

Nearly Tight Bounds for Robust Proper Learning of Halfspaces with a Margin

Ilias Diakonikolas*
University of Wisconsin, Madison
ilias@cs.wisc.edu

Daniel M. Kane†
University of California, San Diego
dakane@cs.ucsd.edu

Pasin Manurangsi‡
University of California, Berkeley
pasin@berkeley.edu

August 30, 2019

Abstract

We study the problem of *properly* learning large margin halfspaces in the agnostic PAC model. In more detail, we study the complexity of properly learning d -dimensional halfspaces on the unit ball within misclassification error $\alpha \cdot \text{OPT}_\gamma + \epsilon$, where OPT_γ is the optimal γ -margin error rate and $\alpha \geq 1$ is the approximation ratio. We give learning algorithms and computational hardness results for this problem, for all values of the approximation ratio $\alpha \geq 1$, that are nearly-matching for a range of parameters. Specifically, for the natural setting that α is any constant bigger than one, we provide an essentially tight complexity characterization. On the positive side, we give an $\alpha = 1.01$ -approximate proper learner that uses $O(1/(\epsilon^2 \gamma^2))$ samples (which is optimal) and runs in time $\text{poly}(d/\epsilon) \cdot 2^{\tilde{O}(1/\gamma^2)}$. On the negative side, we show that *any* constant factor approximate proper learner has runtime $\text{poly}(d/\epsilon) \cdot 2^{(1/\gamma)^{2-o(1)}}$, assuming the Exponential Time Hypothesis.

*Supported by NSF Award CCF-1652862 (CAREER) and a Sloan Research Fellowship. Part of this work was performed at the Simons Institute for the Theory of Computing during the program on Foundations of Data Science.

†Supported by NSF Award CCF-1553288 (CAREER) and a Sloan Research Fellowship.

‡Now at Google Research, Mountain View.

1 Introduction

1.1 Background and Problem Definition

Halfspaces are Boolean functions $h_{\mathbf{w}} : \mathbb{R}^d \rightarrow \{\pm 1\}$ of the form $h_{\mathbf{w}}(\mathbf{x}) = \text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle)$, where $\mathbf{w} \in \mathbb{R}^d$ is the associated weight vector. (The function $\text{sign} : \mathbb{R} \rightarrow \{\pm 1\}$ is defined as $\text{sign}(u) = 1$ if $u \geq 0$ and $\text{sign}(u) = -1$ otherwise.) The problem of learning an unknown halfspace with a margin condition (in the sense that no example is allowed to lie too close to the separating hyperplane) is as old as the field of machine learning — starting with Rosenblatt’s Perceptron algorithm [Ros58] — and has arguably been one of the most influential problems in the development of the field, with techniques such as SVMs [Vap98] and AdaBoost [FS97] coming out of its study.

In this paper, we study the problem of learning γ -margin halfspaces in the *agnostic* PAC model [Hau92, KSS94]. Specifically, there is an unknown distribution \mathcal{D} on $\mathbb{B}_d \times \{\pm 1\}$, where \mathbb{B}_d is the unit ball on \mathbb{R}^d , and the learning algorithm \mathcal{A} is given as input a training set $S = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^m$ of i.i.d. samples drawn from \mathcal{D} . The goal of \mathcal{A} is to output a hypothesis whose error rate is competitive with the γ -margin error rate of the optimal halfspace. In more detail, the *error rate* (misclassification error) of a hypothesis $h : \mathbb{R}^d \rightarrow \{\pm 1\}$ (with respect to \mathcal{D}) is $\text{err}_{0-1}^{\mathcal{D}}(h) \stackrel{\text{def}}{=} \Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[h(\mathbf{x}) \neq y]$. For $\gamma \in (0, 1)$, the γ -margin error rate of a halfspace $h_{\mathbf{w}}(\mathbf{x})$ with $\|\mathbf{w}\|_2 \leq 1$ is $\text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w}) \stackrel{\text{def}}{=} \Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[y \langle \mathbf{w}, \mathbf{x} \rangle \leq \gamma]$. We denote by $\text{OPT}_{\gamma}^{\mathcal{D}} \stackrel{\text{def}}{=} \min_{\|\mathbf{w}\|_2 \leq 1} \text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w})$ the minimum γ -margin error rate achievable by any halfspace. We say that \mathcal{A} is an α -agnostic learner, $\alpha \geq 1$, if it outputs a hypothesis h that with probability at least $1 - \tau$ satisfies $\text{err}_{0-1}^{\mathcal{D}}(h) \leq \alpha \cdot \text{OPT}_{\gamma}^{\mathcal{D}} + \epsilon$. (For $\alpha = 1$, we obtain the standard notion of agnostic learning.) If the hypothesis h is itself a halfspace, we say that the learning algorithm is *proper*. This work focuses on proper learning algorithms.

1.2 Related and Prior Work

In this section, we summarize the prior work that is directly related to the results of this paper. First, we note that the sample complexity of our learning problem (ignoring computational considerations) is well-understood. In particular, the ERM that minimizes the number of γ -margin errors over the training set (subject to a norm constraint) is known to be an agnostic learner ($\alpha = 1$), assuming the sample size is $\Omega(\log(1/\tau)/(\epsilon^2 \gamma^2))$. Specifically, $\Theta(\log(1/\tau)/(\epsilon^2 \gamma^2))$ samples¹ are known to be sufficient and necessary for this learning problem (see, e.g., [BM02, McA03]). In the realizable case ($\text{OPT}_{\gamma}^{\mathcal{D}} = 0$), i.e., if the data is linearly separable with margin γ , the ERM rule above can be implemented in $\text{poly}(d, 1/\epsilon, 1/\gamma)$ time using the Perceptron algorithm. The non-realizable setting ($\text{OPT}_{\gamma}^{\mathcal{D}} > 0$) is much more challenging computationally.

The agnostic version of our problem ($\alpha = 1$) was first considered in [BS00], who gave a *proper* learning algorithm with runtime $\text{poly}(d) \cdot (1/\epsilon)^{\tilde{O}(1/\gamma^2)}$. It was also shown in [BS00] that agnostic proper learning with runtime $\text{poly}(d, 1/\epsilon, 1/\gamma)$ is NP-hard. A question left open by their work was characterizing the computational complexity of proper learning as a function of $1/\gamma$.

Subsequent works focused on *improper* learning. The $\alpha = 1$ case was studied in [SSS09, SSS10] who gave a learning algorithm with sample complexity $\text{poly}(1/\epsilon) \cdot 2^{\tilde{O}(1/\gamma)}$ — i.e., *exponential* in $1/\gamma$ — and computational complexity $\text{poly}(d/\epsilon) \cdot 2^{\tilde{O}(1/\gamma)}$. The increased sample complexity is inherent in their approach, as their algorithm works by solving a convex program over an expanded feature space. [BS12] gave an α -agnostic learning algorithm for all $\alpha \geq 1$ with sample complexity

¹To avoid clutter in the expressions, we will henceforth assume that the failure probability $\tau = 1/10$. Recall that one can always boost the confidence probability with an $O(\log(1/\tau))$ multiplicative overhead in the sample complexity.

$\text{poly}(1/\epsilon) \cdot 2^{\tilde{O}(1/(\alpha\gamma))}$ and computational complexity $\text{poly}(d/\epsilon) \cdot 2^{\tilde{O}(1/(\alpha\gamma))}$. (We note that the Perceptron algorithm is known to achieve $\alpha = 1/\gamma$ [Ser01]. Prior to [BS12], [LS11] gave a $\text{poly}(d, 1/\epsilon, 1/\gamma)$ time algorithm achieving $\alpha = \Theta((1/\gamma)/\sqrt{\log(1/\gamma)})$.) [BS12] posed as an open question whether their upper bounds for improper learning can be achieved with a proper learner.

A related line of work [KLS09, ABL17, DKK⁺16, LRV16, DKK⁺17, DKK⁺18, DKS18, KKM18, DKS19, DKK⁺19] has given polynomial time robust estimators for a range of learning tasks. Specifically, [KLS09, ABL17, DKS18, DKK⁺19] obtained efficient PAC learning algorithms for halfspaces with malicious noise [Val85, KL93], under the assumption that the uncorrupted data comes from a “tame” distribution, e.g., Gaussian or isotropic log-concave. It should be noted that the class of γ -margin distributions considered in this work is significantly broader and can be far from satisfying the structural properties required in the aforementioned works.

A growing body of theoretical work has focused on *adversarially robust learning* (e.g., [BLPR19, MHS19, DNV19, Nak19]). In adversarially robust learning, the learner seeks to output a hypothesis with small γ -robust misclassification error, which for a hypothesis h and a norm $\|\cdot\|$ is typically defined as $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[\exists \mathbf{x}' \text{ with } \|\mathbf{x}' - \mathbf{x}\| \leq \gamma \text{ s.t. } h(\mathbf{x}') \neq y]$. Notice that when h is a halfspace and $\|\cdot\|$ is the Euclidean norm, the γ -robust misclassification error coincides with the γ -margin error in our context. (It should be noted that most of the literature on adversarially robust learning focuses on the ℓ_∞ -norm.) However, the objectives of the two learning settings are slightly different: in adversarially robust learning, the learner would like to output a hypothesis with small γ -robust misclassification error, whereas in our context the learner only has to output a hypothesis with small zero-one misclassification error. Nonetheless, as we point out in Remark 1.3, our algorithms can be adapted to provide guarantees in line with the adversarially robust setting as well.

Finally, in the distribution-independent agnostic setting without margin assumptions, there is compelling complexity-theoretic evidence that even weak learning of halfspaces is computationally intractable [GR06, FGKP06, DOSW11, Dan16, BGS18].

1.3 Our Contributions

We study the complexity of *proper* α -agnostic learning of γ -margin halfspaces on the unit ball. Our main result nearly characterizes the complexity of constant factor approximation to this problem:

Theorem 1.1. *There is an algorithm that uses $O(1/(\epsilon^2\gamma^2))$ samples, runs in time $\text{poly}(d/\epsilon) \cdot 2^{\tilde{O}(1/\gamma^2)}$ and is an $\alpha = 1.01$ -agnostic proper learner for γ -margin halfspaces with confidence probability $9/10$. Moreover, assuming the Randomized Exponential Time Hypothesis, any proper learning algorithm that achieves any constant factor approximation has runtime $\text{poly}(d/\epsilon) \cdot \Omega(2^{(1/\gamma)^{2-o(1)}})$.*

The reader is referred to Theorems 2.4 and 3.1 for detailed statements of the upper and lower bound respectively. A few remarks are in order: First, we note that the approximation ratio of 1.01 in the above theorem statement is not inherent. Our algorithm achieves $\alpha = 1 + \delta$, for any $\delta > 0$, with runtime $\text{poly}(d/\epsilon) \cdot 2^{\tilde{O}(1/(\delta\gamma^2))}$. The runtime of our algorithm significantly improves the runtime of the best known agnostic proper learner [BS00], achieving fixed polynomial dependence on $1/\epsilon$, independent of γ . This gain in runtime comes at the expense of losing a small constant factor in the error guarantee. It is natural to ask whether there exists a 1-agnostic proper learner matching the runtime of our Theorem 1.1. In Theorem 3.2, we establish a computational hardness result implying that such an improvement is unlikely.

The runtime dependence of our algorithm scales as $2^{\tilde{O}(1/\gamma^2)}$ (which is nearly best possible for proper learners), as opposed to $2^{\tilde{O}(1/\gamma)}$ in the best known improper learning algorithms [SSS09, BS12]. In addition to the interpretability of proper learning, we note that the sample complexity of our algorithm is quadratic in $1/\gamma$ (which is information-theoretically optimal), as opposed to

exponential for known improper learners. Moreover, for moderate values of γ , our algorithm may be faster than known improper learners, as it only uses spectral methods and ERM, as opposed to convex optimization. Finally, we note that the lower bound part of Theorem 1.1 implies a computational separation between proper and improper learning for our problem.

In addition, we explore the complexity of α -agnostic learning for large $\alpha > 1$. The following theorem summarizes our results in this setting:

Theorem 1.2. *There is an algorithm that uses $\tilde{O}(1/(\epsilon^2\gamma^2))$ samples, runs in time $\text{poly}(d) \cdot (1/\epsilon)^{\tilde{O}(1/(\alpha\gamma)^2)}$ and is an α -agnostic proper learner for γ -margin halfspaces with confidence probability $9/10$. Moreover, assuming $NP \neq RP$ and the Sliding Scale Conjecture, there exists an absolute constant $c > 0$, such that no $(1/\gamma)^c$ -agnostic proper learner runs in $\text{poly}(d, 1/\epsilon, 1/\gamma)$ time.*

The reader is referred to Theorem 2.7 for the upper bound and Theorem 3.3 for the lower bound. In summary, we give an α -agnostic proper learning algorithm with runtime exponential in $1/(\alpha\gamma)^2$, as opposed to $1/\gamma^2$, and we show that achieving $\alpha = (1/\gamma)^{\Omega(1)}$ is computationally hard. (Assuming only $NP \neq RP$, we can rule out polynomial time α -agnostic proper learning for $\alpha = (1/\gamma)^{\frac{1}{\text{poly}(\log(1/\gamma))}}$.)

Remark 1.3. While not stated explicitly in the subsequent analysis, our algorithms (with a slight modification to the associated constant factors) not only give a halfspace \mathbf{w}^* with zero-one loss at most $\alpha \cdot \text{OPT}_\gamma^{\mathcal{D}} + \epsilon$, but this guarantee holds for the 0.99γ -margin error² of \mathbf{w}^* as well. Thus, our learning algorithms also work in the adversarially robust setting (under the Euclidean norm) with a small loss in the “robustness parameter” (margin) from the one used to compute the optimum (i.e., γ) to the one used to measure the error of the output hypothesis (i.e., 0.99γ).

1.4 Our Techniques

Overview of Algorithms. For the sake of this intuitive explanation, we provide an overview of our algorithms when the underlying distribution \mathcal{D} is explicitly known. The finite sample analysis of our algorithms follows from standard generalization bounds (see Section 2).

Our constant factor approximation algorithm relies on the following observation: Let \mathbf{w}^* be the optimal weight vector. The assumption that $|\langle \mathbf{w}^*, \mathbf{x} \rangle|$ is large for almost all \mathbf{x} (by the margin property) implies a relatively strong condition on \mathbf{w}^* , which will allow us to find a relatively small search space containing a near-optimal solution. A first idea is to consider the matrix $\mathbf{M} = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[\mathbf{x}\mathbf{x}^T]$ and note that $\mathbf{w}^{*T}\mathbf{M}\mathbf{w}^* = \Omega(\gamma^2)$. This in turn implies that \mathbf{w}^* has a large component on the subspace spanned by the largest $O(1/(\epsilon\gamma^2))$ eigenvalues of \mathbf{M} . This idea suggests a basic algorithm that computes a net over unit-norm weight vectors on this subspace and outputs the best answer. This basic algorithm has runtime $\text{poly}(d) \cdot 2^{\tilde{O}(1/(\epsilon\gamma^2))}$ and is analyzed in Section 2.1.

To obtain our $\text{poly}(d/\epsilon) \cdot 2^{\tilde{O}(1/\gamma^2)}$ time constant factor approximation algorithm (establishing the upper bound part of Theorem 1.1), we use a refinement of the above idea. Instead of trying to guess the projection of \mathbf{w}^* onto the space of large eigenvectors *all at once*, we will do so in stages. In particular, it is not hard to see that \mathbf{w}^* has a non-trivial projection onto the subspace spanned by the top $O(1/\gamma^2)$ eigenvalues of \mathbf{M} . If we guess this projection, we will have some approximation to \mathbf{w}^* , but unfortunately not a sufficiently good one. However, we note that the difference between \mathbf{w}^* and our current hypothesis \mathbf{w} will have a large average squared inner product with the misclassified points. This suggests an iterative algorithm that in the i -th iteration considers the second moment matrix $\mathbf{M}^{(i)}$ of the points not correctly classified by the current hypothesis $\text{sign}(\langle \mathbf{w}^{(i)}, \mathbf{x} \rangle)$, guesses

²Here the constant 0.99 can be replaced by any constant less than one, with an appropriate increase to the algorithm’s running time.

a vector \mathbf{u} in the space spanned by the top few eigenvalues of $\mathbf{M}^{(i)}$, and sets $\mathbf{w}^{(i+1)} = \mathbf{u} + \mathbf{w}^{(i)}$. This procedure can be shown to produce a candidate set of weights with cardinality $2^{\tilde{O}(1/\gamma^2)}$ one of which has the desired misclassification error. This algorithm and its analysis are given in Section 2.2.

Our general α -agnostic algorithm (upper bound in Theorem 1.2) relies on approximating the *Chow parameters* of the target halfspace $f_{\mathbf{w}^*}$, i.e., the d numbers $\mathbf{E}[f_{\mathbf{w}^*}(\mathbf{x})\mathbf{x}_i]$, $i \in [d]$. A classical result [Cho61] shows that the exact values of the Chow parameters of a halfspace (over any distribution) uniquely define the halfspace. Although this fact is not very useful under an arbitrary distribution, the margin assumption implies a strong *approximate identifiability* result (Lemma 2.10). Combining this with an algorithm of [DDFS14], we can efficiently compute an approximation to the halfspace $f_{\mathbf{w}^*}$ given an approximation to its Chow parameters. In particular, if we can approximate the Chow parameters to ℓ_2 -error $\nu \cdot \gamma$, we can approximate $f_{\mathbf{w}^*}$ within error $\text{OPT}_\gamma^{\mathcal{D}} + \nu$.

A naive approach to approximate the Chow parameters would be via the empirical Chow parameters, namely $\mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[y\mathbf{x}]$. In the realizable case, this quantity indeed corresponds to the vector of Chow parameters. Unfortunately however, this method does not work in the agnostic case and it can introduce an error of $\omega(\text{OPT}_\gamma^{\mathcal{D}})$. To overcome this obstacle, we note that in order for a small fraction of errors to introduce a large error in the empirical Chow parameters, it must be the case that there is some direction \mathbf{w} in which many of these erroneous points introduce a large error. If we can guess some error that correlates well with \mathbf{w} and also guess the correct projection of our Chow parameters onto this vector, we can correct a decent fraction of the error between the empirical and true Chow parameters. We show that making the correct guesses $\tilde{O}(1/(\gamma\alpha)^2)$ times, we can reduce the empirical error sufficiently so that it can be used to find an accurate hypothesis. Once again, we can compute a hypothesis for each sequence of guesses and return the best one. See Section 2.3 for a detailed analysis.

Overview of Computational Lower Bounds. Our hardness results are shown via two reductions. These reductions take as input an instance of a computationally hard problem and produce a distribution \mathcal{D} on $\mathbb{B}_d \times \{\pm 1\}$. If the starting instance is a YES instance of the original problem, then $\text{OPT}_\gamma^{\mathcal{D}}$ is small for an appropriate value of γ . On the other hand, if the starting instance is a NO instance of the original problem, then $\text{OPT}_{0-1}^{\mathcal{D}}$ is large³. As a result, if there is a “too fast” (α -)agnostic proper learner for γ -margin halfspaces, then we would also get a “too fast” algorithm for the original problem as well, which would violate the corresponding complexity assumption.

To understand the margin parameter γ we can achieve, we need to first understand the problems we start with. For our reductions, the original problems can be viewed in the following form: select k items from v_1, \dots, v_N that satisfy certain “local constraints”. For instance, in our first construction, the reduction is from the k -Clique problem: Given a graph G and an integer k , the goal is to determine whether G contains a k -clique as a subgraph. For this problem, v_1, \dots, v_N correspond to the vertices of G and the “local” constraints are that every pair of selected vertices induces an edge.

Roughly speaking, our reduction produces a distribution \mathcal{D} on $\mathbb{B}_d \times \{\pm 1\}$ in dimension $d = N$, with the i -th dimension corresponding to v_i . The “ideal” solution in the YES case is to set $\mathbf{w}_i = \frac{1}{\sqrt{k}}$ iff v_i is selected and set $\mathbf{w}_i = 0$ otherwise. In our reductions, the local constraints are expressed using “sparse” sample vectors (i.e., vectors with only a constant number of non-zero coordinates all having the same magnitude). For example, in the case of k -Clique, the constraints can be expressed as follows: For every non-edge (i, j) , we must have $(\frac{1}{\sqrt{2}}\mathbf{e}^i + \frac{1}{\sqrt{2}}\mathbf{e}^j) \cdot \mathbf{w} \leq \frac{1}{\sqrt{2k}}$, where \mathbf{e}^i and \mathbf{e}^j denote the i -th and j -th vectors in the standard basis. A main step in both of our proofs is

³We use $\text{OPT}_{0-1}^{\mathcal{D}} \stackrel{\text{def}}{=} \min_{\mathbf{w} \in \mathbb{R}^d} \text{err}_{0-1}^{\mathcal{D}}(\mathbf{w})$ to denote the minimum error rate achievable by any halfspace.

to show that the reduction still works even when we “shift” the right hand side by a small multiple of $\frac{1}{\sqrt{k}}$. For instance, in the case of k -Clique, it is possible to show that, even if we replace $\frac{1}{\sqrt{2k}}$ with, say, $\frac{0.99}{\sqrt{2k}}$, the correctness of the construction remains, and we also get the added benefit that now the constraints are satisfied with a margin of $\gamma = \Theta(\frac{1}{\sqrt{k}})$ for our ideal solution in the YES case.

In the case of k -Clique, the above idea yields a reduction to 1-agnostic learning γ -margin halfspaces with margin $\gamma = \Theta(\frac{1}{\sqrt{k}})$, where the dimension d is N (and $\varepsilon = \frac{1}{\text{poly}(N)}$). As a result, if there is an $f(\frac{1}{\gamma})\text{poly}(d, \frac{1}{\varepsilon})$ -time algorithm for the latter for some function f , then there also exists a $g(k)\text{poly}(N)$ -time algorithm for k -Clique for some function g . The latter statement is considered unlikely, as it would break a widely-believed hypothesis in the area of parameterized complexity.

Ruling out α -agnostic learners, for $\alpha > 1$, is slightly more complicated, since we need to produce the “gap” of α between OPT_γ^D in the YES case and OPT_{0-1}^D in the NO case. To create such a gap, we appeal to the PCP Theorem [AS98, ALM⁺98], which can be thought of as an NP-hardness proof of the following “gap version” of 3SAT: given a 3CNF formula as input, distinguish between the case that the formula is satisfiable and the case that the formula is not even 0.9-satisfiable⁴. Moreover, further strengthened versions of the PCP Theorem [Din07, MR10] actually implies that this Gap-3SAT problem cannot even be solved in time $O(2^{n^{0.999}})$, where n denotes the number of variables in the formula, assuming the Exponential Time Hypothesis (ETH)⁵. Once again, (Gap-)3SAT can be viewed in the form of “item selection with local constraints”. However, the number of selected items k is now equal to n , the number of variables of the formula. With a similar line of reasoning as above, the margin we get is now $\gamma = \Theta(\frac{1}{\sqrt{k}}) = \Theta(\frac{1}{\sqrt{n}})$. As a result, if there is, say, a $2^{(1/\gamma)^{1.99}} \text{poly}(d, \frac{1}{\varepsilon})$ -time α -agnostic proper learner for γ -margin halfspaces (for an appropriate α), then there is an $O(2^{n^{0.995}})$ -time algorithm for Gap-3SAT, which would violate ETH.

Unfortunately, the above described idea only gives the “gap” α that is only slightly larger than 1, because the gap that we start with in the Gap-3SAT problem is already pretty small. To achieve larger gaps, our actual reduction starts from a generalization of 3SAT, called constraint satisfaction problems (CSPs), whose gap problems are hard even for very large gap. This concludes the outline of the main intuitions in our reductions. The detailed proofs are given in Section 3.

1.5 Preliminaries

For $n \in \mathbb{Z}_+$, we denote $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$. We will use small boldface characters for vectors and capital boldface characters for matrices. For a vector $\mathbf{x} \in \mathbb{R}^d$, and $i \in [d]$, \mathbf{x}_i denotes the i -th coordinate of \mathbf{x} , and $\|\mathbf{x}\|_2 \stackrel{\text{def}}{=} (\sum_{i=1}^d \mathbf{x}_i^2)^{1/2}$ denotes the ℓ_2 -norm of \mathbf{x} . We will use $\langle \mathbf{x}, \mathbf{y} \rangle$ for the inner product between $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. For a matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$, we will denote by $\|\mathbf{M}\|_2$ its spectral norm and by $\text{tr}(\mathbf{M})$ its trace. Let $\mathbb{B}_d = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 \leq 1\}$ be the unit ball and $\mathbb{S}_{d-1} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 = 1\}$ be the unit sphere in \mathbb{R}^d .

An origin-centered halfspace is a Boolean-valued function $h_{\mathbf{w}} : \mathbb{R}^d \rightarrow \{\pm 1\}$ of the form $h_{\mathbf{w}}(\mathbf{x}) = \text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle)$, where $\mathbf{w} \in \mathbb{R}^d$. (Note that we may assume w.l.o.g. that $\|\mathbf{w}\|_2 = 1$.) Let $\mathcal{H}_d = \{h_{\mathbf{w}}(\mathbf{x}) = \text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle), \mathbf{w} \in \mathbb{R}^d\}$ denote the class of all origin-centered halfspaces on \mathbb{R}^d . Finally, we use \mathbf{e}^i to denote the i -th standard basis vector, i.e., the vector whose i -th coordinate is one and the remaining coordinates are zero.

⁴In other words, for any assignment to the variables, at least 0.1 fraction of the clauses are unsatisfied.

⁵ETH states that the *exact* version of 3SAT cannot be solved in $2^{o(n)}$ time.

2 Efficient Proper Agnostic Learning of Halfspaces with a Margin

2.1 Warm-Up: Basic Algorithm

In this subsection, we present a basic algorithm that achieves $\alpha = 1$ and whose runtime is $\text{poly}(d)2^{\tilde{O}(1/(\epsilon\gamma^2))}$. Despite its slow runtime, this algorithm serves as a warm-up for our more sophisticated constant factor approximation algorithm in the next subsection.

We start by establishing a basic structural property of this setting which motivates our basic algorithm. We start with the following simple claim:

Claim 2.1. *Let $\mathbf{M}^{\mathcal{D}} = \mathbf{E}_{(\mathbf{x},y) \sim \mathcal{D}}[\mathbf{x}\mathbf{x}^T]$ and \mathbf{w}^* be a unit vector such that $\text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w}^*) \leq \text{OPT}_{\gamma}^{\mathcal{D}} \leq 1/2$. Then, we have that $\|\mathbf{M}^{\mathcal{D}}\|_2 \geq \mathbf{w}^{*T}\mathbf{M}^{\mathcal{D}}\mathbf{w}^* \geq \gamma^2/2$.*

Proof. By assumption, $\Pr_{(\mathbf{x},y) \sim \mathcal{D}}[|\langle \mathbf{w}^*, \mathbf{x} \rangle| \geq \gamma] \geq 1/2$, which implies that $\mathbf{E}_{(\mathbf{x},y) \sim \mathcal{D}}[(\langle \mathbf{w}^*, \mathbf{x} \rangle)^2] \geq \gamma^2/2$. The claim follows from the fact that $\mathbf{v}^T\mathbf{M}^{\mathcal{D}}\mathbf{v} = \mathbf{E}_{(\mathbf{x},y) \sim \mathcal{D}}[(\langle \mathbf{v}, \mathbf{x} \rangle)^2]$, for any $\mathbf{v} \in \mathbb{R}^d$, and the definition of the spectral norm. \square

Claim 2.1 allows us to obtain an approximation to the optimal halfspace by projecting on the space of large eigenvalues of $\mathbf{M}^{\mathcal{D}}$. We will need the following terminology: For $\delta > 0$, let $V_{\geq \delta}^{\mathcal{D}}$ be the space spanned by the eigenvalues of $\mathbf{M}^{\mathcal{D}}$ with magnitude at least δ and $V_{< \delta}^{\mathcal{D}}$ be its complement. Let $\text{Proj}_V(\mathbf{v})$ denote the projection operator of vector \mathbf{v} on subspace V . Then, we have the following:

Lemma 2.2. *Let $\delta > 0$ and $\mathbf{w}' = \text{Proj}_{V_{\geq \delta}^{\mathcal{D}}}(\mathbf{w}^*)$. Then, we have that $\text{err}_{\gamma/2}^{\mathcal{D}}(\mathbf{w}') \leq \text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w}^*) + 4\delta/\gamma^2$.*

Proof. Let $\mathbf{w}^* = \mathbf{w}' + \mathbf{w}''$, where $\mathbf{w}'' = \text{Proj}_{V_{< \delta}^{\mathcal{D}}}(\mathbf{w}^*)$. Observe that for any (\mathbf{x}, y) , if $y\langle \mathbf{w}', \mathbf{x} \rangle \leq \gamma/2$ then $y\langle \mathbf{w}^*, \mathbf{x} \rangle \leq \gamma$, unless $|\langle \mathbf{w}'', \mathbf{x} \rangle| \geq \gamma/2$. Hence, $\text{err}_{\gamma/2}^{\mathcal{D}}(\mathbf{w}') \leq \text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w}^*) + \Pr_{(\mathbf{x},y) \sim \mathcal{D}}[|\langle \mathbf{w}'', \mathbf{x} \rangle| \geq \gamma/2]$. By definition of \mathbf{w}'' and $\mathbf{M}^{\mathcal{D}}$, we have that $\mathbf{E}_{(\mathbf{x},y) \sim \mathcal{D}}[(\langle \mathbf{w}'', \mathbf{x} \rangle)^2] \leq \delta$. By Markov's inequality, we thus obtain $\Pr_{(\mathbf{x},y) \sim \mathcal{D}}[(\langle \mathbf{w}'', \mathbf{x} \rangle)^2 \geq \gamma^2/4] \leq 4\delta/\gamma^2$, completing the proof of the lemma. \square

Motivated by Lemma 2.2, the idea is to enumerate over $V_{\geq \delta}^{\mathcal{D}}$, for $\delta = \Theta(\epsilon\gamma^2)$, and output a vector \mathbf{v} with smallest empirical $\gamma/2$ -margin error. To turn this into an actual algorithm, we work with a finite sample set and enumerate over an appropriate cover of the space $V_{\geq \delta}^{\mathcal{D}}$. The pseudocode is as follows:

Algorithm 1 Basic 1-Agnostic Proper Learning Algorithm

- 1: Draw a multiset $S = \{(\mathbf{x}^{(i)}, y^{(i)})\}$ of i.i.d. samples from \mathcal{D} , where $m = \Omega(\log(1/\tau)/(\epsilon^2\gamma^2))$.
 - 2: Let $\hat{\mathcal{D}}_m$ be the empirical distribution on S .
 - 3: Let $\mathbf{M}^{\hat{\mathcal{D}}_m} = \mathbf{E}_{(\mathbf{x},y) \sim \hat{\mathcal{D}}_m}[\mathbf{x}\mathbf{x}^T]$.
 - 4: Set $\delta = \epsilon\gamma^2/16$. Use SVD to find a basis of $V_{\geq \delta}^{\hat{\mathcal{D}}_m}$.
 - 5: Compute a $\delta/2$ -cover, $C_{\delta/2}$, in ℓ_2 -norm, of $V_{\geq \delta}^{\hat{\mathcal{D}}_m} \cap \mathbb{S}_{d-1}$.
 - 6: Let $\mathbf{v} \in \arg\min_{\mathbf{w} \in C_{\delta/2}} \text{err}_{\gamma/4}^{\hat{\mathcal{D}}_m}(\mathbf{w})$.
 - 7: **return** $h_{\mathbf{v}}(\mathbf{x}) = \text{sign}(\langle \mathbf{v}, \mathbf{x} \rangle)$.
-

First, we analyze the runtime of our algorithm. The SVD of $\mathbf{M}^{\hat{\mathcal{D}}_m}$ can be computed in $\text{poly}(d/\delta)$ time. Note that $V_{\geq \delta}^{\hat{\mathcal{D}}_m}$ has dimension at most $1/\delta$. This follows from the fact that $\mathbf{M}^{\hat{\mathcal{D}}_m}$ is PSD and its trace is $\sum_{i=1}^d \lambda_i = \text{tr}(\mathbf{M}^{\hat{\mathcal{D}}_m}) = \mathbf{E}_{(\mathbf{x},y) \sim \hat{\mathcal{D}}_m}[\text{tr}(\mathbf{x}\mathbf{x}^T)] \leq 1$, where we used that $\|\mathbf{x}\|_2 \leq 1$

with probability 1 over $\widehat{\mathcal{D}}_m$. Therefore, the unit sphere of $V_{\geq \delta}^{\widehat{\mathcal{D}}_m}$ has a $\delta/2$ -cover $C_{\delta/2}$ of size $(2/\delta)^{O(1/\delta)} = 2^{\tilde{O}(1/(\epsilon\gamma^2))}$ that can be computed in output polynomial time.

We now prove correctness. The main idea is to apply Lemma 2.2 for the empirical distribution $\widehat{\mathcal{D}}_m$ combined with the following statistical bound:

Fact 2.3 ([BM02, McA03]). *Let $S = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^m$ be a multiset of i.i.d. samples from \mathcal{D} , where $m = \Omega(\log(1/\tau)/(\epsilon^2\gamma^2))$, and $\widehat{\mathcal{D}}_m$ be the empirical distribution on S . Then with probability at least $1 - \tau$ over S , simultaneously for all unit vectors \mathbf{w} and margins $\gamma > 0$, if $h_{\mathbf{w}}(\mathbf{x}) = \text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle)$, we have that $\text{err}_{0-1}^{\widehat{\mathcal{D}}_m}(h_{\mathbf{w}}) \leq \text{err}_{\gamma}^{\widehat{\mathcal{D}}_m}(\mathbf{w}) + \epsilon$.*

We proceed with the formal proof. First, we claim that for $m = \Omega(\log(1/\tau)/\epsilon^2)$, with probability at least $1 - \tau/2$ over S , we have that $\text{err}_{\gamma}^{\widehat{\mathcal{D}}_m}(\mathbf{w}^*) \leq \text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w}^*) + \epsilon/8$. To see this, note that $\text{err}_{\gamma}^{\widehat{\mathcal{D}}_m}(\mathbf{w}^*)$ can be viewed as a sum of Bernoulli random variables with expectation $\text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w}^*)$. Hence, the claim follows by a Chernoff bound. By an argument similar to that of Lemma 2.2, we have that $\text{err}_{\gamma/4}^{\widehat{\mathcal{D}}_m}(\mathbf{v}) \leq \text{err}_{\gamma/2}^{\widehat{\mathcal{D}}_m}(\mathbf{w}') + \epsilon/2$. Indeed, we can write $\mathbf{v} = \mathbf{w}' + \mathbf{r}$, where $\|\mathbf{r}\|_2 \leq \delta/2$, and follow the same argument.

In summary, we have the following sequence of inequalities:

$$\begin{aligned} \text{err}_{\gamma/4}^{\widehat{\mathcal{D}}_m}(\mathbf{v}) &\leq \text{err}_{\gamma/2}^{\widehat{\mathcal{D}}_m}(\mathbf{w}') + \epsilon/2 \leq \text{err}_{\gamma}^{\widehat{\mathcal{D}}_m}(\mathbf{w}^*) + \epsilon/2 + \epsilon/4 \\ &\leq \text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w}^*) + \epsilon/2 + \epsilon/4 + \epsilon/8, \end{aligned}$$

where the second inequality uses Lemma 2.2 for $\widehat{\mathcal{D}}_m$. Finally, we use Fact 2.3 for $\gamma/4$ and $\epsilon/8$ to obtain that $\text{err}_{0-1}^{\mathcal{D}}(h_{\mathbf{v}}) \leq \text{err}_{\gamma/4}^{\widehat{\mathcal{D}}_m}(\mathbf{v}) + \epsilon/8 \leq \text{OPT}_{\gamma}^{\mathcal{D}} + \epsilon$. The proof follows by a union bound.

2.2 Main Algorithm: Near-Optimal Constant Factor Approximation

In this section, we establish the following theorem, which gives the upper bound part of Theorem 1.1:

Theorem 2.4. *Fix $0 < \delta \leq 1$. There is an algorithm that uses $O(1/(\epsilon^2\gamma^2))$ samples, runs in time $\text{poly}(d/\epsilon) \cdot 2^{\tilde{O}(1/(\delta\gamma^2))}$ and is a $(1+\delta)$ -agnostic proper learner for γ -margin halfspaces with confidence probability $9/10$.*

Our algorithm in this section produces a finite set of candidate weight vectors and outputs the one with the smallest empirical $\gamma/2$ -margin error. For the sake of this intuitive description, we will assume that the algorithm knows the distribution \mathcal{D} in question supported on $\mathbb{B}_d \times \{\pm 1\}$. By assumption, there is a unit vector \mathbf{w}^* so that $\text{err}_{\gamma}^{\mathcal{D}}(\mathbf{w}^*) \leq \text{OPT}_{\gamma}^{\mathcal{D}}$.

We note that if a hypothesis $h_{\mathbf{w}}$ defined by vector \mathbf{w} has $\gamma/2$ -margin error at least a $(1+\delta)\text{OPT}_{\gamma}^{\mathcal{D}}$, then there must be a large number of points correctly classified with γ -margin by $h_{\mathbf{w}^*}$, but not correctly classified with $\gamma/2$ -margin by $h_{\mathbf{w}}$. For all of these points, we must have that $|\langle \mathbf{w}^* - \mathbf{w}, \mathbf{x} \rangle| \geq \gamma/2$. This implies that the $\gamma/2$ -margin-misclassified points of $h_{\mathbf{w}}$ have a large covariance in the $\mathbf{w}^* - \mathbf{w}$ direction. In particular, we have:

Claim 2.5. *Let $\mathbf{w} \in \mathbb{R}^d$ be such that $\text{err}_{\gamma/2}^{\mathcal{D}}(\mathbf{w}) > (1 + \delta)\text{OPT}_{\gamma}^{\mathcal{D}}$. Let \mathcal{D}' be \mathcal{D} conditioned on $y\langle \mathbf{w}, \mathbf{x} \rangle \leq \gamma/2$. Let $\mathbf{M}^{\mathcal{D}'} = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}'}[\mathbf{x}\mathbf{x}^T]$. Then $(\mathbf{w}^* - \mathbf{w})^T \mathbf{M}^{\mathcal{D}'} (\mathbf{w}^* - \mathbf{w}) \geq \delta\gamma^2/8$.*

Proof. We claim that with probability at least $\delta/2$ over $(\mathbf{x}, y) \sim \mathcal{D}'$ we have that $y\langle \mathbf{w}, \mathbf{x} \rangle \leq \gamma/2$ and $y\langle \mathbf{w}^*, \mathbf{x} \rangle \geq \gamma$. To see this, we first note that $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}'}[y\langle \mathbf{w}, \mathbf{x} \rangle > \gamma/2] = 0$ holds by definition of \mathcal{D}' . Hence, we have that

$$\Pr_{(\mathbf{x}, y) \sim \mathcal{D}'}[y\langle \mathbf{w}^*, \mathbf{x} \rangle \leq \gamma] \leq \frac{\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[y\langle \mathbf{w}^*, \mathbf{x} \rangle \leq \gamma]}{\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[y\langle \mathbf{w}, \mathbf{x} \rangle \leq \gamma/2]} < \frac{\text{OPT}_{\gamma}^{\mathcal{D}}}{(1 + \delta)\text{OPT}_{\gamma}^{\mathcal{D}}} = \frac{1}{(1 + \delta)}.$$

By a union bound, we obtain $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}'}[(y \langle \mathbf{w}, \mathbf{x} \rangle > \gamma/2) \cup (y \langle \mathbf{w}^*, \mathbf{x} \rangle \leq \gamma)] \leq \frac{1}{(1+\delta)}$.

Therefore, with probability at least $\delta/(1+\delta) \geq \delta/2$ (since $\delta \leq 1$) over $(\mathbf{x}, y) \sim \mathcal{D}'$ we have that $y \langle \mathbf{w}^* - \mathbf{w}, \mathbf{x} \rangle \geq \gamma/2$, which implies that $\langle \mathbf{w}^* - \mathbf{w}, \mathbf{x} \rangle^2 \geq \gamma^2/4$. Thus, $(\mathbf{w}^* - \mathbf{w})^T \mathbf{M}^{\mathcal{D}'} (\mathbf{w}^* - \mathbf{w}) = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}'}[\langle \mathbf{w}^* - \mathbf{w}, \mathbf{x} \rangle^2] \geq \delta\gamma^2/8$, completing the proof. \square

Claim 2.5 says that $\mathbf{w}^* - \mathbf{w}$ has a large component on the large eigenvalues of $\mathbf{M}^{\mathcal{D}'}$. Building on this claim, we obtain the following result:

Lemma 2.6. *Let $\mathbf{w}^*, \mathbf{w}, \mathbf{M}^{\mathcal{D}'}$ be as in Claim 2.5. There exists $k \in \mathbb{Z}_+$ so that if V_k is the span of the top k eigenvectors of $\mathbf{M}^{\mathcal{D}'}$, we have that $\|\text{Proj}_{V_k}(\mathbf{w}^* - \mathbf{w})\|_2^2 \geq k\delta\gamma^2/8$.*

Proof. Note that the matrix $\mathbf{M}^{\mathcal{D}'}$ is PSD and let $0 > \lambda_{\max} = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d \geq 0$ be its set of eigenvalues. We will denote by $V_{\geq t}$ the space spanned by the eigenvectors of $\mathbf{M}^{\mathcal{D}'}$ corresponding to eigenvalues of magnitude at least t . Let $d_t = \dim(V_{\geq t})$ be the dimension of $V_{\geq t}$, i.e., the number of $i \in [d]$ with $\lambda_i \geq t$. Since \mathbf{x} is supported on the unit ball, for $(\mathbf{x}, y) \sim \mathcal{D}'$, we have that $\text{tr}(\mathbf{M}^{\mathcal{D}'}) = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}'}[\text{tr}(\mathbf{x}\mathbf{x}^T)] \leq 1$. Since $\mathbf{M}^{\mathcal{D}'}$ is PSD, we have that $\text{tr}(\mathbf{M}^{\mathcal{D}'}) = \sum_{i=1}^d \lambda_i$ and we can write

$$1 \geq \text{tr}(\mathbf{M}^{\mathcal{D}'}) = \sum_{i=1}^d \lambda_i = \sum_{i=1}^d \int_0^{\lambda_i} 1 dt = \sum_{i=1}^d \int_0^{\lambda_{\max}} \mathbf{1}_{\lambda_i \geq t} dt = \int_0^{\lambda_{\max}} d_t dt, \quad (1)$$

where the last equality follows by changing the order of the summation and the integration. If the projection of $(\mathbf{w}^* - \mathbf{w})$ onto the i -th eigenvector of $\mathbf{M}^{\mathcal{D}'}$ has ℓ_2 -norm a_i , we have that

$$\delta\gamma^2/8 \leq (\mathbf{w}^* - \mathbf{w})^T \mathbf{M}^{\mathcal{D}'} (\mathbf{w}^* - \mathbf{w}) = \sum_{i=1}^d \lambda_i a_i^2 = \sum_{i=1}^d \int_0^{\lambda_{\max}} a_i^2 \mathbf{1}_{\lambda_i \geq t} dt = \int_0^{\lambda_{\max}} \|\text{Proj}_{V_{\geq t}}(\mathbf{w}^* - \mathbf{w})\|_2^2 dt, \quad (2)$$

where the first inequality uses Claim 2.5, the first equality follows by the Pythagorean theorem, and the last equality follows by changing the order of the summation and the integration.

Combining (1) and (2), we obtain $\int_0^{\lambda_{\max}} \|\text{Proj}_{V_{\geq t}}(\mathbf{w}^* - \mathbf{w})\|_2^2 dt \geq (\delta\gamma^2/8) \int_0^{\lambda_{\max}} d_t dt$. By an averaging argument, there exists $0 \leq t \leq \lambda_{\max}$ such that $\|\text{Proj}_{V_{\geq t}}(\mathbf{w}^* - \mathbf{w})\|_2^2 \geq (\delta\gamma^2/8)d_t$. Letting $k = d_t$ and noting that $V_{\geq t} = V_k$ completes the proof. \square

Lemma 2.6 suggests a method for producing an approximation to \mathbf{w}^* , or more precisely a vector that produces empirical $\gamma/2$ -margin error at most $(1+\delta)\text{OPT}_{\gamma}^{\mathcal{D}}$. We start by describing a non-deterministic procedure, which we will then turn into an actual algorithm.

The method proceeds in a sequence of stages. At stage i , we have a hypothesis weight vector $\mathbf{w}^{(i)}$. (At stage $i = 0$, we start with $\mathbf{w}^{(0)} = \mathbf{0}$.) At any stage i , if $\text{err}_{\gamma/2}^{\mathcal{D}}(\mathbf{w}^{(i)}) \leq (1+\delta)\text{OPT}_{\gamma}^{\mathcal{D}}$, then $\mathbf{w}^{(i)}$ is a sufficient estimator. Otherwise, we consider the matrix $\mathbf{M}^{(i)} = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}^{(i)}}[\mathbf{x}\mathbf{x}^T]$, where $\mathcal{D}^{(i)}$ is \mathcal{D} conditioned on $y \langle \mathbf{w}^{(i)}, \mathbf{x} \rangle \leq \gamma/2$. By Lemma 2.6, we know that for some positive integer value $k^{(i)}$, we have that the projection of $\mathbf{w}^* - \mathbf{w}^{(i)}$ onto $V_{k^{(i)}}$ has squared norm at least $\delta k^{(i)}\gamma^2/8$.

Let $\mathbf{p}^{(i)}$ be this projection. We set $\mathbf{w}^{(i+1)} = \mathbf{w}^{(i)} + \mathbf{p}^{(i)}$. Since the projection of $\mathbf{w}^* - \mathbf{w}^{(i)}$ and its complement are orthogonal, we have

$$\|\mathbf{w}^* - \mathbf{w}^{(i+1)}\|_2^2 = \|\mathbf{w}^* - \mathbf{w}^{(i)}\|_2^2 - \|\mathbf{p}^{(i)}\|_2^2 \leq \|\mathbf{w}^* - \mathbf{w}^{(i)}\|_2^2 - \delta k^{(i)}\gamma^2/8, \quad (3)$$

where the inequality uses the fact that $\|\mathbf{p}^{(i)}\|_2^2 \geq k^{(i)}\delta\gamma^2/8$ (as follows from Lemma 2.6). Let s be the total number of stages. We can write

$$1 \geq \|\mathbf{w}^* - \mathbf{w}^{(0)}\|_2^2 - \|\mathbf{w}^* - \mathbf{w}^{(s)}\|_2^2 = \sum_{i=0}^{s-1} \left(\|\mathbf{w}^* - \mathbf{w}^{(i)}\|_2^2 - \|\mathbf{w}^* - \mathbf{w}^{(i+1)}\|_2^2 \right) \geq (\delta\gamma^2/8) \sum_{i=0}^{s-1} k^{(i)},$$

where the first inequality uses that $\|\mathbf{w}^* - \mathbf{w}^{(0)}\|_2^2 = 1$ and $\|\mathbf{w}^* - \mathbf{w}^{(s)}\|_2^2 \geq 0$, the second notes the telescoping sum, and the third uses (3). We thus have that $s \leq \sum_{i=0}^{s-1} k^{(i)} \leq 8/(\delta\gamma^2)$. Therefore, the above procedure terminates after at most $8/(\delta\gamma^2)$ stages at some $\mathbf{w}^{(s)}$ with $\text{err}_{\gamma/2}^{\mathcal{D}}(\mathbf{w}^{(s)}) \leq (1 + \delta)\text{OPT}_{\gamma}^{\mathcal{D}}$.

We now describe how to turn the above procedure into an actual algorithm. Our algorithm tries to simulate the above described procedure by making appropriate guesses. In particular, we start by guessing a sequence of positive integers $k^{(i)}$ whose sum is at most $8/(\delta\gamma^2)$. This can be done in $2^{O(1/(\delta\gamma^2))}$ ways. Next, given this sequence, our algorithm guesses the vectors $\mathbf{w}^{(i)}$ over all s stages in order. In particular, given $\mathbf{w}^{(i)}$, the algorithm computes the matrix $\mathbf{M}^{(i)}$ and the subspace $V_{k^{(i)}}$, and guesses the projection $\mathbf{p}^{(i)} \in V_{k^{(i)}}$, which then gives $\mathbf{w}^{(i+1)}$. Of course, we cannot expect our algorithm to guess $\mathbf{p}^{(i)}$ exactly (as there are infinitely many points in $V_{k^{(i)}}$), but we can guess it to within ℓ_2 -error $\text{poly}(\gamma)$, by taking an appropriate net. This involves an additional guess of size $(1/\gamma)^{O(k^{(i)})}$ in each stage. In total, our algorithm makes $2^{\tilde{O}(1/(\delta\gamma^2))}$ many different guesses.

We note that the sample version of our algorithm is essentially identical to the idealized version described above, by replacing the distribution \mathcal{D} by its empirical version and leveraging Fact 2.3. The pseudo-code is given in Algorithm 2 below.

Algorithm 2 Near-Optimal $(1 + \delta)$ -Agnostic Proper Learner

- 1: Draw a multiset $S = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^m$ of i.i.d. samples from \mathcal{D} , where $m = \Omega(\log(1/\tau)/(\epsilon^2\gamma^2))$.
 - 2: Let $\widehat{\mathcal{D}}_m$ be the empirical distribution on S .
 - 3: **for** all sequences $k^{(0)}, k^{(1)}, \dots, k^{(s-1)}$ of positive integers with sum at most $8/(\delta\gamma^2) + 2$ **do**
 - 4: Let $\mathbf{w}^{(0)} = \mathbf{0}$.
 - 5: **for** $i = 0, 1, \dots, s - 1$ **do**
 - 6: Let $\mathcal{D}^{(i)}$ be $\widehat{\mathcal{D}}_m$ conditioned on $y\langle \mathbf{w}^{(i)}, \mathbf{x} \rangle \leq \gamma/2$.
 - 7: Let $\mathbf{M}^{(i)} = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}^{(i)}}[\mathbf{x}\mathbf{x}^T]$.
 - 8: Use SVD on $\mathbf{M}^{(i)}$ to find a basis for $V_{k^{(i)}}$, the span of the top $k^{(i)}$ eigenvectors.
 - 9: Let $C^{(i)}$ be a $\delta\gamma^3$ -cover, in ℓ_2 -norm, of $V_{k^{(i)}} \cap \mathbb{B}_d$ of size $(1/(\delta\gamma))^{O(k^{(i)})}$.
 - 10: For each $\mathbf{p}^{(i)} \in C^{(i)}$ repeat the next step of the for loop with $\mathbf{w}^{(i+1)} = \mathbf{w}^{(i)} + \mathbf{p}^{(i)}$.
 - 11: **end for**
 - 12: **end for**
 - 13: Let C denote the set of all $\mathbf{w}^{(i)}$ generated in the above loop.
 - 14: Let $\mathbf{v} \in \text{argmin}_{\mathbf{w} \in C} \text{err}_{\gamma/2}^{\widehat{\mathcal{D}}_m}(\mathbf{w})$.
 - 15: **return** $h_{\mathbf{v}}(\mathbf{x}) = \text{sign}(\langle \mathbf{v}, \mathbf{x} \rangle)$.
-

To show the correctness of the algorithm, we begin by arguing that the set C of candidate weight vectors produced has size $2^{\tilde{O}(1/(\delta\gamma^2))}$. This is because there are only $2^{O(1/(\delta\gamma^2))}$ many possibilities for the sequence of $k^{(i)}$, and for each such sequence the product of the sizes of the $C^{(i)}$ is $(1/(\delta\gamma))^{O(\sum k^{(i)})} = 2^{\tilde{O}(1/(\delta\gamma^2))}$. We note that, by the aforementioned analysis, for any choice of $k^{(0)}, \dots, k^{(i-1)}$ and $\mathbf{w}^{(i)}$, we either have that $\text{err}_{\gamma/2}^{\widehat{\mathcal{D}}_m}(\mathbf{w}^{(i)}) \leq (1 + \delta)\text{OPT}_{\gamma}^{\widehat{\mathcal{D}}_m}$ or there is a choice of $k^{(i)}$ and $\mathbf{p}^{(i)} \in C^{(i)}$ such that

$$\|\mathbf{w}^* - \mathbf{w}^{(i)} - \mathbf{p}^{(i)}\|_2^2 \leq \|\mathbf{w}^* - \mathbf{w}^{(i)}\|_2^2 - \delta k^{(i)}\gamma^2/8 + O(\delta^2\gamma^6),$$

where we used (3) and the fact that $C^{(i)}$ is a $\delta\gamma^3$ -cover of $V_{k^{(i)}}$. Following the execution path of

the algorithm, we either find some $\mathbf{w}^{(i)}$ with $\text{err}_{\gamma/2}^{\widehat{\mathcal{D}}_m}(\mathbf{w}^{(i)}) \leq (1 + \delta)\text{OPT}_{\gamma}^{\widehat{\mathcal{D}}_m}$, or we find a $\mathbf{w}^{(i)}$ with

$$\|\mathbf{w}^* - \mathbf{w}^{(i)}\|_2^2 \leq 1 - \left(\sum_{j=0}^{i-1} k^{(j)} \right) \delta \gamma^2 / 8 + O(\delta \gamma^4),$$

where the last term is an upper bound for $\left(\sum_{j=0}^{i-1} k^{(j)} \right) \cdot O(\delta^2 \gamma^6)$. Note that this sequence terminates in at most $O(1/(\delta \gamma^2))$ stages, when it becomes impossible that $\sum k^{(j)} > 8/(\delta \gamma^2) + 1$. Thus, the output of our algorithm must contain some weight vector \mathbf{v} with $\text{err}_{\gamma/2}^{\widehat{\mathcal{D}}_m}(\mathbf{v}) \leq (1 + \delta)\text{OPT}_{\gamma}^{\widehat{\mathcal{D}}_m}$. The proof now follows by an application of Fact 2.3. This completes the proof of Theorem 2.4.

2.3 α -Agnostic Proper Learning Algorithm

In this section, we show that if one wishes to obtain an α -agnostic proper learner for some large $\alpha \gg 1$, one can obtain runtime exponential in $1/(\alpha \gamma)^2$ rather than $1/\gamma^2$. Formally, we prove:

Theorem 2.7. *There is an algorithm that uses $\tilde{O}(1/(\epsilon^2 \gamma^2))$ samples, runs in time $\text{poly}(d) \cdot (1/\epsilon)^{\tilde{O}(1/(\alpha \gamma)^2)}$ and is an α -agnostic proper learner for γ -margin halfspaces with probability 9/10.*

Let \mathcal{D} be a distribution over $\mathbb{B}_d \times \{-1, 1\}$. Suppose that there exists a unit vector $\mathbf{w}^* \in \mathbb{R}^d$ such that $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[y \langle \mathbf{w}^*, \mathbf{x} \rangle \geq \gamma] \geq 1 - \text{OPT}_{\gamma}^{\mathcal{D}}$ for some $\text{OPT}_{\gamma}^{\mathcal{D}} > 0$. Suppose additionally that $\gamma, \epsilon > 0$ and $\alpha > 1$. We will describe an algorithm that given sample access to \mathcal{D} along with γ, α, ϵ and $\text{OPT}_{\gamma}^{\mathcal{D}}$, draws $O(\log(\alpha/\epsilon)/(\gamma \epsilon)^2)$ samples, runs in time $\text{poly}(d) \cdot (1/\gamma \epsilon)^{\tilde{O}(1/(\alpha \gamma)^2)}$ and with probability at least 9/10 returns a \mathbf{w} with

$$\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[\text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle) \neq y] < O(\alpha \cdot \text{OPT}_{\gamma}^{\mathcal{D}} + \epsilon).$$

We begin by giving an algorithm that works if the distribution \mathcal{D} is known explicitly. We will be able to reduce to this case by using the empirical distribution over a sufficiently large set of samples. That is, we start by establishing the following:

Proposition 2.8. *Let \mathcal{D} be an explicit distribution over $\mathbb{B}_d \times \{-1, 1\}$. Suppose there exists a unit vector \mathbf{w}^* so that $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[y \langle \mathbf{w}^*, \mathbf{x} \rangle \geq \gamma] \geq 1 - \text{OPT}_{\gamma}^{\mathcal{D}}$ for some $\text{OPT}_{\gamma}^{\mathcal{D}} > 0$. Additionally, let $\gamma > 0$ and $\alpha > 1$. There exists an algorithm that given \mathcal{D} along with $\gamma, \alpha, \text{OPT}_{\gamma}^{\mathcal{D}}$, runs in time $\text{poly}(d) \cdot (|\text{supp}(\mathcal{D})|/(\alpha \gamma \cdot \text{OPT}_{\gamma}^{\mathcal{D}}))^{\tilde{O}(1/(\alpha \gamma)^2)}$ and returns a weight vector \mathbf{w} with $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[\text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle) \neq y] < O(\alpha \cdot \text{OPT}_{\gamma}^{\mathcal{D}})$.*

Our main technical tool here will be the vector of Chow parameters [Cho61, OS11, DDFS14], i.e., vector of degree-1 “Fourier coefficients”, of the target halfspace:

Definition 2.9. Given a Boolean function $f : \mathbb{B}_d \rightarrow \{\pm 1\}$, and a distribution $\mathcal{D}_{\mathbf{x}}$ on \mathbb{B}_d the *Chow parameters vector* of f , is the vector $\mathbf{Chow}(f)$ given by the expectation $\mathbf{E}_{\mathbf{x} \sim \mathcal{D}_{\mathbf{x}}}[f(\mathbf{x})\mathbf{x}]$.

It is well-known [Cho61] that the vector of Chow parameters uniquely identifies any halfspace within the class of all Boolean functions. Several robust versions of this fact are known (see, e.g., [Gol06, OS11, DS09, DDFS14, DKS18, DK19]) under various structural assumptions on the underlying distribution. Here we leverage the margin assumption to obtain a robust version of this fact. Specifically, we show that learning the Chow parameters of the halfspace $f_{\mathbf{w}^*}(\mathbf{x}) = \text{sign}(\langle \mathbf{w}^*, \mathbf{x} \rangle)$ determines the function $f_{\mathbf{w}^*}$ up to small error.

In the following, we will denote by $\mathcal{D}_{\mathbf{x}}$ the marginal distribution of \mathcal{D} on \mathbb{B}_d . We have the following simple lemma:

Lemma 2.10. *Let $g : \mathbb{B}_d \rightarrow \{\pm 1\}$ be any Boolean function that satisfies $\Pr_{\mathbf{x} \sim \mathcal{D}_x}[f_{\mathbf{w}^*}(\mathbf{x}) \neq g(\mathbf{x})] \geq \nu + \text{OPT}_\gamma^{\mathcal{D}}$, for some $\nu > 0$. Then, we have that $\|\mathbf{Chow}(f_{\mathbf{w}^*}) - \mathbf{Chow}(g)\|_2 \geq \nu \cdot \gamma$.*

Proof. We can write

$$\begin{aligned} \|\mathbf{Chow}(f_{\mathbf{w}^*}) - \mathbf{Chow}(g)\|_2 &\geq \langle \mathbf{w}^*, \mathbf{Chow}(f_{\mathbf{w}^*}) - \mathbf{Chow}(g) \rangle \\ &= \mathbf{E}_{\mathbf{x} \sim \mathcal{D}_x}[\langle \mathbf{w}^*, \mathbf{x} \rangle (f_{\mathbf{w}^*}(\mathbf{x}) - g(\mathbf{x}))] \\ &= 2\mathbf{E}_{\mathbf{x} \sim \mathcal{D}_x}[\langle \mathbf{w}^*, \mathbf{x} \rangle \cdot \mathbf{1}_{f_{\mathbf{w}^*}(\mathbf{x}) \neq g(\mathbf{x})}]. \end{aligned}$$

Recalling our assumptions $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[y \langle \mathbf{w}^*, \mathbf{x} \rangle \geq \gamma] \geq 1 - \text{OPT}_\gamma^{\mathcal{D}}$ and $\Pr_{\mathbf{x} \sim \mathcal{D}_x}[f_{\mathbf{w}^*}(\mathbf{x}) \neq g(\mathbf{x})] \geq \nu + \text{OPT}_\gamma^{\mathcal{D}}$, we note that there is at least a ν probability over $(\mathbf{x}, y) \sim \mathcal{D}$ that $f_{\mathbf{w}^*}(\mathbf{x}) \neq g(\mathbf{x})$ and $y \langle \mathbf{w}^*, \mathbf{x} \rangle \geq \gamma$, which implies that $|\langle \mathbf{w}^*, \mathbf{x} \rangle| \geq \gamma$. Therefore, the above expectation is at least $\nu \cdot \gamma$. \square

Lemma 2.10, combined with the algorithms in [TTV08, DDFS14], implies that learning an approximation to $\mathbf{Chow}(f_{\mathbf{w}^*})$ is sufficient to learn a good hypothesis.

Lemma 2.11. *There is a polynomial time algorithm that given an explicit distribution \mathcal{D} and a vector \mathbf{c} with $\|\mathbf{Chow}(f_{\mathbf{w}^*}) - \mathbf{c}\|_2 \leq \nu \cdot \gamma$, returns a vector \mathbf{w} that with high probability satisfies $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[f_{\mathbf{w}}(\mathbf{x}) \neq f_{\mathbf{w}^*}(\mathbf{x})] \leq O(\nu + \text{OPT}_\gamma^{\mathcal{D}})$. In particular, for this \mathbf{w} we have that $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[\text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle) \neq y] = O(\nu + \text{OPT}_\gamma^{\mathcal{D}})$.*

Thus, it will suffice to approximate the Chow parameters of $f_{\mathbf{w}^*}$ to error $\alpha\gamma \cdot \text{OPT}_\gamma^{\mathcal{D}}$. One might consider using the empirical Chow parameters, namely $P = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[y\mathbf{x}]$ for this purpose. In the realizable case, this would be the right thing to do, but this naive approach fails in the agnostic setting. Instead, our approach hinges on the following observation: Since $y = f_{\mathbf{w}^*}(\mathbf{x})$ for all but an $\text{OPT}_\gamma^{\mathcal{D}}$ -fraction of \mathbf{x} 's, and since the \mathbf{x} 's are supported in the unit ball, the error has ℓ_2 -norm at most $\text{OPT}_\gamma^{\mathcal{D}}$. In fact, if we have some vector \mathbf{w} so that $\langle \mathbf{w}, P - \mathbf{Chow}(f_{\mathbf{w}^*}) \rangle \geq \alpha\gamma \cdot \text{OPT}_\gamma^{\mathcal{D}}$, then there must be some (\mathbf{x}, y) in the domain of \mathcal{D} with $|\langle \mathbf{x}, \mathbf{w} \rangle| \geq \alpha\gamma$. The idea is to guess this \mathbf{w} and then guess the true projection of $\mathbf{Chow}(f_{\mathbf{w}^*})$ onto \mathbf{w} .

We present the pseudo-code for the algorithm establishing Proposition 2.8 as Algorithm 3 below.

Algorithm 3 α -Agnostic Proper Learner of Proposition 2.8

- 1: Let $m = \lceil \log(1/\alpha\gamma) / (\alpha\gamma)^2 \rceil$.
 - 2: Let $P = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[y\mathbf{x}]$
 - 3: **for** every sequence $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}$ from $\text{supp}(\mathcal{D})$ **do**
 - 4: Let V be the span of $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}$.
 - 5: Let \mathcal{C} be a $(\alpha\gamma \cdot \text{OPT}_\gamma^{\mathcal{D}})$ -cover of the unit ball of V .
 - 6: **for** each $g \in \mathcal{C}$ **do**
 - 7: Let P' be obtained by replacing the projection of P onto V with g . In particular, $P' = P - \text{Proj}_V(P) + g$.
 - 8: Run the algorithm of Lemma 2.11 to find a hypothesis \mathbf{w} .
 - 9: **end for**
 - 10: **end for**
 - 11: **return** The hypothesis that produces smallest empirical error among all \mathbf{w} 's in Line 8.
-

Proof of Proposition 2.8. Firstly, note that the runtime of this algorithm is clearly $\text{poly}(d) \left(\frac{|\text{supp}(\mathcal{D})|}{\text{OPT}_\gamma^{\mathcal{D}} \cdot \alpha\gamma} \right)^{\tilde{O}(1/(\alpha\gamma)^2)}$. It remains to show correctness. We note that by Lemma 2.11 it suffices to show that some P'

is within $O(\alpha\gamma \cdot \text{OPT}_\gamma^{\mathcal{D}})$ of $\mathbf{Chow}(f_{\mathbf{w}^*})$. For this it suffices to show that there is a sequence $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}$ so that $\|\text{Proj}_{V^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2 = O(\alpha\gamma \cdot \text{OPT}_\gamma^{\mathcal{D}})$.

To show this, let V_i be the span of $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(i)}$. We claim that if $\|\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2 \geq \alpha\gamma \cdot \text{OPT}_\gamma^{\mathcal{D}}$, then there exists an $\mathbf{x}^{(i+1)}$ in the support of \mathcal{D} such that

$$\|\text{Proj}_{V_{i+1}^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2^2 = \|\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2^2 \cdot (1 - (\alpha\gamma)^2).$$

To show this, we let \mathbf{w} be the unit vector in the direction of $\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)$. We note that

$$\|\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2 = \langle \mathbf{w}, \mathbf{Chow}(f_{\mathbf{w}^*}) - P \rangle = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[\langle \mathbf{w}, \mathbf{x} \rangle (\text{sign}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)].$$

Since $\text{sign}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y$ is 0 for all but an $\text{OPT}_\gamma^{\mathcal{D}}$ -fraction of (\mathbf{x}, y) , we have that there must be some $\mathbf{x}^{(i+1)}$ so that $\langle \mathbf{x}^{(i+1)}, \mathbf{w} \rangle \geq \|\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2 / \text{OPT}_\gamma^{\mathcal{D}} \geq \alpha\gamma$. If we chose this $\mathbf{x}^{(i+1)}$, we have that

$$\begin{aligned} \|\text{Proj}_{V_{i+1}^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2^2 &\leq \|\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2^2 - \langle \mathbf{x}^{(i+1)}, \mathbf{Chow}(f_{\mathbf{w}^*}) - P \rangle^2 \\ &= \|\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2^2 \cdot (1 - \langle \mathbf{x}^{(i+1)}, \mathbf{w} \rangle^2) \\ &= \|\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2^2 \cdot (1 - (\alpha\gamma)^2). \end{aligned}$$

Therefore, unless $\|\text{Proj}_{V_i^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2^2 < \alpha\gamma \cdot \text{OPT}_\gamma^{\mathcal{D}}$ already for some $i < m$, there exists a sequence $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(m)}$ such that

$$\begin{aligned} \|\text{Proj}_{V_m^\perp}(\mathbf{Chow}(f_{\mathbf{w}^*}) - P)\|_2^2 &\leq \|P - \mathbf{Chow}(f_{\mathbf{w}^*})\|_2^2 \cdot (1 - (\alpha\gamma)^2)^{-m} \\ &\leq \|P - \mathbf{Chow}(f_{\mathbf{w}^*})\|_2^2 \cdot \exp(-m \cdot (\alpha\gamma)^2) \\ &\leq \text{OPT}_\gamma^{\mathcal{D}} \cdot \exp(\log(\alpha\gamma)) \\ &= \text{OPT}_\gamma^{\mathcal{D}} \cdot \alpha\gamma. \end{aligned}$$

So in either case, we have some sequence of \mathbf{x} 's so that the projection onto V^\perp of $\mathbf{Chow}(f_{\mathbf{w}^*}) - P$ is sufficiently small. This completes our analysis. \square

In order to extend this to a proof of Theorem 2.7, we will need to reduce to solving the problem on a finite sample set. This result can be obtained from Proposition 2.8 by some fairly simple reductions.

Firstly, we note that we can assume that $\text{OPT}_\gamma^{\mathcal{D}} \geq \epsilon/\alpha$, as increasing it to this value does not change the problem.

Secondly, we note that if we let $\widehat{\mathcal{D}}$ be the empirical distribution over a set of $\Omega(d/\epsilon^2)$ random samples, then with at least $2/3$ probability we have the following:

- $\Pr_{(\mathbf{x}, y) \sim \widehat{\mathcal{D}}} [y \langle \mathbf{w}^*, \mathbf{x} \rangle \geq \gamma] \geq 1 - O(\text{OPT}_\gamma^{\mathcal{D}})$.
- For any vector \mathbf{w} , $\Pr_{(\mathbf{x}, y) \sim \mathcal{D}} [\text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle) \neq y] = \Pr_{(\mathbf{x}, y) \sim \widehat{\mathcal{D}}} [\text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle) \neq y] + O(\epsilon)$.

The first statement here is by applying the Markov inequality to the probability that $y \langle \mathbf{w}^*, \mathbf{x} \rangle < \gamma$, and the second is by the VC-inequality [DL01]. We note that if the above hold, applying the algorithm from Proposition 2.8 to $\widehat{\mathcal{D}}$ will produce an appropriate \mathbf{w} . This produces an algorithm that uses $O(d/\epsilon^2)$ samples and has runtime $O(d/\gamma\epsilon)^{\tilde{O}(1/(\alpha\gamma)^2)}$.

Unfortunately, this algorithm is not quite satisfactory as the runtime and sample complexity scale poorly with the dimension d . In order to fix this, we will make use of an idea from [KS04]. Namely, we will first apply dimension reduction to a smaller number of dimensions before applying our algorithm. In particular, we will make use of the Johnson-Lindenstrauss lemma:

Lemma 2.12 ([JL84]). *There exists a probability distribution over linear transformations $A : \mathbb{R}^d \rightarrow \mathbb{R}^m$ with $m = O(\log(1/\delta)/\epsilon^2)$ so that for any unit vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^d$, $\Pr_A[|\langle \mathbf{v}, \mathbf{w} \rangle - \langle A\mathbf{v}, A\mathbf{w} \rangle| > \epsilon] < \delta$. Additionally, there are efficient algorithms to sample from such distributions over A .*

We note that this implies in particular that $\|A\mathbf{v}\|_2 = 1 \pm \epsilon$ except with probability δ . Thus, by tweaking the parameters a little bit and letting $h_A(\mathbf{v}) = A\mathbf{v}/\|A\mathbf{v}\|_2$, we have that $h_A(\mathbf{v})$ is always a unit vector and that $\langle h_A(\mathbf{v}), h_A(\mathbf{w}) \rangle = \langle \mathbf{v}, \mathbf{w} \rangle \pm \epsilon$ except with probability δ .

Next, we note that by taking $\epsilon = \gamma/2$ and $\delta = \text{OPT}_\gamma^{\mathcal{D}}$ in the above we have that

$$\begin{aligned} & \Pr_{A, (\mathbf{x}, y) \sim \mathcal{D}}[y \langle h_A(\mathbf{w}^*), h_A(\mathbf{x}) \rangle < \gamma/2] \\ & \leq \Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[y \langle \mathbf{w}^*, \mathbf{x} \rangle < \gamma] + \Pr_{A, (\mathbf{x}, y) \sim \mathcal{D}}[|\langle h_A(\mathbf{w}^*), h_A(\mathbf{x}) \rangle - \langle \mathbf{w}^*, \mathbf{x} \rangle| > \gamma/2] \\ & = O(\text{OPT}_\gamma^{\mathcal{D}}). \end{aligned}$$

Thus, by the Markov inequality, with large constant probability over A we have that

$$\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[y \langle h_A(\mathbf{w}^*), h_A(\mathbf{x}) \rangle < \gamma/2] = O(\text{OPT}_\gamma^{\mathcal{D}}).$$

But this means that the distribution $(h_A(\mathbf{x}), y)$ satisfies the assumptions for our algorithm (with γ replaced by $\gamma/2$ and $\text{OPT}_\gamma^{\mathcal{D}}$ by $O(\text{OPT}_\gamma^{\mathcal{D}})$), but in dimension $m = O(\log(\alpha/\epsilon)/\gamma^2)$. Running the algorithm described above on this set will find us a vector \mathbf{w} so that

$$\Pr_{(\mathbf{x}, y) \sim \mathcal{D}}[\text{sign}(\langle \mathbf{w}, h_A(\mathbf{x}) \rangle) \neq y] = O(\alpha \cdot \text{OPT}_\gamma^{\mathcal{D}} + \epsilon).$$

However, it should be noted that

$$\text{sign}(\langle \mathbf{w}, h_A(\mathbf{x}) \rangle) = \text{sign}(\langle \mathbf{w}, A\mathbf{x} \rangle / \|A\mathbf{x}\|_2) = \text{sign}(\langle \mathbf{w}, A\mathbf{x} \rangle) = \text{sign}(\langle A^T \mathbf{w}, \mathbf{x} \rangle).$$

Thus, $A^T \mathbf{w}$ satisfies the necessary conditions.

Our final algorithm is given below:

Algorithm 4 α -Agnostic Proper Learner of Theorem 2.7

- 1: Pick $A : \mathbb{R}^d \rightarrow \mathbb{R}^m$ with $m = O(\log(\alpha/\epsilon)/\gamma^2)$ from an appropriate Johnson-Lindenstrauss family and define f_A appropriately.
 - 2: Take $O(m/\epsilon^2)$ random samples and let $\widehat{\mathcal{D}}$ be the uniform distribution over $(A\mathbf{x}/\|A\mathbf{x}\|_2, y)$ for samples (\mathbf{x}, y) from this set.
 - 3: Run the algorithm from Proposition 2.8 on $\widehat{\mathcal{D}}$ using $\gamma/2$ instead of γ to find a vector \mathbf{w} .
 - 4: **return** $A^T \mathbf{w}$.
-

3 Computational Hardness Results

In this section, we provide several computational lower bounds for agnostic learning of halfspaces with a margin. To clarify the statements below, we note that we say “there is no algorithm that runs in time $T(d, \frac{1}{\gamma}, \frac{1}{\epsilon})$ ” to mean that no $T(d, \frac{1}{\gamma}, \frac{1}{\epsilon})$ -time algorithm works for *all* combinations of parameters d, γ and ϵ . (Note that we discuss the lower bounds with stronger quantifiers in

Section 3.1.) Moreover, we also ignore the dependency on τ (the probability that the learner can be incorrect), since we only use a fixed τ (say $1/3$) in all the bounds below.

First, we show that, for any constant $\alpha > 1$, α -agnostic learning of γ -margin halfspaces requires $2^{(1/\gamma)^{2-o(1)}} \text{poly}(d, 1/\varepsilon)$ time. Up to the lower order term $\gamma^{o(1)}$ in the exponent, this matches the runtime of our algorithm (in Theorem 2.4). In fact, we show an even stronger result, namely that if the dependency of the running time on the margin is, say, $2^{(1/\gamma)^{1.99}}$, then one has to pay a nearly exponential dependence on d , i.e., $2^{d^{1-o(1)}}$.

This result holds assuming the so-called (randomized) exponential time hypothesis (ETH) [IP01, IPZ01], which postulates that there is no (randomized) algorithm that can solve 3SAT in time $2^{o(n)}$, where n denotes the number of variables. ETH is a standard hypothesis used in proving (tight) running time lower bounds. We do not discuss ETH further here, but interested readers may refer to a survey by Lokshantov et al. [LMS11] for an in-depth discussion and several applications of ETH.

Our first lower bound can be stated more precisely as follows:

Theorem 3.1. *Assuming the (randomized) ETH, for any universal constant $\alpha \geq 1$, there is no proper α -agnostic learner for γ -margin halfspaces that runs in time $O(2^{(1/\gamma)^{2-o(1)}} 2^{d^{1-o(1)}}) f(\frac{1}{\varepsilon})$ for any function f .*

Secondly, we address the question of whether we can achieve $\alpha = 1$ (standard agnostic learning) while retaining running time similar to that of our algorithm. We answer this in the negative (assuming a standard parameterized complexity assumption): there is no $f(\frac{1}{\gamma}) \text{poly}(d, \frac{1}{\varepsilon})$ -time 1-agnostic learner for any function f (e.g., even for $f(\frac{1}{\gamma}) = 2^{2^{1/\gamma}}$). This demonstrates a stark contrast between what we can achieve with and without approximation.

Theorem 3.2. *Assuming $W[1]$ is not contained in randomized FPT, there is no proper 1-agnostic learner for γ -margin halfspaces that runs in time $f(\frac{1}{\gamma}) \text{poly}(d, \frac{1}{\varepsilon})$ for any function f .*

Finally, we explore the other extreme of the trade-off between the running time and approximation ratio, by asking: what is the best approximation ratio we can achieve if we only consider proper learners that run in $\text{poly}(d, \frac{1}{\varepsilon}, \frac{1}{\gamma})$ -time? On this front, it is known [Ser01] that the perceptron algorithm achieves $1/\gamma$ -approximation. We show that a significant improvement over this is unlikely, by showing that $(1/\gamma)^{\frac{1}{\text{polyloglog}(1/\gamma)}}$ -approximation is not possible unless $\text{NP} = \text{RP}$. If we additionally assume the so-called Sliding Scale Conjecture [BGLR94], this ratio can be improved to $(1/\gamma)^c$ for some constant $c > 0$.

Theorem 3.3. *Assuming $\text{NP} \neq \text{RP}$, there is no proper $(1/\gamma)^{1/\text{polyloglog}(1/\gamma)}$ -agnostic learner for γ -margin halfspaces that runs in time $\text{poly}(d, \frac{1}{\varepsilon}, \frac{1}{\gamma})$. Furthermore, assuming $\text{NP} \neq \text{RP}$ and the Sliding Scale Conjecture (Conjecture 1), there is no proper $(1/\gamma)^c$ -agnostic learning for γ -margin halfspaces that runs in time $\text{poly}(d, \frac{1}{\varepsilon}, \frac{1}{\gamma})$ for some constant $c > 0$.*

We note here that the constant c in Theorem 3.3 is not explicit, i.e., it depends on the constant from the Sliding Scale Conjecture (SSC). Moreover, even when assuming the most optimistic parameters of SSC, the constant c we can get is still very small. For instance, it is still possible that a say $\sqrt{1/\gamma}$ -agnostic learning algorithm that runs in polynomial time exists, and this remains an interesting open question. We remark that Daniely et al. [DLS14] have made partial progress in this direction by showing that, any $\text{poly}(d, \frac{1}{\varepsilon}, \frac{1}{\gamma})$ -time learner that belongs to a “generalized linear family” cannot achieve approximation ratio α better than $\Omega\left(\frac{1/\gamma}{\text{polylog}(1/\gamma)}\right)$. We note that the inapproximability ratio of [DLS14] is close to being tight for a natural, yet restricted, family of

improper learners. On the other hand, our proper hardness result holds against *all* proper learners under a widely believed worst-case complexity assumption.

3.1 Lower Bounds with Stronger Quantifiers on Parameters

Before we proceed to our proofs, let us first state a running time lower bound with stronger quantifiers. Recall that previously we only rule out algorithms that work for *all* combinations of d, γ, ε . Below we relax the quantifier so that we need the *for all* quantifier only for d .

Lemma 3.4. *Assuming the (randomized) ETH, for any universal constant $\alpha \geq 1$, there exists $\varepsilon_0 = \varepsilon_0(\alpha)$ such that there is no α -agnostic learner for γ -margin halfspaces that runs in time $O(2^{(1/\gamma)^{2-o(1)}})\text{poly}(d)$ for all d and for some $0 < \varepsilon < \varepsilon_0$ and $\frac{1}{d^{0.5+o(1)}} \leq \gamma = \gamma(d) \leq \frac{1}{(\log d)^{0.5+o(1)}}$ that satisfies $\frac{\gamma(d+1)}{\gamma(d)} \geq \Omega(1)$.*

We remark here that the lower and upper bounds on γ are essentially (i.e., up to lower order terms) the best possible. On the upper bound front, if $\gamma \geq \tilde{O}\left(\frac{1}{\sqrt{\log d}}\right)$, then our algorithmic result (Theorem 2.4) already give a $\text{poly}(d, \frac{1}{\varepsilon})$ -time α -agnostic learner for γ -margin halfspaces (for all constant $\alpha > 1$). On the other hand, if $\gamma \leq O\left(\frac{1}{d^{0.5+o(1)}}\right)$, then the trivial algorithm that exactly solves ERM for $m = O\left(\frac{d}{\varepsilon^2}\right)$ samples only takes $2^{O(d/\varepsilon^2)}$ time, which is already asymptotically faster than $2^{(1/\gamma)^{2-o(1)}}$. The last condition that $\frac{\gamma(d+1)}{\gamma(d)}$ is not too small is a sanity-check condition that prevents “sudden jumps” in $\gamma(d)$ such as $\gamma(d) = \frac{1}{(\log d)^{0.1}}$ and $\gamma(d+1) = \frac{1}{(d+1)^{0.1}}$; note that the condition is satisfied by “typical functions” such as $\gamma(d) = \frac{1}{d^c}$ or $\gamma(d) = \frac{1}{(\log d)^c}$ for some constant c .

As for ε , we only require the algorithm to work for any ε that is not *too large*, i.e., no larger than $\varepsilon_0(\alpha)$. This latter number is just a constant (when α is a constant). We note that it is still an interesting open question to make this requirement as mild as possible; specifically, is it possible to only require the algorithm to work for any $\varepsilon < 1/2$?

3.2 Reduction from k -Clique and Proof of Theorem 3.2

We now proceed to the proofs of our results, starting with Theorem 3.2.

To prove Theorem 3.2, we reduce from the k -Clique problem. In k -Clique, we are given a graph G and an integer k , and the goal is to determine whether the graph G contains a k -clique (as a subgraph).

We take the perspective of *parameterized complexity*. Recall that a parameterized problem with parameter k is said to be *fixed parameter tractable (FPT)* if it can be solved in time $f(k)\text{poly}(n)$ for some computable function f , where n denotes the input size.

It is well-known that k -Clique is complete for the class $W[1]$ [DF95]. In other words, under the (widely-believed) assumption that $W[1]$ does not collapse to FPT (the class of fixed parameter tractable problems), we cannot solve k -Clique in time $f(k)\text{poly}(n)$ for any computable function f . We shall not formally define the class $W[1]$ here; interested readers may refer to the book of Downey and Fellows for an in-depth treatment of the topic [DF13].

Our reduction starts with an instance of k -Clique and produces an instance of agnostic learning with margin γ such that $\gamma = \Omega(1/k)$ (and the dimension is polynomial):

Lemma 3.5. *There exists a polynomial-time algorithm that takes as input an n -vertex graph instance G and an integer k , and produces a distribution \mathcal{D} over $\mathbb{B}_d \times \{\pm 1\}$ and $\gamma, \kappa \in [0, 1]$ such that*

- (Completeness) If G contains a k -clique, then $\text{OPT}_\gamma^{\mathcal{D}} \leq \kappa$.
- (Soundness) If G does not contain a k -clique, then $\text{OPT}_{0-1}^{\mathcal{D}} > \kappa + \frac{0.001}{n^3}$.
- (Margin Parameter) $\gamma \geq \Omega(\frac{1}{\sqrt{k}})$.

We remark here that, in Lemma 3.5 and throughout the remainder of the section, we say that an algorithm produces a distribution \mathcal{D} over $\mathbb{B}_d \times \{\pm 1\}$ to mean that it outputs the set of samples $\{(\mathbf{x}^{(i)}, y^{(i)})\}_{i \in [m]}$ and numbers d_i for each $i \in [m]$ representing the probability of $(\mathbf{x}^{(i)}, y^{(i)})$ with respect to \mathcal{D} . Note that this is stronger than needed since, to prove hardness of learning, it suffices to have an oracle that can sample from \mathcal{D} , but here we actually explicitly produce a full description of \mathcal{D} . Moreover, note that this implicitly implies that the support of \mathcal{D} is of polynomial size (and hence, for any given h , $\text{err}_\gamma^{\mathcal{D}}(h)$ and $\text{err}_{0-1}^{\mathcal{D}}(h)$ can be efficiently computed).

As stated above, Lemma 3.5 immediately implies Theorem 3.2 because, if we can agnostically learn γ -margin halfspaces in time $f(\frac{1}{\gamma})\text{poly}(d, \frac{1}{\varepsilon})$, then we can solve k -Clique in $f(O(\sqrt{k}))\text{poly}(n)$ time, which would imply that $\text{W}[1]$ is contained in (randomized) FPT. This is formalized below.

Proof of Theorem 3.2. Suppose that we have an $f(\frac{1}{\gamma})\text{poly}(d, \frac{1}{\varepsilon})$ -time agnostic learner for γ -margin halfspaces. Given an instance (G, k) of k -Clique, we run the reduction from Lemma 3.5 to produce a distribution \mathcal{D} . We then run the learner on \mathcal{D} with $\varepsilon = \frac{0.001}{n^3}$ (and with $\delta = 1/3$). Note that the learner runs in time $f(O(\sqrt{k}))\text{poly}(n)$ and produces a halfspace h . We then compute $\text{err}_{0-1}^{\mathcal{D}}(h)$; if it is no more than $\kappa + \frac{0.001}{n^3}$, then we output YES. Otherwise, we output NO.

The algorithm described above solves k -Clique (correctly with probability $2/3$) in FPT time. Since k -Clique is $\text{W}[1]$ -complete, this implies that $\text{W}[1]$ is contained in randomized FPT. \square

We now move on to prove Lemma 3.5. Before we do so, let us briefly describe the ideas behind it. The dimension d will be set to n , the number of vertices of G . Each coordinate \mathbf{w}_i is associated with a vertex $i \in V(G)$. In the completeness case, we would like to set $\mathbf{w}_i = \frac{1}{\sqrt{k}}$ iff i is in the k -clique and $\mathbf{w}_i = 0$ otherwise. To enforce a solution to be of this form, we add two types of samples that induces the following constraints:

- *Non-Edge Constraint:* for every *non*-edge (i, j) , we should have $\mathbf{w}_i + \mathbf{w}_j \leq \frac{1}{\sqrt{k}}$. That is, we should “select” at most one vertex among i, j .
- *Vertex Selection Constraint:* each coordinate of \mathbf{w} is at least $\frac{1}{\sqrt{k}}$. Note that we will violate such constraints for all vertices, except those that are “selected”.

If we select the probabilities in \mathcal{D} so that the non-edge constraints are weighted much larger than the vertex selection constraints, then it is always better to not violate any of the first type of constraints. When this is the case, the goal will now be to violate as few vertex selection constraints as possible, which is the same as finding a maximum clique, as desired.

While the above paragraph describes the core idea of the reduction, there are two additional issues we have to resolve:

- *Constant Coordinate:* first, notice that we cannot actually quite write a constraint of the form $\mathbf{w}_i + \mathbf{w}_j \leq \frac{1}{\sqrt{k}}$ using the samples because there is no way to express a value like $\frac{1}{\sqrt{k}}$ directly. To overcome this, we have a “constant coordinate” \mathbf{w}_* , which is supposed to be a constant, and replace the right hand side of non-edge constraints by $\frac{\mathbf{w}_*}{\sqrt{k}}$ (instead of $\frac{1}{\sqrt{k}}$). The new constraint can now be represented by a sample.

- *Margin:* in the above reduction, there was no margin at all! To get the appropriate margin, we “shift” the constraint slightly so that there is a margin. For instance, instead of $\frac{\mathbf{w}_*}{\sqrt{k}}$ for a non-edge constraint, we use $\frac{1.1\mathbf{w}_*}{\sqrt{k}}$. We now have a margin of $\approx \frac{0.1}{\sqrt{k}}$ and it is still possible to argue that the best solution is still to select a clique.

The reduction, which follows the above outline, is formalized below.

Proof of Lemma 3.5. Given a graph $G = (V, E)$, we use n to denote the number of vertices $|V|$ and we rename its vertices so that $V = [n]$. We set $d = n + 1$; we name the first coordinate $*$ and each of the remaining coordinates $i \in [n]$. For brevity, let us also define $\beta = 1 - \frac{0.01}{n^2}$. The distribution \mathcal{D} is defined as follows:

- Add a labeled sample $(-\mathbf{e}^*, -1)$ with probability $\frac{\beta}{2}$ in \mathcal{D} . We refer to this as the *positivity constraint for $*$* .
- For every pair of distinct vertices i, j that do not induce an edge in E , add a labeled sample $(\frac{1}{2}(\frac{1.1}{\sqrt{k}}\mathbf{e}^* - \mathbf{e}^i - \mathbf{e}^j), 1)$ with probability $\frac{\beta}{2\binom{n}{2} - |E|}$ in \mathcal{D} . We refer to this as the *non-edge constraint for (i, j)* .
- For every vertex i , add a labeled sample $(\frac{1}{2}(\mathbf{e}^i - \frac{0.9}{\sqrt{k}}\mathbf{e}^*), 1)$ with probability $\frac{0.01}{n^3}$ in \mathcal{D} . We refer to this as the *vertex selection constraint for i* .

Finally, let $\gamma = \frac{0.1}{2\sqrt{2k}}$, $\kappa = (n - k) \left(\frac{0.01}{n^3}\right)$. It is obvious that the reduction runs in polynomial time.

Completeness. Suppose that G contains a k -clique; let $S \subseteq V$ denote the set of its vertices. We define \mathbf{w} by $\mathbf{w}_* = \frac{1}{\sqrt{2}}$ and, for every $i \in V$, $\mathbf{w}_i = \frac{1}{\sqrt{2k}}$ if $i \in S$ and $\mathbf{w}_i = 0$ otherwise. It is clear that $\|\mathbf{w}\|_2 = 1$ and that, for every $(\mathbf{x}, y) \in \text{supp}(\mathcal{D})$, we have $|\langle \mathbf{w}, \mathbf{x} \rangle| \geq \frac{0.1}{2\sqrt{2k}}$. Finally, observe that all the first two types of constraints are satisfied, and a vertex selection constraint for i is unsatisfied iff $i \notin S$. Thus, we have $\text{err}_\gamma^{\mathcal{D}}(\mathbf{w}) = (n - k) \left(\frac{0.01}{n^3}\right) = \kappa$, which implies that $\text{OPT}_\gamma^{\mathcal{D}} \leq \kappa$ as desired.

Soundness. Suppose contrapositively that $\text{OPT}_{0-1}^{\mathcal{D}} \leq \kappa + \frac{0.001}{n^3}$; that is, there exists \mathbf{w} such that $\text{err}_{0-1}^{\mathcal{D}}(\mathbf{w}) \leq \kappa + \frac{0.001}{n^3}$. Observe that each labeled sample of the first two types of constraints has probability more than $\frac{\beta}{2n^2} > \kappa + \frac{0.001}{n^3}$. As a result, \mathbf{w} must correctly classify these samples. Since \mathbf{w} correctly classifies $(-\mathbf{e}^*, -1)$, it must be that $w_* > 0$.

Now, let T be the set of vertices i such that \mathbf{w} mislabels the vertex selection constraint for i . Observe that $|T| < \frac{(\kappa + \frac{0.001}{n^3})}{\frac{0.01}{n^3}} < n - k + 1$. In other words, $S := V \setminus T$ is of size at least k . We claim that S induces a k -clique in G . To see that this is true, consider a pair of distinct vertices $i, j \in S$. Since \mathbf{w} satisfies the vertex selection constraints for i and for j , we must have $\mathbf{w}_i, \mathbf{w}_j \geq \frac{0.9}{\sqrt{k}}$. This implies that (i, j) is an edge, as otherwise \mathbf{w} would mislabel the non-edge constraint for (i, j) .

As a result, G contains a k -clique as desired. \square

3.3 Reduction from k -CSP and Proofs of Theorems 3.1, 3.3 and Lemma 3.4

In this section, we will prove Theorems 3.1 and 3.3, by reducing from the hardness of approximation of constraint satisfaction problems (CSPs), given by PCP Theorems.

3.3.1 CSPs and PCP Theorem(s)

Before we can state our reductions, we have to formally define CSPs and state the PCP theorems we will use more formally. We start with the definition of k -CSP:

Definition 3.6 (k -CSP). For any integer $k \in \mathbb{N}$, a k -CSP instance $\mathcal{L} = (V, \Sigma, \{\Pi_q\}_{q \in \mathcal{Q}})$ consists of

- The variable set V ,
- The alphabet Σ , which we sometimes refer to as labels,
- Constraints set $\{\Pi_S\}_{S \in \mathcal{Q}}$, where $\mathcal{Q} \subseteq \binom{V}{k}$ is a collection of k -size subset of V . For each subset $S = \{v_1, \dots, v_k\}$, $\Pi_S \subseteq \Sigma^S$ is the set of accepting answers for the constraint Π_S .

Here we think of each $f \in \Sigma^S$ as a function from $f : S \rightarrow \Sigma$.

A k -CSP instance is said to be *regular* if each variable appears in the same number of constraints.

An assignment ϕ is a function $\phi : V \rightarrow \Sigma$. Its value, denoted by $\text{val}_{\mathcal{L}}(\phi)$, is the fraction of constraints $S \in \mathcal{Q}$ such that⁶ $\phi|_S \in \Pi_S$. Such constraints are said to be *satisfied* by ϕ . The value of \mathcal{L} , denoted by $\text{val}(\mathcal{L})$, is the maximum value among all assignments, i.e., $\text{val}(\mathcal{L}) := \max_{\phi} \text{val}_{\mathcal{L}}(\phi)$.

In the ν -GAP- k -CSP problem, we are given a regular instance \mathcal{L} of k -CSP, and we want to distinguish between $\text{val}(\mathcal{L}) = 1$ and $\text{val}(\mathcal{L}) < \nu$.

Throughout this subsection, we use n to denote the instance size of k -CSP, that is $n = \sum_{S \in \mathcal{Q}} |\Pi_S|$.

The celebrated PCP theorem [AS98, ALM⁺98] is equivalent to the proof of NP-hardness of approximating ν -Gap- k -CSP for some constant k and $\nu < 1$. Since we would like to prove (tight) running time lower bounds, we need the versions of PCP Theorems that provides strong running time lower bounds as well. For this task, we turn to the Moshkovitz-Raz PCP theorem, which can not only achieve arbitrarily small constant $\nu > 0$ but also almost exponential running time lower bound.

Theorem 3.7 (Moshkovitz-Raz PCP [MR10]). *Assuming ETH, for any $0 < \nu < 1$, ν -Gap-2-CSP cannot be solved in time $O(2^{n^{1-o(1)}})$, even for instances with $|\Sigma| = O_{\nu}(1)$.*

As for our hardness of approximation result (Theorem 3.3), we are aiming to get as large a ratio as possible. For this purpose, we will use a PCP Theorem of Dinur, Harsha and Kindler, which achieves $\nu = \frac{1}{\text{poly}(n)}$ but need k to be $\text{polyloglog}(n)$.

Theorem 3.8 (Dinur-Harsha-Kindler PCP [DHK15]). *$n^{-\Omega(1)}$ -Gap-polyloglog(n)-CSP is NP-hard.*

Finally, we state the Sliding Scale Conjecture (SSC) of Bellare et al. [BGLR94], which says that the NP-hardness with $\nu = \frac{1}{\text{poly}(n)}$ holds even in the case where k is constant:

Conjecture 1 (Sliding Scale Conjecture [BGLR94]). *For some constant k , $n^{-\Omega(1)}$ -Gap- k -CSP is NP-hard.*

⁶We use $\phi|_S$ to denote the restriction of ϕ on the domain S .

3.3.2 Reducing from k -CSP to Agnostically Learning Halfspaces with Margin

Having set up the notation, we now move on to the reduction from k -CSP to agnostic learning of halfspaces with margin. Our reduction can be viewed as a modification of the reduction from [ABSS97]; compared to [ABSS97], we have to (1) be more careful so that we can get the margin in the completeness and (2) modify the reduction to work even for $k > 2$.

Before we precisely state the formal properties of the reduction, let us give a brief informal intuition behind the reduction. Given an instance $\mathcal{L} = (V, \Sigma, \{\Pi_S\}_{S \in \mathcal{Q}})$ of k -CSP, we will create a distribution \mathcal{D} on $\mathbb{B}_d \times \{\pm 1\}$, where the dimension d is equal to n . Each coordinate is associated with an accepting answer of each constraint; that is, each coordinate is (S, f) where $S \in \mathcal{Q}$ and $f \in \Pi_S$. In the completeness case where we have a perfect assignment ϕ , we would like the halfspace's normal vector to set $\mathbf{w}_{(S,f)} = 1$ iff f is the assignment to predicate S in ϕ (i.e., $f = \phi|_S$), and zero otherwise. To enforce this, we add three types of constraints:

- *Non-negativity Constraint*: that each coordinate of \mathbf{w} should be non-negative.
- *Satisfiability Constraint*: that for each $S \in \mathcal{Q}$, $\mathbf{w}_{(S,f)}$ is positive for at least one $f \in \Pi_S$.
- *Selection Constraint*: for each variable $v \in V$ and label $\sigma \in \Sigma$, we add a constraint that the sum of all $\mathbf{w}_{(S,f)}$, for all S that v appears in and all f that assigns σ to v , is non-positive.

Notice that, for the completeness case, we satisfy the first two types of constraints, and we violate the selection constraints only when $\phi(v) = \sigma$. Intuitively, in the soundness case, we will not be able to “align” the positive $\mathbf{w}_{(S,f)}$ from different S 's together, and we will have to violate a lot more selection constraints.

Of course, there are many subtle points that the above sketch does not address, such as the margin; on this front, we add one more special coordinate \mathbf{w}_* , which we think of as being equal to 1, and we add/subtract δ times this coordinate to each of the constraints, which will create the margin for us. Another issue is that the normal vector of the halfspace (and samples) as above have norm more than one. Indeed, our assignment in the completeness case has norm $O(\sqrt{n})$. Hence, we have to scale the normal vector down by a factor of $O(\sqrt{n})$, which results in a margin of $\gamma = \Omega(1/\sqrt{n})$. This is the reason why we arrive at the running time lower bound of the form $2^{\gamma^{2-o(1)}}$.

The properties and parameter dependencies of the reduction are encapsulated in the following theorem:

Theorem 3.9. *There exists a polynomial time reduction that takes as input a regular instance $\mathcal{L} = (V, \Sigma, \{\Pi_S\}_{S \in \mathcal{Q}})$ of k -CSP and a real number $\nu > 0$, and produces a distribution \mathcal{D} over $\mathbb{B}_d \times \{\pm 1\}$ and positive real numbers $\gamma, \kappa, \varepsilon, \alpha$ such that*

- (Completeness) *If \mathcal{L} is satisfiable, then $\text{OPT}_\gamma^{\mathcal{D}} \leq \kappa$.*
- (Soundness) *If $\text{val}(\mathcal{L}) < \nu$, then $\text{OPT}_{0-1}^{\mathcal{D}} > \alpha \cdot \kappa + \varepsilon$.*
- (Margin Parameter) $\gamma = \Omega\left(\frac{1}{\Delta|\Sigma|^{3k}\sqrt{|\mathcal{Q}|}}\right)$, *where Δ denotes the number of constraints each variable appears in.*
- (Approximation Ratio) $\alpha = \Omega\left(\frac{(1/\nu)^{1/k}}{k}\right)$.
- (Error Parameter) $\varepsilon = \Omega\left(\frac{1}{\Delta|\Sigma|^k}\right) \cdot \alpha$.
- (Dimension) $d = n + 1$.

Proof. Before we define \mathcal{D} , let us specify the parameters:

- First, we let d be $1 + n$. We name the first coordinate as $*$ and each of the remaining coordinates are named (S, f) for a constraint $S \in \mathcal{Q}$ and $f \in \Pi_S$.
- Let $Z := 2(|V| \cdot |\Sigma| + 2k|\mathcal{Q}| + 2k \sum_{e \in E} |\Pi_e|)$ be our “normalizing factor”, which will be used below to normalized the probability.
- Let $\delta := \frac{0.1}{\Delta|\Sigma|^{2k}}$ be the “shift parameter”. Note that this is not the margin γ (which will be defined below).
- Let $s := 10\Delta|\Sigma|^k$ be the scaling factor, which we use to make sure that all our samples lie within the unit ball.
- Let the gap parameter α be $\frac{(1/\nu)^{1/k}}{40k}$.
- Finally, let $\kappa = \frac{|V|}{Z}$ and $\varepsilon = \kappa \cdot \alpha$.

Note that α as defined above can be less than one. However, this is not a problem: in the subsequent proofs of Theorems 3.1 and 3.3, we will always choose the settings of parameters so that $\alpha > 1$.

We are now ready to define the distribution \mathcal{D} on $\mathbb{B}_d \times \{\pm 1\}$, as follows:

1. Add a labeled sample $(-\mathbf{e}^*, -1)$ with probability $1/2$ to \mathcal{D} . This corresponds to the constraint $\mathbf{w}_* > 0$; we refer to this as the *positivity constraint for $*$* .
2. Next, for each coordinate (S, f) , add a labeled sample $\left(\frac{1}{s} \left(\mathbf{e}^{(S,f)} + \delta \cdot \mathbf{e}^*\right), 1\right)$ with probability $2k/Z$ to \mathcal{D} . This corresponds to $\mathbf{w}_{(S,f)} + \delta \cdot \mathbf{w}_* \geq 0$ scaled down by a factor of $1/s$ so that the vector is in the unit ball; we refer to this as the *non-negativity constraint for (S, f)* .
3. For every $S \in \mathcal{Q}$, add a labeled sample $\left(\frac{1}{s} \left(\sum_{f \in \Pi_S} \mathbf{e}^{(S,f)} - (1 - \delta)\mathbf{e}^*\right), 1\right)$ with probability $2k/Z$ to \mathcal{D} . This corresponds to the constraint $\sum_{f \in \Pi_S} \mathbf{w}_{(S,f)} \geq (1 - \delta)\mathbf{w}_*$, scaled down by a factor of $1/s$. We refer to this constraint as the *satisfiability constraint for S* .
4. For every variable $v \in V$ and $\sigma \in \Sigma$, add a labeled sample $\left(\frac{1}{s} \left(\sum_{S \in \mathcal{Q}: v \in S} \sum_{f \in \Pi_S: f(v)=\sigma} \mathbf{e}^{(S,f)} - \delta\mathbf{e}^*\right), -1\right)$ with probability $1/Z$ to \mathcal{D} . This corresponds to the constraint $\sum_{S \in \mathcal{Q}: v \in S} \sum_{f \in \Pi_S: f(v)=\sigma} \mathbf{w}_{(S,f)} < \delta \cdot \mathbf{w}_*$, scaled down by a factor of $1/s$. We refer to this as the *selection constraint for (v, σ)* .

Completeness. Suppose that there exists an assignment $\phi : V \rightarrow \Sigma$ that satisfies all the constraints of \mathcal{L} . Consider the halfspace with normal vector \mathbf{w} defined by $\mathbf{w}_* = \zeta$ and

$$\mathbf{w}_{(S,f)} = \begin{cases} \zeta & \text{if } f = \phi|_S, \\ 0 & \text{otherwise,} \end{cases}$$

where $\zeta := \frac{1}{\sqrt{1+|\mathcal{Q}|}}$ is the normalization factor. It is easy to see that the positivity constraints and the satisfiability constraints are satisfied with margin at least $\gamma = \zeta \cdot \delta/s = \Omega\left(\frac{1}{\Delta|\Sigma|^{3k}\sqrt{|\mathcal{Q}|}}\right)$. Finally, observe that the sum $\sum_{S \in \mathcal{Q}: v \in S} \sum_{f \in \Pi_S: f(v)=\sigma} \mathbf{w}_{(S,f)}$ is zero if $f(v) \neq \sigma$; in this case, the selection constraint for (v, σ) is also satisfied with margin at least γ . As a result, we only incur an error (with respect to margin γ) for the selection constraint for $(v, \phi(v))$ for all $v \in V$; hence, we have $\text{err}_\gamma^{\mathcal{D}}(\mathbf{w}) \leq \frac{1}{2} \cdot |V| = \kappa$ as desired.

Soundness. Suppose contrapositively that there exists \mathbf{w} with $\text{err}_{0-1}^{\mathcal{D}}(\mathbf{w}) \leq \alpha \cdot \kappa + \varepsilon = 2\alpha\kappa$. We will “decode” back an assignment with value at least ν of the CSP from \mathbf{w} .

To do so, first observe that from the positivity constraint for $*$, we must have $\mathbf{w}_* > 0$, as otherwise we would already incur an error of $1/2 > 2\alpha\kappa$ with respect to \mathcal{D} . Now, since scaling (by a positive factor) does not change the fraction of samples violated, we may assume w.l.o.g. that $\mathbf{w}_* = 1$.

Next, we further claim that we may assume without loss of generality that \mathbf{w} does not violate any non-negativity constraints (2) or satisfiability constraints (3). The reason is that, if \mathbf{w} violates a non-negativity constraint for $(S = \{v_1, \dots, v_k\}, f)$, then we may simply change $\mathbf{w}_{(S,f)}$ to zero. This reduces the error by $2k/Z$, while it may only additionally violate k additional selection constraints for $(v_1, f(v_1)), \dots, (v_k, f(v_k))$ which weights k/Z in total with respect to \mathcal{D} . As a result, this change only reduces the error in total. Similarly, if the satisfiability constraint of S is unsatisfied, we may change $\mathbf{w}_{(S,f)}$ for some $f \in \Pi_S$ to a sufficiently large number so that this constraint is satisfied; once again, in total the error decreases. Hence, we may assume that the non-negativity constraints (2) and satisfiability constraints (3) all hold.

Now, for every vertex v , let $L_v \subseteq \Sigma$ denote the set of labels $\sigma \in \Sigma$ such that the selection constraint for (v, σ) is violated. Since we assume that $\text{err}_{0-1}^{\mathcal{D}}(\mathbf{w}) \leq 2\alpha\kappa$, we must have $\sum_{v \in V} |L_v| \leq (2\alpha\kappa)/(1/Z) = 2\alpha|V|$.

Next, let V_{small} denote the set of all variables $v \in V$ such that $|L_v| \leq 20\alpha k$. From the bound we just derived, we must have $|V_{\text{small}}| \geq \left(1 - \frac{1}{10k}\right) |V|$.

Another ingredient we need is the following claim:

Claim 3.10. *For every constraint $S = \{v_1, \dots, v_k\} \in \mathcal{Q}$, there exist $\sigma_1 \in L_{v_1}, \dots, \sigma_k \in L_{v_k}$ that induces an accepting assignment for Π_S (i.e., $f \in \Pi_S$ where f is defined by $f(v_i) = \sigma_i$).*

Proof. Suppose for the sake of contradiction that no such $\sigma_1 \in L_{v_1}, \dots, \sigma_k \in L_{v_k}$ exists. In other words, for every $f \in \Pi_S$, there must exist $i \in [k]$ such that the selection constraint for $(v_i, f(v_i))$ is not violated. This means that

$$\begin{aligned} \delta = \delta \cdot \mathbf{w}_* &> \sum_{S' \in \mathcal{Q}: v \in S'} \sum_{f' \in \Pi_{S'}: f'(v) = \sigma} \mathbf{w}_{(S', f')} \\ &\geq \mathbf{w}_{(S, f)} + \sum_{S' \in \mathcal{Q}: v \in S'} \sum_{f' \in \Pi_{S'}: f'(v) = \sigma} -\delta \cdot \mathbf{w}_* \\ &\geq \mathbf{w}_{(S, f)} - \delta \cdot \Delta |\Sigma|^k, \end{aligned}$$

where the second inequality comes from our assumption, that the non-negativity constraints are satisfied.

Hence, by summing this up over all $f \in \Pi_S$, we get

$$\sum_{f \in \Pi_S} \mathbf{w}_{(S, f)} \leq \delta \cdot (\Delta |\Sigma|^k + 1) \cdot |\Sigma|^k < (1 - \delta),$$

which means that the satisfiability constraint for S is violated, a contradiction. \square

We can now define an assignment $\phi : V \rightarrow \Sigma$ for \mathcal{L} as follows. For every $v \in V$, let $\phi(v)$ be a random label in L_v . Notice here that, by Claim 3.10, the probability that a constraint $S = \{v_1, \dots, v_k\}$ is satisfied is at least $\prod_{i \in [k]} |L_{v_i}|^{-1}$. Hence, the expected total number of satisfied

constraints is at least

$$\begin{aligned} \sum_{S=\{v_1, \dots, v_k\} \in \mathcal{Q}} \prod_{i \in [k]} |L_{v_i}|^{-1} &\geq \sum_{S=\{v_1, \dots, v_k\} \in \mathcal{Q}: v_1, \dots, v_k \in V_{\text{small}}} \prod_{i \in [k]} |L_{v_i}|^{-1} \\ &\geq \sum_{S=\{v_1, \dots, v_k\} \in \mathcal{Q}: v_1, \dots, v_k \in V_{\text{small}}} (20\alpha k)^{-k}. \end{aligned}$$

Recall that we have earlier bound $|V_{\text{small}}|$ to be at least $(1 - \frac{1}{10k})|V|$. Hence, the fraction of constraints that involves some variable outside of V_{small} is at most $(\frac{1}{10k}) \cdot (k) = 0.1$. Plugging this into the above inequality, we get that the expected total number of satisfied constraints is at least

$$0.9|\mathcal{Q}| \cdot (20\alpha k)^{-k} > |\mathcal{Q}| \cdot \nu,$$

where the equality comes from our choice of α . In other words, we have $\text{val}(\mathcal{L}) > \nu$ as desired. \square

3.3.3 Proofs of Theorems 3.1, 3.3 and Lemma 3.4

We now prove Theorem 3.1, by simply applying Theorem 3.9 with appropriate parameters on top of the Moshkovitz-Raz PCP.

Proof of Theorem 3.1. Suppose contrapositively that, for some constant $\tilde{\alpha} \geq 1$ and $\zeta > 0$, we have an $O(2^{(1/\gamma)^{2-\zeta}} 2^{d^{1-\zeta}}) f(\frac{1}{\varepsilon})$ time $\tilde{\alpha}$ -agnostic proper learner for γ -margin halfspaces.

Let $\nu > 0$ be a sufficiently small constant so that the parameter α (when $k = 2$) from Theorem 3.9 is at least $\tilde{\alpha}$. (In particular, we pick $\nu = \frac{1}{C(\tilde{\alpha})^k}$ for some sufficiently large constant C .)

Given an instance \mathcal{L} of ν -Gap-2-CSP, we run the reduction from Theorem 3.9 to produce a distribution \mathcal{D} . We then run the learner on \mathcal{D} with error parameter ε as given by Theorem 3.9 (and with $\delta = 1/3$). Note that the learner runs in $O(2^{(1/\gamma)^{2-\zeta}} 2^{d^{1-\zeta}}) f(\frac{1}{\varepsilon}) = 2^{O(n^{1-\zeta/2})}$ time, and produces a halfspace h . We compute $\text{err}_{0-1}^{\mathcal{D}}(h)$; if it is no more than $\alpha \cdot \kappa + \varepsilon$, then we output YES. Otherwise, output NO.

The algorithm describe above solves ν -Gap-2-CSP (correctly with probability $2/3$) in $2^{O(n^{1-\zeta/2})}$ time, which, by Theorem 3.7, violates (randomized) ETH. \square

Next, we prove Lemma 3.4. The main difference from the above proof is that, since the algorithm works only *for some* margin $\gamma = \gamma(d)$. We will select the dimension d to be as large as possible so that $\gamma(d)$ is still smaller than the margin given by Theorem 3.9. This dimension d will be larger than the dimension given by Theorem 3.9; however, this is not an issue since we can simply “pad” the remaining dimensions by setting the additional coordinates to zeros. This is formalized below.

Proof of Lemma 3.4. Let $\tilde{\alpha} \geq 1$ be any constant. Let $\nu > 0$ be a sufficiently small constant so that the parameter α (when $k = 2$) from Theorem 3.9 is at least $\tilde{\alpha}$. (In particular, we pick $\nu = \frac{1}{C(\tilde{\alpha})^k}$ for some sufficiently large constant C .) Let $\varepsilon_0 = \varepsilon_0(\tilde{\alpha})$ be the parameter ε given by Theorem 3.9.

Suppose contrapositively that, for some $\zeta > 0$, there is an $\tilde{\alpha}$ -agnostic learner \mathcal{A} for $\gamma(\tilde{d})$ -margin halfspaces that runs in time $O(2^{(1/\gamma)^{2-\zeta}}) \text{poly}(\tilde{d})$ for all dimensions \tilde{d} and for some $0 < \varepsilon^* < \varepsilon_0(\alpha)$ and $\gamma(\tilde{d})$ that satisfies

$$\frac{1}{\tilde{d}^{0.5-\zeta}} \leq \gamma(\tilde{d}) \leq \frac{1}{(\log \tilde{d})^{0.5+\zeta}} \quad (4)$$

and

$$\frac{\gamma(\tilde{d} + 1)}{\gamma(\tilde{d})} \geq \zeta. \quad (5)$$

We may assume without loss of generality that $\zeta < 0.1$.

We create an algorithm \mathcal{B} for ν -Gap-2-CSP as follows:

- Given an instance \mathcal{L} of ν -Gap-2-CSP of size n , we first run the reduction from Theorem 3.9 with ν as selected above to produce a distribution \mathcal{D} on $\mathbb{B}_d \times \{\pm 1\}$ (where $d = n + 1$). Let the margin parameter γ be as given in Theorem 3.9; observe that $\gamma = \Omega_\nu(1/\sqrt{n})$.
- Let \tilde{d} be the largest integer so that $\gamma(\tilde{d}) \geq \gamma$. Observe that, from the lower bound in (5), we have $\gamma(d) \geq \frac{1}{d^{0.5-\zeta}}$. Hence, for a sufficiently large d , $\gamma(d)$ is larger than γ (which is $O_\nu(1/\sqrt{d})$). In other words, we have $\tilde{d} \geq d$.
- Create a distribution \mathcal{D}' as follows: for each $(\mathbf{x}, y) \in \text{supp}(\mathcal{D})$, we create a sample (\mathbf{x}', y) in \mathcal{D}' with the same probability and where $\mathbf{x}' \in \mathbb{B}_{\tilde{d}}$ is \mathbf{x} concatenated with 0s in the last $\tilde{d} - d$ coordinates.
- Run the learner \mathcal{A} on \mathcal{D}' with parameter $\gamma(\tilde{d})$ and ε . Suppose that it outputs a halfspace h . We compute $\text{err}_{0-1}^{\mathcal{D}'}(h)$; if this is no more than $\alpha \cdot \kappa + \varepsilon_0(\alpha)$, then output YES. Otherwise, output NO.

It is simple to see that, in the completeness case, we must have $\text{OPT}_{\gamma(\tilde{d})}^{\mathcal{D}'} \leq \text{OPT}_\gamma^{\mathcal{D}'} = \text{OPT}_\gamma^{\mathcal{D}} \leq \kappa$; hence, \mathcal{A} would (with probability 2/3) output a halfspace h with 0-1 error at most $\alpha \cdot \kappa + \varepsilon_0(\alpha)$, and we output YES. On the other hand, in the soundness case, we have $\text{OPT}_{0-1}^{\mathcal{D}'} = \text{OPT}_{0-1}^{\mathcal{D}} > \alpha \cdot \kappa + \varepsilon_0(\tilde{\alpha})$, and we always output NO. Hence, the algorithm is correct with probability 2/3.

Next, to analyze the running time of \mathcal{B} , let us make a couple additional observations. First, from (5), we have

$$\gamma(\tilde{d}) \leq \gamma/\zeta \leq O(1/\sqrt{n}). \quad (6)$$

Furthermore, from the upper bound in (5), we have

$$\tilde{d} \leq 2^{(1/\gamma(\tilde{d}))^{0.5+\zeta}} \leq 2^{O(n^{1+2\zeta})} \leq 2^{O(n^{1-\zeta})}, \quad (7)$$

where the last inequality follows from $\zeta < 0.1$.

As a result, the algorithm runs in time $O(2^{(1/\gamma(d))^{2-\zeta}}) \text{poly}(\tilde{d}) \leq 2^{O(n^{1-\zeta/2})}$, which from Theorem 3.7 would break the (randomized) ETH. \square

Finally, we prove Theorem 3.3, which is again by simply applying Theorem 3.9 to the Dinur-Harsha-Kindler PCP and the Sliding Scale Conjecture.

Proof of Theorem 3.3. By plugging in our reduction from Theorem 3.3 to Theorem 3.8, we get that it is NP-hard to, given a distribution \mathcal{D} , distinguish between $\text{OPT}_\gamma^{\mathcal{D}} \leq \kappa$ or $\text{OPT}_{0-1}^{\mathcal{D}} > \alpha \cdot \kappa + \Omega(\frac{1}{\text{poly}(d)})$, where $\gamma = \frac{1}{d^{\text{poly}(\log \log d)}}$ and $\alpha = d^{1/\text{poly}(\log \log d)} = (1/\gamma)^{1/\text{poly}(\log \log(1/\gamma))}$. In other words, if we have a polynomial time α -agnostic learner for γ -margin halfspaces for this parameter regime, then NP = RP.

Similarly, by plugging in our reduction the Sliding Scale Conjecture, we get that it is NP-hard to, given a distribution \mathcal{D} , distinguish between $\text{OPT}_\gamma^{\mathcal{D}} \leq \kappa$ or $\text{OPT}_{0-1}^{\mathcal{D}} > \alpha \cdot \kappa + \Omega(\frac{1}{\text{poly}(d)})$, where $\gamma = 1/d^{O(1)}$ and $\alpha = d^{\Omega(1)} = (1/\gamma)^{\Omega(1)}$. In other words, if we have a polynomial time α -agnostic learner for γ -margin halfspaces for this parameter regime, then NP = RP. \square

4 Conclusions and Open Problems

This work gives nearly tight upper and lower bounds for the problem of α -agnostic proper learning of halfspaces with a margin, for $\alpha = O(1)$. Our upper and lower bounds for $\alpha = \omega(1)$ are far from tight. Closing this gap is an interesting open problem. Characterizing the fine-grained complexity of the problem for improper learning algorithms remains a challenging open problem.

More broadly, an interesting direction for future work would be to generalize our agnostic learning results to broader classes of geometric functions. Finally, we believe that finding further connections between the problem of agnostic learning with a margin and adversarially robust learning is an intriguing direction to be explored.

References

- [ABL17] P. Awasthi, M. F. Balcan, and P. M. Long. The power of localization for efficiently learning linear separators with noise. *J. ACM*, 63(6):50:1–50:27, 2017.
- [ABSS97] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *J. Comput. Syst. Sci.*, 54(2):317–331, 1997.
- [ALM⁺98] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [AS98] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [BGLR94] M. Bellare, S. Goldwasser, C. Lund, and A. Russell. Efficient probabilistic checkable proofs and applications to approximation. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, page 820, 1994.
- [BGS18] A. Bhattacharyya, S. Ghoshal, and R. Saket. Hardness of learning noisy halfspaces using polynomial thresholds. In *Conference On Learning Theory, COLT 2018*, pages 876–917, 2018.
- [BLPR19] S. Bubeck, Y. T. Lee, E. Price, and I. P. Razenshteyn. Adversarial examples from computational constraints. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, pages 831–840, 2019.
- [BM02] P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- [BS00] S. Ben-David and H. Ulrich Simon. Efficient learning of linear perceptrons. In *Advances in Neural Information Processing Systems (NIPS) 2000*, pages 189–195, 2000.
- [BS12] A. Birnbaum and S. Shalev-Shwartz. Learning halfspaces with the zero-one loss: Time-accuracy tradeoffs. In *Advances in Neural Information Processing Systems 25: NIPS 2012*, pages 935–943, 2012.
- [Cho61] C.K. Chow. On the characterization of threshold functions. In *Proceedings of the Symposium on Switching Circuit Theory and Logical Design (FOCS)*, pages 34–38, 1961.

- [Dan16] A. Daniely. Complexity theoretic limitations on learning halfspaces. In *Proceedings of the 48th Annual Symposium on Theory of Computing, STOC 2016*, pages 105–117, 2016.
- [DDFS14] A. De, I. Diakonikolas, V. Feldman, and R. A. Servedio. Nearly optimal solutions for the chow parameters problem and low-weight approximation of halfspaces. *J. ACM*, 61(2):11:1–11:36, 2014.
- [DF95] R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness II: on completeness for W[1]. *Theor. Comput. Sci.*, 141(1&2):109–131, 1995.
- [DF13] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
- [DHK15] I. Dinur, P. Harsha, and G. Kindler. Polynomially low error PCPs with polyloglog n queries via modular composition. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015*, pages 267–276, 2015.
- [Din07] I. Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007.
- [DK19] I. Diakonikolas and D. M. Kane. Degree- d chow parameters robustly determine degree- d ptf’s (and algorithmic applications). In *Proceedings of the 51st ACM Symposium on Theory of Computing, STOC 2019*, pages 804–815, 2019.
- [DKK⁺16] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Robust estimators in high dimensions without the computational intractability. In *Proceedings of FOCS’16*, pages 655–664, 2016.
- [DKK⁺17] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Being robust (in high dimensions) can be practical. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, pages 999–1008, 2017.
- [DKK⁺18] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Robustly learning a gaussian: Getting optimal error, efficiently. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 2683–2702, 2018.
- [DKK⁺19] I. Diakonikolas, G. Kamath, D. Kane, J. Li, J. Steinhardt, and A. Stewart. Sever: A robust meta-algorithm for stochastic optimization. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, pages 1596–1606, 2019.
- [DKS18] I. Diakonikolas, D. M. Kane, and A. Stewart. Learning geometric concepts with nasty noise. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 1061–1073, 2018.
- [DKS19] I. Diakonikolas, W. Kong, and A. Stewart. Efficient algorithms and lower bounds for robust linear regression. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*, pages 2745–2754, 2019.
- [DL01] L. Devroye and G. Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics, Springer, 2001.

- [DLS14] A. Daniely, N. Linial, and S. Shalev-Shwartz. The complexity of learning halfspaces using generalized linear methods. In *Proceedings of The 27th Conference on Learning Theory, COLT 2014*, pages 244–286, 2014.
- [DNV19] A. Degwekar, P. Nakkiran, and V. Vaikuntanathan. Computational limitations in robust classification and win-win results. In *Conference on Learning Theory, COLT 2019*, pages 994–1028, 2019.
- [DOSW11] I. Diakonikolas, R. O’Donnell, R. Servedio, and Y. Wu. Hardness results for agnostically learning low-degree polynomial threshold functions. In *SODA*, pages 1590–1606, 2011.
- [DS09] I. Diakonikolas and R. Servedio. Improved approximation of linear threshold functions. In *Proc. 24th Annual IEEE Conference on Computational Complexity (CCC)*, pages 161–172, 2009.
- [FGKP06] V. Feldman, P. Gopalan, S. Khot, and A. Ponnuswami. New results for learning noisy parities and halfspaces. In *Proc. FOCS*, pages 563–576, 2006.
- [FS97] Y. Freund and R. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [Gol06] P. Goldberg. A Bound on the Precision Required to Estimate a Boolean Perceptron from its Average Satisfying Assignment. *SIAM Journal on Discrete Mathematics*, 20:328–343, 2006.
- [GR06] V. Guruswami and P. Raghavendra. Hardness of learning halfspaces with noise. In *Proc. 47th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 543–552. IEEE Computer Society, 2006.
- [Hau92] D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100:78–150, 1992.
- [IP01] R. Impagliazzo and R. Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [IPZ01] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
- [JL84] W. Johnson and J. Lindenstrauss. Extensions of Lipschitz mapping into Hilbert space. *Contemporary Mathematics*, 26:189–206, 1984.
- [KKM18] A. R. Klivans, P. K. Kothari, and R. Meka. Efficient algorithms for outlier-robust regression. In *Conference On Learning Theory, COLT 2018*, pages 1420–1430, 2018.
- [KL93] M. J. Kearns and M. Li. Learning in the presence of malicious errors. *SIAM Journal on Computing*, 22(4):807–837, 1993.
- [KLS09] A. Klivans, P. Long, and R. Servedio. Learning halfspaces with malicious noise. To appear in *Proc. 17th Internat. Colloq. on Algorithms, Languages and Programming (ICALP)*, 2009.
- [KS04] A. Klivans and R. Servedio. Learning intersections of halfspaces with a margin. In *Proceedings of the 17th Annual Conference on Learning Theory*, pages 348–362, 2004.

- [KSS94] M. Kearns, R. Schapire, and L. Sellie. Toward Efficient Agnostic Learning. *Machine Learning*, 17(2/3):115–141, 1994.
- [LMS11] D. Lokshtanov, D. Marx, and S. Saurabh. Lower bounds based on the exponential time hypothesis. *Bulletin of the EATCS*, 105:41–72, 2011.
- [LRV16] K. A. Lai, A. B. Rao, and S. Vempala. Agnostic estimation of mean and covariance. In *Proceedings of FOCS’16*, 2016.
- [LS11] P. Long and R. Servedio. Learning large-margin halfspaces with more malicious noise. *NIPS*, 2011.
- [McA03] D. A. McAllester. Simplified PAC-bayesian margin bounds. In *16th Annual Conference on Computational Learning Theory*, pages 203–215, 2003.
- [MHS19] O. Montasser, S. Hanneke, and N. Srebro. VC classes are adversarially robustly learnable, but only improperly. In *Conference on Learning Theory, COLT 2019*, pages 2512–2530, 2019.
- [MR10] D. Moshkovitz and R. Raz. Two-query PCP with subconstant error. *J. ACM*, 57(5):29:1–29:29, 2010.
- [Nak19] P. Nakkiran. Adversarial robustness may be at odds with simplicity. *CoRR*, abs/1901.00532, 2019.
- [OS11] R. O’Donnell and R. Servedio. The Chow Parameters Problem. *SIAM J. on Comput.*, 40(1):165–199, 2011.
- [Ros58] F. Rosenblatt. The Perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65:386–407, 1958.
- [Ser01] R. Servedio. Smooth boosting and learning with malicious noise. In *Proceedings of the Fourteenth Annual Conference on Computational Learning Theory*, pages 473–489, 2001.
- [SSS09] S. Shalev-Shwartz, O. Shamir, and K. Sridharan. Agnostically learning halfspaces with margin errors. In *Technical report, Toyota Technological Institute*, 2009.
- [SSS10] S. Shalev-Shwartz, O. Shamir, and K. Sridharan. Learning kernel-based halfspaces with the zero-one loss. In *The 23rd Conference on Learning Theory, COLT 2010*, pages 441–450, 2010.
- [TTV08] L. Trevisan, M. Tulsiani, and S. Vadhan. Regularity, Boosting and Efficiently Simulating every High Entropy Distribution. Technical Report 103, ECCO, 2008. Conference version in Proc. CCC 2009.
- [Val85] L. Valiant. Learning disjunctions of conjunctions. In *Proc. 9th IJCAI*, pages 560–566, 1985.
- [Vap98] V. Vapnik. *Statistical Learning Theory*. Wiley-Interscience, New York, 1998.