

SALSA PICANTE: a machine learning attack on LWE with binary secrets

Cathy Li *
Meta AI

Jana Sotáková *
Meta AI

Emily Wenger
University of Chicago

Mohamed Malhou
Meta AI

Evrard Garcelon
ENSAE - CREST

Francois Charton †
Meta AI

Kristin Lauter †
Meta AI

Abstract

The *Learning With Errors* (LWE) problem is one of the major hard problems in post-quantum cryptography. For example, 1) the only *Key Exchange Mechanism* KEM standardized by NIST [14] is based on LWE; and 2) current publicly available Homomorphic Encryption (HE) libraries are based on LWE. NIST KEM schemes use random secrets, but homomorphic encryption schemes use binary or ternary secrets, for efficiency reasons. In particular, sparse binary secrets have been proposed, but not standardized [2], for HE.

Prior work SALSA [49] demonstrated a new machine learning attack on sparse binary secrets for the LWE problem in small dimensions (up to $n = 128$) and low Hamming weights (up to $h = 4$). However, this attack assumed access to millions of LWE samples, and was not scaled to higher Hamming weights or dimensions.

Our attack, PICANTE, reduces the number of samples required to just $m = 4n$ samples. Moreover, it can recover secrets with much larger dimensions (up to 350) and Hamming weights (roughly $n/10$, or $h = 33$ for $n = 300$). To achieve this, we introduce a *preprocessing step* which allows us to generate the training data from a linear number of samples and changes the distribution of the training data to improve transformer training. We also improve the distinguisher/secret recovery methods of SALSA and introduce a novel cross-attention recovery mechanism which allows us to read-off the secret directly from the trained models.

1 Introduction

The race for post-quantum cryptography (PQC) is in full steam. A large-scale quantum computer could solve the hard math problems underpinning most deployed public-key cryptographic systems, like RSA [40], in polynomial time, and small-scale quantum computers have already been built. Consequently, new post-quantum cryptographic systems were

proposed and considered for standardization by US National Institute of Standards and Technology (NIST) in the 5-year PQC competition.

In July 2022, NIST standardized 4 schemes from the post-quantum crypto competition [14]. The only key encapsulation mechanism selected—CRYSTALS-Kyber [6]—and one of the three signature schemes—CRYSTALS-Dilithium [25]—are based on a mathematical hardness assumption known as Learning With Errors (LWE) [39].

LWE works as follows: given an integer modulus q , a dimension n , and a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$, the *Learning With Errors problem* is to recover the secret vector, given many instances of a random vector along with their noisy inner products with the secret vector. These noisy inner products are computed by taking a random vector $\mathbf{a} \in \mathbb{Z}_q^n$ and producing $b := \mathbf{a} \cdot \mathbf{s} + e \pmod q$, where e is an “error” term sampled from a narrow discrete Gaussian distribution (i.e. taking small values). The adversary is then given the *samples* (\mathbf{a}, b) . LWE-based encryption schemes encrypt messages by *blinding* them via noisy inner products. Small secrets, such as binary or ternary secrets, are used for homomorphic encryption.

The LWE problem is assumed to be hard for both classical and quantum adversaries [9, 31, 33, 38, 39]. Classical attacks on LWE typically rely on algebraic techniques for lattice reduction to recover the secret \mathbf{s} from pairs (\mathbf{a}, b) [15, 30]. However, the error e added to $\mathbf{a} \cdot \mathbf{s}$ to compute noisy inner products makes algebraic solutions difficult. At its core, the LWE hardness assumption is that it is hard to learn from noisy data. Ironically, the success of the whole field of Machine Learning (ML) depends on the fact that it is possible to train machines to learn from noisy data. One additional difficulty in lattice-based cryptography is that all operations are done modulo an integer q . [49] showed that it was possible to train ML algorithms to handle modular arithmetic.

Current advances in AI systems rely on Deep Learning. One of the most common large model architectures are transformers. Transformers were introduced in [47] for natural language processing (NLP) and machine translation. In recent years, they have been applied to a wide range of problems,

*Co-first authors

†Co-senior authors

Date of the document: March 7, 2023.

from text and image generation [35–37] to image processing [10] and speech recognition [23], where they now achieve state-of-the-art performance [24, 48]. Transformers have also been proposed for problems in mathematics, like symbolic integration [29], theorem proving [34], and numerical computations [11]. Transformers process sequences of tokens (in NLP, sequences of words, making up sentences). They combine a multi-head attention mechanism [7] that takes care of relations between different tokens in the sequence, essentially “decorrelating” it, and a fully-connected neural network (FCNN), which processes the decorrelated sequences.

The first successful attempt to use ML to attack hard lattice problems was presented in SALSA [49], showing the feasibility of recovering sparse binary secrets for LWE problems with relatively small parameters. SALSA trains ML models to learn the underlying structure of the LWE problem given many LWE samples, then leverages the trained model to recover the LWE secret. If the models learn to predict b from \mathbf{a} (even with low accuracy), SALSA can recover the secret \mathbf{s} .

Although promising, SALSA has significant limitations. For the largest dimension $n = 128$, SALSA can only recover Hamming weight $h < 3$. In comparison, real-world binary HE schemes start at dimension $n = 512$ or $n = 1024$. SALSA also requires millions of LWE samples (\mathbf{a}, b) for model training, but a real-world attacker would likely only have access to a few samples. Making SALSA’s approach realistic requires scaling up the parameters of solvable LWE problems (n, q, h) while reducing the number of required samples.

Contributions. In this work, we propose PICANTE, an enhanced ML based attack on the LWE hardness assumption. PICANTE leverages basic principles of the original SALSA attack [49]—transformer training, secret recovery—while introducing several novel techniques. This enhanced attack allows recovery of high-dimensional binary secrets with Hamming weight roughly $n/10$, up to 33, requiring only $4n$ samples for training. Table 1 shows the largest Hamming weights we recover for each dimension.

Dimension	80	150	200	256	300	350
$\log q$	7	13	17	23	27	32
highest h	9	13	22	31	33	26

Table 1. SALSA PICANTE’s highest recovered secret Hamming weights h

As in SALSA [49, Table 7], we observe that it is easier to learn from vectors with a skewed distribution on the entries. Therefore, we introduce a data preprocessing step that uses lattice-reduction methods to produce samples with smaller coefficients. Second, we start with a linear number of samples, $m = 4n$. To produce more samples for training, it was suggested in SALSA [49, Appendix E] to take linear combinations to produce more. However, combining samples changes the error distribution, and larger errors are quickly amplified with the lattice reduction step. To avoid the error explosion,

we simply take subsets of n out of the m initial samples and then apply lattice reduction. This generates large quantities of non-duplicate samples with skewed entries, and we show that transformers can learn from such data equally well or better than if we start from independent LWE samples.

Specifically, this work makes the following contributions:

- **Linear number of samples:** our method only requires a linear number of samples, $m = 4n$ in practice.
- **Data preprocessing:** we preprocess data with classical lattice reduction techniques (e.g. LLL [30] or BKZ [15]), using small block size. This helps transformers learn.
- **Novel secret recovery:** we recover secret bits from the trained transformer, using its *cross-attention* mechanism.

We also improve the data encoding method, introduce rounding to reduce the size of the vocabulary the transformer must learn, improve the distinguisher secret recovery method, and introduce novel combined secret recovery methods.

Mastermind. One way to think about the role of our novel preprocessing step is in analogy with the game Mastermind. In Mastermind, a secret made up of 4 pegs of 6 possible colors is hidden from the guesser. The guesser makes queries of 4 pegs of different colors, where responses indicate how many pegs were a correct match of color and position to the secret. Binary secret LWE can be thought of as Mastermind with n positions and 2 colors, ignoring error.

In our approach to binary secret LWE, the trained model serves as an engine for answering queries, assuming the transformer has started to learn. Consider two extreme types of queries. If you submit a vector with all entries constant, (f, f, \dots, f) , (i.e. very low entropy), you only get the Hamming weight—no information about the position of the 1s. At the other extreme, the Direct secret recovery approach of SALSA makes queries of the form $(0, \dots, 0, K_i, 0, \dots, 0)$, which gives information only about the i^{th} bit of the secret. Submitting queries with random entries (maximal entropy) does not clearly give any particular type of information.

In our preprocessing step, we attempt to reduce the entropy of our samples, thus making it more likely that queries such as those in the Direct secret recovery method bear some similarity to the training samples. So in some sense we can think of our approach as ML for Mastermind (or Wordle).

Acknowledgements. We would like to give our thanks to Mark Tygert, Zeyuan Allen-Zhu, Matteo Pirodda, Mingjie Chen, and Antonio Orvieto for their help with the project.

2 Background on LWE

2.1 Lattice-based cryptography

Lattice-based cryptography is a major field in post-quantum cryptography. Three out of the four schemes selected by NIST [14] are lattice-based, and two are based on a variant of

the LWE problem [39], which will be our focus. Moreover, homomorphic encryption libraries are based on LWE.

Lattices. An n -dimensional integer lattice is the set of all integer linear combinations of n linearly independent vectors in \mathbb{Z}^n . More formally, given n vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{Z}^n$, a lattice is the integer span $\Lambda = \Lambda(\mathbf{v}_1, \dots, \mathbf{v}_n) := \{\sum_{i=1}^n a_i \mathbf{v}_i \mid a_i \in \mathbb{Z}\}$. The vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ are called a *basis* for the lattice Λ . The lattice Λ inherits a norm simply by restriction of the Euclidean norm from \mathbb{R}^n to Λ : any vector $v \in \Lambda$ has norm $\|v\| = \sqrt{v \cdot v}$.

Lattices give rise to several *hard* problems —problems for which the best known algorithms require exponential time in the dimension n for both classical and quantum computers. The most famous and widely-studied is the *Shortest Vector Problem* (SVP): for a lattice Λ , find a nonzero vector $\mathbf{v} \in \Lambda$ with minimal norm. Currently, the best algorithms for SVP take exponential space and time in n [32]. This makes lattices attractive building blocks for post-quantum cryptography.

Learning with Errors. Fix a dimension n , modulus q , number of samples m and a narrow Gaussian probability distribution χ . The “Learning with Errors” (LWE) problem is to recover a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$ given a collection of m noisy samples (\mathbf{a}_i, b_i) , where $\mathbf{a}_1, \dots, \mathbf{a}_m \leftarrow_R \mathbb{Z}_q^n$ are random vectors, and $b_i = \mathbf{a}_i \cdot \mathbf{s} + e_i \bmod q$ are noisy inner products, with the $e_i \in \mathbb{Z}_q$ sampled independently from the error distribution χ .

A *LWE instance* is given by a matrix $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ with \mathbf{A} uniformly random in $\mathbb{Z}_q^{m \times n}$ and $\mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \bmod q$ a column vector where $\mathbf{s} \in \mathbb{Z}_q^n$ is the secret vector and $\mathbf{e} \in \mathbb{Z}_q^m$ with entries sampled from the probability distribution χ . We call any of the pairs (\mathbf{a}_i, b_i) , or equivalently any of the rows of the matrix (\mathbf{A}, \mathbf{b}) , an *LWE-sample* or a *sample*.

Hardness of LWE In 2005, Regev demonstrated a worst-case quantum reduction from the SVP to LWE [39]. Regev also showed that LWE-based cryptographic schemes were far more efficient than existing lattice cryptography methods. Later work demonstrated that LWE is classically as hard as worst-case SVP-like problems [9, 31, 33]. Hence, LWE is considered a solid foundation for (post-quantum) lattice cryptography.

LWE-based schemes LWE-based schemes are not only standardized for Post-Quantum Cryptography [6, 25] and Homomorphic Encryption [2], but also allow for a range of cryptographic constructions beyond key exchange and signatures, including group signatures, secret sharing, and multi-party computation. The NIST standardization competition received 23 entries proposing schemes based on lattice assumptions such as LWE. In CRYSTALS-Kyber [6], the dimension is $n = k \times 256$ for $k = 2, 3, 4$. The LWE-based signature scheme CRYSTALS-Dilithium [25] uses similar size of n . Both use secret vectors with random entries $\bmod q$. Another LWE-based NIST submission, LIZARD, suggests secure LWE dimensions n ranging from 544 to 736 [17, Table 2].

Homomorphic encryption schemes [2] in publicly available libraries such as SEAL use dimension $n = 512$ only for small computations, and generally require dimensions $n =$

1024, 2048 and other powers of 2 up to 2^{15} . But HE implementations commonly use binary or ternary secrets for efficiency reasons (see the standard HES [2]), and many implementations propose to use a sparse (binary) secret with Hamming weight $h \ll n$. For instance, HEAAN uses $n = 2^{15}$, $q = 2^{628}$, ternary secret and Hamming weight 64 [13]. For more on the use of sparse binary secrets in LWE, see [3, 16, 20].

We focus on the case of a binary secret with Hamming weight h and the error distribution χ a centered Gaussian with $\sigma = 3$. $\sigma = 3.2$ is the typical choice for homomorphic encryption [2, 16, 20, 46].

2.2 Attacks on LWE

The LWE problem is assumed to be exponentially hard to solve with classical algorithms [9, 31, 33] or with quantum algorithms [38]. Due to its prominence as a hard problem in post-quantum cryptography, a significant body of work has been devoted to attacking it.

Classical Attacks. Most existing classical attacks leverage *lattice reduction* techniques, and reduce the problem to that of recovering the shortest vector in a lattice. The LLL [30] algorithm runs in polynomial time in the dimension of the lattice (specifically the optimized *fpLLL* [22] implementation runs in time $n^4 \log(q)^2$), but it recovers an exponentially bad approximation to the shortest vector. LLL can be improved using the *block Korkine-Zolotarev* method (BKZ) by Schnorr [43] and Schnorr-Euchner [44]. The BKZ approach relies on an exponential time sub-algorithm applied for increasing block sizes, but can recover shorter vectors than LLL. The main 3 attacks used to estimate secure parameters for lattice-based cryptography are: the uSVP, dual, decoding attacks, all of which require finding a short vector in a particular lattice arising from the LWE instance. The uSVP attack uses Kannan’s embedding [26] to embed the problem instance into a lattice such that the (unique) shortest vector reveals the secret s . For concrete choices for this embedding, see [12]. The Homomorphic Encryption Standard [2, Section 2.1.2] describes the uSVP, dual, and decoding attacks in detail.

SALSA: a machine learning based attack. SALSA [49] demonstrated the possibility of using machine learning models to attack the LWE hardness assumption when sparse binary secrets are used. SALSA trains universal transformers to predict b from input \mathbf{a} and develops secret recovery techniques to extract the secret that is implicitly learned by the model. This work successfully recovered secrets for LWE problems with dimension $n \leq 128$ and Hamming weight ≤ 4 .

SALSA pioneered the use of ML models in cryptanalysis of LWE, but it only recovers secrets for small LWE problems. Much work remains to be done to scale up this work to attack real-world LWE problems with cryptographic sizes of parameters. In this work, we extend the ideas of [49] and demonstrate that ML models can be used to successfully recover secrets from LWE problems with near real-world parameter settings.

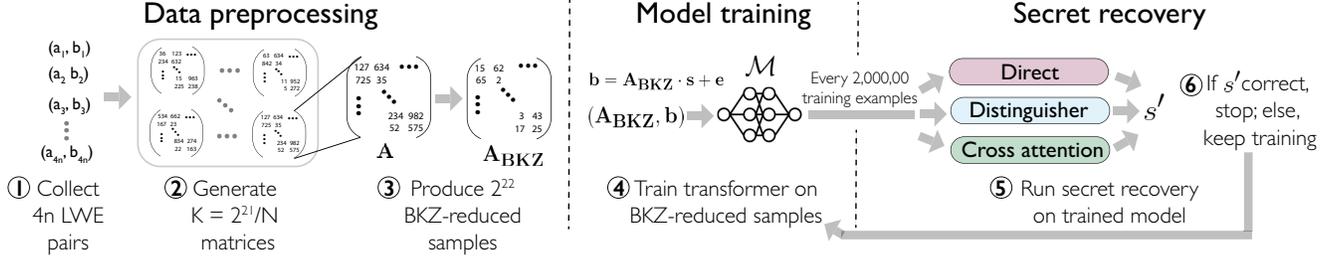


Figure 1. An end-to-end overview of SALSA PICANTE’s attack methodology.

3 Introducing SALSA PICANTE

High level overview. SALSA PICANTE builds upon SALSA and progresses in three stages: (1) data preprocessing, (2) model training, and (3) secret recovery (see Figure 1). Each run of SALSA PICANTE targets an LWE instance with fixed dimension n , modulus q , and binary secret s with Hamming weight h and error distribution χ with $\sigma = 3$.

SALSA PICANTE requires m original LWE pairs, sharing the same secret s . These are of form (\mathbf{a}, b) , with $b = \mathbf{a} \cdot \mathbf{s} + e$. In real world situations, these samples must be collected. In experimental settings we choose $m = 4n$ and generate these samples randomly.

(1) Data preprocessing. During this step, n LWE pairs are randomly selected from the set of original samples and stacked into an $n \times n$ matrix \mathbf{A} and vector \mathbf{b} of length n . The matrix \mathbf{A} is processed by a basis-reduction algorithm (BKZ), and the same linear operation is performed on \mathbf{b} . This creates reduced LWE pairs with smaller norms but larger errors. This step is repeated to produce 2^{22} reduced LWE pairs.

(2) Model Training. The reduced LWE samples (\mathbf{a}, b) are encoded as sequences of numbers, represented in base B , and used to train a deep learning model \mathcal{M} , a transformer. The model \mathcal{M} learns to predict b from \mathbf{a} . Model training proceeds in *epochs*, each using 2 million samples. The 4 million training data are shuffled randomly every 2 epochs.

(3) Secret Recovery. At the end of each epoch, SALSA PICANTE attempts to recover the secret using 3 techniques: direct, distinguisher, and cross-attention. The methods are used separately and can be combined to provide more secret guesses. Secret guesses are evaluated. Model training stops if the secret is recovered; else, another epoch begins.

4 Methods

4.1 LWE data reduction

Whereas SALSA could attack problems of dimension as high as 128, it was limited to binary secrets with low Hamming weights: up to 4 1-bits in the secret. SALSA authors observed [49, Table 7] that, for $n = 50$, binary secrets with Hamming weights up to 15 could be fully or partially recovered, if the coordinates of the samples \mathbf{a} used to train the

model were bounded by αq , with $\alpha < 0.6$.

Our initial experiments confirmed this for larger dimensions, and different restrictions on the coordinates of \mathbf{a} (e.g. $a_i > \alpha q$). Unfortunately, this property cannot be directly exploited to attack LWE. In all practical implementations, the coordinates of \mathbf{a} are sampled from a uniform distribution over \mathbb{Z}_q . Therefore, vectors \mathbf{a} with no coordinates exceeding αq become exponentially rare as the dimension increases.

By taking linear combinations of samples, we can transform LWE pairs $(\mathbf{a}, b = \mathbf{a} \cdot \mathbf{s} + e \bmod q)$ with \mathbf{a} sampled uniformly from \mathbb{Z}_q^n into LWE pairs $(\mathbf{a}', b' = \mathbf{a}' \cdot \mathbf{s} + e' \bmod q)$ using the same secret, but such that the coordinates of \mathbf{a}' take smaller values. Note that taking linear combinations of samples changes the error distribution.

Specifically, given n LWE samples, stored as the rows of an $n \times n$ matrix \mathbf{A} , and a corresponding vector \mathbf{b} of noisy inner products with a fixed secret \mathbf{s} , we can create a matrix \mathbf{A}' with smaller entries than \mathbf{A} , by applying standard basis-reduction algorithms, such as LLL [30] and BKZ [43], to Λ , the n -dimensional lattice defined by the rows of \mathbf{A} . In PICANTE, we run BKZ (from the *fplll* package [22]) on the matrix:

$$\begin{bmatrix} \omega \cdot \mathbf{1}_n & \mathbf{A}_{n \times n} \\ \mathbf{0} & q \cdot \mathbf{1}_n \end{bmatrix},$$

with $\omega \in \mathbb{Z}$ an error penalization parameter, discussed below. Since the BKZ reduction is a change of basis, it is a linear transformation, which we can represent as $[\mathbf{R}_{2n \times n} \quad \mathbf{C}_{2n \times n}]$. The BKZ reduction can be written as a matrix multiplication

$$[\mathbf{R}_{2n \times n} \quad \mathbf{C}_{2n \times n}] \begin{bmatrix} \omega \cdot \mathbf{1}_n & \mathbf{A}_{n \times n} \\ \mathbf{0} & q \cdot \mathbf{1}_n \end{bmatrix} = [\omega \cdot \mathbf{R} \quad \mathbf{R}\mathbf{A} + q\mathbf{C}],$$

with matrices \mathbf{R} and \mathbf{C} are chosen so that $[\omega \cdot \mathbf{R} \quad \mathbf{R}\mathbf{A} + q\mathbf{C}]$ has $2n$ rows with small norms, and the matrix $q\mathbf{C}$ adds an integer multiple of q to each entry in $\mathbf{R}\mathbf{A}$, so that all entries are in the range $(-q/2, q/2)$.

Applying the linear transformation \mathbf{R} to \mathbf{b} , we create a new LWE instance $(\mathbf{R}\mathbf{A}, \mathbf{R}\mathbf{b})$ with the same secret \mathbf{s} , smaller coordinates of $\mathbf{R}\mathbf{A}$ but a *different error distribution*. Let $\mathbf{e} = \mathbf{b} - \mathbf{A} \cdot \mathbf{s}$ be the initial LWE error. After reduction, the error becomes $\mathbf{e}' = \mathbf{R}\mathbf{b} - (\mathbf{R}\mathbf{A}) \cdot \mathbf{s} = \mathbf{R}(\mathbf{b} - \mathbf{A} \cdot \mathbf{s}) = \mathbf{R}\mathbf{e}$: the larger the entries of \mathbf{R} , the more LWE errors are amplified by reduction. All these computations are performed mod q .

Error amplification can be controlled by the error penalization parameter ω . Recall that BKZ computes \mathbf{R} and \mathbf{C} so that the norms of the rows of $[\omega \cdot \mathbf{R} \quad \mathbf{R}\mathbf{A} + q\mathbf{C}]$ are small. A large ω encourages small entries in the rows of \mathbf{R} . On the other hand, it hinders the norm reduction of $\mathbf{R}\mathbf{A} + q\mathbf{C}$, and therefore limits the extent to which the coordinates of \mathbf{A} are reduced. The choice of ω controls a trade-off between the amount of reduction of \mathbf{a} we can achieve, and the amount of additional noise which gets injected in the transformed samples. In practice, we set $\omega = 15$.

The previous paragraphs describe the reduction of n LWE pairs. Larger number of samples (e.g. the million of pairs needed for SALSA) would be divided into batches of n , and processed as above. In this setting, the n LWE pairs are transformed into a matrix $\mathbf{R}\mathbf{A}$ with $2n$ rows, which produces a little less than $2n$ reduced LWE samples (for $n \geq 256$, we observe about 1% zero rows; this fraction is larger for smaller n).

Algorithm choice. We experimented with the two most standard basis-reduction algorithms: LLL and BKZ. Note that we are not pursuing the same goals as mainstream cryptographic applications of BKZ and LLL. Our objective here is not to find the shortest vector in the lattice defined by \mathbf{A} , but to transform \mathbf{A} into a matrix with smaller coefficients. Experimentally, we find that BKZ with small block size ($\beta = 16 - 20$) achieves better reduction than LLL. We also note that BKZ speed-ups, such as BKZ2.0 [15], do not seem to result in improved reduction. For BKZ, the block sizes needed to achieve reduction in PICANTE are significantly smaller than the block sizes that would be required to perform a lattice-reduction attack on problems of the same dimensions (see also § 7).

4.2 TINYLWE: linear number of samples

Transformers, the deep learning architecture used in SALSA PICANTE, typically train on millions of examples. SALSA assumed access to $4,000,000 \approx 2^{22}$ LWE pairs (\mathbf{A}, \mathbf{b}) with the same secret – unrealistic in practice. To mitigate this, PICANTE introduces TINYLWE, a technique that only requires $m = 4n$ LWE pairs – linear in the dimension n .

The goal of TINYLWE is to produce the large set of 4 million samples required to train our models, from a small initial set of $m = 4n$ LWE pairs. Previous works [5, 49] have observed that a set of m LWE pairs $(\mathbf{a}_1, b_1) \dots (\mathbf{a}_m, b_m)$ can always be expanded by considering the linear combinations $(\mathbf{a}, b) = (\sum_i c_i \mathbf{a}_i, \sum_i c_i b_i)$, with $c_i \in \mathbb{Z}$ and $\sum_i |c_i|$ small. We could, therefore, generate a large sample from a small initial set of LWE pairs by creating many such linear combinations, and reducing them using BKZ.

Unfortunately, LWE error is amplified by linear combinations: $e' = b - \mathbf{a} \cdot \mathbf{s} = \sum_i c_i e_i$, with e_i the error in the original LWE sample. Assuming that the c_i are centered, the standard deviation of error grows as the square root of the number of terms in the combination (\sqrt{n} in the general case) times the standard deviation of the distribution of c_i (which is $\sqrt{C/3}$ if

we assume the c_i are uniformly distributed in $[-C, C]$). LWE data reduction further amplifies errors, making transformer training, and secret recovery, very difficult.

Instead of linear combinations, PICANTE uses *subsampling*: we simply use subsets of n out of the m original LWE samples, $(\mathbf{a}_{j_1}, b_{j_1}), \dots (\mathbf{a}_{j_n}, b_{j_n})$, are randomly selected, and arranged in a matrix \mathbf{A} , with rows \mathbf{a}_{j_1} to \mathbf{a}_{j_n} , which is then reduced as per § 4.1. Because the original LWE pairs are merely copied into \mathbf{A} , the associated noisy inner products have the same error distribution as the original samples. This technique allows us to generate up to $\binom{4n}{n}$ unique matrices ($\approx 9.48^n \cdot 0.46/\sqrt{n}$). In PICANTE, we use subsampling to generate about $2^{21}/n$ matrices, which, after reduction, will result in about 2^{22} reduced LWE pairs. Even though subsampled matrices often have rows in common, we experimentally observe that, after reduction, there are almost no duplicate vectors. For $n = 80$, we counted one duplicate in 50,000 examples; for $n \geq 150$, we found no duplicates in 4 million examples.

4.3 Encoding LWE pairs

Prior to model training, PICANTE encodes the LWE samples (i.e. (\mathbf{a}, b) pairs) into sequences of tokens that the transformer can process. After encoding, the integer coordinates of \mathbf{a} and b are represented using two digits in base B (with $B \geq \sqrt{q}$). Our experiments with different values of B (see § 6.1) suggest that large values of B , which limit the most significant digit of a_i and b to a small number of values (i.e. $B = O(q)$), allow for better performance. In our experiments, we use $B = \lfloor q/k \rfloor$ with $k = 2 \cdot \lceil \frac{n}{100} \rceil + 2$.

Using large B creates a problem: for large values of q (for $n \geq 200$, we have $q > 100,000$), the size of the vocabulary becomes too large to be learned by a transformer trained on only 4 million LWE pairs. To mitigate this, we encode the lowest digits of \mathbf{a} and b into B/r buckets of size r (i.e. divide them by r). The value r is chosen so that the overall vocabulary size $B/r < 10,000$ (see Table 2). The use of buckets helps train models for large n but it also causes a loss of precision in the values of \mathbf{a} and b . We believe the impact on performance is limited, because the low bits of \mathbf{a} and b that are rounded off by buckets are those most corrupted by LWE error.

4.4 Models and training

Model architecture. Our transformer architecture, summarized in Figure 2, is strongly inspired by SALSA [49]. Following [47], we use a sequence-to-sequence (seq2seq) model [18], composed of two transformer stacks – an encoder and a decoder – connected by a cross-attention mechanism. The encoder processes the input sequence, the coordinates of \mathbf{a} , represented as sequences of digits. The discrete input tokens are first projected over a high-dimensional space (we use dimension $d = 1024$) by a Linear Embedding Layer with trainable

weights (i.e. embedding is learned during training). The resulting sequence is then processed by a single-layer transformer: a self-attention layer with 4 attention heads, and a FFNN with one hidden layer of 4096 neurons.

The decoder is an auto-regressive model. It predicts the next token in the output sequence, given already decoded output and the input sequence processed by the encoder. At the beginning, the decoder is provided with a conventional beginning of sequence token (BOS), and predicts b_1^* , the first digit of b . It is then fed the sequence BOS, b_1^* , and decoding proceeds until a specific token (EOS) is output.

Decoder input tokens are encoded as 512-dimensional vectors, by a trainable embedding (which also serves to decode the transformer output). The decoder has two layers. First it uses a shared layer (as in [21]), which is iterated through 8 times, feeding layer output back into its input. This recurrent process is controlled by a copy-gate mechanism [19], which decides whether a specific token should be processed by the shared layer or just copied *as is*, skipping the next iteration. After 8 iterations, the output of the shared layer is fed into a “regular” transformer layer. Finally, the output of the decoder is processed by a linear layer, which computes the probabilities that every word in the vocabulary is the next token. The largest probability is selected by applying a softmax function (a differentiable counterpart of the max function).

Both decoder layers are connected with the encoder via a **cross-attention** mechanism with 4 attention heads. In each cross-attention head, the output of the encoder $E = (E_i)_{i \in \mathbb{N}_l}$ (l the input sequence length) is multiplied by two trainable matrices, W_K and W_V , yielding the *Keys* $K = W_K E$ and *Values* $V = W_V E$. The 512-dimensional vector to be decoded, D , is multiplied by a matrix W_Q , yielding the *Query* $Q = W_Q D$. The l scores are then calculated from the query and keys:

$$\text{scores}(E, D) = \text{Softmax}((W_Q D)(W_K E)^T).$$

The scores measure how important each element in the encoder input is when decoding D (i.e. computing b). The cross-attention value for this head is computed as the dot product of scores and values. The values of different heads are then processed by a final linear layer. Cross-attention scores quantify the relation between input positions and output values. PICANTE uses them to recover the secret bit by bit.

Model training. We train the model to predict b from \mathbf{a} , and frame this task as a supervised multi-classification problem:

$$\min_{\theta \in \Theta} \sum_{i=1}^N \sum_{j=1}^K \sum_{k=1}^V \mathbf{1}[y_i[j] = k - 1] \frac{e^{f_{\theta}(x_i)[j,k]}}{\sum_{k'=1}^V e^{f_{\theta}(x_i)[j,k']}}, \quad (1)$$

where $f_{\theta}(x_i) \in \mathbb{R}^{K \times V}$ are model logits evaluated at x_i , $\theta \in \Theta$ are the model parameters, N the training sample size, $K = 2$ the output sequence length and $V = B/r$ the vocabulary size.

To solve (1), we minimize the cross entropy between model predictions $\mathcal{M}(\mathbf{a})$ and the ground truth b , over all tokens in

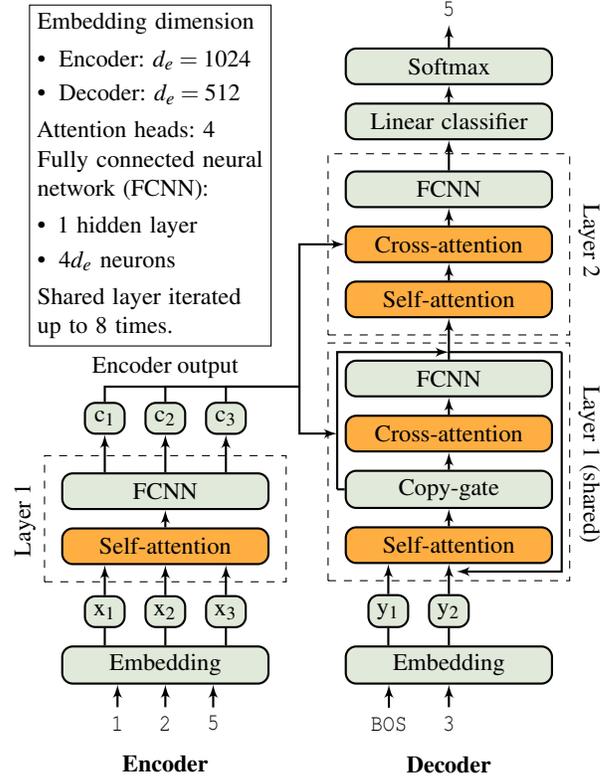


Figure 2. Our transformer architecture.

the output sequence. An alternative approach would set this as a regression problem. However, we believe classification is better adapted to the modular case, and prior works report that reformulating regression as classification leads to state-of-the-art performance [1, 41, 42, 45].

Training is performed by batches of $n_b = 128$ examples. We compute the cross-entropy loss $\mathcal{L}(\mathcal{M}, \mathbf{a}, b)$ over all batch examples, and calculate its gradients $\nabla \mathcal{L}$ with respect to the model parameters (via *back-propagation*). Model parameters are then updated, using the Adam optimizer [27], by $\text{lr} \nabla \mathcal{L}$. The learning rate, lr is set to 10^{-5} , except during the 1000 first optimizer steps, where it is increased linearly from 10^{-8} to 10^{-5} . Every 2 million examples (an *epoch*), model performance is evaluated on a held-out sample, and we attempt to recover the secret. If we fail, another epoch begins.

4.5 Secret Recovery

The intuition behind secret recovery goes as follows: if a model \mathcal{M} can predict b from \mathbf{a} with larger than chance-level accuracy, then \mathcal{M} must somehow “know” the secret key \mathbf{s} , and we can recover \mathbf{s} from \mathcal{M} . SALSA PICANTE uses three methods—*cross attention*, *direct*, and *distinguisher* recovery—to recover the secret. They can be combined for greater accuracy.

Checking correctness. Recovery methods make guesses s' about the (unknown) secret \mathbf{s} . The test whether $s' = \mathbf{s}$ was

introduced in SALSA. On a test sample of N_{test} LWE pairs $(\mathbf{a}_i, b_i)_{1 \leq i \leq N_{\text{test}}}$, compute $b'_i = \mathbf{a}_i \cdot \mathbf{s}'$, and consider the distribution r of $r_i = b'_i - b_i \pmod q$. If $\mathbf{s}' = \mathbf{s}$, then $r \approx e$, the LWE error, with standard deviation σ . If $\mathbf{s}' \neq \mathbf{s}$, then r will be approximately uniformly distributed over \mathbb{Z}_q , with standard deviation $\sigma' = q/\sqrt{12}$. By estimating σ' on a large enough test sample, one can verify $\mathbf{s}' = \mathbf{s}$ to any confidence level.

Such a test can be done on the original LWE samples, that is, with $N_{\text{test}} = m = 4n$. In § A.2 we provide a statistical analysis of this verification technique, which demonstrates that this sample size is sufficient for all dimensions $n \geq 80$.

Cross-Attention. In this new recovery method, we guess the secret from the parameters of a trained model, leveraging the cross-attention scores of the first decoder layer (see Figure 2 and § 4.4). Intuitively, the cross-attention score measures the relevance of input tokens (i.e. coordinates of \mathbf{a}) for the computation of b . Since $b = \mathbf{a} \cdot \mathbf{s} + e$, the coordinates of \mathbf{a} that correspond to the 0 bits of \mathbf{s} have no impact on b . On the other hand, the coordinates associated to the 1 bits in \mathbf{s} have an impact proportional to their value. Therefore, we expect that high cross-attention scores will be found for the input positions that correspond to 1s in the secret.

In practice, we run the trained transformer on a test set of 10,000 reduced LWE samples, and sum the cross-attention scores of all heads. Since \mathbf{a} has n coordinates encoded on 2 tokens, this results a $2n$ -dimensional vector, from which we keep the odd positions (i.e. the high digits of the coordinates of \mathbf{a}), to generate an n -dimensional score vector V . A secret guess \mathbf{s}' is then produced by setting the h largest coordinates of V to one, and the rest to 0.

Direct recovery. PICANTE uses the same direct recovery method as SALSA. This technique leverages the fact that trained transformers can generalize to input they have not seen at training. The trained model is run for special vectors \mathbf{a} , with only one non-zero coordinate: $\mathbf{a} = K\mathbf{e}_i$ with \mathbf{e}_i the i -th basis vector of the n -dimension space, and $K \in \mathbb{Z}_q$. For these vectors, since $b = \mathbf{a} \cdot \mathbf{s} + e$, and e is small, we should have $b \approx 0$ if the i -th bit in the secret $s_i = 0$, and $b \approx K$ if $s_i = 1$. Details can be found in [49]. In practice, we choose different values K_j and run the transformer on $K_j \cdot \mathbf{e}_i$ for every index $i = 1, \dots, n$, identifying potential 1-bits in the secret as above. This gives us a guess for a secret for each K_j .

To obtain a score for each bit to be used in combination methods, for each index i , we compute a score by summing the resulting values of $\mathcal{M}(K_j \cdot \mathbf{e}_i)$ (or, equivalently, taking the mean). We then guess the secret \mathbf{s}' by assuming that the h coordinates with the largest scores are 1 and the rest are 0.

Distinguisher. PICANTE’s version of distinguisher recovery improves upon that of SALSA. This technique is predicated on trained model consistency. The general idea is that if the i -th bit of the secret $s_i = 0$ and \mathbf{e}_i is the i -th standard basis vector, then the model should predict close values for \mathbf{a} and $\mathbf{a} + K \cdot \mathbf{e}_i$.

SALSA’s distinguisher took a LWE sample (\mathbf{a}, b) and com-

pared b with the model prediction $b' = \mathcal{M}(\mathbf{a} + K \cdot \mathbf{e}_i)$ for some random $K \in \mathbb{Z}_q$. This presupposed good model accuracy, i.e. $\mathcal{M}(\mathbf{a}) \approx b$, which often does not happen in practice. In PICANTE, instead of comparing $b' = \mathcal{M}(\mathbf{a} + K \cdot \mathbf{e}_i)$ to b , we compare it to $\mathcal{M}(\mathbf{a})$. The rest of the method is unchanged. The secret guess is determined as in the Direct recovery method.

This new method brings two benefits. First it exploits trained model consistency without requiring prediction accuracy. In practice, this means recovery can happen earlier during training, and even in cases when the model never learns to predict accurately. Second, the SALSA distinguisher needed LWE samples (\mathbf{a}, b) . Since PICANTE’s method no longer uses b , the distinguisher can use randomly generated \mathbf{a} . This reduces the number of LWE samples necessary for the attack.

This recovery method relies on a large number of model evaluations, which makes it very slow. To improve its speed, we use the same \mathbf{a} across all secret bits s_i , therefore halving the number of model evaluations.

Combined secret recovery. In our three recovery methods, a score is computed for every bit in the secret. The secret guess is computed by setting to one the h bits with the largest scores. By combining the scores from different methods, we create four additional techniques, which improve secret recovery (i.e. can recover secrets when all three methods fail).

- *Aggregated rank.* The scores of each method are sorted from largest to smallest, and replaced by their rank. For each bit, the aggregated rank is 1) the highest of the ranks (Highest Rank) or 2) the sum of the ranks (Sum Rank). The h bits with the highest aggregated ranks are set to 1.
- *Aggregated normalized scores.* The scores of each method are normalized to $[0, 1]$. The h bits with the maximum normalized scores (Max Normalized) or the highest sum of normalized scores (Sum Normalized) are set to 1.

These combination rules amount to setting secret bits to 1 bit positions where all, some, or any of the secret recovery methods have a high score. We use aggregated scores from all subsets of the secret recovery methods. Other combination rules could be considered. These mixing techniques are cheap to implement, because they do not require additional transformer evaluation.

4.6 Experimental setting

All PICANTE experiments use the following assumptions.

- For each n , the modulus q is selected after consulting Table 1 in [12]. We set our q such that $\log_2 q$ is smaller than the smallest successful attack reported there (see Table 2).
- For BKZ reduction parameters, we set $\omega = 15$ for the error penalization parameter, for all n . Block size and the LLL parameter δ in *fpLLL* are set to 20 and 0.99 for all $n \leq 200$. To keep preprocessing time reasonable, we decrease these values for $n = 256, 300$ and 350 (see Table 8).

- The error in the original LWE pairs is sampled from a discrete Gaussian distribution, centered at 0, and with $\sigma = 3$.
- We aim to recover binary secrets with sparsity $h/n \approx 10\%$. We usually present recovery data for seven consecutive Hamming weights around $n/10$.
- TINYLWE uses $4n$ randomly generated samples (\mathbf{a}, b) .
- For each n and q , we perform the preprocessing step on random matrices A only once. Then we use that reduced data for experiments with many different secrets and Hamming weights. The samples are generated from the reduced matrices \mathbf{RA} by first generating $\mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ and then computing \mathbf{Rb} . In TINYLWE, the matrices \mathbf{A} do not come from independent LWE samples, but are subsampled from the m original LWE pairs (\mathbf{a}, b) . To generate the values \mathbf{Rb} correctly, we must keep track of which pairs we subsampled. We first generate the values b for the original LWE pairs, and then compute \mathbf{Rb} for \mathbf{b} subsampled the same way as \mathbf{A} .
- All transformers have the same architecture and number of parameters (see Figure 2), except for ablation results on dimension and attention heads are presented in § 6.2.

n	q	$\log_2 q$	δ	β	base	r
80	113	7	0.99	20	29	1
150	6421	13	0.99	20	1071	1
200	130769	17	0.99	20	21795	2^2
256	6139999	23	0.96	18	767500	2^7
300	94056013	27	0.96	16	11757002	2^{11}
350	3831165139	32	0.93	14	383116514	2^{16}

Table 2. PICANTE parameters. n : dimension, q : modulus, δ : delta-LLL (BKZ), β : block-size (BKZ), base: encoding base, r : bucket size for encoding.

5 SALSA PICANTE’s Performance

5.1 Overall Performance

For $n < 350$, PICANTE recovers LWE secrets with sparsity $h/n \approx 10\%$, a significant improvement over SALSA. Performance is slightly worse for $n = 350$ so far: the best Hamming weight PICANTE has achieved is $h = 25$ (sparsity 7%), or $h = 26$ with a larger architecture (Table 11). We believe this is due to the BKZ parameters selected for preprocessing (see the last line of Table 2), which may not have been aggressive enough. Better reduction parameters, and hence longer preprocessing times, may allow recovery of secrets with higher h . Results for a range of Hamming weights in many dimensions n are presented in Table 3.

PICANTE recovers binary secrets with the TinyLWE approach (§4.2) using a linear number of samples for various Hamming weights in dimensions $n = 80, 150, 200, 256, 300, 350$. For each dimension n , we fix a modulus q which is smaller than the q for which the concrete lattice attacks in [12, Table 1] can recover

$n, \log_2 q$	Hamming weight h						
80, 7	4	5	6	7	8	9	10
success	3/5	3/5	2/5	2/5	1/20	1/20	0/20
epoch	2,5,6	0,1,4	0,3	0,8	3	4	
150, 13	9	10	11	12	13	14	15
success	4/5	2/5	3/5	1/5	1/20	0/20	0/20
epoch	1,1,3,6	2,2	8,8,11	9	13		
200, 17	17	18	19	20	21	22	23
success	3/5	2/5	3/5	1/5	2/5	2/20	0/20
epoch	1,1,8	2,11	2,3,9	7	7,10	12,17	
256, 23	26	27	28	29	30	31	32
success	4/5	1/5	1/5	3/5	3/5	4/20	0/20
epoch	2,3,4,7	10	5	5,9,11	17,20,32	6,12,26,27	
300, 27	28	29	30	31	32	33	34
success	2/5	2/5	1/5	1/5	2/5	1/5	0/20
epoch	6,7	6,13	11	11	21,31	39	
350, 32	20	21	22	23	24	25	26
success	2/5	4/5	2/5	2/5	3/20	1/20	0/20
epoch	4,5	5,6,6,9	5,9	9,15	18,21,38	19	

Table 3. Success rates and number of epochs. Highest recovered Hamming weight is in bold.

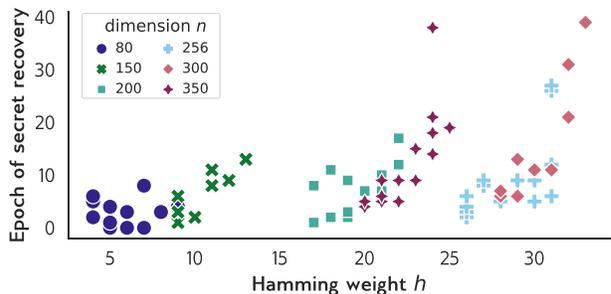


Figure 3. Number of training epochs before secret recovery. For different dimensions and Hamming weights.

secrets with blocksize roughly 40. The size of q is listed as $\log_2 q$ in Table 3 for each n . For each dimension n , we test a range of Hamming weights for the secrets to test. In each case, we show the boundary of Hamming weights where secrets are recovered, along with the trends for the number of successes in each trial, and in which training epoch the secret is recovered. For example, in dimension $n = 350$, we recovered the secret for $h = 25$ in one out of 20 trials, each trial with a different secret. In that case, the success was achieved in epoch 19 of the training. Experiments running trials for the same secret, but different initialization seeds for the model training are presented in § 6.3.

Figure 3 presents the number of epochs needed for secret recovery, for different n and h . Although 66% of the successful secret recoveries occurred in the first 10 epochs for the dimensions and hamming weights we tested so far, we see the number of epochs go up for larger n and h . There is a large variance and it is unclear how the number of epochs grows as we vary n, h , and q . Specifically, about 75% of the successful

experiments succeeded by epoch 4 for $n = 80$; by epoch 8 for $n = 150$; by epoch 13 for $n = 200, 256, 300$; and by epoch 18 for $n = 350$. Training for more epochs does not seem to improve the success probability for lower Hamming weights, but may lead to recovery of higher Hamming weights. We observed this for $n = 300$, for which we reached the highest $h = 32, 33$ with more than 30 epochs.

Note that for a specific value of h , the number of epochs needed for recovery can vary from one experiment to another. This might be due to two factors: different secrets, or different random initializations of the transformers, see § 6.3. Running more trials with different initializations may achieve success in an earlier epoch. However, since training time is not the dominant cost in our approach, we are not currently focused on optimizing for minimizing the number of epochs required.

5.2 TinyLWE vs LWE: few or many samples

PICANTE uses the TINYLWE subsampling technique introduced in §4.2 to recover secrets from only $4n$ collected LWE samples. By comparison, SALSA used 4 million LWE samples. Table 1 compares the performance of PICANTE using TINYLWE, with PICANTE using 2.2 million collected LWE samples, which we call LWE.

Dimension	80	150	200	256	300	350
TINYLWE max h	9	13	22	31	33	26
LWE max h	9	12	21	32	32	25

Table 4. TINYLWE vs LWE. Values: highest h recovered for each n .

There is little difference between the highest Hamming weight of recovered secrets for TINYLWE and LWE (Table 4). Reducing the initial LWE sample size seems to result in no performance loss. In fact, we observe that TINYLWE sometimes recovers larger Hamming weights than LWE. See more detailed comparison results in Table 16 in Appendix A.1.

5.3 Resources needed for PICANTE

The total cost of PICANTE is the sum of the resources needed to preprocess data, train the model, and recover the secret.

n	$\log_2(q)$	Cost per matrix CPU.hours	Matrices needed	Total cost CPU.years
80	7	0.01	34,800	0.05
150	13	3.1	14,600	5.3
200	17	15.9	10,800	19.4
256	23	51.9	8,300	48.1
300	27	105.8	7,100	85.6
350	32	152.0	6,000	105

Table 5. Resources needed for preprocessing. Total resources needed to produce 2^{22} reduced samples, by reducing $2^{21}/n$ matrices. This operation can be run in parallel for each matrix.

Data preprocessing is the most resource intensive part of PICANTE. To generate 2^{22} reduced samples, $2^{21}/n$ matrices must be reduced (one $n \times n$ matrix produces $2n$ reduced samples, see § 4.1). As the dimension increases, the number of matrices scales down linearly. To avoid the exponential cost of BKZ-reduction [43], we fix the block size at most $\beta = 20$ so that the preprocessing step scales polynomially with n and $\log q$. See discussion of the parameter choices for preprocessing in § 5.4 below. In practice, to save resources, we choose smaller β for larger dimensions.

Table 5 reports the preprocessing resources (in cpu hours) for each n . Our preprocessing is fully parallelizable. Using the number of CPUs equal to the number of matrices needed ($2^{21}/n$), one can complete the preprocessing in the time required to reduce one matrix (e.g. 152 hours for $n = 350$).

Model training and secret recovery. The cost of training and recovery is proportional to the number of epochs needed to recover the secret. In § 5.1, we observe large variations of this number, for fixed dimension and Hamming weight. Table 6 reports the average duration of one training epoch and associated secret recovery. All models use the same number of parameters, batch size (128) and epoch size (2 million examples), and are trained on one NVIDIA V100 GPU.

Training time increases with dimension. This is expected, as input sequence length is $2n$, i.e. linear in the dimension n . For secret recovery, the time required for each method is proportional to the number of transformer inferences needed, multiplied by the time required for each inference. The cross-attention method uses a constant number of inferences. Direct recovery uses $15n$ inferences; distinguisher recovery uses $200n$. Like training, the time for a single inference scales linearly with n . Overall, cross-attention recovery scales linearly with n , and direct and distinguisher scale quadratically.

As a result, whereas for $n = 80$, secret recovery only represents 10% of the overall cost of an epoch (i.e. training and recovery), it accounts for almost 50% of the total cost for $n = 350$, and the distinguisher method is responsible for more than 45%. Overall cost could be reduced by not running the distinguisher for all epochs, e.g. waiting until the model has reached a certain accuracy, or running the distinguisher every other epoch. We report the total cost of training and recovery on one GPU, but we can reduce the time significantly by parallelizing across many GPUs. Secret recovery could be run in parallel on a separate machine, and even on different machines for different epochs.

5.4 Effect of preprocessing

Through extensive experimentation in many dimensions, we observed that preprocessing was a critical step to help the transformer to learn and to recover high Hamming weight secrets. Preprocessing using BKZ changes the distribution of both the size of the entries of \mathbf{A} modulo q (shown for $n = 150$ in Figure 4) and the norm of its rows (see Figure 5). Note that

n	80	150	200	256	300	350
$\log q$	7	13	17	23	27	32
Training	41.9	52.2	67.8	81.5	91.8	104.5
Total rec.	4.6	10.1	28.0	54.5	73.3	99.9
CA	0.4	0.8	1.1	1.6	2.0	2.5
Direct	0.4	0.8	2.0	3.0	5.0	6.4
Dist.	3.8	8.5	24.9	49.9	66.3	91.0
Total	46.5	62.3	95.8	135.9	165.0	204.4

Table 6. Training and recovery time per epoch (minutes). All models are trained on a single NVIDIA V100 GPU. Total rec.: total time for Secret recovery; CA: Cross attention; Dist.: Distinguisher method.

we are not trying to obtain the shortest vector in the lattice like the classical uSVP, decoding, and dual attacks.

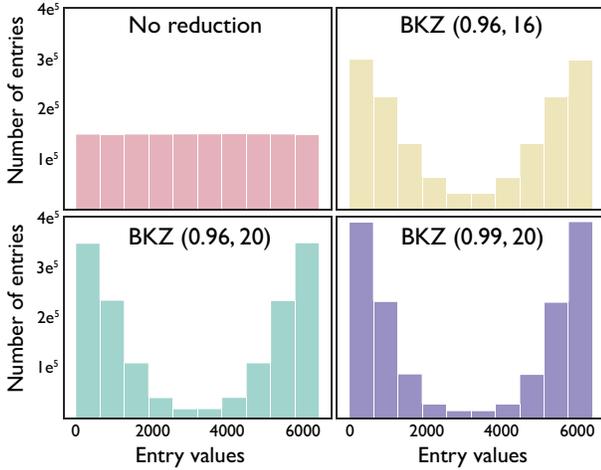


Figure 4. Distribution of sample entry values as strength of norm reduction increases ($n = 150, q = 6421$). BKZ parameters listed as BKZ (β, δ).

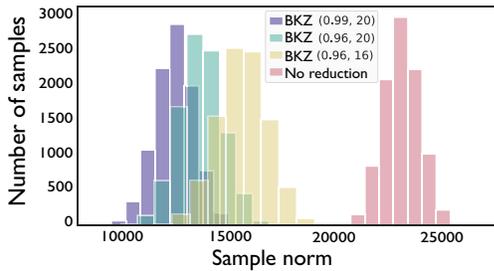


Figure 5. Distribution of sample norms as strength of norm reduction increases ($n = 150$). BKZ parameters listed as BKZ (β, δ).

To determine how much preprocessing to do, we experimented with setting targets for the standard deviation of the entries of the rows. For random matrices \mathbf{A}_{random} , the standard deviation of the entries is $\text{std}(\mathbf{A}_{random}) \approx q/\sqrt{12} \approx 0.29q$. Empirically we observed that if we reduce the standard deviation of the entries to about 35% of that of a random matrix, i.e. $\text{std}(\mathbf{A}) \approx 0.35\text{std}(\mathbf{A}_{random}) \approx 0.1q$, we achieved success in transformer learning and secret recovery. In practice we

decide which parameters β and δ to use in BKZ in *fpLLL* by processing a single matrix \mathbf{A} with various choices and selecting the parameters which give us $\text{std}(\mathbf{A})/\text{std}(\mathbf{A}_{random}) \leq 0.35$ to process the rest of the $2^{21}/n$ matrices.

For $n = 150$ and $q = 6421$ we show the effect of preprocessing — starting with no preprocessing and increasing δ, β parameters until we reach those used in PICANTE. Table 7 shows the highest h secret recovered. With no preprocessing, we were not able to recover any of the secrets for $h \in \{3, \dots, 6\}$; PICANTE succeeds in recovering secrets with $h = 12$.

	δ	-	0.96	0.96	0.99
	β	-	16	20	20
highest h	-	5	8	12	
$\text{std}(\mathbf{A})/\text{std}(\mathbf{A}_{random})$	1	0.667	0.578	0.526	
$\text{norm}(\mathbf{A})/\text{norm}(\mathbf{A}_{random})$	1	0.669	0.581	0.528	
$\text{entropy}(\mathbf{A})/\text{entropy}(\mathbf{A}_{random})$	1	0.973	0.957	0.946	
cost / matrix (minutes)	0	30	54	188	
time out (minutes)	-	60	120	333	

Table 7. Highest Hamming weight secret recovered for varying δ and β ($n = 150, q = 6421$). $\text{std}(\mathbf{A})$: standard deviation of \mathbf{A} 's coefficients post-reduction; $\text{norm}(\mathbf{A})$: average norm of \mathbf{A} 's rows post-reduction; $\text{entropy}(\mathbf{A})$: entropy of \mathbf{A} 's coefficients. The last column are PICANTE's parameters.

Table 8 presents the highest Hamming weight recovered and standard deviations of reduced \mathbf{A} for different n . We observe that, contrary to BKZ-based lattice attacks on LWE, we can **reduce** block size and δ for larger dimensions ($n = 256$ and 300), yet recover secrets with $\geq 10\%$ sparsity. On the other hand, the $n = 350$ results (sparsity 7%) may be caused by low β and δ values, chosen to reduce preprocessing time.

n	$\log_2 q$	δ	β	$\text{std}(\mathbf{A})/\text{std}(\mathbf{A}_{random})$	$\max h$
80	7	0.99	20	0.776	9
150	13	0.99	20	0.527	13
200	17	0.99	20	0.399	22
256	23	0.96	18	0.331	31
300	27	0.96	16	0.320	33
350	32	0.93	14	0.335	25

Table 8. Preprocessing parameters for BKZ in *fpLLL* for PICANTE's best secret recoveries. β : block size, δ : LLL_DELTA.

6 Additional Results

6.1 Effect of encoding base

In this section, we discuss the impact of B , the base we use to encode b and the coordinates of \mathbf{a} (§ 4.3). Table 9 presents the impact of different base choices, for $n = 150$. For this dimension, no buckets are used, i.e. $r = 1$. To keep input sequences short, we want all integers up to q to be encoded on two tokens, i.e. $B \geq \sqrt{q}$. However, we note that values of B close to \sqrt{q} result in worse secret recovery. Experimentally,

base	$h = 9$	$h = 10$	$h = 11$	$h = 12$
$81 \approx \sqrt{q}$	1/5	0/5	0/5	0/5
$402 \approx q/16$	3/5	2/5	0/5	0/5
$803 \approx q/8$	3/5	2/5	2/5	2/5
$1071 \approx q/6$	4/5	2/5	3/5	1/5
$1606 \approx q/4$	3/5	2/5	0/5	0/5

Table 9. Secret recovery rate for different bases B . $n = 150$, $q = 6421$. PICANTE parameters and results are in bold.

Encoding base B	r	Hamming weight h			
		26	27	28	29
$2478 \approx \sqrt{q}$	1	0/5	0/5	0/5	0/5
$383750 \approx q/16$	32	4/5	1/5	1/5	0/5
	128	3/5	2/5	1/5	3/5
	512	4/5	2/5	1/5	3/5
$767500 \approx q/8$	32	4/5	1/5	1/5	1/5
	128	4/5	1/5	1/5	3/5
	512	4/5	2/5	1/5	3/5
$1535000 \approx q/4$	32	0/5	0/5	0/5	0/5
	128	0/5	0/5	0/5	0/5
	512	1/5	0/5	0/5	0/5

Table 10. Secret recovery rates for different bases B and bucket sizes r . $n = 256$, $q = 6139999$. PICANTE parameters are in bold.

we observe that the highest recovery rates are achieved when $B = \lfloor q/k \rfloor$ with $k = 6$ or 8 .

We now study the impact of base B and bucket size r for $n = 256$ and $q = 6139999$. Our choice of base $B = \lfloor q/k \rfloor$ would result in a vocabulary that is too large for the transformer. Hence, we need to tokenize the low order digits in buckets of size r as in § 4.3. We notice that a low value of $B \approx \sqrt{q}$ prevents secret recovery, as does $B = q/4$. For $B = q/8$ and $q/16$, we note the comparable performance for $r = 128$ (our choice on the basis of vocabulary size) and $r = 512$.

6.2 Effect of model architecture

All PICANTE experiments use the same transformer architecture (see § 4.4 for details). For large problem dimension, SALSA reported improved performance with larger models, specifically increased embedding dimensions. Also, the number of attention heads used in PICANTE, 4 in the encoder and decoder, is low, compared to common transformer architectures. Most transformers with 512 dimensions use 8 heads. In this section, we investigate the impact of larger dimensions and number of heads in the encoder and decoder, on secret recovery rate (Table 11), for $n = 350$. We note that increasing dimension and heads do not result in better performance. This contrasts with results in NLP, where performance usually increases with model size.

embedding size encoder/decoder	number of attention heads encoder/decoder/cross attention				
	4/4/4	4/4/8	4/4/16	8/8/8	8/8/16
1024 / 512	25	23	24	-	-
1024 / 768	-	26	24	24	23
1280 / 512	-	24	23	26	22
1280 / 768	-	24	23	23	23

Table 11. Effect of model architecture on PICANTE’s performance. Data shown is the highest Hamming weight recovered over $n = 350$ TINYLWE. PICANTE’s parameters are in bold.

Experiment	Success	Mean epoch	Min, max epochs
1	20/20	4.2	2, 7
2	12/20	12.1	8, 25
3	15/20	9.1	6, 16
4	0/20	-	-

Table 12. Effect of model initialization on secret recovery. $n = 200$, $h = 19$.

6.3 Effect of model initialization

Transformer parameters are randomly initialized before training, and their initial values may impact their performance, a phenomenon known as “lottery tickets”: models learning better, or faster, for some initial parameter values. We explore this effect in 4 experiments, with different secrets, for $n = 200$ and $h = 19$. Each experiment in Table 12 has a different secret; for each secret we train 20 transformers, each initialized with a different seed.

In experiment 1, the secret is recovered for all 20 different seeds, in epochs 2 to 7. For experiments 2 and 3, the secret is recovered about 3/4 of the time, between 6 and 25 epochs. In experiment 4, the secret is never recovered.

This sheds light on results from § 5.1: we observed that, for given n and h , the number of epochs needed for secret recovery varied a lot. Initialization seems an important factor.

This suggests a possible improvement to PICANTE in situations where large computing resources are available. By training several transformers, with different initialization, on the same reduced data, we not only increase our chances of secret recovery, but can also improve training speed, by ending training on all models as soon as one recovers the secret.

6.4 Secret recovery methods

PICANTE leverages 4 secret recovery methods (§ 4.5): distinguisher, direct, cross-attention, and combined. The first three methods output *bit scores* and secret guesses s' . The *bit scores* rank the likelihood of individual secret bits having value 1. The *combined secret recovery* method allows PICANTE to create additional secret guesses by aggregating the scores of the previous methods. Table 13 reports the successes/attempts of all secret recovery methods for TINYLWE. We only report

the method(s) that succeed first: we terminate the experiment after successful recovery. We say that the combined method is successful if and only if it recovered the secret when no individual method would. If an individual method succeeds, the combined method typically succeeds as well.

$n, \log_2 q$	Hamming weight h						
80, 7	4	5	6	7	8	9	10
success	3/5	3/5	2/5	2/5	1/20	1/20	0/20
Distinguisher	3/5	3/5	2/5	2/5	1/20	1/20	0/20
Direct	2/5	1/5	1/5	1/5	1/20	1/20	0/20
Cross-attention	3/5	1/5	2/5	0/5	1/20	0/20	0/20
Combined	0/5	0/5	0/5	0/5	0/20	0/20	0/20
150, 13	9	10	11	12	13	14	15
success	4/5	2/5	3/5	1/5	1/20	0/20	0/20
Distinguisher	2/5	1/5	2/5	0/5	0/20	0/20	0/20
Direct	2/5	0/5	0/5	0/5	1/20	0/20	0/20
Cross-attention	1/5	0/5	1/5	1/5	0/20	0/20	0/20
Combined	0/5	1/5	0/5	0/5	0/20	0/20	0/20
200, 17	17	18	19	20	21	22	23
success	3/5	2/5	3/5	1/5	2/5	2/20	0/20
Distinguisher	2/5	1/5	2/5	1/5	0/5	0/20	0/20
Direct	0/5	0/5	1/5	0/5	0/5	0/20	0/20
Cross-attention	1/5	0/5	0/5	0/5	0/5	1/20	0/20
Combined	0/5	1/5	1/5	0/5	2/5	1/20	0/20
256, 23	26	27	28	29	30	31	32
success	4/5	1/5	1/5	3/5	3/5	4/20	0/20
Distinguisher	3/5	1/5	1/5	1/5	1/5	0/20	0/20
Direct	0/5	0/5	0/5	0/5	0/5	0/20	0/20
Cross-attention	2/5	0/5	1/5	3/5	2/5	2/20	0/20
Combined	1/5	0/5	0/5	0/5	0/5	2/20	0/20
300, 27	28	29	30	31	32	33	34
success	2/5	2/5	1/5	1/5	2/5	1/5	0/20
Distinguisher	2/5	0/5	0/5	0/5	0/5	0/5	0/20
Direct	0/5	0/5	0/5	0/5	0/5	0/5	0/20
Cross-attention	1/5	2/5	1/5	0/5	1/5	1/5	0/20
Combined	0/5	0/5	0/5	1/5	1/5	0/5	0/20
350, 32	20	21	22	23	24	25	26
success	2/5	4/5	2/5	2/5	3/20	1/20	0/20
Distinguisher	1/5	2/5	2/5	1/5	0/20	0/20	0/20
Direct	0/5	0/5	0/5	0/5	0/20	0/20	0/20
Cross-attention	1/5	1/5	1/5	1/5	1/20	1/20	0/20
Combined	1/5	2/5	0/5	0/5	2/20	0/20	0/20

Table 13. Secret recovery successes/attempts for PICANTE’s four recovery methods. For each (n, q, h) setting, we report the number of secrets PICANTE recovers out of attempted attacks (“success” row) and the number of recoveries by each individual method (“Distinguisher” through “Combined” rows). If two methods succeed in the same training epoch, we report both successes, so individual recoveries may exceed the number of total successes. Combined method successes are only reported when all other methods fail.

Two trends are evident in Table 13. First, the direct recovery method is outperformed by other methods as n increases. It works well at $n = 80$, but for $n \geq 256$, it is either slower than other methods or fails to recover the secret. Recall that direct recovery works when for every bit of the secret, the model predictions $\mathcal{M}(K \cdot \mathbf{e}_i)$ corresponds to the secret bit: large when the secret bit is 1 and small otherwise. This happens with lower probability as n grows. Second, the combined recovery method performs better as n increases. Probably for larger n , individual methods cannot glean information about all secret bits, but each gains some information about certain bits. Thus, combining their scores may lead to additional recoveries.

7 Comparison to Existing Attacks on LWE

In practice, the training stage of PICANTE takes less time than pre-processing the data (see Table 5 and Table 6 in § 5.3), so we focus on comparing the cost of preprocessing with the cost of *classical lattice reduction attacks* such as uSVP, decoding, and dual attacks. As PICANTE uses the *fpLLL* package for lattice reduction algorithms, we compare the running times of PICANTE with the uSVP attack, run using *fpLLL*.

The LWE Estimator software package [5] is used to estimate the cost of classical lattice reduction attacks. The LWE Estimator uses theoretical formulas and heuristic estimates to predict which block size will be required for BKZ to recover the secret for a given lattice parameter size. These estimates are widely used to set parameters and estimate security at parameter sizes for which it is impossible to actually run these classical attacks (they would not terminate in our lifetimes). Concrete running times for actual successful attacks can be found in a few places in the literature, e.g. in [4, 8, 12, 28], and we find those useful for comparison here. In particular, for dimensions $n \leq 200$, we chose values of $\log q$ strictly smaller than those used in [12]; for dimensions $n = 256, 300$ and 350 , we use much smaller values of $\log q$ than [28]: for instance, for $n = 350$, we use $\log q = 32$, much smaller than the value $\log q = 52$ in [28].

We present here 3 ways to quantify, estimate, and compare with pure lattice reduction attacks: the LWE estimator, concrete timings for running the uSVP attack at small sizes, and theoretical and heuristic formulas.

LWE Estimator. Table 14 presents the block size and estimated cost for classical attacks, to compare against PICANTE’s successful secret recoveries (using the highest h achieved by PICANTE). LWE Estimator [5] costs are listed in terms of the number of operations in \mathbb{Z}_q , the cost of which can be approximated by $(\log q)^2$.

For example, for $n = 256$, the LWE estimator predicts that the uSVP attack should succeed with block size 40 and cost about $2^{41.8}$ operations in \mathbb{Z}_q . PICANTE uses block size 18.

Concrete running times. Table 15 presents concrete running times for the following attack: We run the primal uSVP attack,

n	q	best attack	cost	block size
80	113	BDD	$2^{48.0}$	$\beta = 63$
150	6421	BDD	$2^{42.7}$	$\beta = 44$
200	130769	BDD	$2^{41.8}$	$\beta = 41$
256	6139999	uSVP/BDD	$2^{41.8}$	$\beta = 40$
300	94056013	uSVP	$2^{41.9}$	$\beta = 40$
350	3831165139	uSVP	$2^{42.0}$	$\beta = 40$

Table 14. LWE Estimator [5] estimates for PICANTE’s most successful recoveries (see Table 1). Cost: number of operations in \mathbb{Z}_q .

using Kannan’s embedding and BKZ2.0 [15] with different block sizes. The dimension for the Kannan’s embedding is determined as in [12]. We choose block sizes close to the block size predicted by the LWE Estimator and compare PICANTE against attacks with similar success probability.

We ran the classical attacks for dimensions up to $n = 256$ with the block size predicted by the LWE Estimator ($\beta = 40$). For $n = 300, 350$, already the first loop in BKZ takes longer than 3 days; we expect the full attack to take several weeks.

The uSVP attack was run using the *fpLLL* package on the same machine as the norm-reduction step of PICANTE. We did not use any optimization for either of the attacks. We see that for $n = 80$ and $h = 9$, PICANTE with $\beta = 20$ achieves similar success to uSVP with $\beta = 60$. The uSVP attack takes about 10 hours to succeed; the time spent on data preprocessing for PICANTE is negligible with enough parallelization, and the training (run on 1 GPU) for successful recoveries took 5 epochs of about 0.8 hours each. The minimal time spent by PICANTE is therefore about 4 hours. For $n = 256$ and $h = 31$, the uSVP attack with block size $\beta = 35$ (smaller than predicted by the LWE estimator) takes about 280 hours to succeed on average; PICANTE with sufficient parallelization needs 52 hours for data pre-processing and about 19 hours for training, so the total time is about 71 hours.

We stress that these timings are rough estimates. Note that optimizations to lattice-reduction for the uSVP attack could also speed up the data preprocessing of PICANTE. We did not include any possible savings from parallelizing the training and secret recovery methods (see § 5.3).

Theoretical analysis. Denote by $BKZ(d, \beta)$ the (classical) cost of BKZ reduction in dimension d with block size β . It can be estimated [2] as

$$2^{0.292\beta+c} \ll BKZ(d, \beta) < 8d \cdot 2^{0.292\beta+c},$$

where the cost $SVP(\beta) = 2^{0.292\beta+c}$ is the cost of the SVP oracle in dimension β (a major step in the BKZ-reduction algorithm). The constant c depends on the attack model—16.4 for sieving and 0 for other analyses. The upper bound arises from the estimate that 8 runs (full loops) of the BKZ-reduction, and hence $8d$ calls to the $SVP(\beta)$ oracle, are needed.

The uSVP attack solves the shortest vector problem in dimension $d > n$; PICANTE applies the BKZ-reduction to

lattices of dimension $d = 2n$ but keeps the block size close to constant ($\beta = 20$, and decreases the block size in larger dimensions for efficiency). We make this choice because the cost of BKZ reduction scales exponentially with the block size. While we do not know whether we can use constant block size $\beta = 20$ for all dimensions, we do expect our block size to grow slower than the block sizes required for lattice-reduction attacks.

High level comparison. PICANTE compares with classical lattice reduction attacks as follows: PICANTE succeeds in recovering the secret vector using much smaller block size than pure lattice reduction attacks, at the expense of processing many more matrices (2.2 million/ n matrices). Because this step is run in parallel, PICANTE recovers secrets faster than the uSVP attack but uses many more CPUs for parallel processing. As the dimension increases and/or $\log q$ decreases, we expect the advantage of PICANTE to grow, due to the exponential cost of the lattice reduction attacks based on BKZ. Future work may produce a more efficient way to preprocess the data or reduce the amount of data needed for training.

8 Discussion

Our attack, PICANTE, demonstrates a dramatic improvement over SALSA, the only prior work on attacking LWE with Machine Learning. PICANTE successfully recovers LWE binary secrets with sparsity up to 10%, for dimensions up to 350. It does so using only $4n$ samples, a realistic assumption in practice. PICANTE’s performance is competitive with that of known state-of-the-art attacks on LWE, particularly when sufficient compute resources are available, as PICANTE’s novel data preprocessing step can be run in parallel on many cores.

Ethical considerations. Although PICANTE demonstrates significant progress towards attacking real-world LWE problems with sparse binary secrets, it cannot yet break problems with real-world-size parameters. In particular, the LWE schemes standardized by NIST use random secrets and are not vulnerable to the attacks presented here. Hence, we do not believe our paper raises any ethical concerns. Nonetheless, we will share a copy of the current paper with the NIST Cryptography group, to inform them of this progress.

Future directions. More work is needed to better understand the effect of the data preprocessing step, since we observe that we only need a 5% reduction of data entropy to succeed (Table 7). Additionally, there may be better ways to preprocess the data to improve transformer learning, which are less costly than using BKZ. In the future, the model training and secret recovery components of the attack could benefit from parallel runs across multiple GPUs, given our observation that different transformer initializations may result in different speeds of secret recovery (§ 6.3). Furthermore, improvements to transformer architecture and secret recovery methods may enable recovery of secrets with more complex parameter settings.

$n, \log_2 q$	h	PICANTE					uSVP attack with BKZ 2.0 and early-abort			
		β	success	Preprocessing CPU.hrs per matrix	# matrices	Training CPU.hrs per epoch	β	success	success time (CPU.hrs)	fail time (CPU.hrs)
80, 7	6, 7	20	4/10	0.01	34800	0.8	60	2/10	8, 12	7.8
		65	6/10				65	6/10	9, 10, 14, 18, 40, 85	72.1
	8, 9	20	2/40	0.01	34800	0.8	55	0/10	—	2.4
		60	1/10				60	1/10	12	6.9
150, 13	9, 10	20	6/10	3.1	14600	1.0	50	5/10	26, 30, 31, 35, 35	22.9
		55	8/10				55	8/10	19, 19, 23, 23, 23, 23, 28, 28	22.7
	11, 12	20	4/10	3.1	14600	1.0	50	2/10	22, 39	9.5
		55	4/10				55	4/10	14, 19, 24, 33	5.6
200, 17	18, 19	20	5/10	16	10800	1.6	45	6/10	12, 13, 18, 21, 21, 21	7.7
		20	3/10	16	10800	1.6	45	3/10	13, 21, 25	12.9
	20, 21	20	3/10	16	10800	1.6	45	3/10	13, 21, 25	12.9
		20	3/10	16	10800	1.6	45	3/10	13, 21, 25	12.9
256, 23	26, 27	18	5/10	52	8300	2.7	40	4/10	203, 221, 243, 265	189.1
		35	7/10	52	8300	2.7	35	7/10	238, 246, 249, 269, 284, 303, 348	241.9
	28, 29	18	4/10	52	8300	2.7	35	7/10	238, 246, 249, 269, 284, 303, 348	241.9
		35	7/10	52	8300	2.7	35	7/10	238, 246, 249, 269, 284, 303, 348	241.9
	30, 31	18	4/10	52	8300	2.7	35	5/10	231, 255, 263, 330, 336	171.7
		35	5/10	52	8300	2.7	35	5/10	231, 255, 263, 330, 336	171.7

Table 15. Concrete running times (CPU.hrs) for uSVP attacks and corresponding PICANTE costs. Training cost includes secret recovery time. PICANTE uses BKZ, the uSVP attack uses BKZ2.0 [15], see the discussion in § 4.1. In all uSVP attacks, we use $\omega = \text{round}(\sqrt{2}\sigma) = 4$. Legend: CPU.hrs: CPU hours. fail time: average time for the failed experiments.

References

- [1] AKKAYA, I., ANDRYCHOWICZ, M., CHOCIEJ, M., LITWIN, M., MCGREW, B., PETRON, A., PAINO, A., PLAPPERT, M., POWELL, G., RIBAS, R., ET AL. Solving rubik’s cube with a robot hand, 2019. <https://arxiv.org/abs/1910.07113>.
- [2] ALBRECHT, M., CHASE, M., CHEN, H., DING, J., GOLDWASSER, S., GORBUNOV, S., HALEVI, S., HOFFSTEIN, J., LAINE, K., AND LAUTER, K. E. A. Homomorphic encryption standard. In *Protecting Privacy through Homomorphic Encryption*. Springer, 2021, pp. 31–62. <https://eprint.iacr.org/2019/939>.
- [3] ALBRECHT, M. R. On Dual Lattice Attacks Against Small-Secret LWE and Parameter Choices in HELib and SEAL. In *Advances in Cryptology – EUROCRYPT 2017* (Cham, 2017), J.-S. Coron and J. B. Nielsen, Eds., Springer International Publishing, pp. 103–129. <https://eprint.iacr.org/2017/047>.
- [4] ALBRECHT, M. R., GÖPFERT, F., VIRIDIA, F., AND WUNDERER, T. Revisiting the expected cost of solving usvp and applications to lwe. In *Proc. of ASIACRYPT (2017)*, Springer, pp. 297–322. <https://eprint.iacr.org/2017/815.pdf>.
- [5] ALBRECHT, M. R., PLAYER, R., AND SCOTT, S. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology* 9, 3 (2015), 169–203. <https://eprint.iacr.org/2015/046>.
- [6] AVANZI, R., BOS, J., DUCAS, L., KILTZ, E., LEPOINT, T., LYUBASHEVSKY, V., SCHANCK, J. M., SCHWABE, P., SEILER, G., AND STEHLÉ, D. CRYSTALS-Kyber (version 3.02) – Submission to round 3 of the NIST post-quantum project. Available at <https://pq-crystals.org/>.
- [7] BAHDANAU, D., CHO, K., AND BENGIO, Y. Neural machine translation by jointly learning to align and translate. In *Proc. of ICLR* (2014).
- [8] BAI, S., MILLER, S., AND WEN, W. A refined analysis of the cost for solving LWE via uSVP. In *Proc. of ASIACRYPT (2019)*, Springer, pp. 181–205. <https://eprint.iacr.org/2019/502.pdf>.
- [9] BRAKERSKI, Z., LANGLOIS, A., PEIKERT, C., REGEV, O., AND STEHLÉ, D. Classical Hardness of Learning with Errors. In *Proc. of the Forty-Fifth Annual ACM Symposium on Theory of Computing* (2013). <https://arxiv.org/abs/1306.0281>.
- [10] CARION, N., MASSA, F., SYNNAEVE, G., USUNIER, N., KIRILLOV, A., AND ZAGORUYKO, S. End-to-end object detection with transformers, 2020. <https://arxiv.org/abs/2005.12872>.
- [11] CHARTON, F. Linear algebra with transformers, 2021. <https://arxiv.org/abs/2112.01898>.
- [12] CHEN, H., CHUA, L., LAUTER, K., AND SONG, Y. On the Concrete Security of LWE with Small Secret. *Cryptology ePrint Archive*, Paper 2020/539, 2020. <https://eprint.iacr.org/2020/539>.
- [13] CHEN, H., AND HAN, K. Homomorphic lower digits removal and improved fhe bootstrapping. In *Advances in Cryptology – EUROCRYPT 2018* (Cham, 2018), J. B. Nielsen and V. Rijmen, Eds., Springer International Publishing, pp. 315–337. <https://eprint.iacr.org/2018/067>.
- [14] CHEN, L., MOODY, D., LIU, Y.-K., ET AL. PQC Standardization Process: Announcing Four Candidates to be Standardized, Plus Fourth Round Candidates. *US Department of Commerce, NIST* (2022). <https://csrc.nist.gov/News/2022/pqc-candidates-to-be-standardized-and-round-4>.
- [15] CHEN, Y., AND NGUYEN, P. Q. BKZ 2.0: Better Lattice Security Estimates. In *Proc. of ASIACRYPT 2011* (2011).
- [16] CHEON, J. H., KIM, A., KIM, M., AND SONG, Y. Homomorphic encryption for arithmetic of approximate numbers. In *Proc. of ASIACRYPT (2017)*.
- [17] CHEON, J. H., KIM, D., LEE, J., AND SONG, Y. Lizard: Cut Off the Tail! A Practical Post-quantum Public-Key Encryption from LWE and LWR. In *Security and Cryptography for Networks* (Cham, 2018), D. Catalano and R. De Prisco, Eds., Springer International Publishing, pp. 160–177. <https://eprint.iacr.org/2016/1126.pdf>.
- [18] CHO, K., VAN MERRIENBOER, B., GULCEHRE, C., BAHDANAU, D., BOUGARES, F., SCHWENK, H., AND BENGIO, Y. Learning phrase

- representations using rnn encoder-decoder for statistical machine translation. In *Proc. of EMNLP* (2014).
- [19] CSORDÁS, R., IRIE, K., AND SCHMIDHUBER, J. The Neural Data Router: Adaptive Control Flow in Transformers Improves Systematic Generalization. In *Proc. of ICML* (2022).
- [20] CURTIS, B. R., AND PLAYER, R. On the feasibility and impact of standardising sparse-secret LWE parameter sets for homomorphic encryption. In *Proc. of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography* (2019). <https://eprint.iacr.org/2019/1148>.
- [21] DEGHANI, M., GOUWS, S., VINYALS, O., USZKOREIT, J., AND KAISER, L. Universal transformers. In *Proc. of ICLR* (2019).
- [22] DEVELOPMENT TEAM, T. F. fp1ll, a lattice reduction library, Version: 5.4.4. Available at <https://github.com/fp1ll/fp1ll>, 2023.
- [23] DONG, L., XU, S., AND XU, B. Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition. In *Proc. of ICASSP* (2018).
- [24] DOSOVITSKIY, A., BEYER, L., KOLESNIKOV, A., WEISSENBORN, D., ZHAI, X., UNTERTHINER, T., DEGHANI, M., MINDERER, M., HEIGOLD, G., GELLY, S., USZKOREIT, J., AND HOULSBY, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. of ICLR* (2021).
- [25] DUCAS, L., KILTZ, E., LEPOINT, T., LYUBASHEVSKY, V., SCHWABE, P., SEILER, G., AND STEHLÉ, D. CRYSTALS-Dilithium – Algorithm Specifications and Supporting Documentation (Version 3.1). Available at <https://pq-crystals.org/>.
- [26] KANNAN, R. Minkowski’s Convex Body Theorem and Integer Programming. *Mathematics of Operations Research* 12 (1987), 415–440.
- [27] KINGMA, D. P., AND BA, J. Adam: A method for stochastic optimization. In *Proc. of ICLR* (2015).
- [28] LAINE, K., AND LAUTER, K. Key recovery for lwe in polynomial time. *Cryptology ePrint Archive* (2015). <https://eprint.iacr.org/2015/176.pdf>.
- [29] LAMPLE, G., AND CHARTON, F. Deep learning for symbolic mathematics. In *Proc. of ICLR* (2020).
- [30] LENSTRA, H. J., LENSTRA, A., AND LOVÁSZ, L. Factoring polynomials with rational coefficients. *Mathematische Annalen* 261 (1982), 515–534.
- [31] LYUBASHEVSKY, V., AND MICCIANCIO, D. On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In *Proc. of Advances in Cryptology - CRYPTO 2009* (2009), S. Halevi, Ed. https://doi.org/10.1007/978-3-642-03356-8_34.
- [32] MICCIANCIO, D., AND VOULGARIS, P. Faster exponential time algorithms for the shortest vector problem. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms* (2010), p. 1468–1480.
- [33] PEIKERT, C. Public-Key Cryptosystems from the Worst-Case Shortest Vector Problem: Extended Abstract. In *Proc. of the Forty-First Annual ACM Symposium on Theory of Computing* (2009). <https://eprint.iacr.org/2008/481>.
- [34] POLU, S., AND SUTSKEVER, I. Generative language modeling for automated theorem proving, 2020. <https://arxiv.org/abs/2009.03393>.
- [35] RADFORD, A., NARASIMHAN, K., SALIMANS, T., AND SUTSKEVER, I. Improving language understanding by generative pre-training. *OpenAI blog* (2018). https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf.
- [36] RADFORD, A., WU, J., CHILD, R., LUAN, D., AMODEI, D., SUTSKEVER, I., ET AL. Language models are unsupervised multitask learners. *OpenAI blog* (2019). https://d4mucfpksywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.
- [37] RAMESH, A., PAVLOV, M., GOH, G., GRAY, S., VOSS, C., RADFORD, A., CHEN, M., AND SUTSKEVER, I. Zero-shot text-to-image generation, 2021. <https://arxiv.org/abs/2102.12092>.
- [38] REGEV, O. Quantum computation and lattice problems. *SIAM Journal on Computing* 33, 3 (2004), 738–760.
- [39] REGEV, O. On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In *Proc. of the Thirty-Seventh Annual ACM Symposium on Theory of Computing* (2005). <https://dblp.org/rec/journals/corr/cs-DS-0304005.bib>.
- [40] RIVEST, R. L., SHAMIR, A., AND ADLEMAN, L. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM* (1978).
- [41] ROGEZ, G., WEINZAEPFEL, P., AND SCHMID, C. Lcr-net: Localization-classification-regression for human pose. In *Proc. of CVPR* (2017).
- [42] ROTHE, R., TIMOFTE, R., AND VAN GOOL, L. Dex: Deep expectation of apparent age from a single image. In *Proc. of ICCV* (2015).
- [43] SCHNORR, C. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science* 53, 2 (1987), 201–224.
- [44] SCHNORR, C. P., AND EUCHNER, M. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical Programming* 66, 1-3 (Aug. 1994), 181–199.
- [45] SCHRITTWIESER, J., ANTONOGLIOU, I., HUBERT, T., SIMONYAN, K., SIFRE, L., SCHMITT, S., GUEZ, A., LOCKHART, E., HASSABIS, D., GRAEPEL, T., ET AL. Mastering ATARI, Go, Chess and Shogi by planning with a learned model. *Nature* 588, 7839 (2020), 604–609.
- [46] Microsoft SEAL (release 4.1). <https://github.com/Microsoft/SEAL>, Jan. 2023. Microsoft Research, Redmond, WA.
- [47] VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L., GOMEZ, A. N., KAISER, L., AND POLOSUKHIN, I. Attention is all you need. In *Proc. of NeurIPS* (2017).
- [48] WANG, Y., MOHAMED, A., LE, D., LIU, C., XIAO, A., MAHADEOKAR, J., HUANG, H., TJANDRA, A., ZHANG, X., ZHANG, F., AND ET AL. Transformer-based acoustic modeling for hybrid speech recognition. *Proc. of ICASSP* (2020).
- [49] WENGER, E., CHEN, M., CHARTON, F., AND LAUTER, K. Salsa: Attacking lattice cryptography with transformers, 2022. <https://arxiv.org/abs/2207.04785>.

A Appendix

A.1 Comparison of TINYLWE and LWE

Table 16 compares secret recovery performance of models trained using samples generated via PICANTE’s TinyLWE approach ($4n$ initial samples) vs. a baseline approach (2^{22} initial samples).

Setting	Hamming weight h							
$n = 80$	4	5	6	7	8	9	10	
TinyLWE	3/5	3/5	2/5	2/5	1/20	1/20	0/20	
LWE	5/5	4/5	3/5	3/5	0/20	1/20	0/20	
$n = 150$	9	10	11	12	13	14	15	
TinyLWE	4/5	2/5	3/5	1/5	1/20	0/20	0/20	
LWE	5/5	2/5	2/5	1/20	0/20	0/20	0/20	
$n = 200$	17	18	19	20	21	22	23	
TinyLWE	3/5	2/5	3/5	1/5	2/5	2/20	0/20	
LWE	3/5	2/5	1/5	1/5	2/20	0/20	0/20	
$n = 256$	26	27	28	29	30	31	32	
TinyLWE	4/5	1/5	1/5	3/5	3/5	4/20	0/20	
LWE	3/5	2/5	3/5	3/5	1/5	3/20	2/20	
$n = 300$	28	29	30	31	32	33	34	
TinyLWE	2/5	2/5	1/5	1/5	2/5	1/5	0/20	
LWE	1/5	3/5	2/5	1/5	1/5	0/5	0/20	
$n = 350$	20	21	22	23	24	25	26	
TinyLWE	2/5	4/5	2/5	2/5	3/20	1/20	0/20	
LWE	3/5	5/5	1/5	4/5	5/20	2/20	0/20	

Table 16. Secret recovery performance: TINYLWE vs. LWE. Reported values are successes/attempts across different n and h settings.

A.2 Statistical properties of secret verification

At the end of the secret recovery phase, we are provided a secret guess \mathbf{s}' , that we need to check. To do so, we use the original $m = 4n$ LWE samples (\mathbf{a}_i, b_i) , and compute the m residuals $r_i = b_i - \mathbf{a}_i \cdot \mathbf{s}'$. If the secret is recovered, we expect the r_i to have the same standard deviation as a LWE sample, i.e. σ . Otherwise, we expect the standard deviation to be that of the uniform distribution, i.e. $q/\sqrt{12}$.

The standard deviation of residuals is estimated by the classical formula:

$$\sigma_{\text{emp}} = \sqrt{\frac{1}{m-1} \sum_{i=0}^{m-1} (r_i - \bar{r})^2}$$

The lower and higher confidence intervals, with level $100(1 - \alpha)\%$ are

$$\sigma_{\text{emp}} \sqrt{\frac{m-1}{\chi_{\alpha/2, m-1}^2}}$$

and

$$\sigma_{\text{emp}} \sqrt{\frac{m-1}{\chi_{(1-\alpha)/2, m-1}^2}}$$

Since $m > 100$, we approximate the chi-square distribution with $m - 1$ degrees of freedom by the normal distribution $\mathcal{N}(m - 1, 2m - 2)$. Table 17 provides estimates of the confidence intervals at level 0.001% for different values of n , and around $\sigma = 3$ and $q/\sqrt{12}$.

n	m	Right ($\sigma = 3$)	Wrong ($q/\sqrt{12}$)
80	320	[2.58, 3.72]	[28.08, 40.45]
150	600	[2.68, 3.48]	$[1.65 \times 10^3, 2.15 \times 10^3]$
200	800	[2.71, 3.40]	$[3.42 \times 10^4, 4.28 \times 10^4]$
256	1024	[2.74, 3.34]	$[1.62 \times 10^6, 1.98 \times 10^6]$
300	1200	[2.76, 3.31]	$[2.50 \times 10^7, 3.00 \times 10^7]$
350	1400	[2.78, 3.29]	$[1.02 \times 10^9, 1.21 \times 10^9]$

Table 17. Confidence intervals (0.001%) for secret verification. Confidence level 0.001%. Right: the secret is correctly predicted ($\sigma = 3$). Wrong: the secret is incorrectly predicted ($\sigma = q/\sqrt{12}$). q from Table 2.

For instance, for $n = 80$, we have $m = 320$ and $q = 113$. The 0.001% level confidence interval for a correct secret prediction (i.e. measuring $\sigma = 3$) is [2.58, 3.72]. For an incorrect prediction (measuring $\sigma = q/\sqrt{12} = 32.62$), it is [28.08, 40.45]. Since the two intervals do not overlap, the sample size we use (m) is large enough to verify secret guesses (with quasi-certitude). As dimension increases, the confidence intervals are further apart. This proves our claim that the original $4n$ LWE samples are sufficient to verify model predictions.

A.3 Understanding secret recovery

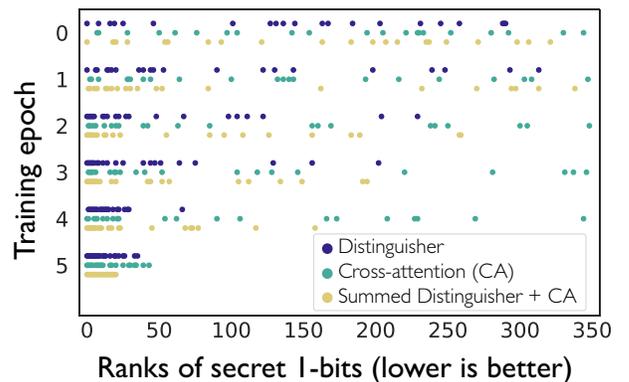


Figure 6. Change in secret bit ranking as model training progresses ($n = 350$ recovery).

We illustrate PICANTE’s secret recovery in Figure 6 for a successful $n = 350$ experiment, in which the combined method recovers the secret in epoch 5. Figure 6 shows how the

rankings of the 1-bits of the secret change throughout training. Our recovery methods guess as 1-bits the first h ranks, so a successful secret recovery occurs when all 1-bits appear in the first h slots along the x -axis.

Over time, distinguisher and CA methods learn better ranks for true secret 1-bits. By epoch 5, the combined method, which sums the ranks of distinguisher and CA methods, can correctly guess the secret. We do not include direct recovery results because of its poor performance for large n .

Plotting Figure 6 requires knowledge of the secret s , and we leverage this knowledge here for illustrative purposes only. We reiterate that PICANTE can validate secret guesses without the knowledge of the secret, using verification as in § 4.5.