

Order-Revealing Encryption and the Hardness of Private Learning

Mark Bun* Mark Zhandry†

May 2, 2015

Abstract

An order-revealing encryption scheme gives a public procedure by which two ciphertexts can be compared to reveal the ordering of their underlying plaintexts. We show how to use order-revealing encryption to separate computationally efficient PAC learning from efficient (ϵ, δ) -differentially private PAC learning. That is, we construct a concept class that is efficiently PAC learnable, but for which every efficient learner fails to be differentially private. This answers a question of Kasiviswanathan et al. (FOCS '08, SIAM J. Comput. '11).

To prove our result, we give a generic transformation from an order-revealing encryption scheme into one with strongly correct comparison, which enables the consistent comparison of ciphertexts that are not obtained as the valid encryption of any message. We believe this construction may be of independent interest.

Keywords: differential privacy, learning theory, order-revealing encryption

*School of Engineering & Applied Sciences, Harvard University. mbun@seas.harvard.edu. Supported by an NDSEG fellowship and NSF grant CNS-1237235.

†Stanford University. mzhandry@gmail.com. Supported by the DARPA PROCEED program.

1 Introduction

Many agencies hold sensitive information about individuals, where statistical analysis of this data could yield great societal benefit. The line of work on differential privacy [DMNS06] aims to enable such analysis while giving a strong formal guarantee on the privacy afforded to individuals. Noting that the framework of computational learning theory captures many of these statistical tasks, Kasiviswanathan et al. [KLN⁺11] initiated the study of *differentially private learning*. Roughly speaking, a differentially private learner is required to output a classification of labeled examples that is accurate, but does not change significantly based on the presence or absence of any individual example.

The early positive results in private learning established that, ignoring computational complexity, any concept class is privately learnable with a number of samples logarithmic in the size of the concept class [KLN⁺11]. Since then, a number of works have improved our understanding of the sample complexity – the minimum number of examples – required by such learners to simultaneously achieve accuracy and privacy. Some of these works showed that privacy incurs an inherent additional cost in sample complexity; that is, some concept classes require more samples to learn privately than they require to learn without privacy [BKN10, CH11, BNS13, FX14, CHS14, BNSV15]. In this work, we address the complementary question of whether there is also a *computational* price of differential privacy for learning tasks, for which much less is known. The initial work of Kasiviswanathan et al. [KLN⁺11] identified the important question of whether any efficiently PAC learnable concept class is also efficiently privately learnable, but only limited progress has been made on this question since then [BKN10, Nis14].

Our main result gives a strong negative answer to this question. We exhibit a concept class that is efficiently PAC learnable, but under plausible cryptographic assumptions cannot be learned efficiently and privately. To prove this result, we establish a connection between private learning and *order-revealing encryption*. We construct a new order-revealing encryption scheme with strong correctness properties that may be of independent learning-theoretic and cryptographic interest.

1.1 Differential Privacy and Private Learning

We first recall Valiant’s (distribution-free) PAC model for learning [Val84]. Let \mathcal{C} be a *concept class* consisting of concepts $c : X \rightarrow \{0, 1\}$ for a data universe X . A learner L is given n samples of the form $(x_i, c(x_i))$ where the x_i ’s are drawn i.i.d. from an unknown distribution, and are labeled according to an unknown concept c . The goal of the learner is to output a *hypothesis* $h : X \rightarrow \{0, 1\}$ from a hypothesis class \mathcal{H} that approximates c well on the unknown distribution. That is, the probability that h disagrees with c on a fresh example from the unknown distribution should be small – say, less than 0.05. The hypothesis class \mathcal{H} may be different from \mathcal{C} , but in the case where $\mathcal{H} \subseteq \mathcal{C}$ we call L a *proper* learner. Moreover, we say a learner is *efficient* if it runs in time polynomial in the description size of c and the size of its examples.

Kasiviswanathan et al. [KLN⁺11] defined a private learner to be a PAC learner that is also differentially private. Two samples $S = \{(x_1, b_1), \dots, (x_n, b_n)\}$ and $S' = \{(x'_1, b'_1), \dots, (x'_n, b'_n)\}$ are said to be *neighboring* if they differ on exactly one example, which we think of as corresponding to one individual’s information. A randomized learner $L : (X \times \{0, 1\})^n \rightarrow \mathcal{H}$ is (ϵ, δ) -*differentially private* if for all neighboring datasets S and S' and all sets $T \subseteq \mathcal{H}$,

$$\Pr[L(S) \in T] \leq e^\epsilon \Pr[L(S') \in T] + \delta.$$

The original definition of differential privacy [DMNS06] took $\delta = 0$, a case which is called *pure* differential privacy. The definition with positive δ , called *approximate* differential privacy, first appeared in [DKM⁺06] and has since been shown to enable substantial accuracy gains. Throughout this introduction, we will think of ε as a small constant, e.g. $\varepsilon = 0.1$, and $\delta = o(1/n)$.

Kasiviswanathan et al. [KLN⁺11] gave a generic “Private Occam’s Razor” algorithm, showing that any concept class \mathcal{C} can be privately (properly) learned using $O(\log |\mathcal{C}|)$ samples. Unfortunately, this algorithm runs in time $\Omega(|\mathcal{C}|)$, which is exponential in the description size of each concept. With an eye toward designing efficient private learners, Blum et al. [BDMN05] made the powerful observation that any efficient learning algorithm in the *statistical queries* (SQ) framework of Kearns [Kea98] can be efficiently simulated with differential privacy. Moreover, Kasiviswanathan et al. [KLN⁺11] showed that the efficient learner for the concept class of parity functions based on Gaussian elimination can also be implemented efficiently with differential privacy. These two techniques – SQ learning and Gaussian elimination – are essentially the only methods known for computationally efficient PAC learning. The fact that these can both be implemented privately led Kasiviswanathan et al. [KLN⁺11] to ask whether *all* efficiently learnable concept classes could also be efficiently learned with differential privacy.

Beimel et al. [BKN10] made partial progress toward this question in the special case of pure differential privacy with proper learning, showing that the sample complexity of efficient learners can be much higher than that of inefficient ones. Specifically, they showed that assuming the existence of pseudorandom generators with exponential stretch, there exists for any $\ell(d) = \omega(\log d)$ a concept class over $\{0, 1\}^d$ for which every efficient proper private learner requires $\Omega(d)$ samples, but an inefficient proper private learner only requires $O(\ell(d))$ examples. Nissim [Nis14] strengthened this result substantially for “representation learning,” where a proper learner is further restricted to output a canonical representation of its hypothesis. He showed that, assuming the existence of one-way functions, there exists a concept class that is efficiently representation learnable, but not efficiently privately representation learnable (even with approximate differential privacy). With Nissim’s kind permission, we give the details of this construction in Section 5.

Despite these negative results for proper learning, one might still have hoped that any efficiently learnable concept class could be efficiently *improperly* learned with privacy. Indeed, a number of works have shown that, especially with differential privacy, improper learning can be much more powerful than proper learning. For instance, Beimel et al. [BKN10] showed that under pure differential privacy, the simple class of Point functions (indicators of a single domain element) requires $\Omega(d)$ samples to privately learn properly, but only $O(\log d)$ samples to privately learn improperly. Moreover, computational separations are known between proper and improper learning even without privacy considerations. Pitt and Valiant [PV88] showed that unless $\mathbf{NP} = \mathbf{RP}$, k -term DNF are not efficiently properly learnable, but they are efficiently improperly learnable [Val84].

Under plausible cryptographic assumptions, we resolve the question of Kasiviswanathan et al. [KLN⁺11] in the negative, even for improper learners. The assumption we need is the existence of “strongly correct” order-revealing encryption (ORE) schemes, described in Section 1.3.

Theorem 1.1 (Informal). *Assuming the existence of strongly correct ORE, there exists an efficiently computable concept class EncThresh that is efficiently PAC learnable, but not efficiently learnable by any (ε, δ) -differentially private algorithm.*

We stress that this result holds even for improper learners and for the relaxed notion of approximate differential privacy. We remark that cryptography has played a major role in shaping our understanding of the computational complexity of learning in a number of models (e.g.

[Val84, KV94, Kha95, Ser00]). It has also been used before to show separations between what is efficiently learnable in different models (e.g. [Blu94, SG04]).

1.2 Our Techniques

We give an informal overview of the construction and analysis of the concept class `EncThresh`.

We first describe the concept class of thresholds `Thresh` and its simple PAC learning algorithm. Consider the domain $[N] = \{1, \dots, N\}$. Given a number $t \in [N]$, a threshold concept c_t is defined by $c_t(x) = 1$ if and only if $x \leq t$. The concept class of thresholds admits a simple and efficient proper PAC learning algorithm L_{Thresh} . Given a sample $\{(x_1, c_t(x_1)), \dots, (x_n, c_t(x_n))\}$ labeled by an unknown concept c_t , the learner L_{Thresh} identifies the largest positive example x_{i^*} and outputs the hypothesis $h = c_{x_{i^*}}$. That is, L_{Thresh} chooses the threshold concept that minimizes the empirical error on its sample. To achieve a small constant error on *any* underlying distribution on examples, it suffices to take $n = O(1)$ samples.

A simple but important observation about L_{Thresh} is that it is completely oblivious to the actual numeric values of its examples, or even to the fact that the domain is $[N]$. In fact, L_{Thresh} works equally well on any totally-ordered domain on which it can efficiently compare examples. In an extreme case, the learner L_{Thresh} still works when its examples are encrypted under an *order-revealing encryption* (ORE) scheme, which guarantees that L_{Thresh} is able to learn the order of its examples, but nothing else about them. Up to small technical modifications, our concept class `EncThresh` is exactly the class `Thresh` where examples are encrypted under an ORE scheme.

For `EncThresh` to be efficiently PAC learnable, it must be learnable even under distributions that place arbitrary weight on examples corresponding to invalid ciphertexts. To this end, we require a “strong correctness” condition on our ORE scheme. The strong correctness condition ensures that all ciphertexts, even those that are not obtained as encryptions of messages, can be compared in a consistent fashion. This condition is not met by current constructions of ORE, and one of the technical contributions of this work is a generic transformation from weakly correct ORE schemes to strongly correct ones.

While a learner similar to L_{Thresh} is able to efficiently PAC learn the concept class `EncThresh`, we argue that it cannot do so while preserving differential privacy with respect to its examples. Intuitively, the security of the ORE scheme ensures that essentially the only thing a learner for `EncThresh` can do is output a hypothesis that compares an example to one it already has. We make this intuition precise by giving an algorithm that traces the hypothesis output by any efficient learner back to one of the examples used to produce it. This formalization builds conceptually on the connection between differential privacy and traitor-tracing schemes (see Section 1.4), but requires new ideas to adapt to the PAC learning model.

1.3 Order-Revealing Encryption

Motivated by the task of answering range queries on encrypted databases, an *order-revealing encryption* (ORE) scheme [BCO11, BLR⁺15] is a special type of symmetric key encryption scheme where it is possible to publicly sort *ciphertexts* according to the order of the *plaintexts*. More precisely, the plaintext space of the scheme is the set of integers $[N] = \{1, \dots, N\}$,¹ and in addition to the *private* encryption and decryption procedures `Enc`, `Dec`, there is a public comparison procedure `Comp` that takes as input two ciphertexts, and reveals the order of the corresponding plaintexts.

¹More generally, any totally-ordered plaintext space can be considered

The notion of *best-possible semantic security*, defined in Boneh et al. [BLR⁺15], intuitively captures the requirement that, given a collection of ciphertexts, no information about the plaintexts is learned, *except* for the ordering.

Known constructions of order-revealing encryption. Order-revealing encryption can be seen as a special case of 2-input *functional encryption*. In such a scheme, there are several functions f_1, \dots, f_k , and given two ciphertexts c_0, c_1 encrypting m_0, m_1 , it is possible to learn $f_i(m_0, m_1)$ for all $i \in [k]$. General *multi-input* functional encryption schemes can be obtained from indistinguishability obfuscation [GGG⁺14] or multilinear maps [BLR⁺15]. It is also possible to build ORE from *single-input* functional encryption with function privacy, which means that f is kept secret. Such schemes can be built from regular single-input schemes without function privacy by work of Brakerski and Segev [BS15], and such single-input schemes can also be built from obfuscation [GGH⁺13b] or multilinear maps [GGHZ14].

Unfortunately, the above constructions are insufficient for our purposes. The issue arises from the fact that our learner needs to work for *any* distribution on ciphertexts, even distributions whose support includes malformed ciphertexts. Unfortunately, previous constructions only achieve a weak form of correctness, which guarantees that encrypting two messages and then comparing the ciphertexts using **Comp** produces the same result (with overwhelming probability) as comparing the plaintexts directly. This requirement only specifies how **Comp** works on *valid* ciphertexts, namely actual encryptions of messages. Moreover, correctness is only guaranteed for these messages with overwhelming probability, meaning even some valid ciphertexts may cause **Comp** to misbehave.

For our learner, this weak form of correctness means, for some distributions that place significant weight on bad ciphertexts, the comparison procedure is completely useless, and thus the learner will fail for these distributions.

We therefore need a stronger correctness guarantee. We need that, for any two *ciphertexts*, the comparison procedure is consistent with decrypting the two ciphertexts and comparing the resulting plaintexts. This correctness guarantee is meaningful even for improperly generated ciphertexts.

We note that none of the existing constructions of order-revealing encryption outlined above satisfy this stronger notion. For the obfuscation-based schemes, ciphertexts consist of obfuscated programs. In these schemes, it is easy to describe invalid ciphertexts where the obfuscated program performs incorrectly, causing the comparison procedure to output the wrong result. In the multilinear map-based schemes, the underlying instantiation use current “noisy” multilinear maps, such as [GGH13a]. An invalid ciphertext could, for example, have too much noise, which will cause the comparison procedure to behave unpredictably.

Obtaining strong correctness. We first argue that, for all existing ORE schemes, the scheme can be modified so that **Comp** is correct for all *valid* ciphertexts. We then give a generic conversion from any ORE scheme with weakly correct comparison, including the tweaked existing schemes, into a strongly correct scheme. We simply modify the ciphertext by adding a non-interactive zero-knowledge (NIZK) proof that the ciphertext is well-formed, with the common reference string added to the public comparison key. Then the decryption and comparison procedures check the proof(s), and only output the result (either decryption or comparison) if the proof(s) are valid. The (computational) zero-knowledge property of the NIZK implies that the addition of the proof to the ciphertext does not affect security. Meanwhile, NIZK soundness implies that any ciphertext

accepted by the decryption and comparison procedures must be valid, and the weak correctness property of the underlying ORE implies that for valid ciphertexts, decryption and comparison are consistent. The result is that comparisons are consistent with decryption *for all* ciphertexts, giving strong correctness.

As we need strong correctness for every ciphertext, even hard-to-generate ones, we need the NIZK proofs to have perfect soundness, as opposed to computational soundness. Such NIZK proofs were built in [GOS12].

We note also that the conversion outlined above is not specific to ORE, and applies more generally to functional encryption schemes.

1.4 Related Work

Hardness of Private Query Release. One of the most basic and well-studied statistical tasks in differential privacy is the problem of releasing answers to *counting queries*. A counting query asks, “what fraction of the records in a dataset D satisfy the predicate q ?”. Given a collection of k counting queries q_1, \dots, q_k from a family \mathcal{Q} , the goal of a query release algorithm is to release approximate answers to these queries while preserving differential privacy. A remarkable result of Blum et al. [BLR08], with subsequent improvements by [DNR⁺09, DRV10, RR10, HR10, GRU12, HLM12], showed that an arbitrary sequence of counting queries can be answered accurately with differential privacy even when k is exponential in the dataset size n . Unfortunately, all of these algorithms that are capable of answering more than n^2 queries are inefficient, running in time exponential in the dimensionality of the data. Moreover, several works [DNR⁺09, Ull13, BZ14] have gone on to show that this inefficiency is likely inherent.

These computational lower bounds for private query release rely on a connection between the hardness of private query release and *traitor-tracing schemes*, which was first observed by Dwork et al. [DNR⁺09]. Traitor-tracing schemes were introduced by Chor, Fiat, and Naor [CFN94] to help digital content producers identify pirates as they illegally redistribute content. Traitor-tracing schemes are conceptually analogous to the example reidentification scheme we use to obtain our hardness result for private learning. Instantiating this connection with the traitor-tracing scheme of Boneh, Sahai, and Waters [BSW06], which relies on certain assumptions in bilinear groups, Dwork et al. [DNR⁺09] exhibited a family of $2^{\tilde{O}(\sqrt{n})}$ queries for which no efficient algorithm can produce a data structure which could be used to answer all queries in this family. Very recently, Boneh and Zhandry [BZ14] constructed a new traitor-tracing scheme based on indistinguishability obfuscation that yields the same infeasibility result for a family of $n \cdot 2^{O(d)}$ queries on records of size d . Extending this connection, Ullman [Ull13] constructed a specialized traitor-tracing scheme to show that no efficient private algorithm can answer more than $\tilde{O}(n^2)$ arbitrary queries that are given as input to the algorithm.

Dwork et al. [DNR⁺09] also showed strong lower bounds against private algorithms for producing *synthetic data*. Synthetic data generation algorithms produce a new “fake” dataset, whose rows are of the same type as those in the original dataset, with the promise that the answers to some restricted set of queries on the synthetic dataset well-approximate the answers on the original dataset. Assuming the existence of one-way functions, Dwork et al. [DNR⁺09] exhibited an efficiently computable collection of queries for which no efficient private algorithm can produce useful synthetic data. Ullman and Vadhan [UV11] refined this result to hold even for extremely simple classes of queries.

Nevertheless, the restriction to synthetic data is significant to these results, and they do not rule

out the possibility that other privacy-preserving data structures can be used to answer large families of restricted queries. In fact, when the synthetic data restriction is lifted, there are algorithms (e.g. [HRS12, TUV12, CTUW14, DNT14]) that answer queries from certain exponentially large families in subexponential time. One can view the problem of synthetic data generation as analogous to proper learning. In both cases, placing natural syntactic restrictions on the output of an algorithm may in fact come at the expense of utility or computational efficiency.

Efficiency of SQ Learning. Feldman and Kanade [FK12] addressed the question of whether information-theoretically efficient SQ learners – i.e., those making polynomially many queries – could be made computationally efficient. One of their main negative results showed that unless $\mathbf{NP} = \mathbf{RP}$, there exists a concept class with polynomial query complexity that is not efficiently SQ learnable. Moreover, this concept class is efficiently PAC learnable, which suggests that the restriction to SQ learning can introduce an inherent computational cost.

We show that the concept class `EncThresh` can be learned (inefficiently) with polynomially many statistical queries. The result of Blum et al. [BDMN05] discussed above, showing that SQ learning algorithms can be efficiently simulated by differentially private algorithms, thus shows that `EncThresh` also separates SQ learners making polynomially many queries from computationally efficient SQ learners.

Corollary 1.2 (Informal). *Assuming the existence of strongly correct ORE, the concept class `EncThresh` is efficiently PAC learnable and has polynomial SQ query complexity, but is not efficiently SQ learnable.*

While our proof relies on much stronger hardness assumptions, it reveals ORE as a new barrier to efficient SQ learning. As discussed in more detail in Section 3.3, even though their result is about computational hardness, Feldman and Kanade’s choice of a concept class relies crucially on the fact that parities are hard to learn in the SQ model even information-theoretically. By contrast, our concept class `EncThresh` is computationally hard to SQ learn for a reason that appears fundamentally different than the information-theoretic hardness of SQ learning parities.

Learning from Encrypted Data. Several works have developed schemes for training, testing, and classifying machine learning models over encrypted data (e.g. [GLN13, BPTG14]). In a model use case, a client holds a sensitive dataset, and uploads an encrypted version of the dataset to a cloud computing service. The cloud service then trains a model over the encrypted data and produces an encrypted classifier it can send back to the client, ideally without learning anything about the examples it received. The notion of privacy afforded to the individuals in the dataset here is complementary to differential privacy. While the cloud service does not learn anything about the individuals in the dataset, its output might still depend heavily on the data of certain individuals.

In fact, our non-differentially private PAC learner for the class `EncThresh` exactly performs the task of learning over encrypted data, producing a classifier without learning anything about its examples beyond their order (this addresses the difficulty of implementing comparisons from prior work [GLN13]). Thus one can interpret our results as showing that not only are these two notions of privacy for machine learning training complementary, but that they may actually be in conflict. Moreover, the strong correctness guarantee we provide for ORE (which applies more generally to multi-input functional encryption) may help enable the theoretical study of learning from encrypted data in other PAC-style settings.

2 Preliminaries and Definitions

2.1 PAC Learning and Private PAC Learning

For each $k \in \mathbb{N}$, let X_k be an instance space (such as $\{0, 1\}^k$), where the parameter k represents the size of the elements in X_k . Let \mathcal{C}_k be a set of boolean functions $\{c : X_k \rightarrow \{0, 1\}\}$. The sequence $(X_1, \mathcal{C}_1), (X_2, \mathcal{C}_2), \dots$ represents an infinite sequence of learning problems defined over instance spaces of increasing dimension. We will generally suppress the parameter k , and refer to the problem of learning \mathcal{C} as the problem of learning \mathcal{C}_k for every k .

A learner L is given examples sampled from an unknown probability distribution \mathcal{D} over X , where the examples are labeled according to an unknown *target concept* $c \in \mathcal{C}$. The learner must select a hypothesis h from a hypothesis class \mathcal{H} that approximates the target concept with respect to the distribution \mathcal{D} . More precisely,

Definition 2.1. The *generalization error* of a hypothesis $h : X \rightarrow \{0, 1\}$ (with respect to a target concept c and distribution \mathcal{D}) is defined by $\text{error}_{\mathcal{D}}(c, h) = \Pr_{x \sim \mathcal{D}}[h(x) \neq c(x)]$. If $\text{error}_{\mathcal{D}}(c, h) \leq \alpha$ we say that h is an α -*good* hypothesis for c on \mathcal{D} .

Definition 2.2 (PAC Learning [Val84]). Algorithm $L : (X \times \{0, 1\})^n \rightarrow \mathcal{H}$ is an (α, β) -*accurate PAC learner* for the concept class \mathcal{C} using hypothesis class \mathcal{H} with sample complexity n if for all target concepts $c \in \mathcal{C}$ and all distributions \mathcal{D} on X , given an input of n samples $S = ((x_1, c(x_1)), \dots, (x_n, c(x_n)))$, where each x_i is drawn i.i.d. from \mathcal{D} , algorithm L outputs a hypothesis $h \in \mathcal{H}$ satisfying $\Pr[\text{error}_{\mathcal{D}}(c, h) \leq \alpha] \geq 1 - \beta$. The probability here is taken over the random choice of the examples in S and the coin tosses of the learner L .

The learner L is *efficient* if it runs in time polynomial in the size parameter k , the representation size of the target concept c , and the accuracy parameters $1/\alpha$ and $1/\beta$. Note that a necessary (but not sufficient) condition for L to be efficient is that its sample complexity n is polynomial in the learning parameters.

If $\mathcal{H} \subseteq \mathcal{C}$ then L is called a *proper* learner. Otherwise, it is called an *improper* learner.

Kasiviswanathan et al. [KLN⁺11] defined a *private learner* as a PAC learner that is also differentially private. Recall the definition of differential privacy:

Definition 2.3. A learner $L : (X \times \{0, 1\})^n \rightarrow \mathcal{H}$ is (ϵ, δ) -*differentially private* if for all sets $T \subseteq \mathcal{H}$, and neighboring sets of examples $S \sim S'$,

$$\Pr[L(S) \in T] \leq e^\epsilon \Pr[L(S') \in T] + \delta.$$

The technical object that we will use to show our hardness results for differential privacy is what we call an *example reidentification scheme*. It is analogous to the hard-to-sanitize database distributions [DNR⁺09, UV11] and re-identifiable database distributions [BUV14] used in prior works to prove hardness results for private query release, but is adapted to the setting of computational learning. In the first step, an algorithm Gen_{ex} chooses a concept and a sample S labeled according to that concept. In the second step, a learner L receives either the sample S or the sample S_{-i} where an appropriately chosen example i is replaced by a junk example, and learns a hypothesis h . Finally, an algorithm Trace_{ex} attempts to use h to identify one of the rows given to L . If Trace_{ex} succeeds at identifying such a row with high probability, then it must be able to distinguish $L(S)$ from $L(S_{-i})$, showing that L cannot be differentially private. We formalize these ideas below.

Definition 2.4. An (α, ξ) -example reidentification scheme for a concept class \mathcal{C} consists of a pair of algorithms, $(\text{Gen}_{\text{ex}}, \text{Trace}_{\text{ex}})$ with the following properties.

$\text{Gen}_{\text{ex}}(k, n)$ Samples a concept $c \in \mathcal{C}_k$ and an associated distribution \mathcal{D} . Draws i.i.d. examples $x_1, \dots, x_n \leftarrow_{\text{R}} \mathcal{D}$, and a fixed value x_0 . Let S denote the labeled sample $((x_1, c(x_1)), \dots, (x_n, c(x_n)))$, and for any index $i \in [n]$, let S_{-i} denote the sample with the pair $(x_i, c(x_i))$ replaced with $(x_0, c(x_0))$.

$\text{Trace}_{\text{ex}}(h)$ Takes state shared with Gen_{ex} as well as a hypothesis h and identifies an index in $[n]$ (or \perp if none is found).

The scheme obeys the following ‘‘completeness’’ and ‘‘soundness’’ criteria on the ability of Trace_{ex} to identify an example given to a learner L .

Completeness. A good hypothesis can be traced to some example. That is, for every efficient learner L ,

$$\Pr[\text{error}_{\mathcal{D}}(c, h) \leq \alpha \wedge \text{Trace}_{\text{ex}}(h) = \perp] \leq \xi.$$

Here, the probability is taken over $h \leftarrow_{\text{R}} L(S)$ and the coins of Gen_{ex} and Trace_{ex} .

Soundness. For every efficient learner L , Trace_{ex} cannot trace i from the sample S_{-i} . That is, for all $i \in [n]$,

$$\Pr[\text{Trace}_{\text{ex}}(h) = i] \leq \xi$$

for $h \leftarrow_{\text{R}} L(S_{-i})$.

We may sometimes relax the completeness condition to hold only under certain restrictions on L 's output (e.g. L is a proper learner or L is a representation learner). In this case, we say the $(\text{Gen}_{\text{ex}}, \text{Trace}_{\text{ex}})$ is an example reidentification scheme for (properly, representation) learning a class \mathcal{C} .

Theorem 2.5. Let $(\text{Gen}_{\text{ex}}, \text{Trace}_{\text{ex}})$ be an (α, ξ) -example reidentification scheme for a concept class \mathcal{C} . Then for every $\beta > 0$ and polynomial $n(k)$, there is no efficient (ε, δ) -differentially private (α, β) -PAC learner for \mathcal{C} using n samples when

$$\delta < \left(\frac{1 - \beta - \xi}{n} \right) - e^{\varepsilon} \xi.$$

In a typical setting of parameters, we will take $\alpha, \beta, \varepsilon = O(1)$ and $\delta, \xi = o(1/n)$, in which case the inequality in Theorem 2.5 will be satisfied for sufficiently large n .

Proof. Suppose instead that there were a computationally efficient (ε, δ) -differentially private (α, β) -PAC learner L for \mathcal{C} using n samples. Then there exists an $i \in [n]$ such that $\Pr[\text{Trace}_{\text{ex}}(L(S)) = i] \geq (1 - \beta - \xi)/n$. However, since L is differentially private,

$$\Pr[\text{Trace}_{\text{ex}}(L(S_{-i})) = i] \geq e^{-\varepsilon} \left(\frac{1 - \beta - \xi}{n} - \delta \right) > \xi(n),$$

which contradicts the soundness of $(\text{Gen}_{\text{ex}}, \text{Trace}_{\text{ex}})$. \square

2.2 Order-Revealing Encryption

Definition 2.6. An Order-Revealing Encryption (ORE) scheme is a tuple $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$ of algorithms where:

- $\text{Gen}(1^\lambda, 1^\ell)$ is a randomized procedure that takes as inputs a security parameter λ and plaintext length ℓ , and outputs a secret encryption/decryption key sk and public parameters params .
- $\text{Enc}(\text{sk}, m)$ is a potentially randomized procedure that takes as input a secret key sk and a message $m \in \{0, 1\}^\ell$, and outputs a ciphertext c .
- $\text{Dec}(\text{sk}, c)$ is a deterministic procedure that takes as input a secret key sk and a ciphertext c , and outputs a plaintext message $m \in \{0, 1\}^\ell$ or a special symbol \perp .
- $\text{Comp}(\text{params}, c_0, c_1)$ is a deterministic procedure that “compares” two ciphertexts, outputting either “>”, “<”, “=”, or \perp .

Correctness. An ORE scheme must satisfy two separate correctness requirements:

- **Correct Decryption:** This is the standard notion of correctness for an encryption scheme, which says that decryption succeeds. We will only consider *strongly* correct decryption, which requires that decryption *always* succeeds. For all security parameters λ and message lengths ℓ ,

$$\Pr[\text{Dec}(\text{sk}, \text{Enc}(\text{sk}, m)) = m : (\text{sk}, \text{params}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)] = 1.$$

- **Correct Comparison:** We require that the comparison function succeeds. We will consider two notions, namely *strong* and *weak*. In order to define these notions, we first define two auxiliary functions:

- $\text{Comp}_{\text{plain}}(m_0, m_1)$ is just the plaintext comparison function. That is, for $m_0 < m_1$, $\text{Comp}_{\text{plain}}(m_0, m_1) = “<”$, $\text{Comp}_{\text{plain}}(m_1, m_0) = “>”$, and $\text{Comp}_{\text{plain}}(m_0, m_0) = “=”$.
- $\text{Comp}_{\text{ciph}}(\text{sk}, c_0, c_1)$ is a ciphertext comparison function which uses the secret key. It first computes $m_b = \text{Dec}(\text{sk}, c_b)$ for $b = 0, 1$. If either $m_0 = \perp$ or $m_1 = \perp$ (in other words, if either decryption failed), then $\text{Comp}_{\text{ciph}}$ outputs \perp . If both $m_0, m_1 \neq \perp$, then the output is $\text{Comp}_{\text{plain}}(m_0, m_1)$.

Now we define our comparison correctness notions:

- **Weakly Correct Comparison:** This informally requires that comparison is consistent with encryption. For all security parameters λ , message lengths ℓ , and messages $m_0, m_1 \in \{0, 1\}^\ell$,

$$\Pr \left[\text{Comp}(\text{params}, c_0, c_1) = \text{Comp}_{\text{plain}}(m_0, m_1) : \begin{array}{l} (\text{sk}, \text{params}) \leftarrow \text{Gen}(1^\lambda, 1^\ell) \\ c_b \leftarrow \text{Enc}(\text{sk}, m_b) \end{array} \right] = 1.$$

In particular, for correctly generated ciphertexts, Comp never outputs \perp .

- **Strongly Correct Comparison:** This informally requires that comparison is consistent with *decryption*. For all security parameters λ , message lengths ℓ , and ciphertexts c_0, c_1 ,

$$\Pr \left[\text{Comp}(\text{params}, c_0, c_1) = \text{Comp}_{\text{ciph}}(\text{sk}, c_0, c_1) : (\text{sk}, \text{params}) \leftarrow \text{Gen}(1^\lambda, 1^\ell) \right] = 1.$$

Security. For security, we will consider a relaxation of the “best possible” security notion of Boneh et al. [BLR⁺15]. Namely, we only consider static adversaries that submit all queries at once. “Best possible” security is a modification of the standard notion of CPA security for symmetric key encryption to block trivial attacks. That is, since the comparison function always leaks the order of the plaintexts, the left and right sets of challenge messages must have the same order. In our relaxation where all challenge messages are queried at once, we can therefore assume without loss of generality that the left and right sequences of messages are sorted in ascending order. For simplicity, we will also disallow the adversary from querying on the same message more than once. This gives the following definition:

Definition 2.7. An ORE scheme $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$ is *statically secure* if, for all efficient adversaries \mathcal{A} , $|\Pr[W_0] - \Pr[W_1]|$ is negligible, where W_b is the event that \mathcal{A} outputs 1 in the following experiment:

- \mathcal{A} produces two message sequences $m_1^{(L)} < m_2^{(L)} < \dots < m_q^{(L)}$ and $m_1^{(R)} < m_2^{(R)} < \dots < m_q^{(R)}$
- The challenger runs $(\text{sk}, \text{params}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)$. It then responds to \mathcal{A} with params , as well as c_1, \dots, c_q where

$$c_i = \begin{cases} \text{Enc}(\text{sk}, m_i^{(L)}) & \text{if } b = 0 \\ \text{Enc}(\text{sk}, m_i^{(R)}) & \text{if } b = 1 \end{cases}$$

- \mathcal{A} outputs a guess b' for b .

We also consider a weaker definition, which only allows the sequences $m_i^{(L)}$ and $m_i^{(R)}$ to differ at a single point:

Definition 2.8. An ORE scheme $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$ is *statically single-challenge secure* if, for all efficient adversaries \mathcal{A} , $|\Pr[W_0] - \Pr[W_1]|$ is negligible, where W_b is the event that \mathcal{A} outputs 1 in the following experiment:

- \mathcal{A} produces a sequence of messages $m_1 < m_2 < \dots < m_q$, and challenge messages m_L, m_R such that $m_i < m_L < m_R < m_{i+1}$ for some $i \in [q - 1]$.
- The challenger runs $(\text{sk}, \text{params}) \leftarrow \text{Gen}(1^\lambda, 1^\ell)$. It then responds to \mathcal{A} with params , as well as c_1, \dots, c_q where $c_i = \text{Enc}(\text{sk}, m_i)$ and

$$c^* = \begin{cases} \text{Enc}(\text{sk}, m_L) & \text{if } b = 0 \\ \text{Enc}(\text{sk}, m_R) & \text{if } b = 1 \end{cases}$$

- \mathcal{A} outputs a guess b' for b .

We now argue that these two definitions are equivalent up to some polynomial loss in security.

Theorem 2.9. (Gen, Enc, Dec, Comp) is statically secure if and only if it is statically single-challenge secure.

Proof. We prove that single-challenge security implies many-challenge security through a sequence of hybrids. Each hybrid will only differ in the messages m_i that are encrypted, and each adjacent hybrid will only differ in a single message. The first hybrid will encrypt $m_i^{(L)}$, and the last hybrid will encrypt $m_i^{(R)}$. Thus, by applying the single-challenge security for each hybrid, we conclude that the first and last hybrids are indistinguishable, thus showing many-challenge security.

Hybrid j for $j \leq q$.

$$m_i = \begin{cases} \min(m_i^{(L)}, m_i^{(R)}) & \text{if } i \leq j \\ m_i^{(L)} & \text{if } i > j \end{cases}$$

First, notice that all the m_i are in order since both sequences $m_i^{(L)}$ and $m_i^{(R)}$ are in order. Second, the only difference between **Hybrid** $(j - 1)$ and **Hybrid** j is that $m_j = m_j^{(L)}$ in **Hybrid** $(j - 1)$ and $m_j = \min(m_j^{(L)}, m_j^{(R)})$ in **Hybrid** j . Thus, single-challenge security implies that each adjacent hybrid is indistinguishable. Moreover, for j where $m_j^{(L)} < m_j^{(R)}$, the two hybrids are actually identical.

Hybrid j for $j > q$.

$$m_i = \begin{cases} \min(m_i^{(L)}, m_i^{(R)}) & \text{if } i \leq 2q - j \\ m_i^{(R)} & \text{if } i > 2q - j \end{cases}$$

Again, notice that all the m_i are in order. Moreover, the only difference between **Hybrid** $(2q - j)$ and **Hybrid** $(2q - j + 1)$ is that $m_j = \min(m_j^{(L)}, m_j^{(R)})$ in **Hybrid** $(2q - j)$ and $m_j = m_j^{(R)}$ in **Hybrid** $(2q - j + 1)$. Thus, single-challenge security implies that each adjacent hybrid is indistinguishable. Moreover, for j where $m_j^{(L)} > m_j^{(R)}$, the two hybrids are actually identical. □

3 The Concept Class EncThresh and its Learnability

Let (Gen, Enc, Dec, Comp) be a statically secure ORE scheme with strongly correct comparison. We define a concept class **EncThresh**, which intuitively captures the class of threshold functions where examples are encrypted under the ORE scheme. Throughout this discussion, we will take $N = 2^\ell$ and regard the plaintext space of the ORE scheme to be $[N] = \{1, \dots, N\}$. Ideally we would like, for each threshold $t \in [N + 1]$ and each $(\text{sk}, \text{params}) \leftarrow \text{Gen}(1^\lambda)$, to define a concept

$$f_{t, \text{sk}, \text{params}}(c) = \begin{cases} 1 & \text{if } \text{Dec}_{\text{sk}}(c) < t \\ 0 & \text{otherwise.} \end{cases}$$

However, we need to make a few technical modifications to ensure that **EncThresh** is efficiently PAC learnable.

1. In order for the learner to be able to use the comparison function `Comp`, it must be given the public parameters `params` generated by the ORE scheme. We address this in the natural way by attaching a set of public parameters to each example. Moreover, we define `EncThresh` so that each concept is supported on the single set of public parameters that corresponds to the secret key used for encryption and decryption.
2. Only a subset of binary strings form valid $(\text{sk}, \text{params})$ pairs that are actually produced by `Gen` in the ORE scheme. To represent concepts, we need a reasonable encoding scheme for these valid pairs. The encoding scheme we choose is the polynomial-length sequence of random coin tosses used by the algorithm `Gen` to produce $(\text{sk}, \text{params})$.

We now formally describe the concept class `EncThresh`. Each concept is parameterized by a string r , representing the coin tosses of the algorithm `Gen`, and a threshold $t \in [N + 1]$ for $N = 2^\ell$. In what follows, let $(\text{sk}^r, \text{params}^r)$ be the keys output by `Gen` $(1^\lambda, 1^\ell)$ when run on the sequence of coin tosses r . Let

$$f_{t,r}(\text{params}, c) = \begin{cases} 1 & \text{if } (\text{params} = \text{params}^r) \wedge (\text{Dec}(\text{sk}^r, c) \neq \perp) \wedge (\text{Dec}(\text{sk}^r, c) < t) \\ 0 & \text{otherwise.} \end{cases}$$

Notice that given t and r , the concept $f_{t,r}$ can be efficiently evaluated. The description length k of the instance space $X_k = \{0, 1\}^k$ is polynomial in the security parameter λ and plaintext length ℓ .

3.1 An Efficient PAC Learner for `EncThresh`

We argue that `EncThresh` is efficiently PAC learnable by formalizing the argument from the introduction. Because we need to include the ORE public parameters in each example, the PAC learner L (Algorithm 3) for `EncThresh` actually works in two stages. In the first stage, L determines whether there is significant probability mass on examples corresponding to some public parameters `params`. Recall that each concept in `EncThresh` is supported on exactly one such set of parameters. If there is no significant mass on any `params`, then the all-zeroes hypothesis is a good hypothesis. On the other hand, if there is a heavy set of parameters, the learner L applies `Comp` using those parameters to learn a good comparator.

Theorem 3.1. *Let $\alpha, \beta > 0$. There exists a PAC learning algorithm L for the concept class `EncThresh` achieving error α and confidence $1 - \beta$. Moreover, L is efficient (running in time polynomial in the parameters $k, 1/\alpha, \log(1/\beta)$).*

Proof. Fix a target concept $f_{t,r} \in \text{EncThresh}_k$ and a distribution \mathcal{D} on examples. First observe that the learner L always outputs a hypothesis with one-sided error, i.e. we always have $h \leq f_{t,r}$ pointwise. Also observe that $f_{t',r} \leq f_{t,r}$ pointwise for any $t' < t$. These both follow from the strong correctness of the ORE scheme. Let $(\text{sk}^r, \text{params}^r)$ denote the keys output by `Gen` $(1^\lambda, 1^\ell)$ when run on the sequence of coin tosses r . Let `POS` denote the set of examples (params, c) on which $f_{t,r}(\text{params}, c) = 1$. We divide the analysis of the learner in to two cases based on the weight \mathcal{D} places on `POS`.

Algorithm 1 Learner L for EncThresh

1. Request examples $\{(\text{params}_1, c_1, b_1), \dots, (\text{params}_n, c_n, b_n)\}$ for $n = \lceil \log(1/\beta)/\alpha \rceil$.
2. Identify an i for which $b_i = 1$ and set $\text{params}^* = \text{params}_i$; if no such i exists, return $h \equiv 0$.
3. Let $G = \{j : \text{params}_j = \text{params}^*, b_j = 1\}$. Let $j^* \in G$ be an index with $\text{Comp}(\text{params}^*, c_j, c_{j^*}) \in \{<, =, \perp\}$ for all $j \in G$.
4. Return h defined by

$$h(\text{params}, c) = \begin{cases} 1 & \text{if } (\text{params} = \text{params}^*) \wedge (\text{Comp}(\text{params}^*, c, c_{j^*}) \in \{<, =\}) \\ 0 & \text{otherwise.} \end{cases}$$

Case 1: \mathcal{D} places weight at least α on POS. Define $\hat{t} \in [N + 1]$ as the largest $\hat{t} \leq t$ such that $\text{error}_{\mathcal{D}}(f_{\hat{t},r}, f_{t,r}) \geq \alpha$. Such a \hat{t} is guaranteed to exist since $f_{0,r}$ is the all-zeros function, and therefore $\text{error}_{\mathcal{D}}(f_{0,r}, f_{t,r})$ is equal to the weight \mathcal{D} places on POS, which is at least α .

Suppose $f_{\hat{t}+1,r} \leq h$ pointwise. Since h has one-sided error (that is, $h \leq f_{t,r}$ pointwise), we have $\text{error}_{\mathcal{D}}(f_{\hat{t}+1,r}, f_{t,r}) = \text{error}_{\mathcal{D}}(f_{\hat{t}+1,r}, h) + \text{error}_{\mathcal{D}}(h, f_{t,r})$, or

$$\text{error}_{\mathcal{D}}(h, f_{t,r}) = \text{error}_{\mathcal{D}}(f_{\hat{t}+1,r}, f_{t,r}) - \text{error}_{\mathcal{D}}(f_{\hat{t}+1,r}, h) \leq \text{error}_{\mathcal{D}}(f_{\hat{t}+1,r}, f_{t,r}) < \alpha.$$

Therefore, it suffices to show that $f_{\hat{t}+1,r} \leq h$ with probability at least $1 - \beta$. This is guaranteed as long as L receives a sample $(\text{params}^r, c_i, 1)$ with $\hat{t} \leq \text{Dec}(\text{sk}^r, c_i) < t$. In other words, $f_{t,r}(\text{params}^r, c_i) = 1$ and $f_{\hat{t},r}(\text{params}^r, c_i) = 0$. Since $f_{\hat{t},r} \leq f_{t,r}$ pointwise, such samples exactly account for the error between $f_{\hat{t},r}$ and $f_{t,r}$. Thus since $\text{error}_{\mathcal{D}}(f_{\hat{t},r}, f_{t,r}) \geq \alpha$, for each i it must be that $\hat{t} \leq \text{Dec}(\text{sk}^r, c_i) < t$ with probability at least α . The learner L therefore receives *some* sample c_i with $\hat{t} \leq \text{Dec}(\text{sk}^r, c_i) < t$ with probability at least $1 - (1 - \alpha)^n \geq 1 - \beta$ (since we took $n \geq \log(1/\beta)/\alpha$).

Case 2: \mathcal{D} places less than α weight on POS. Then the identically zero hypothesis has error at most α , so the claim holds because $0 \leq h \leq f_{t,r}$. □

3.2 Hardness of Privately Learning EncThresh

We now prove the hardness of privately learning EncThresh by constructing an example reidentification scheme for this concept class. Recall that an example reidentification scheme consists of two algorithms, Gen_{ex} , which selects a distribution, a concept, and examples to give to a learner, and Trace_{ex} which attempts to identify one of the examples the learner received.

Our example reidentification scheme yields a hard distribution even for *weak-learning*, where the error parameter α is taken to be inverse-polynomially close to $1/2$.

Theorem 3.2. *Let $\gamma(n)$ and $\xi(n)$ be noticeable functions. Let $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$ be a statically single-challenge secure ORE scheme. Then there exists an (efficient) $(\alpha = \frac{1}{2} - \gamma, \xi)$ -example reidentification scheme $(\text{Gen}_{\text{ex}}, \text{Trace}_{\text{ex}})$ for the concept class EncThresh.*

We start with an informal description of the scheme $(\text{Gen}_{\text{ex}}, \text{Trace}_{\text{ex}})$. The algorithm Gen_{ex} sets up the parameters of the ORE scheme, chooses the “middle” threshold concept corresponding to $t = N/2$, and sets the distribution on examples to be encryptions of uniformly random messages (together with the correct public parameters needed for comparison). Let $m_1 < m_2 < \dots < m_n$ denote the sorted sequence of messages whose encryptions make up the sample produced by Gen_{ex} (with overwhelming probability, they are indeed distinct). We can thus break the plaintext space up into buckets of the form $B_i = [m_i, m_{i+1})$. Suppose L is a (weak) learner that produces a hypothesis h with advantage γ over random guessing. Such a hypothesis h must be able to distinguish encryptions of messages $m \leq t$ from encryptions of messages $m > t$ with advantage γ . Thus, there must be a pair of adjacent buckets B_{i-1}, B_i for which h can distinguish encryptions of messages from B_{i-1} from encryptions from B_i with advantage $\frac{\gamma}{n}$.

This observation leads to a natural definition for Trace_{ex} : locate a pair of adjacent buckets B_{i-1}, B_i that h distinguishes, and output the identity i of the example separating those buckets. Completeness of the resulting scheme, i.e. the fact that some example is reidentified when L succeeds, follows immediately from the preceding discussion. We argue soundness, i.e. that an example absent from L 's sample is not identified, by reducing to the static security of the ORE scheme. The intuition is that if L is not given example i , then it should not be able to distinguish encryptions from bucket B_{i-1} from encryptions from bucket B_i .

To make the security reduction somewhat more precise, suppose for the sake of contradiction that there is an efficient algorithm L that violates the soundness of $(\text{Gen}_{\text{ex}}, \text{Trace}_{\text{ex}})$ with noticeable probability ξ . That is, there is some i such that even without example i , the algorithm L manages to produce (with probability ξ) a hypothesis h that distinguishes B_{i-1} from B_i . A natural first attempt to violate the security of the ORE is to construct an adversary that challenges on the message sequences $m_1 < \dots < m_{i-1} < m_i^{(L)} < m_{i+1}, \dots, m_n$ and $m_1 < \dots < m_{i-1} < m_i^{(R)} < m_{i+1} < \dots < m_n$, where $m_i^{(L)}$ is randomly chosen from B_{i-1} and $m_i^{(R)}$ is randomly chosen from B_i . Then if h can distinguish B_{i-1} from B_i , the adversary can distinguish the two sequences. Unfortunately, this approach fails for a somewhat subtle reason. The hypothesis h is only guaranteed to distinguish B_{i-1} from B_i with probability ξ . If h fails to distinguish the buckets – or distinguishes them in the opposite direction – then the adversary's advantage is lost.

To overcome this issue, we instead rely on the security of the ORE for sequences that differ on *two* messages. For the “left” challenge, our adversary samples two messages from the same randomly chosen bucket, B_{i-1} or B_i (in addition to requesting encryptions of $m_1, \dots, m_{i-1}, m_i, \dots, m_n$). For the “right” challenge, it samples one message from each bucket B_{i-1} and B_i . Let c^0 and c^1 be the ciphertexts corresponding to these challenge messages. If h agrees on c^0 and c^1 , then this suggests the messages are from the same bucket, and the adversary should guess “left”. On the other hand, if h disagrees on c^0 and c^1 , then the adversary should guess “right”. If h distinguishes the buckets B_{i-1} and B_i , this adversary does strictly better than random guessing. On the other hand, even if h fails to distinguish the buckets, the adversary does at least as well as random guessing. So overall, it still has a noticeable advantage at the ORE security game.

We now give the formal proof of Theorem 3.2.

Proof. We construct an example reidentification scheme for EncThresh as follows. The algorithm Gen_{ex} fixes the threshold $t = N/2$ and samples $(\text{sk}^r, \text{params}^r) \leftarrow_{\text{R}} \text{Gen}(1^\lambda, 1^\ell)$, yielding a concept $f_{t,r}$. Let \mathcal{D} be the distribution of $(\text{params}^r, \text{Enc}(\text{sk}^r, m))$ for uniformly random $m \in [N]$. Let $m'_1, \dots, m'_n \leftarrow_{\text{R}} [N]$, and let $m_1 \leq \dots \leq m_n$ be the result of sorting the m'_i . Let $m_0 = 0$ and

$m_{n+1} = N$. Since $n = \text{poly}(k) \ll N$, these random messages will be well-spaced. In particular, with overwhelming probability, $|m_{i+1} - m_i| > 1$ for every i , so we assume this is the case in what follows. Gen_{ex} then sets the samples to be $(x_1 = (\text{params}^r, \text{Enc}(\text{sk}^r, m'_1)), \dots, x_n = (\text{params}^r, \text{Enc}(\text{sk}^r, m'_n)))$. Let $x_0 = (\text{params}^r, \text{Enc}(\text{sk}^r, m_0))$ be a “junk” example.

The algorithm Trace_{ex} creates buckets $B_i = [m_i, m_{i+1})$. For each i , let

$$p_i = \Pr_{m \in B_i, \text{coins of Enc}} [h(\text{params}^r, \text{Enc}(\text{sk}, m)) = 1].$$

By sampling random choices of m in each bucket, Trace_{ex} can efficiently compute a good estimate $\hat{p}_i \approx p_i$ for each i (Lemma 3.3). It then accuses the least i for which $\hat{p}_{i-1} - \hat{p}_i \geq \frac{\gamma}{n}$, and \perp if none is found.

Lemma 3.3. *Let $K = \frac{8n^2}{\gamma^2} \log(9n/\xi)$. For each $i = 0, \dots, n$, let*

$$\hat{p}_i = \frac{1}{K} \sum_{j=1}^K h(x_j)$$

where $x_j = (\text{params}^r, \text{Enc}(\text{sk}^r, m_j))$ for i.i.d. $m_1, \dots, m_K \leftarrow_R B_i$. Then $|\hat{p}_i - p_i| \leq \frac{\gamma}{4n}$ for every i with probability at least $1 - \xi/4$.

Proof. By a Chernoff bound, the probability that any given \hat{p}_i deviates from p_i by more than $\frac{\gamma}{4n}$ is at most $2 \exp(-K\gamma^2/8n^2) \leq \frac{\xi}{4(n+1)}$. The lemma follows by a union bound. \square

We first verify completeness for this scheme. Let L be a learner for EncThresh using n examples. If the hypothesis h produced by L is $(\frac{1}{2} - \gamma)$ -good, then there exists $i_0 < i_1$ such that $p_{i_0} - p_{i_1} \geq 2\gamma$. If this is the case, then there must be an i for which $p_{i-1} - p_i \geq \frac{2\gamma}{n}$. Then with probability all but $\xi(n)/2$ over the estimates \hat{p}_i , we have $\hat{p}_{i-1} - \hat{p}_i \geq \frac{\gamma}{n}$, so some index is accused.

Now we verify soundness. Fix a PPT L , and let $j^* \in [n]$. Suppose L violates the soundness of the scheme with respect to j^* , i.e.

$$\Pr_{h \leftarrow_R L(S_{-j^*}), \text{coins of Gen}_{\text{ex}}} [\text{Trace}_{\text{ex}}(h) = j^*] > \xi.$$

We will use L to construct an adversary \mathcal{A} for the ORE scheme that succeeds with noticeable advantage. It suffices to build an adversary for the static (many-challenge) security of ORE, with Theorem 2.9 showing how to convert it to a single-challenge adversary. This many-challenge adversary is presented as Algorithm 2. (While not explicitly stated, the adversary should halt and output a random guess whenever the messages it samples are not well-spaced.)

Let i^* be such that $m_{i^*} = m'_{j^*}$. With probability at least ξ over the parameters $(\text{sk}^r, \text{params}^r)$, the choice of messages, the choice of the hypothesis h , and the coins of Trace_{ex} , there is a gap $\hat{p}_{i^*-1} - \hat{p}_{i^*} \geq \frac{\gamma}{n}$. Hence, by Lemma 3.3, there is a gap $p_{i^*-1} - p_{i^*} \geq \frac{\gamma}{2n}$ with probability at least $\frac{\xi}{2}$.

We now calculate the advantage of the adversary \mathcal{A} . Fix a hypothesis h . For notational simplicity, let $p = p_{i^*-1}$ and let $q = p_{i^*}$. Let $y_0 = h(\text{params}^r, c_{i^*}^0)$ and $y_1 = h(\text{params}^r, c_{i^*}^1)$. Then the adversary’s success probability is:

Algorithm 2 ORE adversary \mathcal{A}

1. Sample $m'_1, \dots, m'_n \leftarrow_{\text{R}} [N]$, and let $m_1 \leq \dots \leq m_n$ be the result of sorting the m'_j . Let π be the permutation on $\{1, \dots, n\}$ such that $m_{\pi(j)} = m'_j$. Let $m_0 = 0$. Let $i^* = \pi(j^*)$ so that $m_{i^*} = m'_{j^*}$.
2. Construct pairs (m_L^0, m_L^1) and (m_R^0, m_R^1) as follows. Let $B_0 = (m_{i^*-1}, m_{i^*})$ and $B_1 = (m_{i^*}, m_{i^*+1})$. Sample $m_L^0 \leq m_L^1$ at random from the same B_j , for a random choice of $j \in \{0, 1\}$. Sample $m_R^0 \leftarrow_{\text{R}} B_0$ and $m_R^1 \leftarrow_{\text{R}} B_1$.
3. Challenge on the pair of sequences $m_0, m_1, \dots, m_{i^*-1}, m_L^1, m_L^2, m_{i^*}, \dots, m_n$ and $m_0, m_1, \dots, m_{i^*-1}, m_R^1, m_R^2, m_{i^*}, \dots, m_n$, receiving ciphertexts $c_1, \dots, c_{i^*}^0, c_{i^*}^1, \dots, c_n$. For $j \neq j^*$, let $c'_j = c_{\pi(j)}$ so that c'_j is an encryption of m'_j .
4. Set $t = N/2$ and let

$$\begin{aligned}
S_{-j^*} &= \{(\text{params}^r, c'_1, \chi(m'_1 \leq t)), \dots, (\text{params}^r, c'_{j^*-1}, \chi(m'_{j^*-1} \leq t)), \\
&\quad (\text{params}^r, c_0, 1), (\text{params}^r, c'_{j^*+1}, \chi(m'_{j^*+1} \leq t)), \dots, (\text{params}^r, c'_n, \chi(m'_n \leq t))\} \\
&= \{(\text{params}^r, c_{\pi(1)}, \chi(m_{\pi(1)} \leq t)), \dots, (\text{params}^r, c_{\pi(j^*-1)}, \chi(m_{\pi(j^*-1)} \leq t)), \\
&\quad (\text{params}^r, c_0, 1), (\text{params}^r, c_{\pi(j^*+1)}, \chi(m_{\pi(j^*+1)} \leq t)), \dots, (\text{params}^r, c_{\pi(n)}, \chi(m_{\pi(n)} \leq t))\}
\end{aligned}$$

Obtain $h \leftarrow_{\text{R}} L(S_{-j^*})$.

5. Guess $b' = 0$ if $h(\text{params}^r, c_{i^*}^0) = h(\text{params}^r, c_{i^*}^1)$. Otherwise guess $b' = 1$.
-

$$\begin{aligned}
\Pr[b' = b] &= \frac{1}{2}(\Pr[y_0 = y_1 | b = 0] + \Pr[y_0 \neq y_1 | b = 1]) \\
&= \frac{1}{2}\left(\frac{1}{2}(p^2 + (1-p)^2 + q^2 + (1-q)^2) + (1-pq - (1-p)(1-q))\right) \\
&= \frac{1}{2} + \frac{1}{2}(p-q)^2.
\end{aligned}$$

Thus if $p - q \geq \frac{\gamma}{2n}$, then the adversary's advantage is at least $\frac{\gamma^2}{4n^2}$. On the other hand, even for arbitrary values of p, q , the advantage is still nonnegative. Therefore, the advantage of the strategy is at least $\frac{\xi\gamma^2}{8n^2} - \text{negl}(k)$ (the $\text{negl}(k)$ term coming from the assumption that the m'_i sampled were distinct), which is a noticeable function of the parameter k . This contradicts the static security of the ORE scheme. \square

3.3 The SQ Learnability of EncThresh

The statistical query (SQ) model is a natural restriction of the PAC model by which a learner is able to measure statistical properties of its examples, but cannot see the individual examples themselves. We recall the definition of an SQ learner.

Definition 3.4 (SQ learning [Kea98]). Let $c : X \rightarrow \{0, 1\}$ be a target concept and let \mathcal{D} be a distribution over X . In the SQ model, a learner is given access to a *statistical query oracle* $\text{STAT}(c, \mathcal{D})$. It may make queries to this oracle of the form (ψ, τ) , where $\psi : X \times \{0, 1\} \rightarrow \{0, 1\}$ is a query function and $\tau \in (0, 1)$ is an error tolerance. The oracle $\text{STAT}(c, \mathcal{D})$ responds with a value v such that $|v - \Pr_{x \in \mathcal{D}}[\psi(x, c(x)) = 1]| \leq \tau$. The goal of a learner is to produce, with probability at least $1 - \beta$, a hypothesis $h : X \rightarrow \{0, 1\}$ such that $\text{error}_{\mathcal{D}}(c, h) \leq \alpha$. The query functions must be efficiently evaluable, and the tolerance τ must be lower bounded by an inverse polynomial in k and $1/\alpha$.

The *query complexity* of a learner is the worst-case number of queries it issues to the statistical query oracle. An SQ learner is efficient if it also runs in time polynomial in $k, 1/\alpha, 1/\beta$.

Feldman and Kanade [FK12] investigated the relationship between query complexity and computational complexity for SQ learners. They exhibited a concept class \mathcal{C} which is efficiently PAC learnable and SQ learnable with polynomially many queries, but assuming $\mathbf{NP} \neq \mathbf{RP}$, is not efficiently SQ learnable. Concepts in this concept class take the form

$$g_{\phi, y}(x, x') = \begin{cases} \text{PAR}_y(x') & \text{if } x = \phi \\ 0 & \text{otherwise.} \end{cases}$$

Here, $\text{PAR}_y(x')$ is the inner product of y and x' modulo 2. The concept class \mathcal{C} consists of $g_{\phi, y}$ where ϕ is a satisfiable 3-CNF formula and y is the lexicographically first satisfying assignment to ϕ . The efficient PAC learner for parities based on Gaussian elimination shows that \mathcal{C} is also efficiently PAC learnable. It is also (inefficiently) SQ learnable with polynomially many queries: either the all-zeroes hypothesis is good, or an SQ learner can recover the formula ϕ bit-by-bit and determine the satisfying assignment y by brute force. On the other hand, because parities are information-theoretically hard to SQ learn, the satisfying assignment y remains hidden to an SQ learner unless it is able to solve 3-SAT.

In this section, we show that the concept class EncThresh shares these properties with \mathcal{C} . Namely, we know that EncThresh is efficiently PAC learnable and because it is not efficiently privately learnable, it is not efficiently SQ learnable [BDMN05]. We can also show that EncThresh has an SQ learner with polynomial query complexity. Making this observation about EncThresh is of interest because the hardness of SQ learning EncThresh does not seem to be related to the (information-theoretic) hardness of SQ learning parities.

Proposition 3.5. *The concept class EncThresh is (inefficiently) SQ learnable with polynomially many queries.*

As with \mathcal{C} there are two cases. In the first case, the target distribution places nearly zero weight on examples with $\text{params} = \text{params}^r$, and so the all-zeroes hypothesis is good. In the second case, the target distribution places noticeable weight on these examples, and our learner can use statistical queries to recover the comparison parameters params^r bit-by-bit. Once the public parameters are recovered, our learner can determine a corresponding secret key by brute force. Lemma 3.6 below shows that any corresponding secret key – even one that is not actually sk^r – suffices. The learner can then use binary search to determine the threshold value t .

Proof. Let $f_{t,r}$ be the target concept, \mathcal{D} be the target distribution, and α be the target error rate. With the statistical query $(x \times b \mapsto b, \alpha/4)$, we can determine whether the all-zeroes hypothesis is

accurate. That is, if we receive a value that is less than $\alpha/2$, then $\Pr_{x \in \mathcal{D}}[f_{t,r}(x) = 1] \leq \alpha$. If not, then we know that $\Pr_{x \in \mathcal{D}}[f_{t,r}(x) = 1] \geq \alpha/4$, so \mathcal{D} places significant weight on examples prefixed with params^r . Suppose now that we are in the latter case.

Let $m = |\text{params}|$. For $i = 1, \dots, m$, define $\psi_i(\text{params}, c, b) = 1$ if $\text{params}_i = 1$ and $b = 1$, and $\psi_i(\text{params}, c, b) = 0$ otherwise. Then by asking the queries $(\psi_i, \alpha/16)$, we can determine each bit params_i^r of params^r .

Now by brute force search, we determine a secret key sk for which $(\text{sk}, \text{params}^r) \in \text{Range}(\text{Gen})$. The recovered secret key sk may not necessarily be the same as sk^r . However, the following lemma shows that sk and sk^r are functionally equivalent:

Lemma 3.6. *Suppose $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$ is a strongly correct ORE scheme. Then for any pair $(\text{sk}_1, \text{params}), (\text{sk}_2, \text{params}) \in \text{Range}(\text{Gen})$, we have that $\text{Dec}_{\text{sk}_1}(c) = \text{Dec}_{\text{sk}_2}(c)$ for all ciphertexts c .*

With the secret key sk in hand, we now conduct a binary search for the threshold t . Recall that we have an estimate v for the weight that $f_{t,r}$ places on positive examples, i.e. $|v - \Pr_{x \in \mathcal{D}}[f_{t,r}(x) = 1]| \leq \alpha/4$. Starting at $t_1 = N/2$, we issue the query $(\varphi_1, \alpha/4)$ where $\varphi_1(\text{params}, c, b) = 1$ iff $\text{params} = \text{params}^r$ and $\text{Dec}(\text{sk}, c) < t$. Let h_{t_1} denote the hypothesis

$$h_{t_1}(\text{params}, c) = \begin{cases} 1 & \text{if } (\text{params} = \text{params}^r) \wedge (\text{Dec}(\text{sk}, c) \neq \perp) \wedge (\text{Dec}(\text{sk}, c) < t_1) \\ 0 & \text{otherwise.} \end{cases}$$

Thus, the query $(\varphi_1, \alpha/4)$ approximates the weight h_{t_1} places on positive examples. Let the answer to this query be v_1 . If $|v_1 - v| \leq \alpha/2$, then we can halt and output the good hypothesis h_{t_1} . Otherwise, if $v_1 < v - \alpha/2$, we set the next threshold to $t_2 = 3N/4$, and if $v_1 > v + \alpha/2$, we set the next threshold to $t_2 = N/4$. We recurse up to $\log N = \ell = \text{poly}(k)$ times, yielding a good hypothesis for $f_{t,r}$. \square

Proof of Lemma 3.6. Suppose the lemma is not true. First suppose that there exists a ciphertext c such that $\text{Dec}(\text{sk}_1, c) = p_1 < p_2 = \text{Dec}(\text{sk}_2, c)$. Let $c' \in \text{Enc}(\text{sk}_1, p_2)$. Then by strong correctness applied to the parameters $(\text{sk}_1, \text{params})$, we must have $\text{Comp}(\text{params}, c, c') = "<".$ Now by strong correctness applied to $(\text{sk}_2, \text{params})$, we must have $\text{Dec}(\text{sk}_2, c') > p_2$. Thus, $p_1 < \text{Dec}(\text{sk}_1, c') = p_2 < \text{Dec}(\text{sk}_2, c')$. Repeating this argument, we obtain a contradiction because the message space is finite.

Now suppose instead that there is a ciphertext c for which $\text{Dec}(\text{sk}_1, c) = p \in [N]$, but $\text{Dec}(\text{sk}_2, c) = \perp$. Let $c' \in \text{Enc}(\text{sk}_1, p')$ for some $p' > p$. Then $\text{Comp}(\text{params}, c, c') = "<"$ by strong correctness applied to $(\text{params}, \text{sk}_1)$. But $\text{Comp}(\text{params}, c, c') = "\perp"$ by strong correctness applied to $(\text{params}, \text{sk}_2)$, again yielding a contradiction. \square

4 ORE with Strong Correctness

We now explain how to obtain ORE with strongly correct comparison, as all prior ORE schemes only satisfy the weaker notion of correctness. The lack of strong correctness is easiest to see with the scheme of Boneh et al. [BLR⁺15]. The protocol is built from current multilinear map constructions, which are noisy. If the noise terms grow too large, the correctness of the multilinear map is not guaranteed. The comparison function in [BLR⁺15] is computed by performing multilinear operations, and for correctly generated ciphertexts, the operations will give the right answer. However, there exist ciphertexts, namely those with very large noise, for which the comparison function

gives an incorrect output. The result is that the comparison operation is not guaranteed to be consistent with decrypting the ciphertexts and comparing the plaintexts.

As described in the introduction, we give a generic conversion from any ORE scheme with weakly correct comparison into a strongly correct scheme. We simply modify the encryption algorithm by adding a non-interactive zero-knowledge (NIZK) proof that the resulting ciphertext is well-formed. Then the decryption and comparison procedures check the proof(s), and only output a non- \perp result (either decryption or comparison) if the proof(s) are valid.

Instantiating our scheme. In our construction, we need the (weak) correctness of the underlying ORE scheme to hold with probability one. However, the existing protocols only have correctness with overwhelming probability, so some minor adjustments need to be made to the protocols. This is easiest to see in the ORE scheme of Boneh et al. [BLR⁺15]. The Boneh et al. scheme uses noisy multilinear maps [GGH13a] which may introduce errors. Therefore, the protocol described in [BLR⁺15] only achieves the (weak) correctness property with overwhelming probability, whereas we will require (weak) correctness with probability 1 for the conversion. However, it is straightforward to generate the parameters for the protocol in such a way as to completely eliminate errors. Essentially, the parameters in the protocol have an error term that is generated by a (discrete) Gaussian distribution, which has unbounded support. Instead, we truncate the Gaussian, resulting in a noise distribution with bounded support. By truncating sufficiently far from the center, the resulting distribution is also statistically close to the full Gaussian, so security of the protocol with truncated noise follows from the security of the protocol with un-truncated noise. By truncating the noise distribution, it is straightforward to set parameters so that no errors can occur.

It is similarly straightforward to modify current obfuscation candidates, which are also built from multilinear maps, to obtain perfect (weak) correctness by truncating the noise distributions. Thus, our scheme has instantiations using multilinear maps or iO.

4.1 Conversion from Weakly Correct ORE

We describe our generic conversion from an order-revealing encryption scheme with weak correctness using NIZKs. We will need the following additional tools:

Perfectly binding commitments. A perfectly binding commitment Com is a randomized algorithm with two properties. The first is perfect binding, which states that if $\text{Com}(m; r) = \text{Com}(m'; r')$, then $m = m'$. The second requirement is computational hiding, which states that the distributions $\text{Com}(m)$ and $\text{Com}(m')$ are computationally indistinguishable for any messages m, m' . Such commitments can be built, say, from any injective one-way function.

Perfectly sound NIZK. A NIZK protocol consists of three algorithms:

- $\text{Setup}(1^\lambda)$ is a randomized algorithm that outputs a common reference string crs .
- $\text{Prove}(\text{crs}, x, w)$ takes as input a common reference string crs , an NP statement x , and a witness w , and produces a proof π .
- $\text{Ver}(\text{crs}, x, \pi)$ takes as input a common reference string crs , statement x , and a proof π , and outputs either `accept` or `reject`.

We make three requirements for a NIZK:

- **Perfect Completeness.** For all security parameters λ and any true statement x with witness w ,

$$\Pr[\text{Ver}(\text{crs}, x, \pi) = \text{accept} : \text{crs} \leftarrow \text{Setup}(1^\lambda); \pi \leftarrow \text{Prove}(\text{crs}, x, w)] = 1.$$

- **Perfect Soundness.** For all security parameters λ , any *false* statement x and any (invalid) proof π ,

$$\Pr[\text{Ver}(\text{crs}, x, \pi) = \text{accept} : \text{crs} \leftarrow \text{Setup}(1^\lambda)] = 0.$$

- **Computational Zero Knowledge.** There exists a simulator $\mathcal{S}_1, \mathcal{S}_2$ such that for any computationally bounded adversary \mathcal{A} , the quantity

$$\| \Pr[\mathcal{A}^{\text{Prove}(\text{crs}, \cdot, \cdot)}(\text{crs}) = 1 : \text{crs} \leftarrow \text{Setup}(1^\lambda)] - \Pr[\mathcal{A}^{\text{Sim}(\text{crs}, \tau, \cdot, \cdot)}(\text{crs}) = 1 : (\text{crs}, \tau) \leftarrow \mathcal{S}_1(1^\lambda)] \|$$

is negligible, where $\text{Sim}(\text{crs}, \tau, x, w)$ outputs $\mathcal{S}_2(\text{crs}, \tau, x)$ if w is a valid witness for x , and $\text{Sim}(\text{crs}, \tau, x, w) = \perp$ if w is invalid.

NIZKs satisfying these requirements can be built from bilinear maps [GOS12].

4.1.1 The Construction

We now give our conversion. Let $(\text{Setup}, \text{Prove}, \text{Ver})$ be a perfectly sound NIZK and $(\text{Gen}', \text{Enc}', \text{Dec}', \text{Comp}')$ and ORE with *weakly* correct comparison. We will assume that Enc' is deterministic; if not, we can derandomize Enc' using a pseudorandom function. Let Com be a perfectly binding commitment. We construct a new ORE scheme $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$ with *strongly* correct comparison:

- $\text{Gen}(1^\lambda, 1^\ell)$: run $(\text{sk}', \text{params}') \leftarrow \text{Gen}'(1^\lambda, 1^\ell)$. Let $\sigma = \text{Com}(\text{sk}; r)$ for randomness r , and run $\text{crs} \leftarrow \text{Setup}(1^\lambda)$. Then the secret key is $\text{sk} = (\text{sk}', r, \text{crs})$ and the public parameters are $\text{params} = (\text{params}', \sigma, \text{crs})$.
- $\text{Enc}(\text{sk}, m)$: Compute $c' = \text{Enc}'(\text{sk}', m)$. Let $x_{c'}$ be the statement $\exists \hat{m}, \hat{\text{sk}}', \hat{r} : \sigma = \text{Com}(\hat{\text{sk}}', \hat{r}) \wedge c' = \text{Enc}'(\hat{\text{sk}}', \hat{m})$. Run $\pi_{c'} = \text{Prove}(\text{crs}, x_{c'}, (m, \text{sk}', r))$. Output the ciphertext $c = (c', \pi_{c'})$.
- $\text{Dec}(\text{sk}, c)$: Write $c = (c', \pi_{c'})$. If $\text{Ver}(\text{crs}, x_{c'}, \pi_{c'}) = \text{reject}$, output \perp . Otherwise, output $m = \text{Dec}'(\text{sk}', c')$.
- $\text{Comp}(\text{params}, c_0, c_1)$: write $c_b = (c'_b, \pi_{c'_b})$ and $\text{params} = (\text{params}', \sigma, \text{crs})$. If $\text{Ver}(\text{crs}, x_{c'_b}, \pi_{c'_b}) = \text{reject}$ for either $b = 0, 1$, then output \perp . Otherwise, output $\text{Comp}'(\text{params}', c'_0, c'_1)$.

Correctness. Notice that, for each plaintext m , the ciphertext component $c' = \text{Enc}'(\text{sk}', m)$ is the *unique* value such that $\text{Dec}(\text{sk}, (c', \pi)) = m$ for some proof π . Moreover, the completeness of the zero knowledge proof implies that $\text{Enc}(\text{sk}, m)$ outputs a valid proof. Decryption correctness follows.

For strong comparison correctness, consider two ciphertexts c_0, c_1 where $c_b = (c'_b, \pi_{c'_b})$. Suppose both proofs $\pi_{c'_b}$ are valid, which means that verification passes when running Comp and so

$\text{Comp}(\text{params}, c_0, c_1) = \text{Comp}'(\text{params}', c'_0, c'_1)$. Verification also passes when decrypting c_b , and so $\text{Dec}(\text{sk}, c_b) = \text{Dec}'(\text{sk}', c'_b)$.

Since the proofs are valid, $c'_b = \text{Enc}'(\text{sk}', m_b)$ for some m_b for both $b = 0, 1$. The weak correctness of comparison for $(\text{Gen}', \text{Enc}', \text{Dec}', \text{Comp}')$ implies that $\text{Comp}'(\text{params}', c'_0, c'_1) = \text{Comp}_{\text{plain}}(m_0, m_1)$. The decryption correctness of $(\text{Gen}', \text{Enc}', \text{Dec}', \text{Comp}')$ then implies that $\text{Dec}(\text{sk}', c'_b) = m_b$, and therefore $\text{Dec}(\text{sk}, c_b) = m_b$. Thus $\text{Comp}_{\text{ciph}}(\text{sk}, c_0, c_1) = \text{Comp}_{\text{plain}}(m_0, m_1)$. Putting it all together, $\text{Comp}(\text{params}, c_0, c_1) = \text{Comp}_{\text{ciph}}(\text{sk}, c_0, c_1)$, as desired.

Now suppose one of the proofs $\pi_{c'_b}$ are invalid. Then $\text{Comp}(\text{params}, c_0, c_1) = \perp$ and $\text{Dec}(\text{sk}, c_b) = \perp$. This means $\text{Comp}_{\text{ciph}}(\text{sk}, c_0, c_1) = \perp = \text{Comp}(\text{params}, c_0, c_1)$, as desired.

Security. To prove security, we first use the zero-knowledge simulator to simulate the proofs π'_c without using a witness (namely, the secret decryption key). Then we use the hiding property of the commitment to replace σ with a commitment to 0. At this point, the entire game can be simulated using an Enc' oracle, and so the security reduces to the security of Enc' .

Theorem 4.1. *If $(\text{Gen}', \text{Enc}', \text{Dec}', \text{Comp}')$ is a (statically) secure ORE, $(\text{Setup}, \text{Prove}, \text{Ver})$ is computationally zero knowledge, and Com is computationally hiding, then $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$ is a statically secure ORE.*

Proof. We will prove security through a sequence of hybrids. Let \mathcal{A} be an adversary with advantage ϵ in breaking the static security of $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$.

Hybrid 0. This is the real experiment, where $\sigma \leftarrow \text{Com}(\text{sk})$, $\text{crs} \leftarrow \text{Setup}(1^\lambda)$, and the proofs $\pi_{c'}$ are answered using Prove and valid witnesses. \mathcal{A} has advantage ϵ in distinguishing the left and right ciphertexts.

Hybrid 1. This is the same as **Hybrid 0**, except that crs is generated as $(\text{crs}, \tau) \leftarrow \mathcal{S}_1(1^\lambda)$, and all proofs are generated using $\mathcal{S}_2(\text{crs}, \tau, \cdot)$. The zero knowledge property of $(\text{Setup}, \text{Prove}, \text{Ver})$ shows that this is indistinguishable from **Hybrid 0**.

Hybrid 2. This is the same as **Hybrid 1**, except that $\sigma \leftarrow \text{Com}(0)$. Since the randomness for computing σ is not needed for simulation, this change is undetectable using the hiding of Com .

Thus the advantage of \mathcal{A} in **Hybrid 2** is at least $\epsilon - \text{negl}$ for some negligible function negl . Now consider the following adversary $c\mathcal{B}$ that attempts to break the security of $(\text{Gen}', \text{Enc}', \text{Dec}', \text{Comp}')$. \mathcal{B} simulates \mathcal{A} , and forwards the message sequences $m_1^{(L)} < m_2^{(L)} < \dots < m_q^{(L)}$ and $m_1^{(R)} < m_2^{(R)} < \dots < m_q^{(R)}$ produced by \mathcal{A} to its own challenger. In response, it receives params' , and ciphertexts c'_i , where c'_i encrypts either $m_i^{(L)}$ if $b = 0$ or $m_i^{(R)}$ if $b = 1$, for a random bit b chosen by the challenger.

\mathcal{B} now generates $\sigma \leftarrow \text{Com}(0)$ and $(\text{crs}, \tau) \leftarrow \mathcal{S}_1(1^\lambda)$, and lets $\text{params} = (\text{params}', \sigma, \text{crs})$. It also computes $\pi_{c'_i} \leftarrow \mathcal{S}_2(\text{crs}, \tau, x_{c'_i})$, and defines $c_i = (c'_i, \pi_{c'_i})$, and gives params and the c_i to \mathcal{A} . Finally when \mathcal{A} outputs a guess b' for b , \mathcal{B} outputs the same guess b' .

We see that the view of \mathcal{A} as a subroutine of \mathcal{B} is exactly the same view as in **Hybrid 2**. Thus, $b' = b$ with probability at least $\epsilon - \text{negl}$. The security of $(\text{Gen}', \text{Enc}', \text{Dec}', \text{Comp}')$ implies that this quantity, and hence ϵ , must be negligible. Thus \mathcal{A} must have negligible advantage in breaking the security of $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Comp})$. □

5 A Separation for Representation Learning

In this section, we show how to construct a concept class `ValidSig` that separates efficient *representation* learning from efficient private representation learning, assuming only the existence of one-way functions. Here by “representation learning” we mean a restricted form of proper learning where a learner must output a particular representation (i.e. encoding) of a hypothesis h in the concept class \mathcal{C} . As with proper learning, this is a natural syntactic restriction to place on a learner: for instance, if one wants to learn linear threshold functions (LTF), it makes sense to require a learner to produce the actual coefficients of an LTF, rather than an arbitrary circuit that happens to compute an LTF.

The construction is based on the following elegant idea due to Kobbi Nissim [Nis14]. Suppose $H : D \rightarrow R$ is a cryptographic hash function with the property that given x_1, \dots, x_n with $y = H(x_1) = \dots = H(x_n)$, it is infeasible for an efficient adversary to find another x for which $H(x) = y$. Consider the concept class `HashPoint` consisting of the concepts

$$f_x(x') = \begin{cases} 1 & \text{if } H(x) = H(x') \\ 0 & \text{otherwise.} \end{cases}$$

for every $x \in R$. The representation of a concept f_x is the point x . The concept class `HashPoint` is very easy to learn (by representation) without privacy: a learner can identify any positive example x_i and output the representation x_i . Since $H(x_i) = H(x)$, the concept f_{x_i} is actually equal to the target concept f_x . On the other hand, a learner that identifies an index x^* for which $f_{x^*} = f_x$ cannot be differentially private, since the security of the hash function means it is infeasible to produce such an x^* that is not present in the sample.

Note that this argument breaks down if one tries to show that `HashPoint` is not privately properly learnable. While it is infeasible to privately produce a representation x^* for which f_{x^*} is a good hypothesis, the hypothesis $h(x) = \chi(H(x) = h(x_i))$ is equal as a function to every good f_{x^*} . Moreover, this hypothesis can be constructed privately as long as the sample contains sufficiently many positive examples.

We make this discussion formal by constructing a concept class `ValidSig` based on *super-secure digital signature schemes*, which can be constructed from one-way functions. Our use of signatures to derive hardness results for private proper learning is very analogous to prior hardness results for synthetic data generation [DNR⁺09, UV11].

Definition 5.1. A *digital signature scheme* is a triple of algorithms $(\text{Gen}, \text{Sign}, \text{Ver})$ where

- $\text{Gen}(1^\lambda)$ produces a key pair (sk, vk) .
- $\text{Sign}(\text{sk}, m)$ takes the private signing key sk and a message $m \in \{0, 1\}^*$ and produces a signature σ for the message m .
- $\text{Ver}(\text{vk}, m, \sigma)$ takes the public verification key vk , a message m , and a signature σ , and (deterministically) outputs a bit indicating whether σ is a valid signature for m .

The correctness property of a digital signature scheme is that for every $(\text{sk}, \text{vk}) \leftarrow_{\mathcal{R}} \text{Gen}(1^\lambda)$, every message $m \in \{0, 1\}^*$, and every signature $\sigma \leftarrow_{\mathcal{R}} \text{Sign}(\text{sk}, m)$, we have $\text{Ver}(\text{vk}, m, \sigma) = 1$.

Definition 5.2. A digital signature scheme is *super-secure under adaptive chosen-plaintext attacks* if all efficient adversaries \mathcal{A} win the following weak forgery game with negligible probability:

- The challenger samples $(\text{sk}, \text{vk}) \leftarrow_{\text{R}} \text{Gen}(1^\lambda)$.
- The adversary \mathcal{A} is given vk and oracle access to $\text{Sign}(\text{sk}, \cdot)$. It adaptively queries the signing oracle, obtaining a sequence of message-signature pairs A . It then outputs a forgery (m^*, σ^*) .
- The value of the game is 1 iff $\text{Ver}(\text{vk}, m^*, \sigma^*) = 1$ and $(m^*, \sigma^*) \notin A$.

It is known that super-secure digital signature schemes can be constructed from one-way functions [NY89, Rom90, KK05, Gol04].

We now describe our concept class ValidSig . Let $(\text{Gen}, \text{Sign}, \text{Ver})$ be a super-secure digital signature scheme. We define a concept class ValidSig as follows. Fix the message length ℓ . For every (vk, m, σ) with $m \in \{0, 1\}^\ell$ and $\text{Ver}(\text{vk}, m, \sigma) = 1$, define the concept

$$f_{\text{vk}, m, \sigma}(\text{vk}', m', \sigma') = \begin{cases} 1 & \text{if } (\text{vk} = \text{vk}') \wedge (\text{Ver}(\text{vk}, m', \sigma') = 1) \\ 0 & \text{otherwise.} \end{cases}$$

For convenience, we also include the all-zeroes hypothesis in ValidSig , with representation \perp .

Theorem 5.3. *Let $\alpha, \beta > 0$. There exists a proper PAC learning algorithm L for the concept class ValidSig achieving error α and confidence $1 - \beta$. Moreover, L is efficient (running in time polynomial in the parameters $k, 1/\alpha, \log(1/\beta)$).*

Algorithm 3 Learner L for ValidSig

1. Request examples $\{((\text{vk}'_1, m'_1, \sigma'_1), b_1), \dots, ((\text{vk}'_n, m'_n, \sigma'_n), b_n)\}$ for $n = \lceil \log(1/\beta)/\alpha \rceil$.
 2. Identify an i for which $b_i = 1$ and return the representation $(\text{vk}'_i, m'_i, \sigma'_i)$. If no such i exists, return \perp representing the all-zeroes hypothesis.
-

Proof. Fix a target concept $f_{\text{vk}, m, \sigma} \in \text{ValidSig}_k$ and a distribution \mathcal{D} on examples. Let POS denote the set of examples $(\text{vk}', m', \sigma')$ on which $f_{\text{vk}, m, \sigma}(\text{vk}', m', \sigma') = 1$. We divide the analysis of the learner into three cases based on the weight \mathcal{D} places on the sets POS .

Case 1: \mathcal{D} places at least α weight on POS . Then L receives a positive example with probability at least $1 - (1 - \alpha)^n \geq 1 - \beta$, and is thus able to identify a concept that equals the target concept.

Case 2: \mathcal{D} places less than α weight on POS . If L gets a positive example, then the analysis of Case 1 applies. Otherwise, the all-zeroes hypothesis is α -good. □

We now prove the hardness of properly privately learning ValidSig by constructing an example reidentification scheme for properly learning this concept class. Our example reidentification scheme yields a hard distribution even when the error parameter α is taken to be inverse-polynomially close to 1.

Theorem 5.4. *Let $\gamma(n)$ and $\xi(n)$ be noticeable functions. Let $(\text{Gen}, \text{Sign}, \text{Ver})$ be a super-secure digital signature scheme. Then there exists an (efficient) $(\alpha = 1 - \gamma, \xi)$ -example reidentification scheme $(\text{Gen}_{\text{ex}}, \text{Trace}_{\text{ex}})$ for representation learning the concept class ValidSig .*

We now give the proof of Theorem 5.4.

Proof. We construct an example reidentification scheme for **ValidSig** as follows. The algorithm Gen_{ex} samples $(\text{sk}, \text{vk}) \leftarrow_{\text{R}} \text{Gen}(1^\lambda)$, a message $m \in \{0, 1\}^\ell$, and a signature $\sigma \leftarrow_{\text{R}} \text{Sign}(\text{sk}, m)$, yielding a concept $f_{\text{vk}, m, \sigma}$. Let \mathcal{D} be the distribution of $(\text{vk}, m, \text{Sign}(\text{sk}, m))$ for random $m \leftarrow_{\text{R}} \{0, 1\}^\ell$. Gen_{ex} then samples x_0, x_1, \dots, x_n i.i.d. from \mathcal{D} . Given a representation $(\text{vk}^*, m^*, \sigma^*)$, the algorithm Trace_{ex} simply identifies an index i for which $x_i = (\text{vk}^*, m^*, \sigma^*)$, and outputs \perp if none is found.

We first verify completeness for this scheme. Let L be a learner for **ValidSig** using n examples. If the representation $(\text{vk}^*, m^*, \sigma^*)$ produced by L represents an $(1 - \gamma)$ -good hypothesis, then it must be the case that $\text{vk}^* = \text{vk}$ and $\text{Ver}(\text{vk}, m^*, \sigma^*) = 1$. Thus, if L violates the completeness condition, it can be used to construct the weak forgery adversary \mathcal{A} (Figure 4) that succeeds with noticeable probability ξ .

Algorithm 4 Weak forgery adversary \mathcal{A}

1. Query the signing oracle on random messages $m'_1, \dots, m'_n \leftarrow_{\text{R}} \{0, 1\}^\ell$, obtaining signatures $\sigma'_1, \dots, \sigma'_n$.
 2. Run L on the labeled examples $((\text{vk}, m'_1, \sigma'_1), 1), \dots, ((\text{vk}, m'_n, \sigma'_n), 1)$, obtaining a representation (m^*, σ^*) .
 3. Output the forgery (m^*, σ^*) .
-

Now we verify soundness for the scheme. Observe that for any i , the sample S_{-i} contains no information about message m_i . Therefore, the learner has a $2^{-\ell} = \text{negl}(k)$ probability at producing a representation containing message m_i , proving soundness. □

Acknowledgements. We gratefully acknowledge Kobbi Nissim and Salil Vadhan for helpful discussions about this work, and also thank Salil Vadhan for suggestions on its presentation.

References

- [BCO11] Alexandra Boldyreva, Nathan Chenette, and Adam O’Neill. Order-preserving encryption revisited: Improved security analysis and alternative solutions. In *CRYPTO*, 2011.
- [BDMN05] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the SuLQ framework. In Chen Li, editor, *PODS*, pages 128–138. ACM, 2005.
- [BKN10] Amos Beimel, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. In *TCC*, pages 437–454, 2010.
- [BLR08] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to non-interactive database privacy. In Cynthia Dwork, editor, *STOC*, pages 609–618. ACM, 2008.

- [BLR⁺15] Dan Boneh, Kevin Lewi, Mariana Raykova, Amit Sahai, Mark Zhandry, and Joe Zimmerman. Semantically secure order-revealing encryption: Multi-input functional encryption without obfuscation. In *Proc. of EUROCRYPT*, 2015.
- [Blu94] Avrim Blum. Separating distribution-free and mistake-bound learning models over the boolean domain. *SIAM J. Comput.*, 23(5):990–1000, 1994.
- [BNS13] Amos Beimel, Kobbi Nissim, and Uri Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. In Prasad Raghavendra, Sofya Raskhodnikova, Klaus Jansen, and José D. P. Rolim, editors, *APPROX-RANDOM*, volume 8096 of *Lecture Notes in Computer Science*, pages 363–378. Springer, 2013.
- [BNSV15] Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil P. Vadhan. Differentially private release and learning of threshold functions. *CoRR*, abs/1504.07553, 2015.
- [BPTG14] Raphael Bost, Raluca Ada Popa, Stephen Tu, and Shafi Goldwasser. Machine learning classification over encrypted data. *IACR Cryptology ePrint Archive*, 2014:331, 2014.
- [BS15] Zvika Brakerski and Gil Segev. Function-private functional encryption in the private-key setting. In *TCC*, 2015.
- [BSW06] Dan Boneh, Amit Sahai, and Brent Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In *Proceedings of the 24th Annual International Conference on The Theory and Applications of Cryptographic Techniques*, EUROCRYPT’06, pages 573–592, Berlin, Heidelberg, 2006. Springer-Verlag.
- [BUV14] Mark Bun, Jonathan Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 1–10, 2014.
- [BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 480–499, 2014.
- [CFN94] Benny Chor, Amos Fiat, and Moni Naor. Tracing traitors. In Yvo Desmedt, editor, *CRYPTO*, volume 839 of *Lecture Notes in Computer Science*, pages 257–270. Springer, 1994.
- [CH11] Kamalika Chaudhuri and Daniel Hsu. Sample complexity bounds for differentially private learning. In Sham M. Kakade and Ulrike von Luxburg, editors, *COLT*, volume 19 of *JMLR Proceedings*, pages 155–186. JMLR.org, 2011.
- [CHS14] Kamalika Chaudhuri, Daniel Hsu, and Shuang Song. The large margin mechanism for differentially private maximization. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 1287–1295, 2014.
- [CTUW14] Karthekeyan Chandrasekaran, Justin Thaler, Jonathan Ullman, and Andrew Wan. Faster private release of marginals on small databases. *ITCS 2014*, 2014.

- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaude-
nay, editor, *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages
486–503. Springer, 2006.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise
to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *TCC*,
volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.
- [DNR⁺09] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vad-
han. On the complexity of differentially private data release: efficient algorithms and
hardness results. In Michael Mitzenmacher, editor, *STOC*, pages 381–390. ACM, 2009.
- [DNT14] Cynthia Dwork, Aleksandar Nikolov, and Kunal Talwar. Using convex relaxations for
efficiently and privately releasing marginals. In *Proceedings of the Thirtieth Annual
Symposium on Computational Geometry*, SOCG’14, pages 261:261–261:270, New York,
NY, USA, 2014. ACM.
- [DRV10] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential
privacy. In *FOCS*, pages 51–60. IEEE Computer Society, 2010.
- [FK12] Vitaly Feldman and Varun Kanade. Computational bounds on statistical query learn-
ing. In *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27,
2012, Edinburgh, Scotland*, pages 16.1–16.22, 2012.
- [FX14] Vitaly Feldman and David Xiao. Sample complexity bounds on differentially private
learning via communication complexity. *CoRR*, abs/1402.6278, 2014.
- [GGG⁺14] Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-
Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional en-
cryption. In *EUROCRYPT*, pages 578–602, 2014.
- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal
lattices. In *Proc. of EUROCRYPT*, 2013.
- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent
Waters. Candidate indistinguishability obfuscation and functional encryption for all
circuits. In *Proc. of FOCS*, 2013.
- [GGHZ14] Sanjam Garg, Craig Gentry, Shai Halevi, and Mark Zhandry. Fully secure functional
encryption without obfuscation, 2014.
- [GLN13] Thore Graepel, Kristin Lauter, and Michael Naehrig. ML confidential: Machine learn-
ing on encrypted data. In Taekyoung Kwon, Mun-Kyu Lee, and Daesung Kwon, edi-
tors, *Information Security and Cryptology ICISC 2012*, volume 7839 of *Lecture Notes
in Computer Science*, pages 1–21. Springer Berlin Heidelberg, 2013.
- [Gol04] Oded Goldreich. *Foundations of cryptography: volume 2, basic applications*. Cam-
bridge university press, 2004.

- [GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 59(3):11:1–11:35, June 2012.
- [GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In *TCC*, pages 339–356, 2012.
- [HLM12] Moritz Hardt, Katrina Ligett, and Frank McSherry. A simple and practical algorithm for differentially private data release. In Peter L. Bartlett, Fernando C. N. Pereira, Christopher J. C. Burges, Léon Bottou, and Kilian Q. Weinberger, editors, *NIPS*, pages 2348–2356, 2012.
- [HR10] Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *FOCS*, pages 61–70. IEEE Computer Society, 2010.
- [HRS12] Moritz Hardt, Guy N. Rothblum, and Rocco A. Servedio. Private data release via learning thresholds. In Dana Randall, editor, *SODA*, pages 168–187. SIAM, 2012.
- [Kea98] Michael Kearns. Efficient noise-tolerant learning from statistical queries. *J. ACM*, 45(6):983–1006, November 1998.
- [Kha95] Michael Kharitonov. Cryptographic lower bounds for learnability of boolean functions on the uniform distribution. *J. Comput. Syst. Sci.*, 50(3):600–610, 1995.
- [KK05] Jonathan Katz and Chiu-Yuen Koo. On constructing universal one-way hash functions from arbitrary one-way functions. *IACR Cryptology ePrint Archive*, 2005:328, 2005.
- [KLN⁺11] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM J. Comput.*, 40(3):793–826, 2011.
- [KV94] Michael Kearns and Leslie Valiant. Cryptographic limitations on learning boolean formulae and finite automata. *J. ACM*, 41(1):67–95, January 1994.
- [Nis14] Kobbi Nissim. Personal communication, July 2014.
- [NY89] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing*, STOC '89, pages 33–43, New York, NY, USA, 1989. ACM.
- [PV88] Leonard Pitt and Leslie G. Valiant. Computational limitations on learning from examples. *J. ACM*, 35(4):965–984, October 1988.
- [Rom90] J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing*, STOC '90, pages 387–394, New York, NY, USA, 1990. ACM.
- [RR10] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In *STOC*, pages 765–774, 2010.
- [Ser00] Rocco A. Servedio. Computational sample complexity and attribute-efficient learning. *J. Comput. Syst. Sci.*, 60(1):161–178, 2000.

- [SG04] Rocco A. Servedio and Steven J. Gortler. Equivalences and separations between quantum and classical learnability. *SIAM J. Comput.*, 33(5):1067–1092, 2004.
- [TUV12] Justin Thaler, Jonathan Ullman, and Salil P. Vadhan. Faster algorithms for privately releasing marginals. In *ICALP (1)*, pages 810–821, 2012.
- [Ull13] Jonathan Ullman. Answering $n^{2+o(1)}$ counting queries with differential privacy is hard. In *STOC*, pages 361–370, 2013.
- [UV11] Jonathan Ullman and Salil P. Vadhan. PCPs and the hardness of generating private synthetic data. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 400–416. Springer, 2011.
- [Val84] L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, November 1984.