

Bagging is an Optimal PAC Learner

Kasper Green Larsen*
Aarhus University

Abstract

Determining the optimal sample complexity of PAC learning in the realizable setting was a central open problem in learning theory for decades. Finally, the seminal work by Hanneke (2016) gave an algorithm with a provably optimal sample complexity. His algorithm is based on a careful and structured sub-sampling of the training data and then returning a majority vote among hypotheses trained on each of the sub-samples. While being a very exciting theoretical result, it has not had much impact in practice, in part due to inefficiency, since it constructs a polynomial number of sub-samples of the training data, each of linear size.

In this work, we prove the surprising result that the practical and classic heuristic *bagging* (a.k.a. bootstrap aggregation), due to Breiman (1996), is in fact also an optimal PAC learner. Bagging pre-dates Hanneke’s algorithm by twenty years and is taught in most undergraduate machine learning courses. Moreover, we show that it only requires a logarithmic number of sub-samples to reach optimality.

1 Introduction

PAC learning, or probably approximately correct learning (Valiant [15]), is the most classic theoretical model for studying classification problems in supervised learning. For binary classification in the *realizable* setting, the goal is to design a learning algorithm that with probability $1 - \delta$ over a random training data set, outputs a hypothesis that mispredicts the label of a new random sample with probability at most ε . More formally, one assumes that samples come from an input domain \mathcal{X} and that there is an unknown *concept* $c : \mathcal{X} \rightarrow \{-1, 1\}$ that we are trying to learn. The realizable setting means that c belongs to a predefined concept class $\mathbb{C} \subseteq \mathcal{X} \rightarrow \{-1, 1\}$ and that the correct label of any $x \in \mathcal{X}$ is always $c(x)$.

For the above learning task, a learning algorithm \mathcal{A} receives a training data set \mathbf{S} of m i.i.d. samples $(\mathbf{x}_1, c(\mathbf{x}_1)), \dots, (\mathbf{x}_m, c(\mathbf{x}_m))$ where each \mathbf{x}_i is drawn independently from an *unknown* data distribution \mathcal{D} over \mathcal{X} . From this data set, the learning algorithm must output a hypothesis $h_{\mathbf{S}} : \mathcal{X} \rightarrow \{-1, 1\}$. The algorithm \mathcal{A} is a PAC learner, if for any distribution \mathcal{D} and any concept $c \in \mathbb{C}$, it holds that if \mathcal{A} is given enough i.i.d. training samples \mathbf{S} , then with probability at least $1 - \delta$, the hypothesis $h_{\mathbf{S}}$ that it outputs satisfies $\mathcal{L}_{\mathcal{D}}(h_{\mathbf{S}}) = \Pr_{\mathbf{x} \sim \mathcal{D}}[h(\mathbf{x}) \neq c(\mathbf{x})] \leq \varepsilon$. We remark that the algorithm \mathcal{A} knows the concept class \mathbb{C} , but not the data distribution \mathcal{D} . Determining the minimum number of samples $\mathcal{M}(\varepsilon, \delta)$, as a function of ε , δ and the VC-dimension [17] d of \mathbb{C} (see Section 2 for a formal definition of VC-dimension) needed for this learning task, is one of the fundamental problems in PAC learning.

The most natural learning algorithm for the above is *empirical risk minimization* (ERM). Here a learning algorithm simply outputs an arbitrary hypothesis/concept $h_{\mathbf{S}} \in \mathbb{C}$ that correctly predicts the labels of the training data, i.e. it has $h_{\mathbf{S}}(\mathbf{x}_i) = c(\mathbf{x}_i)$ for all $(\mathbf{x}_i, c(\mathbf{x}_i)) \in \mathbf{S}$. Clearly such a hypothesis exists since $c \in \mathbb{C}$. Such a learning algorithm is referred to as a *proper* learner as it outputs a hypothesis/concept from the concept class \mathbb{C} . ERM is known to obtain a sample complexity of $O(\varepsilon^{-1}(d \lg(1/\varepsilon) + \lg(1/\delta)))$ [16, 3]. Moreover, it can be shown that this analysis cannot be tightened, i.e. there are distributions \mathcal{D} and concept classes \mathbb{C} where any proper learner needs $\Omega(\varepsilon^{-1}(d \lg(1/\varepsilon) + \lg(1/\delta)))$ samples [1]. However, a PAC learning

*larsen@cs.au.dk, Supported by Independent Research Fund Denmark (DFR) Sapere Aude Research Leader grant No 9064-00068B.

algorithm is not necessarily required to output a hypothesis $h \in \mathbb{C}$. That better strategies might exist may seem counter-intuitive at first, since we are promised that the unknown concept c lies in \mathbb{C} . Nonetheless, the strongest known lower bounds for arbitrary PAC learning algorithms only show that $\Omega(\varepsilon^{-1}(d + \lg(1/\delta)))$ samples are necessary [3, 6]. This leaves a gap of a factor $\lg(1/\varepsilon)$ between ERM and the lower bound for arbitrary algorithms.

Despite its centrality, closing this gap remained a big open problem for more than thirty years. Finally, in 2016, Hanneke [9] built on ideas by Simon [14] and presented an algorithm with an asymptotically optimal sample complexity of $\mathcal{M}(\varepsilon, \delta) = O(\varepsilon^{-1}(d + \lg(1/\delta)))$. His algorithm is based on constructing a number of subsets $\mathbf{S}_i \subset \mathbf{S}$ of the training data \mathbf{S} with carefully designed overlaps between the \mathbf{S}_i 's (see Section 2). He then runs ERM on each \mathbf{S}_i to obtain hypotheses $h_{\mathbf{S}_i} \in \mathbb{C}$ and finally outputs the hypothesis $f_{\mathbf{S}}$ taking the majority vote $f_{\mathbf{S}}(x) = \text{sign}(\sum_i h_{\mathbf{S}_i}(x))$ among the $h_{\mathbf{S}_i}$'s.

While being a major theoretical breakthrough, Hanneke's algorithm has unfortunately not had any significant practical impact. One explanation is that it requires a rather large number of sub-samples \mathbf{S}_i . Concretely, with optimal $m = \Theta(\varepsilon^{-1}(d + \lg(1/\delta)))$ samples, it requires $m^{\lg 4} \approx m^{0.79}$ sub-samples of linear size $|\mathbf{S}_i| = \Omega(m)$, resulting in a somewhat slow learning algorithm.

Our Contribution. In this work, we present an alternative optimal PAC learner in the realizable setting. Surprisingly, our algorithm is not new, but actually pre-dates Hanneke's algorithm by twenty years. Concretely, we show that the heuristic known as bagging (bootstrap aggregation) by Breiman [4], also gives an optimal PAC learner. Bagging, and its slightly more involved extension known as *random forest* [5], have proved very efficient in practice and are classic topics in introduction to machine learning courses.

In bagging, for t iterations, we sample a subset \mathbf{S}_i of n independent and uniform samples with replacement from \mathbf{S} . We then run ERM on each \mathbf{S}_i to produce hypotheses $h_{\mathbf{S}_i} \in \mathbb{C}$ and finally output the hypothesis $f_{\mathbf{S}_1, \dots, \mathbf{S}_t}$ taking the majority vote $f_{\mathbf{S}_1, \dots, \mathbf{S}_t}(x) = \text{sign}(\sum_i h_{\mathbf{S}_i}(x))$ among the $h_{\mathbf{S}_i}$'s. The sub-samples \mathbf{S}_i are referred to as *bootstrap samples*.

While being similar to Hanneke's algorithm, it is simpler to construct the subsets \mathbf{S}_i , and also, we show that it suffices with just $t = O(\lg(m/\delta))$ bootstrap samples of size n for any $0.02m \leq n \leq m$, which should be compared to Hanneke's algorithm requiring $m^{0.79}$ subsets (where m is optimal $\Theta(\varepsilon^{-1}(d + \lg(1/\delta)))$). For ease of notation, let \mathcal{D}_c denote the distribution of a pair $(\mathbf{x}, c(\mathbf{x}))$ with $\mathbf{x} \sim \mathcal{D}$. Our result is then formalized in the following theorem

Theorem 1. *There is a universal constant $a > 0$ such that for every $0 < \delta < 1$, every distribution \mathcal{D} over an input domain \mathcal{X} , every concept class $\mathbb{C} \subseteq \mathcal{X} \rightarrow \{-1, 1\}$ of VC-dimension d and every $c \in \mathbb{C}$, if $t \geq 18 \ln(2m/\delta)$ and $0.02m \leq n \leq m$, then it holds with probability at least $1 - \delta$ over the random choice of a training set $\mathbf{S} \sim \mathcal{D}_c^m$ and t bootstrap samples $\mathbf{S}_1, \dots, \mathbf{S}_t \subset \mathbf{S}$ of size n , that the hypothesis $f_{\mathbf{S}_1, \dots, \mathbf{S}_t}$ produced by bagging satisfies*

$$\mathcal{L}_{\mathcal{D}}(f_{\mathbf{S}_1, \dots, \mathbf{S}_t}) = \Pr_{\mathbf{x} \sim \mathcal{D}} [f_{\mathbf{S}_1, \dots, \mathbf{S}_t}(\mathbf{x}) \neq c(\mathbf{x})] \leq a \cdot \frac{d + \ln(1/\delta)}{m}.$$

Solving for m such that $\varepsilon \leq \mathcal{L}_{\mathcal{D}}(g)$ gives a sample complexity of $O(\varepsilon^{-1}(d + \lg(1/\delta)))$ as claimed. We remark that the constants 18 and 0.02 can be reduced at the cost of increasing the unspecified constant a .

In addition to providing an alternative and simpler algorithm for optimal PAC learning, we also believe there is much value in providing further theoretical justification for the wide practical success of bagging.

As an interesting secondary result, if we combine our proof with a recent work by Larsen and Ritzert [12], we get that bagging combined with a version of AdaBoost gives an optimal weak to strong learner [10, 11]. We find it quite remarkable that the combination of the classic techniques bagging and boosting yields this result.

In Section 2, we first present Hanneke's optimal PAC learner and highlight the main ideas in his analysis. We then proceed to present a high-level overview of our proof of Theorem 1, which re-uses some of the ideas from Hanneke's proof. In Section 3 we proceed to give the formal proof of Theorem 1. Finally, in Section 4, we comment further on the weak to strong learning result, as well as discuss some directions for further research.

2 Proof overview

In this section, we first present Hanneke’s PAC learning algorithm and discuss the main ideas in his analysis. We then proceed to give a high-level overview of the keys ideas in our proof that bagging is also an optimal PAC learning algorithm. For completeness, we start by recalling the definition of Vapnik-Chervonenkis dimension [17], or VC-dimension for short.

A concept class $\mathbb{C} \subseteq \mathcal{X} \rightarrow \{-1, 1\}$ has VC-dimension d , where d is the largest integer such that there exists a set of d samples $x_1, \dots, x_d \in \mathcal{X}$ for which any labeling of the d samples can be realized by a concept $c \in \mathbb{C}$. That is, $|\{(c(x_1), \dots, c(x_d)) : c \in \mathbb{C}\}| = 2^d$. Throughout the paper, we assume $d \geq 1$ which is always true when \mathbb{C} contains at least two distinct concepts. Also, as the reader may have observed, we consistently use **bold** face letters to denote random variables.

Hanneke’s Algorithm and Analysis. As mentioned in Section 1, Hanneke’s algorithm constructs a carefully selected collection of sub-samples of the training set \mathbf{S} . These sub-samples are constructed by invoking Algorithm 1 as $Sub\text{-}Sample(\mathbf{S}, \emptyset)$.

Algorithm 1: $Sub\text{-}Sample(U, V)$

Input: Two sets of training samples U, V .
Result: Collection consisting of sub-samples.

- 1 **if** $|U| < 4$ **then**
- 2 | **return** $\{U \cup V\}$.
- 3 **else**
- 4 | Partition U into 4 disjoint sets U_0, U_1, U_2, U_3 of $|U|/4$ samples each.
- 5 | **return** $\bigcup_{i=1}^3 Sub\text{-}Sample(U_0, V \cup (\bigcup_{j \in \{1,2,3\} \setminus \{i\}} U_j))$.

For simplicity, we have presented the $Sub\text{-}Sample$ algorithm assuming that m is a power of 4.

Since $|U|$ is reduced by a factor 4 in each recursive call and there are 3 such calls, we get that the total number of sub-samples produced is $3^{\lg_4 m} = m^{\lg_4 3} \approx m^{0.79}$.

To analyse the hypothesis produced by invoking $Sub\text{-}Sample(\mathbf{S}, \emptyset)$, running ERM on each produced sub-sample \mathbf{S}_i to produce hypotheses $h_{\mathbf{S}_i}$ and taking a majority vote $f_{\mathbf{S}}(x) = \text{sign}(\sum_i h_{\mathbf{S}_i}(x))$, we zoom in on the recursive invocations of $Sub\text{-}Sample(\mathbf{U}, \mathbf{V})$. Such an invocation produces a number of sub-samples $\mathbf{S}_1, \dots, \mathbf{S}_t$ of $\mathbf{U} \cup \mathbf{V} \subseteq \mathbf{S}$. Letting $h_{\mathbf{S}_1}, \dots, h_{\mathbf{S}_t}$ denote the hypotheses obtained by running ERM on these sub-samples and $f_{(\mathbf{U}, \mathbf{V})}(x) = \text{sign}(\sum_i h_{\mathbf{S}_i}(x))$ the majority vote among them, we prove by induction (with the base being the leaves of the recursion) that with probability at least $1 - \delta$ over \mathbf{U} , it holds that $\mathcal{L}_{\mathcal{D}}(f_{(\mathbf{U}, \mathbf{V})}) \leq a \cdot (d + \ln(1/\delta))/|\mathbf{U}|$ for a universal constant $a > 0$. If we can complete this inductive step, then the conclusion follows by examining the root invocation $Sub\text{-}Sample(\mathbf{S}, \emptyset)$.

The base case in the inductive proof is simply when $a \cdot (d + \ln(1/\delta))/|\mathbf{U}| > 1$. Here the conclusion follows trivially as we always have $\mathcal{L}_{\mathcal{D}}(f_{(\mathbf{U}, \mathbf{V})}) \leq 1$. For the inductive step, let $f_{1,(\mathbf{U}, \mathbf{V})}, f_{2,(\mathbf{U}, \mathbf{V})}$ and $f_{3,(\mathbf{U}, \mathbf{V})}$ denote the majority voters produced by the three recursive calls $Sub\text{-}Sample(\mathbf{U}_0, \mathbf{V} \cup \mathbf{U}_2 \cup \mathbf{U}_3)$, $Sub\text{-}Sample(\mathbf{U}_0, \mathbf{V} \cup \mathbf{U}_1 \cup \mathbf{U}_3)$ and $Sub\text{-}Sample(\mathbf{U}_0, \mathbf{V} \cup \mathbf{U}_1 \cup \mathbf{U}_2)$. Each of these have $\mathcal{L}_{\mathcal{D}}(f_{i,(\mathbf{U}, \mathbf{V})}) \leq a \cdot (d + \ln(1/\delta))/|\mathbf{U}_i| = 4a \cdot (d + \ln(1/\delta))/|\mathbf{U}|$ by the induction hypothesis (except with some probability δ which we ignore here for simplicity). For short, we say that a hypothesis h errs on x if $h(x) \neq c(x)$. The crux of the argument is now to show that it is very unlikely that $f_{i,(\mathbf{U}, \mathbf{V})}$ errs on a sample $\mathbf{x} \sim \mathcal{D}$ at the same time as a hypothesis $h_{\mathbf{S}'}$ trained on a sub-sample \mathbf{S}' produced by a recursive call $j \neq i$ also errs on \mathbf{x} .

For this step, let us wlog. consider the majority vote $f_{1,(\mathbf{U}, \mathbf{V})}$ (over hypotheses produced by ERM on $Sub\text{-}Sample(\mathbf{U}_0, \mathbf{V} \cup \mathbf{U}_2 \cup \mathbf{U}_3)$). Intuitively, if $\mathcal{L}_{\mathcal{D}}(f_{1,(\mathbf{U}, \mathbf{V})}) \ll a \cdot (d + \ln(1/\delta))/|\mathbf{U}|$ then the hypotheses in $f_{1,(\mathbf{U}, \mathbf{V})}$ contribute little to $\mathcal{L}_{\mathcal{D}}(f_{(\mathbf{U}, \mathbf{V})})$. So assume instead $\mathcal{L}_{\mathcal{D}}(f_{1,(\mathbf{U}, \mathbf{V})}) \approx 4a \cdot (d + \ln(1/\delta))/|\mathbf{U}|$. Consider now some hypothesis $h_{\mathbf{S}'}$ obtained by running ERM on a sub-sample \mathbf{S}' produced by the recursive call $Sub\text{-}Sample(\mathbf{U}_0, \mathbf{V} \cup \mathbf{U}_1 \cup \mathbf{U}_3)$. The key observation and property of the $Sub\text{-}Sample$ algorithm, is that *all* hypotheses in the majority vote $f_{1,(\mathbf{U}, \mathbf{V})}$ have been trained on sub-samples that *exclude* all of \mathbf{U}_1 . This means

that the samples \mathbf{U}_1 are independent of $f_{1,(\mathbf{U},\mathbf{V})}$. When $\mathcal{L}_{\mathcal{D}}(f_{1,(\mathbf{U},\mathbf{V})}) \approx 4a \cdot (d + \ln(1/\delta))/|\mathbf{U}|$, we will now see about $|\mathbf{U}_1|4a \cdot (d + \ln(1/\delta))/|\mathbf{U}| = a \cdot (d + \ln(1/\delta))$ samples $(\mathbf{x}, c(\mathbf{x}))$ in \mathbf{U}_1 for which $f_{1,(\mathbf{U},\mathbf{V})}(\mathbf{x}) \neq c(\mathbf{x})$. The observation is that, conditioned on $f_{1,(\mathbf{U},\mathbf{V})}$, these samples are i.i.d. from the conditional distribution $\mathcal{D}(\cdot \mid f_{1,(\mathbf{U},\mathbf{V})} \text{ errs})$. The second key observation is that $h_{\mathbf{S}'}$ is obtained by ERM on a sub-sample \mathbf{S}' that *includes* all of \mathbf{U}_1 (we add \mathbf{U}_1 to \mathbf{V} in both of the other recursive calls). Moreover, since we are in the realizable setting, we have $h_{\mathbf{S}'}(\mathbf{x}) = c(\mathbf{x})$ for every $\mathbf{x} \in \mathbf{U}_1$. In particular, this holds for all the samples where $f_{1,(\mathbf{U},\mathbf{V})}(\mathbf{x}) \neq c(\mathbf{x})$. The classic sample complexity bounds for proper PAC learning in the realizable setting then implies that $h_{\mathbf{S}'}$ has $\mathcal{L}_{\mathcal{D}(\cdot \mid f_{1,(\mathbf{U},\mathbf{V})} \text{ errs})}(h_{\mathbf{S}'}) = O((d + \ln(1/\delta))/(a \cdot (d + \ln(1/\delta)))) \leq 1/200$ for a sufficiently large. Note that this is under the conditional distribution $\mathcal{D}(\cdot \mid f_{1,(\mathbf{U},\mathbf{V})} \text{ errs})$. That is, $h_{\mathbf{S}'}$ rarely errs when $f_{1,(\mathbf{U},\mathbf{V})}$ errs. We thus get that $\Pr[f_{1,(\mathbf{U},\mathbf{V})}(\mathbf{x}) \neq c(\mathbf{x}) \wedge h_{\mathbf{S}'}(\mathbf{x}) \neq c(\mathbf{x})] = \Pr[f_{1,(\mathbf{U},\mathbf{V})}(\mathbf{x}) \neq c(\mathbf{x})] \cdot \Pr[h_{\mathbf{S}'}(\mathbf{x}) \neq c(\mathbf{x}) \mid f_{1,(\mathbf{U},\mathbf{V})}(\mathbf{x}) \neq c(\mathbf{x})] \leq (a/50) \cdot (d + \ln(1/\delta))/|\mathbf{U}|$. Since this holds for every $f_{i,(\mathbf{U},\mathbf{V})}$ and $h_{\mathbf{S}'}$ from a recursive call $j \neq i$, we can now argue that $\mathcal{L}_{\mathcal{D}}(f_{(\mathbf{U},\mathbf{V})}) \ll a \cdot (d + \ln(1/\delta))/|\mathbf{U}|$. To see this, note first that for $f_{(\mathbf{U},\mathbf{V})}$ to err on an $x \in \mathcal{X}$, it must be the case that at least one $f_{i,(\mathbf{U},\mathbf{V})}$ also errs. Even in this case, since one recursive call only contributes a third of the hypotheses in the majority vote $f_{(\mathbf{U},\mathbf{V})}$, there must be many hypotheses $h_{\mathbf{S}'}$ (at least a $1/2 - 1/3 = 1/6$ fraction of all hypotheses) trained from sub-samples \mathbf{S}' produced by the recursive calls $j \neq i$ that also err on x . But we have just argued that it is very unlikely that both $f_{i,(\mathbf{U},\mathbf{V})}$ and such an $h_{\mathbf{S}'}$ err at the same time. Formalizing this intuition completes the inductive proof.

Bagging Analysis. We now turn to presenting the key ideas in our proof of Theorem 1, i.e. that bagging is an optimal PAC learner. Along the way, we also discuss the issues we encounter towards establishing the result. Recall that in bagging with a training set $\mathbf{S} \sim \mathcal{D}_c^m$, we randomly sub-sample t bootstrap samples $\mathbf{S}_1, \dots, \mathbf{S}_t \subset \mathbf{S}$ each consisting of n i.i.d. samples with replacement from \mathbf{S} . We then run ERM on each \mathbf{S}_i to produce hypotheses $h_{\mathbf{S}_1}, \dots, h_{\mathbf{S}_t}$ and finally return the majority vote $f_{\mathbf{S}_1, \dots, \mathbf{S}_t}(x) = \text{sign}(\sum_i h_{\mathbf{S}_i}(x))$. It will be convenient for us to think of $f_{\mathbf{S}_1, \dots, \mathbf{S}_t}$ in a slightly different way. Concretely, we instead let $f_{\mathbf{S}_1, \dots, \mathbf{S}_t}(x) = (1/t) \sum_i h_{\mathbf{S}_i}(x)$ be a *voting classifier*. Then $\text{sign}(f_{\mathbf{S}_1, \dots, \mathbf{S}_t}(x)) \neq c(x)$ if and only if $f_{\mathbf{S}_1, \dots, \mathbf{S}_t}(x)c(x) \leq 0$. We thus seek to bound $\mathcal{L}_{\mathcal{D}}(f_{\mathbf{S}_1, \dots, \mathbf{S}_t}) = \Pr_{\mathbf{x} \sim \mathcal{D}}[f_{\mathbf{S}_1, \dots, \mathbf{S}_t}(\mathbf{x})c(\mathbf{x}) \leq 0]$. The motivation for thinking about $f_{\mathbf{S}_1, \dots, \mathbf{S}_t}$ as a voting classifier, is that it allows us to re-use some of the ideas that appear in proving generalization bounds for AdaBoost [7] and other voting classifiers.

We first observe that similarly to Hanneke's sub-sampling, if we look at just two hypotheses $h_{\mathbf{S}_i}$ and $h_{\mathbf{S}_j}$ with $i \neq j$, then $h_{\mathbf{S}_i}$ is trained on a bootstrap sample \mathbf{S}_i leaving out a rather large portion of \mathbf{S} . Furthermore, $h_{\mathbf{S}_j}$ will be trained on most of these left-out samples and thus one could again argue that it is unlikely that $h_{\mathbf{S}_i}$ and $h_{\mathbf{S}_j}$ err at the same time. Unfortunately, this line of argument fails when we start combining a non-constant number of hypotheses. In particular, with high probability over the bootstrap samples, the union of any set of ℓ bootstrap samples contains all but an $\exp(-\Omega(\ell))$ -fraction of \mathbf{S} . This leaves very few samples that are independent of the hypotheses trained on such ℓ bootstrap samples. Trying to repeat Hanneke's argument unfortunately requires $\Omega(m)$ independent samples towards the last steps of an inductive proof. In a nutshell, what saves Hanneke's construction is that a third all sub-samples together still leave out a quarter of the training data. For bagging, such a property is just not true if we have more than a constant number of bootstrap samples.

Abandoning the hope of directly applying Hanneke's line of reasoning, we instead start by relating the performance of bagging to that of a particular voting classifier that is deterministically determined from a training set S , i.e. we get rid of the bootstrap samples. To formalize this, we first introduce some notation. From a training set S of m samples $(x_1, c(x_1)), \dots, (x_m, c(x_m))$ and a vector of n not necessarily distinct integers $I = (i_1, \dots, i_n) \in [m]^n$, let $S(I)$ denote the bootstrap sample $(x_{i_1}, c(x_{i_1})), \dots, (x_{i_n}, c(x_{i_n}))$. Then a random bootstrap sample \mathbf{S}_i from S has the same distribution as if we draw \mathbf{I} uniformly from $[m]^n$ and let $\mathbf{S}_i = S(\mathbf{I})$. Also, let $h_{S(I)} \in \mathbb{C}$ denote the hypothesis resulting from running ERM on $S(I)$. Finally, for a list of t vectors $B = (I_1, \dots, I_t) \in [m]^{n \times t}$, we let $f_{S,B} = (1/t) \sum_{i=1}^t h_{S(I_i)}$ denote the voting classifier produced by bagging with bootstrap samples $S(I_1), \dots, S(I_t)$. Using the notation $\mathbf{B} \sim [m]^{n \times t}$ to denote a uniform random \mathbf{B} from $[m]^{n \times t}$ we thus have that the hypothesis produced by bagging on S has the same distribution as $f_{S,\mathbf{B}}$.

Now consider the following voting classifier

$$g_S(x) := \frac{1}{m^n} \sum_{I \in [m]^n} h_{S(I)}(x).$$

That is, g_S is the voting classifier averaging the predictions over all m^n possible bootstrap samples of S . Of course one would never compute g_S . Nonetheless, the performance of the random $f_{S, \mathbf{B}}$ with $\mathbf{B} \sim [m]^{n \times t}$ is closely related to that of g_S . To see this, we introduce the notion of *margins* which are typically used in the study of generalization of voting classifiers, see e.g. the works [2, 8, 12]. For a sample $x \in \mathcal{X}$ and voting classifier $f(x) = (1/t) \sum_{i=1}^t h_i(x)$, we say that f has margin $f(x)c(x)$ on the sample x . Since each $h_i(x)$ is in $\{-1, 1\}$, we have the margin is a number between -1 and 1 . Intuitively, 1 represents that all the hypotheses h_i agree on the label of x and those predictions are correct. In general, a margin of γ implies that an α -fraction of the hypotheses h_i are correct, where $\alpha - (1 - \alpha) = \gamma \Rightarrow \alpha = 1/2 + \gamma/2$. For a margin $0 \leq \gamma \leq 1$, define $\mathcal{L}_{\mathcal{D}}^{\gamma}(f) = \Pr_{\mathbf{x} \sim \mathcal{D}}[f(\mathbf{x})c(\mathbf{x}) \leq \gamma]$. That is, $\mathcal{L}_{\mathcal{D}}^{\gamma}(f)$ is the probability over a random sample \mathbf{x} from \mathcal{D} that f has margin at most γ on \mathbf{x} . We have $\mathcal{L}_{\mathcal{D}}(f) = \mathcal{L}_{\mathcal{D}}^0(f)$. With margins defined, we show that for every training set $S = \{(x_i, c(x_i))\}_{i=1}^m$, if $t = \Omega(\ln(m/\delta))$, then with probability $1 - \delta$ over $\mathbf{B} \sim [m]^{n \times t}$, we have

$$\mathcal{L}_{\mathcal{D}}(f_{S, \mathbf{B}}) \leq \mathcal{L}_{\mathcal{D}}^{1/3}(g_S) + 1/m. \quad (1)$$

What this gives us, is that it suffices to understand how often the voting classifier that averages over all possible bootstrap samples has margin at most $1/3$. To see why (1) is true, notice that every hypothesis $h_{S(\mathbf{I}_i)}$ in $f_{S, \mathbf{B}} = (1/t) \sum_{i=1}^t h_{S(\mathbf{I}_i)}$ is uniform random among the hypotheses averaged by g_S . Hence for any $x \in \mathcal{X}$ where g_S has margin more than $1/3$, we have $\mathbb{E}_{\mathbf{I}_i \sim [m]^n}[h_{S(\mathbf{I}_i)}(x)c(x)] > 1/3$. A Chernoff bound and independence of the bootstrap samples implies that $\Pr_{\mathbf{B} \sim [m]^{n \times t}}[f_{S, \mathbf{B}}(x)c(x) \leq 0] \leq \exp(-\Omega(t)) \leq \delta/m$. Using that this holds for every x with $g_S(x)c(x) > 1/3$ establishes (1).

Our next step is to show that $\mathcal{L}_{\mathcal{D}}^{1/3}(g_S)$ is small with high probability over $\mathbf{S} \sim \mathcal{D}_c^m$. For this, our key idea is to create groups of bootstrap samples $\mathbf{S}(I)$ with $I \in [m]^n$. These groups have a structure similar to those produced by Hanneke's *Sub-Sample* procedure. We remark that these groups are only for the sake of analysis and are not part of the bagging algorithm.

For simplicity, let us for now assume that bagging produced samples without replacement instead of with replacement. To indicate this, we slightly abuse notation and let $\binom{m}{n}$ denote all vectors $I \in [m]^n$ where all entries are distinct. Also, let us assume that n precisely equals the number of samples in each sub-sample created by Hanneke's *Sub-Sample* (technically, this is $n = m - \sum_{i=0}^{\lg_4 m - 1} 4^i$). For a set $S = (x_1, c(x_1)), \dots, (x_m, c(x_m))$, let \mathcal{I} denote the collection of all vectors $I \in \binom{m}{n}$ such that $S(I)$ is one of the sub-samples produced by *Sub-Sample*(S, \emptyset). Note that \mathcal{I} only depends on m , not on S itself. We now define *buckets* \mathcal{C}_i of vectors $I \in \binom{m}{n}$ (i.e. of vectors corresponding to bootstrap samples). For every permutation π of the indices $1, \dots, m$, we create a bucket \mathcal{C}_{π} . We add a vector $I = (i_1, \dots, i_n) \in \binom{m}{n}$ to \mathcal{C}_{π} if and only if $\pi(I) = (\pi(i_1), \dots, \pi(i_n))$ is in \mathcal{I} .

With these buckets defined, we now make several crucial observations. First, for any bucket \mathcal{C}_{π} , if $\mathbf{S} \sim \mathcal{D}_c^m$, then the joint distribution of the bootstrap samples $\mathbf{S}(I)$ with $I \in \mathcal{C}_{\pi}$ is precisely the same as the joint distribution of the sub-samples produced by *Sub-Sample*(\mathbf{S}, \emptyset). This holds since permuting the samples in \mathbf{S} does not change their distribution. Hence for any bucket, Hanneke's analysis shows that the majority of hypotheses $h_{\mathbf{S}(I)}$ with $I \in \mathcal{C}_{\pi}$ rarely errs. More precisely, if we let $f_{\mathbf{S}, \pi} = (1/|\mathcal{C}_{\pi}|) \sum_{I \in \mathcal{C}_{\pi}} h_{\mathbf{S}(I)}$ then $\mathcal{L}_{\mathcal{D}}(f_{\mathbf{S}, \pi}) = O((d + \ln(1/\delta))/m)$ with probability $1 - \delta$ over \mathbf{S} . Here we need something slightly stronger, namely that $\mathcal{L}_{\mathcal{D}}^{5/6}(f_{\mathbf{S}, \pi}) = O((d + \ln(1/\delta))/m)$. Assume for now that this holds.

Next observe that if $x \in \mathcal{X}$ has $g_S(x)c(x) \leq 1/3$ for a training set S , then at least one third of the hypotheses $h_{S(I)}$ with $I \in \binom{m}{n}$ err on x , i.e. $h_{S(\mathbf{I})}$ for \mathbf{I} uniform in $\binom{m}{n}$ errs on x with probability at least $1/3$ (here we assume that g_S averages over $h_{S(I)}$ with $I \in \binom{m}{n}$, i.e. sampling without replacement instead of with). Symmetry now implies that every $I \in \binom{m}{n}$ is included in equally many buckets and all buckets contain equally many vectors I . This observation implies that the uniform $\mathbf{I} \sim \binom{m}{n}$ has the same distribution as if we first sample a uniform random bucket \mathbf{C} and then sample a uniform random \mathbf{I} from \mathbf{C} . But then

if $h_{S(\mathbf{I})}$ errs on x with probability at least $1/3$, it must be the case that a constant fraction of the buckets \mathcal{C}_π have $f_{S,\pi}(x)c(x) \leq 5/6$. This intuitively gives us that with high probability over \mathbf{S} , $\mathcal{L}_D^{1/3}(g_S)$ can only be a constant factor larger than $\mathcal{L}_D^{5/6}(f_{S,\pi})$. But $\mathcal{L}_D^{5/6}(f_{S,\pi})$ is $O((d + \ln(1/\delta))/m)$ with high probability, establishing the same thing for $\mathcal{L}_D^{1/3}(g_S)$ as desired.

The above are the main ideas in the proof, though making the steps completely formal requires some care. Also, let us briefly comment on some of the assumptions made above. First, we assumed that Hanneke's proof could establish the stronger claim $\mathcal{L}_D^{5/6}(f_{S,\pi}) = O((d + \ln(1/\delta))/m)$, not just that $\mathcal{L}_D(f_{S,\pi}) = O((d + \ln(1/\delta))/m)$. To formally carry out this argument, we have to change *Sub-Sample* to partition U into 20 subsets instead of 4 and make 19 recursive calls instead of 3. Next, we assumed above that \mathbf{I} was sampled without replacement. This was used to establish symmetry of the buckets. To handle sampling with replacement, let $D(I)$ denote the set of distinct values appearing in the entries of a vector $I \in [m]^n$. We now create a family of buckets similarly to above for every cardinality of $|D(I)|$. This gives symmetry among buckets corresponding to the same cardinality. We then show that these families of buckets may be combined by allowing a non-uniform sampling over the families. Finally, this step also requires us to handle bootstrap samples corresponding to I with very small $|D(I)|$, e.g. as small as just 1. Since such I are extremely unlikely, this can be charged to the failure probability δ .

In the next section, we proceed to give the formal proof of Theorem 1.

3 Bagging is an optimal PAC learner

We are ready to give the formal proof of Theorem 1. We have restated the theorem here in the notation introduced in Section 2

Theorem 2. *There is a universal constant $a > 0$ such that for every $0 < \delta < 1$, every distribution \mathcal{D} over an input domain \mathcal{X} , every concept class $\mathbb{C} \subseteq \mathcal{X} \rightarrow \{-1, 1\}$ of VC-dimension d and every $c \in \mathbb{C}$, if $t \geq 18 \ln(2m/\delta)$ and $0.02m \leq n \leq m$, then it holds with probability at least $1 - \delta$ over the random choice of a training set $\mathbf{S} \sim \mathcal{D}_c^m$ and t bootstrap samples $\mathbf{B} \sim [m]^{n \times t}$ that*

$$\mathcal{L}_D(f_{S,\mathbf{B}}) \leq a \cdot \frac{d + \ln(1/\delta)}{m}.$$

From hereon, we let \mathcal{D} be an arbitrary distribution over an input domain \mathcal{X} , let $0 < \delta < 1$ and let \mathbb{C} be an arbitrary concept class of some fixed VC-dimension $d \geq 1$ and let $c \in \mathbb{C}$ be arbitrary. For a training set S and a bootstrap sample $S(I)$ with $I \in [m]^n$, we let $h_{S(I)} \in \mathbb{C}$ be the hypothesis returned by the ERM algorithm used for bagging.

To prove Theorem 2, and hence Theorem 1, we start by relating the performance of $f_{S,\mathbf{B}}$ with $\mathbf{B} \sim [m]^{n \times t}$ to another voting classifier g_S that is a deterministic function of S . Concretely, for any training set S of m samples, let g_S denote the voting classifier:

$$g_S(x) = \frac{1}{m^n} \sum_{I \in [m]^n} h_{S(I)}(x).$$

That is, g_S is the voting classifier averaging the predictions over all m^n possible bootstrap samples. Recalling the definition of margins from Section 2, we show the following relation between $\mathcal{L}_D(f_{S,\mathbf{B}})$ and $\mathcal{L}_D^{1/3}(g_S)$:

Lemma 1. *For every $0 < \delta < 1$ and training set S of m samples, if $t \geq 18 \ln(m/\delta)$, then it holds with probability at least $1 - \delta$ over $\mathbf{B} \sim [m]^{n \times t}$ that*

$$\mathcal{L}_D(f_{S,\mathbf{B}}) \leq \mathcal{L}_D^{1/3}(g_S) + 1/m.$$

Proof. Consider a fixed training set S of m samples, a fixed set of bootstrap samples $B \in [m]^{n \times t}$ and an $\mathbf{x} \sim \mathcal{D}$. Let E denote the event that $g_S(\mathbf{x})c(\mathbf{x}) \leq 1/3$ and \bar{E} the complement event. Then

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}(f_{S,B}) &= \Pr_{\mathbf{x} \sim \mathcal{D}}[E] \Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|E)}[f_{S,B}(\mathbf{x})c(\mathbf{x}) \leq 0] + \Pr_{\mathbf{x} \sim \mathcal{D}}[\bar{E}] \Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|\bar{E})}[f_{S,B}(\mathbf{x})c(\mathbf{x}) \leq 0] \\ &\leq \Pr_{\mathbf{x} \sim \mathcal{D}}[E] + \Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|\bar{E})}[f_{S,B}(\mathbf{x})c(\mathbf{x}) \leq 0] \\ &= \mathcal{L}_{\mathcal{D}}^{1/3}(g_S) + \Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|\bar{E})}[f_{S,B}(\mathbf{x})c(\mathbf{x}) \leq 0]. \end{aligned}$$

By linearity of expectation, we have

$$\mathbb{E}_{\mathbf{B} \sim [m]^{n \times t}} \left[\Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|\bar{E})}[f_{S,\mathbf{B}}(\mathbf{x})c(\mathbf{x}) \leq 0] \right] = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}(\cdot|\bar{E})} \left[\Pr_{\mathbf{B} \sim [m]^{n \times t}}[f_{S,\mathbf{B}}(\mathbf{x})c(\mathbf{x}) \leq 0] \right].$$

Now observe that every \mathbf{I}_i in $\mathbf{B} = (\mathbf{I}_1, \dots, \mathbf{I}_t)$ is a uniform random sample from $[m]^n$, thus for x satisfying $g_S(x)c(x) > 1/3$, each \mathbf{I}_i has $\mathbb{E}_{\mathbf{I}_i \sim [m]^n}[h_{S(\mathbf{I}_i)}(x)c(x)] > 1/3$. By independence, and using $h_{S(\mathbf{I}_i)}(x) \in \{-1, 1\}$, it follows from Hoeffding's inequality that for every x with $g_S(x)c(x) > 1/3$, we have $\Pr_{\mathbf{B} \sim [m]^{n \times t}}[f_{S,\mathbf{B}}(x)c(x) \leq 0] \leq \exp(-(1/3)^2 t/2) = \exp(-t/18)$. Thus for $t \geq 18 \ln(m/\delta)$, we have $\mathbb{E}_{\mathbf{B} \sim [m]^{n \times t}}[\Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|\bar{E})}[f_{S,\mathbf{B}}(\mathbf{x})c(\mathbf{x}) \leq 0]] \leq \delta/m$. Markov's inequality implies that with probability at least $1 - \delta$ over the choice of \mathbf{B} , we have $\Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|\bar{E})}[f_{S,\mathbf{B}}(\mathbf{x})c(\mathbf{x}) \leq 0] \leq 1/m$ and thus $\mathcal{L}_{\mathcal{D}}(f_{S,\mathbf{B}}) \leq \mathcal{L}_{\mathcal{D}}^{1/3}(g_S) + 1/m$. \square

In light of Lemma 1 it thus suffices to understand $\mathcal{L}_{\mathcal{D}}^{1/3}(g_S)$ for the deterministic g_S , i.e. we need to understand how often a third of the hypotheses $h_{S(I)}$ with $I \in [m]^n$ mispredict the label of a sample $\mathbf{x} \sim \mathcal{D}$. For this, we prove the following:

Lemma 2. *There is a universal constant $a > 0$, such that for every $0 < \delta < 1$, it holds with probability at least $1 - \delta$ over the choice of a training set $\mathbf{S} \sim \mathcal{D}_c^m$ that*

$$\mathcal{L}_{\mathcal{D}}^{1/3}(g_S) \leq a \cdot \frac{d + \ln(1/\delta)}{m}$$

Before we prove Lemma 2, let us see that it suffices to finish the proof of Theorem 2:

Proof of Theorem 2. Assume $t \geq 18 \ln(2m/\delta)$. Let $\mathbf{S} \sim \mathcal{D}_c^m$ be a random training set of m samples and $\mathbf{B} \sim [m]^{n \times t}$. Define the event E_f that occurs if $\mathcal{L}_{\mathcal{D}}(f_{S,\mathbf{B}}) > \mathcal{L}_{\mathcal{D}}^{1/3}(g_S) + 1/m$ and define E_g as the event that $\mathcal{L}_{\mathcal{D}}^{1/3}(g_S) > a(d + \ln(2/\delta))/m$, where $a > 0$ is the constant from Lemma 2. By Lemma 1 and the constraint on t , we have $\Pr[E_f] \leq \delta/2$. By Lemma 2, we have $\Pr[E_g] \leq \delta/2$. By a union bound, it holds with probability at least $1 - \delta$ that none of the events E_f and E_g occur. In this case, we have $\mathcal{L}_{\mathcal{D}}(f_{S,\mathbf{B}}) \leq \mathcal{L}_{\mathcal{D}}^{1/3}(g_S) + 1/m \leq a \cdot \frac{d + \ln(2/\delta)}{m} + 1/m \leq (a + 2) \frac{d + \ln(1/\delta)}{m}$ (using $\ln(2) \leq 1 \leq d$). We have thus proved Theorem 2 with the constant $a + 2$. \square

3.1 Inducing Structure

What remains is thus to prove Lemma 2. To bound $\mathcal{L}_{\mathcal{D}}^{1/3}(g_S)$, we show that we can partition the hypotheses in $\{h_{S(I)} : I \in [m]^n\}$ into structured groups that are easier to analyse. Concretely, let \mathcal{F} be a collection of buckets $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_N \subseteq [m]^n$. We think of each $I \in \mathcal{C}_i$ as specifying the indices of a bootstrap sample. A bucket is thus a collection of vectors of indices.

The buckets \mathcal{C}_i constituting \mathcal{F} need not be disjoint. For a training set S and a bucket \mathcal{C} , we let $g_{S(\mathcal{C})}$ denote the voting classifier

$$g_{S(\mathcal{C})}(x) = \frac{1}{|\mathcal{C}|} \sum_{I \in \mathcal{C}} h_{S(I)}(x).$$

Together with the family \mathcal{F} , we have a probability distribution \mathcal{P} over the buckets $\mathcal{C}_0, \dots, \mathcal{C}_N$. Using $\mathcal{P}(i)$ to denote $\Pr_{\mathbf{C} \sim \mathcal{P}}[\mathbf{C} = \mathcal{C}_i]$, we thus have $\sum_{i=0}^N \mathcal{P}(i) = 1$ and $\mathcal{P}(i) \geq 0$ for all i . We call the pair $(\mathcal{F}, \mathcal{P})$ ε -representative if the following holds:

1. If we first draw a bucket \mathbf{C} according to the distribution \mathcal{P} and then draw a uniform $\mathbf{I} \in \mathbf{C}$, then \mathbf{I} is uniform random in $[m]^n$.
2. The probability $\mathcal{P}(0)$ is at most ε .

The following lemma shows that if we have an ε -representative pair $(\mathcal{F}, \mathcal{P})$, then it suffices to show that each of the voting classifiers $g_{\mathbf{S}(c_i)}$ with $i \neq 0$ perform well on a random training set $\mathbf{S} \sim \mathcal{D}_c^m$

Lemma 3. *Let $(\mathcal{F}, \mathcal{P})$ be $1/6$ -representative. Then for any $\tau \geq 0$, we have*

$$\Pr_{\mathbf{S} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{1/3}(g_S) > 24 \cdot \tau] \leq 23 \cdot \max_{i \neq 0} \Pr_{\mathbf{S} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{5/6}(g_{\mathbf{S}(c_i)}) > \tau].$$

Proof of Lemma 3. Consider an arbitrary training set S of m samples and an $x \in \mathcal{X}$ for which $g_S(x)c(x) \leq 1/3$. If we let \mathbf{I} be uniform random in $[m]^n$, then $\Pr_{\mathbf{I} \sim [m]^n}[h_{\mathbf{S}(\mathbf{I})}(x) \neq c(x)] \geq 1/3$. At the same time, since $(\mathcal{F}, \mathcal{P})$ is $1/6$ -representative, \mathbf{I} has the same distribution as if we first draw a bucket \mathbf{C} according to \mathcal{P} , then sample a uniform \mathbf{I} in \mathbf{C} . Using $\mathbf{I} \sim \mathbf{C}$ to denote that \mathbf{I} is drawn uniformly from \mathbf{C} , we thus have $\mathbb{E}_{\mathbf{C} \sim \mathcal{P}}[\Pr_{\mathbf{I} \sim \mathbf{C}}[h_{\mathbf{S}(\mathbf{I})}(x) \neq c(x)]] \geq 1/3$. On the other hand,

$$\begin{aligned} \mathbb{E}_{\mathbf{C} \sim \mathcal{P}}[\Pr_{\mathbf{I} \sim \mathbf{C}}[h_{\mathbf{S}(\mathbf{I})}(x) \neq c(x)]] &\leq \\ \Pr_{\mathbf{C} \sim \mathcal{P}}[\mathbf{C} = \mathcal{C}_0] + \Pr_{\mathbf{C} \sim \mathcal{P}}[g_{\mathbf{S}(\mathbf{C})}(x)c(x) \leq 5/6 \wedge \mathbf{C} \neq \mathcal{C}_0] + \mathbb{E}_{\mathbf{C} \sim \mathcal{P}}[\Pr_{\mathbf{I} \sim \mathbf{C}}[h_{\mathbf{S}(\mathbf{I})}(x) \neq c(x)] \mid g_{\mathbf{S}(\mathbf{C})}(x)c(x) > 5/6] &\leq \\ \mathcal{P}(0) + \Pr_{\mathbf{C} \sim \mathcal{P}}[g_{\mathbf{S}(\mathbf{C})}(x)c(x) \leq 5/6 \wedge \mathbf{C} \neq \mathcal{C}_0] + (1/2 - 5/12) &\leq \\ 1/6 + \Pr_{\mathbf{C} \sim \mathcal{P}}[g_{\mathbf{S}(\mathbf{C})}(x)c(x) \leq 5/6 \wedge \mathbf{C} \neq \mathcal{C}_0] + 1/12. \end{aligned}$$

We thus have $\Pr_{\mathbf{C} \sim \mathcal{P}}[g_{\mathbf{S}(\mathbf{C})}(x)c(x) \leq 5/6 \wedge \mathbf{C} \neq \mathcal{C}_0] \geq 1/3 - 1/6 - 1/12 = 1/12$ for such S, x .

For a fixed S and an $\mathbf{x} \sim \mathcal{D}$, let E denote the event that $g_S(\mathbf{x})c(\mathbf{x}) \leq 1/3$. Then the above implies that

$$\begin{aligned} 1/12 &\leq \mathbb{E}_{\mathbf{x} \sim \mathcal{D}(\cdot|E)}[\Pr_{\mathbf{C} \sim \mathcal{P}}[g_{\mathbf{S}(\mathbf{C})}(\mathbf{x})c(\mathbf{x}) \leq 5/6 \wedge \mathbf{C} \neq \mathcal{C}_0]] \\ &= \mathbb{E}_{\mathbf{C} \sim \mathcal{P}}[\Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|E)}[g_{\mathbf{S}(\mathbf{C})}(\mathbf{x})c(\mathbf{x}) \leq 5/6 \wedge \mathbf{C} \neq \mathcal{C}_0]] \end{aligned}$$

For short, let $Z(\mathbf{C})$ take the value $\Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|E)}[g_{\mathbf{S}(\mathbf{C})}(\mathbf{x})c(\mathbf{x}) \leq 5/6 \wedge \mathbf{C} \neq \mathcal{C}_0]$ (which is trivially 0 for $\mathbf{C} = \mathcal{C}_0$). Then the above is $\mathbb{E}_{\mathbf{C} \sim \mathcal{P}}[Z(\mathbf{C})] \geq 1/12$. Moreover, $1 - Z(\mathbf{C})$ is non-negative and $\mathbb{E}_{\mathbf{C} \sim \mathcal{P}}[1 - Z(\mathbf{C})] \leq 11/12$. Hence by Markov's inequality, we have $\Pr_{\mathbf{C} \sim \mathcal{P}}[Z(\mathbf{C}) \leq 1/24] = \Pr_{\mathbf{C} \sim \mathcal{P}}[1 - Z(\mathbf{C}) \geq 23/24] \leq (11/12)/(23/24) = 22/23$. Thus $\Pr_{\mathbf{C} \sim \mathcal{P}}[Z(\mathbf{C}) \geq 1/24] \geq 1/23$. Now observe that if \mathcal{C} satisfies $Z(\mathcal{C}) \geq 1/24$, then $\mathcal{C} \neq \mathcal{C}_0$ and

$$\begin{aligned} \mathcal{L}_D^{5/6}(g_{\mathcal{C}}) &= \\ \Pr_{\mathbf{x} \sim \mathcal{D}}[g_{\mathcal{C}}(\mathbf{x})c(\mathbf{x}) \leq 5/6] &\geq \\ \Pr_{\mathbf{x} \sim \mathcal{D}}[E] \cdot \Pr_{\mathbf{x} \sim \mathcal{D}(\cdot|E)}[g_{\mathcal{C}}(\mathbf{x})c(\mathbf{x}) \leq 5/6] &\geq \\ \Pr_{\mathbf{x} \sim \mathcal{D}}[E]/24 &= \\ \mathcal{L}_D^{1/3}(g_S)/24. \end{aligned}$$

We therefore conclude that

$$\begin{aligned} \Pr_{\mathbf{C} \sim \mathcal{P}}[\mathcal{L}_D^{5/6}(g_{\mathbf{S}(\mathbf{C})}) \geq \mathcal{L}_D^{1/3}(g_S)/24 \wedge \mathbf{C} \neq \mathcal{C}_0] &\geq \\ \Pr_{\mathbf{C} \sim \mathcal{P}}[Z(\mathbf{C}) \geq 1/24] &\geq \\ 1/23. \end{aligned}$$

This implies that for any $\tau \geq 0$ and any S with $\mathcal{L}_D^{1/3}(g_S) > 24 \cdot \tau$, we also have $\Pr_{\mathbf{C} \sim \mathcal{P}}[\mathcal{L}_D^{5/6}(g_{S(\mathbf{C})}) > \tau \wedge \mathbf{C} \neq \mathcal{C}_0] \geq 1/23$. Therefore, by Markov's inequality

$$\begin{aligned} \Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{1/3}(g_{\mathbf{s}}) > 24 \cdot \tau] &\leq \\ \Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\Pr_{\mathbf{C} \sim \mathcal{P}}[\mathcal{L}_D^{5/6}(g_{S(\mathbf{C})}) > \tau \wedge \mathbf{C} \neq \mathcal{C}_0] \geq 1/23] &\leq \\ 23 \cdot \mathbb{E}_{\mathbf{s} \sim \mathcal{D}_c^m}[\Pr_{\mathbf{C} \sim \mathcal{P}}[\mathcal{L}_D^{5/6}(g_{S(\mathbf{C})}) > \tau \wedge \mathbf{C} \neq \mathcal{C}_0]] &= \\ 23 \cdot \mathbb{E}_{\mathbf{C} \sim \mathcal{P}}[\Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{5/6}(g_{S(\mathbf{C})}) > \tau \wedge \mathbf{C} \neq \mathcal{C}_0]] &\leq \\ 23 \cdot \max_{i \neq 0} \Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{5/6}(g_{S(c_i)}) > \tau]. & \end{aligned}$$

□

We next show that it is possible to construct a representative pair in which all buckets except \mathcal{C}_0 behave nicely

Lemma 4. *There is a universal constant $a > 0$, such that for any m and n with $0.02m \leq n \leq m$, there is a 1/6-representative pair $(\mathcal{F} = \mathcal{C}_0, \dots, \mathcal{C}_N, \mathcal{P})$ satisfying that for every $0 < \delta < 1$, it holds that*

$$\max_{i \neq 0} \Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{5/6}(g_{S(c_i)}) > a(d + \ln(1/\delta))/m] < \delta.$$

Before proving Lemma 4, let us see that it suffices to complete the proof of Lemma 2.

Proof of Lemma 2. Let $(\mathcal{F}, \mathcal{P})$ be as in Lemma 4 and let a be the constant in Lemma 4. For any $0 < \delta < 1$, Lemma 3 gives us that

$$\begin{aligned} \Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{1/3}(g_{\mathbf{s}}) > 76a(d + \ln(1/\delta))/m] &\leq \Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{1/3}(g_{\mathbf{s}}) > 24a(d + \ln(23/\delta))/m] \\ &\leq 23 \cdot \max_{i \neq 0} \Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\mathcal{L}_D^{5/6}(g_{S(c_i)}) > a(d + \ln(23/\delta))/m] \\ &\leq 23 \cdot (\delta/23) \\ &= \delta. \end{aligned}$$

Thus we have proved Lemma 2 with the constant $76a$. □

In the next section, we define the pair $(\mathcal{F}, \mathcal{P})$ having the properties claimed in Lemma 4.

3.2 Representative Pair

We now define a 1/6-representative pair $(\mathcal{F}, \mathcal{P})$ with $\mathcal{F} = \mathcal{C}_0, \dots, \mathcal{C}_N$ for some N , such that $(\mathcal{F}, \mathcal{P})$ has the properties claimed in Lemma 4.

We construct a number of buckets for each possible number of distinct values among the entries of a vector $I \in [m]^n$. Recall that we assume $0.02m \leq n \leq m$.

Bucket \mathcal{C}_0 . We use the special bucket \mathcal{C}_0 to handle all vectors $I \in [m]^n$ with fewer than $0.01m$ distinct values among their entries as well as those I with more than $0.9m$ distinct values. That is, if we let $D(I) \subseteq [m]$ denote the set of distinct values occurring in I , then $\mathcal{C}_0 = \{I \in [m]^n : |D(I)| \notin [0.01m, 0.9m]\}$.

Remaining Buckets. For every integer $d \in \{[0.01m], \dots, [0.9m]\}$, we construct a number of buckets $\mathcal{C}_{d,1}, \dots, \mathcal{C}_{d,N_d}$. The final set of buckets \mathcal{F} is simply $\{\mathcal{C}_0\} \cup \left(\bigcup_d \bigcup_{i \in [N_d]} \{\mathcal{C}_{d,i}\}\right)$.

Let $\ell = 20^j$ be the largest power of 20 such that $\ell \leq 0.01m$. For every d , we construct a bucket for every list L of ℓ distinct indices in $[m]$. We thus have $N_d = m(m-1) \cdots (m-\ell+1) = m!/(m-\ell)!$. We only add vectors $I \in [m]^n$ with $|D(I)| = d$ to these buckets. Given a list $L = i_1, \dots, i_\ell$ of distinct indices, invoking Algorithm 2 as $BuildBucket_d(L, \emptyset, \emptyset)$ constructs the bucket corresponding to L .

Algorithm 2: $BuildBucket_d(L, U, V)$

Input: List L of $\ell = 20^j$ distinct indices i_1, \dots, i_ℓ for an integer $j \geq 0$
Set of indices $U \subseteq [m]$ that must be contained in I
Set of indices $V \subseteq [m]$ that must be disjoint from I
Result: Bucket \mathcal{C} containing vectors $I \in [m]^n$

- 1 **if** $j = 0$ **then**
- 2 $U \leftarrow U \cup \{i_1\}$
- 3 **return** $\{I \in [m]^n : |D(I)| = d \wedge U \subseteq D(I) \wedge V \cap D(I) = \emptyset\}$
- 4 **else**
- 5 Partition L into 20 consecutive groups L_0, \dots, L_{19} of 20^{j-1} indices each.
- 6 **return** $\bigcup_{i=1}^{19} BuildBucket_d(L_0, U \cup (\bigcup_{j \in \{1, \dots, 19\} \setminus \{i\}} L_j), V \cup L_i)$

Examining the algorithm, we observe that the recursive procedure computes disjoint sets $U \subset L$ and $V \subset L$. Whenever the recursion bottoms out and we return a subset of $[m]^n$ in step 3 of Algorithm 2, we have $U \cup V = L$ and we output all $I \in [m]^n$ with $|D(I)| = d$, $U \subseteq D(I)$ and $V \cap D(I) = \emptyset$. Since $|L| = \ell \leq 0.01m$ and $d \in \{[0.01m], \dots, [0.9m]\}$, it follows that the output is always non-empty as $|[m] \setminus V| \geq 0.99m \geq d$ and $d \geq 0.01m \geq |U|$.

Probability Distribution. We next argue that we can define a probability distribution \mathcal{P} over the buckets $\mathcal{C}_0, \dots, \mathcal{C}_N$ defined above, such that if we draw a bucket \mathbf{C} according to \mathcal{P} and then a uniform \mathbf{I} from \mathbf{C} , then the distribution of \mathbf{I} is uniform in $[m]^n$ and $\mathcal{P}(0) \leq 1/6$.

Consider first the set of buckets $\mathcal{C}_{d,1}, \dots, \mathcal{C}_{d,N_d}$ added for a particular value of $d \in \{[0.01m], \dots, [0.9m]\}$. These buckets only contain vectors $I \in [m]^n$ with $|D(I)| = d$. By symmetry, since we have a bucket for every list L of distinct indices i_1, \dots, i_ℓ , we have that every I with $|D(I)| = d$ is contained in equally many buckets $\mathcal{C}_{d,i}$, and every bucket contains equally many vectors I . Thus if we draw a uniform random bucket \mathbf{C} among $\mathcal{C}_{d,1}, \dots, \mathcal{C}_{d,N_d}$ and then a uniform random \mathbf{I} in \mathbf{C} , then the distribution of \mathbf{I} is uniform among all I with $|D(I)| = d$.

For every such bucket $\mathcal{C}_{d,i}$, we now define the corresponding probability under \mathcal{P} to be $\Pr_{\mathbf{I} \sim [m]^n}[|D(\mathbf{I})| = d]N_d^{-1}$. Similarly, we define $\mathcal{P}(0)$ to equal $\Pr_{\mathbf{I} \sim [m]^n}[|D(\mathbf{I})| \notin [0.01m, 0.9m]]$. With these probabilities, we have that if we first draw a bucket \mathbf{C} according to \mathcal{P} and then draw a uniform random \mathbf{I} from \mathbf{C} , then the distribution of \mathbf{I} is uniform in $[m]^n$ as required.

Lemma 5. *For the above $(\mathcal{F}, \mathcal{P})$, if $m \geq 100$, then it holds that $\mathcal{P}(0) \leq 1/6$.*

The proof is a standard counting argument for sampling with replacement and can be found in the appendix.

In the following section, we prove that the pair $(\mathcal{F}, \mathcal{P})$ constructed above satisfies the properties in Lemma 4 (if $m < 100$, we do not need to prove anything as $\Pr_{\mathbf{s} \sim \mathcal{D}_c^m}[\mathcal{L}_{\mathcal{D}}^{5/6}(g_{\mathbf{s}(c_i)}) > a(d + \ln(1/\delta))/100] \leq 0$ provided $a \geq 100$).

3.3 Performance in a Bucket

What remains is to understand the behaviour of a voting classifier constructed from the indices in a bucket, i.e. we complete the proof of Lemma 4. As mention above, we may assume $m \geq 100$ as otherwise the claim

in Lemma 4 is trivially true for $a \geq 100$. The proof follows that of Hanneke rather closely.

For the $1/6$ -representative pair $(\mathcal{F} = \mathcal{C}_0, \dots, \mathcal{C}_N, \mathcal{P})$ defined in the previous section, we consider the buckets \mathcal{C} with $\mathcal{C} \neq \mathcal{C}_0$. By construction, any such bucket was obtained from a choice of $d = |D(I)|$ and a list $L = i_1, \dots, i_\ell$ of $\ell = 20^j$ distinct indices. So consider a bucket \mathcal{C} corresponding to some $d \in \{[0.01m], \dots, [0.9m]\}$ and a list L . We show that for any distribution \mathcal{D} , with high probability over the choice of $\mathbf{S} \sim \mathcal{D}_c^m$, we have $\mathcal{L}_\mathcal{D}^{5/6}(g_{\mathbf{S}(\mathcal{C})}) \leq a(d + \ln(1/\delta))/m$.

Recall that from L and d , we construct \mathcal{C} by invoking Algorithm 2 as $\text{BuildBucket}_d(L, \emptyset, \emptyset)$. Now consider the recursion tree \mathcal{T} corresponding to the invocation of $\text{BuildBucket}_d(L, \emptyset, \emptyset)$. The root corresponds to the initial call $\text{BuildBucket}_d(L, \emptyset, \emptyset)$. Each internal node has 19 children w_1, \dots, w_{19} , one for each of the 19 recursive calls $\text{BuildBucket}_d(L_0, U \cup (\bigcup_{j \in \{1, \dots, 19\} \setminus \{i\}} L_j), V \cup L_i)$. We associate with each node $v \in \mathcal{T}$ the sets L_v, U_v and V_v corresponding to the arguments given as parameters to the corresponding call to BuildBucket . Similarly, for every node $v \in \mathcal{T}$, we define \mathcal{W}_v as the collection of all $I \in [m]^n$ added to the output \mathcal{C} during recursive calls starting at v . That is, if v is a leaf, then $\mathcal{W}_v = \{I \in [m]^n : D(I) = d \wedge (U_v \cup \{i_1\}) \subseteq D(I) \wedge V_v \cap D(I) = \emptyset\}$ and if v is internal with children w_1, \dots, w_{19} , then $\mathcal{W}_v = \bigcup_{i=1}^{19} \mathcal{W}_{w_i}$.

With these definitions, we consider the intermediate voting classifiers obtained from the sets \mathcal{W}_v :

$$g_{\mathbf{S}(\mathcal{W}_v)}(x) = \frac{1}{|\mathcal{W}_v|} \sum_{I \in \mathcal{W}_v} h_{S(I)}(x).$$

Importantly, if r is the root of \mathcal{T} , then $g_{\mathbf{S}(\mathcal{W}_r)} = g_{\mathbf{S}(\mathcal{C})}$.

For a node $v \in \mathcal{T}$, we say that its height is the distance from v to the leaves in its subtree. Hence a leaf has height 0 and the root has height j . We prove the following

Lemma 6. *There is a universal constant $a > 0$ such that for any distribution \mathcal{D} , any $0 < \delta < 1$, it holds for every node $v \in \mathcal{T}$ of height $\kappa \in \{0, \dots, j\}$ that*

$$\Pr_{\mathbf{S} \sim \mathcal{D}_c^m} [\mathcal{L}_\mathcal{D}^{5/6}(g_{\mathbf{S}(\mathcal{W}_v)}) > a(d + \ln(1/\delta))/20^\kappa] < \delta.$$

Using Lemma 6, we now prove Lemma 4:

Proof of Lemma 4. Let a be the constant in Lemma 6. Using that $g_{\mathbf{S}(\mathcal{W}_r)} = g_{\mathbf{S}(\mathcal{C})}$ when r is the root of \mathcal{T} shows that $\Pr_{\mathbf{S} \sim \mathcal{D}_c^m} [\mathcal{L}_\mathcal{D}^{5/6}(g_{\mathbf{S}(\mathcal{C})}) > a(d + \ln(1/\delta))/20^j] < \delta$. Using that 20^j was defined to be the largest power of 20 such that $20^j \leq 0.01m$, we have $20^j \geq m/2000$. Thus $\Pr_{\mathbf{S} \sim \mathcal{D}_c^m} [\mathcal{L}_\mathcal{D}^{5/6}(g_{\mathbf{S}(\mathcal{C})}) > 2000 \cdot a(d + \ln(1/\delta))/m] < \delta$. Since this holds for all buckets $\mathcal{C} \neq \mathcal{C}_0$, we have thus proved Lemma 4 with the constant $2000a$. \square

What remains is to prove Lemma 6. We prove lemma by induction in κ . For the base case $\kappa = 0$, the proof is trivial as we always have $\mathcal{L}_\mathcal{D}^{5/6}(g_{\mathbf{S}(\mathcal{W}_v)}) \leq 1 < a(d + \ln(1/\delta))$ provided $a \geq 1$.

For the inductive step, assume the lemma holds for all $\kappa' < \kappa$. Now consider a fixed value of $0 < \delta < 1$, a distribution \mathcal{D} and a fixed node $v \in \mathcal{T}$ of height κ with children w_1, \dots, w_{19} . To complete the inductive step, we define a number of undesirable properties of a training set S and show that a random \mathbf{S} rarely has any of these properties. The first property we define is the following

- For $i = 1, \dots, 19$, let \mathcal{E}_i denote the set of all S for which $\mathcal{L}_\mathcal{D}^{5/6}(g_{\mathbf{S}(\mathcal{W}_{w_i})}) > a(d + \ln(57/\delta))/20^{\kappa-1}$.

Here we trivially have

Lemma 7. *For any i , we have $\Pr_{\mathbf{S} \sim \mathcal{D}_c^m} [S \in \mathcal{E}_i] \leq \delta/57$.*

Proof. The lemma follows immediately from the induction hypothesis. \square

Next, for a child w_i and a training set S , let $E_{S,i} \subseteq \mathcal{X} \times \{-1, 1\}$ denote the subset of all $(x, c(x))$ with $x \in \mathcal{X}$ for which $g_{\mathbf{S}(\mathcal{W}_{w_i})}(x)c(x) \leq 5/6$. Also for the node v let L_0, \dots, L_{19} be the partitioning of L_v by BuildBucket . The second property is

- For $i = 1, \dots, 19$, let \mathcal{E}'_i denote the set of all training sets S for which $\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_{w_i})}) \geq (a/627)(d + \ln(1/\delta))/20^\kappa$ and $|S(L_i) \cap E_{S,i}| < 10^9(d + \ln(57/\delta))$.

Note in the above that L_i is a set of distinct indices and $S(L_i)$ is the subset of samples in S indexed by L_i . The property intuitively states that $g_{S(\mathcal{W}_{w_i})}$ often has a margin less than $5/6$, yet we see few such samples in $S(L_i)$.

For the third and last property, define $\mathcal{D}(\cdot | E_{S,i})$ for a set of samples S , as the distribution of an \mathbf{x} drawn from \mathcal{D} , conditioned on $(\mathbf{x}, c(\mathbf{x})) \in E_{S,i}$. The final property is

- For $i = 1, \dots, 19$, let \mathcal{E}''_i denote the set of all training sets S for which $|S(L_i) \cap E_{S,i}| \geq 10^9(d + \ln(57/\delta))$ and there is a hypothesis $h \in \mathbb{C}$ such that $h(x) = c(x)$ for all $(x, c(x)) \in S(L_i) \cap E_{S,i}$ and $\mathcal{L}_{\mathcal{D}(\cdot | E_{S,i})}(h) > 1/62700$.

We first argue that every training set that does not lie in any of the sets $\mathcal{E}_i, \mathcal{E}'_i, \mathcal{E}''_i$ behaves well

Lemma 8. *Let S satisfy $S \notin \bigcup_i \{\mathcal{E}_i \cup \mathcal{E}'_i \cup \mathcal{E}''_i\}$. Then*

$$\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_v)}) \leq a(d + \ln(1/\delta))/20^\kappa.$$

Combining Lemma 8 with the following completes the inductive proof of Lemma 6

Lemma 9. *Let $\mathcal{S} \sim \mathcal{D}_c^m$. Then if the constant $a > 0$ is large enough, it holds that $\Pr_{\mathcal{S} \sim \mathcal{D}_c^m}[\mathcal{S} \notin \bigcup_i \{\mathcal{E}_i \cup \mathcal{E}'_i \cup \mathcal{E}''_i\}] \geq 1 - \delta$.*

Thus what remains is to prove Lemma 8 and Lemma 9. We start with the proof of Lemma 8.

3.4 Well-Behaved S

Proof of Lemma 8. Let S be a training set of m samples satisfying $\forall i : S \notin (\mathcal{E}_i \cup \mathcal{E}'_i \cup \mathcal{E}''_i)$. Consider any child w_i of v and a vector $I \in \mathcal{W}_{w_j}$ for a child $w_j \neq w_i$. We focus on the hypothesis $h_{S(I)}$ produced by ERM on $S(I)$. Since we are in the realizable case, we know that $h_{S(I)}$ correctly predicts the label of every sample it is trained on, i.e. $h(x) = c(x)$ for all $(x, c(x)) \in S(I)$. Furthermore, by the recursive procedure *BuildBucket_a*, we see that all of L_i is added to U_{w_j} in the recursive call corresponding to the child w_j . Hence every sample in $S(L_i)$ is included in $S(I)$. We now aim to show that it is unlikely that both $h_{S(I)}$ mispredict the label of an x while $g_{S(\mathcal{W}_{w_i})}$ also has a margin of no more than $5/6$ on x

$$\Pr_{\mathbf{x} \sim \mathcal{D}}[g_{S(\mathcal{W}_{w_i})}(\mathbf{x})c(\mathbf{x}) \leq 5/6 \wedge h_{S(I)}(\mathbf{x}) \neq c(\mathbf{x})] \leq (a/627)(d + \ln(1/\delta))/20^\kappa. \quad (2)$$

We consider two cases. In the first case, we have $\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_{w_i})}) < (a/627)(d + \ln(1/\delta))/20^\kappa$. Here the conclusion (2) follows trivially as

$$\begin{aligned} \Pr_{\mathbf{x} \sim \mathcal{D}}[g_{S(\mathcal{W}_{w_i})}(\mathbf{x})c(\mathbf{x}) \leq 5/6 \wedge h(\mathbf{x}) \neq c(\mathbf{x})] &\leq \\ \Pr_{\mathbf{x} \sim \mathcal{D}}[g_{S(\mathcal{W}_{w_i})}(\mathbf{x})c(\mathbf{x}) \leq 5/6] &= \\ \mathcal{L}_D^{5/6}[g_{S(\mathcal{W}_{w_i})}] & \end{aligned}$$

In the second case, we have $\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_{w_i})}) \geq (a/627)(d + \ln(1/\delta))/20^\kappa$. Since $S \notin \mathcal{E}'_i$, this implies $|S(L_i) \cap E_{S,i}| \geq 10^9(d + \ln(57/\delta))$. Since $S \notin \mathcal{E}''_i$, this gives us that $\mathcal{L}_{\mathcal{D}(\cdot | E_{S,i})}(h_{S(I)}) \leq 1/62700$. Using that $S \notin \mathcal{E}_i$, we at the same time have $\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_{w_i})}) \leq a(d + \ln(57/\delta))/20^{\kappa-1}$. We therefore have

$$\begin{aligned} \Pr_{\mathbf{x} \sim \mathcal{D}}[g_{S(\mathcal{W}_{w_i})}(\mathbf{x})c(\mathbf{x}) \leq 5/6 \wedge h_{S(I)}(\mathbf{x}) \neq c(\mathbf{x})] &= \\ \mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_{w_i})}) \cdot \mathcal{L}_{\mathcal{D}(\cdot | E_{S,i})}(h_{S(I)}) &\leq \\ (a(d + \ln(57/\delta))/20^{\kappa-1}) \cdot (1/62700) &\leq \\ (100a(d + \ln(1/\delta))/20^\kappa) \cdot (1/62700) &\leq \\ (a/627)(d + \ln(1/\delta))/20^\kappa & \end{aligned}$$

We have thus established (2). Using this, we now show that $\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_v)}) \leq a(d + \ln(1/\delta))/20^\kappa$. Define E_S as the set of $(x, c(x))$ with $x \in \mathcal{X}$ for which $g_{S(\mathcal{W}_v)}(x)c(x) \leq 5/6$. For any such $(x, c(x)) \in E_S$, consider drawing two vectors $\mathbf{I}_1, \mathbf{I}_2$ uniformly and independently from \mathcal{W}_v . Since each \mathcal{W}_{w_i} contains equally many vectors I , we have that \mathbf{I}_1 and \mathbf{I}_2 have the same distribution as if we first pick a uniform random child $w_{\mathbf{i}_1}$ and then a uniform random \mathbf{I} from $\mathcal{W}_{w_{\mathbf{i}_1}}$. Let \mathbf{i}_1 denote the index of the child containing \mathbf{I}_1 and let \mathbf{i}_2 denote the index of the child containing \mathbf{I}_2 .

We first observe that since $(x, c(x)) \in E_S$, there must be a child index z for which $g_{S(\mathcal{W}_{w_z})}(x)c(x) \leq 5/6$. Fix an arbitrary such $z(x)$ for every $(x, c(x)) \in E_S$. Since $(x, c(x)) \in E_S$, we have that at least an α -fraction of the vectors I in \mathcal{W}_v have $h_{S(I)}(x) \neq c(x)$ where α satisfies $1 - 2\alpha \leq 5/6 \Rightarrow \alpha \geq 1/12$. Hence even outside $\mathcal{W}_{w_z(x)}$ there is still another $|\mathcal{W}_v|/12 - |\mathcal{W}_{w_z(x)}| = |\mathcal{W}_v|/12 - |\mathcal{W}_v|/19 = (7/228)|\mathcal{W}_v| > |\mathcal{W}_v|/33$ vectors $I \in \mathcal{W}_v \setminus \mathcal{W}_{w_z(x)}$ for which $h_{S(I)}(x) \neq c(x)$. We therefore have for $(x, c(x)) \in E_S$ that

$$\Pr_{\mathbf{I}_1, \mathbf{I}_2 \sim \mathcal{W}_v} [\mathbf{i}_1 = z(x) \wedge \mathbf{i}_2 \neq z(x) \wedge h_{S(\mathbf{I}_2)}(x) \neq c(x)] > (1/19) \cdot (1/33) = 1/627.$$

It follows that

$$\Pr_{\mathbf{I}_1, \mathbf{I}_2 \sim \mathcal{W}_v, \mathbf{x} \sim \mathcal{D}} [\mathbf{i}_1 = z(\mathbf{x}) \wedge \mathbf{i}_2 \neq z(\mathbf{x}) \wedge h_{S(\mathbf{I}_2)}(\mathbf{x}) \neq c(\mathbf{x})] > \Pr_{\mathbf{x} \sim \mathcal{D}} [(\mathbf{x}, c(\mathbf{x})) \in E_S] / 627.$$

By averaging, there must exist a fixed choice of I_1^* and I_2^* such that

$$\Pr_{\mathbf{x} \sim \mathcal{D}} [i_1^* = z(\mathbf{x}) \wedge i_2^* \neq z(\mathbf{x}) \wedge h_{S(I_2^*)}(\mathbf{x}) \neq c(\mathbf{x})] > \Pr_{\mathbf{x} \sim \mathcal{D}} [(\mathbf{x}, c(\mathbf{x})) \in E_S] / 627.$$

If $i_1^* = i_2^*$, then the left hand side is 0 and we conclude $\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_v)}) = \Pr_{\mathbf{x} \sim \mathcal{D}} [(\mathbf{x}, c(\mathbf{x})) \in E_S] = 0$. Otherwise, we have $i_1^* \neq i_2^*$ and

$$\Pr_{\mathbf{x} \sim \mathcal{D}} [i_1^* = z(\mathbf{x}) \wedge i_2^* \neq z(\mathbf{x}) \wedge h_{S(I_2^*)}(\mathbf{x}) \neq c(\mathbf{x})] \leq \Pr_{\mathbf{x} \sim \mathcal{D}} [i_1^* = z(\mathbf{x}) \wedge h_{S(I_2^*)}(\mathbf{x}) \neq c(\mathbf{x})].$$

Since I_2^* comes from a child $i_2^* \neq i_1^*$, it follows from (2) that

$$\Pr_{\mathbf{x} \sim \mathcal{D}} [i_1^* = z(\mathbf{x}) \wedge h_{S(I_2^*)}(\mathbf{x}) \neq c(\mathbf{x})] \leq (a/627)(d + \ln(1/\delta))/20^\kappa$$

This finally implies

$$\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_v)}) = \Pr_{\mathbf{x} \sim \mathcal{D}} [(\mathbf{x}, c(\mathbf{x})) \in E_S] < 627(a/627)(d + \ln(1/\delta))/20^\kappa = a(d + \ln(1/\delta))/20^\kappa.$$

□

3.5 Most S are Well-Behaved

In this section, we prove Lemma 9, i.e. that $\mathbf{S} \sim \mathcal{D}_c^m$ rarely lies in $\bigcup_i \{\mathcal{E}_i \cup \mathcal{E}'_i \cup \mathcal{E}''_i\}$. Lemma 7 already handles \mathcal{E}_i . For the remaining, we show the following two

Lemma 10. *If the constant $a > 0$ is large enough, then for any i , we have $\Pr_{\mathbf{S} \sim \mathcal{D}_c^m} [\mathbf{S} \in \mathcal{E}'_i] \leq \delta/57$.*

Lemma 11. *For any i , we have $\Pr_{\mathbf{S} \sim \mathcal{D}_c^m} [\mathbf{S} \in \mathcal{E}''_i] \leq \delta/57$.*

Before we prove these last lemmas, observe that Lemma 7, Lemma 10 and Lemma 11 together imply Lemma 9 by a union bound over all the $3 \cdot 19 = 57$ events $\mathbf{S} \in \mathcal{E}_i, \mathbf{S} \in \mathcal{E}'_i$ and $\mathbf{S} \in \mathcal{E}''_i$. We now give the two remaining proofs. For ease of reading, we first restate the definitions of \mathcal{E}'_i and \mathcal{E}''_i . Also, recall that we defined $E_{S,i}$ as the set of $(x, c(x))$ with $x \in \mathcal{X}$ for which $g_{S(\mathcal{W}_{w_i})}(x)c(x) \leq 5/6$.

- For $i = 1, \dots, 19$, let \mathcal{E}'_i denote the set of all training sets S for which $\mathcal{L}_D^{5/6}(g_{S(\mathcal{W}_{w_i})}) \geq (a/627)(d + \ln(1/\delta))/20^\kappa$ and $|S(L_i) \cap E_{S,i}| < 10^9(d + \ln(57/\delta))$.

- For $i = 1, \dots, 19$, let \mathcal{E}_i'' denote the set of all training sets S for which $|S(L_i) \cap E_{S,i}| \geq 10^9(d + \ln(57/\delta))$ and there is a hypothesis $h \in \mathbb{C}$ such that $h(x) = c(x)$ for all $(x, c(x)) \in S(L_i) \cap E_{S,i}$ and $\mathcal{L}_{\mathcal{D}(\cdot|E_{S,i})}(h) > 1/62700$.

Proof of Lemma 10. Observe that all $I \in \mathcal{W}_{w_i}$ are disjoint from the indices in L_i by definition of *Build-Bucket_d*. Hence for all $I \in \mathcal{W}_{w_i}$, we have that $\mathbf{S}(I)$ is completely determined from $\mathbf{S}([m] \setminus L_i)$. Let us use the notation $\mathbf{S}_1 = \mathbf{S}([m] \setminus L_i)$ and $\mathbf{S}_2 = \mathbf{S}(L_i)$. We will also write $g_{\mathbf{S}(\mathcal{W}_{w_i})} = g_{\mathbf{S}_1(\mathcal{W}_{w_i})}$ to make clear that $g_{\mathbf{S}(\mathcal{W}_{w_i})}$ is completely determined from \mathbf{S}_1 . We likewise write $E_{\mathbf{S},i} = E_{\mathbf{S}_1,i}$ as this set also depends only on \mathbf{S}_1 .

Now let S_1 be any outcome of \mathbf{S}_1 for which $\mathcal{L}_{\mathcal{D}}^{5/6}(g_{S_1(\mathcal{W}_{w_i})}) \geq (a/627)(d + \ln(1/\delta))/20^\kappa$. The samples \mathbf{S}_2 are independent of the outcome $\mathbf{S}_1 = S_1$. Thus each sample in \mathbf{S}_2 belongs to $E_{S_1,i}$ with probability at least $(a/627)(d + \ln(1/\delta))/20^\kappa$ and this is independent across the samples. Thus $\mathbb{E}_{\mathbf{S}_2 \sim \mathcal{D}_c^{|L_i|}}[|\mathbf{S}_2 \cap E_{S_1,i}|] \geq |L_i|(a/627)(d + \ln(1/\delta))/20^\kappa = 20^{\kappa-1}(a/627)(d + \ln(1/\delta))/20^\kappa = (a/12540)(d + \ln(1/\delta))$. For $a \geq 10^{16}$, this is at least $2 \cdot 10^9(d + \ln(57/\delta))$. A Chernoff bound implies that

$$\Pr_{\mathbf{S}_2 \sim \mathcal{D}_c^{|L_i|}}[|\mathbf{S}_2 \cap E_{S_1,i}| < 10^9(d + \ln(57/\delta))] < \exp(-10^9(d + \ln(57/\delta))/8) < \delta/57.$$

Since this holds for every outcome S_1 of \mathbf{S}_1 where $\mathcal{L}_{\mathcal{D}}^{5/6}(g_{S_1(\mathcal{W}_{w_i})}) \geq (a/627)(d + \ln(1/\delta))/20^\kappa$, the conclusion follows. \square

To prove the last lemma, we need the following classic result by Vapnik [16] with slightly better constants due to Blumer et al. [3]

Theorem 3 (Blumer et al. [3]). *For any $0 < \delta < 1$ and any distribution \mathcal{D} over \mathcal{X} , it holds with probability at least $1 - \delta$ over a set $\mathcal{S} \sim \mathcal{D}_c^m$ that every $h \in \mathbb{C}$ with $h(x) = c(x)$ for all $(x, c(x)) \in \mathcal{S}$ satisfies*

$$\mathcal{L}_{\mathcal{D}}(h) \leq (2/m)(d \lg_2(2em/d) + \lg_2(2/\delta)).$$

Corollary 1. *For any $0 < \delta < 1$ and any distribution \mathcal{D} over \mathcal{X} , it holds with probability at least $1 - \delta$ over a set $\mathcal{S} \sim \mathcal{D}_c^m$ with $m \geq 10^9(d + \ln(1/\delta))$ that every $h \in \mathbb{C}$ with $h(x) = c(x)$ for all $(x, c(x)) \in \mathcal{S}$ satisfies*

$$\mathcal{L}_{\mathcal{D}}(h) \leq 1/62700.$$

Proof. The proof can be found in the appendix. It is merely a matter of inserting the value of m in Theorem 3. \square

Proof of Lemma 11. We re-use the notation from the proof of Lemma 10 where we let $\mathbf{S}_1 = \mathbf{S}([m] \setminus L_i)$ and $\mathbf{S}_2 = \mathbf{S}(L_i)$, again exploiting that $E_{\mathbf{S},i} = E_{\mathbf{S}_1,i}$ and $g_{\mathbf{S}(\mathcal{W}_{w_i})} = g_{\mathbf{S}_1(\mathcal{W}_{w_i})}$ are fully determined from \mathbf{S}_1 .

Now fix any possible outcome S_1 of \mathbf{S}_1 . This also fixes $E_{S_1,i}$ and $g_{S_1(\mathcal{W}_{w_i})}$. For $\mathbf{S}_2 \sim \mathcal{D}_c^{|L_i|}$, let $\mathbf{R} \subseteq L_i$ be the subset of indices $j \in L_i$ for which $\mathbf{S}(\{j\}) \in E_{S_1,i}$ (note $\mathbf{S}(\{j\}) \in \mathbf{S}_2$). Now consider any outcome R of \mathbf{R} for which $|R| \geq 10^9(d + \ln(57/\delta))$. Conditioned on $\mathbf{R} = R$, the samples $\mathbf{S}(R) \subseteq \mathbf{S}_2$ are i.i.d. from the conditional distribution $\mathcal{D}(\cdot | E_{S_1,i})$. It follows from Corollary 1 that, conditioned on $\mathbf{R} = R$, it holds with probability at least $1 - \delta/57$ over $\mathbf{S}(R)$ that every $h \in \mathbb{C}$ with $h(x) = c(x)$ for all $(x, c(x)) \in \mathbf{S}(R)$ have $\mathcal{L}_{\mathcal{D}(\cdot|E_{S_1,i})}(h) \leq 1/62700$. Since $\mathbf{S}(R) \subseteq \mathbf{S}_2$, the same holds for all $h \in \mathbb{C}$ with $h(x) = c(x)$ for every $(x, c(x)) \in \mathbf{S}_2 = \mathbf{S}(L_i)$.

Since this holds for all S_1 and R with $|R| \geq 10^9(d + \ln(57/\delta))$, we conclude that $\Pr_{\mathbf{S} \sim \mathcal{D}_c^m}[\mathbf{S} \in \mathcal{E}_i''] \leq \delta/57$. \square

4 Conclusion

In this work, we have shown that the classic bagging heuristic is an optimal PAC learner in the realizable setting if we sample just a logarithmic number of bootstrap samples.

As mentioned in Section 1, we can also combine our proof with the recent work by [12] to obtain a simpler optimal weak to strong learner. Let us comment further on this here. Said briefly, a γ -weak learner, is a learning algorithm that from a large enough constant number of samples from an unknown data distribution \mathcal{D} , with constant probability outputs a hypothesis h with $\mathcal{L}_{\mathcal{D}}(h) \leq 1/2 - \gamma$, i.e. it has an advantage of γ over random guessing. On the other hand, an (ϵ, δ) -strong learner, is a learning algorithm that given enough samples $m(\epsilon, \delta)$, with probability $1 - \delta$ outputs a hypothesis h with $\mathcal{L}_{\mathcal{D}}(h) \leq \epsilon$ [10, 11]. In the recent work [12], the authors showed that if we use the *Sub-Sample* procedure from [9] (see Algorithm 1), but instead of empirical risk minimization, run a version of AdaBoost called AdaBoost $_{\nu}^*$ by [13] on the sub-samples, then this combination gives a weak to strong learner with an optimal sample complexity of $O(d/(\epsilon\gamma^2) + \lg(1/\delta)/\epsilon)$. Here d is the VC-dimension of the hypothesis set from which the weak learner outputs. Examining their proof, we can directly substitute Hanneke’s *Sub-Sample* by bagging, as our new proof boils down to applying Hanneke’s reasoning on the buckets \mathcal{C}_i . We thus get an optimal weak to strong learner from the combination of the two classic concepts of bagging and boosting.

Let us also make another comparison to Hanneke’s algorithm. In his work, the number of sub-samples is independent of δ , whereas our result for bagging requires $\Omega(\ln(m/\delta))$ sub-samples. Indeed, as bagging is a randomized algorithm, there will be some non-zero contribution to the failure probability due to the number of sub-samples (with some non-zero probability, all sub-samples are the same). We find it an interesting open problem whether the analysis can be tightened to yield a better dependency on δ in the number of sub-samples needed. We also mention that $O(m)$ sub-samples suffice for an optimal dependency on δ for any $0 < \delta < 1$. This is because $O(m)$ sub-samples suffice for any $\delta \geq \exp(-O(m))$, and for even smaller δ , we have that $\ln(1/\delta)/m > 1$.

Finally, let us comment on the unspecified constant $a > 0$ in Theorem 1. In our proof, we did not attempt to minimize a and indeed it is rather ridiculous. With some care, one could certainly shave several orders of magnitude, but at the end of it, we still rely on Hanneke’s proof which also does not have small constants. We believe that bagging actually provides quite good constants, but this would require a different proof strategy. At least our reduction to understanding $\mathcal{L}_{\mathcal{D}}^{1/3}(g_S)$ is very tight and incurs almost no loss in the constant a . A good starting point for improvements would be to find an alternative proof that $\mathcal{L}_{\mathcal{D}}^{1/3}(g_S)$ is small with high probability over \mathbf{S} .

References

- [1] P. Auer and R. Ortner. A new pac bound for intersection-closed concept classes. In *Learning Theory*, pages 408–414. Springer Berlin Heidelberg, 2004.
- [2] P. Bartlett, Y. Freund, W. S. Lee, and R. E. Schapire. Boosting the margin: a new explanation for the effectiveness of voting methods. *The Annals of Statistics*, 26(5):1651 – 1686, 1998.
- [3] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *Journal of the ACM (JACM)*, 36(4):929–965, 1989.
- [4] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [5] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- [6] A. Ehrenfeucht, D. Haussler, M. Kearns, and L. Valiant. A general lower bound on the number of examples needed for learning. *Information and Computation*, 82(3):247–261, 1989.
- [7] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [8] W. Gao and Z.-H. Zhou. On the doubt about margin explanation of boosting. *Artificial Intelligence*, 203:1–18, 2013.

- [9] S. Hanneke. The optimal sample complexity of pac learning. *The Journal of Machine Learning Research*, 17(1):1319–1333, 2016.
- [10] M. Kearns. Learning boolean formulae or finite automata is as hard as factoring. *Technical Report TR-14-88 Harvard University Aikem Computation Laboratory*, 1988.
- [11] M. Kearns and L. Valiant. Cryptographic limitations on learning boolean formulae and finite automata. *Journal of the ACM (JACM)*, 41(1):67–95, 1994.
- [12] K. G. Larsen and M. Ritzert. Optimal weak to strong learning. *Advances in Neural Information Processing Systems*, 2022. To appear.
- [13] G. Rätsch, M. K. Warmuth, and J. Shawe-Taylor. Efficient margin maximizing with boosting. *Journal of Machine Learning Research*, 6(12), 2005.
- [14] H. U. Simon. An almost optimal pac algorithm. In *Conference on Learning Theory*, pages 1552–1563. PMLR, 2015.
- [15] L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- [16] V. Vapnik. *Estimation of Dependences Based on Empirical Data: Springer Series in Statistics (Springer Series in Statistics)*. Springer-Verlag, Berlin, Heidelberg, 1982.
- [17] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.

A Deferred Proofs

Here we given the omitted proofs of the claims only needing standard arguments.

Proof of Lemma 5. We need to show that $1/6 \geq \mathcal{P}(0) = \Pr_{\mathbf{I} \sim [m]^n} [|D(\mathbf{I})| \notin [0.01m, 0.9m]]$. Consider first $\Pr_{\mathbf{I} \sim [m]^n} [|D(\mathbf{I})| < 0.01m]$. If $|D(\mathbf{I})| < 0.01m$, then there must be a set $Q \subseteq [m]$ of cardinality $0.01m$ with $D(I) \subseteq Q$. For a fixed Q , this happens with probability precisely $(|Q|/m)^n \leq (0.01)^{0.02m}$. A union bound over all $\binom{m}{0.01m} \leq (100e)^{0.01m}$ choices for Q implies $\Pr_{\mathbf{I} \sim [m]^n} [|D(\mathbf{I})| < 0.01m] \leq (e/100)^{0.01m}$. For $m \geq 100$, this is at most $1/30$. Consider next $\Pr_{\mathbf{I} \sim [m]^n} [|D(\mathbf{I})| \geq 0.9m]$. This probability is only non-zero for $n \geq 0.9m$ so we assume this from hereon. If $|D(\mathbf{I})| \geq 0.9m$, then there must be a set Q of cardinality $0.9m$ such that $Q \subseteq D(\mathbf{I})$. For a fixed such Q , any I with $Q \subseteq D(I)$ can be uniquely described by first specifying the first occurrence of each $i \in Q$ as an index into I and then specifying the remaining indices in I one at a time. There are thus no more than $\binom{n}{0.9m} (0.9m)! m^{n-0.9m}$ such I . Hence $\Pr_{\mathbf{I} \sim [m]^n} [Q \subseteq D(I)] \leq \binom{n}{0.9m} (0.9m)! m^{-0.9m} \leq \binom{m}{0.9m} (0.9m)! m^{-0.9m} = m! / ((0.1m)! m^{0.9m})$. Using Stirling’s approximation, this is no more than $\sqrt{10e} \cdot m^m e^{-m} / ((0.1m)^{0.1m} e^{-0.1m} m^{0.9m}) \leq 6 \cdot 10^{0.1m} e^{-0.9m} \leq 6 \cdot e^{0.3m} e^{-0.9m} = 6 \cdot e^{-0.6m}$. A union bound over all $\binom{m}{0.9m} = \binom{m}{0.1m} \leq (10e)^{0.1m} \leq e^{0.4m}$ choices for Q implies $\Pr_{\mathbf{I} \sim [m]^n} [|D(\mathbf{I})| \geq 0.9m] \leq 6 \cdot e^{-0.2m}$. For $m \geq 100$, this is much less than $1/30$. We finally conclude $\mathcal{P}(0) \leq 1/30 + 1/30 < 1/6$ as claimed. \square

Proof of Corollary 1. We see that for $m \geq 10^9(d + \ln(1/\delta))$, we have both

$$\begin{aligned}
\frac{2(d \lg_2(2em/d) + 1)}{m} &\leq \frac{2d \lg_2(4e10^9 d/d)}{10^9 d} \\
&= \frac{2 \lg_2(4e10^9)}{10^9} \\
&\leq 1/125400
\end{aligned}$$

and

$$\begin{aligned}\frac{2\lg_2(1/\delta)}{m} &\leq \frac{2\lg_2(1/\delta)}{10^9 \ln(1/\delta)} \\ &\leq \frac{2\lg_2(e)}{10^9} \\ &\leq 1/125400.\end{aligned}$$

Thus $(2/m)(d\lg_2(2em/d) + \lg_2(2/\delta)) = (2/m)(d\lg_2(em/d) + \lg_2(1/\delta) + 1) \leq 2/125400 = 1/62700$. \square