

Fine-Grained Distribution-Dependent Learning Curves

Olivier Bousquet

Google, Brain Team

OBOUSQUET@GOOGLE.COM

Steve Hanneke

Purdue University

STEVE.HANNEKE@GMAIL.COM

Shay Moran

Technion – Israel Institute of Technology; Google Research

SMORAN@TECHNION.AC.IL

Jonathan Shafer

University of California, Berkeley

SHAFFERJO@BERKELEY.EDU

Ilya Tolstikhin

Google, Brain Team

TOLSTIKHIN@GOOGLE.COM

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

Learning curves plot the expected error of a learning algorithm as a function of the number of labeled samples it receives from a target distribution. They are widely used as a measure of an algorithm’s performance, but classic PAC learning theory cannot explain their behavior.

As observed by [Antos and Lugosi \(1996, 1998\)](#), the classic ‘No Free Lunch’ lower bounds only trace the upper envelope above all learning curves of specific target distributions. For a concept class with VC dimension d the classic bound decays like d/n , yet it is possible that the learning curve for *every* specific distribution decays exponentially. In this case, for each n there exists a different ‘hard’ distribution requiring d/n samples. [Antos and Lugosi](#) asked which concept classes admit a ‘strong minimax lower bound’ – a lower bound of d'/n that holds for a fixed distribution for infinitely many n .

We solve this problem in a principled manner, by introducing a combinatorial dimension called VCL that characterizes the best d' for which d'/n is a strong minimax lower bound. Conceptually, the VCL dimension determines the asymptotic rate of decay of the minimax learning curve, which we call the ‘distribution-free trail’ of the class. Our characterization strengthens the lower bounds of [Bousquet, Hanneke, Moran, van Handel, and Yehudayoff \(2021\)](#), and it refines their analysis of learning curves, by showing that for classes with finite VCL the learning rate can be decomposed into a linear component that depends only on the hypothesis class and a faster (e.g., exponential) component that depends also on the target distribution. As a corollary, we recover the lower bound of [Antos and Lugosi \(1996, 1998\)](#) for half-spaces in \mathbb{R}^d .

Finally, to provide another viewpoint on our work and how it compares to traditional PAC learning bounds, we also present an alternative formulation of our results in a language that is closer to the PAC setting.

Keywords: PAC Learning, Online Learning, Universal Learning, Strong Minimax Lower Bounds.

1. Introduction

The most fundamental question in learning theory is arguably “*what can be learned, and what amount of resources (such as data and computation) is necessary for learning when learning is possible?*” The classic and definitive mathematical treatment of this question for supervised learning has traditionally been provided by the PAC framework, due to [Vapnik and Chervonenkis \(1968\)](#) and [Valiant \(1984\)](#). However, it has become increasingly clear that the PAC model does not accurately capture the reality of learning; VC bounds are overly pessimistic, and modern machine learning algorithms routinely outperform them. This is partially because the PAC model constitutes a worst-case analysis over all distributions. In contrast, machine learning practitioners are typically faced with one (or a few) target distributions, they are interested in optimizing performance only with respect to these specific distributions, and therefore they can vastly outdo the worst-case analysis.

Indeed, [Antos and Lugosi \(1996, 1998\)](#) observed that while the classic PAC bounds decay like $\Omega(d/n)$ for a class of VC dimension d , there exist hypothesis classes with arbitrarily large VC dimension that are learnable such that for every realizable distribution the expected loss decays exponentially fast.

They wrote:

“[I]n some sense, these [VC] lower bounds are not satisfactory. They do not tell us anything about the way the error decreases as the sample size is increased for a given classification problem. These bounds, for each n , give information about the maximal error within the class, but not about the behavior of the error for a single fixed [distribution] as the sample size n increases. In other words, the ‘bad’ [distribution], causing the largest error for a learning rule, may be different for each n .”¹

This led them to study the following question:

Question 1 (Strong Minimax Lower Bound) *For a VC class, what is the largest $d' \geq 0$ such that for every learning algorithm there exists a realizable distribution for which the expected 0-1 loss is at least d'/n infinitely often?*

They were able to answer this question for a number of specific hypothesis classes. Furthermore, they showed that “it is neither the VC dimension, nor the rate of increase of the shatter coefficients of the class” that determine the answer. The general case, however, has remained open.

In this paper we solve Question 1. We do so in a principled manner, by contributing to the nascent study of distribution-dependent learning curves. We build upon the recent results of [Bousquet, Hanneke, Moran, van Handel, and Yehudayoff \(2021\)](#), who offered a characterization of these curves.

For each instance, consisting of a hypothesis class and a target distribution, the distribution-dependent learning curve is the expected 0-1 loss of a learning algorithm as a function of the number of i.i.d. samples from the distribution (see Appendix A for formal definitions).

1. From [Antos and Lugosi \(1996\)](#), edited for clarity.




Figure 1: Illustration of the difference between distribution-dependent and PAC rates. Each red curve shows exponential decay of the error $\text{Opt} = \mathbb{E}_{S \sim \mathcal{D}^n} [L_{\mathcal{D}}^{0-1}(\hat{h}_S)]$ for a different data distribution \mathcal{D} ; but the PAC rate only captures the pointwise supremum of these curves (the blue curve) which decays linearly at best.

Source: Bousquet et al. (2021), adapted with permission.

The lay of the land when viewed from the perspective of distribution-dependent learning curves is remarkably structured, and remarkably different from that of the PAC model, as captured by the following crisp trichotomy.

Theorem 2 (Bousquet et al. (2021), Theorem 1.6) *For every concept class \mathcal{H} with $|\mathcal{H}| \geq 3$, exactly one of the following holds:*

- \mathcal{H} is learnable with optimal rate e^{-n} .
- \mathcal{H} is learnable with optimal rate $\frac{1}{n}$.
- \mathcal{H} requires arbitrarily slow rates.

This differs markedly from PAC learning bounds because, for example, it is possible for a class \mathcal{H} to have infinite VC dimension but still be learnable with an exponential rate; and ERM algorithms, which are optimal in the PAC setting, can perform arbitrarily worse than the best learning algorithm in the distribution-dependent setting (see Example 2.3 and Example 2.6 respectively in Bousquet et al., 2021). Bousquet et al. (2021) also provide a combinatorial characterization (via infinite trees) that determines for each hypothesis class \mathcal{H} which of the three prongs of the trichotomy it belongs to.

While the trichotomy of Theorem 2 is an important characterization, it is far from constituting a complete distribution-dependent theory of supervised learning. To see this, we recall the definition of *learning at rate $R(n)$* for some function $R : \mathbb{N} \rightarrow [0, 1]$, as used in the trichotomy. Roughly (see Definition 25 below), a class \mathcal{H} is learnable at rate R if there exists a learning algorithm such that for any realizable distribution there exist parameters $C, c \geq 0$ (that depend on the distribution) such that the 0-1 loss of the algorithm after seeing n i.i.d. samples from the distribution is at most $C \cdot R(c \cdot n)$. In other words, each instance, consisting of a hypothesis class and a distribution, determines a pair of parameters $C, c \geq 0$ which together specify the shape of the learning curve.

The characterization of Theorem 2 explains the general shape of the learning curve (exponential, linear, or arbitrarily slow decay), but it is silent with regard to the parameters C, c that specify its

precise shape. In particular, it is not clear in what manner the class and the distribution ‘interact’ to produce these parameters. This is where the present paper comes in.

1.1. Main Results

Our main contributions are:

1. We solve the main question left open by [Antos and Lugosi \(1998, 1996\)](#). We define a new combinatorial dimension that we call the Vapnik–Chervonenkis–Littlestone dimension, or VCL (see [Definition 24](#)), and show that it characterizes the magnitude of the strong minimax lower bound up to universal constants, as follows.

There exist universal constants $\alpha, \beta > 0$ such that for any VC class \mathcal{H} , the dimension $d = \text{VCL}(\mathcal{H}) \leq \text{VC}(\mathcal{H})$, and the number d' defined in [Question 1](#) satisfies $\alpha \cdot d \leq d' \leq \beta \cdot d$.

2. Conceptually, we introduce a notion of *distribution-free tail*, and prove a characterization of the optimal rate of decay for distribution-free tails via the VCL dimension (see [Definition 4](#) and [Theorem 5](#)).
3. Technically, we present a fine-grained analysis of distribution-dependent learning curves. For any class \mathcal{H} , if the dimension $d = \text{VCL}(\mathcal{H}) \geq 0$ is finite, then the optimal expected loss Opt can be bounded by

$$\frac{\Omega(d)}{n} \leq \text{Opt} \leq \frac{O(d)}{n} + \lambda(n), \quad (1)$$

where $n \in \mathbb{N}$ is the number of i.i.d. samples used, $\lambda(n) \in o(1/n)$, and the inequalities hold as follows: for any learning algorithm there exists a distribution such that the lower bound holds for infinitely many values of n ; the upper bound holds for a learning algorithm that we present, for all distributions and all $n \in \mathbb{N}$; the $\Omega(\cdot)$ and $O(\cdot)$ notations hide universal multiplicative constants that are independent of \mathcal{H} and of the distribution; the function $\lambda(n)$ depends on \mathcal{H} and on the distribution. Specifically, we use $\lambda(n) = C \cdot e^{-c \cdot n}$ where the parameters $C, c \geq 0$ depend on \mathcal{H} and on the distribution, but other choices of λ are possible (see [Remark 3](#)).

This bound captures both linear rates (when $d > 0$) and exponential rates (when $d = 0$). We call this type of bound the *fine-grained* rate of \mathcal{H} , to distinguish it from the notion of *coarse* rate used in [Theorem 2](#). See [Definition 26](#) and [Lemma 6](#) for a formal statement of this result.

4. Furthermore, for the hard distribution that satisfies the lower bound in [Equation \(1\)](#), the marginal on the domain \mathcal{X} depends only on the class \mathcal{H} . In particular, in contrast to the lower bounds of [Bousquet et al. \(2021\)](#), the marginal on the domain does not depend on the learning algorithm. This means that in the distribution-dependent learning curve setting, access to unlabeled data is not helpful for learning classes with finite VCL dimension. Namely, semi-supervised learning and supervised learning require the same number of labeled samples.

We note that this is a non-trivial result, employing a sophisticated application of Fatou’s lemma which enables reversing the order of quantifiers, as well as an argument from Ramsey theory.

5. For any class \mathcal{H} , if $\text{VCL}(\mathcal{H}) = \infty$ and \mathcal{H} does not shatter an infinite strong VCL tree (see [Definitions 22](#) and [23](#) below), then \mathcal{H} has a *strongly distribution-dependent linear rate*. Namely,

for every $c \geq 0$ there exists a distribution such that $\text{Opt} \geq c/n$ for infinitely many $n \in \mathbb{N}$ (see Definition 28 and Lemma 6).²

6. We offer an equivalent formulation of our results in a language that is closer to the traditional PAC framework. This provides another viewpoint on our work and how it compares to traditional PAC bounds. See Theorem 10.
7. As a special case, we recover the lower bound of Antos and Lugosi (1998) for half-spaces. We do so by introducing a technique for proving strong lower bounds via a ‘fractal’ argument, which may be useful for other classes as well. (See Theorem 8 and Appendix C.)

1.2. Benefits of the New Characterization

Our upper bound in Equation (1) offers a refinement and reinterpretation of Theorem 2:

- **Upper bound refinement.** For the case of linear rates, Bousquet et al. (2021) showed an upper bound of $\frac{c}{n}$, where $c \geq 0$ depends both on the class and on the distribution. In contrast, our expression in the upper bound constitutes a *decomposition* of the error rate into a linear component that depends only on the class, and a faster (e.g., exponential) component that depends on the class and on the distribution. We view this as a step towards a complete characterization of the optimal distribution-dependent learning rate for supervised learning.
- **Upper bound reinterpretation.** Whereas Theorem 2 depicts exponential rates and linear rates as being two entirely different beasts, Equation (1) presents them in a more unified light, with exponential rates constituting the special case of d/n where $d = 0$.

Our lower bound in Equation (1) offers meaningful improvements over both the previous distribution-dependent lower bound and over the classic ‘no free lunch’ lower bounds from PAC learning, and also constitutes a partial unification of these two results.

- **Quantitative strengthening of distribution-dependent lower bounds.** For classes with linear learning rates, the best previously known distribution-dependent lower bound that applies to general classes was $\Omega(1/n)$ (Bousquet et al., 2021). This applies equally to all classes that have linear rates, and does not distinguish between different degrees of hardness within that broad set of classes. In contrast, we are able to prove a lower bound of $\text{Opt} \geq \Omega(d/n)$, for d that depends only on the class and is tight up to a universal multiplicative constant (independent of the class and of the distribution).
- **Qualitative strengthening of distribution-dependent lower bounds.** Classic PAC lower bounds (discussed further in the next bullet) have the following formulation:

There exists a distribution $\mathcal{D}_{\mathcal{X}} \in \Delta(\mathcal{X})$ such that for any learning algorithm A there exists a hard distribution $\mathcal{D} \in \Delta(\mathcal{X} \times \{0, 1\})$ such that the marginal distribution of \mathcal{D} on \mathcal{X} equals $\mathcal{D}_{\mathcal{X}}$, and the loss of A on distribution \mathcal{D} is large.³ (★)

In contrast, the linear lower bound of Bousquet et al. (2021) offered a weaker statement:

-
2. In the remaining case where $\text{VCL}(\mathcal{H}) = \infty$ and \mathcal{H} has an infinite strong VCL tree, \mathcal{H} requires arbitrarily slow rates, as shown by Bousquet et al. (2021).
 3. The marginal distribution $\mathcal{D}_{\mathcal{X}}$ is simply a uniform distribution over a subset of the domain of cardinality $\text{VC}(\mathcal{H})$ that is shattered by \mathcal{H} in the VC sense.

For any learning algorithm A there exists a hard distribution $\mathcal{D} \in \Delta(\mathcal{X} \times \{0, 1\})$ such that the loss of A on distribution \mathcal{D} is large.

In particular, this weaker formulation left open the possibility that an algorithm that has access to unlabeled samples (as in the semi-supervised learning setting) could beat the lower bound. We show that that is not the case. We strengthen the lower bound of [Bousquet et al. \(2021\)](#), obtaining the stronger formulation as in (\star) .

- **Qualitative strengthening of PAC lower bounds.** Classic PAC learning theory features ‘no free lunch’ lower bounds (e.g., [Shalev-Shwartz and Ben-David, 2014](#), Theorem 5.1), which imply the VC lower bound appearing in the fundamental theorem of PAC learning (e.g., [Shalev-Shwartz and Ben-David, 2014](#), Theorem 6.8, Item 3). These lower bounds leave something to be desired.

To see this, fix a hypothesis class of VC dimension d . The VC bound states that for every $\varepsilon > 0$ there exists a hard (worst-case) distribution \mathcal{D}_ε such that achieving loss at most ε with high probability requires at least $\Omega(d/\varepsilon)$ i.i.d. samples from \mathcal{D}_ε . Namely, for a fixed hypothesis class and a sequence of positive values $\varepsilon_1, \varepsilon_2, \dots$ there exists a sequence of *distinct* hard distributions $\mathcal{D}_1, \mathcal{D}_2, \dots$ such that each \mathcal{D}_i is a hard distribution for achieving loss ε_i — but it is typically not a hard distribution for other values of ε .

Clearly, the type of lower bound studied in VC bounds is strictly weaker than the distribution-dependent lower bounds studied in this paper, where there exists a specific hard distribution such that the lower bound holds for infinitely many values of ε . And as we argued above, instance specific lower bounds are a better match to the reality of most machine learning practitioners, who typically face a specific (fixed) unknown distribution, and would like to calculate how many samples are necessary for obtaining loss ε_1 , or loss ε_2 , or loss ε_3 , etc. — all with respect to the *same* fixed unknown distribution.

Thus, an interesting open question is “for which VC classes is it possible to obtain distribution-dependent linear lower bounds of $\Omega(d/n)$?” (where $d = \text{VC}(\mathcal{H})$ and the bound holds for a single distribution for infinitely many $n \in \mathbb{N}$). This question, which was studied by [Antos and Lugosi \(1996, 1998\)](#), is answered by our characterization as follows. Let \mathcal{H} be a VC class with $0 \leq d' = \text{VCL}(\mathcal{H}) \leq \text{VC}(\mathcal{H}) = d$. If $d' > 0$ then \mathcal{H} has a distribution-dependent linear lower bound of $\Omega(d'/n)$. Otherwise, if $d' = 0$ then \mathcal{H} does not have a distribution-dependent linear lower bound; rather, each learning curve decays exponentially and the upper envelope of all the learning curves decays linearly as $\Theta(d/n)$. In this sense, our results offer a unified perspective of PAC and distribution-dependent lower bounds.

Remark 3 (On the exponential term in the decomposition) *The upper bound in Equation (1) decomposes into $O(d/n) + \lambda(n)$, where λ is a function that decays quickly, at a rate that depends on the distribution. Taking $\lambda(n) = C \cdot e^{-c \cdot n}$ to be an exponential function, as we do in this paper, is a somewhat arbitrary choice. For classes with positive VCL dimension, the same results (e.g., Theorem 5 and Lemma 6) would hold also, say, for a double-exponential function $\lambda(n) = C \cdot e^{-c \cdot e^n}$, albeit with different constants C, c and β in those results. The essential insight expressed by these decompositions is captured by the notion of a distribution-free tail (Definition 4), discussed below.*

1.3. Related Works

Universal Learning. Our work explores the distribution-dependent setting, also called the *universal learning* setting, which was recently formalized by [Bousquet, Hanneke, Moran, van Handel, and Yehudayoff \(2021\)](#). However, it is worthwhile to note that this framework has been studied by earlier works as well.

[Schuurmans \(1997\)](#) revealed the distinction between exponential and linear rates in the universal setting. In more detail, [Schuurmans \(1997\)](#) characterized the optimal learning rate for classes \mathcal{H} that are *concept chains*, meaning that every $h, h' \in \mathcal{H}$ have either $h \leq h'$ everywhere or $h' \leq h$ everywhere.

[van Handel \(2013\)](#) studied the uniform convergence property via the universal lens. He characterizes those hypothesis classes \mathcal{H} satisfying that the empirical losses of *all* hypotheses in the class simultaneously and uniformly converge to the corresponding population losses as the number of examples tends to infinity. The difference with the (more common) distribution-free uniform convergence is that in the universal variant, the rate of the uniform convergence can depend on the source distribution.

An extreme notion of universal learnability is *universal consistency*: a learning rule is universally consistent if its expected loss converges to the Bayes optimal risk for every target distribution. In other words, such algorithms learn every distribution (but at a distribution-dependent rate). The first proof that such learning is possible was provided by [Stone \(1977\)](#) who established the universal consistency of several algorithms such as k -nearest neighbor predictors, kernel rules, and histogram rules; see [Devroye, Györfi, and Lugosi \(1996\)](#) for a thorough discussion of such results.

No Free Lunch. One of the technical contributions in this work is the identification of the VCL dimension as the combinatorial parameter which characterizes when a strong form of the ‘no free lunch’ theorem holds. That is, for which classes is it the case that there exists a single fixed distribution which witnesses the strongest lower bound on the error rate for infinitely many sample sizes n .

The work by [Antos and Lugosi \(1998\)](#) explored this question for VC classes; that is, they asked for which VC classes such a strong ‘no free lunch’ theorem holds. [Antos and Lugosi \(1998\)](#) showed that d -dimensional half-spaces satisfy such a strong ‘no free lunch’ theorem by proving a lower bound of d/n on the learning rate. ([Schuurmans 1997](#) also established such a bound in the 1-dimensional case.) However, a characterization of VC classes with this property remained open; in fact, [Antos and Lugosi \(1998\)](#) explicitly concluded that it is “neither the VC dimension nor the rate of increase of shatter coefficients that determine the asymptotic behavior of the concept class”. Our work resolves this question by showing that the VCL dimension determines this behavior.

Strong Minimax. A recent work by [Ben-David and Blais \(2020\)](#) studies a similar type of lower bounds for the task of computing boolean functions up to error ε . They introduce a new type of minimax theorem which provides a single hard distribution for arbitrarily small ε .

2. Technical Overview

This section provides an overview of our results. All technical details are deferred to the appendices. We start with the following definition.

Definition 4 Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let $f : \mathbb{N} \rightarrow \mathbb{R}$ be a function. We say that f is a distribution-free tail for \mathcal{H} if $\lim_{n \rightarrow \infty} f(n) = 0$ and there exists a (possibly randomized) learning algorithm h such that for any realizable distribution $\mathcal{D} \in \text{Realizable}(\mathcal{H})$, the inequality

$$\mathbb{E}_{S \sim \mathcal{D}^n} [\mathbb{L}_{\mathcal{D}}^{0-1}(h_S)] \leq f(n)$$

holds for all n large enough. We say that \mathcal{H} admits a distribution-free tail if there exists a function that is a distribution-free tail for \mathcal{H} .

In other words, a vanishing f is not a distribution-free tail for \mathcal{H} if for every learning algorithm there exists a realizable distribution such that the expected loss is strictly greater than f infinitely often.

If \mathcal{H} is a class with $\text{VC}(\mathcal{H}) = d < \infty$, then $O(d/n)$ is distribution-free tail for \mathcal{H} , by the VC theorem. However, some class with VC dimension d might admit distribution-free tails that decay much faster than d/n .

Our first result characterizes the optimal rate of decay for distribution-free tails via the VCL dimension, as follows.

Theorem 5 *There exist constants $\alpha, \beta > 0$ as follows. Let \mathcal{X} be a set and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class.*

1. *If $\text{VCL}(\mathcal{H}) = \infty$ then \mathcal{H} does not admit a distribution-free tail.*
2. *If $\text{VCL}(\mathcal{H}) = d \in (0, \infty)$ then for any $\eta > 0$, $f(n) = \beta \cdot d/n + \eta$ is a distribution-free tail for \mathcal{H} . Furthermore, $\alpha \cdot d/n$ is not a distribution-free tail for \mathcal{H} , so the above choice of f is tight up to a multiplicative factor.*
3. *If $\text{VCL}(\mathcal{H}) = 0$, then any function $f(n)$ such that $f(n) = \omega(C \cdot e^{-c \cdot n})$ for all $C, c > 0$ is a distribution-free tail for \mathcal{H} . Furthermore, any function $f(n) = C \cdot e^{-c \cdot n}$ for fixed $C, c > 0$ is not a distribution-free tail for \mathcal{H} .*

Proof Idea for Theorem 5 Items 1 and 2 and the first part of Item 3 follow from Lemma 6. The ‘furthermore’ part of Item 3 follows from a simple modification of the proof of Proposition 2 in [Schuermans \(1997\)](#).⁴ ■

Theorem 5 is a corollary of Lemma 6, which provides a more complete picture.

Lemma 6 (Main Technical Lemma) *There exist constants $\alpha, \beta > 0$ as follows. Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let $d = \text{VCL}(\mathcal{H})$.*

1. *If $d < \infty$ then \mathcal{H} is learnable with optimal fine-grained rate d/n with gap factor $\gamma = \beta/\alpha$; furthermore, the marginal of the hard distribution on \mathcal{X} depends only on \mathcal{H} . Namely, there exists $\mathcal{D}_{\mathcal{X}} \in \Delta(\mathcal{X})$ such that for any (possibly randomized) learning algorithm \hat{h} there exists $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ such that the marginal distribution of \mathcal{D} on \mathcal{X} is $\mathcal{D}_{\mathcal{X}}$, and the inequality*

$$\alpha \cdot \frac{d}{n} \leq \mathbb{E}_{S \sim \mathcal{D}^n} [\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S)]$$

4. See also Lemma 4.2 in [Bousquet et al. \(2021\)](#)

holds for infinitely many $n \in \mathbb{N}$; and there exists a learning algorithm h^* such that for any $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ there exist parameters $C, c > 0$ such that

$$\forall n \in \mathbb{N} : \mathbb{E}_{S \sim \mathcal{D}^n} [\mathbb{L}_{\mathcal{D}}^{0,1}(h_S^*)] \leq \beta \cdot \frac{d}{n} + C \cdot e^{-cn}.$$

2. Otherwise, if \mathcal{H} does not shatter an infinite strong VCL tree, then \mathcal{H} is learnable with a strongly distribution-dependent linear rate. Namely, for any (possibly randomized) learning algorithm \hat{h} and any $c > 0$ there exists $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ such that the inequality

$$\frac{c}{n} \leq \mathbb{E}_{S \sim \mathcal{D}^n} [\mathbb{L}_{\mathcal{D}}^{0,1}(\hat{h}_S)]$$

holds for infinitely many $n \in \mathbb{N}$; and there exists a learning algorithm h^* such that for any $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ there exists $c > 0$ such that

$$\forall n \in \mathbb{N} : \mathbb{E}_{S \sim \mathcal{D}^n} [\mathbb{L}_{\mathcal{D}}^{0,1}(h_S^*)] \leq \frac{c}{n}.$$

3. Otherwise, \mathcal{H} requires arbitrarily slow rates.

Remark 7 Our proofs use $\alpha = 1/100$, $\beta = 8$, and $\gamma = 800$.


All proofs appear in Appendices B and C. Additionally, we provide a brief overview of the main ideas in each proof.

The proof of Lemma 6 is similar to the proof of Theorem 2 from Bousquet et al. (2021). One of the main differences is that we use d -VCL trees for the characterization. Our d -VCL trees (introduced in this paper, see Definition 19 and Figure 2(b)), are an intermediary refinement that lies between the 1-VCL trees and the strong VCL trees that were used in their proof. Identifying that this particular combinatorial structure characterizes the fine-grained rate is a non-trivial contribution of this paper.


Proving the lower bound of Lemma 6 requires some technical improvements upon the technique of Bousquet et al. (2021). For each learning algorithm, they constructed a hard distribution that is concentrated on an infinite ‘target’ branch chosen at random in the d -VCL tree, and argued that if a test point is deeper in the tree than all points in the training set, then the learner will make an incorrect prediction on that test point with probability $1/2$. That approach is not suitable for constructing a single marginal distribution $\mathcal{D}_{\mathcal{X}} \in \Delta(\mathcal{X})$ that is hard for all learning algorithms (because for every target branch there exists an algorithm that returns a hypothesis with low loss on that branch). Instead, we choose a marginal $\mathcal{D}_{\mathcal{X}}$ that is distributed roughly evenly over all branches in the tree, and construct a distribution $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ that has marginal $\mathcal{D}_{\mathcal{X}}$, and has labels corresponding to an infinite target branch in the tree.

In a general d -VCL tree, our approach would be problematic, because even if a test point is deeper in the tree than all the points in the training set, the labels for points in the training set *that do not belong to the target branch* could provide information about the target branch.⁵ We overcome this problem as described in the following proof idea.

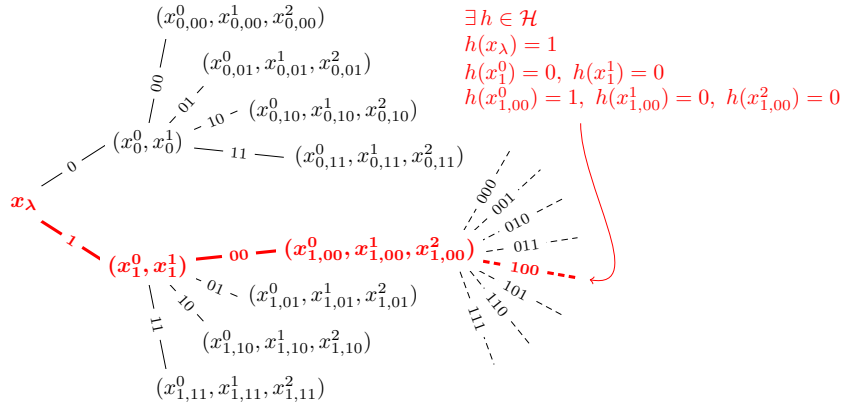
5. Consider the case where for some $x \in \mathcal{X}$ there exist infinite branches $\mathbf{y}^{(0)}$ and $\mathbf{y}^{(1)}$ in the tree, both of which do not contain x , such that for all $b \in \{0, 1\}$, it holds that all $h \in \mathcal{H}$ that are consistent with $\mathbf{y}^{(b)}$ satisfy $h(x) = b$. Then knowing the label for x allows the learner to eliminate one of the branches.



(a) The first few layers of an infinite 1-VCL tree (also called an *infinite Littlestone tree*).



(b) The first few layers of an infinite d -VCL tree. Every arc represents 2^d children, one for each possible labeling of the d points in the preceding node. Introduced in this paper, d -VCL trees form the basis for our novel characterization.



(c) The first few layers of an infinite strong VCL tree. (In Bousquet et al. (2021) this structure was called an *infinite VCL tree*. We add the modifier *strong* to distinguish it from d -VCL trees.)

Figure 2: VCL trees. Every finite branch is consistent with a concept $h \in \mathcal{H}$. This is illustrated here for one branch in each tree, shown in red.

Source: Bousquet et al. (2021), adapted with permission.

Proof Idea for Lemma 6 For the upper bound, we show that the nonexistence of d -VCL trees is equivalent to the existence of a winning strategy for the learner in an infinite two-player game called the ‘forbidden pattern game’. The equivalence is established via an intermediary online-learning game that is easier to analyze because it is a Gale-Stewart game (whereas the forbidden pattern game is not). A winning strategy for the learner in the forbidden pattern game can be converted into a learning algorithm for the distribution-dependent learning setting, by way of the one-inclusion algorithm of [Haussler, Littlestone, and Warmuth \(1994\)](#). The resulting algorithm has the desired rate of at most d/n .

For the lower bound, we use an elementary lemma from Ramsey theory to show that if \mathcal{H} shatters an infinite d -VCL tree, then it also shatters an infinite d -VCL tree that satisfies an additional *indifference* property that we define. This property implies that for any infinite branch in the d -VCL tree, labels for points that do not belong to the branch provide no information about the labels for points that appear lower down in the tree along the branch. Therefore, when the target branch is chosen randomly, we can argue that if the test point appears on the target branch and is lower than all the training samples, then the learner will make an incorrect prediction with probability $1/2$. The lower bound also involves a specific choice of parameters for the hard marginal distribution that enables a delicate application of Fatou’s lemma. ■

As a corollary of our characterization, we recover the lower bound of [Antos and Lugosi \(1998\)](#) for half-spaces in \mathbb{R}^d , up to a constants factor.

Theorem 8 (Antos and Lugosi 1998, Corollary 1) *There exists a constant $\alpha > 0$ as follows. Let $d \in \mathbb{N}$ and $\mathcal{X} = \mathbb{R}^d$. Let $\mathcal{H}_d \subseteq \{0, 1\}^{\mathcal{X}}$ be the set of closed half-spaces in \mathbb{R}^d . For any learning algorithm \hat{h} there exists a distribution $\mathcal{D} \in \text{Realizable}(\mathcal{H}_d)$ such that the inequality*

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[L_{\mathcal{D}}^{0-1}(\hat{h}_S) \right] \geq \alpha \cdot \frac{d}{n}$$

holds for infinitely many values $n \in \mathbb{N}$.

Proof Idea It suffices to show that $\text{VCL}(\mathcal{H}_d) = d$. We consider the dual class for \mathcal{H}_d , and show via a neat ‘fractal’ argument that one can construct an infinite d -VCL tree for \mathcal{H}_d . ■

We believe that the ‘fractal’ argument from this proof could be used to construct d -VCL trees for other classes as well.

Finally, we present another, mostly equivalent viewpoint on our results. Stated in a language that is closer to standard PAC learning, it enables a comparison between PAC bounds and distribution-dependent bounds.

Definition 9 *Let \mathcal{X} be a set and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class. Let $m : (0, 1)^2 \rightarrow \mathbb{N}$ and $k : \text{Realizable}(\mathcal{H}) \rightarrow \mathbb{N}$ be functions. We say that \mathcal{H} is eventually PAC learnable with sample complexity m and kick-in time k if there exist an algorithm \hat{h} such that for any distribution $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ and for any $\varepsilon, \delta \in (0, 1)$ the inequality*

$$\mathbb{P}_{S \sim \mathcal{D}^n} \left[L_{\mathcal{D}}^{0-1}(\hat{h}_S) \leq \varepsilon \right] \geq 1 - \delta$$

holds for all $n \geq \max\{m(\varepsilon, \delta), k(\mathcal{D})\}$.

Theorem 10 *There exist constants $\alpha, \beta > 0$ as follows. Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let $d \in \mathbb{N}$.*

1. *If $d = \text{VCL}(\mathcal{H}) < \infty$ then \mathcal{H} is eventually PAC learnable with sample complexity $m(\varepsilon, \delta) \leq \alpha d \log(1/\delta)/\varepsilon$.*
2. *If \mathcal{H} is eventually PAC learnable with sample complexity*

$$m(\varepsilon, \delta) = d \log(1/\delta)/\varepsilon,$$

then $\text{VCL}(\mathcal{H}) \leq \beta d$.

Proof Idea This follows from Lemma 6, together with a standard amplification argument for converting an algorithm with bounded expected error to a PAC learner. ■

3. Directions for Future Work

We have shown a characterization of fine-grained learning rates in the instance specific setting. Directions for future work include characterizing the precise parameters $C, c \geq 0$ in Equation (1), and obtaining an optimal gap factor (or equivalently, optimal parameters $\alpha, \beta \geq 0$ in Lemma 6).

Like the results of Bousquet et al. (2021), our results describe the *asymptotic* rate at which learning curves decay – but the results are silent as to the properties of learning curves for any finite number of samples. Devising a theory of learning curves that explains both asymptotic and non-asymptotic behavior in a unified way would be valuable.

Our result on semi-supervised learning (Item 4 in Section 1.1, Main Results) suggests that unlabeled data is not helpful in the setting of distribution-dependent learning curves. However, there are good reasons to believe that unlabeled data is helpful for learning in some real-world scenarios. We wonder how this tension could be resolved.

Acknowledgments

SM is a Robert J. Shillman Fellow. His research is supported in part by ISF grant 1225/20, by BSF grant 2018385, by an Azrieli Faculty Fellowship, by the Israel PBC-VATAT, by the Technion Center for Machine Learning and Intelligent Systems (MLIS), and by the European Union (ERC, GENERALIZATION, 101039692). 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 Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

JS would like to thank Robert D. Kleinberg and SM for their mentorship and support while hosting him at Cornell University and at the Technion, respectively. Part of this work was conducted while visiting those institutions. JS was supported in part by DARPA (Defense Advanced Research Projects Agency) contract #HR001120C0015, and the Simons Collaboration on the Theory of Algorithmic Fairness, and by SM’s Azrieli Faculty Fellowship. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Simons Foundation, DARPA, or the Azrieli Foundation.

IT would like to thank Konstantin Golubev for useful discussions.

References

- Reverse Fatou’s Lemma. *ProofWiki*, August 2022. URL https://proofwiki.org/wiki/Reverse_Fatou%27s_Lemma.
- András Antos and Gábor Lugosi. Strong minimax lower bounds for learning. In Avrim Blum and Michael J. Kearns, editors, *Proceedings of the Ninth Annual Conference on Computational Learning Theory, COLT 1996, Desenzano del Garda, Italy, June 28-July 1, 1996*, pages 303–309. ACM, 1996. doi: 10.1145/238061.238160. URL <https://doi.org/10.1145/238061.238160>.
- András Antos and Gábor Lugosi. Strong minimax lower bounds for learning. *Mach. Learn.*, 30(1):31–56, 1998. doi: 10.1023/A:1007454427662. URL <https://doi.org/10.1023/A:1007454427662>.
- Shalev Ben-David and Eric Blais. A new minimax theorem for randomized algorithms (extended abstract). In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 403–411. IEEE, 2020. doi: 10.1109/FOCS46700.2020.00045. URL <https://doi.org/10.1109/FOCS46700.2020.00045>.
- Olivier Bousquet, Steve Hanneke, Shay Moran, Ramon van Handel, and Amir Yehudayoff. A theory of universal learning. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC ’21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 532–541. ACM, 2021. doi: 10.1145/3406325.3451087. URL <https://doi.org/10.1145/3406325.3451087>.
- Andrew Browder. *Mathematical analysis: an introduction*. Springer Science & Business Media, 1996. doi: doi.org/10.1007/978-1-4612-0715-3. URL <https://doi.org/10.1007/978-1-4612-0715-3>.
- Luc Devroye, László Györfi, and Gábor Lugosi. *A Probabilistic Theory of Pattern Recognition*, volume 31 of *Stochastic Modelling and Applied Probability*. Springer, 1996. ISBN 978-1-4612-6877-2. doi: 10.1007/978-1-4612-0711-5. URL <https://doi.org/10.1007/978-1-4612-0711-5>.
- David Gale and F. M. Stewart. Infinite games with perfect information. In *Contributions to the theory of games, vol. 2*, Annals of Mathematics Studies, no. 28, pages 245–266. Princeton University Press, Princeton, N. J., 1953. doi: doi.org/10.1515/9781400881970-014. URL <https://doi.org/10.1515/9781400881970-014>.
- David Haussler, Nick Littlestone, and Manfred K. Warmuth. Predicting $\{0, 1\}$ -functions on randomly drawn points. *Inf. Comput.*, 115(2):248–292, 1994. doi: 10.1006/inco.1994.1097. URL <https://doi.org/10.1006/inco.1994.1097>.
- Dale Schuurmans. Characterizing rational versus exponential learning curves. *J. Comput. Syst. Sci.*, 55(1):140–160, 1997. doi: 10.1006/jcss.1997.1505. URL <https://doi.org/10.1006/jcss.1997.1505>.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning – From Theory to Algorithms*. Cambridge University Press, 2014. ISBN 978-1-10-705713-5. doi: doi.org/10.1017/CBO9781107298019. URL <https://doi.org/10.1017/CBO9781107298019>.

Charles J. Stone. Consistent nonparametric regression. *The Annals of Statistics*, 5(4):595–620, 1977. ISSN 00905364. URL <http://www.jstor.org/stable/2958783>.

Leslie G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, 1984. doi: 10.1145/1968.1972. URL <https://doi.org/10.1145/1968.1972>.

Ramon van Handel. The universal Glivenko-Cantelli property. *Probability and Related Fields*, 155:911–934, 2013. doi: doi.org/10.1007/s00440-012-0416-5. URL <https://doi.org/10.1007/s00440-012-0416-5>.

Vladimir N. Vapnik and Alexey Y. Chervonenkis. The uniform convergence of frequencies of the appearance of events to their probabilities. *Dokl. Akad. Nauk SSSR*, 181(4):781–783, 1968. URL <http://www.ams.org/mathscinet-getitem?mr=0231431>.

Appendix A. Preliminaries

Notation 11 $\mathbb{N} = \{1, 2, 3, \dots\}$, i.e., $0 \notin \mathbb{N}$. For any $n \in \mathbb{N}$, we denote $[n] = \{1, 2, 3, \dots, n\}$.

Notation 12 Let \mathcal{X} be a set. We write $\mathcal{X}^* = \cup_{t=0}^{\infty} \mathcal{X}^t$ to denote the set of all finite strings or finite vectors with elements from \mathcal{X} . \mathcal{X}^* includes the empty string, which we denote by λ .

Notation 13 For a set \mathcal{X} , we write $\Delta(\mathcal{X})$ to denote the set of all distribution with support contained in \mathcal{X} (with respect to some fixed σ -algebra).

Notation 14 For a (finite or infinite) vector $\mathbf{x} = (x_1, x_2, \dots)$, we write $\mathbf{x}_{\leq t}$ to denote the finite prefix (x_1, x_2, \dots, x_t) ; $((\mathbf{x}_t, \mathbf{y}_t))_{t \in \mathbb{N}}$ denotes an infinite sequence of pairs of vectors, where for each t , $(\mathbf{x}_t, \mathbf{y}_t)$ is a pair of vectors; for a (finite or infinite) sequence of pairs of vectors, we denote a finite prefix of the sequence by $(\mathbf{x}, \mathbf{y})_{\leq t} = ((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2) \dots, (\mathbf{x}_t, \mathbf{y}_t))$.

A.1. Traditional Learning Theory

Definition 15 Let \mathcal{X} be a set, and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a set of functions. Let $k \in \mathbb{N}$, $X = \{x_1, x_2, \dots, x_k\} \subseteq \mathcal{X}$. We say that \mathcal{H} shatters X if for any $y_1, y_2, \dots, y_k \in \{0, 1\}$ there exists $h \in \mathcal{H}$ such that $h(x_i) = y_i$ for all $i \in [k]$. The Vapnik–Chervonenkis (VC) dimension of \mathcal{H} , denoted $\text{VC}(\mathcal{H})$, is the largest $d \in \mathbb{N}$ for which there exist a set $X \subseteq \mathcal{X}$ of cardinality d that is shattered by \mathcal{H} . If \mathcal{H} shatters sets of cardinality arbitrarily large, we say that $\text{VC}(\mathcal{H}) = \infty$.

Definition 16 Let \mathcal{X} be a set. A learning algorithm for functions $\mathcal{X} \rightarrow \{0, 1\}$ is an algorithm \hat{h} that takes a sample $S \in (\mathcal{X} \times \{0, 1\})^*$ and outputs a function $\hat{h}_S : \mathcal{X} \rightarrow \{0, 1\}$. The mapping $S \mapsto \hat{h}_S$ may be randomized.

Definition 17 Let \mathcal{X} be a set, let $\mathcal{D} \in \Delta(\mathcal{X} \times \{0, 1\})$, and let $h : \mathcal{X} \rightarrow \{0, 1\}$ be a function. The 0-1 loss of \mathcal{H} with respect to \mathcal{D} is $L_{\mathcal{D}}^{0-1}(\mathcal{H}) = \mathbb{P}_{(x,y) \sim \mathcal{D}}[h(x) \neq y]$.

Definition 18 Let \mathcal{X} be a set, and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a class of functions. The set of realizable distributions for \mathcal{H} is

$$\text{Realizable}(\mathcal{H}) = \left\{ \mathcal{D} \in \Delta(\mathcal{X} \times \{0, 1\}) : \inf_{h \in \mathcal{H}} L_{\mathcal{D}}^{0-1}(h) = 0 \right\}.$$

A.2. The VCL Dimension

Definition 19 Let \mathcal{X} be a set, let $d \in \mathbb{N}$ and $\ell \in \mathbb{N} \cup \{0, \infty\}$. A d -VCL tree of depth ℓ with respect to \mathcal{X} is a set

$$T = \{ \mathbf{x}_{\mathbf{u}} \in \mathcal{X}^d : \mathbf{u} \in \{0, 1\}^{ds}, s \in \mathbb{N} \cup \{0\}, s \leq \ell \}. \quad (2)$$

We say that T is infinite if it has depth $\ell = \infty$.

Note that a d -VCL tree of depth 0 is not empty, rather it contains a single node \mathbf{x}_{λ} where λ denotes the empty string.

Definition 20 Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let $d \in \mathbb{N}$. For each $s \in \mathbb{N}$, let $\mathbf{x}_s = (x_s^1, \dots, x_s^d) \in \mathcal{X}^d$ and $\mathbf{y}_s = (y_s^1, \dots, y_s^d) \in \{0, 1\}^d$. Let $h \in \mathcal{H}$. For any $t \in \mathbb{N}$, we say that the finite sequence $(\mathbf{x}, \mathbf{y})_{\leq t} = ((\mathbf{x}_s, \mathbf{y}_s))_{s=1}^t$ is consistent with h if

$$\forall s \in [t] \forall i \in [d] : h(x_s^i) = y_s^i. \quad (3)$$

We say that the infinite sequence $((\mathbf{x}_s, \mathbf{y}_s))_{s \in \mathbb{N}}$ is consistent with \mathcal{H} if for any $t \in \mathbb{N}$ there exists $h \in \mathcal{H}$ such that $(\mathbf{x}, \mathbf{y})_{\leq t}$ is consistent with h .

Definition 21 Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let T be a d -VCL tree as in Equation (2). We say that \mathcal{H} shatters T if for every $t \in \mathbb{N}$, $t \leq \ell$, and every $\mathbf{y} \in \{0, 1\}^{dt}$ there exists a hypothesis $h \in \mathcal{H}$ that is consistent with $((\mathbf{x}_{\mathbf{y}_{\leq s-1}}, \mathbf{y}_s))_{s=1}^t$ in the sense that

$$\forall s \in [t] \forall j \in [d] : h(x_{\mathbf{y}_{\leq s-1}}^j) = y_s^j, \quad (4)$$

where we use the notation

$$\mathbf{y}_{\leq s} = \left((y_1^1, \dots, y_1^d), \dots, (y_s^1, \dots, y_s^d) \right) \in \{0, 1\}^{ds}$$

to denote a prefix of \mathbf{y} , and

$$\mathbf{x}_{\mathbf{y}_{\leq s}} = \left(x_{\mathbf{y}_{\leq s}}^1, \dots, x_{\mathbf{y}_{\leq s}}^d \right) \in \{0, 1\}^d$$

to denote the members of $\mathbf{x}_{\mathbf{y}_{\leq s}}$.

The d -VCL trees used in this paper are a variant of the trees used in [Bousquet et al. \(2021\)](#). To distinguish the two, we call their construction *strong VCL trees*.

Definition 22 Let \mathcal{X} be a set, let $d \in \mathbb{N}$. An infinite strong VCL tree with respect to \mathcal{X} is a set

$$T = \{ \mathbf{x}_{\mathbf{u}} \in \mathcal{X}^{s+1} : s \in \mathbb{N} \cup \{0\} \wedge \mathbf{u} \in \{0, 1\}^1 \times \{0, 1\}^2 \times \dots \times \{0, 1\}^s \}.$$

Definition 23 Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let T be an infinite strong VCL tree as in Definition 22. We say that \mathcal{H} shatters T if for every $t \in \mathbb{N}$, and every $\mathbf{y} \in \{0, 1\}^1 \times \{0, 1\}^2 \times \cdots \times \{0, 1\}^t$ there exists a hypothesis $h \in \mathcal{H}$ that is consistent with $((\mathbf{x}_{\mathbf{y}_{\leq s-1}}, \mathbf{y}_s))_{s=1}^t$ in the sense that

$$\forall s \in [t] \forall j \in [s] : h(x_{\mathbf{y}_{\leq s-1}}^j) = y_s^j, \quad (5)$$

where we use the notation

$$\mathbf{y}_{\leq s} = \left((y_1^1), (y_2^1, y_2^2), (y_3^1, y_3^2, y_3^3), \dots, (y_s^1, \dots, y_s^s) \right) \in \{0, 1\}^{(\sum_{k=1}^s k)}$$

to denote a prefix of \mathbf{y} , and

$$\mathbf{x}_{\mathbf{y}_{\leq s}} = \left(x_{\mathbf{y}_{\leq s}}^1, \dots, x_{\mathbf{y}_{\leq s}}^{s+1} \right) \in \{0, 1\}^{s+1}$$

to denote the members of $\mathbf{x}_{\mathbf{y}_{\leq s}}$.

Definition 24 Let \mathcal{X} be a set and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$. The Vapnik–Chervonenkis–Littlestone dimension of \mathcal{H} , denoted $\text{VCL}(\mathcal{H})$, is the largest integer $d \geq 0$ such that \mathcal{H} shatters an infinite d -VCL tree. If \mathcal{H} does not shatter any infinite 1-VCL tree, we say that $\text{VCL}(\mathcal{H}) = 0$. If \mathcal{H} shatters infinite d -VCL trees for d arbitrarily large, we say that $\text{VCL}(\mathcal{H}) = \infty$.

A.3. Learning Rates

Bousquet et al. (2021) used the following definition of distribution-dependent learning rates.

Definition 25 (Bousquet et al. 2021, Definition 1.4.) Let \mathcal{H} be a concept class, and let $R : \mathbb{N} \rightarrow [0, 1]$ with $R(n) \rightarrow 0$ be a rate function.

- \mathcal{H} is learnable at rate R if there exists a learning algorithm \hat{h} such that for every $\mathcal{D} \in \text{Realizable}(\mathcal{H})$, there exist $C, c \geq 0$ such that $\mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0,1}(\hat{h}_S) \right] \leq C \cdot R(c \cdot n)$ for all $n \in \mathbb{N}$.
- \mathcal{H} is learnable with rate no faster than R if for every learning algorithm \hat{h} , there exists a $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ and $C, c > 0$ for which $\mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0,1}(\hat{h}_S) \right] \geq C \cdot R(c \cdot n)$ for infinitely many $n \in \mathbb{N}$.
- \mathcal{H} is learnable with optimal rate R if \mathcal{H} is learnable at rate R and \mathcal{H} is not learnable faster than R .
- \mathcal{H} requires arbitrarily slow rates if, for every $R(n) \rightarrow 0$, \mathcal{H} is learnable at rate no faster than R .

In this paper we refine the notion of learning rates, introducing the following more nuanced expressions for linear rates, as follows. Note that our definitions are strictly special cases in the sense that if a class is learnable at rate (learnable at rate no faster than) d/n according to our definition, then it is learnable at rate (learnable at rate no faster than) d/n according to Definition 25 as well.

Definition 26 Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let $d \geq 0$ and $\gamma \geq 1$. We say that:

- \mathcal{H} is learnable with fine-grained rate d/n if there exists a learning algorithm \hat{h} such that for any distribution $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ there exist real numbers $C, c \geq 0$ such that for all $n \in \mathbb{N}$:

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0.1}(\hat{h}_S) \right] \leq \frac{d}{n} + C \cdot \exp(-cn).$$

- \mathcal{H} is learnable with fine-grained rate no faster than d/n if for any learning algorithm \hat{h} there exists a distribution $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ such that the inequality

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0.1}(\hat{h}_S) \right] \geq \frac{d}{n}$$

holds for infinitely many $n \in \mathbb{N}$.

- \mathcal{H} is learnable with optimal fine-grained rate d/n with gap factor γ if \mathcal{H} is learnable with rate no faster than d/n , and is learnable with rate d'/n , where $d' \leq \gamma d$.

To distinguish the two notions of rate, we will refer to the rates of Definition 25 as *coarse rates*.

Remark 27 Ideally, we would like to obtain a gap factor γ that is as close as possible to 1, so that $d = d'$ (see Definition 26). The extent to which this is possible is a topic for further research. Throughout this paper we use $\gamma = 800$.

Definition 28 Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class. We say that \mathcal{H} is learnable with a strongly distribution-dependent linear rate if for any (possibly randomized) learning algorithm \hat{h} and any $c \geq 0$ there exists $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ such that the inequality

$$\frac{c}{n} \leq \mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0.1}(\hat{h}_S) \right] \tag{6}$$

holds for infinitely many $n \in \mathbb{N}$.

Remark 29 There are various technical issues related to measure theory that arise in the distribution-dependent learning setting and are germane to our results. We use the same assumptions as [Bousquet et al. \(2021\)](#), and refer the interested reader to their work for an in-depth discussion (e.g., Section 3.3 and Appendices B and C).

A.4. Gale–Stewart Games

We will use some basic concepts and results concerning infinite games. We refer the reader to Appendix A.1 of [Bousquet et al. \(2021\)](#) for additional references and discussion. Briefly, we consider infinite full information two-player games, in which there exists a set Ω and a subset $W \subseteq \Omega^{\mathbb{N}}$, and at each time $t = 1, 2, 3, \dots$, Player 1 selects an item $x_t \in \Omega$, and then Player 2 selects an item $y_t \in \Omega$. Player 1 wins if and only if the resulting infinite sequence $\mathbf{z} = (x_1, y_1, x_2, y_2, \dots)$ satisfies $\mathbf{z} \in W$; otherwise, Player 2 wins.

We say that Player i has a *winning strategy* if there exists a function $f : \Omega^* \rightarrow \Omega$ such that if in every time $t \in \mathbb{N}$, Player i selects item $f(\mathbf{z}')$ where \mathbf{z}' is the finite sequence of all items selected so far (by both players), then Player i wins the game (regardless of the selections made by the other player).

A game is called *determined* if precisely one of the players has a winning strategy. An infinite game is called *Gale–Stewart* (or *finitely-decidable*) if for every $\mathbf{w} \in W$ there exists $t \in \mathbb{N}$ such that for any

infinite suffix $\mathbf{s} \in \Omega^{\mathbb{N}}$, $\mathbf{w}_{\leq t} \circ \mathbf{s} \in W$, where ‘ \circ ’ denotes concatenation. Namely, every member of W has a finite prefix that certifies its membership in W . We will use the following result.

Theorem 30 (Gale and Stewart 1953) *Every Gale–Stewart game is determined.*

Appendix B. Proof of the Fine-Grained Characterization

B.1. Upper Bound

Throughout this section, let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let $d \in \mathbb{N}$.

Definition 31 *The online learning game for \mathcal{H} of size d , denoted $G_d^{\text{online}}(\mathcal{H})$, is an infinite full information game played between two players, a learner and an adversary. At each time step $t = 1, 2, 3, \dots$:*

1. *The adversary chooses $\mathbf{x}_t = (x_t^1, \dots, x_t^d) \in \mathcal{X}^d$.*
2. *The learner chooses $\mathbf{y}_t = (y_t^1, \dots, y_t^d) \in \{0, 1\}^d$.*

For each $t \in \mathbb{N}$, the version space is defined by

$$\mathcal{H}_t = \mathcal{H}_{\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_t, \mathbf{y}_t} = \{h \in \mathcal{H} : (\forall s \in [t] \forall i \in [d] : h(x_s^i) = y_s^i)\}.$$

If there exists a time step $t \in \mathbb{N}$ such that $\mathcal{H}_t = \emptyset$ then the learner wins the game. Otherwise, the adversary wins the game.

Definition 32 *The forbidden pattern game for \mathcal{H} of size d , denoted $G_{\mathcal{H}}^{\text{forbidden}}(d)$, is an infinite full information game played between two players, a learner and an adversary. At each time step $t = 1, 2, 3, \dots$:*

1. *The adversary chooses $\mathbf{x}_t \in \mathcal{X}^d$.*
2. *The learner chooses $\hat{\mathbf{y}}_t \in \{0, 1\}^d$.*
3. *The adversary chooses $\mathbf{y}_t \in \{0, 1\}^d$.*

The adversary wins the game if the adversary’s infinite sequence $((\mathbf{x}_t, \mathbf{y}_t))_{t \in \mathbb{N}}$ is consistent with \mathcal{H} and $\hat{\mathbf{y}}_t = \mathbf{y}_t$ for infinitely many $t \in \mathbb{N}$. Otherwise, the learner wins the game.

In other words, the learner wins the forbidden pattern game if $\hat{\mathbf{y}}_t$ is eventually a ‘forbidden pattern’ that is not consistent with \mathcal{H} .

We show that the existence of d -VCL trees characterizes the winner in the forbidden pattern game. Note that while the online game is a Gale–Stewart game, the forbidden pattern game is not. This makes the online game a convenient stepping stone towards this characterization, as in the following claim.

Lemma 33 *The following conditions are equivalent:*

1. *There does not exist an infinite d -VCL tree with respect to \mathcal{X} that is shattered by \mathcal{H} .*
2. *There exists a winning strategy for the learner in the online game $G_d^{\text{online}}(\mathcal{H})$.*
3. *There exists a winning strategy for the learner in the forbidden pattern game $G_d^{\text{forbidden}}(\mathcal{H})$.*

The proof of Lemma 33 is divided between Theorems 34 to 37.

Claim 34 *There exists an infinite d -VCL tree with respect to \mathcal{X} that is shattered by \mathcal{H} if and only if there exists a winning strategy for the adversary in $G_d^{\text{online}}(\mathcal{H})$.*

Proof Assume that there exists an infinite d -VCL tree

$$T = \{ \mathbf{x}_{\mathbf{u}} \in \mathcal{X}^d : \mathbf{u} \in (\{0, 1\}^d)^* \} \quad (7)$$

that is shattered by \mathcal{H} . This implies the existence of a winning strategy for the adversary as following. At each time step $t \in \mathbb{N}$, the adversary selects $\mathbf{x}_t = \mathbf{x}_{\mathbf{y}_{\leq t-1}}$. For any choice $\mathbf{y}_t \in \{0, 1\}^d$ made by the learner the version space remains not empty, i.e., $\mathcal{H}_t \neq \emptyset$. This holds because \mathcal{H} shatters T , and so in particular there exists a hypothesis $h \in \mathcal{H}$ that is consistent with $((\mathbf{x}_{\mathbf{y}_{\leq s-1}}, \mathbf{y}_s))_{s \in [t]}$. Hence, the adversary wins the game when playing according to this strategy.

Conversely, assume that there exists a winning strategy for the adversary defined by a function $f : \{0, 1\}^* \rightarrow \mathcal{X}^d$ such that at any time step $t \in \mathbb{N}$, the adversary chooses $\mathbf{x}_t = f(\mathbf{y}_{\leq t-1})$, where $\mathbf{y}_{\leq t-1}$ is the sequence of choices the learner has made so far. The function f defines an infinite d -VCL tree T as in Equation (7) given by $\mathbf{x}_{\mathbf{u}} = f(\mathbf{u})$. Seeing as this is a winning strategy for the adversary, $\mathcal{H}_t \neq \emptyset$ for all $t \in \mathbb{N}$ and all possible choices of $\mathbf{y}_{\leq t}$, and this implies that the tree T is shattered by \mathcal{H} . ■

Claim 35 *In the context of Lemma 33, Item 1 \iff Item 2.*

Proof

$$\begin{aligned} \text{(Item 1)} &\iff \left(\nexists \text{ winning strategy for the adversary in } G_d^{\text{online}}(\mathcal{H}) \right) \\ &\iff \text{(Item 2)}, \end{aligned}$$

where the first equivalence is by Theorem 34, and the second equivalence states that the online learning game is determined, which is true by Theorem 30 because it is a Gale–Stewart game. ■

Claim 36 *In the context of Lemma 33, Item 2 \implies Item 3.*

Proof Idea Use Algorithm 1. A winning strategy for the learner in the online game empties the version space. So eventually, for any \mathbf{x}_t chosen by the adversary, the learner can choose a \mathbf{y}_t such that $(\mathbf{x}, \mathbf{y})_{\leq t}$ is not consistent with \mathcal{H} . ■

Proof Let $f : (\bigcup_{s=1}^{\infty} \mathcal{X}^{ds}) \rightarrow \{0, 1\}^d$ be a function that defines a winning strategy for the learner in the online game $G_{\mathcal{H}}^{\text{online}}(d)$. Namely, in the online game, if in each time step $t \in \mathbb{N}$ the adversary chooses \mathbf{x}_t and the learner chooses $\mathbf{y}_t = f(\mathbf{x}_1, \dots, \mathbf{x}_t)$, then after a finite number of steps $\mathcal{H}_t = \emptyset$.

Given such a function f , Algorithm 1 defines a winning strategy for the learner in the forbidden pattern game. To see this, assume for contradiction that the strategy of Algorithm 1 is not a winning strategy for the learner, namely, assume that there exists a sequence $((\mathbf{x}_t, \mathbf{y}_t))_{t \in \mathbb{N}}$ that is consistent with \mathcal{H} and also $\hat{\mathbf{y}}_t = \mathbf{y}_t$ for infinitely many t when $\hat{\mathbf{y}}$ is chosen by the learner according to Algorithm 1. This implies that the sequences ξ and η defined by the algorithm are infinite.

Assumption: $f : (\bigcup_{s=1}^{\infty} \mathcal{X}^{ds}) \rightarrow \{0, 1\}^d$ is a function that defines a winning strategy for the learner in the online game $G_{\mathcal{H}}^{\text{online}}(d)$.

FORBIDDENPATTERNLEARNER:

```

 $\xi \leftarrow$  empty sequence
 $\eta \leftarrow$  empty sequence
for  $t \leftarrow 1, 2, \dots$  :
  Receive  $\mathbf{x}_t$  from the adversary
  Choose  $\hat{\mathbf{y}}_t \leftarrow f(\xi \circ \mathbf{x}_t)$ 
  Receive  $\mathbf{y}_t$  from the adversary
  if  $\hat{\mathbf{y}}_t = \mathbf{y}_t$ :
     $\xi \leftarrow \xi \circ \mathbf{x}_t$ 
     $\eta \leftarrow \eta \circ \mathbf{y}_t$ 

```

Algorithm 1: A reduction from a winning strategy for the forbidden pattern game to a winning strategy for the online game.

We show that if the adversary in the online game plays this infinite sequence ξ and the learner plays according to the strategy f , then the adversary wins the game, in contradiction to the assumption that f defines a winning strategy for the learner in the online game.

Let $\xi = \xi_1, \xi_2, \dots$ and $\eta = \eta_1, \eta_2, \dots$ where $\xi_t \in \mathcal{X}^d$ and $\eta_t \in \{0, 1\}^d$ for all $t \in \mathbb{N}$. By construction, $((\xi_t, \eta_t))_{t \in \mathbb{N}}$ is consistent with \mathcal{H} because it is a subsequence of $((\mathbf{x}_t, \mathbf{y}_t))_{t \in \mathbb{N}}$. In particular, for any finite prefix $(\xi, \eta)_{\leq t}$ there exists a hypothesis $h \in \mathcal{H}$ that is consistent with $(\xi, \eta)_{\leq t}$. This implies that for any $t \in \mathbb{N}$, the version space $\mathcal{H}_t = \mathcal{H}_{\xi_1, \eta_1, \dots, \xi_t, \eta_t}$ is not empty. However, the sequence $((\xi_t, \eta_t))_{t \in \mathbb{N}}$ is constructed by playing according to the strategy f , namely $\eta_t = f(\xi_1, \dots, \xi_t)$ for all $t \in \mathbb{N}$. We conclude that when the adversary in the online game plays ξ and the learner plays according to f , then $\mathcal{H}_t \neq \emptyset$ for all $t \in \mathbb{N}$, yielding the desired contradiction to the choice of f . ■

Claim 37 *In the context of Lemma 33, Item 3 \implies Item 1.*

Proof We show the contrapositive, namely, if there exists an infinite d -VCL tree shattered by \mathcal{H} then there does not exist a winning strategy for the learner in the forbidden pattern game $G_d^{\text{forbidden}}(\mathcal{H})$ (this is similar to one of the directions in Theorem 34). Indeed, let T be an infinite shattered tree as in Equation (7). Then there exists a winning strategy for the adversary: at each time step $t \in \mathbb{N}$, the adversary chooses $\mathbf{x}_t = \mathbf{x}_{\mathbf{y}_{\leq t-1}}$, and chooses $\mathbf{y}_t \in \{0, 1\}^d$ to be any value such that $\mathbf{y}_t \neq \hat{\mathbf{y}}_t$. Because the tree is shattered, for every $t \in \mathbb{N}$ and every possible $\mathbf{y}_t \in \{0, 1\}^d$ there exists $h \in \mathcal{H}$ that is consistent with $(\mathbf{x}, \mathbf{y})_{\leq t}$. Hence, the resulting sequence $((\mathbf{x}_t, \mathbf{y}_t))_{t \in \mathbb{N}}$ is consistent with \mathcal{H} while also satisfying $\mathbf{y}_t \neq \hat{\mathbf{y}}_t$ for all $t \in \mathbb{N}$, and therefore the adversary wins the game. ■

Notation 38 *Fix a function f as in Algorithm 1, and consider an execution of that algorithm using f in which the adversary plays the sequence $((\mathbf{x}_t, \mathbf{y}_t))_{t \in \mathbb{N}}$. For each $t \in \mathbb{N}$ let $\xi^{(t)}$ denote the value*

of ξ at the beginning of time step t . We write

$$\hat{\mathbf{y}}_t : \mathcal{X}^d \rightarrow \{0, 1\}^d$$

to denote the function given by

$$\hat{\mathbf{y}}_t(x) = \hat{\mathbf{y}}_{(\mathbf{x}, \mathbf{y})_{\leq t-1}}(x) = f(\xi^{(t)} \circ x)$$

that determines the learner's choice at time t , such that $\hat{\mathbf{y}}_t = \hat{\mathbf{y}}_t(\mathbf{x}_t)$ for all $t \in \mathbb{N}$.

Definition 39 Let $\mathcal{D} \in \Delta(\mathcal{X}^d \times \{0, 1\}^d)$ be a distribution, and let $g : \mathcal{X}^d \rightarrow \{0, 1\}^d$ be a function. The forbidden pattern loss of g with respect to \mathcal{D} is

$$\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(g) = \mathbb{P}_{(X, Y) \sim \mathcal{D}}[g(X) = Y] = 1 - \mathbb{L}_{\mathcal{D}}^{0-1}(g).$$

The forbidden pattern loss simply captures the learners objective in the forbidden pattern game, which is to avoid having $\hat{\mathbf{y}}_t = \mathbf{y}_t$.

Claim 40 Assume $\text{VCL}(\mathcal{H}) < \infty$. Let $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ be a distribution, and let $S = ((X_1, Y_1), (X_2, Y_2), \dots)$ be an infinite sequence of i.i.d. samples from \mathcal{D} . Consider an instance of the forbidden pattern game where the adversary plays the sequence S , and the learner plays according to the function $\hat{\mathbf{y}}_t = \hat{\mathbf{y}}_{S_{\leq t-1}}$ as in Theorem 38. Then the forbidden pattern loss satisfies

$$\lim_{t \rightarrow \infty} \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} \left[\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(\hat{\mathbf{y}}_t) > 0 \right] = 0.$$

Proof By the proof of Lemma 33 and the assumption that $\text{VCL}(\mathcal{H}) < \infty$, the strategy $\hat{\mathbf{y}}_t$ is a winning strategy for the learner in the forbidden pattern game.

First, assume that S is consistent with \mathcal{H} . Then there exists a random variable $T \in \mathbb{N}$ that depends on S , such that

$$\mathbb{P}[\forall t \geq T : \hat{\mathbf{y}}_t(X_t) \neq Y_t] = 1. \quad (8)$$

This is true because $\hat{\mathbf{y}}_t$ is a winning strategy for the learner. Furthermore, by construction of the strategy $\hat{\mathbf{y}}$, the function $\hat{\mathbf{y}}_t(x)$ only changes if the learner made a mistake, namely

$$\mathbb{P}[\forall t, t' \geq T \forall x \in \mathcal{X} : \hat{\mathbf{y}}_t(x) = \hat{\mathbf{y}}_{t'}(x)] = 1. \quad (9)$$

Hence,

$$\begin{aligned} & \lim_{t \rightarrow \infty} \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} \left[\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(\hat{\mathbf{y}}_t) = 0 \right] \\ &= \lim_{t \rightarrow \infty} \mathbb{P} \left[\left(\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{1}(\hat{\mathbf{y}}_t(X_{t+k}) = Y_{t+k}) \right) = 0 \right] \\ &\geq \lim_{t \rightarrow \infty} \mathbb{P} \left[\left(\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{1}(\hat{\mathbf{y}}_t(X_{t+k}) = Y_{t+k}) \right) = 0 \wedge t \geq T \right] \\ &= \lim_{t \rightarrow \infty} \mathbb{P} \left[\left(\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{1}(\hat{\mathbf{y}}_{t+k}(X_{t+k}) = Y_{t+k}) \right) = 0 \wedge t \geq T \right] \quad (\text{By Equation (9)}) \\ &= \lim_{t \rightarrow \infty} \mathbb{P}[t \geq T] = 1. \quad (\text{By Equation (8)}) \end{aligned}$$

So

$$\lim_{t \rightarrow \infty} \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} \left[\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(\hat{\mathbf{y}}_t) > 0 \right] = 1 - \lim_{t \rightarrow \infty} \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} \left[\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(\hat{\mathbf{y}}_t) = 0 \right] = 0$$

as desired.

It remains to show that $\mathbb{P}[S \text{ is consistent with } \mathcal{H}] = 1$. This is a consequence of the Borel–Cantelli lemma. Seeing as $\mathcal{D} \in \text{Realizable}(\mathcal{H})$, there exists a sequence $h_1, h_2, \dots \in \mathcal{H}$ such that $\mathbb{L}_{\mathcal{D}}^{0-1}(h_k) \leq 2^{-k}$ for all $k \in \mathbb{N}$. For every $t, k \in \mathbb{N}$ let $G_{t,k} = \{\forall i \in [t] : h_k(X_i) = Y_i\}$ be the event in which $S_{\leq t}$ is consistent with h_k . Then for every $t \in \mathbb{N}$,

$$\sum_{k \in \mathbb{N}} \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}}[-G_{t,k}] \leq \sum_{k \in \mathbb{N}} t \cdot \mathbb{L}_{\mathcal{D}}^{0-1}(h_k) \leq t < \infty.$$

By Borel–Cantelli, this implies that

$$\forall t \in \mathbb{N} : \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}}[\exists k \in \mathbb{N} : G_{t,k}] = 1.$$

In words, for every $t \in \mathbb{N}$, with probability 1 over the choice of S , there exists $k \in \mathbb{N}$ such that h_k is consistent with $S_{\leq t}$. Finally,

$$\mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}}[S \text{ is consistent with } \mathcal{H}] \geq \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} \left[\bigcap_{t \in \mathbb{N}} \{\exists k \in \mathbb{N} : G_{t,k}\} \right] = 1,$$

because a countable intersection of probability 1 events has probability 1. ■

Definition 41 *In the context of Theorem 40, let*

$$t^* = t^*(\mathcal{D}) = \inf \left(\left\{ t \in \mathbb{N} : \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} \left[\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(\hat{\mathbf{y}}_t) > 0 \right] \leq \frac{1}{8} \right\} \cup \{\infty\} \right).$$

The set of good sample sizes for \mathcal{D} is

$$\mathcal{T}_{\mathcal{D}}^{\text{good}} = \left\{ t \in [t^*] : \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} \left[\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(\hat{\mathbf{y}}_t) > 0 \right] \leq \frac{1}{4} \right\}.$$

Claim 42 *There exists a function $\hat{t} : (\mathcal{X} \times \{0, 1\})^* \rightarrow \mathbb{N}$ as follows. Let $\mathcal{D} \in \text{Realizable}(\mathcal{H})$. There exist parameters $C, c \geq 0$ such that for any $n \in \mathbb{N}$,*

$$\mathbb{P}_{S \sim \mathcal{D}^n} \left[\hat{t}(S) \in \mathcal{T}_{\mathcal{D}}^{\text{good}} \right] \geq 1 - Ce^{-cn}.$$

Proof Fix $\mathcal{D} \in \text{Realizable}(\mathcal{H})$. We show that Algorithm 2 satisfies the requirements of the claim. By Theorem 40, $t^* = t^*(\mathcal{D})$ is finite and $\mathcal{T}_{\mathcal{D}}^{\text{good}} \neq \emptyset$.

For each $t \in \mathbb{N}$ let $e_t = \mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} \left[\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(\hat{\mathbf{y}}_t) > 0 \right]$. Hoeffding’s inequality implies that there exist $C_t, c_t \geq 0$ such that $\mathbb{P}[|\hat{e}_t - e_t| > 1/16] \leq C_t \cdot e^{-c_t \cdot n}$. By a union bound,

$$\mathbb{P}_{S \sim \mathcal{D}^{\mathbb{N}}} [\exists t \in [t^*] : |\hat{e}_t - e_t| > 1/16] \leq \sum_{t \in [t^*]} C_t \cdot e^{-c_t \cdot n} \leq C' \cdot e^{-c' \cdot n}, \quad (10)$$

Assumption:

- $S = ((X_1, Y_1), \dots, (X_n, Y_n)) \sim \mathcal{D}^n$ is a labeled training set.
- $\hat{\mathbf{y}}_t = \hat{\mathbf{y}}_{S_{\leq t-1}}$ is as in Theorem 38.
- $m = \lfloor n/2 \rfloor$.

SAMPLESIZEESTIMATOR:

$S^{\text{train}}, S^{\text{test}} \leftarrow$ independent disjoint subsets of S of size m
for $t \in [m]$:
 $k \leftarrow \lfloor m/t \rfloor$
 $S_1^{\text{train}}, \dots, S_k^{\text{train}} \leftarrow$ independent disjoint subsets of S^{train} of size t
 for $i \in [k]$:
 $\hat{e}_{t,i} \leftarrow \mathbb{1} \left(\exists (X, Y) \in S^{\text{test}} : \hat{\mathbf{y}}_{S_i^{\text{train}}}(X) = Y \right)$
 $\hat{e}_t \leftarrow \frac{1}{k} \sum_{i \in [k]} \hat{e}_{t,i}$
 $\hat{t} \leftarrow \inf (\{t \in [m] : \hat{e}_t \leq 3/16\} \cup \{\infty\})$
output \hat{t}

Algorithm 2: An algorithm for finding \hat{t} such that with high probability, $\hat{t} \in \mathcal{T}_{\mathcal{D}}^{\text{good}}$.

for some suitable $C', c' \geq 0$.

Assume that $m \geq t^*$ and $\forall t \in [t^*] : |\hat{e}_t - e_t| \leq 1/16$. Then in particular, $\hat{e}_{t^*} \leq e_{t^*} + 1/16 \leq 1/8 + 1/16 = 3/16$, and therefore the output \hat{t} selected by Algorithm 2 satisfies $\hat{t} \leq t^*$. Additionally, the selected output satisfies $e_{\hat{t}} \leq \hat{e}_{\hat{t}} + 1/16 \leq 3/16 + 1/16 = 1/4$.

Combining the last paragraph with Equation (10), we conclude that there exist $C, c \geq 0$ such that with probability at least $1 - Ce^{-cn}$, \hat{t} satisfies $\hat{t} \leq t^*$ and $e_{t^*} \leq e_{\hat{t}} \leq 1/4$, so in particular $\hat{t} \in \mathcal{T}_{\mathcal{D}}^{\text{good}}$, as desired. ■

Theorem 43 (Haussler, Littlestone, and Warmuth 1994, Theorem 2.3) *Let $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class. There exists a function*

$$A : (\mathcal{X} \times \{0, 1\})^* \times \mathcal{X} \rightarrow \{0, 1\}$$

such that for any target function $f \in \mathcal{F}$, any $n \in \mathbb{N}$, and any $(x_1, \dots, x_n) \in \mathcal{X}^n$, A satisfies

$$\frac{1}{|S_n|} \sum_{\sigma \in S_n} L_{\sigma, f}(A) \leq \frac{\text{VC}(\mathcal{F})}{n},$$

where S_n is the set of all permutation functions $[n] \rightarrow [n]$, and $L_{\sigma, f}(A)$ is the 0-1 loss of A with respect to f and the permutation σ , namely,

$$L_{\sigma, f}(A) = \mathbb{1} \left(A(x_{\sigma(1)}, f(x_{\sigma(1)}), \dots, x_{\sigma(n-1)}, f(x_{\sigma(n-1)}), x_{\sigma(n)}) \neq f(x_{\sigma(n)}) \right).$$

Claim 44 For any pattern avoidance function $g : \mathcal{X}^d \rightarrow \{0, 1\}^d$ there exists a function A_g given by

$$A_g : (\mathcal{X} \times \{0, 1\})^* \rightarrow \{0, 1\}^{\mathcal{X}}$$

such that for any distribution $\mathcal{D} \in \Delta(\mathcal{X}^d \times \{0, 1\}^d)$ for which $L_{\mathcal{D}}^{\text{forbidden}}(g) = 0$ and for any $n \in \mathbb{N}$,

$$\mathbb{E}_{S \sim \mathcal{D}^n} [L_{\mathcal{D}}^{0-1}(A_g(S))] \leq \frac{d}{n}.$$

Proof This follows from Theorem 43, along with an appropriate definition of a VC class from g .

Let

$$\mathcal{F} = \left\{ f \in \{0, 1\}^{\mathcal{X}} : \left(\forall \mathbf{x} \in \mathcal{X}^d : (f(\mathbf{x}_1), \dots, f(\mathbf{x}_d)) \neq g(\mathbf{x}) \right) \right\}$$

be the set of all functions that avoid the pattern $g(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}^d$. Note that $\text{VC}(\mathcal{F}) \leq d$ because there does not exist a shattered subset of \mathcal{X} of cardinality d . Let A_g be the function A corresponding to \mathcal{F} whose existence is guaranteed by Theorem 43. Then

$$\begin{aligned} & \mathbb{E}_{S \sim \mathcal{D}^n} [L_{\mathcal{D}}^{0-1}(A_g(S))] \\ &= \mathbb{E}_{S \sim \mathcal{D}^n, (X, Y) \sim \mathcal{D}} [\mathbb{1}((A_g(S))(X) \neq Y)] \\ &= \mathbb{E}_{((X_1, Y_1), \dots, (X_{n+1}, Y_{n+1})) \sim \mathcal{D}^{n+1}, \sigma \sim \text{U}(S_{n+1})} [L_{\sigma, f}(A_g)] \tag{11} \\ &= \mathbb{E}_{((X_1, Y_1), \dots, (X_{n+1}, Y_{n+1})) \sim \mathcal{D}^{n+1}} \left[\frac{1}{|S_{n+1}|} \sum_{\sigma \in S_{n+1}} L_{\sigma, f}(A_g) \right] \\ &\leq \frac{\text{VC}(\mathcal{F})}{n+1} \leq \frac{d}{n+1}. \tag{By Theorem 43} \end{aligned}$$

In Equation (11), f a function in \mathcal{F} that is consistent with $(X_1, Y_1), \dots, (X_{n+1}, Y_{n+1})$. f is chosen deterministically as a function of $(X_1, Y_1), \dots, (X_{n+1}, Y_{n+1})$. Such an f exists because $L_{\mathcal{D}}^{\text{forbidden}}(g) = 0$, and \mathcal{F} contains all functions that avoid g . In Equation (11), we have used the fact that

$$(S, (X, Y)) \stackrel{d}{=} ((X_{\sigma(1)}, Y_{\sigma(1)}), \dots, (X_{\sigma(n+1)}, Y_{\sigma(n+1)})).$$

■

Lemma 45 If there exists a winning strategy for the learner in the forbidden pattern game $G_d^{\text{forbidden}}(\mathcal{H})$, then \mathcal{H} is learnable with rate d/n .

Proof of Lemma 45 Let $\mathcal{D} \in \text{Realizable}(\mathcal{H})$. We need to show that there exist $C, c \geq 0$ as follows. For any $n \in \mathbb{N}$, let $\hat{h}_S = \text{OPTIMALRATELEARNER}(S)$ with $S \sim \mathcal{D}^n$. Then

$$\mathbb{E}_{S \sim \mathcal{D}^n} [L_{\mathcal{D}}^{0-1}(\hat{h}_S)] \leq d/n + C e^{-cn}.$$

This is established via the following analysis of Algorithm 3. By Theorem 42, there exist $C_0, c_0 \geq 0$ such that

$$\mathbb{P}_{S \sim \mathcal{D}^n} [\hat{t} \in \mathcal{T}_{\mathcal{D}}^{\text{good}}] \geq 1 - C_0 \cdot e^{-c_0 \cdot n}. \tag{12}$$

Assumptions:

- $n \in \mathbb{N}$, $m = \lfloor \frac{n}{2} \rfloor$.
- $\mathcal{D} \in \text{Realizable}(\mathcal{H})$.
- $S = ((X_1, Y_1), \dots, (X_n, Y_n)) \sim \mathcal{D}^n$ is a labeled training set.
- x is the input that needs to be labeled.
- \hat{t} is a function as in Theorem 42.
- For any sequence $\mathbf{z} \in (\mathcal{X}^d \times \{0, 1\}^d)^*$, $\hat{\mathbf{y}}_{\mathbf{z}} : \mathcal{X}^d \rightarrow \{0, 1\}^d$ is a pattern avoidance function as in Theorem 38.
- A_g is a learning algorithm that uses pattern avoidance function g , as in Theorem 44.

OPTIMALRATELEARNER(S):

```

 $\hat{t} \leftarrow \hat{t}(S)$ 
 $k \leftarrow \lfloor m/\hat{t} \rfloor$ 
 $S_g, S_a \leftarrow$  partition of  $S$  into two disjoint sets of size at least  $m$ 
 $S_1, \dots, S_k \leftarrow$  partition of  $S_g$  into  $k$  disjoint sets of size at least  $\hat{t}$ 
for  $i \in [k]$ :
     $g_i \leftarrow \hat{\mathbf{y}}_{S_i}$ 
     $a_i \leftarrow A_{g_i}(S_a)$ 
 $\hat{h} \leftarrow \left( x \mapsto \text{Majority}(a_1(x), \dots, a_k(x)) \right)$   $\triangleright$  Defining a function  $\hat{h} : \mathcal{X} \rightarrow \{0, 1\}$ 
output  $\hat{h}$ 
    
```

Algorithm 3: An algorithm that achieves the optimal learning rate for any class with finite VCL dimension.

From the definition of $\mathcal{T}_{\mathcal{D}}^{\text{good}}$, if $\hat{t} \in \mathcal{T}_{\mathcal{D}}^{\text{good}}$ then for every $i \in [k]$,

$$\mathbb{P}_{S \sim \mathcal{D}^n} \left[\mathbf{L}_{\mathcal{D}}^{\text{forbidden}}(g_i) > 0 \right] \leq \frac{1}{4}.$$

By Hoeffding's inequality, there exist $C_1, c_1 \geq 0$ such that

$$\mathbb{P}_{S \sim \mathcal{D}^n} \left[\frac{|\{i \in [k] : \mathbf{L}_{\mathcal{D}}^{\text{forbidden}}(g_i) > 0\}|}{k} \geq \frac{3}{8} \mid \hat{t} \in \mathcal{T}_{\mathcal{D}}^{\text{good}} \right] \leq C_1 \cdot e^{-c_1 \cdot n}, \quad (13)$$

where we have used the fact that if $\hat{t} \in \mathcal{T}_{\mathcal{D}}^{\text{good}}$ then $k = \Omega(n)$. Applying the inequality $\mathbb{P}[E] \leq \mathbb{P}[E|F] + \mathbb{P}[\neg F]$ to Equations (12) and (13) implies that there exist $C, c \geq 0$ such that

$$\mathbb{P}_{S \sim \mathcal{D}^n} \left[\frac{|\{i \in [k] : \mathbf{L}_{\mathcal{D}}^{\text{forbidden}}(g_i) > 0\}|}{k} \geq \frac{3}{8} \right] \leq C e^{-cn}, \quad (14)$$

From Theorem 44, for any $i \in [k]$, if $\mathbf{L}_{\mathcal{D}}^{\text{forbidden}}(g_i) = 0$ then

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbf{L}_{\mathcal{D}}^{0-1}(A_{g_i}(S)) \right] \leq \frac{d}{n}. \quad (15)$$

Let B be the bad event whose probability is bounded by Equation (14). Then

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S) \right] &= \mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S) \cdot \mathbf{1}(B) \right] + \mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S) \cdot \mathbf{1}(\neg B) \right] \\ &\leq \mathbb{P}_{S \sim \mathcal{D}^n}[B] + \mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S) \cdot \mathbf{1}(\neg B) \right] \\ &\leq C e^{-cn} + \mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S) \cdot \mathbf{1}(\neg B) \right], \end{aligned} \quad (16)$$

where the final inequality follows by Equation (14).

The expectation in the previous line can be bounded by

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S) \cdot \mathbf{1}(\neg B) \right] &= \mathbb{E}_{S \sim \mathcal{D}^n, (X, Y) \sim \mathcal{D}} \left[\mathbf{1}(\text{Majority}(a_1(X), \dots, a_k(X)) \neq Y) \mathbf{1}(\neg B) \right] \\ &\leq \mathbb{P}_{S \sim \mathcal{D}^n, (X, Y) \sim \mathcal{D}} [\text{Majority}(a_1(X), \dots, a_k(X)) \neq Y \wedge \neg B] \\ &\leq \mathbb{P} \left[\left(\frac{|\{i : a_i(X) \neq Y\}|}{k} \geq \frac{1}{2} \right) \wedge \left(\frac{|\{i : \mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(g_i) = 0\}|}{k} \geq \frac{5}{8} \right) \right] \\ &\leq \mathbb{P} \left[\left(\frac{|\{i : (a_i(X) \neq Y) \wedge (\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(g_i) = 0)\}|}{k} \geq \frac{1}{8} \right) \right] \\ &\leq \frac{8}{k} \cdot \mathbb{E} \left[|\{i : (a_i(X) \neq Y) \wedge (\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(g_i) = 0)\}| \right] \quad (\text{Markov's inequality}) \\ &= \frac{8}{k} \cdot \sum_{i \in [k]} \mathbb{P} \left[(a_i(X) \neq Y) \wedge (\mathbb{L}_{\mathcal{D}}^{\text{forbidden}}(g_i) = 0) \right] \\ &\leq \frac{8}{k} \cdot \sum_{i \in [k]} \frac{d}{n} = \frac{8d}{n}, \end{aligned} \quad (17)$$

where the last inequality follows from Equation (15).

Finally, plugging the bound of Equation (Markov's inequality) into Equation (16) yields

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S) \right] \leq C e^{-cn} + \frac{8d}{n},$$

as desired. ■

B.2. Lower Bound

Lemma 46 *For any set \mathcal{X} and any hypothesis class $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ satisfying $d = \text{VCL}(\mathcal{H})$ with $1 \leq d < \infty$, there exists a distribution $\mathcal{D}_{\mathcal{X}} \in \Delta(\mathcal{X})$ such that for any (possibly randomized) learning algorithm \hat{h} there exists $\mathcal{D} \in \text{Realizable}(\mathcal{H})$ such that the marginal distribution of \mathcal{D} on \mathcal{X} is $\mathcal{D}_{\mathcal{X}}$, and the inequality*

$$\mathbb{E}_{S \sim \mathcal{D}^n} \left[\mathbb{L}_{\mathcal{D}}^{0-1}(\hat{h}_S) \right] \geq \frac{d}{100 \cdot n} \quad (18)$$

holds for infinitely many $n \in \mathbb{N}$.

B.2.1. INGREDIENTS

The proof employs a claim about *indifferent* d -VCL trees, which is proved using a simple lemma from Ramsey theory.

Notation 47 For any $\mathbf{u} \in (\{0, 1\}^d)^*$, let $\text{index}(\mathbf{u}) \in \mathbb{N}$ denote the index of \mathbf{u} in the lexicographical ordering of $(\{0, 1\}^d)^*$.

Definition 48 Let $d \in \mathbb{N}$, let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let

$$T = \left\{ \mathbf{x}_{\mathbf{u}} \in \mathcal{X}^d : \mathbf{u} \in (\{0, 1\}^d)^* \right\}$$

be an infinite d -VCL tree that is shattered by \mathcal{H} . Recall that this implies the existence of a collection

$$\mathcal{H}_T = \left\{ h_{\mathbf{u}} \in \mathcal{H} : \mathbf{u} \in (\{0, 1\}^d)^* \right\}$$

of consistent functions, namely, for each $\mathbf{u} \in (\{0, 1\}^d)^*$, $h_{\mathbf{u}}$ is consistent with the path from the root to node \mathbf{u} , as in the definition of shattering a VCL tree (Definition 21).

We say that such a collection \mathcal{H}_T is indifferent if for every $\mathbf{v}, \mathbf{u}, \mathbf{w} \in (\{0, 1\}^d)^*$, if $\text{index}(\mathbf{v}) < \text{index}(\mathbf{u})$, and \mathbf{w} is a descendant of \mathbf{u} in the tree T , then $h_{\mathbf{u}}(\mathbf{x}_{\mathbf{v}}^j) = h_{\mathbf{w}}(\mathbf{x}_{\mathbf{v}}^j)$ for every $j \in [d]$. In words, the functions for all the descendants of a node that appears after \mathbf{v} agree on \mathbf{v} .

We say that T is indifferent if it has a set \mathcal{H}_T of consistent functions that are indifferent.

Intuitively, if T is indifferent, then the labels for a node \mathbf{v} provide no information on the labels of a node \mathbf{u} that appears after \mathbf{v} in the lexicographical order.

The claim about indifferent d -VCL trees is as follows.

Claim 49 Let $d \in \mathbb{N}$, let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let T be an infinite d -VCL tree that is shattered by \mathcal{H} . Then there exists an infinite d -VCL tree T' that is shattered by \mathcal{H} that is indifferent.

Following is the lemma from Ramsey theory used for proving Theorem 49, and a generalized notion of a trees and subtrees used in that lemma.

Definition 50 Let (X, \preceq) be a partial order relation. For $a, b \in X$, we say that b is a child of a if $a \preceq b$ and there does not exist $c \in X$ such that $a \preceq c \preceq b$. For $k \in \mathbb{N}$, we say that (X, \preceq) is an infinite k -ary tree if every $a \in X$ has precisely k distinct children. We say that a partial order (X', \preceq') is a subtree of (X, \preceq) if $X' \subseteq X$, and $\forall a, b \in X' : a \preceq' b \iff a \preceq b$.

Lemma 51 Let $T = (X, \preceq)$ be an infinite k -ary tree, and let $g : X \rightarrow \{0, 1\}$ be a two-coloring of T . Then T has a monochromatic infinite k -ary subtree $T' = (X', \preceq')$, namely there exists T' such that T' is a subtree of T , T' is an infinite k -ary tree, and $|g(X')| = |\{g(a) : a \in X'\}| = 1$.

Proof of Lemma 51 If there exists $a \in X$ such that the set X' consisting of a and all its descendants satisfies $g(X') = \{1\}$, then we are done (take T' to be the subtree consisting of a and all its descendants). Otherwise, every $a \in X$ has a descendant $b \in X$ such that $g(b) = 0$. This implies that one can construct an infinite k -ary subtree that is 0-monochromatic using the following recursive procedure. Let r be any member of X such that $g(r) = 0$. Let T' be an empty tree, and add r to T' .

Subsequently, for each node n added to T' (including r), for each child a of n , add to T' an arbitrary descendant b of a such that $g(b) = 0$. \blacksquare

Proof of Claim 49 First, observe that if $T = \{x_{\mathbf{u}} : \mathbf{u} \in (\{0, 1\}^d)^*\}$ is an infinite d -VCL tree that is shattered by \mathcal{H} with a collection $\{h_{\mathbf{u}} : \mathbf{u} \in (\{0, 1\}^d)^*\}$ of consistent functions, then for any $x \in \mathcal{X}$ there exists an infinite d -VCL tree that is shattered by \mathcal{H} that is a subtree of T and has a collection of consistent functions that agree on x . Indeed, this follows from Lemma 51 by choosing a two-coloring $g : (\{0, 1\}^d)^* \rightarrow \{0, 1\}$ of T given by $g(\mathbf{u}) = h_{\mathbf{u}}(x)$.

Second, we use the above observation to construct an infinite d -VCL tree $T' = \{x'_{\mathbf{u}} : \mathbf{u} \in (\{0, 1\}^d)^*\}$ that is shattered by \mathcal{H} and is indifferent. The construction works by starting with $T' := T$ and then repeatedly modifying T' , as specified in Algorithm 4. Each modification step replaces a subtree $T'_{\mathbf{u}}$ of T' with one of its own infinite d -VCL subtrees, which is obtained by invoking the above observation on $T'_{\mathbf{u}}$ and $x = \mathbf{x}'_{\mathbf{v}}$ for some $j \in [d]$ and some \mathbf{v} that precedes \mathbf{u} . In each step, the set of nodes of T' decreases (is replaced by one of its subsets), and the collection of consistent functions can be decreased in a corresponding manner (be replaced by a subset of itself that corresponds to the new set of nodes).

```

 $T' \leftarrow T$ 
for  $\mathbf{u} \in (\{0, 1\}^d)^*$  in lexicographic order:
  for  $\mathbf{v} \in (\{0, 1\}^d)^*$  such that  $\text{index}(\mathbf{v}) < \text{index}(\mathbf{u})$ :
    for  $j \in [d]$ :
      replace  $T'_{\mathbf{u}}$  with an infinite  $2^d$ -ary subtree of  $T'_{\mathbf{u}}$  that has
        a collection of consistent functions that agree on  $\mathbf{x}'_{\mathbf{v}}$ 

```

Algorithm 4: Construction of an indifferent d -VCL tree. ($T'_{\mathbf{u}}$ denotes the infinite 2^d -ary subtree of T' rooted at node \mathbf{u} .)

Algorithm 4 never terminates, but it defines an infinite d -VCL tree T' . T' is well-defined because for every $\mathbf{r} \in (\{0, 1\}^d)^*$, the value of $\mathbf{x}'_{\mathbf{r}}$ never changes after the outer loop advances past \mathbf{r} (i.e., $\text{index}(\mathbf{u}) > \text{index}(\mathbf{r})$), and so $\mathbf{x}'_{\mathbf{r}}$ is eventually fixed. T' is an infinite 2^d -ary subtree of T (each replacement maintains that T' is an infinite 2^d -ary subtree of T , so the resulting tree defined by this process is also an infinite 2^d -ary subtree of T). This implies that it is a d -VCL tree that is shattered by \mathcal{H} . T' is indifferent by construction, because for each $\mathbf{q}, \mathbf{r}, \mathbf{s} \in (\{0, 1\}^d)^*$ and $k \in [d]$, if $\text{index}(\mathbf{q}) < \text{index}(\mathbf{r})$, and \mathbf{s} is a descendant of \mathbf{r} , then during the iteration of the innermost loop in which $\mathbf{u} = \mathbf{r}$, $\mathbf{v} = \mathbf{q}$, and $j = k$, the subtree $T'_{\mathbf{r}}$ was replaced with a subtree that has a collection of consistent functions that agree on $(\mathbf{x}'_{\mathbf{q}})^k$. In particular this implies that $h_{\mathbf{r}}((\mathbf{x}'_{\mathbf{q}})^k) = h_{\mathbf{s}}((\mathbf{x}'_{\mathbf{q}})^k)$. This agreement continues to hold from that point onwards, because the collection of consistent functions for descendants of $h_{\mathbf{r}}$ can only decrease at each step. \blacksquare

When a tree is indifferent, it admits a notion of *branch functions*, as follows.

Notation 52 $\mathcal{Y} = (\{0, 1\}^d)^{\mathbb{N}}$.

Definition 53 Let $d \in \mathbb{N}$, let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, and let

$$T = \left\{ \mathbf{x}_{\mathbf{u}} \in \mathcal{X}^d : \mathbf{u} \in \left(\{0, 1\}^d \right)^* \right\}$$

be an infinite d -VCL tree that is shattered by \mathcal{H} with a collection

$$\mathcal{H}_T = \left\{ h_{\mathbf{u}} \in \mathcal{H} : \mathbf{u} \in \left(\{0, 1\}^d \right)^* \right\}$$

of consistent functions that are indifferent. Let

$$\mathcal{X}_T = \left\{ \mathbf{x}_{\mathbf{u}}^i : \mathbf{u} \in \left(\{0, 1\}^d \right)^* \wedge i \in [d] \right\}.$$

For every $\mathbf{y} \in \mathcal{Y}$, the branch function for \mathbf{y} is the unique function $f_{\mathbf{y}} : \mathcal{X}_T \rightarrow \{0, 1\}$ such that for each $\mathbf{v} \in \left(\{0, 1\}^d \right)^*$ and $j \in [d]$,

$$f_{\mathbf{y}}(\mathbf{x}_{\mathbf{v}}^j) = h_{\mathbf{u}}(\mathbf{x}_{\mathbf{v}}^j)$$

for a node \mathbf{u} such that $\mathbf{y}_{\leq |\mathbf{u}|} = \mathbf{u}$ and $\text{index}(\mathbf{u}) > \text{index}(\mathbf{v})$. In words, $f_{\mathbf{y}}(\mathbf{x}_{\mathbf{v}}^j)$ is the value assigned to $\mathbf{x}_{\mathbf{v}}^j$ by the consistent function of any node on the infinite branch \mathbf{y} that appears after \mathbf{v} in lexicographic order. (Due to the indifference property, $h_{\mathbf{u}}(\mathbf{x}_{\mathbf{v}}^j)$ is the same for any such node \mathbf{u} .)

We note some consequences of the definitions of indifference and branch functions.

Claim 54 Let T be an indifferent infinite d -VCL tree with a collection of branch functions $\{f_{\mathbf{y}}\}_{\mathbf{y} \in \mathcal{Y}}$. Then:

1. Every branch function $f_{\mathbf{y}}$ is finitely realizable, meaning that for any finite set $\{x_1, \dots, x_m\} \subseteq \mathcal{X}_T$, there exists a function $h \in \mathcal{H}$ such that for all $i \in [m]$, $f_{\mathbf{y}}(x_i) = h(x_i)$.
2. Each element in T is unique. Namely, for every $\mathbf{u}, \mathbf{v} \in \left(\{0, 1\}^d \right)^*$ and every $i, j \in [d]$, if $\mathbf{u} \neq \mathbf{v}$ or $i \neq j$ then $\mathbf{x}_{\mathbf{u}}^i \neq \mathbf{x}_{\mathbf{v}}^j$.
3. Let $\mathbf{v}, \mathbf{u} \in \left(\{0, 1\}^d \right)^*$. If $\text{index}(\mathbf{u}) > \text{index}(\mathbf{v})$ then there exists $\mathbf{b} \in \{0, 1\}^d$ such that for any $\mathbf{y} \in \mathcal{Y}$, if $\mathbf{u} = \mathbf{y}_{\leq |\mathbf{u}|}$ then $f_{\mathbf{y}}(\mathbf{x}_{\mathbf{v}}^j) = \mathbf{b}_j$ for all $j \in [d]$. In words, if \mathbf{v} precedes \mathbf{u} in lexicographical order, then all the branch functions for branches that pass through node \mathbf{u} agree on node \mathbf{v} .

We think of Item 3 as an indifference property for branch functions. Intuitively, it means that knowing the labels for \mathbf{v} does not provide any information on which of the branch functions for branches that pass through \mathbf{u} is more likely to be the correct labeling function. The branch functions that pass through \mathbf{u} are *indifferent* to the labels of \mathbf{v} .

Proof of Claim 54 Item 1 is immediate from the definition of $f_{\mathbf{y}}$. For Item 2, clearly if $\mathbf{u} = \mathbf{v}$ then $\mathbf{x}_{\mathbf{u}}^i \neq \mathbf{x}_{\mathbf{v}}^j$, since otherwise node \mathbf{u} could not have 2^d children, in contradiction to T being a d -VCL tree. Assume for contradiction that $\text{index}(\mathbf{v}) < \text{index}(\mathbf{u})$ and $\mathbf{x}_{\mathbf{u}}^i = \mathbf{x}_{\mathbf{v}}^j$. Then all consistent functions for the children of \mathbf{u} must agree on $\mathbf{x}_{\mathbf{v}}^j$, but that implies that they agree on $\mathbf{x}_{\mathbf{u}}^i$ as well, which is again a contradiction to \mathbf{u} having 2^d children. Finally, Item 3 is immediate from the definition of $f_{\mathbf{y}}$ and from the indifference of T . \blacksquare

The proof of the lower bound also employs the reverse Fatou's lemma.

Lemma 55 (Reverse Fatou; e.g., Browder 1996, Theorem 10.17⁶) *Let $(\Omega, \mathcal{F}, \mu)$ be a measure space. Let $g : \Omega \rightarrow \mathbb{R}$ be a non-negative measurable function such that $\int_{\Omega} g d\mu < \infty$. For each $n \in \mathbb{N}$ let $f_n : \Omega \rightarrow \mathbb{R}$ be a measurable function such that $\forall \omega \in \Omega : f_n(\omega) \leq g(\omega)$. Then*

$$\int_{\Omega} \limsup_{n \rightarrow \infty} f_n d\mu \geq \limsup_{n \rightarrow \infty} \int_{\Omega} f_n d\mu.$$

B.2.2. PROOF OF LOWER BOUND

Proof of Lemma 46 We will define a set of distributions

$$\{P_{\mathbf{y}}\}_{\mathbf{y} \in \mathcal{Y}} \subseteq \text{Realizable}(\mathcal{H})$$

that depends on \mathcal{H} such that all the distributions in the set have the same marginal distribution over \mathcal{X} . The proof uses the probabilistic method to show that for every learning algorithm \hat{h} for \mathcal{H} there exists $\mathbf{y}^* \in \mathcal{Y}$ (that depends on \hat{h}) such that $P_{\mathbf{y}^*}$ is a hard distribution for \hat{h} , namely, that Equation (18) holds for $\mathcal{D} = P_{\mathbf{y}^*}$ for infinitely many values of n .

The set $\{P_{\mathbf{y}}\}_{\mathbf{y} \in \mathcal{Y}}$ is defined as follows. By Theorem 49 and the assumption that $\text{VCL}(\mathcal{H}) = d$, there exist an indifferent infinite d -VCL tree

$$T = \left\{ \mathbf{x}_{\mathbf{u}} \in \mathcal{X}^d : \mathbf{u} \in \left(\{0, 1\}^d \right)^* \right\}$$

with a corresponding collection of branch functions

$$\mathcal{F} = \{f_{\mathbf{y}} \in \{0, 1\}^{\mathcal{X}^d} : \mathbf{y} \in \mathcal{Y}\}.$$

Fix such a pair (T, \mathcal{F}) . For each $\mathbf{y} \in \mathcal{Y}$ let

$$P_{\mathbf{y}}((x, y)) = \sum_{\mathbf{u} \in (\{0, 1\}^d)^*} (d-1)d^{-\text{index}(\mathbf{u})-1} \sum_{i=1}^d \mathbb{1}(x = \mathbf{x}_{\mathbf{u}}^i \wedge y = f_{\mathbf{y}}(\mathbf{x}_{\mathbf{u}}^i)).$$

In words, $P_{\mathbf{y}}$ corresponds to the following sampling procedure:

1. Sample an index $k \in \mathbb{N}$ such that $\forall s \in \mathbb{N} : \mathbb{P}[k = s] = (d-1)d^{-s}$.⁷
2. Let $\mathbf{u} \in (\{0, 1\}^d)^*$ be the k -th string in the lexicographical ordering of $(\{0, 1\}^d)^*$.
3. Sample $j \in [d]$ independently and uniformly at random.
4. Output $\left(\mathbf{x}_{\mathbf{u}}^j, f_{\mathbf{y}}(\mathbf{x}_{\mathbf{u}}^j) \right)$.

Note that the marginal distribution of $P_{\mathbf{y}}$ on \mathcal{X} (the distribution of $\mathbf{x}_{\mathbf{u}}^i$) is the same for all $\mathbf{y} \in \mathcal{Y}$; this is the marginal distribution $\mathcal{D}_{\mathcal{X}}$ mentioned in the statement.

To see that $P_{\mathbf{y}}$ is realizable, note that for every $\varepsilon > 0$ there exists $k_{\varepsilon} \in \mathbb{N}$ such that in Step 1 of the sampling procedure, $\mathbb{P}[k > k_{\varepsilon}] \leq \varepsilon$. $f_{\mathbf{y}}$ is finitely-realizable by \mathcal{H} (Item 1 in Theorem 54), so in

6. See also “Reverse Fatou’s Lemma” in ProofWiki (2022).

7. Recall that for a geometric series, $\sum_{s=1}^{\infty} d^{-s} = \frac{1}{d-1}$ when $d > 1$, and therefore $\sum_{s=1}^{\infty} \mathbb{P}[k = s] = \sum_{s=1}^{\infty} (d-1)d^{-s} = 1$.

particular there exists $h_\varepsilon \in \mathcal{H}$ that is consistent with $Z_\varepsilon = \{(\mathbf{x}_u^j, f_y(\mathbf{x}_u^j)) : \text{index}(\mathbf{u}) \leq k_\varepsilon \wedge j \in [d]\}$. Hence, $L_{P_y}^{0,1}(h_\varepsilon) \leq \mathbb{P}_{(x,y) \sim P_y}[(x,y) \notin Z_\varepsilon] \leq \mathbb{P}[k > k_\varepsilon] \leq \varepsilon$.

For a fixed algorithm \hat{h} and for each $n \in \mathbb{N}$, consider the following experiment:

- A value $\mathbf{y} \in \mathcal{Y}$ is sampled from the uniform distribution $U(\mathcal{Y})$, namely each bit in \mathbf{y} is chosen independently and uniformly at random.
- An i.i.d. training set $S = ((X_1, Y_1, K_1), (X_2, Y_2, K_2), \dots, (X_n, Y_n, K_n)) \sim P_y^n$ is generated according to the sampling procedure of Steps 1 to 4, where for each $i \in [n]$, $K_i \in \mathbb{N}$ is the index selected at Step 1, and (X_i, Y_i) is the output at Step 4.
- An additional test sample $(X, Y, K) \sim P_y$ is generated in the same manner.
- A randomness value ρ is sampled for the algorithm \hat{h} , and then \hat{h} is executed with training set S and randomness ρ and produces a hypothesis \hat{h}_S .
- \hat{h}_S is used to predict a label $\hat{h}_S(X)$ for X .

This experiment defines a joint distribution

$$(\mathbf{y}, S, X, Y, K, \rho) \tag{19}$$

that is used throughout the remainder of the proof.

For any $\kappa \in \mathbb{N}$, let $G(\kappa)$ denote the event in which the following conditions hold:

- $K = \kappa \geq \max\{K_1, \dots, K_n\}$.
- $|\{i \in [n] : K_i = \kappa\}| < d/2$.
- $X \notin \{X_i : i \in [n]\}$.

We make two observations concerning $G(\kappa)$. The first observation is that

$$\mathbb{P}[G(\kappa)] \geq (d-1)d^{-\kappa}/4 \tag{20}$$

when $n = n_\kappa = \lfloor \frac{d^{\kappa+1}}{8(d-1)} \rfloor$. To see this, let

$$C_{=\kappa} = |\{i \in [n_\kappa] : K_i = \kappa\}|, \quad C_{>\kappa} = |\{i \in [n_\kappa] : K_i > \kappa\}|.$$

Then

$$\mathbb{E}[C_{=\kappa}] = n_\kappa \cdot (d-1)d^{-\kappa} \leq \frac{d^{\kappa+1}}{8(d-1)} \cdot (d-1)d^{-\kappa} = \frac{d}{8},$$

and

$$\begin{aligned} \mathbb{E}[C_{>\kappa}] &= n_\kappa \cdot \sum_{s=\kappa+1}^{\infty} (d-1)d^{-s} \\ &\leq \frac{d^{\kappa+1}}{8(d-1)} \cdot 2(d-1)d^{-\kappa-1} = \frac{1}{4}. \end{aligned}$$

By Markov's inequality,

$$\mathbb{P}\left[C_{=\kappa} \geq \frac{d}{2}\right] \leq \frac{1}{4}, \quad \text{and} \quad \mathbb{P}[C_{>\kappa} \geq 1] \leq \frac{1}{4}.$$

By a union bound,

$$\mathbb{P}\left[C_{=\kappa} < \frac{d}{2} \wedge C_{>\kappa} = 0\right] \geq \frac{1}{2}. \quad (21)$$

Hence, for any $\boldsymbol{\eta} \in \mathcal{Y}$,

$$\begin{aligned} & \mathbb{P}_{\mathbf{y}, S, X, Y, K}[G(\kappa) \mid \mathbf{y} = \boldsymbol{\eta}] \\ &= \mathbb{P}[K = \kappa \mid \mathbf{y} = \boldsymbol{\eta}] \cdot \mathbb{P}\left[C_{=\kappa} < \frac{d}{2} \wedge C_{>\kappa} = 0 \mid \mathbf{y} = \boldsymbol{\eta}\right] \\ & \quad \cdot \mathbb{P}\left[X \notin \{X_i : i \in [n_\kappa]\} \mid C_{=\kappa} < \frac{d}{2} \wedge C_{>\kappa} = 0 \wedge K = \kappa \wedge \mathbf{y} = \boldsymbol{\eta}\right] \\ & \geq (d-1)d^{-\kappa} \cdot \frac{1}{2} \\ & \quad \cdot \mathbb{P}\left[X \notin \{X_i : i \in [n_\kappa]\} \mid C_{=\kappa} < \frac{d}{2} \wedge C_{>\kappa} = 0 \wedge K = \kappa \wedge \mathbf{y} = \boldsymbol{\eta}\right] \\ & \hspace{15em} \text{(By Equation (21), } (C_{=\kappa}, C_{>\kappa}) \perp \mathbf{y}) \\ & \geq (d-1)d^{-\kappa} \cdot \frac{1}{2} \cdot \frac{1}{2}. \end{aligned}$$

For the last inequality, recall that the elements in T are unique (Item 2 in Theorem 54). Consequently, for every $i \in [n_\kappa]$, if $K_i < \kappa = K$ then $X_i \neq X$. The conditions $C_{=\kappa} < d/2$ and $K = \kappa$, and the sampling of $j \sim U([d])$ in Step 3 imply that with probability at least $1/2$, $X \notin \{X_i : i \in [n_\kappa] \wedge K_i = \kappa\}$. This establishes Equation (20), which is our first observation about $G(\kappa)$.

Our second observation is that for any κ corresponding to a node on the branch \mathbf{y} , if $G(\kappa)$ occurs then \hat{h} makes an incorrect prediction with probability $1/2$.

Formally, for any $t \in \mathbb{N}$, let $\kappa_{\mathbf{y}, t} = \text{index}(\mathbf{y}_{<t})$, where $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots)$ and $\mathbf{y}_{<t} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{t-1})$. In words, $\kappa_{\mathbf{y}, t}$ is the index in the lexicographic ordering of $(\{0, 1\}^d)^*$ corresponding to the t -th node in the branch \mathbf{y} . Let $n_{\mathbf{y}, t} = n_{\kappa_{\mathbf{y}, t}}$. The second observation states that for any $t \in \mathbb{N}$,

$$\begin{aligned} & \mathbb{E}_{\mathbf{y} \sim U(\mathcal{Y})} \left[\mathbb{P}_{S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}, (X, Y, K) \sim P_{\mathbf{y}, \rho}} \left[\hat{h}_S(X) \neq Y \mid G(\kappa_{\mathbf{y}, t}) \right] \right] \\ &= \mathbb{P}_{\mathbf{y} \sim U(\mathcal{Y}), S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}, (X, Y, K) \sim P_{\mathbf{y}, \rho}} \left[\hat{h}_S(X) \neq Y \mid G(\kappa_{\mathbf{y}, t}) \right] = \frac{1}{2}. \end{aligned} \quad (22)$$

This probability pertains to the special case of the experiment of Equation (19) in which the number n of samples in S depends on \mathbf{y} , satisfying $n = n_{\mathbf{y}, t}$. It is a conditional probability given that $G(\kappa_{\mathbf{y}, t})$ occurred, where $G(\kappa_{\mathbf{y}, t})$ is an event involving $(\mathbf{y}, X, X_1, \dots, X_n, K, K_1, \dots, K_n)$. To establish Equation (22), it suffices to show that for any $t \in \mathbb{N}$,

$$\mathbb{P}_{\mathbf{y} \sim U(\mathcal{Y}), S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}, (X, Y, K) \sim P_{\mathbf{y}, \rho}} \left[Y = 1 \mid X, \{X_i, Y_i\}_{i \in [n_{\mathbf{y}, t}]}, G(\kappa_{\mathbf{y}, t}) \right] = \frac{1}{2}, \quad (23)$$

because the prediction $\hat{h}_S(X)$ depends only on $(X, \{X_i, Y_i\}_{i \in [n_{\mathbf{y}, t}], \rho})$. Roughly, Equation (23) follows from the indifference of $\{f_{\mathbf{y}}\}_{\mathbf{y} \in \mathcal{Y}}$ (Item 3 in Theorem 54), which states that if X is a member of the K -th node in the tree T , then for any X_i with $K_i < K$ there exists a bit $b \in \{0, 1\}$ such that for all branches $\mathbf{y} \in \mathcal{Y}$ that contain node K , $f_{\mathbf{y}}(X_i) = b$. In particular, $Y = f_{\mathbf{y}}(X)$ is a uniformly random bit independent of $\{X_i, Y_i = f_{\mathbf{y}}(X_i)\}_{i \in [n_{\mathbf{y}, t}]} \cup \{X\}$ given $G(\kappa_{\mathbf{y}, t})$.

To flesh out the argument for Equation (23) in further detail, fix $\kappa \in \mathbb{N}$, $(\kappa_1, \dots, \kappa_{n_\kappa}) \in \mathbb{N}^{n_\kappa}$, $(\xi, \xi_1, \dots, \xi_{n_\kappa}) \in \mathcal{X}^{n_\kappa+1}$, and $(\eta_1, \dots, \eta_{n_\kappa}) \in \{0, 1\}^{n_\kappa}$. Consider the following conditional probability of Y for a fixed $t \in \mathbb{N}$, assuming the event being conditioned upon has a positive probability.

$$\begin{aligned}
 & \mathbb{P} \left[Y = 1 \mid \begin{array}{l} \kappa_{\mathbf{y}, t} = \kappa \\ \forall i \in [n_{\mathbf{y}, t}] : X_i = \xi_i \wedge Y_i = \eta_i \wedge K_i = \kappa_i \\ K = \kappa_{\mathbf{y}, t} \geq \max\{K_i : i \in [n_{\mathbf{y}, t}]\} \\ X = \xi \notin \{X_i : i \in [n_{\mathbf{y}, t}]\} \end{array} \right] \\
 &= \mathbb{P} \left[f_{\mathbf{y}}(X) = 1 \mid \begin{array}{l} \kappa_{\mathbf{y}, t} = \kappa \\ \forall i \in [n_{\mathbf{y}, t}] : X_i = \xi_i \wedge f_{\mathbf{y}}(X_i) = \eta_i \wedge K_i = \kappa_i \\ K = \kappa_{\mathbf{y}, t} \geq \max\{K_i : i \in [n_{\mathbf{y}, t}]\} \\ X = \xi \notin \{X_i : i \in [n_{\mathbf{y}, t}]\} \end{array} \right] \\
 & \hspace{15em} \text{(Choice of } Y \text{ and } Y_i) \\
 &= \mathbb{P} \left[f_{\mathbf{y}}(X) = 1 \mid \begin{array}{l} \kappa_{\mathbf{y}, t} = \kappa \\ \forall i \in [n_{\mathbf{y}, t}] : X_i = \xi_i \wedge K_i = \kappa_i \\ K = \kappa_{\mathbf{y}, t} \geq \max\{K_i : i \in [n_{\mathbf{y}, t}]\} \\ X = \xi \notin \{X_i : i \in [n_{\mathbf{y}, t}]\} \end{array} \right] \\
 & \hspace{15em} \text{(Indifference of } \{f_{\mathbf{y}}\}_{\mathbf{y} \in \mathcal{Y}} \text{ – Item 3 in Theorem 54)} \\
 &= \mathbb{P} \left[\mathbf{y}_t^j = 1 \mid \begin{array}{l} \kappa_{\mathbf{y}, t} = \kappa \\ \forall i \in [n_{\mathbf{y}, t}] : X_i = \xi_i \wedge K_i = \kappa_i \\ K = \kappa_{\mathbf{y}, t} \geq \max\{K_i : i \in [n_{\mathbf{y}, t}]\} \\ X = \xi \notin \{X_i : i \in [n_{\mathbf{y}, t}]\} \end{array} \right] = \frac{1}{2},
 \end{aligned}$$

where j is the index of X in the K -th node in the tree. In the last line we have used the fact that $K = \kappa_{\mathbf{y}, t}$ implies that X is on the branch corresponding to \mathbf{y} , and the final equality holds because \mathbf{y}_t is a vector of uniformly random bits chosen independently of $\{X_i, K_i\}_{i \in [n_{\mathbf{y}, t}]} \cup \{X, K, \kappa_{\mathbf{y}, t}\}$ (note that $\kappa_{\mathbf{y}, t}$ and $n_{\mathbf{y}, t}$ are fully determined by t and $\mathbf{y}_{< t}$). This establishes Equation (22), our second observation.

The first observation is used as follows. For every $\mathbf{y} \in \mathcal{Y}$,

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} n \cdot \mathbb{E}_{\rho, S \sim P_{\mathbf{y}}^n} \left[\mathbb{L}_{P_{\mathbf{y}}}^{0-1}(\hat{h}_S) \right] \\
& \geq \limsup_{t \rightarrow \infty} n_{\mathbf{y}, t} \cdot \mathbb{E}_{\rho, S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}} \left[\mathbb{L}_{P_{\mathbf{y}}}^{0-1}(\hat{h}_S) \right] \\
& \quad \text{(If } b_j \text{ is a subsequence of } a_j \text{ then } \limsup a_j \geq \limsup b_j) \\
& = \limsup_{t \rightarrow \infty} n_{\mathbf{y}, t} \cdot \mathbb{P}_{\rho, S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}, (X, Y, K) \sim P_{\mathbf{y}, \rho}} \left[\hat{h}_S(X) \neq Y \right] \\
& \geq \limsup_{t \rightarrow \infty} n_{\mathbf{y}, t} \cdot \mathbb{P}_{\rho, S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}, (X, Y, K) \sim P_{\mathbf{y}}} \left[\left(\hat{h}_S(X) \neq Y \right) \wedge G(\kappa_{\mathbf{y}, t}) \right] \\
& = \limsup_{t \rightarrow \infty} n_{\mathbf{y}, t} \cdot \mathbb{P}[G(\kappa_{\mathbf{y}, t})] \cdot \mathbb{P} \left[\hat{h}_S(X) \neq Y \mid G(\kappa_{\mathbf{y}, t}) \right] \\
& \geq \limsup_{t \rightarrow \infty} \frac{d^{\kappa_{\mathbf{y}, t} + 1}}{9(d-1)} \cdot \frac{(d-1)d^{-\kappa_{\mathbf{y}, t}}}{4} \cdot \mathbb{P} \left[\hat{h}_S(X) \neq Y \mid G(\kappa_{\mathbf{y}, t}) \right] \\
& \quad \text{(By Equation (20) and choice of } n_{\mathbf{y}, t}) \\
& = \limsup_{t \rightarrow \infty} \frac{d}{36} \cdot \mathbb{P}_{\rho, S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}, (X, Y, K) \sim P_{\mathbf{y}}} \left[\hat{h}_S(X) \neq Y \mid G(\kappa_{\mathbf{y}, t}) \right]. \tag{24}
\end{aligned}$$

To complete the proof we use our second observation and Fatou's lemma as follows.

$$\begin{aligned}
& \mathbb{E}_{\mathbf{y} \sim \mathcal{U}(\mathcal{Y})} \left[\limsup_{n \rightarrow \infty} n \cdot \mathbb{E}_{\rho, S \sim P_{\mathbf{y}}^n} \left[\mathbb{L}_{P_{\mathbf{y}}}^{0-1}(\hat{h}_S) \right] \right] \\
& \geq \frac{d}{36} \cdot \mathbb{E}_{\mathbf{y} \sim \mathcal{U}(\mathcal{Y})} \left[\limsup_{t \rightarrow \infty} \mathbb{P}_{\rho, S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}, (X, Y, K) \sim P_{\mathbf{y}}} \left[\hat{h}_S(X) \neq Y \mid G(\kappa_{\mathbf{y}, t}) \right] \right] \\
& \quad \text{(By Equation (24))} \\
& \geq \frac{d}{36} \cdot \limsup_{t \rightarrow \infty} \mathbb{E}_{\mathbf{y} \sim \mathcal{U}(\mathcal{Y})} \left[\mathbb{P}_{\rho, S \sim P_{\mathbf{y}}^{n_{\mathbf{y}, t}}, (X, Y, K) \sim P_{\mathbf{y}}} \left[\hat{h}_S(X) \neq Y \mid G(\kappa_{\mathbf{y}, t}) \right] \right] \\
& \quad \text{(Fatou's lemma (Lemma 55), } \mathbb{P}[\cdot] \leq 1) \\
& = \frac{d}{36} \cdot \frac{1}{2} = \frac{d}{72}. \tag{By Equation (22)}
\end{aligned}$$

This implies that there exists $\mathbf{y} \in \mathcal{Y}$ such that

$$\limsup_{n \rightarrow \infty} n \cdot \mathbb{E}_{S \sim P_{\mathbf{y}}^n} \left[\mathbb{L}_{P_{\mathbf{y}}}^{0-1}(\hat{h}_S) \right] \geq \frac{d}{72}.$$

By the definition of lim sup, the inequality

$$\mathbb{E}_{S \sim P_{\mathbf{y}}^n} \left[\mathbb{L}_{P_{\mathbf{y}}}^{0-1}(\hat{h}_S) \right] \geq \frac{d}{73 \cdot n}$$

holds for infinitely many values of $n \in \mathbb{N}$, as desired. \blacksquare

Appendix C. Result for Half-Spaces

Notation 56 Let $d \in \mathbb{N}$. We write $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\|_2 = 1\}$ to denote the unit sphere in \mathbb{R}^d .

Definition 57 Let $d \in \mathbb{N}$. For any $\mathbf{w} \in \mathbb{S}^{d-1}$, let $h_{\mathbf{w}} : \mathbb{R}^d \rightarrow \{0, 1\}$ be the half-space given by $h_{\mathbf{w}}(\mathbf{x}) = \mathbb{1}(\langle \mathbf{w}, \mathbf{x} \rangle > 0)$. The class of homogeneous half-spaces in \mathbb{R}^d is $\mathcal{H}_d = \{h_{\mathbf{w}} : \mathbf{w} \in \mathbb{S}^{d-1}\}$.

Definition 58 Let $d \in \mathbb{N}$, let $H \subseteq \mathbb{S}^{d-1}$ be a set. We say that a set of points $\{x_1, \dots, x_m\} \subseteq \mathbb{R}^d$ is openly shattered by H if for every vector $\mathbf{y} = (y_1, \dots, y_m) \in \{0, 1\}^m$, there exists an open set $W_{\mathbf{y}} \subseteq H$ such that

$$\forall \mathbf{w} \in W_{\mathbf{y}} \forall i \in [m] : h_{\mathbf{w}}(x_i) = y_i. \quad (25)$$

Lemma 59 Let $d \in \mathbb{N}$, and let $H \subseteq \mathbb{S}^{d-1}$ be an open set. Then there exists a set $X \subseteq \mathbb{S}^{d-1}$ such that $|X| = d - 1$ and X is openly shattered by H .

Proof Fix a point x_0 in the interior of the H . Let $x_1, \dots, x_{d-1} \in \mathbb{S}^{d-1}$ be points such that x_0, x_1, \dots, x_{d-1} is an orthonormal basis of \mathbb{R}^d .

For each $\mathbf{y} = (y_1, \dots, y_{d-1}) \in \{0, 1\}^{d-1}$, let

$$w'_{\mathbf{y}} = x_0 + \varepsilon \cdot \sum_{i \in [d-1]} \text{sign}(y_i - 1/2) \cdot x_i$$

be a point with $\varepsilon > 0$ small enough such that $w_{\mathbf{y}}$ is in the interior of H , where $w_{\mathbf{y}}$ is the projection of $w'_{\mathbf{y}}$ onto \mathbb{S}^{d-1} . From the orthogonality of $\{x_0, \dots, x_{d-1}\}$,

$$\forall i \in [d-1] : h_{w_{\mathbf{y}}}(x_i) = y_i.$$

For each $i \in [d-1]$ and $y \in \{0, 1\}$, let $Q_{i,y} \subseteq \mathbb{S}^{d-1}$ be the set of \mathbf{w} such that $h_{\mathbf{w}}(x_i) = y$. Because we use open half-spaces, $Q_{i,y}$ is open. Observe that for each $\mathbf{y} \in \{0, 1\}^{d-1}$,

$$W_{\mathbf{y}} = H \cap \bigcap_{i \in [d-1]} Q_{i,y_i}$$

is open (as a finite intersection of open sets), and is non-empty because it contains $w_{\mathbf{y}}$. ■

Lemma 60 Let $d \in \mathbb{N}$. Then $\text{VCL}(\mathcal{H}_d) \geq d - 1$.

Proof We recursively construct an infinite $(d-1)$ -VCL tree that is shattered by \mathcal{H}_d . Let $H_{\lambda} = \mathbb{S}^{d-1}$. For every $s \in 0, 1, 2, \dots$ do the following. For every $\mathbf{u} \in \{0, 1\}^{ds}$, note that $H_{\mathbf{u}} \subseteq \mathbb{S}^{d-1}$ is open. Therefore, by Lemma 59, there exists $\mathbf{x}_{\mathbf{u}} = (x_{\mathbf{u}}^1, \dots, x_{\mathbf{u}}^{d-1}) \subseteq \mathbb{S}^{d-1}$ of cardinality $d - 1$ that is openly shattered by $H_{\mathbf{u}}$. Namely, for each $\mathbf{y} \in \{0, 1\}^d$ there exists an open set $W_{\mathbf{y}} \subseteq H_{\mathbf{u}}$ such that Equation (25) holds (for $x_i = x_{\mathbf{u}}^i$ and $m = d$). For each $\mathbf{y} \in \{0, 1\}^d$, define $H_{\mathbf{u} \circ \mathbf{y}} = W_{\mathbf{y}}$.

We claim that $T = \{\mathbf{x}_{\mathbf{u}} : \mathbf{u} \in (\{0, 1\}^d)^*\}$ is a $(d-1)$ -VCL tree that is shattered by \mathcal{H}_d . Indeed, fix $t \in \mathbb{N}$ and $\mathbf{y} \in \{0, 1\}^{td}$. Let $\mathbf{w} \in H_{\mathbf{u}}$. Then the choice of $\mathbf{x}_{\mathbf{u}}$ and $H_{\mathbf{u}}$ implies that

$$\forall s \in [t] \forall j \in [d] : h_{\mathbf{w}}(x_{\mathbf{y}_{\leq s-1}}^j) = y_s^j,$$

as desired. ■