

Distributional PAC-Learning from Nisan’s Natural Proofs

Ari Karchmer   

Boston University, MA, USA

Abstract

Do natural proofs imply efficient learning algorithms? Carmosino et al. (2016) demonstrated that natural proofs of circuit lower bounds for Λ imply efficient algorithms for learning Λ -circuits, but only over *the uniform distribution*, with *membership queries*, and provided $\text{AC}^0[p] \subseteq \Lambda$. We consider whether this implication can be generalized to $\Lambda \not\subseteq \text{AC}^0[p]$, and to learning algorithms which use only random examples and learn over arbitrary example distributions (Valiant’s PAC-learning model).

We first observe that, if, for *any* circuit class Λ , there is an implication from natural proofs for Λ to PAC-learning for Λ , then standard assumptions from lattice-based cryptography do not hold. In particular, we observe that depth-2 majority circuits are a (conditional) counter example to this fully general implication, since Nisan (1993) gave a natural proof, but Klivans and Sherstov (2009) showed hardness of PAC-Learning under lattice-based assumptions. We thus ask: what learning algorithms can we reasonably expect to follow from Nisan’s natural proofs?

Our main result is that all natural proofs arising from a type of communication complexity argument, including Nisan’s, imply PAC-learning algorithms in a new *distributional* variant (i.e., an “average-case” relaxation) of Valiant’s PAC model. Our distributional PAC model is stronger than the average-case prediction model of Blum et al. (1993) and the heuristic PAC model of Nanashima (2021), and has several important properties which make it of independent interest, such as being *boosting-friendly*. The main applications of our result are new distributional PAC-learning algorithms for depth-2 majority circuits, polytopes and DNFs over natural target distributions, as well as the nonexistence of encoded-input weak PRFs that can be evaluated by depth-2 majority circuits.

2012 ACM Subject Classification Theory of computation

Keywords and phrases PAC-learning, average-case complexity, communication complexity, natural proofs

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.68

Related Version *Full Version*: <https://arxiv.org/abs/2310.03641> [19]

Funding The author is supported by the DARPA SIEVE program, Agreement Nos. HR00112020021 and HR00112020023.

Acknowledgements Thanks to Mark Bun, Ran Canetti, Russell Impagliazzo, and Emanuele Viola for thoughtful conversations about this research. Thank you to Mauricio Karchmer for advice on presentational aspects of this paper. Finally, special thanks to Marco Carmosino for helpful comments on a draft of this paper, as well as many discussions pertaining to this research. Part of this research was completed while I was visiting the Simons Institute for the theory of computing.

1 Introduction

Razborov and Rudich [30] introduced the concept of *natural proofs* of circuit lower bounds. Informally, a natural proof of a lower bound for a circuit class Λ encodes an efficient algorithm that can be used to distinguish between the truth tables of *simple* Boolean functions (those with “small” Λ -circuit complexity), and *random* Boolean functions. Razborov and Rudich essentially showed that natural proofs for a circuit class Λ rule out the existence of a cryptographic pseudorandom function (PRF) computable by Λ .



© Ari Karchmer;

licensed under Creative Commons License CC-BY 4.0

15th Innovations in Theoretical Computer Science Conference (ITCS 2024).

Editor: Venkatesan Guruswami; Article No. 68; pp. 68:1–68:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Carmosino et al. [5] strengthened the result of [30] by demonstrating that, provided $\text{AC}^0[p] \subseteq \Lambda$, natural proofs of circuit lower bounds for Λ -circuits of size up to $u(n)$ imply algorithms for *learning* $\text{poly}(n)$ size Λ -circuits with *membership queries* over the *uniform distribution*, in time exponential in $u^{-1}(\text{poly}(n))$.

As a corollary, [5] obtained a state-of-the-art *quasipolynomial* time learning algorithm for $\text{AC}^0[p]$ -circuits (with membership queries, over the uniform distribution), using the natural proofs for subexponential size $\text{AC}^0[p]$ -circuits, for any prime p , of Razborov and Smolensky [29, 35].

Since the result of [5], whether or not there exists a fully general implication from natural proofs to learning algorithms in Valiant’s *original* PAC model [36], even for $\Lambda \not\subseteq \text{AC}^0[p]$, has remained open (see e.g. [12]). In Valiant’s original model, learning algorithms are forced to utilize *random examples*, and learn over *unknown example distributions*.

Question 1. Let Λ be any circuit class. Do natural circuit lower bounds for size $u(n)$ Λ -circuits imply $\exp(u^{-1}(\text{poly}(n)))$ time learning algorithms for $\text{poly}(n)$ size Λ -circuits in Valiant’s PAC model?

Aside theoretical interest in complexity and learning theory, Question 1 is motivated by the prospect of implicitly extending the nonexistence of PRFs in low circuit classes (derived from [29, 35, 30, 5]) to the nonexistence of *weak* PRFs. A weak PRF is a PRF that is only required to be secure if the adversary can inspect uniformly random points, as opposed to to chosen points (see Section 2 of the full version [19] for a formal definition). Weak PRFs suffice for a variety of important cryptographic applications such as symmetric-key encryption (see e.g. [4] for more commentary). Therefore, understanding the minimum complexity needed to evaluate weak PRFs is of significant practical importance.

1.1 Our Contributions

We begin by observing that if the answer to Question 1 is essentially “*yes, for every Λ ,*” then this implies algorithmic breakthroughs for several important and well-studied computational problems. These breakthroughs include a classical polynomial time solution to the unique Shortest Vector Problem (uSVP), and quantum polynomial time algorithms for the Shortest Vector Problem (SVP) and Shortest Independent Vector Problem (SIVP) on lattices.¹ More specifically, we observe that majority-of-threshold circuits ($\text{MAJ} \circ \text{THR}$) cannot realize the implication from natural proofs to polynomial time PAC-learning in Valiant’s model, assuming polynomial time hardness of each of those problems.

► **Theorem 1.** *Suppose that a natural proof against $\text{MAJ} \circ \text{THR}$ -circuits of size $u(n)$ implies that the class of $\text{MAJ} \circ \text{THR}$ -circuits of size $\text{poly}(n)$ is PAC-learnable in Valiant’s model, in time $\exp(u^{-1}(\text{poly}(n)))$. Then, there is a polynomial time classical solution to $\tilde{O}(n^{1.5})$ -uSVP, and polynomial time quantum solutions to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP.*

To argue Theorem 1, we combine two observations. First, natural circuit lower bounds for $\exp(\Omega(n))$ size $\text{MAJ} \circ \text{THR}$ -circuits were proved by Nisan [27]. Second, Klivans and Sherstov [22] showed hardness of polynomial time PAC-learning in Valiant’s model for $\text{MAJ} \circ \text{THR}$, assuming classical hardness of uSVP and quantum hardness SVP and SIVP.

¹ We will not try to discuss the huge literature on lattice problems (and lattice based cryptography). See Section 2 of the full version [19] for a short description of uSVP, SVP, and SIVP, and refer to [31, 32] for more information on complexity of lattice problems.

Taken together, we have both natural proofs against exponential-size $\text{MAJ} \circ \text{THR}$, and hardness of Valiant’s PAC-learning for each. Therefore, we have a natural circuit class that resist an implication between natural lower bounds and Valiant’s PAC-learning. A formal argument is presented in Section 3; to the best of our knowledge, the natural property underlying Nisan’s circuit lower bounds has never been explicitly formalized (until Section 3), though it was acknowledged briefly by Raz [28] and considered implicitly by Viola [37].

Theorem 1 indicates a barrier to a general implication from a natural proof *for any* Λ to a PAC-learning algorithm for Λ in Valiant’s model. Essentially, the natural proofs of [27] confound the hardness result of [22]. In light of this, we shift our focus to the following more specific question.

Question 2. What learning algorithms *are* implied by Nisan’s natural circuit lower bounds?

Answering Question 2 is important if we want to gain understanding of a *general* implication between natural proofs for any class Λ , and some kind of learning algorithms for Λ .

Towards an answer to Question 2, we will focus specifically on the possibility of learning algorithms that utilize only random examples, learn over unknown example distributions, and run in polynomial time. We briefly summarize our contributions towards an answer to Question 2, before digging into the specifics.

- In Section 1.1.1, we present a new learning model called *distributional* PAC-learning, which relaxes Valiant’s model, in order to try to sidestep Theorem 1. The new learning model is like Valiant’s except it essentially removes the requirement of guarantees for the worst-case concept in the class. In Section 2, we illustrate that distributional PAC-learning is independently motivated for both technically and practically oriented reasons.
- In Section 3, we prove Theorem 1. To do so, we give the first (to the best of our knowledge) explicit formalization of Nisan’s natural property for exponential size $\text{MAJ} \circ \text{THR}$ -circuits.
- In Section 4, we prove our main theorem, which discovers a relationship between the computational complexity of distributional PAC-learning and the *communication complexity* of a simple communication game that is associated with a given concept class. This theorem serves as a “technical centerpiece” for extracting distributional PAC-learning algorithms from Nisan’s natural proofs.
- In Section 5, we use the main theorem and Nisan’s natural proof to obtain new distributional PAC-learning algorithms for $\text{MAJ} \circ \text{THR}$ -circuits (similar results for polytopes and DNFs appear in the full version [19]). Additionally, we show how to derive attacks on any weak PRF evaluated by $\text{MAJ} \circ \text{THR}$ -circuits, even when the weak PRF is allowed an encoding procedure for the inputs.

We note that, because the result of [5] only applies to circuit classes that contain $\text{AC}^0[p]$, prior to this work there were no known learning algorithms following directly from Nisan’s natural proofs, in any nontrivial learning model.

Subsequent work

In subsequent work, [18] demonstrated how to construct algorithms from Nisan’s natural proofs in several other learning models. For example, [18] obtained “nontrivial time” agnostic membership query learning algorithms over the uniform distribution and “nontrivial time” distribution-independent membership query algorithms, for classes of sublinear size circuits made up of polynomial threshold function (PTF) gates and SYM^+ gates.

1.1.1 Distributional PAC-Learning

With the goal of obtaining learning algorithms from Nisan’s natural proofs in mind, this paper introduces the distributional PAC-learning model (distPAC-learning). However, the distPAC-learning model is also independently motivated as a relaxation of Valiant’s PAC-learning, which we discuss after defining the model next.

The starting point for the distPAC-learning model is the heuristic PAC-learning (heurPAC-learning) model of Nanashima [26]. Following Nanashima, we define a Boolean concept class by a corresponding *evaluation rule* $\phi = \{\phi_n : \{0, 1\}^* \times \{0, 1\}^n \rightarrow \{-1, 1\}\}_{n \in \mathbb{N}}$. The first input to the evaluation rule is a *binary representation* π_f of a concept f , and the second is an *input* to the concept x . The evaluation rule is defined so that for every $n \in \mathbb{N}$, $\phi_n(\pi_f, x) = f(x)$. An evaluation rule ϕ induces a Boolean concept class $\mathfrak{C} = \{\mathfrak{C}_n\}_{n \in \mathbb{N}}$ defined by

$$\mathfrak{C}_n = \{f(x) := \phi_n(\pi_f, x) : \pi_f \in \{0, 1\}^*\} \quad (1)$$

We refer to \mathfrak{C} as the ϕ -induced concept class.

For a function $s : \mathbb{N} \rightarrow \mathbb{N}$, we say that the ϕ -induced concept class \mathfrak{C} is $s(n)$ -*represented* if, for every $n \in \mathbb{N}$, under the evaluation rule ϕ_n , every $f \in \mathfrak{C}_n$ has a binary representation of length at most $s(n)$. Considering evaluation rules helps for formalizing learning using a distribution over a concept class. We let μ denote a *target distribution* over concepts $f \in \mathfrak{C}_n$, or, equivalently, over binary representations $\pi_f \in \{0, 1\}^{s(n)}$.

What access to the concept algorithms in the distPAC-learning model are allowed? We continue following the heurPAC-learning model (Valiant’s, too), and allow access to only random examples sampled from an *unknown* example distribution ρ over $\{0, 1\}^n$. We denote by $\text{EX}(f, \rho)$ the *example oracle* that returns labelled examples $\langle x, f(x) \rangle$ for $x \sim \rho$.

A distPAC-learning algorithm takes three confidence parameters as input. The accuracy parameter ϵ , the failure parameter δ , and the heuristic parameter η . Essentially, the distPAC-learning model requires that, for a fixed evaluation rule ϕ , there exists some large probability mass of the ϕ -induced concept class \mathfrak{C} , as determined by μ and η , that is learnable in Valiant’s model.

► **Definition 2** (Distributional PAC-learning). *Let ϕ be an evaluation rule, and let the ϕ -induced concept class \mathfrak{C} be $s(n)$ -represented. The pair (\mathfrak{C}, μ) is distributionally PAC-learnable if there exists an algorithm A such that, for any $n \in \mathbb{N}$, $\epsilon, \delta, \eta > 0$,*

$$\Pr_{f \sim \mu} \left[\Pr_A \left[\forall \rho : \Pr_{x \sim \rho} \left[h(x) \neq f(x) : h \leftarrow A^{\text{EX}(f, \rho)}(n, \epsilon, \delta, \eta) \right] \leq \epsilon \right] \geq 1 - \delta \right] \geq 1 - \eta \quad (2)$$

When A runs in time $\text{poly}(n, s(n), \epsilon^{-1}, \delta^{-1}, \eta^{-1})$, we say that (\mathfrak{C}, μ) is efficiently distPAC-learnable.

Distributional PAC-learning is a clear relaxation of Valiant’s PAC-learning, since it no longer requires good learning guarantees for “worst-case” concepts.

DistPAC-learning vs. heurPAC-learning

The essential difference between distPAC-learning and heurPAC-learning is the requirement that there exists a *single* large probability mass of concepts that is learnable with respect to *any* example distribution (see the location of the universal quantification over ρ in equation 2). In heurPAC-learning, the order of quantifiers is different: it is only required that for each example distribution ρ , a large *but possibly different* probability mass of the concept class is learnable. This independently motivates our model for technical and practical reasons: for

example, distPAC-learning allows “black-box” use of *boosting algorithms*. In other words, the equivalence of weak² and strong learning is preserved in our model [33, 9] (see Section 2 for a formal statement on this). This is not necessarily true in heurPAC-learning. In a nutshell, distPAC-learning is stronger than both heurPAC-learning and the Blum et al. average-case prediction model. Therefore distPAC-learning also inherits the well-founded motivation behind the theory of heuristic PAC-learning (see [26]). See Section 1.3 for a continued discussion. Towards further motivating the distPAC-learning model, in Section 2 we give formal statements on useful properties of the distPAC-learning model, including the equivalence between weak and strong learning, an equivalence between hardness of distribution-specific variant of distPAC-learning and the existence of one-way functions, and finally on using the classic technique of Kearns and Valiant [20] for proving hardness of distPAC-learning with respect to *specific* target distributions.

1.1.2 DistPAC-Learning Algorithms from Nisan’s Natural Proofs

We design new learning algorithms in the distributional PAC model. These learning algorithms arise from Nisan’s natural proof technique, which we now describe in more detail.

Nisan’s technique

Nisan [27] used the following communication complexity argument for proving circuit lower bounds against circuits with threshold functions as gates. First, identify a function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, which requires *high* 2-party communication complexity in some model (e.g. randomized, deterministic, distributional, etc.). Then, identify a circuit class \mathcal{C} such that for every $g \in \mathcal{C}$, g is computable by a *low-cost* 2-party communication protocol in that model. Finally, conclude that f requires large \mathcal{C} -circuits (see Section 2 of the full version [19] for essential definitions of 2-party communication complexity in various models, and [24] for further reference). To provide an example, let $f(x, y) = \text{IP2}(x, y) = \sum_{i=1}^n x_i y_i \bmod 2$ be the inner product mod 2 function. It is known that IP2 requires $\Omega(n)$ bits to be transmitted in any randomized communication complexity protocol; therefore, as shown by [27], since $\text{MAJ} \circ \text{THR}$ circuits are computed by randomized communication complexity protocols with cost $O(\log n)$, IP2 must require $\text{MAJ} \circ \text{THR}$ circuits of exponential size. This lower bound remains one of the strongest known – as of now it is still not ruled out that $\text{NEXP} \subseteq \text{THR} \circ \text{THR}$. In fact, proving NEXP is not contained in $\text{THR} \circ \text{THR}$ is considered a “major frontier” in complexity theory [6].

Main Theorem

We now introduce our main theorem, which is used as a “technical centerpiece” for obtaining distPAC-learning algorithms from Nisan’s natural lower bounds for $\text{MAJ} \circ \text{THR}$ -circuits, which are presented after. Roughly speaking, the main theorem presents a relationship between the computational complexity of distPAC-learning, and the communication complexity of a simple *communication game* associated with a given evaluation rule.

For any evaluation rule ϕ and ϕ -induced $s(n)$ -represented concept class \mathcal{C} , we define the associated *communication game* \mathfrak{G} over the product distribution (μ, ρ) , played as follows. A binary representation π_f of a function $f \in \mathcal{C}_n$ is sampled according to μ , and an input x is

² A weak learning algorithm is only required to output a hypothesis that has a predictive advantage only slightly better than a coin toss.

sampled from ρ . Player one is given the binary representation π_f of f , and player two is given the input x . The two parties communicate until they are ready to output a value b , and win the game if $b = \phi_n(\pi_f, x) = f(x)$. We say that ϕ is *evaluated* by a 2-party distributional communication protocol with cost $c(n)$ and bias $\gamma(n)$ over (μ, ρ) , if for every $n \in \mathbb{N}$, the two parties can communicate at most $c(n)$ bits before winning \mathfrak{G} with probability at least $1/2 + \gamma(n)$ over the random inputs drawn from (μ, ρ) .

Now we are ready to state the main theorem.

► **Theorem 3.** *Let ϕ be an evaluation rule. Suppose that, for any product distribution (μ, ρ) , ϕ is evaluated by a 2-party distributional communication protocol with cost $c := c(n)$ and bias $\gamma := \gamma(n)$ over (μ, ρ) . Then, for the ϕ -induced $s(n)$ -represented concept class \mathfrak{C} , and a time $t(n)$ -samplable distribution μ , the pair (\mathfrak{C}, μ) is distributionally PAC-learnable. The learning algorithm runs in time polynomial in $n, s(n), t(n), \epsilon^{-1}, \delta^{-1}, \eta^{-1}, \gamma^{-1}$ and 2^c .*

We give an overview of the proof of this theorem in Section 1.2. We remark that slight improvement of the exponential dependency of c requires a significant breakthrough in computational learning theory. For example, improving the efficiency to $O(2^{c^{1/2}})$ gives a polynomial time distinguishing algorithm for the long-time weak PRF candidate of [2]. We refer to Section 7 of the full version [19] for details.

DistPAC-learning from Nisan's natural proofs

Nisan's technique clearly encodes a randomized communication complexity *upper bound*, which, by an averaging argument, can be converted to a distributional protocol over any distribution (without increased cost or decreased bias). Hence, Theorem 4 directly follows from a combination of Nisan's lower bounds and Theorem 3 since, as indicated in Nisan's lower bounds, every function in $\text{MAJ} \circ \text{THR}$ has a randomized communication protocol of cost $O(\log n)$ and large bias.

► **Theorem 4.** *Let $\phi \in \text{MAJ} \circ \text{THR}$ be any evaluation rule, and let μ be any polynomial time samplable target distribution. For the ϕ -induced $s(n)$ -represented concept class \mathfrak{C} , the pair (\mathfrak{C}, μ) is efficiently distPAC-learnable.*

Theorem 4 considers concept classes by the complexity of their evaluation rule. This is weaker than the learning-theoretic standard of considering the complexity of concepts directly. Any $s(n)$ -represented concept class \mathfrak{C} ϕ -induced by the rule $\phi \in \mathcal{C}$ must satisfy $\mathfrak{C} \subseteq \mathcal{C}$ (assuming \mathcal{C} -circuits can be $\text{poly}(s(n))$ -size). For \mathcal{C} containing a universal function (e.g. P/poly , or NC^1), $\mathcal{C} = \mathfrak{C}$, but not necessarily for lower circuit classes.

By non-black-box inspection of the result of [22], which proves hardness of learning $\text{MAJ} \circ \text{THR}$ in Valiant's PAC model, we find that it actually provides a polynomial time samplable *distribution* over $\text{MAJ} \circ \text{THR}$ -circuits that is hard to learn, even weakly (rather than just worst-case hardness). In other words, it shows a target distribution μ^* such that hardness of $\tilde{O}(n^{1.5})$ -SVP and its variants implies $(\text{MAJ} \circ \text{THR}, \mu^*)$ is not efficiently distPAC-learnable! In Section 2, we actually show that the *entire family* of proof techniques for showing hardness of Valiant's PAC-learning (used here by [22], and due originally to [20]), can be used to prove hardness of distPAC-learning with respect to *specific* polynomial time samplable target distributions.

Therefore, in Theorem 4, considering the complexity of the evaluation rule $\phi \in \text{MAJ} \circ \text{THR}$ and distPAC-learning of the ϕ -induced concept class \mathfrak{C} is likely a needed relaxation, since the target distribution in the theorem can be *any* polynomial time samplable distribution. Hence, in order to get a distPAC-learning algorithm for $\text{MAJ} \circ \text{THR}$ -circuits, we need to restrict the target distribution somehow.

Natural Target Distributions

In light of this, we show that MAJ \circ THR-circuits are distPAC-learnable, with respect to the following more natural families of target distributions. By more natural, we mean that the target distribution is not designed by a cryptographer (as opposed to μ^*). Indeed, this highlights another feature of the distPAC-learning model: efficient learnability of (\mathfrak{C}, μ) for “organic” target distributions μ can coexist with hardness for “inorganic” target distributions like μ^* .

Let $L = (T_1, \dots, T_m)$ be a list of $m := \text{poly}(n)$ linear threshold functions, and let μ be a $\text{poly}(n)$ time samplable distribution over $\{0, 1\}^m$. We define the distribution μ_L over MAJ \circ THR-circuits is sampled as follows. First, sample $\theta \sim \mu$. Then, output the MAJ \circ THR-circuit that is the majority vote over each $T_i \in L$ such that $\theta_i = 1$.

► **Theorem 5.** *Let $L = (T_1, \dots, T_m)$ be any list of $m := \text{poly}(n)$ linear threshold functions, and let μ be any $\text{poly}(n)$ time samplable distribution over $\{0, 1\}^m$. The pair (MAJ \circ THR, μ_L) is efficiently distPAC-learnable.*

Previously, no polynomial time distPAC-learning algorithms were known, for any reasonable type of target distributions over MAJ \circ THR-circuits. An interesting and important feature of our distPAC-learning algorithm is that it does not need to know μ or L to work (see Section 5 for details).

We additionally consider slightly modified target distributions, that correspond to interesting and natural distributions over subclasses of MAJ \circ THR-circuits: polytopes (that is, and-of-thresholds circuits (AND \circ THR)) and DNFs. Note that, in distributional PAC-learning, subclasses are not necessarily distPAC-learnable if their superclass is, since it is possible that the subclass is hard-core, and consisting of functions that are hard for the superclass distPAC-learning algorithm.

For the polytope distribution, let $L = (T_1, \dots, T_m)$ be a list of $m := \text{poly}(n)$ linear threshold functions, and let μ be a $\text{poly}(n)$ time samplable distribution over $\{0, 1\}^m$. The distribution μ_L^\wedge over polytopes is sampled as follows. First, sample $\theta \sim \mu$. Then, output the polytope that is the conjunction of all $T_i \in L$ such that $\theta_i = 1$.

► **Theorem 6.** *Let $L = (T_1, \dots, T_m)$ be any list of $m := \text{poly}(n)$ linear threshold functions, and let μ be any $\text{poly}(n)$ time samplable distribution over $\{0, 1\}^m$. The pair (AND \circ THR, μ_L^\wedge) is efficiently distPAC-learnable.*

For the DNF distribution, let $L = (T_1, \dots, T_m)$ be a list of $m := \text{poly}(n)$ disjunctions on n -bit inputs, and let μ be a $\text{poly}(n)$ time samplable distribution over $\{0, 1\}^m$. The distribution $\mu_L^{\wedge\vee}$ over DNFs is sampled as follows. First, sample $\theta \sim \mu$. Then, output the DNF that is a conjunction of all disjunctions $T_i \in L$ such that $\theta_i = 1$.

► **Theorem 7.** *Let $L = (T_1, \dots, T_m)$ be any list of $m := \text{poly}(n)$ disjunctions on n -bit inputs, and let μ be any $\text{poly}(n)$ time samplable distribution over $\{0, 1\}^m$. The pair (DNF, $\mu_L^{\wedge\vee}$) is efficiently distPAC-learnable.*

Even though distPAC-learning is stronger than heurPAC-learning, Theorem 5, 6 and 7 are formally incomparable to the heurPAC-learning algorithm for $O(\log n)$ -juntas due to [26]. This is for the following reasons. On one hand, Theorem 5, 6 and 7 are stronger because MAJ \circ THR-circuits, polytopes, and DNFs are strictly more powerful than $O(\log n)$ -juntas, and we handle learning over arbitrary example distributions, while [26] only handles the uniform example distribution. However, the confounding variable is that the heurPAC

algorithm works with respect to the uniform distribution over $O(\log n)$ -juntas, while for any L, μ , the target distributions that we learn are not uniform over their support. We thus cannot show that our algorithm is stronger than Nanashima's in a formal sense.

1.1.3 Impossibility of Encoded-Input Weak PRFs

Although we do not obtain any PAC-learning algorithm in Valiant's model from Nisan's natural proofs, we show that distributional PAC-learning is still enough to rule out weak PRFs (which was one of the initial motivations of studying Question 1). In fact, we show that this is true even when the weak PRF is allowed an arbitrary input encoding.

► **Theorem 8.** *There exists no encoded-input weak PRF that is evaluated by a MAJ ◦ THR-circuit.*

Our notion of encoded-input weak PRF is the natural weak analogue of the encoded-input PRF introduced by [3]. Loosely speaking, an encoded-input weak PRF is a PRF that is only required to be secure when the adversary sees random points, where the inputs are taken uniformly at random from a predefined multi-subset of the input space. We refer to Section 5 for details.

1.2 Proof Overview of Theorem 3

We now overview the ideas behind the proof of Theorem 3. The most important tool we use is the 2-party norm of a function, $R_2(f)$, which is defined to be the expected product of a function computed on a list of correlated inputs.

► **Definition 9** (2-party norm). *For $f : (\{0, 1\}^n)^2 \rightarrow \{-1, 1\}$, the 2-party norm of f is defined as*

$$R_2(f) := \mathbb{E}_{x_1^0, x_2^0, x_1^1, x_2^1 \sim U_n} \left[\prod_{\epsilon_1, \epsilon_2 \in \{0, 1\}} f(x_1^{\epsilon_1}, x_2^{\epsilon_2}) \right] \quad (3)$$

Throughout the paper, we use U_n to denote the uniform distribution over $\{0, 1\}^n$. The 2-party norm is a special case of the k -party norm (sometimes called the cube-measure or box-norm), which was introduced by [1] for obtaining lower bounds in k -party Number-on-Forehead communication complexity.

The crucial property about $R_2(f)$ is that, up to parameters, it upper bounds the correlation of f with functions computable by deterministic 2-party communication protocols. We denote by $\Pi[2, c]$ the set of all $f : (\{0, 1\}^n)^2 \rightarrow \{-1, 1\}$ that have deterministic 2-party communication protocols with cost at most c . For a definition of the deterministic 2-party communication model, see Section 2 of the full version [19].

Implicit in all three of [7, 28, 38] (who showed a related theorem in the more general k -party case), is the following bound:

► **Theorem 10** (*The correlation bound* – [7, 28, 38]). *For every function $f : (\{0, 1\}^n)^2 \rightarrow \{-1, 1\}$,*

$$\text{Cor}(f, \Pi[2, c]) = \max_{\pi \in \Pi[2, c]} \left| \mathbb{E}_x [f(x) \cdot \pi(x)] \right| \leq 2^c \cdot R_2(f)^{1/4} \quad (4)$$

for x uniformly distributed over $(\{0, 1\}^n)^2$.

Equation 4 implies that $(2^{-c} \cdot \text{Cor}(f, \Pi[2, c]))^4 \leq R_2(f)$.

The construction of the learning algorithm of Theorem 3 uses the lower bound on $R_2(f)$ to distinguish structure from randomness. In other words, hypothetically consider functions $f : (\{0, 1\}^n)^2 \rightarrow \{-1, 1\}$ such that the quantity $(2^{-c} \cdot \text{Cor}(f, \Pi[2, c]))^4$ is relatively *large* (greater than $1/\text{poly}(n)$, say). Such functions can be distinguished from uniformly random functions, by taking a random sample from the distribution over the value inside the expectation in equation 3. This follows from the fact that $R_2(\psi)$ for a uniformly random function $\psi : (\{0, 1\}^n)^2 \rightarrow \{-1, 1\}$ is bounded from above by a negligible function of n .

Using this idea, we have the following proof outline. First, we can try to prove a “distinguisher-to-predictor” lemma, in the style of [39], in order to obtain a weak randomized predictor for f (a weak predictor requires accuracy of a prediction for an unseen example to be only slightly more accurate than a coin toss). Second, we could apply standard averaging arguments to construct a weak PAC-learning algorithm. Finally, we could apply celebrated boosting results from learning theory [33, 9] to produce a full-blown PAC-learning algorithm.

However, this proof outline remains incomplete. First, the 2-party norm of the function is the expectation of a product of *correlated* inputs, so we have not given any way of using independent random examples. Second, we have said nothing of how to handle arbitrary example distributions (the inputs to f on the right hand side of equation 4 should be uniformly random). We handle both of these problems simultaneously, roughly by thinking of f as the *evaluation rule*, and not the concept itself.

First, let us describe how f should be viewed in more detail. There are two inputs to f , x_1 and x_2 . Without loss of generality, identify x_1 as a random string for sampling the target distribution μ , and identify x_2 as a random string for sampling the example distribution ρ , with $|x_1| = |x_2| = m$. We abuse the notation and let $z = \rho(x_2)$ to denote a point z sampled according to ρ with the random bits x_2 . Similarly, we let g be the function represented by $\pi_g = \mu(x_1)$. Next, fix the evaluation rule ϕ , which is the map that takes as input the concept representation π_g , plus the input z , and outputs $\phi(\pi_g, z) = g(z) = y$. As a function of x_1, x_2 , we can thus write the process of generating a labelled example as $\langle \rho(x_2), \phi(\mu(x_1), \rho(x_2)) \rangle = \langle z, g(z) \rangle = \langle z, y \rangle$. We let $f(x_1, x_2) = \phi(\mu(x_1), \rho(x_2))$. This allows us to write:

$$R_2(f) = \mathbb{E}_{x_1^0, x_2^0, x_1^1, x_2^1} [v(x_1^0, x_2^0, x_1^1, x_2^1)] \text{ for } v(x_1^0, x_2^0, x_1^1, x_2^1) := \prod_{\epsilon_1, \epsilon_2 \in \{0, 1\}} \phi(\mu(x_1^{\epsilon_1}), \rho(x_2^{\epsilon_2}))$$

Now, we describe how we construct a weak randomized predictor which only uses random examples from an arbitrary ρ . At the core, we will use the example oracle to sample a single instance of $v(x_1^0, x_2^0, x_1^1, x_2^1)$, over uniformly random $x_1^0, x_2^0, x_1^1, x_2^1 \in \{0, 1\}^m$. To see the significance of this, observe that by definition $v(x_1^0, x_2^0, x_1^1, x_2^1)$ has expected value $R_2(f)$. Hence, the process of sampling this value distinguishes examples labelled by uniformly random functions from examples labelled by concepts sampled according to μ – as long as μ samples representations of concept that are evaluated by ϕ . This claim is justified because whenever it is possible to win the communication game \mathfrak{G} associated with ϕ with high bias and low communication, Theorem 10 implies that $R_2(f)$ is large. In other words, $R_2(f)$ is guaranteed to be large whenever it is possible to efficiently (probabilistically) communicate the evaluation rule ϕ (because this implies winning \mathfrak{G} with good bias). At this point, we use a simple hybrid argument to construct a randomized prediction algorithm for examples sampled according to ρ .

It remains to verify that the randomized prediction algorithm can actually sample $v(x_1^0, x_2^0, x_1^1, x_2^1)$, using only access to $\text{EX}(g, \rho)$, where g is the concept sampled according to the target distribution μ . To see this, observe that the distribution over $v(x_1^0, x_2^0, x_1^1, x_2^1)$ is

identical to the distribution over $g(z_1)g(z_2)h(z_1)h(z_2)$, for $\langle z_1, g(z_1) \rangle, \langle z_2, g(z_2) \rangle \sim \text{EX}(g, \rho)$, and $h \sim \mu$. The value $h(z_1)h(z_2)$ can be computed because $h(z_1)$ and $h(z_2)$ can be *queried*, since h is sampled *locally* by the algorithm. Therefore, we only need $\text{EX}(g, \rho)$.

We also need to verify that ρ need not be efficiently samplable. To argue this this, we observe that communicating parties participating in \mathfrak{G} have unbounded computational power. This means that, even if ρ is an arbitrary distribution, there is no effect on the communication cost of \mathfrak{G} . Indeed, the process of sampling ρ can be viewed as a local pre-processing step in the protocol for party two. Therefore, $R_2(f)$ does not decrease when ρ is arbitrary.

1.3 Discussion

Distributional PAC-Learning vs. Related Models

As mentioned previously, distPAC-learning strengthens heurPAC-learning. This is due to the stronger quantification over example distributions. The main benefit of this is that it facilitates boosting of weak learning algorithms, which needs worst-case guarantees over the example distribution (see Section 2 for a formal statement). We encourage the reader to visit Section 1.2 of [26], as their points regarding the motivations of heurPAC-learning as a relaxation of Valiant’s PAC model, largely apply to distPAC-learning as well. Additionally, see Section 1.2 of [26] for commentary of the differences with previous “implicit” definitions of average-case learning, such as in [16, 15, 34] also apply to distPAC-learning.

In comparison to the seminal work of [2], distPAC-learning also differs on the order of quantifiers. In the definition of the average-case prediction considered by [2], both the target distribution μ and the example distribution ρ are fixed. This means that there can be a different prediction algorithm, for each pair of μ and ρ . This model is weaker than both the heurPAC-learning and distPAC-learning models.

Other Related Work

Many other relationships between learning theory and communication complexity have been studied. Some notable examples include [23, 25, 11, 17] (also see the references therein). All of these works study relationships between communication complexity and notions of learning complexity, such as sample complexity [23, 17], differentially private sample complexity [11], margin complexity [25], VC dimension [23, 11] and Littlestone dimension [11]. These works are all incomparable to ours, as they do not directly study relationships between communication complexity and the *computational* complexity of learning.

Learning intersections of halfspaces (i.e., ands of linear threshold functions) was considered by [21]. Using Fourier-analytic techniques, [21] showed a polynomial time learning algorithm for any function of a constant number of halfspaces with respect to the uniform distribution over examples. Additionally, [21] gave a quasi-polynomial time algorithm for learning any Boolean function of a *polylogarithmic* number of bounded-weights linear threshold functions, under *any* distribution over examples. Our learning results (Theorem 4, and Theorem 5) are at the moment similar but incomparable; we get polynomial time *distributional* PAC-learning of concepts *evaluated* by majorities of linear threshold functions over *any example distribution*.

Additional Remarks and Future Work

In this work, we began by observing that if Question 1 resolves to “yes, for every Λ ,” then cryptographic assumptions such as quantum polynomial time hardness of $\tilde{O}(n^{1.5})$ -SVP (and its variants) do not hold. To continue our study, we shifted the focus to understanding what

kind of learning algorithms *are* implied, specifically by Nisan’s natural proofs. To this end, we introduced the distPAC-learning model as a relaxation of Valiant’s PAC model. Using Nisan’s natural proofs, we at least obtained novel learning algorithms in the distributional PAC model for $\text{MAJ} \circ \text{THR}$, over natural target distributions, and more generally, any concept class induced by evaluation rules contained in $\text{MAJ} \circ \text{THR}$.

Towards this result, we exploited the specific aspects of Nisan’s lower bound method. Therefore, it remains open whether or not *other* natural proofs imply efficient distPAC-learning algorithms for *other* concept classes, such as $\text{AC}^0[p]$. In particular, the natural proofs for $\text{AC}^0[p]$ of [29, 35] are not affected by Theorem 1, so presumably they could even imply algorithms in Valiant’s model. At present, the difficulty in proving a similar barrier to Theorem 1 for $\text{AC}^0[p]$ is that, while we are able to prove super-polynomial hardness of PAC-learning in Valiant’s model, we do not have any exponentially strong natural proofs. We regard this question as the primary direction for future research. We note that works such as [4] have conjectured weak PRF candidates that can be evaluated by $\text{AC}^0[2]$, with considerable evidence to support *subexponential* security of the candidates. A weak PRF evaluated by $\text{AC}^0[2]$ with subexponential security would preclude any (even quasipolynomial time) distPAC-learning algorithm for concept classes induced by a $\text{AC}^0[2]$ evaluation rule. That being said, we hope that our results shed light on what aspects of natural proofs are useful for learning algorithms that cannot query the concept.

Finally, an interesting direction is to obtain distPAC-learning algorithms for other natural distributions over $\text{MAJ} \circ \text{THR}$ -circuits, polytopes, and DNFs.

2 The Distributional PAC-Learning Model

In this section, we give some useful properties of distPAC-learning, which motivate the definition independently. All proofs are omitted due to space constraints. We refer to the full version [19] for all proofs.

Equivalence Between Weak and Strong DistPAC-Learning

Celebrated results of computational learning theory, indicate that efficient weak and strong PAC-learning are equivalent [33] (in the “filtering” setting, this is shown by e.g. [9]).

Recall that $\mathfrak{C} = \{\mathfrak{C}_n\}_{n \in \mathbb{N}}$ is a Boolean concept class that is $s(n)$ -represented and induced by the evaluation rule $\phi = \{\phi_n\}_{n \in \mathbb{N}}$. We define a weak version of distPAC-learning. The only difference with Definition 2 is that the accuracy of the hypothesis only needs to achieve error below $1/2 - \epsilon$ with high probability.

► **Definition 11** (Weak distPAC-learning). *Let ϕ be an evaluation rule, and let the ϕ -induced concept class \mathfrak{C} be $s(n)$ -represented. The pair (\mathfrak{C}, μ) is weakly distributionally PAC-learnable if there exists an algorithm A such that, for any $n \in \mathbb{N}, \delta, \eta > 0$,*

$$\Pr_{f \sim \mu} \left[\Pr_A \left[\forall \rho : \Pr_{x \sim \rho} \left[h(x) \neq f(x) : h \leftarrow A^{\text{EX}(f, \rho)}(n, \delta, \eta) \right] \leq 1/2 - 1/\text{poly}(n) \right] \geq 1 - \delta \right] \geq 1 - \eta \quad (5)$$

When A runs in time $\text{poly}(n, s(n), \delta^{-1}, \eta^{-1})$, we say that (\mathfrak{C}, μ) is efficiently distPAC-learnable.

We now demonstrate that the boosting results from Valiant’s PAC model carry over to the distPAC-learning model.

► **Theorem 12** (distPAC Boosting). *Let ϕ be an evaluation rule, and let \mathfrak{C} be the ϕ -induced $s(n)$ -represented concept class. For any target distribution μ , (\mathfrak{C}, μ) is efficiently weakly distributionally PAC-learnable if and only if (\mathfrak{C}, μ) is efficiently distributionally PAC-learnable.*

DistPAC-Learning vs HeurPAC-Learning

► **Theorem 13.** *Let ϕ be an evaluation rule, and let \mathfrak{C} be the ϕ -induced $s(n)$ -represented concept class. Let U be the uniform distribution over $\{0, 1\}^{s(n)}$. If (\mathfrak{C}, U) is distPAC-learnable in time t , then \mathfrak{C} is heurPAC-learnable in time t .*

Distribution-Specific Learning

To obtain cryptographic primitives from the hardness of learning, the literature often considers hardness of *distribution-specific* learning (as in [2, 26]). In distribution-specific distPAC-learning, we intentionally fix the example distribution ρ . This limits the robustness of distPAC-learning, and entirely collapses it into Nanashima’s notion of distribution-specific heurPAC-learning.

► **Definition 14** (Distribution-specific distPAC-learning). *Let ϕ be an evaluation rule, and let the ϕ -induced concept class \mathfrak{C} be $s(n)$ -represented. The pair (\mathfrak{C}, μ) is distributionally PAC-learnable over the example distribution ρ if there exists an algorithm A such that, for any $n \in \mathbb{N}, \epsilon, \delta, \eta > 0$,*

$$\Pr_{f \sim \mu} \left[\Pr_A \left[\Pr_{x \sim \rho} \left[h(x) \neq f(x) : h \leftarrow A^{\text{Ex}(f, \rho)}(n, \epsilon, \delta, \eta) \right] \leq \epsilon \right] \geq 1 - \delta \right] \geq 1 - \eta \quad (6)$$

When A runs in time $\text{poly}(n, s(n), \epsilon^{-1}, \delta^{-1}, \eta^{-1})$, we say that (\mathfrak{C}, μ) is efficiently distPAC-learnable over ρ .

Cryptography from Hardness of Distribution-Specific DistPAC-Learning

The collapse of distPAC-learning and heurPAC-learning in the distribution-specific setting is a feature of distPAC-learning: using results of [26] on one-way functions from hardness of distribution-specific heurPAC-learning, we obtain the same in distribution-specific distPAC-learning.

► **Theorem 15** (OWFs from hardness of ρ -specific distPAC-learning). *Suppose that P/poly is not efficiently distPAC-learnable with respect to the uniform distribution over concept representations, on a polynomial time samplable example distribution ρ . Then, there exists an (infinitely-often) one-way function.*

Hardness of DistPAC-Learning from Cryptography

In the other direction, we can use the classic results [13, 14] to easily see that the existence of a one-way function implies that there exists a polynomial time samplable μ such that $(\text{P/poly}, \mu)$ is not efficiently distPAC-learnable on any polynomial time samplable ρ .

Additionally, we now demonstrate that the technique for proving hardness of PAC-learning in Valiant’s model (by Kearns and Valiant [20]), also works for distPAC-learning. The technique gives concrete “hard target distributions” for distPAC-learners. Note, it does not imply that distPAC-learning is hard *for every* target distribution.

We sketch the proof technique of Kearns and Valiant, but direct the reader to [20] for the technicalities and also [22] for a more in-depth overview. The proof technique is built around the idea that learning the decryption function of a public key cryptosystem should be hard.

Indeed, this is true by the following reasoning. For any given public key cryptosystem, an attacker can simulate an example oracle for the decryption function, by sampling a public and private key pair, and then generating encryptions of ones and zeros (with equal probability), using the randomized encryption function. Then, if there is a PAC-learning algorithm (in Valiant’s model) for the decryption function, the attacker can apply it to the dataset collected as described, and then predict the messages of new ciphertexts, which breaks security.

We show that it suffices to have a distPAC-learning algorithm, rather than a PAC-learning algorithm in Valiant’s model, to break security of a public key cryptosystem in this way.

► **Theorem 16.** *Suppose there exists a secure public key encryption scheme $(\text{gen}, \text{enc}, \text{dec})$ that generates n -bit encryptions of single bit messages, where decryption error is $\epsilon := \epsilon(n) \leq 1/n^{\omega(1)}$. Let μ be the distribution over decryption functions $\text{dec}(\text{sk}, \cdot)$ with $\text{sk}, \text{pk} \sim \text{gen}$ and hard-wired. If for every sk , $\text{dec}(\text{sk}, \cdot) \in \Lambda$, then (Λ, μ) is not efficiently distPAC-learnable.*

3 No General Implication from Natural Proofs to PAC-Learning

In this section, we show that for circuit classes that can compute $\text{MAJ} \circ \text{THR}$ -circuits, there can be no general implication from natural proofs to PAC-learning algorithms in Valiant’s model, unless $\tilde{O}(n^{1.5})$ -uSVP has a polynomial time solution, and $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP have polynomial time quantum solutions. We sketch an analogous fact for natural proofs for DNFs, and a natural assumption on the polynomial time hardness of refuting random k -SAT instances. Additionally, we will explain that there exists a special target distribution such that there cannot even be an implication from natural proofs to distPAC-learning algorithms for that distribution, under the same lattice assumptions.

3.1 Nisan’s Natural Proof

We will formally define the natural property that arises from Nisan’s natural proofs. We begin by defining what is a natural property. Let F_n be the set of all Boolean functions on n inputs. Typically, the Boolean functions in the context of circuit complexity are $f : \{0, 1\}^n \rightarrow \{0, 1\}$. We will continue to use $f : \{0, 1\}^n \rightarrow \{-1, 1\}$ to maintain continuity with the previous sections.

► **Definition 17 (Natural Property [30]).** *A Natural Property is a sequence of subsets $Q = \{Q_n\}_{n \in \mathbb{N}}$ of $F = \{F_n\}_{n \in \mathbb{N}}$, if it satisfies the following conditions.*

1. **Constructivity.** *The predicate “is f contained in Q_n ” can be computed in polynomial time.*
2. **Largeness.** $|Q_n| \geq \delta_n \cdot |F_n|$.
3. **Usefulness.** *For any sequence of functions $f_n \in F_n$ (for $n \in \mathbb{N}$), if $f_n \in \Lambda$ then $f_n \notin Q_n$ almost everywhere.*

When Q_n satisfies these conditions we say that it is a natural property for Λ with density δ_n . In general, we fix $\delta_n = 1/2$.³

Now we will define the property that arises from Nisan’s lower bounds for $\text{MAJ} \circ \text{THR}$ -circuits [27], and then demonstrate that it is natural. To the best of our knowledge, the natural property has never been explicitly formalized prior to this paper, but the fact that it

³ See Lemma 2.7 of [5] for an explanation for why this is reasonable.

is natural is essentially credited to the sequence of works [27, 7, 28, 38]; [27] provided the lower bound, and [7, 28, 38] all implicitly proved it was natural. The property is defined by the following algorithm that classifies whether or not a truth table is in the property.

► **Definition 18** (Nisan's Natural Property).

1. **Input.** $T_{f_n} \in \{0, 1\}^{2^n}$ a truth table of a function $f_n \in F_n$.
2. Choose any partition of the n inputs of f_n into two sets A, B with $|A| = \lceil n/2 \rceil$ and $|B| = \lfloor n/2 \rfloor$. Let $x_1 || x_2$ denote the partitioned input according to A, B .
3. Given T_{f_n} , compute

$$\alpha = 2^n \cdot R_2(f_n) = \sum_{\substack{x_1^0, x_1^1 \in \{0, 1\}^{\lceil n/2 \rceil} \\ x_2^0, x_2^1 \in \{0, 1\}^{\lfloor n/2 \rfloor}}} \prod_{b_1, b_2 \in \{0, 1\}} f_n(x_1^{b_1} || x_2^{b_2})$$

4. **Output.** If $|\alpha| \leq 2^{2n/3}$, print 1, otherwise print 0.

Let $\text{MAJ}_s \circ \text{THR}$ denote the class of majority-of-thresholds circuits where the top majority gate has fan-in s .

► **Theorem 19** ([27, 7, 28, 38]). *There is a constant c such that there is a natural property for $\text{MAJ}_{2^{n/c}} \circ \text{THR}$ with density $1/2$.*

To prove this theorem, one uses the result of Nisan [27]. The notation $\text{rII}[2, c, \gamma]$ denotes the set of functions $f : \{0, 1\}^n \rightarrow \{-1, 1\}$ that have 2-party randomized communication complexity, with cost at most c and bias at least γ , for every partition of the inputs to the function. See Section 2 of the full version [19] for formal definitions of communication complexity classes and protocols in this work.

► **Theorem 20** ([27]). $\text{MAJ}_s \circ \text{THR} \subseteq \text{rII}[2, O(\log(s)), O(1/s)]$

Now we prove Theorem 19.

Proof of Theorem 19. We prove each of the three necessary properties individually. First, observe that Definition 18 defines a polynomial time algorithm, in the size of the truth table (this proves the Constructivity property).

For the Largeness property, we will analyze the probability that the algorithm outputs 1 given a random truth table as input.

Consider $\alpha/2^{-n} = R_2(f_n)$. We have that

$$\alpha/2^{-n} = \mathbb{E}_{x_1^0, x_1^1, x_2^0, x_2^1} \left[\prod_{b_1, b_2 \in \{0, 1\}} f_n(x_1^{b_1}, x_2^{b_2}) \right]$$

This expected value over the uniformly random choice of $x_1^0, x_1^1, x_2^0, x_2^1$ is 0 except for when either $x_1^0 = x_2^0$ or $x_1^1 = x_2^1$. By a union bound, we have then that $\alpha/2^{-n} \leq 2^{1-(n+1)/2}$. Therefore, $\alpha \leq 2^{n/2+1}$, and the probability over a random truth table T_{f_n} that $\alpha \leq 2^{2n/3}$ is at least $1/2$. This proves density is at least $1/2$.

Finally, we prove Usefulness. By Theorem 10 and Theorem 20, we know that when T_{f_n} is the truth table of $f_n \in \text{MAJ}_s \circ \text{THR}$, then for $s \leq 2^{n/c}$ and some sufficiently large constant c , we have that $\alpha \geq 2^n / \text{poly}(2^{n/c}) \geq 2^{2n/3} + 1$. Thus, whenever $f_n \in \text{MAJ}_{2^{n/c}} \circ \text{THR}$, the algorithm defining the natural property outputs 0, as desired. ◀

3.2 Proof of Theorem 1

In the last section, we demonstrated that $\text{MAJ}_s \circ \text{THR}$ has a dense natural property for s up to $2^{n/c}$ for some constant c . Now we will use this fact, together with hardness of learning $\text{MAJ} \circ \text{THR}$ in Valiant’s PAC model [22] to prove Theorem 1.

► **Theorem 21** (Theorem 1.3 in [22]). *Assume that $\text{MAJ} \circ \text{THR}$ is PAC-learnable in time $\text{poly}(n)$. Then there is a polynomial time solution to $\tilde{O}(n^{1.5})$ -uSVP, and polynomial time quantum solutions to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP.*

► **Theorem 22** (Theorem 1 restated). *Let $u : \mathbb{N} \rightarrow \mathbb{N}$. Suppose that a natural proof against $\text{MAJ}_{u(n)} \circ \text{THR}$ -circuits (and density $1/2$) implies that $\text{poly}(n)$ size $\text{MAJ} \circ \text{THR}$ -circuits are PAC-learnable in Valiant’s model, in time $\exp(u^{-1}(\text{poly}(n)))$. Then, then there is a polynomial time solution to $\tilde{O}(n^{1.5})$ -uSVP, and polynomial time quantum solutions to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP.*

Proof. By Theorem 19, there is a constant c such that there is a natural property for $\text{MAJ}_{2^{n/c}} \circ \text{THR}$ with density $1/2$. Thus, by the condition of the current theorem, we conclude that the class of $\text{MAJ} \circ \text{THR}$ -circuits is PAC-learnable in Valiant’s model. By Theorem 21, we then obtain a polynomial time solution to $\tilde{O}(n^{1.5})$ -uSVP, and polynomial time quantum solutions to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP. ◀

Hardness of DistPAC-Learning for Majority of Threshold Circuits for all Target Distributions

► **Corollary 23.** *There exists a polynomial time samplable distribution μ over polynomial size $\text{MAJ} \circ \text{THR}$ -circuits, such that if $(\text{MAJ} \circ \text{THR}, \mu)$ is efficiently distPAC-learnable, then there is a polynomial time solution to $\tilde{O}(n^{1.5})$ -uSVP, and polynomial time quantum solutions to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP.*

Proof. Implicitly from Theorem 21, we know that the decryption function of Regev’s public key cryptosystem can be implemented by a $\text{MAJ} \circ \text{THR}$ -circuit, with negligible decryption error. Therefore, by Theorem 16, the statement follows. ◀

A Similar Barrier via Hardness of Learning DNFs

A similar statement to Theorem 22 can be observed for DNFs using a different hardness of PAC-learning result of [8]. Specifically, [8] prove that PAC-learning DNFs in Valiant’s model implies the ability to refute random k -SAT instances. Feige [10] first introduce hardness assumptions regarding refuting k -SAT instances.

An algorithm refutes random k -SAT instances with $c(n) \geq \Omega(n)$ clauses if on $1 - o(1)$ fraction of the k -SAT formulas with $c(n)$ constraints, it outputs “unsatisfiable”, while whenever it encounters a satisfiable k -SAT formula with $c(n)$ constraints, it outputs “satisfiable”. A random k -SAT instance $I = \{C_1, \dots, C_{c(n)}\}$ is sampled in the way that each clause/constraint C_i is chosen uniformly at random from the set of k -SAT clauses/constraints with n variables.

Thus, it follows directly from the assumption regarding hardness of refuting random k -SAT formulas from [8] that there is no implication from a Natural Proof for exponential size DNFs, to PAC-learning for DNFs. This is because Nisan’s natural proof clearly works for exponential size DNFs, since the set is clearly contained by the set of exponential size majority-of-threshold circuits.

4 Main Result

In this section, we will prove Theorem 3, which obtains distPAC-learning algorithms for concept classes that have low-cost associated communication games. To do so, we will first define the communication game, and then obtain a weak distPAC-learning algorithm (see Definition 11). Finally, we will conclude Theorem 3 using the equivalence between weak and strong distPAC-learning (see Theorem 12).

In the following sections, we will use Theorem 3 to derive distPAC-learning algorithms for natural distributions over MAJ \circ THR-circuits, polytopes, and DNFs, and impossibility results for weak PRFs that can be evaluated with MAJ \circ THR-circuits with an arbitrary input encoding.

4.1 Communication Games

Recall that $\mathfrak{C} = \{\mathfrak{C}_n\}_{n \in \mathbb{N}}$ is the $s(n)$ -represented Boolean concept class that is induced by the evaluation function $\phi = \{\phi_n\}_{n \in \mathbb{N}}$. We define the communication game associated with \mathfrak{C} :

► **Definition 24** (2-party distributional communication game). *With respect to a evaluation rule ϕ and product distribution (μ, ρ) , the 2-party communication game $\mathfrak{G}[\phi, n, (\mu, \rho)]$ is the following:*

- *Setup: $\pi_f \sim \mu, x \sim \rho$.*
- *Player 1 gets as input $\pi_f \in \{0, 1\}^{s(n)}$ for concept $f \in \mathfrak{C}_n$.*
- *Player 2 gets as input a string $x \in \{0, 1\}^n$.*
- *The object of the game is for the parties to output the value $\phi_n(\pi_f, x) = f(x)$, using as few bits of communication as possible.*

We say that $\mathfrak{G}[\phi, n, (\mu, \rho)]$ is (c, γ) -evaluated if the parties can communicate at most c bits, and win the game with probability $1/2 + \gamma$ (over the random sample of inputs according to (μ, ρ)).

Other definitions

We direct the reader to Section 2 of the full version [19] for the necessary definitions of communication complexity.

4.2 Weak Learning

Towards Theorem 3, we will start by first obtaining a weak learning algorithm, which only requires prediction accuracy marginally better than a coin toss.

Notation

In the following, we discuss boolean functions $f : \{0, 1\}^n \rightarrow \{-1, 1\}$, and denote by U_n the uniform distribution over $\{0, 1\}^n$. For shorthand, we will write $c := c(n), \gamma := \gamma(n)$, to denote number of bits of communication and protocol bias, which are dependent on n , the input length of a concept. Also, in the rest of the paper we will streamline notation by eliding the subscripts on distributions coming from ensembles indexed by $n \in \mathbb{N}$.

► **Theorem 25.** *Let ϕ be an evaluation rule. Suppose that, for every $n \in \mathbb{N}$, and product distribution (μ, ρ) , $\mathfrak{G}[\phi, n, (\mu, \rho)]$ is (c, γ) -evaluated. Then there exists an algorithm A such that, for any $n \in \mathbb{N}, \delta, \eta > 0$,*

$$\Pr_{f \sim \mu} \left[\Pr_A \left[\forall \rho : \Pr_{x \sim \rho} [h(x) \neq f(x) : h \leftarrow A^{Ex(f, \rho)}(n, \delta, \eta)] \leq \frac{1}{2} - \text{poly}(\gamma \cdot 2^{-c}) \right] \geq 1 - \delta \right] \geq 1 - \eta$$

For μ samplable in time $t(n)$, A runs in time $\text{poly}(n, t(n), s(n), \gamma^{-1}, \delta^{-1}, 2^c)$.

Proof. To construct A , we will follow three steps:

1. Construct a weak randomized predictor L .
2. Argue that many good non-uniform but deterministic predictors exist, by fixing coins and samples for L .
3. Construct a deterministic predictor by sampling and then testing enough non-uniform predictors.

Steps 2 and 3 follow from standard techniques (i.e., “constructive averaging”), and are omitted from this version due to space constraints. The full version [19] contains steps 2 and 3. ◀

■ **Algorithm 1** Randomized predictor $L^{\text{EX}(f,\rho)}$.

-
- 1: **Input:** $z \in \{0, 1\}^n, n \in \mathbb{N}, \delta, \eta > 0$
 - 2: Pick uniformly random values $b_1, b_2 \in \{0, 1\}$.
 - 3: Pick uniformly random values $r_{00}, r_{01}, r_{10}, r_{11} \in \{-1, 1\}$
 - 4: Sample $\pi_g \sim \mu$.
 - 5: Sample $(w, y) \sim \text{EX}(f, \rho)$.
 - 6: **if** $b_1 = b_2 = 0$ **then**
 - 7: $v \leftarrow \prod_{i,j \in \{0,1\}} r_{ij}$
 - 8: **if** $b_1 = 0, b_2 = 1$ **then**
 - 9: $v \leftarrow y \cdot r_{00} \cdot \prod_{i,j \in \{0,1\}} r_{ij}$
 - 10: **if** $b_1 = 1, b_2 = 0$ **then**
 - 11: $v \leftarrow \phi(\pi_g, w) \cdot \phi(\pi_g, z) \cdot \prod_{j \in \{0,1\}} r_{1j}$
 - 12: **if** $b_1 = 1, b_2 = 1$ **then**
 - 13: $v \leftarrow y \cdot \phi(\pi_g, w) \cdot \phi(\pi_g, z) \cdot r_{11}$
 - 14: $b \leftarrow r_{b_1 b_2}$
 - 15: **Output** $b \cdot v$
-

▷ **Claim 26 (Weak randomized predictor).** Let ϕ be an evaluation rule. Under the conditions of Theorem 25, there exists a randomized algorithm L , running in time

$$\text{poly}(n, t(n), s(n), \gamma^{-1}, \delta^{-1}, \eta^{-1}, 2^c)$$

such that for any $n \in \mathbb{N}, \delta, \eta > 0$, the following equation is satisfied:

$$\Pr_{f \sim \mu} \left[\Pr_L \left[\forall \rho : \Pr_{z \sim \rho} \left[L^{\text{EX}(f,\rho)}(z, n, \delta, \eta) \neq f(z) \right] \leq \frac{1}{2} - \text{poly}(\gamma \cdot 2^{-c}) \right] \geq 1 - \delta \right] \geq 1 - \eta \quad (7)$$

Proof of Claim 26. We will abuse notation and write $\pi_g \sim \mu$ to denote the binary representation of a concept g , distributed appropriately according to the target distribution μ over the ϕ -induced $s(n)$ -represented concept class \mathfrak{C} . See the randomized predictor L as Algorithm 1.

Consider the distribution \mathcal{M} over 2×2 matrices

$$C = \begin{bmatrix} \phi(\pi_f, z) & \phi(\pi_f, w) \\ \phi(\pi_g, z) & \phi(\pi_g, w) \end{bmatrix}$$

where $\pi_f, \pi_g \sim \mu$ and $z, w \sim \rho$. We now claim that, under the conditions of Theorem 25, this distribution is efficiently *distinguishable* from the distribution \mathcal{R} over random 2×2 matrices,

$$R = \begin{bmatrix} r_{00} & r_{01} \\ r_{10} & r_{11} \end{bmatrix}$$

$$\begin{aligned}
 H_1 &= \mathcal{R} \\
 H_2 &= \begin{cases} C_{k\ell} & \text{when } k = 0, \ell = 0 \\ R_{k\ell} & \text{otherwise} \end{cases} \\
 H_3 &= \begin{cases} C_{k\ell} & \text{when } k = 0, \ell \leq 1 \\ R_{k\ell} & \text{otherwise} \end{cases} \\
 H_4 &= \begin{cases} C_{k\ell} & \text{when } k = 0 \text{ or } \ell \leq 0 \\ R_{k\ell} & \text{otherwise} \end{cases} \\
 H_5 &= \mathcal{M}
 \end{aligned}$$

■ **Figure 1** Hybrid sequence.

To see this, observe that the distribution over C is identical to the following distribution over 2×2 matrices (abusing notation, $z = \rho(y)$ denotes a point z sampled according to ρ with the random bits y)

$$D = \begin{bmatrix} \phi(\mu(x), \rho(y)) & \phi(\mu(x), \rho(y')) \\ \phi(\mu(x'), \rho(y)) & \phi(\mu(x'), \rho(y')) \end{bmatrix}$$

Here, x, x', y, y' are uniformly random strings. We can assume that without loss of generality that x, x', y, y' are all the same length, by defaulting to the maximum necessary length for sampling μ, ρ (and padding the shorter strings with useless random bits). Therefore, identifying $\xi(x, y) := \phi(\mu(x), \rho(y))$, we can now see that

$$R_2(\xi) = \mathbb{E}_{\substack{\pi_f, \pi_g \\ z, w}} \left[\prod_{i, j \in \{0, 1\}} C_{ij} \right]$$

It now readily follows that when $\mathfrak{G}[\phi, n, (\mu, \rho)]$ is (c, γ) -evaluated (which is true by assumption), then

$$R_2(\xi) = \mathbb{E}_{\substack{\pi_f, \pi_g \\ z, w}} \left[\prod_{i, j \in \{0, 1\}} C_{ij} \right] \geq (\gamma \cdot 2^{-c})^4$$

On the other hand,

$$\mathbb{E}_R \left[\prod_{i, j \in \{0, 1\}} R_{ij} \right] = 0$$

Now that we have established this, we may proceed by a hybrid argument. Define the neighboring hybrid distributions H_1, H_2, H_3, H_4, H_5 over 2×2 matrices, as in Figure 1.

It then follows that for random hybrid neighbors H_i, H_{i+1} ($i \in [4]$),

$$\mathbb{E}_{i \sim [4]} \left[\mathbb{E}_{H' \sim H_{i+1}} \left[\prod_{k, j \in \{0, 1\}} H'_{kj} = 1 \right] - \mathbb{E}_{H \sim H_i} \left[\prod_{k, j \in \{0, 1\}} H_{kj} = 1 \right] \right] \geq (\gamma \cdot 2^{-c})^4 / 4 \quad (8)$$

To ease notation, let $D(H) = \prod_{k,j \in \{0,1\}} H_{kj}$, and let V_i denote the event that $D(H_i) = 1$. Intuitively, the function D stands for “distinguisher,” and can be thought of as such.

We continue by observing that, by definition, the value stored as v in L (Algorithm 1) is $D(H_i)$ for a random $i \in [4]$. Hence, the output of L , which is written as $D(H_i) \cdot b$, is interpreted as a prediction, where $b = r_{b_1 b_2}$ is the “guess bit.” Note that, the string $b_1 b_2$ is the binary representation of i .

Now, conditioning on correctness of this guess bit, we have that for all ρ , and probabilities taken over $z \sim \rho, f \sim \mu$ and the randomness of L :

$$\begin{aligned} \Pr[L^{\text{Ex}(f,\rho)}(z, n, \delta, \eta) = f(z)] &= \Pr[L^{\text{Ex}(f,\rho)}(z, n, \delta, \eta) = f(z) \mid b = f(z)] \cdot \Pr[b = f(z)] \\ &\quad + \Pr[L^{\text{Ex}(f,\rho)}(z, n, \delta, \eta) = f(z) \mid b \neq f(z)] \cdot \Pr[b \neq f(z)] \\ &= \frac{1}{2} \left(\Pr[b \cdot D(H_i) = f(z) \mid b = f(z)] \right. \\ &\quad \left. + \Pr[b \cdot D(H_i) = f(z) \mid b \neq f(z)] \right) \end{aligned}$$

Indeed, when V_i is unsatisfied, this means that the output of L is b . The case analysis follows:

$$\begin{aligned} \Pr[L^{\text{Ex}(f,\rho)}(z, n, \delta, \eta) = f(z)] &= \frac{1}{2} \left(\Pr[V_i \mid b = f(z)] + \Pr[\neg V_i \mid b \neq f(z)] \right) \\ &= \frac{1}{2} + \frac{1}{2} \left(\Pr[V_i \mid b = f(z)] - \Pr[V_i \mid b \neq f(z)] \right) \end{aligned}$$

By conditioning, we know that:

$$\Pr[V_i] = \frac{1}{2} \Pr[V_i \mid b = f(z)] + \frac{1}{2} \Pr[V_i \mid b \neq f(z)]$$

rearranging the terms, we get:

$$\frac{1}{2} \Pr[V_i \mid b \neq f(z)] = \Pr[V_i] - \frac{1}{2} \Pr[V_i \mid b = f(z)]$$

We thus conclude:

$$\Pr[L^{\text{Ex}(f,\rho)}(z, n, \delta, \eta) = f(z)] = \frac{1}{2} + \underbrace{\Pr[V_i \mid b = f(z)]}_{(\alpha)} - \underbrace{\Pr[V_i]}_{(\beta)}$$

The term (α) corresponds to the case that L computes the 2-party norm on a sample from H_{i+i} (i.e., the product of the entries of a matrix sampled from H_{i+i}), while term (β) is the case that L computes the 2-party norm on a sample from H_i (the product of the entries of a matrix sampled from H_i). Thus, by equation (8),

$$\begin{aligned} \Pr[L^{\text{Ex}(f,\rho)}(z, n, \delta, \eta) = f(z)] &= \frac{1}{2} + (1 - \Pr[D(H_{i+1}) = 1]) - (1 - \Pr[D(H_i) = 1]) \\ &\geq \frac{1}{2} + (\gamma \cdot 2^{-c})^4 / 8 \quad \triangleleft \end{aligned}$$

Establishing Claim 26 suffices for step 1. We refer the reader to the full version [19] for steps 2 and 3.

► **Remark.** Using the 2-party norm as we do is a *universal* distinguisher. That is, it distinguishes any evaluation rule that is (c, γ) -evaluated from a random function (using the lower bound of $(\gamma \cdot 2^{-c})^4$). Therefore it holds that for *arbitrary* choice of distribution ρ , we obtain the desired guarantee. Indeed, the choice of ρ can be adversarial with respect to $f \sim \mu$, and it never needs to be known by A .

4.3 Main Theorem

Theorem 25 is enough to prove Theorem 3. Recall that $\mathfrak{C} = \{\mathfrak{C}_n\}_{n \in \mathbb{N}}$ is a boolean concept class that is $s(n)$ -representable by the evaluation function $\phi = \{\phi_n\}_{n \in \mathbb{N}}$.

► **Theorem 27** (Theorem 3, restated). *Let ϕ be an evaluation rule, and suppose that for every $n \in \mathbb{N}$, and product distribution (μ, ρ) , $\mathfrak{G}[\phi, n, (\mu, \rho)]$ is (c, γ) -evaluated. Then, (\mathfrak{C}, μ) is distPAC-learnable. For a time $t(n)$ -samplable μ and $s(n)$ -represented \mathfrak{C} , the learning algorithm runs in time polynomial in $n, t(n), s(n), \gamma^{-1}, \epsilon^{-1}, \delta^{-1}, \eta^{-1}$, and 2^c .*

Proof. Immediate from Theorem 12 and Theorem 25. ◀

► **Remark.** The exponential dependency of c is possibly necessary (i.e., necessary assuming exponentially secure one-way functions exist), since the theorem does not restrict \mathfrak{C} (e.g., it can be P/poly), and no matter what $c \leq n$. Also, see the full version [19] for further discussion.

► **Theorem 28** (distPAC-learning – Theorem 4 restated). *Let $\phi \in \text{MAJ} \circ \text{THR}$ be any evaluation rule, and let μ be any polynomial time samplable target distribution. Then, for the ϕ -induced $s(n)$ -represented concept class \mathfrak{C} , the pair (\mathfrak{C}, μ) is efficiently distPAC-learnable.*

Proof of Theorem 28. By Theorem 20, $\text{MAJ} \circ \text{THR}$ has a randomized communication protocol with cost $O(\log n)$ and bias at least $1/\text{poly}(n)$. This implies that for any product distribution (μ, ρ) it is the case that $\mathfrak{G}[\phi, n, (\mu, \rho)]$ is $(O(\log n), 1/\text{poly}(n))$ -evaluated, since $\phi \in \text{MAJ} \circ \text{THR}$. This is enough to conclude the theorem by Theorem 27. ◀

5 DistPAC-Learning on Natural Target Distributions and the Impossibility of Encoded-Input Weak PRFs

Using Theorem 28 we can derive distributional PAC-learning algorithms for interesting and natural distributions over majority-of-threshold circuits, as well as intersections of thresholds, and DNFs. We state the result for distributional PAC-learning of majority-of-threshold circuits below, but refer to the full version [19] for the proof, and the results on intersections of thresholds, and DNFs.

A target distribution over $\text{MAJ} \circ \text{THR}$

Let $L = (T_1, \dots, T_m)$ be a list of $m := \text{poly}(n)$ linear threshold functions. Also, let μ be any $\text{poly}(n)$ time samplable distribution over $\{0, 1\}^m$. We define the distribution μ_L over $\text{MAJ} \circ \text{THR}$ -circuits as follows.

- Sample $\theta \sim \mu$.
 - Output the $\text{MAJ} \circ \text{THR}$ -circuit that is the majority vote over each $T_i \in L$ such that $\theta_i = 0$.
- We will show that, for any μ, L as described above, $(\text{MAJ} \circ \text{THR}, \mu_L)$ is efficiently distPAC-learnable.

► **Theorem 29.** *Let $L = (T_1, \dots, T_m)$ be any list of $m := \text{poly}(n)$ linear threshold functions over n -bit inputs, and let μ be any $\text{poly}(n)$ time samplable distribution over $\{0, 1\}^m$. The pair $(\text{MAJ} \circ \text{THR}, \mu_L)$ is efficiently distPAC-learnable.*

Non-black-box access to μ_L

The distPAC-learning algorithm does not need descriptions of μ or L as input. Instead, it suffices to have black-box access to a sampling machine that outputs appropriately distributed polynomial size circuits that are functionally equivalent to the appropriate MAJ \circ THR-circuit, under some fixed encoding scheme. To see this, observe that the randomized predictor L (Algorithm 1) uses the evaluation rule ϕ and a concept representation π_g to basically obtain some labels of g . But it is not important what the representation is, or how ϕ is implemented. For the accuracy of the learning algorithm, it only matters that, it is *possible* to implement ϕ by a MAJ \circ THR circuit (under some concept representation). Hence, the learner can work with some far messier representation of concepts, such as by arbitrary polynomial size circuits.

The fact that the learner does not need to know μ or L is a good property that enhances the convenience of the learning algorithm. Arguably, knowledge of μ and L would be a prohibitive assumption in practice; rather, black-box sample access to a more complex and messy form of μ_L (provided that the learner can still interpret the encoding well enough to evaluate it) is a significantly weaker assumption, corresponding to a scenario where the learner has some non-explicit knowledge about the possible concepts it may encounter.

Impossibility of Encoded-Input Weak PRFs

We can additionally apply Theorem 27 to prove that weak PRFs augmented with a keyless encoding procedure cannot exist, when evaluating the encoded input on any key is forced to be done by a MAJ \circ THR-circuit. More specifically, we show that *encoded-input weak PRFs* cannot be evaluated by MAJ \circ THR. Our definition of *encoded-input weak PRFs*, omitted here and written in the full version [19], is the natural relaxation of the definition of encoded-input *strong* PRFs of Boneh et al. [3].

The proof of our result is a straightforward application of Theorem 27, and appears in the full version [19].

► **Theorem 30** (Theorem 8 restated). *There exists no encoded-input weak PRF that is evaluated by a MAJ \circ THR-circuit.*

References

- 1 László Babai, Noam Nisan, and Mária Szegedy. Multipart protocols, pseudorandom generators for logspace, and time-space trade-offs. *Journal of Computer and System Sciences*, 45(2):204–232, 1992.
- 2 Avrim Blum, Merrick Furst, Michael Kearns, and Richard J Lipton. Cryptographic primitives based on hard learning problems. In *Annual International Cryptology Conference*, pages 278–291. Springer, 1993.
- 3 Dan Boneh, Yuval Ishai, Alain Passelègue, Amit Sahai, and David J Wu. Exploring crypto dark matter: New simple prf candidates and their applications. In *Theory of Cryptography: 16th International Conference, TCC 2018, Panaji, India, November 11–14, 2018, Proceedings, Part II*, pages 699–729. Springer, 2018.
- 4 Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Low-complexity weak pseudorandom functions in $ac^0[\text{mod}2]$. In *Annual International Cryptology Conference*, pages 487–516. Springer, 2021.
- 5 Marco L Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. Learning algorithms from natural proofs. In *31st Conference on Computational Complexity (CCC 2016)*. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2016.

- 6 Lijie Chen. Toward super-polynomial size lower bounds for depth-two threshold circuits. *arXiv preprint*, 2018. [arXiv:1805.10698](#).
- 7 Fan RK Chung and Prasad Tetali. Communication complexity and quasi randomness. *SIAM Journal on Discrete Mathematics*, 6(1):110–123, 1993.
- 8 Amit Daniely and Shai Shalev-Shwartz. Complexity theoretic limitations on learning dnf’s. In *Conference on Learning Theory*, pages 815–830. PMLR, 2016.
- 9 Carlos Domingo, Osamu Watanabe, et al. Madaboost: A modification of adaboost. In *COLT*, pages 180–189, 2000.
- 10 Uriel Feige. Relations between average case complexity and approximation complexity. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 534–543, 2002.
- 11 Vitaly Feldman and David Xiao. Sample complexity bounds on differentially private learning via communication complexity. In *Conference on Learning Theory*, pages 1000–1019. PMLR, 2014.
- 12 Halley Goldberg and Valentine Kabanets. Improved learning from kolmogorov complexity. *ECCC Report*, 2023.
- 13 Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.
- 14 Johan Håstad, Russell Impagliazzo, Leonid A Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- 15 Jeffrey C Jackson, Homin K Lee, Rocco A Servedio, and Andrew Wan. Learning random monotone dnf. *Discrete Applied Mathematics*, 159(5):259–271, 2011.
- 16 Jeffrey C Jackson and Rocco A Servedio. Learning random log-depth decision trees under uniform distribution. *SIAM Journal on Computing*, 34(5):1107–1128, 2005.
- 17 Daniel Kane, Roi Livni, Shay Moran, and Amir Yehudayoff. On communication complexity of classification problems. In *Conference on Learning Theory*, pages 1903–1943. PMLR, 2019.
- 18 Ari Karchmer. Agnostic membership query learning with nontrivial savings: New results, techniques. *arXiv preprint*, 2023. [arXiv:2311.06690](#).
- 19 Ari Karchmer. Distributional pac-learning from nisan’s natural proofs, 2023. [arXiv:2310.03641](#).
- 20 Michael Kearns and Leslie Valiant. Cryptographic limitations on learning boolean formulae and finite automata. *Journal of the ACM (JACM)*, 41(1):67–95, 1994.
- 21 Adam R Klivans, Ryan O’Donnell, and Rocco A Servedio. Learning intersections and thresholds of halfspaces. *Journal of Computer and System Sciences*, 68(4):808–840, 2004.
- 22 Adam R Klivans and Alexander A Sherstov. Cryptographic hardness for learning intersections of halfspaces. *Journal of Computer and System Sciences*, 75(1):2–12, 2009.
- 23 Ilan Kremer, Noam Nisan, and Dana Ron. On randomized one-round communication complexity. *Computational Complexity*, 8:21–49, 1999.
- 24 Eyal Kushilevitz and Noam Nisan. Communication complexity, 1996.
- 25 Nati Linial and Adi Shraibman. Learning complexity vs communication complexity. *Combinatorics, Probability and Computing*, 18(1-2):227–245, 2009.
- 26 Mikito Nanashima. A theory of heuristic learnability. In *Conference on Learning Theory*, pages 3483–3525. PMLR, 2021.
- 27 Noam Nisan. The communication complexity of threshold gates. *Combinatorics, Paul Erdos is Eighty*, 1:301–315, 1993.
- 28 Ran Raz. The bns-chung criterion for multi-party communication complexity. *Computational Complexity*, 9(2):113–122, 2000.
- 29 Alexander A Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.
- 30 Alexander A Razborov and Steven Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55(1):24–35, 1997.

- 31 Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009.
- 32 Oded Regev. On the complexity of lattice problems with polynomial approximation factors. In *The LLL Algorithm: Survey and Applications*, pages 475–496. Springer, 2009.
- 33 Robert E. Schapire. The strength of weak learnability. *Mach. Learn.*, 5:197–227, 1990. doi:10.1007/BF00116037.
- 34 Linda Sellie. Exact learning of random dnf over the uniform distribution. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 45–54, 2009.
- 35 Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 77–82, 1987.
- 36 Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- 37 Emanuele Viola. The communication complexity of addition. *Combinatorica*, 35:703–747, 2015.
- 38 Emanuele Viola and Avi Wigderson. Norms, xor lemmas, and lower bounds for gf(2) polynomials and multiparty protocols. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 141–154. IEEE, 2007.
- 39 Andrew C Yao. Theory and application of trapdoor functions. In *23rd Annual Symposium on Foundations of Computer Science (SFCS 1982)*, pages 80–91. IEEE, 1982.