
Towards a Unified Information-Theoretic Framework for Generalization

Mahdi Haghifam*
University of Toronto,
Vector Institute

Gintare Karolina Dziugaite†
Mila

Shay Moran
Technion,
Google Research

Daniel M. Roy
University of Toronto,
Vector Institute

Abstract

In this work, we investigate the expressiveness of the “conditional mutual information” (CMI) framework of Steinke and Zakynthinou [1] and the prospect of using it to provide a unified framework for proving generalization bounds in the realizable setting. We first demonstrate that one can use this framework to express non-trivial (but sub-optimal) bounds for any learning algorithm that outputs hypotheses from a class of bounded VC dimension. We then explore two directions of strengthening this bound: (i) Can the CMI framework express optimal bounds for VC classes? (ii) Can the CMI framework be used to analyze algorithms whose output hypothesis space is unrestricted (i.e. has an unbounded VC dimension)?

With respect to Item (i) we prove that the CMI framework yields the optimal bound on the expected risk of *Support Vector Machines* (SVMs) for learning halfspaces. This result is an application of our general result showing that *stable* compression schemes [2] of size k have uniformly bounded CMI of order $O(k)$.

We further show that an inherent limitation of proper learning of VC classes contradicts the existence of a proper learner with constant CMI, and it implies a negative resolution to an open problem of Steinke and Zakynthinou [3]. We further study the CMI of empirical risk minimizers (ERMs) of class \mathcal{H} and show that it is possible to output all consistent classifiers (version space) with bounded CMI *if and only if* \mathcal{H} has a bounded star number [4].

With respect to Item (ii) we prove a general reduction showing that “leave-one-out” analysis is expressible via the CMI framework. As a corollary we investigate the CMI of the one-inclusion-graph algorithm proposed by Haussler et al. [5]. More generally, we show that the CMI framework is universal in the sense that for *every* consistent algorithm and data distribution, the expected risk vanishes as the number of samples diverges *if and only if* its evaluated CMI has sublinear growth with the number of samples.

1 Introduction

In this work, we study the expressiveness of generalization bounds in terms of information-theoretic measures of dependence between the output of a learning algorithm and input data. Information-theoretic techniques for proving generalization bounds are powerful; they can provide generalization

*Part of the work was done the author was an intern at Element AI, a ServiceNow company.

†This work was carried out while the author was at ServiceNow. It was finalized at Google Brain.

bounds that are algorithm-dependent, data-dependent, and distribution-dependent. This approach was initiated by Russo and Zou [6, 7] and Xu and Raginsky [8] and has since been extended by a number of authors [9–19]. More recently, attention has shifted to whether these techniques can also characterize worst-case (minimax) rates for certain learning problems.

Let \mathcal{D} be an unknown distribution on a space \mathcal{Z} , and let \mathcal{H} be a set of classifiers. Consider a (randomized) learning algorithm $\mathcal{A} = (\mathcal{A}_n)_{n \geq 1}$ that selects an element \hat{h} in \mathcal{H} , based on n i.i.d. samples $S_n \sim \mathcal{D}^{\otimes n}$, i.e., $\hat{h} = \hat{h}_n = \mathcal{A}_n(S_n)$. The initial focus of this line of work was on the mutual information $I(\mathcal{A}_n(S_n); S_n)$ between the input and the output of a learning problem. This quantity is sometimes referred to as the *input–output mutual information (IOMI)* of an algorithm and denoted by $\text{IOMI}_{\mathcal{D}}(\mathcal{A}_n)$. A natural question is whether the IOMI framework can provide a sharp characterization of the learnability of Vapnik–Chervonenkis (VC) classes, for which we have strong generalization guarantees. A negative resolution was provided by Bassily et al. [12] for the concept class of thresholds in one dimension. Follow up work by Nachum et al. [13] extended the argument in [12], proving the following result:

Theorem 1.1 (Thm. 1, [13]). *For every $d \in \mathbb{N}$ and every $n \geq 2d^2$, there exists a finite input space \mathcal{X} and a concept class $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ of VC-dimension d such that, for all proper and consistent learning algorithm \mathcal{A}_n , there exists a realizable distribution \mathcal{D} such that $\text{IOMI}_{\mathcal{D}}(\mathcal{A}_n) = \Omega(d \log \log(|\mathcal{X}|/d))$.*

Livni and Moran [18] extended this result even further, showing that, for the class of one-dimensional thresholds over $\{1, \dots, m\}$, $m \in \mathbb{N}$,³ for every learning algorithm \mathcal{A} there exists a realizable distribution such that either the risk (population loss) is large or the $\text{IOMI}_{\mathcal{D}}(\mathcal{A}_n)$ scales with the cardinality of the space, m . These results highlight an important limitation of the IOMI framework: given an unbounded input space, for any “good” learning algorithm there are always scenarios in which $\text{IOMI}_{\mathcal{D}}(\mathcal{A}_n)$ is unbounded. Therefore, the distribution-free learnability of VC classes cannot be expressed using the IOMI framework.

In this paper, we focus on the “*conditional mutual information*” (CMI) framework, proposed by Steinke and Zakynthinou [1]. In order to reason about the generalization error of a learning algorithm, they introduce a super sample that contains the training sample as a random subset and compute the mutual information between the input and output conditional on the super sample (formal definitions are provided in Section 2.1). Improvements of this framework and its application in studying the generalization of specific learning algorithms have been studied in [20–25].

The current paper revolves around the following fundamental question: For which learning problems and learning algorithms is the CMI framework expressive enough to accurately estimate the generalization error? We will focus in particular on whether we can recover optimal worst case (minimax) rates for VC classes satisfying certain properties. The answer to these question provide evidence that the CMI framework provides a unifying framework for studying generalization.

For VC classes, Steinke and Zakynthinou [1] revealed a stark separation between the CMI framework and IOMI framework. They showed the existence of an empirical risk minimization (ERM) algorithm whose CMI is no larger than $d \log n + 2$ for learning every VC class of dimension d given n i.i.d. training samples. In contrast to Theorem 1.1, CMI does not scale with the cardinality of the space. However, the bound on the CMI combined with Steinke and Zakynthinou’s CMI-based generalization bound, leads to a bound on the expected excess risk that is suboptimal in some cases by a $\log n$ factor. (For an overview of the known bounds for learning VC classes, please refer to Appendix A.) The suboptimality of their bound prompted Steinke and Zakynthinou [3] to conjecture that the CMI bound for proper learners of VC classes can be improved to $O(d)$. Moreover, Steinke and Zakynthinou connected CMI framework to the sample compression framework of [26] by showing that a sample compression \mathcal{A}_n of size k has $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq k \log 2n$. Their bound for sample compression schemes is also suboptimal in some cases by a $\log n$ factor.

1.1 Contributions

In this paper we extend the reach of the CMI framework by demonstrating its unifying nature for obtaining optimal or near-optimal bounds for the expected excess risk of the various algorithms in the realizable setting.

³This concept class can be defined as follows. Let $\mathcal{X} = \{1, \dots, m\}$. Let $k \in \mathbb{N}$ and $h_k : \mathcal{X} \rightarrow \{0, 1\}$ define as $h_k(x) = \mathbb{1}[x > k]$. Then, the class of one-dimensional thresholds over $\{1, \dots, m\}$ is $\mathcal{H}_m = \{h_k | k \in \mathbb{N}\}$.

1. We demonstrate that one can use the CMI framework to express non-trivial (but sub-optimal) bounds for every improper learning algorithm that outputs a hypothesis from a class with a bounded VC dimension. This is achieved by an empirical variant of CMI defined by [1].
2. We study the CMI of SVMs for learning half spaces and show that the CMI framework yields optimal bounds on the expected excess risk. Our bound on the CMI of SVM is an application of our general result giving optimal CMI bounds for stable sample compression schemes [2, 27], which improve on CMI bounds for general sample compression schemes [1] by a $\log n$ factor.
3. In the context of proper learning of VC classes, we exhibit VC classes for which the CMI of any proper learner cannot be bounded by any real-valued function of the VC dimension. Then, we consider VC classes with finite star number [4], and prove the existence of a learner with bounded CMI. Finally, we show that the release of the set of all consistent classifiers in \mathcal{H} has bounded CMI *if and only if* \mathcal{H} has finite star number.
4. We show that CMI framework is *universal* in the realizable setting. More precisely, for every data distribution and consistent learner, the bound on excess risk obtained by the CMI framework vanishes if and only if the excess risk also vanishes as the number training samples diverges. We then show that any learning algorithm with a “leave-one-out” bound of order $O(1/n)$ yields an evaluated-CMI bound of order $O(\log n)$. As an application, we study the classical one-inclusion graph algorithm of Haussler et al. [5] for improper learning of VC classes, and provide a nearly optimal bound on its expected excess risk using the CMI framework. We also prove there exists a randomized one-inclusion graph which learns point functions (singleton) with bounded CMI.

Our results indicate that CMI is a very expressive generalization framework, and one that can tie together existing frameworks. Although most of our results are stated for binary classification in the distribution-free setting, it is interesting to note that the CMI framework is known to provide numerically non-vacuous generalization error guarantees for some modern deep learning models and datasets in the distribution-dependent setting [21, 25]. These developments in a range of different problem settings highlight the importance of understanding the expressiveness of the CMI framework.

2 Preliminaries

We consider the problem of binary classification, with inputs in some space \mathcal{X} assigned labels in $\mathcal{Y} = \{0, 1\}$. A concept (or hypothesis) class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ is a set of functions $h : \mathcal{X} \rightarrow \mathcal{Y}$. We say \mathcal{H} *shatters* $(x_1, \dots, x_m) \in \mathcal{X}^m$ if for all $(y_1, \dots, y_m) \in \{0, 1\}^m$, there exists $h \in \mathcal{H}$, such that, for all $i \in [m]$, we have $h(x_i) = y_i$. The *VC dimension* of \mathcal{H} , denoted by d , is the largest $m \in \mathbb{N}$ for which there exists $(x_1, \dots, x_m) \in \mathcal{X}^m$ shattered by \mathcal{H} . If no such finite m exists, then $d = \infty$.

Let \mathcal{D} be a distribution on $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. The *empirical (classification) risk* of a classifier $h : \mathcal{X} \rightarrow \mathcal{Y}$ on a sample $s = ((x_1, y_1), \dots, (x_n, y_n)) \in \mathcal{Z}^n$ is $\hat{R}_s(h) = n^{-1} \sum_{i \in [n]} \ell(h, (x_i, y_i))$, where $\ell(h, (x, y)) = \mathbb{1}[h(x) \neq y]$. Let $S_n \sim \mathcal{D}^n$, i.e., let S_n be a sequence of i.i.d. random elements in \mathcal{Z} with common distribution \mathcal{D} . (We can view S_n itself as a random element in \mathcal{Z}^n .) The *risk* of h is $R_{\mathcal{D}}(h) = \mathbb{E} \hat{R}_{S_n}(h)$, where \mathbb{E} denotes the expectation operator. (The risk has, of course, no dependence on n due to the data being i.i.d.)

A distribution \mathcal{D} is *realizable by a class* $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ if there exists $h \in \mathcal{H}$ such that $R_{\mathcal{D}}(h) = 0$. A sequence $((x_1, y_1), \dots, (x_n, y_n))$ is said to be *realizable by* \mathcal{H} , if for some $h \in \mathcal{H}$, $h(x_i) = y_i$ for all $i \in [n] = \{1, \dots, n\}$. Note that if a distribution is realizable by \mathcal{H} , it implies that with probability one over $S_n \sim \mathcal{D}^n$, the training sample S_n is realizable by \mathcal{H} .

Let $\mathcal{A} = (\mathcal{A}_n)_{n \geq 1}$ denote a (potentially randomized) learning algorithm, which, for any positive integer n , maps S_n to an element of $\mathcal{X} \rightarrow \mathcal{Y}$. We say that \mathcal{A} is a *proper learner for a class* $\mathcal{H} \subseteq \mathcal{X} \rightarrow \mathcal{Y}$ if the codomain of \mathcal{A}_n is a subset of \mathcal{H} for every n . We say \mathcal{A}_n is a consistent algorithm (learner) if $\hat{R}_{S_n}(\mathcal{A}_n(S_n)) = 0$ a.s. Our primary interest in this paper is the *expected generalization error* of \mathcal{A}_n with respect to \mathcal{D} , defined as $\text{EGE}_{\mathcal{D}}(\mathcal{A}_n) = \mathbb{E}[R_{\mathcal{D}}(\mathcal{A}_n(S_n)) - \hat{R}_{S_n}(\mathcal{A}_n(S_n))]$, where we average over both the choice of training sample and the randomness within the algorithm \mathcal{A}_n .

2.1 Conditional mutual information (CMI) of an algorithm

In order to study generalization, and avoid some of the pitfalls of earlier approaches based on mutual information, Steinke and Zakynthinou [1] propose to study the information contained in a

“supersample” Z , a training sample S_n taken from the supersample, and the hypothesis $\mathcal{A}_n(S_n)$ output by a possibly randomized learning algorithm, given S_n as input. Formally, let $Z = (Z_{i,j})_{i \in \{0,1\}, j \in [n]}$ to be an array of i.i.d. random elements in the space \mathcal{Z} of labeled examples, with a common distribution \mathcal{D} . In order to choose a training sample S_n of size n from Z , let $U = (U_1, U_2, \dots, U_n)$ be a sequence of i.i.d. Bernoulli random variables in $\{0, 1\}$, independent from Z , with $\mathbb{P}(U_i = 0) = 1/2$. Define $S_n = Z_U = (Z_{U_j, j})_{j=1}^n$. The *conditional mutual information (CMI)* of \mathcal{A}_n , denoted $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n)$, is defined to be the mutual information between $\mathcal{A}_n(S_n)$ and U given Z , denoted $I(\mathcal{A}_n(S_n); U|Z)$. This quantity is equivalent to $I(\mathcal{A}_n(S_n); S_n|Z)$ when \mathcal{D} is atomless, since (U_1, \dots, U_n) is a.s. measurable with respect to S_n and Z . Because Z and U are independent, $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq H(U|Z) = H(U) = n \log 2$. We now pause to introduce this and other information-theoretic quantities formally.

2.2 Measures of divergence and information

Let P, Q be probability measures on a measurable space. (We ignore measure-theoretic pathologies for clarity.) For a P -integrable or nonnegative function f , let $P[f] = \int f dP$. When Q is absolutely continuous with respect to P , denoted $Q \ll P$, write $\frac{dQ}{dP}$ for (an arbitrary version of) the Radon–Nikodym derivative (or density) of Q with respect to P . The *KL divergence* (or *relative entropy*) of Q with respect to P , denoted $\text{KL}(Q \| P)$, is defined as $Q[\log \frac{dQ}{dP}]$ when $Q \ll P$ and infinity otherwise.

For a random element X in some measurable space \mathcal{X} , let $\mathbb{P}[X]$ denote its distribution, which lives in the space $\mathcal{M}_1(\mathcal{X})$ of all probability measures on \mathcal{X} . Given another random element, say Y in \mathcal{T} , let $\mathbb{P}^Y[X]$ denote the conditional distribution of X given Y . If X and Y are independent, $\mathbb{P}^Y[X] = \mathbb{P}[X]$ a.s. For an event, say $X \in A$, $\mathbb{P}^Y[X \in A]$ denotes the event’s conditional probability given Y , which is defined to be the conditional expectation of the indicator random variable $\mathbb{1}[X \in A]$ given Y , denoted $\mathbb{E}^Y \mathbb{1}[X \in A]$.⁴ By the chain (aka tower) rule, $\mathbb{E}\mathbb{E}^{\mathcal{F}} = \mathbb{E}$ for any σ -algebra \mathcal{F} .

The *mutual information between X and Y* is $I(X; Y) = \text{KL}(\mathbb{P}[(X, Y)] \| \mathbb{P}[X] \otimes \mathbb{P}[Y])$, where \otimes forms the product measure. Writing $\mathbb{P}^Z[(X, Y)]$ for the conditional distribution of the pair (X, Y) given a random element Z , the *disintegrated mutual information between X and Y given Z* , is

$$I^Z(X; Y) = \text{KL}(\mathbb{P}^Z[(X, Y)] \| \mathbb{P}^Z[X] \otimes \mathbb{P}^Z[Y]).$$

Then the *conditional mutual information* of X and Y given Z is $I(X, Y|Z) = \mathbb{E}I^Z(X, Y)$.

Let $\mu = \mathbb{P}[X]$ and let $\kappa(Y) = \mathbb{P}^Y[X]$ a.s. If X concentrates on a countable set V with counting measure ν , the (*Shannon*) *entropy of X* is $H(X) = -\mu[\log \frac{d\mu}{d\nu}] = -\sum_{x \in V} \mathbb{P}(X = x) \log \mathbb{P}(X = x)$. The *disintegrated entropy of X given Y* is defined by $H^Y(X) = -\kappa(Y)[\log \frac{d\kappa(Y)}{d\nu}]$, while the *conditional entropy of X given Y* is $H(X|Y) = \mathbb{E}[H^Y(X)]$. Note that $H(X|Y) \leq H(X)$. We will make use of the following lemma whose proof can be found in [28].

Lemma 2.1. *Let (X_1, X_2, \dots, X_n) be a discrete random vector, and Y be an arbitrary random variable. Then, $H(X_1, \dots, X_n|Y) \geq \sum_{i=1}^n H(X_i|X_{-i}, Y)$, where $X_{-i} = (X_j : j \in [n], j \neq i)$.*

Steinke and Zakyntinou establish a range of generalization bounds in terms of CMI. Our primary interest is in bounds for algorithms that have vanishing empirical risk. For $[0, 1]$ -bounded loss, Steinke and Zakyntinou show that

$$\mathbb{E}R_{\mathcal{D}}(\mathcal{A}_n(S_n)) \leq 2\mathbb{E}\hat{R}_S(\mathcal{A}_n(S_n)) + \frac{3\text{CMI}_{\mathcal{D}}(\mathcal{A}_n)}{n}. \quad (1)$$

For consistent learners (i.e., those that achieve zero empirical error a.s.), they also establish

$$\mathbb{E}R_{\mathcal{D}}(\mathcal{A}_n(S_n)) \leq \frac{\text{CMI}_{\mathcal{D}}(\mathcal{A}_n)}{n \log 2}. \quad (2)$$

Steinke and Zakyntinou also introduce a variant of CMI based on the information revealed by the learner’s losses on Z , rather than by the output hypothesis, $\mathcal{A}_n(S_n)$, directly.

Definition 2.2 (Evaluated CMI, [1, §6.2.2]). Let $L \in \{0, 1\}^{2 \times n}$ be the array with entries $L_{i,j} = \ell(\mathcal{A}_n(S_n), Z_{i,j})$ for $i \in \{0, 1\}, j \in [n]$. The *evaluated conditional mutual information of \mathcal{A}_n with respect to \mathcal{D}* , denoted by $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))$, is the conditional mutual information $I(L; U|Z)$.

⁴By definition, $\mathbb{P}^Y[X]$ is a $\sigma(Y)$ -measurable random element in $\mathcal{M}_1(\mathcal{X})$, i.e., $\mathbb{P}^Z[U] = \kappa(Z)$ a.s. for some measurable map $\kappa : \mathcal{T} \rightarrow \mathcal{M}_1(\mathcal{X})$. More generally, if, say $\mathcal{F} = \sigma(Y, Z)$ is the σ -algebra generated by Y and Z , then a conditional distribution/probability/expectation given \mathcal{F} is a measurable function of Y and Z .

By the data processing inequality, $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) \leq \text{CMI}_{\mathcal{D}}(\mathcal{A}_n)$. Therefore, $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))$ is also bounded above by $n \log 2$. For consistent learners, Steinke and Zakynthinou show

$$\mathbb{E}R_{\mathcal{D}}(\mathcal{A}_n(S_n)) \leq 1.5 \frac{\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))}{n}. \quad (3)$$

For consistent learners \mathcal{A}_n with bounded CMI or eCMI, these results imply their expected excess risk is of order $O(1/n)$. The following result gives a nearly optimal bound for the generalization error for VC classes in term of the evaluated CMI. The proof (Appendix B) uses standard arguments, controlling the cardinality of the support of L using the Sauer–Shelah lemma.

Theorem 2.3. *For every n , let $\mathcal{A}_n : \mathcal{Z}^n \rightarrow \mathcal{H}_n$, where \mathcal{H}_n is a concept class with VC dimension d_n . Then, for every n and distribution on Z , $I^Z(L; U) \leq d_n \log 6n$ a.s. In particular, $\sup_{\mathcal{D}} \text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) \in O(d_n \log n)$.*

Remark 2.4. Markov’s inequality and Eq. (2) imply $\mathbb{P}(R_{\mathcal{D}}(\mathcal{A}_n(S_n)) \geq \epsilon) \leq \text{CMI}_{\mathcal{D}}(\mathcal{A}_n)/(\log(2)n\epsilon)$ for consistent learners. By [21, Thm. 2.1], $I(\mathcal{A}_n(S_n); U|Z) \leq I(\mathcal{A}_n(S_n); S_n)$. This observation, combined with [12, Prop. 11], implies there is an input space, data distribution, and consistent learning algorithm for which this tail bound’s dependence on n is *tight*. If one were to obtain sample complexity bounds via such tail bounds, one would only prove that $O(1/(\epsilon\delta))$ samples suffice to find a hypothesis with ϵ estimation error with probability at least $1 - \delta$. The linear dependence on $1/\delta$ is, however, suboptimal. As such, it seems that the CMI framework cannot be used to obtain optimal sample complexity bounds in the PAC framework. Recent proposals for disintegrated notions of CMI in [22] might provide a framework for studying the sample complexity of PAC learning using an information-theoretic framework. \triangleleft

3 Optimal CMI Bound for SVM and Stable Compression Schemes

In this section, we show that the CMI framework can be used to derive an optimal excess risk bound for the SVM algorithm learning half spaces in \mathbb{R}^d . To show this, we establish optimal CMI bounds for the subclass of *stable* sample compression schemes, which imply this section’s main result:

Theorem 3.1. *Let \mathcal{A}_n be the SVM algorithm for learning the class of half spaces in \mathbb{R}^d . Then, for every $n > d/2$ and realizable distribution \mathcal{D} in \mathbb{R}^d , we have $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq 2(d+1) \log 2$.*

Combining this result with Eq. (2) gives $\mathbb{E}R_{\mathcal{D}}(\mathcal{A}_n(S_n)) \leq 2(d+1)/n$. The lower bound for expected excess risk of linear classifiers in [29] shows this bound is optimal up to a constant factor.

3.1 CMI of Stable Compression Schemes

Littlestone and Warmuth [26] introduced compression schemes, which capture the idea that a consistent hypothesis can be defined in terms of a fixed number of samples. Formally, for a concept class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, a *sample compression scheme* of size $k \in \mathbb{N}$ is a pair (κ, ρ) of maps such that, for all samples $s = ((x_i, y_i))_{i=1}^n$ of size $n \geq k$, the map κ compresses the sample into a length- k subsequence $\kappa(s) \subseteq s$ which the map ρ uses to reconstruct an empirical risk minimizer $\hat{h} = \rho(\kappa(s))$. Steinke and Zakynthinou prove the following upper bound on the CMI of a sample compression scheme.

Theorem 3.2 ([1, Thm. 4.1]). *Let \mathcal{H} be a hypothesis class that has a sample compression scheme (κ, ρ) of size k . Then, $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq k \log(2n)$ where $\mathcal{A}_n(\cdot) = \rho(\kappa(\cdot))$.*

Note that the bound in Theorem 3.2 *cannot* be improved from $O(k \log n)$ to $O(k)$ for *every* sample compression scheme, and so the bound in Theorem 3.2 is tight, and cannot be improved without further information about the compression scheme. The proof of the optimality stems from the fact that there exists compression schemes of size k and data distributions \mathcal{D} such that there is a lower bound $\mathbb{E}[R_{\mathcal{D}}(\mathcal{A}_n)] = \Omega(k \log(n)/n)$ where $\mathcal{A}_n(\cdot) = \rho(\kappa(\cdot))$ [30, 31]. Combining this lower bound with Eq. (2) proves the optimality of Theorem 3.2.

Nevertheless, we can circumvent this lower bound by considering an important subclass of the sample compression schemes. Many natural compression schemes are also *stable* in the sense that removing any training example that was not in the compressed sequence does not alter the resulting classifier. To give a formal definition, we write $s \subseteq s'$ for two sequences s, s' if, under some permutation, s is a subsequence of s' .

Definition 3.3 (Stable sample compression scheme; [2]). A sample compression scheme (κ, ρ) of size k is said to be *stable* if κ is symmetric (i.e., invariant to permutation of its input) and, for every realizable sample s of size $n \geq k$, and every sequence s' such that $\kappa(s) \subseteq s' \subseteq s$, we have $\rho(\kappa(s)) = \rho(\kappa(s'))$. Due to the symmetry of κ , we refer to its output as the *compression set*, although the equivalence class of sequence under permutations is the structure of a multiset, not a set.

The concept of a stable compression scheme has its roots in the analysis of the SVM for learning half-spaces in \mathbb{R}^d [32], which is the quintessential example of a stable sample compression scheme. For SVMs, the compression (multi)set contains at most $d + 1$ distinct “support vectors” for any given training set. The reconstruction map outputs the max-margin classifier over the set of support vectors. By stability, removing any training example that is not a support vector does not change the resulting classifier [33, Sec. 5.3.2]. In the next theorem, we present a uniform CMI bound over realizable distributions for every stable sample compression scheme. Our bound removes the $\log n$ factor from Theorem 3.2 and is *optimal* up to a constant factor in the distribution-free setting.

Theorem 3.4. *Let \mathcal{H} be a concept class with a stable compression scheme (κ, ρ) of size k . Then, for every realizable data distribution \mathcal{D} and $n \geq k$, $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq 2k \log 2$, where $\mathcal{A}_n = \rho(\kappa(\cdot))$.*

Remark 3.5. Steinke and Zakynthinou [1, Sec. 4.4] propose an algorithm for learning threshold functions (positive rays) in the realizable setting over \mathbb{R} that achieves $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq 2 \log 2$. It is interesting to note that their algorithm can be viewed as a stable compression scheme. Specifically, for a realizable training set s , let $x^* = \min\{x \in \mathbb{R} : (x, 1) \in s\}$ if s has any sample with label 1, otherwise let $x^* = \infty$. Then the algorithm proposed by Steinke and Zakynthinou is $\mathcal{A}_n(S_n) = \hat{h}$, where $\hat{h}(x) = \mathbb{1}[x \geq x^*]$. Steinke and Zakynthinou present a bespoke analysis of this special algorithm. It is straightforward to see that the algorithm is a stable compression scheme of size one and the compression map here is symmetric. Therefore, the result of Theorem 3.4 gives $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq 2 \log 2$. \triangleleft

Proof of Theorem 3.4. Let $W = \mathcal{A}_n(Z_U) = \rho(\kappa(Z_U))$ and note that \mathcal{A}_n is deterministic. We have $\text{H}(U|Z) = n \log 2$ due to independence of U and Z and the independence of components of U . Then, by the definition of mutual information in terms of entropy, and Lemma 2.1,

$$\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) = \text{H}(U|Z) - \text{H}(U|W, Z) \leq n \log 2 - \sum_{i=1}^n \text{H}(U_i|W, Z, U_{-i}). \quad (4)$$

Fix $i \in [n]$, and define $U_{i \rightarrow b} \triangleq (U_1, \dots, U_{i-1}, b, U_{i+1}, \dots, U_n)$ for $b \in \{0, 1\}$. Using this notation, we can define two training sets $S_{i \rightarrow b} = Z_{U_{i \rightarrow b}}$ for $b \in \{0, 1\}$. Let \mathcal{F}_i be the σ -algebra $\sigma(W, Z, U_{-i})$ and let E be the event $\rho(\kappa(S_{i \rightarrow 0})) = \rho(\kappa(S_{i \rightarrow 1}))$. Then, by the non-negativity of entropy,

$$\text{H}(U_i|W, Z, U_{-i}) = \mathbb{E}[\text{H}^{\mathcal{F}_i}(U_i)] \geq \mathbb{E}[\text{H}^{\mathcal{F}_i}(U_i)\mathbb{1}[E]]. \quad (5)$$

Note that, conditional on the sub- σ -algebra $\mathcal{G}_i = \sigma(Z, U_{-i})$, W takes on at most two values. However, on the event E (or equivalently, conditioning further on the event E , since E is \mathcal{G}_i -measurable), W is now nonrandom because it takes on a single value. It follows that, conditional on \mathcal{G}_i and the event E , W is trivially independent of every random variable, including U_i . Ergo, on E , $\mathbb{E}^{\mathcal{G}_i}[U_i] = \mathbb{E}^{\mathcal{F}_i}[U_i] = \mathbb{P}^{\mathcal{F}_i}[U_i = 1]$. But U_i is independent of \mathcal{G}_i , and so $\mathbb{E}^{\mathcal{G}_i}[U_i] = \mathbb{E}[U_i] = \frac{1}{2}$. Thus, on E , $\mathbb{P}^{\mathcal{F}_i}[U_i = 1] = \frac{1}{2}$ and so $\text{H}^{\mathcal{F}_i}(U_i) = \text{H}_b(\frac{1}{2}) = \log 2$. Therefore,

$$\text{H}(U_i|W, Z, U_{-i}) \geq \log 2 \cdot \mathbb{P}(\rho(\kappa(S_{i \rightarrow 0})) = \rho(\kappa(S_{i \rightarrow 1}))). \quad (6)$$

We can bound the probability of E from below using the stability property of the compression scheme. For any $(x_1, \tilde{x}_1, x_2, \dots, x_n) \in \mathcal{X}^{n+1}$ and $h \in \mathcal{H}$, consider two multisets $S = \{(x_1, h(x_1)), (x_2, h(x_2)), \dots, (x_n, h(x_n))\}$ and $\tilde{S} = \{(\tilde{x}_1, h(\tilde{x}_1)), (x_2, h(x_2)), \dots, (x_n, h(x_n))\}$, where S and \tilde{S} differ only in the first element. Define the multiset $S \cup \tilde{S} = \{(x_1, h(x_1)), (\tilde{x}_1, h(\tilde{x}_1)), (x_2, h(x_2)), \dots, (x_n, h(x_n))\}$. We claim that if $(x_1, h(x_1))$ and $(\tilde{x}_1, h(\tilde{x}_1))$ are not the members of the compression set $S \cup \tilde{S}$, then $(x_1, h(x_1))$ and $(\tilde{x}_1, h(\tilde{x}_1))$ are not in the compression set of S and \tilde{S} , respectively. To prove this claim, since $(x_1, h(x_1))$ is not in the compression set $S \cup \tilde{S}$ by the stability of κ , we have $\rho(\kappa(S \cup \tilde{S})) = \rho(\kappa(S \cup \tilde{S} \setminus \{(x_1, h(x_1))\}))$. By the definition of S and \tilde{S} , $S \cup \tilde{S} \setminus \{(x_1, h(x_1))\} = \tilde{S}$. Thus, combining facts that $(x_1, h(x_1))$ is not in the compression set $S \cup \tilde{S}$ and $\kappa(S \cup \tilde{S}) = \kappa(S \cup \tilde{S} \setminus \{(x_1, h(x_1))\}) = \kappa(\tilde{S})$, we obtain $(\tilde{x}_1, h(\tilde{x}_1))$ is not in the

compression set \tilde{S} . Similarly, we can prove $(x_1, h(x_1))$ is not a member of the compression set S by switching x_1 with \tilde{x}_1 in the argument. By this argument,

$$\mathbb{P}(\rho(\kappa(S_{i \rightarrow 0})) = \rho(\kappa(S_{i \rightarrow 1}))) \geq \mathbb{P}(Z_{0,i} \notin \kappa(S_{i \rightarrow 0} \cup S_{i \rightarrow 1}) \wedge Z_{1,i} \notin \kappa(S_{i \rightarrow 0} \cup S_{i \rightarrow 1})). \quad (7)$$

Recall that the elements of Z are i.i.d., hence exchangeable. Since the size of the sample compression is k and κ is symmetric, we have

$$\mathbb{P}(Z_{0,i} \notin \kappa(S_{i \rightarrow 0} \cup S_{i \rightarrow 1}) \wedge Z_{1,i} \notin \kappa(S_{i \rightarrow 0} \cup S_{i \rightarrow 1})) \geq \binom{n-1}{k} / \binom{n+1}{k}. \quad (8)$$

Combining Eqs. (6) to (8) yields $H(U_i|W, Z, U_{-i}) \geq \log 2 \cdot \binom{n-1}{k} / \binom{n+1}{k} \geq (1 - 2k/n) \log 2$. Finally, the result follows by substitution of this bound into Eq. (4). \square

The result for SVMs (Theorem 3.1) follows immediately from Theorem 3.4 and the fact that the SVM may be expressed as a stable compression scheme of size $d + 1$.

4 CMI of Proper Learning of VC classes

Following their paper introducing CMI, Steinke and Zakynthinou posed several open problems asking whether VC classes under realizability admit learners with bounded CMI. We will restate their conjectures, and then showing that there exist some VC classes for which it is not possible to find a proper learner with bounded CMI under realizability. We then consider a subset of VC classes, namely VC classes with finite star number, and show that for such concept classes, there exists an ERM with bounded CMI.

We first state the main result of [1] on the CMI of proper learners.

Theorem 4.1 (Thm. 4.12, [1]). *Let \mathcal{H} be a concept class with VC dimension d . Then for all $n \in \mathbb{N}$, there exists a proper ERM algorithm \mathcal{A}_n for learning \mathcal{H} such that for every realizable distribution \mathcal{D} , $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) = O(d \log n)$.*

Remark 4.2 (Comparison of Theorem 4.1 and Theorem 2.3). First, note that Theorem 4.1 does not hold for every ERM algorithm. As discussed in [1], we can construct pathological ERMs with nearly maximal CMI by simply encoding the information U into the “lower-order” bits of W .

It is also worth noting that our result in Theorem 2.3 is more general. There we show that a bound $O(d \log n)$ holds for *evaluated* CMI of *any* algorithm that outputs a hypothesis from VC class, whereas Theorem 4.1 holds for a *specific* proper algorithm. \triangleleft

4.1 A Limitation of Proper Learning

Steinke and Zakynthinou [3] propose two conjectures regarding CMI for proper learning of VC classes under the realizability assumption, both of which can be seen as special cases of the following statement:

Statement 1. *There exists a real-valued function f and constant $c \geq 0$ such that, for every nonnegative integer d and VC class $\mathcal{H} \subseteq \mathcal{X} \rightarrow \mathcal{Y}$ of dimension d , there exists a proper learning algorithm \mathcal{A} for \mathcal{H} such that, for every $n \geq d$, $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq f(d)$ for all \mathcal{D} and, for every realizable $s \in \mathcal{Z}^n$, $\mathbb{E} \hat{R}_s(\mathcal{A}_n(s)) \leq c d/n$, where the expectation is taken only over the randomness in \mathcal{A}_n .*

Steinke and Zakynthinou [3] conjecture that Statement 1 holds for f linear. In this section, we show that Statement 1 is false in general: it is not possible to find a proper learning algorithm for *every* VC class that removes the $\log(n)$ factor from Theorem 4.1. For a class $\mathcal{H} \subseteq \mathcal{X} \rightarrow \mathcal{Y}$, let $\mathcal{M}_{\text{prop}}^{\mathcal{H}}(\epsilon, \delta)$ denote the *proper optimal sample complexity* of (ϵ, δ) -PAC learning \mathcal{H} , i.e., $\mathcal{M}_{\text{prop}}^{\mathcal{H}}(\epsilon, \delta)$ is the least integer n , for which there exists a proper learning algorithm \mathcal{A} such that, for every realizable distribution \mathcal{D} , $\mathbb{P}(R_{\mathcal{D}}(\mathcal{A}_n(S_n)) \geq \epsilon) \leq \delta$. The following result provides a lower-bound on the sample complexity of proper learning:

Theorem 4.3 (Thm. 11, [2]). *Let $\epsilon \in (0, 1/8)$ and $\delta \in (0, 1/100)$. There exists a concept class with VC dimension d for which we have $\mathcal{M}_{\text{prop}}^{\mathcal{H}}(\epsilon, \delta) \geq \frac{\tilde{c}}{\epsilon} (d \text{Log} \frac{1}{\epsilon} + \text{Log} \frac{1}{\delta})$ for a fixed numerical constant $\tilde{c} > 0$, where $\text{Log}(x) = \max\{1, \log(x)\}$ for $x \geq 0$.*

We now present the main result: for VC classes, we show that the existence of a learning algorithm with bounded CMI contradicts the lower bound on the sample complexity in Theorem 4.3. The proof can be found in Appendix C.

Theorem 4.4. *Statement 1 is false.*

Remark 4.5. Consider a modified Statement 1, seeking a proper learner with bounded eCMI instead. We can show that this modified statement is also false. \triangleleft

4.2 VC Classes with Finite Star Number

Theorem 4.4 states that it is not possible to find a proper learning algorithm with bounded CMI for *every* VC class. Note that this limitation does not imply a failure of the CMI framework for characterizing the expected excess risk of learning VC classes. Instead, the impossibility can be attributed to an inherent limitation of proper learning algorithms, since there exist VC classes such that no proper learning algorithm \mathcal{A}_n satisfies $\mathbb{E}[R_{\mathcal{D}}(\mathcal{A}_n)] = O(1/n)$ [2]. In this section, we consider a family of VC classes for which we *can* show the existence of a learner with bounded CMI. We begin with some definitions.

Two sequences $((x_1, y_1), \dots, (x_n, y_n))$ and $((x'_1, y'_1), \dots, (x'_n, y'_n))$ are *neighbours* if $x_i = x'_i$ for all $i \in [n]$, and $y_i = y'_i$ for all but exactly one $i \in [n]$. Fix any concept class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$. *Star number* of \mathcal{H} [4, Def. 2], denoted by \mathfrak{s} , is the largest integer n such that there exists a realizable $s \in (\mathcal{X} \times \mathcal{Y})^n$, and every neighbour of s is realizable by \mathcal{H} . If no such largest integer n exists, then $\mathfrak{s} = \infty$. Hanneke and Yang [4, Sec. 4.1] calculate the star number of some common concept classes. It is straightforward to see that $d \leq \mathfrak{s}$. For any $n \in \mathbb{N}$, and $s = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathcal{X} \times \mathcal{Y})^n$, define a *version space* of s with respect to \mathcal{H} as $V_{\mathcal{H}}[s] = \{h \in \mathcal{H} : \hat{R}_s(h) = 0\}$, a set of classifiers that are consistent with s .

4.2.1 Star Number, Version Space, and CMI

Fix any concept class \mathcal{H} , and assume that, after observing a training sample S_n , we want to output the version space $V_{\mathcal{H}}[S_n]$, i.e., the set of all classifiers consistent with S_n . We are interested in the following question: for which concept classes does the version space carry little information about the training samples conditioned on the supersample? More precisely, for which classes is $I(V_{\mathcal{H}}[S_n]; U|Z) = O(1)$? Note that bounding the ‘‘CMI’’ of the version space provides a bound on the CMI of a broad class of algorithms that choose a particular ERM based solely on the version space, potentially under further constraints, such as privacy, fairness, etc.

In this section, we give a complete characterization of when $I(V_{\mathcal{H}}[S_n]; U|Z) = O(1)$, and show that it is possible if and only if \mathcal{H} has finite star number. In particular, given a class with infinite star number, we demonstrate that $I(V_{\mathcal{H}}[S_n]; U|Z) = \Omega(n)$. We begin with an upper bound, whose proof can be found in Appendix D.

Theorem 4.6. *Let $n \in \mathbb{N}$, \mathcal{H} be a concept class with star number \mathfrak{s} , and \mathcal{D} be a realizable distribution. Let Z , U , and S_n be as defined in the beginning of this section. Then for every $n \geq \mathfrak{s}$, we have $I(V_{\mathcal{H}}[S_n]; U|Z) \leq 2\mathfrak{s} \log 2$.*

We can use the data processing inequality and Theorem 4.6 to obtain the following:

Corollary 4.7. *Let \mathcal{H} be a concept class with the star number \mathfrak{s} . Consider any ERM algorithm \mathcal{A}_n for which the Markov chain $S_n - V_{\mathcal{H}}[S_n] - \mathcal{A}_n(S_n)$ holds; in other words, the output of the algorithm and the training set are conditionally independent given the version space. Then, for any such an algorithm, for every $n \geq \mathfrak{s}$, and every realizable distribution \mathcal{D} , we have $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq 2\mathfrak{s} \log 2$.*

In Corollary 4.7, by assuming the Markov structure $S_n - V_{\mathcal{H}}[S_n] - \mathcal{A}_n(S_n)$ we restrict the information of the ERM algorithm $\mathcal{A}_n(S_n)$. One might try to extend our result in Corollary 4.7 such that it holds for *any* ERM without any constraints. However, for the class of one-dimensional threshold over \mathbb{R} , whose star number is two, one can construct an ERM with maximal CMI [1, Sec. 4.3]. Therefore, the Markov chain assumption cannot be removed. The next theorem shows $\mathfrak{s} < \infty$ is a necessary condition, for otherwise, there exist learning scenarios under which we cannot output the version space, even with merely sublinear CMI.

Theorem 4.8. *For every $n \in \mathbb{N}$, $n \geq 2$ and for every concept class \mathcal{H} with star number \mathfrak{s} with $\mathfrak{s} \geq 2$ over input space \mathcal{X} , there exists a realizable data distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$ such that $I(V_{\mathcal{H}}[S_n]; U|Z) = \Omega(\min\{\mathfrak{s}, n\})$.*

Proof sketch: Let $\mathcal{X} = [n]$ and consider the concept class $\mathcal{H} = \{h_0, h_1, \dots, h_n : \mathcal{X} \rightarrow \mathcal{Y}\}$, where $h_0(x) = 0$ is the zero function and $h_t(x) = \mathbb{1}[x = t]$, for $t \in [n]$, are point functions. It is easy to

see that this concept class has star number n on \mathcal{X} . Let \mathcal{D} correspond to the uniform distribution on \mathcal{X} and target function h_0 . Consider the bijection between \mathcal{H} and $\{0, 1, \dots, n\} \supseteq \mathcal{X}$. For every training sequence, the version space contains 0 and every point in \mathcal{X} not observed in $S_n = Z_U$. The key observation is that, in each column of Z , one point was *not* selected for training, and so each column contains zero or one points in the version space. Whenever there is one point, the value of U_i is revealed for that column. We show that the number of columns with this property is a lower bound on $I(\mathcal{V}_{\mathcal{H}}[S_n]; U|Z)$. A coupon collector's argument yields a lower bound the number of such columns. The formal proof can be found in Appendix E. \square

4.2.2 An ERM whose CMI is logarithmic in star number

In the next theorem, we show that there exists an ERM for learning VC classes with a finite star number for which the CMI is upper bounded by a constant and its dependence on star number is logarithmic. The proof is provided in Appendix F.

Theorem 4.9. *Let \mathcal{H} be a concept class with VC dimension d and star number s . Then, there exists an ERM \mathcal{A}_n for learning \mathcal{H} such that for every $n \geq s$ and for every realizable distribution \mathcal{D} , we have $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) = O(d \log(s/d))$.*

Note that Theorem 4.9 shows the existence of a specific ERM with constant CMI, whereas in Corollary 4.7 we show a broad class of ERMs has bounded CMI.

5 Universality of eCMI and Improper Learning of VC Classes

The eCMI, introduced in Definition 2.2, is an appropriate information-theoretic notion for analyzing learning algorithms when there is no natural parameterization of the set of possible predictors, such as for improper or transductive algorithms. In this section, we show that eCMI is *universal* in the realizable setting. Then, we show that the CMI framework can be used to obtain a near-optimal bound on the expected excess risk of any algorithm with a leave-one-out error guarantee. As an application, we study CMI of the classical *one-inclusion graph prediction* algorithm, which was first proposed by Haussler et al. [5] as an optimal improper learner for VC classes. The next theorem is the main result of this section, whose proof can be found in Appendix G.

Theorem 5.1. *Let $n \geq 2 \in \mathbb{N}$, let \mathcal{A}_n be a learning algorithm, and let \mathcal{D} be a distribution on \mathcal{Z} . Assume with probability one $\hat{R}_{S_n}(\mathcal{A}_n(S_n)) = 0$. Then,*

$$2/3 R_{\mathcal{D}}(\mathcal{A}_n) \stackrel{(a)}{\leq} \text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))/n \stackrel{(b)}{\leq} H_b(R_{\mathcal{D}}(\mathcal{A}_n)) + R_{\mathcal{D}}(\mathcal{A}_n) \log(2), \quad (9)$$

where $H_b(\cdot)$ is the binary entropy function, and $R_{\mathcal{D}}(\mathcal{A}_n) = \mathbb{E}[R_{\mathcal{D}}(\mathcal{A}_n(S_n))]$.

The inequality (a) in Eq. (9) implies that, if $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))/n$ vanishes as n diverges, then $R_{\mathcal{D}}(\mathcal{A}_n)$ vanishes as well. The inequality (b) is more interesting: it implies that, if $R_{\mathcal{D}}(\mathcal{A}_n)$ vanishes as n diverges, then $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))/n$ also vanishes.

Assume that a consistent algorithm \mathcal{A} satisfies $R_{\mathcal{D}}(\mathcal{A}_n) = \theta/n$ for $\theta \in \mathbb{R} \geq 1$. Then, it is straightforward to see from Direction (b) in Eq. (9) that $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))/n = O(\theta \log(n))$. Also, for an algorithm with $R_{\mathcal{D}}(\mathcal{A}_n) = \theta \log(n)/n$ the upper bound in Eq. (9) is given by $O(\theta (\log(n))^2)$. This observation suggests that our upper bound for $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))/n$ in Eq. (9) provides a bound on the expected excess risk which is sub-optimal by a $\log(n)$ factor in some interesting cases.

Remark 5.2. Note that the result in Theorem 5.1 *does not* imply our results in former sections. In particular our results in Theorem 2.3, Theorem 3.4, Corollary 4.7, and (later in) Theorem 5.6 show that CMI framework provides *optimal* characterization of the expected excess risk in the considered scenarios. \triangleleft

The following corollary summarizes our result for the consistent algorithms with a leave-one-out error guarantee.

Corollary 5.3. *Let $n \in \mathbb{N}$ and $\theta \in \mathbb{R}_+$, such that $n \geq 2\theta$. Let \mathcal{A}_n be a consistent learning algorithm. Let \mathcal{D} be a distribution on \mathcal{Z} and assume that, with probability one over a sequence $S = (Z_1, \dots, Z_{n+1}) \sim \mathcal{D}^{n+1}$, we have $\frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{E}^S[\ell(\mathcal{A}_n(S_{-i}), Z_i)] \leq \frac{\theta}{n+1}$, where the expectation is taken only over the randomness in \mathcal{A}_n . Then,*

$$\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) \leq \theta \log((n+1)/\theta) + 2\theta \log 2.$$

5.1 The One-Inclusion Graph Prediction Strategy

Haussler et al. [5] proposed an improper learning rule for learning VC classes based on the one-inclusion graph [34]. We provide a description of this algorithm in Appendix I. The deterministic version of this prediction rule satisfies the following property. Let \mathcal{H} be a concept class with VC dimension d . For every $n \in \mathbb{N}$, $h \in \mathcal{H}$, and $(x_1, \dots, x_{n+1}) \in \mathcal{X}^{n+1}$, let $S = ((x_1, h(x_1)), \dots, (x_{n+1}, h(x_{n+1})))$. Then $\frac{1}{n+1} \sum_{i=1}^{n+1} \ell(\mathcal{A}_n(S_{-i}), (x_i, h(x_i))) \leq \frac{d}{n+1}$. A direct application of Corollary 5.3 gives the following results.

Corollary 5.4. *Let \mathcal{A}_n denote the deterministic one-inclusion graph for learning class \mathcal{H} with VC dimension d . Then, for every realizable distribution \mathcal{D} and $n \geq 2d$, we have $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) \leq d \log((n+1)/d) + 2d \log 2$.*

Remark 5.5. In Theorem 2.3 we provide a bound on eCMI of any proper ERM. However, for improper learners, we can construct a consistent algorithm with maximal eCMI. For instance, consider $\mathcal{X} = [0, 1]$, $\mathcal{D}_{\mathcal{X}} = \text{Unif}([0, 1])$, the concept class of threshold with target function $h^*(x) = \mathbb{1}[x \geq 1/2]$. Consider a learning algorithm that gives the correct predictions on the points that are in the training set, and for a point that is not in the training set it always predicts one. One can show that eCMI of this consistent algorithm is $\Omega(n)$. \triangleleft

Haussler et al. [5] showed that the one-inclusion graph algorithm achieves $\mathbb{E}R_{\mathcal{D}}(\mathcal{A}_n(S_n)) \leq d/n$ for learning a class \mathcal{H} with VC dimension d . Corollary 5.4 implies that $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) = O(d \log(n))$ for every *deterministic* one-inclusion graph prediction rule. Combining this result with Eq. (3) provides a bound on the excess risk which is suboptimal by a $\log n$ factor. In the next theorem, we show that, in at least one interesting special case, it is possible to remove the logarithmic factor from eCMI by exploiting a *randomized* one-inclusion graph prediction algorithm.

Theorem 5.6. *Let \mathcal{H} denote the class of singletons (point functions) on $\mathcal{X} = \mathbb{R}$. There exists a randomized one-inclusion graph prediction rule \mathcal{A}_n for learning class \mathcal{H} such that for every realizable distribution \mathcal{D} and $n \geq 2$, we have $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) = O(1)$.*

6 Remaining Gaps and Open Questions

For proper learning of VC classes, Hanneke [35] showed the assumption $\mathfrak{s} < \infty$ is a necessary and sufficient condition for the existence of a distribution-free bound on the expected risk of all ERMs converging at a rate $O(1/n)$. In Corollary 4.7 and Theorem 4.9, we showed the same rate for the expected risk of a broad class of ERMs can be obtained using the CMI framework. It is an open question to show that for a class with finite star number, every ERM has bounded eCMI.

An important open problem is to show that for every VC class with finite dual Helly number [2] there exists a proper learning algorithm such that for every data distribution its expected empirical risk converges at a rate of $O(1/n)$ and it has bounded CMI. Combining the generalization guarantees that one can retrieve from Eq. (2) the expected excess risk of the learner with these properties matches the optimal rate from Bousquet et al. [2].

For improper learning of VC classes, we showed a general result for the deterministic one-inclusion graph prediction rule which is suboptimal by a $\log n$ factor. We conjecture that for every VC class with dimension d there exists a probability assignment for the randomized one-inclusion graph for which eCMI is $O(d)$. In Theorem 5.6, we showed this claim holds for the class of point functions.

We also remark that if the answers to the above questions are affirmative, then it can be argued that the CMI framework is expressive enough so that it can explain generalization properties of VC classes. Otherwise, a negative answer to any of the questions implies that there is gap between CMI framework and VC theory.

In Theorem 5.1 we proved that eCMI is *universal* in the realizable setting. A fundamental question to ask is whether for every data distribution \mathcal{D} and consistent learner \mathcal{A} , $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))/n$ vanishes as the number training samples n diverges *at the same rate* with the excess risk, i.e., $R_{\mathcal{D}}(\mathcal{A}_n)$.

Acknowledgments

The authors would like to thank Blair Bilodeau, Mufan Bill Li, and Jeffery Negrea for feedback on drafts of this work.

Funding

M. Haghifam is supported by the Vector Institute, University of Toronto, and a MITACS Accelerate Fellowship with Element AI. S. Moran is a Robert J. Shillman Fellow and is supported by the ISF, grant no. 1225/20, by an Azrieli Faculty Fellowship, and by BSF grant 2018385. D. M. Roy was supported, in part, by an NSERC Discovery Grant, Ontario Early Researcher Award, and a stipend provided by the Charles Simonyi Endowment. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute www.vectorinstitute.ai/partners.

References

- [1] T. Steinke and L. Zakythinou. “Reasoning About Generalization via Conditional Mutual Information”. In: *Proceedings of the 33rd Conference On Learning Theory*. Ed. by J. Abernethy and S. Agarwal. Vol. 125. Proceedings of Machine Learning Research. PMLR, July 2020, pp. 3437–3452. arXiv: 2001.09122.
- [2] O. Bousquet, S. Hanneke, S. Moran, and N. Zhivotovskiy. “Proper Learning, Helly Number, and an Optimal SVM Bound”. In: *Proceedings of the 33rd Conference On Learning Theory*. Ed. by J. Abernethy and S. Agarwal. Vol. 125. Proceedings of Machine Learning Research. PMLR, July 2020, pp. 582–609.
- [3] T. Steinke and L. Zakythinou. “Open Problem: Information Complexity of VC Learning”. In: *Proceedings of the 33rd Conference On Learning Theory*. Ed. by J. Abernethy and S. Agarwal. Vol. 125. Proceedings of Machine Learning Research. PMLR, July 2020, pp. 3857–3863.
- [4] S. Hanneke and L. Yang. “Minimax analysis of active learning.” *J. Mach. Learn. Res.* 16.12 (2015), pp. 3487–3602.
- [5] D. Haussler, N. Littlestone, and M. K. Warmuth. “Predicting $\{0, 1\}$ -functions on randomly drawn points”. *Information and Computation* 115.2 (1994), pp. 248–292.
- [6] D. Russo and J. Zou. *How much does your data exploration overfit? Controlling bias via information usage*. 2015. arXiv: 1511.05219.
- [7] D. Russo and J. Zou. “Controlling Bias in Adaptive Data Analysis Using Information Theory”. In: *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*. Ed. by A. Gretton and C. C. Robert. Vol. 51. Proceedings of Machine Learning Research. Cadiz, Spain: PMLR, 2016, pp. 1232–1240.
- [8] A. Xu and M. Raginsky. “Information-theoretic analysis of generalization capability of learning algorithms”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 2524–2533.
- [9] M. Raginsky, A. Rakhlin, M. Tsao, Y. Wu, and A. Xu. “Information-theoretic analysis of stability and bias of learning algorithms”. In: *2016 IEEE Information Theory Workshop (ITW)*. IEEE. 2016, pp. 26–30.
- [10] A. Asadi, E. Abbe, and S. Verdú. “Chaining mutual information and tightening generalization bounds”. In: *Advances in Neural Information Processing Systems* 32. 2018.
- [11] A. Pensia, V. Jog, and P.-L. Loh. “Generalization error bounds for noisy, iterative algorithms”. In: *2018 IEEE International Symposium on Information Theory (ISIT)*. 2018, pp. 546–550.
- [12] R. Bassily, S. Moran, I. Nachum, J. Shafer, and A. Yehudayoff. “Learners that Use Little Information”. In: *Algorithmic Learning Theory*. 2018, pp. 25–55.
- [13] I. Nachum, J. Shafer, and A. Yehudayoff. “A direct sum result for the information complexity of learning”. In: *Conference On Learning Theory*. PMLR. 2018, pp. 1547–1568.
- [14] I. Nachum and A. Yehudayoff. “Average-case information complexity of learning”. In: *Algorithmic Learning Theory*. PMLR. 2019, pp. 633–646.
- [15] I. Issa, A. R. Esposito, and M. Gastpar. “Strengthened Information-theoretic Bounds on the Generalization Error”. In: *2019 IEEE International Symposium on Information Theory (ISIT)*. 2019, pp. 582–586.
- [16] J. Negrea, M. Haghifam, G. K. Dziugaite, A. Khisti, and D. M. Roy. “Information-Theoretic Generalization Bounds for SGLD via Data-Dependent Estimates”. In: *Advances in Neural Information Processing Systems*. 2019, pp. 11013–11023.
- [17] Y. Bu, S. Zou, and V. V. Veeravalli. “Tightening Mutual Information-Based Bounds on Generalization Error”. *IEEE Journal on Selected Areas in Information Theory* 1.1 (2020), pp. 121–130.

- [18] R. Livni and S. Moran. “A Limitation of the PAC-Bayes Framework”. In: *Advances in Neural Information Processing Systems*. Vol. 33. 2020, pp. 20543–20553.
- [19] G. Neu, G. K. Dziugaite, M. Haghifam, and D. M. Roy. “Information-Theoretic Generalization Bounds for Stochastic Gradient Descent”. In: *Proceedings of Thirty Fourth Conference on Learning Theory*. Ed. by M. Belkin and S. Kpotufe. Vol. 134. Proceedings of Machine Learning Research. PMLR, 15–19 Aug 2021, pp. 3526–3545.
- [20] H. Hafez-Kolahi, Z. Golgooni, S. Kasaei, and M. Soleymani. “Conditioning and Processing: Techniques to Improve Information-Theoretic Generalization Bounds”. In: *Advances in Neural Information Processing Systems 34*. 2020.
- [21] M. Haghifam, J. Negrea, A. Khisti, D. M. Roy, and G. K. Dziugaite. “Sharpened Generalization Bounds based on Conditional Mutual Information and an Application to Noisy, Iterative Algorithms”. In: *Advances in Neural Information Processing Systems*. 2020, pp. 9925–9935.
- [22] F. Hellström and G. Durisi. “Generalization Bounds via Information Density and Conditional Information Density”. *IEEE Journal on Selected Areas in Information Theory* 1.3 (2020), pp. 824–839.
- [23] B. Rodríguez-Gálvez, G. Bassi, R. Thobaben, and M. Skoglund. “On Random Subset Generalization Error Bounds and the Stochastic Gradient Langevin Dynamics Algorithm”. In: *IEEE Information Theory Workshop (ITW)*. IEEE. 2020.
- [24] R. Zhou, C. Tian, and T. Liu. “Individually Conditional Individual Mutual Information Bound on Generalization Error”. In: *2021 IEEE International Symposium on Information Theory (ISIT)*. 2021, pp. 670–675.
- [25] F. Hellström and G. Durisi. “Fast-Rate Loss Bounds via Conditional Information Measures with Applications to Neural Networks”. In: *2021 IEEE International Symposium on Information Theory (ISIT)*. 2021, pp. 952–957.
- [26] N. Littlestone and M. Warmuth. “Relating data compression and learnability” (1986).
- [27] S. Hanneke and A. Kontorovich. “Stable Sample Compression Schemes: New Applications and an Optimal SVM Margin Bound”. In: *Algorithmic Learning Theory*. PMLR. 2021, pp. 697–721.
- [28] M. Madiman and P. Tetali. “Sandwich bounds for joint entropy”. In: *2007 IEEE International Symposium on Information Theory*. IEEE. 2007, pp. 511–515.
- [29] P. M. Long and R. J. Long. “On the complexity of proper distribution-free learning of linear classifiers”. In: *Algorithmic Learning Theory*. PMLR. 2020, pp. 583–591.
- [30] S. Hanneke and A. Kontorovich. “A sharp lower bound for agnostic learning with sample compression schemes”. In: *Algorithmic Learning Theory*. PMLR. 2019, pp. 489–505.
- [31] S. Floyd and M. Warmuth. “Sample compression, learnability, and the Vapnik-Chervonenkis dimension”. *Machine learning* 21.3 (1995), pp. 269–304.
- [32] V. Vapnik and A. Chervonenkis. *Theory of pattern recognition*. 1974.
- [33] M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [34] N. Alon, D. Haussler, and E. Welzl. “Partitioning and geometric embedding of range spaces of finite Vapnik-Chervonenkis dimension”. In: *Proceedings of the third annual symposium on Computational geometry*. 1987, pp. 331–340.
- [35] S. Hanneke. “Refined error bounds for several learning algorithms”. *The Journal of Machine Learning Research* 17.1 (2016), pp. 4667–4721.
- [36] S. Hanneke. “The optimal sample complexity of PAC learning”. *The Journal of Machine Learning Research* 17.1 (2016), pp. 1319–1333.
- [37] S. Moran and A. Yehudayoff. “Sample compression schemes for VC classes”. *Journal of the ACM (JACM)* 63.3 (2016), pp. 1–10.
- [38] S. Hanneke. “Teaching dimension and the complexity of active learning”. In: *International Conference on Computational Learning Theory*. Springer. 2007, pp. 66–81.
- [39] R. El-Yaniv and Y. Wiener. “On the Foundations of Noise-free Selective Classification”. *Journal of Machine Learning Research* 11.5 (2010).
- [40] E. Zermelo. “Untersuchungen über die Grundlagen der Mengenlehre. I”. *Mathematische Annalen* 65.2 (1908), pp. 261–281.

A Known Bounds for Learning VC Classes

In this part, we provide a landscape of the known results for the learning VC classes. One key distinction is proper learning versus improper learning. In particular, for every VC class with dimension d , there exists a consistent and improper learning algorithm that achieves $O(d/n)$ risk under realizability, and this bound is optimal [36, 5]. The situation for proper learning is much more complicated. In general, the achievable rate for the proper learning of VC classes is off by a log factor, i.e., $O(d \log(n)/n)$. Bousquet et al. [2] show that when the dual Helly and hollow star number, which are combinatorial complexity measures of the class, agree, then they characterize the existence of an optimal proper learner. Also, a subclass of VC for which the log factor provably cannot be removed using proper learners is characterized in [2, Thm. 11]. Moreover, for general ERMs, Hanneke [35] shows the finiteness of star number is a necessary and sufficient condition under which we can remove the log factor using any arbitrary ERMs.

It is interesting to note that the results of Moran and Yehudayoff [37] reveal a connection between the general sample compression schemes and VC classes. However, it is not known whether, in general, the optimal rates for VC classes are always witnessed by compression schemes.

B Proof of Theorem 2.3

Fix $n \in \mathbb{N}$. We have $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) = \mathbb{E}[I^Z(L; U)]$ by definition, and $I^Z(L; U) \leq H^Z(L)$ a.s., by the nonnegativity of (conditional) entropy. It then suffices to bound the cardinality, C , of the support of the conditional distribution of L , given Z , because $H^Z(L) \leq \log C$ a.s. For all $i \in \{0, 1\}$ and $j \in [n]$, write $Z_{i,j} = (X_{i,j}, Y_{i,j})$, and consider the set of possible predictions on Z ,

$$P = \{p \in \{0, 1\}^{\{0,1\} \times [n]} : \exists h \in \mathcal{H}_n, \forall i \in \{0, 1\}, j \in [n], p_{i,j} = h(X_{i,j})\}. \quad (10)$$

The set P is precisely the set of possible labellings of the $2n$ inputs $(X_{i,j})$, which is bounded by the growth function of \mathcal{H}_n evaluated at $2n$ points. By the Sauer–Shelah lemma, the cardinality of P is thus bounded above by $6n^{d_n}$. Note that the support of the conditional distribution of L given Z is

$$\{c \in \{0, 1\}^{\{0,1\} \times [n]} : \exists p \in P, \forall i \in \{0, 1\}, \forall j \in [n], \text{if and only if } c_{i,j} = \mathbb{1}[p_{i,j} \neq Y_{i,j}]\}. \quad (11)$$

Therefore, the cardinality of the support is no greater than that of P , hence $C \leq 6n^{d_n}$.

C Proof of Theorem 4.4

We prove the claim by contradiction. Pick f and $c \geq 0$. Let \mathcal{H} be a concept class with finite VC dimension d as shown to exist by Theorem 4.3.

Let \mathcal{A} be a proper learning algorithm for \mathcal{H} , let $n \geq d$, and assume, for the eventual purpose of obtaining a contradiction, that $\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq f(d)$ for all \mathcal{D} and, for all $s \in \mathcal{Z}^n$, $\mathbb{E}\hat{R}_s(\mathcal{A}_n(s)) \leq cd/n$ if there exists $h \in \mathcal{H}$ such that $\hat{R}_s(h) = 0$. Pick a realizable distribution \mathcal{D} . It follows from the above assumption and Eq. (1) that $\mathbb{E}[R_{\mathcal{D}}(\mathcal{A}_n(S_n))] \leq 2cd/n + 3f(d)/n = (3f(d) + 2cd)/n$. By Markov’s inequality, $\mathbb{P}(R_{\mathcal{D}}(\mathcal{A}_n(S_n)) \geq \epsilon) \leq \frac{1}{n\epsilon}(3f(d) + 2cd)$. It follows that the sample complexity of proper learning \mathcal{H} satisfies

$$\mathcal{M}_{\text{prop}}^{\mathcal{H}}(\epsilon, \delta) \leq \frac{1}{\epsilon\delta}(3f(d) + 2cd). \quad (12)$$

Now, fix $\delta \in (0, 1/100)$ and fix a convergent sequence of $\epsilon_i \downarrow 0$. There exists J such that, for all $i \geq J$,

$$\frac{1}{\tilde{c}\delta}(3f(d) + 2cd) < d \text{Log} \frac{1}{\epsilon_i} + \text{Log} \frac{1}{\delta}, \quad (13)$$

for \tilde{c} as in Theorem 4.3. Combining Eq. (13) with Eq. (12), $\mathcal{M}_{\text{prop}}^{\mathcal{H}}(\epsilon_i, \delta) < \frac{\tilde{c}}{\epsilon_i}(d \text{Log} \frac{1}{\epsilon_i} + \text{Log} \frac{1}{\delta})$ for $i \geq J$. Simultaneously, from Theorem 4.3, it follows that $\mathcal{M}_{\text{prop}}^{\mathcal{H}}(\epsilon_i, \delta) \geq \frac{\tilde{c}}{\epsilon_i}(d \text{Log} \frac{1}{\epsilon_i} + \text{Log} \frac{1}{\delta})$, a contradiction.

D Proof of Theorem 4.6

Let $n \geq 5$. First we begin with two definitions: letting S_n be the random element in $(\mathcal{X} \times \mathcal{Y})^n$ representing our training sample, the *empirical teaching dimension* [38, 39], denoted ETD_n , is size of the smallest subset of S that produces the same version space, i.e.,

$$\text{ETD}_n = \min\{|S'| : S' \subseteq S_n, V_{\mathcal{H}}[S'] = V_{\mathcal{H}}[S_n]\}. \quad (14)$$

An important fact about the empirical teaching dimension is that ETD_n is bounded by star number almost surely [35]. An *empirical teaching set* is any subset of S_n that achieves the minimum in the definition of the empirical teaching dimension. Consider a realizable training set $S_n \in (\mathcal{X} \times \mathcal{Y})^n$. Let S' denote an empirical teaching set of S_n . Let $\hat{S} \subseteq S_n \setminus S'$. By the definition of the version space we have $V_{\mathcal{H}}[S_n] \subseteq V_{\mathcal{H}}[S_n \setminus \hat{S}] \subseteq V_{\mathcal{H}}[S']$. Also, by the definition of S' , we have $V_{\mathcal{H}}[S_n] = V_{\mathcal{H}}[S']$; therefore, $V_{\mathcal{H}}[S_n] = V_{\mathcal{H}}[S_n \setminus \hat{S}]$. This argument shows that the version space is *stable*, in the sense that removing any point that is not in S' does not alter the version space. Let $\text{ETS}(S_n)$ denote the empirical teaching set S_n .

We also need the definition of the *region of disagreement* denoted by $\text{DIS}(V_{\mathcal{H}}[S_n]) = \{x \in \mathcal{X} \mid \exists h, h' \in V_{\mathcal{H}}[S_n] \text{ such that } h(x) \neq h'(x)\}$.

Following the same line of reasoning as in Eq. (4) we obtain $I(V_{\mathcal{H}}[Z_U]; U|Z) \leq n \log 2 - \sum_{i=1}^n H(U_i|U_{-i}, V_{\mathcal{H}}[Z_U], Z)$. Since the order of the training set does not change the version space we get $\sum_{i=1}^n H(U_i|U_{-i}, V_{\mathcal{H}}[Z_U], Z) = nH(U_1|U_{-1}, V_{\mathcal{H}}[Z_U], Z)$. Fix $i \in [n]$, and define $U_{i \rightarrow b} \triangleq (U_1, \dots, U_{i-1}, b, U_{i+1}, \dots, U_n)$ for $b \in \{0, 1\}$. Using this notation we can define training sets $S_{i \rightarrow b} \triangleq Z_{U_{i \rightarrow b}}$ for $b \in \{0, 1\}$ and $i \in [n]$. Let $\mathcal{F}_1 = \sigma(U_{-1}, V_{\mathcal{H}}[Z_U], Z)$. Then, we have

$$H(U_1|U_{-1}, V_{\mathcal{H}}[Z_U], Z) \geq \mathbb{E}[H^{\mathcal{F}_1}(U_1) \mathbb{1}[V_{\mathcal{H}}[S_{1 \rightarrow 0}] = V_{\mathcal{H}}[S_{1 \rightarrow 1}]]]. \quad (15)$$

Using the same techniques as in the proof Theorem 3.4, we can show that on the event $\mathbb{1}[V_{\mathcal{H}}[S_{1 \rightarrow 0}] = V_{\mathcal{H}}[S_{1 \rightarrow 1}]]$, $H^{\mathcal{F}_1}(U_1) = \log(2)$. Thus, $H(U_1|U_{-1}, V_{\mathcal{H}}[Z_U], Z) \geq \mathbb{P}(V_{\mathcal{H}}[S_{1 \rightarrow 0}] = V_{\mathcal{H}}[S_{1 \rightarrow 1}]) \log 2$.

For $i \in [n]$, define the training set $S_{-i} \triangleq \{Z_{U_1}, \dots, Z_{U_{i-1}}, Z_{U_{i+1}}, \dots, Z_{U_n}\}$. Recall that $Z_{i,j} = (X_{i,j}, Y_{i,j})$. We claim that

$$V_{\mathcal{H}}[S_{1 \rightarrow 0}] \neq V_{\mathcal{H}}[S_{1 \rightarrow 1}] \Rightarrow (X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])) \vee (X_{1,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])). \quad (16)$$

We prove this claim by contraposition. Given that $(X_{0,1} \notin \text{DIS}(V_{\mathcal{H}}[S_{-1}])) \wedge (X_{1,1} \notin \text{DIS}(V_{\mathcal{H}}[S_{-1}]))$, we have that the concepts in $\text{DIS}(V_{\mathcal{H}}[S_{-1}])$ agree for prediction of $X_{0,1}$ and $X_{1,1}$. As \mathcal{D} is a realizable distribution, we conclude $V_{\mathcal{H}}[S_{1 \rightarrow 0}] = V_{\mathcal{H}}[S_{-1}] = V_{\mathcal{H}}[S_{1 \rightarrow 1}]$.

In the next step, we provide an upper bound on $\mathbb{P}((X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])) \vee (X_{1,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])))$ as follows

$$\begin{aligned} & \mathbb{P}((X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])) \vee (X_{1,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}]))) \\ & \leq \mathbb{P}((X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}]))) + \mathbb{P}((X_{1,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}]))) \\ & = 2\mathbb{P}(X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])). \end{aligned} \quad (17)$$

Here, we have used the union bound and the points in Z are i.i.d.. Then, we can write

$$\begin{aligned} \mathbb{P}(X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])) &= \mathbb{E}[\mathbb{1}[X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])]] \\ &= \mathbb{E}[\mathbb{1}[X_1 \in \text{DIS}(V_{\mathcal{H}}[\{Z_2, \dots, Z_n\}])]]. \end{aligned} \quad (18)$$

The last step follows from U and Z are independent, and the points in Z are i.i.d.. Therefore, in the last step we consider the expectation over $(Z_1, \dots, Z_n) \sim \mathcal{D}^{\otimes n}$. Note that $Z_i = (X_i, Y_i)$. By the exchangeability of the points in (Z_1, \dots, Z_n) we have

$$\mathbb{E}[\mathbb{1}[X_1 \in \text{DIS}(V_{\mathcal{H}}[\{Z_2, \dots, Z_n\}])]] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathbb{1}[X_i \in \text{DIS}(V_{\mathcal{H}}[\{Z_1, \dots, Z_n\} \setminus \{Z_i\}])]]. \quad (19)$$

Then, we claim that given $X_i \in \text{DIS}(V_{\mathcal{H}}[\{Z_1, \dots, Z_n\} \setminus \{Z_i\}])$ then $X_i \in S'$ where S' is any teaching set of $\{Z_1, \dots, Z_n\}$. We can easily prove this claim by contradiction and the stability of the

version space shown in the beginning of this section. Therefore,

$$\begin{aligned}
\frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathbb{1}[X_i \in \text{DIS}(V_{\mathcal{H}}[\{Z_1, \dots, Z_n\} \setminus \{Z_i\}])]] &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathbb{1}[X_i \in S']] \\
&= \frac{1}{n} \mathbb{E} \left[\sum_{i=1}^n \mathbb{1}[X_i \in S'] \right] \\
&\leq \frac{\mathfrak{s}}{n}.
\end{aligned} \tag{20}$$

Here, we have used the linearity of the expectation, and the last step follows from the fact that the cardinality of S' is at most \mathfrak{s} almost surely. By Eq. (17)-Eq. (20), we have $\mathbb{P}((X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])) \vee (X_{1,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}]))) \leq 2\mathfrak{s}/n$. Then, by Eq. (16) we have

$$\begin{aligned}
\text{H}(U_1|U_{-1}, V_{\mathcal{H}}[Z_U], Z) &\geq \mathbb{P}(V_{\mathcal{H}}[S_{1 \rightarrow 0}] = V_{\mathcal{H}}[S_{1 \rightarrow 1}]) \log 2 \\
&\geq [1 - \mathbb{P}((X_{0,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])) \vee (X_{1,1} \in \text{DIS}(V_{\mathcal{H}}[S_{-1}])))] \log 2 \\
&\geq (1 - \frac{2\mathfrak{s}}{n}) \log 2.
\end{aligned} \tag{21}$$

Finally, combining this results and $I(V_{\mathcal{H}}[Z_U]; U|Z) \leq n \log 2 - n \text{H}(U_1|U_{-1}, V_{\mathcal{H}}[Z_U], Z)$, we obtain $I(V_{\mathcal{H}}[Z_U]; U|Z) \leq 2\mathfrak{s} \log 2$ which was to be shown.

E Proof of Theorem 4.8

We begin the proof by a lemma, which can be seen as the generalization of the well-known coupon collector's problem [2, Lem. 19].

Lemma E.1. *Let $M, m \in \mathbb{N}$ and $1 \leq m \leq M$. Assume we take k samples X_1, \dots, X_k uniformly at random with replacement from $[M]$. Then $\mathbb{P}(|[M] \setminus \{X_1, \dots, X_k\}| \geq m) \geq 1/2$ if $k \leq (M/2) \log \frac{M}{m}$.*

Note that $|[M] \setminus \{X_1, \dots, X_k\}|$ is the number of unseen elements from $[M]$ after observing samples X_1, \dots, X_k .

Consider a concept class \mathcal{H} over the input space \mathcal{X} whose star number is \mathfrak{s} . Let $M = \min\{n, \mathfrak{s}\}$. By the definition of star number, there exists $(x_1, \dots, x_M) \in \mathcal{X}^M$, such that $((x_1, y_1), \dots, (x_M, y_M))$ is realizable by $h_0^* \in \mathcal{H}$, and every neighbour of this sequence is also realizable by a classifier in \mathcal{H} . Let $h_{x_i} \in \mathcal{H}$ be any classifier such that $\{j \in [M] | h_{x_i}(x_j) \neq y_j\} = \{i\}$.

For the case $M = 1$, our lower bound is simply $\min\{n, \mathfrak{s}\} - 1 = M - 1 = 0$ which is trivial since the mutual information is non-negative. Therefore, in the rest of the proof we assume $M \geq 2$.

For the case $M \geq 2$, consider the following distribution on the input space

$$\mathcal{D}_X(x_1) = 1 - \frac{M-1}{n} \text{ and } \mathcal{D}_X(x_i) = \frac{1}{n}, \text{ for } i \in \{2, \dots, M\}.$$

Also let the target function (labelling function) be h_0^* . Let Z, U , and S be defined as usual, based on a sample from \mathcal{D}_X labeled by h_0^* . Since \mathcal{D}_X has zero measure on $\mathcal{X} \setminus \{x_1, \dots, x_M\}$, we can, without any loss of generality, assume that $\mathcal{H} = \{h_0^*, h_{x_1}, \dots, h_{x_M}\}$, as every other classifier is equivalent to one of these \mathcal{D}_X -almost everywhere. When we release the version space, we can agree in advance that elements stand for their equivalence classes, which does not affect the information content.

Let $X = (X_{i,j})_{i \in \{0,1\}, j \in [n]}$ denote the inputs observed in the supersample Z . Define X_U as the sequence of the inputs observed in Z_U . Similarly, let $X_{\bar{U}}$ denote the sequence of inputs observed in the ‘‘ghost sample’’ $Z_{\bar{U}}$, where \bar{U} denotes the sequence U but with every entry flipped. In the following, we write $x \in X_U$ to mean that there exists $i \in [n]$ such that $X_{U,i} = x$.

We claim that $V_{\mathcal{H}}[Z_U] = \{h_{x_i} | i \in [M], x_i \notin X_U\} \cup \{h_0^*\}$.⁵ To see this, note that, if $x_i \notin X_U$, then $h_{x_i} = h_0^*$ on S and so both are ERM. In the other direction, if $x_i \in X_U$, then h_{x_i} makes at least one mistake, and so is not an ERM. Thus, if $V_{\mathcal{H}}[Z_U]$ has a non-empty intersection with $\{h_{X_{0,i}}, h_{X_{1,i}}\}$,

⁵Formally, this statement holds almost surely. We will skip such declarations for the remainder of the proof.

we can perfectly recover U_i from $V_{\mathcal{H}}[Z_U]$, since in each column of the supersample Z , only one point is selected for the training set. We use this observation to lower bound the conditional entropy of U given Z and the version space.

To that end, let $J = \{j \in [n] | V_{\mathcal{H}}[Z_U] \cap \{h_{X_{0,j}}, h_{X_{1,j}}\} \neq \emptyset\}$ be the set that contains the index of every column j for which we can perfectly recover U_j from the version space. Note that we can represent J in an equivalent form as follows. For two sequences $v, w \in \{x_1, \dots, x_M\}^n$, define $K(v, w) = \{j \in [n] | w_j \neq v_i \text{ for all } i \in [n]\}$. By the definition of the version space, $x_i \notin X_U$ is equivalent to $h_{x_i} \in V_{\mathcal{H}}[Z_U]$. Let $j \in [n]$, then

$$j \in J \Leftrightarrow V_{\mathcal{H}}[Z_U] \cap \{h_{X_{1-U_j,j}}\} = \{h_{X_{1-U_j,j}}\} \Leftrightarrow X_{1-U_j,j} \notin X_U \Leftrightarrow j \in K(X_U, X_{\bar{U}}).$$

Therefore, $J = K(X_U, X_{\bar{U}})$.

By the definition of the mutual information, and the fact that U and Z are independent, we have

$$I(V_{\mathcal{H}}[Z_U]; U|Z) = n \log(2) - H(U|Z, V_{\mathcal{H}}[Z_U]). \quad (22)$$

Then

$$H(U|Z, V_{\mathcal{H}}[Z_U]) = H(U|Z, V_{\mathcal{H}}[Z_U], J) \quad (23)$$

$$= H(U_J, U_{J^c}|Z, V_{\mathcal{H}}[Z_U], J) \quad (24)$$

$$= H(U_{J^c}|Z, V_{\mathcal{H}}[Z_U], J) \quad (25)$$

$$\leq \mathbb{E}[n - |J|] \log(2), \quad (26)$$

where Eq. (23) follows from J being known from Z and $V_{\mathcal{H}}[Z_U]$; Eq. (25) follows from U_J being known from $J, Z, V_{\mathcal{H}}[Z_U]$; and Eq. (26) follows from the cardinality of the support of the distribution of U_{J^c} being no more than $2^{(n-|J|)}$. Therefore, by Eq. (22) and Eq. (26), we have

$$I(V_{\mathcal{H}}[Z_U]; U|Z) \geq \mathbb{E}[|J|] \log 2. \quad (27)$$

In the next step of the proof, we lower bound $\mathbb{E}[|J|]$. Because U and Z are independent and Z is an i.i.d. array, X_U and $X_{\bar{U}}$ are independent and identically distributed sequences of i.i.d. elements with common distribution $\mathcal{D}_{\mathcal{X}}$. Thus, $\mathbb{E}[|J|] = \mathbb{E}[|K(X, \bar{X})|]$, where X and \bar{X} are i.i.d. copies of X_U (equivalently, $X_{\bar{U}}$).

Let \hat{M} be the number of elements of $\{x_2, \dots, x_M\}$ *not* appearing in X , i.e., $\hat{M} = |\{i \in \{2, \dots, M\} | x_i \neq X_j \text{ for all } j \in [n]\}|$. Conditional on X , if $x_1 \in X$, then $|K(X, \bar{X})|$ is a Binomial random variable with n trials, each succeeding with probability $\frac{\hat{M}}{n}$. If $x_1 \notin X$, then $|K(X, \bar{X})|$ is a Binomial random variable with n trials, each succeeding with probability $\frac{\hat{M}}{n} + 1 - \frac{M-1}{n} \geq \frac{\hat{M}}{n}$. Therefore,

$$\mathbb{E}[|J|] = \mathbb{E}[|K(X, \bar{X})|] \quad (28)$$

$$= \mathbb{E}[\mathbb{E}^X[|K(X, \bar{X})|]] \quad (29)$$

$$\geq \mathbb{E}[\hat{M}]. \quad (30)$$

Here, Eq. (29) follows from the tower rule and Eq. (30) follows from the fact that the mean of the binomial distribution with n trials, each succeeding with probability p , is np .

Fix $\beta = \exp(-3)$ and $m = \beta(M-1)$. Let \hat{N} denote the number of samples in X falling in $\{x_2, \dots, x_M\}$. Then, using the tower rule and Markov's inequality we have

$$\mathbb{E}[\hat{M}] = \mathbb{E}[\mathbb{E}^{\hat{N}}[\hat{M}]] \quad (31)$$

$$\geq \mathbb{E}[m \mathbb{P}^{\hat{N}}[\hat{M} \geq m]] \quad (32)$$

$$\geq \mathbb{E}[m \mathbb{P}^{\hat{N}}[\hat{M} \geq m] \mathbb{1}[\hat{N} \leq ((M-1)/2) \log \beta^{-1}]] \quad (33)$$

Note that, conditional on \hat{N} , the \hat{N} samples falling in $\{x_2, \dots, x_M\}$ are conditionally independent, with conditional distribution uniform on this set. Using this fact, from Lemma E.1 we obtain

$$\mathbb{E}[m \mathbb{P}^{\hat{N}}[\hat{M} \geq m] \mathbb{1}[\hat{N} \leq ((M-1)/2) \log \beta^{-1}]] \geq \frac{1}{2} m \mathbb{P}(\hat{N} \leq ((M-1)/2) \log \beta^{-1}). \quad (34)$$

We have $\mathbb{E}[\hat{N}] = M - 1$ and, by Markov's inequality, $\mathbb{P}(\hat{N} \leq \frac{M-1}{2} \log \beta^{-1}) \geq 1 - 2/\log(\beta^{-1}) = 1/3$. By combining Eq. (27), Eq. (30), Eq. (34), and $\mathbb{P}(\hat{N} \leq \frac{M-1}{2} \log \beta^{-1}) \geq 1/3$ we get

$$\begin{aligned} I(\mathcal{V}_{\mathcal{H}}[Z_U]; U|Z) &\geq \frac{\beta \log 2}{6}(M - 1) \\ &= \Omega(\min\{n, \mathfrak{s}\} - 1), \end{aligned}$$

which was to be shown.

F Proof of Theorem 4.9

By the well-ordering theorem [40], there exists a binary relation \ll on \mathcal{H} that is transitive, total, antisymmetric, and well-ordered. In particular, the well-ordered property implies that every nonempty subset \mathcal{H} has the least element. The proposed learning algorithm is given by

$$\mathcal{A}_n(S_n) = \text{LE}(\mathcal{V}_{\mathcal{H}}[S_n]),$$

where LE of a nonempty set denotes its least element with respect to \ll .⁶ Note that \mathcal{A}_n is *deterministic*, *consistent*, and *permutation-invariant*, i.e., the order of the points in S_n does not impact the output. Let $W = \mathcal{A}_n(Z_U)$. By the definition of the mutual information and Lemma 2.1, we obtain

$$\begin{aligned} \text{CMI}_{\mathcal{D}}(\mathcal{A}_n) &= H(U|Z) - H(U|W, Z) \\ &\leq n \log(2) - \sum_{i=1}^n H(U_i|U_{-i}, Z, W). \end{aligned} \quad (35)$$

The last step follows since U is independent of Z and the independence of indices of U and Lemma 2.1. By the permutation-invariance of the algorithm, we have $\sum_{i=1}^n H(U_i|U_{-i}, Z, W) = nH(U_1|U_{-1}, Z, W)$. Fix $i \in [n]$, and define $U_{i \rightarrow j} \triangleq (U_1, \dots, U_{i-1}, j, U_{i+1}, \dots, U_n)$ for $j \in \{0, 1\}$. Using this notation we can define training set $S_{i \rightarrow j} = Z_{U_{i \rightarrow j}}$ for $j \in \{0, 1\}$ and $U_{-i} \in \{0, 1\}^{n-1}$. Let $\mathcal{F}_1 = \sigma(W, U_{-1}, Z)$. Define the \mathcal{F}_1 -measurable event $\mathcal{E} = \{\mathcal{A}_n(S_{1 \rightarrow 0}) = \mathcal{A}_n(S_{1 \rightarrow 1})\}$. Then, we can write $H(U_1|U_{-1}, Z, W) = \mathbb{E}[H^{\mathcal{F}_1}(U_1)(\mathbb{1}[\mathcal{E}] + \mathbb{1}[\mathcal{E}^c])]$. We claim the following facts:

1. On the event \mathcal{E}^c , $H^{\mathcal{F}_1}(U_1) = 0$ a.s..
2. On the event \mathcal{E} , $H^{\mathcal{F}_1}(U_1) = \log 2$ a.s..

The reason is that since the algorithm is deterministic, on the event \mathcal{E}^c we can perfectly recover U_1 since W is either $\mathcal{A}_n(S_{1 \rightarrow 0})$ or $\mathcal{A}_n(S_{1 \rightarrow 1})$ which shows that $H^{\mathcal{F}_1}(U_1) = 0$. Then, on the event \mathcal{E} , using the Bayes rule we can show that $H^{\mathcal{F}_1}(U_1) = \log 2$. Therefore, we have $H(U_1|U_{-1}, Z, W) = \mathbb{E}[\mathbb{1}[\mathcal{E}] \log(2)]$. We can further lower bound

$$H(U_1|U_{-1}, Z, W) \geq \mathbb{E}[\mathbb{1}[\mathcal{E}] \mathbb{1}[\ell(\mathcal{A}_n(S_{1 \rightarrow 0}), Z_{1,1}) = 0 \wedge \ell(\mathcal{A}_n(S_{1 \rightarrow 1}), Z_{0,1}) = 0]] \log(2). \quad (36)$$

Next, we claim that on the event $\{\ell(\mathcal{A}_n(S_{1 \rightarrow 0}), Z_{1,1}) = 0\} \wedge \{\ell(\mathcal{A}_n(S_{1 \rightarrow 1}), Z_{0,1}) = 0\}$, we have $\mathcal{A}_n(S_{1 \rightarrow 0}) = \mathcal{A}_n(S_{1 \rightarrow 1})$ a.s.. This claim can be proved by contradiction. Assume $\mathcal{A}_n(S_{1 \rightarrow 0}) \neq \mathcal{A}_n(S_{1 \rightarrow 1})$. Consider the version space $\mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 0} \cup S_{1 \rightarrow 1}]$. By the assumptions $\ell(\mathcal{A}_n(S_{1 \rightarrow 0}), Z_{1,1}) = 0$ and $\ell(\mathcal{A}_n(S_{1 \rightarrow 1}), Z_{0,1}) = 0$, we have $\mathcal{A}_n(S_{1 \rightarrow 0}) \in \mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 0} \cup S_{1 \rightarrow 1}]$ and $\mathcal{A}_n(S_{1 \rightarrow 1}) \in \mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 0} \cup S_{1 \rightarrow 1}]$. Also, it is immediate to see that $\mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 0} \cup S_{1 \rightarrow 1}] \subseteq \mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 0}]$ and $\mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 0} \cup S_{1 \rightarrow 1}] \subseteq \mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 1}]$. Therefore, we have $\mathcal{A}_n(S_{1 \rightarrow 0}) \in \mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 1}]$ and $\mathcal{A}_n(S_{1 \rightarrow 1}) \in \mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 0}]$. By the definition of the algorithm, we choose the least hypothesis from the version space. Thus, $\mathcal{A}_n(S_{1 \rightarrow 0}) \ll \mathcal{A}_n(S_{1 \rightarrow 1})$ since $\mathcal{A}_n(S_{1 \rightarrow 1}) \in \mathcal{V}_{\mathcal{H}}[S_{1 \rightarrow 0}]$. Similarly, we can show $\mathcal{A}_n(S_{1 \rightarrow 1}) \ll \mathcal{A}_n(S_{1 \rightarrow 0})$. Therefore, considering $\mathcal{A}_n(S_{1 \rightarrow 1}) \ll \mathcal{A}_n(S_{1 \rightarrow 0})$ and $\mathcal{A}_n(S_{1 \rightarrow 0}) \ll \mathcal{A}_n(S_{1 \rightarrow 1})$, we conclude that the assumption $\mathcal{A}_n(S_{1 \rightarrow 0}) \neq \mathcal{A}_n(S_{1 \rightarrow 1})$ is false, a contradiction.

Having proved that $\mathcal{A}_n(S_{1 \rightarrow 0}) = \mathcal{A}_n(S_{1 \rightarrow 1})$ on the event $\{\ell(\mathcal{A}_n(S_{1 \rightarrow 0}), Z_{1,1}) = 0\} \wedge \{\ell(\mathcal{A}_n(S_{1 \rightarrow 1}), Z_{0,1}) = 0\}$, we can further simplify Eq. (36) as

$$\begin{aligned} H(U_1|U_{-1}, Z, W) &\geq \mathbb{E}[\mathbb{1}[\ell(\mathcal{A}_n(S_{1 \rightarrow 1}), Z_{0,1}) = 0] \mathbb{1}[\ell(\mathcal{A}_n(S_{1 \rightarrow 0}), Z_{1,1}) = 0]] \log 2 \\ &\geq (1 - \mathbb{E}[\mathbb{1}[\ell(\mathcal{A}_n(S_{1 \rightarrow 1}), Z_{0,1}) = 1]]) - \mathbb{E}[\mathbb{1}[\ell(\mathcal{A}_n(S_{1 \rightarrow 1}), Z_{0,1}) = 1]]) \log 2 \\ &= (1 - 2\mathbb{E}[R_{\mathcal{D}}(\mathcal{A}_n(S_n))]) \log 2. \end{aligned}$$

⁶In some cases this classifier may not be measurable. We will assume it is. To avoid measure-theoretic issues, one may assume \mathcal{X} is countably infinite or finite, or design a well-ordering by hand to guarantee measurability if possible.

The last step follows since the elements of $S_{1 \rightarrow 0}$ and $S_{1 \rightarrow 1}$ are i.i.d.. By plugging this lower bound into Eq. (35), we obtain

$$\text{CMI}_{\mathcal{D}}(\mathcal{A}_n) \leq 2n \log(2) \mathbb{E}[R_{\mathcal{D}}(\mathcal{A}_n(S_n))]. \quad (37)$$

Then, we use the result from Corollary 12 of [35] which states the following bound holds uniformly for the expected risk of *all* consistent and proper learners for learning a concept class with VC dimension d and star number \mathfrak{s} :

$$\mathbb{E}[R_{\mathcal{D}}(\mathcal{A}_n(S_n))] = O\left(\frac{d}{n} \log\left(\frac{\min\{\mathfrak{s}, n\}}{d}\right)\right).$$

Finally, the stated result follows by combining this bound with Eq. (37).

G Proof of Theorem 5.1

Let $L = (L_1, \dots, L_n)$ where $L_i \in \{0, 1\}^2$ denotes the vector at Column i of the loss array L . By the definition of mutual information, we have $I(L; U|Z) \leq H(L|Z)$. Since conditioning decreases the entropy, we have $H(L|Z) \leq H(L)$. Finally, by the chain rule, $H(L) \leq \sum_{i=1}^n H(L_i)$.

Note that L_i takes values in $\{[1, 0]^\top, [0, 0]^\top, [0, 1]^\top\}$, as it is assumed that \mathcal{A}_n is consistent and, by construction, at each column, one of the points is selected for the training set. If $U_i = 0$, then $Z_{0,i}$ is in the training set and so, because the algorithm is consistent, $L_{0,i} = 0$ a.s. Therefore,

$$\begin{aligned} \mathbb{P}[L_i = [1, 0]^\top] &= \mathbb{E}[\mathbb{P}^U[L_i = [1, 0]^\top]] \\ &= \mathbb{E}[\mathbb{1}[U_i = 0]\mathbb{P}^U[L_i = [1, 0]^\top]] + \mathbb{E}[\mathbb{1}[U_i = 1]\mathbb{P}^U[L_i = [1, 0]^\top]] \\ &= \mathbb{E}[\mathbb{1}[U_i = 1]\mathbb{P}^U[L_i = [1, 0]^\top]] \end{aligned}$$

Conditioning on event $U_i = 1$, we have $L_i = [1, 0]^\top = \ell(\mathcal{A}_n(Z_U), Z_{0,i}) = 1$ and therefore

$$\mathbb{P}^U[L_i = [1, 0]^\top] = \mathbb{P}^U[\ell(\mathcal{A}_n(Z_U), Z_{0,i}) = 1] = R_{\mathcal{D}}(\mathcal{A}_n),$$

where the final equality follows from the definition of the expected risk and the fact that $Z_{0,i}$ is not in Z_U . Therefore, $\mathbb{P}[L_i = [1, 0]^\top] = \mathbb{E}[\mathbb{1}[U_i = 1]]R_{\mathcal{D}}(\mathcal{A}_n) = (1/2)R_{\mathcal{D}}(\mathcal{A}_n)$. The same idea establishes that $\mathbb{P}[L_i = [0, 1]^\top] = (1/2)R_{\mathcal{D}}(\mathcal{A}_n)$.

Therefore, by the definition of the entropy

$$\begin{aligned} H(L_i) &= -(1 - R_{\mathcal{D}}(\mathcal{A}_n)) \log(1 - R_{\mathcal{D}}(\mathcal{A}_n)) - R_{\mathcal{D}}(\mathcal{A}_n) \log\left(\frac{R_{\mathcal{D}}(\mathcal{A}_n)}{2}\right) \\ &= H_b(R_{\mathcal{D}}(\mathcal{A}_n)) + R_{\mathcal{D}}(\mathcal{A}_n) \log(2). \end{aligned}$$

The stated results follows from $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) \leq \sum_{i=1}^n H(L_i)$.

H Proof of Theorem 5.6

We begin by presenting a theorem that will be used later to prove Theorem 5.6. This theorem, which might be of independent interest, shows the *average leave-one-error over supersample Z* can be used to upper bound $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))$.

Lets begin by introducing some notations. For $n \in \mathbb{N}$, let $[2n]_{n+1}$ denote the set of all size- $n+1$ subsets of $[2n]$ and let $H_b(\cdot)$ denote the binary entropy function.

Theorem H.1. *Let \mathcal{A} denote a consistent and symmetric algorithm. Let $P_e : \mathcal{Z}^n \times \mathcal{Z} \rightarrow [0, 1]$. For $s = ((x_1, y_1), \dots, (x_n, y_n)) \in \mathcal{Z}^n$ and $(x, y) \in \mathcal{Z}$, $P_e(s; (x, y))$ denotes the probability of error \mathcal{A}_n for predicting the label x where the randomness is over the internal randomness in \mathcal{A} . Then, for every distribution \mathcal{D} , we have*

$$\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) \leq n \mathbb{E} \left[H_b(\kappa(\tilde{Z})) + \kappa(\tilde{Z}) \log 2 - \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} H_b(P_e(\tilde{Z}_{J-\{j\}}; \tilde{Z}_j)) \right] \quad (38)$$

where $\tilde{Z} = (Z_1, \dots, Z_{2n}) \sim \mathcal{D}^{\otimes(2n)}$ and $\kappa(\tilde{Z}) = \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} P_e(\tilde{Z}_{J-\{j\}}; \tilde{Z}_j)$ which takes values in $[0, 1]$ almost surely. Note, in Eq. (38), the outer expectation is over \tilde{Z} .

The proof of Theorem H.1 is deferred to Appendix H.2.




Figure 1: One-inclusion graph of point functions for a set of distinct points.

H.1 Proof of Theorem 5.6

First of all, note that given *distinct* points $(x_1, \dots, x_{n+1}) \in \mathbb{R}^{n+1}$, the vertex set of the one-inclusion graph is given by

$$\{(h(x_1), \dots, h(x_{n+1})) | h \in \mathcal{H}\} = \{(a_1, \dots, a_{n+1}) | a_i \in \{0, 1\} \text{ for all } i \in [n+1], \sum_{i \in [n+1]} a_i \leq 1\}.$$

As an example Fig. 1 illustrates the structure of the one-inclusion graph for a sequence of *distinct* points. The prediction rule is as follows. Given a training set $((x_1, y_1), \dots, (x_n, y_n))$ and test example x , we have two cases. If $y_i = 0$ for all $i \in [n]$, then the label of x is predicted to be 0 with probability $\frac{n}{n+1}$, and it is predicted to be one with probability $\frac{1}{n+1}$. The second case is that there exists a point with label 1. Denote its index by $i^* \in [n]$. In this case, the target function is known to be $\mathbb{1}[x = x_{i^*}]$.

To upper bound eCMI for this prediction rule, we consider the upper bound provided in Theorem H.1. For a fixed supersample $\tilde{Z} = ((x_1, y_1), \dots, (x_{2n}, y_{2n}))$, we will consider two cases:

Case I: For all $i \in [2n]$ we have $y_i = 0$. Let $J \in [2n]_{n+1}$ and let $\deg(\tilde{Z}_J)$ denote the degree of the vertex $(0, 0, \dots, 0, 0)$ in the one-inclusion graph \tilde{Z}_J . Note that if there is no repetition in \tilde{Z}_J then, the degree is $n+1$. If repetition exists, it is less than $n+1$. Then, according to the prediction rule we have

$$\begin{aligned} \kappa(\tilde{Z}) &= \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} P_e(\tilde{Z}_{J-\{j\}}, \tilde{Z}_j) \\ &= \frac{1}{n+1} \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{\deg(\tilde{Z}_J)}{n+1}. \end{aligned} \quad (39)$$

Also,

$$\mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} \text{H}_b(P_e(\tilde{Z}_{J-\{j\}}, \tilde{Z}_j)) = \text{H}_b\left(\frac{1}{n+1}\right) \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{\deg(\tilde{Z}_J)}{n+1}. \quad (40)$$

let $\alpha = \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{\deg(\tilde{Z}_J)}{n+1}$. If $\alpha = 0$, then $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) = 0 = O(1)$. Therefore, assume $\alpha > 0$, then we can use Eq. (38), Eq. (39), and Eq. (40) to obtain

$$\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) \leq n \mathbb{E} \left[\text{H}_b\left(\frac{\alpha}{n+1}\right) + \frac{\alpha \log 2}{n+1} - \alpha \text{H}_b\left(\frac{1}{n+1}\right) \right] \quad (41)$$

$$\begin{aligned} &\leq n \mathbb{E} \left[-\frac{\alpha}{n+1} \log \frac{\alpha}{n+1} - \left(1 - \frac{\alpha}{n+1}\right) \log \left(1 - \frac{\alpha}{n+1}\right) + \frac{\alpha \log 2}{n+1} + \right. \\ &\quad \left. \frac{\alpha}{n+1} \log \left(\frac{1}{n+1}\right) + \alpha \left(1 - \frac{1}{n+1}\right) \log \left(1 - \frac{1}{n+1}\right) \right] \quad (42) \end{aligned}$$

$$\leq n \mathbb{E} \left[-\frac{\alpha}{n+1} \log(\alpha) + \frac{2\alpha \log 2}{n+1} \right] \quad (43)$$

$$\leq n \left[\frac{2 \log 2}{n+1} + \frac{\exp(-1)}{n+1} \right] \quad (44)$$

$$= O(1). \quad (45)$$

Eq. (43) is obtained by removing the negative terms and the inequality $-(1-x) \log(1-x) \leq x \log 2$ for $x \in [0, 1]$. Eq. (44) follows from $\max_{x \in [0, 1]} -x \log(x) = \exp(-1)$ and $\alpha \leq 1$.

Case II: There exists $i \in [2n]$ such that $y_i = 1$. Let $m = |\{i \in [2n] | y_i = 1\}|$, $\beta = \frac{\binom{2n-m}{n+1}}{\binom{2n}{n+1}}$, and $\alpha = \mathbb{E}_J[\frac{\text{deg}(Z_J)}{n+1} | \text{no point with label 1 in } \tilde{Z}_J]$. Then, we have

$$\begin{aligned} \kappa(\tilde{Z}) &= \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} P_e(\tilde{Z}_{J-\{j\}}, \tilde{Z}_j) \\ &= \beta \mathbb{E}_J \left[\frac{1}{n+1} \sum_{j \in J} P_e(\tilde{Z}_{J-\{j\}}, \tilde{Z}_j) \mid \text{no point with label 1 in } \tilde{Z}_J \right] \\ &\quad + (1-\beta) \mathbb{E}_J \left[\frac{1}{n+1} \sum_{j \in J} P_e(\tilde{Z}_{J-\{j\}}, \tilde{Z}_j) \mid \text{there exists points with label 1 in } \tilde{Z}_J \right] \\ &\leq \alpha \beta \frac{1}{n+1} + (1-\beta) \frac{1}{(n+1)^2}, \end{aligned} \quad (46)$$

where in the last step we have used the fact that when in \tilde{Z}_J there is a point with label 1, then the leave-one-out error is at most $\frac{1}{(n+1)^2}$ and $\beta = \mathbb{P}(\text{no point with label 1 in } \tilde{Z}_J) = \frac{\binom{2n-m}{n+1}}{\binom{2n}{n+1}}$.

The binary entropy function is concave and increasing in the interval $[0, 1/2]$. The concavity of the binary entropy function implies that $H_b(v_2) \leq H_b(v_1)'(v_2 - v_1) + H_b(v_1)$ for all v_1 and v_2 in $(0, 1)$. Then,

$$\begin{aligned} H_b(\kappa(\tilde{Z})) &\leq H_b\left(\alpha \beta \frac{1}{n+1} + (1-\beta) \frac{1}{(n+1)^2}\right) \\ &\leq H_b\left(\frac{\alpha \beta}{n+1}\right) + \frac{1-\beta}{(n+1)^2} \log\left(\frac{n+1-\alpha \beta}{\alpha \beta}\right). \end{aligned} \quad (47)$$

Then, by considering the summation only over \tilde{Z}_J without any point with label 1, we obtain

$$\mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} H_b(P_e(\tilde{Z}_{J-\{j\}}, \tilde{Z}_j)) \geq \alpha \beta H_b\left(\frac{1}{n+1}\right). \quad (48)$$

Then, we can substitute Eq. (47) and Eq. (48) into Eq. (38) to obtain $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) = O(1)$ following the same line of reasoning as in Eq. (41)–Eq. (45).

H.2 Proof of Theorem H.1

We begin with introducing some notations. Let Γ_{2n} be the set of all bijective mappings from $[2n]$ to $\{0, 1\} \times [n]$. Let $\sigma \in \Gamma_{2n}$ and $x = (x_1, \dots, x_{2n}) \in \mathcal{X}^{2n}$ be a vector of length $2n$. Then, x^σ denotes a matrix of size $2 \times n$ where $x_{i,j}^\sigma = x_{\sigma^{-1}(i,j)}$ for $i \in \{0, 1\}$ and $j \in [n]$. Also, for every $m, k \in \mathbb{N}$ and $1 \leq k \leq m$, let $[m]_k$ denote the set of all subsets of size k of $[m]$.

For every $n \in \mathbb{N}$ let $\pi \sim \text{Unif}(\Gamma_{2n})$, $\tilde{Z} = (Z_1, \dots, Z_{2n}) \sim \mathcal{D}^{\otimes(2n)}$, and $U = (U_1, \dots, U_n) \sim (\text{Ber}(\{0, 1\}))^{\otimes n}$ where π , \tilde{Z} and U are mutually independent. Let $S_n = (\tilde{Z}_{\bar{U}_j, j}^\pi)_{j=1}^n$ and $\mathcal{A}_n(S_n)$ be a learning algorithm. Let $L \in \{0, 1\}^{2 \times n}$ be a matrix with entries $L_{i,j} = \ell(\mathcal{A}_n(S_n), \tilde{Z}_{i,j}^\pi)$ for $i \in \{0, 1\}$ and $j \in [n]$. Then, using these random variables, we can define $I(L; U | \tilde{Z}, \pi)$.

Lemma H.2. $\text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n)) = I(L; U | \tilde{Z}, \pi)$

Proof. \tilde{Z}^π is a $\sigma(\tilde{Z}, \pi)$ -measurable random variable therefore we have $I(L; U | \tilde{Z}, \pi) = I(L; U | \tilde{Z}^\pi, \tilde{Z}, \pi)$. Then, conditioned on \tilde{Z}^π , L and U are independent from \tilde{Z} and π . Finally, note that the samples in \tilde{Z} are i.i.d., therefore we have $I(L; U | \tilde{Z}^\pi) = \text{eCMI}_{\mathcal{D}}(\ell(\mathcal{A}_n))$. \square

Lemma H.3. Let π , U , and \tilde{Z} be as defined in the beginning of this section. Let $n \in \mathbb{N}$ and $f : \mathcal{Z}^n \times \mathcal{Z} \rightarrow \mathbb{R}$ be a real-valued function where f is permutation-invariant with respect to its first input. Then

$$\mathbb{E}^{\tilde{Z}} \left[f((\tilde{Z}_{\bar{U}_1, 1}^\pi, \dots, \tilde{Z}_{\bar{U}_n, n}^\pi); \tilde{Z}_{\bar{U}_1, 1}^\pi) \right] = \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} f(\tilde{Z}_{J-\{j\}}; \tilde{Z}_j), \quad (49)$$

where $\tilde{Z}_J = (\tilde{Z}_{J_1}, \dots, \tilde{Z}_{J_{n+1}})$ and $\bar{U}_i = 1 - U_i$.

Proof. write

$$\mathbb{E}^{\tilde{Z}} [f(\{\tilde{Z}_{U_1,1}^\pi, \dots, \tilde{Z}_{U_n,n}^\pi\}; \tilde{Z}_{U_1,1}^\pi)] = \sum_{(i_1, \dots, i_{n+1}) \in [2n]_{n+1}} f(\{\tilde{Z}_{i_1}, \dots, \tilde{Z}_{i_n}\}; \tilde{Z}_{i_{n+1}}) \mathbb{P}^{\tilde{Z}}[(\{\tilde{Z}_{U_1,1}^\pi, \dots, \tilde{Z}_{U_n,n}^\pi\}, \tilde{Z}_{U_1,1}^\pi) = (\{\tilde{Z}_{i_1}, \dots, \tilde{Z}_{i_n}\}, \tilde{Z}_{i_{n+1}})].$$

Then,

$$\begin{aligned} \mathbb{P}^{\tilde{Z}}[(\{\tilde{Z}_{U_1,1}^\pi, \dots, \tilde{Z}_{U_n,n}^\pi\}, \tilde{Z}_{U_1,1}^\pi) = (\{\tilde{Z}_{i_1}, \dots, \tilde{Z}_{i_n}\}, \tilde{Z}_{i_{n+1}})] &= \frac{n!(n-1)!}{(2n)!} \\ &= \frac{1}{\binom{2n}{n+1}(n+1)} \end{aligned}$$

where the last line follows since U and π are independent. It is easy to verify that is exactly equal to the RHS of Eq. (49). \square

Let $L = (L_1, \dots, L_n)$ where $L_i \in \{0, 1\}^2$ denotes the vector at Column i of the loss vector L . By the definition and the chain rule for mutual information we have

$$I(L; U | \tilde{Z}, \pi) = \sum_{i=1}^n I(L_i; U | L_1, \dots, L_{i-1}, \tilde{Z}, \pi) \quad (50)$$

$$= \sum_{i=1}^n H(L_i | L_1, \dots, L_{i-1}, \tilde{Z}, \pi) - H(L_i | L_1, \dots, L_{i-1}, \tilde{Z}, \pi, U) \quad (51)$$

$$\leq \sum_{i=1}^n H(L_i | \tilde{Z}, \pi) - H(L_i | \tilde{Z}, \pi, U) \quad (52)$$

$$= n(H(L_1 | \tilde{Z}, \pi) - H(L_1 | \tilde{Z}, \pi, U)). \quad (53)$$

Here, Eq. (52) due to the Markov chain $L_i - (U, \tilde{Z}, \pi) - L_{-i}$ and removing conditions increases the entropy. Then, the last step follows since the algorithm is permutation-invariant. Note that if \mathcal{A} is a deterministic algorithm the second term on the RHS of Eq. (53) is zero.

Next we provide an upper bound for $H(L_1 | \tilde{Z}, \pi)$. Note that L_1 can take values in $\{[1, 0]^\top, [0, 0]^\top, [0, 1]^\top\}$ as it is assumed that \mathcal{A}_n is consistent and, by construction, at each column of \tilde{Z}^π one of the points is selected for the training set. Due to the symmetry imposed by π and U we have with probability one

$$\mathbb{P}^{\tilde{Z}}[L_1 = [1, 0]^\top] = \mathbb{P}^{\tilde{Z}}[L_1 = [0, 1]^\top]. \quad (54)$$

Then, note that there exists a function κ , depending only on \mathcal{A}_n , such that

$$\kappa(\tilde{Z}) = \mathbb{P}^{\tilde{Z}}[L_{U_1,1} = 1]. \quad (55)$$

We claim that the common value in Eq. (54) is given by $\frac{\kappa(\tilde{Z})}{2}$. This claim can be easily proved by taking the expectation with respect to U_1 in Eq. (54) and Eq. (55).

Then, we will show that using Lemma H.3, $\kappa(\tilde{Z})$ is given by the average leave-one-out error of \mathcal{A} over \tilde{Z} . Given a training set $S = ((x_1, y_1), \dots, (x_n, y_n)) \in \mathcal{Z}^n$ and a test point $z = (x, y) \in \mathcal{Z}$, let $P_e(S; z)$ denote the probability that \mathcal{A}_n makes a mistake in predicting the label x . Using this notation, we have

$$\kappa(\tilde{Z}) = \mathbb{E}^{\tilde{Z}} [P_e(\{\tilde{Z}_{U_1,1}^\pi, \dots, \tilde{Z}_{U_n,n}^\pi\}; \tilde{Z}_{U_1,1}^\pi)] \quad (56)$$

$$= \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} P_e(\tilde{Z}_{J-\{j\}}; \tilde{Z}_j), \quad (57)$$

where in the last step we have used Lemma H.3. Then, by the fact that conditioning reduces the entropy we have $H(L_1 | \pi, \tilde{Z}) \leq H(L_1 | \tilde{Z})$. As shown above,

$$\begin{aligned} \mathbb{P}^{\tilde{Z}}[L_1 = [1, 0]^\top] &= \mathbb{P}^{\tilde{Z}}[L_1 = [0, 1]^\top] \\ &= \frac{1}{2} \kappa(\tilde{Z}) \\ &= \frac{1}{2} \mathbb{E}_{J \sim \text{Unif}([2n]_{n+1})} \frac{1}{n+1} \sum_{j \in J} P_e(\tilde{Z}_{J-\{j\}}; \tilde{Z}_j). \end{aligned} \quad (58)$$

Then, by the fact that the L_1 can take values in $\{[1, 0]^\top, [0, 0]^\top, [0, 1]^\top\}$ and Eq. (58) we obtain

$$\begin{aligned} \mathbb{H}(L_1|\tilde{Z}) &= \mathbb{E}[-(1 - \kappa(\tilde{Z})) \log(1 - \kappa(\tilde{Z})) - \kappa(\tilde{Z})/2 \log(\kappa(\tilde{Z})/2) - \kappa(\tilde{Z})/2 \log(\kappa(\tilde{Z})/2)] \\ &= \mathbb{E}[\mathbb{H}_b(\kappa(\tilde{Z})) + \kappa(\tilde{Z}) \log(2)], \end{aligned} \quad (59)$$

where $\mathbb{H}_b(\cdot)$ denotes the binary entropy function.

Finally, we find a closed form expression for $\mathbb{H}(L_1|\tilde{Z}, \pi, U)$. Since it is assumed that \mathcal{A}_n is consistent we have $L_{U_i, i} = 0$ a.s.. Therefore $\mathbb{H}(L_1|\tilde{Z}, \pi, U) = \mathbb{H}(L_{\bar{U}_1, 1}|\tilde{Z}, \pi, U)$. Then, we can write

$$\begin{aligned} \mathbb{H}(L_{\bar{U}_1, 1}|\tilde{Z}, \pi, U) &= \mathbb{E}[\mathbb{H}^{\tilde{Z}, \pi, U}(L_{\bar{U}_1, 1})] \\ &= \mathbb{E}[\mathbb{E}^{\tilde{Z}}[\mathbb{H}^{\tilde{Z}, \pi, U}(L_{\bar{U}_1, 1})]] \\ &= \mathbb{E}[\mathbb{E}^{\tilde{Z}}[\mathbb{H}_b([P_e((\tilde{Z}_{\bar{U}_1, 1}^\pi, \dots, \tilde{Z}_{\bar{U}_n, n}^\pi); \tilde{Z}_{\bar{U}_1, 1}^\pi))]]] \\ &= \mathbb{E}[\mathbb{E}_J \sim \text{Unif}(\{2n\}_{n+1}) \frac{1}{n+1} \sum_{j \in J} \mathbb{H}_b(P_e(\tilde{Z}_{J-\{j\}}; \tilde{Z}_j))], \end{aligned} \quad (60)$$

where in the last line we have used Lemma H.3. Finally by combining Eq. (59), Eq. (60), and Eq. (53) we obtain the stated result in Eq. (38).

I Description of the One-Inclusion Graph Prediction Algorithm

In this part we provide a short description of the one-inclusion transductive learning algorithm of Haussler et al. [5]. Let $\bar{X} = (x_1, \dots, x_n) \in \mathcal{X}^n$ and let \mathcal{H} be a concept class with VC dimension d . Let $\mathcal{H}_{|\bar{X}}$ be the equivalence class induced by \mathcal{H} on the instances given in \bar{X} . Similarly for a classifier $h \in \mathcal{H}$, we can define $h_{|\bar{X}}$ as the restriction of h to the instances in \bar{X} . For every $h \in \mathcal{H}_{|\bar{X}}$, let $v_h = (h(x_1), \dots, h(x_n)) \in \{0, 1\}^n$. Haussler et al. [5] defined the one-inclusion graph of \bar{X} denoted by $\mathcal{G}_{\mathcal{H}}(\bar{X}) = (V, E)$ as follows. $\mathcal{G}_{\mathcal{H}}(\bar{X})$ has the vertex set $V = \{v_h : h \in \mathcal{H}_{|\bar{X}}\}$, and $(v_h, v_{h'}) \in E$ if and only if the hamming distance of v_h and $v_{h'}$ is one. For an example of $\mathcal{G}_{\mathcal{H}}(\bar{X})$ for the concept class of intervals in one dimension, see Fig. 1 of [5]. Consider *probability assignment mapping* $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})} : E \times V \rightarrow [0, 1]$ such that for each edge e incident to v_h and $v_{h'}$ the following two conditions holds.

1. for all $h'' \in \mathcal{H}_{|\bar{X}}$ with $h'' \neq h$ and $h'' \neq h'$, we have $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_{h''}) = 0$.
2. $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_h) \geq 0$, $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_{h'}) \geq 0$, and $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_h) + f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_{h'}) = 1$.

For every $i \in [n]$ and $h \in \mathcal{H}_{|\bar{X}}$, let $c_{i, h} \subset \mathcal{H}_{|\bar{X}}$ be the set of all the hypotheses in $\mathcal{H}_{|\bar{X}}$ whose restriction to $\bar{X} \setminus \{x_i\}$ equals to $h_{|\bar{X} \setminus \{x_i\}}$. Let $h^* \in \mathcal{H}$, consider a realizable $S = ((x_1, y_1), \dots, (x_n, y_n))$ where $y_i = h^*(x_i)$ for $i \in [n]$. Assume $S_{n-1} = ((x_1, y_1), \dots, (x_{n-1}, y_{n-1}))$ is given to a learner and the learner aims to predict the label of x_n . It is immediate to see that the set of hypotheses consistent with S_{n-1} is c_{n, h^*} . Clearly, $|c_{n, h^*}| \in \{1, 2\}$ and if $|c_{n, h^*}| = 1$, we know that the target hypothesis is h^* . But, what should the learner do when $|c_{n, h^*}| = 2$?

Using $\mathcal{G}_{\mathcal{H}}(\bar{X})$ we can think of the case $|c_{n, h^*}| = 2$ as there is a vertex $v_{h'}$ adjacent to v_{h^*} , and $v_{h'}$ and v_{h^*} differ in n -th position. Assume $|c_{n, h^*}| = 2$ and $e_{n, h^*} = \{h^*, h'\}$. Using the probability assignment mapping $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}$ of $\mathcal{G}_{\mathcal{H}}(\bar{X})$, the strategy proposed by Haussler et al. [5] predicts the label x_n to be $h'(x_n)$ with probability $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_{h^*})$, and to be $h^*(x_n)$ with probability $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_{h'})$. By identifying a deep combinatorial property of $\mathcal{G}_{\mathcal{H}}(\bar{X})$, Thm. 2.3 in [5] shows that there exists a probability assignment mapping $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}$ with the mentioned properties such that $\sum_{e \in E} f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_h) \leq d$ for all $h \in \mathcal{G}_{\mathcal{H}}(\bar{X})$, and finding such a mapping is computationally easy. Moreover, Haussler et al. [5] show there exists *deterministic* probability assignment for every one-inclusion graph such that $f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_h) \in \{0, 1\}$ for all $e \in E$ and $v_h \in V$, and $\sum_{e \in E} f_{\mathcal{G}_{\mathcal{H}}(\bar{X})}(e, v_h) \leq d$ for all $h \in \mathcal{G}_{\mathcal{H}}(\bar{X})$.