

Cryptographic Hardness for Learning Intersections of Halfspaces

Adam R. Klivans*

University of Texas at Austin, Department of Computer Sciences, Austin, TX 78712 USA

Alexander A. Sherstov

University of Texas at Austin, Department of Computer Sciences, Austin, TX 78712 USA

Abstract

We give the first representation-independent hardness results for PAC learning intersections of halfspaces, a central concept class in computational learning theory. Our hardness results are derived from two public-key cryptosystems due to Regev, which are based on the worst-case hardness of well-studied lattice problems. Specifically, we prove that a polynomial-time algorithm for PAC learning intersections of n^ε halfspaces (for a constant $\varepsilon > 0$) in n dimensions would yield a polynomial-time solution to $\tilde{O}(n^{1.5})$ -uSVP (unique shortest vector problem). We also prove that PAC learning intersections of n^ε low-weight halfspaces would yield a polynomial-time quantum solution to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP (shortest vector problem and shortest independent vector problem, respectively). Our approach also yields the first representation-independent hardness results for learning polynomial-size depth-2 neural networks and polynomial-size depth-3 arithmetic circuits.

Key words: Cryptographic hardness results, intersections of halfspaces, computational learning theory, lattice-based cryptography

1 Introduction

A halfspace in n dimensions is a Boolean function of the form $a_1x_1 + \dots + a_nx_n \geq \theta$, where a_1, \dots, a_n, θ are integers. Halfspace-based learning methods have important applications in almost every area of computer science, including data mining, artificial intelligence, and computer vision. A natural and important extension of the

* Supported by NSF CAREER Award CCF-0643829.

Email addresses: klivans@cs.utexas.edu (Adam R. Klivans),
sherstov@cs.utexas.edu (Alexander A. Sherstov).

concept class of halfspaces is the concept class of *intersections* of halfspaces. While many efficient algorithms exist for PAC learning a *single* halfspace, the problem of learning the intersection of even two halfspaces remains a central challenge in computational learning theory, and a variety of efficient algorithms have been developed for natural restrictions of the problem [17,18,21,31] (for a definition of the PAC model see Section 2). Attempts to prove that the problem is hard have been met with limited success: all known hardness results for the general problem of PAC learning intersections of halfspaces apply only to the case of *proper* learning, where the output hypothesis must be of the same form as the unknown concept.

1.1 Our Results

We obtain the first *representation-independent* hardness results for PAC learning intersections of halfspaces. By “representation-independent,” we mean that we place no restrictions on the learner’s output hypothesis other than polynomial-time computability. Assuming the intractability of the lattice problems uSVP (unique shortest vector problem), SVP (shortest vector problem), or SIVP (shortest independent vector problem), we prove that there is no polynomial-time PAC learning algorithm for intersections of n^ϵ halfspaces (for any $\epsilon > 0$). The above lattice problems are widely believed to be hard [24].

In this work, we will use cryptosystems based on the hardness of these lattice problems as a black box in order to obtain our hardness-of-learning results. We will therefore not attempt to summarize the vast literature on the complexity of lattice problems and instead refer the reader to several works by Regev [26,27] and Aharonov and Regev [1]. We sketch the lattice problems briefly in Section 2.

Our hardness results apply even to intersections of *light* halfspaces, i.e., halfspaces whose weight $|\theta| + \sum_{i=1}^n |a_i|$ is bounded by a polynomial in n . We first state our hardness results for intersections of *arbitrary* halfspaces. Throughout this paper, “PAC learnable” stands for “learnable in the PAC model in polynomial time.”

Theorem 1.1 *Assume that intersections of n^ϵ halfspaces in n dimensions are PAC-learnable for some constant $\epsilon > 0$. Then there is a polynomial-time solution to $\tilde{O}(n^{1.5})$ -uSVP.*

With a different (incomparable) hardness assumption, we obtain an intractability result for learning intersections of *light* halfspaces, a less powerful concept class:

Theorem 1.2 *Assume that intersections of n^ϵ light halfspaces in n dimensions are PAC-learnable for some constant $\epsilon > 0$. Then there is a polynomial-time quantum solution to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP.*

Oded Regev has informed us that Theorem 1.1 also applies to light halfspaces; see

Remark 5.1 for details.

We note here that we can prove something slightly stronger than what is stated in Theorem 1.2. That is, if intersections of n^ϵ light halfspaces in n dimensions are PAC-learnable, we obtain a polynomial-time solution to the LWE (“Learning With Errors”) problem, a version of the noisy parity learning problem over larger fields (see Regev [24] for details).

These hardness results extend to polynomial-size depth-2 neural networks as follows:

Theorem 1.3 *Assume that depth-2 polynomial-size circuits of majority gates are PAC 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.*

Finally, we prove a hardness result for learning depth-3 arithmetic circuits:

Theorem 1.4 *Assume that depth-3 polynomial-size arithmetic circuits are PAC-learnable in polynomial time. Then there is a polynomial-time quantum solution to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP.*

We are not aware of any previous representation-independent hardness results for learning small-depth arithmetic circuits.

A natural question to ask is whether our approach can yield hardness results for other classes such as AC^0 or, more ambitiously, polynomial-size DNF formulas. In Section 6 we show that the decryption functions of the cryptosystems we use contain PARITY as a subfunction, so we cannot directly apply this approach.

Remark. In a recent work, Feldman *et al.* [7] have independently obtained a result very similar to Theorem 1.3. They show that a polynomial-time algorithm for learning depth-2 polynomial-size majority circuits would break the Ajtai-Dwork cryptosystem. In contrast, our work makes use of more recent cryptosystems due to Regev (the security of Regev’s cryptosystems is based on weaker assumptions than the ones used by Ajtai and Dwork).

1.2 Previous Results

In his fundamental paper on learning, Valiant [30] established a cryptographic hardness result for learning polynomial-size circuits. Kearns and Valiant [12] used number-theoretic problems (inverting the RSA function, deciding quadratic residuosity, and factoring Blum integers) to obtain hardness results for NC^1 cir-

circuits, constant-depth threshold circuits TC^0 , and deterministic finite automata. Kharitonov [15] obtained hardness results for AC^1 and NC^1 circuits based on the conjectured hardness of the subset sum problem. Kharitonov [14] later used the Blum-Blum-Shub pseudorandom generator [6] to obtain a hardness result for learning AC^0 and TC^0 that holds even under the uniform distribution and if membership queries are allowed.

Hardness results of any kind for learning intersections of halfspaces, by contrast, have seen quite limited progress. Until recently, the problem was known to be hard only for *proper* learning: if the learner’s output hypothesis must be from a restricted class of functions (e.g., intersections of halfspaces), then the learning problem is NP-hard with respect to randomized reductions [5,3]. Klivans and Sherstov [19] have since obtained a $2^{\Omega(\sqrt{n})}$ lower bound on the sample complexity of learning intersections of \sqrt{n} halfspaces in the *statistical query* (SQ) model, an important restriction of the PAC model. Since the SQ model is a restriction of PAC, the lower bounds in [19] do not imply hardness in the PAC model, the subject of this paper. We are not aware of any other results on the difficulty of learning intersections of halfspaces.

We are also not aware of any representation-independent hardness results for PAC learning small-depth arithmetic circuits. There is a long line of research establishing lower bounds on the query complexity of polynomial *interpolation* algorithms over various fields, but these do not imply hardness results for the problem of PAC learning polynomials with small representations as arithmetic circuits (see Section 5.1 for more details).

1.3 Our Techniques

Our results exploit recent cryptosystems due to Regev [23,24], which improve on the security of the Ajtai-Dwork cryptosystem [2]. These cryptosystems are based on the hardness of the well-studied lattice problems $uSVP$, SVP , and $SIVP$. As pointed out in [24], an advantage of these problems is the equivalence of their worst-case and average-case complexity. In other words, an efficient algorithm for solving these problems on a nonnegligible (inverse-polynomial) fraction of instances yields an efficient algorithm for solving *every* instance. This contrasts with common number-theoretic problems such as factoring or deciding quadratic residuosity. Furthermore, lattice-based cryptosystems feature decryption functions that are completely different from modular exponentiation $d(Y) = Y^D \bmod N$, the decryption function that is at the heart of virtually every number-theoretic cryptosystem. As a result, lattice-based cryptosystems imply hardness results that number-theoretic cryptosystems have not yielded.

An established method [12] for obtaining hardness results for a concept class \mathcal{C}

is to demonstrate that \mathcal{C} can compute the decryption function of a public-key cryptosystem. Intersections of a polynomial number of halfspaces, however, cannot compute the decryption functions of the cryptosystems that we use. In fact, the decryption functions in question contain PARITY as a subfunction (see Section 6), which cannot be computed by intersections of a polynomial number of any unate functions [19]. Furthermore, the decryption functions for Regev’s cryptosystems perform a division or an iterated addition, which require threshold circuits of depth 3 and 2, respectively [32,29]. Threshold circuits of depth 2 and higher are known to be more powerful than intersections of halfspaces.

To overcome these difficulties, we use non-uniform distributions on $\{0, 1\}^n$ to help us with the computation. This technique allows us to use intersections of *degree-2 polynomial threshold functions* to compute the decryption function while still obtaining a hardness result for intersections of *halfspaces*.

The remainder of this article is organized as follows. Section 2 covers technical preliminaries and provides a detailed overview of the cryptosystems that we use. The crucial connection between learning and cryptography is the subject of Section 3. The main ingredient of our proof is presented in Section 4 and is concerned with the construction of efficient circuits for the decryption functions of the cryptosystems. Section 5 establishes our main results, with further discussion in Section 6.

2 Preliminaries

A *halfspace* in n dimensions is a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ of the form

$$f(x) = \begin{cases} 1 & \text{if } a_1x_1 + a_2x_2 + \cdots + a_nx_n \geq \theta, \\ 0 & \text{otherwise,} \end{cases}$$

where a_1, \dots, a_n, θ are some fixed integers. It is well known that the absolute values of a_1, \dots, a_n, θ can be assumed to be at most $2^{O(n \log n)}$. The intersection of k halfspaces is a Boolean function $g = \bigwedge_{i=1}^k h_i$, where each h_i is a halfspace. A *polynomial threshold function* (PTF) of degree d is a Boolean function of the form

$$f(x) = \begin{cases} 1 & \text{if } p(x) \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$

where p is a degree- d polynomial in x_1, x_2, \dots, x_n with integer coefficients. Note that a halfspace is a PTF of degree 1. The *weight* of a PTF f is the sum of the absolute values of the integer coefficients of the associated polynomial p . A PTF is called *light* if its weight is bounded by a polynomial in n . (Strictly speaking, this definition concerns not a single PTF but rather an infinite sequence $f_1, f_2, \dots, f_n, \dots$ of PTFs, one for each input length. For brevity, however, we will follow the general

convention of identifying a sequence of functions $f_1, f_2, \dots, f_n, \dots$ with its typical n th representative, f_n).

We adopt the *probably approximately correct* (PAC) model of learning, due to Valiant [30]. An overview of this model is as follows. A *concept class* \mathcal{C} is any set of Boolean functions $\{0, 1\}^n \rightarrow \{0, 1\}$. In the PAC model, one fixes an arbitrary *target function* $f \in \mathcal{C}$ and a distribution μ on $\{0, 1\}^n$. The learner, who does not know f or μ , receives labeled examples $(x^1, f(x^1)), (x^2, f(x^2)), \dots$, where $x^1, x^2, \dots \in \{0, 1\}^n$ are chosen independently at random according to μ . The learner is said to learn \mathcal{C} if, given $\varepsilon \in (0, 1)$ and $\text{poly}(n, \frac{1}{\varepsilon})$ labeled examples, it outputs a hypothesis h that with high probability has $\Pr_{x \sim \mu}[f(x) \neq h(x)] < \varepsilon$. We will be using a looser requirement called *weak learning*, which relaxes the success criterion to $\Pr_{x \sim \mu}[f(x) \neq h(x)] < \frac{1}{2} - \frac{1}{n^c}$ for a constant c ; for contrast, the original framework is known as *strong learning*. Throughout this paper, ‘‘PAC learning’’ is a shorthand for PAC learning in polynomial time and under arbitrary distributions μ . In this arbitrary-distribution setting, weak PAC learning is equivalent to strong PAC learning, and we will sometimes not draw the distinction between the two in the development to follow. For further background on computational learning theory, see [13].

2.1 Lattice-based Cryptography

This subsection describes lattice-based cryptography and presents two relevant lattice-based cryptosystems due to Regev [23,24]. A *lattice* in n dimensions is the set $\{a_1 \mathbf{v}_1 + \dots + a_n \mathbf{v}_n : a_1, \dots, a_n \in \mathbb{Z}\}$ of all integral linear combinations of a given basis $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^n$. The primary problems on lattices are the *unique shortest vector problem* $f(n)$ -uSVP, *shortest vector problem* $f(n)$ -SVP, and *shortest independent vector problem* $f(n)$ -SIVP. In $f(n)$ -uSVP, the goal is to find a shortest nonzero vector in the lattice, provided that it is shorter by a factor of at least $f(n)$ than any other non-parallel vector. In $f(n)$ -SVP, the goal is to approximate the length of a shortest nonzero vector within a factor of $f(n)$. Thus, uSVP is a special case of SVP, distinguished by the ‘‘uniqueness’’ condition. Finally, in $f(n)$ -SIVP, the goal is to output a set of n linearly independent lattice vectors of length at most $f(n) \cdot \text{opt}$, where opt is the minimum length over all sets of n linearly independent vectors from the lattice (the length of a set is the length of its longest vector). Note that all three problems become harder as the approximation factor $1 \leq f(n) \leq \text{poly}(n)$ decreases. We will be working with $f(n) = \tilde{O}(n^{1.5})$, an approximation factor for which these three problems are believed to be hard (none of the above lattice problems are known to admit a subexponential time solution for any setting of $f(n)$ we consider in this paper).

We note here that there is a large body of work examining the hardness of these lattice problems depending on the choice of $f(n)$. Roughly speaking, certain variants

of the shortest vector problem are known to be NP-hard if $f(n)$ is chosen to be a small constant. On the other hand, it is known that for larger values of $f(n)$, such as \sqrt{n} , some lattice problems are unlikely to be NP-hard (i.e., if they were NP-hard, the polynomial-time hierarchy would collapse). We refer the reader to the excellent survey by Regev [27] for a detailed description of the hardness of these problems.

The cryptosystems below encrypt one-bit messages (0 and 1). Encryption is randomized; decryption is deterministic. Let $e_{K,r} : \{0, 1\} \rightarrow \{0, 1\}^{\text{poly}(n)}$ denote the encryption function corresponding to a choice of private and public keys $K = (K_{\text{priv}}, K_{\text{pub}})$ and a random string r . In discussing security, we will need the following notion.

Definition 2.1 (Distinguisher) *An algorithm \mathcal{A} is said to distinguish between the encryptions of 0 and 1 if for some universal constant c ,*

$$\left| \Pr_{K,r}[\mathcal{A}(K_{\text{pub}}, e_{K,r}(1)) = 1] - \Pr_{K,r}[\mathcal{A}(K_{\text{pub}}, e_{K,r}(0)) = 1] \right| \geq \frac{1}{n^c}.$$

We focus on those aspects of the cryptosystems that are relevant to the hardness proofs in this paper. For example, we state the numeric ranges of public and private keys without describing the key generation procedure. We follow the established convention of denoting polynomially-bounded quantities (in n) by lowercase letters, and superpolynomial ones by capital letters.

2.2 The uSVP-based Cryptosystem

We start with a cryptosystem, due to Regev [23], whose security is based on the worst-case hardness of uSVP. Let n be the security parameter. Denote $N = 2^{8n^2}$ and $m = cn^2$, where c is a universal constant. Let $\gamma(n)$ be any function with $\gamma(n) = \omega(n\sqrt{\log n})$, where faster-growing functions γ correspond to worse security guarantees but also a lower probability of decryption error.

Private key: A real number H with $\sqrt{N} \leq H < 2\sqrt{N}$.

Public key: A vector (A_1, \dots, A_m, i_0) , where $i_0 \in \{1, \dots, m\}$ and each $A_i \in \{0, \dots, N-1\}$.

Encryption: To encrypt 0, pick a random set $S \subseteq [m]$ and output $\sum_{i \in S} A_i \bmod N$.

To encrypt 1, pick a random set $S \subseteq [m]$ and output $\lfloor A_{i_0}/2 \rfloor + \sum_{i \in S} A_i \bmod N$.

Decryption: On receipt of $W \in \{0, \dots, N-1\}$, decrypt 0 if $\text{frac}(WH/N) < 1/4$, and 1 otherwise. Here $\text{frac}(a) \equiv \min\{\lceil a \rceil - a, a - \lfloor a \rfloor\}$ denotes the distance from $a \in \mathbb{R}$ to the closest integer. By a standard argument, the security and correctness of the cryptosystem are unaffected if we change the decryption function to $\text{frac}(AW) < 1/4$, where A is a representation of H/N to within $\text{poly}(n)$ fractional bits.

Correctness: The probability of decryption error (over the choice of private and public keys and the randomness in the encryption) is $2^{-\Omega(\gamma(n)^2/m)}$.

Regev [23] showed that breaking the above cryptosystem would yield a polynomial-time algorithm for **uSVP**. A more detailed statement follows (see Theorem 4.5 and Lemma 5.4 of [23]):

Theorem 2.2 (Regev [23]) *Assume that there is a polynomial-time distinguisher between the encryptions of 0 and 1. Then there is a polynomial-time solution to every instance of $(\sqrt{n} \cdot \gamma(n))$ -uSVP.*

We will set $\gamma(n) = n \log n$ to make the probability of decryption error negligible (inverse-superpolynomial) while guaranteeing $\tilde{O}(n^{1.5})$ -uSVP security. Regev's cryptosystem thus improves on the public-key cryptosystem of Ajtai and Dwork [2] whose security is based on the worst-case hardness of $O(n^8)$ -uSVP, an easier problem than $\tilde{O}(n^{1.5})$ -uSVP.

2.3 SVP- and SIVP-based Cryptosystem

The second cryptosystem [24] is based on the worst-case *quantum* hardness of **SVP** and **SIVP**. Let n be the security parameter. Denote by p a prime with $n^2 < p < 2n^2$, and let $m = 5(n+1)(1+2\log n)$. Let $\gamma(n)$ be any function with $\gamma(n) = \omega(\sqrt{n} \log n)$, where faster-growing functions γ correspond to worse security guarantees but also a lower probability of decryption error.

Private key: A vector $\mathbf{s} \in \mathbb{Z}_p^n$.

Public key: A sequence of pairs $(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_m, b_m)$, where each $\mathbf{a}_i \in \mathbb{Z}_p^n$ and $b_i \in \mathbb{Z}_p$.

Encryption: To encrypt 0, pick $S \subseteq [m]$ randomly and output $(\sum_{i \in S} \mathbf{a}_i, \sum_{i \in S} b_i)$. To encrypt 1, pick $S \subseteq [m]$ randomly and output $(\sum_{i \in S} \mathbf{a}_i, \lfloor p/2 \rfloor + \sum_{i \in S} b_i)$. (All arithmetic is modulo p .)

Decryption: On receipt of $(\mathbf{a}, b) \in \mathbb{Z}_p^n \times \mathbb{Z}_p$, decrypt 0 if $b - \langle \mathbf{a}, \mathbf{s} \rangle$ is closer to 0 than to $\lfloor p/2 \rfloor$ modulo p . Decrypt 1 otherwise. (All arithmetic is modulo p .)

Correctness: The probability of decryption error (over the choice of private and public keys and the randomness in the encryption) is $2^{-\Omega(\gamma(n)^2/m)}$.

Regev [24] showed that breaking the above cryptosystem would imply a polynomial-time quantum algorithm for solving **SVP** and **SIVP**. A more precise statement is as follows (see Theorem 3.1 and Lemmas 4.4, 5.4 of [24]):

Theorem 2.3 (Regev [24]) *Assume that there is a polynomial-time (possibly quantum) algorithm for distinguishing between the encryptions of 0 and 1. Then there is a polynomial-time quantum solution to $\tilde{O}(n \cdot \gamma(n))$ -SVP and $\tilde{O}(n \cdot \gamma(n))$ -SIVP.*

We adopt the setting $\gamma(n) = \sqrt{n} \log^2 n$ to make the probability of decryption error negligible while guaranteeing $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP security. Observe that this second cryptosystem is preferable to the first in that it is based on the worst-case hardness of a more general lattice problem (SVP vs. uSVP). The disadvantage of the second cryptosystem is that breaking it would only yield a *quantum* algorithm for SVP, as opposed to the first cryptosystem which would yield a classical algorithm for uSVP.

3 Learning Decryption Functions vs. Breaking Cryptosystems

In their seminal paper [12], Kearns and Valiant established a key relationship between the security of a public-key cryptosystem and the hardness of learning an associated concept class. We re-derive it below for completeness and extend it to allow for errors in the decryption process. This link is a natural consequence of the ease of encrypting messages with the public key. A large pool of such encryptions can be viewed as a set of training examples for learning the decryption function. But learning the decryption function to a nonnegligible advantage would mean breaking the cryptosystem. Assuming that the cryptosystem is secure, we can thus conclude that it is not feasible to learn the decryption function. We formalize this observation in the following lemma:

Lemma 3.1 (Cryptography and learning; cf. Kearns & Valiant [12]) *Consider a public-key cryptosystem for encrypting individual bits by n -bit strings. Let \mathcal{C} be a concept class that contains all the decryption functions $d_K : \{0, 1\}^n \rightarrow \{0, 1\}$ of the cryptosystem, one for each choice of key $K = (K_{\text{priv}}, K_{\text{pub}})$. Let $\varepsilon(n) = \Pr_{K,r}[d_K(e_{K,r}(0)) \neq 0 \text{ or } d_K(e_{K,r}(1)) \neq 1]$ be the probability of decryption error (over the choice of keys and randomization in the encryption). If \mathcal{C} is weakly PAC-learnable in time $t(n)$ with $t(n)\varepsilon(n) = 1/n^{\omega(1)}$, then there is a distinguisher between the encryptions of 0 and 1 that runs in time $O(t(n))$.*

Proof. For a pair of keys $K = (K_{\text{priv}}, K_{\text{pub}})$, let $e_{K,r} : \{0, 1\} \rightarrow \{0, 1\}^n$ be the randomized encryption function (indexed by the choice of random string r). Let $d_K : \{0, 1\}^n \rightarrow \{0, 1\}$ denote the matching decryption function. We will use the assumed learnability of \mathcal{C} to exhibit an algorithm \mathcal{A} that runs in time $O(t(n))$ and has

$$\Pr_{K,r}[\mathcal{A}(K_{\text{pub}}, e_{K,r}(1)) = 1] - \Pr_{K,r}[\mathcal{A}(K_{\text{pub}}, e_{K,r}(0)) = 1] \geq \frac{1}{n^c}$$

for some universal constant c , as long as $t(n)\varepsilon(n) = 1/n^{\omega(1)}$. The probability is taken over the choice of keys, randomness in the encryption, and any internal randomization in \mathcal{A} . It follows that \mathcal{A} is the desired distinguisher.

Algorithm \mathcal{A} takes as input a pair (K_{pub}, w) , where $w \in \{0, 1\}^n$ is the encryption of an unknown bit. First, \mathcal{A} draws $t(n)$ independent training examples, choosing each

as follows:

- (1) Pick $b = 0$ or $b = 1$, with equal probability.
- (2) Pick r , an unbiased random string.
- (3) Create a training example $\langle e_{K,r}(b), b \rangle$.

Next, \mathcal{A} passes the training examples to the assumed algorithm for learning \mathcal{C} . Assume no decryption error has occurred, i.e., the decryption function d_K is consistent with all the generated examples. Then the learning algorithm outputs a hypothesis h that approximates d_K with a nonnegligible advantage:

$$\Pr_{b,r}[h(e_{K,r}(b)) = d_K(e_{K,r}(b))] \geq \frac{1}{2} + \frac{1}{n^c}, \quad (3.1)$$

for some constant c . With this hypothesis in hand, algorithm \mathcal{A} outputs $h(w)$ and exits.¹

It remains to show that \mathcal{A} is indeed a distinguisher. We will first handle the case in which no decryption error occurs; call this event $\overline{\mathcal{E}}$. Then:

$$\begin{aligned} & \Pr_{K,r}[\mathcal{A}(K_{\text{pub}}, e_{K,r}(1)) = 1 \mid \overline{\mathcal{E}}] - \Pr_{K,r}[\mathcal{A}(K_{\text{pub}}, e_{K,r}(0)) = 1 \mid \overline{\mathcal{E}}] \\ &= \Pr_{K,r}[h(e_{K,r}(1)) = 1] - \Pr_{K,r}[h(e_{K,r}(0)) = 1] \\ &= 2 \Pr_{K,b,r}[h(e_{K,r}(b)) = b] - 1 \\ &\geq 2 \left(\Pr_{K,b,r}[h(e_{K,r}(b)) = d_K(e_{K,r}(b))] - \Pr_{K,b,r}[d_K(e_{K,r}(b)) \neq b] \right) - 1 \\ &\geq 1 + \frac{2}{n^c} - 2\varepsilon(n) - 1 \\ &= \frac{2}{n^c} - 2\varepsilon(n). \end{aligned}$$

¹ We have assumed that, given consistent training examples, the learner is guaranteed to succeed in finding a hypothesis h that satisfies (3.1). This makes for a shorter and simpler proof. In reality, we need only assume that the learner succeeds with probability $1/\text{poly}(n)$, and outputs “FAIL” otherwise. To accommodate this more general setting, it suffices to have \mathcal{A} output a random value (0 or 1) whenever the learner fails.

We now extend the analysis to account for possible decryption errors. Observe that the likelihood of a decryption error on a run of \mathcal{A} is small:

$$\begin{aligned}
\Pr[\mathcal{E}] &= \mathbf{E}_K [\Pr[\mathcal{E} \mid K]] \\
&\leq \mathbf{E}_K \left[t(n) \cdot \Pr_{b,r} [d_K(e_{K,r}(b)) \neq b \mid K] \right] && \text{(by union bound)} \\
&= t(n) \cdot \Pr_{K,b,r} [d_K(e_{K,r}(b)) \neq b] \\
&\leq t(n)\varepsilon(n).
\end{aligned}$$

This upper bound on $\Pr[\mathcal{E}]$, along with the above analysis of the error-free case, allows us to complete the proof of the desired claim (for all n large enough):

$$\begin{aligned}
&\Pr_{K,r} [\mathcal{A}(K_{\text{pub}}, e_{K,r}(1)) = 1] - \Pr_{K,r} [\mathcal{A}(K_{\text{pub}}, e_{K,r}(0)) = 1] \\
&\geq \left(\Pr_{K,r} [\mathcal{A}(K_{\text{pub}}, e_{K,r}(1)) = 1 \mid \overline{\mathcal{E}}] - \Pr_{K,r} [\mathcal{A}(K_{\text{pub}}, e_{K,r}(0)) = 1 \mid \overline{\mathcal{E}}] \right) - 2\Pr[\mathcal{E}] \\
&\geq \frac{2}{n^c} - 2\varepsilon(n) - 2t(n)\varepsilon(n) \\
&\geq \frac{1}{n^c}.
\end{aligned}$$

□

4 Implementing the Decryption Functions

Section 3 demonstrated that if a public-key cryptosystem is secure, then no concept class that can implement its decryption function is efficiently PAC-learnable. In what follows, we obtain implementations of the decryption functions from Section 2 by *intersections of degree-2 PTFs*. This will lead to a hardness result for learning intersections of degree-2 PTFs. We will obtain the main result of the paper by noting that intersections of degree-2 PTFs are no harder to learn than are intersections of halfspaces, a claim we formalize next.

Lemma 4.1 *Assume that intersections of n^ε arbitrary (respectively, light) halfspaces are weakly PAC-learnable. Then for any constant $c > 0$, intersections of n^c arbitrary (respectively, light) degree-2 PTFs are weakly PAC-learnable.*

Proof. We will prove the “light” case only; the “arbitrary” case is analogous. Consider the following concept classes:

- \mathcal{C} : intersections of n^ε light halfspaces;
- \mathcal{C}' : intersections of n^ε light degree-2 PTFs;
- \mathcal{C}'' : intersections of n^c light degree-2 PTFs.

First observe that a polynomial-time PAC-learning algorithm for \mathcal{C} implies one for \mathcal{C}' . This is because a degree-2 PTF in the n variables x_1, \dots, x_n is a halfspace in the $n + \binom{n}{2}$ variables $x_1, \dots, x_n, x_1x_2, x_1x_3, \dots, x_{n-1}x_n$, which yields a polynomial-time map from training/testing examples for a degree-2 PTF to those for a halfspace. This map is naturally viewed as a change of distribution: a given distribution of (x_1, \dots, x_n) will induce another, non-uniform distribution in the $n + \binom{n}{2}$ new variables.

Finally, a polynomial-time learning algorithm for \mathcal{C}' implies one for \mathcal{C}'' : by a standard padding argument, the problem of PAC learning the intersection of n^c halfspaces reduces to n^ε halfspaces for any constant $c > 0$. \square

4.1 The uSVP-based Cryptosystem

Recall that $\text{frac}(a) \doteq \min\{\lceil a \rceil - a, a - \lfloor a \rfloor\}$ denotes the distance from $a \in \mathbb{R}$ to the closest integer. Throughout this section, $\{a\}$ stands for the fractional part of $a \in \mathbb{R}$. Define the Boolean predicate

$$\text{NEAR-INT}(a) = 1 \iff \text{frac}(a) < 1/4.$$

This predicate ignores the integral part of a , meaning that $\text{NEAR-INT}(a) = \text{NEAR-INT}(\{a\})$.

The decryption function in the uSVP-based cryptosystem (Section 2) is $d_A(W) = \text{NEAR-INT}(AW)$, where A is a fixed real number and W is an integer input, both with a polynomial number of bits. We will demonstrate how to implement $\text{NEAR-INT}(AW)$ with intersections of degree-2 PTFs. A critical ingredient of our implementation is the ‘‘interval trick’’ of Siu and Roychowdhury [29], an insightful idea that was used in [29] to obtain a depth-2 light-weight threshold circuit for iterated addition.

Lemma 4.2 (Implementing the uSVP-based decryption function) *Let $A > 0$ be a real number with k fractional bits. Then the function $f(x) = \text{NEAR-INT}(A \sum_{j=0}^{n-1} x_j 2^j)$ can be computed by the intersection of k PTFs with degree 2 and weight $O(k^4 4^k)$.*

Proof. Let $\{A\} = .b_1 b_2 \dots b_k$ be the fractional part of A in binary, with $b_i \in \{0, 1\}$

for all i . The integral part of A is irrelevant. Then

$$\left\{ A \sum_{j=0}^{n-1} x_j 2^j \right\} = \left\{ \sum_{i=1}^k \sum_{j=0}^{n-1} b_i x_j 2^{j-i} \right\} = \left\{ \sum_{i=1}^k \sum_{j=0}^{\min\{n-1, i-1\}} b_i x_j 2^{j-i} \right\},$$

where the last equation follows by dropping those terms $b_i x_j 2^{j-i}$ that are whole numbers. Denote

$$S(x) \Rightarrow \sum_{i=1}^k \sum_{j=0}^{\min\{n-1, i-1\}} b_i x_j 2^{j-i}$$

so that $\{A \sum_{j=0}^{n-1} x_j 2^j\} = \{S(x)\}$. Observe that $S(x)$ is a multiple of $1/2^k$ and ranges between 0 and k . We will use degree-2 PTFs to identify intervals in $[0, k]$ on which $\text{NEAR-INT}(S(x)) = 1$. A listing of the first few such intervals is as follows:

Value of $S(x)$ in binary	NEAR-INT($S(x)$)
. 0 0 0 0 ... 0 0 ⋮ . 0 0 1 1 ... 1 1	1
. 0 1 0 0 ... 0 0 ⋮ . 1 1 0 0 ... 0 0	0
. 1 1 0 0 ... 0 1 ⋮ 1 . 0 0 1 1 ... 1 1	1
1 . 0 1 0 0 ... 0 0 ⋮ 1 . 1 1 0 0 ... 0 0	0
1 . 1 1 0 0 ... 0 1 ⋮ 1 0 . 0 0 1 1 ... 1 1	1

Each interval $[a, b]$ can be recognized by the PTF

$$\left(S(x) - \frac{a+b}{2} \right)^2 \leq \left(\frac{b-a}{2} \right)^2,$$

whose integral representation has weight $O(k^4 4^k)$. To compute the negation of an interval, we replace the inequality sign by “ $>$ ”. Finally, there are at most $2k + 1$ intervals because every two consecutive intervals, starting at the second, cover a distance of 1 on the interval $[0, k]$. By AND’ing the negations of the k intervals on which $\text{NEAR-INT}(S(x)) = 0$, we obtain the desired f as an AND of k weight- $O(k^4 4^k)$ degree-2 PTFs. \square

4.2 SVP- and SIVP-based Cryptosystems

For an integer a , define the Boolean predicate

$$\text{NEAR-MID}_p(a) \iff |b - \lfloor p/2 \rfloor| \leq \min\{b, p - b\},$$

where $b \in \{0, 1, \dots, p - 1\}$ is the integer with $a \equiv b \pmod{p}$. Recall that the decryption function in the SVP- and SIVP-based cryptosystem (Section 2) is $d_{s_1, \dots, s_n}(b, a_1, \dots, a_n) = \text{NEAR-MID}_p(b - \sum a_i s_i)$, where all s_i, a_i , and b are integers in $\{0, \dots, p - 1\} = \mathbb{Z}_p$. We will show how to compute d_{s_1, \dots, s_n} with intersections of degree-2 PTFs.

Lemma 4.3 (Implementing the SVP- and SIVP-based decryption function)

Let $d_{s_1, \dots, s_n} : (\{0, 1\}^{\log p})^{n+1} \rightarrow \{0, 1\}$ be the Boolean function defined by

$$d_{s_1, \dots, s_n}(x) = \text{NEAR-MID}_p \left(\sum_{i=0}^{\log p - 1} 2^i x_{0,i} - \sum_{j=1}^n s_j \sum_{i=0}^{\log p - 1} 2^i x_{j,i} \right),$$

where all s_i are integers in $\{0, \dots, p - 1\}$. Then d_{s_1, \dots, s_n} can be computed by the intersection of $n \log p$ PTFs with degree 2 and weight $O((pn \log p)^2)$.

Proof. Denote

$$S(x) = \sum_{i=0}^{\log p - 1} 2^i x_{0,i} - \sum_{j=1}^n \sum_{i=0}^{\log p - 1} (2^i s_j \bmod p) x_{j,i}.$$

Thus, $S(x)$ is the original weighted sum $\left(\sum_{i=0}^{\log p - 1} 2^i x_{0,i} - \sum_{j=1}^n s_j \sum_{i=0}^{\log p - 1} 2^i x_{j,i} \right)$ with the coefficients reduced modulo p .

Using the definition of NEAR-MID_p , we have $d_{s_1, \dots, s_n}(x) = \text{NEAR-MID}_p(S(x))$. The integer $S(x)$ ranges between $-(p - 1)n \log p$ and $p - 1$, a total range of length $< pn \log p$. As in the proof of Lemma 4.2, this range can be divided into consecutive intervals on which $d_{s_1, \dots, s_n}(x)$ is constant (i.e., does not change value within an interval).

Every two consecutive intervals cover a length of p units. Thus, there are a total of $\leq 2(pn \log p)/p = 2n \log p$ consecutive intervals. By picking out the $n \log p$ intervals on which $d_{s_1, \dots, s_n}(x) = 0$ and AND'ing their negations, we can compute d_{s_1, \dots, s_n} exactly. It remains to note that the negation of an interval $[a, b]$ can be computed by a degree-2 weight- $O((pn \log p)^2)$ PTF of the form $(S(x) - \frac{a+b}{2})^2 > (\frac{b-a}{2})^2$. \square

We additionally observe that the decryption function in the SVP- and SIVP-based cryptosystem can be computed by a depth-3 arithmetic circuit.

Lemma 4.4 (Extension to arithmetic circuits) *Let $d_{s_1, \dots, s_n} : (\{0, 1\}^{\log p})^{n+1} \rightarrow \{0, 1\}$ be the Boolean function defined by*

$$d_{s_1, \dots, s_n}(x) = \text{NEAR-MID}_p \left(\sum_{i=0}^{\log p - 1} 2^i x_{0,i} - \sum_{j=1}^n s_j \sum_{i=0}^{\log p - 1} 2^i x_{j,i} \right),$$

where all s_i are integers in $\{0, \dots, p-1\}$. Then d_{s_1, \dots, s_n} can be computed by a depth-3 arithmetic circuit of size $\text{poly}(p, n)$.

Proof. Set $S(x)$ as in the proof of Lemma 4.3. Then $S(x)$ is an integer in the range $R \equiv [-(p-1)n \log p, p-1] \cap \mathbb{Z}$ and completely determines the target function: $d_{s_1, \dots, s_n}(x) = \text{NEAR-MID}_p(S(x))$.

Let g be a polynomial such that $g(S(x)) = d_{s_1, \dots, s_n}(x)$ for all Boolean inputs x . It can be constructed by interpolating d_{s_1, \dots, s_n} on the range of $S(x)$ via the Lagrange formula:

$$g(y) = \sum_{r \in R} \text{NEAR-MID}_p(r) \cdot \prod_{r' \in R, r' \neq r} \frac{y - r'}{r - r'}.$$

Since the range R contains $\text{poly}(p, n)$ integers, $g(S(x))$ can be computed by a depth-3 arithmetic circuit of size $\text{poly}(p, n)$ with input $S(x)$ and summation gates at the bottom. But $S(x)$ is a sum of $\text{poly}(p, n)$ terms, each a singleton variable x_i or a constant. Thus, d_{s_1, \dots, s_n} can be computed directly by a depth-3 arithmetic circuit of size $\text{poly}(p, n)$ with inputs x . \square

5 Main Results

Based on the assumed hardness of the cryptosystems in Section 2 and the learning-to-cryptography reductions of Sections 3 and 4, we are in a position to prove the desired hardness results for learning intersections of halfspaces.

Theorem 1.1. (Restated from page 2.) *Assume that intersections of n^ε halfspaces in n dimensions are PAC-learnable for some constant $\varepsilon > 0$. Then there is a polynomial-time solution to $\tilde{O}(n^{1.5})$ -uSVP.*

Proof. Let \mathcal{C} denote the concept class of intersections of n^ε halfspaces, and let \mathcal{C}' denote the concept class of intersections of n^c degree-2 PTFs (for a large enough constant $c > 0$). By Lemma 4.1, the assumed PAC-learnability of \mathcal{C} implies the PAC-learnability of \mathcal{C}' . By Lemma 4.2, all the decryption functions in the uSVP-based cryptosystem are in \mathcal{C}' . A PAC-learning algorithm for \mathcal{C}' would thus yield a distinguisher between the encryptions of 0 and 1 (by Lemma 3.1) and hence an efficient solution to $O(\sqrt{n} \cdot \gamma(n))$ -uSVP for $\gamma(n) = n \log n$ (by Theorem 2.2). \square

Remark 5.1 Oded Regev has informed us [25] that Theorem 1.1 is also valid for light halfspaces, rather than arbitrary ones as stated. To see this, note that in Regev's first cryptosystem (Lemma 5.2 of [23]), except with probability exponentially small in n , the quantity $\text{frac}(AW)$ is bounded away from $\frac{1}{4}$ by a small constant. Therefore, with extremely high probability, we can ignore many of the least significant bits of AW , as these bits can only change the value of AW by $o(1)$. In Lemma 4.2, this allows one to restrict the sum $S(x)$ to contain only terms $b_{ix_j} 2^{j-i}$ with $j - i > -C \log n$ (for a sufficiently large constant $C > 0$), since the remaining terms contribute at most $o(1)$. The integral representation of the resulting PTF would have polynomial weight, leading to hardness for intersections of light halfspaces.

Theorem 1.2. (Restated from page 2.) *Assume that intersections of n^ε light halfspaces in n dimensions are PAC-learnable for some constant $\varepsilon > 0$. Then there is a polynomial-time quantum solution to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP.*

Proof. Let \mathcal{C} denote the concept class of intersections of n^ε light halfspaces, and let \mathcal{C}' denote the concept class of intersections of n^c light degree-2 PTFs (for a large enough constant $c > 0$). By Lemma 4.1, the assumed PAC-learnability of \mathcal{C} implies the PAC-learnability of \mathcal{C}' . By Lemma 4.3, the decryption function in the uSVP-based cryptosystem is in \mathcal{C}' . A PAC-learning algorithm for \mathcal{C}' would thus yield a distinguisher between the encryptions of 0 and 1 (by Lemma 3.1) and, as a result, an efficient quantum solution to $\tilde{O}(n \cdot \gamma(n))$ -SVP and $\tilde{O}(n \cdot \gamma(n))$ -SIVP for $\gamma(n) = \sqrt{n} \log^2 n$ (by Theorem 2.3). \square

Theorems 1.1 and 1.2 both imply a hardness result for learning polynomial-size depth-2 circuits of majority gates, a concept class commonly denoted by $\widehat{\text{LT}}_2$. To prove this, we will need a result regarding light threshold circuits, due to Goldmann, Håstad, and Razborov [9] and Goldmann and Karpinski [10]. Let $\widehat{\text{LT}}_d$ denote the class of depth- d polynomial-size circuits of threshold gates with polynomially-bounded weights. Let $\widetilde{\text{LT}}_d$ denote the class of depth- d polynomial-size threshold circuits in which only the output gate is required to have polynomially-bounded weights.

Theorem 5.2 [9,10] *For any fixed integer d , $\widehat{\text{LT}}_d = \widetilde{\text{LT}}_d$.*

We are now in a position to prove the desired hardness result for depth-2 neural networks.

Theorem 1.3. (Restated from page 3.) *Assume that depth-2 polynomial-size circuits of majority gates are PAC 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. Let $\widehat{\wedge\text{LT}}_1$ (respectively, $\wedge\text{LT}_1$) denote the concept classes of intersections of polynomially many light (respectively, arbitrary) halfspaces. By Theorems 1.1 and 1.2, it suffices to show that $\widehat{\wedge\text{LT}}_1 \subseteq \widehat{\text{LT}}_2$ and $\wedge\text{LT}_1 \subseteq \widehat{\text{LT}}_2$. The first statement is obvious: each halfspace is already a majority gate (with the inputs suitably negated/replicated), and the top gate $\text{AND}(f_1, f_2, \dots, f_t)$ can be replaced by a majority gate $\text{MAJ}(-t, f_1, f_2, \dots, f_t)$. To prove that $\wedge\text{LT}_1 \subseteq \widehat{\text{LT}}_2$, observe that $\wedge\text{LT}_1 \subseteq \widetilde{\text{LT}}_2$ (by an argument similar to the first case) and $\widetilde{\text{LT}}_2 = \widehat{\text{LT}}_2$ (by Theorem 5.2). \square

5.1 Hardness for PAC Learning Arithmetic Circuits

Here we give a hardness result for PAC learning depth-3 arithmetic circuits over the integers. Many researchers have constructed efficient, sparse polynomial interpolation algorithms where the learner has *query* access to the unknown polynomial [20,22,28]. If, in addition to membership queries, the learner can make equivalence queries, Klivans and Shpilka [16] have shown how to exactly learn restricted types of depth-3 arithmetic circuits via multiplicity automata techniques [4]. We show that if the learner receives random examples only, then learning depth-3 polynomial-size arithmetic circuits is as hard as solving $\tilde{O}(n^{1.5})$ -SVP in quantum polynomial-time:

Theorem 1.4. (Restated from page 3.) *Assume that depth-3 polynomial-size arithmetic circuits are PAC-learnable in polynomial time. Then there is a polynomial-time quantum solution to $\tilde{O}(n^{1.5})$ -SVP and $\tilde{O}(n^{1.5})$ -SIVP.*

Proof. Invoke Lemma 4.4 and argue as before (see the proofs of Theorems 1.1 and 1.2). \square

6 Hardness for AC^0 ?

A natural question to ask is whether our approach could yield hardness results for other concept classes. Particularly interesting candidates are AC^0 and, more

ambitiously, polynomial-size DNF formulas. Here we prove that the decryption functions of Regev's cryptosystems contain PARITY as a subfunction and thus are not computable in AC^0 .

We start with the easier proof. Recall that the decryption function of the SVP- and SIVP-based cryptosystem is $f_{s_1, \dots, s_n}(a_1, \dots, a_n, b) = \text{NEAR-MID}_p(b - \sum a_i s_i)$, where all s_i, a_i , and b are integers in $\{0, \dots, p-1\} = \mathbb{Z}_p$ with $n^2 < p < 2n^2$.

Proposition 6.1 (SVP-, SIVP-based cryptosystem and AC^0) *The decryption function of the SVP- and SIVP-based cryptosystem, $f_{s_1, \dots, s_n}(a_1, \dots, a_n, b) = \text{NEAR-MID}_p(b - \sum a_i s_i)$, is not in AC^0 .*

Proof. Let $x_1, x_2, \dots, x_n \in \{0, 1\}^n$. Note that

$$\text{NEAR-MID}_p\left(\frac{p-1}{2} \sum x_i\right) = \text{NEAR-MID}_p\left(\frac{p}{2} \sum x_i\right) = \text{PARITY}(x_1, \dots, x_n).$$

The first equality holds because $\frac{1}{2} \sum x_i \leq \frac{n}{2} \ll p$. Thus, $\text{PARITY}(x_1, \dots, x_n)$ is a subfunction of $\text{NEAR-MID}_p(b - \sum a_i s_i)$. Since AC^0 cannot compute PARITY [8,11], the claim follows. \square

Recall now that the decryption function in the uSVP-based cryptosystem is $d_A(X) = \text{NEAR-INT}(AX)$, where A is a fixed real number and X is an integer input. For convenience, we assume that X has $n+1$ bits rather than n .

Proposition 6.2 (uSVP-based cryptosystem and AC^0) *The decryption function of the uSVP-based cryptosystem, $d_A(X) = \text{NEAR-INT}(AX)$, is not in AC^0 .*

Proof. We will show that $d_A(X)$ computes PARITY on a subset of $\Theta(n/\log n)$ bits from among x_1, \dots, x_n (when the other bits are set to 0). The claim will follow.

Let $\Delta = 3 + \log n$ and $A = \sum_{i=0}^{n/\Delta} 2^{-i\Delta-1}$. For convenience of notation, we assume that $\Delta \mid n$. In what follows, we show that $d_A(X) = \text{PARITY}(x_0, x_\Delta, x_{2\Delta}, \dots, x_n)$ when $x_i = 0$ for all $i \notin \{0, \Delta, 2\Delta, \dots, n\}$. Namely,

$$\begin{aligned} d_A(X) &= \text{NEAR-INT}(AX) \\ &= \text{NEAR-INT}\left(\left(\sum_{i=0}^{n/\Delta} \frac{1}{2^{i\Delta+1}}\right) \left(\sum_{j=0}^{n/\Delta} x_{j\Delta} 2^{j\Delta}\right)\right) \\ &= \text{NEAR-INT}\left(\sum_i \sum_{j>i} \frac{x_{j\Delta} 2^{j\Delta}}{2^{i\Delta+1}} + \sum_i \frac{x_{i\Delta} 2^{i\Delta}}{2^{i\Delta+1}} + \sum_i \sum_{j<i} \frac{x_{j\Delta} 2^{j\Delta}}{2^{i\Delta+1}}\right). \end{aligned}$$

The first summation features only whole numbers and can thus be dropped. The second summation is precisely $\frac{1}{2}(x_0 + x_\Delta + x_{2\Delta} + \dots + x_n)$, a multiple of $\frac{1}{2}$. The

third summation does not exceed $1/8$ (by the choice of Δ and the geometric series) and thus does not affect the result. We obtain:

$$d_A(X) = \text{NEAR-INT} \left(\frac{x_0 + x_\Delta + x_{2\Delta} + \dots + x_n}{2} \right).$$

The latter expression is clearly $\text{PARITY}(x_0, x_\Delta, x_{2\Delta}, \dots, x_n)$. \square

References

- [1] D. Aharonov, O. Regev, Lattice problems in $\text{NP} \cap \text{coNP}$, *J. ACM* 52 (5) (2005) 749–765.
- [2] M. Ajtai, C. Dwork, A public-key cryptosystem with worst-case/average-case equivalence, in: *Proceedings of the 29th Symposium on Theory of Computing (STOC)*, 1997.
- [3] M. Alekhnovich, M. Braverman, V. Feldman, A. Klivans, T. Pitassi, Learnability and automatizability, in: *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, 2004.
- [4] A. Beimel, F. Bergadano, N. H. Bshouty, E. Kushilevitz, S. Varricchio, Learning functions represented as multiplicity automata, *J. ACM* 47 (3) (2000) 506–530.
- [5] A. L. Blum, R. L. Rivest, Training a 3-node neural network is NP-complete, *Neural Networks* 5 (1992) 117–127.
- [6] L. Blum, M. Blum, M. Shub, A simple unpredictable pseudo-random number generator, *SIAM J. Comput.* 15 (2) (1986) 364–383.
- [7] V. Feldman, P. Gopalan, S. Khot, A. K. Ponnuswami, New results for learning noisy parities and halfspaces, in: *Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS)*, 2006.
- [8] M. L. Furst, J. B. Saxe, M. Sipser, Parity, circuits, and the polynomial-time hierarchy, *Mathematical Systems Theory* 17 (1) (1984) 13–27.
- [9] M. Goldmann, J. Håstad, A. A. Razborov, Majority gates vs. general weighted threshold gates, *Computational Complexity* 2 (1992) 277–300.
- [10] M. Goldmann, M. Karpinski, Simulating threshold circuits by majority circuits, *SIAM J. Comput.* 27 (1) (1998) 230–246.
- [11] J. Håstad, *Computational limitations of small-depth circuits*, MIT Press, Cambridge, MA, USA, 1987.
- [12] M. Kearns, L. Valiant, Cryptographic limitations on learning Boolean formulae and finite automata, *J. ACM* 41 (1) (1994) 67–95.

- [13] M. J. Kearns, U. V. Vazirani, *An Introduction to Computational Learning Theory*, MIT Press, Cambridge, MA, USA, 1994.
- [14] M. Kharitonov, Cryptographic hardness of distribution-specific learning, in: *Proceedings of the 25th Symposium on Theory of Computing (STOC)*, 1993.
- [15] M. Kharitonov, Cryptographic lower bounds for learnability of Boolean functions on the uniform distribution, *J. Comput. Syst. Sci.* 50 (3) (1995) 600–610.
- [16] A. Klivans, A. Shpilka, Learning arithmetic circuits via partial derivatives, in: *Proceedings of the Workshop on Computational Learning Theory (COLT)*, 2003.
- [17] A. R. Klivans, R. O’Donnell, R. A. Servedio, Learning intersections and thresholds of halfspaces, *J. Comput. Syst. Sci.* 68 (4) (2004) 808–840.
- [18] A. R. Klivans, R. A. Servedio, Learning intersections of halfspaces with a margin, *J. Comput. Syst. Sci.* 74 (1) (2008) 35–48.
- [19] A. R. Klivans, A. A. Sherstov, Unconditional lower bounds for learning intersections of halfspaces, *Machine Learning* 69 (2–3) (2007) 97–114.
- [20] A. R. Klivans, D. A. Spielman, Randomness efficient identity testing of multivariate polynomials, in: *Proceedings of the 33rd Symposium on Theory of Computing (STOC)*, 2001.
- [21] S. Kwek, L. Pitt, PAC learning intersections of halfspaces with membership queries, *Algorithmica* 22 (1/2) (1998) 53–75.
- [22] Y. Mansour, Randomized interpolation and approximation of sparse polynomials, *SIAM J. Comput.* 24 (2) (1995) 357–368.
- [23] O. Regev, New lattice based cryptographic constructions, in: *Proceedings of the 35th Symposium on Theory of Computing (STOC)*, 2003, final version in *J. ACM* 51(6):899-942, 2004.
- [24] O. Regev, On lattices, learning with errors, random linear codes, and cryptography, in: *Proceedings of the 37th Symposium on Theory of Computing (STOC)*, 2005.
- [25] O. Regev, Personal communication (2006).
- [26] O. Regev, Lattice-based cryptography, in: *Proceedings of the 26th Annual International Cryptology Conference (CRYPTO)*, vol. 4117, 2006.
- [27] O. Regev, On the complexity of lattice problems with polynomial approximation factors, survey prepared for the LLL+25 conference. Available at <http://www.cs.tau.ac.il/~odedr> (June 2007).
- [28] R. E. Schapire, L. Sellie, Learning sparse multivariate polynomials over a field with queries and counterexamples, *J. Comput. Syst. Sci.* 52 (2) (1996) 201–213.
- [29] K.-Y. Siu, V. P. Roychowdhury, On optimal depth threshold circuits for multiplication and related problems, *SIAM J. Discrete Math.* 7 (2) (1994) 284–292.
- [30] L. G. Valiant, A theory of the learnable, *Commun. ACM* 27 (11) (1984) 1134–1142.

- [31] S. Vempala, A random sampling based algorithm for learning the intersection of halfspaces, in: Proceedings of the 38th Annual Symposium on Foundations of Computer Science, 1997.
- [32] I. Wegener, Optimal lower bounds on the depth of polynomial-size threshold circuits for some arithmetic functions, *Inf. Process. Lett.* 46 (2) (1993) 85–87.