geq \frac{cd}{\alpha^2 \varepsilon^2 m}$ and $c = \frac{16C}{\check{c}^2}$, and we may invoke Lemma 7 to obtain

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}, \alpha\gamma)) \leq \check{C}d' \ln^{1+\varepsilon}\left(\frac{2m}{d'\alpha}\right).$$

Next, we analyze the term $d' \ln^{1+\varepsilon} \left(\frac{2m}{d'\alpha} \right)$. We now notice that the function $x \ln^{1+\varepsilon} \left(\frac{a}{x} \right)$ has derivative $\ln^\varepsilon \left(\frac{a}{x} \right) \left(\ln \left(\frac{a}{x} \right) - 1 - \varepsilon \right)$, thus the derivative is positive for $\ln \left(\frac{a}{x} \right) \geq 1 + \varepsilon$ or equivalently $\frac{a}{\exp(1+\varepsilon)} \geq x$, whereby we conclude that $x \ln^{1+\varepsilon} \left(\frac{a}{x} \right)$ is increasing for $\frac{a}{\exp(1+\varepsilon)} \geq x$. Now using this with $a = \frac{2m}{\alpha} \geq \frac{64Cd}{\check{c}^2 \alpha^2 \varepsilon^2 \gamma^2}$, where we used that $\gamma^2 \geq \frac{cd}{\alpha^2 \varepsilon^2 m}$, $c = \frac{16C}{\check{c}^2}$ and $\alpha < \frac{1}{2}$, and having the upper bound on $d' = \text{fat}_{\check{c}\alpha\varepsilon\gamma}(\Delta(\mathcal{H})_{[\gamma]})$ of $\frac{4Cd}{\check{c}^2 \alpha^2 \varepsilon^2 \gamma^2} \leq \frac{\frac{32Cd}{\check{c}^2 \alpha^2 \varepsilon^2 \gamma^2}}{\exp(1+\varepsilon)}$, since $\frac{32}{\exp(1+\varepsilon)} \geq 4$, we conclude that $d' = \text{fat}_{\check{c}\alpha\varepsilon\gamma}(\Delta(\mathcal{H})_{[\gamma]}) \leq \frac{4Cd}{\check{c}^2 \alpha^2 \varepsilon^2 \gamma^2} \leq \frac{a}{\exp(1+\varepsilon)}$. This then implies by the above argued monotonicity of $x \ln^{1+\varepsilon} \left(\frac{a}{x} \right)$, with $a = \frac{2m}{\alpha}$ that

$$\check{C}d' \ln^{1+\varepsilon} \left(\frac{2m}{d'\alpha} \right) \leq \check{C} \frac{4Cd}{\check{c}^2 \alpha^2 \varepsilon^2 \gamma^2} \ln^{1+\varepsilon} \left(\frac{2m}{\alpha} \frac{\check{c}^2 \alpha^2 \varepsilon^2 \gamma^2}{4Cd} \right) \leq \frac{c'd}{\alpha^2 \varepsilon^2 \gamma^2} \ln^{1+\varepsilon} \left(\frac{8\alpha m \varepsilon^2 \gamma^2}{cd} \right),$$

where the last equivalently follows from $c = \frac{16C}{\check{c}^2}$ and $c' = \frac{16\check{C}C}{\check{c}^2}$ (we notice that if c' is not at least 1, we could set it equal to 1 and still have an upper bound) which concludes the proof. \square

With the proof of [Lemma 5](#) done, we now provide the proof of [Lemma 8](#). The last part of the proof of [Lemma 8](#), which lower and upper bounds the Rademacher complexity, uses an idea similar to [\[Larsen and Ritzert, 2022\]](#) [[Lemma 9](#)].

Proof of [Lemma 8](#). First if $\beta > \gamma$, then by hypotheses in $\Delta(\mathcal{H})_{[\gamma]}$ attaining values in $[-\gamma, \gamma]$, it follows that $\text{fat}_\beta(\Delta(\mathcal{H})_{[\gamma]}) = 0$, since no hypothesis in $\Delta(\mathcal{H})_{[\gamma]}$ can be above $r_1 + \beta$ or below $r_1 - \beta$ for any $r_1 \in \mathbb{R}$. As $\frac{Cd}{\beta} > 0$ this proves the case $\beta > \gamma$, thus, we consider for now the case that $0 < \beta \leq \gamma$. Let $x_1, \dots, x_{d'}$ and $r \in [\beta - \gamma, \gamma - \beta]^{d'}$ be β -shattered by $\Delta(\mathcal{H})_{[\gamma]}$. We note that we may assume that r is in $[\beta - \gamma, \gamma - \beta]^{d'}$, as $[-\gamma, \gamma]$ contains the image of $\Delta(\mathcal{H})_{[\gamma]}$. If an entry of r_i is outside the interval $[\beta - \gamma, \gamma - \beta]^{d'}$, it must either be the case that no function in $\Delta(\mathcal{H})_{[\gamma]}$ can be β below or above r_i , depending on r_i being positive or negative. Now by $x_1, \dots, x_{d'}$ and $r \in [\beta - \gamma, \gamma - \beta]^{d'}$ being β -shattered by $\Delta(\mathcal{H})_{[\gamma]}$, we have that for any $b \in \{-1, 1\}^{d'}$, that there exists $f_{[\gamma]} \in \mathcal{F}_{[\gamma]}$ such that

$$f_{[\gamma]}(x_i) \leq r_i - \beta \quad \text{if } b_i = -1 \quad (36)$$

$$f_{[\gamma]}(x_i) \geq r_i + \beta \quad \text{if } b_i = 1. \quad (37)$$

We now recall that $f_{[\gamma]} \in \mathcal{F}_{[\gamma]}$ is generated by a function $f \in \Delta(\mathcal{H})$ in the following way

$$f_{[\gamma]}(x) = \begin{cases} \gamma & \text{if } f(x) > \gamma \\ -\gamma & \text{if } f(x) < -\gamma \\ f(x) & \text{else,} \end{cases}$$

i.e. f always being below $f_{[\gamma]}$ if $f_{[\gamma]}$ is less strictly less than γ and f always being above $f_{[\gamma]}$ if $f_{[\gamma]}$ is strictly larger than $-\gamma$. This implies that the function $f \in \Delta(\mathcal{H})$ generating $f_{[\gamma]}$ also satisfies:

$$\begin{aligned} f(x_i) &\leq r_i - \beta & \text{if } b_i = -1 \\ f(x_i) &\geq r_i + \beta & \text{if } b_i = 1, \end{aligned} \quad (38)$$

as $r_i \in [\beta - \gamma, \gamma - \beta]$, we have, if $b_i = -1$, that $\gamma > r_i - \beta \geq f_{[\gamma]}(x_i) \geq f(x_i)$, and if $b_i = 1$, then $-\gamma < r_i + \beta \leq f_{[\gamma]}(x_i) \leq f(x_i)$. Thus, we conclude that $\Delta(\mathcal{H})$ also β -shatters $x_1, \dots, x_{d'}$ and $r_1, \dots, r_{d'}$. We now notice that Eq. (38) implies that for all $b \in \{-1, 1\}^{d'}$, there exists $f \in \Delta(\mathcal{H})$ such that:

$$b_i(f(x_i) - r_i) \geq \beta.$$

Using this, and adding $\mathbb{E}_{\sigma \sim \{-1, 1\}^{d'}} \left[\sum_{i=1}^{d'} \sigma_i(-r_i) \right] = 0$ we have the following lower bound on the Rademacher complexity of $\Delta(\mathcal{H})$ on $x_1, \dots, x_{d'}$:

$$\mathbb{E}_{\sigma \sim \{-1, 1\}^{d'}} \left[\frac{1}{d'} \sup_{f \in \Delta(\mathcal{H})} \sum_{i=1}^{d'} \sigma_i f(x_i) \right] = \mathbb{E}_{\sigma \sim \{-1, 1\}^{d'}} \left[\frac{1}{d'} \sup_{f \in \Delta(\mathcal{H})} \sum_{i=1}^{d'} \sigma_i (f(x_i) - r_i) \right] \geq \beta.$$

Furthermore, since $f \in \Delta(\mathcal{H})$ is a convex combination of hypotheses in \mathcal{H} , the Rademacher complexity of $\Delta(\mathcal{H})$ is the same as that of \mathcal{H} . Due to [Dudley, 1978] [See, for instance, [Hajek and Raginsky, 2021], Theorem 7.2], we have that since \mathcal{H} has VC dimension d , the Rademacher complexity of \mathcal{H} is at most $\sqrt{\frac{Cd}{d'}}$, where $C \geq 1$ is a universal constant. Thus, we conclude:

$$\mathbb{E}_{\sigma \sim \{-1, 1\}^{d'}} \left[\sup_{f \in \Delta(\mathcal{H})} \sum_{i=1}^{d'} \sigma_i f(x_i) \right] \leq \mathbb{E}_{\sigma \sim \{-1, 1\}^{d'}} \left[\sup_{h \in \mathcal{H}} \sum_{i=1}^{d'} \sigma_i h(x_i) \right] \leq \sqrt{\frac{Cd}{d'}},$$

which implies that $d' \leq \frac{Cd}{\beta^2}$, as claimed and concludes the proof \square

To prove Lemma 7, we need the following lemma due to [Rudelson and Vershynin, 2006]. [Rudelson and Vershynin, 2006] bounds the size of a maximal α -packing, which upper bounds the size of a minimal α -cover, since an α -packing is also an α -cover.

Theorem 9 ([Rudelson and Vershynin, 2006] Theorem 4.4). *There exist universal constants $\check{C} \geq 1$ and $\check{c} > 0$ such that: For a point set $X = \{x_1, \dots, x_m\} \subseteq \mathcal{X}$ of size m and a function class \mathcal{F} defined on \mathcal{X} , if $\forall f \in \mathcal{F}$, we have that $\sum_{x \in X} |f(x)|/m \leq 1$, then for $0 < \varepsilon < 1$, $0 < \alpha < \frac{1}{2}$, and $d = \text{fat}_{\check{c}\alpha\varepsilon}(\mathcal{F})$*

$$\ln(\mathcal{N}_\infty(X, \mathcal{F}, \alpha)) \leq \check{C}d \ln\left(\frac{2m}{d\alpha}\right) \ln^\varepsilon\left(\frac{2m}{d}\right).$$

We now give the proof of Lemma 7.

Proof of Lemma 7. Let $0 < \varepsilon < 1$. We consider the function class $\Delta(\mathcal{H})_{[\gamma]}/\gamma = \{f : f = f'/\gamma \text{ for } f' \in \Delta(\mathcal{H})_{[\gamma]}\}$. We note that any function $f \in \Delta(\mathcal{H})_{[\gamma]}/\gamma$ is bounded in absolute value by 1, since any $f \in \Delta(\mathcal{H})_{[\gamma]}$ is bounded in absolute value by γ . Thus, we conclude that for $f \in \Delta(\mathcal{H})_{[\gamma]}/\gamma$

$$\sum_{x \in X} \frac{|f(x)|}{m} \leq 1.$$

We now invoke [Theorem 9](#) (using the notation $d' = \text{fat}_{\tilde{c}\varepsilon\alpha}(\Delta(\mathcal{H})_{[\gamma]}/\gamma)$) to obtain that there exist universal constants $\check{C} \geq 1$ and $\check{c} > 0$ such that

$$\begin{aligned} \ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}/\gamma, \alpha)) &\leq \check{C}d' \ln\left(\frac{2m}{d'\alpha}\right) \cdot \ln^\varepsilon\left(\frac{2m}{d}\right) \\ &\leq \check{C}d' \ln^{1+\varepsilon}\left(\frac{2m}{d'\alpha}\right). \end{aligned} \quad (\text{by } \alpha < \frac{1}{2})$$

Now since $\text{fat}_{\tilde{c}\varepsilon\alpha}(\Delta(\mathcal{H})_{[\gamma]}/\gamma) = \text{fat}_{\tilde{c}\varepsilon\alpha\gamma}(\Delta(\mathcal{H})_{[\gamma]})$ (so $d' = d = \text{fat}_{\tilde{c}\varepsilon\alpha\gamma}(\Delta(\mathcal{H})_{[\gamma]})$), we obtain

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}/\gamma, \alpha)) \leq \check{C}d \ln^{1+\varepsilon}\left(\frac{2m}{d\alpha}\right).$$

Furthermore, we note that for any $f \in \Delta(\mathcal{H})_{[\gamma]}$ and $f' \in N$, where N is a minimal (in terms of size) α -cover of $\Delta(\mathcal{H})_{[\gamma]}$ on X in infinity norm, such that f' is α -close to $\frac{f}{\gamma} \in \Delta(\mathcal{H})_{[\gamma]}/\gamma$, it holds that

$$\max_{x \in X} \left| \frac{f(x)}{\gamma} - f'(x) \right| \leq \alpha \Rightarrow \max_{x \in X} |f - \gamma f'(x)| \leq \alpha\gamma. \quad (39)$$

Thus, we conclude that $\{\gamma \cdot f'\}_{f' \in N}$, forms an $\alpha\gamma$ cover of $\Delta(\mathcal{H})_{[\gamma]}$. Consequently, we have

$$\ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}, \alpha\gamma)) \leq \ln(\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}/\gamma, \alpha)) \leq \check{C}d \ln^{1+\varepsilon}\left(\frac{2m}{d\alpha}\right) \quad (40)$$

which concludes the proof. \square

5 Final Upper bound

In this section, we use [Lemma 4](#) and [Corollary 6](#) to give our generalization bound for hypotheses in $\Delta(\mathcal{H})$, for all margin levels simultaneously. We now present our generalization bound and then proceed with the proof.

Theorem 10. *There exists a universal constant c such that: For $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ a hypothesis class of VC-dimension d , distribution \mathcal{D} over $\mathcal{X} \times \{-1, 1\}$, failure parameter $0 < \delta < 1$, and training sequence size $m \in \mathbb{N}$, it holds with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}^m$ that: for any margin $0 < \gamma \leq 1$ and any function $f \in \Delta(\mathcal{H})$*

$$\mathcal{L}_{\mathcal{D}}(f) \leq \mathcal{L}_{\mathbf{S}}^\gamma(f) + \sqrt{c \cdot \mathcal{L}_{\mathbf{S}}^\gamma(f) \left(\frac{d\Gamma\left(\frac{m\gamma^2}{d}\right)}{m\gamma^2} + \frac{\ln\left(\frac{\varepsilon}{\delta}\right)}{m} \right)} + c \left(\frac{d\Gamma\left(\frac{m\gamma^2}{d}\right)}{m\gamma^2} + \frac{\ln\left(\frac{\varepsilon}{\delta}\right)}{m} \right).$$

where $\Gamma(x) = \text{Ln}^2(\text{Ln}(x)) \text{Ln}(x)$.

Proof. Let $c' \geq c > 0$ and $c' \geq 1$ be the universal constants from [Corollary 6](#) and $\tilde{c} = \max(1/c, c')$. We will show in the following, that with probability at least $1 - \delta$ over \mathbf{S} , it

holds for all $0 < \gamma \leq 1$ and all $f \in \Delta(\mathcal{H})$ that

$$\mathcal{L}_{\mathcal{D}}(f) \leq \mathcal{L}_{\mathbf{S}}^{\gamma}(f) + 64 \left(\sqrt{\frac{\mathcal{L}_{\mathbf{S}}^{\gamma}(f) \cdot 2 \left(\ln \left(\frac{1640e}{\delta} \right) + \frac{9600\tilde{c}^2 \min(c,1)d}{32\gamma^2} \Gamma \left(\frac{32m\gamma^2}{\min(c,1)d} \right) \right)}{m}} + \frac{5 \left(\ln \left(\frac{1640e}{\delta} \right) + \frac{9600\tilde{c}^2 \min(c,1)d}{32\gamma^2} \Gamma \left(\frac{32m\gamma^2}{\min(c,1)d} \right) \right)}{m} \right). \quad (41)$$

We first notice that if $\frac{9600\tilde{c}^2 \min(c,1)d}{32m} \geq \gamma^2$, the right-hand side of the above is greater than 1, and we are done by the left-hand side always being at most 1, i.e. the above inequality holds with probability 1. Thus, we now consider the case where $1 \geq \gamma^2 \geq \frac{9600\tilde{c}^2 \min(c,1)d}{32m}$. We now further restrict the values we have to consider for the bound in Eq. (41) to be non vacuuous, this is done to the end of employing Corollary 6 later in the argument.

We will soon argue that the function $\Gamma(x)/x$ is decreasing for $x \geq e^4$, we notice this is also a continuous function in x . Now consider the function $9600\tilde{c}^2 \Gamma \left(\frac{32m\gamma^2}{\min(c,1)d} \right) / \left(\frac{32m\gamma^2}{\min(c,1)d} \right)$, which by the above is decreasing in $\frac{32m\gamma^2}{\min(c,1)d} \geq e^e$, which is the case since we consider γ such that $1 \geq \frac{32m\gamma^2}{\min(c,1)d} \geq 9600\tilde{c}^2 \geq 9600$. This also implies that for the right hand side of Eq. (41) being less than 1 for any $1 \geq \gamma^2 \geq \frac{9600\tilde{c}^2 \min(c,1)d}{32m}$ it must be the case that m is such that $9600\tilde{c}^2 \Gamma \left(\frac{32m}{\min(c,1)d} \right) / \left(\frac{32m}{\min(c,1)d} \right) < 1$ which we assume is the case from now on. Furthermore, by the above argued continuity and the function being decreasing, let γ' be such that $1 \geq \gamma'^2 \geq \frac{9600\tilde{c}^2 \min(c,1)d}{32m}$ and that $9600\tilde{c}^2 \Gamma \left(\frac{32m\gamma'^2}{\min(c,1)d} \right) / \left(\frac{32m\gamma'^2}{\min(c,1)d} \right) = 1$, notice that γ' must be within the interval since the function was decreasing in $\left(\frac{32m\gamma'^2}{\min(c,1)d} \right)$ and for the value $\gamma^2 = \frac{9600\tilde{c}^2 \min(c,1)d}{32m}$ the function is greater than 1. Notice that for the values of γ , between $\gamma'^2 \geq \gamma \geq \frac{9600\tilde{c}^2 \min(c,1)d}{32m}$, the bound on the right hand side of Eq. (41) is again larger than 1, so we consider from now on the case $1 \geq \gamma^2 \geq \gamma'^2$, and m such that $9600\tilde{c}^2 \Gamma \left(\frac{32m}{\min(c,1)d} \right) / \left(\frac{32m}{\min(c,1)d} \right) < 1$. Furthermore, again by the above argued monotonicity we notice that for $1 \geq \gamma^2 \geq \gamma'^2$ we have that

$$9600\tilde{c}^2 \Gamma \left(\frac{32m\gamma^2}{\min(c,1)d} \right) / \left(\frac{32m\gamma^2}{\min(c,1)d} \right) \leq 9600\tilde{c}^2 \Gamma \left(\frac{32\gamma'^2 m}{\min(c,1)d} \right) / \left(\frac{32\gamma'^2 m}{\min(c,1)d} \right) = 1 \quad (42)$$

which gonna be important later when we want to invoke Corollary 6.

To argue that $\Gamma(x)/x$ is decreasing for $x \geq e^4$, we notice that for $x \geq e^4$ we have that $\Gamma(x) = \text{Ln}^2(\text{Ln}(x)) \text{Ln}(x) = \ln^2(\ln(x)) \ln(x)$, thus it suffices to show that the function $f(x) = \ln^2(\ln(x)) \ln(x)/x$ is decreasing for $x \geq e^4$. f has derivative $(2 \ln(\ln(x)) - \ln(x)(\ln(\ln(x)))^2 + (\ln(\ln(x)))^2)/x^2$. We observe that for $x \geq e^2$ we have that $-\ln(x)(\ln(\ln(x)))^2/2 + (\ln(\ln(x)))^2 \leq 0$ and for $x \geq e^4$ we have that $2 \ln(\ln(x)) - \ln(x)(\ln(\ln(x)))^2/2 \leq 0$. Whereby we conclude that $f(x) = \ln^2(\ln(x)) \ln(x)/x$ is decreasing for $x \geq e^4$. We are now ready to setup the parameters for the union bound over the different level sets of γ .

Let $I = \{0, \dots, \lfloor \log_2(1/\gamma') \rfloor\}$ and define $\gamma_0^i = \sqrt{2^i \gamma'}$ and $\gamma_1^i = \sqrt{2} \gamma_0^i$, for $i \in I$, except $i = \lfloor \log_2(1/\gamma') \rfloor$, where we define $\gamma_1^{\lfloor \log_2(1/\gamma') \rfloor} = 1$. For $i \in I$ we also define $N_i = \exp \left(\frac{72c'd}{(\gamma_0^i)^2} \Gamma \left(\frac{16m(\gamma_0^i)^2}{cd} \right) \right)$ and $\delta_i = \frac{\delta}{1640N_i^2} \frac{d}{m(\gamma_0^i)^2}$. Furthermore, for each $i \in I$ we define $J_i = \{0, 1, \dots, \lfloor \frac{m}{\ln(N_i)} \rfloor\}$ and for $j \in J_i$ define $\tau_0^{i,j} = j \frac{\ln(N_i)}{m}$ and $\tau_1^{i,j} = (j+1) \frac{\ln(N_i)}{m}$. Lastly

for $i \in I$ and $j \in J_i$ we define

$$\beta_{i,j} = 64 \left(\sqrt{\frac{\tau_1^{i,j} \cdot 2 \ln\left(\frac{e}{\delta_i}\right)}{m}} + \frac{2 \ln\left(\frac{e}{\delta_i}\right)}{m} \right).$$

For now let $i \in I$, then with the above notation, we get by the union bound and an invocation of [Lemma 4](#) that

$$\begin{aligned} & \mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m} \left[\exists j \in J_i, \exists \gamma \in [\gamma_0^i, \gamma_1^i], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0^{i,j}, \tau_1^{i,j}], \mathcal{L}_{\mathcal{D}}(f) \geq \tau_1^{i,j} + \beta_{i,j} \right] \\ & \leq \sum_{j \in J_i} \mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m} \left[\exists \gamma \in [\gamma_0^i, \gamma_1^i], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0^{i,j}, \tau_1^{i,j}], \mathcal{L}_{\mathcal{D}}(f) \geq \tau_1^{i,j} + \beta_{i,j} \right] \\ & \leq \left(\left\lfloor \frac{m}{\ln(N_i)} \right\rfloor + 1 \right) \delta_i \sup_{X \in \mathcal{X}^{2m}} \mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[2\gamma_1^i]}, \frac{\gamma_0^i}{2}). \end{aligned} \quad (43)$$

We remark that if $i \in I$ is such that $\frac{\ln(N_i)}{m} > 1$, then we have that $\tau_1^{i,j} = (j+1)\frac{\ln(N_i)}{m} > 1$, which implies that we could have upper bounded the above with 0. Thus, we assume for now that $i \in I$ such that $\frac{\ln(N_i)}{m} \leq 1$, which implies that $\left(\left\lfloor \frac{m}{\ln(N_i)} \right\rfloor + 1\right)$ is upper bounded by $\frac{2m}{\ln(N_i)}$.

To the end of employing [Corollary 6](#) with $\gamma = \min(2\gamma_1^i, 1)$ and $\alpha = \frac{\gamma_0^i}{4\gamma_1^i}$, we notice that for $\gamma = \min(2\gamma_1^i, 1)$ and $\alpha = \frac{\gamma_0^i}{4\gamma_1^i}$ we have that

$$\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[2\gamma_1^i]}, \frac{\gamma_0^i}{2}) \leq \mathcal{N}_\infty(X, \Delta(\mathcal{H})_{[\gamma]}, \alpha\gamma) \quad (44)$$

where the inequality follows from in the case that $\gamma = \min(2\gamma_1^i, 1) = 1$, then since the functions in $\Delta(\mathcal{H})$ only attains values in $[-1, 1]$, we have that $\Delta(\mathcal{H})_{[2\gamma_1^i]} = \Delta(\mathcal{H})_{[1]} = \Delta(\mathcal{H})$ and that the covering number being decreasing in the precision argument and $\gamma_0^i/2 \geq \alpha\gamma$. Furthermore, to the end of invoking [Corollary 6](#), we now argue that $\gamma^2 \geq \frac{4cd \text{Ln}^2(\text{Ln}(\frac{8\alpha m \gamma^2}{cd}))}{\alpha^2 m}$ for the above γ and α . To this end we notice that since γ_0^i and $\gamma\sqrt{2}\gamma_0^i$, except $i = \lfloor \log_2(1/\gamma') \rfloor$, where we define $\gamma_1^{\lfloor \log_2(1/\gamma') \rfloor} = 1$, we have the following $\frac{1}{4\sqrt{2}} \leq \alpha = \frac{\gamma_0^i}{4\gamma_1^i} \leq \frac{1}{4}$. Furthermore, we argued previously that it suffices to consider the case that $\frac{9600\tilde{c}^2 \min(c,1)d}{32} \Gamma\left(\frac{32m}{\min(c,1)d}\right)/m < 1$. Now if $\gamma = 1$ we get that

$$\frac{4cd \text{Ln}^2(\text{Ln}(\frac{8\alpha m \gamma^2}{cd}))}{\alpha^2 m} \leq \frac{128cd \text{Ln}^2(\text{Ln}(\frac{2m}{cd}))}{m} \leq \frac{9600\tilde{c}^2 \min(c,1)d}{32m} \Gamma\left(\frac{32m}{\min(c,1)d}\right) < 1 = \gamma^2$$

where the first inequality uses that $\frac{1}{4\sqrt{2}} \leq \alpha \leq \frac{1}{4}$, and the second inequality that $c \leq \tilde{c} = \max(c', 1/c)$, where $c' \geq c$ and $9600/32 \geq 128$ and the last inequality the case $\gamma = 1$ we consider.

Now if $\gamma = 2\gamma_1^i$ we get

$$\begin{aligned} \frac{4cd \text{Ln}^2(\text{Ln}(\frac{8\alpha m \gamma^2}{cd}))}{\alpha^2 m} &\leq \frac{128cd \text{Ln}^2(\text{Ln}(\frac{8m(\gamma_1^i)^2}{cd}))}{m} \leq \frac{9600\tilde{c}^2 \min(c, 1)d}{32m} \Gamma\left(\frac{32m(\gamma_1^i)^2}{\min(c, 1)d}\right) \\ &\leq (\gamma_1^i)^2 \leq \gamma^2 \end{aligned}$$

where the first inequality uses that $\frac{1}{4\sqrt{2}} \leq \alpha \leq \frac{1}{4}$, and the second inequality that $c \leq \tilde{c} = \max(c', 1/c)$, where $c' \geq c$, that $\gamma = 2\gamma_1^i$ and $9600/32 \geq 128$ where the second to last inequality follows from [Eq. \(42\)](#) and $1 \geq \gamma_1^i \geq \gamma_0^i \geq \gamma'$ and the last inequality by $\gamma = 2\gamma_1^i$, whereby we have argued that $\gamma^2 \geq \frac{4cd \text{Ln}^2(\text{Ln}(\frac{8\alpha m \gamma^2}{cd}))}{\alpha^2 m}$, and since $0 < \gamma \leq 1$ and $0 < \alpha < 1/2$ we get by invoking [Corollary 6](#) that

$$\begin{aligned} \mathcal{N}_\infty(X, \Delta(\mathcal{H})_{\lfloor 2\gamma_1^i \rfloor}, \frac{\gamma_0^i}{2}) &\leq \mathcal{N}_\infty(X, \Delta(\mathcal{H})_{\lfloor \gamma \rfloor}, \alpha\gamma) \leq \exp\left(\frac{c'd}{\alpha^2 \gamma^2} \Gamma\left(\frac{8\alpha m \gamma^2}{cd}\right)\right) \\ &\leq \exp\left(\frac{16c'd}{(\gamma_0^i)^2} \Gamma\left(\frac{16m(\gamma_0^i)^2}{cd}\right)\right) \leq N_i, \end{aligned}$$

where the first inequality follows from [Eq. \(44\)](#), the second inequality by [Corollary 6](#), the third inequality by using on the term outside of $\Gamma(\cdot)$ that $\frac{1}{\alpha^2 \gamma^2} = \frac{16(\gamma_1^i)^2}{(\gamma_0^i)^2 \min(4(\gamma_1^i)^2, 1)} \leq 16/(\gamma_0^i)^2$ and using on the term inside of $\Gamma(\cdot)$ that $\gamma_1^i \leq \sqrt{2}(\gamma_0^i)^2$ to upper bound $\alpha\gamma^2 = \frac{\gamma_0^i}{4\gamma_1^i} (\min(2\gamma_1^i, 1))^2 \leq \sqrt{2}(\gamma_0^i)^2$. Thus, by using that we concluded that $\mathcal{N}_\infty(X, \Delta(\mathcal{H})_{\lfloor 2\gamma_1^i \rfloor}, \frac{\gamma_0^i}{2}) \leq N_i$ and that $(\lfloor \frac{m}{\ln(N_i)} \rfloor + 1) \leq \frac{2m}{\ln(N_i)}$, and $\delta_i = \frac{\delta}{1640N_i^2} \frac{d}{m(\gamma_0^i)^2}$, we get by plugging into the last expression of [Eq. \(43\)](#), that

$$\left(\left\lfloor \frac{m}{\ln(N_i)} \right\rfloor + 1\right) \delta_i \sup_{X \in \mathcal{X}^{2m}} \mathcal{N}_\infty(X, \Delta(\mathcal{H})_{\lfloor 2\gamma_1^i \rfloor}, \frac{\gamma_0^i}{2}) \leq \frac{m}{\ln(N_i)} \frac{\delta}{820N_i} \frac{d}{m(\gamma_0^i)^2}. \quad (45)$$

Furthermore, we notice that by $N_i = \exp\left(\frac{72c'd}{(\gamma_0^i)^2} \Gamma\left(\frac{16m(\gamma_0^i)^2}{cd}\right)\right)$, we have that

$$\frac{m}{\ln(N_i)} \leq \frac{m(\gamma_0^i)^2}{72c'd \Gamma\left(\frac{16m(\gamma_0^i)^2}{cd}\right)} \leq \frac{m(\gamma_0^i)^2}{d}, \quad (46)$$

where the last inequality follows by $c' \geq 1$ and $\Gamma\left(\frac{16m(\gamma_0^i)^2}{cd}\right) \geq 1$. Now, combining [Eq. \(43\)](#), [Eq. \(45\)](#) and [Eq. \(46\)](#) we get that

$$\begin{aligned} &\mathbb{P}_{\mathbf{s} \sim \mathcal{D}^m} \left[\exists j \in J_i, \exists \gamma \in [\gamma_0^i, \gamma_1^i], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathcal{S}}^\gamma(f) \in [\tau_0^{i,j}, \tau_1^{i,j}], \mathcal{L}_{\mathcal{D}}(f) \geq \tau_1^{i,j} + \beta_{i,j} \right] \\ &\leq \frac{\delta}{820N_i} = \frac{\delta}{820} \exp\left(-\frac{72c'd}{(\gamma_0^i)^2} \Gamma\left(\frac{16m(\gamma_0^i)^2}{cd}\right)\right). \end{aligned}$$

We showed the above for any $i \in I = \{0, \dots, \lfloor \log_2(1/\gamma') \rfloor\}$; thus, by an application of the

union bound, and $\gamma_0^i = \sqrt{2^i \gamma'}$, we conclude that:

$$\begin{aligned}
& \mathbb{P}_{\mathbf{S} \sim \mathcal{D}^m} \left[\exists i \in I, \exists j \in J_i, \exists \gamma \in [\gamma_0^i, \gamma_1^i], \exists f \in \Delta(\mathcal{H}) : \mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0^{i,j}, \tau_1^{i,j}], \mathcal{L}_{\mathcal{D}}(f) \geq \tau_1^{i,j} + \beta_{i,j} \right] \\
& \leq \sum_{i=0}^{\lfloor \log_2(1/\gamma') \rfloor} \frac{\delta}{820} \exp \left(-\frac{72c'd}{(\gamma_0^i)^2} \Gamma \left(\frac{16m(\gamma_0^i)^2}{cd} \right) \right) \leq \frac{\delta}{820} \sum_{i=0}^{\lfloor \log_2(1/\gamma') \rfloor} \exp \left(-\frac{72c'd}{2^i \gamma'} \right) \\
& \leq \frac{\delta}{820} \sum_{i=0}^{\lfloor \log_2(1/\gamma') \rfloor} \exp \left(-\frac{72c'd}{2^{\lfloor \log_2(1/\gamma') \rfloor - i} \gamma'} \right) \leq \frac{\delta}{820} \sum_{i=0}^{\lfloor \log_2(1/\gamma') \rfloor} \exp \left(-\frac{72c'd 2^i}{2^{\log_2(1/\gamma')}} \right) \\
& \leq \frac{\delta}{820} \sum_{i=0}^{\lfloor \log_2(1/\gamma') \rfloor} \exp(-72c'd 2^i) \leq \frac{\delta}{820},
\end{aligned}$$

where the first inequality follows from the union bound, the second inequality by $\Gamma \left(\frac{16m(\gamma_0^i)^2}{cd} \right) \geq 1$, the third inequality by summing in the reverse order, the fourth inequality by $2^{\lfloor \log_2(1/\gamma') \rfloor} \leq 2^{\log_2(1/\gamma')}$, so $-1/2^{\lfloor \log_2(1/\gamma') \rfloor} \leq -1/2^{\log_2(1/\gamma')}$, and the last inequality by the sum being less than 1. Thus, we conclude that with probability at least $1 - \delta$, we have that for all $i \in I$, for all $j \in J_i$, for all $\gamma \in [\gamma_0^i, \gamma_1^i]$, for all $f \in \Delta(\mathcal{H})$, that either

$$\mathcal{L}_{\mathbf{S}}^\gamma(f) \notin [\tau_0^{i,j}, \tau_1^{i,j}]$$

or

$$\mathcal{L}_{\mathcal{D}}(f) < \tau_1^{i,j} + \beta_{i,j} = \tau_1^{i,j} + 64 \left(\sqrt{\frac{\tau_1^{i,j} \cdot 2 \ln \left(\frac{e}{\delta_i} \right)}{m}} + \frac{2 \ln \left(\frac{e}{\delta_i} \right)}{m} \right),$$

where the last equality uses that $\beta_{i,j} = 64 \left(\sqrt{\frac{\tau_1^{i,j} \cdot 2 \ln \left(\frac{e}{\delta_i} \right)}{m}} + \frac{2 \ln \left(\frac{e}{\delta_i} \right)}{m} \right)$. Furthermore, since $\cup_{i \in I} [\gamma_0^i, \gamma_1^i] = [\gamma_0^0, \gamma_1^{\lfloor \log_2(1/\gamma') \rfloor}] = [\gamma', 1]$ and for each $i \in I$, $\cup_{j \in J_i} [\tau_0^{i,j}, \tau_1^{i,j}]$ contains the interval $[0, 1]$ and $\mathcal{L}_{\mathbf{S}}^\gamma(f) \in [0, 1]$, the above implies that with probability at least $1 - \delta$ over \mathbf{S} , for any $\gamma \in [\gamma', 1]$ and for any $f \in \Delta(\mathcal{H})$ there exists $i \in I$, such that $\gamma \in [\gamma_0^i, \gamma_1^i]$, and $j \in J_i$ such that $\mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0^{i,j}, \tau_1^{i,j}]$, and it holds that

$$\mathcal{L}_{\mathcal{D}}(f) < \tau_1^{i,j} + 64 \left(\sqrt{\frac{\tau_1^{i,j} \cdot 2 \ln \left(\frac{e}{\delta_i} \right)}{m}} + \frac{2 \ln \left(\frac{e}{\delta_i} \right)}{m} \right). \quad (47)$$

Now in the case of $\gamma \in [\gamma_0^i, \gamma_1^i]$, $\mathcal{L}_{\mathbf{S}}^\gamma(f) \in [\tau_0^{i,j}, \tau_1^{i,j}]$ and Eq. (47) holding, we get, by the definition of $\tau_0^{i,j} = j \frac{\ln(N_i)}{m}$ and $\tau_1^{i,j} = (j+1) \frac{\ln(N_i)}{m}$, $\ln(N_i) \leq \ln \left(\frac{e}{\delta_i} \right)$ and $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$,

that

$$\begin{aligned}
\mathcal{L}_{\mathcal{D}}(f) &\leq \tau_1^{i,j} + 64 \left(\sqrt{\frac{\tau_1^{i,j} \cdot 2 \ln\left(\frac{e}{\delta_i}\right)}{m}} + \frac{2 \ln\left(\frac{e}{\delta_i}\right)}{m} \right) \\
&\leq \tau_0^{i,j} + \frac{\ln\left(\frac{e}{\delta_i}\right)}{m} + 64 \left(\sqrt{\frac{\tau_0^{i,j} \cdot 2 \ln\left(\frac{e}{\delta_i}\right)}{m}} + \sqrt{2} \frac{\ln\left(\frac{e}{\delta_i}\right)}{m} + \frac{2 \ln\left(\frac{e}{\delta_i}\right)}{m} \right) \\
&\leq \mathcal{L}_{\mathbf{S}}^\gamma(f) + 64 \left(\sqrt{\frac{\mathcal{L}_{\mathbf{S}}^\gamma(f) \cdot 2 \ln\left(\frac{e}{\delta_i}\right)}{m}} + \frac{5 \ln\left(\frac{e}{\delta_i}\right)}{m} \right).
\end{aligned} \tag{48}$$

Furthermore, we have that $N_i = \exp\left(\frac{72c'd}{(\gamma_0^i)^2} \Gamma\left(\frac{16m(\gamma_0^i)^2}{cd}\right)\right)$, $\delta_i = \frac{\delta}{1640N_i^2} \frac{d}{m(\gamma_0^i)^2}$, and by $\gamma \in [\gamma_0^i, \gamma_1^i]$, with $\gamma_1^i \leq \sqrt{2}\gamma_0^i$, it implies that $\gamma \leq \sqrt{2}\gamma_0^i$, which combined gives that $\ln(e/\delta_i)$ can be bounded by

$$\begin{aligned}
\ln\left(\frac{e}{\delta_i}\right) &= \ln\left(\frac{1640e}{\delta}\right) + 2 \ln(N_i) + \ln\left(\frac{m(\gamma_0^i)^2}{d}\right) \\
&\leq \ln\left(\frac{1640e}{\delta}\right) + \frac{144c'd}{(\gamma_0^i)^2} \Gamma\left(\frac{16m(\gamma_0^i)^2}{cd}\right) + \ln\left(\frac{m\gamma^2}{d}\right) \quad (\text{by } \gamma \geq \gamma_0^i) \\
&\leq \ln\left(\frac{1640e}{\delta}\right) + \frac{288c'd}{\gamma^2} \Gamma\left(\frac{16m\gamma^2}{\min(c,1)d}\right) + \text{Ln}\left(\frac{m\gamma^2}{d}\right). \quad (\text{by } \gamma \leq \sqrt{2}\gamma_0^i \text{ and } \gamma_0^i \leq \gamma)
\end{aligned} \tag{49}$$

We notice that since $\ln(x)/x$ has derivative $\frac{1-\ln(x)}{x^2} = \frac{\ln(e/x)}{x^2}$ which is less than 0 for $x \geq e$, we have that $\text{Ln}(x)/x$ is decreasing for $x \geq e$ and since $\frac{1}{x}$ is decreasing for $0 < x < e$ we conclude that $\text{Ln}(x)/x$ is decreasing for $x > 0$. Thus, we conclude that

$$\frac{\text{Ln}\left(\frac{m\gamma^2}{d}\right)}{m} \leq \frac{d \text{Ln}\left(\frac{m\gamma^2}{d}\right)}{m\gamma^2} \leq \frac{d \text{Ln}\left(\frac{m\gamma^2}{d}\right)}{m\gamma^2} \leq \frac{d \Gamma\left(\frac{16m\gamma^2}{\min(c,1)d}\right)}{m\gamma^2}.$$

Now, using this with [Eq. \(48\)](#) and [Eq. \(49\)](#), we conclude that

$$\begin{aligned}
\mathcal{L}_{\mathcal{D}}(f) &\leq \mathcal{L}_{\mathbf{S}}^\gamma(f) + 64 \left(\sqrt{\frac{\mathcal{L}_{\mathbf{S}}^\gamma(f) \cdot 2 \left(\ln\left(\frac{1640e}{\delta}\right) + \frac{288c'd}{\gamma^2} \Gamma\left(\frac{16m\gamma^2}{\min(c,1)d}\right) + \text{Ln}\left(\frac{m\gamma^2}{d}\right) \right)}{m}} \right. \\
&\quad \left. + \frac{5 \left(\ln\left(\frac{1640e}{\delta}\right) + \frac{288c'd}{\gamma^2} \Gamma\left(\frac{16m\gamma^2}{\min(c,1)d}\right) + \text{Ln}\left(\frac{m\gamma^2}{d}\right) \right)}{m} \right) \\
&\leq \mathcal{L}_{\mathbf{S}}^\gamma(f) + 64 \left(\sqrt{\frac{\mathcal{L}_{\mathbf{S}}^\gamma(f) \cdot 2 \left(\ln\left(\frac{1640e}{\delta}\right) + \frac{289c'd}{\gamma^2} \Gamma\left(\frac{16m\gamma^2}{\min(c,1)d}\right) \right)}{m}} \right. \\
&\quad \left. + \frac{5 \left(\ln\left(\frac{1640e}{\delta}\right) + \frac{289c'd}{\gamma^2} \Gamma\left(\frac{16m\gamma^2}{\min(c,1)d}\right) \right)}{m} \right).
\end{aligned}$$

Thus we conclude that with probability at least $1 - \delta$ over \mathbf{S} it holds for any $\gamma \in [\gamma', 1]$ and any $f \in \Delta(\mathcal{H})$ that

$$\mathcal{L}_{\mathcal{D}}(f) \leq \mathcal{L}_{\mathbf{S}}^{\gamma}(f) + 64 \left(\sqrt{\frac{\mathcal{L}_{\mathbf{S}}^{\gamma}(f) \cdot 2 \left(\ln \left(\frac{1640e}{\delta} \right) + \frac{289c'd}{\gamma^2} \Gamma \left(\frac{16m\gamma^2}{\min(c,1)d} \right) \right)}{m}} \right) \quad (50)$$

$$+ \frac{5 \left(\ln \left(\frac{1640e}{\delta} \right) + \frac{289c'd}{\gamma^2} \Gamma \left(\frac{16m\gamma^2}{\min(c,1)d} \right) \right)}{m}, \quad (51)$$

and since we start by concluding that with probability 1 over \mathbf{S} for any $\gamma \in [0, \gamma']$ and any $f \in \Delta(\mathcal{H})$ Eq. (41) holds, i.e.

$$\mathcal{L}_{\mathcal{D}}(f) \leq \mathcal{L}_{\mathbf{S}}^{\gamma}(f) + 64 \left(\sqrt{\frac{\mathcal{L}_{\mathbf{S}}^{\gamma}(f) \cdot 2 \left(\ln \left(\frac{1640e}{\delta} \right) + \frac{9600\tilde{c}^2 \min(c,1)d}{32\gamma^2} \Gamma \left(\frac{32m\gamma^2}{\min(c,1)d} \right) \right)}{m}} \right) \quad (52)$$

$$+ \frac{5 \left(\ln \left(\frac{1640e}{\delta} \right) + \frac{9600\tilde{c}^2 \min(c,1)d}{32\gamma^2} \Gamma \left(\frac{32m\gamma^2}{\min(c,1)d} \right) \right)}{m}$$

and the right hand side of the Eq. (52) being larger than Eq. (50) (we recall that $\tilde{c} = \max(c', 1/c)$), this implies that Eq. (52) holds with probability at least $1 - \delta$ over \mathbf{S} for any $0 < \gamma \leq 1$ and $f \in \Delta(\mathcal{H})$, which concludes the proof of Theorem 10. \square

6 Majority of Three Large Margin Classifiers

In this section we give the proof of Theorem 11, which implies Corollary 3.

To this end recall some notation. For a target concept $t \in \{-1, 1\}^{\mathcal{X}}$ we use $(\mathcal{X} \times \{-1, 1\})_t^*$ for the set of all possible training sequences on \mathcal{X} , labelled by t , that is for $S \in (\mathcal{X} \times \{-1, 1\})_t^*$, any train example $(x, y) \in S$, is such that $y = t(x)$. We remark that S is seen as a vector/sequence so it may have repetitions of similar train examples. Furthermore, for a distribution \mathcal{D} over \mathcal{X} we write \mathcal{D}_t for the distribution over $\mathcal{X} \times \{-1, 1\}$, defined by $\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_t} [(\mathbf{x}, \mathbf{y}) \in A] = \mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [(\mathbf{x}, t(\mathbf{x})) \in A]$ for any $A \subseteq \mathcal{X} \times \{-1, 1\}$. Furthermore, for $R \subset \mathcal{X}$, such that $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \in R] \neq 0$, we define $\mathcal{D}_t | R$ as

$$\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_t | R} [(\mathbf{x}, \mathbf{y}) \in A] = \mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [(\mathbf{x}, t(\mathbf{x})) \in A | \mathbf{x} \in R] = \frac{\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [(\mathbf{x}, t(\mathbf{x})) \in A, \mathbf{x} \in R]}{\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \in R]}.$$

We define a learning algorithm L as a mapping from $(\mathcal{X} \times \{-1, 1\})_t^*$ to $\mathbb{R}^{\mathcal{X}}$, for $S \in (\mathcal{X} \times \{-1, 1\})_t^*$, we write L_S for short of $L(S) \in \mathbb{R}^{\mathcal{X}}$. Furthermore, if $L_S \in \Delta(\mathcal{H})$ for any $S \in (\mathcal{X} \times \{-1, 1\})_t^*$ we write $L \in \Delta(\mathcal{H})$. For $0 < \gamma < 1$ and target concept t we define a γ -margin algorithm L for t as a mapping from $(\mathcal{X} \times \{-1, 1\})_t^*$ to $\mathbb{R}^{\mathcal{X}}$, such that for a $S \in (\mathcal{X} \times \{-1, 1\})_t^*$, we have that $L_S(x)y \geq \gamma$ for all $(x, y) \in S$. Furthermore, for three functions f_1, f_2, f_3 we define $\text{Maj}(f_1, f_2, f_3) = \text{sign}(\text{sign}(f_1) + \text{sign}(f_2) + \text{sign}(f_3))$, with $\text{sign}(0) = 0$.

With the above notation introduced, we can now state Theorem 11, which is saying that the majority vote of a margin classifier algorithm run on 3 independent training sequences implies the following error bound.

Theorem 11. For distribution \mathcal{D} over \mathcal{X} , target concept t , hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ with VC-dimension d , training sequence size m , margin $0 < \gamma < 1$ and i.i.d. training sequences $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3 \sim \mathcal{D}_t^m$, it holds for any γ -margin learning algorithm $L \in \Delta(\mathcal{H})$ for t that

$$\mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3 \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t}(\text{Maj}(L_{\mathbf{S}_1}, L_{\mathbf{S}_2}, L_{\mathbf{S}_3}))] = O\left(\frac{d}{\gamma^2 m}\right).$$

Now as AdaBoost is a $\Omega(\gamma)$ -margin learning algorithm when given access to a γ -weak learner \mathcal{W} , the above [Theorem 11](#) implies [Corollary 3](#).

To give the proof of [Theorem 11](#) we need the following lemma, which bounds the expected value of two outputs of a γ -margin learning algorithm trained on two independent training sequences erring simultaneously.

Lemma 12. There exists a universal constant $C \geq 1$ such that: For distribution \mathcal{D} over \mathcal{X} , target concept t , hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ with VC-dimension d , training sequence size m , margin $0 < \gamma < 1$, i.i.d. training sequences $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3 \sim \mathcal{D}_t^m$, it holds for any γ -margin learning algorithm $L \in \Delta(\mathcal{H})$ for t that

$$\mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2 \sim \mathcal{D}_t^m} \left[\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [\text{sign}(L_{\mathbf{S}_1}(\mathbf{x})) \neq t(\mathbf{x}), \text{sign}(L_{\mathbf{S}_2}(\mathbf{x})) \neq t(\mathbf{x})] \right] = \frac{96Cd}{\gamma^2 m}.$$

We postpone the proof of [Lemma 12](#) for later in this section and now give the proof of [Theorem 11](#).

Proof of [Theorem 11](#). We observe for $\text{Maj}(L_{\mathbf{S}_1}, L_{\mathbf{S}_2}, L_{\mathbf{S}_3})$ to fail on an example (x, y) it must be the case that two of the classifiers err, i.e. there exists $i, j \in \{1, 2, 3\}$, where $i \neq j$ such that $\text{sign}(L_{\mathbf{S}_i}(x)) \neq y, \text{sign}(L_{\mathbf{S}_j}(x)) \neq y$. Thus by a union bound, $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3$ being i.i.d. and [Lemma 12](#), we conclude that

$$\begin{aligned} & \mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3 \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t}(\text{Maj}(L_{\mathbf{S}_1}, L_{\mathbf{S}_2}, L_{\mathbf{S}_3}))] \\ & \leq \sum_{i>j} \mathbb{E}_{\mathbf{S}_i, \mathbf{S}_j \sim \mathcal{D}_t^m} \left[\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [\text{sign}(L_{\mathbf{S}_i}(\mathbf{x})) \neq t(\mathbf{x}), \text{sign}(L_{\mathbf{S}_j}(\mathbf{x})) \neq t(\mathbf{x})] \right] = \frac{288Cd}{\gamma^2 m} \end{aligned}$$

which concludes the proof. \square

We now prove [Lemma 12](#). To the end of showing [Lemma 12](#) we need the following lemma which bounds the conditional error of a large margin learning algorithm under $\mathcal{D}_t|R$.

Lemma 13. There exists a universal constant $c \geq 1$ such that: For distribution \mathcal{D} over \mathcal{X} , target concept t , hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ with VC-dimension d , training sequence size m , margin $0 < \gamma < 1$, i.i.d. training sequence $\mathbf{S} \sim \mathcal{D}_t^m$, subset $R \subseteq \mathcal{X}$ such that $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}}[R] := \mathbb{P}_{\mathbf{x} \sim \mathcal{D}}[\mathbf{x} \in R] \neq 0$, it holds for any γ -margin learning algorithm $L \in \Delta(\mathcal{H})$ for t that

$$\mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}})] \leq \frac{28cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}_{\mathbf{x} \sim \mathcal{D}}[R] \gamma^2 m}{2d}) \right)}{\mathbb{P}_{\mathbf{x} \sim \mathcal{D}}[R] \gamma^2 m}.$$

We postpone the proof of [Lemma 13](#) to after the proof of [Lemma 12](#), which we give now.

Proof of [Lemma 12](#). For $i \in \{0, 1, \dots\}$ we define the disjoint regions $R_i \subseteq \mathcal{X}$ $R_i = \{x \in \mathcal{X} : 2^{-i-1} < \mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} [\text{sign}(L_{\mathbf{S}}(x)) \neq t(x)] \leq 2^{-i}\}$ of \mathcal{X} (we will write $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [R_i]$ for short for $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [\mathbf{x} \in R_i]$). Then by the law of total probability \mathbf{S}_1 and \mathbf{S}_2 being independent we have

$$\begin{aligned} \mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2 \sim \mathcal{D}_t^m} \left[\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [\text{sign}(L_{\mathbf{S}_1}(\mathbf{x})) \neq t(\mathbf{x}), \text{sign}(L_{\mathbf{S}_2}(\mathbf{x})) \neq t(\mathbf{x})] \right] &= \mathbb{E} \left[\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [\mathbf{S} \sim \mathcal{D}_t^m [\text{sign}(L_{\mathbf{S}}(\mathbf{x})) \neq t(\mathbf{x})]^2] \right] \\ &= \sum_{i=0}^{\infty} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} [\text{sign}(L_{\mathbf{S}}(\mathbf{x})) \neq t(\mathbf{x})]^2 \mid R_i \right] \mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [R_i] \leq \sum_{i=0}^{\infty} 2^{-2i} \mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [R_i], \end{aligned} \quad (53)$$

where the last inequality follows from the definition of R_i . We will show that for each $i \in \{0, 1, 2, \dots\}$ we have that $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [R_i] \leq \frac{8C(i+1)^2 2^i d}{\gamma^2 m}$ for some universal constant $C \geq 1$. Using this, the above gives us that

$$\mathbb{E}_{\mathbf{S}_1, \mathbf{S}_2 \sim \mathcal{D}_t^m} \left[\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [\text{sign}(L_{\mathbf{S}_1}(\mathbf{x})) \neq t(\mathbf{x}), \text{sign}(L_{\mathbf{S}_2}(\mathbf{x})) \neq t(\mathbf{x})] \right] \leq \sum_{i=0}^{\infty} 2^{-2i} \frac{8C(i+1)^2 2^i d}{\gamma^2 m} = \frac{96Cd}{\gamma^2 m},$$

where the second inequality follows from $\sum_{i=0}^{\infty} 2^{-i}(i+1)^2 = 12$ and this gives the claim of [Lemma 12](#). We thus proceed to show that for each $i \in \{0, 1, 2, \dots\}$, we have that $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [R_i] \leq \frac{8C(i+1)^2 2^i d}{\gamma^2 m}$. To this end let $i \in \{0, 1, 2, \dots\}$. If $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [R_i] = 0$ then we are done, thus we consider the case that $\mathbb{P}_{\mathbf{x} \sim \mathcal{D}} [R_i] \neq 0$. We first observe that:

$$\mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t | R_i}(L_{\mathbf{S}})] = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} [\text{sign}(L_{\mathbf{S}}(\mathbf{x})) \neq t(\mathbf{x})] \mid \mathbf{x} \in R_i \right] \geq 2^{-i-1},$$

where the equality follows by the definition of $\text{sign}(0) = 0$, and the inequality follows by the definition of R_i . Furthermore, by [Lemma 13](#) we have that

$$\mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t | R_i}(L_{\mathbf{S}})] \leq \frac{28cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}[R] \gamma^2 m}{2d}) \right)}{\mathbb{P}[R] \gamma^2 m},$$

where $c \geq 1$ is a universal constant. Thus, we conclude that

$$2^{-i-1} \leq \frac{28cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}[R] \gamma^2 m}{2d}) \right)}{\mathbb{P}[R] \gamma^2 m}. \quad (54)$$

Now the function $\frac{\ln^2(\max(e^2, x))}{x}$ for $x > e^2$ is decreasing since it has derivative $\frac{2 \ln(x) - \ln^2(x)}{x^2}$, which is negative for $x > e^2$. Furthermore, we have that $\frac{\ln^2(\max(e^2, x))}{x}$ is decreasing for $0 < x \leq e^2$, so $\frac{\ln^2(\max(e^2, x))}{x}$ is decreasing for $x > 0$. Now assume for a contradiction that $\mathbb{P}[R] \geq \frac{8C(i+1)^2 2^i d}{\gamma^2 m}$ or equivalently that $\frac{\mathbb{P}[R] \gamma^2 m}{2d} \geq C(i+1)^2 2^i$ for some $C \geq e^2$ to be chosen large enough later. Thus, since we concluded that $\frac{\ln^2(\max(e^2, x))}{x}$ is decreasing for $x > 0$, and we assumed for contradiction that $\frac{\mathbb{P}[R] \gamma^2 m}{2d} \geq C(i+1)^2 2^i$ we get that

$$\frac{28cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}[R] \gamma^2 m}{2d}) \right)}{\mathbb{P}[R] \gamma^2 m} \leq \frac{14c \ln^2 \left(\max(e^2, C(i+1)^2 2^i) \right)}{C(i+1)^2 2^i} \leq \frac{14c \ln^2 \left(C(i+1)^2 2^i \right)}{C(i+1)^2 2^i}$$

where the last inequality follows from $C \geq e^2$, and $i \geq 0$. Now since it holds that $\ln(C(i+1)^{2^i}) \leq 3 \max(\ln(C), 2 \ln(i+1), i \ln(2))$ we get that the following inequality holds $\ln^2(C(i+1)^{2^i}) \leq 9 \max(\ln^2(C), 4 \ln^2(i+1), i^2 \ln^2(2))$. Thus, we conclude that

$$\frac{28cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}[R]\gamma^2 m}{2d}) \right)}{\mathbb{P}[R] \gamma^2 m} \leq \frac{126c \max(\ln^2(C), 4 \ln^2(i+1), i^2 \ln^2(2))}{C(i+1)^{2^i}}.$$

Furthermore since $\ln^2(i+1)/((i+1)^2) \leq 1, i^2/((i+1)^2) \leq 1$ and $\ln(C^{\frac{1}{4}}) \leq \ln(1+C^{\frac{1}{4}}) \leq C^{\frac{1}{4}}$ we conclude that

$$\frac{28cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}[R]\gamma^2 m}{2d}) \right)}{\mathbb{P}[R] \gamma^2 m} \leq \frac{126c \max(\ln^2(C), 4, \ln^2(2))}{C2^i} \leq \frac{126c \max(16C^{\frac{1}{2}}, 4)}{C2^i} \leq \frac{2016c}{C^{1/2}2^i}.$$

where the last inequality follows from $C \geq e^2$. Since the above is decreasing in C we get that for $C = (4 \cdot 2016)^2$, it holds that

$$\frac{28cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}[R]\gamma^2 m}{2d}) \right)}{\mathbb{P}[R] \gamma^2 m} < \frac{1}{2^{i+2}},$$

which is a contradiction with [Eq. \(54\)](#), so it must be the case that $\mathbb{P}[R] \leq \frac{8C(i+1)^{2^i}d}{\gamma^2 m}$, as claimed below [Eq. \(53\)](#) which concludes the proof. \square

We now give the proof of [Lemma 13](#).

Proof of Lemma 13. If $\frac{d}{\mathbb{P}[R]\gamma^2 m} \geq 1$ then we are done by $\mathcal{L}_{\mathcal{D}_t|R}$ always being at most 1, thus for the remainder of the proof we consider the case $\frac{d}{\mathbb{P}[R]\gamma^2 m} < 1$.

Define $\mathbf{N} = \sum_{(x,y) \in \mathbf{S}} \mathbb{1}\{x \in R\}$, i.e. the number of examples in \mathbf{S} , that has its input point in R . We notice that \mathbf{N} has expectation $\mathbb{P}[R]m$. Thus, by \mathbf{N} being a sum of i.i.d. $\{0, 1\}$ -random variables it follows by an application of Chernoff that

$$\mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathbf{N} \leq \mathbb{P}[R]m/2] \leq \exp(-\mathbb{P}[R]m/8) \leq \frac{8}{\mathbb{P}[R]m},$$

where the last inequality follows from $\exp(-x) \leq \frac{1}{x}$ for $x > 0$. Thus, from the above and the law of total probability, we conclude that

$$\mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}})] \leq \sum_{i=\lceil \mathbb{P}[R]m/2 \rceil}^m \mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}})|\mathbf{N} = i] \mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathbf{N} = i] + \frac{8}{\mathbb{P}[R]m}. \quad (55)$$

For each $i \in \{\lceil \mathbb{P}[R]m/2 \rceil, \dots, m\}$, we will show that $\mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}})|\mathbf{N} = i]$ is upper bound by $\frac{20cd \ln^2(\max(e^2, \frac{\mathbb{P}[R]\gamma^2 m}{2d}))}{\mathbb{P}[R]\gamma^2 m}$, where $c \geq 1$ is a universal constant, which implies that

$$\mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}})] \leq \frac{20cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}[R]\gamma^2 m}{2d}) \right)}{\mathbb{P}[R] \gamma^2 m} + \frac{8}{\mathbb{P}[R]m} \leq \frac{28cd \ln^2 \left(\max(e^2, \frac{\mathbb{P}[R]\gamma^2 m}{2d}) \right)}{\mathbb{P}[R] \gamma^2 m}$$

as claimed. We now show for each $i \in \{\lceil \mathbb{P}[R]m/2 \rceil, \dots, m\}$ that $\mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}})|\mathbf{N} = i] \leq \frac{20cd \ln^2(\max(e^2, \frac{\mathbb{P}[R]\gamma^2 m}{2d}))}{\mathbb{P}[R]\gamma^2 m}$.

Let for now $i \in \{\lceil \mathbb{P}[R]m/2 \rceil, \dots, m\}$. First since $\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}})$ is nonnegative we have that its expectation can be calculated in terms of its cumulative distribution function,

$$\begin{aligned} \mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}}) | \mathbf{N} = i] &= \int_0^\infty \mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}}) > x | \mathbf{N} = i] dx \\ &\leq \frac{4cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i} + \int_{\frac{4cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i}}^\infty \mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}}) > x | \mathbf{N} = i] dx. \end{aligned} \quad (56)$$

We now notice that under the conditional distribution $\mathbf{N} = i$, it is the case that the γ -margin learning algorithm $L_{\mathbf{S}}$, contains i labelled examples from $\mathcal{D}_t|R$. Furthermore, since $L_{\mathbf{S}}$ has $\mathcal{L}_{\mathbf{S}}^\gamma(L_{\mathbf{S}}) = 0$ it also has zero margin-loss on the examples drawn from $\mathcal{D}_t|R$. Thus, by invoking [Theorem 10](#), it holds with probability at least $1 - \delta$ over $\mathbf{S} \sim \mathcal{D}_t^m$ conditioned on $\mathbf{N} = i$, that

$$\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}}) \leq \frac{cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i} + \frac{c \ln\left(\frac{\epsilon}{\delta}\right)}{i} \leq \max\left(\frac{2cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i}, \frac{2c \ln\left(\frac{\epsilon}{\delta}\right)}{i}\right) \quad (57)$$

where we have upper bounded Γ by Ln^2 , which holds since for $x \leq e^e$ we have that $\Gamma(x) = \text{Ln}^2(\text{Ln}(x)) \text{Ln}(x) = \text{Ln}(x) \leq \text{Ln}^2(x)$ and for $x > e^e$ $\Gamma(x) = \text{Ln}^2(\text{Ln}(x)) \text{Ln}(x) \leq \text{Ln}^2(x)$, and used that $a + b \leq 2 \max(a, b)$ for $a, b \geq 0$, and $c \geq 1$ being the universal constant of [Theorem 10](#). For $x > 2c/i$, we now notice that if we set $\delta = e \cdot \exp(-ix/2c)$, which is strictly less than 1 by $x > 2c/i$, we conclude from [Eq. \(57\)](#) that

$$\mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} \left[\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}}) > \max\left(\frac{2cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i}, x\right) \mid \mathbf{N} = i \right] \leq e \cdot \exp(-ix/2c),$$

Furthermore, we notice that if $0 < x \leq 2c/i$ then the above right-hand side is at least 1, which is also upper bounding the left-hand side since it is at most 1, thus the above holds for any $x > 0$. Thus, plugging this into [Eq. \(56\)](#) we get that

$$\begin{aligned} \mathbb{E}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}}) | \mathbf{N} = i] &\leq \frac{4cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i} + \int_{\frac{4cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i}}^\infty \mathbb{P}_{\mathbf{S} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t|R}(L_{\mathbf{S}}) > x | \mathbf{N} = i] dx \\ &\leq \frac{4cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i} + \int_{\frac{4cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i}}^\infty e \cdot \exp(-ix/2c) dx \\ &\leq \frac{4cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i} + \frac{2ec}{i} \cdot \exp\left(-\frac{4cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{2c\gamma^2}\right) \\ &\leq \frac{10cd \text{Ln}^2\left(\frac{\gamma^2 i}{d}\right)}{\gamma^2 i} \leq \frac{10cd \ln^2\left(\max(e^2, \frac{\gamma^2 i}{d})\right)}{\gamma^2 i}, \end{aligned}$$

where we have used in the third inequality that $\int \exp(-ax) dx = -\exp(-ax)/a + C$, in the second to last inequality that $4 + 2e \leq 10$, and in the last that $\text{Ln}(x) = \ln(\max(e, x)) \leq$

$\ln(\max(e^2, x))$. Now for $\frac{\gamma^2 i}{d} \leq e^2$, $\frac{10cd \ln^2(\max(e^2, \frac{\gamma^2 i}{d}))}{\gamma^2 i}$ is a decreasing function in $\frac{\gamma^2 i}{d}$. Furthermore, since $\ln^2(x)/x$ has derivative $\frac{2 \ln(x) - \ln^2(x)}{x^2}$ we conclude that $\ln^2(x)/x$ is decreasing for $x \geq e^2$, whereby we conclude that $\frac{10cd \ln^2(\max(e^2, \frac{\gamma^2 i}{d}))}{\gamma^2 i}$ is decreasing in $\frac{\gamma^2 i}{d}$ for $\frac{\gamma^2 i}{d} \geq e^2$, so we conclude that for $\frac{\gamma^2 i}{d} > 0$ the function $\frac{10cd \ln^2(\max(e^2, \frac{\gamma^2 i}{d}))}{\gamma^2 i}$ is decreasing in $\frac{\gamma^2 i}{d}$. Now $i \geq \lceil \mathbb{P}[R] m/2 \rceil \geq \mathbb{P}[R] m/2 > 0$ thus we have $\frac{\gamma^2 i}{d} \geq \frac{\mathbb{P}[R] \gamma^2 m}{2d}$, which by the above argued monotonicity implies that $\frac{10cd \ln^2(\max(e^2, \frac{\gamma^2 i}{d}))}{\gamma^2 i} \leq \frac{2 \cdot 10cd \ln^2(\max(e^2, \frac{\mathbb{P}[R] \gamma^2 m}{2d}))}{\mathbb{P}[R] \gamma^2 m}$. Thus, we have argued that

$$\mathbb{E}_{\mathbf{s} \sim \mathcal{D}_t^m} [\mathcal{L}_{\mathcal{D}_t | R}(L_{\mathbf{s}}) | \mathbf{N} = i] \leq \frac{10cd \ln^2(\max(e^2, \frac{\mathbb{P}[R] \gamma^2 m}{2d}))}{\mathbb{P}[R] \gamma^2 m},$$

which shows the claim below [Eq. \(55\)](#) and concludes the proof. \square

7 Acknowledgments

Mikael Møller Høgsgaard is funded by a DFF Sapere Aude Research Leader Grant No. 9064-00068B by the Independent Research Fund Denmark. Kasper Green Larsen is co-funded by DFF Grant No. 9064-00068B and co-funded by the European Union (ERC, TUCCLA, 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.

References

- [Aden-Ali et al., 2024] Aden-Ali, I., Høgsgaard, M. M., Larsen, K. G., and Zhivotovskiy, N. (2024). Majority-of-three: The simplest optimal learner? In Agrawal, S. and Roth, A., editors, *The Thirty Seventh Annual Conference on Learning Theory, June 30 - July 3, 2023, Edmonton, Canada*, volume 247 of *Proceedings of Machine Learning Research*, pages 22–45. PMLR.
- [Breiman, 1996] Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24:123–140.
- [Breiman, 1999] Breiman, L. (1999). Prediction games and arcing algorithms. *Neural Comput.*, 11(7):1493–1517.
- [Dudley, 1978] Dudley, R. M. (1978). Central Limit Theorems for Empirical Measures. *The Annals of Probability*, 6(6):899 – 929.
- [Freund and Schapire, 1997] Freund, Y. and Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139.
- [Gao and Zhou, 2013] Gao, W. and Zhou, Z.-H. (2013). On the doubt about margin explanation of boosting. *Artificial Intelligence*, 203:1–18.

- [Grønlund et al., 2020] Grønlund, A., Kamma, L., and Larsen, K. G. (2020). Margins are insufficient for explaining gradient boosting. In *NeurIPS*.
- [Grønlund et al., 2019] Grønlund, A., Kamma, L., Larsen, K. G., Mathiasen, A., and Nelson, J. (2019). Margin-based generalization lower bounds for boosted classifiers. In *NeurIPS*, pages 11940–11949.
- [Hajek and Raginsky, 2021] Hajek, B. and Raginsky, M. (2021). ECE 543: Statistical Learning Theory. *Department of Electrical and Computer Engineering and the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign*. Last updated March 18, 2021.
- [Hanneke, 2016] Hanneke, S. (2016). The optimal sample complexity of PAC learning. *J. Mach. Learn. Res.*, 17:38:1–38:15.
- [Hoeffding, 1963] Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30.
- [Høgsgaard et al., 2024] Høgsgaard, M. M., Larsen, K. G., and Mathiasen, M. E. (2024). The many faces of optimal weak-to-strong learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- [Høgsgaard et al., 2023] Høgsgaard, M. M., Larsen, K. G., and Ritzert, M. (2023). AdaBoost is not an optimal weak to strong learner. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 13118–13140. PMLR.
- [Kearns and Valiant, 1988] Kearns, M. and Valiant, L. (1988). *Learning Boolean Formulae Or Finite Automata is as Hard as Factoring*. Technical report (Harvard University. Aiken Computation Laboratory). Harvard University, Center for Research in Computing Technology, Aiken Computation Laboratory.
- [Kearns and Valiant, 1994] Kearns, M. and Valiant, L. (1994). Cryptographic limitations on learning boolean formulae and finite automata. *J. ACM*, 41(1):67–95.
- [Larsen, 2023] Larsen, K. G. (2023). Bagging is an optimal PAC learner. In Neu, G. and Rosasco, L., 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.
- [Larsen and Ritzert, 2022] Larsen, K. G. and Ritzert, M. (2022). Optimal weak to strong learning.
- [Rudelson and Vershynin, 2006] Rudelson, M. and Vershynin, R. (2006). Combinatorics of random processes and sections of convex bodies. *Annals of Mathematics*, 164.
- [Schapire and Freund, 2012] Schapire, R. E. and Freund, Y. (2012). *Boosting: Foundations and Algorithms*. The MIT Press.

[Schapire et al., 1998] Schapire, R. E., Freund, Y., Bartlett, P., and Lee, W. S. (1998). Boosting the margin: A new explanation for the effectiveness of voting methods. *Annals of Statistics*, 26(5):1651 – 1686.