

The Extended Littlestone’s Dimension for Learning with Mistakes and Abstentions

Chicheng Zhang

Kamalika Chaudhuri

*Department of Computer Science and Engineering
University of California San Diego
9500 Gilman Drive 0404, La Jolla, CA, 92093, USA*

CHICHENGZHANG@UCSD.EDU

KAMALIKA@CS.UCSD.EDU

Abstract

This paper studies classification with an abstention option in the online setting. In this setting, examples arrive sequentially, the learner is given a hypothesis class \mathcal{H} , and the goal of the learner is to either predict a label on each example or abstain, while ensuring that it does not make more than a pre-specified number of mistakes when it does predict a label.

Previous work on this problem has left open two main challenges. First, not much is known about the optimality of algorithms, and in particular, about what an optimal algorithmic strategy is for any individual hypothesis class. Second, while the realizable case has been studied, the more realistic non-realizable scenario is not well-understood. In this paper, we address both challenges. First, we provide a novel measure, called the Extended Littlestone’s Dimension, which captures the number of abstentions needed to ensure a certain number of mistakes. Second, we explore the non-realizable case, and provide upper and lower bounds on the number of abstentions required by an algorithm to guarantee a specified number of mistakes.

1. Introduction

Many machine learning applications, such as fraud detection in credit card transactions and medical diagnosis, involve high misclassification penalties. In these cases, it is often desirable to design a predictor which is guaranteed to make no more than a certain number of prediction mistakes, at the expense of a few abstentions; the examples on which the learner abstains are then passed on to a human, who can then take a closer look.

In this paper, we study this problem – classification with an option to abstain – in an online setting. Examples arrive sequentially, the learner is given a hypothesis class \mathcal{H} , and the goal of the learner is to either predict a label on each example or abstain, while ensuring that it does not make more than a pre-specified number of mistakes. The first formal theoretical framework to address this problem was the Knows What It Knows (KWIK) model (Li et al., 2011); assuming that labels are generated by a hypothesis in \mathcal{H} (the so-called *realizability assumption*), this framework requires the learner to *always* predict correctly. Li et al. (2011) also provided an algorithm in this model that makes no mistakes and requires $|\mathcal{H}| - 1$ abstentions for finite \mathcal{H} . Sayedi et al. (2010) further extended this model to allow learners that do not always have to be correct, but are permitted to make upto a bounded number of mistakes. They showed that when k mistakes are allowed, the number of abstentions can be reduced to $(k + 1)|\mathcal{H}|^{1/(k+1)}$.

While previous work has looked at designing generic online learning algorithm in this setting, there are two remaining challenges. First not much is known about the optimality of these algorithms, and in particular, about what an optimal algorithmic strategy would be for any individual hypothesis class. The second challenge is to understand what happens in a more realistic scenario where the realizability assumption does not hold. While this has been studied in a regression setting (Szita and Szepesvári, 2011), not much is known about the classification case.

In this paper, we address both challenges. We first provide a new measure that, given a hypothesis class \mathcal{H} , captures how many abstentions are needed to ensure a certain number of mistakes, and we provide an optimal algorithm that achieves this number. Our measure is closely related to the notion of Littlestone’s dimension for online learning with no abstentions, and we call it the Extended Littlestone’s dimension. Formalizing this notion additionally allows us to extend our algorithm to infinite hypothesis classes; while algorithms were previously known for some specific infinite classes (Sayedi et al., 2010), no generic algorithm was known.

Next, we focus our attention on the non-realizable case. In this case, we make an l -bias assumption, which ensures that the labels are generated by a function that disagrees with some (unknown) hypothesis $h \in \mathcal{H}$ on at most l examples. We show that (at least some form of) this assumption is necessary; there exists a finite hypothesis class \mathcal{H} , such that when the l -bias assumption holds, any algorithm that abstains a finite number of times must make at least l mistakes. Moreover, there also exists an infinite hypothesis class \mathcal{H} with Littlestone’s dimension d such that any algorithm that abstains a finite number of times must make at least $l + d$ mistakes. To complement these lower bounds, we show that we can run a version of our algorithm when the l -bias assumption holds, and provide an upper bound on the number of abstentions it makes.

2. The Setting

We consider the problem of online classification in the model of Sayedi et al. (2010), where the learner is allowed to occasionally abstain from prediction. The precise setting is as follows. At time t , the adversary presents an example x_t in some instance space \mathcal{X} . The learner makes its prediction \hat{y}_t , which can be either -1 , $+1$, or \perp (I don’t know). The adversary then reveals an outcome $y_t \in \{-1, +1\}$. The interaction between the learner and the adversary continues, and the performance of the learner is measured by the total number of mistakes and abstentions made throughout the process.

To help make decisions, the learner has access to a hypothesis class \mathcal{H} . Each hypothesis h in \mathcal{H} is a prediction rule mapping from \mathcal{X} to $\{-1, +1\}$.

Basic Notations. Given two hypotheses h_1 and h_2 , their *product* is defined as a new hypothesis $h_1 \cdot h_2$ which is a function that takes x as input, and outputs $h_1(x) \cdot h_2(x)$. Given two hypothesis classes \mathcal{H}_1 and \mathcal{H}_2 , we define $\mathcal{H}_1 \cdot \mathcal{H}_2$ to be the class of functions achievable by taking the product between a function in \mathcal{H}_1 and a function in \mathcal{H}_2 . Formally,

$$\mathcal{H}_1 \cdot \mathcal{H}_2 = \{h_1 \cdot h_2 : h_1 \in \mathcal{H}_1, h_2 \in \mathcal{H}_2\}$$

Define function $I(A) = 1$ if A is true, $I(A) = 0$ if A is false.

Denote by \mathcal{C}^l the class of union of at most l singletons in instance domain \mathcal{X} . That is, hypotheses that take value $+1$ on \mathcal{X} , except for at most l points:

$$\mathcal{C}^l = \{1 - 2I(x = x_1 \vee \dots \vee x = x_i) : x_1, \dots, x_i \in \mathcal{X}, i \leq l\}$$

In this paper, we will address both realizable and nonrealizable cases, defined below.

Realizable Case. In the realizable case, we assume there is a hypothesis h in \mathcal{H} that makes no mistakes over time. Formally, a sequence $S = (x_1, y_1), \dots, (x_n, y_n)$ is called \mathcal{H} -realizable, if and only if

$$\exists h \in \mathcal{H}, \quad |\{t : h(x_t) \neq y_t\}| = 0$$

Non-realizable Case. In the non-realizable case, we assume that the label of examples are generated by a function that disagrees with some hypothesis on at most l examples, which we call l -bias assumption.¹ Formally, define \mathcal{H}^l as the set of classifiers where its prediction differs from some classifier in \mathcal{H} on at most l points, i.e. $\mathcal{H}^l = \mathcal{H} \cdot \mathcal{C}^l$. A sequence $S = (x_1, y_1), \dots, (x_n, y_n)$ is said to have l -bias with respect to \mathcal{H} , if and only if it is \mathcal{H}^l -realizable, i.e.

$$\exists h \in \mathcal{H}^l, \quad |\{t : h(x_t) \neq y_t\}| = 0$$

Version Space and Disagreement Region. In the realizable case, it is often convenient to consider the set of hypotheses that agree with the labeled examples revealed so far. Given an labeled set S and a set of hypotheses V , $V[S] \subseteq V$ is defined as the set of all classifiers that classify S correctly:

$$V[S] = \{h \in V : \text{for all } (x, y) \in S, h(x) = y\}$$

At the start of time t , the version space is defined as the set of hypotheses in \mathcal{H} that agree with the examples $(x_1, y_1), \dots, (x_{t-1}, y_{t-1})$ seen so far.

We say that an example x is in the disagreement region of a hypothesis set V , denoted by $\text{DIS}(V)$, if both $V[(x, +1)]$ and $V[(x, -1)]$ are nonempty.

Nontrivial Rounds. For deterministic learners, it is always suboptimal for the adversary to present an example on which the learner predicts correctly, since this will only impose additional constraints on examples shown in the future without changing the number of mistakes and abstentions made. A round t is called *nontrivial* if and only if the learner makes a mistake or abstains on that round.

The Mistake Bound Model. The mistake bound model (Barzdin and Freivald, 1972; Littlestone, 1987; Angluin, 1987) is a central model in online classification. An online learning algorithm \mathcal{A} achieves a mistake bound M with respect to a set of sequences $\mathcal{S} \subseteq (\mathcal{X} \times \{-1, +1\})^*$ if and only if for any adversary showing sequences $S = ((x_1, y_1), \dots, (x_n, y_n))$ in \mathcal{S} , \mathcal{A} 's prediction $\hat{y}_1, \dots, \hat{y}_n \in \{-1, +1\}$ satisfies

$$\sum_{t=1}^n I(\hat{y}_t \neq y_t) \leq M$$

The k -SZB Model. In this paper, we consider the k -SZB model² studied by Sayedi et al. (2010); Demaine and Zadimoghaddam (2013), which extends the mistake bound model by additionally allowing the learner to say ‘‘Don't Know’’(\perp). An online learning algorithm \mathcal{A} achieves a (k, m) -SZB bound with respect to a set of sequences $\mathcal{S} \subseteq (\mathcal{X} \times \{-1, +1\})^*$, if and only if for any adversary that presents sequences $S = ((x_1, y_1), \dots, (x_n, y_n))$ in \mathcal{S} , \mathcal{A} 's prediction $\hat{y}_1, \dots, \hat{y}_n \in \{-1, +1, \perp\}$ satisfies

$$\sum_{t=1}^n I(\hat{y}_t \neq -y_t) \leq k$$

1. This is similar to the l -mistake assumption in the expert problem (Cesa-Bianchi et al., 1996; Abernethy et al., 2006).

2. The name SZB is from the author names of the paper Sayedi et al. (2010)

$$\sum_{t=1}^n (I(\hat{y}_t = -y_t) + I(\hat{y}_t = \perp)) \leq m$$

In other words, the number of mistakes is at most k , and number of nontrivial rounds (where the algorithm makes a mistake or abstains) is at most m . When $k = 0$, this is exactly the KWIK model (Li et al., 2011). We will look at this model in both the realizable and non-realizable cases.

3. Extended Littlestone’s Dimension

We begin with the realizable case and the definition of the Extended Littlestone’s Dimension. We first define an extended mistake tree, which is a natural generalization of the mistake tree, and then use it to characterize the optimal number of non-trivial rounds (abstentions + mistakes) for any algorithm in the k -SZB model. We finally present an optimal algorithm (Algorithm 2) for this model, and a recursive formulation of Extended Littlestone’s dimension.

3.1. Background: Mistake Bound, Littlestone’s Dimension and Standard Optimal Algorithm

Littlestone (1987) provides a characterization of the optimal mistake bound in the realizable case, which is measured by Littlestone’s dimension. We begin by describing this characterization.

Mistake Trees. Littlestone’s dimension is closely related to the notion of a mistake tree. A mistake tree³ of a hypothesis class \mathcal{H} is a complete binary tree⁴, whose leaves are classifiers in \mathcal{H} and whose internal nodes correspond to examples in \mathcal{X} . A mistake tree may have no internal nodes, in which case it only contains a leaf corresponding to a classifier h in \mathcal{H} – we call it a zeroth order mistake tree. Given an internal node, the edge connecting it and its left (resp. right) child is labeled -1 (resp. $+1$).

A root to leaf path p in mistake tree T is a sequence of nodes and edges $v_1 e_1 v_2 e_2 \dots v_n e_n v_{n+1}$, where v_1, \dots, v_n are internal nodes in T corresponding to examples in \mathcal{X} , v_1 is the root node of T , each e_i is an edge in T that connects v_i and v_{i+1} , $v_{n+1} = h$ is a classifier in \mathcal{H} corresponding to a leaf in T . For each i , edge e_i connects v_i and v_{i+1} . The length of a path $l(p)$ is defined as the number of edges in p . For each leaf, the associated classifier agrees with the internal nodes and edges along the path up to the root. That is, if each node v_i corresponds to example x_i and each edge e_i has label y_i , then h agrees with examples $\{(x_1, y_1), \dots, (x_n, y_n)\}$. See Figure 1 for an illustration.

A mistake tree T succinctly represents a strategy of the adversary in response to a deterministic learner. At $t = 1$, the adversary picks the example x_1 corresponding to the root node to show to the learner. If the learner predicts $\hat{y}_1 = -1$, the adversary reveals label $y_1 = +1$, and follows the downward edge labeled $+1$; otherwise it follows the other edge. If at time $t \geq 2$, the adversary reaches a node with example x_t , then x_t is shown to the learner, and one of the downward edges adjacent to this node is followed. The interaction comes to an end when a leaf is reached. It can be seen that the adversary forces the learner to make a mistake at each node of the mistake tree; this implies that if every root-to-leaf path of the mistake tree has depth d , then the adversary can force the learner to make d mistakes using the associated strategy.

We are now ready to define Littlestone’s dimension.

3. In Littlestone (1987) this is instead called a “complete mistake tree”; in Shalev-Shwartz (2012) this is called a \mathcal{H} -shattered tree.

4. A complete binary tree is one in which every level is completely filled with nodes.




Figure 1: A mistake tree with respect to set of threshold classifiers $\mathcal{H} = \{h_i = 2I(x \leq i) - 1 : i = 1, 2, 3, 4\}$.

Definition 1 *The Littlestone's dimension of hypothesis class \mathcal{H} , $\text{Ldim}(\mathcal{H})$, is the maximum depth of any mistake tree of \mathcal{H} .*

Theorem 2 ((Littlestone, 1987)) *For a hypothesis class \mathcal{H} , the optimal mistake bound of any deterministic algorithm with respect to adversaries showing \mathcal{H} -realizable sequences is equal to $\text{Ldim}(\mathcal{H})$.*

Standard Optimal Algorithm. Algorithm 1 presents the Standard Optimal Algorithm, which is an optimal deterministic algorithm for online classification in the realizable case. It maintains a version space V over time. At each time t , it predicts a label y_t such that each mistake will force the version space's Littlestone's dimension to drop by at least 1. Therefore, the number of mistakes made by Algorithm 1 is at most $\text{Ldim}(\mathcal{H})$.

Algorithm 1 Standard Optimal Algorithm (Littlestone, 1987)

- 1: Input: hypothesis class \mathcal{H} .
 - 2: Initialize version space $V \leftarrow \mathcal{H}$.
 - 3: **for** $t = 1, 2, \dots$, **do**
 - 4: Receive example $x_t \in \mathcal{X}$.
 - 5: Make prediction $\hat{y}_t = \arg \max_y \text{Ldim}(V[(x, y)])$.
 - 6: Receive label y_t .
 - 7: **if** $\hat{y}_t = -y_t$ **then**
 - 8: Update version space $V \leftarrow V[(x_t, y_t)]$.
 - 9: **end if**
 - 10: **end for**
-

Recursive Definition. For finite \mathcal{H} ⁵ one also has the following recurrence for its Littlestone's dimension:

$$\text{Ldim}(\mathcal{H}) = \begin{cases} 0 & |\mathcal{H}| = 1 \\ 1 + \max_{x \in \text{DIS}(\mathcal{H})} \min_{y \in \{\pm 1\}} \text{Ldim}(\mathcal{H}[(x, y)]) & |\mathcal{H}| > 1 \end{cases}$$

3.2. Extended Littlestone's Dimension

We now define extended Littlestone's Dimension, which measures the difficulty of online learning a hypothesis class in the k -SZB model.

5. For an infinite \mathcal{H} , the recurrence may not reach the base case.

Extended Mistake Trees. An adversary’s strategy in response to a deterministic learner in the k -SZB model can be succinctly represented by extended mistake trees. An extended mistake tree for \mathcal{H} is a full⁶ binary tree, whose leaves are classifiers in \mathcal{H} and whose internal nodes are examples in \mathcal{X} . An extended mistake tree may have no internal node, in which case it only contains a leaf corresponding to a classifier h in \mathcal{H} – we call it a zeroth order extended mistake tree. Unlike mistake trees, now, there are two type of edges: solid and dashed, representing mistakes and abstentions, respectively. Each node is associated with two downward solid edges, one to each child. Additionally, each node is associated with exactly one downward dashed edge connecting to one of its two children. For a downward edge of a node, whether solid or dashed, if it is connected with the node’s left child, then it is labeled -1 , and vice versa. Just as in mistake trees, for each leaf, the associated classifier agrees with the internal nodes and edges along the path up to the root. See Figure 2 for an illustration.

A root to leaf path p of an extended mistake tree T is a sequence of nodes and edges $v_1 e_1 v_2 e_2 \dots v_n e_n v_{n+1}$, where v_1, \dots, v_n are internal nodes in T corresponding to examples in \mathcal{X} , v_1 is the root node of T , each e_i is an edge in T that connects v_i and v_{i+1} , $v_{n+1} = h$ is a classifier in \mathcal{H} corresponding to a leaf in T . Here if there are multiple edges between v_i and v_{i+1} , any one of them can be used by p .

Given an extended mistake tree T , the associated adversarial strategy can be described as follows. At $t = 1$, the adversary chooses the example x_1 corresponding to the root node to show to the learner. If the learner predicts $\hat{y}_1 = -1$, it reveals label $y_1 = +1$, follows the downward solid edge labeled $+1$, and vice versa. Otherwise, if $\hat{y}_1 = \perp$, it reveals y_1 as the label on the dashed edge and follows the downward dashed edge. At time $t \geq 2$, if the adversary reaches a node with example x_t , then x_t is shown to the learner, and one of its adjacent downward edges is followed. The interaction comes to an end when a leaf is reached. It can be seen that with this strategy, the adversary forces every round to be nontrivial. If the depth of the leaf reached is d , then the number of nontrivial rounds is d .

As an example, the extended mistake tree in Figure 2 can be used by the adversary as follows. Initially $x_1 = 2$ is presented to the learner. If the learner predicts $\hat{y}_1 = -1$, the adversary reveals label $y_1 = +1$ and follows the right downward solid edge to reach node $x_2 = 3$. At time $t = 2$, the learner now shows example $x_2 = 3$; If the learner predicts $\hat{y}_2 = \perp$, the adversary reveals $y_2 = +1$ according to the label on the dashed edge and follows the edge to reach node $x_3 = 4$. At time $t = 3$, the learner shows example $x_3 = 4$; If the learner predicts $\hat{y}_2 = +1$, the adversary reveals label $y_2 = -1$ and follows the left downward solid edge to reach a leaf containing hypothesis h_3 . This concludes the interaction, and the learner makes a total of 3 nontrivial rounds: 2 mistakes and 1 abstentions. Note that realizability assumption is maintained, as $h_3 \in \mathcal{H}$ agrees with the examples $(2, +1)$, $(3, +1)$, $(4, -1)$ shown. More generally, one can show that if the learner is not allowed to make any mistakes, then the adversary is able to force 3 nontrivial rounds by following this strategy. This motivates the definition below.

Definition 3 We say that an extended mistake tree T is (k, m) -difficult for integers $k, m \geq 0$, if all its root to leaf paths in T using at most k solid edges have length at least m .

For example, the extended mistake tree in Figure 2 is $(0, 3)$ -difficult.

Extended Standard Optimal Algorithm. Algorithm 2 presents the Extended Standard Optimal Algorithm (SOA.DK), which, as we will show, is an optimal deterministic algorithm for online

6. A full binary tree is one in which every internal node has exactly two children.




Figure 2: An extended mistake tree with respect to set of threshold classifiers $\mathcal{H} = \{h_i = 2I(x \leq i) - 1 : i = 1, 2, 3, 4\}$.

prediction in the k -SZB model in the realizable case. Note that it works even when the hypothesis class \mathcal{H} is infinite. Similar to the Standard Optimal Algorithm, it maintains a version space V . For a new example x_t , it predicts $\hat{y}_t \in \{-1, +1, \perp\}$ by computing function ELdim over subsets of V . The function ELdim is defined as follows.

Definition 4 (Extended Littlestone's Dimension) For a hypothesis class V and integer $k \geq 0$, the extended Littlestone's dimension $\text{ELdim}(V, k)$ is defined as:

$$\text{ELdim}(V, k) := \sup\{m \in \mathbb{N} : \text{There exists a } (k, m)\text{-difficult extended mistake tree for } V\}$$

We remark that if for every integer m , V has a (k, m) -extended mistake tree, then $\text{ELdim}(V, k) = \infty$; If $V = \emptyset$, then $\text{ELdim}(V, k) = -\infty$. Since for $k' < k$, a (k, m) -difficult extended mistake tree is also (k', m) -difficult, $\text{ELdim}(V, k)$ is nonincreasing with respect to k .

We first show that when $\text{ELdim}(V, k)$ is high, then an adversary can force a large number of nontrivial rounds by showing a V -realizable sequence, to any deterministic algorithms that guarantees at most k mistakes.

Lemma 5 Suppose we are given a hypothesis set V and integers $k \geq 0, m \geq 0$. If $\text{ELdim}(V, k) \geq m$, then there is a strategy of the adversary that presents a V -realizable sequence and that can force any deterministic algorithm that guarantees $\leq k$ mistakes to have $\geq m$ nontrivial rounds.

In the following lemma, we show that given a mistake budget k , if the extended Littlestone's dimension of V is small, then SOA.DK has a small number of nontrivial rounds for V -realizable sequences.

Lemma 6 (Performance Guarantees of SOA.DK) Suppose we are given a hypothesis class V and integers $k \geq 0, m \geq 0$. If $\text{ELdim}(V, k) \leq m$, then Algorithm 2, when run on V with mistake budget k , achieves a (k, m) -SZB bound with respect to any adversary that shows V -realizable sequences.

An immediate consequence of Lemma 6 is that SOA.DK is optimal, in the sense that it has the smallest number of worst case nontrivial rounds, amongst all *deterministic* algorithms that work in k -SZB model.

Algorithm 2 Extended Standard Optimal Algorithm: SOA.DK

```

1: Input: hypothesis class  $\mathcal{H}$ , mistake budget  $k$ .
2: Initialize version space  $V \leftarrow \mathcal{H}$ .
3: for  $t = 1, 2, \dots$ , do
4:   Receive example  $x_t \in \mathcal{X}$ .
5:   if  $x_t \notin \text{DIS}(V)$  then                                     # All classifiers in  $V$  predict unanimously
6:     Predict  $\hat{y}_t = h(x_t)$ , where  $h$  is an arbitrary hypothesis in  $V$ .
7:   else                                                         # There is disagreement among  $V$ 
8:     if  $k = 0$  then                                           # Zero mistake budget, must output  $\perp$ 
9:       Predict  $\hat{y}_t = \perp$ .
10:    else                                                         # Predict by minimizing the ELdim of future version space
11:      Compute  $m_{+1} = \text{ELdim}(V[(x_t, -1)], k - 1)$ ,  $m_{-1} = \text{ELdim}(V[(x_t, +1)], k - 1)$ , and
12:       $m_{\perp} = \max(\text{ELdim}(V[(x_t, -1)], k), \text{ELdim}(V[(x_t, +1)], k))$ 
13:      Predict  $\hat{y}_t = \arg \min\{m_y : y \in \{-1, +1, \perp\}\}$ .
14:    end if
15:  end if
16:  Receive label  $y_t$ .
17:  if  $\hat{y}_t = -y_t$  or  $\hat{y}_t = \perp$  then  $V \leftarrow V[(x_t, y_t)]$  end if           # Update version space
18:  if  $\hat{y}_t = -y_t$  then  $k \leftarrow k - 1$  end if           # Update mistake budget
19: end for
    
```

Theorem 7 (Optimality of SOA.DK) *Suppose we are given a hypothesis class \mathcal{H} and integers $k \geq 0$, $m \geq 1$ such that $\text{ELdim}(\mathcal{H}, k) = m$. Then:*

- (a) *SOA.DK achieves a (k, m) -SZB bound for any adversary that shows \mathcal{H} -realizable sequences.*
- (b) *There exists an adversary showing \mathcal{H} -realizable sequences, such that no deterministic algorithm \mathcal{A} can achieve a $(k, m - 1)$ -SZB bound.*

The following simple property relates extended Littlestone’s dimension to Littlestone’s dimension.

Theorem 8 (Relating Ldim to ELdim) *Suppose we are given a hypothesis class \mathcal{H} . If $\text{Ldim}(\mathcal{H}) = d < \infty$, then*

$$\text{ELdim}(\mathcal{H}, d) = d$$

Recursive Definition. We provide a recursive characterization of Extended Littlestone’s dimension. For finite \mathcal{H} ⁷, the following recurrence holds for its extended Littlestone’s dimension:

$$\text{ELdim}(\mathcal{H}, k) = \begin{cases} 0 & |\mathcal{H}| = 1 \\ 1 + \max_{x \in \text{DIS}(\mathcal{H})} \max_{y \in \{\pm 1\}} \text{ELdim}(\mathcal{H}[(x, y)], 0) & |\mathcal{H}| > 1, k = 0 \\ 1 + \max_{x \in \text{DIS}(\mathcal{H})} \max_{y \in \{\pm 1\}} \min(\text{ELdim}(\mathcal{H}[(x, y)], k - 1), \text{ELdim}(\mathcal{H}[(x, -y)], k)) & |\mathcal{H}| > 1, k \geq 1 \end{cases}$$

7. Just as with Littlestone’s dimension, for an infinite \mathcal{H} , the recurrence may not reach the base case.




Figure 3: A depth-3 \mathcal{X} -valued tree \mathbf{x} .

The recurrence is an immediate consequence of Lemma 32 in Appendix C.

4. Properties of Extended Littlestone's Dimension

We next present upper bounds on the Extended Littlestone's Dimension of a hypothesis class \mathcal{H} . Our upper bounds depend on the *tree shattering coefficient*, a notion analogous to the growth function, which is implicit in Ben-David et al. (2009). We also present some examples of Extended Littlestone's Dimension.

4.1. Tree Shattering Coefficient

The shattering coefficient (also known as the growth function), initially studied in (Vapnik and Chervonenkis, 1971), is a key notion in PAC learnability.

Definition 9 Given a hypothesis \mathcal{H} , the shattering coefficient of \mathcal{H} , $\Pi(\mathcal{H}, t)$ is defined as the maximum number of labelings achievable by \mathcal{H} over t points. Formally,

$$\Pi(\mathcal{H}, t) := \max_{x_1, \dots, x_t} |\{(h(x_1), \dots, h(x_t)) : h \in \mathcal{H}\}|$$

Inspired by the shattering coefficient, in online learning, we define the notion of *tree shattering coefficient* below, implicit in (Ben-David et al., 2009). As we will see, this notion is crucial to online learnability in both the mistake bound and the k -SZB models. First we set up our notation by adopting the notion of trees in (Rakhlin et al., 2010).

Definition 10 (\mathcal{X} -valued Trees, see (Rakhlin et al., 2010)) A depth- t \mathcal{X} -valued tree \mathbf{x} is a series of mappings $(\mathbf{x}_1, \dots, \mathbf{x}_t)$, where $\mathbf{x}_i : \{\pm 1\}^{i-1} \rightarrow \mathcal{X}$. The root of the tree \mathbf{x} is the constant function $\mathbf{x}_1 \in \mathcal{X}$. For integer t , the mapping $\mathbf{x}_t(\cdot)$ takes care of the nodes in level t .

To see why a series of mappings corresponds to a tree, we first note that a tuple $(\epsilon_1, \dots, \epsilon_{s-1})$ in $\{\pm 1\}^{s-1}$ can be thought of as a left/right sequence of length $s-1$, where -1 stands for left and $+1$ stands for right, respectively. The node reached from the root following the path corresponding to the left/right sequence corresponds to $\mathbf{x}_s(\epsilon_1, \dots, \epsilon_{s-1}) \in \mathcal{X}$. For example, the root node corresponds to $\mathbf{x}_1 \in \mathcal{X}$, the left child of the root corresponds to $\mathbf{x}_2(-1) \in \mathcal{X}$, the right child of the left child of the root corresponds to $\mathbf{x}_3(-1, +1) \in \mathcal{X}$, etc. See Figure 3 for an illustration. We slightly abuse the notation to let $\mathbf{x}_t(\epsilon)$ denote $\mathbf{x}_t(\epsilon_1, \dots, \epsilon_{t-1})$.




Figure 4: A concrete depth-3 \mathcal{X} -valued tree \mathbf{x} , where $\mathbf{x}_1 = z_1$, $\mathbf{x}_2(-1) = z_2$, $\mathbf{x}_2(+1) = z_3$, $\mathbf{x}_3(-1, -1) = z_4$, $\mathbf{x}_3(-1, +1) = z_5$, $\mathbf{x}_3(+1, -1) = z_6$, $\mathbf{x}_3(+1, +1) = z_7$. There are 4 root to leaf paths that agrees with some hypothesis in \mathcal{H} (\times in a leaf indicates that no hypothesis in \mathcal{H} agree with the path from root to it), i.e. $|S(\mathcal{H}, \mathbf{x})| = 4$.

Note that a \mathcal{X} -valued tree is not a mistake tree or an extended mistake tree, since it does not have leaves corresponding to hypotheses in \mathcal{H} .

Given a \mathcal{X} -valued tree \mathbf{x} , we add an extra level of edges at the bottom. Specifically for each leaf $\mathbf{x}_t(\epsilon)$, we attach a left and a right downward edge onto it, labeled -1 and $+1$ respectively. Now, consider every root to leaf path in the tree. If there is some classifier h in \mathcal{H} that agrees with the path, we label the leaf with h ; otherwise we label the leaf with symbol \times . We count the number of leaves not labeled \times , denoted by function $S(\mathcal{H}, \mathbf{x})$. See Figure 4 for an example.

Definition 11 For a depth- t \mathcal{X} -valued tree \mathbf{x} , and a hypothesis class \mathcal{H} , define function $S(\mathcal{H}, \mathbf{x})$ as the maximum number of labelings achievable by \mathcal{H} on \mathbf{x} . Formally,

$$S(\mathcal{H}, \mathbf{x}) := \{(\epsilon_1, \epsilon_2, \dots, \epsilon_t) \in \{\pm 1\}^t : \epsilon_1 = h(\mathbf{x}_1(\epsilon)), \epsilon_2 = h(\mathbf{x}_2(\epsilon)), \dots, \epsilon_t = h(\mathbf{x}_t(\epsilon)), \text{ for some } h \in \mathcal{H}\}$$

Definition 12 Given hypothesis class \mathcal{H} , and integer $t \geq 1$, the tree shattering coefficient of \mathcal{H} , $\mathcal{S}(\mathcal{H}, t)$ is defined as the maximum number of labelings achievable by \mathcal{H} over all depth- t trees. Formally,

$$\mathcal{S}(\mathcal{H}, t) := \max_{\mathbf{x}} |S(\mathcal{H}, \mathbf{x})|$$

Additionally, define $\mathcal{S}(\mathcal{H}, 0) := 1$ if \mathcal{H} is nonempty, $\mathcal{S}(\mathcal{H}, 0) := 0$ if \mathcal{H} is empty.

In other words, given hypothesis class \mathcal{H} and a depth- t tree \mathbf{x} with internal nodes only, there are at most $\mathcal{S}(\mathcal{H}, t)$ distinct paths in T consistent with some classifier $h \in \mathcal{H}$. Since a depth- t has at most 2^t root to leaf paths, $\mathcal{S}(\mathcal{H}, t) \leq 2^t$. Note that if \mathcal{H} has a depth- t mistake tree, then $\mathcal{S}(\mathcal{H}, t) = 2^t$. If we constrain the trees chosen to be constant among nodes in the same depth, then the tree shattering coefficient is equivalent to the shattering coefficient. In Appendix A, we show that the tree shattering coefficient is connected with the sequential growth function (maximal sequential zero covering number), defined in (Rakhlin et al., 2010).

The following two lemmas give bounds on tree shattering coefficients, implicit in [Rakhlin et al. \(2010\)](#); [Ben-David et al. \(2009\)](#). For finite hypothesis class \mathcal{H} , its tree shattering coefficient is at most the size of $|\mathcal{H}|$.

Lemma 13 *If \mathcal{H} is finite, then for any $t \geq 0$, $\mathcal{S}(\mathcal{H}, t) \leq |\mathcal{H}|$.*

Furthermore, if an infinite hypothesis class \mathcal{H} has Littlestone's dimension $d < \infty$, its tree shattering coefficient is polynomial in t , that is, $O(t^d)$.

Lemma 14 *If \mathcal{H} has Littlestone's dimension $d < \infty$, then for any $t \geq 0$, $\mathcal{S}(\mathcal{H}, t) \leq \binom{t}{\leq d}$.*

4.2. Upper Bound on Extended Littlestone's Dimension

We present [Theorem 15](#), the main result of this section, which upper bounds the extended Littlestone's dimension in terms of tree shattering coefficient. Intuitively, if \mathcal{H} is not expressive, then it has small tree shattering coefficient, and a tighter upper bound on its extended Littlestone's dimension can be established. Note that the bound is valid even if \mathcal{H} is infinite, and hence it is a strict generalization of [Sayedi et al. \(2010\)](#).

Theorem 15 *For any hypothesis class \mathcal{H} and integer $k \geq 0$,*

$$\text{ELdim}(\mathcal{H}, k) \leq \sup \left\{ t : \binom{t}{\leq k+1} \leq \mathcal{S}(\mathcal{H}, t) \right\}$$

For finite hypothesis classes one has the following corollary.

Corollary 16 *For any hypothesis class \mathcal{H} such that $|\mathcal{H}| < \infty$ and integer $k \geq 0$,*

$$\text{ELdim}(\mathcal{H}, k) \leq \max \left\{ t : \binom{t}{\leq k+1} \leq |\mathcal{H}| \right\}$$

Since $\binom{t}{\leq k+1} \geq \left(\frac{t}{k+1}\right)^{k+1}$, this implies $\text{ELdim}(\mathcal{H}, k) \leq \max \left\{ t : \left(\frac{t}{k+1}\right)^{k+1} \leq |\mathcal{H}| \right\} \leq (k+1)|\mathcal{H}|^{\frac{1}{k+1}}$, which recovers the result of [Sayedi et al. \(2010\)](#).⁸

4.3. Case Study: Thresholds (Finite Class)

We give a precise characterization of the Extended Littlestone's dimension for the class of thresholds. In this case, the bound given by [Theorem 15](#) is tight.

Consider the instance domain \mathcal{X} being \mathbb{R} and the hypothesis class \mathcal{H} being the set of n distinct threshold functions $\{2I(x \leq t) - 1 : t \in \{t_1, \dots, t_n\}\}$.⁹

8. Although it is implicit in [Sayedi et al. \(2010\)](#) that the result can be refined by using the optimal solution of the Egg Dropping Game ([Gasarch and Fletcher, 2008](#); [Boardman, 2004](#)), here we give an alternative proof using a more general technique.

9. Note that for an infinite set of thresholds, e.g. $\mathcal{H} = \{2I(x \leq t) - 1 : t \in [0, 1]\}$, $\text{Ldim}(\mathcal{H}) = \infty$, hence $\text{ELdim}(\mathcal{H}, k) = \infty$ for any finite k .

Theorem 17 Consider \mathcal{H} a set of threshold classifiers $\mathcal{H} = \{2I(x \leq t) - 1 : t \in \{t_1, \dots, t_n\}\}$. Then

$$\text{ELdim}(\mathcal{H}, k) = \max \left\{ t : \binom{t}{\leq k+1} \leq n \right\}$$

The proof of Theorem 17 is provided in Appendix D. The upper bound follows immediately from Corollary 16. The lower bound comes from an explicit construction of optimal extended mistake trees by exploiting structure in the class of threshold classifiers.

4.4. Case Study: Union of Singletons (Infinite Class)

We give a precise characterization of the Extended Littlestone’s dimension for the class of unions of singletons. In this case the bound given by Theorem 15 is tight. Consider the concept class of union of at most l singletons \mathcal{C}^l , with instance domain \mathcal{X} such that $|\mathcal{X}| = \infty$. Note that $\text{Ldim}(\mathcal{C}^l) = l$ and $S(\mathcal{C}^l, t) = \binom{t}{\leq l}$ (See Lemma 37 for a proof). We have the following result.

Theorem 18 Consider the hypothesis class \mathcal{C}^l , the class of union of at most l singletons. Then,

$$\text{ELdim}(\mathcal{C}^l, k) = \sup \left\{ t : \binom{t}{\leq k+1} \leq \binom{t}{\leq l} \right\} = \begin{cases} \infty, & k \leq l-1 \\ l, & k \geq l \end{cases}$$

Note that Theorem 18 involves infinite hypothesis classes and is broader than the results of Sayedi et al. (2010). The proof of Theorem 18 is provided in Appendix D. The upper bound follows immediately from Theorem 15. The lower bound comes from an explicit construction of optimal extended mistake trees by exploiting structures in the class of union of singleton classifiers.

5. Non-Realizable Case

We now consider the non-realizable case. For the rest of the section, we assume the l -bias assumption holds, i.e. the sequence $(x_1, y_1), \dots, (x_n, y_n)$ presented by the adversary is \mathcal{H}^l -realizable. Recall that $\mathcal{H}^l = \mathcal{H} \cdot \mathcal{C}^l$, the class of hypothesis that disagrees with \mathcal{H} on at most l points.

5.1. Lower Bounds for Deterministic Prediction

A natural question is, when the l -bias assumption holds, is it possible to derive algorithms with a small number of abstentions in k -SZB model? Perhaps surprisingly, the answer depends on whether \mathcal{H} is finite or not. We show next that there is a finite hypothesis class \mathcal{H} , such that for any $k < l$, and any integer m , any algorithm which is guaranteed to make k or less mistakes can be forced to abstain at least m times. Moreover, for any m , there is a infinite hypothesis class with Littlestone’s dimension d , such that for any $k < l+d$, any algorithm that is guaranteed to make k or less mistakes can be forced to abstain at least m times.

5.1.1. LOWER BOUNDS FOR FINITE HYPOTHESIS CLASSES

We first show that, for finite hypothesis classes \mathcal{H} , when $k < l$, no algorithm can guarantee a (k, m) -SZB bound with finite m under the l -bias assumption.

Theorem 19 *There exists an instance domain \mathcal{X} , a single-element hypothesis class \mathcal{H} , such that the following holds. If $k < l$, then for any integer $m \geq 0$, there exists a strategy of the adversary satisfying the l -bias assumption that forces any deterministic algorithm guaranteeing at most k mistakes to have at least $m + 1$ nontrivial rounds.*

5.1.2. LOWER BOUNDS FOR INFINITE HYPOTHESIS CLASSES

We show that, given a hypothesis class \mathcal{H} with $\text{Ldim}(\mathcal{H}) = d$, when $k < l + d$, no algorithm can guarantee a (k, m) -SZB bound for finite m under the l -bias assumption.

Theorem 20 *There exists an instance domain \mathcal{X} , a hypothesis class \mathcal{H} with Littlestone's dimension $d < \infty$, such that the following holds. If $k < l + d$, then for any integer $m \geq 0$, there exists a strategy of the adversary satisfying the l -bias assumption that forces any deterministic algorithm guaranteeing at most k mistakes to have at least $m + 1$ nontrivial rounds.*

5.2. Upper Bounds

5.2.1. UPPER BOUNDS FOR FINITE HYPOTHESIS CLASSES

Since a sequence satisfying the l -bias assumption is \mathcal{H}^l -realizable, to provide an upper bound on the number of non-trivial rounds under this assumption, we need to provide an upper bound on $\text{ELdim}(\mathcal{H}^l, k)$. We now provide such upper bounds on arbitrary finite hypothesis classes \mathcal{H} . Note that since the hypothesis class \mathcal{H}^l is infinite, this result is more general than the kind of results in (Sayedi et al., 2010).

Lemma 21 *Suppose we are given a finite hypothesis class \mathcal{H} , integer $k \geq 0$, $l \geq 0$ such that $k \geq l$. Then,*

$$\text{ELdim}(\mathcal{H}^l, k) \leq e(k + 1) \cdot |\mathcal{H}|^{\frac{1}{k+1-l}}$$

Corollary 22 *Suppose we are given a finite hypothesis class \mathcal{H} and integers $k, l \geq 0$ such that $k \geq l$. If Algorithm SOA.DK is run with input hypothesis class \mathcal{H}^l and mistake budget k , then for any adversary that shows sequences satisfying the l -bias assumption with respect to \mathcal{H} , SOA.DK makes at most k mistakes and has at most $e(k + 1) \cdot |\mathcal{H}|^{\frac{1}{k+1-l}}$ nontrivial rounds.*

5.2.2. UPPER BOUNDS FOR INFINITE HYPOTHESIS CLASSES

We now derive a corresponding upper bound for infinite hypothesis classes \mathcal{H} with finite Littlestone's dimension.

Lemma 23 *Suppose we are given a hypothesis class \mathcal{H} with Littlestone's dimension $d < \infty$, integer $k \geq 0$, $l \geq 0$ such that $k \geq l + d$. Then,*

$$\text{ELdim}(\mathcal{H}^l, k) \leq (k + 1) \cdot e^{\frac{2k+2}{k+1-l-d}}$$

Corollary 24 *Suppose we are given a hypothesis class \mathcal{H} with Littlestone’s dimension d and integer $k, l \geq 0$ such that $k \geq l + d$. If Algorithm SOA.DK is run with input hypothesis class \mathcal{H}^l and mistake budget k , then for any adversary that shows sequences satisfying the l -bias assumption with respect to \mathcal{H} , SOA.DK makes at most k mistakes and has at most $(k + 1) \cdot e^{\frac{2k+2}{k+1-l-d}}$ nontrivial rounds.*

5.3. Lower Bounds for Randomized Prediction

We show that the results in Section 5.1 hold even when the learner makes soft predictions.

Randomized Prediction Model. Consider the following randomized variant of online classification model. At time t , the adversary presents example x_t in \mathcal{X} , and the learner outputs a tuple $(p_{t,-}, p_{t,+}, 1 - p_{t,-} - p_{t,+})$, with $p_{t,-} \geq 0$, $p_{t,+} \geq 0$ and $1 - p_{t,-} - p_{t,+} \geq 0$. The tuple $(p_{t,-}, p_{t,+}, 1 - p_{t,-} - p_{t,+})$ represents the learner’s strategy of predicting $+1$ with probability $p_{t,+}$, -1 with probability $p_{t,-}$ and abstaining with probability $1 - p_{t,+} - p_{t,-}$. The adversary then reveals an outcome $y_t \in \{-1, +1\}$, and the learner incurs a mistake penalty of $p_{t,+}$ if $y_t = -1$ and $p_{t,-}$ if $y_t = 1$; it also incurs an abstention penalty of $1 - p_{t,+} - p_{t,-}$. When $p_{t,+}$ and $p_{t,-}$ take values in $\{0, 1\}$, observe that this is equivalent to our prediction model in Section 2.

So given examples $(x_1, y_1), \dots, (x_n, y_n)$, the cumulative mistake penalty upto time n is as $\sum_{t=1}^n I(y_t = -1)p_{t,+} + I(y_t = +1)p_{t,-}$ and the cumulative abstention penalty is $\sum_{t=1}^n (1 - p_{t,+} - p_{t,-})$. We have the following result for finite hypothesis classes.

Theorem 25 *There exists an instance domain \mathcal{X} , a single-element hypothesis class \mathcal{H} , such that the following holds. If $k < l$, then for any $a \geq 0$, there exists a strategy of the adversary satisfying l -bias assumption, such that any algorithm guaranteeing a cumulative mistake penalty at most k in the randomized prediction model must have cumulative abstention penalty at least a .*

For infinite hypothesis classes with Littlestone’s dimension d , we have the following result.

Theorem 26 *There exists an instance domain \mathcal{X} and a hypothesis class \mathcal{H} with Littlestone’s dimension d such that the following holds. If $k < l + d$, then for any $a \geq 0$, there exists a strategy of the adversary satisfying the l -bias assumption, such that any algorithm guaranteeing a mistake penalty of at most k in the randomized prediction model must have cumulative abstention penalty at least a .*

6. Related Work

The first formal framework for online learning with abstentions is the Knows What It Knows (KWIK) model (Li et al., 2011), which works as follows. At time t , the learner is given an example x_t , and is expected to output either the correct label for x_t or abstain from prediction. Li et al. (2011) formalizes versions of this model for both classification and regression settings, and provides a classification algorithm that achieves no mistakes and $|\mathcal{H}| - 1$ abstentions when the sequences provided are realizable for a finite hypothesis class \mathcal{H} . Szita and Szepesvári (2011) provides algorithms for online regression that apply even when the realizability assumption is relaxed.

Perhaps the most related to our work is Sayedi et al. (2010), where the KWIK model is relaxed in the classification setting by allowing the learner to make $\leq k$ mistakes. This work presents an algorithm that, given a finite hypothesis \mathcal{H} , can make at most k mistakes with at most $(k + 1)|\mathcal{H}|^{\frac{1}{k+1}}$ abstentions. Our work extends Sayedi et al. (2010) in that we provide an optimal algorithm that

exploits finer structures in \mathcal{H} , and also in that we allow the hypothesis class \mathcal{H} to be infinite. Demaine and Zadimoghaddam (2013) extends Sayedi et al. (2010) by providing efficient algorithms for the class of disjunctions. Finally, another important line of work for online classification with abstentions is conformal prediction (Shafer and Vovk, 2008), which, given a conformity measure R and an error probability measure δ , shows a strategy for constructing confidence sets in an online manner that contain the correct label with probability $1 - \delta$. Our framework differs from this line of work in that the conformity measure for us is not specified.

There is a large volume of literature on online classification when no abstentions are allowed. The mistake bound model, initially proposed by (Littlestone, 1987; Angluin, 1987), considers online binary classification in the realizable case. Littlestone (1987) also introduces the standard optimal algorithm and optimal mistake bound (aka Littlestone’s dimension $\text{Ldim}(\mathcal{H})$). There has been much literature on developing algorithms for specific hypothesis classes in the mistake bound model; see Shalev-Shwartz and Ben-David (2014); Cesa-Bianchi and Lugosi (2006) for examples. Ben-David et al. (2009) considers online classification (with no abstentions) in the agnostic case; they show that if the hypothesis class \mathcal{H} has finite Littlestone’s dimension, then it is possible to design an online prediction algorithm that makes $l + \tilde{O}(\sqrt{\text{Ldim}(\mathcal{H})T} + \text{Ldim}(\mathcal{H}))$ mistakes over T rounds, where l is the minimum error of any hypotheses in \mathcal{H} . In follow-up work, (Rakhlin et al., 2010, 2012, 2015a,b) have developed a rich theory of online learning, and defined complexity measures such as sequential Rademacher complexity, and sequential covering number that characterize the complexity of online learning. However, this theory does not apply to online learning with abstentions.

In the batch setting, the problem of classification with an abstention option has been both empirically and theoretically studied since the pioneering work of Chow (1970). It is however unclear how to directly apply the results in the batch setting to the online setting, because of the adversarial nature of the examples. Herbei and Wegkamp (2006); Bartlett and Wegkamp (2008); Yuan and Wegkamp (2010) consider classification where the decision to abstain is made based on thresholding a real-valued function that belongs to a fixed function class. Freund et al. (2004) provides an algorithm that performs weighted majority style aggregation over a hypothesis class and abstains when the aggregate is close to zero. Kalai et al. (2012); Kanade and Thaler (2014) study a related problem called reliable learning, and gives a predictor that achieves low error at the expense of abstentions. Balsubramani (2016) considers the problem in transductive setting, where the goal is to make aggregated predictions with abstention based on an ensemble of classifiers, where some error upper bounds on individual classifiers are known. Finally, inspired by the active learning algorithm of (Cohn et al., 1994), El-Yaniv and Wiener (2010) proposes a abstention principle in the realizable case, which guarantees a zero error. El-Yaniv and Wiener (2011) shows how to extend the idea to nonrealizable case, where the predictor has zero error with respect to the optimal hypothesis and Zhang and Chaudhuri (2014) gives an improved predictor when a nonzero amount of error is allowed.

Acknowledgements. We thank NSF under IIS 1162581 for research support. CZ would like to thank Akshay Balsubramani and Haipeng Luo for helpful discussions.

References

Jacob Abernethy, John Langford, and Manfred K. Warmuth. Continuous experts and the binning algorithm. In *19th Annual Conference on Learning Theory, COLT 2006*, pages 544–558, 2006.

- Dana Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1987. doi: 10.1007/BF00116828.
- Akshay Balsubramani. Learning to abstain from binary prediction. *arXiv preprint arXiv:1602.08151*, 2016.
- P. L. Bartlett and M. H. Wegkamp. Classification with a reject option using a hinge loss. *JMLR*, 9, 2008.
- Ya M Barzdin and RV Freivald. Prediction of general recursive functions. *Doklady Akademii Nauk SSSR*, 206(3):521, 1972.
- Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. Agnostic online learning. In *COLT 2009 - The 22nd Conference on Learning Theory, Montreal, Quebec, Canada, June 18-21, 2009*, 2009.
- Michael Boardman. The egg drop number. *Mathematics Magazine*, 77(5):368–372, 2004.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Nicolò Cesa-Bianchi, Yoav Freund, David P. Helmbold, and Manfred K. Warmuth. On-line prediction and conversion strategies. *Machine Learning*, 25(1):71–110, 1996.
- C.K. Chow. On optimum error and reject trade-off. *IEEE Trans. on Information Theory*, 1970.
- D. A. Cohn, L. E. Atlas, and R. E. Ladner. Improving generalization with active learning. *Machine Learning*, 15(2), 1994.
- Erik D. Demaine and Morteza Zadimoghaddam. Learning disjunctions: Near-optimal trade-off between mistakes and ”i don’t know’s”. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013*, pages 1369–1379, 2013. doi: 10.1137/1.9781611973105.99.
- R. El-Yaniv and Y. Wiener. On the foundations of noise-free selective classification. *JMLR*, 11, 2010.
- R. El-Yaniv and Y. Wiener. Agnostic selective classification. In *NIPS*, 2011.
- Y. Freund, Y. Mansour, and R. E. Schapire. Generalization bounds for averaged classifiers. *The Ann. of Stat.*, 32, 2004.
- William Gasarch and Stuart Fletcher. The egg game. 2008.
- R. Herbei and M. H. Wegkamp. Classification with reject option. *Canadian J. of Stat.*, 4, 2006.
- Adam Tauman Kalai, Varun Kanade, and Yishay Mansour. Reliable agnostic learning. *J. Comput. Syst. Sci.*, 78(5):1481–1495, 2012. doi: 10.1016/j.jcss.2011.12.026.
- Varun Kanade and Justin Thaler. Distribution-independent reliable learning. In *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, pages 3–24, 2014.

- Lihong Li, Michael L. Littman, Thomas J. Walsh, and Alexander L. Strehl. Knows what it knows: a framework for self-aware learning. *Machine Learning*, 82(3):399–443, 2011. doi: 10.1007/s10994-010-5225-4.
- Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 1987. doi: 10.1007/BF00116827.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Random averages, combinatorial parameters, and learnability. In *Advances in Neural Information Processing Systems*, pages 1984–1992, 2010.
- Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Relax and randomize : From value to algorithms. In *Advances in Neural Information Processing Systems 25*, pages 2150–2158, 2012.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Sequential complexities and uniform martingale laws of large numbers. *Probability Theory and Related Fields*, 161(1-2):111–153, 2015a.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning via sequential complexities. *The Journal of Machine Learning Research*, 16(1):155–186, 2015b.
- Amin Sayedi, Morteza Zadimoghaddam, and Avrim Blum. Trading off mistakes and don’t-know predictions. In *Advances in Neural Information Processing Systems 23*, pages 2092–2100, 2010.
- Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9:371–421, 2008.
- Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012. doi: 10.1561/22000000018.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge University Press, 2014.
- István Szita and Csaba Szepesvári. Agnostic kwik learning and efficient approximate reinforcement learning. In *COLT*, pages 739–772, 2011.
- Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.
- M. Yuan and M. H. Wegkamp. Classification methods with reject option based on convex risk minimization. *JMLR*, 11, 2010.
- C. Zhang and K. Chaudhuri. Beyond disagreement-based agnostic active learning. In *NIPS*, 2014.

Appendix A. Tree Shattering Coefficient and Sequential Growth Function

In this section, we show that the tree shattering coefficient $\mathcal{S}(\mathcal{H}, t)$ is at most the size of the sequential growth function (also known as maximal sequential zero covering number) of \mathcal{H} (Rakhlin et al., 2010). We start with some notations.

Definition 27 (Sequential Zero Cover and Sequential Zero Covering Number, see [Rakhlin et al. \(2010\)](#))

A set V of depth- t trees is a sequential zero cover of \mathcal{H} on a depth- t tree \mathbf{x} , if

$$\forall h \in \mathcal{H}, \forall \epsilon \in \{\pm 1\}^t, \exists \mathbf{v} \in V, \text{ s.t. } \mathbf{v}_s(\epsilon) = h(\mathbf{x}_s(\epsilon)), s = 1, 2, \dots, t$$

The sequential zero covering number of a hypothesis class \mathcal{H} on a given tree \mathbf{x} is defined as

$$\mathcal{N}(0, \mathcal{H}, \mathbf{x}) := \min\{ |V| : V \text{ is a zero-cover of } \mathcal{H} \text{ on } \mathbf{x} \}$$

The maximal sequential zero covering number is the maximum sequential zero covering number of \mathcal{H} over all depth- t \mathcal{X} -valued trees \mathbf{x} , that is,

$$\mathcal{N}(0, \mathcal{H}, t) := \max_{\mathbf{x}} \mathcal{N}(0, \mathcal{H}, \mathbf{x})$$

Theorem 28 For a given hypothesis class \mathcal{H} and integer $t \geq 0$,

$$\mathcal{S}(\mathcal{H}, t) \leq \mathcal{N}(0, \mathcal{H}, t)$$

Proof This is an immediate consequence of [Lemma 29](#). ■

Lemma 29 Suppose we are given a \mathcal{X} -valued tree \mathbf{x} and a hypothesis class \mathcal{H} . If V is a sequential zero cover of \mathcal{H} on \mathbf{x} , then the size of $\mathcal{S}(\mathcal{H}, \mathbf{x})$ is at most $|V|$.

Proof Recall that

$$\mathcal{S}(\mathcal{H}, \mathbf{x}) = \{ (\epsilon_1, \epsilon_2, \dots, \epsilon_t) \in \{\pm 1\}^t : \epsilon_1 = h(\mathbf{x}_1(\epsilon)), \epsilon_2 = h(\mathbf{x}_2(\epsilon)), \dots, \epsilon_t = h(\mathbf{x}_t(\epsilon)), h \in \mathcal{H} \}$$

Given an element $(\epsilon_1, \epsilon_2, \dots, \epsilon_t)$ in $\mathcal{S}(\mathcal{H}, \mathbf{x})$, there exists some h in \mathcal{H} such that

$$\epsilon_1 = h(\mathbf{x}_1(\epsilon)), \epsilon_2 = h(\mathbf{x}_2(\epsilon)), \dots, \epsilon_t = h(\mathbf{x}_t(\epsilon))$$

Since V is a zero-cover of \mathcal{H} , there exists a depth- t tree $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_t)$ in V such that

$$\mathbf{v}_1(\epsilon) = h(\mathbf{x}_1(\epsilon)), \mathbf{v}_2(\epsilon) = h(\mathbf{x}_2(\epsilon)), \dots, \mathbf{v}_t(\epsilon) = h(\mathbf{x}_t(\epsilon))$$

Hence,

$$\mathbf{v}_1(\epsilon) = \epsilon_1, \mathbf{v}_2(\epsilon) = \epsilon_2, \dots, \mathbf{v}_t(\epsilon) = \epsilon_t$$

More explicitly,

$$\mathbf{v}_1 = \epsilon_1, \mathbf{v}_2(\epsilon_1) = \epsilon_2, \mathbf{v}_3(\epsilon_1, \epsilon_2) = \epsilon_3, \dots, \mathbf{v}_t(\epsilon_1, \dots, \epsilon_{t-1}) = \epsilon_t \quad (1)$$

To summarize, for every $(\epsilon_1, \epsilon_2, \dots, \epsilon_t)$ in $\mathcal{S}(\mathcal{H}, \mathbf{x})$, there is a tree \mathbf{v} in V such that Equation (1) holds. Since for each tree \mathbf{v} there can be at most one $(\epsilon_1, \epsilon_2, \dots, \epsilon_t)$ such that Equation (1) holds, this implies that $|\mathcal{S}(\mathcal{H}, \mathbf{x})| \leq |V|$. ■

Appendix B. Reducing the Expert Problem to Online Classification with Finite Class

In this section, we show that the problem of Prediction with Expert Advice (abbrev. Expert Problem) with l -mistake assumption (Cesa-Bianchi et al., 1996; Abernethy et al., 2006) can be cast to the problem studied in this paper, i.e. online classification with a finite hypothesis class with l -bias assumption. Specifically, in the expert problem, at each time t , the algorithm is given experts' advice $(x_{1,t}, \dots, x_{N,t}) \in \{-1, 1\}^N$, and predicts $\hat{y}_t \in \{-1, 1, \perp\}$. Then adversary reveals label $y_t \in \{-1, 1\}$. The l -mistake assumption states that there is an expert i that makes at most l mistakes throughout the process, i.e.

$$\exists i, |\{t : x_{i,t} \neq y_t\}| \leq l$$

For $i = 1, 2, \dots, N$, define hypothesis $h_i : \mathbb{R}^{N+1} \rightarrow \mathbb{R}$ as mapping a $(N+1)$ -dimensional vector to its i th coordinate. Define hypothesis class $\mathcal{H}_N := \{h_i : i = 1, \dots, N\}$. We have the following result relating the l -mistake assumption to l -bias assumption; the intuition is to concatenate a new coordinate at the end of the experts' advice to make all the examples shown distinct.

Proposition 30 *The following are equivalent:*

- (a) *The sequence of expert advice and labels $(x_{1,t}, \dots, x_{N,t}), y_t, t = 1, 2, \dots$ satisfies l -mistake assumption.*
- (b) *The sequence $x_t = (x_{1,t}, \dots, x_{N,t}, t), y_t, t = 1, 2, \dots$ satisfies l -bias assumption with respect to \mathcal{H}_N .*

Proof ((a) \Rightarrow (b))

If $(x_{1,t}, \dots, x_{N,t}), y_t, t = 1, 2, \dots$ satisfies l -mistake assumption, then there is $i \in \{1, \dots, N\}$ such that

$$M_i = |\{t : x_{i,t} \neq y_t\}| \leq l$$

Hence, h_i is correct on all but the rounds t in M_i , i.e. on examples $\{(x_{1,t}, \dots, x_{N,t}, t) : t \in M_i\}$, which are distinct and has size at most l . Therefore, the sequence $\{(x_{1,t}, \dots, x_{N,t}, t), t = 1, 2, \dots\}$ satisfies l -bias assumption with respect to \mathcal{H}_N .

((b) \Rightarrow (a))

If the sequence $(x_{1,t}, \dots, x_{N,t}, t), y_t, t = 1, 2, \dots$ satisfies l -bias assumption with respect to \mathcal{H}_N , then there exists h_i that is correct on all but $p \leq l$ examples shown. That is, p , the size of the set

$$M_i = |\{t : x_{i,t} \neq y_t\}|$$

is at most l . This immediately implies (a). ■

An immediate consequence of the above proposition is that, for an instance of the expert problem with l -mistake assumption, we can convert it to an instance of online classification in \mathcal{H}_N under l -bias assumption, and apply SOA.DK on \mathcal{H}_N to get mistake-abstention tradeoffs.

Appendix C. Proofs from Section 3

We first provide some auxiliary lemmas regarding properties of extended mistake trees and extended Littlestone's dimension. This will serve as the basis of the proof of Lemma 6.

We state a property about subtrees of a (k, m) -difficult extended mistake tree.

Lemma 31 (Recursive Property of Extended Mistake Trees) *Suppose we are given hypothesis class \mathcal{H} that has an extended mistake tree T with root x , left subtree T_{-1} , right subtree T_{+1} and integers $k \geq 0, m \geq 1$. For the root node x , denote by e_l its downward left solid edge, e_r its downward right solid edge, and e_d its downward dashed edge. Denote by y the label of e_d .*

- (i) *The following statements are equivalent: (a) T is $(0, m)$ -difficult. (b) T_y is $(0, m-1)$ -difficult.*
- (ii) *For $k \geq 1$, the following statements are equivalent: (a) T is (k, m) -difficult. (b) T_{-y} is $(k-1, m-1)$ -difficult, and T_y is $(k, m-1)$ -difficult.*

Proof [Proof of Lemma 31]

Without loss of generality, suppose $y = +1$. The case of $y = -1$ can be shown symmetrically.

Proof of item (i):

((a) \Rightarrow (b))

Consider a root to leaf path p in T_{+1} that uses no solid edges. Now consider path p_+ , the result of prepending the root node x and the downward dashed edge from root x to its right child onto p , i.e. $p_+ = xe_dp$. It can be seen that p_+ uses no solid edges, and $l(p_+) = l(p) + 1$. Since T is $(0, m)$ -difficult, $l(p_+) \geq m$, therefore $l(p) \geq m - 1$, thus showing T_{+1} is $(0, m-1)$ -difficult.

((b) \Rightarrow (a))

Consider a root to leaf path $p = v_1e_1v_2e_2 \dots v_n e_n v_{n+1}$ in T that uses no solid edges. The first edge of p must be the downward dashed edge e_d . Define path p_- as the result of deleting the first node $v_1 = x$ and the first edge e_1 from p , i.e. $p_- = v_2e_2 \dots v_n e_n v_{n+1}$. Since T_{-1} is $(0, m-1)$ -difficult, we get that $l(p_-) \geq m - 1$. Therefore $l(p) = l(p_-) + 1 \geq m$. Therefore, any path p in T that uses no solid edges must be of length at least m . Thus, T is $(0, m)$ -difficult.

Proof of item (ii):

((a) \Rightarrow (b))

(1) Consider a root to leaf path p in T_{-1} that uses at most $k-1$ solid edges. Now consider path p_+ , the result of prepending the root node x and the downward edge from root x to its left child onto p , i.e. $p_+ = xe_l p$. It can be seen that p_+ uses at most k solid edges, and $l(p_+) = l(p) + 1$. Since T is (k, m) -difficult, $l(p_+) \geq m$, therefore $l(p) \geq m - 1$, thus showing T_{-1} is $(k-1, m-1)$ -difficult.

(2) Consider a root to leaf path p in T_{+1} that uses at most k solid edges. Now consider path p_+ , the result of prepending the root node x and the downward dashed edge from root x to its right child onto p , i.e. $p_+ = xe_d p$. It can be seen that p_+ uses at most k solid edges, and $l(p_+) = l(p) + 1$. Since T is (k, m) -difficult, $l(p_+) \geq m$, therefore $l(p) \geq m - 1$, thus showing T_{+1} is $(k, m-1)$ -difficult.

((b) \Rightarrow (a))

Consider a root to leaf path $p = v_1e_1v_2e_2 \dots v_n e_n v_{n+1}$ in T that uses at most k solid edges. Define path p_- as the result of deleting the first node $v_1 = x$ the first edge e_1 from p , i.e. $p_- = v_2e_2 \dots v_n e_n v_{n+1}$.

(1) If the first edge of p is a downward edge from root x to its left child, then p_- is a root to leaf path in T_{-1} , and uses at most $k-1$ solid edges, since the first edge e_1 has to be a solid edge. Since T_{+1} is $(k-1, m-1)$ -difficult, we get that $l(p_-) \geq m - 1$.

(2) If the first edge of p is the downward edge from root x to its right child, then p_- is a root to leaf path in T_{+1} , and uses at most k solid edges. Since T_{-1} is $(k, m-1)$ -difficult, we get that $l(p_-) \geq m - 1$.

In both cases, $l(p_-) \geq m - 1$. Hence $l(p) = l(p_-) + 1 \geq m$.

In summary, any path p in T that uses at most k solid edges must be of length m . Thus, T is (k, m) -difficult. ■

Built upon Lemma 31, we obtain the following result regarding \mathcal{H} 's extended Littlestone's dimension.

Lemma 32 (Recursive Property of ELdim) *Suppose we are given hypothesis class \mathcal{H} and integers $k \geq 0, m \geq 1$.*

- (i) *The following statements are equivalent: (a) $\text{ELdim}(\mathcal{H}, 0)$ is at least m . (b) There exists (x, y) such that x is in $\text{DIS}(\mathcal{H})$, and $\text{ELdim}(\mathcal{H}[(x, y)], 0)$ is at least $m - 1$.*
- (ii) *If $k \geq 1$, the following statements are equivalent: (a) $\text{ELdim}(\mathcal{H}, k)$ is at least m . (b) There exists (x, y) such that both $\text{ELdim}(\mathcal{H}[(x, y)], k)$ and $\text{ELdim}(\mathcal{H}[(x, -y)], k - 1)$ are at least $m - 1$.*

Proof [Proof of Lemma 32]

Proof of item (i):

((b) \Rightarrow (a))

Suppose there exists an example (x, y) such that $x \in \text{DIS}(\mathcal{H})$ and $\text{ELdim}(\mathcal{H}[(x, y)], 0) \geq m - 1$. Then, $\mathcal{H}[(x, y)]$ has a $(0, m - 1)$ -difficult extended mistake tree T_y and $\mathcal{H}[(x, -y)]$ has a zeroth order mistake tree T_{-y} . Construct a new tree T , where its root is x , and its subtrees are T_y and T_{-y} respectively. The dashed downward edge is connected to the subtree T_y . By Lemma 31, T is a $(0, m)$ -difficult extended mistake tree with respect to \mathcal{H} . The result follows.

((a) \Rightarrow (b))

Suppose $\text{ELdim}(\mathcal{H}, 0) \geq m$. Then \mathcal{H} has a $(0, m)$ -difficult mistake tree T . Let x be the root of T , and $y \in \{-1, +1\}$ be the label of the root's downward dashed edge. Since T is a full binary tree, there must be leaves in both the left subtree and the right subtree of the root, i.e. there exist h_1, h_2 in \mathcal{H} , $h_1(x) = -1$ and $h_2(x) = -1$. Thus, x is in $\text{DIS}(\mathcal{H})$. By Lemma 31, T_y is a $(0, m - 1)$ -difficult extended mistake tree with respect to $\mathcal{H}[(x, y)]$. The result follows.

Proof of item (ii):

((b) \Rightarrow (a))

Suppose there exists an example (x, y) such that both $\text{ELdim}(\mathcal{H}[(x, y)], k) \geq m - 1$ and $\text{ELdim}(\mathcal{H}[(x, -y)], k - 1) \geq m - 1$. Then, $\mathcal{H}[(x, y)]$ has a $(k, m - 1)$ -difficult extended mistake tree T_y and $\mathcal{H}[(x, -y)]$ has a $(k - 1, m - 1)$ -difficult extended mistake tree T_{-y} . Now construct a new tree T , where its root is x , and its subtrees are T_y and T_{-y} respectively. The dashed downward edge is connected to the subtree T_y . By Lemma 31, T is a (k, m) -difficult extended mistake tree with respect to \mathcal{H} . The result follows.

((a) \Rightarrow (b))

Suppose $\text{ELdim}(\mathcal{H}, k) \geq m$. Then \mathcal{H} has a (k, m) -difficult mistake tree T . Let x be the root of T , and $y \in \{-1, +1\}$ be the root's downward dashed edge label. By Lemma 31, T_{-y} is a $(k - 1, m - 1)$ -difficult extended mistake tree with respect to $\mathcal{H}[(x, -y)]$, and T_y is a $(k, m - 1)$ -difficult extended mistake tree with respect to $\mathcal{H}[(x, y)]$. The result follows. ■

Proof [Proof of Lemma 5] Since $\text{ELdim}(\mathcal{H}, k) \geq m$, there is a (k, m) -difficult extended mistake tree $T_{\mathcal{H}}$ with respect to \mathcal{H} . We consider the strategy of the adversary associated with $T_{\mathcal{H}}$. Now consider any deterministic learning algorithm \mathcal{A} that guarantees at most k mistakes. Since \mathcal{A} is deterministic, the interaction between \mathcal{A} and the adversary follows some path p from root to leaf. The number of mistakes is equal to the number of solid edges in p , and the number of abstentions is equal to the number of dashed edges in the p . Since \mathcal{A} guarantees k mistakes, p must contain at most k solid edges, thus it must be of length at least m , as $T_{\mathcal{H}}$ is (k, m) -difficult. Therefore, the number of nontrivial rounds of \mathcal{A} is at least m . \blacksquare

Proof [Proof of Lemma 6] We prove the lemma by joint induction on (k, m) .

Base Case. Consider pairs (k, m) , where $k = 0$ or $m = 0$.

(1) For $m = 0$ and $k \geq 0$, if there is no $(k, 1)$ -difficult extended mistake tree, then for all $x \in \mathcal{X}$, V predicts unanimously on x . Otherwise, there are two hypotheses h_1 and h_2 and an example x such that $h_1(x) = -1$ and $h_2(x) = +1$. Consider extended mistake tree T as follows. T has x as its root, and h_1 and h_2 are leaves directly connecting to the root, where h_1 is on the left and h_2 is on the right. The downward dashed edge is connected to the right, i.e. has label $+1$. It can be seen that T is $(k, 1)$ -difficult for any $k \geq 0$. Therefore, Algorithm 2 always predicts correctly, and there will be no nontrivial rounds subsequently.

(2) For $k = 0$ and $m \geq 0$, we show the result by induction on m . The base case $m = 0$ has been shown in (1). For the inductive case, assume the inductive hypothesis holds for $m' \leq m - 1$. Now, given a hypothesis class V such that $\text{ELdim}(V, 0)$ is at most m . Consider the first nontrivial round t when running Algorithm 2 with version space V . The example x_t must be in $\text{DIS}(V)$, and the algorithm outputs $\hat{y}_t = \perp$. We claim that the resulting version space $V[(x_t, y_t)]$ is such that $\text{ELdim}(V[(x_t, y_t)], 0) \leq m - 1$. Indeed, suppose $\text{ELdim}(V[(x_t, y_t)], 0) \geq m$, then by Lemma 32, $\text{ELdim}(V, 0) \geq m + 1$, which is a contradiction.

Note that from time $t + 1$ on, the adversary is only allowed to show $V[(x_t, y_t)]$ -realizable sequences. By inductive hypothesis, Algorithm 2 runs on $V[(x_t, y_t)]$ and achieves $(0, m - 1)$ -SZB bound from time $t + 1$ on. Therefore, Algorithm 2 achieves $(0, m)$ -SZB bound throughout the process.

Inductive Case. Consider pairs (k, m) where $k \geq 1$ and $m \geq 1$. Assume for all $k' \leq k$, $m' \leq m$ and $k' + m' \leq k + m - 1$, the inductive hypothesis holds. Now, consider a hypothesis class V such that $\text{ELdim}(V, k) \leq m$. Consider the first nontrivial round t when we run Algorithm 2 on V . The example x_t must be in $\text{DIS}(V)$. According to Algorithm 2's prediction \hat{y}_t , we consider three cases separately,

Case 1: $\hat{y}_t = -1$. In this case, since round t is nontrivial, $y_t = -\hat{y}_t = +1$. We claim that $\text{ELdim}(V[(x_t, +1)], k - 1) \leq m - 1$. Indeed, assume (for the sake of contradiction) that $m_{-1} \geq \text{ELdim}(V[(x_t, +1)], k - 1) \geq m$. By definition of Algorithm 2, $\text{ELdim}(V[(x_t, -1)], k - 1) = m_{+1} \geq m_{-1} \geq m$. Hence, for any $y \in \{-1, +1\}$, $\text{ELdim}(V[(x_t, y)], k - 1) \geq m$.

Also by definition of Algorithm 2, $\max(\text{ELdim}(V[(x_t, +1)], k), \text{ELdim}(V[(x_t, -1)], k)) = m_{\perp} \geq m$. Thus, there exists some $\hat{y} \in \{-1, +1\}$ such that

$$\text{ELdim}(V[(x_t, \hat{y})], k) \geq m$$

Therefore, for \hat{y} , we have $\text{ELdim}(V[(x_t, \hat{y})], k) \geq m$ and $\text{ELdim}(V[(x_t, -\hat{y})], k-1) \geq m$. By Lemma 32, $\text{ELdim}(V, k) \geq m+1$, which is a contradiction.

Note that from time $t+1$ on, the adversary is only allowed to show $V[(x_t, y_t)]$ -realizable sequences. By inductive hypothesis, Algorithm 2 runs on $V[(x_t, y_t)]$ with mistake budget $k-1$ and achieves $(k-1, m-1)$ -SZB bound from round $t+1$ on. Therefore, Algorithm 2 achieves (k, m) -SZB bound throughout the process.

Case 2: $\hat{y}_t = +1$. This case is symmetric to Case 1.

Case 3: $\hat{y}_t = \perp$. We first claim that $\text{ELdim}(V[(x_t, -1)], k) \leq m-1$. Indeed, assume (for the sake of contradiction) that $\text{ELdim}(V[(x_t, -1)], k) \geq m$. By definition of Algorithm 2, $m_{-1} \geq m_{\perp} \geq m$, that is

$$\text{ELdim}(V[(x_t, +1)], k-1) \geq m$$

By Lemma 32, $\text{ELdim}(V, k) \geq m+1$, contradiction. Symmetrically, one also has $\text{ELdim}(V[(x_t, +1)], k) \leq m-1$.

Hence, irrespective of the outcome $y_t \in \{-1, +1\}$, the resulting version space $V[(x_t, y_t)]$ satisfies that $\text{ELdim}(V[(x_t, y_t)], k) \leq m-1$. Note that from time $t+1$ on, the adversary is only allowed to show $V[(x_t, y_t)]$ -realizable sequences. By inductive hypothesis, Algorithm 2 runs on $V[(x_t, y_t)]$ with mistake budget k , and achieves $(k, m-1)$ -SZB bound from round $t+1$ on. Therefore, Algorithm 2 achieves (k, m) -SZB bound throughout the process.

In summary, Algorithm 2, when run on V , achieves (k, m) -SZB bound. This completes the induction. \blacksquare

Proof [Proof of Theorem 7]

(a) This is an immediate consequence of Lemma 6.

(b) By Lemma 5, there is a strategy of the adversary such that any deterministic learner guaranteeing at most k mistakes must have at least m nontrivial rounds. Therefore, no deterministic learner can achieve a $(k, m-1)$ -SZB bound. \blacksquare

Proof [Proof of Theorem 8]

Recall that $\text{Ldim}(\mathcal{H}) = d < \infty$.

We first show $\text{ELdim}(\mathcal{H}, d) \leq d$. Indeed, SOA is guaranteed to make at most d mistakes and no abstentions for \mathcal{H} -realizable sequences. This has a total of at most d nontrivial rounds. Now, by Lemma 5, if $\text{ELdim}(\mathcal{H}, d) \geq d+1$, SOA must have at least $d+1$ nontrivial rounds, contradiction.

On the other hand, since $\text{Ldim}(\mathcal{H}) = d$, there is a depth- d mistake tree T with respect to \mathcal{H} . Consider the following modification of T : for each internal node, add a dashed downward edge to its right child. It can be seen that the resulting tree, \tilde{T} , is a (d, d) -difficult extended mistake tree. Therefore $\text{ELdim}(\mathcal{H}, d) \geq d$.

In summary, $\text{ELdim}(\mathcal{H}, d) = d$. \blacksquare

Appendix D. Proofs from Section 4

Proof [Proof of Lemma 13] For any depth- t tree \mathbf{x} , note that

$$|S(\mathcal{H}, \mathbf{x})| \leq |\mathcal{H}|$$

Therefore,

$$S(\mathcal{H}, t) = \max_{\mathbf{x}} |S(\mathcal{H}, \mathbf{x})| \leq |\mathcal{H}|$$

■

Lemma 33 (Recursive Formula) For a hypothesis class \mathcal{H} and $t \geq 1$, we have

$$S(\mathcal{H}, t) = \max_{x \in \mathcal{X}} (\mathcal{S}(\mathcal{H}[(x, -1)], t-1) + \mathcal{S}(\mathcal{H}[(x, +1)], t-1))$$

We need the following notation of subtrees to give the proof of Lemma 33.

Definition 34 (Subtrees, see Rakhlin et al. (2010)) Given a depth- t tree \mathbf{x} , the left subtree \mathbf{x}^l of \mathbf{x} at the root is defined as $t-1$ mappings $(\mathbf{x}_1^l, \dots, \mathbf{x}_{t-1}^l)$, where $\mathbf{x}_i^l(\epsilon) = \mathbf{x}(\{-1\} \times \epsilon)$, for $\epsilon \in \{\pm 1\}^{t-1}$. The right subtree \mathbf{x}^r of \mathbf{x} at the root is defined as $t-1$ mappings $(\mathbf{x}_1^r, \dots, \mathbf{x}_{t-1}^r)$, where $\mathbf{x}_i^r(\epsilon) = \mathbf{x}(\{+1\} \times \epsilon)$, for $\epsilon \in \{\pm 1\}^{t-1}$.

Proof [Proof of Lemma 33] Consider the definition of $S(\mathcal{H}, t)$:

$$\max_{\mathbf{x}} |\{(\epsilon_1, \epsilon_2, \dots, \epsilon_t) \in \{\pm 1\}^t : \epsilon_1 = h(\mathbf{x}_1(\epsilon)), \epsilon_2 = h(\mathbf{x}_2(\epsilon)), \dots, \epsilon_t = h(\mathbf{x}_t(\epsilon)), h \in \mathcal{H}\}|$$

This can be alternatively written as

$$\begin{aligned} \max_{\mathbf{x}} |\{(-1, \sigma_1, \dots, \sigma_{t-1}) \in \{\pm 1\}^{t-1} : \sigma_1 = h(\mathbf{x}_1^l(\sigma)), \dots, \sigma_{t-1} = h(\mathbf{x}_{t-1}^l(\sigma)), h \in \mathcal{H}[(\mathbf{x}_1, -1)]\} \\ \cup \{(+1, \sigma_1, \dots, \sigma_{t-1}) \in \{\pm 1\}^{t-1} : \sigma_1 = h(\mathbf{x}_1^r(\sigma)), \dots, \sigma_{t-1} = h(\mathbf{x}_{t-1}^r(\sigma)), h \in \mathcal{H}[(\mathbf{x}_1, +1)]\}| \end{aligned}$$

The above is equal to

$$\begin{aligned} \max_{\mathbf{x}_1 \in \mathcal{X}} \{ \max_{\mathbf{x}^l} |\{(-1, \sigma_1, \dots, \sigma_{t-1}) \in \{\pm 1\}^{t-1} : \sigma_1 = h(\mathbf{x}_1^l(\sigma)), \dots, \sigma_{t-1} = h(\mathbf{x}_{t-1}^l(\sigma)), h \in \mathcal{H}[(\mathbf{x}_1, -1)]\}| \\ + \max_{\mathbf{x}^r} |\{(+1, \sigma_1, \dots, \sigma_{t-1}) \in \{\pm 1\}^{t-1} : \sigma_1 = h(\mathbf{x}_1^r(\sigma)), \dots, \sigma_{t-1} = h(\mathbf{x}_{t-1}^r(\sigma)), h \in \mathcal{H}[(\mathbf{x}_1, +1)]\}| \} \end{aligned}$$

Note that the right hand side is precisely $\max_{x_1 \in \mathcal{X}} (\mathcal{S}(\mathcal{H}[(x_1, -1)], t-1) + \mathcal{S}(\mathcal{H}[(x_1, +1)], t-1))$. The lemma follows. ■

Now we are ready to prove Lemma 14.

Proof [Proof of Lemma 14] We prove the result by joint induction on (t, d) .

Base Case: Consider $t = 0$ or $d = 0$. If $t = 0$, then $S(\mathcal{H}, 0) \leq 1 = \binom{0}{\leq d}$. If $d = 0$, then $S(\mathcal{H}, t) \leq 1 = \binom{t}{\leq 0}$.

Inductive Case: For $t \geq 1$ and $d \geq 1$, assume the result holds for (t', d') such that $t' \leq t$, $d' \leq d$ and $t' + d' \leq t + d - 1$. First by Lemma 33, for some x in \mathcal{X} , $\mathcal{S}(\mathcal{H}, t) \leq \mathcal{S}(\mathcal{H}[(x, -1)], t - 1) + \mathcal{S}(\mathcal{H}[(x, +1)], t - 1)$.

Second, Since $\text{Ldim}(\mathcal{H}) = d$, for x , there exists $y \in \{-1, +1\}$ such that $\text{Ldim}(\mathcal{H}[(x, y)]) \leq d - 1$ and $\text{Ldim}(\mathcal{H}[(x, -y)]) \leq d$. Hence by inductive hypothesis, there exists $y \in \{-1, +1\}$ such that $\mathcal{S}(\mathcal{H}[(x, y)], t - 1) \leq \binom{t-1}{\leq d-1}$ and $\mathcal{S}(\mathcal{H}[(x, -y)], t - 1) \leq \binom{t-1}{\leq d}$. Therefore

$$\mathcal{S}(\mathcal{H}, t) \leq \mathcal{S}(\mathcal{H}[(x, -1)], t - 1) + \mathcal{S}(\mathcal{H}[(x, +1)], t - 1) \leq \binom{t-1}{\leq d-1} + \binom{t-1}{\leq d} \leq \binom{t}{\leq d}$$

This completes the induction. ■

Proof [Proof of Theorem 15] For any integer m , if $m \leq \text{ELdim}(\mathcal{H}, k)$, then by Lemma 35,

$$\mathcal{S}(\mathcal{H}, m) \geq \binom{m}{\leq k+1}$$

This implies that

$$m \leq \sup \left\{ t : \binom{t}{\leq k+1} \leq \mathcal{S}(\mathcal{H}, t) \right\}$$

Taking $m = \text{ELdim}(\mathcal{H}, k)$, we get the theorem. ■

Lemma 35 Suppose k, t are nonnegative integers. If $\text{ELdim}(\mathcal{H}, k) \geq t$, then $\mathcal{S}(\mathcal{H}, t) \geq \binom{t}{\leq k+1}$.

Proof [Proof of Lemma 35] By joint induction on (k, t) .

Base Case: We consider (k, t) pairs where $k = 0$ or $t = 0$.

(1) For $t = 0$, $\text{ELdim}(\mathcal{H}, k) \geq 0$ implies that \mathcal{H} is nonempty. Thus, $\mathcal{S}(\mathcal{H}, 0) = 1 \geq \binom{0}{\leq k+1}$.

(2) For $k = 0$, we prove the result by induction on t . The case of $t = 0$ has been shown in (1). For the inductive case, by Lemma 32, there exists (x, y) such that $x \in \text{DIS}(\mathcal{H})$ and $\text{ELdim}(\mathcal{H}[(x, y)], 0) \geq t - 1$. Thus, by inductive hypothesis, $\mathcal{S}(\mathcal{H}[(x, y)], t - 1) \geq t$. Also, since $\mathcal{H}[(x, -y)]$ is nonempty, we get $\mathcal{S}(\mathcal{H}[(x, -y)], t - 1) \geq 1$. Thus,

$$\mathcal{S}(\mathcal{H}, t) \geq \mathcal{S}(\mathcal{H}[(x, y)], t - 1) + \mathcal{S}(\mathcal{H}[(x, -y)], t - 1) \geq t + 1$$

This completes the proof for $k = 0$.

Inductive Case: For $t \geq 1$ and $k \geq 1$, suppose the inductive hypothesis holds for any (k', t') such that $k' \leq k$, $t' \leq t$, $k' + t' \leq k + t + 1$.

Now suppose $\text{ELdim}(\mathcal{H}, k) \geq t$. By Lemma 32, there exists (x, y) such that $\text{ELdim}(\mathcal{H}[(x, y)], k) \geq t - 1$ and $\text{ELdim}(\mathcal{H}[(x, -y)], k - 1) \geq t - 1$. Thus by inductive hypothesis, $\mathcal{S}(\mathcal{H}[(x, y)], t - 1) \geq \binom{t-1}{\leq k+1}$ and $\mathcal{S}(\mathcal{H}[(x, -y)], t - 1) \geq \binom{t-1}{\leq k}$. Therefore,

$$\mathcal{S}(\mathcal{H}, t) \geq \mathcal{S}(\mathcal{H}[(x, y)], t - 1) + \mathcal{S}(\mathcal{H}[(x, -y)], t - 1) \geq \binom{t}{\leq k+1}$$

This completes the induction. ■




Figure 5: Construction of $T_{0,m}$, an extended mistake tree given parameters $k = 0$ and $m \geq 0$. For each i , h_i is defined as $h_i(x) := 2I(x \leq t_i) - 1$.

Proof [Proof of Theorem 17] Note that $\mathcal{S}(\mathcal{H}, t) \leq |\mathcal{H}| = n$, therefore by Lemma 35, $\text{ELdim}(\mathcal{H}, k) \leq \max\{t : \binom{t}{\leq k+1} \leq n\}$.

On the other hand, Lemma 36 implies that for all m such that $\binom{m}{\leq k+1} \leq n$, there is a (m, k) -difficult extended mistake tree with respect to \mathcal{H} . Hence $\text{ELdim}(\mathcal{H}, k) \geq \max\{t : \binom{t}{\leq k+1} \leq n\}$. Combining the lower and upper bound, we get the theorem. \blacksquare

Lemma 36 Consider the set of threshold classifiers $\mathcal{H} = \{2I(x \leq t) - 1 : t \in \{t_1, \dots, t_n\}\}$. If integers $k \geq 0$ and $m \geq 0$ are such that $\binom{m}{k+1} \leq n$, then \mathcal{H} has a (k, m) -difficult mistake tree.

Proof We prove the lemma by joint induction on (k, m) .

Base Case. Consider $k = 0$ or $m = 0$.

(1) For $k = 0$, $\binom{m}{\leq k} = m + 1$. We show a construction of $T_{0,m}$, a $(0, m)$ -difficult extended mistake tree in Figure 5. It can be seen that the resulting tree $T_{0,m}$ is $(0, m)$ -difficult, as the only root to leaf path using no solid edges corresponds to examples $(t_2, +1), \dots, (t_{m+1}, +1)$, which has length m .

(2) For $m = 0$ and integer k , $\binom{m}{\leq k} = 1$. The zeroth order extended mistake tree containing h_{t_1} is a $(k, 0)$ -difficult extended mistake tree.

Inductive Case. For $k \geq 1$ and $m \geq 1$, assume the inductive hypothesis holds for (k', m') such that $k' \leq k$, $m' \leq m$ and $k' + m' \leq k + m - 1$.

We now construct $T_{k,m}$, a (k, m) -difficult extended mistake tree, using hypotheses in \mathcal{H} . Let $r_- = \binom{m-1}{\leq k}$, $r_+ = \binom{m-1}{\leq k+1}$. Consider hypothesis class $\mathcal{H}_- = \{2I(x \leq t) - 1 : t \in \{t_1, \dots, t_{r_-}\}\}$ and $\mathcal{H}_+ = \{2I(x \leq t) - 1 : t \in \{t_{r_-+1}, \dots, t_{r_-+r_+}\}\}$. Note that $r_- + r_+ \leq \binom{m-1}{\leq k} + \binom{m-1}{\leq k-1} \leq \binom{m}{\leq k} \leq n$, thus \mathcal{H}_- and \mathcal{H}_+ are well defined.

Since $|\mathcal{H}_-| \geq \binom{m-1}{\leq k}$, by inductive hypothesis, there is a $(k-1, m-1)$ difficult extended mistake tree $T_{k-1, m-1}$ with respect to \mathcal{H}_- . Similarly, since $|\mathcal{H}_+| \geq \binom{m-1}{\leq k+1}$, by inductive hypothesis, there is a $(k, m-1)$ difficult extended mistake tree $T_{k, m-1}$ with respect to \mathcal{H}_+ .




Figure 6: Construction of $T_{k,m}$, an extended mistake tree given parameters $k \geq 1$ and $m \geq 1$, from $T_{k,m-1}$ and $T_{k-1,m-1}$.

Now Let x be a real number in (t_{r-}, t_{r-+1}) , it can be seen that all hypotheses in \mathcal{H}_- classifies x as -1 and all hypothesis in \mathcal{H}_+ classifies x as $+1$. We construct $T_{k,m}$ as in Figure 6, where x is at the root, and its downward left solid edge connect to $T_{k-1,m-1}$; its downward right solid edge and downward dashed edge connects to $T_{k,m-1}$. Note that $T_{k,m}$ is a valid extended mistake tree, since all hypotheses at the leaves in $T_{k-1,m-1}$ (resp. $T_{k-1,m}$) classifies x as -1 (resp. $+1$). By Lemma 31, T is (k, m) -difficult. ■

Lemma 37 Let \mathcal{C}^l be the class of unions of at most l singletons. Then

$$\mathcal{S}(\mathcal{C}^l, t) = \binom{t}{\leq l}$$

Proof (1) $\mathcal{S}(\mathcal{C}^l, t) \leq \binom{t}{\leq l}$ From Lemma 14.

(2) Consider a \mathcal{X} -valued tree \mathbf{x} with all its elements distinct. Then, consider the set

$$\mathcal{S}(\mathcal{C}^l, \mathbf{x}) = \{(\epsilon_1, \dots, \epsilon_t) : \exists h \in \mathcal{C}^l, h(\mathbf{x}_s(\epsilon)) = \epsilon_s, s = 1, 2, \dots, t\}$$

We claim that $\mathcal{S}(\mathcal{C}^l, \mathbf{x})$ contains $\{\epsilon = (\epsilon_1, \dots, \epsilon_t) : |\{s : \epsilon_s = -1\}| \leq l\}$. Indeed, for any element in $\{\epsilon = (\epsilon_1, \dots, \epsilon_t) : |\{s : \epsilon_s = -1\}| \leq l\}$, the hypothesis $h = 1 - 2I(x \in \{\mathbf{x}_s(\epsilon) : \epsilon_s = -1\}) \in \mathcal{C}^l$ satisfies that $h(\mathbf{x}_s(\epsilon)) = \epsilon_s$, for $s = 1, 2, \dots, t$. Hence $\mathcal{S}(\mathcal{C}^l, \mathbf{x}) \geq \binom{t}{\leq l}$, implying $\mathcal{S}(\mathcal{C}^l, t) \geq \binom{t}{\leq l}$.

In summary, $\mathcal{S}(\mathcal{C}^l, t) = \binom{t}{\leq l}$. ■

Proof [Proof of Theorem 18]

(1) Consider the case that $k \leq l - 1$. By Lemma 38, for any integer m , there is a (k, m) -difficult extended mistake tree with respect to \mathcal{C}^l . Thus, $\text{ELdim}(\mathcal{C}^l, k) = \infty$.

(2) Consider the case that $k \geq l$. By Lemma 37, $\mathcal{S}(\mathcal{C}^l, t) = \binom{t}{\leq l}$. By Theorem 15,

$$\text{ELdim}(\mathcal{C}^l, k) \leq \max \left\{ t : \binom{t}{\leq k+1} \leq \binom{t}{\leq l} \right\} = l$$

This gives that $\text{ELdim}(\mathcal{C}^l, k) \leq l$.

On the other hand, \mathcal{C}^l has a mistake tree T of depth l . Consider the following modification of T : for each internal node, add a dashed downward edge to its right child. It can be seen that the resulting tree, \tilde{T} , is an (l, l) -difficult extended mistake tree. Therefore T' is also a (k, l) -difficult mistake tree, which gives that $\text{ELdim}(\mathcal{C}^l, k) \geq l$.

Hence, we conclude that $\text{ELdim}(\mathcal{C}^l, k) = l$. ■

Recall that \mathcal{C}^l is the class of union of at most l singletons in instance domain \mathcal{X} . That is, hypotheses that take value $+1$ on \mathcal{X} , except for at most l points.

Lemma 38 *Suppose we are given an infinite domain \mathcal{X} and an integer $l \geq 1$. Then for any integer $m \geq 0$, there exists a $(l - 1, m)$ -difficult extended mistake tree with respect to hypothesis class \mathcal{C}^l , such that all its dashed edges are labeled $+1$.*

Proof By induction on l .

Base Case. For $l = 1$, the construction of the required extended mistake tree with respect to \mathcal{C}^1 is given in Figure 7. Note that the tree is $(0, m)$ -difficult, and all its dashed edges are labeled $+1$.

Inductive Case. Suppose the inductive hypothesis holds for any $l' \leq l - 1$. Now pick an arbitrary $x \in \mathcal{X}$. Fix integer m . Consider $(\mathcal{X}_1, \mathcal{X}_2)$, a partition of $\mathcal{X} \setminus \{x\}$, where both $|\mathcal{X}_1|$ and $|\mathcal{X}_2|$ are infinite.

By inductive hypothesis, there is a $(l - 1, m)$ -difficult extended mistake tree T_+ with respect to \mathcal{C}^l on domain \mathcal{X}_1 , such that all its dashed edges are labeled $+1$. Since for any $h \in \mathcal{C}^{l-1}$, there exists $h' \in \mathcal{C}^l[(x, +1)]$ such that $h \equiv h'$ on \mathcal{X}_1 , we can modify T_+ 's leaves such that they all correspond to hypotheses in $\mathcal{C}^l[(x, +1)]$, getting a new extended mistake tree \tilde{T}_+ .

Similarly, by inductive hypothesis, there is a $(l - 2, m)$ -difficult extended mistake tree T_- with respect to \mathcal{C}^{l-1} on domain \mathcal{X}_2 , such that all its dashed edges are labeled $+1$. Since for any $h \in \mathcal{C}^l$, there exists $h' \in \mathcal{C}^l[(x, -1)]$ such that $h \equiv h'$ on \mathcal{X}_2 , we can modify T_- 's leaves such that they all correspond to hypotheses in $\mathcal{C}^l[(x, -1)]$, getting a new extended mistake tree \tilde{T}_- .

Now consider the extended mistake tree T rooted at x , with its left subtree as \tilde{T}_- and right subtree as \tilde{T}_+ . The dashed downward edge of root is linked to its right child, i.e. has label $+1$. Note that T is a valid extended mistake tree, since all hypotheses at the leaves in \tilde{T}_- (resp. \tilde{T}_+) classifies x as -1 (resp. $+1$). By Lemma 31, T is $(l - 1, m + 1)$ -difficult, hence $(l - 1, m)$ -difficult. Additionally, all its dashed edges are labeled $+1$. Since the choice of m is arbitrary, this completes the induction. ■

Appendix E. Proofs from Section 5

Proof [Proof of Theorem 19] Let \mathcal{X} be an infinite set. Let \mathcal{H} be the hypothesis class containing only one hypothesis $h \equiv +1$. Note that $\mathcal{H}^l = \mathcal{C}^l$ and by Theorem 18, $\text{ELdim}(\mathcal{H}^l, k) = \infty$ for $k < l$. By Lemma 5, the theorem follows. ■




Figure 7: A $(0, m)$ -difficult extended mistake tree with respect to \mathcal{C}^1 . x_1, \dots, x_m are distinct elements in \mathcal{X} , and for each i , h_i is defined as $h_i(x) := 1 - 2I(x = x_i)$. h_+ is the constant function $+1$.

Proof [Proof of Theorem 20] Let \mathcal{X} be an infinite set. Let \mathcal{H} be the \mathcal{C}^d , the class of unions of at most d singletons. Note that $\mathcal{H}^l = \mathcal{C}^{l+d}$ and by Theorem 18, $\text{ELdim}(\mathcal{C}^{l+d}, k) = \infty$ for $k < l + d$. By Lemma 5, the theorem follows. \blacksquare

We will need the following result regarding the tree shattering coefficient of the “product” of two hypothesis classes.

Lemma 39 *Suppose $\mathcal{H}_1, \mathcal{H}_2$ are two hypothesis classes. If $\mathcal{H} = \mathcal{H}_1 \cdot \mathcal{H}_2$, then for all integers $t \geq 0$,*

$$\mathcal{S}(\mathcal{H}, t) \leq \mathcal{S}(\mathcal{H}_1, t) \cdot \mathcal{S}(\mathcal{H}_2, t)$$

We first show a basic property of tree shattering coefficient.

Lemma 40 *For hypothesis classes $\mathcal{H}_1, \mathcal{H}_2$, $\mathcal{S}(\mathcal{H}_1 \cup \mathcal{H}_2, t) \leq \mathcal{S}(\mathcal{H}_1, t) + \mathcal{S}(\mathcal{H}_2, t)$.*

Proof For any depth- t tree \mathbf{x} , we have that

$$\mathcal{S}(\mathcal{H}_1 \cup \mathcal{H}_2, \mathbf{x}) \subseteq \mathcal{S}(\mathcal{H}_1, \mathbf{x}) \cup \mathcal{S}(\mathcal{H}_2, \mathbf{x})$$

Therefore,

$$|\mathcal{S}(\mathcal{H}_1 \cup \mathcal{H}_2, \mathbf{x})| \leq |\mathcal{S}(\mathcal{H}_1, \mathbf{x})| + |\mathcal{S}(\mathcal{H}_2, \mathbf{x})| \leq \mathcal{S}(\mathcal{H}_1, t) + \mathcal{S}(\mathcal{H}_2, t)$$

Since the choice of \mathbf{x} is arbitrary, we get

$$\mathcal{S}(\mathcal{H}_1 \cup \mathcal{H}_2, t) \leq \mathcal{S}(\mathcal{H}_1, t) + \mathcal{S}(\mathcal{H}_2, t)$$

\blacksquare

Proof [Proof of Lemma 39] By induction on t .

Base Case: Consider $t = 0$. If one of $\mathcal{S}(\mathcal{H}_1, 0)$, $\mathcal{S}(\mathcal{H}_2, 0)$ is 0, this implies $\mathcal{H}_1 = \emptyset$ or $\mathcal{H}_2 = \emptyset$. Therefore, $\mathcal{H} = \emptyset$, the result holds. Otherwise, both $\mathcal{S}(\mathcal{H}_1, 0)$ and $\mathcal{S}(\mathcal{H}_2, 0)$ are at least 1. In this case \mathcal{H} is nonempty, thus $1 = \mathcal{S}(\mathcal{H}, 0) \leq 2^0 = 1$, the result also hold.

Inductive Case: Given $t \geq 1$, assume the inductive hypothesis $\mathcal{S}(\mathcal{F}_1 \cdot \mathcal{F}_2, t - 1) \leq \mathcal{S}(\mathcal{F}_1, t - 1) \cdot \mathcal{S}(\mathcal{F}_2, t - 1)$ holds for any hypothesis classes $\mathcal{F}_1, \mathcal{F}_2$. Fix $x \in \mathcal{X}$. Note that $\mathcal{H}[(x, +1)] = (\mathcal{H}_1[(x, +1)] \cdot \mathcal{H}_2[(x, +1)]) \cup (\mathcal{H}_1[(x, -1)] \cdot \mathcal{H}_2[(x, -1)])$. Therefore,

$$\begin{aligned} & \mathcal{S}(\mathcal{H}[(x, +1)], t - 1) \\ & \leq \mathcal{S}(\mathcal{H}_1[(x, +1)] \cdot \mathcal{H}_2[(x, +1)], t - 1) + \mathcal{S}(\mathcal{H}_1[(x, -1)] \cdot \mathcal{H}_2[(x, -1)], t - 1) \\ & \leq \mathcal{S}(\mathcal{H}_1[(x, +1)], t - 1)\mathcal{S}(\mathcal{H}_2[(x, +1)], t - 1) + \mathcal{S}(\mathcal{H}_1[(x, -1)], t - 1)\mathcal{S}(\mathcal{H}_2[(x, -1)], t - 1) \end{aligned}$$

where the first inequality is Lemma 40, the second inequality is by inductive hypothesis.

Likewise, we have

$$\mathcal{S}(\mathcal{H}[(x, -1)], t - 1) \leq \mathcal{S}(\mathcal{H}_1[(x, -1)], t - 1)\mathcal{S}(\mathcal{H}_2[(x, +1)], t - 1) + \mathcal{S}(\mathcal{H}_1[(x, +1)], t - 1)\mathcal{S}(\mathcal{H}_2[(x, -1)], t - 1)$$

Therefore,

$$\begin{aligned} & \mathcal{S}(\mathcal{H}[(x, -1)], t - 1) + \mathcal{S}(\mathcal{H}[(x, +1)], t - 1) \\ & \leq (\mathcal{S}(\mathcal{H}_1[(x, -1)], t - 1) + \mathcal{S}(\mathcal{H}_1[(x, +1)], t - 1))(\mathcal{S}(\mathcal{H}_2[(x, -1)], t - 1) + \mathcal{S}(\mathcal{H}_2[(x, +1)], t - 1)) \\ & \leq \mathcal{S}(\mathcal{H}_1, t)\mathcal{S}(\mathcal{H}_2, t) \end{aligned}$$

where the second inequality is from Lemma 32. Since the choice of x is arbitrary, we get

$$\mathcal{S}(\mathcal{H}, t) = \max_{x \in \mathcal{X}} (\mathcal{S}(\mathcal{H}[(x, -1)], t - 1) + \mathcal{S}(\mathcal{H}[(x, +1)], t - 1)) \leq \mathcal{S}(\mathcal{H}_1, t)\mathcal{S}(\mathcal{H}_2, t)$$

■

Proof [Proof of Lemma 21] Note that by Lemmas 14 and 39,

$$\mathcal{S}(\mathcal{H}^l, t) \leq \mathcal{S}(C^l, t) \cdot \mathcal{S}(\mathcal{H}, t) = \binom{t}{\leq l} \cdot \mathcal{S}(\mathcal{H}, t) \leq |\mathcal{H}| \binom{t}{\leq l}$$

Hence,

$$\text{ELdim}(\mathcal{H}^l, k) \leq \max \left\{ t : \binom{t}{\leq k+1} \leq |\mathcal{H}| \binom{t}{\leq l} \right\}$$

Now, consider any t such that $t \geq 2l$ and

$$\binom{t}{\leq k+1} \leq |\mathcal{H}| \binom{t}{\leq l}$$

Since $\binom{t}{\leq k+1} \geq \binom{t}{k+1} \geq \left(\frac{t}{k+1}\right)^{k+1}$, and $\binom{t}{\leq l} \leq \left(\frac{et}{l}\right)^l$ for $t \geq 2l$, we get

$$\left(\frac{t}{k+1}\right)^{k+1} \leq |\mathcal{H}| \left(\frac{et}{l}\right)^l$$

Hence,

$$t^{k+1-l} \leq |\mathcal{H}| \frac{(k+1)^{k+1}}{l^l}$$

Since $\frac{(k+1)^{k+1}}{l^l} \leq (e(k+1))^{k+1-l}$, we get

$$t^{k+1-l} \leq |\mathcal{H}| (e(k+1))^{k+1-l}$$

That is, $t \leq e(k+1) |\mathcal{H}|^{\frac{1}{k+1-l}}$.

In summary,

$$\text{ELdim}(\mathcal{H}^l, k) \leq \max(2l, e(k+1) |\mathcal{H}|^{\frac{1}{k+1-l}}) = e(k+1) |\mathcal{H}|^{\frac{1}{k+1-l}}$$

where the equality uses the fact that $k \geq l$. ■

Proof [Proof of Lemma 23] Note that by Lemmas 14 and 39,

$$\mathcal{S}(\mathcal{H}^l, t) \leq \mathcal{S}(\mathcal{C}^l, t) \cdot \mathcal{S}(\mathcal{H}, t) = \binom{t}{\leq l} \cdot \mathcal{S}(\mathcal{H}, t) \leq \binom{t}{\leq d} \binom{t}{\leq l}$$

Hence,

$$\text{ELdim}(\mathcal{H}^l, k) \leq \max \left\{ t : \binom{t}{\leq k+1} \leq \binom{t}{\leq d} \binom{t}{\leq l} \right\}$$

Now, consider any t such that $t \geq 2l$, $t \geq 2d$ and

$$\binom{t}{\leq k+1} \leq \binom{t}{\leq d} \binom{t}{\leq l}$$

Since $\binom{t}{\leq k+1} \geq \binom{t}{k+1} \geq \left(\frac{t}{k+1}\right)^{k+1}$, $\binom{t}{\leq d} \leq \left(\frac{et}{d}\right)^d$ for $t \geq 2d$, and $\binom{t}{\leq l} \leq \left(\frac{et}{l}\right)^l$ for $t \geq 2l$, we get

$$\left(\frac{t}{k+1}\right)^{k+1} \leq \left(\frac{et}{d}\right)^d \left(\frac{et}{l}\right)^l$$

Hence,

$$t^{k+1-l-d} \leq e^{l+d} \frac{(k+1)^{k+1}}{l^l d^d}$$

Since

$$\begin{aligned} & \frac{(k+1)^{k+1}}{l^l d^d} \\ &= \left(1 + \frac{k+1-l}{l}\right)^l \left(1 + \frac{k+1-d}{d}\right)^d (k+1)^{k+1-l-d} \\ &\leq e^{2k+2-l-d} (k+1)^{k+1-l-d} \end{aligned}$$

we get

$$t \leq (k+1) \cdot e^{\frac{2k+2}{k+1-l-d}}$$

In summary,

$$\text{ELdim}(\mathcal{H}^l, k) \leq \max(2l, 2d, (k+1) \cdot e^{\frac{2k+2}{k+1-l-d}}) = (k+1) \cdot e^{\frac{2k+2}{k+1-l-d}}$$

where the equality uses the fact that $k \geq l + d$. ■

Proof [Proof of Theorem 25] Let \mathcal{X} be an infinite set. x_1, x_2, \dots is a sequence of distinct elements from \mathcal{X} . Let \mathcal{H} be the hypothesis class containing only one hypothesis $h \equiv +1$. Note that $\mathcal{H}^l = \mathcal{C}^l$. Let $\epsilon = 1 - k/l > 0$, thus $k = l(1 - \epsilon)$. Fix integer $m = \lceil \frac{2}{\epsilon}(a + l) + 2l \rceil$. By Lemma 38, \mathcal{C}^l has a $(l - 1, m)$ -difficult extended mistake tree T .

We define the following strategy by the adversary based on T . At time $t = 1$, the adversary chooses the example x_1 corresponding to the root of T , and shows it to the learner. If $p_{1,+} > 1 - \epsilon$, then it reveals label $y_1 = -1$ and follows the downward solid edge labeled -1 to reach the left child of the root; otherwise it reveals label $y_1 = +1$ and follows the downward dashed edge labeled $+1$ to reach the right child of the root. At time $t \geq 2$, suppose the adversary reaches a node with example x_t , then x_t is shown to the learner, and one of the downward edges adjacent to this node is followed by the same rule. The interaction comes to an end when a leaf is reached. It can be seen that the realizability assumption is maintained.

Consider an Algorithm \mathcal{A} that guarantees a cumulative mistake penalty at most k .

(1) We claim that the interaction between the learner and the adversary lasts for at least m rounds. To see this, note that \mathcal{A} predicts at most $l - 1$ times such that $p_{t,+} > 1 - \epsilon$. Assume this is not the case, that is,

$$|\{t \in [m] : p_{t,+} > 1 - \epsilon\}| \geq l$$

Suppose the first l times \mathcal{A} predicts $p_{t,+} > 1 - \epsilon$ are $1 \leq t_1 < \dots < t_l \leq m$. Then, according to the adversary's strategy, $y_{t_1} = \dots = y_{t_l} = -1$. Thus, the cumulative mistake penalty made by \mathcal{A} up to time t_l is at least

$$\sum_{i=1}^l p_{t_i,+} > l(1 - \epsilon) = k$$

This implies that \mathcal{A} has a cumulative mistake penalty $> k$, contradiction. Therefore throughout the interaction, the number of solid edges used is at most $l - 1$. Since T is $(l - 1, m)$ -difficult, any path that going downward from the root using $l - 1$ solid edges must be of length at least m , hence the interaction between the learner and the adversary lasts for at least m rounds.

(2) We claim that over the first m rounds, there are at most $\frac{2l}{\epsilon}$ rounds such that \mathcal{A} predicts $p_{t,-} > \epsilon/2$ and $p_{t,+} \leq 1 - \epsilon$. Assume this is not the case, that is,

$$|\{t \in [m] : p_{t,-} > \epsilon/2 \wedge p_{t,+} \leq 1 - \epsilon\}| \geq \frac{2l}{\epsilon}$$

Suppose the first $g = \lceil \frac{2l}{\epsilon} \rceil$ times \mathcal{A} predicts -1 are $1 \leq s_1 < \dots < s_g \leq m$. Then, according to the adversary's strategy, $y_{s_1} = \dots = y_{s_g} = +1$. Thus the cumulative mistake penalty made by \mathcal{A} up to time s_g is at least

$$\sum_{i=1}^g p_{s_i,-} > g \cdot \frac{\epsilon}{2} \geq l(1 - \epsilon) = k$$

This implies that \mathcal{A} has a cumulative mistake penalty $> k$ over time, contradiction.

Therefore, among the first m rounds, there are at most $l + (l + \frac{2l}{\epsilon}) = 2l + \frac{2l}{\epsilon}$ rounds such that $p_{t,+} > 1 - \epsilon$ or $p_{t,-} > \epsilon/2$. Thus there are at least $(m - \frac{2l}{\epsilon} - 2l)$ rounds such that $p_{t,+} \leq 1 - \epsilon$ and

$p_{t,-} \leq \epsilon/2$, implying $1 - p_{t,+} - p_{t,-} \geq \epsilon/2$. Thus, the cumulative abstention penalty up to time m is at least

$$\left(m - \frac{2l}{\epsilon} - 2l\right) \cdot \frac{\epsilon}{2} \geq a$$

■

Proof [Proof of Theorem 26] Let \mathcal{X} be an infinite set. Let \mathcal{H} be the \mathcal{C}^d , the class of unions of at most d singletons. Note that $\mathcal{H}^l = \mathcal{C}^{l+d}$. Hence l -bias assumption with respect to \mathcal{H} is equivalent to \mathcal{C}^{l+d} -realizability. The rest of the proof is analogous to the proof of Theorem 25. ■