Near-Optimal Time-Sparsity Trade-Offs for Solving Noisy Linear Equations

Kiril Bangachev[*]　　Guy Bresler[†]　　Stefan Tiegel[‡]　　Vinod Vaikuntanathan[§]

MIT　　　　　　　　MIT　　　　　　　ETH Zürich　　　　　　　MIT

Abstract

We present a polynomial-time reduction from solving noisy linear equations over $\mathbb{Z}/q\mathbb{Z}$ in dimension $\Theta(k \log n / \mathsf{poly}(\log k, \log q, \log \log n))$ with a uniformly random coefficient matrix to noisy linear equations over $\mathbb{Z}/q\mathbb{Z}$ in dimension n where each row of the coefficient matrix has uniformly random support of size k. This allows us to deduce the hardness of sparse problems from their dense counterparts. In particular, we derive hardness results in the following canonical settings:

1. Assuming the ℓ-dimensional (dense) learning with errors (LWE) problem over a polynomial-size field takes time $2^{\Omega(\ell)}$, k-sparse LWE in dimension n takes time

$$n^{\Omega(k/(\log k \cdot (\log k + \log \log n)))} .$$

2. Assuming the ℓ-dimensional (dense) learning parity with noise (LPN) problem over $\mathbb{F}_2$ takes time $2^{\Omega(\ell/\log \ell)}$, k-sparse LPN in dimension n takes time

$$n^{\Omega(k/(\log k \cdot (\log k + \log \log n)^2))} .$$

These running time lower bounds are nearly tight as both sparse problems can be solved in time $n^{O(k)}$, given sufficiently many samples.

Our reduction allows us to derive several consequences in cryptography and the computational complexity of statistical problems. In addition, as a new application, we give a reduction from k-sparse LWE to noisy tensor completion. Concretely, composing the two reductions implies that order-k rank-2^{k-1} noisy tensor completion in $\mathbb{R}^{n^{\otimes k}}$ takes time $n^{\Omega(k/\log k \cdot (\log k + \log \log n))}$, assuming the exponential hardness of standard worst-case lattice problems.

Our reduction is (nearly) lossless and extremely versatile: it preserves the number of samples up to a $1 - o(1)$ multiplicative factor with high probability and it preserves the noise distribution. In particular, it also applies to the learning with rounding problem. The same reduction works for a wide range of ring sizes from $q = 2$ to q that is super-polynomial in the initial dimension ℓ and for varying support sizes between samples. Finally, our reduction is compatible with decision (testing), search, and strong refutation and, hence, reduces hardness in the sparse setting to that of the standard, dense, setting for all three tasks.

Contents

1 Introduction

Random noisy linear equations over $\mathbb{Z}/q\mathbb{Z}$ play a fundamental role in average-case complexity, cryptography, and learning theory. Two canonical instances are the *learning parity with noise* (LPN) and the *learning with errors* (LWE) problems, that we will describe shortly.

Noisy linear equation problems are defined by a distribution over pairs $(A, y) \in (\mathbb{Z}/q\mathbb{Z})^{m \times n} \times (\mathbb{Z}/q\mathbb{Z})^m$, where A is the *public coefficient matrix* and y the *labels*. The rows of the public matrix A are often taken to be independent and identically distributed, in which case m is viewed as the number of independent examples in supervised learning. The basic goal is to determine whether there is an approximate linear relation between A and y, i.e., does there exist an $s \in (\mathbb{Z}/q\mathbb{Z})^n$ such that $y \approx As$ mod q. Three different tasks have received attention:

(Decision) Distinguish whether A and y are independent or $y \approx As$ for some s.

(Search) Under the promise that $y \approx As$ for some s, recover s.

(Refutation) Under the promise that A and y are independent, certify (i.e., prove) that y is not close to As for *any* s.

For the search and decision tasks, $y \approx As$ means that $y = As + e$ for some "small" error vector e. For the refutation task, one is most often interested in certifying an upper bound on the maximum number of equations that are approximately satisfiable simultaneously. Importantly, the upper bound should be correct *for all* inputs and "significantly better than trivial" on the average. (We refer to [Section 2.7](#) for rigorous definitions of the three tasks.) We are interested in how many equations (i.e., how large m needs to be) and how much computation time is necessary to solve each of these tasks.

The setup described above is nearly identical to that of standard linear regression, studied extensively in statistics, for which a variety of efficient algorithms can achieve nearly optimal sample complexity under certain assumptions on A. The primary difference lies in the modular arithmetic, which in the presence of even small errors makes the above tasks algorithmically challenging.

Depending on the specific choice of q and distributions of A, s, and e, researchers have put forth conjectures on how much time and equations are necessary to solve the problem. Such *hardness conjectures* form the basis of numerous cryptographic primitives and also have implications towards understanding the computational complexities of other problems. One of the goals of average-case complexity theory is to simplify the landscape and relate the computational complexities of different problems. Can one show that the various hardness conjectures follow from just one of them?

In this paper, we relate two canonical settings of noisy linear equations. We reduce the standard (dense) setting, in which each row of the public matrix A is independent and uniform over $(\mathbb{Z}/q\mathbb{Z})^L$, to the k-sparse setting, in which each row of the public matrix A is n-dimensional and has an independent uniformly random support of size k. In both cases, the secret s is uniform over $(\mathbb{Z}/q\mathbb{Z})^L$, respectively $(\mathbb{Z}/q\mathbb{Z})^n$. For concreteness, we state our theorem below in terms of LPN and LWE, which will be defined shortly.

Theorem 1.1 (Informal, see [Theorem 4.1](#) and [Corollaries 6.2](#) to [6.5](#)). *There is a* $\mathsf{poly}(n)$-*time reduction from standard LPN (respectively standard LWE) in dimension*

$$L = \frac{k \log n}{\mathsf{poly}(\log q, \log k, \log \log n)}$$

to k-LPN (respectively k-LWE) in dimension n. The reduction preserves the number of samples up to a $1 − o(1)$ multiplicative factor. Furthermore, for decision and search, the reduction preserves the noise distribution[1] and for refutation it preserves the strength of the refutation algorithm up to an additive $o(1)$ term.

All four settings of {dense, k-sparse}, {LPN, LWE} are common in the literature and have found important applications to cryptography, statistics, and learning theory. Yet, the relationship between the standard dense public matrix setting versus the sparse public matrix setting has remained elusive so far (although, see Sections 3.3.2 and 3.4.2 for a discussion of prior work that made progress in this direction).

This setting should not be confused with the setting in which the rows of the coefficient matrix are dense but the secret is sparse. The relation of the sparse secrete versions to the standard (dense) settings is much better understood; we refer to Section 3.6 for a more in-depth discussion.

1.1 LPN and LWE

In LPN, the ring is binary, i.e. $q = 2$, and the entries of the error vector are independently distributed as $\text{Bern}(\delta)$, where δ is a noise parameter (small δ yields a sparse error vector, with most entries equal to zero). Some works also consider LPN over large rings but again with sparse noise [IPS09; JLS21; JLS22; CHK+24a]. Unless otherwise stated, we will understand that $q = 2$ when we refer to LPN. In LWE, the modulus q is "large" (typically at least polynomial in L) and the noise is dense (most commonly drawn from a discrete Gaussian distribution of standard deviation at least $2\sqrt{L}$, in which case a reduction from worst-case lattice problems is known [Reg09]).

1.1.1 The Standard Setting: Uniformly Random (Dense) Public Matrix

In the standard variants of both LPN and LWE, the public matrix A and secret s are drawn independently and uniformly over $(\mathbb{Z}/q\mathbb{Z})^L$ and, hence, both are dense. Under suitable choices of the noise parameters, the decision versions of both problems are widely believed to be nearly exponentially hard.[2] In particular,[3]

the fastest known algorithm for LPN runs in time $2^{\Omega(L/\log L)}$ [BKW03],

and any improvement to this would be considered a breakthrough.

For LWE, several algorithms improve on the most naïve brute-force enumeration over the q^L possible secrets, but for large enough q and noise rate,

the fastest known algorithm for LWE takes time $2^{\Omega(L)}$.

There is compelling evidence that this cannot be improved upon, as such an algorithm would yield a $2^{o(L)}$-time (quantum) algorithm for several *worst-case* lattice problems (e.g. the approximate shortest vector problem with polynomial approximation factors) believed to require exponential time; see [Reg09; ADR+15; ABB+22] and the references therein.

The assumed near-exponential hardness in conjunction with the simple linear structure make both standard LPN and LWE (and variants of LPN and LWE known to inherit their hardness from the standard variants) extremely useful for a variety of applications, which we discuss in Sections 3.3 and 3.4.

---

[1]In particular, it also to the learning with rounding problem.

[2]Note that search is as least as hard as decision, and so is the natural refutation task, as discussed in Section 2.7.

[3]Oftentimes in statistical problems, there is a trade-off between the number of samples required and the running time of the algorithm. In this work, we mostly focus on a running-time lower bound. In that case, the lower bound holds regardless of the number of samples used.

1.1.2 The Setting of a Sparse Public Matrix

In many applications, additional structure of the noisy linear equations is very useful. Starting with the work of [Ale03], researchers have considered noisy linear equations in which each row a_i (in dimension n) of the public matrix A is k-sparse, i.e., has k non-zero entries. Typically, the locations of the supports are uniform and independent between samples. We denote the respective LPN and LWE variants by k-LPN and k-LWE[4]. One reason the sparse public matrix variant is useful is that the projections $a_i^\top s$ of s on a_i capture local interactions of the coordinates of the secret s — each $a_i^\top s$ depends only on k of the coordinates of s. Hence, one can use k-LPN and k-LWE to argue about problems related to learning from "local" structure, which are ubiquitous in computer science.

For both k-LPN and k-LWE, there is a simple brute-force algorithm that succeeds (for decision/search) in time $n^{O(k)}$ given $n^{1.01k}$ equations. With access to that many random k-sparse equations, each of the possible $\binom{n}{k}$ supports will appear $n^{\Omega(k)}$ times. Restricting to a given support of size k, one obtains a k-dimensional problem which can be solved in time $2^{O(k)}$ via the respective algorithms for dense k-dimensional LPN/LWE. Many other $n^{O(k)}$-time algorithms are known provided a sufficient number of samples, for example via gradient-descent dynamics in neural networks [BEG+24]. Perhaps surprisingly, for small values of k, no better running time than the naive $n^{\Theta(k)}$ is known. Further, with access to fewer than $n^{k/2}$ samples, there exists some evidence for $\exp(\mathsf{poly}(n))$-hardness of refuting k-sparse LPN [KMO+17].

The study of the sparse variant k-LWE was initiated much more recently [JLS24], but again no improvement to the naive $n^{\Omega(k)}$ sample and time requirement is known (and in [JLS24], the authors show that natural strategies fail). Altogether,

$$\text{the fastest known algorithms for } k\text{-LPN and } k\text{-LWE take } n^{\Omega(k)} \text{ time.}$$

Applications of k-LPN and k-LWE. Assuming that the $n^{\Omega(k)}$ runtime for k-LPN is optimal has become a standard and very useful assumption in learning theory and cryptography. In [Fei02; Ale03], a natural hardness assumption for 3-LPN was shown to imply new hardness of approximation results for min bisection, densest subgraph, and max bipartite clique. Hardness of k-LPN has also been instrumental in deriving hardness results for statistical tasks: agnostically learning halfspaces [Dan16], learning DNFs [DV21; BCM24], and noisy tensor completion [BM16] (see also our result described in Section 1.3).

Beyond applications to other problems, the refutation task on k-LPN (also known as refuting random k-XOR formulas) is a canonical problem in average-case complexity that is of interest in its own right [Fei02; Ale03; FKO06; BM16; AOW15; RRS17; KMO+17; GKM22; GHK+23].

In cryptography, sparse LPN has found applications in private- and public-key encryption [IKO+08; ABW10; DMN12], signature schemes [IKO+08], construction of pseudorandom generators with linear stretch in $\mathbf{NC^0}$ [AIK06], oblivious transfer [CRR21], multi-party homomorphic secret sharing [DIJ+23], indistinguishability obfuscation [RVV24] and, concurrent to our work, somewhat homomorphic encryption [CHK+24b].

As noted above, the study of k-LWE was only recently initiated in [JLS24]. The authors use k-LWE to construct state-of-the-art linear (and constant degree) homomorphic encryption schemes. In the current work, we give a further application of k-LWE — we derive a lower bound for tensor completion based on the worst-case hardness of lattice problems.

---

[4]We note that "k-LWE" does not refer to a unique model as different distributions for the values on the support are possible when $q > 2$. The results in the current work apply to *any* distribution on the support. In [JLS24], the authors only consider the uniform distribution on the support.

1.1.3 An Unsatisfactory State of Affairs: Motivation For This Paper

In light of the continuous use of the hardness of sparse noisy linear equations over the past more than 20 years and recent developments ([JLS24] for LWE), we anticipate further applications of the assumptions to fundamental problems in cryptography, learning theory, and statistics. Of central importance to this endeavour are the following questions:

What are reasonable hardness assumptions for problems on k-sparse linear equations?
How can we be confident in these assumptions?

These questions are the main motivation for our work. We address them via a reduction: we exhibit an algorithm that transforms standard noisy linear equations into sparse noisy linear equations. Our reduction shows that significantly improving over the brute-force running times of the sparse version would improve the state-of-the-art algorithms for the dense version, which would be a major breakthrough. We remark that our aim is to obtain *tight* hardness for the sparse problems under standard assumptions on the dense problem, essentially converting between brute-force running times. That is, obtaining $n^{\Omega(k)}$-type lower bounds for the sparse problem assuming $2^{\Omega(n/\mathsf{poly}(\log n))}$-hardness for the dense problem.

We note that the prior work of [JLS24] also gives a reduction from k-dimensional standard (dense) LWE to k-sparse n-dimensional LWE, showing that if the standard version takes exponential time, then the k-sparse version takes time at least $2^{\Omega(k)}$. This lower bound, however, is suboptimal, as it does not depend on the dimension n and does not show hardness when $k = O(\log n)$, a sparsity that is desirable for the applications in [JLS24] and crucial for the applications of k-LPN to hardness of learning. In particular, the level of sparsity often naturally corresponds to some measure of complexity in the learning problem (such as the number of terms in a DNF formula), and the goal is to rule out polynomial-time algorithms for as simple settings as possible (e.g., DNFs with as few terms as possible). Further, it is far from the running time of the naive $n^{\Omega(k)}$ algorithm. We discuss this further in Section 2.2.

1.2 Main Result: Near-Optimal Time-Sparsity Tradeoffs

Our main result, Theorem 1.1, is a nearly tight reduction from the standard (dense) setting to the sparse setting for solving noisy linear equations. Our hardness results in the sparse problem nearly match known algorithmic guarantees under standard assumptions on the dense problem.

Our reduction is flexible and addresses a variety of scenarios. It applies to all three of the decision, search, and refutation problems (see Section 2.7 for definitions). Second, our reduction applies simultaneously to a wide range of ring sizes, capturing both the LPN setting of $q = 2$ and polynomial q, and standard settings of LWE such as q being a large polynomial (or even larger) in the dimension of the dense problem. Third, the reduction is agnostic to the error distribution in the case of decision and search problems. Finally, our reduction can accommodate other variations, such as an arbitrary distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ on the support of the rows of A (see Section 5.4) and varying support sizes across samples (see Section 5.5).

Implied Hardness Results and Comparison to Known Algorithms. Our result implies that, assuming (near-)exponential hardness of the dense problems, we achieve hardness of the sparse problems that (nearly) matches algorithmic upper bounds.

Corollary 1.1 (k-LPN hardness). *If there is no $2^{o(L/\log L)}$-time algorithm for the decision/search/ refutation problem for standard LPN in dimension L, then there is no $n^{o(\frac{k}{\log k(\log k + \log\log n)^2})}$-time algorithm for the decision/search/refutation problem for k-LPN in dimension n.*

This result is nearly tight, since as noted earlier, given $O(n^{1.01k})$ samples there are $n^{O(k)}$-time algorithms.

Our lower bounds do not restrict the number of samples accessible to the algorithm. It is interesting to compare this with the more fine-grained picture for refutation when given fewer than $O(n^{k/2})$ samples. Given only $O(n^{k(1-\delta)/2+\delta})$ samples, for some constant δ, the fastest known algorithm runs in time $2^{\tilde{O}(n^\delta)}$ [RRS17]. Furthermore, in the case of binary LPN there is evidence, based on lower bounds against the powerful sum of squares hierarchy, that this might be necessary [Gri01; Sch08; KMO+17]. In Appendix A, we provide similar evidence for any finite ring $\mathbb{Z}/q\mathbb{Z}$ and for arbitrary distribution on the support for the rows of the public matrix A within the framework of *low-degree polynomial tests* [BHK+19; HS17; Hop18]. We leave open the problem of showing an analogous result using a reduction-based approach as in this paper.

In a similar vein to our k-LPN result above, we obtain the following hardness in the case of LWE:

Corollary 1.2 (k-LWE hardness). *If there is no $2^{o(L)}$-time algorithm for the decision/search/refutation problem for standard LWE in dimension L with $q = \mathsf{poly}(L)$, there is no $n^{o(\frac{k}{\log k \cdot (\log k + \log \log n)})}$ time algorithm for the decision/search/refutation problem for k-LWE in dimension n.*

For all three of decision/search/refutation for both LWE and LPN, our results imply that under standard near-exponential hardness assumptions on the dense problem, there is no polynomial-time algorithm for the sparse problem when $k = \omega((\log \log n)^2 (\log \log \log n))$ for LPN and $k = \omega((\log \log n)(\log \log \log n))$ for LWE (and $k = o(\sqrt{n})$ always). For both LWE and LPN, we also derive similar hardness results in the sparse setting assuming only sub-exponential or quasi-polynomial hardness of the standard problems. Sections 3.3.1 and 3.4.1 describe these and other results.

Downstream Applications to Cryptography and Learning Theory. As a direct consequence of Corollary 1.1, many (but not all) previous results based on k-LPN can now be based on standard LPN. We refer to Sections 4.5 and 4.6 for an in-depth discussion.

1.3   Application to Tensor Completion

Besides the downstream applcations obtained by combining Theorem 1.1 with previous reductions, we also show a new reduction from k-sparse LWE to noisy tensor completion (see, e.g., [BM16]) that we discuss next. Together with Theorem 1.1 this gives the first hardness for this problem based on a well-established worst-case assumption (when combined with the standard reduction from the gap-shortest vector problem to the dense version of LWE [Reg09; Pei09]).

We define the problem rigorously in Section 3.5 and only give an informal overview here. In the tensor completion problem, there is an unknown order k tensor that is approximately rank r,

$$T = \sum_{i=1}^{r} \sigma_i \cdot (u_i^1 \otimes u_i^2 \otimes \cdots \otimes u_i^k) + \Delta \tag{1}$$

in $\mathbb{R}^{n \times \cdots \times n}$ and we observe m uniformly random entries. Here, Δ is some appropriately constrained noise tensor. For normalization, we assume that the components $u_i^j \in \mathbb{R}^n$ satisfy $\|u_i^j\|_2 = \sqrt{n}$. The σ_i are scalars (note that they can make the norms arbitrary large or small). The goal is to construct an estimate $\hat{T}$ such that $\|T - \hat{T}\|_1 \le o(\|T\|_1)$ using as few observed entries as possible.

There exist algorithms [BM16] (and [MS18] for the noiseless version) for this problem running in time $n^{O(k)}$ with access to $m = \tilde{O}(rn^{k/2})$ observed entries under two assumptions which are *necessary even for*

inefficient algorithms. First, *the noise Δ has small norm relative to T*, for instance it holds that $\|\Delta\|_1 = o(\|T\|_1)$. Second, *the components are incoherent*, i.e. *no component u_i^j is concentrated on a small set of coordinates.* We discuss these assumptions more in Section 3.5.

Of course, if one aims to explicitly store $\widehat{T} \in \mathbb{R}^{n^{\otimes k}}$, there is no hope for an algorithm running in time $o(n^k)$. Yet, one may still hope that one can learn a data-structure $\widehat{T}$ which on input $(i_1, i_2, \ldots, i_k)$ returns $\widehat{T}_{i_1, i_2, \ldots, i_k} \in \mathbb{R}$ such that $\mathbb{E}_{(i_1, \ldots, i_k)} |T_{i_1, \ldots, i_k} - \widehat{T}_{i_1, \ldots, i_k}| = o\big(\mathbb{E}_{(i_1, \ldots, i_k)} |T_{i_1, \ldots, i_k}|\big)$. For example, if the algorithm finds scalars $(\widehat{\sigma}_i)_{i=1}^r$ and components $(\widehat{u}_i^{\,j})_{1 \leq i \leq k, 1 \leq j \leq r}$ such that

$$\sum_{i=1}^r \sigma_i \cdot (u_i^1 \otimes u_i^2 \otimes \cdots \otimes u_i^k) \approx \sum_{i=1}^r \widehat{\sigma}_i \cdot (\widehat{u}_i^{\,1} \otimes \widehat{u}_i^{\,2} \otimes \cdots \otimes \widehat{u}_i^{\,k}),$$

then the data structure can compute $\widehat{T}_{i_1, i_2, \ldots, i_k}$ as $\sum_{\ell=1}^r \widehat{\sigma}_\ell (\widehat{u}_\ell^{\,1})_{i_1} (\widehat{u}_\ell^{\,2})_{i_2} \cdots (\widehat{u}_\ell^{\,k})_{i_k}$ in time $\tilde{O}(kr)$, which is much faster than n^k when the rank r is small.

Such an algorithm running in time $n^{o(k)}$, however, is not known. We give the first, to the best of our knowledge, hardness evidence via a reduction from the decisional binary k-sparse LWE problem (in which all values on the support are equal to 1). Combining this with Theorem 1.1 and the worst-to-average-case reduction from the gap shortest vector problem to standard (dense) LWE [Reg09] gives hardness evidence based on worst-case lattice problems.

Theorem 1.2 (Lattice-Based Lower-Bound For Tensor Completion, Informal). *Assume that the gap-shortest vector problem in dimension L takes time at least $2^{\Omega(L)}$ in the worst-case (for quantum algorithms). Then, order-k noisy tensor completion with rank $r \geq 2^{k-1}$ cannot be solved in time $n^{o\left(\frac{k \log n}{\log k \cdot (\log k + \log \log n)}\right)}$ even under the small noise and incoherence assumptions.*

We remark that the small noise and incoherence assumption hold with exponentially high probability. Our lower bound will follow by showing that we cannot solve the following tensor distinguishing problem: Either we observe i.i.d. sampled entries of a noisy low-rank tensor, or we observe i.i.d. sampled entries of a tensor with independent, mean $o(1)$, and constant variance entries.

2 Overview of Our Reduction

We next give an overview of the proof of our main theorem, namely Theorem 1.1.

2.1 The High-Level Approach: Linear Mapping from Dense to Sparse Equations

In what follows, we denote the input to the dense problem by (B, y) and the input to the sparse problem by (A, y'). Suppose we are given (B, y) such that B is uniform over $(\mathbb{Z}/q\mathbb{Z})^{m \times L}$ and we want to solve one of the decision/search/refutation problems given access to an oracle for the same problem in the sparse setting. The high-level idea is to transform each dense equation to a sparse one in a way that preserves the approximate linear relation $b_i^\top s \approx y$ (or absence thereof) between "coefficients" b_i in row i of B and the label y_i.

A common approach in the literature towards this end (e.g. [BLM+13] for binary LWE, and [JLS24] for sparse LWE in certain regimes, discussed in Section 2.2) is to map the coefficient vector of each equation. If we can find a linear mapping, this will preserve the linear relation between the coefficient vector and the label. More specifically, we aim to find a matrix $G \in (\mathbb{Z}/q\mathbb{Z})^{n \times L}$ (we sometime refer to this as a *gadget*

or *decoding matrix*) such that given a sample $(b_i, y_i) \in (\mathbb{Z}/q\mathbb{Z})^L \times \mathbb{Z}/q\mathbb{Z}$, we can efficiently find a k-sparse $a_i \in (\mathbb{Z}/q\mathbb{Z})^n$ such that $a_i^\top G = b_i^\top$. Note that given this, it holds that

$$y_i \approx b_i^\top s \text{ for some } s \quad \Leftrightarrow \quad y_i \approx a_i^\top(Gs) \text{ for some } s. \tag{2}$$

Further, it should hold that the sparse equation has the correct distribution, i.e., that a_i is uniform over k-sparse vectors in $(\mathbb{Z}/q\mathbb{Z})^n$. That is, we require the following property:

$$\text{If } b_i \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), \text{ then } a_i \text{ is a uniformly random } k\text{-sparse vector in } (\mathbb{Z}/q\mathbb{Z})^n. \tag{3}$$

We refer to the process of computing a given b as the *decoding algorithm*. (Indeed, a watchful reader would already have seen the analogy with syndrome decoding in coding theory). We claim that such a primitive can imply a reduction between all of the decision/search/refutation versions. We illustrate this here for the search task, and refer to Section 2.6 for the other two tasks.[5] Assume that $y_i = b_i^\top s + e_i$ for some error term e_i. Then, we can compute the mapped equation $y_i' = y_i + a_i^\top z$, where $z \in (\mathbb{Z}/q\mathbb{Z})^n$ is drawn uniformly at random. It holds that

$$y_i' = y_i + a_i^\top z = b_i^\top s + e_i + a_i^\top z = a_i^\top(Gs + z) + e_i,$$

where we use the same z for each sample. The only role of z is that the final secret $Gs + z$ is uniform over $(\mathbb{Z}/q\mathbb{Z})^n$ and independent of everything else. Thus, given a search oracle for the sparse version we can recover $Gs + z$ and hence also Gs. If G is full-rank we can then recover s.

2.2 Prior Gadget Matrix Constructions and Challenges

The only prior work that we are aware of which reduces standard (dense) noisy linear equations to sparse ones is [JLS24] in the case of LWE (and large-modulus LPN) with q prime. The authors follow the gadget matrix approach outlined above with $L = k$ and $G \in (\mathbb{Z}/q\mathbb{Z})^{n \times k}$ the Vandermonde matrix, i.e., whose i-th row is equal to $(i^0, i^1, i^2, \ldots, i^{k-1})$.

To solve $a^\top G = b^\top$ on input $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^k)$ such that a is a uniformly random k-sparse vector, they first draw a uniformly random support $S \in \binom{[n]}{k}$ *independently* of b. Then, they restrict a and G to coordinates S and solve the equation $a_S^\top G_{S \times [k]} = b^\top$. Whenever $q \geq n$, *every $k \times k$ minor of G is of full rank* and, hence, a unique solution for a_S exists. Since b is uniform over $(\mathbb{Z}/q\mathbb{Z})^k$, it is easy to check that a_S is also uniform over $(\mathbb{Z}/q\mathbb{Z})^k$, so a has the correct distribution. Overall, the reduction of [JLS24] shows that n-dimensional k-sparse LWE is at least as hard as k-dimensional LWE; assuming exponential hardness of LWE implies that the k-sparse version takes time at least $2^{\Omega(k)}$.

We outline two ways in which one might hope to improve over this reduction:

Desideratum 1: Small Values of k. As noted just above, the hardness implied by [JLS24] for n-dimensional k-LWE is at best only $2^{\Omega(k)}$, which does not yield any hardness against polynomial-time (in n) attacks when $k = O(\log n)$. It is instructive to analyze where their reduction is loose. As an intermediate step of transforming a dense $b \in (\mathbb{Z}/q\mathbb{Z})^L$ to a k-sparse $a \in (\mathbb{Z}/q\mathbb{Z})^n$, they compute a dense vector $\tilde{a} \in (\mathbb{Z}/q\mathbb{Z})^k$ such that $\tilde{a}^\top \tilde{G} = b$, for some deterministic matrix $\tilde{G}$. Note that the entropy of b is equal to $L \log q$ and this must be smaller than the entropy of $\tilde{a}$. Since the entropy of $\tilde{a}$ is at most $k \log q$, we must have $k \log q \geq L \log q$ and hence $k \geq L$. This approach does not exploit the entropy of the support location of $a \in (\mathbb{Z}/q\mathbb{Z})^n$. There are

---

[5]Crucially, for the refutation task we need that *almost all* of the equations are mapped successfully with high probability, while this is not necessary for decision and search. See Section 2.6 for details.

$\binom{n}{k}$ choices for the support of a, which yields $\log \binom{n}{k} \approx k \log n$ additional bits of entropy. Thus, if one simply requires $a^\top G = b^\top$ (instead of $\tilde{a}^\top \tilde{G} = b$), the corresponding entropy calculation is $k \log n + k \log q \geq L \log q$, allowing for larger values of L. We record this as a guiding principle:

> For small k, when solving $a^\top G = b^\top$, one must choose the support of a based on b. (P1)

Note that a similar entropy does give *some* limitations for general decoding matrices when $L = \omega(k)$. Since the entropy of a is at most $k \log n$ it must hold that $L \log q \leq k \log n$, i.e., $L \leq k \log n / \log q$.[6] We will show that this upper bound can nearly be achieved. In particular, we will be able to choose $L = k \log n / \mathsf{poly}(\log k, \log \log n, \log q)$.

Desideratum 2: Small Moduli. We would like our reduction to work also for small moduli, and in particular for $q = 2$ as in the case of LPN. Note that a key step in the reduction of [JLS24] is the inversion of an arbitrary $k \times k$ sub-matrix of G, which requires every $k \times k$ minor of G to be full rank. Such matrices G are known as *maximum distance separable* and in the regime $n \geq k$, conjecturally, cannot exist if $q < n - 2$ (e.g. [Huf03]). Hence, one cannot hope for a reduction in the important case of binary LPN using the approach of [JLS24].

We would also like to be more versatile, e.g. obtain sparse vectors a_i with distribution on the support different from uniform over $(\mathbb{Z}/q\mathbb{Z})^k$. As one example, our application to the hardness of tensor completion from LWE requires *binary* k-sparse coefficients a_i. This does not seem possible using the approach in [JLS24].

2.3   A Toy Reduction for Block-wise Sparsity

A gadget mapping dense L-dimensional vectors to *blockwise k-sparse* n-dimensional vectors is not hard to construct. A blockwise sparse vector is one whose coordinates are divided into k blocks, each of which has a single non-zero coordinate. Indeed, for a parameter $t = L/k$, let $H \in (\mathbb{Z}/q\mathbb{Z})^{t \times q^t}$ denote a matrix whose columns are all the vectors in $(\mathbb{Z}/q\mathbb{Z})^t$. Then, we will let the decoding gadget $G = I_k \otimes H \in (\mathbb{Z}/q\mathbb{Z})^{tk \times q^t k}$. First of all, since $L = tk$ and $n = q^t k$, notice that

$$L = \frac{k \log(n/k)}{\log q} \ .$$

Secondly, given an L-dimensional vector b, one can decode each block of t coordinates of b into a unit vector in q^t dimensions. Concatenating the unit vectors gives us a block-wise sparse n-dimensional vector a such that $a^\top G = b^\top$.

Permuting the coordinates might seem to suffice to turn blockwise sparsity into true sparsity. However, since the same permutation has to be used for all m samples, the blockwise structure starts becoming apparent from even a small number of samples. For example, given $\omega(n^2)$ samples, every two coordinates i, j appear in some row of the public matrix A when the supports are uniformly random k-tuples. This property is not shared by the blockwise model for any number of samples under any permutation.

To decode to random sparse vectors, we need a cleverer idea which we describe next. We begin with an informal non-mathematical overview.

---

[6]In this argument we implicitly assumed that G is deterministic. Since we use the same G for all equations, even for a random G the average entropy added on any single equation is at most $\frac{nL \log q}{m}$ which is negligible for all parameter settings we consider.

2.4 Motivation for Our Reduction: Translation of Natural Languages

The problem at hand is reminiscent of an age-old task: translating from one language (dense coefficient vectors) to another (sparse coefficient vectors). Translation of natural languages motivates our approach. First, we split each sentence (coefficient vector) into words and translate word-by-word using a dictionary. Second, we use the grammar of the two languages to combine the respective words into meaningful sentences (coefficient vectors). The intuition why such an approach is efficient is that the small number of possible words (as opposed to possible sentences or even paragraphs) allows for efficient translation. Imagine if one had to print a dictionary in the real world containing translations of all possible sentences!

For us, the grammar of the "dense language" is simple. Each sentence is simply a concatenation of words. Our dictionary allows us to translate each of these words into a sequence of k indices over some small domain. These new sequences of length k are the words in the "sparse language". One of our contributions is to specify a grammar that allows us to combine k indices over a small domain into a k-sparse equation. It should already be clear that the approach taken is more sophisticated than simple concatenation, which would blow-up the sparsity k.

2.5 Overview of Our Reduction to Sparse Equations

We begin by explaining the reduction in the case that the distribution on the support is deterministically equal to 1, which captures both sparse problems k-LPN and *binary k-LWE*. At the end of this section, we explain how a simple modification addresses arbitrary distributions over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$, the integers coprime with q, on the support.

Recall that as in [JLS24], we want to construct a gadget matrix $G \in (\mathbb{Z}/q\mathbb{Z})^{n \times L}$ and, on input a uniform random $b \in (\mathbb{Z}/q\mathbb{Z})^L$, solve $a^\top G = b^\top$ for a k-sparse $a \in \{0, 1\}^n$. However, as observed in (P1), in doing so, we need to use the randomness of b in choosing the location of the support of a. In particular, we need to sample from (some large enough subset of) the possible supports. This task can be interpreted as a sampling version of the popular k-SUM problem, in which given n vectors in $(\mathbb{Z}/q\mathbb{Z})^L$ and a target b one needs to find k of them that sum to b. Evidence for hardness of this problem exists both in the average-case setting (when the n vectors are random) based on lattice problems [BSV20] and in the setting of worst-case vectors based on the exponential-time hypothesis [PW10; ABH+22]. Essentially, these results state that the time required is $\min(n^{\Omega(k)}, 2^{\Omega(L/\log k)})$, both of which are too slow for our purposes. Instead, we need a reduction that runs in time $n^{o(k)}$ (to conclude $n^{o(k)}$ hardness of k-sparse LWE) and $2^{L/\log L} = 2^{L/\log \omega(k)}$ in the relevant regime $L = \omega(k)$ in the case of LPN (as our reduction needs to be faster than the strongest plausible $2^{L/\log L}$ hardness of LPN). This step of solving k-sum poses a serious obstacle to our approach. Unfortunately, our task is actually even harder – we furthermore need to *sample* from the k-SUM solutions.

We circumvent the hardness of k-sum by endowing G with a convenient structure that makes finding solutions and sampling from them easy.

Word Decoding: Splitting into Small Sub-Problems That Can Be Solved Efficiently. We will decompose the problem into t sub-problems by splitting the string b into t equal parts (similar to the celebrated BKW algorithm [BKW03]) of size $h := L/t$. Namely, represent each $b \in (\mathbb{Z}/q\mathbb{Z})^L$ as a concatenation of t uniform independent strings $b = (b_1, b_2, \ldots, b_t)$ in $(\mathbb{Z}/q\mathbb{Z})^h$. We will decode each b_j into a k-tuple of indices over $[hq]^k$ via a smaller gadget matrix $V \in (\mathbb{Z}/q\mathbb{Z})^{hq \times h}$ (this is the dictionary mentioned above). As we will need to solve k-SUM over this V, it is natural that we want V to have a very simple explicit structure. We

take V to contain all multiples of the identity:

$$V = \begin{pmatrix} 0 \times I_h \\ 1 \times I_h \\ \vdots \\ (q-1) \times I_h \end{pmatrix}.$$

We call the hq rows $V_{1,:}, V_{2,:}, \dots, V_{hq,:} \in \mathbb{Z}_q^h$ words. We want to find $\xi_j = (\xi_j(1), \dots, \xi_j(k)) \in [hq]^k$ such that

$$\sum_{\psi=1}^{k} V_{\xi_j(\psi),:} = b_j.$$

Furthermore, if b_j is uniform over $(\mathbb{Z}/q\mathbb{Z})^h$, we will require the marginal distribution of ξ_j to be uniform over $[hq]^k$. We achieve this by outputting a uniformly random ξ_j that satisfies $\sum_{\psi=1}^{k} w_{\xi_j(\psi)} = b_j$.

Here, we take the opportunity to introduce some useful notation:

$$\mathsf{Preimage}_V(b) := \left\{ (\xi(1), \xi(2), \dots, \xi(k)) \in [hq]^k \,\Big|\, \sum_{\psi=1}^{k} V_{\xi(\psi),:} = b \right\}. \tag{4}$$

The simple structure of V allows for extremely efficient sampling from the preimage of b.

Lemma 2.1 (Efficient preimage sampling). *On input $k > h + 2$ and $b \in (\mathbb{Z}/q\mathbb{Z})^h$, Algorithm 2 outputs a uniformly random sample from $\mathsf{Preimage}_V(b)$. The algorithm runs in time $\mathsf{poly}(k, h, \log q)$.*

On a high level, this can (roughly) be achieved by sampling indices uniformly at random (from $[hq]$) until the number of indices left matches the number of coordinates that differ between the current state, i.e., the sum of rows of V indexed by the current indices, and b. The choice of the remaining coordinates is then fixed. Permuting them randomly will result in a uniformly random element from the preimage. We give full detail in Section 5.1.1.

When combining the different words into sentences, it will be crucial that $\xi_j(1), \dots, \xi_j(k)$ are independent and uniform over $[hq]$. To guarantee this, a key property will be that each $\mathsf{Preimage}_V(b)$ is of nearly equal size. Otherwise, k-tuples $\xi_j(1), \dots, \xi_j(k)$ corresponding to b with large (respectively, small) $\mathsf{Preimage}_V(b)$ will be under-represented (respectively, over-represented). Again, the simple structure of V can be utilized to show the following fact.

Lemma 2.2 (Uniformity of preimage size). *Suppose that $k \geq 4h(\log h + \log q + \log(1/\psi))$ where $\psi \in (0, 1)$ is any real number. Then, for each $u \in (\mathbb{Z}/q\mathbb{Z})^h$,*

$$\frac{(1-\psi)(hq)^k}{q^h} \leq |\mathsf{Preimage}_V(u)| \leq \frac{(1+\psi)(hq)^k}{q^h}.$$

For reasons that will become clear, we will choose $\psi = o(1/t)$, say $\psi = (ht)^{-2}$.

We give the proof of Lemma 2.2 in Section 5.1.3, which follows by analyzing the point-wise convergence of the following random walk on $(\mathbb{Z}/q\mathbb{Z})^h$. At each step j, the random walk chooses independently uniformly random $i_j \in [h]$ and $m_j \in \{0, 1, 2, \dots, q-1\}$ and moves m_j units in the direction e_{i_j}. Recalling the definition of $\mathsf{Preimage}_V(b)$ in (4), we interpret this as $\xi(j) = hm_j + i_j$ so that $V_{\xi(j),:} = m_j(e_{i_j})^\top$.

In the case $q = 2$, this is just the standard lazy random walk over $\{0, 1\}^h$. It is well-known that this walk converges to an ϵ- total variation distance from the uniform distribution after roughly $h(\log h + \log(1/\epsilon))$

steps (e.g. [LPW06]). For our applications, a stronger point-wise convergence guarantee is crucial. Trying to deduce a point-wise convergence bound from the total variation bound would require $k \gtrsim h^2$, which would result in a much looser lower bound. Instead, we directly verify the needed point-wise convergence.

In order to make the output tuples exactly equal, we will use a simple rejection-based method that rejects each b_j with probability $1 - \frac{|\mathsf{Preimage}_V(b_j)|}{(1+\psi)(hq)^k/q^h}$, which is a well-defined probability and is at most 2ψ under the guarantee of Lemma 2.2. To make this algorithmically efficient, we need to be able to compute $|\mathsf{Preimage}_V(b_j)|$. Again, the simple structure of V allows such a primitive:

Lemma 2.3 (Computing preimage sizes)**.** *There exists an algorithm, Algorithm 3, which on input k and $b \in (\mathbb{Z}/q\mathbb{Z})^h$, outputs $|\mathsf{Preimage}_V(b)|$. The algorithm runs in time $\mathsf{poly}(k, h, \log q)$.*

The key observation for the algorithm is that due to symmetry, the size of the set $\mathsf{Preimage}_V(b)$ only depends on the number of non-zero entries of b. This size can be efficiently computed via dynamic programming. We give full detail in Section 5.1.2.

Remark 1. Lemma 2.2 explains why at least when using the gadget matrix V, we need to split the original L-dimensional problem into t smaller problems. For the desired mixing to occur (even in total variation), we need $k \geq h\log(h)$. Hence, if we directly use V to decode the entire dense vector, we will reduce the hardness of k-sparse LPN/LWE in dimension n from standard LPN/LWE in dimension $h = O(k/\log k)$. Thus, the obtained hardness will be at most $2^{O(k/\log k)}$, which is significantly worse than the desired $2^{k\log n/\mathsf{poly}(\log k, \log\log n)}$ for small values of k. We overcome this issue via our novel "grammar" of sparse sentences.

Sentence Decoding: Combining the Solutions of Small Sub-Problems. We will think of concatenations of words as sentences. Each sentence will have exactly t words. The rows of our final gadget matrix will be all possible sentences, of which there are $(hq)^t$.

Let $n = (hq)^t$. We will use the decodings $\{\xi_j\}_{j=1}^t$ of $\{b_j\}_{j=1}^t$ obtained via the matrix V to produce a k-tuple $\xi \in [n]^k = ([hq]^t)^k$ decoding of b. The key insight is that $\{b_j\}_{j=1}^t$ are t i.i.d. vectors in $(\mathbb{Z}/q\mathbb{Z})^h$ and, hence, $\{\xi_j\}_{j=1}^t$ are t i.i.d. tuples in $[hq]^k$. In particular, $(\xi_j(\psi))_{1\leq j\leq t, 1\leq \psi\leq k}$ are kt independent uniform integers in $[hq]$.

By definition, ξ should be uniformly random. We define $\xi(\psi) := (\xi_1(\psi), \xi_2(\psi), \ldots, \xi_t(\psi))$ for $1 \leq \psi \leq k$ and $\xi = (\xi(1), \xi(2), \ldots, \xi(k))$. Visually, we represent ξ as the following matrix in $[hq]^{t\times k}$.

$$\xi = \begin{pmatrix} \xi_1(1) & \xi_1(2) & \ldots & \xi_1(k) \\ \xi_2(1) & \xi_2(2) & \ldots & \xi_2(k) \\ & & \vdots & \\ \xi_t(1) & \xi_t(2) & \ldots & \xi_t(k) \end{pmatrix}.$$

In it, each row corresponds to a decoded word (that is, $\xi_j(1), \xi_j(2), \ldots, \xi_j(k)$ is the k-tuple corresponding to b_j). Each column $\xi(\psi) = (\xi_1(\psi), \xi_2(\psi), \ldots, \xi_t(\psi))^\top$ corresponds to a t-tuple in $[hq]^t$ and, hence, can be treated as an index in $[(hq)^t] = [n]$. This is the k-sparse grammar we referred to. Crucially, each of the k indices (columns) are independent uniform. Thus, we have managed to translate $b \in (\mathbb{Z}/q\mathbb{Z})^{ht}$ into k uniform independent indices in $[(hq)^t] = [n]$, as desired.

To finish, we put this in the language of gadget matrices and demonstrate a linear relationship between b and ξ. Let $G \in (\mathbb{Z}/q\mathbb{Z})^{n\times L}$ be such that its rows are indexed by $(\phi_1, \ldots, \phi_t) \in [hq]^t$ and each row is defined by

$$G_{(\phi_1, \phi_2, \ldots, \phi_t), :} = (V_{\phi_1, :}, V_{\phi_2, :}, \ldots, V_{\phi_t, :}) \in (\mathbb{Z}/q\mathbb{Z})^{ht} = (\mathbb{Z}/q\mathbb{Z})^L . \tag{5}$$

In other words, G contains all possible concatenations of t words from V.

Since $\sum_{\psi=1}^{k} V_{\xi_j(\psi),:} = b_j^\top$ holds for each j, it must be the case that

$$\sum_{\psi=1}^{k} G_{\xi(\psi)} = \sum_{\psi=1}^{k} \left(V_{\xi_1(\psi),:}, V_{\xi_2(\psi),:}, \ldots, V_{\xi_t(\psi),:} \right) = \left(\sum_{\psi=1}^{k} V_{\xi_1(\psi),:}, \sum_{\psi=1}^{k} V_{\xi_2(\psi),:}, \ldots, \sum_{\psi=1}^{k} V_{\xi_t(\psi),:} \right)$$
$$= \left(b_1^\top, b_2^\top, \ldots, b_t^\top \right) = b^\top . \tag{6}$$

Under a mild condition $k = o(\sqrt{n})$, the birthday paradox shows that with high probability, all of the $\xi(\psi)$ are distinct. When this event holds we can represent ξ as a (binary) k-sparse vector a that is uniformly random (the non-zero coordinates of a are indexed by $\xi(1), \xi(2), \ldots, \xi(k)$). (6) is equivalent to $a^\top G = b^\top$ as desired.

Example Choice of Parameters for LPN. We formally choose the parameters in Section 6 but here we give a brief intuition about how large they are for concreteness. For simplicity, we focus on the decision version of LPN, i.e., $q = 2$. We assume $2^{\Omega(L/\log L)}$ hardness of standard LPN and will derive a $n^{\Omega(\frac{k}{\log k(\log k + \log \log n)^2})}$ lower bound for the k-sparse decision version.[7]

Note that $t = \frac{\log n}{\log h + \log q} = \frac{\log n}{1 + \log h}$. Hence, $L = ht = \frac{h \log n}{1 + \log h}$. This value is clearly increasing in h. Thus, we want to choose the largest h possible in order to get the strongest runtime lower bound.

However, to apply Lemma 2.2 we require $k \geq 4h(\log h + \log q + \log(ht)^2)$ and $4h(\log h + \log q + \log(ht)^2) = \Theta(h(\log h + \log \log n))$. Hence, we can choose h at most $h = c\frac{k}{\log k + \log \log n}$ for some small absolute constant $c > 0$. Thus, $t = \frac{\log n}{\log h + \log q} = \Theta(\frac{\log n}{\log k})$. Altogether, $L = ht = \Theta(\frac{k \log n}{\log k(\log k + \log \log n)})$. We obtain

$$\exp\left(\Omega\left(\frac{L}{\log L} \right) \right) = \exp\left(\Omega\left(\frac{k \log n}{\log k(\log k + \log \log n)^2} \right) \right) \quad \text{hardness.}$$

Other Distributions on the Support. Given a vector $b \in (\mathbb{Z}/q\mathbb{Z})^L$, we first solve the binary problem. That is, find a k-tuple $(i_1, \ldots, i_k) \in [n]^k$ such that the corresponding rows of G sum to b. If b is uniform, then so is $(i_1, \ldots, i_k)$, as already mentioned. Next, sample $\rho = (\rho_1, \rho_2, \ldots, \rho_k) \in ((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from $\mathcal{D}$. Suppose that the row of G indexed by i_j corresponds to

$$(\kappa_1(j)e_{\Delta_1(j)}^\top, \kappa_2(j)e_{\Delta_2(j)}^\top, \ldots, \kappa_t(j)e_{\Delta_t(j)}^\top),$$

where $\kappa_i(j) \in \mathbb{Z}/q\mathbb{Z}$ and $\Delta_i(j) \in [h]$. We use the convention that the first h (zero) rows of V are $0e_1^\top, 0e_2^\top, \ldots, 0e_s^\top$. Now, replace this index with the index i_j' corresponding to

$$(\kappa_1(j)\rho_j^{-1}e_{\Delta_1(j)}^\top, \kappa_2(j)\rho_j^{-1}e_{\Delta_2(j)}^\top, \ldots, \kappa_t(j)\rho_j^{-1}e_{\Delta_t(j)}^\top).$$

Since ρ_j and q are coprime, i_j' is still uniform over $[n]$. Suppose again that all i_j' are distinct, we can show that this happens with high probability. Let a' be the vector with support $(i_1', i_2', \ldots, i_k')$ such that $a_{i_j'}' = \rho_j$. Then, $(a')^\top G = a^\top G = b^\top$ and clearly the support of a' is uniform and the distribution on the support equal to $\mathcal{D}$. We give full detail in Section 5.4.

---

[7]For technical reasons we require the condition that $k = o(\sqrt{n})$.

Variable Support Sizes. Finally, we can also accommodate for varying support sizes. The reason is that all of our primitives - Lemmas 2.1 to 2.3 only rely on k being sufficiently large and, thus, it is enough to accommodate for the effectively smallest possible support size k.

More concretely, suppose that the distribution of support sizes is $\mathcal{K}$ where with probability $1 - \eta$ with $\eta = o(1/m)$ the support size is at least $k_{\min}$. Then, for each sample b_i, we first draw $k_i \sim \mathcal{K}$ and then use the word-decoding and sentence-decoding procedures for k_i. We can show that the k-tuples are still uniform for all $k_i \geq k_{\min}$ if the input b is random. With high probability, we never see $k_i < k_{\min}$. More detail is given in Section 5.5.

2.6 Reduction for Decision, Search, and Refutation

As before, Eqs. (2) and (3) are sufficient to obtain a reduction for any of the decision, search, and refutation tasks. Thus, given a search oracle we can recover $Gs + z$ and hence also Gs. We already discussed this for the search task. Recall that given (b_i, y_i), our reduction outputs $(a_i, y_i' := y_i + a_i^\top z)$, where z is uniform over $(\mathbb{Z}/q\mathbb{Z})^n$. Note that this implies that the "planted" distribution of the decision task will be mapped correctly. Similarly, if in the original input y_i and b_i are independent then so will y_i' and a_i be since decoding b_i into a_i does not depend on y_i and $y_i' = a_i^\top z + y_i$. Note that for both decision and search, the reduction preserves the error of each equation.

Lastly, we discuss refutation. For simplicity of exposition, we focus on the LPN setting here, but the same arguments apply to LWE (see Section 5.6). Given (A, y), we want two guarantees. For all approximately-satisfiable systems, i.e. $As \approx y$ for some s, we want to output *approximately satisfiable* with probability 1. For the case when A, y are drawn independently (and hence the system is far from satisfiable with high probability), we want to output *strongly unsatisfiable*, with high probability.

In the discussion on search and testing, we considered the slightly idealized setting in which every equation is decoded successfully. Yet, many reductions (including ours) lose samples. Note that for decision and search even only preserving a small $1/100$ (or less) fraction of the samples would be sufficient. For refutation, this is not sufficient since our algorithm has to map approximately satisfiable inputs to approximately satisfiable outputs. Yet, consider the following extreme case. A very large $99/100$ fraction of the clauses in the input are satisfiable due to the fact that they are all the same (a_1, y_1). The rest $1/100$ of instances are highly unsatisfiable. A decoding process which always fails on a_1 and succeeds on all the other samples would lose the $99/100$ fraction of satisfiable equations and could produce a highly unsatisfiable instance.

Luckily, we are able to construct gadget matrices and a decoding algorithm that only lose a $o(1)$-fraction of the samples with high probability over worst-case inputs. Indeed, recall that each $b_j \in (\mathbb{Z}/q\mathbb{Z})^h$ is rejected with probability at most $O(\psi) = O((ht)^{-2})$ and, hence, b is rejected with probability at most $O(t(ht)^{-2}) = o(1)$. This resolves the above-described issue. Hence, our reduction maps approximately satisfiable instances to approximately satisfiable instances and instances with B, y independent to instances with A, y independent. Thus, it also works for refutation.

2.7 Organization of the Rest of the Paper

In Section 3 we give formal definitions and summarize relevant prior work. In Section 4 we formally state our results. The following section Section 5 contains the full reduction. Our reduction as well as its runtime are parametrized. In Section 6 we describe several specific parameter instantiations which lead to hardness results for (sparse) k-LPN and k-LWE based on popular assumptions for the hardness of standard LPN and LWE. We end with our application to noisy tensor completion in Section 7.

3 Background on Noisy Linear Equations and Tensor Completion

3.1 Noisy Linear Equations over $\mathbb{Z}/q\mathbb{Z}$

In this current paper, we focus on the following family of noisy linear equations over $\mathbb{Z}/q\mathbb{Z}$ unless otherwise stated. Recall that $\mathbb{Z}/q\mathbb{Z}$ is the ring of integers modulo q, which is furthermore a field when q is a prime. Let q be a given modulus and $\mathsf{D}^{\mathrm{coeff}}$ be a distribution over $(\mathbb{Z}/q\mathbb{Z})^n$ and $\mathsf{D}^{\mathrm{error}} = (\mathsf{D}^{\mathrm{error}}(y))_{y \in \mathbb{Z}/q\mathbb{Z}}$ be a family of distributions. We generate samples $\{(a_i, y_i)\}_{i=1}^m$ as follows:

1. Sample $a_1, a_2, \ldots, a_m \overset{\text{i.i.d.}}{\sim} \mathsf{D}^{\mathrm{coeff}}$.

2. Sample $s \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^n)$.

3. Sample $e_j \sim \mathsf{D}^{\mathrm{error}}(a_j^\top s)$ independently for each $j \in [m]$.

4. Output $\{(a_i, y_i := a_i^\top s + e_i)\}$.

We will denote an input consisting of m samples from this distribution by

$$\mathsf{NLE}(\text{samples} = m, \text{modulus} = q, \mathsf{D}^{\mathrm{coeff}}, \mathsf{D}^{\mathrm{error}}). \tag{7}$$

For example, $\mathsf{NLE}(\text{samples} = m, \text{modulus} = 2, \mathsf{Unif}(\mathbb{Z}_2^n), \mathsf{Bern}(\frac{1}{2} - \gamma))$ captures standard LPN and $\mathsf{NLE}(\text{samples} = m, \text{modulus} = q, \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^n), D_\sigma)$ captures standard LWE. Most frequently, the noise $\mathsf{D}^{\mathrm{error}}(y)$ is independent of $y \in \mathbb{Z}/q\mathbb{Z}$, e.g. Discrete Gaussian. We state in this more general case to include problems like learning with rounding, which are also captured by our reduction [BPR12].

The following notation will be useful for sparse linear equations. Suppose that $\mathcal{D}$ is a distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$, where $(\mathbb{Z}/q\mathbb{Z})^\times$ denotes the set of residues relative prime with q. Let $\mathsf{randomsupport}(n, k, \mathcal{D})$ be the following distribution over k-sparse vectors in $(\mathbb{Z}/q\mathbb{Z})^n$: First, draw a uniformly random set $\{i_1, \ldots, i_k\}$ of size k. Then, draw $\rho = (\rho_1, \rho_2, \ldots, \rho_k) \sim \mathcal{D}$ and set $a_{i_j} = \rho_j$ for all $1 \le j \le k$ and $a_i = 0$ otherwise.

We are interested in instances of $\mathsf{NLE}(m, q, \mathsf{randomsupport}(n, k, \mathcal{D}), \mathsf{D}^{\mathrm{error}})$. Concrete examples of interest include:

1. *k-LPN*: Let $q = 2$ and $\mathcal{D}$ output the vector 1^k deterministically.

2. *Binary k-LWE*: Let $\mathcal{D}$ output the vector 1^k deterministically. This is the basis for our reduction to tensor completion, see Section 7.

3. *k-LWE with (Almost) Uniform Values on The Support*: Let $\mathcal{D}$ be $\mathsf{Unif}\left(((\mathbb{Z}/q\mathbb{Z})^\times)^k\right)$.

The last example almost recovers the sparse LWE instance from [JLS24]. The only difference is that in [JLS24], the distribution $\mathcal{D}$ on the support is $\mathsf{Unif}\left((\mathbb{Z}/q\mathbb{Z})^k\right)$ instead. This seems to make little difference in the regime of a large prime q and k as in [JLS24]. However, our work focuses on the case when both k and q are very small with respect to n. In particular, this means that under $\mathcal{D} = \mathsf{Unif}\left((\mathbb{Z}/q\mathbb{Z})^k\right)$, there would be many samples a with support of size 1. This makes both testing and search easy. We refer the reader to Section 5.5 for a further discussion.

3.2 Statistical Tasks for Noisy Linear Equations over $\mathbb{Z}/q\mathbb{Z}$

Associated to noisy linear equations are several different tasks. The simplest one is *decision.*

Definition 1 (Decision). On input m pairs $\{(a_i, y_i)\}_{i=1}^m$ as well as the parameters m, q, $\mathsf{D}^{\mathrm{coeff}}, \mathsf{D}^{\mathrm{error}}$, one needs to solve the following hypothesis testing problem.

$$
\begin{aligned}
\mathsf{H}_0 : \quad & a_1, \ldots, a_m \overset{\mathrm{i.i.d.}}{\sim} \mathsf{D}^{\mathrm{coeff}} \text{ and independently } y_1, \ldots, y_m \overset{\mathrm{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}) \quad \text{versus} \\
\mathsf{H}_1 : \quad & \mathsf{NLE}(\mathrm{samples} = m, \mathrm{modulus} = q, \mathsf{D}^{\mathrm{coeff}}, \mathsf{D}^{\mathrm{error}}).
\end{aligned}
\tag{8}
$$

In words, the goal is to decide whether there is an approximate linear dependence between the samples $a_1, a_2, \ldots, a_m$ and labels $y_1, y_2, \ldots, y_m$. Another task is to recover the (approximate) linear relation under the promise that there exists one:

Definition 2 (Search). On input m pairs $\{(a_i, y_i)\}_{i=1}^m$ from

$$
\mathsf{NLE}(\mathrm{samples} = m, \mathrm{modulus} = q, \mathsf{D}^{\mathrm{coeff}}, \mathsf{D}^{\mathrm{error}})
$$

as well as the parameters m, q, $\mathsf{D}^{\mathrm{coeff}}, \mathsf{D}^{\mathrm{error}}$, recover s.

Finally, of interest is the refutation task. This task arises naturally in certain frameworks such as the sum-of-squares hierarchy when one tries to solve the search or detection task (see, e.g., [BM16; RRS17]). In words, the task is to construct a proof that the uncorrelated H_1 distribution in (8) does not correspond to *any* secret vector s. One way to formalized this would be to output an upper bound on the number of simultaneously satisfiable equations. We will use the following more general formalization: Suppose that $\mu : \mathbb{Z}/q\mathbb{Z} \longrightarrow [0,1]$ is some weight function such that $\mu(x) = 0$ if and only if $x = 0$ and $\mu(y) = \mu(-y)$. For example, $\mu(x)$ could be the indicator $\mathbb{1}[x \neq 0]$ (which will correspond to the number of unsatisfiable equations). Or, the normalized distance to the circle defined by $\mu(x) = \left(\frac{\min(|x|, q-|x|)}{q}\right)^p$ for some $p \in [1, +\infty)$(which will correspond to an L_p weight of the noise vector in LWE).

Definition 3 (Strong Refutation). Suppose that with high probability, for some $1 \geq \Delta > \delta > 0$ and[8]weight function μ, it is the case that

$$
\frac{1}{m} \sum_{i=1}^m \mu(y_i - \langle a_i, s \rangle) \geq \Delta \text{ for all } s
$$

$$
\text{if } a_1, a_2, \ldots, a_m \overset{\mathrm{i.i.d.}}{\sim} \mathsf{D}^{\mathrm{coeff}}, y_1, y_2, \ldots, y_m \overset{\mathrm{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}).
$$

An algorithm strongly (μ, δ, Δ, p)-refutes the approximate satisfiability of a random instance on input m pairs $\{(a_i, y_i)\}_{i=1}^m$ if it has the following guarantee.

1. It outputs STRONGLY UNSATISFIABLE with probability at least p if $a_1, a_2, \ldots, a_m \overset{\mathrm{i.i.d.}}{\sim} \mathsf{D}^{\mathrm{coeff}}$ and $y_1, y_2, \ldots, y_m \overset{\mathrm{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$.

2. It outputs APPROXIMATELY SATISFIABLE with probability 1 for any $\{(a_i, y_i)\}_{i=1}^m$ for which there exists some $s \in (\mathbb{Z}/q\mathbb{Z})^n$ such that $\frac{1}{m} \sum_{i=1}^m \mu(y_i - \langle a_1, s \rangle) \leq \delta$.

---

[8]The setting of $\delta = 0$ corresponds to *weak refutation* and its analogue is also studied in the literature for other natural problems such as the planted clique [BHK+19]. In the case of linear equations over finite fields, however, weak refutation is trivially solved by Gaussian elimination as also noted in [RRS17].

We call an m-tuple $\{(a_i, y_i)\}_{i=1}^m$ with the last property δ-APPROXIMATELY SATISFIABLE.

Remark 2 (On the Relative Hardness of Testing, Search, and Refutation). An algorithm solving search also solves the testing task. Similarly, under the additional guarantee for refutation that

$$\frac{1}{m} \sum_{i=1}^m \mu(y_i - \langle a_1, s \rangle) \leq \delta \text{ for some } s$$

$$\text{if } \{(a_i, y_i)\}_{i=1}^m \sim \mathsf{NLE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}}, \mathsf{D}^{\mathrm{error}})$$

with high probability, a refutation algorithm also solves the testing task. Thus, a lower-bound for testing also implies a lower bound for the other two problems.

We now give concrete instantiations of the above for the case of LPN and LWE.

3.3 Learning Parity With Noise

Based on the hardness of LPN, one can construct one-way functions, secret key identification schemes [Pie12], bit commitment schemes [JKP+12], private key encryption schemes [GRS08; ACP+09], public key encryption schemes [ABW10; KMP14; YZ16], signatures schemes with constant overhead [IKO+08], two-round oblivious transfer [DGH+20], and others. In learning theory, LPN is perhaps the most canonical problem with a statistical-computational gap and has found applications to lower bound constructions for learning probabilistic finite automata with an evaluator [KMR+94], certain classes of graphical models [MR06; BKM19], agnostically learning halfspaces [KKM+08; KK14], learning ReLU functions [GKK+17; GKK19], and others .

3.3.1 Standard LPN

As discussed before, we set $q = 2, \mathsf{D}^{\mathrm{coeff}} = \mathsf{Unif}(\{0,1\}^n)$ and $\mathsf{D}^{\mathrm{error}} = \mathsf{Bern}(\frac{1}{2} - \gamma)$ for some $\gamma > 0$. When $\gamma > 0$ is a constant, n samples are sufficient and necessary to solve the search problem *information theoretically* with high probability. The celebrated Blum-Kalai-Wasserman algorithm [BKW03] solves the search problem in time $2^{O(n/\log n)}$ with access to $2^{O(n/\log n)}$ samples. In [Lyu05], Lyubashevsky builds on the BKW algorithm and gives a $2^{O(n/\log\log n)}$-time algorithm with access to only $n^{1+\epsilon}$ samples for any constant $\epsilon > 0$. The algorithms of Blum-Kalai-Wasserman and Lyubashevsky are sate-of-the-art for standard learning parity with noise and the respective running-times are oftentimes assumed as optimal. Formal evidence for the hardness of learning parity with noise is a $\exp(\Omega(n))$ lower bound in the statistical-query model [KMR+94]. While the reader might be at first surprised that this suggests stronger hardness, it is no mistake. The statistical-query model is a restricted model of computation and does not capture the combinatorial algorithms of [BKW03; Lyu05]. Finally, [FGK+06] provides hardness evidence for LPN via reductions from learning DNF formulas and juntas.

Our Hardness Assumptions. In this current work, we will analyse the implications of the following standard hardness assumptions of standard LPN:

(LPN A1) *Near-Exponential Hardness of LPN:* Detection/Search/Refutation for standard LPN is hard in time $2^{o(n/\log n)}$.

(LPN A2) *Sub-Exponential Hardness of LPN:* Detection/Search/ Refutation for standard LPN is hard in time $2^{n^{1-\alpha}}$ for some absolute constant $\alpha \in (0, 1)$.

(LPN A3) *Quasi-Polynomial Hardness of LPN:* Detection/Search/ Refutation for standard LPN is hard in time $2^{(\log n)^{1+c}}$ for some absolute constant $c > 0$.

3.3.2   Sparse LPN

The problem of sparse LPN was introduced in [Ale03; Fei02]. In the k-sparse LPN setting, we set $q = 2$, $\mathsf{D}^{\text{coeff}} = \mathsf{randomsupport}(n, k, 1^k)$ and $\mathsf{D}^{\text{error}} = \mathsf{Bern}(\frac{1}{2} - \gamma)$ for some $\gamma > 0$. Note that such equations can equivalently be interpreted as random noisy k-XOR formulas. Concretely, viewing each k-sparse sample a_i as the indicator of $(t_i^1, t_i^2, \dots, t_i^k) \in [n]^k$, one is provided with m uniformly drawn samples

$$\left\{ \left((t_i^1, t_i^2, \dots, t_i^k), (\bigotimes_{j=1}^{k} s_{t_i^j}) \otimes e_i \right) \right\}_{i=1}^{m}. \tag{9}$$

The refutation task is at least as hard as the decision task. Further, there are search-to-decision reductions [App12; BSV19]. In particular, [BSV19] shows that under the assumption that one can solve decisional k-sparse LPN in time $\mathcal{T}$ with m samples, one can also solve the search version in time $\mathsf{poly}(\mathcal{T}, m, n)$ using $O(m^{1+2/k})$ samples.

The task of refuting noisy k-XOR has received significant attention. In this task, we take $\mu(x) = \mathbb{1}[x \neq 0]$ in Definition 3 and, hence, one aims to refute the fact that a non-trivial fraction of clauses can be satisfied. Spectral and sum-of-squares refutation algorithms achieve the following guarantees. Using $m = \tilde{O}(n^{k/2})$ samples, there exists an algorithm running in time $n^{O(k)}$ which refutes the noisy k-XOR instance [BM16; AOW15]. Further, there are sub-exponential time algorithms that work even when given less samples. For any $\delta \in (0, 1)$, there exists an algorithm running in time $\exp(\tilde{O}(n^\delta))$ which refutes the noisy k-XOR instance using $m = \tilde{O}(n^{(k/2-1)(1-\delta)+1})$ samples [RRS17].

Known Hardness Results   The following hardness results are known for k-sparse LPN. Within the sum-of-squares hierarchy, the algorithms of [BM16; MS18; RRS17] are tight up to logarithmic factors (by setting $D = n^\delta$ in the following). In particular, the canonical degree-D SoS relaxation, fails to refute an instance of the noisy k-XOR with $m = n \times \tilde{O}(n/D)^{\frac{k-1}{2}}$ samples. Further, in certain regimes, there exists further evidence of hardness against other restricted models of computation such as $\mathbf{AC^0}$ circuits and myopic algorithms, see [App12].

The only hardness evidence (conditionally) ruling out any polynomial time algorithm for refuting noisy k-XOR that we are aware of is the following win-win result. In [LYY+17], the authors show that one of the following two statements is true. Either there is a polynomial-time reduction from standard (dense) LPN with noise rate μ to a version of k-sparse LPN in which each entry of the coefficent vectors are distributed as $\mathsf{Bern}(\mu)$. Or, one can build a public key encryption scheme from standard LPN. The authors of [LYY+17] view this as evidence for the hardness of sparse LPN as public-key encryption from standard LPN with low noise has remained elusive over the years and would be a major breakthrough.

3.4   Learning With Errors and Lattice Problems

LWE is a cornerstone of modern cryptography that enables fully homomorphic encryption [BV11a] (building on Gentry's original work [Gen09; Gen10] and improvements [SV10; DGH+10; BV11b] which all rely on variants of LWE), non-interactive zero-knowledge proofs [CCH+19], public-key encryption [Reg09], succinct non-interactive argument systems [JKL+24; MPV24] and many others (see [Pei16]). More recently, the LWE assumption has also been used to establish hardness of several fundamental statistical problems

such as learning Gaussian mixture models [BRS+21; GVV22], learning (intersections of) halfspaces [KS06; DKM+22; Tie23; DKR23; Tie24], learning periodic neurons [SZB24], and others.

3.4.1 Standard LWE

In the standard learning with errors problem, one takes $D^{\text{coeff}} = \text{Unif}((\mathbb{Z}/q\mathbb{Z})^n)$ and $D^{\text{error}} = \mathcal{D}_\sigma$, where $\mathcal{D}_\sigma$ is a discrete Gaussian over the integer lattice centered at 0 (see [Reg09] for a formal definition).

There are several known algorithms for the search variant, achieving different guarantees in different regimes. First, one can adapt the Blum-Kalai-Wasserman algorithm [BKW03; Wag02] and obtain a running time of $q^{\Omega(n/\log(q/B))}$ with $q^{\Omega(n/\log(q/B))}$ samples where B is the support-size of the noise[9]. Next, an algebraic algorithm of Arora and Ge [AG11] works in time $O(n^B)$ with $n^{O(B)}$ samples. Finally, the Lenstra-Lenstra-Lovasz algorithm [LLL82] and improvements [Sch87] based on a reduction to lattice problems succeed whenever $q/B = \exp(\Omega(n))$ or $m = \exp(\Omega(n))$.

In particular, for the setting of large polynomial modulus $q = \Theta(n^C)$ (where C is a large constant) and large enough Gaussian noise D_σ (with $\sigma = n^c$ for some $c \geq 1, c \ll C$) the fastest algorithms run in time $2^{\Omega(n)}$. This is no coincidence, and is known to be tight under standard assumptions on *worst-case* lattice probblems. We do not define these problems as we will not need them in the current paper and reference [Pei16]. In particular [Reg09], let $q \leq \exp(\text{poly}(n))$ and $\beta \in (0, 1)$ be such that $\beta q \geq 2\sqrt{n}$. Then, the decisional problem for

$$\text{LWE}(\text{samples} = m, \text{modulus} = q, D^{\text{coeff}} = \text{Unif}((\mathbb{Z}/q\mathbb{Z})^n), D^{\text{error}} = D_{\beta q})$$

is at least as hard as the gap shortest vector problem and the shortest independent vectors problem on n-dimensional lattices within approximation ratio $\tilde{O}(n/\beta)$ for quantum algorithms. In [Pei09], this result is dequantized for a variation of the gap shortest vector problem and under the additional assumption $q \geq 2^{n/2}$ for the standard gap shortest vector problem. This assumption was removed in [BLP+13], however, they only showed hardness under the assumption that $\sqrt{n}$-dimensional lattice problems are hard.

Our Hardness Assumptions.   Still, a near- or sub-exponential time algorithm for n-dimensional LWE would be considered a breakthrough. Together with the known best algorithms for LWE and lattice problems, this gives rise to the following common hardness assumptions:

(**LWE A1**) *Exponential Hardness of LWE:* Detection/Search/Refutation for standard LWE

$$\text{LWE}(\text{samples} = m, \text{modulus} = q, D^{\text{coeff}} = \text{Unif}((\mathbb{Z}/q\mathbb{Z})^n), D^{\text{error}} = D_{\beta q})$$

is hard in time $2^{o(n)}$ when $\beta q > 2\sqrt{n}$.

(**LWE A2**) *Sub-Exponential Hardness of LWE:* For some constant $\beta \in (0, 1)$, Detection/Search/ Refutation for standard LWE

$$\text{LWE}(\text{samples} = m, \text{modulus} = q, D^{\text{coeff}} = \text{Unif}((\mathbb{Z}/q\mathbb{Z})^n), D^{\text{error}} = D_{\beta q})$$

is hard in time $2^{n^{1-\alpha}}$ when $\beta q > 2\sqrt{n}$ for some absolute constant $\alpha \in (0, 1)$.

---

[9]That is, there exists some set $\mathcal{B} \subseteq \mathbb{Z}/q\mathbb{Z}$ of size B such that $e_1, e_2, \dots, e_m \in \mathcal{B}$ with high probability.

(LWE A3) *Quasi-Polynomial Hardness of LWE:* For some constant $\beta \in (0, 1)$, Detection/Search/ Refutation for standard LWE

$$\mathsf{LWE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^n), \mathsf{D}^{\mathrm{error}} = D_{\beta q})$$

is hard in time $2^{(\log n)^{1+c}}$ when $\beta q > 2\sqrt{n}$ for some absolute constant $c > 0$.

3.4.2 Sparse LWE

The study of LWE with a sparse public matrix is much more recent and was initiated in [JLS24]. Specifically, they consider the case of

$$\mathsf{LWE}\left(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{randomsupport}\big(n, k, \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^k)\big), D_\sigma\right).$$

They show that when $q \geq n$, k-sparse LWE in dimension n is at least as hard as standard k-dimensional LWE. Beyond that, the authors of [JLS24] show that natural strategies based on lattice reductions, finding short vectors in the kernel of the public matrix, and variants of the BKW algorithm do not seem to break k-sparse LWE.

3.5 The Tensor Completion Problem

A noisy rank-r order-k tensor is defined by the equation

$$T = \sum_{i=1}^{r} \sigma_i \times u_i^1 \otimes u_i^2 \otimes \cdots \otimes u_i^k + \Delta \in \mathbb{R}^{n \times n \times \cdots \times n}. \tag{10}$$

where σ_i are scalars (note that they can make the norms arbitrary large), Δ is some noise tensor and $u_i^j \in \mathbb{R}^n$ for all $i \in [r], j \in [k]$ with $\|u_i^j\|_2 = \sqrt{n}$ for normalization.

Tensor Completion. In the tensor completion problem, one observes a set of randomly chosen entries of an unknown tensor. That is, for a uniformly random subset S of $[n]^k$, one observes $\{T_{i_1,i_2,\ldots,i_k}\}_{(i_1,i_2,\ldots,i_k) \in S}$. The goal is to output a tensor $\widehat{T} \in (\mathbb{R}^n)^{\otimes k}$ such that T and $\widehat{T}$ are close in some metric. This problem has numerous applications throughout the sciences including image impainting [LMW+09], hyperspectral imaging [LL10; SVD+11], geophysical imaging [KSS13]. A standard choice is the (normalized) ℓ_1 norm [BM16], i.e.

$$\frac{1}{n^k} \sum_{i_1,i_2,\ldots,i_k \in [n]^k} |T_{i_1,i_2,\ldots,i_k} - \widehat{T}_{i_1,i_2,\ldots,i_k}| = \frac{1}{n^k}\|T - \widehat{T}\|_1.$$

Of course, insofar as $\widehat{T} \in (\mathbb{R}^n)^{\otimes k}$, explicitly outputting a tensor $\widehat{T}$ requires time and space $\Omega(n^k)$. Yet, in the low-rank regime this information is compressible. A noiseless rank-r order-k tensor is uniquely specified by rnk numbers, the entries of $\{u_i^j\}_{1 \leq i \leq r, 1 \leq j \leq k}$. Given these numbers, one can evaluate an entry in $\tilde{O}(rk)$ time.

Thus, we formalize the tensor completion task via the following evaluation oracle.

Problem 1 (Tensor Completion − Search Problem). In the noisy tensor completion problem, there is a hidden noisy order-k, rank-r tensor T as in (10) where $\|\Delta\|_1 = o(\|T\|_1)$. On input a uniformly random subset S of $[n]^k$ of size m, one observes $\{T_{i_1,i_2,\dots,i_k}\}_{(i_1,i_2,\dots,i_k)\in S}$. The goal is to compute in time $\mathcal{T}$ a data structure $\widehat{T}$ with the following guarantees. First, $\widehat{T}$ should produce a real number on input $(i_1, i_2, \dots, i_k) \in [n]^k$ in time $\mathcal{T}$. Second,

$$\mathop{\mathbb{E}}_{(i_1,i_2,\dots,i_k)\sim\mathsf{Unif}([n]^k)}\left[\left|\widehat{T}[i_1, i_2, \dots, i_k] - T_{i_1,i_2,\dots,i_k}\right|\right] \leq o\left(\frac{1}{n^k}\|T\|_1\right).$$

One can also ask for the analogous refutation (with, say, $\mu(x, y) = \|x - y\|_1$) and testing tasks for tensor completion introduced in Section 2.7.

Typical Assumptions. For the problem to be solvable even by inefficient algorithms, several assumptions are necessary:

1. *Incoherence.* One needs to impose some condition that the vectors are spread-out. Otherwise, if the mass concentrates on a small set, given only random samples one may not be able to observe any non-trivial entries. The following condition is sufficient: A vector $a \in \mathbb{R}^n$ is C-incoherent if $\|a\|_\infty \leq \frac{C}{\sqrt{n}}\|a\|_2$. Then, assume that all components are incoherent for some small C [BM16; MS18].

2. *Signal-to-Noise Ratio.* Finally, one needs to impose a condition on the noise Δ. As a minimal assumption, one needs $\|\Delta\|_1 = o(\|T\|_1)$.

Algorithms. Perhaps surprisingly, without any further distributional assumptions on Δ and $\{u_i^j\}_{1\leq i\leq r, 1\leq j\leq k}$, the above three conditions are sufficient for efficient algorithms [BM16] (the related [MS18] is on noiseless tensor completion). Here we give the respective guarantees.

Theorem 3.1 (Informal, [BM16; MS18]). *Consider an instance of the tensor completion problem Problem 1 with $\mathsf{polylog}(n)$-incoherent components and $r \leq n^{k/2-1}$. Then, there exists an algorithm solving the search tensor completion problem with access to $\mathsf{polylog}(n) \times (rn^{k/2})$ samples in $\mathsf{poly}(n^{k/2})$ time.*

Lower Bounds. One may recognize that the decisional version of noisy k-XOR (sparse LPN) problem is directly related to a tensor distinguishing problem that can be solved via a tensor completion algorithm. Indeed, in the planted case of noise k-XOR, let the secret be $s \in \{\pm 1\}^n$. Then, one observes entries of the tensor $T = s^{\otimes k} + \Delta$ where each entry of Δ is in $\{-2, 0, 2\}$. Namely, $\Delta_{i_1,i_2,\dots,i_k} = 0$ in the case of noise equal to 0 and $\Delta_{i_1,i_2,\dots,i_k} = -2s_{i_1}s_{i_2}\cdots s_{i_k}$ in the case of noise equal to 1. Furthermore, the tensor is 1-incoherent. In the null case, one just observes completely independent entries that are uniform over $\{-1, 1\}$. Hence, one can translate the lower bounds for noisy k-XOR into lower-bounds for tensor completion.

In the current work, we will develop a similar bound but instead based on sparse LWE (and, via our Theorem 1.1) based on standard LWE. The advantage of an LWE lower-bound (over an LPN lower bound) is that standard LWE inherits the worst-case hardness of lattice problems as in [Reg09; Pei09]. Additionally, we obtain a lower bound based on *standard* LPN as we know it implies the hardness of *sparse* LPN.

3.6 Other Works

3.6.1 Other "Non-Standard" Variants of LPN and LWE

We remark that besides the sparse variants of LPN and LWE in this paper, other settings beyond the standard ones have also been studied. Again, the motivation is similar − additional structure in the noisy

linear equations is useful towards constructing cryptographic protocols and computational lower bounds. Different variants are useful towards different goals. In all cases, a worthwhile pursuit is to relate to the hardness of the standard variants via reductions. Such reductions are well-understood in some cases and entirely open in others.

Structured Public Matrices. The main object of study in the current work is noisy linear equations with sparse public matrices. Other structures of the public matrix have also proven useful in applications. In [AK23], the authors study LPN with a public matrix with k-sparse rows but correlated support locations which ensure a certain structure of the kernel. In [DJ24], the authors study a dense-sparse LPN in which the public matrix is a product of a sparse and dense matrix. For these two variants, to the best of our knowledge, no reduction is known from standard LPN.

The authors of [BLM+13] show a reduction from standard LWE to LWE when each entry of the public matrix A is an independent uniform bit. We will say more about the high-level idea of this reduction briefly.

Structured Secrets. Perhaps the most canonical instance of a structured secret is the case when the secret is k-sparse. But the rows of the coefficient matrix are dense (and uniform over $(\mathbb{Z}/q\mathbb{Z})^L$. Works using a sparse secret in LPN abound [Kea98; FGK+06; KK14; Val15; GKK+17; GKK19; BKM19; YYL+21; BEG+24; ABB+24; BH24]. Similarly, common is LWE with a secret drawn from a k-sparse or otherwise low-entropy distribution [GKP+10; Mic19; BD20; GVV22]. Reductions from standard LWE to LWE with a different secret distribution are somewhat well understood. It turns out that the key quantity determining the hardness is the min-entropy of the distribution [GKP+10; Mic19; BD20; GVV22]. For example, [GVV22] show that modulus-q LWE with k-sparse binary secret in dimension n is as hard as standard modulus-q LWE in dimension $\Theta(k \log(n/k)/\log q)$. In addition, for LPN, the reduction of Applebaum, Cash, Peikert and Sahai [ACP+09] can be repurposed to show that choosing the LPN secret from the error distribution has the same hardness as choosing it at random, modulo a slight difference in the number of samples.

We note that the reduction techniques for structured secrets do not seem applicable to our case of structured public matrices – as the secret is unknown in both the starting (standard) problem and target problem, no step of explicit decoding a dense secret vector into a sparse secret vector is needed and existential arguments based on the leftover hash lemma suffice. In our case, we need to explicitly decode dense row vectors of the public matrix into sparse vectors, so we need constructive rather than existential arguments.

Structured Noise. The key difference between LPN and LWE is the sparsity of the noise distribution. While LWE is typically studied with dense (discrete Gaussian) noise, LPN is often of interest with sparse noise. Yet, in certain applications it turns out that LWE with a sparse noise distribution is useful, e.g. [MP13; STA20; JLS21; JLS22; RVV24]. It turns out that the noise distribution has a profound effect on the hardness of LWE. For example, if the noise is binary, the celebrated Arora-Ge algorithm breaks LWE in polynomial time with access to quadratically many samples [AG11]. In [MP13; STA20], the authors give a direct reduction from worst-case lattice problems to LWE with binary (or otherwise small) noise, but those only apply to a regime of $n(1 + \Omega(1/\log n))$ samples.

Another setting is that of learning with rounding (LWR) [BPR12] where the samples are of the form $(a, \lfloor \frac{p \cdot a^T s}{q} \rfloor \mod p)$. One could think of LWR as LWE with a deterministic, sample-dependent, noise. Our reductions are agnostic of the noise distributions and apply to the LWR setting as well: that is, they reduce (dense) LWR to sparse LWR in the same way as they reduce LWE to sparse LWE.

Continuous LWE (CLWE).   CLWE was recently introduced in [BRS+21]. In it, the public matrix A has i.i.d. Gaussian entries, the secret s is a random unit vector, and the noise is Gaussian. The modular arithmetic is enforced in the labels, explcitly $y_i \equiv a_i^\top s + e_i \pmod 1$. In [GVV22], the authors show that CLWE and LWE are equivalent. CLWE is the starting point of many recent hardness results of statistical tasks [BRS+21; GVV22; DKM+22; DKR23; Tie23; Tie24; SZB24].

3.6.2   Memory Lower Bounds in the Streaming Model

Researchers have also considered other computational resources than running time and number of samples. For example, in [Raz17; MM17; GRT18; GKL+21] the authors study the streaming model in which the pairs (a_i, y_i) arrive sequentially. As usual in streaming algorithms, space is the main resource of interest. The state-of-the-art results show that the search problem for standard LPN takes either memory $\Omega(n^2/\epsilon)$ or time $\exp(\Omega(n))$ [GKL+21]. In the case of k-LPN for $k \leq n^9$, one needs either memory $nk^{.99}/\epsilon$ or $\exp(\Omega(k \log k))$ samples [GKL+21]. Results for larger k and other variants of LPN such as the case of a sparse secret are also covered in the aforementioned works.

4   Our Contributions

4.1   Main Reduction

Here, we state our main reduction guarantee. It holds under the following assumption on the set of parameters.

Parameter Relation 1. *Let $L, q, n, k, m, h, t \in \mathbb{N}$ are such that*

1. $n = (hq)^t, L = ht$.

2. $k \geq 4h(\log h + \log q + 2\log(ht))$ and $k = o(\sqrt{n})$.

   The reduction is slightly different depending on the specific task — testing, search, or refutation. Nevertheless, we can summarize as follows.

Theorem 4.1 (Full statements in Corollaries 5.13 to 5.15)**.** *Let $L, q, n, k, m, h, t \in \mathbb{N}$ satisfy the conditions in Parameter Relation 1. Let $\mathcal{D}$ be any distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from which we can produce a sample in time $\mathcal{T}_{\mathsf{sample}}$. Assume that there exists an algorithm that, given $m_1 := \lfloor m(1 - 3/L - k^2/n) \rfloor$ samples, runs in time $\mathcal{T}$ and solves problem $\in \{testing, search, refutation\}$ with probability p in the k-sparse setting. I.e., where the relevant distributions are*

$$\mathsf{NLE}(\mathsf{samples} = m_1, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{randomsupport}(n, k, \mathcal{D}), \mathsf{D}^{\mathrm{error}})$$

and/or $(\mathsf{randomsupport}(n, k, \mathcal{D}) \times \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}))^{\otimes m_1}$. Then, there exists an algorithm which solves in time $\mathcal{T} + O(\mathsf{poly}(n, m, \mathcal{T}_{\mathsf{sample}}))$ problem with probability $p - \exp(-2mL^{-2})$ in the standard setting. I.e., where the relevant distributions are

$$\mathsf{NLE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), \mathsf{D}^{\mathrm{error}})$$

and/or $(\mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L) \times \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}))^{\otimes m}$.

We note that our reduction is polynomial-time given that we can sample from $\mathcal{D}$ in polynomial time. For example, if $\mathcal{D}$ is just deterministically 1^k, we have $\mathcal{T}_{\mathsf{sample}} = O(1)$ deterministically.

We also discuss the probability of success. We always take $L = n^{o(1)}$ and $m = \Omega(n)$, sometimes as large as $m = n^{o(k/\mathsf{poly}(\log k, \log \log n))}$. Thus, $\exp(-2mL^{-2})$ is only a small inverse-exponential additive factor. The number of samples is preserved up to $1 - o(1)$ multiplicative factor.

The key primitive in this reduction is decoding individual equation from the standard (dense) setting to the sparse setting.

Theorem 4.2. *Suppose that Parameter Relation 1 holds. Then, there exists an algorithm (Algorithm 5) that performs the following operation in time $\mathsf{poly}(n, \mathcal{T}_{\mathsf{sample}})$.*

1. *Given $b \in (\mathbb{Z}/q\mathbb{Z})^L$, Algorithm 5 either outputs* FAIL *or outputs a k-sparse $a \in (\mathbb{Z}/q\mathbb{Z})^n$ such that $a^\top G = b^\top$. Furthermore, the algorithm outputs* FAIL *with probability at most $2/L + k^2/n$.*

2. *If $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L)$, then, conditioned on not failing, the output a of Algorithm 5 is distributed according to $\mathsf{randomsupport}(n, k, \mathcal{D})$.*

4.2 Implied Hardness

We summarize with tables the implied hardness of sparse LWE and sparse LPN based on multiple standard assumptions of the hardness of standard variants. The full statements and proofs (based on concrete choices of the relevant parameters) are in Corollaries 5.13 to 5.15.

Hardness of Sparse LPN for Decision and Search		
Assumed time complexity of standard LPN in dimension L with error distribution Bernoulli(γ).	Implied time complexity of k-sparse LPN in dimension n with error distribution Bernoulli(γ).	Restriction on sparsity
Corollary 6.3: $\exp(\Omega(L/\log L))$	$\exp(\Omega(\frac{k\log n}{\log k(\log k + \log\log n)^2}))$	$k = \omega((\log\log n)^2 \log\log\log n)$, $k = o(\sqrt{n})$
Corollary 6.4: $\exp(\Omega(L^{1-\alpha})), \alpha \in (0,1)$	$\exp(\Omega((\frac{k(\log n)}{\log k(\log k + \log\log n)})^{1-\alpha}))$	$k = \omega((\log n)^{\alpha/(1-\alpha)} \log\log n)$, $k = o(\sqrt{n})$
Corollary 6.5: $\exp(\Omega((\log L)^{1+c})), c > 0$	$\exp(\Omega((\log k)^{1+c}))$	$k = \exp(\omega((\log n)^{1/(1+c)}))$, $k = o(\sqrt{n})$

The results for LWE are similar looking.

Hardness of Sparse LWE		
Assumed time complexity of standard LWE in dimension L with modulus $q = L^\kappa$ (and error distribution $\mathsf{D}^{\mathrm{error}}$ in case of decision and search)	Implied time complexity of k-sparse LWE in dimension n with modulus q, distribution on the support $\mathcal{D}$ (and error distribution $\mathsf{D}^{\mathrm{error}}$ in case of decision and search)	Restriction on sparsity
Corollary 6.2: $\exp(\Omega(L))$	$\exp(\Omega(\frac{k\log n}{\log k(\log k + \log\log n)}))$	$k = \omega((\log\log n)(\log\log\log n))$, $k = o(\sqrt{n})$
Corollary 6.4: $\exp(\Omega(L^{1-\alpha})), \alpha \in (0,1)$	$\exp(\Omega((\frac{k(\log n)}{\log k(\log k + \log\log n)})^{1-\alpha}))$	$k = \omega((\log n)^{\alpha/(1-\alpha)}\log\log n)$, $k = o(\sqrt{n})$
Corollary 6.5: $\exp(\Omega((\log L)^{1+c})), c > 0$	$\exp(\Omega((\log k)^{1+c}))$	$k = \exp(\omega((\log n)^{1/(1+c)}))$, $k = o(\sqrt{n})$

4.3 Tightness of Results

Of course, one may wonder whether our hardness results for sparse LPN and sparse LWE are tight even under the most aggressive (near)-exponential assumptions, in which case we get a $n^{\Omega(\frac{k}{\mathsf{poly}(\log k, \log\log n)})}$ time complexity. To the best of our knowledge, even for the decision problem (which is the easiest among the three) the best known algorithm requires $O(\sqrt{q} \cdot (qn)^{k/2})$ time.

Folklore Algorithm for the Decision Problem.   Observe that there are $\binom{n}{k}(q-1)^k$ distinct k-sparse vectors over $(\mathbb{Z}/q\mathbb{Z})^n$. Recall that under both null and planted, $a_i \sim \mathsf{randomsupport}(n, k, \mathcal{D})$. Split them into $O(\sqrt{q})$ groups of size $O((qn)^{k/2})$ each. Then, with large constant probability, in at least a (say) $1/10$ fraction of the groups, one can find two indices $i_t \neq j_t$ such that $a_{i_t} = a_{j_t}$. Without loss of generality, assume this happens for groups $1, \ldots, s'$ for some ($s' = O(\sqrt{q})$). Next, in each of these groups, compute $\psi_t := y_{i_t} - y_{j_t}$. Note that under null, we have that $\psi_t \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$, while under planted $\mathcal{E}^* := \mathcal{L}(e_1 - e_2)$ where $e_1, e_2 \overset{\text{i.i.d.}}{\sim} \mathcal{E}$.[10] Under the mild condition that $\mathcal{E}^*$ is at least constantly far in total variation to $\mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$, we can now successfully test between the two cases with constant success probability since they are supported on q elements and we have $t = O(\sqrt{q})$ samples (see, e.g., [Can20]). We can amplify the success probability to $1 - \eta$ by standard amplification, if we increase the number of equations to $O(\sqrt{q}(qn)^{k/2} \cdot \log(1/\eta))$ (and runs in the same time).

Limitation of Sample-to-Sample Decoding.   Of course, one should observe that we get lower bounds of the form $n^{\frac{k}{\mathsf{poly}(\log k, \log\log n)}}$ instead of n^k. While this gap does not seem significant, one may still hope to close it. It turns out that at least in the case of LPN, the gap is unavoidable using the gadget matrix approach of decoding each equation individually. Indeed, suppose that we start from L-dimensional standard LPN and decode each $b \in \{0,1\}^L$ individually (for simplicity, say with probability 1) into some a k-sparse $a \in \{0,1\}^n$ such that $a^\top G = b^\top$ for some gadget matrix G. Since b can be recovered from a, we know that the entropy

---

[10]For a random variable X, we denote by $\mathcal{L}(X)$ the distribution of X.

of a is at least the entropy of b. That is, $\log \binom{n}{k} \geq L$, hence $k \log n \geq L$. However, the hardness of the starting L-dimensional problem is at most $\exp(L/\log L)$ due to [BKW03]. Thus, the hardness of the final problem is at most $\exp(L/\log L) \leq \exp(k \log n / \log(k \log n)) = n^{\frac{k}{\log k + \log \log n}}$. In particular, this also means that one cannot hope to get hardness using this approach when $k = O(\log \log n)$. Thus, our reduction is close to optimal within the class of reductions using gadget matrices and decoding equations individually.

Low-Degree Polynomial Lower Bounds. We give evidence for the tightness of the $n^{k/2}$ behaviour in the low-degree polynomial testing framework. The proof is standard for the low-degree polynomial method, so we defer it to Appendix A.2. It recovers the refutation lower-bound [Sch08; KMO+17] when $k = \log(n)^{O(1)}$. We remark that this result suggests hardness even when the error distribution is constantly 0. This is not a mistake, as the low-degree framework is known not to capture algorithms for the noiseless case, such as Gaussian elimination or the LLL algorithm [ZSW+22; DK22b].

Proposition 4.3. *The detection problem with any error distribution $\mathcal{E}$ and distribution on the support $\mathcal{D}$ for LWE (modulus $= q$, samples $\sim$ randomsupport$(n, k, \mathcal{D})$, secret $\sim$ Unif$((\mathbb{Z}/q\mathbb{Z})^n)$, error $\sim \mathcal{E}$) is hard for degree-D polynomial tests when $m = o\left(\frac{n}{k} \times (ne^3/kD)^{k/2-1}\right)$.*

Interpreting degree-$O(k \log n)$ polynomial tests as a proxy for algorithms running in time $n^{O(k)}$, we get $n^{O(k)}$ time hardness with $(n/k^2\mathsf{polylog}(n))^{k/2}$ samples. One must note that this result suggests that when the number of samples is significantly lower than $n^{k/2}$, say $m = \tilde{O}(n)^{(k/2-1)(1-\delta)+1}$, the distinguishing problem is hard even for degree $D = \tilde{O}(n^\delta)$ polynomials. We pose the following question as an open problem:

Problem 2. Find a reduction from standard LPN to k-sparse LPN suggesting $\exp(\tilde{O}(n^\delta))$ hardness with access to only $m = \tilde{O}(n)^{(k/2-1)(1-\delta)+1}$ samples.

In light of the recent work [CSZ24], such a reduction should be noise-aware. In the low noise-regime of Bernoulli(γ) where $\gamma \leq n^{-\frac{1+\delta}{2}}$ and $m \gg n^{\frac{1+\delta}{2}+k\frac{1-\delta}{2}}$, the work [CSZ24] exhibits a polynomial-time algorithm.

Recalling our discussion of the $n^{\frac{k}{\mathsf{poly}(\log k, \log \log n)}}$ versus n^k gap, we again note that a reduction based on separately decoding individual equations cannot give the desired result.

4.4 Application to Noisy Tensor Completion

Our lower bound for tensor completion will follow from the following theorem. It says that we can map samples from the decisional k-sparse LWE problem to samples from the following tensor distinguishing problem: Samples from the planted distributions will be transformed into i.i.d. sampled entries of a noisy low-rank tensor, and samples from the planted distributions will be transformed into i.i.d. sampled entries of a tensor with independent entries with mean $o(1)$ and constant variance. Note that a completion algorithm as outlined in Section 7 can be used to solve this distinguishing problem, by splitting the set of input samples into two, running the algorithm on the first half, and computing the empirical error on the second half.

Theorem 4.4. *There exist an algorithm running in time $O(m \log q)$ which on input*

$$\mathsf{NLE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{randomsupport}(n, k, 1^k), \mathsf{D}^{\mathrm{error}} = \mathcal{E})$$

produces m uniformly and independently sampled entries from an order-k rank-2^{k-1} real noisy tensor T with noise Δ. Furthermore, T, Δ with high probability satisfy the following assumptions.

1. Incoherence: *The noiseless tensor is $O(1)$-incoherent.*

2. Noise: *The noise Δ satisfies $\|\Delta\|_1/\|T\|_1 = O(\alpha)$, where $\alpha := \mathbb{E}_{e\sim\mathcal{E}}[\min(|e|,|q-e|)/q]$.*

Further, the same algorithm maps m i.i.d. samples of the form $(b_i, y_i) \sim \mathsf{randomsupport}(n, k, 1^k) \times \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$ to m uniformly and independently sampled entries from a tensor in which each entry is independent, has mean $o(1)$, and constant variance (assuming that $q = \omega(1)$).

We summarize the implied hardness based on standard LWE by composing Theorem 4.4 with our reductions to sparse LWE. In the case of tensors over $\mathbb{C}$, we can reduce the rank above to 1, see Theorem 7.1.

Hardness of Noisy Tensor Completion based on Standard LWE		
Assumed time complexity of standard LWE in dimension L with modulus $q = L^\kappa = \mathsf{poly}(L)$ and error distribution $D_{\beta q}, \beta q > 2\sqrt{L}$.	Implied time complexity of rank-2^{k-1} order-k noisy tensor completion with $O(1)$ incoherent entries	Restriction on the order of the tensor
(LWE A1): $\exp(\Omega(L))$	$\exp(\Omega(\frac{k\log n}{\log k(\log k+\log\log n)}))$, with error $\frac{\|\Delta\|_1}{\|T\|_1} = O\left(\left(\frac{\log k(\log k+\log\log n)}{k\log n}\right)^{\kappa-1/2}\right)$	$k = \omega((\log\log n)(\log\log\log n))$, $k = o(\sqrt{n})$
(LWE A2): $\exp(\Omega(L^{1-\alpha}))$	$\exp(\Omega((\frac{k\log n}{\log k(\log k+\log\log n)})^{1-\alpha}))$, with error $\frac{\|\Delta\|_1}{\|T\|_1} = O\left(\left(\frac{\log k(\log k+\log\log n)}{k\log n}\right)^{\kappa-1/2}\right)$	$k = \omega((\log n)^{\alpha/(1-\alpha)}(\log\log n)^2)$, $k = o(\sqrt{n})$
(LWE A3): $\exp(\Omega((\log L)^{1+c}))$	$\exp(\Omega((\log k)^{1+c}))$ with error $\frac{\|\Delta\|_1}{\|T\|_1} = O\left(\left(\frac{\log k(\log k+\log\log n)}{k\log n}\right)^{\kappa-1/2}\right).$	$k = \exp(\omega((\log n)^{1/(1+c)})),$ $k = o(\sqrt{n}).$

Perhaps surprisingly, we derive the essentially same time and sample-complexity (albeit with different noise and rank) from standard LPN. These follow directly from our reduction from standard LPN to k-sparse LPN and the discussion at the end of Section 3.5.

Hardness of Noisy Tensor Completion based on Standard LPN		
Assumed time complexity of standard LPN in dimension L with error distribution Bernoulli$(1/2 - \gamma)$.	Implied time complexity of rank-1 order-k noisy tensor completion with $O(1)$ incoherent entries.	Restriction on the order of the tensor
(LPN A1): $\exp(\Omega(L/\log L))$	$\exp(\Omega(\frac{k \log n}{\log k (\log k + \log \log n)^2}))$, with error $\frac{\|\Delta\|_1}{\|T\|_1} = 1 - 2\gamma + \tilde{O}(\frac{1}{\sqrt{m}})$	$k = \omega((\log \log n)^2 \log \log \log n)$ $k = o(\sqrt{n})$
(LPN A2): $\exp(\Omega(L^{1-\alpha}))$	$\exp(\Omega((\frac{k \log n}{\log k (\log k + \log \log n)})^{1-\alpha}))$, with error $\frac{\|\Delta\|_1}{\|T\|_1} = 1 - 2\gamma + \tilde{O}(\frac{1}{\sqrt{m}})$	$k = \omega((\log n)^{\alpha/(1-\alpha)} (\log \log n)^2)$, $k = o(\sqrt{n})$.
(LPN A3): $\exp(\Omega((\log L)^{1+c}))$	$\exp(\Omega((\log k)^{1+c}))$, with error $\frac{\|\Delta\|_1}{\|T\|_1} = 1 - 2\gamma + \tilde{O}(\frac{1}{\sqrt{m}})$	$k = \exp(\omega((\log n)^{1/(1+c)})$, $k = o(\sqrt{n})$.

4.5 Downstream Applications to Learning Theory

The decision version of sparse LPN, or equivalently, distinguishing noisy planted k-XOR instances from fully random ones, is a popular assumption to show hardness result for PAC learning and related models. It is particularly appealing for the following reason: Since sparse LPN is a parameterized assumption, often the sparsity naturally maps to some other parameter in the learning problem, so you can understand the dependence on this parameter via the dependence of sparse LPN on the sparsity. Ruling out algorithms for smaller sparsities then naturally leads to hardness results with better dependence on the relevant parameters in the learning problem.

We briefly state the relevant definitions. In PAC learning, or the related agnostic learning model, we have the following set-up: Let C be a (known) class of Boolean functions over some domain $\mathcal{X}$. In PAC learning, there is an unknown function $f^* \in C$ and we observe i.i.d. samples $(x_i, f^*(x_i))$, where $x_i \sim \mathcal{D}_x$ for some arbitrary (unknown) distribution $\mathcal{D}_x$ over $\mathcal{X}$. The task is to output an arbitrary function $\hat{f}$ (not necessarily in C) such that $\mathbb{P}_{x \sim \mathcal{D}_x}(\hat{f}(x) \neq f^*(x)) \leq \varepsilon$ for some target accuracy ε. The fact that we do not require that $\hat{f} \in C$, is referred to as *improper* learning. A weaker notion of learning, called weak learning, is to output $\hat{f}$ that achieves error better than $1/2$ (or random guessing in the case of biased labels).

The agnostic model aims to incorporate noise in the PAC model. Instead of receiving clean i.i.d. samples $(x_i, f^*(x_i))$, we observe i.i.d. samples $(x_i, y_i) \sim \mathcal{D}$ drawn from some arbitrary *joint* distribution over $\mathcal{X} \times \{0, 1\}$, and it is not necessarily the case that for all $y_i = f^*(x_i)$ for a function $f^* \in C$. The goal is to output $\hat{f}$ (again, not necessarily in C) that competes with the best classifier in C. Formally, let $\mathsf{OPT} = \inf_{f \in C} \mathbb{P}_{(x,y) \sim \mathcal{D}}(f(x) \neq y)$. We ask to output $\hat{f}$ such that $\mathbb{P}_{(x,y) \sim \mathcal{D}}(\hat{f}(x) \neq y) \leq \mathsf{OPT} + \varepsilon$. Sometimes (in both settings) we also consider the case where the marginal distribution of x is known, e.g., standard Gaussian or uniform over the hypercube.

PAC Learning DNFs. A DNF formula over n (boolean) variables is a Boolean function that is a disjunction of conjunctions. We say a DNF formula has s *terms*, if there are at most s disjunctions (there can be

arbitrarily many conjunctions in each term). DNF formulas are one of the simplest yet most fundamental function classes in learning theory. In particular, learning DNFs reduces to many other natural learning problems such as learning intersections of halfspaces, agnostically learning halfspaces, and learning sparse polynomial threshold functions. DNFs with only a constant number of terms can be learned in polynomial time [Val84]. Without restricting the number of terms, the fastest known algorithm runs in time $n^{O(n^{1/3} \log s)}$ [KS04].

To-date, the strongest known lower bound, ruling out efficient algorithms to learn $\omega(1)$-term DNFs, was shown in [DV21] under the assumption that there exist pseudo-random generators (short PRGs) with constant locality and arbitrarily large polynomial stretch. In two very recent works [BCM24; RVV24], it has been observed that in some applications, this assumption can be replaced by the sparse LPN assumption (by introducing the so-called notion of a *structured-seed local PRG*). The work of [BCM24] shows that this is also the case of the application to learning DNFs in [DV21]. In particular, letting k be large enough, combining the construction of [DV21] with their structured-seed local PRG, [BCM24, Section 6.4] show that, if there is no polynomial-time algorithm for k-sparse LPN, there is no polynomial-time algorithm to learn DNFs with $\exp(k^{\omega(1)})$ terms. Assuming that the assumption holds with $k = \omega(1)$, implies that there is no polynomial time algorithm to learn DNFs with $\omega(1)$ terms. Using Corollary 6.3, we can conclude the following, setting $k = \omega((\log \log n)^2 \log \log n)$:

Theorem 4.5. *Let $L \in \mathbb{N}$ and $\alpha \in (0, 1/2)$ be an absolute constant. Assuming the decision version of standard (dense) LPN in dimension L with constant noise rate requires time $2^{\Omega(L/\log L)}$, there is no polynomial-time algorithm for learning $\exp((\log \log n)^{\omega(1)})$-term DNFs over n variables.*

Agnostically Learning Halfspaces over Arbitrary Distributions. Daniely [Dan16] shows that under the assumption that k-sparse LPN with constant noise rates for $k = \log^s n$ for some large enough absolute constant s, takes time at least $n^{\Omega(k)}$, we cannot agnostically learn halfspaces up to any error better than $1/2$ in polynomial time (in dimension N), even if $\mathsf{OPT} = \exp(-\log^{1-\nu}(N))$ for some arbitrarily small constant ν. Let $\alpha > 0$ be a small enough constant (depending on ν). Upon careful examination of their proof, one can see that their reduction still goes through if we only assume there is no $\exp(\Omega((k \log n)^{1-\alpha}))$-time algorithm for k-sparse LPN when $k = \log^s n$.[11] Together with Corollary 6.4 we obtain the following theorem:[12]

Theorem 4.6. *Let $\nu > 0$ be an absolute constant. Then, there exists $\alpha = \alpha(\nu) < 1$ such that assuming the decision version of standard (dense) LPN with constant noise rate in dimension L takes time at least $2^{\Omega(L^{1-\alpha})}$, there is no polynomial-time algorithm to agnostically learn halfspaces, even if $\mathsf{OPT} = \exp(-\log^{1-\nu}(N))$.*

4.6 Downstream Applications to Cryptography

PRGs with Linear Stretch and small locality. Following the work of [AIK06], one can create a pseudorandom generator with linear stretch and locality $\max(k, 5)$ from the hardness of k-LPN in polynomial time. Thus, assuming near-exponential hardness **(LPN A1)** of standard LPN, we obtain a PRG with linear stretch and locality k for any $k = \omega((\log \log n)^2 \log \log \log n)$, i.e. slightly super constant locality. Under the weaker subexponential hardness **(LPN A2)**, we obtain a PRG with linear stretch and locality k for any $k = \omega((\log n)^{(1-\alpha)/\alpha}(\log \log n)^2)$.

---

[11]Specifically, set $r = \frac{k^{1-\alpha}}{\log^\alpha(n)}$ and $d = \frac{r}{\log \log k}$ in section "Connecting the dots" in Appendix 3 of [Dan16].

[12]We do not make the dependence of α on ν explicit, but offer the following explanation: It stems from the fact that for smaller ν we need to choose s larger. Since our reduction from dense to sparse LPN restricts how large k can be, we need to choose α small enough, so that we can allow $k = \log^s n$. The exact dependence can in principle be extracted from the proofs.

Multi-party Homomorphic Secret Sharing [DIJ+23] and Somewhat Homomorphic Encryption [CHK+24b].
The multiparty homomorphic secret sharing protocol in [DIJ+23] is based on the hardness of the decision
problem for k-LPN with $k = \mathsf{polylog}(n)$ against poly-time distinguishers. Our results show that such hardness follows from near-exponential hardness **(LPN A1)** or the weaker sub-exponential hardness **(LPN A2)**
of standard LPN. The same holds for the recent work of [CHK+24b] who construct a somewhat homomorphic encryption scheme from sparse LPN and the decisional Diffie-Hellman assumption.

In contrast, some other applications, most notably indistinguishability obfuscation [RVV24] requires the
hardness of k-sparse LPN for constant k. As of now, our reductions do not have any implications to this
work.

5 Full Reduction from Dense to Sparse Equations

Here, we make the outline in Section 2.5 rigorous. Recall that the key ingredient is decoding each dense
noisy linear equation individually into a sparse noisy linear equation. Formally, our procedure achieves
the following guarantee.

Theorem (Restatement of Theorem 4.2). *Suppose that Parameter Relation 1 holds. Then, there exists an
algorithm (Algorithm 5) that performs the following operation in time* $\mathsf{poly}(n, \mathcal{T}_{\mathsf{sample}}))$.

1. *Given* $b \in (\mathbb{Z}/q\mathbb{Z})^L$, *Algorithm 5 either outputs* FAIL *or outputs a k-sparse* $a \in (\mathbb{Z}/q\mathbb{Z})^n$ *such that*
   $a^\top G = b^\top$. *Furthermore, the algorithm outputs* FAIL *with probability at most* $2/L + k^2/n$.

2. *If* $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L)$, *then, conditioned on not failing, the output a of Algorithm 5 is distributed according to* $\mathsf{randomsupport}(n, k, \mathcal{D})$.

By performing this operation for all input samples, we obtain our full reduction from dense to sparse
noisy linear equations:

Algorithm 1 Full Reduction to Sparse Noisy Linear Equations

Parameters: k, L, n and Algorithm 5.
  Input: m samples (b_i, y_i) from either $\mathsf{NLE}(q, \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), \mathcal{E})$ or $\mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L) \times$
$\mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$.

Procedure:
  1. Sample $z \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^n)$.
  2. For $i = 1, 2 \ldots, m$:
    a. Run Algorithm 5 on input b_i and parameters k, L, n.
    b. If Algorithm 5 outputs FAIL, discard the sample.
    c. Else, let $a_i \in (\mathbb{Z}/q\mathbb{Z})^n$ be its output. Output the sample $(a_i, \langle a_i, z \rangle + y_i)$.

We analyze Algorithm 1 in Section 5.6 separately for the testing, search, and refutation tasks. We
remark that for the refutation task it would be sufficient to output the sample (a_i, y_i) in Step 1.c) above (but
it still is correct as written).

Before that, we give the proof of Theorem 4.2. The proof follows the steps in Section 2.5 and is organized as follows. Just like in Section 2.5, we begin with the case of binary k-sparse vectors.

1. In Section 5.1, we define the properties of the word decoding V and sentence decoding G matrices
   needed.

2. In Section 5.2 we define and analyze the word-decoding procedure for binary k-sparse vectors.

3. In Section 5.3 we define and analyze the sentence-decoding procedure for binary k-sparse vectors.

4. In Section 5.4 we show how to use the sentence decoding for binary k-sparse vectors for arbitrary distributions $\mathcal{D}$ on the support.

5. In Section 5.5 we show how to handle distributions with variable support sizes.

5.1 Word Decoding Primitives and Properties

Here we prove the three main claims about word decoding – Lemmas 2.1 to 2.3 as well as the fact that our sentence-decoding matrix is invertible.

5.1.1 Efficient Preimage Sampling: Proof of Lemma 2.1

For a vector $b \in (\mathbb{Z}/q\mathbb{Z})^h$, recall that we defines its preimage with respect to a matrix V as

$$\mathsf{Preimage}_V(b) := \left\{ (\xi(1), \xi(2), \dots, \xi(k)) \in [hq]^k \ \Big| \ \sum_{\psi=1}^{k} V_{\xi(\psi),:} = b \right\}. \tag{11}$$

Further, recall that our word-decoding matrix was defined as

$$V = \begin{pmatrix} 0 \times I_h \\ 1 \times I_h \\ \vdots \\ (q-1) \times I_h \end{pmatrix}.$$

We now restate and prove Lemma 2.1.

Lemma (Restatement of Lemma 2.1). *On input $k > h+2$ and $b \in (\mathbb{Z}/q\mathbb{Z})^h$, Algorithm 2 outputs a uniformly random sample from $\mathsf{Preimage}_V(b)$. The algorithm runs in time $\mathsf{poly}(k, h, \log q)$.*

The running time of Algorithm 2 is clearly $\mathsf{poly}(k, h, \log q)$ We proceed to correctness, which follows from the following two claims.

Claim 5.1. *The algorithm always returns. When it returns $\xi(1), \xi(2), \dots, \xi(k)$, it holds that*

$$\sum_{\ell=1}^{k} V_{\xi(\ell),:} = b.$$

Proof. **If the algorithm returns, it does so correctly.** Suppose that case 1 occurs. That is, the algorithm reaches a point where $|A_r| = k - r$. As $|A_r| \leq h < k$, it follows that $r \geq 1$.

By the definition of $\xi(r+1), \xi(r+2), \dots, \xi(k)$, it holds that

$$\sum_{\ell=r+1}^{k} V_{\xi(\ell),:} = \sum_{\ell=r+1}^{k} V_{\xi'(\ell),:} = \sum_{m \in A_r} (b - \beta_r)_j \times e_j^\top = b - \beta_r,$$

30

Algorithm 2 Sampling From $\mathsf{Preimage}_V(b)$

Parameters: k, h such that $k > h + 2$.
Input: $b \in (\mathbb{Z}/q\mathbb{Z})^h$.

Initialize: $\beta_0 = \underbrace{(0, 0, \dots, 0)}_{h} \in (\mathbb{Z}/q\mathbb{Z})^h$, $A_0 = [h]$.

Procedure: For $r = 0$ to k :

  1. If $|A_r| = k - r$, compute the $k - r$ vectors $\xi'(r+1), \xi'(r+2), \dots, \xi'(k)$ corresponding
  to the indices of $\{(b - \beta_r)_j \times e_j^\top\}_{j \in A_r}$ in V. Let π be a random permutation of $r+1, r+2, \dots, k$.
  Set $\xi(j) = \xi'(\pi(j))$. Return $\xi(1), \xi(2), \dots, \xi(k)$.
  2. Else if $|A_r| = k - r - 1$. Take $\xi(r+1)$ to be a uniformly random element indexing one of
  the zero rows in V and rows corresponding to coordinates where β_r and b differ:

$$\{1, 2, \dots, h\} \bigcup \bigcup_{m \in A_r, j \in \{1, 2, \dots, q-1\}} \{jh + m\}.$$

  3. Else if $|A_r| \le k - r - 2$, sample $\xi(r+1) \sim \mathsf{Unif}([hq])$.
  4. Compute $\beta_{r+1} = \beta_r + V_{\xi(r+1),:}$ and $A_{r+1} = \{m \in [h] \ : \ (\beta_{r+1})_m \ne b_m\}$.

---

as A_r is defined as $\{m \in [h] \ : \ (\beta_r)_m \ne b_m\}$ for $r \ge 1$. Now, observe that case 1 can only occur once as the algorithm returns when it occurs. Hence, by the definition of β_r, it holds that $\beta_r = \sum_{\ell=1}^r V_{\xi(\ell),:}$. Altogether,

$$\sum_{\ell=1}^k V_{\xi(\ell),:} = \sum_{\ell=1}^r V_{\xi(\ell),:} + \sum_{\ell=r+1}^k V_{\xi(\ell),:} = \beta_r + (b - \beta_r) = b.$$

The algorithm always returns. Now, we need to show that the algorithm always reaches case 1. Track the quantity $k - r - |A_r|$. When $r = 0$, this quantity is $k - h > 0$. When $r = k$ (if this point is ever reached), $k - r - |A_r| = -|A_k| \le 0$.

  We make the following simple observation that $k - r - |A_r| \ge 0$ at any point of time during the algorithm. This follows by induction and case analysis of which of 2 and 3 is invoked:

1. If case 2 is invoked, that is, $k - r - |A_r| = 1$, then $k - (r+1) - |A_{r+1}| \in \{0, 1\}$. The vector β_{r+1} coincides with β_r outside of coordinates indexed by $A_r = \{m \in [h] \ : \ (\beta_r)_m \ne b_m\}$, so $A_{r+1} \subseteq A_r$ and $k - (r+1) - |A_{r+1}| \ge k - (r+1) - |A_r| = 0$. Note also that β_r, β_{r+1} differ on at most one coordinate, so $|A_{r+1}| \ge |A_r| - 1$. Thus, $k - (r+1) - |A_{r+1}| \le k - (r+1) - (|A_r| - 1) = k - r - |A_r| = 1$.

2. If case 3 is invoked, that is, $k - r - |A_r| \ge 2$, then $k - (r+1) - |A_{r+1}| \ge 0$. Indeed, note that β_r and β_{r+1} differ on at most one coordinate. Thus, $|A_{r+1}| \le |A_r| + 1$ and, so, $k - (r+1) - |A_{r+1}| \ge k - (r+1) - (|A_r| + 1) \ge 0$.

In particular, this means that if the algorithm has not returned by step k yet, $0 \le k - k - |A_k| = -|A_k| \le 0$. Hence, $|A_k| = 0$. Thus, it will return at step k. $\qquad\square$

  So far, we know that the algorithm always returns correctly. We need to show that its output is uniform over $\mathsf{Preimage}_V(b)$. This follows immediately from the following claim.

Claim 5.2. *Consider any $(\xi''(1), \xi''(2), \dots, \xi''(k)) \in \mathsf{Preimage}_V(b)$. There exists a unique realization of the randomness used in [Algorithm 2](#) such that the algorithm outputs $\xi''(1), \dots, \xi''(k)$.*

Proof. We define the following quantities. Let $\beta_r'' = \sum_{j=1}^r V_{\xi''(j),:}$ and $A_r'' = \{m \in [h] \; : \; (\beta_r)_m \neq b_m\}$. Note that $\beta_k'' = b, A_k'' = \emptyset$. We make the following observation:

Observation 5.3. *For any r, the vectors β_r'' and $b = \beta_k''$ differ on at most $k - r$ coordinates.*

Proof. $\beta_k'' - \beta_r'' = \sum_{j=r+1}^k V_{\xi''(j),:}$, which is a sum of at most $k - r$ unit vectors. $\qquad\square$

Let also
$$Z = \min \left\{ z \in [k] \;\middle|\; k - r - |A_r''| = 0 \right\}.$$
This is a well-defined quantity since clearly $A_k'' = \emptyset$ by the definition of $\mathsf{Preimage}_V(b)$ and, thus, for $k = r$, $k - r - |A_r''| = 0$. Define also
$$T = \min \left\{ t \in [k] \;\middle|\; k - t - |A_t''| \leq 1 \right\}.$$
Again, T is well-defined and $T \leq Z$. We now claim that the unique randomness that leads to $(\xi(1), \xi(2), \dots, \xi(k)) = (\xi''(1), \xi''(2), \dots, \xi''(k))$ in Algorithm 2 is the following.

1. On steps $r = 0, 1, 2, \dots, T - 1$, we are in case 3 and $\xi(r + 1) = \xi''(r + 1)$.

2. On (the potentially empty set of) steps $r = T, T + 1, \dots, Z - 1$, we are in case 2 and $\xi(r+1) = \xi''(r+1)$.

3. On steps $r = Z$, we are in case 1. Then, the sampled $\xi'(r + 1), \dots, \xi'(k)$ are all distinct and are a permutation of $\{\xi''(r + 1), \dots, \xi''(k)\}$. The sampled permutation π is the unique permutation of $\{r + 1, \dots, k\}$ such that $\xi''(j) = \xi'(\pi(j))$.

We prove these claims by analyzing the algorithm step by step.

1. **Steps $r = 0, 1, 2, \dots, T - 1$.** Note that initially, $k - 0 - |A_0| = k - s \geq 2$. Hence, case 1 is invoked. It must be the case that $\xi(1) = \xi'(1)$. As $k - t - |A_t''| \geq 2$ for $t \in \{1, 2, \dots, T - 1\}$ by the definition of T, case 1 must also be invoked at steps $0, 1, 2, \dots, T - 1$.

2. **Steps $r = T, T + 1, \dots, Z - 1$.** Suppose that $T < Z$ (that is, $\{T, T + 1, \dots, Z - 1\}$ is a non-empty set of steps). Consider step T. At this point of time, $k - T - |A_T''| = 1$. Now, we claim that $\xi''(T + 1)$ must be either indexing a zero row in V or a row in which β_T, b differ, i.e.
$$\{1, 2, \dots, h\} \quad \bigcup \quad \bigcup_{m \in A_r, j \in \{1,2,\dots,q-1\}} \{jh + m\}.$$

Indeed, if indexes some other row m', then β_{T+1}'' and b would differ in $A_T'' \cup \{m'\}$ so $|A_{T+1}''| = |A_T''| + 1$. Thus, $k - |A_{T+1}''| - (T + 1) = k - |A_T''| - T - 2 = -1$. This is impossible since β_{T+1}'' and $\beta_k'' = b$ differ on at most $k - (T + 1)$ coordinates by Observation 5.3, meaning that $|A_{T+1}''| \leq k - (T + 1)$.

Altogether, this means that indeed
$$\xi''(T + 1) \in \{1, 2, \dots, h\} \quad \bigcup \quad \bigcup_{m \in A_r, j \in \{1,2,\dots,q-1\}} \{jh + m\},$$
which is exactly the admissible set for $\xi(T + 1)$ by Algorithm 2. Thus, there is a unique choice of $\xi(T + 1)$.

Finally, by the definitions of Z, T, it must be the case that
$$k - r - |A_r| = 1 \text{ for } T \leq r \leq Z - 1,$$
so there is similarly a unique choice at each step $T + 1, T + 2, \dots, Z - 1$.

32

3. **Step** Z. Finally, suppose that $r = Z$. Thus, $0 = k - r - |A_r''|$. Note that at this point of time, β_r'' and $b = \beta_k''$ differ on exactly $k - r$ coordinates. However,

$$\beta_k'' - \beta_r'' = \sum_{j=r+1}^{k} V_{\xi''(j),:}$$

is a sum of at most $k - r$ unit vectors. This means that $\beta_k'' - \beta_r''$ is a sum of *exactly* $k - r$ unit vectors, which are in some arbitrary order

$$\{(b - \beta_r'')_j \times e_j^\top\}_{j \in A_r''}.$$

As in step 1 of Algorithm 2, the algorithm correctly computes indices $\xi'(r+1), \dots, \xi'(k)$ corresponding to these vectors. In particular, $\xi''(r + 1), \dots, \xi''(k)$ is some permutation of $\xi'(r + 1), \dots, \xi'(k)$. As $\xi'(r + 1), \dots, \xi'(k)$ are distinct (since A_r indexes $k - r$ distinct unit vectors $e_j^\top$), there is a unique permutation π that sends $\xi'(r + 1), \dots, \xi'(k)$ to $\xi''(r + 1), \dots, \xi''(k)$.

Altogether, we have shown that indeed there is a unique choice of the random bits used by the algorithm that yield $\xi''(1), \xi''(2), \dots, \xi''(k)$. This is sufficient. $\square$

We also derive the following corollary from the proof.

Corollary 5.4. *Suppose that $k > h + 2$ and let $u \in (\mathbb{Z}/q\mathbb{Z})^h$. Let $K \subseteq [k]$ be such that $|K| \leq k - h - 2$. Then, if $(\xi(1), \xi(2), \dots, \xi(k)) \sim \mathsf{Unif}(\mathsf{Preimage}_V(u))$, it holds that*

$$(\xi(j))_{j \in K} \text{ is uniform over } [hq]^{|K|}.$$

Proof. Due to symmetry, it is enough to consider $K = \{1, 2, \dots, k - h - 2\}$. Note, however, that in Algorithm 2, during the first $k - h - 2$, only case 3 is invoked. This is true since $|A_r| \leq h \leq k - r - 2$ whenever $r \leq k - h - 2$. Hence, $\xi(1), \xi(2), \dots, \xi(k - h - 2)$ are uniform over $[hq]^{k-h-2}$. $\square$

5.1.2 Computing Preimage Sizes: Proof of Lemma 2.3

We next proof Lemma 2.3, restated below. V and $\mathsf{Preimage}_V(b)$ are as in the previous section.

Lemma (Restatement of Lemma 2.3). *There exists an algorithm, Algorithm 3, which on input k and $b \in (\mathbb{Z}/q\mathbb{Z})^h$, outputs $|\mathsf{Preimage}_V(b)|$. The algorithm runs in time $\mathsf{poly}(k, h, \log q)$.*

Again, the algorithm clearly works in time $\mathsf{poly}(k, h, \log q)$. We proceed to correctness. We claim that $f(t, r)$ computes exactly

$$\left| \left\{ \xi(1), \xi(2), \dots, \xi(r) \in [hq]^r \,\middle|\, \sum_{j=1}^{r} V_{\xi(j),:} = \beta \right\} \right|$$

where $\beta \in (\mathbb{Z}/q\mathbb{Z})^h$ is an arbitrary vector with t non-zero coordinates.

We first need to show that this quantity is indeed well-defined.

Lemma 5.5. *Suppose that $\beta_1, \beta_2 \in (\mathbb{Z}/q\mathbb{Z})^h$ are two vectors both of which have exactly t non-zero coordinates. Then,*

$$\left| \left\{ \xi(1), \xi(2), \dots, \xi(r) \in [hq]^r \,\middle|\, \sum_{j=1}^{r} V_{\xi(j),:} = \beta_1 \right\} \right| = \left| \left\{ \xi(1), \xi(2), \dots, \xi(r) \in [hq]^r \,\middle|\, \sum_{j=1}^{r} V_{\xi(j),:} = \beta_2 \right\} \right|.$$

Algorithm 3 Computing $|\mathsf{Preimage}_V(b)|$.

Parameters: k, h such that $k > s$.
Input: $b \in (\mathbb{Z}/q\mathbb{Z})^h$.
Preprocessing: Compute $\Delta = |\{m \in [h] \; : \; b_m \neq 0\}|$.

Initialize: Array f indexed by $\{-1, 0, 1, 2, \ldots, h, h+1\} \times \{0, 1, 2, \ldots, k\}$ where $f[0, 0] = 1$ and all other values are equal to 0.
Procedure:
For $r = 1$ to k :
     For $t = 0$ to h :
          Compute $f(t, r) = t \times f(t-1, r-1) + (h + t(q-2)) \times f(t, r-1) + (h-t)(q-1) \times f(t+1, r-1)$.
Return $f(\Delta, k)$.

Proof. Let

$$\beta_1 = \sum_{\ell=1}^{t} j_\ell^1 \times e_{\kappa(1,\ell)}^\top, \qquad \beta_2 = \sum_{\ell=1}^{t} j_\ell^2 \times e_{\kappa(2,\ell)}^\top,$$

where $j_\ell^u \in (\mathbb{Z}/q\mathbb{Z}) \backslash \{0\}$ are the non-zero coefficients and $\kappa(u, \ell) \in [h]$ are the corresponding coordinates.

We construct the following bijection between $\mathsf{Preimage}_V(\beta_1)$ and $\mathsf{Preimage}_V(\beta_2)$.

Take some $(\xi(1), \xi(2), \ldots, \xi(k)) \in \mathsf{Preimage}_V(\beta_1)$. Suppose that $\xi(i)$ corresponds to the vector $\gamma_i e_{m_i}^\top$. Now, define $\xi'(i)$ as follows:

1. If $m_i \notin \{\kappa(1, 1), \kappa(1, 2), \ldots, \kappa(1, t)\}$, then $\xi'(i) := \xi(i)$.

2. If $m_i = \kappa(1, \ell)$ for some ℓ and $m_j \neq m_i$ for any $j < i$, then define $\xi'(i)$ as the index corresponding to $(\gamma_i + j_\ell^2 - j_\ell^1) \times e_{\kappa(2,\ell)}^\top$.

3. If $m_i = \kappa(1, \ell)$ for some ℓ and there exists some $j < i$ such that $m_j = m_i$, then define $\xi'(i)$ as the index corresponding to $\gamma_i \times e_{\kappa(2,\ell)}^\top$.

One can trivially check that $(\xi'(1), \xi'(2), \ldots, \xi'(k)) \in \mathsf{Preimage}_V(\beta_2)$ and the constructed map $\xi \longrightarrow \xi'$ is a bijection between $\mathsf{Preimage}_V(\beta_1)$ and $\mathsf{Preimage}_V(\beta_2)$. $\qquad \square$

With this in mind, $\left| \left\{ \xi(1), \xi(2), \ldots, \xi(r) \in [hq]^r \; \middle| \; \sum_{j=1}^r V_{\xi(j),:} = \beta \right\} \right|$ where $\beta \in (\mathbb{Z}/q\mathbb{Z})^h$ is an arbitrary vector with t non-zero coordinates is a well-defined quantity. Now, we can show by induction on r that $f(t, r)$ computes this quantity. The base case $r = 0$ is clear. Now, we proceed to induction. Take any $\beta \in (\mathbb{Z}/q\mathbb{Z})^h$ with t non-zero coordinates and

$$\xi(1), \xi(2), \ldots, \xi(r) \in [hq]^r \text{ such that } \sum_{j=1}^r V_{\xi(j),:} = \beta.$$

Consider $\xi(r)$. There are 4 different possibilities:

1. If $V_{\xi(r),:} = \underbrace{0, 0, \ldots, 0}_{r}$. There are h such choices for $\xi(r)$. For each of them, note that $\sum_{j=1}^{r-1} V_{\xi(j),:} = \beta$.

By induction, there are $f(t, r-1)$ for $\xi(1), \xi(2), \ldots, \xi(r-1)$.

2. $V_{\xi(r),:}$ is a unit vector with support disjoint from β. Note that since β has support of size t, there are $h - t$ choices for the location of the non-zero coordinate of $V_{\xi(r),:}$ and $q - 1$ choices for its value. For each of these $(q-1)(h-t)$ choices, $\sum_{j=1}^{r-1} V_{\xi(j),:}$ has support of size $t+1$. Hence, there are $f(t+1, r-1)$ ways to choose $\xi(1), \xi(2), \ldots, \xi(r-1)$.

3. $V_{\xi(r),:} = me_j^\top$ is a unit vector such that j is in the support of β, but $\beta_j \neq m$. There are t choices for j and, for each of them, $q - 2$ choices for m. Furthermore, $\sum_{j=1}^{r-1} V_{\xi(j),:}$ has support of size t. Hence, there are $f(t, r-1)$ ways to choose $\xi(1), \xi(2), \ldots, \xi(r-1)$.

4. $V_{\xi(r),:} = me_j^\top$ is a unit vector such that j is in the support of β and $\beta_j = m$. There are t choices for j and for each, a unique choice for m. Furthermore, $\sum_{j=1}^{r-1} V_{\xi(j),:}$ has support of size $t - 1$. Hence, there are $f(t - 1, r - 1)$ ways to choose $\xi(1), \xi(2), \ldots, \xi(r-1)$.

Combining all four cases, we conclude that indeed

$$f(t,r) = h \times f(t, r-1) + (h-t)(q-1) \times f(t+1, r-1) + t(q-2) \times f(t, r-1) + t \times f(t-1, r-1)$$
$$= t \times f(t-1, r-1) + (h + t(q-2)) \times f(t, r-1) + (h-t)(q-1) \times f(t+1, r-1).$$

5.1.3 Uniformity of Preimage Size: Proof of Lemma 2.2

We next proof Lemma 2.2, restated below.

Lemma (Restatement of Lemma 2.2). *Suppose that $k \geq 4h(\log h + \log q + \log(1/\psi))$ where $\psi \in (0, 1)$ is any real number. Then, for each $u \in (\mathbb{Z}/q\mathbb{Z})^h$,*

$$\frac{(1-\psi)(hq)^k}{q^h} \leq |\mathsf{Preimage}_V(u)| \leq \frac{(1+\psi)(hq)^k}{q^h}.$$

We will rephrase Lemma 2.2 in terms of a random walk over $(\mathbb{Z}/q\mathbb{Z})^h$. At each time step j, the walk chooses a uniformly random direction $e_i^\top$ and uniformly random $m \in \mathbb{Z}/q\mathbb{Z}$ and takes a step of length m in direction $e_i^\top$. This corresponds to $\xi(j) = V_{\xi(j),:} = me_i^\top = V_{mh+i,:}$. Our goal is to show the following point-wise convergence of the random walk.

Lemma 5.6. *Consider the following random walk over $(\mathbb{Z}/q\mathbb{Z})^h$ starting at $x_0 = (0, 0, \ldots, 0)$. At each step j, a uniformly random $i \in [h]$ and $m \in \{0, 1, 2, \ldots, q-1\}$ are drawn. Then, $x_j = x_{j-1} + me_i$.*

Suppose that $k \geq 4h(\log h + \log q + \log(1/\psi))$ where $\psi \in (0, 1)$ is any number. Then, for each $u \in (\mathbb{Z}/q\mathbb{Z})^h$,

$$\frac{1-\psi}{q^h} \leq \mathbb{P}[x_k = u] \leq \frac{1+\psi}{q^h}.$$

Proof. Denote by $i_1, i_2, \ldots, i_k$ the respective draws of i in the random walk.

Lower Bound. Notice that conditioned on the fact that $\{i_1, i_2, \ldots, i_k\} = [h]$, x_k is uniform over $(\mathbb{Z}/q\mathbb{Z})^h$. This is true since to each coordinate, we have added (at least once) a uniform element from $(\mathbb{Z}/q\mathbb{Z})^h$. Hence,

$$\mathbb{P}[x_k = u] \geq \frac{1}{q^h} \mathbb{P}\big[\{i_1, i_2, \ldots, i_k\} = [h]\big].$$

Now, we bound the right-hand side as follows.

$$\mathbb{P}\big[\{i_1, i_2, \dots, i_k\} = [h]\big] \geq 1 - \sum_{i \in [h]} \mathbb{P}\big[i \notin \{i_1, i_2, \dots, i_k\}\big]$$

$$= 1 - h \times \mathbb{P}\big[1 \notin \{i_1, i_2, \dots, i_k\}\big] = 1 - h \times \prod_{j=1}^{k} \mathbb{P}[i_j \neq 1]$$

$$= 1 - h \times (1 - 1/h)^k \geq 1 - h \times (1 - 1/h)^{2h(\log h + \log \psi^{-1})}$$

$$\geq 1 - h \times e^{-(\log h + \log \psi^{-1})} = 1 - \psi.$$

Upper Bound. Suppose that u has support $A = \{i \in [h] \ : \ u_i \neq 0\}$. Hence, $A \subseteq \{i_1, i_2, \dots, i_k\}$. It follows that

$$\mathbb{P}[x_k = u] = \sum_{B \subseteq [h] \ : \ A \subseteq B} \mathbb{P}[x_k = u | \{i_1, i_2, \dots, i_k\} = B] \times \mathbb{P}[\{i_1, i_2, \dots, i_k\} = B].$$

Now, observe that if $A \subseteq B = \{i_1, i_2, \dots, i_k\}$, then $x_k|_B \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^{|B|})$. Furthermore, over the complement, we have $x_k|_{[h] \setminus B} = 0$. Thus, $\mathbb{P}[x_k = u | \{i_1, i_2, \dots, i_k\} = B] = 1/q^{|B|}$ as the values on $\{i_1, i_2, \dots, i_k\}$ should match.
   On the other hand,

$$\mathbb{P}[\{i_1, i_2, \dots, i_k\} = B] \leq \Big(\frac{|B|}{h}\Big)^k$$

since each i_j needs to be in the set B. Altogether,

$$\mathbb{P}[x_k = u] = \sum_{B \subseteq [h] \ : \ A \subseteq B} \frac{1}{q^{|B|}} \Big(\frac{|B|}{h}\Big)^k \leq \sum_{B \subseteq [h]} \frac{1}{q^{|B|}} \Big(\frac{|B|}{h}\Big)^k = \sum_{t=0}^{h} \binom{h}{t} \frac{1}{q^t} \Big(\frac{t}{h}\Big)^k = \frac{1}{q^h} + \sum_{t=1}^{h-1} \binom{h}{t} \frac{1}{q^t} \Big(\frac{t}{h}\Big)^k.$$

   To bound the last expression, we will first show that $t \longrightarrow \binom{h}{t} \frac{1}{q^t} \big(\frac{t}{h}\big)^k$ is increasing on $\{1, \dots, h-1\}$. Indeed, this is equivalent to

$$\binom{h}{t} \frac{1}{q^t} \Big(\frac{t}{h}\Big)^k \leq \binom{h}{t+1} \frac{1}{q^{t+1}} \Big(\frac{t+1}{h}\Big)^k \iff$$

$$q \leq \Big(\frac{t+1}{t}\Big)^k \times \frac{\binom{h}{t+1}}{\binom{h}{t}} \iff$$

$$q \times (t+1)/(h-t) \leq \Big(\frac{t+1}{t}\Big)^k$$

As $t \in \{1, 2, \dots, h-1\}$, it is enough to show that $qh \leq (1 + 1/h)^k$. This follows immediately as $k \geq 4h(\log h +$

$\log q$). Altogether,

$$
\begin{aligned}
\mathbb{P}[x_k = u] &\leq \frac{1}{q^h} + \sum_{t=1}^{h-1} \binom{h}{t} \frac{1}{q^t} \left(\frac{t}{h}\right)^k \\
&\leq \frac{1}{q^h} + (h-1)\binom{h}{h-1} \frac{1}{q^{h-1}} \left(\frac{h-1}{h}\right)^k \\
&= \frac{1}{q^h}\left(1 + qh(h-1)\left(\frac{h-1}{h}\right)^k\right) \\
&\leq \frac{1}{q^h}\left(1 + qh(h-1)\left(\frac{h-1}{h}\right)^{4h(\log h + \log q + \log(1/\psi))}\right) \\
&\leq \frac{1}{q^h}\left(1 + qh(h-1)e^{-2\log hq/\psi}\right) \leq \frac{1+\psi}{q^h}.
\end{aligned}
$$
$\square$

5.1.4 Full Rank of Sentence-Decoding Matrix

Our last step will be to show that our sentence-decoding matrix G is full rank, which is needed for the search reduction. Recall (cf. (5)) that $G \in (\mathbb{Z}/q\mathbb{Z})^{n \times L}$ was defined by indexing rows by $(\phi_1, \dots, \phi_t) \in [hq]^t$ and each row is defined as $G_{(\phi_1, \phi_2, \dots, \phi_t),:} = (V_{\phi_1, :}, V_{\phi_2, :}, \dots, V_{\phi_t, :})$.

Lemma 5.7. *The matrix $G \in (\mathbb{Z}/q\mathbb{Z})^{n \times L}$ defined as above is of rank L.*

Proof. We simply show that I_L is a submatrix of G. Indeed, note that for any $u \in [L]$, where $u = u_1 h + u_2$ for $0 \leq u_1 \leq t-1, 1 \leq u_2 \leq h$,

$$
G_{\underbrace{1,1,\dots,1}_{u_1},u_2,\underbrace{1,1,\dots,1}_{t-u_1-1},:} = \underbrace{0,0,\dots,0}_{u_1 h + u_2 - 1}, 1, \underbrace{0,0\dots,0}_{th - u_1 h - u_2}.
$$

$\square$

5.2 Word Decoding

Now, we show how to decode a single $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^h)$ using a good word-decoding matrix.

Algorithm 4 Word Decoding

Parameters: k, h.

Input: $b \in (\mathbb{Z}/q\mathbb{Z})^h$.

Procedure:
1. Compute $|\mathsf{Preimage}_V(b)|$ using Algorithm 3.
2. Output FAIL with probability $1 - \frac{|\mathsf{Preimage}_V(b)|}{\frac{(1+(ht)^{-2})(hq)^k}{q^h}}$ and terminate.
3. Output a uniformly random $(\xi(1), \xi(2), \dots, \xi(k))$ from $\mathsf{Preimage}_V(b)$ using Algorithm 2.

Lemma 5.8. *Suppose that $k \geq 4h(\log h + \log q + 2\log(ht))$. Then, Algorithm 4 satisfies the following*

1. *Given $b \in (\mathbb{Z}/q\mathbb{Z})^h$, it runs in time $\mathsf{poly}(k,h)$ and either outputs FAIL or outputs a k-tuple $(\xi(1), \xi(2), \dots, \xi(k)) \in [hq]^k$ such that $\sum_{i=1}^k v_{\xi(i)} = b^\top$. Furthermore, it outputs FAIL with probability at most $2 \times (ht)^{-2}$.*

2. *If $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L)$, conditioned on Algorithm 4 not failing, its output $(\xi(1), \xi(2), \dots, \xi(k))$ is distributed uniformly on $[hq]^k$.*

Proof. The running time is $\mathsf{poly}(h,k)$ by Lemmas 2.1 and 2.3. The property $\sum_{i=1}^k v_{\xi(i)} = b^\top$ follows by definition of $\mathsf{Preimage}_V(b)$.

By Lemma 2.2 for $\psi = (ht)^{-2}$, the failure probability is at most $1 - \frac{\frac{(1-(ht)^{-2})(hq)^k}{q^h}}{\frac{(1+(ht)^{-2})(hq)^k}{q^h}} = 1 - \frac{1-(ht)^{-2}}{1+(ht)^{-2}} \leq 2(ht)^{-2}$.

Note also that $1 - \frac{|\mathsf{Preimage}_V(b)|}{\frac{(1+(ht)^{-2})(hq)^k}{q^h}} \geq 0$, so it is a well-defined probability.

Finally, we analyze the marginal distribution of the output for input $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^h)$. For any fixed $(\tilde{\xi}(1), \tilde{\xi}(2), \dots, \tilde{\xi}(k))$, we compute:

$$\mathbb{P}[\text{output} = (\tilde{\xi}(1), \tilde{\xi}(2), \dots, \tilde{\xi}(k))]$$

$$= \mathbb{P}[\text{output} = (\tilde{\xi}(1), \tilde{\xi}(2), \dots, \tilde{\xi}(k)), \ \text{Algorithm does not fail at 2.}, \ b^\top = \sum_{\psi=1}^k V_{\tilde{\xi}(\psi),:}]$$

$$= \mathbb{P}[\text{output} = (\tilde{\xi}(1), \tilde{\xi}(2), \dots, \tilde{\xi}(k))| \ \text{Algorithm does not fail at 2.}, \ b^\top = \sum_{\psi=1}^k V_{\tilde{\xi}(\psi),:}] \times$$

$$\times \ \mathbb{P}[\text{Algorithm does not fail at 2.}| b^\top = \sum_{\psi=1}^k V_{\tilde{\xi}(\psi),:}] \times$$

$$\times \ \mathbb{P}[b^\top = \sum_{\psi=1}^k V_{\tilde{\xi}(\psi),:}]$$

$$= \frac{1}{|\mathsf{Preimage}_V(b)|} \times \frac{|\mathsf{Preimage}_V(b)|}{\frac{(1+(ht)^{-2})(hq)^k}{q^h}} \times \frac{1}{q^h} = \frac{1}{(1+(ht)^{-2})(hq)^k}.$$

So far, we used the facts that, conditioned on not failing, the algorithm outputs a uniform sample from $\mathsf{Preimage}_V(b)$. The algorithm does not fail with probability $\frac{|\mathsf{Preimage}_V(b)|}{\frac{(1+(ht)^{-2})(hq)^k}{q^h}}$. Finally, for any fixed $(\tilde{\xi}(1), \tilde{\xi}(2), \dots, \tilde{\xi}(k))$, we have $\mathbb{P}[b^\top = \sum_{\psi=1}^k V_{\tilde{\xi}(\psi),:}] = 1/q^h$ over $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^h)$.

As the quantity $\frac{1}{(1+(ht)^{-2})(hq)^k}$ is uniform over all choices of $(\tilde{\xi}(1), \dots, \tilde{\xi}(k))$, the output distribution is uniform over $[hq]^k$. $\qquad\square$

5.3 Sentence Decoding

We can use the word-decoding primitive Algorithm 4 to decode sentences, and prove Theorem 4.2 as follows. As mentioned before, we first describe the case when $\mathcal{D} = 1^k$, i.e., the entries on the support are deterministically 1. We handle general distributions in the next section.

Algorithm 5 Binary Sentence Decoding

Parameters: h, k, t.

Input: $b \in (\mathbb{Z}/q\mathbb{Z})^{ht}$.

Procedure:
1. Split b into $(b_1, b_2, \ldots, b_t)$ where $b_i \in (\mathbb{Z}/q\mathbb{Z})^h$.
2. For $j = 1, 2, \ldots, t$:
  a. Run Algorithm 4 on input b_j with parameters k, h.
  b. If Algorithm 4 outputs FAIL, output FAIL too and terminate.
  c. Else, store the output $(\xi_j(1), \xi_j(2), \ldots, \xi_j(k)) \in [hq]^k$.
3. Set $\xi(i) = (\xi_1(i), \xi_2(i), \ldots, \xi_t(i)) \in [hq]^t$ for $1 \leq i \leq k$.
4. If there exist $i_1 \neq i_2$ such that $\xi(i_1) = \xi(i_2)$, output FAIL and terminate.
5. Else, output $a \in (\mathbb{Z}/q\mathbb{Z})^{(hq)^t}$ such that $a_{\xi(i)} = 1$ for $1 \leq i \leq k$ and $a_j = 0$ for all other j.

Theorem 4.2 for the case of $\mathcal{D} = 1^k$ directly follows from the following lemma.

Lemma 5.9. *Suppose that $k \geq 4h(\log h + \log q + 2\log(ht))$ and $k = o(\sqrt{n}) = o((hq)^{t/2})$. Then Algorithm 5 satisfies the following*

1. *Given $b \in (\mathbb{Z}/q\mathbb{Z})^L$, it runs in time $\mathsf{poly}(n)$ and either outputs FAIL or outputs a k-sparse $a \in (\mathbb{Z}/q\mathbb{Z})^n$ such that $a^\top G = b^\top$. Furthermore, the algorithm outputs FAIL with probability at most $2/L + k^2/n$.*

2. *If $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L)$, conditioned on Algorithm 5 not failing, its output a is distributed according to $\mathsf{randomsupport}(n, k, 1^k)$.*

Proof. The running time follows from Lemma 5.8 as in Step 2 there are t calls to Algorithm 4, each taking time $\mathsf{poly}(k, h) = \mathsf{poly}(n)$ and all other steps can be implemented to run in this time as well. Conditioned on not outputting FAIL, by Step 4, the vector a we output is exactly k sparse. Further, using the guarantees of Algorithm 4 (cf. Lemma 5.8) it holds that

$$a^\top G = \sum_{i=1}^k G_{\xi(i),:} = \sum_{i=1}^k (V_{\xi_1(i),:}, V_{\xi_2(i),:}, \ldots, V_{\xi_t(i),:}) = (b_1^\top, b_2^\top, \ldots, b_t^\top) = b^\top .$$

By Lemma 5.8, each of the t iterations in Step 2 passes independently with probability at least $1 - 2 \times (ht)^{-2}$. By union bound, the algorithm does not terminate before step 3 with probability at least $1 - 2 \times h^{-2}t^{-1} \geq 1 - 2/L$ for any b. Conditioned on not terminating by Step 3, the final support indices are the columns $\xi(1), \xi(2), \ldots, \xi(k)$ of the matrix

$$\xi = \begin{pmatrix} \xi_1(1) & \xi_1(2) & \ldots & \xi_1(k) \\ \xi_2(1) & \xi_2(2) & \ldots & \xi_2(k) \\ & & \vdots & \\ \xi_t(1) & \xi_t(2) & \ldots & \xi_t(k) \end{pmatrix}.$$

By Corollary 5.4, every two columns of this matrix are uniform independent over $[hq]^t$. Hence, the probability that two given rows coincide is $(1/hq)^{-t} = 1/n$. If $k = o(\sqrt{n})$, by union bound over all pairs, it follows that the probability two rows coincide is at most $k^2/n = o(1)$.

Lastly, if $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^{ht})$, then $b_1, b_2, \ldots, b_t \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^h)$. Thus, by Lemma 5.8, the coordinates of ξ are iid numbers distributed uniformly over $[hq]$. Conditioned on $\xi(1), \xi(2), \ldots, \xi(k)$ being pairwise distinct, $\xi(1), \xi(2), \ldots, \xi(k)$ is a uniformly random subset of size k of $[hq]^t$. $\qquad\square$

Remark 3. Both Algorithm 5 and Algorithm 1 are trivially parallelizable. In Algorithm 5, one can fully parallelize the decoding of individual words. In Algorithm 1, one can fully parallelize the decoding of separate samples.

5.4 The Case of Arbitrary Distributions on the Support

We next describe how to prove Theorem 4.2 in the case of arbitrary distributions $\mathcal{D}$ on the support. Our goal is to implement a word-shift transformation that takes a pair (a, b) such that $a^\top G = b^\top$, where a is uniform over k-sparse binary vectors and produces a' such that $(a')^\top G = b^\top$ and $a' \in \mathbb{Z}/q\mathbb{Z}$ has exactly k non-zero values. These values are, furthermore, distributed according to $\mathcal{D}$ and independent of the location of the support of a. Note that this would directly imply the result, by adding it after the last step in Algorithm 5.

Lemma 5.10. *Let $\mathcal{D}$ be an arbitrary distribution on $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from which we can sample in time $\mathcal{T}_{\mathsf{sample}}$. There exists an algorithm (Algorithm 6 below) satisfying*

1. *Given $\xi \in [hq]^k$, Algorithm 6 runs in time $O(\mathcal{T}_{\mathsf{sample}} + \mathsf{poly}(k, h, t, \log q))$ and outputs $\xi' \in [(hq)^t]^k$ and $\rho \in (\mathbb{Z}/q\mathbb{Z} \setminus \{0\})^k$ such that*

$$\sum_{\psi=1}^{k} \rho_\psi \cdot G_{\xi'(\psi),:} = \sum_{\psi=1}^{k} G_{\xi(\psi),:} \, .$$

2. *If $\xi \sim \mathsf{Unif}([(hq)^t]^k)$, then $\xi' \sim \mathsf{Unif}([(hq)^t]^k)$, $\rho \sim \mathcal{D}$ and the two are independent.*

We use the following word-shift operation. Recall that we can identify $\xi_j(\psi) \in [hq]$ with $(\Delta_j(\psi), \kappa_j(\psi)) \in [h] \times \{0, 1, 2 \ldots, q-1\}$ via the transformation $\xi_j(\psi) = h \cdot \kappa_j(\psi) + \Delta_j(\psi)$.

Definition 4 (Word-Shift Transformation). Let $\xi = (\xi(1), \xi(2), \ldots, \xi(k)) \in [(hq)^t]^k$, where $\xi(\psi) = (\xi_1(\psi), \xi_2(\psi), \ldots, \xi_t(\psi)) \in [hq]^t$. Given $\rho \in ((\mathbb{Z}/q\mathbb{Z})^\times)^k$, define $T_\rho[\xi] \in [(hq)^t]^k$ as follows:

1. For each $\psi \in [k]$ and $j \in [t]$ and $\xi_j(\psi) = (\Delta_j(\psi), \kappa_j(\psi))$, define $\xi'_j(\psi) = (\Delta_j(\psi), \rho_\psi^{-1}\kappa_j(\psi))$, where arithmetic in the second coordinate is $\pmod{q}$. Note that ρ_ψ^{-1} is well-defined as it is a residue in $((\mathbb{Z}/q\mathbb{Z})^\times)$.

2. Let $T_\rho[\xi]$ be obtained by applying this transformation to all $\xi_j(\psi)$.

Note that T_ρ is bijective for every ρ since the operation performed on all $\xi_j(\psi)$ is bijective. Further, it can be computed in time $\mathsf{poly}(k, h, \log q)$ by doing the arithmetic over each of the k coordinates separately. The key property of the word-shift operation is the following.

Observation 5.11. *Recall G from Eq. (5). For any $\xi \in [(hq)^t]^k$, $\rho \in (\mathbb{Z}/q\mathbb{Z} \setminus \{0\})^k$ and $\xi' = T_\rho[\xi]$, it holds that*

$$\sum_{\psi=1}^{k} \rho_\psi \cdot G_{\xi'(\psi),:} = \sum_{\psi=1}^{k} G_{\xi(\psi),:} \, .$$

Proof.

$$\sum_{\psi=1}^{k} \rho_\psi \cdot G_{\xi'(\psi),:} = \sum_{\psi=1}^{k} \rho_\psi G_{(\xi_1'(\psi),\xi_2'(\psi),\dots,\xi_t'(\psi)),:}$$

$$= \sum_{\psi=1}^{k} \rho_\psi (V_{\xi_1'(\psi),:}, V_{\xi_2'(\psi),:}, \dots, V_{\xi_t'(\psi),:})$$

$$= \sum_{\psi=1}^{k} \rho_\psi \left(V_{\Delta_1(\psi)+h\kappa_1'(\psi),:}, V_{\Delta_2(\psi)+h\kappa_2'(\psi),:}, \dots, V_{\Delta_t(\psi)+h\kappa_t'(\psi),:} \right)$$

$$= \sum_{\psi=1}^{k} \rho_\psi \left(\kappa_1'(\psi)e_{\Delta_1(\psi)}^\top, \kappa_2'(\psi)e_{\Delta_2(\psi)}^\top, \dots, \kappa_t'(\psi)e_{\Delta_t(\psi)}^\top \right)$$

$$= \sum_{\psi=1}^{k} \left(\rho_\psi\kappa_1'(\psi)e_{\Delta_1(\psi)}^\top, \rho_\psi\kappa_2'(\psi)e_{\Delta_2(\psi)}^\top, \dots, \rho_\psi\kappa_t'(\psi)e_{\Delta_t(\psi)}^\top \right)$$

$$\equiv_q \sum_{\psi=1}^{k} \left(\kappa_1(\psi)e_{\Delta_1(\psi)}^\top, \kappa_2(\psi)e_{\Delta_2(\psi)}^\top, \dots, \kappa_t(\psi)e_{\Delta_t(\psi)}^\top \right)$$

$$= \sum_{\psi=1}^{k} \left(V_{\Delta_1(\psi)+h\kappa_1(\psi)}, V_{\Delta_2(\psi)+h\kappa_2(\psi)}, \dots, V_{\Delta_t(\psi)+h\kappa_t(\psi)} \right)$$

$$= \sum_{\psi=1}^{k} (V_{\xi_1(\psi),:}, V_{\xi_2(\psi),:}, \dots, V_{\xi_t(\psi),:})$$

$$= \sum_{\psi=1}^{k} G_{\xi(\psi),:}. \qquad \square$$

This immediately implies that the following algorithm satisfies the first property of Lemma 5.10 (the run time follows since we can sample from $\mathcal{D}$ in time $\mathcal{T}_{\mathsf{sample}}$).

Algorithm 6 Fixing Distribution on Support of Rows

    Parameters: A distribution $\mathcal{D}$ over $(\mathbb{Z}\backslash\{0\})^k$.
    Input: $\xi \in [(hq)^t]^k$
    Procedure:
    1. Sample $\rho \sim \mathcal{D}$.
    2. Output $\xi' = T_\rho[\xi]$.

It remains to show that the distributions are transformed correctly. Assume $\xi = (\xi(1), \dots, \xi(k)) \sim$ Unif$([(hq)^t]^k)$. Note that $\rho \sim \mathcal{D}$ by construction. Hence, the distribution on the support is correct. Now, we need to show that $\xi' \sim$ Unif$([(hq)^t]^k)$ as well and that it is independent from ρ. This follows since ξ is independent from ρ and the operation T_ρ is a bijection.

Modifying Algorithm 5 and proof of Theorem 4.2 We change Algorithm 5 as follows: Let $\xi \in [(hq)^t]^k$ be the k-tuple computed in Step 3 and suppose that the algorithm does not terminate on Step 4. Run Algorithm 6

with input ξ and let ρ and ξ' be its output. In Step 5, output instead $a \in (\mathbb{Z}/q\mathbb{Z})^{(hq)^t}$ such that $a_{\xi'(\psi)} = \rho_\psi$ for $1 \leq \psi \leq k$ and 0 otherwise.

Combining the first parts of Lemmas 5.9 and 5.10 it follows that this modification runs in time $O(\mathcal{T}_{\mathsf{sample}} + \mathsf{poly}(k, h, t, \log q))$ and its output a is k-sparse and satisfies $a^\top G = b^\top$. From the proof of Lemma 5.9 we know that if the input $b \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L)$, then $\xi \sim \mathsf{Unif}[(hq)^t]^k$. By Lemma 5.10 it follows that $\xi' \sim \mathsf{Unif}[(hq)^t]^k$ as well, $\rho \sim \mathcal{D}$ and the two are independent. It follows that the support of a has size exactly k and is uniformly at random and independent of ρ. On its support it takes values ρ, thus, $a \sim \mathsf{randomsupport}(n, k, \mathcal{D})$. Note that probability of failure is the same as in the binary case as the additional word-shifting operation never fails.

With this, the proof of Theorem 4.2 is complete.

5.5 Variable Support Sizes

In certain applications, one might be interested in sparse noisy linear equations where the support size varies between different samples. This, for example, is the case in the instance of [JLS24] since the distribution of entries on the k selected coordinates is $\mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$ and, hence, some entries might be zero. Similarly, one may consider the "binomial sparsity" setting of LPN where each entry is independently 1 with probability k/n.

There is a simple reduction from the variable-size support to the fixed-size support setting by keeping only the most frequent support size, thus losing at most a $1/n$ fraction of the samples. Yet, a reduction in opposite direction seems less trivial.

Nevertheless, our reduction can also easily accommodate the setting of supports of varying size. Specifically, suppose that $\mathcal{K}$ is some distribution over $\{0, 1, 2, \ldots n\}$ and $(\mathcal{D}_i)_{i=0}^n$ is a family of distributions of values on the support, where $\mathcal{D}_i$ is a distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^i$. Then, the respective problem denoted by

$$\mathsf{NLE}(\mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{randomsupport}(n, \mathcal{K}, (\mathcal{D}_i)_{i=0}^n), \mathsf{D}^{\mathrm{error}})$$

with m samples is defined as follows. Independently for each $i \in [m]$, one first draws $k_i \sim \mathcal{K}$, then draws a uniformly random set of k_i distinct indices $j_1, j_2, \ldots, j_{k_i} \in [n]$. Then, draws $\rho \sim \mathcal{D}_{k_i}$, $\rho \in ((\mathbb{Z}/q\mathbb{Z})^\times)^{k_i}$ and sets $a \in (\mathbb{Z}/q\mathbb{Z})^n$ such that $a_{j_\ell} = \rho_\ell$ for each $\ell \in [k_i]$ and $a_\xi = 0$ otherwise. Then, one draws $s \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^n)$. Finally, one draws an error $e_i \sim \mathsf{D}^{\mathrm{error}}(\langle a_i, s \rangle)$ and computes the sample (a_i, y_i) where $y_i = \langle a_i, s \rangle + e_i$. The corresponding H_0 model is defined similarly by drawing the samples a_i as described and the labels $y_1, y_2, \ldots, y_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$ The public matrix and labels are (A, y).

How to Modify Our Reduction. We modify our reduction to accommodate for varying support sizes as follows: Suppose that the distribution of support sizes is $\mathcal{K}$ where with probability $1 - \eta$ with $\eta = o(1/m)$ the support size is at least k_{min} and at most $o(\sqrt{n})$. Then, for each sample b_i, we first draw $k_i \sim \mathcal{S}$ and then use the word-decoding and sentence-decoding procedures for k_i. Note that if the conditions of our main primitives Lemmas 2.1 to 2.3 hold for k_{min}, hey also hold for any $k \geq k_{\mathsf{min}}$. With high probability, we never see $k_i < k_{\mathsf{min}}$ or $k_i = \Omega(\sqrt{n})$. We illustrate with two examples.

Example 1: The setting of binomial support. Suppose that $\mathcal{K} = \mathsf{Binomial}(k/n, n)$. This is the setting in which each entry is independently in the support with probability k/n. If $k = \omega(\log n)$, Then, by simple Chernoff bound, with probability $1 - 2\exp(-k/12)$, it is the case that $k' \sim \mathcal{K}$ satisfies $k' \in [k/2, 2k]$. Hence, one can get a high-probability reduction for $m = o(\exp(k/12))$ samples.

Note that up to the factor 12 in the exponent, this behaviour is tight. Specifically, in the case of binomial support with probability $k(1 - k/n)^{n-1} = \exp(-\Theta(k))$, a sample has a single non-zero coordinate. Thus, with $\mathsf{poly}(n) \times \exp(\Theta(k)) = \exp(O(k))$ samples, for each $i \in [n]$, there exist $\mathsf{poly}(n)$ samples a_j such that a_j has i as its unique non-zero coordinate. This makes detection and search trivial.

Example 2: The setting of [JLS24]. Similarly, in the set-up of [JLS24], the authors study sparse a_i with uniform "support" of size k, in which each value on the "support" is distributed as $\mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$. Equivalently, when q is a prime, this is the distribution $\mathsf{randomsupport}(n, \mathcal{K}, (\mathcal{D}_i)_{i=0}^{k})$ in which $k \sim \mathcal{K}$ is distributed as $\mathsf{Binomial}(1 - 1/q, k)$ and $\mathcal{D}_i = \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})\backslash\{0\})^i) = \mathsf{Unif}(((\mathbb{Z}/q\mathbb{Z})^\times)^i)$.

Again, note that with probability $k(1 - 1/q)(1/q)^k = \exp(-\Theta(k \log q))$, a sample has a single non-zero coordinate. Thus, the largest number of samples for which one may expect computational hardness is $\exp(O(k \log q))$. At the same time, with probability $1 - 2\exp(-\Omega(k \log q))$, $k' \sim \mathcal{K}$, it is the case that $k' \in [k/2, k]$. This shows that our reduction is tight up to the exact constant in the exponent.

5.6 Decision, Search, and Refutation in Our Reduction

We now show that solving one of the decision, search, refutation problems on the sparse instance solves the same problem on the dense instance. We begin by noting that with high probability our reduction preserves almost all of the input samples.

Lemma 5.12. *Suppose that the input to Algorithm 1 consists of m samples $\{(b_i, y_i)\}_{i=1}^{m}$ such that $b_i \in (\mathbb{Z}/q\mathbb{Z})^L$, $y_i \in \mathbb{Z}/q\mathbb{Z}$. Let the number of samples in the output be m_1. Then, with probability at least $1 - \exp(-2mL^{-2})$, $m_1 \geq m(1 - 3/L - k^2/n)$.*

Proof. For each sample b_i, let X_i be the indicator that Algorithm 5 does not fail on input b_i. The variables $X_1, X_2, \ldots, X_m$ are independent Bernoulli random variables with $X_i \sim \mathsf{Bernoulli}(p_i)$ for $p_i \geq 1 - 2/L - k^2/n$ by Lemma 5.9. Hence, $\mathbb{E}[\sum X_i] \geq m(1 - 2/L - k^2/n)$. As each X_i is bounded in $[0, 1]$, it is $1/4$-subgaussian. Hence, $\sum X_i$ is $m/4$-subgaussian and so it deviates from its mean by m/L with probability at most $\exp((mL^{-1})^2/(m/2))$. $\square$

In order to reduce to a fixed-size setting, our reduction will simply fail if there are less than $m_1 := \lfloor m(1 - 3/L - k^2/n) \rfloor$ successfully decoded samples and keep the first $\lfloor m(1 - 3/L - k^2/n) \rfloor$ samples in case there are at least that many samples.

Decision. We observe that Algorithm 1 maps

$$\mathsf{H}_1 \quad : \quad \mathsf{NLE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{D}^{\mathsf{coeff}} = \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), \mathsf{D}^{\mathsf{error}})$$

to

$$\mathsf{H}_1' \quad : \quad \mathsf{NLE}(\mathsf{samples} = m_1, \mathsf{modulus} = q, \mathsf{D}^{\mathsf{coeff}} = \mathsf{randomsupport}(n, k, \mathcal{D}), \mathsf{D}^{\mathsf{error}})$$

and

$$\mathsf{H}_0 \quad : \quad b_1, b_2, \ldots, b_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), y_1, y_2, \ldots, y_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z})$$

to

$$\mathsf{H}_0' \quad : \quad a_1, a_2, \ldots, a_{m_1} \overset{\text{i.i.d.}}{\sim} \mathsf{randomsupport}(n, k, \mathcal{D}), y_1, y_2, \ldots, y_{m_1} \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}).$$

Indeed, by Theorem 4.2, a sample $b_i \sim \mathsf{Unif}(\mathbb{Z}_q^L)$ is mapped to a sample $a_i \sim \mathsf{randomsupport}(n, k, \mathcal{D})$. Hence, the distribution of the samples is correct both under $\mathsf{H}_0, \mathsf{H}_1$. We need to also analyze the labels.

In the H_0 case, we have that $y := (y_1, y_2, \dots, y_m) \sim \mathsf{Unif}(\mathbb{Z}_q^m)$ and is independent of $b_1, b_2, \dots, b_m$. Let $y' \in \mathbb{Z}_q^{m_1}$ be the subset of labels corresponding to successfully decoded a_i. As the success of Dec_V is independent of labels, the vector y' is independent of $a_1, a_2, \dots, a_{m_1}$ and uniform over $\mathbb{Z}_q^{m_1}$. Hence, so is $y' + Az$. Thus, H_0 is indeed mapped to H_0'.

In the H_1 case, $y_i = b_i^\top s + e_i = a_i^\top Gs + e_i = a_i^\top (Gs) + e_i$. Hence, the final label vector y' is $A(Gs + z) + e'$. Again, note that each entry of e' is independent from A and distributed according to $\mathsf{D}^{\mathrm{error}}$. Furthermore, as $z \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^n)$, $Gs + z \sim \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^n)$ and is independent of A, e. Hence, H_1 is indeed mapped to H_1'. We record these observations as follows:

Corollary 5.13 (Testing Reduction). *Assume that the conditions in [Parameter Relation 1](#) hold and let $\mathcal{D}$ be an arbitrary distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from which we can sample in time $\mathcal{T}_{\mathsf{sample}}$. Suppose that there exists an algorithm which solves in time $\mathcal{T}$ with advantage ϵ the following sparse testing problem with $m_1 := \lfloor m(1 - 3/L - k^2/n) \rfloor$ samples:*

$$\mathsf{H}_0' \quad : \quad a_1, a_2, \dots, a_{m_1} \overset{\text{i.i.d.}}{\sim} \mathsf{randomsupport}(n, k, \mathcal{D}), y_1, y_2, \dots, y_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}),$$

$$\mathsf{H}_1' \quad : \quad \mathsf{NLE}(\mathsf{samples} = m_1, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{randomsupport}(n, k, \mathcal{D}), \mathsf{D}^{\mathrm{error}}).$$

Then, there exists an algorithm which solves in time $\mathcal{T} + \mathsf{poly}(n, m, \mathcal{T}_{\mathsf{sample}})$ with advantage $\epsilon - \exp(-2mL^{-2})$ the following standard testing problem with m samples:

$$\mathsf{H}_0 \quad : \quad b_1, b_2, \dots, b_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), y_1, y_2, \dots, y_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}),$$

$$\mathsf{H}_1 \quad : \quad \mathsf{NLE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), \mathsf{D}^{\mathrm{error}}).$$

Search. As in the case of testing, H_1 is mapped to H_1' with the new secret $s' = Gs + z$. Note, however, that z is chosen by the reduction [Algorithm 1](#) and is, hence, known. Therefore, if one finds $s' = Gs + z$, one can infer Gs. Finally, as G is full rank according to [Lemma 5.7](#), one can infer s from Gs. We record as follows:

Corollary 5.14 (Search Reduction). *Assume that the conditions in [Parameter Relation 1](#) hold and let $\mathcal{D}$ be an arbitrary distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from which we can sample in time $\mathcal{T}_{\mathsf{sample}}$. Suppose that there exists an algorithm which solves in time $\mathcal{T}$ with success probability p the following sparse search problem with $m_1 := \lfloor m(1 - 3/L - k^2/n) \rfloor$ samples:*

$$\mathsf{NLE}(\mathsf{samples} = m_1, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{randomsupport}(n, k, \mathcal{D}), \mathsf{D}^{\mathrm{error}}).$$

Then, there exists an algorithm which solves in time $\mathcal{T} + \mathsf{poly}(n, m, \mathcal{T}_{\mathsf{sample}})$ with probability $p - \exp(-2mL^{-2})$ the following standard search problem with m samples:

$$\mathsf{NLE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), \mathsf{D}^{\mathrm{error}}).$$

Refutation. Finally, consider the refutation problem [Definition 3](#). Observe that for any b that is successfully decoded into a such that $b^\top = a^\top G$, any $y \in \mathbb{Z}/q\mathbb{Z}$ and any $s \in (\mathbb{Z}/q\mathbb{Z})^L$, it is the case that

$$\mu(b^\top s - y) = \mu(a^\top Gs - y) = \mu\big(a^\top(Gs + z) - (y + a^\top z)\big).$$

In particular, this means that for any fixed secret s, the transformation between dense and sparse linear equations preserves the weight $\mu(\cdot)$. As $\mu(\cdot)$ takes values in $[0, 1]$, this means that if $a_1, a_2, \ldots, a_{m_1}$ are the output sparse vectors and $y_1', y_2', \ldots, y_{m_1}'$ the corresponding output labels, for any secret s,

$$(m - m_1) + \sum_{i=1}^{m_1} \mu\big(a_i^\top(Gs + z) - y_i'\big) \geq \sum_{i=1}^{m} \mu(b_i^\top s - y_i) \geq \sum_{i=1}^{m_1} \mu\big(a_i^\top(Gs + z) - y_i'\big).$$

Recalling that $m_1 = \lfloor m(1 - 3/L - k^2/n) \rfloor$,

$$4/L + 2k^2/n + \frac{1}{m_1} \sum_{i=1}^{m_1} \mu\big(a_i^\top(Gs + z) - y_i'\big) \geq (1 + 4/L + 2k^2/n)\frac{1}{m} \sum_{i=1}^{m} \mu(b_i^\top s - y_i)$$

$$\geq \frac{1}{m_1} \sum_{i=1}^{m_1} \mu\big(a_i^\top(Gs + z) - y_i'\big).$$

In particular, this means that a δ-APPROXIMATELY SATISFIABLE instance is mapped to a $\delta(1+4/L+2k^2/n)$-APPROXIMATELY SATISFIABLE instance whenever Algorithm 1 produces at least m_1 samples, which happens with probability at least $1 - \exp(-2mL^{-2})$.

Furthermore, a random dense instance H_0 is mapped to a random sparse instance H_0' with $m_1 = \lfloor m(1 - 3/L - k^2/n) \rfloor$ samples with probability at least $1 - \exp(-2mL^{-2})$. Again, if the starting dense instance is Δ-UNSATISFIABLE, then the target sparse dense instance is $(\Delta + 4/L + 2k^2/n)$-UNSATISFIABLE.

Altogether, this means that under a very small blow-up of parameters, one can use a refutation algorithm for the sparse instance combined with our reduction to refute the dense instance (outputting APPROXIMATELY SATISFIABLE whenever Algorithm 1 fails to produce m_1 samples).

Corollary 5.15 (Refutation Reduction). *Assume that the conditions in Parameter Relation 1 hold and let $\mathcal{D}$ be an arbitrary distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from which we can sample in time $\mathcal{T}_{\mathsf{sample}}$. Suppose that there exists an algorithm which (μ, δ, Δ, p)-refutes in time $\mathcal{T}$ the following sparse random instance with $m_1 := \lfloor m(1 - 3/L - k^2/n) \rfloor$ samples:*

$$a_1, a_2, \ldots, a_{m_1} \overset{\text{i.i.d.}}{\sim} \mathsf{randomsupport}(n, k, \mathcal{D}), \quad y_1, y_2, \ldots, y_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}).$$

Then, there exists an algorithm which $(\mu, \delta(1 - 3/L - 2k^2/n), \Delta + 3/L + 2k^2/n, p - \exp(-2mL^{-2}))$-refutes in time $\mathcal{T} + \mathsf{poly}(n, m, \mathcal{T}_{\mathsf{sample}})$ the following standard random instance with m samples:

$$b_1, b_2, \ldots, b_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}((\mathbb{Z}/q\mathbb{Z})^L), \quad y_1, y_2, \ldots, y_m \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}/q\mathbb{Z}).$$

6 Certain Parameter Instantiations

Recall that we denote the time to produce a sample from $\mathcal{D}$ by $\mathcal{T}_{\mathsf{sample}}$. Our reductions Corollaries 5.13 to 5.15 run in time $\mathcal{T} = \mathsf{poly}(n, m, \mathcal{T}_{\mathsf{sample}})$. We will assume that $\mathcal{T}_{\mathsf{sample}} = \mathsf{poly}(n)$, so the entire reduction runs in time $\mathsf{poly}(n, m)$. Hence, they imply hardness only if we assume that the starting standard (dense) problem is hard in time $\mathsf{poly}(n)$. We show how to choose our parameters based on standard assumptions about how much time is needed for the dense problem. Recall that they need to satisfy Parameter Relation 1, restated below.

Parameter Relation (Restatement of Parameter Relation 1)**.** *Let* $L, q, n, k, m, h, t \in \mathbb{N}$ *are such that*

1. $n = (hq)^t, L = ht$.

2. $k \geq 4h(\log h + \log q + 2\log(ht))$ *and* $k = o(\sqrt{n})$.

In these parameter relations, $t = \frac{\log n}{\log h + \log q}$ and $L = \frac{h \log n}{\log h + \log q}$. As the assumed hardness of standard problems increases with the dimension L, we will aim to choose the largest L possible. Clearly, $h \longrightarrow \frac{h \log n}{\log h + \log q}$ is increasing, so we will try to choose the largest h possible that satisfies Parameter Relation 1.

In LPN, we have $q = 2$ and in LWE, a typical setting is $q = L^{O(1)} = (ht)^{O(1)}$. We show that we can choose the following dimension in that case.

Observation 6.1. *If* $q \leq L^\kappa$ *for some absolute constant* κ, *we can choose*

$$L = \Theta\left(\frac{k \log n}{\log k(\log k + \log \log n)} \right).$$

Proof. In that case the inequality $k \geq 4h(\log h + \log q + 2\log(ht))$ is satisfied if $k \geq (5\kappa + 5)(h \log h + h \log t) = \Theta(h \log h + h \log t)$.

Choose $h = \lceil \frac{k}{((5\kappa+5))^2(\log k + \log \log n)} \rceil$. Hence, $t = \frac{\log n}{\log h} = \Theta(\frac{\log n}{\log k})$. Thus,

$$4h(\log h + \log q + 2(\log ht)) \leq (5\kappa + 5)(h \log h + h \log t)$$

$$\leq (5\kappa + 5)\lceil \frac{k}{(5\kappa + 5)^2(\log k + \log \log n)} \rceil \log k$$

$$+ (5\kappa + 5)\lceil \frac{k}{(5\kappa + 5)^2(\log k + \log \log n)} \rceil \log \log n \leq k.$$

It follows that $L = ht = \Theta(\frac{k \log n}{\log k(\log k + \log \log n)})$. $\qquad\square$

Corollary 6.2 (Exponential Hardness of **LWE**)**.** *Assume (LWE A1): standard testing (respectively, search, refutation) for*

$$\mathsf{LWE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{Unif}(\mathbb{Z}_q^L), \mathsf{D}^{\mathrm{error}})$$

is hard to solve with advantage p (respectively, probability of success) in time $2^{o(L)}$ with access to m samples for some $q = L^{O(1)}$.

Let k, n be such that $k = \omega\big((\log \log n)(\log \log \log n)\big), k = o(\sqrt{n})$. Let $\mathcal{D}$ be any distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from which we can sample in $\mathsf{poly}(n)$ time.

Then, testing (respectively, search and refutation) for

$$\mathsf{LWE}(\mathsf{samples} = m_1, \mathsf{modulus} = q, \mathsf{D}^{\mathrm{coeff}} = \mathsf{randomsupport}(n, k, \mathcal{D}), \mathsf{D}^{\mathrm{error}})$$

is hard to solve with advantage (respectively, probability of success) $p + \exp(-m/(k \log n)^2)$ in time $2^{o\left(\frac{k \log n}{\log k(\log k + \log \log n)}\right)}$ with access to $m_1 = m(1 - O(1/k + k^2/n))$ samples for $q = (k \log n)^{O(1)}$.

Proof. In Observation 6.1, $L = \Theta\left(\frac{k \log n}{\log k(\log k + \log \log n)} \right)$. We need $2^L = \omega(\mathsf{poly}(n))$ for the assumed hardness to be stronger than the time of the reduction. This yields $L = \omega(\log n)$ or $k = \omega((\log \log n)(\log \log \log n))$. $\qquad\square$

Corollary 6.3 (Near-Exponential Hardness of **LPN**). *Assume (LPN A1): standard testing (respectively, search, refutation) for*

$$\mathsf{LPN}(\text{samples} = m, \text{modulus} = 2, \mathsf{D}^{\text{coeff}} = \mathsf{Unif}(\mathbb{Z}_2^L), \mathsf{D}^{\text{error}})$$

is hard to solve with advantage p (respectively, probability of success) in time $2^{o(L/\log L)}$ with access to m samples.

Let k, n be such that $k = \omega\big((\log\log n)^2(\log\log\log n)\big), k = o(\sqrt{n})$.

Then, testing (respectively, search and refutation) for

$$\mathsf{LPN}(\text{samples} = m_1, \text{modulus} = 2, \mathsf{D}^{\text{coeff}} = \mathsf{randomsupport}(n, k, 1^k), \mathsf{D}^{\text{error}})$$

is hard to solve with advantage (respectively, probability of success) $p + \exp(-m/(k\log n)^2)$ in time $2^{o\left(\frac{k\log n}{\log k(\log k + \log\log n)^2}\right)}$ with access to $m_1 = m(1 - O(1/k + k^2/n))$ samples.

Proof. In Observation 6.1, $L/\log L = \Theta\left(\frac{k\log n}{\log k(\log k + \log\log n)^2}\right)$. $\qquad\qquad\square$

Corollary 6.4 (Sub-Exponential Hardness of **LPN/LWE**). *Assume (LPN A2) or (LWE A2): standard testing (respectively, search, refutation) for*

$$\mathsf{LPN}(\text{samples} = m, \text{modulus} = 2, \mathsf{D}^{\text{coeff}} = \mathsf{Unif}(\mathbb{Z}_2^L), \mathsf{D}^{\text{error}})$$

(respectively, LWE for $q = L^{O(1)}$) is hard to solve with advantage p (respectively, probability of success) in time $2^{L^{1-\alpha}}$ with access to m samples.

Let k, n be such that $k = \omega((\log n)^{\alpha/(1-\alpha)}(\log\log n)^2), k = o(\sqrt{n})$. Let $\mathcal{D}$ be any distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from which we can sample in $\mathsf{poly}(n)$ time.

Then, testing (respectively, search and refutation), for

$$\mathsf{LPN}(\text{samples} = m_1, \text{modulus} = 2, \mathsf{D}^{\text{coeff}} = \mathsf{randomsupport}(n, k, \mathcal{D}), \mathsf{D}^{\text{error}})$$

(respectively, LWE with $q = (k\log n)^{O(1)}$) is hard to solve in time $2^{o\left(\frac{k^{1-\alpha}(\log n)^{1-\alpha}}{(\log k)^{1-\alpha}(\log k + \log\log n)^{1-\alpha}}\right)}$ with $m_1 = m(1 - O(1/k + k^2/n))$ samples and advantage $p + \exp(-m/(k\log n)^2)$.

Proof. In Observation 6.1, $L^{1-\alpha} = \Theta\left(\frac{k^{1-\alpha}(\log n)^{1-\alpha}}{(\log k)^{1-\alpha}(\log k + \log\log n)^{1-\alpha}}\right)$. We need this quantity to be at least $\mathsf{poly}(n)$.

Hence, $\qquad \Theta\left(\frac{k^{1-\alpha}(\log n)^{1-\alpha}}{(\log k)^{1-\alpha}(\log k + \log\log n)^{1-\alpha}}\right) \qquad = \qquad \omega(\log n) \qquad$ which $\qquad$ leads $\qquad$ to $k = \omega((\log n)^{\alpha/(1-\alpha)}(\log\log n)^2)$. $\qquad\qquad\square$

Corollary 6.5 (Quasi-Polynomial Hardness of **LPN/LWE**). *Assume (LPN A3) or (LWE A3): standard testing (respectively, search, refutation) for*

$$\mathsf{LPN}(\text{samples} = m, \text{modulus} = 2, \mathsf{D}^{\text{coeff}} = \mathsf{Unif}(\mathbb{Z}_2^L), \mathsf{D}^{\text{error}})$$

(respectively, LWE for $q = L^{O(1)}$) is hard to solve with advantage p (respectively, probability of success) in time $2^{(\log L)^{1+c}}$ for some constant $c > 0$ with access to m samples.

Let k, n be such that $k = \exp(\omega((\log n)^{1/(1+c)})), k = o(\sqrt{n})$. Let $\mathcal{D}$ be any distribution over $((\mathbb{Z}/q\mathbb{Z})^\times)^k$ from which we can sample in $\mathsf{poly}(n)$ time.

Then, testing (respectively, search and refutation), for

$$\mathsf{LPN}(\text{samples} = m_1, \text{modulus} = 2, \mathsf{D}^{\text{coeff}} = \text{randomsupport}(n, k, \mathcal{D}), \mathsf{D}^{\text{error}})$$

(respectively, LWE with $q = (k \log n)^{O(1)}$) is hard to solve in time $2^{o\left((\log k)^{1+c}\right)}$ with $m_1 = m(1 - O(1/k + k^2/n))$ samples and advantage $p + \exp(-m/(k \log n)^2)$.

Proof. In [Observation 6.1](), $(\log L)^{1+c} = \Theta\left((\log k + \log \log n)^{1+c}\right)$. We need this quantity to be $\omega(\log n)$. Hence, $\Theta\left((\log k + \log \log n)^{1+c}\right) = \omega(\log n)$ which leads to $k = \exp(\omega((\log n)^{1/(1+c)}))$. $\qquad\square$

7 Hardness of Noisy Tensor Completion

Theorem (Restatement of [Theorem 4.4](). *There exist an algorithm running in time $O(m \log q)$ which on input*

$$\mathsf{NLE}(\text{samples} = m, \text{modulus} = q, \mathsf{D}^{\text{coeff}} = \text{randomsupport}(n, k, 1^k), \mathsf{D}^{\text{error}} = \mathcal{E})$$

produces m uniformly and independently sampled entries from an order-k rank-2^{k-1} real noisy tensor T with noise Δ. Furthermore, T, Δ with high probability satisfy the following assumptions.

1. Incoherence: *The noiseless tensor is $O(1)$-incoherent.*

2. Noise: *The noise Δ satisfies $\|\Delta\|_1/\|T\|_1 = O(\alpha)$, where $\alpha := \mathbb{E}_{e \sim \mathcal{E}}[\min(|e|, |q - e|)/q]$.*

Further, the same algorithm maps m i.i.d. samples of the form $(b_i, y_i) \sim \text{randomsupport}(n, k, 1^k) \times \text{Unif}(\mathbb{Z}/q\mathbb{Z})$ to m uniformly and independently sampled entries from a tensor in which each entry is independent, has mean $o(1)$, and constant variance (assuming that $q = \omega(1)$).

We note that assuming hardness of the distinguishing problem for LWE, we obtain hardness for the respective distinguishing problem for tensor completion and, hence, for [Problem 1]().

Our algorithm is simple and mimics the relationship between k-LPN and noisy k-XOR in the paragraph on lower bounds in [Section 3.5](). It can be summarized succinctly as follows. For a complex number z, denote by $\Re(z)$ its real part.

Algorithm 7 Reduction from Sparse Binary Noisy Linear Equations to Tensor Completion

---

Input: m samples $(a_j, y_j = \langle a_j, s \rangle + e_j)_{j=1}^m$ from $\mathsf{NLE}(m, q, \text{randomsupport}(n, k, 1^k), \mathcal{E})$.

Procedure: For each $j \in [m]$:
    1. Compute the support $(t_j^1, t_j^2, \dots, t_j^k)$ of a_j.
    2. Compute $\Re(\exp(2\pi i y_j/q))$.

Output: $\left\{ (t_j^1, t_j^2, \dots, t_j^k), \Re\left(\exp(2\pi i y_j/q)\right) \right\}_{j=1}^m$.

7.1 Analysis of the Reduction

The Null Distribution. First consider the case when each sample is drawn from randomsupport$(n, k, 1^k) \times$ Unif$(\mathbb{Z}/q\mathbb{Z})$. Clearly, all entries are independent and the indices $(t_j^1, \ldots, t_j^k)$ have the correct distributions. Further, each tensor entry we output is distributed as $\cos(2\pi \frac{x}{q})$, where x is uniform over $\mathbb{Z}/q\mathbb{Z}$. Note that as $q \to \infty$, the expectation converges to $\mathbb{E}_{\phi \sim \text{Unif}([0,2\pi])} \cos(\phi) = 0$ (since the sum approximates this integral). It follows that the mean of each entry is $o_q(1)$, since q is growing with n, it is also $o(1) = o_n(1)$. An analogous argument shows that the variance is at least $\mathbb{E}_{\phi \sim \text{Unif}([0,2\pi])} \cos^2(\phi) - o(1) = \Omega(1)$.

From Sparse Binary LWE to Complex Tensors. The main idea is that each $\langle a_i, s \rangle$ is a sum of k uniformly random terms. This is the same as the tensor, except that we take a sum instead of product. We can fix this easily by taking the multiplicative group $\mathbb{Z}/q\mathbb{Z}$ rather than the additive, i.e. represent each s_t by $\exp(2\pi i s_t/q)$. Now, complexify s by

$$z := \left(\exp\left(\frac{2\pi i s_1}{q}\right), \exp\left(\frac{2\pi i s_2}{q}\right), \ldots, \exp\left(\frac{2\pi i s_n}{q}\right) \right).$$

Then,

$$\exp\left(\frac{2\pi i y_j}{q}\right) = \exp\left(\frac{2\pi i}{q}\left(e_j + \sum_{\ell=1}^{k} s_{t_j^\ell}\right)\right) = \prod_{\ell=1}^{k} z_{t_j^\ell} \times \exp\left(\frac{2\pi i e_j}{q}\right).$$

Now, let $\Delta_{t_j^1, t_j^2, \ldots, t_j^k} := \prod_{\ell=1}^{k} z_{t_j^\ell} \times \exp(2\pi i e_j/q) - \prod_{\ell=1}^{k} z_{t_j^\ell}$. Note that

$$|\Delta_{A_j}| = \left| \prod_{\ell=1}^{k} z_{t_j^\ell} \times \exp(2\pi i e_j/q) - \prod_{\ell=1}^{k} z_{t_j^\ell} \right| \tag{12}$$
$$= |\exp(2\pi i e_j/q) - 1| \leq |\cos(2\pi i e_j/q) - 1| + |\sin(2\pi i e_j/q)| \leq 4\pi \times \max(|e_j|, |q - e_j|)/q.$$

Hence, under this transformation, $\left\{ ((t_j^1, t_j^2, \ldots, t_j^k), \exp(2\pi i y_j/q) \right\}_{j=1}^{m}$ exactly corresponds to the entries of the noisy complex tensor of order k and rank 1 given by $T_{\mathbb{C}} := z^{\otimes k} + \Delta$.

From a Complex Tensor to a Real Tensor. Allowing to blow-up the rank to 2^{k-1}, we obtain a real tensor by simply taking the real part. I.e., if $z = u + iv, \Delta = \Xi + i\Psi$, we have that

$$\mathfrak{R}(z^{\otimes k} + \Delta) = \mathfrak{R}((u + iv)^{\otimes k} + \Xi + i\Psi)$$
$$= \sum_{h \in \{0,1\}^k \,:\, 1^T h \equiv 0 \pmod 2} (-1)^{1^T h/2}(u + h_1 v) \otimes (u + h_2 v) \otimes \cdots \otimes (u + h_k v) + \Xi =: T_{\mathbb{R}}. \tag{13}$$

Our output will be

$$\left\{ ((t_j^1, t_j^2, \ldots, t_j^k), (T_{\mathbb{R}})_{t_j^1, t_j^2, \ldots, t_j^k} \right\}_{j=1}^{m}.$$

What is left is to verify the incoherence and noise rate assumptions.

Assumption 1: Incoherence. Each u_t, v_t for $1 \leq t \leq n$ has the same distribution as $\cos(2\pi \frac{x}{q})$ (or $\sin(2\pi \frac{x}{q})$ respectively) for $x \sim$ Unif$(\mathbb{Z}/q\mathbb{Z})$. Hence, using the same argument as for the null case, the random variables u_t^2 (respectively v_t^2) are bounded and have mean $\Omega(1)$. Thus, with exponentially high probability, $\|u\| = \Theta(\sqrt{n})$ and $\|v\| = \Theta(\sqrt{n})$. Hence, the incoherence is $O(1)$ with exponentially high probability.

Assumption 2: Error Size. For each $(t_j^1, t_j^2, \ldots, t_j^k)$, by Eq. (12),

$$|\Xi_{t_j^1, t_j^2, \ldots, t_j^k}| \leq |\Delta_{t_j^1, t_j^2, \ldots, t_j^k}| \leq 4\pi \times \max(|e_j|, |q - e_j|)/q.$$

Hence, a straightforward concentration argument shows that

$$\|\Delta\|_1 = O(n^k \times \mathop{\mathbb{E}}_{e \sim \mathcal{E}}[\max(|e|, |q - e|)]) = O(n^k \alpha)$$

with high probability. At the same time, with high probability, $\|T_{\mathbb{R}}\|_1 = \Theta(n^k)$. Altogether, it follows that $\|\Delta\|_1/\|T_{\mathbb{R}}\|_1 = O(\alpha)$.

7.2 Extensions

Complex Tensors. We extract from our proof a reduction to complex tensor completion for $T_{\mathbb{C}}$. Complex tensor completion is analogously defined to the real tensor completion.

Theorem 7.1. *There exist an algorithm running in time $O(m \log q)$ which on input*

$$\mathsf{NLE}(\mathsf{samples} = m, \mathsf{modulus} = q, \mathsf{sampledistribution} = \mathsf{randomsupport}(n, k, 1^k), \mathsf{errordistribution} = \mathcal{E})$$

produces m iid sampled entries from an order-k rank-1 complex noisy tensor T with noise Δ. Furthermore, T, Δ with high probability satisfies the following assumptions:

1. Incoherence: *The tensor is $O(1)$-incoherent.*

2. Noise: *The noise Δ satisfies $\|\Delta\|_1/\|T\|_1 = O(\beta)$, where $\beta := \mathbb{E}_{e \sim \mathcal{E}}[\min(|e|, |q - e|)/q]$.*

Gaussianizing the Components. One can easily turn the components into jointly Gaussian. Namely, one first samples independent Gaussians $g, h \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_n)$ in Algorithm 7. The only difference is in the output — it consists of the samples

$$\left\{(t_j^1, t_j^2, \ldots, t_j^k), \quad \Re\left(\exp(2\pi i y_j/q) \times \prod_{\ell=1}^k (g_{t_j^\ell} + i \times h_{t_j^\ell})\right)\right\}_{j=1}^m.$$

The rotational invariance of the Gaussian distribution shows that $\widetilde{z} = \widetilde{u} + i\widetilde{v}$, where $\widetilde{u}, \widetilde{v} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_n)$. In particular, this means that:

- In the complex case, the single component is drawn from the standard complex Gaussian distribution.

- In the real case, each component is of the form $\widetilde{u} + h\widetilde{v}$ for $h \in \{0, 1\}$ where $\widetilde{u}, \widetilde{v} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_n)$. Thus, the components are jointly Gaussian.

The noise assumption $\|\Delta\|_1/\|T\|_1 = O(\alpha)$, continues to hold but the incoherence increases to $O(\sqrt{\log n})$. This does not seem to have computational or information-theoretic implications.

Symmetrizing the Components. In [BM16, Lemma A.2.], the authors show a black-box reduction from a tensor-completion problem into a symmetric tensor completion problem in the case when k is odd with only a 2^{k-1}-factor blow-up in the dimension. A symmetric tensor is such that $u_j^1 = u_j^2 = \cdots = u_j^k \; \forall j$. This reduction also applies to our setting.

7.3 Implied Hardness

Hardness from Standard LWE: We can now immediately compose Theorem 4.4 with Corollaries 6.2, 6.4 and 6.5 to get hardness for Problem 1 and obtain the results in Section 4.4. We note that in the case of LWE, $\|\Delta\|_1 / \|T\|_1 = O(\beta) = O(\sqrt{\bar{L}}/q)$ as in [Reg09].

Hardness from Standard LPN: In the case of LPN, we simply use the equivalence in Section 3.5 between noisy k-XOR and tensor-completion. Each entry of T is 1. Each entry of Δ is independently 0 with probability $1/2 + \gamma$ and ± 2 with probability $1/2 - \gamma$. Hence, each entry has absolute value $1 - 2\gamma$ in expectation. With high probability, the average of m entries is $1 - 2\gamma + \tilde{O}(1/\sqrt{m})$.

References

[ABB+22]   Divesh Aggarwal, Huck Bennett, Zvika Brakerski, Alexander Golovnev, Rajendra Kumar, Zeyong Li, Spencer Peters, Noah Stephens-Davidowitz, and Vinod Vaikuntanathan. "Lattice Problems Beyond Polynomial Time". In: *CoRR* abs/2211.11693 (2022). DOI: 10.48550/ARXIV.2211.116 arXiv: 2211.11693. URL: https://doi.org/10.48550/arXiv.2211.11693 (cit. on p. 2).

[ABB+24]   Emmanuel Abbe, Enric Boix-Adsera, Matthew Brennan, Guy Bresler, and Dheeraj Nagaraj. "The staircase property: how hierarchical structure can guide deep learning". In: *Proceedings of the 35th International Conference on Neural Information Processing Systems*. Curran Associates Inc., 2024. ISBN: 9781713845393. URL: https://openreview.net/forum?id=fj6rFciApc (cit. on p. 21).

[ABH+22]   Amir Abboud, Karl Bringmann, Danny Hermelin, and Dvir Shabtay. "SETH-based Lower Bounds for Subset Sum and Bicriteria Path". In: *ACM Trans. Algorithms* 18.1 (Jan. 2022). ISSN: 1549-6325. DOI: 10.1145/3450524 (cit. on p. 9).

[ABW10]   Benny Applebaum, Boaz Barak, and Avi Wigderson. "Public-key cryptography from different assumptions". In: *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*. STOC '10. Cambridge, Massachusetts, USA: Association for Computing Machinery, 2010, pp. 171–180. ISBN: 9781450300506. DOI: 10.1145/1806689.1806715 (cit. on pp. 3, 16).

[ACP+09]   Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. "Fast Cryptographic Primitives and Circular-Secure Encryption Based on Hard Learning Problems". In: *Advances in Cryptology - CRYPTO 2009*. Ed. by Shai Halevi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 595–618. ISBN: 978-3-642-03356-8. DOI: 10.1007/978-3-642-03356-835 (cit. on pp. 16, 21).

[ADR+15]   Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. "Solving the Shortest Vector Problem in 2^n Time Using Discrete Gaussian Sampling: Extended Abstract". In: *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*. Ed. by Rocco A. Servedio and Ronitt Rubinfeld. ACM, 2015, pp. 733–742. DOI: 10.1145/2746539.2746606. URL: https://doi.org/10.1145/ (cit. on p. 2).

[AG11]   Sanjeev Arora and Rong Ge. "New Algorithms for Learning in Presence of Errors". In: *Automata, Languages and Programming*. Ed. by Luca Aceto, Monika Henzinger, and Jiří Sgall. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 403–415. ISBN: 978-3-642-22006-7. DOI: 10.1007/978-3-642-22006-734 (cit. on pp. 18, 21).

[AIK06]   Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. "On Pseudorandom Generators with Linear Stretch in NC0". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Ed. by Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 260–271. ISBN: 978-3-540-38045-0. DOI: 10.1007/1183092425 (cit. on pp. 3, 28).

[AK23]   Benny Applebaum and Eliran Kachlon. "Sampling Graphs without Forbidden Subgraphs and Unbalanced Expanders with Negligible Error". In: *SIAM Journal on Computing* 52.6 (2023), pp. 1321–1368. DOI: 10.1137/22M1484134 (cit. on p. 21).

[Ale03]   M. Alekhnovich. "More on average case vs approximation complexity". In: *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.* 2003, pp. 298–307. DOI: 10.1109/SFCS.2003.1238204 (cit. on pp. 3, 17).

[AOW15]   Sarah R. Allen, Ryan O'Donnell, and David Witmer. "How to Refute a Random CSP". In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science* (2015), pp. 689–708. DOI: 10.1109/FOCS.2015.48 (cit. on pp. 3, 17).

[App12]   Benny Applebaum. "Pseudorandom generators with long stretch and low locality from random local one-way functions". In: *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing.* STOC '12. New York, New York, USA: Association for Computing Machinery, 2012, pp. 805–816. ISBN: 9781450312455. DOI: 10.1145/2213977.2214050 (cit. on p. 17).

[BBH+21]   Matthew Brennan, Guy Bresler, Samuel B. Hopkins, Jerry Li, and Tselil Schramm. "Statistical Query Algorithms and Low-Degree Tests Are Almost Equivalent". In: *Conference on Learning Theory* (2021). URL: https://proceedings.mlr.press/v134/brennan21a.html (cit. on p. 63).

[BCM24]   Dung Bui, Geoffroy Couteau, and Nikolas Melissaris. "Structured-Seed Local Pseudorandom Generators and their Applications". In: *Cryptology ePrint Archive* (2024). URL: https://eprint.iacr.c (cit. on pp. 3, 28).

[BD20]   Zvika Brakerski and Nico Döttling. "Hardness of LWE on General Entropic Distributions". In: *Advances in Cryptology – EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part II.* Zagreb, Croatia: Springer-Verlag, 2020, pp. 551–575. ISBN: 978-3-030-45723-5. DOI: 10.1007/978-3-030-45724-2 19 (cit. on p. 21).

[BEG+24]   Boaz Barak, Benjamin L. Edelman, Surbhi Goel, Sham Kakade, Eran Malach, and Cyril Zhang. "Hidden progress in deep learning: SGD learns parities near the computational limit". In: *Proceedings of the 36th International Conference on Neural Information Processing Systems.* New Orleans, LA, USA: Curran Associates Inc., 2024. ISBN: 9781713871088. URL: https://openreview.net/ (cit. on pp. 3, 21).

[BH24]   Nader H. Bshouty and George Haddad. *Approximating the Number of Relevant Variables in a Parity Implies Proper Learning.* 2024. arXiv: 2407.11832 [cs.LG] (cit. on p. 21).

[BHK+19]   Boaz Barak, Samuel Hopkins, Jonathan Kelner, Pravesh K. Kothari, Ankur Moitra, and Aaron Potechin. "A Nearly Tight Sum-of-Squares Lower Bound for the Planted Clique Problem". In: *SIAM Journal on Computing* 48.2 (2019), pp. 687–735. DOI: 10.1137/17M1138236 (cit. on pp. 5, 15, 63).

[BKM19]   Guy Bresler, Frederic Koehler, and Ankur Moitra. "Learning restricted Boltzmann machines via influence maximization". In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing.* STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, pp. 828–839. ISBN: 9781450367059. DOI: 10.1145/3313276.3316372 (cit. on pp. 16, 21).

[BKW03]   Avrim Blum, Adam Kalai, and Hal Wasserman. "Noise-tolerant learning, the parity problem, and the statistical query model". In: *J. ACM* 50.4 (July 2003), pp. 506–519. ISSN: 0004-5411. DOI: 10.1145/792538.792543 (cit. on pp. 2, 9, 16, 18, 25, 64).

[BKW19]   Afonso S Bandeira, Dmitriy Kunisky, and Alexander S Wein. *Computational Hardness of Certifying Bounds on Constrained PCA Problems*. 2019. arXiv: `1902.07324` (cit. on p. 63).

[BLM+13]   Dan Boneh, Kevin Lewi, Hart Montgomery, and Ananth Raghunathan. "Key Homomorphic PRFs and Their Applications". In: *Advances in Cryptology – CRYPTO 2013*. Ed. by Ran Canetti and Juan A. Garay. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 410–428. ISBN: 978-3-642-40041-4. DOI: `10.1007/978-3-642-40041-423` (cit. on pp. 6, 21).

[BLP+13]   Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. "Classical hardness of learning with errors". In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 2013, pp. 575–584. DOI: `10.1145/2488608.2488680` (cit. on p. 18).

[BM16]   Boaz Barak and Ankur Moitra. "Noisy Tensor Completion via the Sum-of-Squares Hierarchy". In: *29th Annual Conference on Learning Theory*. Ed. by Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir. Vol. 49. Proceedings of Machine Learning Research. Columbia University, New York, New York, USA: PMLR, June 2016, pp. 417–445. URL: `https://proceedings.ml` (cit. on pp. 3, 5, 15, 17, 19, 20, 50).

[BPR12]   Abhishek Banerjee, Chris Peikert, and Alon Rosen. "Pseudorandom Functions and Lattices". In: *Advances in Cryptology – EUROCRYPT 2012*. Ed. by David Pointcheval and Thomas Johansson. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 719–737. ISBN: 978-3-642-29011-4 (cit. on pp. 14, 21).

[BRS+21]   Joan Bruna, Oded Regev, Min Jae Song, and Yi Tang. "Continuous LWE". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2021. Virtual, Italy: Association for Computing Machinery, 2021, pp. 694–707. ISBN: 9781450380539. DOI: `10.1145/3406325.3451000` (cit. on pp. 18, 22).

[BSV19]   Andrej Bogdanov, Manuel Sabin, and Prashant Nalini Vasudevan. "XOR codes and sparse learning parity with noise". In: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '19. San Diego, California: Society for Industrial and Applied Mathematics, 2019, pp. 986–1004. DOI: `10.5555/3310435.3310496` (cit. on p. 17).

[BSV20]   Zvika Brakerski, Noah Stephens-Davidowitz, and Vinod Vaikuntanathan. *On the Hardness of Average-case k-SUM*. 2020. arXiv: `2010.08821 [cs.CC]` (cit. on p. 9).

[BV11a]   Zvika Brakerski and Vinod Vaikuntanathan. "Efficient Fully Homomorphic Encryption from (Standard) LWE". In: *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*. 2011, pp. 97–106. DOI: `10.1109/FOCS.2011.12` (cit. on p. 17).

[BV11b]   Zvika Brakerski and Vinod Vaikuntanathan. "Fully Homomorphic Encryption from Ring-LWE and Security for Key Dependent Messages". In: *Advances in Cryptology – CRYPTO 2011*. Ed. by Phillip Rogaway. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 505–524. ISBN: 978-3-642-22792-9. DOI: `10.1007/978-3-642-22792-929` (cit. on p. 17).

[Can20]   Clément L Canonne. "A survey on distribution testing: Your data is big. But is it blue?" In: *Theory of Computing* (2020), pp. 1–100. DOI: `10.4086/toc.gs.2020.009` (cit. on p. 24).

[CCH+19]   Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Roth-blum, and Daniel Wichs. "Fiat-Shamir: from practice to theory". In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, pp. 1082–1090. ISBN: 9781450367059. DOI: 10.1145/3313276.33 (cit. on p. 17).

[CHK+24a] Henry Corrigan-Gibbs, Alexandra Henzinger, Yael Kalai, and Vinod Vaikuntanathan. *Somewhat Homomorphic Encryption from Linear Homomorphism and Sparse LPN*. Cryptology ePrint Archive, Paper 2024/1760. 2024. URL: https://eprint.iacr.org/2024/1760 (cit. on p. 2).

[CHK+24b] Henry Corrigan-Gibbs, Alexandra Henzinger, Yael Kalai, and Vinod Vaikuntanathan. *Somewhat Homomorphic Encryption from Linear Homomorphism and Sparse LPN*. Cryptology ePrint Archive, Paper 2024/1760. 2024. URL: https://eprint.iacr.org/2024/1760 (cit. on pp. 3, 29).

[CRR21]   Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. "Silver: Silent VOLE and Oblivious Transfer from Hardness of Decoding Structured LDPC Codes". In: *Advances in Cryptology – CRYPTO 2021*. Ed. by Tal Malkin and Chris Peikert. Cham: Springer International Publishing, 2021, pp. 502–534. ISBN: 978-3-030-84252-9. DOI: 10.1007/978-3-030-84252-9_17 (cit. on p. 3).

[CSZ24]   Xue Chen, Wenxuan Shu, and Zhaienhe Zhou. *Algorithms for Sparse LPN and LSPN Against Low-noise*. 2024. arXiv: 2407.19215 [cs.CR]. URL: https://arxiv.org/abs/2407.19215 (cit. on p. 25).

[Dan16]   Amit Daniely. "Complexity theoretic limitations on learning halfspaces". In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 2016, pp. 105–117. DOI: 10.1145/2897518.2897520 (cit. on pp. 3, 28).

[DGH+10]   Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. "Fully Homomorphic Encryption over the Integers". In: *Advances in Cryptology – EUROCRYPT 2010*. Ed. by Henri Gilbert. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 24–43. ISBN: 978-3-642-13190-5. DOI: 10.1007/978-3-642-13190-5_2 (cit. on p. 17).

[DGH+20]   Nico Döttling, Sanjam Garg, Mohammad Hajiabadi, Daniel Masny, and Daniel Wichs. "Two-Round Oblivious Transfer from CDH or LPN". In: *Advances in Cryptology – EUROCRYPT 2020*. Ed. by Anne Canteaut and Yuval Ishai. Cham: Springer International Publishing, 2020, pp. 768–797. ISBN: 978-3-030-45724-2. DOI: 10.1007/978-3-030-45724-2_26 (cit. on p. 16).

[DIJ+23]   Quang Dao, Yuval Ishai, Aayush Jain, and Huijia Lin. "Multi-party Homomorphic Secret Sharing and Sublinear MPC from Sparse LPN". In: *Advances in Cryptology – CRYPTO 2023*. Ed. by Helena Handschuh and Anna Lysyanskaya. Cham: Springer Nature Switzerland, 2023, pp. 315–348. ISBN: 978-3-031-38545-2. DOI: 10.1007/978-3-031-38545-2_11 (cit. on pp. 3, 29).

[DJ24]   Quang Dao and Aayush Jain. *Lossy Cryptography from Code-Based Assumptions*. Cryptology ePrint Archive, Paper 2024/175. 2024. URL: https://eprint.iacr.org/2024/175 (cit. on p. 21).

[DK22a]   Ilias Diakonikolas and Daniel Kane. "Non-Gaussian Component Analysis via Lattice Basis Reduction". In: *Proceedings of Thirty Fifth Conference on Learning Theory*. Ed. by Po-Ling Loh and Maxim Raginsky. Vol. 178. Proceedings of Machine Learning Research. PMLR, July 2022, pp. 4535–4547. URL: https://proceedings.mlr.press/v178/diakonikolas22d.html (cit. on p. 64).

[DK22b]   Ilias Diakonikolas and Daniel Kane. "Non-gaussian component analysis via lattice basis reduction". In: *Conference on Learning Theory*. PMLR. 2022, pp. 4535–4547 (cit. on p. 25).

[DKM+22]   Ilias Diakonikolas, Daniel Kane, Pasin Manurangsi, and Lisheng Ren. "Cryptographic hardness of learning halfspaces with massart noise". In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 3624–3636. DOI: 10.5555/3600270.3600532 (cit. on pp. 18, 22).

[DKR23]   Ilias Diakonikolas, Daniel Kane, and Lisheng Ren. "Near-optimal cryptographic hardness of agnostically learning halfspaces and relu regression under gaussian marginals". In: *International Conference on Machine Learning*. PMLR. 2023, pp. 7922–7938. DOI: 10.5555/3618408.3618722 (cit. on pp. 18, 22).

[DMN12]   Nico Döttling, Jörn Müller-Quade, and Anderson C. A. Nascimento. "IND-CCA Secure Cryptography Based on a Variant of the LPN Problem". In: *Advances in Cryptology – ASIACRYPT 2012*. Ed. by Xiaoyun Wang and Kazue Sako. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 485–503. ISBN: 978-3-642-34961-4. DOI: 10.1007/978-3-642-34961-430 (cit. on p. 3).

[DMW23]   Abhishek Dhawan, Cheng Mao, and Alexander S. Wein. *Detection of Dense Subhypergraphs by Low-Degree Polynomials*. 2023. arXiv: 2304.08135 [cs.DS] (cit. on p. 63).

[DV21]   Amit Daniely and Gal Vardi. "From local pseudorandom generators to hardness of learning". In: *Conference on Learning Theory*. PMLR. 2021, pp. 1358–1394. URL: https://proceedings.mlr.pre (cit. on pp. 3, 28).

[Fei02]   Uriel Feige. "Relations between average case complexity and approximation complexity". In: *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*. STOC '02. Montreal, Quebec, Canada: Association for Computing Machinery, 2002, pp. 534–543. ISBN: 1581134959. DOI: 10.1145/509907.509985 (cit. on pp. 3, 17).

[FGK+06]   Vitaly Feldman, Parikshit Gopalan, Subhash Khot, and Ashok Kumar Ponnuswami. "New results for learning noisy parities and halfspaces". In: *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. IEEE. 2006, pp. 563–574. DOI: 10.1109/FOCS.2006.51 (cit. on pp. 16, 21).

[FKO06]   Uriel Feige, Jeong Han Kim, and Eran Ofek. "Witnesses for non-satisfiability of dense random 3CNF formulas". In: *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 2006, pp. 497–508. DOI: 10.1109/FOCS.2006.78 (cit. on p. 3).

[Gen09]   Craig Gentry. "Fully homomorphic encryption using ideal lattices". In: *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*. STOC '09. Bethesda, MD, USA: Association for Computing Machinery, 2009, pp. 169–178. ISBN: 9781605585062. DOI: 10.1145/1536414. (cit. on p. 17).

[Gen10]   Craig Gentry. "Toward Basing Fully Homomorphic Encryption on Worst-Case Hardness". In: *Advances in Cryptology – CRYPTO 2010*. Ed. by Tal Rabin. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 116–137. ISBN: 978-3-642-14623-7. DOI: 10.1007/978-3-642-14623-77 (cit. on p. 17).

[GHK+23]   V. Guruswami, J. Hsieh, P. K. Kothari, and P. Manohar. "Efficient Algorithms for Semirandom Planted CSPs at the Refutation Threshold". In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. Los Alamitos, CA, USA: IEEE Computer Society, Nov. 2023, pp. 307–327. DOI: 10.1109/FOCS57990.2023.00026 (cit. on p. 3).

[GKK+17]   Surbhi Goel, Varun Kanade, Adam Klivans, and Justin Thaler. "Reliably Learning the ReLU in Polynomial Time". In: *Proceedings of the 2017 Conference on Learning Theory*. Ed. by Satyen Kale and Ohad Shamir. Vol. 65. Proceedings of Machine Learning Research. PMLR, July 2017, pp. 1004–1042. URL: https://proceedings.mlr.press/v65/goel17a.html (cit. on pp. 16, 21).

[GKK19]   Surbhi Goel, Sushrut Karmalkar, and Adam Klivans. "Time/accuracy tradeoffs for learning a relu with respect to gaussian marginals". In: *Advances in neural information processing systems* 32 (2019). DOI: 10.5555/3454287.3455057 (cit. on pp. 16, 21).

[GKL+21]   Sumegha Garg, Pravesh K. Kothari, Pengda Liu, and Ran Raz. "Memory-sample lower bounds for learning parity with noise". English (US). In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2021*. Ed. by Mary Wootters and Laura Sanita. Leibniz International Proceedings in Informatics, LIPIcs. Germany: Schloss Dagstuhl- Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, Sept. 2021. DOI: 10.4230/LIPIcs.APPROX/RANDOM.2021.60 (cit. on p. 22).

[GKM22]   Venkatesan Guruswami, Pravesh K. Kothari, and Peter Manohar. "Algorithms and certificates for Boolean CSP refutation: smoothed is no harder than random". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2022. Rome, Italy: Association for Computing Machinery, 2022, pp. 678–689. ISBN: 9781450392648. DOI: 10.1145/3519935.35199 (cit. on p. 3).

[GKP+10]   Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. "Robustness of the Learning with Errors Assumption". In: *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*. Ed. by Andrew Chi-Chih Yao. Tsinghua University Press, 2010, pp. 230–240. URL: http://conference.iiis.tsinghua.edu.c (cit. on p. 21).

[Gri01]   Dima Grigoriev. "Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity". In: *Theoretical Computer Science* 259.1-2 (2001), pp. 613–622. DOI: 10.1016/S0304-3975(00)00 (cit. on p. 5).

[GRS08]   Henri Gilbert, Matthew J. B. Robshaw, and Yannick Seurin. "How to Encrypt with the LPN Problem". In: *Automata, Languages and Programming*. Ed. by Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 679–690. ISBN: 978-3-540-70583-3. DOI: 10.1007/978-3-540-70583-355 (cit. on p. 16).

[GRT18] Sumegha Garg, Ran Raz, and Avishay Tal. "Extractor-based time-space lower bounds for learning". In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2018. Los Angeles, CA, USA: Association for Computing Machinery, 2018, pp. 990–1002. ISBN: 9781450355599. DOI: 10.1145/3188745.3188962 (cit. on p. 22).

[GVV22] Aparna Gupte, Neekon Vafa, and Vinod Vaikuntanathan. "Continuous LWE is as Hard as LWE and Applications to Learning Gaussian Mixtures". In: *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. 2022, pp. 1162–1173. DOI: 10.1109/FOCS54457.2022.001 (cit. on pp. 18, 21, 22).

[HKP+17] Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. "The power of sum-of-squares for detecting hidden structures". In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2017, pp. 720–731 (cit. on p. 63).

[Hop18] Samuel B Hopkins. "Statistical Inference and the Sum of Squares Method". PhD thesis. Cornell University, 2018. URL: https://www.samuelbhopkins.com/thesis.pdf (cit. on pp. 5, 63).

[HS17] Samuel B Hopkins and David Steurer. "Efficient Bayesian estimation from few samples: community detection and related problems". In: *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*. IEEE. 2017, pp. 379–390. DOI: 10.1109/FOCS.2017.42 (cit. on pp. 5, 63).

[Huf03] W. C. (William Cary) Huffman. *Fundamentals of error-correcting codes*. 1st ed. Cambridge, U.K. ; New York : Cambridge University Press, June 2003. ISBN: 0-521-13170-7 (cit. on p. 8).

[IKO+08] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. "Cryptography with constant computational overhead". In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*. STOC '08. Victoria, British Columbia, Canada: Association for Computing Machinery, 2008, pp. 433–442. ISBN: 9781605580470. DOI: 10.1145/1374376.1374438 (cit. on pp. 3, 16).

[IPS09] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. "Secure arithmetic computation with no honest majority". In: *Theory of Cryptography: 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings 6*. Springer. 2009, pp. 294–314. DOI: 10.1007/978-3-642-00457-518 (cit. on p. 2).

[JKL+24] Zhengzhong Jin, Yael Kalai, Alex Lombardi, and Vinod Vaikuntanathan. "SNARGs under LWE via Propositional Proofs". In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. STOC 2024. Vancouver, BC, Canada: Association for Computing Machinery, 2024, pp. 1750–1757. DOI: 10.1145/3618260.3649770 (cit. on p. 17).

[JKP+12] Abhishek Jain, Stephan Krenn, Krzysztof Pietrzak, and Aris Tentes. "Commitments and Efficient Zero-Knowledge Proofs from Learning Parity with Noise". In: *Advances in Cryptology – ASIACRYPT 2012*. Ed. by Xiaoyun Wang and Kazue Sako. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 663–680. ISBN: 978-3-642-34961-4 (cit. on p. 16).

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. "Indistinguishability obfuscation from well-founded assumptions". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 60–73. DOI: 10.1145/3406325.3451093 (cit. on pp. 2, 21).

[JLS22]   Aayush Jain, Huijia Lin, and Amit Sahai. "Indistinguishability obfuscation from LPN over F p, DLIN, and PRGs in NC 0". In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2022, pp. 670–699 (cit. on pp. 2, 21).

[JLS24]   Aayush Jain, Huijia Lin, and Sagnik Saha. "A Systematic Study of Sparse LWE". In: *Advances in Cryptology – CRYPTO 2024*. Springer-Verlag, 2024 (cit. on pp. 3, 4, 6–9, 14, 19, 42, 43).

[Kea98]   Michael Kearns. "Efficient noise-tolerant learning from statistical queries". In: *J. ACM* 45.6 (Nov. 1998), pp. 983–1006. ISSN: 0004-5411 (cit. on p. 21).

[KK14]   Adam Klivans and Pravesh Kothari. "Embedding hard learning problems into gaussian space". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2014 (cit. on pp. 16, 21).

[KKM+08]   Adam Tauman Kalai, Adam R Klivans, Yishay Mansour, and Rocco A Servedio. "Agnostically learning halfspaces". In: *SIAM Journal on Computing* 37.6 (2008), pp. 1777–1805 (cit. on p. 16).

[KMO+17]   Pravesh K. Kothari, Ryuhei Mori, Ryan O'Donnell, and David Witmer. "Sum of squares lower bounds for refuting any CSP". In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2017. Montreal, Canada: Association for Computing Machinery, 2017, pp. 132–145. ISBN: 9781450345286. DOI: 10.1145/3055399.3055485. URL: https://doi.org/10.1145/3055399.3055485 (cit. on pp. 3, 5, 25).

[KMP14]   Eike Kiltz, Daniel Masny, and Krzysztof Pietrzak. "Simple Chosen-Ciphertext Security from Low-Noise LPN". In: *Public-Key Cryptography – PKC 2014*. Ed. by Hugo Krawczyk. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 1–18. ISBN: 978-3-642-54631-0. DOI: 10.1007/978-3-64 (cit. on p. 16).

[KMR+94]   Michael Kearns, Yishay Mansour, Dana Ron, Ronitt Rubinfeld, Robert E. Schapire, and Linda Sellie. "On the learnability of discrete distributions". In: *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*. STOC '94. Montreal, Quebec, Canada: Association for Computing Machinery, 1994, pp. 273–282. ISBN: 0897916638. DOI: 10.1145/195058.195155. URL: https://doi.org/10.1145/195058.195155 (cit. on p. 16).

[KS04]   Adam R. Klivans and Rocco A. Servedio. "Learning DNF in time 2õ(n1/3)". In: *J. Comput. Syst. Sci.* 68.2 (Mar. 2004), pp. 303–318. ISSN: 0022-0000 (cit. on p. 28).

[KS06]   Adam R. Klivans and Alexander A. Sherstov. "Cryptographic Hardness for Learning Intersections of Halfspaces". In: *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 2006, pp. 553–562. DOI: 10.1109/FOCS.2006.24 (cit. on p. 18).

[KSS13]   Nadia Kreimer, Aaron Stanton, and Mauricio D. Sacchi. "Tensor completion based on nuclear norm minimization for 5D seismic data reconstruction". In: *GEOPHYSICS* 78.6 (2013), pp. V273–V284 (cit. on p. 19).

[LL10]   Nan Li and Baoxin Li. "Tensor completion for on-board compression of hyperspectral images". In: *2010 IEEE International Conference on Image Processing*. 2010, pp. 517–520. DOI: 10.1109/ICIP.2010.5651225 (cit. on p. 19).

[LLL82]   A. K. Lenstra, H. W. Lenstra, and L. Lovász. "Factoring polynomials with rational coefficients". eng. In: *Mathematische annalen* 261.4 (1982), pp. 515–534. ISSN: 0025-5831 (cit. on pp. 18, 64).

[LMW+09]   Ji Liu, Przemyslaw Musialski, Peter Wonka, and Jieping Ye. "Tensor completion for estimating missing values in visual data". In: *2009 IEEE 12th International Conference on Computer Vision*. 2009, pp. 2114–2121. DOI: 10.1109/ICCV.2009.5459463 (cit. on p. 19).

[LPW06]   David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2006. URL: http://scholar.google.com/scholar.bib?q=info:3 (cit. on p. 11).

[Lyu05]   Vadim Lyubashevsky. "The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem". In: *Proceedings of the 8th International Workshop on Approximation, Randomization and Combinatorial Optimization Problems, and Proceedings of the 9th International Conference on Randamization and Computation: Algorithms and Techniques*. APPROX'05/RANDOM'05. Berkeley, CA: Springer-Verlag, 2005, pp. 378–389. ISBN: 3540282394. DOI: 10.1007/1153846232. URL: https://doi.org/10.1007/1153846232 (cit. on p. 16).

[LYY+17]   Hanlin Liu, Di Yan, Yu Yu, and Shuoyao Zhao. "On the Hardness of Sparsely Learning Parity with Noise". In: *Provable Security*. Ed. by Tatsuaki Okamoto, Yong Yu, Man Ho Au, and Yannan Li. Cham: Springer International Publishing, 2017, pp. 261–267. ISBN: 978-3-319-68637-0 (cit. on p. 17).

[Mic19]   Daniele Micciancio. "On the hardness of learning with errors with binary secrets". In: *Theory of Computing* 14 (Jan. 2019). DOI: 10.4086/toc.2018.v014a013 (cit. on p. 21).

[MM17]   Dana Moshkovitz and Michal Moshkovitz. "Mixing Implies Lower Bounds for Space Bounded Learning". In: *Proceedings of the 2017 Conference on Learning Theory*. Ed. by Satyen Kale and Ohad Shamir. Vol. 65. Proceedings of Machine Learning Research. PMLR, July 2017, pp. 1516–1566. URL: https://proceedings.mlr.press/v65/moshkovitz17a.html (cit. on p. 22).

[MP13]   Daniele Micciancio and Chris Peikert. "Hardness of SIS and LWE with Small Parameters". In: *Advances in Cryptology – CRYPTO 2013*. Ed. by Ran Canetti and Juan A. Garay. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 21–39. ISBN: 978-3-642-40041-4 (cit. on p. 21).

[MPV24]   Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. "Adaptively Sound Zero-Knowledge SNARKs for UP". In: *Advances in Cryptology – CRYPTO 2024*. Ed. by Leonid Reyzin and Douglas Stebila. Cham: Springer Nature Switzerland, 2024, pp. 38–71. ISBN: 978-3-031-68403-6. DOI: 10.1007/978-3-031-68403-62 (cit. on p. 17).

[MR06]   Elchanan Mossel and Sébastien Roch. "Learning nonsingular phylogenies and hidden Markov models". In: *The Annals of Applied Probability* 16.2 (2006), pp. 583–614. DOI: 10.1214/105051606000000 (cit. on p. 16).

[MS18]   Andrea Montanari and Nike Sun. "Spectral Algorithms for Tensor Completion". In: *Communications on Pure and Applied Mathematics* 71.11 (2018), pp. 2381–2425. DOI: 10.1002/cpa.21748 (cit. on pp. 5, 17, 20).

[MW22]   Andrea Montanari and Alexander S. Wein. *Equivalence of Approximate Message Passing and Low-Degree Polynomials in Rank-One Matrix Estimation*. 2022. arXiv: 2212.06996 [math.ST] (cit. on p. 63).

[MW23]   Ankur Moitra and Alexander S. Wein. *Precise Error Rates for Computationally Efficient Testing*. 2023. arXiv: 2311.00289 [math.ST] (cit. on p. 64).

[ODo14] Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge : Cambridge University Press; 2014. DOI: [10.1017/CBO9781139814782](https://doi.org/10.1017/CBO9781139814782) (cit. on p. 63).

[Pei09] Chris Peikert. "Public-key cryptosystems from the worst-case shortest vector problem: extended abstract". In: *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*. STOC '09. Bethesda, MD, USA: Association for Computing Machinery, 2009, pp. 333–342. ISBN: 9781605585062. DOI: [10.1145/1536414.1536461](https://doi.org/10.1145/1536414.1536461) (cit. on pp. 5, 18, 20).

[Pei16] Chris Peikert. *A Decade of Lattice Cryptography*. 2016. DOI: [10.1561/0400000074](https://doi.org/10.1561/0400000074) (cit. on pp. 17, 18).

[Pie12] Krzysztof Pietrzak. "Cryptography from Learning Parity with Noise". In: *SOFSEM 2012: Theory and Practice of Computer Science*. Ed. by Mária Bieliková, Gerhard Friedrich, Georg Gottlob, Stefan Katzenbeisser, and György Turán. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 99–114. DOI: [10.1007/978-3-642-27660-69](https://doi.org/10.1007/978-3-642-27660-69) (cit. on p. 16).

[PW10] Mihai Patrascu and Ryan Williams. "On the possibility of faster SAT algorithms". In: *Proceedings of the 2010 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2010, pp. 1065–1075. DOI: [10.5555/1873601.1873687](https://doi.org/10.5555/1873601.1873687) (cit. on p. 9).

[Raz17] Ran Raz. "A Time-Space Lower Bound for a Large Class of Learning Problems". In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. 2017, pp. 732–742. DOI: [10.1109/FOCS.2017.73](https://doi.org/10.1109/FOCS.2017.73) (cit. on p. 22).

[Reg09] Oded Regev. "On lattices, learning with errors, random linear codes, and cryptography". In: *J. ACM* 56.6 (Sept. 2009). ISSN: 0004-5411. DOI: [10.1145/1568318.1568324](https://doi.org/10.1145/1568318.1568324) (cit. on pp. 2, 5, 6, 17, 18, 20, 51).

[RRS17] Prasad Raghavendra, Satish Rao, and Tselil Schramm. "Strongly refuting random CSPs below the spectral threshold". In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2017. Montreal, Canada: Association for Computing Machinery, 2017, pp. 121–131. ISBN: 9781450345286. DOI: [10.1145/3055399.3055417](https://doi.org/10.1145/3055399.3055417) (cit. on pp. 3, 5, 15, 17).

[RVV24] Seyoon Ragavan, Neekon Vafa, and Vinod Vaikuntanathan. *Indistinguishability Obfuscation from Bilinear Maps and LPN Variants*. Cryptology ePrint Archive, Paper 2024/856. 2024. URL: https://eprint.iacr.org/2024/856 (cit. on pp. 3, 21, 28, 29).

[Sch08] Grant Schoenebeck. "Linear level Lasserre lower bounds for certain k-CSPs". In: *2008 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE. 2008, pp. 593–602. DOI: [10.1109/FOCS.2008.74](https://doi.org/10.1109/FOCS.2008.74) (cit. on pp. 5, 25).

[Sch87] C.P. Schnorr. "A hierarchy of polynomial time lattice basis reduction algorithms". In: *Theoretical Computer Science* 53.2 (1987), pp. 201–224. ISSN: 0304-3975. DOI: https://doi.org/10.1016/030 (cit. on p. 18).

[STA20] Chao Sun, Mehdi Tibouchi, and Masayuki Abe. "Revisiting the Hardness of Binary Error LWE". In: *Information Security and Privacy*. Ed. by Joseph K. Liu and Hui Cui. Cham: Springer International Publishing, 2020, pp. 425–444. ISBN: 978-3-030-55304-3. DOI: [10.1007/978-3-030-55304](https://doi.org/10.1007/978-3-030-55304) (cit. on p. 21).

[SV10]   N. P. Smart and F. Vercauteren. "Fully Homomorphic Encryption with Relatively Small Key and Ciphertext Sizes". In: *Public Key Cryptography – PKC 2010*. Ed. by Phong Q. Nguyen and David Pointcheval. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 420–443. ISBN: 978-3-642-13013-7. DOI: 10.1007/978-3-642-13013-725 (cit. on p. 17).

[SVD+11]   Marco Signoretto, Raf Van de Plas, Bart De Moor, and Johan A. K. Suykens. "Tensor Versus Matrix Completion: A Comparison With Application to Spectral Data". In: *IEEE Signal Processing Letters* 18.7 (2011), pp. 403–406. DOI: 10.1109/LSP.2011.2151856 (cit. on p. 19).

[SZB24]   Min Jae Song, Ilias Zadik, and Joan Bruna. "On the cryptographic hardness of learning single periodic neurons". In: *Proceedings of the 35th International Conference on Neural Information Processing Systems*. NIPS '21. Red Hook, NY, USA: Curran Associates Inc., 2024. ISBN: 9781713845393. DOI: 10.5555/3540261.3542527 (cit. on pp. 18, 22).

[Tie23]   Stefan Tiegel. "Hardness of Agnostically Learning Halfspaces from Worst-Case Lattice Problems". In: *Proceedings of Thirty Sixth Conference on Learning Theory*. Ed. by Gergely Neu and Lorenzo Rosasco. Vol. 195. Proceedings of Machine Learning Research. PMLR, July 2023, pp. 3029–3064. URL: https://proceedings.mlr.press/v195/tiegel23a/tiegel23a.p (cit. on pp. 18, 22).

[Tie24]   Stefan Tiegel. "Improved Hardness Results for Learning Intersections of Halfspaces". In: *Proceedings of Thirty Seventh Conference on Learning Theory*. Ed. by Shipra Agrawal and Aaron Roth. Vol. 247. Proceedings of Machine Learning Research. PMLR, July 2024, pp. 4764–4786. URL: https://proceedings.mlr.press/v247/tiegel24a/tiegel24a.pdf (cit. on pp. 18, 22).

[Val15]   Gregory Valiant. "Finding Correlations in Subquadratic Time, with Applications to Learning Parities and the Closest Pair Problem". In: *J. ACM* 62.2 (May 2015). ISSN: 0004-5411. DOI: 10.1145/2728167 (cit. on p. 21).

[Val84]   Leslie G Valiant. "A theory of the learnable". In: *Communications of the ACM* 27.11 (1984), pp. 1134–1142 (cit. on p. 28).

[Wag02]   David Wagner. "A Generalized Birthday Problem". In: *Advances in Cryptology — CRYPTO 2002*. Ed. by Moti Yung. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 288–304. ISBN: 978-3-540-45708-4. DOI: 10.1007/3-540-45708-919 (cit. on p. 18).

[YYL+21]   Di Yan, Yu Yu, Hanlin Liu, Shuoyao Zhao, and Jiang Zhang. "An improved algorithm for learning sparse parities in the presence of noise". In: *Theoretical Computer Science* 873 (2021), pp. 76–86. ISSN: 0304-3975. DOI: https://doi.org/10.1016/j.tcs.2021.04.026 (cit. on p. 21).

[YZ16]   Yu Yu and Jiang Zhang. "Cryptography with Auxiliary Input and Trapdoor from Constant-Noise LPN". In: *Advances in Cryptology – CRYPTO 2016*. Ed. by Matthew Robshaw and Jonathan Katz. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 214–243. ISBN: 978-3-662-53018-4. DOI: 10.1007/978-3-662-53018-49 (cit. on p. 16).

[ZSW+22]   Ilias Zadik, Min Jae Song, Alexander S Wein, and Joan Bruna. "Lattice-Based Methods Surpass Sum-of-Squares in Clustering". In: *Proceedings of Thirty Fifth Conference on Learning Theory*. Ed. by Po-Ling Loh and Maxim Raginsky. Vol. 178. Proceedings of Machine Learning Research. PMLR, July 2022, pp. 1247–1248. URL: https://proceedings.mlr.press/v178/zadik2 (cit. on pp. 25, 64).

A Low-Degree Hardness of Sparse Linear Equations

A.1 Preliminaries on Low-Degree Polynomial Hardness

Low-degree polynomial hardness, emerging from high-dimensional statistics [HS17; HKP+17; BHK+19; Hop18], is an increasingly popular framework for providing evidence for hardness of statistical problems.

The central observation is that functions that depend on very few bits are cheap to evaluate. Hence, so are linear combinations of such functions. We formalize this notion by *coordinate degree*.

Definition 5 (e.g. [ODo14]). We say that a function $f : \mathcal{X}^n \longrightarrow \mathbb{R}$ has coordinate degree at most d if it can be represented as:

$$f(x_1, x_2, \ldots, x_n) = \sum_{S \subseteq [n] \ : \ |S| \leq d} f_S((x_i)_{i \in S})$$

Assuming $\exp(\tilde{O}(d))$ evaluation time of a function in at most d variables, a coordinate-degree d function can be evaluated in time $\exp(\tilde{O}(d))$.

The abelian group structure of $\mathbb{Z}_q^n$ gives a nice representation of coordinate-degree-d functions over $\mathbb{Z}_q^n$.

Lemma A.1 (e.g. [ODo14]). *Every coordinate-degree-d function $f : \mathbb{Z}_q^n \longrightarrow \mathbb{R}$ has a unique representation of the form*

$$f(x) = \sum_{v \in \mathbb{Z}_q^n \ : \ |\mathsf{support}(v)| \leq d} c_v \exp \Big(\frac{2\pi i \langle x, v \rangle}{q} \Big) \qquad \forall x \in \mathbb{Z}_q^n.$$

One can use a low-coordinate-degree function to test between two distributions $\mathcal{P}_0$ and $\mathcal{P}_1$ if the means are sufficiently separated. Immediately implied by Chebyshev's inequality is the fact that if $|\mathbb{E}_{X \sim \mathcal{P}_0} f(X) - \mathbb{E}_{X \sim \mathcal{P}_1} f(X)| = \omega(\sqrt{\mathsf{Var}_{X \sim \mathcal{P}_0} f(X) + \mathsf{Var}_{X \sim \mathcal{P}_1} f(X)})$, one can test between $H_0 : \mathcal{P}_0$ and $H_1 : \mathcal{P}_1$ with type I + type II error $o(1)$ by simply evaluating $f(X)$ on input X and comparing to $\frac{1}{2}(\mathbb{E}_{X \sim \mathcal{P}_0} f(X) + \mathbb{E}_{X \sim \mathcal{P}_1} f(X))$. When this condition does not hold for *any* low-degree coordinate function, we say that the distinguish problem is low-degree hard. Formally, we also allow conditioning on a high-probability "well-behaved" event.

Definition 6 (e.g. [DMW23]). Suppose that $\mathcal{P}_0$ and $\mathcal{P}_1$ are two distributions over $\mathcal{X}^n$. Suppose that there exists some event $\mathcal{E}$ such that $\mathbb{P}_{X \sim \mathcal{P}_1}[X \in \mathcal{E}] = 1 - o(1)$ holds and for any coordinate-degree-d function $f : \mathcal{X}^n \longrightarrow \mathbb{R}$,

$$| \mathop{\mathbb{E}}_{X \sim \mathcal{P}_0} [f(X)] - \mathop{\mathbb{E}}_{X \sim \mathcal{P}_1} [f(X)|X \in \mathcal{E}]| = o\Big(\sqrt{\mathsf{Var}_{X \sim \mathcal{P}_0} f(X) + \mathsf{Var}_{X \sim \mathcal{P}_1}[f(X)|X \in \mathcal{E}]} \Big)$$

holds. Then, we say that the hypothesis testing problem $H_0 : \mathcal{P}_0$ vs $H_1 : \mathcal{P}_1$ is hard against coordinate-degree-d polynomial tests.

Typically, hardness against degree $\omega(\log n)$ tests is viewed as evidence for hardness against polynomial-time algorithms and hardness against degree n^C tests for constant C is viewed as evidence of subexponential hardness.

While low-degree polynomial tests certainly do not capture *all* computationally efficient distinguising algorithms, they are an expressive and rich class. They capture certain common instantiations of spectral algorithms [BKW19], SQ algorithms [BBH+21], and approximate message passing algorithms [MW22] among others. Low-degree polynomial tests are also conjectured to capture all polynomial-time algorithms for a well-defined class of sufficiently noisy problems in high dimension [Hop18].

We do note that there are many problems for which there are better than low-degree polynomial algorithms and, in fact, some of them come exactly from problems on lattices and noisy linear equations such as the celebrated Blum-Kalai-Wasserman algorithm [BKW03] or more recent work on "flat" parallel pancakes [ZSW+22; DK22a] based on the celebrated Lenstra-Lenstra-Lovasz algorithm [LLL82].

Remark 4 (Distinguishing Advantage)**.** In this overview we focused on advantage $o(1)$ for the distinguishing problem, which is the most common set-up in the literature on high-dimensional statistics. Recent work [MW23] has also considered a more fine-grained analysis of the advantage. We do not engage in this endeavour since the low-degree hardness results mostly serve as evidence in the current work.

A.2 Proof of Proposition 4.3

With this in hand, we can prove Proposition 4.3 restated below.

Proposition (Restatement of Proposition 4.3)**.** *The detection problem with any error distribution $\mathcal{E}$ and distribution on the support $\mathcal{D}$ for LWE (modulus $= q$, samples $\sim$ randomsupport$(n, k, \mathcal{D})$, secret $\sim$ Unif$((\mathbb{Z}/q\mathbb{Z})^n)$, error $\sim$ $\mathcal{E}$) is hard for degree-D polynomial tests when $m = o\left(\frac{n}{k} \times (ne^3/kD)^{k/2-1}\right)$.*

We begin with a simple lemma which we will use to choose a well-behaved high-probability event in Definition 6.

Lemma A.2. *Suppose that $m = o\left(\frac{n^{k/2}}{k^{k/2}D^{k/2-1}e^{3k}}\right)$ and $a_1, a_2, \dots, a_m \overset{\text{i.i.d.}}{\sim}$ randomsupport$(n, k, \mathcal{D})$. Then, with high probability, every D vectors among $a_1, a_2, \dots, a_m$ are linearly independent.*

Proof. Suppose that there is a set of at most D vectors a_i which are not linearly independent. Let $S = \{i_1, i_2, \dots, i_s\}$ be the smallest such set. Then, we claim that each $j \in \bigcup_{i \in S} \text{support}(a_i)$ appears at least twice in the multiset $\bigcup_{i \in S} \text{support}(a_i)$. Otherwise, if some j appears only once and it appears for vector a_i, clearly a_i is linear independent of $(a_\ell)_{\ell \in S \setminus \{i\}}$. Hence, $S \setminus \{i\}$ is a smaller set of vectors which are not linearly independent, contradiction.

In particular this means that $|\bigcup_{i \in S} \text{support}(a_i)| \leq k \times |S|/2 = ks/2$. We bound the probability of this event:

$$\mathbb{P}\left[\exists S \subseteq [m] \ : \ |S| \leq D \text{ and } |\bigcup_{i \in S} \text{support}(a_i)| \leq k \times |S|/2\right]$$

$$\leq \sum_{S \subseteq [m] \ : \ |S| \leq D} \mathbb{P}\left[|\bigcup_{i \in S} \text{support}(a_i)| \leq k \times |S|/2\right] \tag{14}$$

$$\leq \sum_{s=2}^{D} \binom{m}{s} \mathbb{P}\left[|\bigcup_{i=1}^{s} \text{support}(a_i)| \leq k \times s/2\right].$$

We finish by bounding the quantity in the last line. Let $Z = \bigcup_{i=1}^{s} \text{support}(a_i)$. If Z is contained in a set Z' of size $ks/2$, there are at most $\binom{n}{ks/2}$ choices for this set. Once the set Z' has been chosen, each of the vectors $a_1, a_2, \dots, a_s$ can be chosen in $\binom{ks/2}{k}$ ways. Hence, we bound the resulting probability by $\frac{\binom{n}{ks/2}\binom{ks/2}{k}^s}{\binom{n}{k}^s}$. Going back to Eq. (14) and using $(\frac{a}{b})^b \leq \binom{a}{b} \leq (\frac{ea}{b})^b$, we bound

$$\sum_{s=2}^{D} \binom{m}{s} \frac{\binom{n}{ks/2}\binom{ks/2}{k}^s}{\binom{n}{k}^s}$$

$$\leq \sum_{s=2}^{D} \left(\frac{me}{s}\right)^s \times \left(\frac{ne}{ks/2}\right)^{ks/2} \times \left(\frac{eks/2}{k}\right)^{ks} \left(\frac{n}{k}\right)^{-sk}$$

$$\leq \sum_{s=2}^{D} \left(\frac{me \times (ne)^{k/2} \times (es/2)^k \times k^k}{s \times (ks/2)^{k/2} \times n^k}\right)^s$$

$$\leq \sum_{s=2}^{D} \left(\frac{mk^{k/2}s^{k/2-1}e^{3k}}{n^{k/2}}\right)^s$$

$$\leq \sum_{s=2}^{D} \left(\frac{mk^{k/2}D^{k/2-1}e^{3k}}{n^{k/2}}\right)^s.$$

When $m = o\left(\frac{n^{k/2}}{k^{k/2}D^{k/2-1}e^{3k}}\right)$, the last expression is of order $o(1)$. $\qquad \square$

Lemma A.3. *Suppose that $a_1, a_2, \ldots, a_m \in \mathbb{Z}_q^n$ are fixed known vectors such that any D of these vectors are linearly independent. Let $\mathcal{E}$ be an arbitrary distribution over $\mathbb{Z}_q$. Then, no function (even depending on $a_1, a_2, \ldots, a_m$) of coordinate degree at most D can distinguish between the following hypotheses:*

$$\begin{aligned}
\mathsf{H}_0 &: (y_1, y_2, \ldots, y_m) \stackrel{\text{i.i.d.}}{\sim} \mathsf{Unif}(\mathbb{Z}_q^m) \\
\mathsf{H}_1 &: (y_1, y_2, \ldots, y_m) = (\langle a_1, s\rangle + e_1, \langle a_2, s\rangle + e_2, \ldots \langle a_m, s\rangle + e_m), \\
&\quad \text{where } s \sim \mathsf{Unif}(\mathbb{Z}_q^n), e_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{E}.
\end{aligned} \tag{15}$$

Proof. By Lemma A.1, every coordinate-degree D function g of $(y_i)_{i\in[m]}$ can be represented as

$$g(y_1, y_2, \ldots, y_m) = \sum_{\substack{S \subseteq [m], \\ |S| \leq D}} \sum_{\substack{v \in \mathbb{Z}_q^m, \\ \text{support}(v) = S}} c_v \exp\left(\frac{2\pi i \langle y, v\rangle}{q}\right).$$

To show that coordinate-degree D functions cannot distinguish the two hypothesis, it is enough to show that for each such g, $\mathbb{E}_{\mathsf{H}_0}[g] = \mathbb{E}_{\mathsf{H}_1}[g]$ by Definition 6. For $S = \emptyset$, both expectations are clearly 1. We will shot that for any $S \neq \emptyset$ and $v \in \mathbb{Z}_q^n$: $\text{support}(v) = S$ it holds that

$$\mathbb{E}_{\mathsf{H}_0}\left[\exp\left(\frac{2\pi i \langle y, v\rangle}{q}\right)\right] = 0 \text{ and } \mathbb{E}_{\mathsf{H}_1}\left[\exp\left(\frac{2\pi i \langle y, v\rangle}{q}\right)\right] = 0.$$

The equality follows trivially for H_0. For H_1, observe that

$$\langle y, v\rangle = \sum_{i\in S} y_i v_i = \sum_{i\in S} (\langle a_i, s\rangle + e_i) v_i = \langle s, \sum_{i\in S} v_i a_i\rangle + \sum_{i\in S} v_i e_i.$$

Since $|S| \leq D$ and no D of the vectors $a_1, a_2, \ldots, a_m$ are linearly dependent, $\sum_{i\in S} v_i a_i \neq 0$. Hence,

$$\mathbb{E}_{s\sim\mathsf{Unif}(\mathbb{Z}_q^n)}\left[\exp\left(\frac{2\pi i \left(\langle s, \sum_{i\in S} v_i a_i\rangle + \sum_{i\in S} v_i e_i\right)}{q}\right)\right] = 0,$$

as desired. $\qquad \square$

Putting It All Together. By Lemma A.2, if $m = o\left(\frac{n^{k/2}}{k^{k/2}D^{k/2-1}e^{3k}}\right)$, every D of the k-sparse samples $a_1, a_2, \ldots, a_m$ are linearly independent with high-probability. This is the high-probability well-behaved event we use to invoke Definition 6. By Lemma A.3, this implies hardness against polynomials g which are degree at most D in the labels $y_1, y_2, \ldots, y_m$ but depend arbitrarily on $a_1, a_2, \ldots, a_m$ (recall that g can depend arbitrarily on $a_1, a_2, \ldots, a_m$ in Lemma A.3). In particular, this captures all polynomials in (A, y) of degree at most D.