

The One-Inclusion Graph Algorithm is not Always Optimal

Ishaq Aden-Ali

ADENALI@BERKELEY.EDU

Department of Electrical Engineering and Computer Science, UC Berkeley.

Yeshwanth Cherapanamjeri

YESHWANTH@BERKELEY.EDU

Department of Electrical Engineering and Computer Science, UC Berkeley.

Abhishek Shetty

SHETTY@BERKELEY.EDU

Department of Electrical Engineering and Computer Science, UC Berkeley.

Nikita Zhivotovskiy

ZHIVOTOVSKIY@BERKELEY.EDU

Department of Statistics, UC Berkeley.

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

The one-inclusion graph algorithm of Haussler, Littlestone, and Warmuth achieves an optimal in-expectation risk bound in the standard PAC classification setup. In one of the first COLT open problems, Warmuth conjectured that this prediction strategy *always* implies an optimal high probability bound on the risk, and hence is also an optimal PAC algorithm. We refute this conjecture in the strongest sense: for any practically interesting Vapnik-Chervonenkis class, we provide an in-expectation optimal one-inclusion graph algorithm whose high probability risk bound cannot go beyond that implied by Markov’s inequality. Our construction of these poorly performing one-inclusion graph algorithms uses Varshamov-Tenengolts error correcting codes.

Our negative result has several implications. First, it shows that the same poor high-probability performance is inherited by several recent prediction strategies based on generalizations of the one-inclusion graph algorithm. Second, our analysis shows yet another statistical problem that enjoys an estimator that is provably optimal in expectation via a leave-one-out argument, but fails in the high-probability regime. This discrepancy occurs despite the boundedness of the binary loss for which arguments based on concentration inequalities often provide sharp high probability risk bounds.

Keywords: PAC Learning, Risk Bounds, Sample Complexity, one-inclusion graphs, Error Correcting Codes

1. Introduction and main result

Learning binary classifiers is arguably the oldest problem in the theory of machine learning. The model, which is captured by the PAC (Probably Approximately Correct) model of learning, traces back to the early works of [Vapnik and Chervonenkis \(1964, 1974\)](#) and of [Valiant \(1984\)](#). A PAC learning algorithm is one that satisfies the following. Given a hypothesis class \mathcal{F} (set of binary-valued functions) on a domain \mathcal{X} , we observe a sample of points $(X_1, f^*(X_1)), \dots, (X_n, f^*(X_n))$, called the *training sample*, where X_1, \dots, X_n are independent copies of a random variable $X \in \mathcal{X}$ distributed according to an unknown distribution P , and $f^* \in \mathcal{F}$ is an unknown target hypothesis. The aim is to provide an algorithm that, based on the sample, outputs a hypothesis \hat{f} whose prediction error is as small as possible with high probability with respect to the realization of the training sample. Formally, we define the error $\text{err}_P(\hat{f}) = \Pr_{X \sim P} [\hat{f}(X) \neq f^*(X)]$ as the probability of misclassification on a freshly sampled point. Since the target hypothesis f^* is in the hypothesis class \mathcal{F} , the most natural strategy is to pick *any* hypothesis $\hat{f} \in \mathcal{F}$ that is *sample-consistent*:

$\hat{f}(X_i) = f^*(X_i)$ for all $i = 1, \dots, n$. The hypothesis selected using this strategy is usually referred to as an *empirical risk minimizer* (ERM). Denoting any of these hypotheses by \hat{f}_{ERM} , the standard bound (Vapnik and Chervonenkis, 1968; Blumer et al., 1989) shows that

$$\text{err}_P(\hat{f}_{\text{ERM}}) = O\left(\frac{d}{n} \log\left(\frac{n}{d}\right) + \frac{1}{n} \log\left(\frac{1}{\delta}\right)\right), \quad (1)$$

with probability at least $1 - \delta$, where d is the Vapnik-Chervonenkis (VC) dimension of \mathcal{F} (formally defined in Section 3). Although this bound has been recently sharpened for ERM for some specific concept classes (Hanneke, 2016b; Zhivotovskiy and Hanneke, 2018), it is known that the optimal risk bound

$$\Theta\left(\frac{d}{n} + \frac{1}{n} \log\left(\frac{1}{\delta}\right)\right) \quad (2)$$

can only be achieved in general by *improper algorithms* (Bousquet et al., 2020). These are algorithms that output a hypothesis outside the hypothesis class \mathcal{F} and thus exclude the standard ERM strategy. The question of achieving the optimal sample complexity (2) of PAC learning has been resolved by Hanneke (2016a), whose solution sharpens the majority vote analysis of Simon (2015). The solution uses a recursive majority vote scheme to achieve the optimal sample complexity (2).

However, for many years, the most natural candidate for being an optimal PAC learner was the *one-inclusion graph* algorithm of Haussler, Littlestone, and Warmuth (1994), whose *in-expectation* risk is known to be optimal in this model. We now briefly describe the one-inclusion graph algorithm. Recall that the projection of a hypothesis class \mathcal{F} onto a subset of the domain $S = \{x_1, \dots, x_n\} \subseteq \mathcal{X}$ is defined as $\mathcal{F}|_S = \{(f(x_1), \dots, f(x_n)) : f \in \mathcal{F}\}$. Informally, the one-inclusion graph algorithm pre-determines a strategy that, for all possible realizations S of the set of unique elements in the training sample and candidate test points x , orients/directs the edges of the *one-inclusion graph* $\mathcal{G}(\mathcal{F}|_{S \cup \{x\}}) = (V, E)$ whose vertices are $\mathcal{F}|_{S \cup \{x\}}$ and whose edges connect two hypotheses that only differ on a *single* point in $S \cup \{x\}$. At test time, given a concrete realization of S and test point x , the algorithm finds the edge (a pair of hypotheses) consistent with the training sample in the one-inclusion graph $\mathcal{G}(\mathcal{F}|_{S \cup \{x\}})$, and predicts on x using the label that the head of this edge (a hypothesis) assigns to x . When the maximum out-degree of the vertices is small, one can show that the expected error of this prediction strategy is small (Haussler et al., 1994). Throughout this section we denote the output of a one-inclusion graph algorithm by \hat{f}_{OIG} . We postpone a formal description of this algorithm to Section 3.

The result by Haussler, Littlestone, and Warmuth (1994) shows that the following error bounds hold for *any* distribution P and *any* target concept $f^* \in \mathcal{F}$,

$$\mathbb{E} \text{err}_P(\hat{f}_{\text{OIG}}) \leq \frac{d}{n+1}, \quad \text{and by Markov's inequality:} \quad \text{err}_P(\hat{f}_{\text{OIG}}) \leq \frac{d}{(n+1)\delta}, \quad (3)$$

with probability at least $1 - \delta$. In the bound above, the expectation is taken with respect to the random sample X_1, \dots, X_n . In one of the first COLT open problems Warmuth conjectured (Warmuth, 2004) that the one-inclusion graph algorithm can *always* achieve the optimal sample complexity (2). One of the motivations of Warmuth was that the one-inclusion graph algorithm is almost optimal in expectation (Li et al., 2001), including even the constant factor which is asymptotically tight. In this context, it is worth mentioning that the existing solution of Hanneke (2016a) has a relatively large constant in the risk bound. Therefore, as noticed by Hanneke (2016a), a positive solution to the conjecture of Warmuth could also lead to improved constant factors in the optimal bound (2). Recently,

Larsen (2022) used arguments inspired by Hanneke (2016a) to show that the classical bagging algorithm also achieves the optimal PAC complexity bound, while also bringing up the question about optimal constants. Some additional interest in this question arises from numerous recent generalizations of the one-inclusion graph algorithm, some of which will be mentioned in Section 2. By understanding the basic PAC learning setting, we can possibly improve the bounds for these newer generalizations.

In this paper, we refute the conjecture of Warmuth in a strong sense by showing that the above application of Markov’s inequality is essentially the best one can hope for in general. Our negative result works for almost any practically interesting hypothesis class including the class induced by half-spaces in \mathbb{R}^p for $p \geq 2$. Before stating our result and Warmuth’s question formally, we first introduce a few notions.

A hypothesis class \mathcal{F} is said to contain a *star set* of arbitrary size if for each n , there is a subset $S = \{x_1, \dots, x_n\}$ of the domain and classifiers $f_0, f_1, \dots, f_n \in \mathcal{F}$ such that for all i , f_0 and f_i disagree on a unique point in S . A more formal definition can be found in Definition 2. In particular, such classes are exactly the classes whose *star number* (see the formal definition in (Hanneke and Yang, 2015, Definition 2)) is infinite. Many practically interesting VC classes have arbitrarily large star sets. These examples include the class of intervals on the real line as well as the class of half-spaces in \mathbb{R}^p for $p \geq 2$. We refer to (Hanneke and Yang, 2015, Section 4.1) for a detailed discussion.

Let \mathcal{F} be a hypothesis class defined on a countable set \mathcal{X} with finite VC dimension d . Informally, we say that a one-inclusion graph algorithm is *valid* for \mathcal{F} if for any n -element subset $S \subseteq \mathcal{X}$, the one-inclusion graph of the projection of \mathcal{F} on S is oriented in a way such that the maximum out-degree of the one-inclusion graph is at most $O(d)$. A formal definition can be found in Section 3. The result of Haussler, Littlestone, and Warmuth (1994) implies that every VC class has a valid one-inclusion graph algorithm. Moreover, for any such algorithm both of the bounds in (3) hold up to multiplicative constant factors. Observe that orientations of one-inclusion graphs with out-degree at most $O(d)$ are not necessarily unique, so in general we have a family of valid strategies. Using this definition, we can phrase Warmuth’s conjecture (Warmuth, 2004) as follows.

Conjecture (Warmuth (2004)) *The optimal sample complexity bound (2) holds for any valid one-inclusion graph algorithm.*

Our main result negatively resolves Warmuth’s conjecture.

Theorem 1 *Let \mathcal{F} be a hypothesis class defined on a countable set \mathcal{X} with finite VC dimension d . Assume that \mathcal{F} has a star set of arbitrary size. There are positive absolute constants c_1, c_2 and c_3 such that for any sample size n and confidence parameter δ satisfying $\delta \in (c_1 d/n, c_2)$, there is a valid one-inclusion graph algorithm that outputs the hypothesis \hat{f}_{OIG} such that for some distribution $P = P(n, \delta)$ over \mathcal{X} and target concept $f^* \in \mathcal{F}$ we have, with probability at least δ ,*

$$\text{err}_P(\hat{f}_{\text{OIG}}) \geq c_3 \frac{d}{n\delta}.$$

1.1. Proof overview

We now provide a high-level overview of our proof. For simplicity, we will sketch the argument for the hypothesis class \mathcal{F}^{ind} consisting of functions that take on the value 1 on at most one point in

the domain $\mathcal{X} = \mathbb{Z}$. In other words, \mathcal{F}^{ind} consists of the zero function and the indicator functions of single points on \mathcal{X} . It is easy to verify that the VC dimension of \mathcal{F}^{ind} is 1. Furthermore, notice that every n -element subset $S = \{x_1, \dots, x_n\} \subseteq \mathcal{X}$ is a star set for \mathcal{F} , witnessed by $\mathcal{F}|_S = \{f_0, \dots, f_n\}$, where f_0 is the zero function and f_i is the indicator on x_i for $i \geq 1$.

Recall that the one-inclusion graph algorithm uses the entire training sample S together with the test point X_{n+1} to determine the label of X_{n+1} by building and orienting a one-inclusion graph. For our specific class \mathcal{F}^{ind} , whenever the test point X_{n+1} is not in the training sample, this strategy always boils down to picking between the label $y_{n+1} = 1$ determined by f_{n+1} and the label $y_{n+1} = 0$ determined by f_0 . Our goal will be to show that there is a choice of orientation that is valid for this class, but leads to a one-inclusion graph algorithm that has constant error with decent probability.

We now explain the desired properties that we want the orientations we choose to satisfy. Fix a sample size n . Our hard distribution will be the uniform distribution on the set $[2n]$ and we will pick the target function f^* to be the zero function. Let S be a training sample of size n and assume for simplicity we always sample n unique points. Our goal will be to show that a $\Theta(1/n)$ fraction of the possible training samples can be made into “bad” training samples. Specifically, we will show that for any such “bad” training sample S , when we project \mathcal{F}^{ind} onto S and an unseen test point X_{n+1} , the orientation determined by the one-inclusion graph algorithm directs the edge $e = \{f_0, f_{n+1}\}$ from the zero function f_0 to f_{n+1} for a constant fraction of the realizations of X_{n+1} . (See Fig. 1.) This would immediately imply that with probability at least $\Theta(1/n)$ (we sample the training sample uniformly) we get a one-inclusion graph algorithm that has constant error.

How should we select these “bad” training samples? A naive approach would be to try and pick random orientations, i.e., for each of the possible one-inclusion graphs formed by $n + 1$ sized sets $S' \subseteq [2n]$, pick a random edge $\{f_0, f_i\}$ and direct it towards f_i . Unfortunately, such an approach cannot work since it does not “coordinate” the errors well. Concretely, let S be an n -sized training set, $Z = [2n] \setminus S$ (the test points not observed in the training set) and for any $X_{n+1} \in Z$, $A_{X_{n+1}}$ is the event that the oriented one-inclusion graph for $S \cup \{X_{n+1}\}$ directs $e = \{f_0, f_{n+1}\}$ towards f_{n+1} . Then the probability that more than $t = C \log(n)$, for some large constant C , possible extensions of S are oriented towards f_{n+1} is bounded as

$$\Pr [\text{At least } t \text{ of the } A_{X_{n+1}} \text{ occur}] \leq \binom{n}{t} \left(\frac{1}{n+1}\right)^t \leq \frac{n^t}{t!} \cdot \frac{1}{(n+1)^t} < \frac{1}{t!} < \frac{1}{n}.$$

Hence, this strategy of randomized orientations would result in a set of orientations where at most $1/n$ fraction of the training samples incur error more than $t/n = O(\log(n)/n)$, a bound significantly worse than the desired bound from Theorem 1 which corresponds to constant error in this scenario. The key point of failure in this approach is that for any fixed training sample, S , the orientation of each of its extensions is independently chosen. To overcome this, we use a different approach.

We instead shift our perspective towards the possible *training sets*, S , that one may observe and correlate the orientations of the $(n + 1)$ -sized extensions such that large error is incurred when S is drawn as a training set. Recall that these orientations were previously chosen randomly. Formally, let \mathcal{S} denote the set of $(2n)$ -length binary vectors with exactly n -ones where each element corresponds to a possible training set that we may sample. Our goal now is to identify a subset T such that $|T| \geq \Omega(|\mathcal{S}|/n)$ and for each $S \in T$, orient most of its extensions such that the edge $e = \{f_0, f_{n+1}\}$ is oriented towards f_{n+1} . The core challenge here is that any given $n + 1$ sized set, S' , has $n + 1$ possible training sets which could have been extended to generate it and we need to

ensure that while defining extensions for two distinct $S_1, S_2 \in T$, we do not generate two contradictory orientations for the *same* $n + 1$ sized extension, S' . Our problem now reduces to the task finding a suitable set, T , which simultaneously constitutes a significant fraction of \mathcal{S} and whose elements do not result in clashes when extended to their $n + 1$ sized counterparts. To do this, we exploit a connection to coding theory.

Consider the following family of vectors:

$$\mathcal{C}_a(2n) = \left\{ C \in \{0, 1\}^{2n} : \sum_i^{2n} i \cdot C(i) \equiv a \pmod{2n + 1} \right\},$$

where $a \in \{0, 1, \dots, 2n\}$. These are the celebrated *Varshamov-Tenengolts* (VT) error correcting codes (with parameter a) introduced by [Varshamov and Tenengolts \(1965\)](#). In the context of coding theory, VT codes are able to recover a transmitted message $C \in \mathcal{C}_a(2n)$ from a corrupted message \tilde{C} where a single element in the vector was flipped from a 1 to a 0.¹ A consequence of this property of $\mathcal{C}_a(2n)$ is that, for any two vectors $C_1, C_2 \in \mathcal{C}_a(2n)$ that have the same number of 1s, C_1 and C_2 must differ on more than 2 entries. We can use this “uniqueness” property in the following way: consider the subsets of $\mathcal{C}_a(2n)$ that have an equal number of 0s and 1s and call this set T_a . We will view T_a as a collection of possible training samples we can receive. Notice that any extension of any two distinct training samples $S_1, S_2 \in T_a$ to $n + 1$ sized sets S'_1 and S'_2 must satisfy $S'_1 \neq S'_2$ by the “uniqueness” property and hence, this rules out the possibility of obtaining contradictory orientations for the same $n + 1$ sized extensions from two different training sets. Thus, for any training sample $S \in T_a$, we can coordinate the error of the algorithm by orienting the one-inclusion graphs formed from every extension $S' = S \cup \{x_{n+1}\}$ (where $x_{n+1} \in [2n] \setminus S$) to direct the edge $\{f_0, f_{n+1}\}$ towards f_{n+1} . By picking an a that maximizes the size of T_a , we can conclude that T_a contains at least a $1/(2n + 1)$ fraction of the possible training samples, each of which is “bad” since they induce a prediction error of $1/2$. Our argument easily applies to VC classes that have arbitrarily large star sets. In the actual proof we extend the argument above to use multiple T_a ’s together with a careful application of the probabilistic method. Doing this introduces a tradeoff between prediction error and failure probability that incorporates all three parameters δ , n , and d . The full details of our construction can be found in [Section 4](#).


The remainder of this paper is organized as follows. In [Section 2](#) we survey the relevant literature and discuss related results. In [Section 3](#) we state some preliminary definitions and introduce notation used throughout the paper. [Section 4](#) is devoted to our construction and the full proof of [Theorem 1](#).

2. Related work and discussion


In this section, we present some remarks to help the reader place our result in a broader context. We also discuss some relevant literature.

The leave-one-out/exchangeability argument and high probability risk bounds. The upper bound for any valid one-inclusion graph algorithm is proven in [Haussler et al. \(1994\)](#) using a leave-one-out argument. The earliest theoretical analysis of a leave-one-out argument in our context is sometimes attributed to [Lunts and Brailovsky \(1967\)](#), who noticed that although this method gives sharp in-expectation risk bounds, it does not necessarily lead to high probability/low variance

1. In fact, as shown by [Levenshtein \(1966\)](#), VT codes can even handle single bit deletions. See ([Sloane, 2002](#)).



(a) An orientation that incorrectly predicts using f_{n+1} .



(b) Orientation that always correctly predicts using f_0 .

Figure 1: Projection of the class of indicators \mathcal{F}^{ind} onto a size $n + 1$ star set S . For $i \geq 1$ the function f_i is the function that disagrees with f_0 on x_i . The red edge represents which hypothesis the one-inclusion graph algorithm picks between to determine the label of the point x_{n+1} .

bounds. In fact, Lunts and Brailovsky connected the variance of the leave-one-out bound with stability² properties of the underlying learning algorithm and provided an example where the variance of the prediction error can be large despite a small leave-one-out error. Based on this, Vapnik and Chervonenkis later asked (Vapnik and Chervonenkis, 1974, Chapter VI, Section 7) if, for most practically interesting classification algorithms, a small leave-one-out error also leads to small prediction error with high probability. It is worth mentioning that the recent analysis of *stable compression schemes* in (Zhivotovskiy, 2017; Bousquet et al., 2020; Hanneke and Kontorovich, 2021) shows that a leave-one-out argument can lead to sharp high probability bounds when additional stability-type assumptions are made. The result of Theorem 1 is on the other end of the spectrum: we construct a valid one-inclusion graph algorithm that does not satisfy any of these stability-type properties.

Confidence boosting approach. One of the standard ways to boost low confidence classifiers (corresponding to e.g., $\delta = 1/2$ in (3)) is based on the following approach. One first splits the sample into approximately $\log(1/\delta)$ non-intersecting equal parts to learn roughly $\log(1/\delta)$ independent classifiers, and then aggregates them to pick the best classifier. This idea is exploited in Haussler et al. (1994), where the authors run the one-inclusion graph algorithm multiple times to provide an algorithm whose probability of error is bounded by

$$O\left(\frac{d}{n} \log\left(\frac{1}{\delta}\right)\right).$$

2. By stability we mean sensitivity to small perturbations in the training sample.

While confidence boosting appears to be a general technique to get high probability bounds, we argue that at least in some cases we have to understand the high probability performance of the original algorithm. First, we observe that the term $\log\left(\frac{1}{\delta}\right)$ appears multiplicatively in the risk bound for the boosted algorithm, while the optimal risk bound (2) has an additive $\log\left(\frac{1}{\delta}\right)$ term. Second, it is not even clear that this approach is generally applicable to leave-one-out based estimators as the confidence boosting approach exploits the realizable/boundedness of the loss assumptions in a strong sense. In particular, it has been recently shown in (Mourtada et al., 2021) that in a related setup of agnostic linear regression the algorithm of Forster and Warmuth (2002), whose analysis follows the same leave-one-out argument used in the analysis of the one-inclusion graph algorithm (Haussler et al., 1994), provides a constant risk with at least constant probability. When learning in the presence of noise, Markov’s inequality in the confidence boosting trick cannot be applied as the *excess risk* is potentially negative for improper learners (see (Mourtada et al., 2021) for more details).

Orientations of one-inclusion graph leading to optimal PAC learners. For some specific classes the one-inclusion graph algorithm may correspond to an optimal PAC learner. In particular, for *intersection closed* VC classes the so-called *closure* algorithm achieves the bound (2). We refer to (Darnstädt, 2015) for a detailed description as well as to a sequence of papers (Helmbold et al., 1990; Auer and Ortner, 2007; Hanneke, 2016b; Bousquet et al., 2020), where various results showing the optimality of the closure algorithm are provided. However, as noted by Warmuth (2004), the closure algorithm corresponds to a *specific* family of orientations of the one-inclusion graph (the same observation is explicit in (Auer and Ortner, 2007, Remark 2)). It is not clear if the known arguments (the original proof was based on a network flow argument (Haussler et al., 1994), and another existence argument of Haussler (1995) is based on the result of Alon and Tarsi (1992)) that focus on the *existence* of valid one-inclusion graph algorithms will lead to orientations exactly corresponding to the closure algorithm for intersection-closed classes. In fact, it follows from our analysis that for some intersection-closed classes we have two valid one-inclusion graph prediction strategies such that one achieves the optimal PAC bound (2), while the second cannot bypass the tail of Markov’s inequality (3).

Applications and extensions of the OIG algorithm. Despite some limitations observed in this paper, the importance of the original one-inclusion algorithm of Haussler, Littlestone, and Warmuth (1994) follows from various applications and extensions of this algorithm in the literature. Here we list some of them. The analysis of the one-inclusion graph algorithm is used in the proof of Haussler’s packing lemma (Haussler, 1995) that provides a sharp bound for covering numbers of VC classes. We refer to (Bartlett and Long, 1998; Kupavskii and Zhivotovskiy, 2020) for extensions and simplifications of Haussler’s analysis. Haussler’s packing lemma is widely used in empirical process theory and computational geometry.

A careful inspection of the analysis of Haussler, Littlestone, and Warmuth (1994) shows that the error bound of the one-inclusion graph does not depend on the VC dimension itself, but rather on the average density of subgraphs of the one-inclusion graph. This observation has allowed for a multi-class extension³ of this algorithm leading to new complexity measures and non-trivial risk bounds in this setting (Rubinstein et al., 2009; Simon and Szörényi, 2010; Daniely and Shalev-Shwartz, 2014; Brukhim et al., 2022). Similarly, the one-inclusion graph algorithm can be extended to the

3. We remark that the analysis of Hanneke’s optimal PAC algorithm (Hanneke, 2016a) uses a VC dimension based uniform convergence argument so it does not easily apply to these extensions.

(bounded) real valued regression setup (Bartlett and Long, 1998). The analysis of Long (1998) allows, in particular, to extend the one-inclusion graph algorithm to the agnostic classification setting. Explicit bounds of this sort can be found in (Hanneke et al., 2022). Furthermore, the one-inclusion graphs algorithm and some extensions have proven to be useful in other generalizations of the binary classification setup such as universal learning (Bousquet et al., 2021), robust learning (Attias et al., 2022; Montasser et al., 2022), and several other setups (Haghifam et al., 2022; Alon et al., 2022; Charikar and Pabbaraju, 2022). As we mentioned earlier, our lower bound has implications for some of the settings studied in these papers.

3. Preliminaries and notation

In our problem setting, there is an unknown probability distribution P over some countable *instance space* \mathcal{X} that generates examples, and a known hypothesis class \mathcal{F} which is a collection of binary functions (hypotheses) that map from the instance space \mathcal{X} to the *label space* $\mathcal{Y} = \{0, 1\}$. Furthermore, there is an unknown *target* hypothesis $f^* \in \mathcal{F}$ that labels the examples generated by P . A learning algorithm receives an i.i.d. *training sample* $\bar{S} = (X_1, \dots, X_n)$ sampled from P along with their accompanying *labels* $(f^*(X_1), \dots, f^*(X_n))$ as input, and produces a hypothesis \hat{f} (not necessarily in \mathcal{F}) as its output. The goal of the learning algorithm is to produce a hypothesis \hat{f} that has low error under P which we define to be $\text{err}_P(\hat{f}) = \Pr_{X \sim P} [\hat{f}(X) \neq f^*(X)]$. We will find it convenient to differentiate between the training sample \bar{S} that is potentially a multiset and its corresponding set version S which we will refer to as the *training set*. It will also be convenient to define the labelled training sample $(\bar{S}, f^*(\bar{S})) = ((X_1, f^*(X_1)), \dots, (X_n, f^*(X_n)))$. The labelled training set $(S, f^*(S))$ is defined similarly.

We define the uniform distribution over a finite set A to be $U(A)$. When the choice of the set A is clear from context we will sometimes abbreviate this to U . For a finite set A and any $k \leq |A|$, define $\mathcal{S}_k(A)$ to be the set of subsets of A of size k , i.e., $\mathcal{S}_k(A) = \{A' \subseteq A : |A'| = k\}$. We will often abbreviate this to \mathcal{S}_k when the finite set A we use is clear from the context. For any positive integer n define $[n] = \{1, 2, \dots, n\}$. Given a hypothesis class \mathcal{F} and subset of the instance space $S = \{x_1, \dots, x_n\} \subseteq \mathcal{X}$, we define the projection of \mathcal{F} onto S to be $\mathcal{F}|_S = \{(f(x_1), \dots, f(x_n)) : f \in \mathcal{F}\}$. In words, the projection is the set of all functions the hypothesis class \mathcal{F} realizes on the set S . We say an n -element set $S = \{x_1, \dots, x_n\} \subseteq \mathcal{X}$ is *shattered* by the hypothesis class \mathcal{F} if the projection $\mathcal{F}|_S = \{0, 1\}^n$, i.e., \mathcal{F} realizes every possible function on S . The *Vapnik-Chervonenkis (VC) dimension* of a hypothesis class \mathcal{F} is the largest integer d such that there exists a d -element subset of \mathcal{X} that is shattered by \mathcal{F} . Throughout this paper we will use the variable d to denote the VC dimension of hypothesis class \mathcal{F} and it will always be clear from context which hypothesis class d will correspond to. For a projection $\mathcal{F}|_S$ with $|S| = n$, we define the Hamming distance between two hypotheses f and g in $\mathcal{F}|_S$, denoted by $\rho_n(f, g)$, to be the number of elements in S that f and g differ on: $\rho_n(f, g) = |\{x \in S : f(x) \neq g(x)\}|$.

We next define the notion of star sets discussed earlier. The notion was first explicitly defined in the context of active learning, see (Hanneke and Yang, 2015). This will be the main notion of complexity that will be used in our lower bound.

Definition 2 (Star sets) *A concept class \mathcal{F} defined on an infinite domain \mathcal{X} has a star set of arbitrary size if for any integer n , there exist $S = \{x_1, \dots, x_n\} \subseteq \mathcal{X}$ and classifiers $f_0, f_1, \dots, f_n \in \mathcal{F}$*

such that for all $i = 1, \dots, n$,

$$\{x \in S : f_i(x) \neq f_0(x)\} = \{x_i\}.$$

For any n such a set S is a star set. The corresponding classifiers $\{f_i\}_{i=0}^n$ witness the star set S .

We now formally define one-inclusion graphs, orientations of one-inclusion graphs, and the one-inclusion graph algorithm.

Definition 3 (One-inclusion graph) Fix a hypothesis class \mathcal{F} and an n -element subset of the domain $S \subseteq \mathcal{X}$. The one inclusion graph $\mathcal{G}(\mathcal{F}|_S) = (V, E)$ has its vertex set as $V := \mathcal{F}|_S$ and the edge set as

$$E := \{\{f, g\} : f, g \in \mathcal{F}|_S, \rho_n(f, g) = 1\}.$$

In words, we connect an edge between two vertices (projected hypotheses) if and only if they differ on a single element in the set S .

In order to use the one-inclusion graph in a prediction algorithm, we will “orient” the undirected one-inclusion graph into a directed graph.

Definition 4 (Orientations for one-inclusion graphs) Fix hypothesis class \mathcal{F} and any finite subset $S \subseteq \mathcal{X}$. An orientation of $\mathcal{G}(\mathcal{F}|_S)$ is a function $\sigma_S : E \rightarrow V$ such that for any $e \in E$, $\sigma_S(e) \in e$. Let V be a collection of finite subsets of \mathcal{X} . An orientation rule for V , usually denoted by σ , is a collection of orientations for each element of V , that is,

$$\sigma = \{\sigma_S : S \in V\}.$$

An orientation rule now denotes orientations for all possible finite subsets of \mathcal{X} .

Put differently, an orientation takes an *undirected* one-inclusion graph and defines a corresponding *directed* one-inclusion graph by determining the head of each edge. An orientation rule for hypothesis class \mathcal{F} just tells us how we should orient any of the possible one-inclusion graphs we can obtain from \mathcal{F} and any finite subset of the domain S . For a one-inclusion graph $\mathcal{G}(\mathcal{F}|_S) = (V, E)$ and orientation σ_S , define the *out-degree* of a vertex $v \in V$ to be

$$\text{out}(v; \sigma_S) = |\{e \in E : v \in e, \sigma_S(e) \neq v\}|.$$

The *max out-degree* of $\mathcal{G}(\mathcal{F}|_S)$ is naturally defined as $\text{out}(\sigma_S) = \max_{v \in V} \text{out}(v; \sigma_S)$.

We are now ready to describe the one-inclusion graph algorithm. We present its pseudocode in [Algorithm 1](#). A one-inclusion graph algorithm is defined by the orientation rule σ it uses and the labeled training set $(S, f^*(S))$ it receives. This is reflected in the notation $\hat{f}_{\sigma, S}$ that we use for the hypothesis produced by the one-inclusion graph algorithm (see [Algorithm 1](#)). We now define what it means for a one-inclusion graph algorithm to be valid.

Definition 5 Fix a hypothesis class \mathcal{F} with VC dimension d . We say a one-inclusion graph algorithm for \mathcal{F} that uses orientation rule σ is valid if

$$\max_{\sigma_S \in \sigma} \text{out}(\sigma_S) = O(d).$$

A beautiful result of [Haussler et al. \(1994\)](#) shows that there is always a valid one-inclusion graph algorithm for any hypothesis class with finite VC dimension.

Theorem 6 (Haussler et al. (1994, Theorem 2.2)) *For any hypothesis class \mathcal{F} over instance space \mathcal{X} with VC dimension at most d and any finite $S \subseteq \mathcal{X}$, there exists an orientation σ_S of $\mathcal{G}(\mathcal{F}|_S)$ with*

$$\text{out}(\sigma_S) \leq d.$$

Algorithm 1 One-inclusion graph algorithm.

Inputs: Labelled training set $(S, f^*(S))$, orientation rule σ .

Output: Hypothesis $\hat{f}_{\sigma,S} : \mathcal{X} \rightarrow \mathcal{Y}$.

For any point $x \in \mathcal{X}$ the hypothesis $\hat{f}_{\sigma,S}$ predicts as follows:

- 1: If there is a unique label y for x consistent with $(S, f^*(S))$ and \mathcal{F} , predict y .
 - 2: Let e be the edge in $\mathcal{G}(\mathcal{F}|_{S \cup \{x\}})$ with hypotheses consistent with $(S, f^*(S))$ but not on x .
 - 3: Predict according to $\sigma_{S \cup \{x\}}(e)$, i.e., the hypothesis pointed to in the orientation of e .
-

4. The construction

In this section, we prove our main result. Let $A = \{x_1, \dots, x_m\}$ be a set of size m with a fixed ordering of its elements. For any k -element subset $B \subseteq A$ with $k \leq m$, we will slightly overload notation by simultaneously referring to B as a subset of elements of A , and the length n binary vector that is 1 in the i -th entry if and only if x_i is in B . We will clarify the definition of the set A such that its correspondence to length n binary vectors is clear.

We begin by recalling VT codes and stating a useful property they enjoy. VT codes of length m and parameter $a \in \{0, 1, \dots, m\}$ are given by the following family of vectors:

$$\mathcal{C}_a(m) = \left\{ C \in \{0, 1\}^m : \sum_{i=1}^m i \cdot C(i) \equiv a \pmod{m+1} \right\}.$$

We will find the following notion of coverage to be useful.

Definition 7 (Coverage) *Fix an integer m . We say a vector $S \in \{0, 1\}^m$ is covered by the vector $S' \in \{0, 1\}^m$ if there is an index $i \in [m]$ such that $S(i) = 0$ but $S'(i) = 1$, and for every $j \neq i$ we have $S(j) = S'(j)$. We will denote this by $S \prec S'$.*

With this notation, VT codes satisfy a “uniqueness” property crucial for our analysis.

Lemma 8 (Unique neighborhoods) *Let A be a set of m elements. For any $k \leq m$, binary vector $S' \in \mathcal{S}_k(A)$ and $a \in \{0, 1, \dots, m\}$, there exists at most one binary vector $S \in \mathcal{C}_a(m)$ with $S \prec S'$.*

Proof Towards a contradiction assume that for some $S' \in \mathcal{S}_k(A)$ there were two vectors $S_1, S_2 \in \mathcal{C}_a(m)$ that satisfied the property above with corresponding indices i_1 and i_2 . Then, we have that S_1 and S_2 disagree only on i_1 and i_2 . So

$$0 \equiv \sum_{i=1}^m i \cdot S_1(i) - \sum_{i=1}^m i \cdot S_2(i) \equiv i_2 - i_1 \pmod{m+1}. \quad (4)$$

This is a contradiction since $1 \leq i_1, i_2 \leq m$. ■

Let \mathcal{F} be the hypothesis class with star sets of arbitrary size and VC dimension d for which we would like to construct the bad one-inclusion graph algorithm. Let n denote the sample size and let $\mathcal{W}_{2n} \subseteq \mathcal{X}$ be a star set of size $2n$. Let f_0, f_1, \dots, f_{2n} denote the corresponding functions that witness the star set. Through a simple re-labeling procedure, we may assume that the center of the star set f_0 is the all zeros function; i.e., $f_0(x) = 0$ for all $x \in \mathcal{W}_{2n}$. Our distribution will simply be the uniform distribution over \mathcal{W}_{2n} with the labels generated by $f^* = f_0$. Hence, our labelled training sample consists of points $\bar{S} = ((X_1, f^*(X_1)), \dots, (X_n, f^*(X_n)))$ where each X_i is sampled from the uniform distribution $U := U(\mathcal{W}_{2n})$ and $f^*(X_i) = 0$.

We will first prove a lower bound on the error conditioned on the number of unique elements observed in our training sample $\bar{S} \sim U^n$, i.e., the size of the *training set* S . Let $|S| = k$. Note that S may be associated with an element of $\mathcal{S}_k(\mathcal{W}_{2n})$. Slightly abusing notation we denote $\mathcal{S}_k(\mathcal{W}_{2n})$ by \mathcal{S}_k , and $\mathcal{S}_{k+1}(\mathcal{W}_{2n})$ by \mathcal{S}_{k+1} respectively. We will now define a bipartite graph with one set of vertices a subset of \mathcal{S}_k (and hence, a possible realization of S) and the other \mathcal{S}_{k+1} . This graph will then be used to construct a orientation rule with poor performance. Formally, the vertex sets of the graph are defined below

$$\begin{aligned} V_1 &= \bigcup_{0 \leq i < 4\lceil \delta n \rceil} T_{(i)}, \\ V_2 &= \mathcal{S}_{k+1}, \end{aligned} \tag{VERT-SETS}$$

where $T_i := \mathcal{C}_i(2n) \cap \mathcal{S}_k$ is the intersection of the code $\mathcal{C}_i(2n)$ with the set of vectors containing exactly k ones and $T_{(0)}, \dots, T_{(2n)}$ is a re-ordering such that $|T_{(0)}| \geq |T_{(1)}| \geq \dots \geq |T_{(2n)}|$. In particular, since $\mathcal{C}_i(2n) \cap \mathcal{C}_j(2n) = \emptyset$ for $i \neq j$, we have for any $\ell \in [2n] \cup \{0\}$,

$$\sum_{i=0}^{\ell} |T_{(i)}| \geq \frac{\ell+1}{2n+1} \cdot |\mathcal{S}_k| \geq \frac{\ell+1}{3n} \cdot |\mathcal{S}_k|. \tag{5}$$

We will construct our family of one-inclusion graphs such that for a substantial fraction of the training sets in V_1 , the corresponding one-inclusion strategy incurs large error. We now construct the edge set of our (undirected) bipartite graph, $G = (V_1, V_2, E)$. Define for all $v \in V_1 \cup V_2$,

$$N(v) = \{u \in V_1 \cup V_2 : (u, v) \in E\},$$

that is, the set of neighbors of v in G . We will use the probabilistic method to construct an edge set, E , satisfying certain cardinality constraints. The first constraint ensures that the family of one-inclusion graphs we construct from G satisfies the appropriate out-degree constraints while the second will be used to show that they incur large error.

Lemma 9 *There exists an edge set E such that the bipartite graph $G = (V_1, V_2, E)$ satisfies,*

$$\text{for all } v' \in V_2, |N(v')| \leq d, \quad \text{and} \quad \left| \left\{ v \in V_1 : |N(v)| \geq \frac{d}{8\delta} \right\} \right| \geq \frac{3}{4} |V_1|.$$

Proof Our proof will utilize the probabilistic method. Defining $M(v') = \{v \in V_1 : v \prec v'\}$ for all $v' \in V_2$, E is constructed according to the following random process. For any $v' \in V_2$:

1. With $\ell = \min \{d, |M(v')|\}$, pick $\{v_i\}_{i \in [\ell]}$ uniformly at random without replacement from $M(v')$.
2. Add all edges $\{(v', v_i)\}_{i \in [\ell]}$ to E .

Note that the first claim of the lemma follows immediately from the definition of the probabilistic process. For the second, fix $v \in V_1$ and define

$$Q(v) = \{v' \in \mathcal{S}_{k+1} : v \prec v'\},$$

and for $v' \in V_2$ define

$$p_{v'} = \frac{\min \{d, |M(v')|\}}{|M(v')|}.$$

We have by the above definition and linearity of expectation

$$\mathbb{E}[|N(v)|] = \sum_{v' \in Q(v)} \frac{\min \{d, |M(v')|\}}{|M(v')|} = \sum_{v' \in Q(v)} p_{v'}.$$

From the independence of the indicators $\mathbf{1}\{v' \in N(v)\}$,

$$\text{Var}(|N(v)|) = \text{Var}\left(\sum_{v' \in Q(v)} \mathbf{1}\{v' \in N(v)\}\right) = \sum_{v' \in Q(v)} \text{Var}(\mathbf{1}\{v' \in N(v)\}) \leq \sum_{v' \in Q(v)} p_{v'}.$$

By Lemma 8 and the definition of V_1 , we have $|M(v')| \leq 4\lceil \delta n \rceil$ for all $v' \in V_2$. When c_1 (Theorem 1) is large enough, we have $4\lceil \delta n \rceil \leq 6\delta n$ and $n\delta \geq c_1 d$, which yields

$$p_{v'} \geq \frac{1}{6} \cdot \frac{d}{\delta n}.$$

Additionally, we have when c_2 (Theorem 1) is small enough and using the fact that $|Q(v)| \geq n$,

$$\frac{\sqrt{\text{Var}(|N(v)|)}}{\mathbb{E}[|N(v)|]} \leq \frac{\sqrt{\sum_{v' \in Q(v)} p_{v'}}}{\sum_{v' \in Q(v)} p_{v'}} = \sqrt{\frac{1}{\sum_{v' \in Q(v)} p_{v'}}} \leq \sqrt{\frac{6\delta n}{d|Q(v)|}} \leq \sqrt{\frac{6\delta}{d}} \leq \frac{1}{10}.$$

This yields via Chebyshev's inequality,

$$\Pr\left[|N(v)| \geq \frac{d}{8\delta}\right] \geq \Pr\left[|N(v)| \geq \frac{|Q(v)|}{8} \cdot \frac{d}{\delta n}\right] \geq \Pr\left[|N(v)| \geq \frac{3}{4} \sum_{v' \in Q(v)} p_{v'}\right] \geq \frac{3}{4}.$$

Linearity of expectation (on $v \in V_1$) establishes the second claim by the probabilistic method. \blacksquare

With the bipartite graph, $G = (V_1, V_2, E)$, defined in [VERT-SETS](#) and Lemma 9, we will now construct our family of one-inclusion graphs for the sets in \mathcal{S}_{k+1} . We think of an element of $S' \in \mathcal{S}_{k+1}$ as the union of a training set S and a new test point x_{n+1} . Informally, for any set $S' \in \mathcal{S}_{k+1}$, we pick any low out-degree orientation and then reorient the edges of the zero function. We prove that for most of the k -sized training sets $S \in V_1$ and a significant fraction

of their “extensions” to $(k + 1)$ -sized sets $S' \in V_2$, the one-inclusion graph algorithm predicts incorrectly on $x_{n+1} \in S' \setminus S$. Our construction is formally described in [Algorithm 2](#). Note that [Algorithm 2](#) only constructs an orientation rule for V_2 (which corresponds to \mathcal{S}_{k+1} in this context).

Algorithm 2 Constructing sub-optimal one-inclusion graphs

Input: Bipartite Graph $G = (V_1, V_2, E)$.

Output: Orientation rule σ for V_2 .

- 1: Set $\sigma = \{\}$.
 - 2: **for** $S' \in V_2$ **do**
 - 3: Pick any orientation $\sigma_{S'}$ that has max out-degree at most d .
 - 4: For every $S \in N(S')$, reorient $\sigma_{S'}(\{f_0, f_i\}) = f_i$ where $f_i = \mathbf{1}\{x \in S' \setminus S\}$.
 - 5: For any other other edge $\{f_i, f_0\}$ not modified in the previous step set $\sigma_{S'}(\{f_0, f_i\}) = f_0$.
 - 6: Add $\sigma_{S'}$ to σ .
 - 7: **end for**
 - 8: **return** σ .
-

Before we proceed, we first show that the orientation rule constructed in [Algorithm 2](#) is valid.

Lemma 10 *When G satisfies the conclusion of [Lemma 9](#) the orientation rule for V_2 , σ , constructed in [Algorithm 2](#) satisfies*

$$\max_{\sigma_{S'} \in \sigma} \text{out}(\sigma_{S'}) \leq d + 1.$$

Remark 11 *For the sake of generality, we prove [Lemma 10](#) with an upper bound of $d + 1$. If we consider more structured classes, e.g., the class of functions that take on the value 1 at most d times, we can improve this bound to d . This would match the out-degree bound of [Theorem 6](#).*

Remark 12 *The bound $d + 1$ on the max out-degree only slightly changes the risk bound for the one-inclusion graph algorithm. In particular, our orientation rule gives*

$$\mathbb{E} \text{err}_P(\widehat{f}_{\text{OIG}}) \leq \frac{d+1}{n+1}, \quad \text{and by Markov's inequality:} \quad \text{err}_P(\widehat{f}_{\text{OIG}}) \leq \frac{d+1}{(n+1)\delta}.$$

Proof of Lemma 10: Let $S' \in \mathcal{S}_{k+1}$. The out-degree of the all-zeros hypothesis f_0 is at most d by the first claim of [Lemma 9](#). For any hypothesis $g \in \mathcal{F}|_{S'}$ with $g \neq f_0$, its out-degree is at most d in the original orientation $\sigma_{S'}$ selected in Line 3 of [Algorithm 2](#) (such an orientation exists by [Theorem 6](#)). We add at most one outgoing edge from g to f_0 in Line 5 concluding the proof. ■

Our last lemma will show that the one-inclusion strategy defined by the orientation rule σ constructed in [Algorithm 2](#) incurs large error for most training sets in V_1 . In fact, this will be shown for the subset whose existence was established in [Lemma 9](#). Formally, define W_1^k as

$$W_1^k = \left\{ v \in V_1 : |N(v)| \geq \frac{d}{8\delta} \right\}.$$

The following lemma shows that any training set from W_1^k incurs large error.

Lemma 13 For any $S \in W_1^k$, the one-inclusion graph algorithm that predicts using the orientation rule σ defined by [Algorithm 2](#) on input G satisfying the conclusion of [Lemma 9](#) has

$$\text{err}_P(\hat{f}_{\sigma,S}) \geq \frac{1}{16} \cdot \frac{d}{\delta n}.$$

Proof By the construction of the orientation rule σ in [Algorithm 2](#), we only make an error on x if $S \cup \{x\} \in N(S)$, as we only predict 1 in this scenario. We now have

$$\text{err}_P(\hat{f}_{\sigma,S}) \geq \frac{|N(S)|}{2n} \geq \frac{1}{16} \cdot \frac{d}{\delta n}$$

establishing the lemma. ■

We will now prove [Theorem 1](#). Recall that our distribution is the uniform distribution over the participating elements of a star set of size $2n$ with the center canonically identified with the zero function and the elements with the set $[2n]$. For each $k \in [n]$ denoting the possible number of unique observed elements, consider the family of one-inclusion prediction strategies constructed in [Algorithm 2](#) with the bipartite graph defined in [VERT-SETS](#) and [Lemma 9](#). By [Lemma 13](#), any unique training set S in W_1^k incurs large error. Hence, we only need to lower bound the probability of observing a training set from $W_1 := \cup_{k \in [n]} W_1^k$. Note that conditioned on $|S| = k$, the training set S is uniformly distributed on \mathcal{S}_k . Hence, we have by [Lemma 9](#) and [\(5\)](#) that

$$\Pr_{\tilde{S} \sim U^n} [S \in W_1^k \mid |S| = k] = \frac{|W_1^k|}{|\mathcal{S}_k|} \geq \frac{3}{4} \cdot \frac{|V_1|}{|\mathcal{S}_k|} = \frac{3}{4} \cdot \sum_{i=0}^{4[\delta n]-1} \frac{|T_{(i)}|}{|\mathcal{S}_k|} \geq \frac{3}{4|\mathcal{S}_k|} \cdot \frac{4[\delta n]|\mathcal{S}_k|}{3n} \geq \delta.$$

To conclude the proof of the theorem, we now have

$$\Pr_{\tilde{S} \sim U^n} [S \in W_1] = \sum_{k=1}^n \Pr_{\tilde{S} \sim U^n} [S \in W_1^k \mid |S| = k] \cdot \Pr_{\tilde{S} \sim U^n} [|S| = k] \geq \delta$$

as required. ■

Acknowledgments

The authors would like to thank Omar Alrabiah for fruitful discussions.

References

- Noga Alon and Michael Tarsi. Colorings and orientations of graphs. *Combinatorica*, 12(2):125–134, 1992. [7](#)
- Noga Alon, Steve Hanneke, Ron Holzman, and Shay Moran. A theory of PAC learnability of partial concept classes. In *62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 658–671, 2022. [8](#)
- Idan Attias, Steve Hanneke, and Yishay Mansour. A characterization of semi-supervised adversarially-robust PAC learnability. *arXiv preprint arXiv:2202.05420*, 2022. [8](#)

- Peter Auer and Ronald Ortner. A new PAC bound for intersection-closed concept classes. *Machine Learning*, 66(2):151–163, 2007. [7](#)
- Peter L Bartlett and Philip M Long. Prediction, learning, uniform convergence, and scale-sensitive dimensions. *Journal of Computer and System Sciences*, 56(2):174–190, 1998. [7](#), [8](#)
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36(4):929–965, 1989. [2](#)
- Olivier Bousquet, Steve Hanneke, Shay Moran, and Nikita Zhivotovskiy. Proper learning, Helly number, and an optimal SVM bound. In *Conference on Learning Theory*, pages 582–609. PMLR, 2020. [2](#), [6](#), [7](#)
- Olivier Bousquet, Steve Hanneke, Shay Moran, Ramon Van Handel, and Amir Yehudayoff. A theory of universal learning. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 532–541, 2021. [8](#)
- Nataly Brukhim, Daniel Carmon, Irit Dinur, Shay Moran, and Amir Yehudayoff. A characterization of multiclass learnability. *arXiv preprint arXiv:2203.01550*, 2022. [7](#)
- Moses Charikar and Chirag Pabbaraju. A characterization of list learnability. *arXiv preprint arXiv:2211.04956*, 2022. [8](#)
- Amit Daniely and Shai Shalev-Shwartz. Optimal learners for multiclass problems. In *Conference on Learning Theory*, pages 287–316. PMLR, 2014. [7](#)
- Malte Darnstädt. The optimal PAC bound for intersection-closed concept classes. *Information Processing Letters*, 115(4):458–461, 2015. [7](#)
- Jürgen Forster and Manfred K Warmuth. Relative expected instantaneous loss bounds. *Journal of Computer and System Sciences*, 64(1):76–102, 2002. [7](#)
- Mahdi Haghifam, Shay Moran, Daniel M. Roy, and Gintare Karolina Dziugiate. Understanding generalization via leave-one-out conditional mutual information. In *IEEE International Symposium on Information Theory (ISIT)*, pages 2487–2492, 2022. [8](#)
- Steve Hanneke. The optimal sample complexity of PAC learning. *The Journal of Machine Learning Research*, 17(1):1319–1333, 2016a. [2](#), [3](#), [7](#)
- Steve Hanneke. Refined error bounds for several learning algorithms. *The Journal of Machine Learning Research*, 17(1):4667–4721, 2016b. [2](#), [7](#)
- Steve Hanneke and Aryeh Kontorovich. Stable sample compression schemes: New applications and an optimal SVM margin bound. In *Algorithmic Learning Theory*, pages 697–721. PMLR, 2021. [6](#)
- Steve Hanneke and Liu Yang. Minimax analysis of active learning. *Journal of Machine Learning Research*, 16(1):3487–3602, 2015. [3](#), [8](#)
- Steve Hanneke, Amin Karbasi, Mohammad Mahmoodi, Idan Mehalal, and Shay Moran. On optimal learning under targeted data poisoning. *arXiv preprint arXiv:2210.02713*, 2022. [8](#)

- David Haussler. Sphere packing numbers for subsets of the boolean n-cube with bounded Vapnik-Chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69(2):217–232, 1995. [7](#)
- David Haussler, Nick Littlestone, and Manfred K Warmuth. Predicting $\{0, 1\}$ -functions on randomly drawn points. *Information and Computation*, 115(2):248–292, 1994. [2](#), [3](#), [5](#), [6](#), [7](#), [10](#)
- David Helmbold, Robert Sloan, and Manfred K Warmuth. Learning nested differences of intersection-closed concept classes. *Machine Learning*, 5(2):165–196, 1990. [7](#)
- Andrey Kupavskii and Nikita Zhivotovskiy. When are epsilon-nets small? *Journal of Computer and System Sciences*, 110:22–36, 2020. [7](#)
- Kasper Green Larsen. Bagging is an optimal PAC learner. *arXiv preprint arXiv:2212.02264*, 2022. [3](#)
- Vladimir I Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10(8):707–710, 1966. [5](#)
- Yi Li, Philip M Long, and Aravind Srinivasan. The one-inclusion graph algorithm is near-optimal for the prediction model of learning. *IEEE Transactions on Information Theory*, 47(3):1257–1261, 2001. [2](#)
- Philip M Long. The complexity of learning according to two models of a drifting environment. In *Conference on Computational learning theory*, pages 116–125, 1998. [8](#)
- Alexander Lunts and Victor Brailovsky. Evaluation of attributes obtained in statistical decision rules. *Engineering Cybernetics*, 3:98–109, 1967. [5](#)
- Omar Montasser, Steve Hanneke, and Nathan Srebro. Adversarially robust learning: A generic minimax optimal learner and characterization. *arXiv preprint arXiv:2209.07369*, 2022. [8](#)
- Jaouad Mourtada, Tomas Vaškevičius, and Nikita Zhivotovskiy. Distribution-free robust linear regression. *Mathematical Statistics and Learning*, 4(3-4):253–292, 2021. [7](#)
- Benjamin Rubinstein, Peter Bartlett, and Hyam Rubinstein. Shifting: One-inclusion mistake bounds and sample compression. *Journal of Computer and System Sciences*, 75(1):37–59, 2009. [7](#)
- Hans U Simon. An almost optimal PAC algorithm. In *Conference on Learning Theory*, pages 1552–1563. PMLR, 2015. [2](#)
- Hans U Simon and Balázs Szörényi. One-inclusion hypergraph density revisited. *Information Processing Letters*, 110(8-9):341–344, 2010. [7](#)
- Neil J. A. Sloane. *On Single-Deletion-Correcting Codes*, pages 273–292. De Gruyter, 2002. [5](#)
- Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984. [1](#)
- Vladimir Vapnik and Alexey Chervonenkis. A class of algorithms for pattern recognition learning. *Avtomatika i Telemekhanika*, 25(6):937–945, 1964. [1](#)

Vladimir Vapnik and Alexey Chervonenkis. Algorithms with complete memory and recurrent algorithms in the problem of learning pattern recognition. *Avtomatika i Telemekhanika*, pages 95–106, 1968. [2](#)

Vladimir Vapnik and Alexey Chervonenkis. *Theory of Pattern Recognition*. Nauka, Moscow, 1974. [1](#), [6](#)

Rom Varshamov and Grigory Tenengolts. Codes which correct single asymmetric errors. *Avtomatika i Telemekhanika*, 26(2):288–292, 1965. [5](#)

Manfred K Warmuth. The optimal PAC algorithm. In *International Conference on Computational Learning Theory*, pages 641–642. Springer, 2004. [2](#), [3](#), [7](#)

Nikita Zhivotovskiy. Optimal learning via local entropies and sample compression. In *Conference on Learning Theory*, pages 2023–2065. PMLR, 2017. [6](#)

Nikita Zhivotovskiy and Steve Hanneke. Localization of VC classes: Beyond local Rademacher complexities. *Theoretical Computer Science*, 742:27–49, 2018. [2](#)