

Optimal Mistake Bounds for Transductive Online Learning

Zachary Chase

Kent State University

zchase2@kent.edu

Steve Hanneke

Purdue University

steve.hanneke@gmail.com

Shay Moran

*Departments of Mathematics, Computer Science,
and Data and Decision Sciences*

Technion – Israel Institute of Technology;

Google Research

smoran@technion.ac.il

Jonathan Shafer

MIT

shaferjo@mit.edu

Abstract

We resolve a 30-year-old open problem concerning the power of unlabeled data in online learning by tightly quantifying the gap between transductive and standard online learning. In the standard setting, the optimal mistake bound is characterized by the Littlestone dimension d of the concept class \mathcal{H} (Littlestone, 1987). We prove that in the transductive setting, the mistake bound is at least $\Omega(\sqrt{d})$. This constitutes an exponential improvement over previous lower bounds of $\Omega(\log \log d)$, $\Omega(\sqrt{\log d})$, and $\Omega(\log d)$, due respectively to Ben-David, Kushilevitz, and Mansour (1995, 1997), and Hanneke, Moran, and Shafer (2023). We also show that this lower bound is tight: for every d , there exists a class of Littlestone dimension d with transductive mistake bound $O(\sqrt{d})$. Our upper bound also improves upon the best known upper bound of $(2/3) \cdot d$ from Ben-David et al. (1997). These results establish a quadratic gap between transductive and standard online learning, thereby highlighting the benefit of advance access to the unlabeled instance sequence. This contrasts with the PAC setting, where transductive and standard learning exhibit similar sample complexities.

Contents

1	Introduction	1
1.1	Setting: <i>Standard</i> vs. <i>Transductive</i> Online Learning	2
1.2	Main Result	3
1.3	Related Works	4
2	Technical Overview	4
2.1	Paths in Trees	5
2.2	Proof Ideas for the Lower Bound	5
2.3	Proof Ideas for the Upper Bound	7
2.3.1	Sparse Encodings are Easy to Guess	7
2.3.2	Construction of the Hypothesis Class	8
2.3.3	Naïve Learning Strategy	8
2.3.4	Danger Zone Minimization	10
2.3.5	Splitting Experts	11
2.3.6	Transition to Halving	11
2.4	Some Intuition for the Quantity \sqrt{d}	12
3	Directions for Future Work	12
4	Organization	13
5	Preliminaries	13
5.1	Basic Notation	13
5.2	Standard Online Learning	13
5.3	Transductive Online Learning	14
5.4	Mistake Bounds	14
5.5	Trees	15
5.6	Littlestone Dimension	16
6	Lower Bound	16
6.1	Statement	16
6.2	Analysis of the Adversary	18
6.3	Proof	21
7	Sequence Length	22
7.1	Rigid Adversary	23
7.2	Essential Indices	25
7.3	Proof	26
8	Upper Bound	26
8.1	Statement	26
8.2	Hypothesis Class	27
8.3	Algorithm	29

8.3.1	How Experts Work	29
8.3.2	Pseudocode	31
8.4	Analysis	32
8.4.1	Assumption-Consistent Expert	32
8.4.2	Transition to Halving	34
8.4.3	Performance of Best Expert	35
8.4.4	Multiplicative Weights Mistake Bound	38
8.5	Proof	39
	References	40
	Appendices	42
	A Halving	42

1 Introduction

The transductive model is a basic and well-studied framework in learning theory, dating back to the early works of Vapnik. It has been investigated both in statistical and online settings, and is motivated by the principle that to make good predictions on a specific set of test instances, one need not construct a fully general classifier that performs well on the entire domain — including points that may never actually appear. Rather, it may be sufficient to tailor predictions for a fixed, known set of instances.

This perspective naturally connects to a broader question in learning theory: what is the value of unlabeled data? In the transductive setting, the learner is given the sequence of unlabeled test instances in advance and is then required to predict their labels one by one. Thus, the transductive model can be viewed as a natural formalization of learning with unlabeled data: the test instances are known in advance, but their labels are not. The central question is whether such prior access to the unlabeled sequence can help reduce the number of prediction mistakes — compared to the standard online model, where the instances arrive and are labeled one at a time.

Recall for instance that in the standard PAC¹ model of supervised learning, there are cases where access to unlabeled data is not helpful. Indeed, the “hard population distributions” used to prove the standard VC² lower bound are constructed by taking a fixed and known marginal distribution over a VC-shattered set. Namely, the cases that are hardest to learn in the PAC setting include ones where the learner knows the marginal distribution over the domain, and can therefore generate as much unlabeled data as it wishes. And yet, in those cases, access to unlabeled data provides no acceleration compared to an algorithm (like ERM³) that does not use unlabeled data.

Seeing as unlabeled data is often a lot easier to obtain than labeled data, there have been considerable efforts to understand when and to what extent can access to unlabeled data accelerate learning.⁴

In particular, it is natural to ask, for which plausible models of learning is access to unlabeled data beneficial? Online learning (Littlestone, 1987) is perhaps the model of learning that is most-extensively studied in learning theory after the PAC model and its variants. Therefore, the general question considered in this paper is:

Question 1. *Quantitatively, how much (if at all) is access to unlabeled data beneficial for learning in the online learning setting?*

This question is naturally instantiated by comparing *transductive* online learning — where the learner has advance access to the full sequence x_1, x_2, \dots, x_n of unlabeled instances — with *standard* online learning, where no such access is given. This perspective has also been adopted in prior

¹Probably Approximately Correct. For an exposition of the standard terminology and results mentioned in this paragraph see, e.g., Shalev-Shwartz and Ben-David (2014).

²Vapnik–Chervonenkis.

³Empirical Risk Minimization.

⁴The literature on semi-supervised learning is surveyed in Joachims (1999); Zhu (2005); Zhu and Goldberg (2009); Zhu (2010); Chapelle, Schölkopf, and Zien (2006). Theoretical works on the topic include Benedek and Itai (1991); Blum and Mitchell (1998); Ben-David, Lu, Pál, and Sotáková (2008); Balcan and Blum (2010); Darnstädt, Simon, and Szörényi (2013); Göpfert, Ben-David, Bousquet, Gelly, Tolstikhin, and Urner (2019).

work: for example, Kakade and Kalai (2005), Cesa-Bianchi and Shamir (2013), and Hoi, Sahoo, Lu, and Zhao (2021) (Section 7.3) all describe transductive online learning as a setting in which the learner has access to “unlabeled data”. We thus refine the question above as follows:

Question 2. *Quantitatively, how much (if at all) is learning in the transductive online learning setting easier than learning in the standard online learning setting? Specifically, how much is the optimal number of mistakes in the transductive setting smaller than in the standard setting?*

Addressing this question, our main result (Theorem 1.1) states that the optimal number of mistakes in the transductive setting (with access to unlabeled data) is at most quadratically smaller than in the standard setting (without unlabeled data). Furthermore, there are hypothesis classes for which a quadratic gap is achieved.

1.1 Setting: Standard vs. Transductive Online Learning

Standard online learning (Littlestone, 1987) is a zero-sum, perfect- and complete-information game played over n rounds between two players, a *learner* and an *adversary*. The game is played with respect to a *domain* set \mathcal{X} and a *hypothesis class* $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ (consisting of functions $\mathcal{X} \rightarrow \{0, 1\}$), where n , \mathcal{X} and \mathcal{H} are fixed and known to both players. The game proceeds as follows.

For each round $t = 1, 2, \dots, n$:

- a. The adversary selects an *instance* $x_t \in \mathcal{X}$ and sends it to the learner.
- b. The learner selects a *prediction* $\hat{y}_t \in \{0, 1\}$ and sends it to the adversary.
- c. The adversary selects a *label* $y_t \in \{0, 1\}$ and sends it to the learner. The selected label must be *realizable*, meaning that $\exists h \in \mathcal{H} \forall i \in [t] : h(x_i) = y_i$.

Game 1: The standard online learning setting.

The *number of mistakes* for a learner L and an adversary A is $M_{\text{std}}(\mathcal{H}, n, L, A) = |\{t \in [n] : \hat{y}_t \neq y_t\}|$. We are interested in understanding the *optimal number of mistakes*, which is

$$M_{\text{std}}(\mathcal{H}) = \sup_{n \in \mathbb{N}} \inf_{L \in \mathcal{L}} \sup_{A \in \mathcal{A}} M_{\text{std}}(\mathcal{H}, n, L, A),$$

where \mathcal{A} and \mathcal{L} are the set of all deterministic adversaries and learners, respectively.⁵ It is well known that $M_{\text{std}}(\mathcal{H})$ is characterized by the Littlestone dimension, namely, $M_{\text{std}}(\mathcal{H}) = \text{LD}(\mathcal{H})$ (see Theorem 5.7 and Definition 5.6).

The *transductive* online learning setting (Ben-David et al., 1995, 1997) is similar, except that the learner has access to the full sequence of unlabeled instances in advance. Namely,

⁵In this paper, we only consider deterministic learners and adversaries. As is common in learning theory, we avoid questions of computability and allow the learner and adversary to be any function. See Section 5 for formal definitions of \mathcal{A} and \mathcal{L} .

One could consider an alternative setting where the learner and adversary are randomized and the adversary selects each label y_t *before* the learner tosses the random coins to determine the prediction \hat{y}_t ; the optimal number of mistakes in the alternative setting equals the optimal number of mistakes in our deterministic setting up to a factor of two.

The adversary selects a *sequence* $x_1, x_2, \dots, x_n \in \mathcal{X}$ and sends it to the learner.

For each round $t = 1, 2, \dots, n$:

- a. The learner selects a *prediction* $\hat{y}_t \in \{0, 1\}$ and sends it to the adversary.
- b. The adversary selects a *label* $y_t \in \{0, 1\}$ and sends it to the learner. The selected label must be *realizable*, meaning that $\exists h \in \mathcal{H} \forall i \in [t] : h(x_i) = y_i$.

Game 2: The transductive online learning setting.

The optimal number of mistakes for the transductive setting is defined exactly as before,

$$M_{\text{tr}}(\mathcal{H}, n, L, A) = |\{t \in [n] : \hat{y}_t \neq y_t\}|, \text{ and } M_{\text{tr}}(\mathcal{H}) = \sup_{n \in \mathbb{N}} \inf_{L \in \mathcal{L}} \sup_{A \in \mathcal{A}} M_{\text{tr}}(\mathcal{H}, n, L, A),$$

with the only difference between the standard quantity $M_{\text{std}}(\mathcal{H})$ and the transductive quantity $M_{\text{tr}}(\mathcal{H})$ being in how the game is defined.

1.2 Main Result

Notice that for every hypothesis class \mathcal{H} , $M_{\text{tr}}(\mathcal{H}) \leq M_{\text{std}}(\mathcal{H})$. Indeed, in the transductive setting the adversary declares the sequence x at the start of the game. This reduces the number of mistakes because the transductive adversary is less powerful (it cannot adaptively alter the sequence mid-game), and also because the transductive learner is more powerful (it has more information).⁶

While for some classes $M_{\text{tr}}(\mathcal{H}) = M_{\text{std}}(\mathcal{H})$, we study the largest possible separation. The best previous lower bound on M_{tr} , due to [Hanneke, Moran, and Shafer \(2023\)](#), states that for every class \mathcal{H} ,

$$M_{\text{tr}}(\mathcal{H}) \geq \Omega(\log d),$$

where $d = M_{\text{std}}(\mathcal{H})$. In the other direction, [Ben-David et al. \(1997\)](#) constructed⁷ a class \mathcal{H} such that $M_{\text{std}}(\mathcal{H}) = d$ and

$$M_{\text{tr}}(\mathcal{H}) \leq \frac{2}{3}d.$$

⁶One could also define an intermediate setting, where the adversary is less powerful because it must select the sequence at the start of the game and cannot change it during the gameplay, but the learner does not have more information because the adversary only reveals the instances in the sequence one at a time as in the standard setting. However, this intermediate setting would not model the learner having *access* to unlabeled data.

On a technical level, the intermediate setting is essentially not easier for the learner than the standard online learning setting. If the hypothesis class shatters a Littlestone tree of depth $d \in \mathbb{N}$, then for every deterministic learner there exists a deterministic adversary that forces the learner to make at least d mistakes in the intermediate-transductive setting: The intermediate-transductive adversary simulates an execution of the deterministic learner playing against the standard online deterministic adversary in the *standard* online learning setting to construct a sequence of instances x_i and labels y_i (corresponding to a root-to-leaf branch of the tree). After the simulation is complete, the intermediate-transductive adversary commits to that sequence of instances at the beginning of the intermediate-transductive game, and assigns the same labels that the adversary assigned in the simulation. Because the learner is deterministic and its view in the two games is the same, it is guaranteed to make at least d mistakes on this sequence also in the intermediate-transductive setting.

Similarly, if the learner is randomized and at each round the adversary selects a label before the learner tosses its random coins, then the intermediate-transductive adversary can force an expected number of $d/2$ mistakes by selecting a root-to-leaf branch of the tree uniformly at random.

⁷Their class consists of all disjoint unions of $\Theta(d)$ functions from a specific constant-sized class.

This left an exponential gap between the best known lower and upper bounds on M_{tr} , namely $\Omega(\log d)$ versus $\frac{2}{3}d$. Our main result closes this gap:

Theorem 1.1 (Main result).

- For every hypothesis class $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$,

$$M_{\text{tr}}(\mathcal{H}) = \Omega(\sqrt{d}),$$

where $d = M_{\text{std}}(\mathcal{H})$.

- On the other hand, for every d there exists a hypothesis class \mathcal{H} with $M_{\text{std}}(\mathcal{H}) = d$ and

$$M_{\text{tr}}(\mathcal{H}) = O(\sqrt{d}).$$

This result is stated in considerably greater detail in Theorems 6.1 and 8.1.

1.3 Related Works

The notion of *transductive inference* as a more efficient alternative to *inductive inference* in statistical learning theory was introduced by Vapnik (1979, 2006); Gammernan, Vovk, and Vapnik (1998); Chapelle, Vapnik, and Weston (1999). The *online learning* setting is due to Littlestone (1987), who also proved that the optimal number of mistakes is characterized by the Littlestone dimension (see Theorem 5.7).

The *transductive online learning* setting studied in the current paper, was first defined by Ben-David, Kushilevitz, and Mansour (1995), who used the name *worst sequence off-line model*. Among other results, they showed a lower bound of $\Omega(\log \log d)$ on the number of mistakes required to learn a class with Littlestone dimension d . The authors subsequently presented an exponentially stronger lower bound of $\Omega(\sqrt{\log d})$ in Ben-David, Kushilevitz, and Mansour (1997). However, understanding where the optimal number of mistakes is situated within the range $[\Omega(\sqrt{\log d}), 2d/3]$ remained an open question.


Kakade and Kalai (2005) presented an oracle-efficient algorithm for the transductive online learning setting, and may have been the first to use that name. Their result was subsequently improved upon by Cesa-Bianchi and Shamir (2013).

The present work is most similar to that of Hanneke, Moran, and Shafer (2023) which, among other results, gave a quadratically-stronger mistake lower bound of $\Omega(\log d)$ for classes with Littlestone dimension d in the transductive online setting. The proof of our lower bound utilizes some of their ideas, but yields a quantitative improvement by combining it with some new ideas.


Hanneke, Raman, Shaeiri, and Subedi (2024) studied a setting of *multi-class* transductive online learning where the number of possible labels is unbounded.

2 Technical Overview


In this section we explain some of the main ideas in our proofs. Formal definitions appear in Section 5. Full formal statements of the results, as well as detailed rigorous proofs, appear in Sections 6 to 8.




(a) A *perfect binary tree* of depth 2. Each *node* is labeled by an element of the domain \mathcal{X} . These labels need not be distinct (e.g., it is possible that $x_1 = x_6$). x_0 is the *root* of the tree, x_0 , x_1 and x_2 are *internal nodes*, and x_3, \dots, x_6 are *leaves*.



(b) A function $f : \mathcal{X} \rightarrow \{0, 1\}$ assigns a binary label to each node in the tree, represented here by edges with arrowhead tips. This figure depicts the function $f(x_i) = \mathbb{1}(i \notin \{2, 3\})$. (Note that the gray dots (\bullet) in the figure are purely a pictorial detail. In this paper they are not considered nodes or leaves of the tree.)



(c) Every function $f : \mathcal{X} \rightarrow \{0, 1\}$ defines a *path* in the tree, which is a sequence $u_0, u_1, u_2, \dots, u_{d-1}$, where u_0 is the root, d is the depth of the tree, and for each $i \in [d-1]$, u_i is the b -child of u_{i-1} with $b = f(u_{i-1}) \in \{0, 1\}$. This figure shows that the function f from Figure 1b has $\text{path}(f) = (x_0, x_2, x_5)$, depicted in red. In particular, x_2 is ‘on-path’ for f , but x_6 is ‘off-path’ for f .



(d) In this paper we use a naming convention where, without loss of generality, we identify the domain elements x_i that are assigned to nodes with bit strings. The root is identified with the empty string λ , and for each pair of nodes u, v such that u is the b -child of v (for $b \in \{0, 1\}$), we have $u = v \circ b$, where ‘ \circ ’ denotes string concatenation. (Because the x_i ’s may not be distinct, a domain element may be identified with more than one bit string.)

Figure 1: Paths in trees.

2.1 Paths in Trees

We make extensive use of the following notion. Given a perfect binary tree T_d of depth d , every function $f : T_d \rightarrow \{0, 1\}$ defines a unique *path* in the tree. The path is a sequence of nodes $\text{path}(f) = (x_{i_0}, x_{i_1}, \dots, x_{i_d})$, as explained in Figure 1c. See Section 5 for formal definitions.

2.2 Proof Ideas for the Lower Bound

We start with an elementary observation about the adversary’s dilemma in the transductive online learning setting. Before round t of the game, the adversary selected a full sequence of instances $x_1, x_2, \dots, x_n \in \mathcal{X}$, and assigned some initial labels $y_1, y_2, \dots, y_{t-1} \in \{0, 1\}$. At the start of round t ,

the adversary must consider the *version space*,

$$\mathcal{H}_t = \{h \in \mathcal{H} : (\forall i \in [t-1] : h(x_i) = y_i)\}.$$

If all $h \in \mathcal{H}_t$ assign $h(x_t) = b$ for some $b \in \{0, 1\}$, then the adversary has no choice but to assign the label $y_t = b$. Otherwise, the adversary can *force a mistake* at time t . Namely, after seeing the learner's prediction \hat{y}_t , the adversary can assign $y_t = 1 - \hat{y}_t$, incrementing the number of learner mistakes by 1.

But “just because you can, doesn't mean you should”. If the adversary is greedy and forces a mistake at time t , they may pay dearly for that later. As an extreme example, consider the case where there is a single $h_1 \in \mathcal{H}_t$ that assigns $h_1(x_t) = 1$, and all other functions $h \in \mathcal{H}_t$ assign $h(x_t) = 0$. If the learner selects $\hat{y}_t = 1$ and the adversary forces a mistake at time t , the version space at all subsequent times $s > t$ will be $\mathcal{H}_s = \{h_1\}$, and the adversary will be prevented from forcing any further mistakes.

A natural strategy for the adversary is therefore to be greedy up to a certain limit. Namely, at each time t the adversary computes the ratio⁸

$$r_t = \frac{|\{h \in \mathcal{H}_t : h(x_t) = 1\}|}{|\mathcal{H}_t|}.$$

If $r_t \in [\varepsilon, 1 - \varepsilon]$ for some parameter $\varepsilon > 0$ (“the version space is not too unbalanced”), then the adversary forces a mistake. Otherwise, the adversary assigns the majority label, i.e., $y_t = \mathbb{1}(r_t \geq 1/2)$. This ensures that the version space does not shrink too fast:

- If no mistake is forced, then $|\mathcal{H}_{t+1}| \geq (1 - \varepsilon) \cdot |\mathcal{H}_t|$, and
- If a mistake is forced, $|\mathcal{H}_{t+1}| \geq \varepsilon \cdot |\mathcal{H}_t|$.

In particular, at the end of the game, the version space \mathcal{H}_{n+1} is of size

$$|\mathcal{H}_{n+1}| \geq \varepsilon^M \cdot (1 - \varepsilon)^{n-M} \cdot |\mathcal{H}| \geq \varepsilon^M \cdot (1 - \varepsilon)^n \cdot 2^d, \quad (1)$$

where M is the number of mistakes that the adversary forces and n is the length of the sequence. The class has size $|\mathcal{H}| \geq 2^d$ because $\text{LD}(\mathcal{H}) = d$, and by removing functions from the class if necessary (which can only make learning easier), we may assume without loss of generality that $|\mathcal{H}| = 2^d$. Namely, the class precisely shatters a Littlestone tree of depth $d - 1$ such that for every assignment of labels to a root-to-leaf path in the tree, the class contains exactly one function that agrees with that assignment (see Definition 5.6 for detail).

Notice that we have not yet specified how the adversary selects the sequence x . While the adversary's labeling strategy is extremely simple (determined by the ratio r_t and the prediction \hat{y}_t), constructing of the sequence x requires some care, to ensure that it has the following two properties:

- **Property I:** The length n of the sequence satisfies $n = 2^{\Theta(\sqrt{d})}$, and

⁸For a class \mathcal{H} of Littlestone dimension d , the adversary will use only a subset of \mathcal{H} of cardinality 2^d that shatters a Littlestone tree of depth $d - 1$. So without loss of generality, we assume that \mathcal{H} has cardinality 2^d (in particular, \mathcal{H} is finite), and the ratio is well-defined.

between 0 and $2^6 - 1$. We guess the bits one at a time, and after each guess, an adaptive adversary tells us whether our guess was correct.

Now, if the bit string is a binary encoding, the task is hard. Each bit can either be 0 or 1, regardless of the values of the previous bits, and so the adversary can force a mistake on every bit. On the other hand, if we know that the string is a one-hot encoding, there exists an attractive strategy – always guess 0. This ensures that we will make at most 1 mistake.

Note that at the end of the guessing game we have learned the same amount of *information* (for a number between 0 and $2^n - 1$, we learned n bits of information), but the number of *mistakes* is very different (n mistakes vs. 1 mistake).

2.3.2 Construction of the Hypothesis Class

We now describe a construction of a hypothesis class that is easy to learn in the transductive setting, using the idea of a sparse encoding. Recall that a class \mathcal{H} has Littlestone dimension at least d (Definition 5.6 in Section 5) if there exists a Littlestone tree of depth $d - 1$ such that for every $b \in \{0, 1\}^d$ there exists $h = h_b \in \mathcal{H}$ such that the values on the path of h agree with b . More formally, $\forall i \in [d] : h(b_{<i}) = b_i$, and in particular $\text{path}(h) = (\lambda, b_{\leq 1}, b_{\leq 2}, b_{\leq 3}, \dots, b_{\leq d-1})$. Thus, when constructing a class that shatters a specific Littlestone tree of depth $d - 1$, we need to define 2^d functions $\{h_b : b \in \{0, 1\}^d\}$. For each function h_b , the on-path values of the function are fixed (fully determined by b), while for the remaining values there is complete freedom (for the nodes u that are off-path we may assign any values $h_b(u) \in \{0, 1\}$).

Perhaps the simplest way to construct a class of Littlestone dimension d is simply to assign all on-path values as required, and assign 0 to all other values. Namely, if u is a prefix of b then $h_b(u) = b_{|u|+1}$, and otherwise $h_b(u) = 0$. In a sense, this is the ‘minimal’ class of Littlestone dimension d for a specific Littlestone tree.¹¹

Observe that the ‘minimal’ class does not have the desired property of being easy to learn in the transductive setting.¹² However, a certain variation of the ‘minimal’ class that embeds a sparse encoding does satisfy the requirement. In this variation, on-path value of the function h_b are assigned as they must (as determined by b), while the off-path values are sampled independently using a biased coin, such that each of them is 0 with high probability, but has a small probability of being 1. The probability is chosen carefully so that the class satisfies some simple combinatorial properties, as described further in Section 2.3.6 and Lemma 8.2.

2.3.3 Naïve Learning Strategy

We now explain in broad strokes how the probabilistic construction of the hypothesis class in Section 2.3.2 is useful for learning with few mistakes in the transductive setting.

¹¹More formally, this is a class with a minimal number of nodes labeled 1.

¹²The adversary can declare a sequence x consisting of all the nodes in the tree in breadth-first order, and then force d mistakes – one mistake in each layer (depth) of the tree. Specifically, regardless of how the adversary selects the labels, for each $i \in [d]$ there exists a node u_i at depth i that is on-path. When it is time for the learner to predict a label for this u_i , the learner knows that u_i is on-path because it has seen the correct labels for all the ancestors of u_i . However, the adversary has the freedom to extend the path arbitrarily to the left or to the right, and can therefore force a mistake on u_i .

Notice that when predicting labels for the ‘minimal’ class with nodes in breadth-first order, the learner knows at each step whether they are labeling an on-path or off-path node, because the learner has already seen the correct labels for all ancestors of the current node. For off-path nodes, the learner knows that the true label is 0, so it never makes mistakes on off-path nodes, but it also gains no new information when the true labels for off-path nodes are revealed. No risk, but no reward either. Instead, all the information about the true labeling function is revealed only at on-path nodes, where the adversary has complete freedom to assign labels and force mistakes. That’s why the adversary can force d mistakes.

For the randomly-chosen class, when predicting labels for off-path nodes, the learner may still safely predict a label of 0. But the reasoning for this is quite different. Conceptually, every off-path label is part of a sparse codeword that identifies the correct labeling function.¹³ Because the coin is biased, each bit of the codeword is easy to guess (it is likely to be 0), but every time that the adversary reveals that the true label for an off-path node is indeed 0, the learner gains a small (nonzero) amount of information about the true labeling function. Additionally, when the adversary selects an off-path label of 1, that reveals a lot of information about the true labeling function (such labels are rare in the hypothesis class), and therefore the adversary cannot force many off-path mistakes. Overall, the information about the true labeling function is ‘smeared’ throughout all labels of the tree (0s and 1s, on-path and off-path).¹⁴

Thus, the naïve general strategy for the learner when using the probabilistically-constructed class is to learn most of the information about the true labeling function by observing off-path labels. By the time the learner reaches an on-path node, it hopefully has already learned enough about the true labeling function in order to make a good prediction on that node.

However, making this general strategy work requires overcoming some very substantial obstacles:

1. Recall that in the transductive setting, the adversary can present the nodes of the tree in any order of its choosing – it does not have to present the tree in breadth-first order. The naïve strategy works only if the learner sees many off-path nodes before it sees most on-path nodes. But what happens if the adversary decides to present many on-path nodes near the beginning of the sequence? To handle this, the learner incorporates a strategy we call ‘danger zone minimization’, as described in Section 2.3.4.
2. Another, equally problematic, issue also arises from the fact that the sequence presented by the adversary might not be in breadth-first order. Recall that breadth-first order¹⁵ has the property that for every node u in the sequence, all the ancestors of u appear *before* u in the sequence. This means that by the time the learner needs to predict a label for u , the learner knows whether u is on-path or off-path for the true labeling function. But what happens if

¹³The coin-flips for off-path labels are all independent. For example, if X is a set of nodes all of which are off-path for a subset H of the hypothesis class, then the random variables $\{h(x) : h \in H, x \in X\}$ are i.i.d.

¹⁴Furthermore, the labels for most not-too-small subsets of the nodes reveal a lot of information about the correct labeling function – not just for a particular subset of nodes. These properties led us to code-name this construction while working on the paper as ‘everything everywhere all at once’ (in reference to a 2022 film of that name). This is in contrast to the ‘minimal’ function, where the information is concentrated entirely on the function path. The asymmetry between the ‘minimal’ class and the probabilistic class is similar to that between the binary and one-hot encodings in Section 2.3.1 above.

¹⁵As well as depth-first order.

the adversary presents u before some of u 's ancestors? Or omits some of u 's ancestors from the sequence altogether? In this case the learner doesn't know if u is on-path or off-path, and this presents a double hazard. One hazard is that the learner doesn't know what label to predict for u – if u is off-path, the learner can simply predict 0, but if it is on-path it must do something more elaborate. The second hazard is that, after seeing the correct label for u , it is not clear what the learner can infer from it. If u is off-path, its label should be interpreted as part of a sparse encoding of the labeling function. But if u is on-path, the interpretation must be entirely different. To overcome this challenge, the learner incorporates a strategy we call ‘splitting experts’, described in Section 2.3.5.

3. Limiting off-path mistakes. Thanks to the coin's bias, most off-path nodes have a true label of 0. Nonetheless, each function in the hypothesis class still has an expected number of $2^{\Omega(d)}$ off-path nodes labeled 1, so the learner can afford to misclassify only a vanishing fraction of them! To limit the number of mistakes, the learner extracts information from the sparse encoding and executes a ‘transition to Halving’ strategy, as described in Section 2.3.6.

2.3.4 Danger Zone Minimization

Utilizing information from the ‘sparse encoding’ of the off-path nodes to make good predictions for on-path nodes requires that the learner first see the true labels for many off-path nodes. Until that happens, the learner expects to make many mistakes on on-path nodes. However, whether a node is on-path or off-path is not fixed in advanced – the adversary may decide this adaptively, in response to the learners predictions.

Danger zone minimization is a strategy used by the learner, to force the adversary to assign few nodes in the beginning of the sequence as on-path (otherwise, if initial nodes are assigned to be on-path by the adversary, then the learner will make few mistakes on those nodes). This is analogous to the standard Halving algorithm (Algorithm 7), but instead of minimizing the cardinality of the set of consistent hypotheses (the ‘version space’), the learner minimizes a subset of the domain (the ‘danger zone’).

Concretely, at the beginning of the game the learner initializes a set $S = \{x_1, x_2, \dots, x_{t_{\max}}\}$ consisting of the first $t_{\max} = 2^{\Omega(\sqrt{d})}$ instances in the sequence x selected by the adversary. This set represents the ‘danger zone’ – nodes in the beginning of the sequence that have not been labeled yet, that *might* be on-path, and that are not ancestors of a previously-labeled on-path node.¹⁶ To predict a label for an instance x_i , the learner selects a label \hat{y}_i such that if \hat{y}_i is wrong, the danger zone will shrink by at least $1/3$. That is, for $b \in \{0, 1\}$, if the set S_b of b -descendants of x_i has cardinality $|S_b| \geq |S|/3$, the learner predicts $\hat{y}_i = b$. Then, if the adversary selects $y_i = 1 - b$, that implies that all b -descendants of x_i are off-path for the true labeling functions. Therefore, the learner removes all b -descendants of x_i from the danger zone, and the new cardinality is $|S \setminus S_b| \leq (2/3) \cdot |S|$. This guarantees that the learner can make at most $O(\log t_{\max}) = O(\sqrt{d})$ such mistakes before the danger zone is empty.¹⁷

¹⁶If u is an ancestor of some on-path node v , and v is a b -descendant of u for $b \in \{0, 1\}$, then the true label for u must be b .

¹⁷Once the danger zone is empty, the learner cannot make any further on-path mistakes within the prefix $x_1, x_2, \dots, x_{t_{\max}}$. And it will make at most $O(\sqrt{d})$ mistakes on the remaining nodes $x_{t_{\max}+1}, x_{t_{\max}+2}, \dots$, as explained

If neither S_0 nor S_1 have cardinality at least $|S|/3$, the learner predicts $\hat{y}_i = 0$. If $y_i = 1$ and x_i is on-path for the true labeling function, then the learner updates the danger zone to be $S_0 \cup S_1$,¹⁸ again shrinking the danger zone by a factor of at most $2/3$. Otherwise, if $y_i = 1$ and x_i is off-path, then it was an off-path node labeled 1 (which is rare), and the learner can afford to misclassify it (see Section 2.3.6).

2.3.5 Splitting Experts

The danger zone minimization strategy requires that the learner know whether the node u being classified is on-path or off-path for the true labeling function. However, if u appears in the sequence before some of its ancestors, the learner does not know this. To overcome this difficulty, the learner implements a variant of the standard *multiplicative weights algorithm* using *splitting experts*. This means that initially there is a single expert executing danger zone minimization. When a node u is reached for which danger zone minimization requires knowing whether u is on-path or off-path and that information is not yet evident, each expert is split into two experts, one of which continues the execution of danger zone minimization under the assumption that u is on-path, and the other under the opposite assumption. Thus, at each point in time, there exists precisely one expert for which all path-related assumptions are correct, and therefore that expert will make at most $O(\sqrt{d})$ mistakes. The multiplicative weights algorithm guarantees that the overall number of mistakes will be linear in the the number of mistakes of the best expert, i.e., $O(\sqrt{d})$.

2.3.6 Transition to Halving

The hypothesis class is engineered such that it satisfies the following property: there are at most $2^{O(\sqrt{d})}$ functions in the hypothesis class that agree with any set of $t_{\max} = 2^{\Omega(\sqrt{d})}$ labels, or that agree that a set of $\Theta(\sqrt{d})$ nodes are all off-path and labeled 1 (this follows from Lemma 8.2).

Therefore, once the true labels for the first t_{\max} instances $x_1, x_2, \dots, x_{t_{\max}}$ have been revealed, or once $\Theta(\sqrt{d})$ off-path labels of 1 have been revealed (whichever happens first), the learner can *transition to halving*: stop doing danger zone minimization, and instead predict the labels for the remaining nodes using the standard Halving algorithm (Algorithm 7) on the subset of the hypothesis class that survived. Halving on $2^{O(\sqrt{d})}$ functions is guaranteed to make at most $O(\sqrt{d})$ mistakes (Fact A.1).

However, seeing as the learner lacks information on which nodes are off-path, it uses experts, and each expert maintains different path-related assumptions. Thus, each expert decides separately at which point to transition to Halving. The unique expert that makes only correct assumptions will transition ‘at the right time’. That expert will make at most $O(\sqrt{d})$ mistakes during danger zone minimization, and then at most $O(\sqrt{d})$ additional mistakes during halving.

in Section 2.3.6.

¹⁸Because on-path nodes must be either be descendants or ancestors of x_i , and the definition of the danger zone does not require that it contain ancestors of nodes that have been labeled.

2.4 Some Intuition for the Quantity \sqrt{d}

We briefly sketch where the quantity \sqrt{d} arises from. This is a back-of-the-envelope calculation without proof, intended purely as an aid for intuition. Suppose we assigned off-path labels of 1 with probability 2^{-k} instead of $2^{-\sqrt{d}}$. Consider a sequence x_1, \dots, x_n of $n = d/2k$ leaves. For any sequence of labels $y_1, \dots, y_n \in \{0, 1\}$, taking $s = \sum_{i \in [n]} y_i$, there exist roughly

$$2^d \cdot (2^{-k})^s \cdot (1 - 2^{-k})^{n-s} \geq 2^d \cdot (2^{-k})^n \gg 0$$

functions in the class for which these leaves are off-path and which agree with the labels y_1, \dots, y_n . Therefore, the adversary can force at least $n = \Omega(d/k)$ mistakes.

Similarly, for the sequence x_1, \dots, x_n consisting of all the nodes in the tree of depth at most $k/2$ in breadth-first order, the adversary can force a mistake on every on-path node while assigning a label of 0 to all off-path nodes, for a total of $k/2$ mistakes. This is true because for any assignment of on-path labels, the fraction of functions which agree with the on-path labels that assign a label of 0 to all off-path nodes is roughly

$$(1 - 2^{-k})^{2^{k/2}} \approx 1,$$

so in particular for any labeling of the on-path nodes there exists a function in the class that agrees with that labeling and assigns 0 to all off-path nodes.

Therefore, for any k , we obtain a *lower bound* of $\Omega(\frac{d}{k} + k)$ on the number of mistakes. For any k , $\frac{d}{k} + k \geq \sqrt{d}$, giving a lower bound of $\Omega(\sqrt{d})$. Choosing $k = \sqrt{d}$ to minimize the lower bound will in fact yield a matching upper bound of $O(\sqrt{d})$, as we show in this paper. This completes our overview of the upper bound.

3 Directions for Future Work

Following are some interesting open questions:

1. Does there exist an efficient learning algorithm that achieves the $O(\sqrt{d})$ upper bound of Theorem 8.1? One needs to be careful about the definition of efficiency here, but one possible formalization is as follows. Does there exist a learning algorithm A and a sequence of classes $\mathcal{H}_1, \mathcal{H}_2, \dots$, such that for every $d \in \mathbb{N}$:
 - $\text{LD}(\mathcal{H}_d) = d$, and
 - Given as input the index d and a sequence x_1, \dots, x_n , the algorithm A runs in time $\text{poly}(d, n)$ and makes at most $O(\sqrt{d})$ mistakes assuming the labels are realizable by \mathcal{H}_d .
2. Is there a tradeoff between the cardinality of the domain \mathcal{X} and the upper bound on the number of mistakes? We used a domain of size roughly 2^d in order to obtain our upper bound of $O(\sqrt{d})$. Is it possible to get the same bound with a domain of size $\text{poly}(d)$?
3. Obtaining more precise asymptotics; for example, is there (an explicit) constant $\alpha > 0$ such that the optimal transductive mistake bound is $(\alpha + o(1))\sqrt{d}$?

4 Organization

Formal definitions appear in Section 5. Formal statements and proofs for the lower bound and upper bound appear in Section 6 and Section 8, respectively. Optimal sequence length is discussed in Section 7.

5 Preliminaries

5.1 Basic Notation

Notation 5.1. $\mathbb{N} = \{1, 2, 3, \dots\}$, i.e., $0 \notin \mathbb{N}$. $\log(\cdot)$ and $\ln(\cdot)$ denote logarithm to base 2 and e , respectively.

Notation 5.2 (Sequences). Let \mathcal{X} be a set and $n, k \in \mathbb{N}$. For a sequence $x = (x_1, \dots, x_n) \in \mathcal{X}^n$, we write $x_{\leq k}$ to denote the subsequence (x_1, \dots, x_k) . If $k \leq 0$ then $x_{\leq k}$ denotes the empty sequence, which is also denoted by $\lambda = \mathcal{X}^0$. We use the notation $\mathcal{X}^{\leq n} = \cup_{i=0}^n \mathcal{X}^i$.

5.2 Standard Online Learning

Let \mathcal{X} be a set, and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a collection of functions called a *hypothesis class*. A *learner strategy* or simply *learner* for the standard online learning game (Game 1) is a function

$$L : \prod_{i=0}^{n-1} (\mathcal{X} \times \{0, 1\})^i \times \mathcal{X} \rightarrow \{0, 1\},$$

where $n \in \mathbb{N}$ is the number of rounds in the game. The set of all such learner strategies is denoted \mathcal{L}_n . An *adversary strategy* or simply *adversary* for the standard online learning game is a pair of functions

$$\begin{aligned} A_{\text{instance}} &: \prod_{i=0}^{n-1} (\mathcal{X} \times \{0, 1\} \times \{0, 1\})^i \rightarrow \mathcal{X}, \text{ and} \\ A_{\text{label}} &: \prod_{i=1}^{n-1} (\mathcal{X} \times \{0, 1\} \times \{0, 1\})^i \times \{0, 1\} \rightarrow \{0, 1\}. \end{aligned}$$

The set of all such adversary strategies is denoted \mathcal{A}_n .

Semantically, the interpretation of these strategies is that in each round $t \in [n]$ of Game 1, the adversary selects an instance

$$x_t = A_{\text{instance}}(x_1, \hat{y}_1, y_1, \dots, x_{t-1}, \hat{y}_{t-1}, y_{t-1}) \in \mathcal{X},$$

then the learner makes a prediction

$$\hat{y}_t = L(x_1, y_1, \dots, x_{t-1}, y_{t-1}, x_t) \in \{0, 1\},$$

and finally, the adversary assigns a label

$$y_t = A_{\text{label}}(x_1, \hat{y}_1, y_1, \dots, x_{t-1}, \hat{y}_{t-1}, y_{t-1}, \hat{y}_t) \in \{0, 1\}.$$

The adversary's function A_{label} must satisfy *realizability*, meaning that there exists $h \in \mathcal{H}$ such that

$$\forall t \in [n] : y_t = h(x_t).$$

The number of mistakes in a game with n rounds and hypothesis class \mathcal{H} between learner L and adversary A is

$$M_{\text{std}}(\mathcal{H}, n, L, A) = |\{t \in [n] : \hat{y}_t \neq y_t\}|.$$

5.3 Transductive Online Learning

Given \mathcal{X} and \mathcal{H} as in Section 5.2, a learner strategy for the *transductive online learning setting* (Game 2) is a function

$$L : \mathcal{X}^n \times \bigcup_{i=0}^{n-1} \{0, 1\}^i \rightarrow \{0, 1\},$$

where $n \in \mathbb{N}$ is the number of rounds in the game. An adversary strategy consists of a sequence $x \in \mathcal{X}^n$ and an *adversary labeling strategy*, which is a function

$$A : \left(\bigcup_{i=0}^{n-1} \{0, 1\}^{2^i} \right) \times \{0, 1\} \rightarrow \{0, 1\}.$$

The sets of all such learner and adversary strategies are denoted \mathcal{L}_n and \mathcal{A}_n respectively.

Semantically, the interpretation of these strategies is that at the start of Game 2, the adversary selects the sequence x . Then, in each round $t \in [n]$, the learner makes a prediction

$$\hat{y}_t = L(x, y_1, \dots, y_{t-1}) \in \{0, 1\},$$

and then the adversary assigns a label

$$y_t = A(\hat{y}_1, y_1, \dots, \hat{y}_{t-1}, y_{t-1}, \hat{y}_t) \in \{0, 1\}.$$

Exactly as in Section 5.2, the adversary's function A must satisfy realizability, namely,

$$\exists h \in \mathcal{H} \forall t \in [n] : y_t = h(x_t),$$

and the number of mistakes in a game with sequence length n and hypothesis class \mathcal{H} between learner L and adversary A is

$$M_{\text{tr}}(\mathcal{H}, n, L, A) = |\{t \in [n] : \hat{y}_t \neq y_t\}|.$$

5.4 Mistake Bounds

In this paper, we study *optimal mistake bounds*, or the *optimal number of mistakes*, which is the value of Games 1 and 2. For $M \in \{M_{\text{std}}, M_{\text{tr}}\}$, the optimal number of mistakes in a game with hypothesis class \mathcal{H} and sequence length n is,

$$M(\mathcal{H}, n) = \inf_{L \in \mathcal{L}_n} \sup_{A \in \mathcal{A}_n} M(\mathcal{H}, n, L, A).$$

The optimal number of mistakes for hypothesis class \mathcal{H} is

$$M(\mathcal{H}) = \sup_{n \in \mathbb{N}} M(\mathcal{H}, n).$$

Remark 5.3. As is common in learning theory literature, in both Game 1 and Game 2, we take the sets \mathcal{L}_n and \mathcal{A}_n to be the sets of all (deterministic) functions. In this paper, we do not consider randomized strategies. By allowing arbitrary functions, we ignore issues relating to computability.

5.5 Trees

Definition 5.4 (Notation for binary trees). *Let $d \in \mathbb{N} \cup \{0\}$. A perfect binary tree of depth d is a collection of $2^{d+1} - 1$ nodes, which we identify with the collection of binary strings*

$$T_d = \{\{0, 1\}^k : k \in \{0, 1, 2, \dots, d\}\}.$$

The empty string, denoted $\lambda = \{0, 1\}^0$, is a member of T_d and is called the root of the tree. Every string $u \in \{0, 1\}^d$ is called a leaf. The depth of a node $u \in T_d$, denoted $|u|$, is the length of u as a string, namely, the integer k such that $u \in \{0, 1\}^k$.

For two nodes $u, v \in T_d$, we say that u is a parent of v , and that v is a child of u , if $v = u \circ 0$ or $v = u \circ 1$, where \circ denotes string concatenation. More fully, for $b \in \{0, 1\}$, we say that v is a b -child of u if $v = u \circ b$.

Recursively, we define that u is an ancestor of v and that v is a descendant of u , and write $u \preceq v$, if one of the following holds:

- $u = v$, or
- $\exists w \in T_d \exists b \in \{0, 1\} : (u \preceq w) \wedge (w \circ b = v)$.

For $b \in \{0, 1\}$, we say that v is a b -descendant of u , denoted $u \preceq_b v$, if v is a descendant of the b -child of u .

A function $f : T_d \rightarrow \{0, 1\}$ specifies a particular root-to-leaf path in the tree T_d (see Figure 1). The on-path nodes for f are the set of $d + 1$ nodes on that root-to-leaf path, as in the following definition.

Definition 5.5 (Paths in a binary tree). *Let $d, k \in \mathbb{N}$, $k \leq d$. Let $u \in \{0, 1\}^k$ be a node in T_d . The path to u is the unique sequence $\text{path}(u) = (u_0, u_1, u_2, \dots, u_k)$ such that $u_0 = \lambda$ is the root, $u_k = u$, and u_i is a child of u_{i-1} for all $i \in [k]$.*

Let $f : T_d \rightarrow \{0, 1\}$ be a function. The path of f is the unique sequence $\text{path}(f) = (u_0, u_1, u_2, \dots, u_d)$ such that $u_0 = \lambda$ is the root, and for each $i \in [d]$, $u_i = u_{i-1} \circ f(u_{i-1})$. Namely, u_i is the $f(u_{i-1})$ -child of u_{i-1} .

For a node $v \in T_d$ and a function $f : T_d \rightarrow \{0, 1\}$, we write $v \in \text{path}(f)$ if $\text{path}(f) = (u_0, \dots, u_d)$ and there exists $i \in \{0, \dots, d\}$ such that $u_i = v$. Otherwise, we write $v \notin \text{path}(f)$.

For a node $v \in T_d$ and a set of functions $\mathcal{F} \subseteq \{0, 1\}^{T_d}$, we write $v \in \text{path}(\mathcal{F})$ if

$$\forall f \in \mathcal{F} : v \in \text{path}(f).$$

Otherwise, we write $v \notin \text{path}(\mathcal{F})$.

5.6 Littlestone Dimension

Definition 5.6 (Littlestone, 1987). Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$, and let $d \in \mathbb{N} \cup \{0\}$. We say that \mathcal{H} shatters the binary tree T_d if there exists a mapping $T_d \rightarrow \mathcal{X}$ given by $u \mapsto x_u$ such that for every $u \in \{0, 1\}^{d+1}$ there exists $h_u \in \mathcal{H}$ such that

$$\forall i \in [d + 1] : h(x_{u_{\leq i-1}}) = u_i.$$

The Littlestone dimension of \mathcal{H} , denoted $\text{LD}(\mathcal{H})$, is the supremum over all $d \in \mathbb{N}$ such that there exists a Littlestone tree of depth $d - 1$ that is shattered by \mathcal{H} .

Note that by defining the Littlestone dimension this way, every class with Littlestone dimension $d \in \mathbb{N}$ contains at least 2^d functions.




Figure 2: A shattered Littlestone tree of depth 2. The empty sequence is denoted by λ .

(Source: Bousquet et al., 2021)

Theorem 5.7 (Littlestone, 1987). Let \mathcal{X} be a set and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ such that $d = \text{LD}(\mathcal{H}) < \infty$. Then there exists a strategy for the learner that guarantees that the learner will make at most d mistakes in the standard (non-transductive) online learning setting, regardless of the adversary's strategy and of the number n of instances to be labeled. Furthermore, there exists an adversary that forces every learner to make at least $\min\{n, d\}$ mistakes.

6 Lower Bound

6.1 Statement

Our $\Omega(\sqrt{d})$ lower bound states the following.

Theorem 6.1 (Lower bound). There exists a constant $d_0 \geq 0$ as follows. Let $d \in \mathbb{N}$, $d \geq d_0$, let \mathcal{X} be a set, and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class with $\text{LD}(\mathcal{H}) = d$. Then there exist a sequence $x \in \mathcal{X}^n$ of length $n = O(d \cdot 2^{\sqrt{d}})$ and an adversary A that always selects the sequence x and uses a simple adaptive labeling strategy (as in Algorithm 1), such that for every learning rule L ,

$$M_{\text{tr}}(\mathcal{H}, n, L, A) \geq \sqrt{d}/10. \quad (3)$$

Furthermore, for every integer $n \in \mathbb{N}$,

$$M_{\text{tr}}(\mathcal{H}, n) \geq \min \{ \sqrt{d}/10, \lfloor \log(n+1) \rfloor \}. \quad (4)$$

Remark 6.2. The assumption $\text{LD}(\mathcal{H}) = d$ implies that for all $k \in [d]$, \mathcal{H} shatters a Littlestone tree of depth k . Thus, the lower bound of Eq. (3) in Theorem 6.1 immediately implies that for every $k \in [d]$ there exists a sequence $x^{(k)} \in \mathcal{X}^{n_k}$ of length $n_k = O(k \cdot 2^{\sqrt{k}})$ such that the adversary A_k that presents the sequence $x^{(k)}$ and assigns labels using the simple labeling strategy of Algorithm 1 ensures that for every learner L ,

$$M_{\text{tr}}(\mathcal{H}, n_k, L, A_k) \geq \sqrt{k}/10.$$

See Section 2.2 for a general overview of Theorem 6.1 and the main proof ideas. In the following subsections we prove Theorem 6.1. Algorithm 1 gives an explicit construction of the adversary that witnesses the lower bound, using Algorithm 2 as a subroutine. We start with presenting some initial observations about the behavior of these algorithms in Section 6.2.

Assumptions:

- $d \in \mathbb{N}$, $\varepsilon = 2^{-\sqrt{d}/2}$.
- $T = T_d$ is a perfect binary tree of depth d .
- $\mathcal{H} \subseteq \{0, 1\}^T$ is a class that shatters T .

TRANSDUCTIVEADVERSARY (\mathcal{H}):

$(x_1, x_2, \dots, x_n) \leftarrow \text{CONSTRUCTSEQUENCE}(\mathcal{H})$

▷ See Algorithm 2.

send (x_1, x_2, \dots, x_n) to learner

$\mathcal{H}_0 \leftarrow \mathcal{H}$

for $t \in [n]$:

receive \hat{y}_t from learner

$$r_t \leftarrow \frac{|\{h \in \mathcal{H}_{t-1} : h(x_t) = 1\}|}{|\mathcal{H}_{t-1}|}$$

$$y_{\text{maj}} \leftarrow \mathbb{1}(r_t \geq 1/2)$$

$$y_t \leftarrow \begin{cases} y_{\text{maj}} & r_t \notin [\varepsilon, 1 - \varepsilon] \\ 1 - \hat{y}_t & \text{otherwise} \end{cases}$$

send y_t to learner

$$\mathcal{H}_t \leftarrow \{h \in \mathcal{H}_{t-1} : h(x_t) = y_t\}$$

Algorithm 1: The strategy for the adversary that achieves the lower bound in Theorem 6.1. Note that while the construction of the sequence x is not entirely trivial, the adversary's strategy for labeling this sequence is very simple.

6.2 Analysis of the Adversary

Claim 6.3. *Let $d \in \mathbb{N}$, let $M = \sqrt{d}/10$, and let $\mathcal{H} \subseteq \{0, 1\}^{T_d}$ be a hypothesis class. Consider an execution of `CONSTRUCTSEQUENCE`(\mathcal{H}) as in Algorithm 2 that produces a sequence x_1, x_2, \dots, x_n . Then:*

- (a) *For all $i \in [n]$, $\text{path}(x_i)$ is a subsequence of x_0, x_1, \dots, x_i .*
- (b) *The length n of the sequence satisfies $n < n_d$, where $n_d = (d + 1) \cdot 2^{M+1}$.*

Proof.

- (a) Fix $i \in [n]$. It suffices to show that for all $u \in T_d$, if $u \preceq x_i$ then $u \in (x_1, x_2, \dots, x_i)$. Proceed by induction on i . For the base case $i = 1$, the claim holds because $x_1 = \lambda$.

For the induction step, assume the claim holds for $i \in [n - 1]$. Let $u \preceq x_{i+1}$, we prove that $u \in (x_1, x_2, \dots, x_{i+1})$. Assume $x_{i+1} \neq \lambda$ (otherwise, there is nothing to prove).

Because x_{i+1} appears in the sequence x , it must have been added to Q before it was added to x . The only place where items that are not λ are added to Q is in the line $Q \leftarrow Q \cup \{x_t \circ y\}$. Namely, there exist an index $j \in [i]$ and a bit $y \in \{0, 1\}$ such that $x_{i+1} = x_j \circ y$ (note that $j < i + 1$ because x_j was added to the sequence before x_{i+1}). If $x_j = u$ we are done. Otherwise, note that x_j is the parent of x_{i+1} , and therefore $u \preceq x_j$. By the induction hypothesis, $u \in (x_1, x_2, \dots, x_j)$. This concludes the proof.

- (b) Items are added to the sequence x only if they were previously added to Q . By induction on $i \in [n]$, for each x_i in the sequence, there is at most one iteration of the “while $|Q| > 0$ ” loop in which x_i is added to Q . The base case $i = 1$ holds because $x_1 = \lambda$ is the root, which is added to Q before the while loop, and λ is never added to Q within that loop because the line “ $Q \leftarrow Q \cup \{x_t \circ y\}$ ” can only add non-empty bit strings. For the induction step, if the claim holds for all natural numbers j such that $1 \leq j < i \leq n$ then it holds for i . Indeed, for $i \geq 2$, x_i can be added to Q only via the line “ $Q \leftarrow Q \cup \{x_t \circ y\}$ ”, and only in the iteration of the while loop where x_t is the parent of x_i in the tree T_d . In that iteration, the parent x_t of x_i is popped from Q , which implies that x_t was added to Q in some previous iteration of the while loop ($t < i$), and is no longer in Q after being popped. By the induction hypothesis, x_t will never be added to Q again, and therefore in all subsequent iterations of the while loop x_t will not be the parent of x_i , so x_i cannot be added to Q in subsequent iterations via the line “ $Q \leftarrow Q \cup \{x_t \circ y\}$ ”.

Furthermore, if a node x_i is added to Q in some iteration of the while loop, then it remains in Q for the duration of that iteration. So for all $i \in \{2, 3, \dots, n\}$, there is precisely one execution of the line “ $Q \leftarrow Q \cup \{x_t \circ y\}$ ” that adds x_i to Q . Namely, there is precisely one point in time during the execution of Algorithm 2 in which $x_i = x_t \circ y$, $x_i \notin Q$, and the line “ $Q \leftarrow Q \cup \{x_t \circ y\}$ ” is executed resulting in $x_i \in Q$.

Consider a function f that maps $i \in \{2, 3, \dots, n\}$ to the value of the index b' during the unique execution of the line “ $Q \leftarrow Q \cup \{x_t \circ y\}$ ” that adds x_i to Q . Namely, if b' had some value β when x_i was added to Q , then $f(i) = \beta$.

Notice that “ $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{x_t \circ y\}$ ” is executed only if the condition $x_t \in \text{path}(\mathcal{H}_{b'})$ is satisfied in the previous line. Furthermore, the line “ $\mathcal{H}_{b'} \leftarrow \{h \in \mathcal{H}_b : h(x_t) = y\}$ ” ensures that the node $x_i = x_t \circ y$ being added to \mathcal{Q} satisfies $x_t \circ y \in \text{path}(\mathcal{H}_{b'})$, namely

$$\forall h \in \mathcal{H}_{b'} : x_i \in \text{path}(h).$$

Consequently, $x_i \in \text{path}(\mathcal{G})$ for any class \mathcal{G} that is a subset of $\mathcal{H}_{b'}$; in particular, because the only way that $\mathcal{H}_{b'}$ might be modified later during the execution of Algorithm 2 is by removing elements, it follows that $x_i \in \text{path}(\mathcal{H}_{b'})$ when the line “ $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{x_t \circ y\}$ ” is executed and in all subsequent times.

However, $|\text{path}(\mathcal{G})| = d + 1$ for any class $\mathcal{G} \subseteq \{0, 1\}^{T_d}$. This implies that f maps at most $(d + 1)$ nodes to each bit string. In other words, for any bit string b , the size of the preimage satisfies $|f^{-1}(b)| \leq d + 1$.

The condition “ $|b| < M$ ” in Algorithm 2 ensures that $|b'| \leq M$, namely, $b' \in \{0, 1\}^k$ for $k \in \{0, 1, 2, \dots, M\}$. Thus,

$$\begin{aligned} n &= 1 + |\{2, 3, \dots, n\}| \\ &= 1 + \sum_{\substack{b \in \{0, 1\}^k \\ k \in \{0, \dots, M\}}} |\{i \in \{2, 3, \dots, n\} : f(i) = b\}| \\ &= 1 + \sum_{\substack{b \in \{0, 1\}^k \\ k \in \{0, \dots, M\}}} |f^{-1}(b)| \\ &\leq 1 + \sum_{\substack{b \in \{0, 1\}^k \\ k \in \{0, \dots, M\}}} (d + 1) \\ &\leq 1 + (d + 1) \cdot (2^{M+1} - 1). \\ &< (d + 1) \cdot 2^{M+1}, \end{aligned}$$

as desired. □

Claim 6.4. *Let $d \in \mathbb{N}$, let $M = \sqrt{d}/10$, and let $\mathcal{H} \subseteq \{0, 1\}^{T_d}$ be a hypothesis class. Consider an execution of `TRANSDUCTIVEADVERSARY` (\mathcal{H}) as in Algorithm 1. Let*

$$\mathcal{H}_0, \mathcal{H}_1, \dots, \mathcal{H}_n$$

be the sequence of hypothesis classes created by `TRANSDUCTIVEADVERSARY`, let

$$S = \{t \in [n] : r_t \in [\varepsilon, 1 - \varepsilon]\}$$

be the set of indices where `TRANSDUCTIVEADVERSARY` forces a mistake, and let

$$\mathbb{H}_0, \mathbb{H}_1, \dots, \mathbb{H}_n$$

be the sequence of collections created by the subroutine `CONSTRUCTSEQUENCE` (Algorithm 2). If $|S| \leq M$ then

$$\forall t \in \{0, 1, \dots, n\} : \mathcal{H}_t \in \mathbb{H}_t.$$

Proof. Proceed by induction on $t \in \{0, 1, \dots, n\}$. The base case $t = 0$ is satisfied, because $\mathcal{H}_0 = \mathcal{H} \in \{\mathcal{H}\} = \mathbb{H}_0$. For the induction step, assume that $\mathcal{H}_{i-1} \in \mathbb{H}_{i-1}$ for some $i \in [n]$. We prove that $\mathcal{H}_i \in \mathbb{H}_i$.

Let y_i be the label assigned to x_i by `TRANSDUCTIVEADVERSARY`. Then

$$\mathcal{H}_i = \{h \in \mathcal{H}_{i-1} : h(x_i) = y_i\}.$$

Consider the iteration of the while loop in `CONSTRUCTSEQUENCE` that starts with $t \leftarrow i$. By the induction hypothesis, $\mathcal{H}_{i-1} \in \mathbb{H}_{i-1}$. Therefore, in this iteration of the while loop, there will be an iteration of the “for $\mathcal{H}_b \in \mathbb{H}_{t-1}$ ” loop where $\mathcal{H}_b = \mathcal{H}_{i-1}$. In that iteration, $y_i \in \mathcal{Y}$ by construction of y_i and \mathcal{Y} . Therefore, in the iteration of the “for $y \in \mathcal{Y}$ ” loop in which $y = y_i$,

$$\mathcal{H}_{b'} = \{h \in \mathcal{H}_b : h(x_t) = y\} = \{h \in \mathcal{H}_{i-1} : h(x_i) = y_i\} = \mathcal{H}_i.$$

The class $\mathcal{H}_{b'}$ is then added to $\mathbb{H}_i = \mathbb{H}_t$ in the line “ $\mathbb{H}_t \leftarrow \mathbb{H}_t \cup \{\mathcal{H}_{b'}\}$ ”. Furthermore, no class is ever removed from \mathbb{H}_t . So $\mathcal{H}_i \in \mathbb{H}_i$, as desired. \square

Claim 6.5. *Let $d \in \mathbb{N}$, let $M = \sqrt{d}/10$, and let $\mathcal{H} \subseteq \{0, 1\}^{T_d}$ be a hypothesis class. Consider an execution of `TRANSDUCTIVEADVERSARY`(\mathcal{H}) as in Algorithm 1 where the adversary constructs a sequence of nodes $x_1, x_2, \dots, x_n \in T_d$ and a sequence of classes $\mathcal{H}_0, \mathcal{H}_1, \dots, \mathcal{H}_n \subseteq \{0, 1\}^{T_d}$. Let*

$$S = \{t \in [n] : r_t \in [\varepsilon, 1 - \varepsilon]\}$$

be the set of indices where `TRANSDUCTIVEADVERSARY` forces a mistake, and assume that $|S| \leq M$. Then for all $k \in \{0, 1, \dots, d\}$ there exists $i \in [n]$ such that

1. $|x_i| = k$, and
2. $x_i \in \text{path}(\mathcal{H}_{i-1})$,

Proof. Proceed by induction on k . For the base case $k = 0$, notice that $x_1 = \lambda$, $|\lambda| = 0$, and $\lambda \in \text{path}(\mathcal{H}_{-1})$.

For the induction step, assume the claim holds for some $k \in \{0, 1, \dots, d-1\}$, and take $i_k \in [n]$ such that $|x_{i_k}| = k$ and $x_{i_k} \in \text{path}(\mathcal{H}_{i_k-1})$; we prove that the claim holds for $k+1$ as well.

Consider the iteration of the while loop in `CONSTRUCTSEQUENCE` in which x_{i_k} is added to the sequence (i.e., the iteration starting with $t \leftarrow i_k$). By Claim 6.4 and the assumption $|S| \leq M$, $\mathcal{H}_{i_k-1} \in \mathbb{H}_{i_k-1}$. Hence, within this iteration of the while loop, there is an iteration of the “for $\mathcal{H}_b \in \mathbb{H}_{t-1}$ ” loop such that $\mathcal{H}_b = \mathcal{H}_{i_k-1}$. By construction, the set \mathcal{Y} always contains the label predicted by the adversary, so $y_{i_k} \in \mathcal{Y}$. Consider the iteration of the “for $y \in \mathcal{Y}$ ” loop such that $y = y_{i_k}$. By the induction hypothesis, $x_i \in \text{path}(\mathcal{H}_{i_k-1})$, and since $\mathcal{H}_{b'} \subseteq \mathcal{H}_b = \mathcal{H}_{i_k-1}$, it follows that $x_{i_k} \in \text{path}(\mathcal{H}_{b'})$. Seeing as $|x_{i_k}| < d$, in the last line of this iteration of the “for $y \in \mathcal{Y}$ ” loop, the node $x_{i_{k+1}} := x_{i_k} \circ y_{i_k}$ is added to \mathcal{Q} . This guarantees that $x_{i_{k+1}}$ will eventually be popped from \mathcal{Q} and added to the sequence returned by `CONSTRUCTSEQUENCE`. Once a node has been added to the sequence, it is never removed.

Notice that $|x_{i_{k+1}}| = |x_{i_k}| + 1 = k + 1$, satisfying Item 1. Therefore, it remains to show Item 2, namely, to show that $x_{i_{k+1}} \in \text{path}(\mathcal{H}_{i_{k+1}-1})$.

Indeed, by the induction hypothesis, $x_i \in \text{path}(\mathcal{H}_{i_{k-1}})$, and in the iteration of the “for $y \in \mathcal{Y}$ ” discussed above, $\mathcal{H}_b = \mathcal{H}_{i_{k-1}}$, $\mathcal{H}_{b'} = \mathcal{H}_{i_k}$, and $\mathcal{H}_{b'} = \{h \in \mathcal{H}_b : h(x_{i_k}) = y_{i_k}\}$. Hence,

$$\forall h \in \mathcal{H}_{i_k} : x_{i_k} \in \text{path}(h) \wedge h(x_{i_k}) = y_{i_k}.$$

Seeing as $x_{i_{k+1}} = x_{i_k} \circ y_{i_k}$ This implies that

$$\forall h \in \mathcal{H}_{i_k} : x_{i_{k+1}} \in \text{path}(h).$$

Item 2 follows from the inclusion $\mathcal{H}_{i_{k+1}-1} \subseteq \mathcal{H}_{i_k}$. □

6.3 Proof

Finally, we complete the proof of the lower bound.

Proof of Theorem 6.1. Fix $d_0 = 800$ and assume $d \geq d_0$. Seeing as $\text{LD}(\mathcal{H}) = d$, \mathcal{H} shatters the tree T_d . By replacing \mathcal{H} with a suitable subset of \mathcal{H} of cardinality 2^{d+1} , renaming the elements in the domain of \mathcal{H} to nodes of T_d , and restricting the domain of each function in \mathcal{H} to T_d , assume without loss of generality that $\mathcal{H} \subseteq \{0, 1\}^{T_d}$, $|\mathcal{H}| = 2^{d+1}$, and \mathcal{H} shatters T_d .

Consider the loop “for $t \in [n]$ ” in Algorithm 1, and let

$$S = \{s_1, s_2, \dots, s_m\} = \{t \in [n] : r_t \in [\varepsilon, 1 - \varepsilon]\}$$

be the set of indices where the adversary forces a mistake, such that the learner makes at least $m = |S|$ mistakes. Let $M = \sqrt{d}/10$, and assume for contradiction that $m \leq M$.

By Claim 6.5, there exists $t \in [n]$ such that $|x_t| = d$ (i.e., x_t is a leaf in T_d) and $x_t \in \text{path}(\mathcal{H}_{t-1})$, namely,

$$\forall h \in \mathcal{H}_{t-1} : x_t \in \text{path}(h).$$

Seeing as x_t is a leaf,

$$\forall h \in \mathcal{H}_{t-1} : \text{path}(x_t) = \text{path}(h). \tag{5}$$

By construction,

$$\mathcal{H}_t \subseteq \left\{ h \in \mathcal{H} : (\forall i \in [t] : h(x_i) = y_i) \right\},$$

and \mathcal{H}_t is not empty. Fix some $h^* \in \mathcal{H}_t \subseteq \mathcal{H}_{t-1}$. By Item (a) in Claim 6.3, $\text{path}(x_t) = \text{path}(h^*)$ is a subsequence of x_1, x_2, \dots, x_t , so

$$\forall h \in \mathcal{H}_t \forall x \in \text{path}(h^*) : h(x) = h^*(x).$$

Seeing as \mathcal{H} shatters T_d and $|\mathcal{H}| = 2^{k+1}$, if two functions $h, h^* \in \mathcal{H}$ agree on the labels for all nodes in $\text{path}(h^*)$, then $h = h^*$. We conclude that $\mathcal{H}_t = \{h^*\}$ and $|\mathcal{H}_t| = 1$.

Consider the loop “for $t \in [n]$ ” in Algorithm 1. For each $t \in [n]$,

$$|\mathcal{H}_t| \geq \begin{cases} \varepsilon \cdot |\mathcal{H}_{t-1}| & t \in S \\ (1 - \varepsilon) \cdot |\mathcal{H}_{t-1}| & t \notin S. \end{cases}$$

Hence,

$$\begin{aligned}
1 &= |\mathcal{H}_t| \\
&\geq \varepsilon^m \cdot (1 - \varepsilon)^{n-m} \cdot |\mathcal{H}_0| \\
&= \varepsilon^m \cdot (1 - \varepsilon)^{n-m} \cdot 2^{d+1} \\
&\geq \varepsilon^m \cdot (1 - \varepsilon)^n \cdot 2^{d+1} \\
&\geq \varepsilon^m \cdot (1 - \varepsilon)^{n_d} \cdot 2^{d+1} && \text{(by Item (b) in Claim 6.3.)} \\
&\geq \varepsilon^m \cdot 2^d = 2^{-m\sqrt{d}/2+d},
\end{aligned} \tag{6}$$

where the final line holds because $\varepsilon = 2^{-\sqrt{d}/2}$, $n_d = (d + 1) \cdot 2^{\sqrt{d}/10+1}$, and

$$(1 - \varepsilon)^{n_d} = \left(1 - 2^{-\sqrt{d}/2}\right)^{(d+1) \cdot 2^{\sqrt{d}/10+1}} \geq \frac{1}{2}$$

for our choice of $d \geq 800$. Rearranging Eq. (6) yields

$$2\sqrt{d} \leq m.$$

This is a contradiction to the assumption $m \leq M = \sqrt{d}/10$. We conclude that an adversary A following Algorithm 1 satisfies

$$\inf_{L \in \mathcal{L}_n} M_{\text{tr}}(\mathcal{H}, n, L, A) \geq m > M = \sqrt{d}/10, \tag{7}$$

as desired.

To establish the “furthermore” part of the theorem, fix a length $n \in \mathbb{N}$. Let k be the largest integer such that $2^{\lceil \sqrt{k}/10 \rceil} \leq n + 1$ and $k \leq d$. By Eq. (7), there exists some sequence on which the adversary can force every learning rule to make at least $\lceil \sqrt{k}/10 \rceil$ mistakes. By Theorem 7.2, this implies that there exists a sequence of length $2^{\lceil \sqrt{k}/10 \rceil} - 1 \leq n$ on which the adversary can force every learning rule to make at least $\lceil \sqrt{k}/10 \rceil = \min \{ \lceil \sqrt{d}/10 \rceil, \lceil \log(n + 1) \rceil \}$ mistakes. Namely,

$$M_{\text{tr}}(\mathcal{H}, n) \geq \min \{ \lceil \sqrt{d}/10 \rceil, \lceil \log(n + 1) \rceil \},$$

as in Eq. (4). □

7 Sequence Length

In this section, we show that if there exists a sequence on which the adversary can force M mistakes, then a sequence of length $2^M - 1$ is sufficient, and this upper bound is tight for some classes.¹⁹

¹⁹Of course, there also exist classes for which a shorter sequence is sufficient. For instance, if the class shatters (in the VC sense) a subset of the domain of cardinality M , then a sequence of length M suffices.

Definition 7.1 (Minimal sequence). Let \mathcal{X} be a set, let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a class, and let $M \in \mathbb{N}$. The minimal sequence length for forcing M mistakes for the class \mathcal{H} , denoted $\text{MinLen}(\mathcal{H}, M)$ is

$$\text{MinLen}(\mathcal{H}, M) = \inf \{n \in \mathbb{N} : (\exists x \in \mathcal{X}^n : M_{\text{tr}}(\mathcal{H}, x) \geq M)\}.$$

In words, $\text{MinLen}(\mathcal{H}, M)$ is the smallest integer n for which there exists a sequence of length n on which the adversary can force at least n mistakes; if no such sequence exists, then $\text{MinLen}(\mathcal{H}, M) = \infty$.

Theorem 7.2 (Minimal sequence bound). Let \mathcal{X} be a set, and fix $M \in \mathbb{N}$. Then for any class $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$, if $\text{MinLen}(\mathcal{H}, M) < \infty$ then

$$\text{MinLen}(\mathcal{H}, M) \leq 2^M - 1.$$

Furthermore, there exists a class $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ for which $\text{MinLen}(\mathcal{H}, M) = 2^M - 1$.

Theorem 7.2 is a corollary of the tree rank characterization of M_{tr} from Ben-David et al. (1997). For completeness, we present a direct proof of Theorem 7.2 that does not directly invoke that characterization. Roughly, given an adversary A_0 that forces every learner to make at least M mistakes on a (possibly long) sequence x , we apply two modifications to obtain new adversaries

$$A_0 \rightsquigarrow A_1 \rightsquigarrow A_2.$$

A_1 forces M mistakes and has a specific structure that we call ‘rigidity’, but it still uses the same (possibly long) sequence x . Capitalizing on the rigid structure, A_2 selects a subsequence of x of length at most $2^M - 1$, and forces M mistakes on that subsequence.

7.1 Rigid Adversary

Definition 7.3 (Rigid adversary). Let $n \in \mathbb{N}$, let \mathcal{X} be a set, and let

$$A : \left(\bigcup_{k=0}^{n-1} \{0, 1\}^{2^k} \right) \times \{0, 1\} \rightarrow \{0, 1\}$$

be an adversary strategy for some fixed sequence $x \in \mathcal{X}^n$. We say that A is rigid if there exists a function

$$f : \bigcup_{k=0}^{n-1} \{0, 1\}^k \rightarrow \{0, 1, \star\}$$

such that for all $k \in \{0, 1, \dots, n-1\}$ and all $y, \hat{y} \in \{0, 1\}^k$,

$$A(\hat{y}_1, y_1, \dots, \hat{y}_k, y_k, \hat{y}_{k+1}) = \begin{cases} f(y_1, \dots, y_k) & f(y_1, \dots, y_k) \in \{0, 1\} \\ 1 - \hat{y}_{k+1} & f(y_1, \dots, y_k) = \star \end{cases}.$$

Note that if an adversary is rigid, then the function f that witnesses this is uniquely determined.

Claim 7.4 (Rigid adversary exists). Let $n, M \in \mathbb{N}$, let \mathcal{X} be a set, let $x \in \mathcal{X}^n$, and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a class. Let A be an adversary strategy that forces every learner to make at least M mistakes on x . Then there exists an adversary strategy A^* such that:

1. A^* forces every learner to make at least M mistakes on x and A^* is rigid.
2. Let f be the function that witnesses the rigidity of A^* . Then for every $y \in \{0, 1\}^n$, the sequence

$$f(y_{\leq 0}), f(y_{\leq 1}), f(y_{\leq 2}), \dots, f(y),$$

has at least M members equal to \star .

Proof of Claim 7.4. For Item 1, consider the adversary strategy A^* that simulates an execution of A , as in Algorithm 3. In broad strokes, A^* functions as a middle-man between the learner and A . As the learner makes a sequence of predictions $\hat{y} \in \{0, 1\}^n$, the adversary A^* generates a sequence of (possibly different) predictions $\tilde{y} \in \{0, 1\}^n$, and sends those to the adversary A . Adversary A sees only the predictions \tilde{y} , and assigns labels $y \in \{0, 1\}^n$, which are relayed back to the learner by A^* with no modifications.

First, observe that A^* satisfies the realizability requirement. Indeed, A^* simulates an execution of A such that the sequence of labels y_1, \dots, y_n sent by A^* to the learner is exactly the sequence of labels selected by A . Seeing as A is realizable, every sequence of labels selected by A is realizable, and therefore every sequence of labels selected by A^* must be realizable as well.

Second, observe that A^* forces every learner to make at least M mistakes. To see this, notice that in Algorithm 3,

$$\sum_{t \in [n]} \mathbb{1}(\tilde{y}_t \neq y_t) \geq M. \quad (8)$$

Indeed, A forces every learner to make at least M mistakes, and in particular this applies to a learner that makes predictions \tilde{y} as in the simulation. Furthermore, observe that A^* only alters the predictions it receives from the learner in cases when it selects a label that is accepted by A , namely,

$$\forall t \in [n] : \tilde{y}_t \neq \hat{y}_t \implies \tilde{y}_t = y_t. \quad (9)$$

Therefore, if $E = \{t \in [n] : \tilde{y}_t = \hat{y}_t\}$, then

$$\begin{aligned} \sum_{t \in [n]} \mathbb{1}(\tilde{y}_t \neq y_t) &= \sum_{t \in E} \mathbb{1}(\tilde{y}_t \neq y_t) + \sum_{t \in [n] \setminus E} \mathbb{1}(\tilde{y}_t \neq y_t) \\ &= \sum_{t \in E} \mathbb{1}(\tilde{y}_t \neq y_t) + 0 && \text{(By Eq. (9))} \\ &= \sum_{t \in E} \mathbb{1}(\hat{y}_t \neq y_t) && \text{(Definition of } E\text{)} \\ &\leq \sum_{t \in [n]} \mathbb{1}(\hat{y}_t \neq y_t). && (10) \end{aligned}$$

Combining Eqs. (8) and (10) implies that A forces at least M mistakes.

Third, we show that A^* is rigid. We claim that there exists a function $g : \{0, 1\}^{\leq n-1} \rightarrow \{0, 1\}^{\leq n-1}$ such that for every $t \in \{0, 1, 2, \dots, n-1\}$,

$$(\tilde{y}_1, \dots, \tilde{y}_t) = g(y_1, \dots, y_t).$$

Proceed by induction on t . For the base case $t = 0$ there is nothing to prove. For the induction step, we assume the claim holds for some $t = k < n - 1$, and show that it holds for $t = k + 1$. From Algorithm 3, \tilde{y}_{k+1} satisfies

$$\tilde{y}_{k+1} = \begin{cases} 0 & A(\tilde{y}_1, y_1, \dots, \tilde{y}_k, y_k, 0) = 0 \\ 1 & A(\tilde{y}_1, y_1, \dots, \tilde{y}_k, y_k, 0) = A(\tilde{y}_1, y_1, \dots, \tilde{y}_k, y_k, 1) = 1 \\ 1 - y_{k+1} & \text{otherwise} \end{cases} . \quad (11)$$

The first two cases in Eq. (11) are immediate from Algorithm 3, and the remaining case occurs when A forces a mistake at time $k + 1$, namely, when A selects $y_{k+1} = 1 - \tilde{y}_{k+1}$. Thus, \tilde{y}_{k+1} is a function of $y_{\leq k+1}$ and $\tilde{y}_{\leq k}$. By the induction hypothesis, $\tilde{y}_{\leq k} = g(y_{\leq k})$, so \tilde{y}_{k+1} is simply a function of $y_{\leq k+1}$. This establishes the existence of the desired function g .

Hence, A^* is rigid, as witnessed by the function

$$f(y_1, \dots, y_k) = \begin{cases} 0 & A(\tilde{y}_1, y_1, \dots, \tilde{y}_k, y_k, 0) = 0 \\ 1 & A(\tilde{y}_1, y_1, \dots, \tilde{y}_k, y_k, 0) = A(\tilde{y}_1, y_1, \dots, \tilde{y}_k, y_k, 1) = 1 \\ \star & \text{otherwise} \end{cases} ,$$

where f is a well-defined function because $\tilde{y}_{\leq k} = g(y_{\leq k})$.

We have seen that A^* is a valid (realizable) adversary that forces every learner to make at least M mistakes, and it is rigid. This concludes the proof of Item 1.

Finally, For Item 2, note that $\tilde{y}_t \neq y_t$ only if A forces a mistake at time t in the sense that A selects $y_t = 1 - b$ for any prediction $b \in \{0, 1\}$ provided at time t . If A forces a mistake at time t , then A^* forces a mistake at time t as well. Therefore, if $\tilde{y}_t \neq y_t$, then $f(y_{<t}) = \star$, namely, \tilde{y}_t makes mistakes only when the value of f is \star . By Eq. (8), \tilde{y}_t makes at least M mistakes throughout the game, so there must be at least M rounds where f outputs \star , as desired. \square

7.2 Essential Indices

Definition 7.5. Let $n, M \in \mathbb{N}$, let \mathcal{X} be a set, let $x \in \mathcal{X}^n$, and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a class. Let A be a rigid adversary strategy witnessed by function f . We say that an index $t \in [n]$ is essential for A for forcing M mistakes on x if there exists a sequence $y \in \{0, 1\}^{t-1}$ such that $f(y) = \star$ and the sequence

$$f(y_{\leq 0}), f(y_{\leq 1}), f(y_{\leq 2}), \dots, f(y_{\leq t-1})$$

contains at most $M - 1$ members equal to \star .

Claim 7.6. Let $n, M \in \mathbb{N}$, let \mathcal{X} be a set, let $x \in \mathcal{X}^n$, and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a class. Let A be a rigid adversary strategy. Then $[n]$ contains at most $2^M - 1$ indices that are essential for A for forcing M mistakes on x .

Proof. For each essential index $t \in [n]$, there exists a label sequence $y \in \{0, 1\}^{t-1}$ that witnesses that t is essential, as in Definition 7.5. Each label sequence y is a witness for at most one index (the index $|y| + 1$), so it suffices to show that the set $Y \subseteq \{0, 1\}^{\leq n-1}$ of all witness label sequences is of cardinality at most $2^M - 1$.

Think of Y as a collection of nodes in the binary tree T_{n-1} (Definition 5.4). By Definition 7.5, if $y \in Y$, then the collection of all ancestors of y in Y has cardinality

$$|\{y_{\leq i} : i \in \{0, 1, 2, \dots, |y| - 1\}\} \cap Y| \leq M - 1.$$

Namely, Y is a subtree of depth at most $d = M - 1$ in the binary tree T_{n-1} .²⁰ Hence, the number of nodes in Y is at most

$$2^{d+1} - 1 = 2^M - 1,$$

as desired. □

7.3 Proof

Proof of Theorem 7.2. If $\text{MinLen}(\mathcal{H}, M) < \infty$, then there exist a sequence $x \in \mathcal{X}^n$, and an adversary A_0 that forces every learner to make at least M mistakes on x . By Claim 7.4, there exists a rigid adversary A_1 that causes every learner to make at least M mistakes on x ,²¹ and also satisfies Item 2 in Claim 7.4. Let f be the function that witnesses the rigidity of A_1 . By Claim 7.6, the set $I \subseteq [n]$ of indices that are essential for A_1 for forcing M mistakes on x has cardinality $k = |I| \leq 2^M - 1$.

Algorithm 4 defines a new adversary, A_2 , which forces every learner to make at least M mistakes on a sequence of length k . A_2 is realizable, because A_1 is realizable.²²

To see that adversary A_2 forces every learner to make at least M mistakes, let y_1, \dots, y_n be the sequence of labels assigned by A_2 . Seeing as A_2 assigns the same labels as A_1 , and A_1 satisfies Item 2 in Claim 7.4, it follows that there are at least M indices $j \in [n]$ such that $f(y_{\leq j-1}) = \star$. Fix $J \subseteq [n]$ to be the first M such indices. Then $J \subseteq I$, namely, all the indices in J are essential for A_1 for forcing M mistakes on x (Definition 7.5).

Therefore, for each $j \in J$, A_2 includes the instance x_j in the sequence of length k sent to the learner. Then, in round j of the n rounds simulated by A_2 :

- The learner makes a prediction $\hat{y}_j \in \{0, 1\}$ corresponding to instance x_j .
- Adversary A_2 sends prediction \hat{y}_j to adversary A_1 . Because $f(y_{\leq j-1}) = \star$, adversary A_1 assigns the label $y_j = 1 - \hat{y}_j$. Adversary A_2 then sends that label y_j to the learner. So the learner makes a mistake on x_j .

Hence, the learner makes at least $|J| = M$ mistakes, as desired. □

8 Upper Bound

8.1 Statement

The following result states that the lower bound of Theorem 6.1 is tight for some classes.

²⁰The depth of a subtree is s if the longest root-to-node path contains $s + 1$ nodes from the subtree.

²¹This is Item 1 in Claim 7.4.

²²The argument for realizability is the same as in the proof of Claim 7.4.

Theorem 8.1 (Upper bound, and separation between standard and transductive online learning). *For every integer $d \geq 43$, there exists a hypothesis class $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ with a domain \mathcal{X} of size $|\mathcal{X}| = 2^d - 1$ such that $\text{LD}(\mathcal{H}) = d$ and the following two conditions hold for all $n \in \mathbb{N}$:*

1. $M_{\text{tr}}(\mathcal{H}, n) \leq 48 \cdot \sqrt{d}$.
2. $M_{\text{std}}(\mathcal{H}, n) = \min\{n, d\}$.

8.2 Hypothesis Class

In this section we construct the hypothesis class for Theorem 8.1.

Lemma 8.2. *Let $d \in \mathbb{N}$, $d \geq 42$. Let T_d be a perfect binary tree of depth d , as in Definition 5.4. Then there exists a collection of functions $\mathcal{H} \subseteq \{0, 1\}^{T_d}$ such that $\text{LD}(\mathcal{H}) = d + 1$ and the following two conditions hold for all $H \subseteq \mathcal{H}$ and all $X \subseteq T_d$:*

1. *If $\forall h \in H \forall x \in X : x \notin \text{path}(h) \wedge h(x) = 0$, then $\min\{|H|, |X|\} < 2^{2\sqrt{d}}$.*
2. *If $\forall h \in H \forall x \in X : x \notin \text{path}(h) \wedge h(x) = 1$, then $|H| < 2^{2\sqrt{d}}$ or $|X| < 3\sqrt{d}$.*

The proof employs the probabilistic method, showing that a hypothesis class sampled randomly from a suitable distribution has the desired properties with very high probability.

Proof. Let \mathcal{P} be a probability distribution over hypothesis classes. Formally, $\mathcal{P} \in \Delta(\{0, 1\}^{T_d})^{2^{d+1}}$ is a distribution over vectors of hypotheses. Each vector $\mathcal{H} \in \text{supp}(\mathcal{P})$ consists of 2^{d+1} hypotheses,

$$\mathcal{H} = (h_b)_{b \in \{0, 1\}^{d+1}},$$

where for each $b \in \{0, 1\}^{d+1}$, hypothesis h_b is a function $h_b : T_d \rightarrow \{0, 1\}$ sampled independently as follows:

- For each $i \in [d] \cup \{0\}$: $h_b(b_{\leq i}) = b_{i+1}$. (In particular, with probability 1, $\text{path}(h_b) = (b_{\leq 0}, b_{\leq 1}, \dots, b_{\leq d})$, each entry in the vector \mathcal{H} is unique, and \mathcal{H} shatters T_d .)
- For each $x \in T_d \setminus \text{path}(h_b)$, the bit $h_b(x) \in \{0, 1\}$ is sampled $\text{Ber}(2^{-\sqrt{d}})$ independently of all other bits in \mathcal{H} , i.e., $\mathbb{P}[h_b(x) = 1] = \mathbb{P}[h_b(x) = 1 \mid \{h_{b'}\}_{b' \neq b}, \{h_b(x')\}_{x' \neq x}] = 2^{-\sqrt{d}}$.

In words, for all nodes on the path in the tree corresponding to b , the function h_b assigns a label according to b , and for all other nodes, h_b assigns a label of 1 with probability $2^{-\sqrt{d}}$, and a label of 0 otherwise. In particular, the collection \mathcal{H} Littlestone-shatters the tree T_d .

Fix $B \subseteq \{0, 1\}^{d+1}$ and $X \subseteq T_d$, and let $E(B, X, y)$ denote the event

$$\{\forall b \in B \forall x \in X : x \notin \text{path}(h_b) \wedge h_b(x) = y\}. \quad (12)$$

Seeing as each off-path label $h_b(x) \in \{0, 1\}$ is sampled independently,

$$\begin{aligned} \mathbb{P}_{\mathcal{H} \sim \mathcal{P}}[E(B, X, 0)] &= \prod_{(b, x) \in B \times X} \mathbb{P}_{\mathcal{H} \sim \mathcal{P}}[x \notin \text{path}(h_b) \wedge h_b(x) = 0] \\ &\leq (1 - 2^{-\sqrt{d}})^{|B \times X|}. \end{aligned} \quad (13)$$

Hence,

$$\begin{aligned}
& \mathbb{P}_{\mathcal{H} \sim \mathcal{P}} \left[\exists B \subseteq \{0, 1\}^{d+1} \exists X \subseteq T_d : E(B, X, 0) \wedge \min\{|B|, |X|\} \geq 2^{2\sqrt{d}} \right] \\
&= \mathbb{P}_{\mathcal{H} \sim \mathcal{P}} \left[\exists B \subseteq \{0, 1\}^{d+1} \exists X \subseteq T_d : E(B, X, 0) \wedge |B| = |X| = \left\lceil 2^{2\sqrt{d}} \right\rceil \right] \\
&\leq \binom{|\{0, 1\}^{d+1}|}{\left\lceil 2^{2\sqrt{d}} \right\rceil} \binom{|T_d|}{\left\lceil 2^{2\sqrt{d}} \right\rceil} (1 - 2^{-\sqrt{d}})^{2^{4\sqrt{d}}} && \text{(union bound, Eq. (13))} \\
&< \binom{2^{d+1}}{2^{2\sqrt{d}} + 1}^2 \cdot (1 - 2^{-\sqrt{d}})^{2^{4\sqrt{d}}} && (2^{2\sqrt{d}} \text{ may not be an integer;} \\
&&& \text{formally, we use the general-} \\
&&& \text{ized binomial coefficient, or} \\
&&& \text{simply skip to the next line)} \\
&< 2^{2(d+1) \cdot (2^{2\sqrt{d}} + 1)} \cdot e^{-2^{-\sqrt{d}} \cdot 2^{4\sqrt{d}}} && \left(\binom{n}{k} < n^k \text{ for } k \geq e; 1 + x \leq e^x \text{ for } x \in \mathbb{R} \right) \\
&< 2^{2(d+2) \cdot 2^{2\sqrt{d}}} \cdot 2^{-2^{-\sqrt{d}} \cdot 2^{4\sqrt{d}}} && (2^{2\sqrt{d}} \geq d + 1 \text{ for } d \geq 0) \\
&= 2^{2^{2\sqrt{d}} \cdot (2d + 4 - 2^{2\sqrt{d}})} \\
&< 2^{-2^{2\sqrt{d}}}. && (2d + 4 - 2^{2\sqrt{d}} < -1 \text{ for } d \geq 42) \\
&&& (14)
\end{aligned}$$

Similarly,

$$\mathbb{P}_{\mathcal{H} \sim \mathcal{P}} [\forall b \in B \forall x \in X : x \notin \text{path}(h_b) \wedge h_b(x) = 1] \leq 2^{-\sqrt{d} \cdot |B \times X|}, \quad (15)$$

so

$$\begin{aligned}
& \mathbb{P}_{\mathcal{H} \sim \mathcal{P}} \left[\exists B \subseteq \{0, 1\}^{d+1} \exists X \subseteq T_d : E(B, X, 0) \wedge |H| \geq 2^{2\sqrt{d}} \wedge |X| \geq 3\sqrt{d} \right] \\
&\leq \binom{|\{0, 1\}^{d+1}|}{\left\lceil 2^{2\sqrt{d}} \right\rceil} \binom{|T_d|}{\left\lceil 3\sqrt{d} \right\rceil} \cdot 2^{-\sqrt{d} \cdot 2^{2\sqrt{d}} \cdot 3\sqrt{d}} && \text{(union bound, Eq. (15))} \\
&\leq \binom{2^{d+1}}{2^{2\sqrt{d}} + 1} \binom{2^{d+1}}{\left\lceil 3\sqrt{d} \right\rceil + 1} \cdot 2^{-3d \cdot 2^{2\sqrt{d}}} \\
&< 2^{(d+1) \cdot (2^{2\sqrt{d}} + 1)} \cdot 2^{(d+1) \cdot (3\sqrt{d} + 1)} \cdot 2^{-3d \cdot 2^{2\sqrt{d}}} && \left(\binom{n}{k} < n^k \text{ for } k \geq e \right) \\
&< 2^{(d+1) \cdot (2^{2\sqrt{d}} + 3\sqrt{d} + 2)} \cdot 2^{-3d \cdot 2^{2\sqrt{d}}} \\
&< 2^{2d \cdot 2^{2\sqrt{d}}} \cdot 2^{-3d \cdot 2^{2\sqrt{d}}} && \text{(for } d \geq 4) \\
&< 2^{-d2^{2\sqrt{d}}}. && (16)
\end{aligned}$$

Applying a union bound to Eqs. (14) and (16) gives

$$\mathbb{P}_{\mathcal{H} \sim \mathcal{P}} [\mathcal{H} \text{ satisfies Items 1 and 2}] \geq 1 - 2^{-2^{2\sqrt{d}}} - 2^{-d2^{2\sqrt{d}}} \geq 1 - 10^{-100}.$$

In particular, there exists a collection \mathcal{H} that satisfies Items 1 and 2. Furthermore, this collection has $\text{LD}(\mathcal{H}) = d + 1$ (namely, $\text{LD}(\mathcal{H}) \geq d + 1$ because it shatters T_d ; and $\text{LD}(\mathcal{H}) \leq d + 1$ because $|\mathcal{H}| = 2^{d+1}$). \square

8.3 Algorithm

In this section we describe Algorithms 5, 6a, and 6c, which together constitute the learning algorithm that achieves the $O(\sqrt{d})$ mistake upper bound in the transductive setting, as in Theorem 8.1. See Section 2.3 for a general overview of these algorithms.

8.3.1 How Experts Work

We start with some preliminary remarks about experts in Algorithms 5, 6a, and 6c.

Experts. A tuple $e = (S, u, H)$ defines an expert that can make predictions using the procedure `EXPERT.PREDICT(e, \cdot)`. The tuple e reflects two kinds of information:

1. *Knowledge.* Information that the expert *knows* with certainty. Specifically, this reflects the labels y_1, y_2, \dots sent by the adversary so far. All experts see the labels sent by the adversary, so this knowledge is the same for all experts.
2. *Assumptions.* At certain times, experts make *assumptions* about things that are not known for certain. Specifically, experts assume that certain nodes x are on-path ($x \in \text{path}(h)$) or off-path ($x \notin \text{path}(h)$) with respect to the correct labeling function $h : T_d \rightarrow \{0, 1\}$. Assumptions are simply guesses that may be wrong, and therefore when an expert needs to make such an assumption, it splits into two experts (as described below), with one expert assuming $x \in \text{path}(h)$, and the other expert assuming $x \notin \text{path}(h)$. This ensures that there always exists an expert for which all assumptions are correct.

In greater detail, the contents of the state tuple $e = (S, u, H)$ represents the knowledge and assumptions of the expert as follows:

- $u \in T_d$ – This single node encodes everything the expert knows and assumes about which of the nodes labeled so far are on-path. Observe that if $v_1, v_2, \dots, v_k \in T_d$ are nodes that are assumed to be on-path (and all these assumptions are consistent), then these k assumptions can be represented succinctly by assigning $u = v_{i^*}$ where v_{i^*} is the deepest node among v_1, v_2, \dots, v_k . Therefore, u simply holds the deepest node in the tree that is known or assumed to be on-path. At the start of the algorithm, this value is initialized to be $u = \lambda$, because the root is known to be on-path regardless of the target function.
- $S \subseteq T_d$ – the ‘danger zone’, as described in Section 2.3.4. This is a collection that contains all nodes in the prefix $x_{\leq t_{\max}} = (x_1, x_2, \dots, x_{t_{\max}})$ of the sequence to be classified that have not been labeled yet and *might* be on-path for the true labeling function h given what the expert knows and assumes so far. However, S is not required to contain ancestors of nodes that are assumed to be on-path. Initially, S equals the prefix $x_{\leq t_{\max}}$. As information accumulates, nodes that cannot be on-path are removed from S . For instance, if $x_i \in T_d$ is assigned label $y_i \in \{0, 1\}$ by the adversary, then any $(1 - y_i)$ -descendant of x_i (including x_i itself) may safely be removed from S .
- $H \subseteq \{0, 1\}^{T_d}$ – the version space of the experts, i.e., the collection of all functions that could be the correct labeling function given everything that the expert knows and assumes. Initially, H contains all functions in \mathcal{H} . As information accumulates, some functions are ruled out. Specifically, a function h can be removed from H for two reasons: (i) the adversary assigns a

label $y \neq h(x)$ to some node $x \in T_d$; (ii) the expert makes an assumption that some $x \in T_d$ is on-path for the correct labeling function but $x \notin \text{path}(h)$, or vice versa, the expert assumes that x is off-path for the correct labeling function but $x \in \text{path}(h)$.

Updates and splits. An expert can be modified using the procedure $\text{EXPERT.EXTENDEDUPDATE}(e, \cdot, \cdot)$. This procedure either returns a single modified tuple (S, u, H) (in the first two return statements in the procedure), in which case we think of the expert as being *updated*; or alternatively, the procedure returns two tuples $e_\epsilon = (S_\epsilon, u_\epsilon, H_\epsilon)$ and $e_{\neg\epsilon} = (S_{\neg\epsilon}, u_{\neg\epsilon}, H_{\neg\epsilon})$ (in the third return statement), in which case we think of the expert as being *split* into two experts. The expert e_ϵ corresponds to adding an assumption that the most recently presented node x_t is on-path for the correct labeling function, and $e_{\neg\epsilon}$ corresponds to adding the opposite assumption.

Ancestry. At the end of each iteration of the outer ‘for’ loop in Algorithm 5, for each expert $e \in E_{t+1}$ there exists a unique *ancestry* sequence $\text{ancestry}(e) = (e_1, e_2, \dots, e_{t+1})$ such that $e_1 = (\{x_1, \dots, x_{t_{\max}}\}, \lambda, \mathcal{H})$ is the initial single expert that was created before the start of the outer ‘for’ loop, $e_{t+1} = e$ is the latest version of the expert, and for each $i \in [t]$, the expert e_{i+1} was created by an execution of $\text{EXPERT.BASICUPDATE}(e_i, \cdot, \cdot)$ possibly followed by an execution of $\text{EXPERT.EXTENDEDUPDATE}$.²³

²³Note that in this paper, we use genealogical metaphors in two distinct contexts that should not be confused. First, as is customary, we use “child”, “parent”, “ancestor” and “descendant” to describe relations between nodes in the binary tree T_d , which constitutes the domain of our hypothesis class. Separately from that, we use “ancestor” and “descendant” to describe relations between experts.

This overlap in terminology can partially be excused by the fact that the history of experts also forms a binary tree. Indeed, initially there is a single expert (the root of the tree), and experts can split into two, corresponding to a node having two children as in a binary tree. Seeing as experts cannot merge, the expert history corresponds precisely to a binary tree. (However, the domain T_d is a *perfect* binary tree, whereas the binary tree corresponding to expert genealogy need not be balanced).

To reduce confusion, we use $\text{path}(\cdot)$ only for nodes in T_d , and $\text{ancestry}(\cdot)$ only for experts, even though these operators are mathematically equivalent (however, $\text{path}(\cdot)$ is defined not only for nodes in T_d but also for functions $T_d \rightarrow \{0, 1\}$).

8.3.2 Pseudocode

Assumptions:

- $d, n \in \mathbb{N}$, λ is the empty string.
- $\mathcal{H} \subseteq \{0, 1\}^{T_d}$ is the class that exists by Lemma 8.2.
- $x_1, x_2, \dots, x_n \in T_d$ are points to be classified.

TRANSDUCTIVELEARNER($\mathcal{H}, d, (x_1, x_2, \dots, x_n)$):

```

 $t \leftarrow 0, t_{\max} \leftarrow 2^{4\sqrt{d}}$ 
 $e \leftarrow (\{x_1, \dots, x_{t_{\max}}\}, \lambda, \mathcal{H})$       ▷ The initial expert. An expert is defined by a 3-tuple.
 $w(e) \leftarrow 1$                                 ▷ Assign the initial expert a weight of 1.
 $E_1 \leftarrow \{e\}$                               ▷  $E_t$  is the set of experts used for predicting  $\hat{y}_t$ .
 $E_2, \dots, E_n, E_{n+1} \leftarrow \emptyset$ 

for  $t \leftarrow 1, 2, \dots, n$ :

     $\hat{y}_t \leftarrow \mathbb{1}\left(\sum_{e \in E_t} w(e) \cdot \text{EXPERT.PREDICT}(e, x_t) \geq \frac{1}{2}\right)$   ▷ A weighted majority, using
                                                                                      Algorithm 6a.

    send prediction  $\hat{y}_t$  to adversary
    receive correct label  $y_t \in \{0, 1\}$  from adversary

    for  $e \in E_t$ :                                ▷ Update the experts.
         $e \leftarrow \text{EXPERT.BASICUPDATE}(e, x_t, y_t)$   ▷ Remove functions that disagree with
                                                                                      the label  $y_t$  from the version space.

        if  $\text{EXPERT.PREDICT}(e, x_t) = y_t$ :
             $E_{t+1} \leftarrow E_{t+1} \cup \{e\}$         ▷ If expert  $e$  made a correct prediction,
                                                                                      no further update is needed.

        else:
             $U \leftarrow \text{EXPERT.EXTENDEDUPDATE}(e, x_t, y_t)$   ▷ If  $e$  made a mistake,
                                                                                      update  $e$  using Algo-
                                                                                      rithm 6c. This might
                                                                                      cause  $e$  to be split into
                                                                                      two experts.

            for  $e' \in U$ :
                 $E_{t+1} \leftarrow E_{t+1} \cup \{e'\}$     ▷ Add updated expert(s) to  $E_{t+1}$ .
                 $w(e') \leftarrow w(e)/(2 \cdot |U|)$     ▷ When  $e$  makes a mistake, its weight
                                                                                      is decreased by a factor of 2 and then
                                                                                      split equally between its descendants.

```

Algorithm 5: A transductive online learning algorithm that makes at most $O(\sqrt{d})$ mistakes. It is a variant of the multiplicative weights algorithm that employs splitting experts. Namely, we start with a single expert, and when an expert makes a mistake it may split into two experts. The behavior of the experts is defined in Algorithms 6a and 6c.

8.4 Analysis

In this section we prove our main result, Theorem 8.1.

8.4.1 Assumption-Consistent Expert

Occasionally, when an expert is updated, it makes an assumption about whether the most-recently presented node x_t is on-path or off-path with respect to the true labeling function h . In these updates, the expert is split into two: one expert assumes that $x_t \in \text{path}(h)$, and the other assumes $x_t \notin \text{path}(h)$. Clearly, by splitting into two in this manner, we preserve the invariant that the set of experts always contains a ‘vindicated’ expert e^* such that all the assumptions made by e^* are correct. This simple observation is made formal in the following definition and claim.

Definition 8.3 (Assumption consistency). *For an expert $e \in E_{t+1}$ with $\text{ancestry}(e) = (e_1, e_2, \dots, e_{t+1})$, and an index $i \in [t]$, we say that the $i \rightarrow (i+1)$ update of e was assumption-consistent with a function $h : T_d \rightarrow \{0, 1\}$ if one of the following conditions holds:*

- $e_{i+1} = \text{EXPERT.BASICUPDATE}(e_i, x_i, y_i)$; or
- e_{i+1} was the single expert returned when executing $\text{EXPERT.EXTENDEDUPDATE}(e'_i, x_i, y_i)$ for $e'_i = \text{EXPERT.BASICUPDATE}(e_i, x_i, y_i)$; or
- Executing $\text{EXPERT.EXTENDEDUPDATE}(e'_i, x_i, y_i)$ with $e'_i = \text{EXPERT.BASICUPDATE}(e_i, x_i, y_i)$ returned two experts $(S_\epsilon, u_\epsilon, H_\epsilon)$ and $(S_\notin, u_\notin, H_\notin)$ (as in the third return statement), and furthermore,

$$e_{i+1} = \begin{cases} (S_\epsilon, u_\epsilon, H_\epsilon) & x_i \in \text{path}(h) \\ (S_\notin, u_\notin, H_\notin) & x_i \notin \text{path}(h). \end{cases} \quad (17)$$

We say that an expert $e \in E_{t+1}$ is assumption-consistent with h if for all $i \in [t]$, the $i \rightarrow (i+1)$ update of e was assumption-consistent with h .

Claim 8.4 (Existence of assumption-consistent expert). *Let $d, n, t \in \mathbb{N}$, $t \leq n$, let $\mathcal{H} \subseteq \{0, 1\}^{T_d}$, let $x_1, \dots, x_n \in T_d$, and let $h : T_d \rightarrow \{0, 1\}$. Consider an execution of*

$$\text{TRANSDUCTIVELEARNER}(\mathcal{H}, d, (x_1, x_2, \dots, x_n))$$

as in Algorithm 5. Then, at the end of the t -th iteration of the outer ‘for’ loop in TRANSDUCTIVELEARNER, there exists a unique expert $e_{t+1}^ \in E_{t+1}$ that is assumption-consistent with h .*

Proof. We prove by induction that, for all $s \in [t+1]$, E_s contains a unique expert that is assumption-consistent with h . The base case $s = 1$ is clear, because E_1 contains only a single expert that was never modified. For the induction step, let e_s^* be the unique assumption-consistent expert in E_s , and consider the $s \rightarrow (s+1)$ update. Notice that by Definition 8.3,

- For all $e \in E_s \setminus \{e_s^*\}$, every expert $e' \in E_{s+1}$ such that e' was created from e by executing $\text{EXPERT.BASICUPDATE}(e, x_s, y_s)$ possibly followed by an execution of $\text{EXPERT.EXTENDEDUPDATE}$ is not assumption-consistent with h ; and
- Either $\text{EXPERT.BASICUPDATE}(e_s^*, x_s, y_s) \in E_{s+1}$ and $\text{EXPERT.EXTENDEDUPDATE}(e_s^*, x_s, y_s)$ is not executed (e_s^* is added to E_{s+1} with just a basic update), or precisely one of the experts that were

created from e_s^* by executing `EXPERT.EXTENDEDUPDATE` and added to E_{s+1} is assumption-consistent with h .

Seeing as the $s \rightarrow (s + 1)$ update executes `EXPERT.BASICUPDATE` and `EXPERT.EXTENDEDUPDATE` at most once for each $e \in E_s$, it follows that E_{s+1} contains precisely one expert that is assumption-consistent with h . \square

An expert $e = (S, u, H)$ that is assumption-consistent with the correct labeling function enjoys two simple properties. The first property is that the node u in the expert encodes correct information about which previously seen nodes are on-path for the correct labeling function.

The second property is that the set S contains all future nodes that are on-path for the correct labeling function and are also deeper in the tree than all nodes assumed to be on-path so far. These two properties are formalized in the following claim.

Claim 8.5 (Properties of assumption-consistent expert). *Let $d, n, t \in \mathbb{N}$, $t \leq n + 1$, let $\mathcal{H} \subseteq \{0, 1\}^{T_d}$, let $x_1, \dots, x_n \in T_d$. Consider an execution of*

$$\text{TRANSDUCTIVELEARNER}(\mathcal{H}, d, (x_1, x_2, \dots, x_n))$$

as in Algorithm 5. Assume that the adversary selects labels $y_1, y_2, \dots, y_n \in \{0, 1\}$ that are consistent with some function $h : T_d \rightarrow \{0, 1\}$. Let $e_t^ = (S_t^*, u_t^*, H_t^*) \in E_t$ be the unique expert in E_t that is assumption-consistent with h .²⁴ Then the following two properties hold:*

1. $u_t^* \in \text{path}(h)$.
2. $\{x \in \{x_t, x_{t+1}, \dots, x_{t_{\max}}\} : x \in \text{path}(h) \wedge x \not\preceq u_t^*\} \subseteq S_t^*$.

Proof of Claim 8.5. The proof proceeds by induction on t . For the base case $t = 1$, E_1 contains a single expert $e_1^* = (S_1^*, u_1^*, H_1^*)$ where $u_1^* = \lambda$ is the root of T_d . Indeed, $\lambda \in \text{path}(h)$ for any function $h : T_d \rightarrow \{0, 1\}$. This establishes the base case for Item 1. Additionally, $S_1^* = \{x_1, x_2, \dots, x_{t_{\max}}\}$, satisfying the base case for Item 2.

For the induction step, we assume that the claim holds for some integer $t = i$, and show that it holds for $t = i + 1$ as well. First, we establish Item 1. If $e_{i+1}^* = \text{EXPERT.BASICUPDATE}(e_i^*, x_i, y_i)$, then the claim is immediate because $u_{i+1}^* = u_i^* \in \text{path}(h)$. Otherwise, by Definition 8.3 and the first two return statements in `EXPERT.EXTENDEDUPDATE`, either $e_{i+1}^* = (S_{i+1}^*, u_{i+1}^*, H_{i+1}^*)$ has $u_{i+1}^* = u_i^* \in \text{path}(h)$, in which case the claim is immediate, or else e_{i+1}^* satisfies Eq. (17), namely,

$$e_{i+1}^* = \begin{cases} (S_{\epsilon}, u_{\epsilon}, H_{\epsilon}) & x_i \in \text{path}(h) \\ (S_{\not\epsilon}, u_{\not\epsilon}, H_{\not\epsilon}) & x_i \notin \text{path}(h). \end{cases}$$

As defined in `EXPERT.EXTENDEDUPDATE`, u_{ϵ} is equal either to u_i^* or to x_i , so if $x_i \in \text{path}(h)$ then

$$u_{i+1}^* = u_{\epsilon} \in \{u_i^*, x_i\} \subseteq \text{path}(h).$$

On the other hand, if $x_i \notin \text{path}(h)$ then we get $u_{i+1}^* = u_{\not\epsilon} = u_i^* \in \text{path}(h)$. We see that in all cases, $u_{i+1}^* \in \text{path}(h)$ as desired. This concludes the proof of Item 1.

²⁴Recall that e_i^* exists by Claim 8.4.

For Item 2, again, if $e_{i+1}^* = \text{EXPERT.BASICUPDATE}(e_i^*, x_i, y_i)$, then the claim is immediate because $S_{i+1}^* = S_i^*$ and $u_{i+1}^* = u_i^*$. Otherwise, consider the various ways in which u_{i+1}^* and S_{i+1}^* can be assigned by $\text{EXPERT.EXTENDEDUPDATE}$. In the first return statement, $u_{i+1}^* = u_i^*$ and $S_{i+1}^* = S_i^*$, and the claim is immediate.

The second return statement assigns $u_{i+1}^* = u_i^*$ and $S_{i+1}^* = S_i^* \setminus S_{1-y_i}$, where S_{1-y_i} is the set of $(1 - y_i)$ -descendants of x_i (including x_i itself). Notice that regardless of whether x_i is on-path for the correct labeling function h or not, none of the $(1 - y_i)$ -descendants of x_i (except possibly x_i itself) can be on-path for h , because h assigns a label y_i to x_i . And seeing as Item 2 only requires that S_{i+1}^* contain nodes from $\{x_{i+1}, x_{i+2}, \dots, x_{t_{\max}}\}$, it is also safe to remove x_i . Therefore, removing S_{1-y_i} preserves Item 2.

For the third return statement, there are two possibilities. The first possibility is that $u_{i+1}^* = u_\epsilon = u_i^*$ and $S_{i+1}^* = S_\epsilon = S_i^*$, in which case the claim is immediate. The second possibility assigns $u_{i+1}^* = u_\epsilon$, and $S_{i+1}^* = S_\epsilon = S_0 \cup S_1$, namely, S_{i+1}^* is constructed by removing the non-descendants of x_i from S_i^* . By Eq. (17), this happens when $x_i \in \text{path}(h)$, so all non-descendants of x_i are either off-path for h , or they are ancestors of x_i . Seeing as $x_i \in \text{path}(h)$ and $u_i^* \in \text{path}(h)$, and u_ϵ is the deeper node between these two, any node that is an ancestor of x_i is also an ancestor of $u_{i+1}^* = u_\epsilon$. Thus, all the nodes removed are either off-path for h , or they are ancestors of u_{i+1}^* , satisfying Item 2. (Similarly, any node that is an ancestor of u_i^* is also an ancestor of u_{i+1}^* , so we do not need to add any new nodes to S_{i+1}^* that are not included in S_i^* .)

We see that in all cases, Item 2 is preserved, as desired. \square

8.4.2 Transition to Halving

Claim 8.6. *Let $d, n, t \in \mathbb{N}$, $d \geq 16$, let $\mathcal{H} \subseteq \{0, 1\}^{T_d}$, and let $x_1, \dots, x_n \in T_d$. Consider an execution of*

$$\text{TRANSDUCTIVELEARNER}(\mathcal{H}, (x_1, x_2, \dots, x_n))$$

as in Algorithm 5. Let $t > t_{\max} = 2^{4\sqrt{d}}$ and let $e = (S, u, H) \in E_t$ be an expert. Then

$$|H| \leq 2^{2\sqrt{d}}.$$

Proof of Claim 8.6. Assume for contradiction that $|H| > 2^{2\sqrt{d}}$. Let $H' \subseteq H$ be an arbitrary subset of size $2^{2\sqrt{d}} + 1$. Let

$$P = \cup_{h \in H'} \text{path}(h).$$

Seeing as each root-to-leaf path contains $d + 1$ nodes,

$$|P| \leq |H'| \cdot (d + 1) \leq \left(2^{2\sqrt{d}} + 1\right) \cdot (d + 1) \leq d2^{2\sqrt{d}+1}. \quad (18)$$

Let y_1, y_2, \dots, y_t be the labels provided by the adversary in the first t_{\max} iterations. The line in $\text{EXPERT.BASICUPDATE}$ constructing H using $\text{HALVING.UPDATE}(H, x, y)$ ensures that

$$\forall h \in H \forall i \in [t_{\max}] : h(x_i) = y_i. \quad (19)$$

Consider two cases:

- **Case I.** $\sum_{i=1}^{t_{\max}} y_i \leq t_{\max}/2$. Then the set

$$X_0 = \{x_i : i \in [t_{\max}] \wedge y_i = 0\}$$

has cardinality $|X_0| \geq t_{\max}/2$. Let $X'_0 = X_0 \setminus P$. By Eq. (18),

$$|X'_0| \geq \frac{t_{\max}}{2} - d2^{2\sqrt{d}+1} = 2^{4\sqrt{d}} - d2^{2\sqrt{d}+1}. \quad (20)$$

From the choice of X'_0 , the inclusion $H' \subseteq H$, and Eq. (19),

$$\forall h \in H' \forall x \in X'_0 : x \notin \text{path}(h) \wedge h(x) = 0. \quad (21)$$

Seeing as $|H'| > 2^{2\sqrt{d}}$, Eq. (21) and Item 1 from Lemma 8.2 imply that

$$|X'_0| \leq 2^{2\sqrt{d}}. \quad (22)$$

Combining Eqs. (20) and (22) yields

$$\begin{aligned} 2^{2\sqrt{d}} &\geq |X'_0| \geq 2^{4\sqrt{d}} - d2^{2\sqrt{d}+1} \\ &\geq 2^{4\sqrt{d}-1} \end{aligned} \quad (d \geq 16),$$

which is a contradiction.

- **Case II.** $\sum_{i=1}^{t_{\max}} y_i > t_{\max}/2$. A similar argument gives a contradiction by defining

$$X_1 = \{x_i : i \in [t_{\max}] \wedge y_i = 1\}, \text{ and } X'_1 = X_1 \setminus P.$$

As before,

$$|X'_1| \geq \frac{t_{\max}}{2} - d2^{2\sqrt{d}+1} \geq 2^{4\sqrt{d}} - d2^{2\sqrt{d}+1}. \quad (23)$$

for all $d \in \mathbb{N}$. However, $|H'| > 2^{2\sqrt{d}}$ and Item 2 imply that

$$|X'_1| < 3\sqrt{d}, \quad (24)$$

which is a contradiction. \square

8.4.3 Performance of Best Expert

Claim 8.7 (Existence of expert with large weight). *Let $d, n \in \mathbb{N}$, $d \geq 16$, let $\mathcal{H} \subseteq \{0, 1\}^{T_d}$, and let $x_1, \dots, x_n \in T_d$. Consider an execution of*

$$\text{TRANSDUCTIVELEARNER}(\mathcal{H}, (x_1, x_2, \dots, x_n))$$

as in Algorithm 5. Then, at the end of the execution, there exists $e \in E_{n+1}$ such that

$$w(e) \geq 2^{-48\sqrt{d}}. \quad (25)$$

Note that the lower bound in Eq. (25) does not depend on n .

Proof. Fix a hypothesis $h \in \mathcal{H}$ such that $h(x_t) = y_t$ for all $t \in [n]$ (such an h exists because the adversary must always select a realizable label).

By Claim 8.4, there exists $e_{n+1}^* \in E_{n+1}$ that is assumption-consistent with h . Let $\text{ancestry}(e_{n+1}^*) = (e_1^*, e_2^*, \dots, e_{n+1}^*)$. We argue that this ancestry sequence makes few mistakes. Specifically, for each $t \in [n]$, let $\hat{y}_t^* = \text{EXPERT.PREDICT}(e_t^*, x_t)$. We claim that

$$m := \sum_{t=1}^n \mathbb{1}(\hat{y}_t^* \neq y_t) \leq 24\sqrt{d}.$$

Indeed, let $B = \{t \in [n] : \hat{y}_t^* \neq y_t\}$ be the set of m indices where a mistake was made. For each $t \in B$, let $e_t^* = (S, u, H)$, and note that each $t \in B$ has a corresponding execution of $\text{EXPERT.PREDICT}(e_t^*, x_t)$, and an execution of $e'_t = \text{EXPERT.BASICUPDATE}(e_t^*, x_t, y_t)$ followed by $\text{EXPERT.EXTENDEDUPDATE}(e'_t, x_t, y_t)$ that produces e_{t+1}^* ($\text{EXPERT.EXTENDEDUPDATE}$ is executed because $t \in B$, i.e., a mistake was made). We partition the indices in B into six cases (six disjoint sets), and bound the number of indices that fall in each.

- **Case I.** *The execution of $\text{EXPERT.PREDICT}(e_t^*, x_t)$ exited via the first return statement in that procedure.* This happens once $|H| \leq 2^{2\sqrt{d}}$, and from that point on, the expert and all subsequent experts in the ancestry are exactly simulating the HALVING algorithm (Algorithm 7) in both predictions and updates. Hence, by Fact A.1, B contains at most $m_I = 2\sqrt{d}$ such indices.
- **Case II.** *The execution of $\text{EXPERT.PREDICT}(e_t^*, x_t)$ exited via the second return statement in that procedure.* In particular $x \preccurlyeq u$, and the predicted label was $\hat{y}_t^* = b \in \{0, 1\}$ such that $x_t \preccurlyeq_b u$. Because e_t^* is assumption-consistent with h , Item 1 in Claim 8.5 implies that $u \in \text{path}(h)$. Namely, we see that u is a b -descendant of x_t and $u \in \text{path}(h)$. It follows that $\hat{y}_t^* = b = h(x_t) = y_t$. So no mistakes are made in Case II, and the number of indices $t \in B$ that belong to Case II is simply $m_{II} = 0$.

In the remaining cases, we assume that $\text{EXPERT.PREDICT}(e_t^*, x_t)$ exited via the third return statement in that procedure, so the prediction was

$$\hat{y}_t^* = \mathbb{1}(|S_1| > |S|/3), \tag{26}$$

where $S_1 = \{x' \in S : x_t \preccurlyeq_1 x'\}$. These cases are as follows.

- **Case III.** *The execution of $\text{EXPERT.EXTENDEDUPDATE}(e'_t, x_t, y_t)$ exited via the first return statement in that procedure.* Namely, after the update, the resulting expert e_{t+1}^* has $|H| \leq 2^{2\sqrt{d}}$. However, because we are not in Case I, at the beginning of the iteration expert e_t^* had $|H| > 2^{2\sqrt{d}}$. Seeing as the cardinality of H decreases monotonically throughout the ancestry e_1^*, \dots, e_{n+1}^* , this type of mistake can happen at most $m_{III} = 1$ times.
- **Case IV.** *The execution of $\text{EXPERT.EXTENDEDUPDATE}(e'_t, x_t, y_t)$ exited via the second return statement in that procedure.* In this case, $|S_{(1-y_t)}| > |S|/3$, and $e_{t+1}^* = (S', u, H)$ with $S' = S \setminus S_{1-y_t}$. So $|S'| < 2|S|/3$. Namely, the update causes the cardinality of the set S to be multiplied by a factor of at most $2/3$ and it strictly decreases. Seeing as the

initial cardinality is t_{\max} , and cardinalities are integers, the number of times this can happen is at most

$$m_{\text{IV}} = \frac{\log t_{\max}}{\log(3/2)} + 1 = \frac{4\sqrt{d}}{\log(3/2)} + 1. \quad (27)$$

In the remaining cases, we assume that the execution of `EXPERT.EXTENDEDUPDATE`(e_t^*, x_t, y_t) exited via the third return statement in that procedure. This implies that

$$|S_{\hat{y}_t^*}| \leq |S|/3 \quad (28)$$

Combining this with Eq. (26), it follows $\hat{y}_t^* = 0$ and therefore $y_t = 1$. The remaining cases are as follows.

- **Case V.** $x_t \in \text{path}(h)$. Let $e_t^* = (S, u, H)$. Seeing as $|H| > 2^{2\sqrt{d}}$ (because we are not in Case I), Claim 8.6 (with the assumption $d \geq 16$) implies that $t \leq t_{\max}$. By Item 2 of Claim 8.5, the facts $x_t \not\ll u$ (we are not in Case II) and $x_t \in \text{path}(h)$ imply that $x_t \in S$. In particular, S is not empty.

Because the $t \rightarrow (t+1)$ update of e_{t+1}^* was assumption-consistent with h , Eq. (17) implies that $e_{t+1}^* = (S_\epsilon, u_\epsilon, H_\epsilon)$, with $S_\epsilon = S_0 \cup S_1$. Observe that

- $|S_0| \leq |S|/3$ (plugging $\hat{y}_t^* = 0$ into Eq. (28)); and
- $|S_1| \leq |S|/3$ (because otherwise, by Eq. (26), the prediction would have been $\hat{y}_t^* = 1$).

Therefore,

$$|S_\epsilon| \leq |S_0| + |S_1| \leq 2|S|/3. \quad (29)$$

As in Case IV, combining Eq. (29) and the fact that S is not empty imply an upper bound m_v on the number of times Case V can happen, with the bound being the same number $m_v = m_{\text{IV}}$ as in Eq. (27).

- **Case VI.** $x_t \notin \text{path}(h)$. So (x_t, y_t) is a pair such that $x_t \notin \text{path}(h)$ and $y_t = 1$. Assume for contradiction that this type of mistake can happen strictly more than

$$m_{\text{VI}} = 3\sqrt{d}$$

times. Let $t_1, t_2, \dots, t_{m_{\text{VI}}}$ be the indices of the first m_{VI} iterations of the outer ‘for’ loop of `TRANSDUCTIVELARNER` in which this type of mistake happened. Note that if at the end of iteration $t_{m_{\text{VI}}}$, we had expert $e_{t_{m_{\text{VI}}+1}}^* = (S_{t_{m_{\text{VI}}+1}}, u_{t_{m_{\text{VI}}+1}}, H_{t_{m_{\text{VI}}+1}})$ such that $|H_{t_{m_{\text{VI}}+1}}| \leq 2^{2\sqrt{d}}$, then from that point onwards, the expert would be simulating the halving algorithm, and in particular, it would not make any further mistake of the type in Case VI (all subsequent mistakes would belong to Case I). Hence, by the assumption that strictly more than m_{VI} mistakes were made, it follows that $|H_{t_{m_{\text{VI}}+1}}| > 2^{2\sqrt{d}}$. Let

$$H^* = \left\{ h' \in \mathcal{H} : \left(\forall t \in [m_{\text{VI}}] : h'(x_{t_i}) = 1 \wedge x_{t_i} \notin \text{path}(h') \right) \right\}.$$

Because $e_{t_{m_{VI}+1}}^*$ is assumption-consistent with h , and from the construction of $H_{t_{m_{VI}+1}}$ using H_\in and H_\notin in EXPERT.EXTENDEDUPDATE, it follows that $H_{t_{m_{VI}+1}} \subseteq H^*$. So there exist collections $H^* \subseteq \mathcal{H}$ and $X = \{x_t : t \in [m_{VI}]\} \subseteq T_d$ such that

- $|H^*| \geq |H_{t_{m_{VI}+1}}| > 2^{2\sqrt{d}}$,
- $|X| = m_{VI} = 3\sqrt{d}$,
- $\forall h' \in H^* \forall x \in X : h'(x) = 1$.
- $\forall h' \in H^* \forall x \in X : x \notin \text{path}(h')$.

This is a contradiction to the choice of \mathcal{H} , specifically, to Item 2 in Lemma 8.2.

Thus, combining the analyses of all cases, we see that the number of mistakes made by the ancestry(e_{n+1}^*) is at most

$$\begin{aligned} m &\leq m_I + m_{II} + m_{III} + m_{IV} + m_V + m_{VI} \\ &\leq 2\sqrt{d} + 0 + 1 + \left(\frac{4\sqrt{d}}{\log(3/2)} + 1 \right) + \left(\frac{4\sqrt{d}}{\log(3/2)} + 1 \right) + 3\sqrt{d} \\ &\leq 24\sqrt{d}. \end{aligned}$$

The weights satisfy

$$w(e_{t+1}^*) \begin{cases} = w(e_t^*) & \hat{y}_t^* = y_t \\ \geq \frac{1}{4} \cdot w(e_t^*) & \hat{y}_t^* \neq y_t. \end{cases}$$

This implies that $w(e_{n+1}^*) \geq w(e_1^*) \cdot \prod_{t=1}^n 4^{-\mathbb{1}(\hat{y}_t \neq y_t)} = w(e_1^*) \cdot 4^{-m} \geq 4^{-24\sqrt{d}} = 2^{-48\sqrt{d}}$, as desired. \square

8.4.4 Multiplicative Weights Mistake Bound

Claim 8.8 (Mistake bound for multiplicative weights). *Let $d, n \in \mathbb{N}$, let $\alpha > 0$, let $\mathcal{H} \subseteq \{0, 1\}^{T_d}$, and let $x_1, \dots, x_n \in T_d$. Consider an execution of*

$$\text{TRANSDUCTIVELEARNER}(\mathcal{H}, (x_1, x_2, \dots, x_n))$$

as in Algorithm 5. Assume that at the end of the execution, there exists $e^ \in E_{n+1}$ such that*

$$w(e^*) \geq 2^{-\alpha}.$$

Then TRANSDUCTIVELEARNER makes at most α mistakes.

Proof of Claim 8.8. For all $i \in [n+1]$, let $w(E_i) = \sum_{e \in E_i} w(e)$. For each $i \in [n]$, if $\hat{y}_i \neq y_i$, then $w(E_{i+1}) \leq w(E_i)/2$. Hence, if TRANSDUCTIVELEARNER makes m mistakes, then by induction

$$w(E_{n+1}) \leq w(E_1) \cdot \prod_{t=1}^n 2^{-\mathbb{1}(\hat{y}_t \neq y_t)} = 2^{-m} \cdot w(E_1).$$

So

$$2^{-\alpha} \leq w(e^*) \leq \sum_{e \in E_{n+1}} w(e) = w(E_{n+1}) \leq 2^{-m} \cdot w(E_1) = 2^{-m}.$$

We conclude that

$$m \leq \alpha,$$

as desired. □

8.5 Proof

Proof of Theorem 8.1. Fix an integer $d \geq 43$. Let $\mathcal{H} \subseteq \{0, 1\}^{T_{d-1}}$ be the class constructed by invoking Lemma 8.2 for the integer $d - 1 \geq 42$. We argue that this class satisfies the requirements of Theorem 8.1.

By construction, \mathcal{H} is a class of Littlestone dimension precisely d . By Theorem 5.7, this implies the equality in Item 2.

We now show the upper bound in Item 1. We argue that `TRANSDUCTIVELEARNER` (Algorithm 5) satisfies this upper bound. By Claim 8.7, at the end of the execution of `TRANSDUCTIVELEARNER` there exists an expert $e \in E_{n+1}$ such that $w(e) \geq 2^{-48\sqrt{d}}$. By Claim 8.8, this implies that the number of mistakes made by `TRANSDUCTIVELEARNER` is at most $48\sqrt{d}$, as desired. □

References

- Maria-Florina Balcan and Avrim Blum. A discriminative model for semi-supervised learning. *J. ACM*, 57(3):19:1–19:46, 2010. doi:10.1145/1706591.1706599. URL <https://doi.org/10.1145/1706591.1706599>.
- Shai Ben-David, Eyal Kushilevitz, and Yishay Mansour. Online learning versus offline learning. In Paul M. B. Vitányi, editor, *Computational Learning Theory, Second European Conference, EuroCOLT '95, Barcelona, Spain, March 13-15, 1995, Proceedings*, volume 904 of *Lecture Notes in Computer Science*, pages 38–52. Springer, 1995. doi:10.1007/3-540-59119-2_167. URL https://doi.org/10.1007/3-540-59119-2_167.
- Shai Ben-David, Eyal Kushilevitz, and Yishay Mansour. Online learning versus offline learning. *Mach. Learn.*, 29(1):45–63, 1997. doi:10.1023/A:1007465907571. URL <https://doi.org/10.1023/A:1007465907571>.
- Shai Ben-David, Tyler Lu, Dávid Pál, and Miroslava Sotáková. Learning low-density separators. *CoRR*, abs/0805.2891, 2008. URL <http://arxiv.org/abs/0805.2891>.
- Gyora M. Benedek and Alon Itai. Learnability with respect to fixed distributions. *Theor. Comput. Sci.*, 86(2):377–390, 1991. doi:10.1016/0304-3975(91)90026-X. URL [https://doi.org/10.1016/0304-3975\(91\)90026-X](https://doi.org/10.1016/0304-3975(91)90026-X).
- Avrim Blum and Tom M. Mitchell. Combining labeled and unlabeled data with co-training. In Peter L. Bartlett and Yishay Mansour, editors, *Proceedings of the Eleventh Annual Conference on Computational Learning Theory, COLT 1998, Madison, Wisconsin, USA, July 24-26, 1998*, pages 92–100. ACM, 1998. doi:10.1145/279943.279962. URL <https://doi.org/10.1145/279943.279962>.
- 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 2021: 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>.
- Nicolò Cesa-Bianchi and Ohad Shamir. Efficient transductive online learning via randomized rounding. In Bernhard Schölkopf, Zhiyuan Luo, and Vladimir Vovk, editors, *Empirical Inference - Festschrift in Honor of Vladimir N. Vapnik*, pages 177–194. Springer, 2013. doi:10.1007/978-3-642-41136-6_16. URL https://doi.org/10.1007/978-3-642-41136-6_16.
- Olivier Chapelle, Vladimir N. Vapnik, and Jason Weston. Transductive inference for estimating values of functions. In Sara A. Solla, Todd K. Leen, and Klaus-Robert Müller, editors, *Advances in Neural Information Processing Systems 12, [NIPS Conference, Denver, Colorado, USA, November 29 - December 4, 1999]*, pages 421–427. The MIT Press, 1999. URL <http://papers.nips.cc/paper/1699-transductive-inference-for-estimating-values-of-functions>.
- Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien, editors. *Semi-Supervised Learning*. The

- MIT Press, 2006. ISBN 9780262033589. doi:10.7551/MITPRESS/9780262033589.001.0001. URL <https://doi.org/10.7551/mitpress/9780262033589.001.0001>.
- Malte Darnstädt, Hans Ulrich Simon, and Balázs Szörényi. Unlabeled data does provably help. In Natacha Portier and Thomas Wilke, editors, *30th International Symposium on Theoretical Aspects of Computer Science, STACS 2013, February 27 - March 2, 2013, Kiel, Germany*, volume 20 of *LIPICs*, pages 185–196. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013. doi:10.4230/LIPICs.STACS.2013.185. URL <https://doi.org/10.4230/LIPICs.STACS.2013.185>.
- Alexander Gammerman, Volodya Vovk, and Vladimir N. Vapnik. Learning by transduction. In Gregory F. Cooper and Serafin Moral, editors, *UAI 1998: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, University of Wisconsin Business School, Madison, Wisconsin, USA, July 24-26, 1998*, pages 148–155. Morgan Kaufmann, 1998. URL https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=243&proceeding_id=14.
- Christina Göpfert, Shai Ben-David, Olivier Bousquet, Sylvain Gelly, Ilya O. Tolstikhin, and Ruth Urner. When can unlabeled data improve the learning rate? In Alina Beygelzimer and Daniel Hsu, editors, *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pages 1500–1518. PMLR, 2019. URL <http://proceedings.mlr.press/v99/gopfert19a.html>.
- Steve Hanneke, Shay Moran, and Jonathan Shafer. A trichotomy for transductive online learning. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/3e32af2df2cd13dfbcbe6e8d38111068-Abstract-Conference.html.
- Steve Hanneke, Vinod Raman, Amirreza Shaeiri, and Unique Subedi. Multiclass transductive online learning. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/6f244818d72b2a4be9b1225d1344e950-Abstract-Conference.html.
- Steven C. H. Hoi, Doyen Sahoo, Jing Lu, and Peilin Zhao. Online learning: A comprehensive survey. *Neurocomputing*, 459:249–289, 2021. doi:10.1016/J.NEUCOM.2021.04.112. URL <https://doi.org/10.1016/j.neucom.2021.04.112>.
- Thorsten Joachims. Transductive inference for text classification using support vector machines. In Ivan Bratko and Saso Dzeroski, editors, *Proceedings of the Sixteenth International Conference on Machine Learning (ICML 1999), Bled, Slovenia, June 27 - 30, 1999*, pages 200–209. Morgan Kaufmann, 1999.
- Sham M. Kakade and Adam Kalai. From batch to transductive online learning. In *Advances in Neural Information Processing Systems 18 [Neural Information Processing Sys-*

tems, NIPS 2005, December 5-8, 2005, Vancouver, British Columbia, Canada], pages 611–618, 2005. URL <https://proceedings.neurips.cc/paper/2005/hash/17693c91d9204b7a7646284bb3adb603-Abstract.html>.

Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Mach. Learn.*, 2(4):285–318, 1987. doi:10.1007/BF00116827. URL <https://doi.org/10.1007/BF00116827>.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014. ISBN 978-1-10-705713-5. URL <http://www.cambridge.org/de/academic/subjects/computer-science/pattern-recognition-and-machine-learning/understanding-machine-learning-theory-algorithms>.

Vladimir N. Vapnik. *Estimation of Dependencies Based on Empirical Data*. Nauka, Moscow, 1979. URL <https://www.ipu.ru/node/63854/publications>. In Russian.

Vladimir N. Vapnik. *Estimation of Dependences Based on Empirical Data*. Springer, 2nd edition, 2006. ISBN 978-0-387-30865-4. doi:10.1007/0-387-34239-7. URL <https://doi.org/10.1007/0-387-34239-7>.

Xiaojin Zhu. Semi-supervised learning literature survey. Technical report, Department of Computer Sciences, University of Wisconsin–Madison, 2005.

Xiaojin Zhu. Semi-supervised learning. In Claude Sammut and Geoffrey I. Webb, editors, *Encyclopedia of Machine Learning*, pages 892–897. Springer, 2010. doi:10.1007/978-0-387-30164-8_749. URL https://doi.org/10.1007/978-0-387-30164-8_749.

Xiaojin Zhu and Andrew B. Goldberg. *Introduction to Semi-Supervised Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2009. ISBN 978-3-031-00420-9. doi:10.2200/S00196ED1V01Y200906AIM006. URL <https://doi.org/10.2200/S00196ED1V01Y200906AIM006>.

A Halving

Fact A.1. Let \mathcal{X} be a set, and let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class. Then for all $n \in \mathbb{N}$, all sequences $x \in \mathcal{X}^n$, and all realizable adversaries, HALVING (Algorithm 7) makes at most $\log |\mathcal{H}|$ mistakes in the transductive online learning (Game 2).²⁵ Namely,

$$\sup_{n \in \mathbb{N}} \sup_{A \in \mathcal{A}_n} M_{\text{tr}}(\mathcal{H}, n, \text{HALVING}, A) \leq \log |\mathcal{H}|.$$

²⁵With the suitable syntactic modification, it also makes at most $\log |\mathcal{H}|$ mistakes in the standard online learning (Game 1).

Assumptions:

- $d \in \mathbb{N}$, $M = \sqrt{d}/10$, $\varepsilon = 2^{-\sqrt{d}/2}$.
- $T = T_d$ is a perfect binary tree of depth d .
- λ , the empty string, is the root of T .
- $\mathcal{H} \subseteq \{0, 1\}^T$ is a class that shatters T .

CONSTRUCTSEQUENCE (\mathcal{H}): $\mathcal{H}_\lambda \leftarrow \mathcal{H}$ $\mathbb{H}_0 \leftarrow \{\mathcal{H}_\lambda\}$

▷ A set of classes indexed by bit strings.

 $\mathcal{Q} \leftarrow \{\lambda\}$

▷ A set of nodes to be processed.

 $t \leftarrow 0$ **while** $|\mathcal{Q}| > 0$: $t \leftarrow t + 1$ $x_t \leftarrow$ arbitrary element from \mathcal{Q} ▷ Pop an arbitrary element from \mathcal{Q}
and add it to the output sequence. $\mathcal{Q} \leftarrow \mathcal{Q} \setminus \{x_t\}$ $\mathbb{H}_t \leftarrow \emptyset$ **for** $\mathcal{H}_b \in \mathbb{H}_{t-1}$:

$$r \leftarrow \frac{|\{h \in \mathcal{H}_b : h(x_t) = 1\}|}{|\mathcal{H}_b|}$$

$$\mathcal{Y} \leftarrow \begin{cases} \{0, 1\} & (r \in [\varepsilon, 1 - \varepsilon]) \wedge (|b| < M) \\ \{\mathbb{1}(r \geq 1/2)\} & \text{otherwise} \end{cases}$$

▷ Adversary will
force mistakes on
the first M balanced
nodes.**for** $y \in \mathcal{Y}$:

$$b' \leftarrow \begin{cases} b & |\mathcal{Y}| = 1 \\ b \circ y & |\mathcal{Y}| = 2 \end{cases}$$

▷ Restrict class to agree with y . If splitting the class in two to force a mistake then create new indices.

$$\mathcal{H}_{b'} \leftarrow \{h \in \mathcal{H}_b : h(x_t) = y\}$$

$$\mathbb{H}_t \leftarrow \mathbb{H}_t \cup \{\mathcal{H}_{b'}\}$$

if $x_t \in \text{path}(\mathcal{H}_{b'}) \wedge |x_t| < d$:▷ If x_t is on-path for $\mathcal{H}_{b'}$ and it has a y -child, add that child to \mathcal{Q} .

$$\mathcal{Q} \leftarrow \mathcal{Q} \cup \{x_t \circ y\}$$

return (x_1, x_2, \dots, x_t) Algorithm 2: A subroutine of Algorithm 1 for selecting the sequence x .

Assumptions:

- $n \in \mathbb{N}$, \mathcal{X} is a set, $x \in \mathcal{X}^n$ is a fixed sequence of instances.
- $A : \left(\bigcup_{k=0}^{n-1} \{0, 1\}^{2k}\right) \times \{0, 1\} \rightarrow \{0, 1\}$ is an adversary labeling strategy for x .

RIGIDADVERSARY:

```
send  $x_1, \dots, x_n$  to the learner
for  $t = 1, 2, \dots, n$ :
    receive prediction  $\hat{y}_t$  from learner
    if  $A(\tilde{y}_1, y_1, \dots, \tilde{y}_{t-1}, y_{t-1}, 0) = 0$ :
         $\tilde{y}_t \leftarrow 0$ 
    else if  $A(\tilde{y}_1, y_1, \dots, \tilde{y}_{t-1}, y_{t-1}, 1) = 1$ :
         $\tilde{y}_t \leftarrow 1$ 
    else:
         $\tilde{y}_t \leftarrow \hat{y}_t$ 
    send prediction  $\tilde{y}_t$  to  $A$ 
    receive label  $y_t$  from  $A$ 
    send label  $y_t$  to learner
```

Algorithm 3: Construction of a rigid adversary, by simulating a given adversary A .

Assumptions:

- $n, M \in \mathbb{N}$, \mathcal{X} is a set, $x \in \mathcal{X}^n$ is a fixed sequence of instances.
- $A_1 : \left(\bigcup_{k=0}^{n-1} \{0, 1\}^{2^k}\right) \times \{0, 1\} \rightarrow \{0, 1\}$ is a rigid adversary labeling strategy for x that forces every learner to make at least M mistakes on the sequence x , and satisfies Items 1 and 2 in Claim 7.4.
- $I = \{i_1, i_2, \dots, i_k\} \subseteq [n]$ is the set of indices that are essential for A for forcing M mistakes on x , and $i_1 \leq i_2 \leq \dots \leq i_k$. By Claim 7.6, $k \leq 2^M - 1$.

MINIMALADVERSARY:

send $x_{i_1}, x_{i_2}, \dots, x_{i_k}$ to the learner

for $t = 1, 2, \dots, n$:

if $t \in I$:

receive prediction \hat{y}_t from learner

send prediction \hat{y}_t to A_1

receive label y_t from A_1

send label y_t to learner

else:

send prediction $\hat{y}_t = 0$ to A_1

receive label y_t from A_1

Algorithm 4: Construction of an adversary that forces M mistakes using a sequence x of length at most $2^M - 1$. In the proof of Theorem 7.2, this adversary is A_2 . Internally, it simulates a rigid adversary A_1 .

Assumptions:

- $d \in \mathbb{N}$, $x \in T_d$.
- $e = (S, u, H)$ is a tuple that defines an expert:
 - $S \subseteq T_d$ – a collection of nodes that could be on-path for the true labeling function given what the expert knows and assumes.
 - $u \in T_d$ – the deepest node known or assumed to be on-path by the expert.
 - $H \subseteq \{0, 1\}^{T_d}$ – the collection of all functions that could be the correct labeling function given what the expert knows and assumes.

EXPERT.PREDICT(e, x):

```

( $S, u, H$ )  $\leftarrow e$                                 ▷ Unpack the state that defines the expert.

if  $|H| \leq 2^{2\sqrt{d}}$ :
    return HALVING.PREDICT( $H, x$ )                    ▷ Once  $H$  becomes small enough, simulate the Halving algorithm (Algorithm 7). [Case I]

if  $x \preceq u$ :
    return  $b \in \{0, 1\}$  such that  $x \preceq_b u$         ▷  $u$  is assumed to be on-path. If  $u$  is a  $b$ -descendant of  $x$ , then the correct label for  $x$  must be  $b$ . [Case II]

return  $\mathbb{1}(|\{x' \in S : x \preceq_1 x'\}| > |S|/3)$     ▷ Output some  $b \in \{0, 1\}$  such that more than 1/3 of suspected on-path nodes are  $b$ -descendants of  $x$ , if such a  $b$  exists. Otherwise (when at least 1/3 of  $S$  are non-descendants of  $x$ ), output 0. [Cases III to VI]

```

Algorithm 6a: A subroutine of Algorithm 5 that defines how an expert makes predictions.

Assumptions:

- x, e, S, u, H – as in Algorithm 6a.
- y – the correct label for x , as selected by the adversary.

EXPERT.BASICUPDATE(e, x, y):

```

( $S, u, H$ )  $\leftarrow e$                                 ▷ Unpack the state that defines the expert.

 $H \leftarrow$  HALVING.UPDATE( $H, x, y$ )                ▷ Update the version space, as in the Halving algorithm (Algorithm 7).

return ( $S, u, H$ )

```

Algorithm 6b: A subroutine of Algorithm 5 that defines how an expert is updated each time that a label is selected by the adversary.

Assumptions:

- d, x, e, S, u, H – as in Algorithm 6a.
- y – the correct label for x , as selected by the adversary.

EXPERT.EXTENDEDUPDATE(e, x, y):

$(S, u, H) \leftarrow e$	▷ Unpack the state that defines the expert.
if $ H \leq 2^{2\sqrt{d}}$:	▷ If the version space is small, we just simulate the Halving algorithm, so the update is complete. [Case III]
return $\{(S, u, H)\}$	
for $b \in \{0, 1\}$:	
$S_b \leftarrow \{x' \in S : x \preceq_b x'\}$	▷ Set of suspected on-path nodes that are b -descendant of x .
if $ S_{(1-y)} > S /3$:	
$S' \leftarrow S \setminus S_{(1-y)}$	▷ At least 1/3 of suspected on-path nodes were b -descendants of x , and therefore the expert predicted label $\hat{y} = b$. But the correct label was $y = 1 - b$. Remove all b -descendants of x from S . [Case IV]
return $\{(S', u, H)\}$	
else:	
$S_{\notin} \leftarrow S; u_{\notin} \leftarrow u$	▷ Split e in two. First, construct e_{\notin} to be an updated version of e after adding the assumption that $x \notin \text{path}(h)$ for the correct labeling function h .
$H_{\notin} = \{h \in H : x \notin \text{path}(h)\}$	
$e_{\notin} \leftarrow (S_{\notin}, u_{\notin}, H_{\notin})$	
$S_{\in} \leftarrow S_0 \cup S_1$	▷ Next, construct e_{\in} to be an updated version of e adding the assumption $x \in \text{path}(h)$. S_{\in} contains only nodes that are descendants of x .
$u_{\in} \leftarrow u$	▷ u_{\in} represents updating the prior assumption that u is on path by adding that x is also on path.
if $u_{\in} \preceq x$:	
$u_{\in} \leftarrow x$	
$H_{\in} = \{h \in H : x \in \text{path}(h)\}$	▷ H_{\in} is obtained by updating the version space to include only function where x is on path.
$e_{\in} \leftarrow (S_{\in}, u_{\in}, H_{\in})$	
return $\{e_{\notin}, e_{\in}\}$	▷ [Cases V and VI]

Algorithm 6c: A subroutine of Algorithm 5 that defines how an expert is updated (and possibly split into two) when it makes a mistake.

Assumptions:

- \mathcal{X} a set, $k \in \mathbb{N}$.
- $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ is a finite hypothesis class.
- $x, x_1, \dots, x_k \in \mathcal{X}, y \in \{0, 1\}$.

HALVING($\mathcal{H}, (x_1, x_2, \dots, x_k)$):

$\mathcal{H}_1 \leftarrow \mathcal{H}$

for $i \in [k]$:

$\hat{y}_i \leftarrow \text{HALVING.PREDICT}(\mathcal{H}, x_i)$

send prediction \hat{y}_i to adversary

receive correct label $y_i \in \{0, 1\}$ from adversary

$\mathcal{H}_{i+1} \leftarrow \text{HALVING.UPDATE}(\mathcal{H}_i, x_i, y_i)$

HALVING.PREDICT(\mathcal{H}, x):

return $\mathbb{1}\left(\frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} h(x) \geq \frac{1}{2}\right)$

HALVING.UPDATE(\mathcal{H}, x, y):

return $\{h \in \mathcal{H} : h(x) = y\}$

Algorithm 7: This is the well-known halving algorithm. The experts in Algorithms 6a and 6c simulate this algorithm once their version space becomes small enough.