
A Theory of PAC Learnability under Transformation Invariances

Han Shao

Toyota Technological Institute Chicago
Chicago, 60637
han@ttic.edu

Omar Montasser

Toyota Technological Institute Chicago
Chicago, 60637
omar@ttic.edu

Avrim Blum

Toyota Technological Institute Chicago
Chicago, 60637
avrim@ttic.edu

Abstract

Transformation invariances are present in many real-world problems. For example, image classification is usually invariant to rotation and color transformation: a rotated car in a different color is still identified as a car. Data augmentation, which adds the transformed data into the training set and trains a model on the augmented data, is one commonly used technique to build these invariances into the learning process. However, it is unclear how data augmentation performs theoretically and what the optimal algorithm is in presence of transformation invariances. In this paper, we study PAC learnability under transformation invariances in three settings according to different levels of realizability: (i) A hypothesis fits the augmented data; (ii) A hypothesis fits only the original data and the transformed data lying in the support of the data distribution; (iii) Agnostic case. One interesting observation is that distinguishing between the original data and the transformed data is necessary to achieve optimal accuracy in setting (ii) and (iii), which implies that any algorithm not differentiating between the original and transformed data (including data augmentation) is not optimal. Furthermore, this type of algorithms can even “harm” the accuracy. In setting (i), although it is unnecessary to distinguish between the two data sets, data augmentation still does not perform optimally. Due to such a difference, we propose two combinatorial measures characterizing the optimal sample complexity in setting (i) and (ii)(iii) and provide the optimal algorithms.

1 Introduction

Transformation invariances are present in many real-world learning problems. That is, given a certain set of transformations, the label of an instance is preserved under any transformation from the set. Image classification is often invariant to rotation/flip/color translation. Syntax parsing is invariant to exchange of noun phrases in a sentence. Such invariances are often built into the learning process by two ways. One is designing new architectures in neural networks to learn a transformation invariant feature, which is usually task-specific and challenging. A more universally applicable and easier way is data augmentation (DA)¹, that is, adding the transformed data into the training set and training a model with the augmented data. Although DA performs well empirically, it is unclear whether and when DA “helps”. In this paper, we focus on answering two questions:

¹Throughout the paper, we refer to ERM over the augmented data by DA.

How does data augmentation perform theoretically?

What is the optimal algorithm in terms of sample complexity under transformation invariances?

We formalize the problem of binary classification under transformation invariances in the PAC model. Given instance space \mathcal{X} , label space $\mathcal{Y} = \{0, 1\}$, and hypothesis class \mathcal{H} , we consider the following three settings according to different levels of realizability.

- (i) Invariantly realizable setting: There exists a hypothesis $h^* \in \mathcal{H}$ such that h^* can correctly classify not only the natural data (drawn from the data distribution) but also the transformed data. For example, considering the transformation of rotating images where all natural images are upright, the hypothesis h^* can correctly classify every upright image (natural data) and their rotations (transformed data).
- (ii) Relaxed realizable setting: There exists a hypothesis $h^* \in \mathcal{H}$ such h^* has zero error over the support of the data distribution (and therefore will correctly classify the transformed data that lies in the support of the data distribution), but h^* may not correctly classify transformed data that lies outside the support of the natural data distribution. For example, there exists an h^* classifying all small rotations that lie in the support of the distribution correctly, but misclassifying upside-down cars.
- (iii) Agnostic setting: Every hypothesis in \mathcal{H} might not fit the natural data.

In most of this work, we consider the case where the set of transformations forms a group (e.g., all rotations and all color translations), which is a classic setting studied in literature (e.g., Cohen and Welling, 2016; Bloem-Reddy and Teh, 2020; Chen et al., 2020). Some algorithms and analyses in this work also apply to non-group transformations (e.g., croppings).

Main contributions First, we show that DA outperforms vanilla ERM but is sub-optimal in setting (i) above. We then introduce a complexity measure (see Definition 4) that characterizes the optimal sample complexity of learning in setting (i), and we give an optimal (up to log-factors) algorithm in this setting based on 1-inclusion-graph predictors. Second, we characterize the complexity of learning in setting (ii) when the learner only receives the augmented data (without specifying which are natural). Such a characterization provides us with a sufficient condition under which DA "hurts". Third, we introduce a complexity measure (see Definition 5) that characterizes the optimal sample complexity of learning in settings (ii) and (iii) above, and we give optimal algorithms for these settings. Finally, we also provide adaptive learning algorithms that interpolate between settings (i) and (ii), i.e., when h^* is partially invariant. We want to emphasize that our complexity measures take into account the complexity of both the hypothesis class \mathcal{H} and the set of transformations being considered. The results are formally summarized in Section 3.

Related work Theoretical guarantees of DA has received a lot of attention recently. Chen et al. (2020); Lyle et al. (2020) study theoretical guarantees of DA under the assumption of "equality" in distribution, i.e., for any transformation in the transformation group, the data distribution of the transformed data is approximately the same as that of the natural data (e.g., the upside-down variations of images happen at the same probability as the original upright images). Under this assumption, they show that DA reduces variance and induces better generalization error upper bounds. Our work does not make such an assumption. Dao et al. (2019) models augmentation as a Markov process and shows that for kernel linear classifiers, DA can be approximated by first-order feature averaging and second-order variance regularization components. The concurrent work by Shen et al. (2022) studies the benefit of DA when training a two layer convolutional neural network in a specific multi-view model, showing that DA can alter the relative importance of various features. There is a line of theoretical study on the invariance gain in different models. For example, Elesedy and Zaidi (2021) study the linear model and Elesedy (2021); Mei et al. (2021); Bietti et al. (2021) study the non-parametric regression. The concurrent work by Elesedy (2022) also studies PAC learning under transformation invariances but only provides an upper bound on the sample complexity, while our work provides a complete characterization of learning under this model with optimal algorithms. There is a parallel line of theoretical study on architecture design (e.g., Wood and Shawe-Taylor, 1996; Ravanbakhsh et al., 2017; Kondor and Trivedi, 2018; Bloem-Reddy and Teh, 2020).

Learning under transformation invariances has also been studied a lot empirically. Here we briefly mention a few results. DA has been applied as standard method in modern deep learning, e.g., in Alexnet (Krizhevsky et al., 2012). Gontijo-Lopes et al. (2020) proposes two measures, affinity and diversity, to quantify the performance of the existing DA methods. Fawzi et al. (2016); Cubuk

et al. (2018); Chatzipantazis et al. (2021) study how to automatically search for improved data augmentation policies. For architecture design, one celebrated example is convolutions (Fukushima and Miyake, 1982; LeCun et al., 1989), which are translation equivariant. See Cohen and Welling (2016); Dieleman et al. (2016); Worrall et al. (2017) for more different architectures invariant or equivariant to different symmetries.

Another line of related work is adversarial training, which adds the perturbed data into the training set and can be considered as a special type of data augmentation. Raghunathan et al. (2019); Schmidt et al. (2018); Nakkiran (2019) study the standard accuracy of adversarial training and provide examples showing that adversarial training can sometimes “harm” standard accuracy.

Notation For any $n \in \mathbb{N}$, let e_1, e_2, \dots denote the standard basis vectors in \mathbb{R}^n . For any set \mathcal{V} and any $\mathbf{v} \in \mathcal{V}^n$, let $\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n) \in \mathcal{V}^{n-1}$ denote the remaining part of \mathbf{v} after removing the i -th entry and $(v', \mathbf{v}_{-i}) = (v_1, \dots, v_{i-1}, v', v_{i+1}, \dots, v_n) \in \mathcal{V}^n$ denote the vector after replacing i -th entry of \mathbf{v} with $v' \in \mathcal{V}$. Let \oplus denote the bitwise XOR operator. For any $h \in \mathcal{Y}^{\mathcal{X}}$ and $X = \{x_1, \dots, x_n\} \subset \mathcal{X}$, denote $h|_X = (h(x_1), \dots, h(x_n))$ the restriction of h on X . A data set or a sample is a multiset of $\mathcal{X} \times \mathcal{Y}$. For any sample S , let $S_{\mathcal{X}} = \{x | (x, y) \in S\}$ (with multiplicity) and for any distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$, for $(x, y) \sim \mathcal{D}$, let $\mathcal{D}_{\mathcal{X}}$ denote the marginal distribution of x . For any data distribution \mathcal{D} and any hypothesis h , the expected error $\text{err}_{\mathcal{D}}(h) := \Pr_{(x,y) \sim \mathcal{D}}(h(x) \neq y)$. Denote $\text{err}(h) = \text{err}_{\mathcal{D}}(h)$ when \mathcal{D} is clear from the context. For any sample S of finite size, $\text{err}_S(h) := \frac{1}{|S|} \sum_{(x,y) \in S} \mathbb{1}[h(x) \neq y]$. For any sample S of possibly of infinite size, we say $\text{err}_S(h) = 0$ if $h(x) = y$ for all $(x, y) \in S$.

2 Problem setup

We study binary classification under transformation invariances. We denote by \mathcal{X} the instance space, $\mathcal{Y} = \{0, 1\}$ the label space and \mathcal{H} the hypothesis class.

Group transformations We consider a group \mathcal{G} of transformations acting on the instance space through a mapping $\alpha : \mathcal{G} \times \mathcal{X} \mapsto \mathcal{X}$, which is compatible with the group operation. For convenience, we write $\alpha(g, x) = gx$ for $g \in \mathcal{G}$ and $x \in \mathcal{X}$. For example, consider $\mathcal{G} = \{e, g_1, g_2, g_3\}$ where e is the identify function and g_i is rotation by $90i$ degrees. Given an image x , $ex = x$ is the original image and g_1x is the image rotated by 90 degrees. The orbit of any $x \in \mathcal{X}$ is the subset of \mathcal{X} that can be obtained by acting an element in \mathcal{G} on x , $\mathcal{G}x := \{gx | g \in \mathcal{G}\}$. Note that since \mathcal{G} is a group, for any $x' \in \mathcal{G}x$, we have $\mathcal{G}x' = \mathcal{G}x$. Thus we can divide the instance space \mathcal{X} into a collection of separated orbits, which does not depend on the data distribution. Given a (natural) data set $S \subset \mathcal{X} \times \mathcal{Y}$, we call $\mathcal{G}S := \{(gx, y) | (x, y) \in S, g \in \mathcal{G}\}$ the augmented data set.

Transformation invariant hypotheses and distributions To model transformation invariance, we assume that the true labels are invariant over the orbits of natural data. Formally, for any transformation group \mathcal{G} and $X \subset \mathcal{X}$, we say a hypothesis h is (\mathcal{G}, X) -invariant if

$$h(gx) = h(x), \forall g \in \mathcal{G}, x \in X.$$

That is to say, for every $x \in X$, h predicts every instance in the orbit of x the same as x . For any marginal distribution $\mathcal{D}_{\mathcal{X}}$ over \mathcal{X} , we say a hypothesis h is $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant if $h(gx) = h(x)$ for all $g \in \mathcal{G}$, for all $x \in \text{supp}(\mathcal{D}_{\mathcal{X}})$, i.e., $\Pr_{x \sim \mathcal{D}_{\mathcal{X}}}(\exists x' \in \mathcal{G}x : h(x') \neq h(x)) = 0$. We say a distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ is \mathcal{G} -invariant if there exists a $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant hypothesis f^* (possibly not in \mathcal{H}) with $\text{err}_{\mathcal{D}}(f^*) = 0$. We assume that the data distribution is \mathcal{G} -invariant throughout the paper.

Realizability of hypothesis class We consider three settings according to the different levels of realizability of \mathcal{H} : (i) invariantly realizable setting, where there exists a $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant hypothesis $h^* \in \mathcal{H}$ with $\text{err}_{\mathcal{D}}(h^*) = 0$; (ii) relaxed realizable setting, where there exists a (not necessarily $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant) hypothesis $h^* \in \mathcal{H}$ with $\text{err}_{\mathcal{D}}(h^*) = 0$; and (iii) agnostic setting, where there might not exist a hypothesis in \mathcal{H} with zero error. To understand the difference among the three settings, here is an example.

Example 1. Consider $\mathcal{X} = \{\pm 1, \pm 2\}$, $\mathcal{G} = \{e, -e\}$ being the group generated by flipping the sign (i.e., $\mathcal{G}x = \{x, -x\}$), and the data distribution \mathcal{D} being the uniform distribution over $\{(1, 0), (2, 0)\}$. If $\mathcal{H} = \{h(\cdot) = 0\}$ contains only the all-zero function, then it is in setting (i) as $h(\cdot) = 0$ is $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant and $\text{err}_{\mathcal{D}}(h) = 0$; If $\mathcal{H} = \{\mathbb{1}[x < 0]\}$ contains only the hypothesis predicting $\{-1, -2\}$ as 1 and $\{1, 2\}$ as 0, then it is in setting (ii) as $\mathbb{1}[x < 0]$ is not $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant but $\text{err}_{\mathcal{D}}(\mathbb{1}[x < 0]) = 0$; If $\mathcal{H} = \{\mathbb{1}[x > 0]\}$, it is in setting (iii) as no hypothesis in \mathcal{H} has zero error.

The following definitions formalize the notion of PAC learning in the three settings.

Definition 1 (Invariantly realizable PAC learnability). *For any $\epsilon, \delta \in (0, 1)$, the sample complexity of invariantly realizable (ϵ, δ) -PAC learning of \mathcal{H} with respect to transformation group \mathcal{G} , denoted $\mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$, is defined as the smallest $m \in \mathbb{N}$ for which there exists a learning rule \mathcal{A} such that, for every \mathcal{G} -invariant data distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ where there exists a $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant predictor $h^* \in \mathcal{H}$ with zero error; $\text{err}_{\mathcal{D}}(h^*) = 0$, with probability at least $1 - \delta$ over $S \sim \mathcal{D}^m$,*

$$\text{err}_{\mathcal{D}}(\mathcal{A}(S)) \leq \epsilon.$$

If no such m exists, define $\mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) = \infty$. We say that \mathcal{H} is PAC learnable in the invariantly realizable setting with respect to transformation group \mathcal{G} if $\forall \epsilon, \delta \in (0, 1)$, $\mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ is finite. For any algorithm \mathcal{A} , denote by $\mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}, \mathcal{A})$ the sample complexity of \mathcal{A} .

Definition 2 (Relaxed realizable PAC learnability). *For any $\epsilon, \delta \in (0, 1)$, the sample complexity of relaxed realizable (ϵ, δ) -PAC learning of \mathcal{H} with respect to transformation group \mathcal{G} , denoted $\mathcal{M}_{\text{RE}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$, is defined as the smallest $m \in \mathbb{N}$ for which there exists a learning rule \mathcal{A} such that, for every \mathcal{G} -invariant data distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ where there exists a predictor $h^* \in \mathcal{H}$ with zero error; $\text{err}_{\mathcal{D}}(h^*) = 0$, with probability at least $1 - \delta$ over $S \sim \mathcal{D}^m$,*

$$\text{err}_{\mathcal{D}}(\mathcal{A}(S)) \leq \epsilon.$$

If no such m exists, define $\mathcal{M}_{\text{RE}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) = \infty$. We say that \mathcal{H} is PAC learnable in the relaxed realizable setting with respect to transformation group \mathcal{G} if $\forall \epsilon, \delta \in (0, 1)$, $\mathcal{M}_{\text{RE}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ is finite. For any algorithm \mathcal{A} , denote by $\mathcal{M}_{\text{RE}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}, \mathcal{A})$ the sample complexity of \mathcal{A} .

Definition 3 (Agnostic PAC learnability). *For any $\epsilon, \delta \in (0, 1)$, the sample complexity of agnostic (ϵ, δ) -PAC learning of \mathcal{H} with respect to transformation group \mathcal{G} , denoted $\mathcal{M}_{\text{AG}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$, is defined as the smallest $m \in \mathbb{N}$ for which there exists a learning rule \mathcal{A} such that, for every \mathcal{G} -invariant data distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$, with probability at least $1 - \delta$ over $S \sim \mathcal{D}^m$,*

$$\text{err}_{\mathcal{D}}(\mathcal{A}(S)) \leq \inf_{h \in \mathcal{H}} \text{err}_{\mathcal{D}}(h) + \epsilon.$$

If no such m exists, define $\mathcal{M}_{\text{AG}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) = \infty$. We say that \mathcal{H} is PAC learnable in the agnostic setting with respect to transformation group \mathcal{G} if $\forall \epsilon, \delta \in (0, 1)$, $\mathcal{M}_{\text{AG}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ is finite.

Data augmentation One main goal of this work is to analyze the sample complexity of data augmentation. When we talk of data augmentation (DA) as an algorithm, it actually means ERM over the augmented data. Specifically, given a fixed loss function \mathcal{L} mapping a data set and a hypothesis to $[0, 1]$, and a training set S_{trn} , DA outputs an $h \in \mathcal{H}$ such that $h(x) = y$ for all $(x, y) \in \mathcal{G}S_{\text{trn}}$ if there exists one; outputs a hypothesis $h \in \mathcal{H}$ with the minimal loss $\mathcal{L}(\mathcal{G}S_{\text{trn}}, h)$ otherwise. When we say DA without specifying the loss function, it means DA w.r.t. an arbitrary loss function, which can be defined based on any probability measure on the transformation group.

To characterize sample complexities, we define two measures as follows.

Definition 4 (VC dimension of orbits). *The VC dimension of orbits, denoted $\text{VC}_o(\mathcal{H}, \mathcal{G})$, is defined as the largest integer k for which there exists a set $X = \{x_1, \dots, x_k\} \subset \mathcal{X}$ such that their orbits are pairwise disjoint, i.e., $\mathcal{G}x_i \cap \mathcal{G}x_j = \emptyset, \forall i, j \in [k]$ and every labeling of X is realized by a (\mathcal{G}, X) -invariant hypothesis in \mathcal{H} , i.e., $\forall y \in \{0, 1\}^k$, there exists a (\mathcal{G}, X) -invariant hypothesis $h \in \mathcal{H}$ s.t. $h(x_i) = y_i, \forall i \in [k]$.*

Definition 5 (VC dimension across orbits). *The VC dimension across orbits, denoted $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$, is defined as the largest integer k for which there exists a set $X = \{x_1, \dots, x_k\} \subset \mathcal{X}$ such that their orbits are pairwise disjoint, i.e., $\mathcal{G}x_i \cap \mathcal{G}x_j = \emptyset, \forall i, j \in [k]$ and every labeling of X is realized by a hypothesis in \mathcal{H} , i.e., $\forall y \in \{0, 1\}^k$, there exists a hypothesis $h \in \mathcal{H}$ s.t. $h(x_i) = y_i, \forall i \in [k]$.*

Let $\text{VCdim}(\mathcal{H})$ denote the VC dimension of \mathcal{H} . By definition, it is direct to check that $\text{VC}_o(\mathcal{H}, \mathcal{G}) \leq \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) \leq \text{VCdim}(\mathcal{H})$. For any \mathcal{H}, \mathcal{G} with $\text{VC}_o(\mathcal{H}, \mathcal{G}) = d$, we can supplement \mathcal{H} to a new hypothesis class \mathcal{H}' such that $\text{VC}_o(\mathcal{H}', \mathcal{G})$ is still d while $\text{VC}_{\text{ao}}(\mathcal{H}', \mathcal{G})$ is as large as the total number of orbits with at least two instances, i.e., $\text{VC}_{\text{ao}}(\mathcal{H}', \mathcal{G}) = |\{\mathcal{G}x \mid |\mathcal{G}x| \geq 2, x \in \mathcal{X}\}|$. This can be done by supplementing \mathcal{H} with all hypotheses predicting $\mathcal{G}x$ with two different labels for all x with $|\mathcal{G}x| \geq 2$. Besides, for any \mathcal{H} with $\text{VCdim}(\mathcal{H}) = d$, we can construct a transformation group \mathcal{G} to make all instances lie in one single orbit, which makes $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) \leq 1$. Hence the gap among the three measures can be arbitrarily large. Here are a few examples for better understanding of the gaps.

Example 2. Consider $\mathcal{X} = \{\pm 1, \pm 2, \dots, \pm 2d\}$ for some $d > 0$, $\mathcal{H} = \{\mathbb{1}[x \in A] \mid A \subset [2d] \text{ and } |A| = d\}$ being the set of all hypotheses labeling exact d elements from $[2d]$ as 1 and $\mathcal{G} = \{e, -e\}$ being the group generated by flipping the sign. Then we have $\text{VC}_o(\mathcal{H}, \mathcal{G}) = 0$ since for any $i \in [2d]$, there is no $(\mathcal{G}, \{i\})$ -invariant hypothesis that can label i as 1 (which is due to the fact that $-i$ is labeled as 0 by any hypothesis in \mathcal{H}). It is direct to check that $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = \text{VCdim}(\mathcal{H}) = d$.

Example 3. Consider $\mathcal{X} = \{x \in \mathbb{R}^2 \mid \|x\|_2 = 1\}$ being a circle, the hypothesis class $\mathcal{H} = \{0, 1\}^{\mathcal{X}}$ being all labeling functions and \mathcal{G} being all rotations (thus $\forall x, \mathcal{G}x = \mathcal{X}$). Then we have $\text{VC}_o(\mathcal{H}, \mathcal{G}) = \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = 1$ as there is only one orbit and $\text{VCdim}(\mathcal{H}) = \infty$.

Example 4. Consider the natural data being k upright images and the transformation set \mathcal{G} is rotation by $0, 360/n, 2 \cdot 360/n, \dots, (n-1) \cdot 360/n$ degrees for some integer n . For an expressive hypothesis class \mathcal{H} (e.g., neural networks) that can shatter all rotated versions of these images, we have $\text{VCdim}(\mathcal{H}) = nk$ and $\text{VC}_o(\mathcal{H}, \mathcal{G}) = \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = k$. For a hypothesis class \mathcal{H}' composed of all hypotheses labeling all upright images and their upside-down variations differently, we have $\text{VCdim}(\mathcal{H}') = (n-1)k$, $\text{VC}_{\text{ao}}(\mathcal{H}', \mathcal{G}) = k$ and $\text{VC}_o(\mathcal{H}', \mathcal{G}) = 0$.

3 Main results

We next present and discuss our main results.

- Invariantly realizable setting (Definition 1)
 - **DA “helps” but is not optimal. The sample complexity of DA is characterized by $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$.** For any \mathcal{H}, \mathcal{G} , DA can learn \mathcal{H} with sample complexity $\tilde{O}\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta}\right)$, where \tilde{O} ignores log-factors of $\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon}$ (Theorem 1). For all $d > 0$, there exists \mathcal{H}, \mathcal{G} with $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = d$ and $\text{VC}_o(\mathcal{H}, \mathcal{G}) = 0$ such that DA needs $\Omega\left(\frac{d}{\epsilon}\right)$ samples (Theorem 2).
 - **The optimal sample complexity is characterized by $\text{VC}_o(\mathcal{H}, \mathcal{G})$.** For any \mathcal{H}, \mathcal{G} , we have $\Omega\left(\frac{\text{VC}_o(\mathcal{H}, \mathcal{G})}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta}\right) \leq \mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) \leq O\left(\frac{\text{VC}_o(\mathcal{H}, \mathcal{G})}{\epsilon} \log \frac{1}{\delta}\right)$ (Theorem 4). We propose an algorithm achieving this upper bound based on 1-inclusion graphs, which does not distinguish between the original and transformed data. It is worth noting that the algorithm takes the invariance over the test point into account, which provides some theoretical justification for test-time adaptation such as Wang et al. (2021).
- Relaxed realizable setting (Definition 2)
 - **DA can “hurt”.** DA belongs to the family of algorithms not distinguishing the original data from the transformed data. We show that the optimal sample complexity of this family is characterized by $\mu(\mathcal{H}, \mathcal{G})$ (see Definition 6) (Theorem 5), which can be arbitrarily larger than $\text{VCdim}(\mathcal{H})$. This implies that for any \mathcal{H}, \mathcal{G} with $\mu(\mathcal{H}, \mathcal{G}) > \text{VCdim}(\mathcal{H})$, the sample complexity of DA is higher than that of ERM.
 - **The optimal sample complexity is characterized by $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$.** For any \mathcal{H}, \mathcal{G} , we have $\Omega\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} + \frac{\log(1/\delta)}{\epsilon}\right) \leq \mathcal{M}_{\text{RE}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) \leq \tilde{O}\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta}\right)$ (Theorem 7). We propose two algorithms achieving similar upper bounds, with one based on ERM and one based on 1-inclusion graphs. Both algorithms have to distinguish between the original and the transformed data.
 - **An adaptive algorithm interpolates between two settings.** We present an algorithm that adapts to different levels of invariance of the target function h^* , which achieves $\tilde{O}\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon}\right)$ sample complexity in the relaxed realizable setting and $\tilde{O}\left(\frac{\text{VC}_o(\mathcal{H}, \mathcal{G})}{\epsilon}\right)$ sample complexity in the invariantly realizable setting without knowing it (Theorem 9 in Appendix).
- Agnostic setting (Definition 3)
 - **The optimal sample complexity is characterized by $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$.** For any \mathcal{H}, \mathcal{G} , $\mathcal{M}_{\text{AG}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) = O\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon^2} \log^2\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon}\right) + \frac{1}{\epsilon^2} \log\left(\frac{1}{\delta}\right)\right)$ (Theorem 8). Since $\mathcal{M}_{\text{AG}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) \geq \mathcal{M}_{\text{RE}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$, $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$ characterizes the optimal sample complexity.

4 Invariantly Realizable setting

In this section, we discuss the results in the invariantly realizable setting (see Definition 1).

4.1 DA “helps” but is not optimal

We show that in the invariantly realizable setting, DA indeed “helps” to improve the sample complexity from $\tilde{O}(\frac{\text{VCdim}(\mathcal{H})}{\epsilon})$ (the sample complexity of ERM in standard PAC learning) to $\tilde{O}(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon})$. First, we have the following upper bound on the sample complexity of DA.

Theorem 1. *For any \mathcal{H}, \mathcal{G} with $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) < \infty$, DA satisfies that $\mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}, \text{DA}) = O(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} \log^3 \frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta})$.*

Intuitively, for a set of instances in one orbit that can be labeled by \mathcal{H} in multiple ways, we only need to observe one instance from this orbit to learn the labels of all the instances by applying DA. Thus, DA helps to improve the accuracy. The detailed proof is deferred to Appendix A. However, DA does not fully exploit the transformation invariances as it only utilizes the invariances of the training set. Hence, DA does not perform optimally in presence of the transformation invariances. In fact, besides DA, all proper learners (i.e., learners outputting a hypothesis in \mathcal{H}) have the same problem.

Theorem 2. *For any $d > 0$, there exists a hypothesis class \mathcal{H}_d and a group \mathcal{G}_d with $\text{VC}_{\text{ao}}(\mathcal{H}_d, \mathcal{G}_d) = d$ and $\text{VC}_o(\mathcal{H}_d, \mathcal{G}_d) = 0$ such that $\mathcal{M}_{\text{INV}}(\epsilon, \frac{1}{9}; \mathcal{H}_d, \mathcal{G}_d, \mathcal{A}) = \Omega(\frac{d}{\epsilon})$ for any proper learner \mathcal{A} , including DA and standard ERM.*

The theorem shows that DA is sub-optimal as we will show that the optimal sample complexity is characterized by $\text{VC}_o(\mathcal{H}, \mathcal{G})$ in Theorem 4. We provide an idea of the construction here and defer the detailed proof to Appendix B. Consider the \mathcal{X}, \mathcal{H} and \mathcal{G} in Example 2. Pick the target function $\mathbb{1}[x \in A]$ uniformly at random from \mathcal{H} and let the data distribution only put probability mass on points in $[2d] \setminus A$, the orbits of which are labeled as 0 by the target function. Then any proper learner must predict d unobserved examples of $[2d]$ as 1, which leads to high error if the learner observes fewer than $d/2$ examples. Theorem 2 also implies that for any hypothesis class including \mathcal{H} as a subset, there exists a DA learner (i.e., a proper learner fitting the augmented data) whose sample complexity is $\Omega(\frac{d}{\epsilon})$.

Theorem 1 shows that the sample complexity of DA is $\tilde{O}(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon})$, better than that of ERM in standard PAC learning, $\tilde{O}(\frac{\text{VCdim}(\mathcal{H})}{\epsilon})$. This is insufficient to show that DA outperforms ERM as it might be possible that ERM can also achieve better sample complexity in presence of transformation invariances. To illustrate that DA indeed “helps”, we show that any algorithm without exploiting the transformation invariances still requires sample complexity of $\Omega(\frac{\text{VCdim}(\mathcal{H})}{\epsilon})$.

Theorem 3. *For any \mathcal{H} , there exists a group \mathcal{G} with $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) \leq 5$ s.t. $\mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}, \mathcal{A}) = \Omega(\frac{\text{VCdim}(\mathcal{H})}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta})$ for any algorithm \mathcal{A} not given any information about \mathcal{G} (e.g., ERM).*

The basic idea is that, given a set of k instances that can be shattered by \mathcal{H} for some $k > 0$, \mathcal{G} is uniformly at random picked from a set of 2^k groups, each of which partitions the set into two orbits in a different way. If given \mathcal{G} , the algorithm only need to observe one instance in each orbit to learn the labels of all k instances. If not, the algorithm can only randomly guess the label of every unobserved instance. The detailed construction is included in Appendix C.

4.2 The optimal algorithm

We show that the optimal sample complexity is characterized by $\text{VC}_o(\mathcal{H}, \mathcal{G})$.

Theorem 4. *For any \mathcal{H}, \mathcal{G} with $\text{VC}_o(\mathcal{H}, \mathcal{G}) < \infty$, we have $\Omega(\frac{\text{VC}_o(\mathcal{H}, \mathcal{G})}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta}) \leq \mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) \leq O(\frac{\text{VC}_o(\mathcal{H}, \mathcal{G})}{\epsilon} \log \frac{1}{\delta})$.*

Our algorithm is based on the 1-inclusion-graph predictor by Haussler et al. (1994). Given hypothesis class H and instance space $X = \{x_1, \dots, x_t\}$, the classical 1-inclusion-graph consists of vertices $\{h|_X | h \in H\}$, which are labelings of X realized by H , and two vertices are connected by an edge if and only if they differ at the labeling of exactly a single $x_i \in X$. Haussler et al. (1994) shows that the edges can be oriented such that each vertex has in-degree at most $\text{VCdim}(H)$. This orientation can be translated to a prediction rule. Specifically, for any $i \in [t]$, given the labels of all instances in X except x_i , if there are two hypotheses $h, h' \in H$ such that their labelings are consistent with the labels of $X \setminus \{x_i\}$ and different at x_i , then $h|_X, h'|_X$ are two vertices in the graph and we predict

the label of x_i as the edge between $h|_X, h'|_X$ is oriented against. The average leave-one-out-error is upper bounded by $\frac{\text{VCdim}(H)}{t}$.

Lemma 1 (Theorem 2.3 of Haussler et al. (1994)). *For any hypothesis class H and instance space X with $\text{VCdim}(H) < \infty$, there is a function $Q : (X \times \mathcal{Y})^* \times X \mapsto \mathcal{Y}$ such that, for any $t \in \mathbb{N}$ and sample $\{(x_1, y_1), \dots, (x_t, y_t)\}$ that is realizable w.r.t. H ,*

$$\frac{1}{t!} \sum_{\sigma \in \text{Sym}(t)} \mathbb{1}[Q(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [t-1]}, x_{\sigma(t)}) \neq y_{\sigma(t)}] \leq \frac{\text{VCdim}(H)}{t}, \quad (1)$$

where $\text{Sym}(t)$ denotes the symmetric group on $[t]$. The function Q can be constructed by a 1-inclusion-graph predictor.

Denote by $Q_{H,X}$ the function guaranteed by Eq (1) for hypothesis class H and instance space X . For any $t \in \mathbb{N}$ and $S = \{(x_1, y_1), \dots, (x_t, y_t)\}$, let X_S denote the set of different elements in S_X . Define $\mathcal{H}(X_S) := \{h|_{X_S} | h \in \mathcal{H} \text{ is } (\mathcal{G}, X_S)\text{-invariant}\}$ being the set of all possible (\mathcal{G}, X_S) -invariant labelings of X_S . We then define our algorithm $\mathcal{A}(S)$ by letting $\mathcal{A}(S)(x) = Q_{\mathcal{H}(X_S \cup \{x\}), X_S \cup \{x\}}(S, x)$ if $\mathcal{H}(X_S \cup \{x\}) \neq \emptyset$ and predicting arbitrarily if $\mathcal{H}(X_S \cup \{x\}) = \emptyset$. That is to say, $\mathcal{A}(S)$ needs to construct a function Q for every test example. Given any test example, this 1-inclusion-graph-based algorithm takes into account whether the prediction can be invariant over the whole orbit of the test example and thus benefits from the invariance of test examples. This can provide some theoretical justification for test-time adaptation such as Wang et al. (2021). By definition, we have $\text{VCdim}(\mathcal{H}(X_S \cup \{x\})) \leq \text{VC}_o(\mathcal{H}, \mathcal{G})$ for all S and x . Then the expected error of \mathcal{A} can be bounded by $\text{VC}_o(\mathcal{H}, \mathcal{G})$ through Lemma 1. We defer the details and the proof of Theorem 4 to Appendix D. Note that the results of Theorem 4 also apply to non-group transformations².

5 Relaxed realizable setting

In this section, we discuss the results in the relaxed realizable setting (see Definition 2). As we can see, DA belongs to the family of algorithms not distinguishing between the original and transformed data. In Section 5.1, we provide a tight characterization $\mu(\mathcal{H}, \mathcal{G})$ (Definition 6) on the sample complexity of this family algorithms. This implies that when $\mu(\mathcal{H}, \mathcal{G}) > \text{VCdim}(\mathcal{H})$, there exists a distribution s.t. DA performs worse than ERM. We then show that there exists \mathcal{H}, \mathcal{G} such that $\mu(\mathcal{H}, \mathcal{G}) > \text{VCdim}(\mathcal{H})$ and the gap can be arbitrarily large. In Section 5.2, we provide two optimal algorithms, both of which have to distinguish between the original and transformed data.

5.1 DA can even “hurt”

In the invariantly realizable setting, the optimal algorithm based on 1-inclusion graphs does not need to distinguish between the original and transformed data since $\mathcal{H}(X_S \cup \{x\})$ in the algorithm is fully determined by the augmented data. However, in the relaxed realizable setting, distinguishing between the original and transformed data is crucial. In the following, we will provide a characterization of the sample complexity of algorithms not distinguishing between the original and transformed data, including DA. Such a characterization induces a sufficient condition when DA “hurts”.

Let $\mathcal{M}_{\text{DA}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ be the smallest integer m for which there exists a learning rule \mathcal{A} such that for every \mathcal{G} -invariant data distribution \mathcal{D} , with probability at least $1 - \delta$ over $S_{\text{trn}} \sim \mathcal{D}^m$, $\text{err}_{\mathcal{D}}(\mathcal{A}(S_{\text{trn}})) \leq \epsilon$. The quantity $\mathcal{M}_{\text{DA}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ is the optimal sample complexity achievable if algorithms can only access the augmented data without knowing the original training set. In standard PAC learning, the optimal sample complexity can be characterized by the maximum density of any subgraph of the 1-inclusion graphs, which is actually equal to the VC dimension (Haussler et al., 1994; Daniely and Shalev-Shwartz, 2014). Analogously, we characterize $\mathcal{M}_{\text{DA}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ based on a variant of 1-inclusion graphs, which is constructed as follows.

In the 1-inclusion graph for standard PAC learning (Haussler et al., 1994), given any sequence of instances $\mathbf{x} = (x_1, \dots, x_t)$, the vertices are labelings of \mathbf{x} and two vertices are connected iff. they are different at only one instance in \mathbf{x} and this instance appears once. In our setting, the input is a multiset of labeled orbits and an unlabeled test instance, hence the vertices are pairs of labelings of orbits

²In this case, we only assume that \mathcal{G} contains the identity element.

and unlabeled instances. Specifically, for any $t \in \mathbb{N}$, given a multi-set of orbits $\phi = \{\phi_1, \dots, \phi_t\}$ of some unknown original data, a labeling $\mathbf{f} \in \mathcal{Y}^t$ is possible iff. there exists a sequence of instances $\mathbf{x} = (x_1, \dots, x_t) \in \prod_{i=1}^t \phi_i$ and a hypothesis $h \in \mathcal{H}$ such that $h|_{\mathbf{x}} = \mathbf{f}$ and that instances in the same orbit are labeled the same, i.e., $(\mathcal{G}x_i \times \{1 - f_i\}) \cap \{(x_j, f_j)\}_{j \in [t]} = \emptyset$ for all $i \in [t]$. We denote the set of all possible labelings of ϕ by

$$\Pi_{\mathcal{H}}(\phi) := \{\mathbf{f} \in \mathcal{Y}^t \mid \exists h \in \mathcal{H}, \exists \mathbf{x} \in \prod_{i=1}^t \phi_i, h|_{\mathbf{x}} = \mathbf{f} \text{ and } \cup_{i \in [t]} (\mathcal{G}x_i \times \{1 - f_i\}) \cap \{(x_j, f_j)\}_{j \in [t]} = \emptyset\}. \quad (2)$$

Denote the set of all such sequences $\mathbf{x} \in \prod_{i=1}^t \phi_i$ of instances that can be labeled as \mathbf{f} by

$$\mathcal{U}_{\mathbf{f}}(\phi) := \{\mathbf{x} \in \prod_{i=1}^t \phi_i \mid \exists h \in \mathcal{H}, h|_{\mathbf{x}} = \mathbf{f} \text{ and } \cup_{i \in [t]} (\mathcal{G}x_i \times \{1 - f_i\}) \cap \{(x_j, f_j)\}_{j \in [t]} = \emptyset\}.$$

Denote the set of all pairs of labeling and its corresponding instance sequence by

$$B(\mathcal{H}, \mathcal{G}, \phi) := \cup_{\mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} \{\mathbf{f}\} \times \mathcal{U}_{\mathbf{f}}(\phi). \quad (3)$$

For any $(\mathbf{f}, \mathbf{x}) \in B(\mathcal{H}, \mathcal{G}, \phi)$, \mathbf{x} is a candidate of original data and \mathbf{f} is a candidate of labeling of \mathbf{x} . Now we define a graph $G_{\mathcal{H}, \mathcal{G}}(\phi)$, where the vertices are all pairs of labeling $\mathbf{f} \in \Pi_{\mathcal{H}}(\phi)$ and an element in a instance sequence corresponding to \mathbf{f} . Formally, the vertex set is

$$V = \{(\mathbf{f}, x_i) \mid (\mathbf{f}, \mathbf{x}) \in B(\mathcal{H}, \mathcal{G}, \phi), i \in [t]\}.$$

For every two vertices (\mathbf{f}, x) and (\mathbf{g}, z) , they are connected if and only if (i) $x = z$; (ii) there exists $j \in [t]$ such that $x \in \phi_j$, $f_i = g_i, \forall i \neq j$ and $f_j \neq g_j$; and (iii) ϕ_j only appear once in ϕ . Each edge can be represented by $e = \{\mathbf{f}, \mathbf{g}, x\}$ and we denote E the edge set. If an edge $e = \{\mathbf{f}, \mathbf{g}, x\}$ exists, the edge could be recovered given only \mathbf{f}, x or \mathbf{g}, x , and thus, we also denote by $e(\mathbf{f}, x) = e(\mathbf{g}, x) = \{\mathbf{f}, \mathbf{g}, x\}$. Any algorithm accessing only the augmented data corresponds to an orientation of edges in the graph we constructed, which leads to the following definition.

Definition 6. Let $w : E \times \Pi_{\mathcal{H}}(\phi) \mapsto [0, 1]$ be a mapping such that for every $e = \{\mathbf{f}, \mathbf{g}, x\}$, $w(e, \mathbf{f}) + w(e, \mathbf{g}) = 1$ and $w(e, \mathbf{h}) = 0$ if $\mathbf{h} \notin e$ and let W be the set of all such mappings. Note that w actually defines a randomized orientation of each edge in graph $G_{\mathcal{H}, \mathcal{G}}(\phi)$: the edge $e = \{\mathbf{f}, \mathbf{g}, x\}$ is oriented towards vertex (\mathbf{f}, x) with probability $w(e, \mathbf{f})$. For any $(\mathbf{f}, \mathbf{x}) \in B(\mathcal{H}, \mathcal{G}, \phi)$, it corresponds to a cluster of vertices $\{(\mathbf{f}, x_i) \mid i \in [t]\}$ in $G_{\mathcal{H}, \mathcal{G}}(\phi)$ and $\sum_{i \in [t]: \exists e \in E, \{\mathbf{f}, x_i\} \subseteq e} w(e(\mathbf{f}, x_i), \mathbf{f})$ is the expected in-degree of the cluster. Let $\Delta(B(\mathcal{H}, \mathcal{G}, \phi))$ denote the set of all distributions over $B(\mathcal{H}, \mathcal{G}, \phi)$. For any $P \in \Delta(B(\mathcal{H}, \mathcal{G}, \phi))$, we define

$$\mu(\mathcal{H}, \mathcal{G}, \phi, P) := \min_{w \in W} \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P} \left[\sum_{i \in [t]: \exists e \in E, \{\mathbf{f}, x_i\} \subseteq e} w(e(\mathbf{f}, x_i), \mathbf{f}) \right]. \quad (4)$$

By taking the supremum over P , we define $\mu(\mathcal{H}, \mathcal{G}, \phi) := \sup_{P \in \Delta(B(\mathcal{H}, \mathcal{G}, \phi))} \mu(\mathcal{H}, \mathcal{G}, \phi, P)$. By taking supremum over ϕ , we define $\mu(\mathcal{H}, \mathcal{G}, t) := \sup_{\phi: |\phi|=t} \mu(\mathcal{H}, \mathcal{G}, \phi)$ and

$$\mu(\mathcal{H}, \mathcal{G}) := \sup_{t \in \mathbb{N}} \mu(\mathcal{H}, \mathcal{G}, t). \quad (5)$$

Theorem 5. For any \mathcal{H}, \mathcal{G} , $\mathcal{M}_{\text{DA}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ satisfies the following bounds:

- For all $t \geq 2$ with $\mu(\mathcal{H}, \mathcal{G}, t) < \infty$, $\mathcal{M}_{\text{DA}}(\epsilon, \frac{\mu(\mathcal{H}, \mathcal{G}, t)}{16(t-1)}; \mathcal{H}, \mathcal{G}) = \Omega(\frac{\mu(\mathcal{H}, \mathcal{G}, t)}{\epsilon})$. This implies that if $\mu(\mathcal{H}, \mathcal{G}) < \infty$, there exists a constant c dependent on \mathcal{H}, \mathcal{G} s.t. $\mathcal{M}_{\text{DA}}(\epsilon, c; \mathcal{H}, \mathcal{G}) = \Omega(\frac{\mu(\mathcal{H}, \mathcal{G})}{\epsilon})$.
- For all t with $\frac{1}{6} \leq \frac{\mu(\mathcal{H}, \mathcal{G}, t)}{t} \leq \frac{1}{3}$, $\mathcal{M}_{\text{DA}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) = O(\frac{\mu(\mathcal{H}, \mathcal{G}, t)}{\epsilon} \log^2 \frac{\mu(\mathcal{H}, \mathcal{G}, t)}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta})$.
- If $\mu(\mathcal{H}, \mathcal{G}) < \infty$, $\mathcal{M}_{\text{DA}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) = O(\frac{\mu(\mathcal{H}, \mathcal{G}) \log(1/\delta)}{\epsilon})$.

Theorem 5 implies that when $\mu(\mathcal{H}, \mathcal{G}) > \text{VCdim}(\mathcal{H})$, there exists a distribution such that any algorithm not differentiating between the original and transformed data performs worse than simply applying ERM over the original data. We defer the proof of Theorem 5 to Appendix E. As we can see, the definition of $\mu(\mathcal{H}, \mathcal{G})$ is not intuitive and it might be difficult to calculate $\mu(\mathcal{H}, \mathcal{G})$ as well as to further determine when $\mu(\mathcal{H}, \mathcal{G}) > \text{VCdim}(\mathcal{H})$. We introduce a new dimension as follows, which lower bounds $\mu(\mathcal{H}, \mathcal{G})$ and is easier to calculate.

Definition 7 (VC Dimension of orbits generated by \mathcal{H}). *The VC dimension of orbits generated by \mathcal{H} , denoted $\dim(\mathcal{H}, \mathcal{G})$, is defined as the largest integer k for which there exists a set $X = \{x_1, \dots, x_k\} \subset \mathcal{X}$ such that (i) their orbits $\phi = \{\phi_1, \dots, \phi_k\}$ are pairwise disjoint, (ii) $\Pi_{\mathcal{H}}(\phi) = \mathcal{Y}^k$ (defined in Eq (2)) and (iii) there exists a set $B = \{(\mathbf{f}, \mathbf{x}_{\mathbf{f}})\}_{\mathbf{f} \in \mathcal{Y}^k} \subset B(\mathcal{H}, \mathcal{G}, \phi)$ (defined in Eq (3)) such that $\mathbf{f} \oplus \mathbf{g} = \mathbf{e}_i$ implies $x_{\mathbf{f}, i} = x_{\mathbf{g}, i}$ for all $i \in [k]$, $\mathbf{f}, \mathbf{g} \in \mathcal{Y}^k$.*

Theorem 6. *For any \mathcal{H}, \mathcal{G} , $\mu(\mathcal{H}, \mathcal{G}) \geq \dim(\mathcal{H}, \mathcal{G})/2$.*

The proof is included in Appendix F. Through this dimension, we claim that the gap between $\mu(\mathcal{H}, \mathcal{G})$ and $\text{VCdim}(\mathcal{H})$ can be arbitrarily large. In the following, we give an example of \mathcal{H}, \mathcal{G} with $\dim(\mathcal{H}, \mathcal{G}) \gg \text{VCdim}(\mathcal{H})$, which cannot be learned by DA but can be easily learned by ERM.

Example 5. *For any $d > 0$, let $\mathcal{X} = \{\pm 1\} \times \{\mathbf{e}_i | i \in [2d]\} \subset \mathbb{R}^{2d+1}$, $\mathcal{H} = \{x_1 > 0, x_1 \leq 0\}$ and $\mathcal{G} = \{I_{2d+1}, I_{2d+1} - 2\text{diag}(\mathbf{e}_1)\}$ (i.e., the cyclic group generated by flipping the sign of x_1). It is easy to check that $\text{VCdim}(\mathcal{H}) = 1$. Let $X = \{(1, \mathbf{e}_1), (1, \mathbf{e}_2), \dots, (1, \mathbf{e}_{2d})\}$ and then the orbits generated from X are $\phi = \{(-1, \mathbf{e}_i), (1, \mathbf{e}_i) | i \in [2d]\}$. For every labeling $\mathbf{f} \in \mathcal{Y}^{2d}$, if $\sum_{i \in [2d]} f_i$ is odd, let $\mathbf{x}_{\mathbf{f}} = ((2f_i - 1, \mathbf{e}_i))_{i=1}^{2d}$; if $\sum_{i \in [2d]} f_i$ is even, let $\mathbf{x}_{\mathbf{f}} = ((1 - 2f_i, \mathbf{e}_i))_{i=1}^{2d}$. It is direct to check that $(\mathbf{f}, \mathbf{x}_{\mathbf{f}}) \in B(\mathcal{H}, \mathcal{G}, \phi)$ for all $\mathbf{f} \in \mathcal{Y}^{2d}$. Then for all $i \in [2d]$, $\mathbf{f} \oplus \mathbf{g} = \mathbf{e}_i$ implies $x_{\mathbf{f}, i} = x_{\mathbf{g}, i}$. Hence, $\dim(\mathcal{H}, \mathcal{G}) = 2d$, where d can be an arbitrary positive integer. According to Theorem 6, we have $\mu(\mathcal{H}, \mathcal{G}) \geq d$.*

The above example can be interpreted in a vision scenario. Let's consider an example of classifying land birds versus water birds. The natural data is $2d$ images of land birds with land background and water birds with water background. The transformation set is composed of keeping the current background and changing the background from land (water) to water (land). Consider simple hypotheses depending on backgrounds only. Specifically, $\mathcal{H} = \{h_1, h_2\}$ with h_1 predicting all images with water background as water birds and h_2 predicting all images with water background as land birds. Let the data distribution be the uniform distribution over all the original images. Then given any training data, h_1 and h_2 have the same empirical loss on the augmented training data. Thus, for any unobserved image, DA will make a mistake with constant probability. Hence DA requires at least $\Omega(d)$ sample complexity. It is direct to check that standard ERM only needs one labeled instance to achieve zero error.

Open question: It is unclear whether $\mu(\mathcal{H}, \mathcal{G})$ is upper bounded by $\dim(\mathcal{H}, \mathcal{G})$. If true, then we can tightly characterize $\mathcal{M}_{\text{DA}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ by $\dim(\mathcal{H}, \mathcal{G})$.

5.2 The optimal algorithms

Different from the invariantly realizable setting, the optimal sample complexity in the relaxed realizable setting is characterized by $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$. The optimal (up to log-factors) sample complexity can be achieved by another variant of 1-inclusion-graph predictor. Besides, we propose an ERM-based algorithm, called ERM-INV (see Appendix G for details), achieving the similar guarantee.

Theorem 7. *For any \mathcal{H}, \mathcal{G} with $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) < \infty$, we have $\Omega\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} + \frac{\log(1/\delta)}{\epsilon}\right) \leq \mathcal{M}_{\text{RE}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) \leq O\left(\min\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} \log^3 \frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta}, \frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon} \log \frac{1}{\delta}\right)\right)$.*

We defer the details of algorithms and the proof of Theorem 7 to Appendix G. Usually, ERM-INV is more efficient than the 1-inclusion-graph predictor. But the 1-inclusion-graph predictor as well as the lower bound can apply to non-group transformations. Another advantage of 1-inclusion-graph predictor is allowing us to design an adaptive framework which automatically adjusts to different levels of invariance of h^* . Specifically, for any hypothesis $h \in \mathcal{H}$, we say h is $(1 - \eta)$ -invariant over the distribution $\mathcal{D}_{\mathcal{X}}$ for some $\eta \in [0, 1]$ if $\mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}}(\exists x' \in \mathcal{G}x, h(x') \neq h(x)) = \eta$. When $\eta(h^*) = 0$, it degenerates into the invariantly realizable setting, which implies that we can achieve better bounds when $\eta(h^*)$ is smaller. We propose an adaptive algorithm with sample complexity dependent on $\eta(h^*)$ and the details are included in Appendix I.1.

6 Agnostic setting

In the agnostic setting (Definition 3), $\inf_{h \in \mathcal{H}} \text{err}(h)$ is possibly non-zero. Different from the agnostic setting in the standard PAC learning allowing probabilistic labels, our problem is limited to

deterministic labels because we assume that the data distribution is \mathcal{G} -invariant, i.e., there exists a $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant hypothesis f^* (possibly not in \mathcal{H}) with $\text{err}_{\mathcal{D}}(f^*) = 0$.

Theorem 8. *The sample complexity in the agnostic setting satisfies:*

- For all $d > 0$, there exists \mathcal{H}, \mathcal{G} with $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = d$, $\mathcal{M}_{\text{AG}}(\epsilon, 1/64; \mathcal{H}, \mathcal{G}) = \Omega(\frac{d}{\epsilon^2})$.
- For any \mathcal{H}, \mathcal{G} with $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) < \infty$, $\mathcal{M}_{\text{AG}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) = \tilde{O}\left(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon^2} + \frac{1}{\epsilon^2} \log(\frac{1}{\delta})\right)$.

For upper bound, we show that ERM-INV achieves sample complexity $\tilde{O}(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon^2})$. There is another way of achieving similar upper bound based on applying the reduction-to-realizable technique of David et al. (2016). Note that a direct combination of any reduction-to-realizable technique and any optimal algorithm in relaxed realizable setting does not work in our agnostic setting. This is because the relaxed realizable setting requires not only realizability, but also invariance in the support of the data distribution. For example, the reduction method of Hopkins et al. (2021) needs to run a realizable algorithm over a set labeled by each $h \in \mathcal{H}$, which might label two instances in the same orbit differently and make the realizable algorithm not well-defined. When combining the reduction method of David et al. (2016) and the 1-inclusion-graph-type algorithm, the similar problem also exists but can be fixed by predicting arbitrarily when the invariance property is not satisfied. For lower bound, According to Ben-David and Uner (2014), the sample complexity of agnostic PAC learning under deterministic labels is not fully determined by the VC dimension. Following the construction by Ben-David and Uner (2014), we provide an analogous lower bound in our setting. The algorithm details and the proofs are deferred to Appendix H. Analogous to the realizable setting, we provide one algorithm adapting to different levels of invariance of the optimal hypothesis in \mathcal{H} in Appendix I.2. Similar to the results in the realizable settings, the lower bound and the 1-inclusion-graph predictor in the agnostic setting also apply to non-group transformations.

7 Discussion

Definition of invariance under probabilistic labels In this work, we model invariance by assuming that the data distribution is \mathcal{G} -invariant, which restricts the labels to be deterministic. It is unclear what “invariance under probabilistic labels” means. One option is assuming that the distribution of the labels is invariant over the orbits, $\Pr(y|gx) = \Pr(y|x)$ for all $g \in \mathcal{G}, x \in \mathcal{X}$. However, such a condition may not characterize invariance in real-world scenarios due to classes having different underlying distributions. For example, given a fuzzy image with probability 0.5 being a car and 0.5 being a tree, it is uncertain if the chance of this image being a car is still 0.5 after rotation.

The performance of DA under non-group transformations Most results of DA and ERM-type algorithms only hold when the transformation set is a group. If we regard adversarial training as a special type of data augmentation through a ball around the natural data, then the transformation set is not a group. The appropriate way to formulate theoretical guarantees for DA under arbitrary transformations is still an open question.

Acknowledgements

This work was supported in part by the National Science Foundation under grant CCF-1815011 and by the Defense Advanced Research Projects Agency under cooperative agreement HR00112020003. The views expressed in this work do not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred. Approved for public release; distribution is unlimited.

We thank anonymous reviewers for their valuable suggestions. HS thanks Freda Shi for discussion on the application of DA and suggestions from an applied viewpoint.

References

- Ben-David, S. and Uner, R. (2014). The sample complexity of agnostic learning under deterministic labels. In *Conference on Learning Theory*, pages 527–542. PMLR.
- Bietti, A., Venturi, L., and Bruna, J. (2021). On the sample complexity of learning under geometric stability. *Advances in Neural Information Processing Systems*, 34.

- Bloem-Reddy, B. and Teh, Y. W. (2020). Probabilistic symmetries and invariant neural networks. J. Mach. Learn. Res., 21:90–1.
- Blumer, A., Ehrenfeucht, A., Haussler, D., and Warmuth, M. K. (1989). Learnability and the Vapnik-Chervonenkis dimension. Journal of the ACM (JACM), 36(4):929–965.
- Chatzipantazis, E., Pertigkiozoglou, S., Dobriban, E., and Daniilidis, K. (2021). Learning augmentation distributions using transformed risk minimization. arXiv preprint arXiv:2111.08190.
- Chen, S., Dobriban, E., and Lee, J. H. (2020). A group-theoretic framework for data augmentation. Journal of Machine Learning Research, 21:1–71.
- Cohen, T. and Welling, M. (2016). Group equivariant convolutional networks. In International conference on machine learning, pages 2990–2999. PMLR.
- Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., and Le, Q. V. (2018). Autoaugment: Learning augmentation policies from data. arXiv preprint arXiv:1805.09501.
- Daniely, A. and Shalev-Shwartz, S. (2014). Optimal learners for multiclass problems. In Conference on Learning Theory, pages 287–316. PMLR.
- Dao, T., Gu, A., Ratner, A., Smith, V., De Sa, C., and Ré, C. (2019). A kernel theory of modern data augmentation. In International Conference on Machine Learning, pages 1528–1537. PMLR.
- David, O., Moran, S., and Yehudayoff, A. (2016). Supervised learning through the lens of compression. Advances in Neural Information Processing Systems, 29.
- Dieleman, S., De Fauw, J., and Kavukcuoglu, K. (2016). Exploiting cyclic symmetry in convolutional neural networks. In International conference on machine learning, pages 1889–1898. PMLR.
- Ehrenfeucht, A., Haussler, D., Kearns, M., and Valiant, L. (1989). A general lower bound on the number of examples needed for learning. Information and Computation, 82(3):247–261.
- Elesedy, B. (2021). Provably strict generalisation benefit for invariance in kernel methods. Advances in Neural Information Processing Systems, 34.
- Elesedy, B. (2022). Group symmetry in pac learning. In ICLR 2022 Workshop on Geometrical and Topological Representation Learning.
- Elesedy, B. and Zaidi, S. (2021). Provably strict generalisation benefit for equivariant models. In International Conference on Machine Learning, pages 2959–2969. PMLR.
- Fawzi, A., Samulowitz, H., Turaga, D., and Frossard, P. (2016). Adaptive data augmentation for image classification. In 2016 IEEE international conference on image processing (ICIP), pages 3688–3692. Ieee.
- Fukushima, K. and Miyake, S. (1982). Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition. In Competition and cooperation in neural nets, pages 267–285. Springer.
- Gontijo-Lopes, R., Smullin, S. J., Cubuk, E. D., and Dyer, E. (2020). Affinity and diversity: Quantifying mechanisms of data augmentation. arXiv preprint arXiv:2002.08973.
- Graepel, T., Herbrich, R., and Shawe-Taylor, J. (2005). Pac-bayesian compression bounds on the prediction error of learning algorithms for classification. Machine Learning, 59(1-2):55–76.
- Haussler, D., Littlestone, N., and Warmuth, M. K. (1994). Predicting $\{0, 1\}$ -functions on randomly drawn points. Information and Computation, 115(2):248–292.
- Hopkins, M., Kane, D., Lovett, S., and Mahajan, G. (2021). Realizable learning is all you need. arXiv preprint arXiv:2111.04746.
- Kondor, R. and Trivedi, S. (2018). On the generalization of equivariance and convolution in neural networks to the action of compact groups. In International Conference on Machine Learning, pages 2747–2755. PMLR.

- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 25.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. (1989). Backpropagation applied to handwritten zip code recognition. Neural computation, 1(4):541–551.
- Littlestone, N. and Warmuth, M. (1986). Relating data compression and learnability.
- Lyle, C., van der Wilk, M., Kwiatkowska, M., Gal, Y., and Bloem-Reddy, B. (2020). On the benefits of invariance in neural networks. arXiv preprint arXiv:2005.00178.
- Mei, S., Misiakiewicz, T., and Montanari, A. (2021). Learning with invariances in random features and kernel models. In Conference on Learning Theory, pages 3351–3418. PMLR.
- Nakkiran, P. (2019). Adversarial robustness may be at odds with simplicity. arXiv preprint arXiv:1901.00532.
- Raghunathan, A., Xie, S. M., Yang, F., Duchi, J. C., and Liang, P. (2019). Adversarial training can hurt generalization. arXiv preprint arXiv:1906.06032.
- Ravanbakhsh, S., Schneider, J., and Poczos, B. (2017). Equivariance through parameter-sharing. In International Conference on Machine Learning, pages 2892–2901. PMLR.
- Schapire, R. E. and Freund, Y. (2012). Boosting: Foundations and Algorithms. MIT Press.
- Schmidt, L., Santurkar, S., Tsipras, D., Talwar, K., and Madry, A. (2018). Adversarially robust generalization requires more data. arXiv preprint arXiv:1804.11285.
- Shen, R., Bubeck, S., and Gunasekar, S. (2022). Data augmentation as feature manipulation: a story of desert cows and grass cows. arXiv preprint arXiv:2203.01572.
- Vapnik, V. and Chervonenkis, A. (1974). Theory of Pattern Recognition. Nauka, Moscow.
- Wang, D., Shelhamer, E., Liu, S., Olshausen, B. A., and Darrell, T. (2021). Tent: Fully test-time adaptation by entropy minimization. In International Conference on Learning Representations.
- Wood, J. and Shawe-Taylor, J. (1996). Representation theory and invariant neural networks. Discrete applied mathematics, 69(1-2):33–60.
- Worrall, D. E., Garbin, S. J., Turmukhambetov, D., and Brostow, G. J. (2017). Harmonic networks: Deep translation and rotation equivariance. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 5028–5037.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#)
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#) This is a technical theory work.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#)
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#)
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[N/A\]](#)

- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Proof of Theorem 1

Proof. Let $d = \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$. According to the definition of DA and the invariantly realizable setting, given the input $S \sim \mathcal{D}^m$, the output of DA $\hat{h} \in \mathcal{H}$ satisfies $\text{err}_{\mathcal{G}S}(\hat{h}) = 0$, i.e., $\hat{h}(x) = y$ for all $(x, y) \in \mathcal{G}S$. Consider two sets S and S' of m i.i.d. samples drawn from the data distribution \mathcal{D} each. We denote A_S the event of $\{\exists h \in \mathcal{H}, \text{err}_{\mathcal{D}}(h) \geq \epsilon, \text{err}_{\mathcal{G}S}(h) = 0\}$ and $B_{S, S'}$ the event of $\{\exists h \in \mathcal{H}, \text{err}_{S'}(h) \geq \frac{\epsilon}{2}, \text{err}_{\mathcal{G}S}(h) = 0\}$. By Chernoff bound, we have $\Pr(B_{S, S'}) \geq \Pr(A_S) \cdot \Pr(B_{S, S'} | A_S) \geq \frac{1}{2} \Pr(A_S)$ when $m \geq \frac{8}{\epsilon}$. The sampling process of S and S' is equivalent to drawing $2m$ i.i.d. samples and then randomly partitioning into S and S' of m each. For any fixed S'' , for any h with $\text{err}_{S'}(h) \geq \frac{\epsilon}{2}$ and $\text{err}_{\mathcal{G}S}(h) = 0$, if h misclassifies some $(x, y) \in S''$, then all examples in the orbit of x , i.e., $\mathcal{G}\{(x, y)\} \cap S''$, must go to S' . Now to prove the theorem, we divide S'' into two categories in terms of the number of examples in each orbit. Let $R_1 = \{x | |\mathcal{G}x \cap S''_{\mathcal{X}}| \geq \log^2 m, x \in S''_{\mathcal{X}}\}$ and $R_2 = S''_{\mathcal{X}} \setminus R_1$. For any h making at least $\frac{\epsilon m}{2}$ mistakes in $S''_{\mathcal{X}}$, h either makes at least $\frac{\epsilon m}{4}$ mistakes in R_1 or makes at least $\frac{\epsilon m}{4}$ mistakes in R_2 . Then let $\mathcal{H}_0 \subset \mathcal{H}$ denote the set of hypotheses making at least $\frac{\epsilon m}{2}$ mistakes in $S''_{\mathcal{X}}$ and divide \mathcal{H}_0 into two sub-classes as follows.

- Let $\mathcal{H}_1 = \{h \in \mathcal{H}_0 | h \text{ makes at least } \frac{\epsilon m}{4} \text{ mistakes in } R_1\}$. For any $h \in \mathcal{H}_1$, we let $X(h) \subset R_1$ denote a minimal set of examples in R_1 (breaking ties arbitrarily but in a fixed way) such that h misclassify $X(h)$ and $|(\mathcal{G}X(h)) \cap S''_{\mathcal{X}}| \geq \frac{\epsilon m}{4}$ where $\mathcal{G}X(h) = \{gx | x \in X(h), g \in \mathcal{G}\}$ is the set of all examples lying in the orbits generated from $X(h)$. Let $K(h) = \mathcal{G}X(h)$ and $\mathcal{K} = \{K(h) | h \in \mathcal{H}_1\}$ the collection of all such sets. Notice that each example in $X(h)$ must belong to different orbits, otherwise it is not minimal. Besides, each orbit in $K(h)$ contains at least $\log^2 m$ examples from $S''_{\mathcal{X}}$ according to the definition of R_1 . Hence, we have $|X(h)| \leq \frac{\epsilon m}{4 \log^2 m}$. Since there are at most $\frac{2m}{\log^2 m}$ orbits generated from R_1 , we have $|\mathcal{K}| \leq \sum_{i=1}^{\frac{\epsilon m}{4 \log^2 m}} \binom{2m}{i} \leq \left(\frac{8e}{\epsilon}\right)^{\frac{\epsilon m}{4 \log^2 m}}$. Recall that $\text{err}_{\mathcal{G}S}(h) = 0$ iff. $h(x) = y$ for all $(x, y) \in \mathcal{G}S$. Since h misclassify $X(h)$, all examples in their orbits must go to S' to guarantee $\text{err}_{\mathcal{G}S}(h) = 0$. Thus, we have

$$\begin{aligned} & \Pr(\exists h \in \mathcal{H}_1, \text{err}_{S'}(h) \geq \frac{\epsilon}{2}, \text{err}_{\mathcal{G}S}(h) = 0) \\ & \leq \Pr(\exists h \in \mathcal{H}_1, K(h) \cap S_{\mathcal{X}} = \emptyset) = \Pr(\exists K \in \mathcal{K}, K \cap S_{\mathcal{X}} = \emptyset) \\ & \leq \sum_{K \in \mathcal{K}} 2^{-\frac{\epsilon m}{4}} \leq \left(\frac{8e}{\epsilon}\right)^{\frac{\epsilon m}{4 \log^2 m}} \cdot 2^{-\frac{\epsilon m}{4}} = 2^{-\frac{\epsilon m}{4} \left(1 - \frac{\log(8e/\epsilon)}{\log^2 m}\right)} \leq 2^{-\frac{\epsilon m}{8}}, \end{aligned}$$

when $m \geq \frac{8e}{\epsilon} + 4$.

- Let $\mathcal{H}_2 = \mathcal{H}_0 \setminus \mathcal{H}_1$. That is to say, for all $h \in \mathcal{H}_2$, h will make at least $\frac{\epsilon m}{4}$ mistakes in R_2 . Since $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = d$ and every orbit generated from R_2 contains fewer than $\log^2 m$ examples in $S''_{\mathcal{X}}$, the number of examples in R_2 that can be shattered by \mathcal{H} is no greater than $d \log^2 m$. Thus the number of ways labeling examples in R_2 is upper bounded by $\left(\frac{2em}{d}\right)^{d \log^2 m}$ by Sauer's lemma. Hence, we have

$$\Pr(\exists h \in \mathcal{H}_2, \text{err}_{S'}(h) \geq \frac{\epsilon}{2}, \text{err}_{\mathcal{G}S}(h) = 0) \leq \left(\frac{2em}{d}\right)^{d \log^2 m} \cdot 2^{-\frac{\epsilon m}{4}} = 2^{-\frac{\epsilon m}{4} + d \log^2 m \log(2em/d)}.$$

Combining the results for \mathcal{H}_1 and \mathcal{H}_2 , we have

$$\begin{aligned} \Pr(B_{S, S'}) & \leq \Pr(\exists h \in \mathcal{H}_1, \text{err}_{S'}(h) \geq \frac{\epsilon}{2}, \text{err}_{\mathcal{G}S}(h) = 0) + \Pr(\exists h \in \mathcal{H}_2, \text{err}_{S'}(h) \geq \frac{\epsilon}{2}, \text{err}_{\mathcal{G}S}(h) = 0) \\ & \leq 2^{-\frac{\epsilon m}{8}} + 2^{-\frac{\epsilon m}{4} + d \log^2 m \log(2em/d)} \\ & \leq \frac{\delta}{2}, \end{aligned}$$

when $m \geq \frac{8}{\epsilon} (d \log^2 m \log \frac{2em}{d} + \log \frac{4}{\delta} + e) + 4$. □

B Proof of Theorem 2

Proof. For any $d > 0$, for any $\mathcal{X}, \mathcal{H}, \mathcal{G}$ satisfying that there exists a subset $X = \{x_0, x_1, \dots, x_{2d}\} \subset \mathcal{X}$ such that

- their orbits are pairwise disjoint;
- for all $u \subset \{1, \dots, 2d\}$ with $|u| = d$, there exists an $h_u \in \mathcal{H}$ such that $h_u(gx_i) = 1$ for all $g \in \mathcal{G}, i \in u \cup \{0\}$ and $h_u(x_i) = 0$ for $i \in \{0, \dots, 2d\} \setminus u$,

we will prove the theorem for $\mathcal{H}' = \{h_u | u \subset \{1, \dots, 2d\}, |u| = d\}$. For $\mathcal{X}, \mathcal{H}, \mathcal{G}$ satisfying the above conditions, we have $\text{VC}_{\text{ao}}(\mathcal{H}', \mathcal{G}) \geq d$ and $\text{VC}_o(\mathcal{H}', \mathcal{G})$ between 0 and $\text{VC}_{\text{ao}}(\mathcal{H}', \mathcal{G})$. Then consider that the target function $h^* = h_{u^*}$ is chosen uniformly at random from \mathcal{H}' . The marginal data distribution $\mathcal{D}_{\mathcal{X}}$ puts probability mass $1 - 16\epsilon$ on x_0 and $\frac{16\epsilon}{d}$ on each point in $X_{u^*} := \{x_i | i \in u^*\}$. Then the target function h^* is $(\mathcal{G}, \mathcal{D}_{\mathcal{X}})$ -invariant.

Let the sample size $m = \frac{d}{64\epsilon}$. Given the training set $S_{\text{trn}} \sim \mathcal{D}^m$, the expected number of sampled examples in X_{u^*} is $\frac{d}{4}$. By Markov's inequality, with probability greater than $1/2$, we observed fewer than $\frac{d}{2}$ points of X_{u^*} in S_{trn} (denoted as event B). Let \mathcal{A} be any proper learner, which means \mathcal{A} must output a hypothesis in \mathcal{H}' . For any $h \in \mathcal{H}'$ consistent with $\mathcal{G}S_{\text{trn}}$, h must predict d unobserved points in $\{x_1, \dots, x_{2d}\}$ as 0. Since for each unobserved point in $\{x_1, \dots, x_{2d}\}$ labeled as 0 by h , conditioned on event B , this point has probability greater than $\frac{1}{3}$ to be in X_{u^*} , which implies it is misclassified by h . By following the standard technique, let $\text{err}'(h) = \Pr_{(x,y) \sim \mathcal{D}}(h(x) \neq y \wedge x \in X_{u^*})$, which is no greater than $\text{err}(h)$ for any predictor h . Hence,

$$\begin{aligned} \mathbb{E}_{h^*, S_{\text{trn}}} [\text{err}'(\mathcal{A}(S_{\text{trn}})) | B] &= \mathbb{E}_{S_{\text{trn}}} [\mathbb{E}_{h^*} [\text{err}'(\mathcal{A}(S_{\text{trn}})) | S_{\text{trn}}, B] | B] \\ &= \mathbb{E}_{S_{\text{trn}}} \left[\mathbb{E}_{h^*} \left[\frac{16\epsilon}{d} \sum_{i \in [2d]: \mathcal{A}(S_{\text{trn}}, x_i) = 0} \mathbb{1}[i \in u^*] | S_{\text{trn}}, B \right] | B \right] \\ &= \frac{16\epsilon}{d} \mathbb{E}_{S_{\text{trn}}} \left[\sum_{i \in [2d]: \mathcal{A}(S_{\text{trn}}, i) = 0} \mathbb{E}_{h^*} [\mathbb{1}[i \in u^*] | S_{\text{trn}}, B] | B \right] \\ &> \frac{16\epsilon}{d} \mathbb{E}_{S_{\text{trn}}} \left[\sum_{i \in [2d]: \mathcal{A}(S_{\text{trn}}, i) = 0} \frac{1}{3} | B \right] \\ &= \frac{16\epsilon}{3}. \end{aligned}$$

Then we have

$$\mathbb{E}_{h^*, S_{\text{trn}}} [\text{err}'(\mathcal{A}(S_{\text{trn}}))] \geq \mathbb{E}_{h^*, S_{\text{trn}}} [\text{err}'(\mathcal{A}(S_{\text{trn}})) | B] \cdot \Pr(B) > \frac{8\epsilon}{3}.$$

Thus, for any proper learner \mathcal{A} , there exists a target hypothesis $h^* \in \mathcal{H}'$ and a data distribution s.t. $\mathbb{E}_{S_{\text{trn}}} [\text{err}'(\mathcal{A}(S_{\text{trn}}))] > 8\epsilon/3$. Since $\text{err}'(h) \leq 16\epsilon$ for any predictor h , with probability greater than $\frac{1}{9}$, $\text{err}(\mathcal{A}(S_{\text{trn}})) \geq \text{err}'(\mathcal{A}(S_{\text{trn}})) > \epsilon$.

Here is an example of $\mathcal{X}, \mathcal{H}', \mathcal{G}$ satisfying the above conditions. Let $\mathcal{X} = \{0, \pm 1, \pm 2, \dots, \pm 2d\}$. The group \mathcal{G} is defined as $\mathcal{G}_d = \{e, -e\}$ where e is the identity element. Thus \mathcal{X} can be divided into $2d + 1$ pairwise disjoint orbits, $\{\{0\}, \{\pm 1\}, \dots, \{\pm 2d\}\}$. For any $u \subset [2d]$ with $|u| = d$, define $h_u := 1 - \mathbb{1}[[2d] \setminus u]$, which labels $[2d] \setminus u$ by 0 and the other points by 1. Then we define the hypothesis class $\mathcal{H}' = \{h_u | |u| = d\}$. Since $\{1, \dots, d\}$ can be shattered and no $d + 1$ points can be shattered by \mathcal{H}' , we have $\text{VC}_{\text{ao}}(\mathcal{H}', \mathcal{G}) = d$. And since $\{0, -1, \dots, -2d\}$ can only be labeled as 1 by \mathcal{H}' , we have $\text{VC}_o(\mathcal{H}', \mathcal{G}) = 0$. \square

C Proof of Theorem 3

Proof. Let $d = \text{VCdim}(\mathcal{H})$ and $X = \{x_1, x_2, \dots, x_d\}$ be a set of examples shattered by \mathcal{H} . Let S_{d-1} denote the permutation group acting on $d - 1$ objects. Then for any partition $(A, [d - 1] \setminus A)$

of $[d-1]$, we let $\mathcal{G}(A) := \{\sigma \in \mathcal{S}_{d-1} \mid \forall i \in A, \sigma(i) \in A\}$. By acting $\mathcal{G}(A)$ on X , then X are partitioned into three orbits: $\{x_i \mid i \in A\}$, $\{x_i \mid i \notin A\}$ and $\{x_d\}$.

For convenience, we first consider the case of the instance space being X . Since there are only three orbits and every labeling of X is realized by \mathcal{H} , $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}(A)) = \text{VC}_{\text{o}}(\mathcal{H}, \mathcal{G}(A)) = 3$ for all $A \subset [d-1]$. Consider that we pick a set A^* uniformly at random from $2^{[d-1]}$. Let $h^* = \mathbb{1}[\{x_i\}_{i \in A^*}]$ and $\mathcal{G} = \mathcal{G}(A^*)$. The data distribution put probability mass $1 - 16\epsilon$ on x_d and the remaining 16ϵ uniformly over $\{x_i \mid i \in [d-1]\}$. Then for the training set size $m = \frac{d-1}{64\epsilon}$, with probability at least $1/2$, at most half of $\{x_1, x_2, \dots, x_{d-1}\}$ is sampled in the training set \mathcal{S}_{trn} . For any algorithm \mathcal{A} not knowing \mathcal{G} , \mathcal{A} will output a hypothesis $\hat{h} = \mathcal{A}(\mathcal{S}_{\text{trn}})$, which does not depend on \mathcal{G} . For each unobserved point x in $\{x_1, x_2, \dots, x_{d-1}\}$, \mathcal{A} has probability $1/2$ to misclassify x . Following the standard technique, let $\text{err}'(h) := \mathbb{P}_{(x,y) \sim \mathcal{D}}(h(x) \neq y \wedge x \in \{x_1, x_2, \dots, x_{d-1}\})$ and then we have

$$\mathbb{E}_{A^*, \mathcal{S}_{\text{trn}}} [\text{err}'(\mathcal{A}(\mathcal{S}_{\text{trn}}))] \geq 2\epsilon,$$

which implies $\mathcal{M}_{\text{INV}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}, \mathcal{A}) = \Omega(\frac{d}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta})$ by applying the standard technique in proving a lower bound of sample complexity in standard PAC learning.

In the case where the instance space not being X , we modify $\mathcal{G}(A^*)$ a little by arranging all points in $\{x \mid h^*(x) = 1\} \setminus X$ in one orbit and all points in $\{x \mid h^*(x) = 0\} \setminus X$ in another orbit. Then there are at most 5 orbits, and $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = \text{VC}_{\text{o}}(\mathcal{H}, \mathcal{G}) \leq 5$. \square

D Proof of Theorem 4

Proof. For any $k \leq \text{VC}_{\text{o}}(\mathcal{H}, \mathcal{G})$, let $X_k = \{x_1, \dots, x_k\}$ be a set shattered in the way defined in Definition 4. Then X_k can be shattered by $\mathcal{H}(X_k) = \{h|_{X_k} \mid h \in \mathcal{H} \text{ is } (\mathcal{G}, X_k)\text{-invariant}\}$ and $\text{VCdim}(\mathcal{H}(X_k)) = k$. Since any data distribution \mathcal{D} with $\mathcal{D}_{\mathcal{X}}(X_k) = 1$ is \mathcal{G} -invariant, any lower bound on the sample complexity of PAC learning of $\mathcal{H}(X_k)$ also lower bounds the sample complexity of invariantly realizable PAC learning of \mathcal{H} . Then the lower bound follows by standard arguments from Vapnik and Chervonenkis (1974); Blumer et al. (1989); Ehrenfeucht et al. (1989).

For the upper bound, recall that the algorithm \mathcal{A} is defined by letting $\mathcal{A}(S)(x) = Q_{\mathcal{H}(X_S \cup \{x\}), X_S \cup \{x\}}(S, x)$ if $\mathcal{H}(X_S \cup \{x\}) \neq \emptyset$ and predicting arbitrarily if $\mathcal{H}(X_S \cup \{x\}) = \emptyset$ in Section 4.2. Due to the invariantly-realizable setting, if S and the test point (x, y) are i.i.d. from the data distribution, $h^*_{|X_S \cup \{x\}}$ is in $\mathcal{H}(X_S \cup \{x\})$ a.s. and then, $\mathcal{H}(X_S \cup \{x\})$ is nonempty. Following the analogous proof by Haussler et al. (1994) for standard PAC learning, we have

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}^t} [\text{err}(\mathcal{A}(S))] &= \mathbb{E}_{(x_i, y_i)_{i \in [t+1]} \sim \mathcal{D}^{t+1}} [\mathbb{1}[\mathcal{A}(\{x_i, y_i\}_{i \in [t]}, x_{t+1}) \neq y_{t+1}]] \\ &= \frac{1}{(t+1)!} \sum_{\sigma \in \text{Sym}(t+1)} \mathbb{E} [\mathbb{1}[\mathcal{A}(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [t]}, x_{\sigma(t+1)}) \neq y_{\sigma(t+1)}]] \\ &= \mathbb{E} \left[\frac{1}{(t+1)!} \sum_{\sigma \in \text{Sym}(t+1)} \mathbb{1}[\mathcal{A}(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [t]}, x_{\sigma(t+1)}) \neq y_{\sigma(t+1)}] \right] \\ &\leq \frac{\mathbb{E} [\text{VCdim}(\mathcal{H}(\{x_i \mid i \in [t+1]\}))]}{t+1} \tag{6} \\ &\leq \frac{\text{VC}_{\text{o}}(\mathcal{H}, \mathcal{G})}{t+1}, \tag{7} \end{aligned}$$

where Eq (6) adopts Lemma 1 and Eq (7) holds due to the definition of $\text{VC}_{\text{o}}(\mathcal{H}, \mathcal{G})$. To convert this algorithm, guaranteeing the expected error upper bounded by $\text{VC}_{\text{o}}(\mathcal{H}, \mathcal{G})$, into an algorithm with high probability $1 - \delta$, we again follow an argument of Haussler et al. (1994). Specifically, the algorithm runs \mathcal{A} for $\lceil \log(2/\delta) \rceil$ times, each time using a new sample of size $\lceil 4\text{VC}_{\text{o}}(\mathcal{H}, \mathcal{G})/\epsilon \rceil$. Then the algorithm selects the hypothesis from the outputs with the minimal error on a new sample of size $\lceil 32/\epsilon(\ln(2/\delta) + \ln(\lceil \log(2/\delta) \rceil + 1)) \rceil$. \square

E Proof of Theorem 5

We first introduce a useful lemma about a well-known Boosting algorithm, known as α -Boost. Given access to a weak learning algorithm, it can output a hypothesis with strong learning guarantee. See Schapire and Freund (2012) for a proof.

Lemma 2 (Boosting). *For any $k, n \in \mathbb{N}$ and multiset $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathcal{Y}$, suppose \mathcal{A}_0 is an algorithm that, for any distribution \mathcal{P} on $\mathcal{X} \times \mathcal{Y}$ with $\mathcal{P}(\{(x_1, y_1), \dots, (x_n, y_n)\}) = 1$, there exists $S_{\mathcal{P}} \in \{(x_1, y_1), \dots, (x_n, y_n)\}^k$ with $\text{err}_{\mathcal{P}}(\mathcal{A}_0(S_{\mathcal{P}})) \leq 1/3$. Then there is a numerical constant $c \geq 1$ such that, for $T = \lceil c \log(n) \rceil$, there exists multisets $S_1, \dots, S_T \in \{(x_1, y_1), \dots, (x_n, y_n)\}^k$ such that, for $\hat{h}(\cdot) = \text{Majority}(\mathcal{A}_0(S_1)(\cdot), \dots, \mathcal{A}_0(S_T)(\cdot))$, it holds that $\hat{h}(x_i) = y_i$ for all $i \in [n]$.*

Part of the proof relies on a well-known generalization bound for compression schemes. The following is the classic result due to Littlestone and Warmuth (1986).

Lemma 3 (Consistent compression generalization bound). *There exists a finite numerical constant $c > 0$ such that, for any compression scheme (κ, ρ) , for any $n \in \mathbb{N}$ and $\delta \in (0, 1)$, for any distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$, for $S \sim \mathcal{D}^n$, with probability at least $1 - \delta$, if $\text{err}_S(\rho(\kappa(S))) = 0$, then*

$$\text{err}_{\mathcal{D}}(\rho(\kappa(S))) \leq \frac{c}{n - |\kappa(S)|} (|\kappa(S)| \log(n) + \log(1/\delta)).$$

Proof of the first part of Theorem 5. The proof is inspired by the idea that representing algorithms by an orientation in a 1-inclusion graph in the transductive setting by Daniely and Shalev-Shwartz (2014). We will first prove a lower bound in the transductive setting and then extend the result to the inductive setting. For any $t \geq 2$, denote $\mu = \mu(\mathcal{H}, \mathcal{G}, t)$ and let ϕ with $|\phi| = t$ and $P \in \Delta(B(\mathcal{H}, \mathcal{G}, \phi))$ be given such that $\mu(\mathcal{H}, \mathcal{G}, \phi, P) \geq \frac{\mu}{2}$. A augmented dataset $\mathcal{G}S = \{(gx, y) | (x, y) \in S, g \in \mathcal{G}\}$ is the same as a multiset of labeled orbits $\{\mathcal{G}x \times \{y\} | (x, y) \in S\}$ up to different data formats. For convenience, we overload the notation a little by also allowing \mathcal{A} being a mapping from a multiset of labeled orbits to a hypothesis. Then we construct a 1-inclusion graph $G_{\mathcal{H}, \mathcal{G}}(\phi) = \{V, E\}$ as introduced in Section 5 and define a mapping $w_{\mathcal{A}} \in W$ as follows. For any edge $e = \{\mathbf{f}, \mathbf{g}, x_i\} \in E$, let

$$w_{\mathcal{A}}(\{\mathbf{f}, \mathbf{g}, x_i\}, \mathbf{f}) = \Pr(\mathcal{A}(\phi_{-i}, \mathbf{f}_{-i}, x_i) = g_i),$$

and

$$w_{\mathcal{A}}(\{\mathbf{f}, \mathbf{g}, x_i\}, \mathbf{g}) = \Pr(\mathcal{A}(\phi_{-i}, \mathbf{f}_{-i}, x_i) = f_i),$$

which is well-defined as $\Pr(\mathcal{A}(\phi_{-i}, \mathbf{f}_{-i}, x_i) = g_i) + \Pr(\mathcal{A}(\phi_{-i}, \mathbf{f}_{-i}, x_i) = f_i) = 1$. Suppose our target function and the instance sequence (\mathbf{f}, \mathbf{x}) is drawn from the distribution P . Then the expected number of mistakes in the transductive learning setting is

$$\begin{aligned} & \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P, \mathcal{A}} \left[\sum_{i=1}^t \mathbb{1}[\mathcal{A}(\phi_{-i}, \mathbf{f}_{-i}, x_i) \neq f_i] \right] \\ & \geq \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P} \left[\sum_{i \in [t]: \exists e \in E, \{\mathbf{f}, x_i\} \subset e} w_{\mathcal{A}}(e(\mathbf{f}, x_i), \mathbf{f}) \right] \\ & \geq \min_w \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P} \left[\sum_{i \in [t]: \exists e \in E, \{\mathbf{f}, x_i\} \subset e} w(e(\mathbf{f}, x_i), \mathbf{f}) \right] \\ & = \mu(\mathcal{H}, \mathcal{G}, \phi, P), \end{aligned} \tag{8}$$

where Eq (8) holds due to the definition of $\mu(\mathcal{H}, \mathcal{G}, \phi, P)$.

Now we prove the lower bound in the inductive setting based on the similar idea. Let ϕ and P be the same as those in the transductive setting. For any $\epsilon < \frac{\mu}{16(t-1)}$, we draw $(\mathbf{f}, \mathbf{x}) \sim P$ and then let \mathbf{f} be our target function and let the marginal data distribution be like, putting probability mass $1 - \frac{16(t-1)\epsilon}{\mu}$ on x_t and the remaining probability mass uniformly over $\{x_1, \dots, x_{t-1}\}$. Denote this data distribution by $\mathcal{D}_{\mathbf{f}, \mathbf{x}}$. For any fixed $i \in [t-1]$, $\Pr(x_i \notin S_{\text{trn}, \mathcal{X}}) = (1 - \frac{16\epsilon}{\mu})^m \geq (\frac{1}{4})^{\frac{16m\epsilon}{\mu}} \geq \frac{1}{2}$

when $m \leq \frac{\mu}{32\epsilon}$. For any hypothesis h , let $\text{err}'_{\mathcal{D}}(h) = \Pr_{(x,y) \sim \mathcal{D}}(h(x) \neq y \text{ and } x \in \{x_i\}_{i \in [t-1]})$ and we always have $\text{err}(h) \geq \text{err}'(h)$. Then we have

$$\begin{aligned}
& \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P, S_{\text{trn}} \sim \mathcal{D}_{\mathbf{f}, \mathbf{x}}^m, \mathcal{A}} [\text{err}'(\mathcal{A}(\mathcal{G}S_{\text{trn}}))] \\
&= \frac{16\epsilon}{\mu} \sum_{i=1}^{t-1} \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P, S_{\text{trn}} \sim \mathcal{D}_{\mathbf{f}, \mathbf{x}}^m, \mathcal{A}} [\mathbb{1}[\mathcal{A}(\mathcal{G}S_{\text{trn}}, x_i) \neq f_i]] \\
&\geq \frac{16\epsilon}{\mu} \sum_{i=1}^{t-1} \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P} \left[\Pr_{S_{\text{trn}} \sim \mathcal{D}_{\mathbf{f}, \mathbf{x}}^m, \mathcal{A}} (\mathcal{A}(\mathcal{G}S_{\text{trn}}, x_i) \neq f_i | x_i \notin S_{\text{trn}, \mathcal{X}}) \Pr(x_i \notin S_{\text{trn}, \mathcal{X}}) \right] \\
&\geq \frac{8\epsilon}{\mu} \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P} \left[\sum_{i=1}^{t-1} \Pr_{S_{\text{trn}} \sim \mathcal{D}_{\mathbf{f}, \mathbf{x}}^m, \mathcal{A}} (\mathcal{A}(\mathcal{G}S_{\text{trn}}, x_i) \neq f_i | x_i \notin S_{\text{trn}, \mathcal{X}}) \right]. \tag{9}
\end{aligned}$$

For all $i \in [t-1]$, for all $z \in \phi_i$, if there is an edge $e = \{\mathbf{f}, \mathbf{g}, z\} \in E$, we let

$$\tilde{w}(e, \mathbf{f}) = \Pr_{S_{\text{trn}} \sim \mathcal{D}_{\mathbf{f}, \mathbf{x}}^m, \mathcal{A}} (\mathcal{A}(\mathcal{G}S_{\text{trn}}, z) \neq f_i | z \notin S_{\text{trn}, \mathcal{X}}),$$

where \mathbf{x} is an arbitrary sequence in $\mathcal{U}_{\mathbf{f}}(\phi)$ satisfying that $x_i = z$. Then $\tilde{w}(e, \mathbf{f})$ is well-defined since the distribution of $\mathcal{G}S_{\text{trn}}$ conditioned on $x_i \notin S_{\text{trn}, \mathcal{X}}$ is the same for all $\mathbf{x} \in \mathcal{U}_{\mathbf{f}}(\phi)$ with $x_i = z$. Actually, conditioned on $x_i \notin S_{\text{trn}, \mathcal{X}}$, the distribution of $\mathcal{G}S_{\text{trn}}$ is also the same when S_{trn} is sampled from $\mathcal{D}_{\mathbf{g}, \mathbf{x}'}$ where \mathbf{x}' is an arbitrary sequence in $\mathcal{U}_{\mathbf{g}}(\phi)$ satisfying that $x'_i = z$. Hence, $\tilde{w}(e, \mathbf{f}) + \tilde{w}(e, \mathbf{g}) = 1$. By letting $\tilde{w}(e, \mathbf{h}) = 0$ for all $\mathbf{h} \notin e$, \tilde{w} is in W . Then we have

$$\begin{aligned}
\text{Eq (9)} &\geq \frac{8\epsilon}{\mu} \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P} \left[\sum_{i \in [t-1]: \exists e \in E, \{\mathbf{f}, x_i\} \subset e} \tilde{w}(e, \mathbf{f}) \right] \\
&\geq \frac{8\epsilon}{\mu} \min_{w \in W} \mathbb{E}_{(\mathbf{f}, \mathbf{x}) \sim P} \left[\sum_{i \in [t]: \exists e \in E, \{\mathbf{f}, x_i\} \subset e} w(e, \mathbf{f}) - 1 \right] \\
&\geq 4\epsilon(1 - 2/\mu) \geq 2\epsilon,
\end{aligned}$$

when $\mu \geq 4$. Hence, there exists a labeling function \mathbf{f} (i.e., there exists a target function h^*) and a data distribution $\mathcal{D}_{\mathbf{f}, \mathbf{x}}$ such that $\mathbb{E}_{S_{\text{trn}}, \mathcal{A}} [\text{err}'(\mathcal{A}(\mathcal{G}S_{\text{trn}}))] \geq 2\epsilon$. Since $\text{err}'(h) \leq \frac{16(t-1)\epsilon}{\mu}$ for all hypothesis h , $\Pr(\text{err}(\mathcal{A}(\mathcal{G}S_{\text{trn}})) > \epsilon) \geq \Pr(\text{err}'(\mathcal{A}(\mathcal{G}S_{\text{trn}})) > \epsilon) > \frac{\mu}{16(t-1)}$. \square

Proof of the second and third parts of Theorem 5. For any $n \in \mathbb{N}$, for any given sample $\{(x_1, y_1), \dots, (x_{n+1}, y_{n+1})\}$, let $\phi = \{\phi_1, \dots, \phi_{n+1}\} = \{\mathcal{G}x_1, \dots, \mathcal{G}x_{n+1}\}$ denote the multi-set of $t+1$ orbits and construct the one-inclusion graph $G_{\mathcal{H}, \mathcal{G}}(\phi)$. As mentioned in Definition 6, every $w \in W$ defines a randomized orientation of each edge in graph $G_{\mathcal{H}, \mathcal{G}}(\phi)$. That is, for any fixed $w \in W$, for every edge $e = \{\mathbf{f}, \mathbf{g}, x\}$, w defines a probability over $\{\mathbf{f}, \mathbf{g}\}$. Then we can construct an algorithm \mathcal{A}_w for each $w \in W$.

Given the input $(\phi_{-i}, \mathbf{y}_{-i}, x_i)$, \mathcal{A}_w finds the subset of vertices $\{(\mathbf{f}, x_i) | \mathbf{f}_{-i} = \mathbf{y}_{-i}\}$ whose labelings are consistent with $(\phi_{-i}, \mathbf{y}_{-i})$. If there exist two such vertices, (\mathbf{f}, x_i) and (\mathbf{g}, x_i) , they must be connected by $e = (\mathbf{f}, \mathbf{g}, x_i)$ due to the definition of $G_{\mathcal{H}, \mathcal{G}}(\phi)$. Then \mathcal{A}_w will predict x_i as f_i with probability $w(e, \mathbf{g})$ and as g_i with probability $w(e, \mathbf{f})$. If only one such vertex (\mathbf{f}, x_i) exists, \mathcal{A}_w predicts the label of x_i by f_i . Due to the realizable setting, there must exist at least one such vertex and the algorithm's prediction must be correct when only one vertex exists.

To complete the algorithm, the remaining part is how to choose a good $w \in W$. For any true labeling \mathbf{f}^* and any sequence of natural data $\mathbf{x}^* \in \mathcal{U}_{\mathbf{f}^*}(\phi)$, for each $i \in [n+1]$, if there is an edge $e \supset \{\mathbf{f}^*, x_i^*\}$, it means the algorithm possibly misclassify x_i^* (with the probability dependent on w); if there is no such an edge, it means the algorithm will not misclassify x_i^* no matter what w is. For any labeling $\mathbf{f} \in \Pi_{\mathcal{H}}(\phi)$ and a sequence of natural data $\mathbf{x} \in \mathcal{U}_{\mathbf{f}}(\phi)$, we can represent the subset of the points in \mathbf{x} that the algorithm is uncertain about by a mapping $a_{\mathbf{f}, \mathbf{x}} : E \times \Pi_{\mathcal{H}}(\phi) \mapsto \{0, 1\}$ where $a_{\mathbf{f}, \mathbf{x}}(e, \mathbf{g}) = 1$ iff. $\mathbf{g} = \mathbf{f}$ and there exists $i \in [n+1]$ s.t. $\{\mathbf{f}, x_i\} \in e$. Due to the definition, $a_{\mathbf{f}, \mathbf{x}}$ has at most $n+1$ non-zero entries. Let $A = \{a_{\mathbf{f}, \mathbf{x}} | \mathbf{f} \in \Pi_{\mathcal{H}}(\phi), \mathbf{x} \in \mathcal{U}_{\mathbf{f}}(\phi)\}$ denote the set of

all such mappings. We now first consider the case where $|E| < \infty$. Then for a training set of size n , we can rewrite the expected error as

$$\begin{aligned}
& \mathbb{E}_{S \sim \mathcal{D}^n} [\text{err}(\mathcal{A}_w(\mathcal{G}S))] \\
&= \mathbb{E}_{S \sim \mathcal{D}^n, (x, y) \sim \mathcal{D}, \mathcal{A}_w} [\mathbb{1}[\mathcal{A}_w(\mathcal{G}S, x) \neq y]] \\
&= \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}^{n+1}, \mathcal{A}_w} \left[\frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1}[\mathcal{A}_w(\phi_{-i}, \mathbf{y}_{-i}, x_i) \neq y_i] \right] \\
&= \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}^{n+1}, \mathcal{A}_w} \left[\frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1}[\mathcal{A}_w(\phi_{-i}, \mathbf{y}_{-i}, x_i) \neq y_i] \mathbb{1}[\exists e \supset \{x_i, \mathbf{f}^*\}] \right] \\
&= \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}^{n+1}} \left[\frac{1}{n+1} \sum_{e \in E} w(e, \mathbf{f}^*) a_{\mathbf{f}^*, \mathbf{x}}(e, \mathbf{f}^*) \right] \\
&= \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}^{n+1}} \left[\frac{1}{n+1} \sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w(e, \mathbf{f}) a_{\mathbf{f}^*, \mathbf{x}}(e, \mathbf{f}) \right]. \tag{10}
\end{aligned}$$

where the last equality holds due to $a_{\mathbf{f}^*, \mathbf{x}}(e, \mathbf{f}) = 0$ for all $\mathbf{f} \neq \mathbf{f}^*$. Since \mathbf{f}^* and \mathbf{x} is unknown, our goal is to find a w with $\sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w(e, \mathbf{f}) a_{\mathbf{f}^*, \mathbf{x}}(e, \mathbf{f})$ upper bounded for all \mathbf{f}^* and \mathbf{x} . The algorithm picks $w^* = \arg \min_{w \in W} \max_{a \in A} \sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w(e, \mathbf{f}) a(e, \mathbf{f})$. Then we have

$$\begin{aligned}
& \max_{a \in A} \sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w^*(e, \mathbf{f}) a(e, \mathbf{f}) = \min_{w \in W} \max_{a \in A} \sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w(e, \mathbf{f}) a(e, \mathbf{f}) \\
&= \min_{w \in W} \max_{a \in \text{Conv}(A)} \sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w(e, \mathbf{f}) a(e, \mathbf{f}) = \max_{a \in \text{Conv}(A)} \min_{w \in W} \sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w(e, \mathbf{f}) a(e, \mathbf{f}), \tag{11}
\end{aligned}$$

where the last equality is due to Minimax theorem. Since the optimal solution a^* to Eq. (11) is in the convex hull of A , there is a distribution $P^* \in \Delta(B(\mathcal{H}, \mathcal{G}, \phi))$ such that $a^* = \mathbb{E}_{(\mathbf{f}^*, \mathbf{x}) \sim P^*} [a_{\mathbf{f}^*, \mathbf{x}}]$. Then we have

$$\begin{aligned}
\text{Eq (11)} &= \min_{w \in W} \sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w(e, \mathbf{f}) \mathbb{E}_{(\mathbf{f}^*, \mathbf{x}) \sim P^*} [a_{\mathbf{f}^*, \mathbf{x}}(e, \mathbf{f})] \\
&= \min_{w \in W} \mathbb{E}_{(\mathbf{f}^*, \mathbf{x}) \sim P^*} \left[\sum_{e \in E, \mathbf{f} \in \Pi_{\mathcal{H}}(\phi)} w(e, \mathbf{f}) a_{\mathbf{f}^*, \mathbf{x}}(e, \mathbf{f}) \right] \\
&= \min_{w \in W} \mathbb{E}_{(\mathbf{f}^*, \mathbf{x}) \sim P^*} \left[\sum_{i \in [n+1]: \exists e \in E, \{\mathbf{f}^*, x_i\} \subset e} w(e(\mathbf{f}^*, x_i), \mathbf{f}^*) \right] \\
&= \mu(\mathcal{H}, \mathcal{G}, \phi, P^*).
\end{aligned}$$

Combined with Eq (10), we have $\mathbb{E}_{S \sim \mathcal{D}^n} [\text{err}(\mathcal{A}_{w^*}(\mathcal{G}S))] \leq \frac{\mu(\mathcal{H}, \mathcal{G}, \phi, n+1)}{n+1}$.

For $|E| = \infty$, $E \times \Pi_{\mathcal{H}}(\phi)$ could be infinite dimensional and thus, we need to use Sion's minimax theorem. The details of how to apply Sion's minimax theorem are described as follows. For all $w \in W$, for any edge $e = \{\mathbf{f}, \mathbf{g}, x\}$, we have $w(e, \mathbf{f}) + w(e, \mathbf{g}) = 1$ and thus, there exists a one-to-one mapping $\beta : W \mapsto [0, 1]^E$ where $\beta(w)(e) = w(e, \mathbf{f}_{e,0})$ where $\mathbf{f}_{e,0}$ is the labeling in e predicting x as zero. In the following, we will overload the notation by using w to represent $\beta(w)$ when it is clear from the context that it is in the space $[0, 1]^E$. Then we define a mapping $\text{BL}(\cdot, \cdot) : [0, 1]^E \times A$ by

$$\text{BL}(w, a) := \sum_{e: a(e, \mathbf{f}_{e,0})=1} w(e) + \sum_{e: a(e, \mathbf{f}_{e,1})=1} (1 - w(e)),$$

where $\mathbf{f}_{e,0}, \mathbf{f}_{e,1}$ are labelings in $e = \{\mathbf{f}_{e,0}, \mathbf{f}_{e,1}, x\}$ and they label x as 0 and 1 respectively. Then similar to Eq (10), we can represent the expected error as

$$\mathbb{E}_{S \sim \mathcal{D}^n} [\text{err}(\mathcal{A}_w(\mathcal{G}S))] = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}^{n+1}} \left[\frac{1}{n+1} \text{BL}(w, a_{\mathbf{f}^*, \mathbf{x}}) \right].$$

Now we want to upper bound $\min_{w \in [0,1]^E} \sup_{a \in A} \text{BL}(w, a)$. $\text{BL}(\cdot, \cdot)$ can be extended to $\mathbb{R}^E \times A$ by letting $\text{BL}(w, a) = \sum_{e: a(e, \mathbf{f}_{e,0})=1} w(e) + \sum_{e: a(e, \mathbf{f}_{e,1})=1} (1 - w(e))$ for $w \in \mathbb{R}^E$. Since a has at most $n + 1$ non-zeros entries, $|\text{BL}(w, a)| \leq (n + 1) \max(\|w\|_\infty, \|1 - w\|_\infty)$. Let $\tilde{A} = \{\sum_{i=1}^N c_i a_i \mid \forall i, a_i \in A, c_i > 0, \sum_{i=1}^N c_i = 1, N \in \mathbb{N}\}$ be the set of all finite convex combination of elements in A . We extend BL from $\mathbb{R}^E \times A$ to $\mathbb{R}^E \times \tilde{A}$ by defining $\text{BL}(w, a) = \sum_{i=1}^N c_i \text{BL}(w, a_i)$ for $a = \sum_{i=1}^N c_i a_i \in \tilde{A}$.

We define a metric d in \tilde{A} : for $a = \sum_{i=1}^N c(a_i) a_i$ and $a' = \sum_{i=1}^{N'} c'(a'_i) a'_i$, the distance between a and a' is $d(a, a') := \sum_{\alpha \in A: c(\alpha) \neq 0 \text{ or } c'(\alpha) \neq 0} |c(\alpha) - c'(\alpha)|$. For any fixed $w \in \mathbb{R}^E$, for any $a \in \tilde{A}$, for every open ball $\mathcal{B}_r(\text{BL}(w, a))$ centered at $\text{BL}(w, a)$ with radius $r > 0$, there is an open ball $\mathcal{B}_{r'}(a)$ in the metric space (\tilde{A}, d) with $r' = \frac{r}{(n+1)(\max(\|w\|_\infty, \|1-w\|_\infty))}$ such that for all $a' \in \mathcal{B}_{r'}(a)$, $|\text{BL}(w, a') - \text{BL}(w, a)| \leq \sum_{\alpha \in \{a_i\}_{i \in [N]} \cup \{a'_i\}_{i \in [N']}} |c(\alpha) - c'(\alpha)| |\text{BL}(w, \alpha)| < r' \cdot |\text{BL}(w, \alpha)| \leq r$. Hence, $\text{BL}(w, \cdot)$ is continuous for all $w \in \mathbb{R}^E$.

Consider the the standard topology in \mathbb{R} and then $[0, 1]$ is compact. Then let \mathcal{T} be the product topology of \mathbb{R}^E . Then by Tychonoff theorem, $[0, 1]^E$ is compact in \mathbb{R}^E . For any fixed $a = \sum_{i=1}^N c_i a_i \in \tilde{A}$, there are at most $N(n + 1)$ non-zero entries. Then for any $w \in \mathbb{R}^E$, for every open ball $\mathcal{B}_r(\text{BL}(w, a))$ centered at $\text{BL}(w, a)$ with radius $r > 0$, then there is a neighborhood $U = \prod_{e \in E} S_e$, where $S_e = (w_e - \frac{r}{N(n+1)}, w_e + \frac{r}{N(n+1)})$ if at least one of $a(e, \mathbf{f}_{e,0}), a(e, \mathbf{f}_{e,1})$ is non-zero and $S_e = \mathbb{R}$ for other $e \in E$, such that $|\text{BL}(w', a) - \text{BL}(w, a)| < r$ for all $w' \in U$. That is, $\text{BL}(U) \subset \mathcal{B}_r(\text{BL}(w, a))$. Hence, $\text{BL}(\cdot, a)$ is continuous for all $a \in \tilde{A}$.

It is easy to check that $\text{BL}(\cdot, a)$ and $\text{BL}(w, \cdot)$ are linear for all $a \in \tilde{A}, w \in W$. Then by Sion's minimax theorem, we have

$$\min_{w \in [0,1]^E} \sup_{a \in A} \text{BL}(w, a) \leq \min_{w \in [0,1]^E} \sup_{a \in \tilde{A}} \text{BL}(w, a) = \sup_{a \in \tilde{A}} \min_{w \in [0,1]^E} \text{BL}(w, a) =: v^*.$$

Let $w^* = \arg \min_{w \in [0,1]^E} \sup_{a \in A} \text{BL}(w, a)$. There exists a sequence a_1, a_2, \dots in \tilde{A} such that $\lim_{k \rightarrow \infty} \min_{w \in [0,1]^E} \text{BL}(w, a_k) = v^*$. For each a_k , we let $P_k \in \Delta(B(\mathcal{H}, \mathcal{G}, \phi))$ be the distribution over $B(\mathcal{H}, \mathcal{G}, \phi)$ such that $a_k = \mathbb{E}_{(\mathbf{f}^*, \mathbf{x}) \sim P_k} [a_{\mathbf{f}^*, \mathbf{x}}]$. Due to the definition of \tilde{A} , we know that P_k is a discrete distribution with finite support. Then we have

$$\begin{aligned} \sup_{a \in A} \text{BL}(w^*, a) &\leq \min_{w \in [0,1]^E} \sup_{a \in \tilde{A}} \text{BL}(w, a) \\ &= \sup_{a \in \tilde{A}} \min_{w \in [0,1]^E} \text{BL}(w, a) \\ &= \lim_{k \rightarrow \infty} \min_{w \in [0,1]^E} \text{BL}(w, \mathbb{E}_{(\mathbf{f}^*, \mathbf{x}) \sim P_k} [a_{\mathbf{f}^*, \mathbf{x}}]) \\ &= \lim_{k \rightarrow \infty} \min_{w \in [0,1]^E} \mathbb{E}_{(\mathbf{f}^*, \mathbf{x}) \sim P_k} [\text{BL}(w, a_{\mathbf{f}^*, \mathbf{x}})] \\ &\leq \sup_{P \in \Delta(B(\mathcal{H}, \mathcal{G}, \phi))} \min_{w \in [0,1]^E} \mathbb{E}_{(\mathbf{f}^*, \mathbf{x}) \sim P} [\text{BL}(w, a_{\mathbf{f}^*, \mathbf{x}})] \\ &= \mu(\mathcal{H}, \mathcal{G}, \phi, n + 1). \end{aligned}$$

Hence, $\mathbb{E}_{S \sim \mathcal{D}^n} [\text{err}(\mathcal{A}_{w^*}(\mathcal{G}S))] \leq \frac{\mu(\mathcal{H}, \mathcal{G}, \phi, n+1)}{n+1}$.

The first upper bound We can convert the above bound into a high probability bound by α -Boost. Let $\mathcal{A}_0 = \mathcal{A}_{w^*}$ as defined above. Let t be any positive integer such that $\frac{1}{6} \leq \frac{\mu(\mathcal{H}, \mathcal{G}, t)}{t} \leq \frac{1}{3}$. As established above, for $S_{\text{trn}} = \{(x_1, y_1), \dots, (x_m, y_m)\} \sim \mathcal{D}^m$ and any distribution \mathcal{P} supported on $\{(x_1, y_1), \dots, (x_m, y_m)\}$, for $k = t - 1$ and $S \sim \mathcal{P}^k$, $\mathbb{E}[\text{err}_{\mathcal{P}}(\mathcal{A}_0(S))] \leq 1/3$. Thus given S_{trn} and \mathcal{P} , there exists a deterministic choice of $S_{\mathcal{P}} \in \{(x_1, y_1), \dots, (x_m, y_m)\}^k$ with $\text{err}_{\mathcal{P}}(\mathcal{A}_0(S_{\mathcal{P}})) \leq 1/3$. Then Lemma 2 implies that for a value $T = \lceil c_1 \log m \rceil$ (for numerical constant $c_1 \geq 1$), there exists $S_1, \dots, S_T \in \{(x_1, y_1), \dots, (x_m, y_m)\}^k$ such that, for $\hat{h}(\cdot) = \text{Majority}(\mathcal{A}_0(S_1)(\cdot), \dots, \mathcal{A}_0(S_T)(\cdot))$, it holds that $\hat{h}(x_i) = y_i$ for all $i \in [m]$. Note that \hat{h}

can be expressed as a compression scheme. By Lemma 3, with probability at least $1 - \delta$,

$$\text{err}(\hat{h}) \leq \frac{c_2}{m - kT} \left(kT \log m + \log \frac{1}{\delta} \right),$$

for a numerical constant $c_2 \geq 1$. Thus, for any given $\epsilon \in (0, 1)$, the right hand side can be made less than ϵ for an appropriate choice of

$$m = O\left(\frac{1}{\epsilon}(k \log^2 \frac{k}{\epsilon} + \log \frac{1}{\delta})\right),$$

where $k = t - 1 \leq 6\mu(\mathcal{H}, \mathcal{G}, t) - 1$.

The second upper bound Again, using the same standard technique as we used in Theorem 4 to convert an algorithm with expected error upper bound to an algorithm with high probability guarantee. The algorithm runs \mathcal{A}_{w^*} for $\lceil \log(2/\delta) \rceil$ times, each time using a new sample of size $\lceil 4\mu(\mathcal{H}, \mathcal{G})/\epsilon \rceil$. Then the algorithm selects the hypothesis from the outputs with the minimal error on a new sample of size $\lceil 32/\epsilon(\ln(2/\delta) + \ln(\lceil \log(2/\delta) \rceil + 1)) \rceil$. \square

F Proof of Theorem 6

Proof. Let $d = \dim(\mathcal{H}, \mathcal{G})$. Let ϕ and the corresponding $B = \{(\mathbf{f}, \mathbf{x}_{\mathbf{f}})\}_{\mathbf{f} \in \mathcal{Y}^d}$ be given. Let P be the uniform distribution over B . Then

$$\begin{aligned} \mu(\mathcal{H}, \mathcal{G}, \phi, P) &= \min_{w \in W} \mathbb{E}_{(\mathbf{f}, \mathbf{x}_{\mathbf{f}}) \sim P} \left[\sum_{i \in [d]: \exists e \in E, \{\mathbf{f}, \mathbf{x}_{\mathbf{f}, i}\} \subseteq e} w(e, \mathbf{f}) \right] \\ &= \min_w \frac{1}{2^d} \sum_{\mathbf{f} \in \mathcal{Y}^d} \sum_{i \in [d]: \exists e \in E, \{\mathbf{f}, \mathbf{x}_{\mathbf{f}, i}\} \subseteq e} w(e, \mathbf{f}) \\ &= \frac{1}{2^d} \min_w \sum_{\mathbf{f} \in \mathcal{Y}^d} \sum_{i=1}^d \frac{1}{2} (w(e(\mathbf{f}, \mathbf{x}_{\mathbf{f}, i}), \mathbf{f}) + w(e(\mathbf{f}, \mathbf{x}_{\mathbf{f}, i}), (1 - f_i, \mathbf{f}_{-i}))) \\ &= \frac{1}{2^d} \cdot \frac{2^d \cdot d}{2} = \frac{d}{2}, \end{aligned} \tag{12}$$

where Eq (12) holds since $x_{\mathbf{f}, i} = x_{\mathbf{g}, i}$ if $\mathbf{f} \oplus \mathbf{g} = \mathbf{e}_i$ due to the definition of $\dim(\mathcal{H}, \mathcal{G})$. \square

G Proof of Theorem 7

Proof of the lower bound. Let $d = \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$, let $X_d = \{x_1, \dots, x_d\}$ be a set shattered in the way defined in Definition 5 and let $\mathcal{H}' := \{h|_{X_d} | h \in \mathcal{H}\}$. Then $\text{VCdim}(\mathcal{H}') = d$ and any lower bound on the sample complexity of PAC learning of \mathcal{H}' is also a lower bound on the sample complexity of relaxed realizable PAC learning of \mathcal{H} . The lower bound follow by standard arguments from Vapnik and Chervonenkis (1974); Blumer et al. (1989); Ehrenfeucht et al. (1989). \square

Proof of the upper bound of the algorithm ERM-INV. The algorithm ERM-INV works as follows. ERM-INV first applies ERM over the original data set. Then for every test instance x , if x lies in the orbits generated by the original data, the algorithm predicts x by the label of the training instance in the same orbit; otherwise, the algorithm predicts according to the ERM output. Specifically, given the training set $S_{\text{trn}} = \{(x_1, y_1), \dots, (x_m, y_m)\} \sim \mathcal{D}^m$, the algorithm finds a hypothesis $h \in \mathcal{H}$ consistent with S_{trn} and then outputs $f_{h, S_{\text{trn}}}$ defined by

$$f_{h, S_{\text{trn}}}(x) = \begin{cases} y_i & \text{if there exists } i \in [m] \text{ s.t. } x \in \mathcal{G}x_i, \\ h(x) & \text{o.w.} \end{cases} \tag{13}$$

The function $f_{h, S_{\text{trn}}}$ is well-defined a.s. when the data distribution \mathcal{D} is \mathcal{G} -invariant since if there exists $i \neq j$ such that $x_i \in \mathcal{G}x_j$, then $y_i = y_j = h^*(x_j)$.

The proof idea is similar to that of Theorem 1. Let $d = \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$. Consider two sets S and S' of m i.i.d. samples drawn from the data distribution \mathcal{D} each. We denote A_S the event of $\{\exists h \in \mathcal{H}, \text{err}_S(h) = 0, \text{err}_{\mathcal{D}}(f_{h,S}) \geq \epsilon\}$ and $B_{S,S'} = \{\exists h \in \mathcal{H}, \text{err}_S(h) = 0, \text{err}_{S'}(f_{h,S}) \geq \frac{\epsilon}{2}\}$. By Chernoff bound, we have $\Pr(B_{S,S'}) \geq \Pr(A_S) \cdot \Pr(B_{S,S'}|A_S) \geq \frac{1}{2} \Pr(A_S)$ when $m \geq \frac{8}{\epsilon}$. The sampling process of S and S' is equivalent to drawing $2m$ i.i.d. samples and then randomly partitioning into S and S' of m each. For any fixed S'' , let us divide S'' into two categories in terms of the number of examples in each orbit. Let $R_1 = \{x | |\mathcal{G}x \cap S''_{\mathcal{X}}| \geq \log^2 m\}$ and $R_2 = S''_{\mathcal{X}} \setminus R_1$. Let $\mathcal{H}_0 \subset \mathcal{H}$ denote the set hypotheses making at least $\frac{\epsilon m}{2}$ mistakes in S'' . Now we divide \mathcal{H}_0 into two sub-classes as follows.

- Let $\mathcal{H}_1 = \{h \in \mathcal{H}_0 | h \text{ makes fewer than } \frac{\epsilon m}{4} \text{ mistakes in } R_2\}$. Let $\mathcal{F}_{S''}(\mathcal{H}_1) := \{f_{h,T} | h \in \mathcal{H}_1, T \subset S'', |T| = m\}$. For any h , if S is correctly labeled, then $\text{err}_{\mathcal{G}S}(f_{h,S}) = 0$. Thus we have

$$\begin{aligned} & \Pr(\exists h \in \mathcal{H}_1, \text{err}_S(h) = 0, \text{err}_{S'}(f_{h,S}) \geq \frac{\epsilon}{2}) \\ & \leq \Pr(\exists h \in \mathcal{H}_1, \text{err}_{\mathcal{G}S}(f_{h,S}) = 0, \text{err}_{S'}(f_{h,S}) \geq \frac{\epsilon}{2}) \\ & \leq \Pr(\exists f \in \mathcal{F}_{S''}(\mathcal{H}_1), \text{err}_{\mathcal{G}S}(f) = 0, \text{err}_{S'}(f) \geq \frac{\epsilon}{2}). \end{aligned}$$

Since every $h \in \mathcal{H}_1$ makes fewer than $\frac{\epsilon m}{4}$ mistakes in R_2 , for all $f \in \mathcal{F}_{S''}(\mathcal{H}_1)$, if f makes at least $\frac{\epsilon m}{2}$ mistakes in S'' , it must makes at least $\frac{\epsilon m}{4}$ mistakes in R_1 . Similar to the case 1 in the proof of Theorem 1, for any $f \in \mathcal{F}_{S''}(\mathcal{H}_1)$, we let $X(f) \subset R_1$ denote a minimal set of examples in R_1 (breaking ties arbitrarily but in a fixed way) such that f misclassify $X(f)$ and $|\mathcal{G}X(f) \cap S''_{\mathcal{X}}| \geq \frac{\epsilon m}{4}$ where $\mathcal{G}X(f) = \{gx | x \in X(f), g \in \mathcal{G}\}$ is the set of all examples lying in the orbits generated from $X(f)$. Let $K(f) = \mathcal{G}X(f)$ and $\mathcal{K} = \{K(f) | f \in \mathcal{F}_{S''}(\mathcal{H}_1)\}$. Notice that each example in $X(f)$ must belong to different orbits, otherwise it is not minimal. Besides, each orbit contains at least $\log^2 m$ examples from $S''_{\mathcal{X}}$. Hence, $|X(f)| \leq \frac{\epsilon m}{4 \log^2 m}$. Since there are at most $\frac{2m}{\log^2 m}$ orbits generated from R_1 , we have $|\mathcal{K}| \leq \sum_{i=1}^{\frac{\epsilon m}{4 \log^2 m}} \binom{2m}{i} \leq \left(\frac{8e}{\epsilon}\right)^{\frac{\epsilon m}{4 \log^2 m}}$. Since f misclassify $X(f)$, all examples in $K(f)$ must go to S' to guarantee $\text{err}_{\mathcal{G}S}(f) = 0$ and thus, we have

$$\begin{aligned} & \Pr(\exists f \in \mathcal{F}_{S''}(\mathcal{H}_1), \text{err}_{\mathcal{G}S}(f) = 0, \text{err}_{S'}(f) \geq \frac{\epsilon}{2}) \\ & \leq \Pr(\exists f \in \mathcal{F}_{S''}(\mathcal{H}_1), K(f) \cap S_{\mathcal{X}} = \emptyset) = \Pr(\exists K \in \mathcal{K}, K \cap S_{\mathcal{X}} = \emptyset) \\ & \leq \sum_{K \in \mathcal{K}} 2^{-\frac{\epsilon m}{4}} \leq \left(\frac{8e}{\epsilon}\right)^{\frac{\epsilon m}{4 \log^2 m}} \cdot 2^{-\frac{\epsilon m}{4}} = 2^{-\frac{\epsilon m}{4} \left(1 - \frac{\log(8e/\epsilon)}{\log^2 m}\right)} \leq 2^{-\frac{\epsilon m}{8}}, \end{aligned}$$

when $m \geq \frac{8e}{\epsilon} + 4$.

- Let $\mathcal{H}_2 = \mathcal{H}_0 \setminus \mathcal{H}_1$. Now we will bound $\Pr(\exists h \in \mathcal{H}_2, \text{err}_S(h) = 0, \text{err}_{S'}(f_{h,S}) \geq \frac{\epsilon}{2})$. Similar to the case 2 in Theorem 1, every $h \in \mathcal{H}_2$ will make at least $\frac{\epsilon m}{4}$ mistakes in R_2 . Since $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = d$ and every orbit generated from R_2 contains fewer than $\log^2 m$ examples, the number of examples in R_2 that can be shattered by \mathcal{H} is no greater than $d \log^2 m$. Thus, the number of ways labeling examples in R_2 is upper bounded by $\left(\frac{2em}{d}\right)^{d \log^2 m}$ by Sauer's lemma. For any multi-subset $X \subset S''_{\mathcal{X}}$ and hypothesis h , we denote by $\widehat{M}_X(h)$ the number of instances in X misclassified by h . Hence, we have

$$\begin{aligned} & \Pr(\exists h \in \mathcal{H}_2, \text{err}_S(h) = 0, \text{err}_{S'}(f_{h,S}) \geq \frac{\epsilon}{2}) \\ & \leq \Pr(\exists h \in \mathcal{H}_2, \widehat{M}_{S_{\mathcal{X}} \cap R_2}(h) = 0, \widehat{M}_{S'_{\mathcal{X}} \cap R_2}(h) \geq \frac{\epsilon m}{4}) \\ & \leq \left(\frac{2em}{d}\right)^{d \log^2 m} \cdot 2^{-\frac{\epsilon m}{4}} = 2^{-\frac{\epsilon m}{4} + d \log^2 m \log(2em/d)} \end{aligned}$$

Combining the results for \mathcal{H}_1 and \mathcal{H}_2 , we have

$$\Pr(B_{S,S'}) \leq 2^{-\frac{\epsilon m}{8}} + 2^{-\frac{\epsilon m}{4} + d \log^2 m \log(2em/d)} \leq \frac{\delta}{2},$$

when $m \geq \frac{8}{\epsilon} (d \log^2 m \log \frac{2em}{d} + \log \frac{4}{\delta} + e) + 4$. \square

Proof of the upper bound of the 1-inclusion-graph predictor. The algorithm is similar to the 1-inclusion-graph predictor in Theorem 4. For any $t \in \mathbb{N}$ and $S = \{(x_1, y_1), \dots, (x_t, y_t)\}$, let X_S be the set of different elements in $S_{\mathcal{X}}$ and $\mathcal{H}'(X_S) := \{h_{|X_S} | \forall x', x \in X_S, x' \in \mathcal{G}x \text{ implies } h(x') = h(x)\}$. Here $\mathcal{H}'(X_S)$ is different from $\mathcal{H}(X_S)$ defined in Theorem 4 in the sense that every hypothesis in $\mathcal{H}'(X_S)$ is not (\mathcal{G}, X_S) -invariant but only predict the observed examples in the same orbit in the same way. Note that $h_{|X_S}^*$ is in $\mathcal{H}'(X_S)$ if S is realized by h^* . Let $Q_{\mathcal{H}'(X_S), X_S}$ be the function guaranteed by Eq (1) for the instance space X_S and hypothesis class $\mathcal{H}'(X_S)$. Given a set S of t i.i.d. samples, we let $\mathcal{A}(S)$ be defined as $\mathcal{A}(S, x) = Q_{\mathcal{H}'(X_S \cup \{x\}), X_S \cup \{x\}}(S, x)$ if $\mathcal{H}'(X_S \cup \{x\})$ is nonempty and predicting arbitrarily if it is empty. Following the analogous proof of Theorem 4, we have

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}^t} [\text{err}(\mathcal{A}(S))] &= \mathbb{E}_{(x_i, y_i)_{i \in [t+1]} \sim \mathcal{D}^{t+1}} [\mathbb{1}[\mathcal{A}(\{x_i, y_i\}_{i \in [t]}, x_{t+1}) \neq y_{t+1}]] \\ &= \frac{1}{(t+1)!} \sum_{\sigma \in \text{Sym}(t+1)} \mathbb{E} [\mathbb{1}[\mathcal{A}(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [t]}, x_{\sigma(t+1)}) \neq y_{\sigma(t+1)}]] \\ &= \mathbb{E} \left[\frac{1}{(t+1)!} \sum_{\sigma \in \text{Sym}(t+1)} \mathbb{1}[\mathcal{A}(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [t]}, x_{\sigma(t+1)}) \neq y_{\sigma(t+1)}] \right] \\ &\leq \frac{\mathbb{E} [\text{VCdim}(\mathcal{H}'(\{x_i | i \in [t+1]\}))]}{t+1} \leq \frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{t+1}. \end{aligned}$$

Again, we use the same method as that in Theorem 4 to convert this algorithm, guaranteeing the expected error upper bounded by $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$, into an algorithm with high probability $1 - \delta$. Specifically, the algorithm runs \mathcal{A} for $\lceil \log(2/\delta) \rceil$ times, each time using a new sample of size $\lceil 4\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})/\epsilon \rceil$. Then the algorithm selects the hypothesis from the outputs with the minimal error on a new sample of size $\lceil 32/\epsilon(\ln(2/\delta) + \ln(\lceil \log(2/\delta) \rceil + 1)) \rceil$. \square

H Proof of Theorem 8

To prove the theorem, we will use a generalization bound for agnostic compression scheme by Graepel et al. (2005).

Lemma 4 (Agnostic compression generalization bound). *There exists a finite numerical constant $c > 0$ such that, for any compression scheme (κ, ρ) , for any $n \in \mathbb{N}$ and $\delta \in (0, 1)$, for any distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$, for $S \sim \mathcal{D}^n$, letting $B(S, \delta) := \frac{1}{n}(|\kappa(S)| \log(n) + \log(1/\delta))$, with probability at least $1 - \delta$, then*

$$|\text{err}_{\mathcal{D}}(\rho(\kappa(S))) - \text{err}_S(\rho(\kappa(S)))| \leq c\sqrt{\text{err}_S(\rho(\kappa(S)))B(S, \delta)} + cB(S, \delta).$$

Proof of the lower bound. For the lower bound, our construction follows Ben-David and Uner (2014). For any \mathcal{X}, \mathcal{H} , let A_1, \dots, A_d be subsets of \mathcal{X} such that the orbits of every two different elements $x, x' \in \cup_{i \in [d]} A_i$ are disjoint, i.e., $\forall x, x' \in \cup_{i \in [d]} A_i, \mathcal{G}x \cap \mathcal{G}x' = \emptyset$. We say \mathcal{H} set-shatters A_1, \dots, A_d if for every binary vector $\mathbf{y} \in \mathcal{Y}^d$, there exists some $h_{\mathbf{y}} \in \mathcal{H}$ such that for all $i \in [d]$ and $x \in \mathcal{X}$, if $x \in A_i$ then $h_{\mathbf{y}}(x) = y_i$. Then by following Theorem 7 of Ben-David and Uner (2014), if \mathcal{H} set-shatters A_1, \dots, A_d for some infinite subsets A_1, \dots, A_d of \mathcal{X} , the standard agnostic PAC sample complexity of learning \mathcal{H} under deterministic labels for instance space being $\cup_{i \in [d]} A_i$ is lower bounded by $\frac{d}{\epsilon^2}$ for all $\delta < 1/32$. Since any data distribution \mathcal{D} with $\mathcal{D}_{\mathcal{X}}(\cup_{i \in [d]} A_i) = 1$ is \mathcal{G} -invariant, the above lower bound also lower bounds $\mathcal{M}_{\text{AG}}(\epsilon, \delta; \mathcal{H}, \mathcal{G})$ for all $\delta < 1/32$. Let $h(x) = 0$ for all $x \notin \cup_{i \in [k]} A_i$ and for all $h \in \mathcal{H}$, $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) = d$. Thus, $\mathcal{M}_{\text{AG}}(\epsilon, 1/64; \mathcal{H}, \mathcal{G}) = \Omega(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon^2})$. The construction above works for any $d > 0$. \square

Proof of the upper bound of ERM-INV. Let $d = \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$. Consider two sets S and S' of m i.i.d. samples drawn from the data distribution \mathcal{D} each. We denote A_S the event of $\{\exists h \in \mathcal{H}, \text{err}_{\mathcal{D}}(f_{h,S}) \geq \text{err}_S(h) + \epsilon\}$ and $B_{S,S'} = \{\exists h \in \mathcal{H}, \text{err}_{S'}(f_{h,S}) \geq \text{err}_S(h) + \frac{\epsilon}{2}\}$. By Hoeffding bound, we have $\Pr(B_{S,S'}) \geq \Pr(A_S) \cdot \Pr(B_{S,S'} | A_S) \geq \frac{1}{2} \Pr(A_S)$ when $m \geq \frac{2}{\epsilon^2}$. The sampling process of S and S' is equivalent to drawing $2m$ i.i.d. samples and then randomly partitioning into S and S' of m each. For any fixed S'' , let us divide S'' into two categories in terms of the number of examples

in each orbit. Let $R_1 = \{x \mid |\mathcal{G}x \cap S''_{\mathcal{X}}| \geq \log^2 m\}$ and $R_2 = S''_{\mathcal{X}} \setminus R_1$. For any multi-subset $X \subset S''_{\mathcal{X}}$ and hypothesis h , we denote by $\widehat{M}_X(h)$ the number of instances in X misclassified by h . Let $\mathcal{F}_{S''}(\mathcal{H}) := \{f_{h,T} \mid h \in \mathcal{H}, T \subset S'', |T| = m\}$. Then we have

$$\begin{aligned}
& \Pr(B_{S,S'}) \\
&= \Pr(\exists h \in \mathcal{H}, \text{err}_{S'}(f_{h,S}) \geq \text{err}_S(h) + \frac{\epsilon}{2}) \\
&\leq \Pr(\exists h \in \mathcal{H}, (\widehat{M}_{S'_{\mathcal{X}} \cap R_1}(f_{h,S}) \geq \widehat{M}_{S_{\mathcal{X}} \cap R_1}(h) + \frac{\epsilon m}{4}) \vee (\widehat{M}_{S'_{\mathcal{X}} \cap R_2}(f_{h,S}) \geq \widehat{M}_{S_{\mathcal{X}} \cap R_2}(h) + \frac{\epsilon m}{4})) \\
&\leq \underbrace{\Pr(\exists h \in \mathcal{H}, \widehat{M}_{S'_{\mathcal{X}} \cap R_1}(f_{h,S}) \geq \widehat{M}_{S_{\mathcal{X}} \cap R_1}(h) + \frac{\epsilon m}{4})}_{(a)} \\
&\quad + \underbrace{\Pr(\exists h \in \mathcal{H}, \widehat{M}_{S'_{\mathcal{X}} \cap R_2}(f_{h,S}) \geq \widehat{M}_{S_{\mathcal{X}} \cap R_2}(h) + \frac{\epsilon m}{4})}_{(b)}.
\end{aligned}$$

For the first term, since $\text{err}_{\mathcal{G}S}(f_{h,S}) = 0$ according to the definition of $f_{h,S}$,

$$(a) \leq \Pr(\exists f \in \mathcal{F}_{S''}(\mathcal{H}), \widehat{M}_{S'_{\mathcal{X}} \cap R_1}(f) \geq \frac{\epsilon m}{4}, \text{err}_{\mathcal{G}S}(f) = 0).$$

Again, similar to case 1 in Theorem 1 (also the first upper bound in Theorem 7), for any $f \in \mathcal{F}_{S''}(\mathcal{H})$, we let $X(f) \subset R_1$ denote a minimal set of examples in R_1 (breaking ties arbitrarily but in a fixed way) such that f misclassify $X(f)$ and $|(\mathcal{G}X(f)) \cap S''_{\mathcal{X}}| \geq \frac{\epsilon m}{4}$ where $\mathcal{G}X(f) = \{gx \mid x \in X(f), g \in \mathcal{G}\}$ is the set of all examples lying in the orbits generated from $X(f)$. Let $K(f) = \mathcal{G}X(f)$ and $\mathcal{K} = \{K(f) \mid f \in \mathcal{F}_{S''}(\mathcal{H})\}$. Notice that each example in $X(f)$ must belong to different orbits, otherwise it is not minimal. Besides, each orbit contains at least $\log^2 m$ examples. Hence, $|X(f)| \leq \frac{\epsilon m}{4 \log^2 m}$. Since there are at most $\frac{2m}{\log^2 m}$ orbits generated from R_1 , we have $|\mathcal{K}| \leq \sum_{i=1}^{\frac{\epsilon m}{4 \log^2 m}} \binom{\frac{2m}{\log^2 m}}{i} \leq \left(\frac{8e}{\epsilon}\right)^{\frac{\epsilon m}{4 \log^2 m}}$. Since f misclassify $X(f)$, all examples in $K(f)$ must go to S' to guarantee $\text{err}_{\mathcal{G}S}(f) = 0$ and thus, we have

$$\begin{aligned}
& \Pr(\exists f \in \mathcal{F}_{S''}(\mathcal{H}), \widehat{M}_{S'_{\mathcal{X}} \cap R_1}(f) \geq \frac{\epsilon m}{4}, \text{err}_{\mathcal{G}S}(f) = 0) \\
&\leq \Pr(\exists f \in \mathcal{F}_{S''}(\mathcal{H}), K(f) \cap S_{\mathcal{X}} = \emptyset) = \Pr(\exists K \in \mathcal{K}, K \cap S_{\mathcal{X}} = \emptyset) \\
&\leq \sum_{K \in \mathcal{K}} 2^{-\frac{\epsilon m}{4}} \leq \left(\frac{8e}{\epsilon}\right)^{\frac{\epsilon m}{4 \log^2 m}} \cdot 2^{-\frac{\epsilon m}{4}} = 2^{-\frac{\epsilon m}{4} \left(1 - \frac{\log(8e/\epsilon)}{\log^2 m}\right)} \leq 2^{-\frac{\epsilon m}{8}},
\end{aligned}$$

when $m \geq \frac{8e}{\epsilon}$. For the second term, since $\widehat{M}_{S'_{\mathcal{X}} \cap R_2}(h) \geq \widehat{M}_{S'_{\mathcal{X}} \cap R_2}(f_{h,S})$, then we have

$$\begin{aligned}
(b) &\leq \Pr(\exists h \in \mathcal{H}, \widehat{M}_{S'_{\mathcal{X}} \cap R_2}(h) \geq \widehat{M}_{S_{\mathcal{X}} \cap R_2}(h) + \frac{\epsilon m}{4}) \\
&\leq \left(\frac{2em}{d}\right)^{d \log^2 m} \cdot e^{-2m' \left(\frac{\epsilon m}{8m'}\right)^2} \\
&\leq e^{-\frac{\epsilon^2 m}{32} + d \log^2 m \ln(2em/d)},
\end{aligned} \tag{14}$$

where Eq (14) adopts Hoeffding bound and Sauer's lemma (the number points in R_2 that can be shattered by \mathcal{H} is at most $d \log^2 m$; otherwise $\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G}) > d$) and $m' = \max_{h \in \mathcal{H}} M_{R_2}(h)$. By setting $m \geq \frac{32}{\epsilon^2} (d \log^2 m \ln(2em/d) + \ln(8/\delta) + 1) + \frac{8}{\epsilon} \log(8/\delta)$, we have $\Pr(A_S) \leq \delta/2$. By Hoeffding bound, with probability at least $1 - \delta/2$, $\text{err}_S(h^*) \leq \text{err}_{\mathcal{D}}(h^*) + \frac{\epsilon}{2}$. By a union bound, we have that with probability at least $1 - \delta$, $\text{err}_{\mathcal{D}}(f_{h,S}) \leq \text{err}_S(h) + \epsilon \leq \text{err}_S(h^*) + \epsilon \leq \text{err}_{\mathcal{D}}(h^*) + 3\epsilon/2$. \square

Proof of the upper bound of the 1-inclusion-graph predictor. Here we provide another algorithm based on the technique of reduction-to-realizable of David et al. (2016) and the 1-inclusion-graph predictor in the relaxed realizable setting. Following the argument in Theorem 7, given a sample $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$, let X_S be the set of different elements in $S_{\mathcal{X}}$ and define

$\mathcal{H}'(X_S) := \{h|_{X_S} | \forall x', x \in X_S, x' \in \mathcal{G}x \text{ implies } h(x') = h(x)\}$. Then we define an algorithm \mathcal{A}_0 as

$$\mathcal{A}_0(S, x) = \begin{cases} Q_{\mathcal{H}'(X_S \cup \{x\}), X_S \cup \{x\}}(S, x) & \text{if } \mathcal{H}'(X_S \cup \{x\}) \neq \emptyset, \\ 0 & \text{o.w.} \end{cases} \quad (15)$$

where $Q_{\mathcal{H}'(X_S \cup \{x\}), X_S \cup \{x\}}$ is the function guaranteed by Eq (1) for hypothesis class $\mathcal{H}'(X_S \cup \{x\})$ and instance space $X_S \cup \{x\}$. In the relaxed realizable setting, if $S \sim \mathcal{D}^n$ and $x \sim \mathcal{D}_{\mathcal{X}}$, $\mathcal{H}'(X_S \cup \{x\}) \neq \emptyset$ a.s. as it contains $h^*|_{X_S \cup \{x\}}$. While in the agnostic setting, it is not the case. Even if $S \sim \mathcal{D}^n$ and $x \sim \mathcal{D}_{\mathcal{X}}$, $\mathcal{H}'(X_S \cup \{x\})$ could be empty as there might exist an instance $x' \in X_S \cap \mathcal{G}x$ such that no hypothesis in \mathcal{H} labeling them in the same way.

Let $d = \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$. If a sample $\{(x_1, y_1), \dots, (x_{n+1}, y_{n+1})\}$ is realizable by \mathcal{H} , then by Lemma 1, we have

$$\frac{1}{(n+1)!} \sum_{\sigma \in \text{Sym}(n+1)} \mathbb{1}[\mathcal{A}_0(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [n]}, x_{\sigma(n+1)}) \neq y_{\sigma(n+1)}] \leq \frac{d}{n+1}. \quad (16)$$

Now we use \mathcal{A}_0 as a weak learner to construct a compression scheme by following the construction David et al. (2016). Given a training set $S_{\text{trn}} = \{(x_1, y_1), \dots, (x_m, y_m)\}$, let R denote the largest submultiset of S_{trn} that is realizable w.r.t. \mathcal{H} . If $|R| = 0$, then define \hat{h} as the all-0 function $\hat{h}(x) = 0$. Otherwise, if $|R| > 0$, for any distribution \mathcal{P} on R , by Eq (16), we have that

$$\begin{aligned} & \mathbb{E}_{S \sim \mathcal{P}^{3d}} [\text{err}_{\mathcal{P}}(\mathcal{A}_0(S))] \\ &= \mathbb{E}_{(x_i, y_i)_{i \in [3d+1]} \sim \mathcal{P}^{3d}} \left[\frac{1}{(3d+1)!} \sum_{\sigma \in \text{Sym}(3d+1)} \mathbb{1}[\mathcal{A}_0(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [3d]}, x_{\sigma(3d+1)}) \neq y_{\sigma(3d+1)}] \right] \\ &\leq \frac{1}{3}. \end{aligned}$$

Hence, there exists $S_{\mathcal{P}} \in R^{3d}$ with $\text{err}_{\mathcal{P}}(\mathcal{A}_0(S_{\mathcal{P}})) \leq 1/3$. Thus, the algorithm \mathcal{A}_0 can serve as a weak learner. By Lemma 2, for $T = \lceil c \log |R| \rceil$ (for a numerical constant $c > 0$), there exist $S_1, \dots, S_T \in R^{3d}$ such that, letting $\hat{h}(\cdot) = \text{Majority}(\mathcal{A}_0(S_1)(\cdot), \dots, \mathcal{A}_0(S_T)(\cdot))$, we have $\text{err}_R(\hat{h}) = 0$. Thus, $\text{err}_{S_{\text{trn}}}(\hat{h}) \leq \inf_{h \in \mathcal{H}} \text{err}_{S_{\text{trn}}}(h)$. Here \hat{h} is the output of the compression scheme that selects $\kappa(S_{\text{trn}}) = (S_1, \dots, S_T)$ and $\rho(\kappa(S_{\text{trn}})) = \hat{h}$. By Lemma 4, with probability at least $1 - \delta/2$,

$$\text{err}_{\mathcal{D}}(\hat{h}) \leq \text{err}_{S_{\text{trn}}}(\hat{h}) + c' \sqrt{\text{err}_{S_{\text{trn}}}(\hat{h}) B(S_{\text{trn}}, \delta/2)} + c' B(S_{\text{trn}}, \delta/2),$$

for a numerical constant $c' > 0$. Following the same argument of David et al. (2016) (Lemma 3.2), we have that

$$\Pr_{S_{\text{trn}} \sim \mathcal{D}^m} (\text{err}_S(\hat{h}) \geq \inf_{h \in \mathcal{H}} \text{err}_{\mathcal{D}}(h) + \sqrt{\log(2/\delta)/m}) \leq \delta/2.$$

By taking a union bound, then we have that with probability at least $1 - \delta$,

$$\begin{aligned} \text{err}_{\mathcal{D}}(\hat{h}) &\leq \inf_{h \in \mathcal{H}} \text{err}_{\mathcal{D}}(h) + c' \sqrt{\left(\inf_{h \in \mathcal{H}} \text{err}_{\mathcal{D}}(h) + \sqrt{\log(2/\delta)/m} \right) \frac{1}{m} (3d \lceil c \log |R| \rceil \log m + \log 1/\delta)} \\ &\quad + \frac{c'}{m} (3d \lceil c \log |R| \rceil \log m + \log(1/\delta)). \end{aligned}$$

Hence,

$$\mathcal{M}_{\text{AG}}(\epsilon, \delta; \mathcal{H}, \mathcal{G}) = O\left(\frac{d}{\epsilon^2} \log^2\left(\frac{d}{\epsilon}\right) + \frac{1}{\epsilon^2} \log(1/\delta)\right).$$

□

I Adaptive algorithms

For any hypothesis $h \in \mathcal{H}$, we say h is $(1 - \eta)$ -invariant over the distribution $\mathcal{D}_{\mathcal{X}}$ for some $\eta \in [0, 1]$ if $\mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}}(\exists x' \in \mathcal{G}x, h(x') \neq h(x)) = \eta$. We call $\eta(h) = \eta$ the invariance parameter of h with respect to $\mathcal{D}_{\mathcal{X}}$. In the relaxed realizable setting, the problem degenerates into the invariantly realizable setting when $\eta(h^*) = 0$. This implies that we can benefit from transformation invariances more when $\eta(h^*)$ is smaller. The case is similar in the agnostic setting. In this section, we discuss adaptive learning algorithms for different levels of invariance of the target function.

I.1 An adaptive algorithm in the relaxed realizable setting

There might exist more than one hypotheses in \mathcal{H} with zero error and we let the target function h^* be the one with the smallest invariance parameter (breaking ties arbitrarily). For any multiset X in \mathcal{X} , for any hypothesis h , denote $h|_X$ the restriction of h on the set of different elements in X . Then we introduce a distribution-dependent dimension as follows.

Definition 8 (approximate $(1 - \eta)$ -invariant VC dimension). *For any $\eta \in [0, 1]$ and finite multiset $X \subset \mathcal{X}$, let $\mathcal{H}^\eta(X) := \{h|_X \mid \frac{1}{|X|} \sum_{x \in X} \mathbb{1}[\exists x' \in \mathcal{G}x, h(x') \neq h(x)] \leq \eta \wedge \forall x, x' \in \mathcal{G}x \cap X, h(x') = h(x)\}$. For any $m \in \mathbb{N}$, marginal distribution $\mathcal{D}_{\mathcal{X}}$ and target function h^* , the approximate $(1 - \eta)$ -invariant VC dimension is defined as*

$$\text{VC}_o^\eta(m, h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}) := \mathbb{E}_{X \sim \mathcal{D}_{\mathcal{X}}^{m+1}} \left[\text{VCdim}(\mathcal{H}^\eta(X)) \mid h^*|_X \in \mathcal{H}^\eta(X) \right],$$

when $\Pr(h^*|_X \in \mathcal{H}^\eta(X)) > 0$ and $\text{VC}_o^\eta(m, h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}) = 0$ when $\Pr(h^*|_X \in \mathcal{H}^\eta(X)) = 0$. By taking supremum over m ,

$$\text{VC}_o^\eta(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}) = \sup_{m \in \mathbb{N}} \text{VC}_o^\eta(m, h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}).$$

In this definition, all hypotheses in $\mathcal{H}^\eta(X)$ need to satisfy two constraints: a) the empirical invariance parameter is less than or equal to η , and b) the prediction over instances in X is invariant over orbits. Here the second constraint arises due to the fact that h^* satisfies this constraint. Note that $\text{VC}_o^\eta(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}})$ is monotonic increasing in η . For all $\mathcal{D}_{\mathcal{X}}, h^*$, we have $\text{VC}_o^0(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}) \leq \text{VC}_o(\mathcal{H}, \mathcal{G})$ and $\text{VC}_o^\eta(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}) \leq \text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})$ for all $\eta \in [0, 1]$. We will use the approximate $(1 - \eta)$ -invariant VC dimension to characterize the sample complexity dependent on η . Ideally, it is heuristic to adopt a notion like $\text{VC}_{\text{ao}}(\{h \in \mathcal{H} \mid \eta(h) = \eta(h^*)\})$. But this is impossible to achieve as we cannot obtain an accurate estimate of $\eta(h^*)$ via finite data points.

Proposition 1. *If it is known that $\eta(h^*) \leq \eta$ for some $\eta \in [0, 1]$, for any training sample size $m \in \mathbb{N}$, for any $\Delta \geq \sqrt{\frac{\ln(n+1)}{2(n+1)}}$ with $n = \Theta(\frac{m}{\log(1/\delta)})$, there is an algorithm achieving error $O(\frac{\text{VC}_o^{\eta+\Delta}(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}) \log(1/\delta)}{m})$ with probability at least $1 - \delta$.*

Given a sample $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ and a test instance x , let X_S denote the set of different elements in $S_{\mathcal{X}}$ and $X = \{x_1, \dots, x_n, x\}$ the multiset of all unlabeled instances from both the training set and the test instance. The algorithm \mathcal{A} is defined by

$$\mathcal{A}(S, x) = Q_{\mathcal{H}^{\eta+\Delta}(X), X_S \cup \{x\}}(S, x),$$

where $Q_{\mathcal{H}^{\eta+\Delta}(X), X_S \cup \{x\}}$ is the function guaranteed by Eq (1) for the instance space $X_S \cup \{x\}$ and the concept class $\mathcal{H}^{\eta+\Delta}(X)$. Similar to Theorem 4, algorithm \mathcal{A} can achieve expected error $\frac{\text{VC}_o^{\eta+\Delta}(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}) + 1}{n+1}$. Then by following the same confidence boosting argument, we run \mathcal{A} for $\lceil \log(2/\delta) \rceil$ times on independent new samples and select the output hypothesis with minimum error a new sample. The details of the algorithm and the proof of Proposition 1 are deferred to Appendix J.

In the more general case where η is unknown, we build an algorithm based on the algorithm for known η above. Denote $\mathcal{A}_{\eta, \Delta}$ the algorithm satisfying the guarantee in Proposition 1 with probability $1 - \delta/2$ for hyperparameters η and Δ . Then we divide $[0, 1]$ into uniform intervals and then search for the interval where $\eta(h^*)$ lies in. The detailed algorithm is provided in Algorithm 1.

Theorem 9. *Set $\Delta = \sqrt{\frac{\ln(n+1)}{2(n+1)}}$ for $n = \Theta(\frac{m}{\log(m) \log(1/\delta)})$ and $|S_1| = \Theta(\frac{m}{\log m})$. Let $i^* \geq 0$ be the smallest integer i such that $\max((2i - 1)\Delta, 0) \geq \eta(h^*)$. Then Algorithm 1 achieves error $O(\frac{\text{VC}_o^{2i^* \Delta}(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_{\mathcal{X}}) \log(1/\delta) \log(m)}{m})$ with probability at least $1 - \delta$.*

Algorithm 1 An adaptive algorithm in the relaxed realizable setting

- 1: Input: a labeled sample S_{trn} of size m , $m_1 \in \mathbb{N}$, $\Delta \in [0, 1]$
 - 2: Randomly partition S into S_1 with $|S_1| = m_1$ and S_2 with $|S_2| = m - m_1$
 - 3: **for** $i = 0, 1, \dots, \lceil 1/(2\Delta) \rceil$ **do**
 - 4: Let $h_i = \mathcal{A}_{(2i-1)\Delta, \Delta}(S_1)$
 - 5: Return $\hat{h} = \arg \min_{h \in \{h_i | i \in \{0, \dots, \lceil 1/(2\Delta) \rceil\}\}} \text{err}_{S_2}(h)$
 - 6: **end for**
-

Algorithm 2 An adaptive algorithm in the agnostic case

- 1: Input: a unlabeled data set U of size u and a labeled data set S of size m , $m_1 \in \mathbb{N}$, $\Delta > 0$
 - 2: Randomly divide S into S_1 of size m_1 and S_2 of size $m - m_1$
 - 3: Use U to partition \mathcal{H} into $\hat{\mathcal{H}}_1, \dots, \hat{\mathcal{H}}_K$ with $\hat{\mathcal{H}}_i = \{h \in \mathcal{H} | \frac{1}{u} \sum_{x \in U} \mathbb{1}[\exists x' \in \mathcal{G}x, h(x') \neq h(x)] \in (2(i-1)\Delta, 2i\Delta]\}$ and $K = \lceil \frac{1}{2\Delta} \rceil$
 - 4: **for** $i = 0, 1, 2, \dots, K$ **do**
 - 5: Run the algorithm \mathcal{A} in Theorem 8 over S_1 for hypothesis class $\hat{\mathcal{H}}_i$ and output h_i
 - 6: Return $\hat{h} = \arg \min_{h \in \{h_i | i \in \{0, \dots, K\}\}} \text{err}_{S_2}(h)$
 - 7: **end for**
-

The proof is deferred to Appendix K. The algorithm above perform close to optimally in both the invariant realizable setting and the relaxed realizable setting. Intuitively, the algorithm above outperforms PAC-optimal algorithms in Theorem 7 when $\eta(h^*)$ is small. Below is an example showing the advantage of this adaptive algorithm in the extreme case of $\eta(h^*) = 0$.

Example 6. Consider the construction of $\mathcal{H}_d, \mathcal{G}_d, h^*, \mathcal{D}$ in the proof of Theorem 2. In this example, we have $i^* = 0$ and $\mathcal{H}^0(X) = \{\mathbf{1}\}$ for any X . Thus, $\text{VC}_o^0(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_X) = 0$. Hence, the algorithm above only requires sample complexity $O(1)$ to output a zero-error predictor. However, algorithms only caring the worst case upper bound may require much more samples. For example, *ERM-INV* predicts exactly the same as standard *ERM* in this example and thus, it requires $\Omega(\frac{\text{VC}_{\text{ao}}(\mathcal{H}, \mathcal{G})}{\epsilon})$ samples to achieve ϵ error.

Open question: It is unclear whether this adaptive algorithm is optimal and whether the approximate $(1-\eta)$ -invariant VC dimension is the best η -dependent measure to characterize the sample complexity.

I.2 An adaptive algorithm in the agnostic setting

In the agnostic setting, we assume that there exists an optimal hypothesis $h^* \in \mathcal{H}$ such that $\text{err}(h^*) = \inf_{h \in \mathcal{H}} \text{err}(h)$. Then similar to the realizable setting, we can design algorithms that adapt to $\eta(h^*)$. However, it is more challenging to design an adaptive algorithm in the agnostic setting than in the relaxed realizable setting. One of the most direct ideas is to combine agnostic compression scheme with the adaptive algorithm in the realizable setting. One possible way of combination is finding the largest realizable subset of the data and applying the adaptive algorithm in the relaxed realizable setting. However, this does not work since the realizable subset is not i.i.d. and the empirical invariance parameter calculated based on this subset is biased. Another possible way of combination is calculating the empirical invariance parameter over the whole data set, reducing the hypothesis class based on this empirical value and then run the compression scheme in Theorem 8 based on this reduced hypothesis class. This does not work either because the predictor depends on the whole data set now and the compression size is too large. Hence, there is a significant barrier that has arisen as a result of estimating the invariance parameter while obtaining low error at the same time.

To get around this obstacle, we provide an approach of using two independent data sets. Specifically, we partition the hypothesis class into subclasses with different empirical invariance parameters based on a data set first. Notice that in this step, we only need an unlabeled data set. Then we run the compression scheme in Theorem 8 for each subclass and return the one with the small validation error. The detailed algorithm is presented in Algorithm 2.

Theorem 10. For each $h \in \mathcal{H}$, the invariance indicator function of h is a mapping $\iota_h : \mathcal{X} \mapsto \{0, 1\}$ such that $\iota_h(x) = \mathbb{1}[\exists x' \in \mathcal{G}x, h(x') \neq h(x)]$. Denote $\mathcal{I} = \{\iota_h | h \in \mathcal{H}\}$ the set of invariance indicator functions for all $h \in \mathcal{H}$. Then for any $\Delta \in (0, 1]$, Algorithm 2 can achieve

$$\text{err}(\hat{h}) \leq \text{err}(h^*) + O\left(\sqrt{\frac{\text{VC}_{\text{ao}}(\mathcal{H}^*) \log^2 m + \log(1/\delta) + \log(1/\Delta)}{m}}\right), \text{ where } \mathcal{H}^* = \{h \mid \eta(h) \in (\eta(h^*) - 2\Delta - 2\Delta', \eta(h^*) + 2\Delta + 2\Delta')\} \text{ with } \Delta' = \Theta\left(\sqrt{\frac{\text{VCdim}(\mathcal{I}) \log(u) + \log(1/\delta)}{u}}\right).$$

The detailed proof of Theorem 10 is deferred to Appendix L. The upper bound in Theorem 10 depends on $\text{VCdim}(\mathcal{I})$, which can be arbitrarily larger than $\text{VCdim}(\mathcal{H})$. For example, for any $d > 0$, let $\mathcal{X} = \{0, 1\}^d \times [d]$. For each $\mathbf{b} \in \{0, 1\}^d$, define a hypothesis $h_{\mathbf{b}}$ by letting $h_{\mathbf{b}}((\mathbf{b}, i)) = b_i$ and $h_{\mathbf{b}}(x) = 0$ for all other $x \in \mathcal{X}$. Let the hypothesis class $\mathcal{H} = \{h_{\mathbf{b}} \mid \mathbf{b} \in \{0, 1\}^d\}$. Then it is direct to check that $\text{VCdim}(\mathcal{H}) = 1$ but $\text{VCdim}(\mathcal{I}) = d$. It is unclear how to design an adaptive algorithm with theoretical guarantee independent of $\text{VCdim}(\mathcal{I})$. The η -dependent dimension we adopt in the agnostic setting is different from that in the relaxed realizable setting. It is also unclear what is the best way to characterize the dependence on η in the agnostic setting.

J Proof of Proposition 1

Proof. The proof follows that of Theorem 4. Given a training sample of size m , let n be the largest integer such that $m \geq n \lceil \log(2/\delta) \rceil + \lceil 8(n+1)(\ln(2/\delta) + \ln(\lceil \log(2/\delta) \rceil + 1)) \rceil$. For any sample $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ and a test instance $x_{n+1} = x$, if h^* is in $\mathcal{H}^{\eta+\Delta}(X)$, by Lemma 1, we have that

$$\frac{1}{(n+1)!} \sum_{\sigma \in \text{Sym}(n+1)} \mathbb{1}[\mathcal{A}(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [n]}, x_{\sigma(n+1)}) \neq y_{\sigma(n+1)}] \leq \frac{\text{VCdim}(\mathcal{H}^{\eta+\Delta}(X))}{n+1}.$$

Thus,

$$\begin{aligned} & \mathbb{E}_{S \sim \mathcal{D}^n} [\text{err}(\mathcal{A}(S))] \\ &= \mathbb{E}_{(x_i, y_i)_{i \in [n+1]} \sim \mathcal{D}^{n+1}} [\mathbb{1}[\mathcal{A}(\{x_i, y_i\}_{i \in [n]}, x_{n+1}) \neq y_{n+1}]] \\ &= \frac{1}{(n+1)!} \sum_{\sigma \in \text{Sym}(n+1)} \mathbb{E} [\mathbb{1}[\mathcal{A}(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [n]}, x_{\sigma(n+1)}) \neq y_{\sigma(n+1)}]] \\ &= \mathbb{E} \left[\frac{1}{(n+1)!} \sum_{\sigma \in \text{Sym}(n+1)} \mathbb{1}[\mathcal{A}(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [n]}, x_{\sigma(n+1)}) \neq y_{\sigma(n+1)}] \right] \\ &\leq \mathbb{E} \left[\frac{1}{(n+1)!} \sum_{\sigma \in \text{Sym}(n+1)} \mathbb{1}[\mathcal{A}(\{x_{\sigma(i)}, y_{\sigma(i)}\}_{i \in [n]}, x_{\sigma(n+1)}) \neq y_{\sigma(n+1)}] \mid h_{|X}^* \in \mathcal{H}^{\eta+\Delta}(X) \right] \\ &\quad \cdot \Pr(h_{|X}^* \in \mathcal{H}^{\eta+\Delta}(X)) + \Pr(h_{|X}^* \notin \mathcal{H}^{\eta+\Delta}(X)) \\ &\leq \frac{\mathbb{E}_{X \sim \mathcal{D}_X^{n+1}} [\text{VCdim}(\mathcal{H}^{\eta+\Delta}(X)) \mid h_{|X}^* \in \mathcal{H}^{\eta+\Delta}(X)]}{n+1} \Pr(h_{|X}^* \in \mathcal{H}^{\eta+\Delta}(X)) \\ &\quad + \Pr(h_{|X}^* \notin \mathcal{H}^{\eta+\Delta}(X)) \\ &\leq \frac{\text{VC}_o^{\eta+\Delta}(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_X)}{n+1} \Pr(h_{|X}^* \in \mathcal{H}^{\eta+\Delta}(X)) + \Pr(h_{|X}^* \notin \mathcal{H}^{\eta+\Delta}(X)) \\ &\leq \frac{\text{VC}_o^{\eta+\Delta}(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_X) + 1}{n+1}, \end{aligned} \tag{17}$$

where Eq (17) holds due to $\Pr(h_{|X}^* \notin \mathcal{H}^{\eta+\Delta}(X)) \leq \Pr(\frac{1}{|X|} \sum_{x \in X} \mathbb{1}[\exists x' \in \mathcal{G}, h^*(x') \neq h^*(x)] - \eta(h^*) \geq \Delta) \leq \exp(-2(n+1)\Delta^2)$ by Hoeffding bound.

Then we again follow the classic technique to boost the confidence. The algorithm runs \mathcal{A} for $\lceil \log(2/\delta) \rceil$ times, each time using a new sample S_i of size n for $i = 1, \dots, \lceil \log(2/\delta) \rceil$. Let $h_i = \mathcal{A}(S_i)$ and then selects the hypothesis \hat{h} from $\{h_i \mid i \in [\lceil \log(2/\delta) \rceil]\}$ with the minimal error on a new sample S_0 of size $t = \lceil 8(n+1)(\ln(2/\delta) + \ln(\lceil \log(2/\delta) \rceil + 1)) \rceil$.

Denote $\epsilon = \frac{\text{VC}_o^{\eta+\Delta}(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_X) + 1}{n+1}$. For each i , by Eq (17), we have $\mathbb{E}[\text{err}(h_i)] \leq \epsilon$. By Markov's inequality, with probability at least $\frac{1}{2}$, $\text{err}(h_i) \leq 2\epsilon$. Since h_i are independent, we have that with

probability at least $1 - \frac{\delta}{2}$, at least one of $\{h_i | i \in [\lceil \log(2/\delta) \rceil]\}$ has error smaller than 2ϵ . Then by Chernoff bound, for each i , on the event $\text{err}(h_i) \leq 2\epsilon$,

$$\Pr(\text{err}_{S_0}(h_i) > 3\epsilon | h_i) < e^{-\frac{t\epsilon}{6}}.$$

Also, on the event $\text{err}(h_i) > 4\epsilon$,

$$\Pr(\text{err}_{S_0}(h_i) \leq 3\epsilon | h_i) \leq e^{-\frac{t\epsilon}{8}}.$$

Thus, by the law of total probability and a union bound, with probability at least $1 - \frac{\delta}{2}$, if any i has $\text{err}(h_i) \leq 2\epsilon$, then the returned hypothesis \hat{h} has $\text{err}(\hat{h}) \leq 4\epsilon$. By a union bound, the proof is completed. \square

K Proof of Theorem 9

Proof. Given a training sample of size m , let $|S_1| = m_1$ and $|S_2| = m_2 = m - m_1$. The values of m_1, m_2 are determined later. Let n be the largest integer such that $m_1 \geq n \lceil \log(2/\delta) \rceil + \lceil 8(n+1)(\ln(2/\delta) + \ln(\lceil \log(2/\delta) \rceil + 1)) \rceil$. Let $\Delta = \sqrt{\frac{\ln(n+1)}{2(n+1)}}$ and then the number of rounds we run $\mathcal{A}_{\eta, \Delta}$ as a subroutine is upper bounded by $\lceil 1/(2\Delta) \rceil + 1 = \lceil \sqrt{(n+1)/(2 \ln(n+1))} \rceil + 1 \leq m$. According to Proposition 1, there is a numerical constant $c > 0$ such that with probability $1 - \delta/2$, $\text{err}(h_{i^*}) \leq \frac{c \text{VC}_o^{2i^* \Delta}(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_X) \ln(2/\delta)}{m_1}$. Denote $\epsilon = \frac{c \text{VC}_o^{2i^* \Delta}(h^*, \mathcal{H}, \mathcal{G}, \mathcal{D}_X) \ln(2/\delta)}{m_1}$. By Chernoff bound, for each $i \neq i^*$, on the event $\text{err}(h_i) > 4\epsilon$,

$$\Pr(\text{err}_{S_2}(h_i) \leq 2\epsilon | h_i) \leq e^{-\frac{m_2 \epsilon}{2}}.$$

Also, on the event $\text{err}(h_{i^*}) \leq \epsilon$,

$$\Pr(\text{err}_{S_2}(h_{i^*}) > 2\epsilon | h_{i^*}) < e^{-\frac{m_2 \epsilon}{3}}.$$

Let $m_1 = \frac{cm}{3 \ln m + c}$ and $m_2 = \frac{3m \ln m}{3 \ln m + c}$. Then with probability at least $1 - me^{-m_2 \epsilon/3} \geq 1 - \delta/2$, if $\text{err}(h_{i^*}) \leq \epsilon$, then the returned classifier has error smaller than 4ϵ . By taking a union bound, the proof is completed. \square

L Proof of Theorem 10

Proof. Given a labeled data set S of size m , let $|S_1| = m_1$ and $|S_2| = m_2 = m - m_1$. The values of m_1, m_2 are determined later. Let \hat{i} be the $i \in [K]$ such that $h^* \in \hat{\mathcal{H}}_i$. By Theorem 8, we have that there exists a numerical constant $c > 0$ such that with probability at least $1 - \delta/3$,

$$\text{err}(h_{\hat{i}}) \leq \text{err}(h^*) + c \sqrt{\frac{\text{VC}_{\text{ao}}(\hat{\mathcal{H}}_{\hat{i}}) \log^2 m_1 + \log(1/\delta)}{m_1}}.$$

Then by Hoeffding bound, for any $\epsilon > 0$, for each $i \in K$, with probability at least $1 - \delta/(3K)$,

$$|\text{err}_{S_2}(h_i) - \text{err}(h_i)| \leq \epsilon,$$

when $m_2 \geq \frac{\ln(6K/\delta)}{2\epsilon^2}$. Then by taking a union bound, with probability at least $1 - \delta/3$, for all $i \in K$ we have

$$|\text{err}_{S_2}(h_i) - \text{err}(h_i)| \leq \epsilon.$$

Hence, with probability at least $1 - 2\delta/3$,

$$\text{err}(\hat{h}) \leq \text{err}(h_{\hat{i}}) + 2\epsilon \leq \text{err}(h^*) + c \sqrt{\frac{\text{VC}_{\text{ao}}(\hat{\mathcal{H}}_{\hat{i}}) \log^2 m_1 + \log(1/\delta)}{m_1}} + 2\epsilon.$$

By uniform convergence bound, with probability at least $1 - \delta/3$,

$$\left| \frac{1}{u} \sum_{x \in U} \mathbb{1}[\exists x' \in \mathcal{G}x, h(x') \neq h(x)] - \eta(h) \right| \leq \Delta', \forall h \in \mathcal{H}.$$

It follows that $\hat{\mathcal{H}}_{\hat{i}} \subset \mathcal{H}^*$ and thus, $\text{VC}_{\text{ao}}(\hat{\mathcal{H}}_{\hat{i}}) \leq \text{VC}_{\text{ao}}(\mathcal{H}^*)$. The proof is completed by letting $\epsilon = \sqrt{\frac{\text{VC}_{\text{ao}}(\mathcal{H}^*) + \log(1/\delta)}{m_1}}$ and $m_1 = \frac{2(\text{VC}_{\text{ao}}(\mathcal{H}^*) + \log(1/\delta))m}{2(\text{VC}_{\text{ao}}(\mathcal{H}^*) + \log(1/\delta)) + \ln(6K/\delta)}$. \square