AI for Code-based Cryptography

Mohamed Malhou[1,2], Ludovic Perret[3,2], and Kristin Lauter[1]

[1] FAIR, Meta
[2] Sorbonne Université, CNRS, LIP6, F-75005 Paris, France
[3] EPITA, EPITA Research Lab (LRE), Le Kremlin-Bicêtre, France

Abstract. We introduce the use of machine learning in the cryptanalysis of code-based cryptography. Our focus is on distinguishing problems related to the security of NIST round-4 McEliece-like cryptosystems, particularly for Goppa codes used in `ClassicMcEliece` and Quasi-Cyclic Moderate Density Parity-Check (QC-MDPC) codes used in `BIKE`. We present `DeepDistinguisher`, a new algorithm for distinguishing structured codes from random linear codes that uses a transformer. The results show that the new distinguisher achieves a high level of accuracy in distinguishing Goppa codes, suggesting that their structure may be more recognizable by AI models. Our approach outperforms traditional attacks in distinguishing Goppa codes in certain settings and does generalize to larger code lengths without further training using a puncturing technique. We also present the first distinguishing results dedicated to MDPC and QC-MDPC codes.

Keywords: Classic McEliece · Goppa Codes · QC-MDPC · Code Distinguishability · Deep Learning · Transformers

1 Introduction

In recent years, the cryptographic community has been actively preparing for the cyber-security challenges posed by cryptographic-relevant quantum computers. To address this quantum threat, the National Institute of Standards and Technology (NIST) has started a multi-stage standardization effort [11] to replace current number-theoretic-based cryptographic standards with a new generation of quantum-resistant algorithms. In 2024, NIST has released a first set of post-quantum cryptography (PQC) standards, including the Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM, [36]), the Module-Lattice-Based Digital Signature Algorithm (ML-DSA,[35]) and the Stateless Hash-Based Digital Signature Algorithm (SHB-DSA,[37]).

The standardization of post-quantum cryptography is still ongoing, with NIST currently conducting a fourth round of evaluations to identify additional key encapsulation mechanisms (KEMs) [38]. Candidates that remain in the fourth round all belong to code-based cryptography [8,39,19], a family based on the algorithmic and NP-hardness of decoding random linear codes [6,19]. In particular, two candidates follow the general framework of the McEliece cryptosystem [32]

: `ClassicMcEliece` [1] that uses binary Goppa codes as initially proposed by Robert McEliece in 1978, and `BIKE` [34,2] a variant relying on Quasi-Cyclic Moderate Density Parity-Check (QC-MDPC) codes.

A fundamental question for these schemes, and post-quantum cryptography in general, is the hardness of the underlying algorithmic problems. This issue is both critical and highly challenging, as the security of standardized and (surviving) candidate schemes for NIST has been intensively scrutinized in the past years. Introducing any new cryptanalytic technique can be considered as a notable achievement.

In this context, rcecent advancements in Machine Learning (ML) have introduced new paradigms for accelerating cryptanalysis. In particular, deep learning models – especially transformer-based architectures – have demonstrated remarkable success in pattern recognition, feature extraction, and automated discovery of hidden structures in high-dimensional data. In [47,29,30,48], the authors introduced the use of transformers to attack the Learning With Errors (LWE) problem [42], a central problem in post-quantum cryptography. The capability of these models to learn joint distributions of sequential data makes them a promising tool for identifying latent structures in algebraic constructions.

This paper presents a novel application of ML techniques to assess the security of code-based public-key cryptosystems. Specifically, we consider the problem of distinguishing structured public codes (e.g., Goppa or QC-MDPC) from random codes. To do so, we design a supervised learning framework on finite field data and introduce a transformer-based algorithm, `DeepDistinguisher`, designed to classify structured codes more effectively. Our work does not directly impact the security of Classic McEliece; however, we believe our findings will inspire researchers to further explore this problem. To our knowledge, this is the first application of ML techniques in code-based cryptography.

A fundamental limitation of `DeepDistinguisher`, and ML techniques in general, is the difficulty of explainability. We mitigate this issue by performing extensive experimental validation. Specifically, for Goppa codes, we empirically demonstrate that `DeepDistinguisher` can classify structured codes from random with high accuracy, even outperforming the most recent approaches [40,21,14] for some specific parameters.

1.1 Related works

AI, and more specifically ML is becoming a powerful approach in cryptanalysis, with a growing body of research demonstrating that neural networks can detect patterns. In post-quantum cryptography, a first generation of ML techniques have been used in [24] to attack group-based cryptosystems. More recently, the authors [47,29,30,48] leverage recent advances on ML – in particular the introduction of transformers [46] – for solving Learning With Errors (LWE) problems. Lastly, [27] considers learning-based information-theoretic metrics, leveraging mutual information estimation and binary classification to evaluate the security of cryptographic schemes under a chosen-plaintext attacks (`IND-CPA`). The study

demonstrates that neural networks can efficiently identify cryptosystems that are not `IND-CPA` by modeling the distinguishability of ciphertexts as a classification task.

In this paper, we present the first ML-based attack, `DeepDistinguisher`, on the distinguishing problem arising in the security of McEliece-like cryptosystems. In particular, we consider the *Goppa Code Distinguishing* (GD) problem [13]; probably the most famous example of code distinguishing problem.

*Problem 1 (Goppa Code Distinguishing (*GD*) problem). Given a generator matrix* $\mathbf{G} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ *of a* $[n,k]_q$ *linear code, the Goppa Code Distinguishing (*GD*) problem asks to decide if* $\mathbf{G}$ *is the generator matrix of a Goppa code or a randomly drawn matrix.*

The GD problem is related to the security of McEliece cryptosystems. It was formally introduced in [13] and was initially believed to be hard. Thus, it served as an assumption for reducing the security of the McEliece cryptosystem to the problem of decoding random linear codes [18]. Our understanding of the hardness of GD has significantly shifted in the past ten years, culminating with the so-called *Syzygy distinguisher* [40] that now solves GD for a broad range of parameters with a complexity asymptotically faster than the best generic decoding algorithms.

The Syzygy distinguisher, as well as improved results on GD such as [15,14], are built on the polynomial-time distinguisher presented by Faugère, Gautier, Otmani, Perret and Tillich (FGOPT) [21]. The core idea behind the FGOPT distinguisher is to analyze the behavior of the Gröbner basis computation [10,9] of an algebraic system associated to McEliece's public-key. This computation behaves differently if the algebraic system is generated from a McEliece's public-key or with a randomly generated matrix. FGOPT described specific linear relations occurring in such computations due to the Goppa (or alternant) structure, leading to a polynomial-time distinguisher solving GD for codes whose rate $R = \frac{k}{n}$ is close to 1. Since this result, a major open question has been how to extend the distinguishing rate.

At Asiacrypt'23, Couvreur, Mora and Tillich (CMT, [14]) finally demonstrated that the approach from [21] can be improved. CMT introduced a new algebraic modeling and leveraged more general *algebraic relations*, known as syzygies, arising in the Gröbner basis computations. Whilst FGOPT can distinguish codes with rates extremely close to 1, CMT pushed the boundaries of distinguishing rates in the range $[\frac{2}{3}, 1]$.

A central contribution of the latest distinguisher [40] is to precisely predict the syzygies arising at any step of the Gröbner computations. In particular, such distinguisher allows to solve GD for a broad range of parameters with a complexity sub-exponential in the error-correcting capability. Asymptotically, there is no more limitation on the rate for Syzygy distinguisher. For fixed parameters, the situation is different. Remark that, unlike FGOPT, the CMT and Syzygy distinguishers are not polynomial-time algorithms. The rate remains a limiting factor, and certain code parameters cannot be distinguished by either the Syzygy or CMT

distinguishers due to fundamental theoretical limitations and/or computational complexity constraints..

The code distinguishing problem for Quasi-Cyclic (QC) and general Moderate Density Parity-Check (MDPC) was formally introduced [34]. To the best of our knowledge, no dedicated technique exists for distinguishing such codes. The only known approach, described in [34], relies on finding low-weight codewords in the public code – a problem equivalent to message recovery. This led the authors of [34] to introduce an assumption about the (exponential) hardness of distinguishing MDPC codes.

We emphasize that the hardness of GD has, yet, no direct impact on the security of McEliece. That is, there is currently no generic technique allowing to mount an attack against McEliece using a distinguisher. However, recent results [5,14] demonstrated that the same techniques used for distinguishing alternant/Goppa codes can also be applied to attack a version of McEliece using generic alternant codes with high rates.

For these reasons, developing more efficient distinguishing techniques for structured codes remains a critical and pressing challenge. In this work, we advance in this direction by introducing a novel and flexible ML-based distinguisher, `DeepDistinguisher`, for alternant, Goppa, MDPC, and QC-MDPC codes. Additionally, we introduce a new problem – the Hidden Goppa Code (HGC) problem – , which seeks to recover the generator matrix $\mathbf{G} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ of a Goppa code given a masked version of $\mathbf{G}$. The HGC problem serves as an intermediate step between distinguishing a Goppa code and key-recovery. Although related to code detection/reconstruction problems [12,45], we believe this problem has not been explicitly formulated before.

1.2 Organization of the paper and main results

After this introduction, the paper is organized as follows. Section 2 provides the necessary background and definitions for understanding our work. Section 3 introduces a simple statistical distinguisher, highlighting that the entries of public generator matrices of binary Goppa codes are not exactly from a uniform distribution. However, this approach has strong limitations, motivating to consider advanced ML techniques.

Our distinguisher, `DeepDistinguisher`, is then detailed in Section 4 where we describe the training framework, data generation process, and evaluation strategies. We highlight the importance of the chosen deep learning architecture and the data representation. A key challenge in applying deep learning to this problem lies in the representation of finite field elements. Deep learning models operate on real or complex numbers, necessitating careful transformations to encode finite field elements effectively. Additionally, various strategies exist for processing and representing the input matrices, and we make specific design choices that enhance the efficiency and convergence of model training.

Section 5 presents our experimental results, demonstrating that our model achieves high classification accuracy, and outperforms traditional algebraic distinguishers

such as FGOPT [21], CMT [14], and Syzygy [40] on the toy parameter settings proposed in [40]. These algorithms could only achieve distinguishability of binary Goppa codes for a polynomial degree of $t = 3$ in the setting where the extension degree is $m = 6$ and the code length is maximal: $n = q^m = 64$. In this case, the lowest distinguishable code length was achieved by the Syzygy distinguisher [40] at $n = 50$. First, we significantly push this limit further, distinguishing codes of lengths as low as $n \geq 24$. Secondly, we extend the distinguishable range of t to $t = 2, 3, 4, \ldots, 9$. These settings serve as benchmark cases for evaluating distinguishing attacks. We further push the experimental limits by testing our model on codes of lengths $n = 128, 256, \ldots$ and applying puncturing techniques to evaluate even larger code lengths. Additionally, we present distinguishing results on ternary Goppa codes, and certain binary alternant codes, but also the first specific distinguishing results on (QC) MDPC codes. Although our experimental results are limited, they suggest that distinguishing these codes is easier than finding low-weight codewords, as stated in [34].

Finally, in Section 6, we introduce a more challenging problem: given a public generator matrix of a Goppa code with missing entries, recover the missing values such that the outcome is a valid Goppa code. This is a harder problem than distinguishing and seems impossible without the knowledge of some information about the private key. Our model successfully recovers these missing values, demonstrating that the structure of Goppa codes can be learned and exploited by AI.

2 Preliminaries

2.1 Notation

Finite fields. We consider the finite field $\mathbb{F}_q$ of order q with q a prime power. For some integer $m > 0, \mathbb{F}_{q^m}$ is the field extension of $\mathbb{F}_q$ of degree m: $\mathbb{F}_{q^m} \cong \mathbb{F}_q[x]/g(x) \cong \mathbb{F}_q[\alpha]$ with α a root of an irreducible polynomial $g(x)$ of degree m. Any element $\beta \in \mathbb{F}_{q^m}$ can be naturally associated with its vector form in $\mathbb{F}_q$ as $(c_0, \ldots, c_{m-1}) \in \mathbb{F}_q^m$, where $\beta = \sum_{i=0}^{m-1} c_i \alpha^i$.

Vectors and matrices. We use lowercase letters to represent integers, while integer intervals are expressed as $[\![a; b]\!]$. Matrices are denoted by bold uppercase letters, and vectors by bold lowercase letters. For a vector $\mathbf{v}$, the notation v_i refers to its i-th component, and $\mathbf{v}^\top$ denotes its transpose. $\mathcal{M}_{k\times n}(\mathbb{F})$ will denote the set of $k \times n$ matrices with coefficients over a finite field $\mathbb{F}$.

For a matrix $\mathbf{A} \in \mathcal{M}_{k\times n}(\mathbb{F})$, the element in the i-th row and j-th column is denoted by a_{ij}. A sub-matrix of $\mathbf{A}$, specified by a set of rows $\mathcal{I}$ and a set of columns $\mathcal{J}$, is written as $\mathbf{A}[\mathcal{I}, \mathcal{J}]$. Additionally, a specific row or column of a matrix $\mathbf{A}$ is indicated by $\mathbf{A}[i, :]$ and $\mathbf{A}[:, j]$, respectively.

2.2 Linear Codes and the Bounded Distance Decoding Problem

A $[k, n]_q$ *linear code* $\mathcal{C} \subseteq \mathbb{F}_q^n$ is a k-dimensional subspace in $\mathbb{F}_q^n$. The rate of $\mathcal{C}$ is defined as k/n and elements of $\mathcal{C}$ are called codewords. $\mathcal{C}$ can be specified by

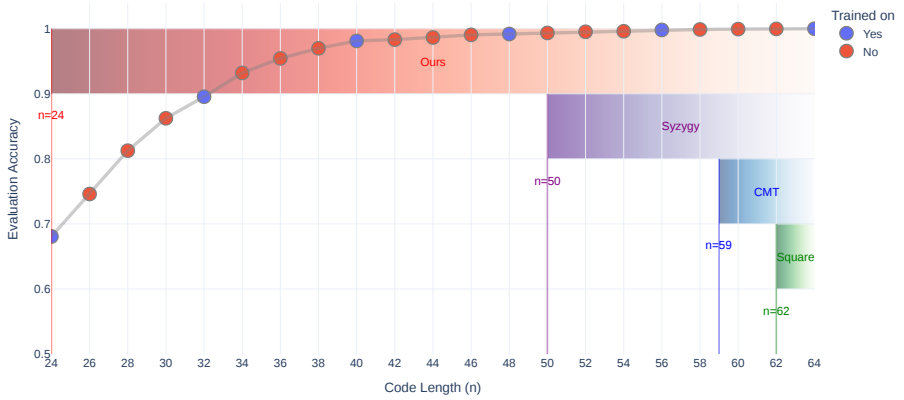

Fig. 1: Model accuracy as a function of code length. The model is trained on Binary Goppa Codes with extension degree $m = 6$ and irreducible polynomials of degree $t = 3$ as in [14] and [40]. The line and scatter points indicate the evaluation accuracy of our model on each tested value of code length. The scatter color indicates the range where the model was trained ($n = 24 + 8k$ for $k = 0, 1, ..$) showing that the model generalizes well to unseen input shapes. For attacks from the literature, we show the smallest code length reported in their respective papers, as no accuracy measures were provided.

a full rank *generator matrix* $\mathbf{G} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ such that $\mathcal{C} = \{\mathbf{m}\,\mathbf{G} \mid \mathbf{m} \in \mathbb{F}_q^k\}$. The *standard form* of $\mathbf{G}$ is $\mathbf{G} = [\mathbf{I}_k \mid \mathbf{A}]$, with $\mathbf{I}_k$ being the $k \times k$ identity matrix. Equivalently, $\mathcal{C}$ can be represented by *parity-check matrix* $\mathbf{H} \in \mathcal{M}_{(n-k)\times n}(\mathbb{F}_q)$ that satisfies $\mathbf{H}\,\mathbf{c}^\top = \mathbf{0}_{(n-k)}, \forall \mathbf{c} \in \mathcal{C}$, and its row space is the dual of $\mathcal{C}$.

We introduce below a general operation on codes that will be used to extend the range of applicability of `DeepDistinguisher`.

Definition 1 (Punctured Code). *Given a code $\mathcal{C} \subseteq \mathbb{F}^n$, and a subset $\mathcal{I} \subset [\![1; n]\!]$, the punctured code over $\mathcal{I}$ is defined as :*

$$\mathcal{P}_{\mathcal{I}}(\mathcal{C}) = \left\{ (c_i)_{i\in[\![1;n]\!]\setminus\mathcal{I}} \,\middle|\, \mathbf{c} \in \mathcal{C} \right\}.$$

Code-based cryptography [8,39,19] is based on the intractability, i.e. NP-Hardness, of the Bounded Distance Decoding (BDD, [6]) problem:

*Problem 2 (Bounded Distance Decoding (*BDD*) problem). Given the generator matrix $\mathbf{G} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ of a $[n, k]_q$ linear code, a codeword $\mathbf{c} \in \mathbb{F}_q^n$, and an integer $t > 0$, the* BDD *problem asks to find – if any – $\mathbf{m} \in \mathbb{F}_q^k$ such that:*

$$w_H(\mathbf{c} - \mathbf{m}\mathbf{G}) \leq t,$$

Table 1: Parameters for Code Distinguishers

Parameter	Description
$\mathscr{F}$	A family of codes
$\mathscr{G}, \mathscr{A}, \mathscr{R}$	Families of Goppa, alternant and random codes
$\mathscr{M}, \mathscr{Q}$	Families of MDPC and QC codes.
n	Code length
m	Extension field degree
t	Alternant order or Goppa polynomial degree
k	Code dimension ($k \geq n - mt$ for Goppa)
q	Base field order (prime power)
$\mathbb{F}_{q^m}$	Galois field of order q^m

with w_H being the Hamming weight of the vector, i.e. the number of its non-zero coordinates.

Solving BDD for random codes, i.e. random generator matrices $\mathbf{G} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$, is a long-standing problem whose most effective algorithms are all exponential [44,7,19].

2.3 McEliece Framework and Code Distinguishing Problem

The McEliece cryptosystem [32] is certainly the most popular code-based public-key cryptosystem. In particular, round-4 NIST candidates `ClassicMcEliece` [1] and `BIKE` [2] follow the general framework described below.

- **Secret-Key.** A structured generator matrix $\mathbf{G}_s \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ of a $[n,k]_q$ linear code with a known decoding algorithm.
- **Public-Key.** A scrambled generator matrix $\mathbf{G} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ derived from the secret-key $\mathbf{G}_s$.
- **Encryption.** Given a message $\mathbf{m} \in \mathbb{F}_q^k$, the ciphertext is computed as $\mathbf{c} = \mathbf{m} \cdot \mathbf{G} + \mathbf{e} \in \mathbb{F}_q^n$, where $\mathbf{e} \in \mathbb{F}_q^n$ is an error vector of small Hamming Weight.
- **Decryption.** Given a ciphertext $\mathbf{c} \in \mathbb{F}_q^n$, the receiver maps $\mathbf{c}$ to a noisy codeword on the secret code and applies the code's decoding algorithm to recover the message.

From this description, it is clear that the security of McEliece (message-recovery) relies on the hardness of BDD. In addition, it is natural to introduce a general distinguishability problem for a structured family of linear codes $\mathscr{F}$.

*Problem 3 (Code Distinguishability (*CD*) problem). Given a generator matrix $\mathbf{G} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ of a $[n,k]_q$ linear code, the* CD *problem asks to decide if $\mathbf{G}$ is the generator matrix of an $\mathscr{F}$-code or randomly drawn.*

In this paper, $\mathscr{F}$ includes Goppa or alternant $(\mathscr{G}, \mathscr{A})$ codes as well as MDPC and QC-MDPC $(\mathscr{M}, \mathscr{Q} \cap \mathscr{M})$ codes.

2.4 Alternant and Goppa Codes

The family of codes used in `ClassicMcEliece` can be conveniently described by introducing Generalized Reed-Solomon codes.

Definition 2 (Generalized Reed-Solomon Code, [41]). *Let $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n) \in \mathbb{F}_q^n$ be an n-tuple of distinct elements in $\mathbb{F}_q$, called a support, and $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_n) \in (\mathbb{F}_q^*)^n$ a n-tuple of nonzero elements in $\mathbb{F}_q$, called multiplier. The Generalized Reed-Solomon code of length n and dimension k, denoted by $\mathrm{GRS}_{n,k}(\boldsymbol{\alpha}, \boldsymbol{\beta})$, is defined as:*

$$\mathrm{GRS}_{q,n,k}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \{(\beta_1 f(\alpha_1), \ldots, \beta_n f(\alpha_n)) \mid f \in \mathbb{F}_q[x], \deg(f) < k\}.$$

Remark that the following weighted Vandermonde matrix is a generator matrix of the GRS code.

$$\mathbf{V}_t[\boldsymbol{\alpha}, \boldsymbol{\beta}] = \begin{pmatrix} \beta_1 & \beta_2 & \cdots & \beta_n \\ \beta_1\alpha_1 & \beta_2\alpha_2 & \cdots & \beta_n\alpha_n \\ \vdots & \vdots & \ddots & \vdots \\ \beta_1\alpha_1^{t-1} & \beta_2\alpha_2^{t-1} & \cdots & \beta_n\alpha_n^{t-1} \end{pmatrix}.$$

Alternant codes can be viewed as subfield subcodes of GRS codes.

Definition 3 (Alternant Code, [31]). *Let $\boldsymbol{\alpha} \in \mathbb{F}_{q^m}^n$ be a support and $\boldsymbol{\beta} \in (\mathbb{F}_{q^m}^*)^n$ be a multiplier as in Definition 2 such that $n \leq q^m$. The alternant code of degree t, denoted by $\mathcal{A}_t(\boldsymbol{\alpha}, \boldsymbol{\beta})$, is given by:*

$$\mathcal{A}_t(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \mathrm{GRS}_{q^m,n,t}(\boldsymbol{\alpha}, \boldsymbol{\beta})^{\perp} \cap \mathbb{F}_q^n.$$

$\mathcal{A}_t(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is $[n, k \geq n - mt]_q$ linear code.

Once the support and multipliers vectors are known, alternant codes of degree t can be decoded in polynomial-time up to errors with Hamming weight $t/2$ [31, Ch. 12]. McEliece cryptosystem relies on a sub-class of alternant codes.

Definition 4 (Goppa Code, [23]). *Let $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n) \in \mathbb{F}_{q^m}^n$ be a support and $g(x) \in \mathbb{F}_{q^m}[x]$ be a degree t irreducible polynomial, called a Goppa polynomial, such that $g(\alpha_i) \neq 0, \forall\ 1 \leq i \leq n$. The Goppa code, denoted $\mathcal{G}(\boldsymbol{\alpha}, g)$, is defined as follows:*

$$\mathcal{G}(\boldsymbol{\alpha}, g) = \mathcal{A}_t\left(\boldsymbol{\alpha}, \frac{1}{g(\boldsymbol{\alpha})}\right).$$

$\mathcal{G}(\boldsymbol{\alpha}, g)$ is a $[n, k \geq n - mt]_q$ linear code.

Goppa codes, viewed as alternant codes, naturally inherit a decoding algorithm that corrects up to $t/2$ errors. For Binary Goppa codes ($q = 2$), we can improve this bound to correct twice as many errors in polynomial-time.

2.5 Codes with Sparse Parity-Check Matrices

BIKE [34,2] relies on linear codes described by compact and sparse matrices.

Definition 5 (Moderate Density Parity-Check codes, [34]). *An* (n, k, w)-MDPC *code is a linear code of length* n*, codimension* k *admitting a parity check matrix with constant row weight* w *which scales in* $O(\sqrt{n \log n})$*.*

BIKE adds a structure to MDPC codes allowing to decrease the size of the public-key.

Definition 6 (Quasi-cyclic codes, [34]). *An* $[n, k]_q$*-linear code is Quasi-Cyclic (*QC*) if there is some integer* n_0 *such that every cyclic shift of a codeword by* n_0 *places is again a codeword.*

2.6 Solving the Code Distinguishing Problem

A well-studied example of a code distinguishing problem occurs when the family $\mathscr{F}$ is restricted to Goppa or alternant codes (Section 2.4). This corresponds to the classical McEliece scheme and the Goppa Code Distinguishing (GD) problem.

The first efficient algorithm for solving this problem, FGOPT [21], relies critically on the code rate k/n. In [21,14], the authors precisely characterize the range of parameters for which FGOPT can distinguish Goppa codes in polynomial time.

Definition 7 (Square–distinguishable Goppa code). *A Goppa code* $\mathcal{G}(\boldsymbol{\alpha}, g)$*, with* $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n) \in \mathbb{F}_{q^m}$ *a support and* $g(x) \in \mathbb{F}_{q^m}[x]$ *a Goppa polynomial of degree* t*, is said to be square–distinguishable if:*

$$n > \binom{tm+1}{2} - \frac{m}{2}(t-1)(t-2), \qquad \text{when } \ t < q-1 \tag{1}$$

$$n > \binom{tm+1}{2} - \frac{m}{2} t \left((2e_{\mathcal{G}}+1)t - 2(q-1)q^{e_{\mathcal{G}}-1} - 1\right), \ \text{otherwise,} \tag{2}$$

where $e_{\mathcal{G}} = \min\{i \in \mathbb{N} \mid t \leq (q-1)^2 q^i\} + 1 = \left\lceil \log_q\left(\frac{t}{(q-1)^2}\right) \right\rceil + 1.$

Note that similar results can be derived for Alternant or Binary Goppa codes.

In [14], Couvreur, Mora and Tillich (CMT) extended the concept of distinguishable codes by introducing a new class called d-distinguishable codes. This concept is based on invariants related to the Hilbert function, a fundamental tool from commutative algebra [16], commonly used to assess the complexity of Gröbner basis computations [10,9]. In particular, it applies to Pfaffian ideals, i.e. ideals generated by symbolic minors of skew-matrices [33,20], modeling specific relations of alternant and Goppa codes.

Definition 8 (d-distinguishable, simplified from [14]). *Let* $\mathcal{C}$ *be a* $[n, tm]_{\mathbb{F}_{q^m}}$ *linear code,* $\mathcal{P}_2^+(\mathcal{C})$ *be the Pfaffian ideal associated to* $\mathcal{C}$ *[14, Sec. 5.2] and* $\mathrm{HF}_{\mathcal{P}_2^+(\mathcal{C})}$

be the corresponding Hilbert function. $\mathcal{C}$ is said to be d-distinguishable from a generic $[n, tm]$ linear code over $\mathbb{F}_{q^m}$ when the following holds:

$$\mathrm{HF}_{\mathcal{P}_2^+(\mathcal{C})}(d) \neq \max\left(0, \sum_{i=0}^{d} \frac{(-1)^i}{tm+d-i-1}\binom{n-tm}{i}\binom{tm+d-i-1}{d-i+1}\binom{tm+d-i-1}{d-i}\right).$$

1-distinguishable codes correspond to square–distinguishable Goppa codes (Definition 7). In [14], the authors demonstrated that d-distinguishability, for $d > 1$, allows to distinguish a broader family of codes than FGOPT, albeit at a higher computational cost. In particular, the complexity of the CMT distinguisher is bounded from above by:

$$\mathcal{O}\left(\left(\binom{tm}{2} - k + 1\right)\binom{\binom{tm}{2} + d_{\mathrm{reg}} - 1}{d_{\mathrm{reg}}}^{\omega}\right), \tag{3}$$

where $2 \leq \omega < 3$ is a feasible linear algebra constant, and d_{reg} is the degree of regularity [4], i.e. the maximum degree reached in the computation of (degree-based) Gröbner basis of the Pfaffian ideal $\mathcal{P}_2^+(\mathcal{C})$. In [14], the authors conjecture that d_{reg} behaves asymptotically as $d_{\mathrm{reg}} \sim c\frac{(tm)^2}{n-tm}$, where c is a constant close to $\frac{1}{4}$. These lead to a new algorithm that distinguishes codes with a rate in the range $[2/3, 1]$. Its complexity interpolates between polynomial-time (square–distinguishable Goppa codes) and super-exponential for constant rates.

The syzygy distinguisher [40] includes and extends previous results. It refines the algebraic modeling from CMT and conducts a more precise analysis of the syzygies occurring during a Gröbner computation. The dimension of these syzygies are related to so-called Betti numbers that depend on the structure of the code considered. These allow the authors to present a new distinguisher that is asymptomatically independent of the rate. Its complexity is bounded from above by

$$\kappa = q^{\left(\omega\frac{R^2}{1-R}+o(1)\right)\frac{(\log_q \log_q(n))^3}{(\log_q(n))^2}n},$$

where R is the rate of the dual code $R = mt/n$, and ω is the linear algebra exponent. The algorithm is not polynomial-time, but remains sub-exponential in the error-correcting capacity.

2.7 Basics of deep learning

Before presenting our approach, we first introduce the fundamental concepts of deep learning to provide the necessary background for a clear understanding of our methodology [22].

A *deep neural network* is a parametric family of functions $F_\theta \colon \mathcal{X} \to \mathcal{Y}$, where $\theta \in \mathbb{R}^p$ represents all trainable parameters (i.e., the entries of weight matrices and bias vectors) [22] and $\mathcal{X}$ and $\mathcal{Y}$ are measurable spaces. Concretely, for an input $x \in \mathcal{X}$, one may write:

$$F_\theta \colon x \;\mapsto\; W_d\,\sigma\big(\cdots\sigma(W_1 x + b_1)\cdots\big) + b_d,$$

with each W_i a weight matrix, b_i a bias vector, and σ a fixed nonlinearity such as ReLU (Rectified Linear Unit $x \to \max(0, x)$) applied component-wise. In this case, the trainable parameters of F are $\theta = \{W_1, b_1, \ldots, W_d, b_d\}$.

Training amounts to minimizing an empirical risk

$$\min_{\theta} \sum_{i} \ell\big(F_\theta;\, x_i,\, y_i\big),$$

where $\{(x_i, y_i)\}$ is a labeled dataset and ℓ is a chosen loss function. A common case is *binary classification*, where $\mathcal{Y} = \{0, 1\}$ and one trains F_θ to output probabilities in $[0, 1]$ by minimizing the *binary cross-entropy* loss [22]:

$$\sum_{i} \Big[- y_i \log\big(F_\theta(x_i)\big) \;-\; (1 - y_i) \log\big(1 - F_\theta(x_i)\big) \Big].$$

Parameters θ are typically updated via gradient-based algorithms (e.g., stochastic gradient descent) that converge to a θ^* that produces accurate predictions on new (held-out/validation) data. By the *universal approximation* theorem, sufficiently large networks can approximate wide classes of continuous functions on compact domains [17,26].

A *Transformer encoder* [46] is a deep, sequence-to-sequence map that takes an ordered collection $\{x_1, \ldots, x_n\} \subseteq \mathcal{X}$—which may be text tokens, image patches, matrix rows, etc.—and outputs a sequence of embeddings $H_L \in \mathbb{R}^{n \times d}$. Each x_i is first embedded (or projected) into $h_{0,i} \in \mathbb{R}^d$. Then, each of the L layers applies *multi-head self-attention* and a small *feed-forward* sub-network, with residual connections and normalization. Formally, for layer ℓ, we form affine queries, keys, and values from $H_{\ell-1}$, compute

$$\mathrm{Att}(Q, K, V) = \mathrm{softmax}\Big(\tfrac{QK^\top}{\sqrt{d_k}}\Big)\, V$$

for each attention head, concatenate the head outputs, and project them. We update

$$H'_\ell = \mathrm{LayerNorm}\big(H_{\ell-1} + \mathrm{MultiHeadAtt}(H_{\ell-1})\big), \quad H_\ell = \mathrm{LayerNorm}\big(H'_\ell + \mathrm{FFN}(H'_\ell)\big).$$

All parameters (in the attention and feed-forward blocks) are trained end-to-end via gradient descent to yield a final encoder representation H_L. If the task is a classification, then we project the hidden state H_L into the output space using a trainable linear layer [46,3].

3 A Simple Statistical Hamming Weight Distinguisher

In this section, we present a simple statistical distinguisher based on an experimental observation about the distribution of generator matrices of Goppa codes. Computing the total Hamming weight of these matrices allows us to distinguish Goppa codes from *uniform* random codes, although with weaker performance compared to other methods.

We now fix the parameters under consideration, namely the code length n, extension degree m, Goppa degree t, and the dimension $k = n - mt$, and we consider random codes with the same dimension. Given a generator matrix $\mathbf{G} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ from a code family $\mathscr{F} \in \{\mathscr{G}, \mathscr{R}\}$ (for Goppa versus Random), we define its total Hamming weight as

$$w_H(\mathbf{G}) = \sum_{i=1}^{k} w_H(\mathbf{g}_i),$$

where $\mathbf{g}_i$ is the i-th row of $\mathbf{G}$. The variable $w_H(\mathbf{G})|_{\mathscr{F}}$ is a random variable that takes discrete values, with the distribution denoted as $p_{\mathscr{F}}$. To check for any differences in the behavior of $w_H(\mathbf{G})|_{\mathscr{F}}$ between the two distributions, we empirically estimate the total variation distance between $f_{\mathscr{G}}$ and $f_{\mathscr{R}}$ using $1M$ samples:

$$D_{TV}(\hat{f}_{\mathscr{G}}, \hat{f}_{\mathscr{R}}) = \frac{1}{2}\sum_{x\in\mathcal{X}} |\hat{f}_{\mathscr{G}}(x) - \hat{f}_{\mathscr{R}}(x)|.$$

The empirical distributions $\hat{f}_{\mathscr{G}}$ and $\hat{f}_{\mathscr{R}}$ are shown as histogram plots in the appendix fig. 6 for codes of length $n = 64$ and extension degree $m = 6$. We can see that the distributions don't match especially for small values of t and exceptionally for $t = 9$ when $n = 64$ and $m = 6$ which is not the case for larger codes. We show this by also varying the extension degree and plotting the empirical TV distance in fig. 2 as a function of m and t. The distance is significant for small values of t but exponentially decreases.

To devise a statistical distinguisher, we can employ a hypothesis test based on the likelihood ratio. In the random case, the distribution of the total Hamming weight is known; a binomial distribution since it's a sum of independent Bernoulli variables with $p_{\mathscr{R}} = 0.5, N_{\mathscr{R}} = mt(n - mt)$ (only codes in standard form are considered). For the Goppa case, we make an assumption that the distribution is also a binomial and empirically estimate $\hat{p}_{\mathscr{G}}$ and $\hat{N}_{\mathscr{G}}$ using some training data.

Given a sample $x = w_H(\mathbf{G})$, we compute the likelihood ratio:

$$\log \Lambda(x) = \log \frac{\binom{\hat{N}_{\mathscr{G}}}{x}}{\binom{N_{\mathscr{R}}}{x}} + x \log \frac{\hat{p}_{\mathscr{G}}}{p_{\mathscr{R}}} + \log \frac{(1-\hat{p}_{\mathscr{G}})^{\hat{N}_{\mathscr{G}}-x}}{(1-p_{\mathscr{R}})^{N_{\mathscr{R}}-x}}$$

The distinguisher can therefore be expressed as follows:

$$\mathscr{D}(\mathbf{G}) = \mathbb{1}_{\{\log \Lambda(x) > \tau\}}(w_H(\mathbf{G}))$$

Using this simple distinguisher with $\tau = 0$, we can achieve a test accuracy of 73% on a balanced 1M dataset with the parameters of the first graph ($t = 2, n = 64, m = 6$) of fig. 6 and 62% accuracy on $t = 3$ and only 57% for $t = 4$.

It's worth noting that by changing the sampling distribution over random linear codes from uniformly random to $\mathcal{B}(p)^{k\times(n-k)}$ (independent Bernoulli entries) for p matching the experimental value for Goppa codes of the same parameters,

this distinguisher will degrade. Interestingly enough, training a linear model such as the logistic regression to distinguish these codes will automatically find this Hamming weight distinguisher as its best solution, this means that a more complex/deep architecture is needed to capture a non linear boundary between the two classes. So let's consider a much more powerful approach.

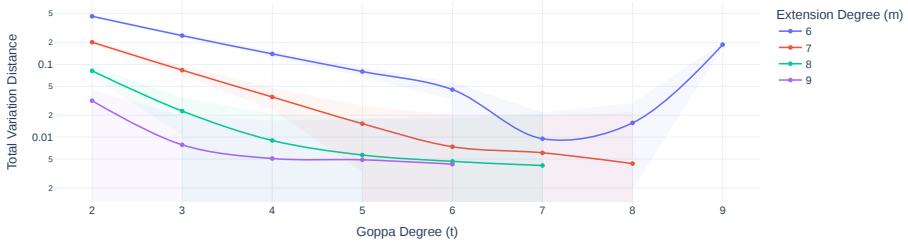

Fig. 2: Total variation distance of empirical distributions $p_{\mathscr{G}}$ (Goppa) and $p_{\mathscr{R}}$ (Random) of the total Hamming weight metric for binary codes of length $n = 64$ with varying extension degree m and polynomial degree t. The distance axis is in log scale showing the distance exponentially converging to 0. The plateauing might be because of the estimation error being high. 1M samples were used for probability estimation.

4 A Transformer-Based Algorithm for Code Distinguishing

In this section, we introduce a novel and natural method for code distinguishing based on deep learning. The motivation behind this approach is that a deep learning model, trained to classify samples from different families of codes, can potentially identify patterns revealed by a public generator matrix. Unlike classical approaches that rely on predefined heuristics or algebraic properties, a deep learning model can adaptively discover hidden structural differences between code families if any.

4.1 Deep Distinguisher

Let $\mathscr{F} \in \{\mathscr{G}, \mathscr{A}, \mathscr{M}, \mathscr{Q} \cap \mathscr{M}\}$ denote the code family of interest (Goppa, Alternant, MDPC, and QC-MDPC codes), we train a deep learning model whose input is a matrix of shape $k \times n$ to output a scalar that represents the probability that the input spans a code from $\mathscr{F}$. Similarly to [43], our approach leverages an encoder-only Transformer model, which has demonstrated strong performance compared to other models we tried, namely, logistic regression, neural nets and convolutional networks. Neural networks of the same size as our model do achieve

some good accuracies after some period of training, but we find that they are very slow at learning and rarely find the best solution that yields 100% accuracy.

The model processes an input sequence of vectors through an embedding layer, followed by four layers consisting of *self-attention* and *feed-forward networks*. Each vector in the sequence has an embedding dimension of $d_{\text{emb}} = 1024$.

In the self-attention mechanism, the model employs multi-head attention mechanism with $h = 4$ heads, where each head operates on a subspace of dimension $d_{\text{head}} = \frac{d_{\text{emb}}}{h} = 256$. The input vectors are first projected into these lower-dimensional subspaces, processed independently by each head, and then recombined to restore the original embedding dimension.

The feed-forward network (FFN) in each block consists of two linear transformations with a *GELU* [25] non-linearity in between. It first expands the dimension to $4 \times d_{\text{emb}} = 4096$ using a fully connected layer, applies the activation function, and then projects the vectors back to the original embedding dimension d_{emb}.

After processing through these layers, the final sequence representation is obtained via *max-pooling* over the sequence length. The resulting pooled vector is then linearly projected into a scalar, which serves as the model's *logit* and is used in the loss function for optimization. We use Adam optimizer [28] with warmup ≈ 1000 steps and set the learning rate to $lr = 10^{-5}$ and weight decay to $\omega = 10^{-3}$. We use a binary cross-entropy loss function to optimize the model during training which is basically maximizing the likelihood of the training batches. To evaluate the model's performance, we measure accuracy and precision on a separate balanced test set.

Data Representation. This is a crucial factor in achieving our distinguishing results. In fact, given that the input is a standard form matrix over a finite field, multiple encoding strategies are possible, including flattening the matrix, patching, or tokenizing field elements. However, the most effective approach is as follows: we bypass the need for a tokenizer and treat each input matrix as a sequence of rows that form a basis of the code, where each row serves as a 'token' input to the Transformer. The embedding representation of each row is simply a trainable linear projection of the row itself to a larger embedding space after lifting the finite field entries to $\mathbb{R}$ using the encoding guidelines described below. We add an absolute positional encoding on the sequence level.

When the base field is not $\mathbb{F}_2$, we encode the field elements differently based on the value of q. If q is prime, we use angular embedding as in [43], which doubles the dimension of the rows. Otherwise, if q is a prime power, we first represent the elements as vectors of polynomial coefficients, then apply the appropriate encoding based on the prime base field.

Example: The field $\mathbb{F}_9$ can be constructed as an extension of the base field $\mathbb{F}_3$ using the irreducible polynomial $x^2 + 1$. Let z be one of its roots. $\mathbb{F}_9$ elements are expressed as $a + bz$ with $a, b \in \mathbb{F}_3$. Therefore, we represent these elements as vectors (a, b). Now to encode $\mathbb{F}_3$ elements, we use angular embedding, resulting in a 4 dimensional vector $(\cos(2\pi\frac{a}{3}), \sin(2\pi\frac{a}{3}), \cos(2\pi i\frac{b}{3}), \sin(2\pi i\frac{b}{3})) \in \mathbb{R}^4$.

4.2 Datasets

We consider Goppa and Alternant codes with a fixed code length n. For a given set of parameters — extension degree m and degree $t \in \mathbb{N}$ — we generate a dataset $\mathcal{D}_{\mathscr{F}}$ uniformly from the family $\mathscr{F}$ of codes, retaining only the codes of rank $k = n - mt$. Each generator matrix is computed in systematic form. Additionally, we generate a dataset $\mathcal{D}_{\mathscr{R}}$ by uniformly sampling random linear codes with the same parameters and size, following the same procedure. We define $\mathcal{D} = \mathcal{D}_{\mathscr{F}} \cup \mathcal{D}_{\mathscr{R}}$, and use the notation $\mathcal{D}[q, n, m, t]$ to explicitly specify the parameters when needed. It is important to note that $\mathscr{F} \subset \mathscr{R}$, meaning that, when generating the dataset for random linear codes, there is a non-zero probability that some samples might belong to $\mathscr{F}$ codes (e.g., Goppa or alternant codes). However, this probability is negligible due to the structure of the family $\mathscr{F}$ and the comparatively vast size of $\mathscr{R}$. As a result, its impact on the dataset is statistically insignificant for our analysis.

4.3 Out-of-distribution Evaluation: Punctured Codes

To evaluate the distinguisher on instances of larger code lengths, we puncture the code by truncating the public generator matrix to fit into the training shape and assess the model on the resulting code. Initially, this approach does not yield satisfactory results when applied just once. However, by repeatedly truncating the original matrix through subsampling of rows and columns, and evaluating the model on each of these subsampled matrices, we can improve performance. By aggregating the results, we determine an optimal decision threshold based on the number of positive classifications or vote counts.

Algorithm 1 Sample Punctured SubCode

procedure SamplePuncturedSubCode($\mathbf{G}$)
  Input: Generator matrix $\mathbf{G} = \left(\mathbf{I}_{k_0} \mid \mathbf{A}\right) \in \mathbb{F}_q^{k_0 \times n_0}$
  Output: Punctured subcode matrix $\mathbf{G}[\mathcal{I}, \mathcal{J}]$
  Sample valid (n, k) such that $k \leq k_0$ and $n - k \leq n_0 - k_0$
  Sample k row indices $\mathcal{I}$ from $[\![0; k_0 - 1]\!]$ without replacement.
  Sample $n - k$ column indices $\mathcal{J}'$ from $[\![k_0; n_0 - 1]\!]$
  Form the set of column indices $\mathcal{J} = \mathcal{J}' \cup \mathcal{I}$
  return submatrix $\mathbf{G}[\mathcal{I}, \mathcal{J}]$
end procedure

More formally, given $\mathbf{G} = \left(\mathbf{I}_{k_0}\ \mathbf{A}\right) \in \mathbb{F}_q^{k_0 \times n_0}$ a standard form generator matrix of a linear code $\mathcal{C} \in \mathscr{F}$, to evaluate the model on punctured subcodes of $\mathcal{C}$, we first sample one of the training code parameters (k, n) such that $k \leq k_0$ & $n - k \leq n_0 - k_0$. Then sample k row indices $\mathcal{I}$ and $n - k$ column indices $\mathcal{J}'$ from the matrix $\mathbf{A}$ in addition to the k columns forming the identity matrix of shape k: $\mathcal{J} = \mathcal{J}' \cup \mathcal{I}$. This is important because during training, the model only sees

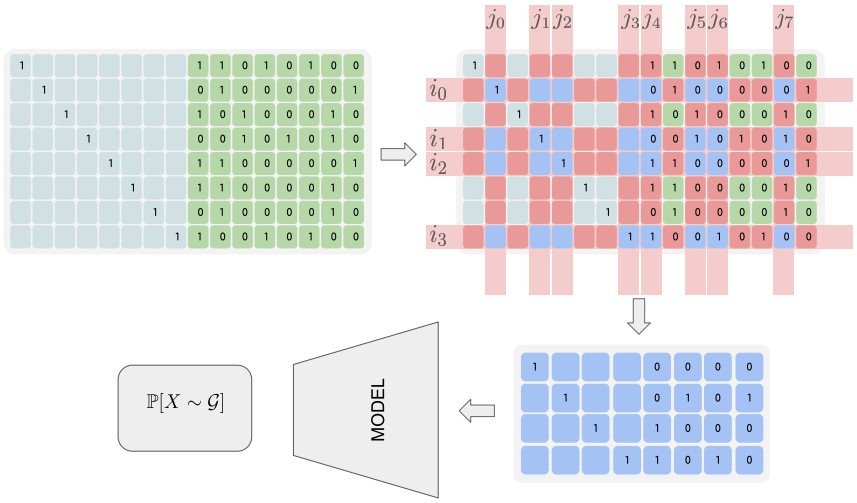

Fig. 3: To assess an $[8, 4]$-model (trained on codes of length $n = 8$ and dimension 4) on a $[16, 8]$-code, we puncture the input code by sampling $i_0 = j_0, \ldots, i_3 = j_3$ and $j_4, \ldots, j_7$ randomly to create a new $[8, 4]$-code in standard form.

standard form matrices so we don't expect it to generalize to unseen input in those first columns. Therefore, the identity matrix acts as a positional encoding of the sequence. Finally, we assess the model on $\mathbf{G}[\mathcal{I}, \mathcal{J}]$.

Algorithm 2 Evaluate distinguisher on larger codes using puncturing.

Require: Generator matrix $\mathbf{G} = \left(\mathbf{I}_{k_0} \mid \mathbf{A}\right) \in \mathbb{F}_q^{k_0 \times n_0}$, number of trials m
Ensure: Aggregated result of model evaluations
1: **function** EvaluateModel($\mathbf{G}, k_0, n_0, m$)
2: Initialize result_sum $\leftarrow 0$
3: **for** $i = 1$ to m **do**
4: $\mathbf{G}_{\text{punc}} \leftarrow$ SamplePuncturedSubCode($\mathbf{G}, k_0, n_0$)
5: result $\leftarrow$ EvaluateDistinguisher($\mathbf{G}_{\text{punc}}$)
6: result_sum $\leftarrow$ result_sum + result
7: **end for**
8: **return** result_sum
9: **end function**

5 Experiments and Results

In this part, we present the experimental results of the `DeepDistinguisher` (Section 4) on alternant/Goppa codes (Section 2.4) and MDPC/QC-MDPC

codes (Section 2.5). In the former case, we follow the methodology introduced in [40] to derive the parameters q, t, m and n (the code dimension is computed as $k = n - mt$). The approach is as follows:

- First, we fix the field size q and the extension degree m. We set the length as $n = q^m$ (full support) and find the largest t for which the code can be distinguished.
- Once such t is identified, we fix its value as well as the corresponding q and m. We then search for the smallest value of n that is still distinguishable.

Our experimental results for `DeepDistinguisher` are presented in two parts. In section 5.1, we analyze a specific set of parameters introduced in [14,40]. The code considered is relatively small (length at most 64). However, this allows explicit comparison of different distinguishers on a common benchmark. In Section 5.2, we present more extensive results; pushing the practical experiments to code of length up to 1024. We conclude this part by providing experimental results for the codes underlying `BIKE`; demonstrating the flexibility of the `DeepDistinguisher` distinguisher (Section 5.3).

5.1 Comparing Distinguishers for Goppa on a Small Benchmark

In [14,40], the authors presented experimental results for their distinguishers on a Binary Goppa code with $q = 2, m = 6$ and $t = 3$. The maximal length is $n = 64$, the FGOPT distinguisher will be able to distinguish up to $n_{\text{square}} = 62$. The CMT distinguisher [14] reported $n_{\text{CMT}} = 59$ and [40] brings down to $n_{\text{syzygy}} = 50$. Below this length, the conditions for distinguishability from [40] are not verified anymore.

As highlighted in Figures 1 and 4, our distinguisher works for any code length tested, with nearly 100% accuracy for most values of $n \geq 40$. We set a lower bound on the values of n such that the rate is no less than 0.20, which, for instance, corresponds to $n > 22$ when $t = 3$. The accuracy tends to drop for very small code rates.

In this case of $q = 2, m = 6$, our distinguisher works for all values of $t \in [\![2 : 9]\!]$, and achieves perfect accuracy when $t \leq 6$. Figure 4 shows a heatmap of our model's accuracy across different values of n and t. This visualization provides a comprehensive overview of how accuracy varies with these parameters, serving as a benchmark for further investigations and comparisons with other approaches and future works.

Inference Complexity. Since we are using a standard model size throughout our work, we can give an estimate of the time complexity of our distinguisher. In fact, the complexity of this distinguisher is determined by its inference time, which corresponds to performing a forward pass through the trained model times the number of calls to the model which is usually once. This cost is proportional to the model size (with at most $\leq 50M$ parameters) and scales polynomially with the input parameters k, n. In practice, calling our model takes about 10 milliseconds (ms) on CPU (or $100ms$ in one CPU thread) and less than $\approx 1ms$

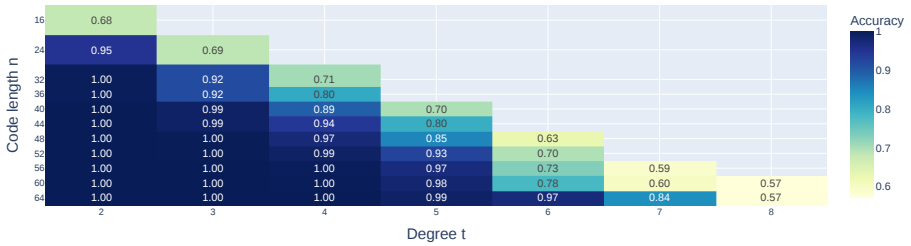

Fig. 4: Heatmap of model accuracy for $q = 2, m = 6$ as a function of code length n and degree parameter t. In this experiment, one model is trained per value of t only. Meaning the model is trained on codes with varying lengths n, allowing the model to interpolate well to unseen code lengths.

on a GPU. While training requires several hours, it is a one-time, offline process whose cost will be amortized.

5.2 Distinguishing Goppa and Alternant Codes

Goppa codes demonstrate distinguishability across a range of parameters with some specific configurations that achieve perfect accuracy. Binary codes of length $n = 128$ and extension degree $m = 7$ can be distinguished up to polynomial degree $t = 8$, with 100% accuracy for $t \leq 4$ as shown in table 3. For larger codes such as $n = 512, m = 9$, the model - of the same size as the 128-model - distinguishes codes up to $t = 4$.

In general, we observe that the accuracy of the model is lower as the degree of the Goppa polynomial t and the extension degree m increase. Experiments show that alternant codes are harder to distinguish from random codes, achieving accuracy better than random only when $t \leq 3$, for codes of length $n = 64$ and extension degree $m = 6$, as demonstrated in Table 2.

Goppa Codes with $q = 3, m = 4, n = 64$. When considering *ternary* Goppa codes with $m = 4$, we observe that the distinguishing task is more challenging compared to the binary case. Nevertheless, our distinguisher remains effective up to degree $t = 6$ (corresponding to a code rate of $R = 0.63$) as shown in Table 2. As the degree increases, accuracy drops considerably. For $t = 4$, the accuracy decreases to 90.34%, and for $t = 5$, it drops sharply to 54.71%, indicating that distinguishing becomes significantly harder.

Goppa Codes with $q = 2, m = 7$. We train the distinguisher on a 12M dataset of binary codes of extension degree $m = 7$ while varying the code lengths and degrees t. But first, we train on maximal code length $n = 128$ to figure out

Table 2: Distinguishing Results for Goppa/Alternant Codes with $n = 64$. In this experiment, for each setting $[q, n, m, t]$, the model is trained on a dataset $\mathcal{D}[q, n, m, t]$ of total size $\leq 40m$ samples and evaluated on $10k$ unseen samples.

Code	(q, m)	Degree (t)	Rate (R)	Accuracy (%)	Training Steps
Goppa	$(2, 6)$	2	0.81	99.12	1K
		3	0.72	98.88	8.5K
		4	0.63	98.52	22K
		5	0.53	98.24	48.5K
		6	0.44	96.68	243K
		7	0.34	84.60	848.5K
		8	0.25	57.42	90K
		9	0.16	75.92	296.5K
	$(3, 4)$	2	0.88	98.25	9.8K
		3	0.81	98.02	82K
		4	0.75	90.34	154.2K
		5	0.69	54.71	113.8K
		6	0.63	52.52	44.8K
Alternant	$(2, 6)$	2	0.81	57.82	15.6K
		3	0.72	53.06	14.4K
		4	0.63	51.80	18.8K

the highest distinguishable value of polynomial degree t. As illustrated in Table 3, the model perfectly distinguishes codes up to $t = 4$ which corresponds to a code rate of $R = 0.78$; a rate that is not square-distinguishable due to the condition in eq. (1). The accuracy starts to drop beyond that value of t but still does better than random. The highest degree t we can distinguish is $t = 8$ with a rate of $R = 0.56$ but only with a accuracy 0.52%. This rate is beyond the CMT[14] distinguishable range ($R \geq 2/3$). More details are provided in Table 3.

Next, we systematically vary the code length for each degree t to identify the point at which our distinguisher fails. Figure 5 presents a heatmap of accuracies for different pairs (n, t), illustrating the performance across various code lengths and degrees. These results serve as complementary benchmarks to the $m = 6$ case, where performance appears to be saturated, providing additional insights into our distinguisher's behavior. Notably, no public implementations of classical attacks are available for direct comparison, making our results a standalone reference for this setting.

Goppa Codes $n = 256$. To test the limits of our model on larger codes, we generate datasets of $8M$ samples ($4M$ for each class) for codes of length 256. We train the model of the same size on these datasets and report the accuracies obtained in Table 4. We notice the performance degrading fast with the degree t and only do better than random for $t \leq 5$. Further efforts and resources in terms

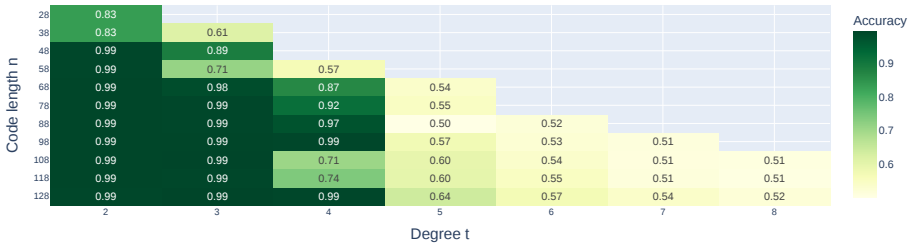

Fig. 5: Heatmap of model accuracy for classifying Goppa codes vs random codes $q = 2, m = 7$ as a function of code length n and degree parameter t.

Table 3: Distinguishing Results for Binary Goppa Codes with $n = 128$, $m = 7$. Balanced dataset $\mathcal{D}[q, n, m, t]$ of total size $\leq 40M$ samples and evaluated on $10k$ unseen samples. We also show the time complexity of the classical attacks [14] $\mathbb{C}_{CMT}^{\text{sparse}} = 3\left(\binom{tm}{2} - k + 1\right)\binom{\binom{tm}{2}+d_{\text{reg}}-1}{d_{\text{reg}}}^2$. Our attack's cost is the inference cost which is $\mathcal{O}(10ms)$ for most experiments.

Goppa Degree (t)	**Rate (R)**	**Accuracy (%)**	**Training Steps**	$\mathbb{C}_{CMT}$
2	0.89	98.14	2K	-
3	0.84	99.48	91K	2^{24}
4	0.78	98.88	36K	2^{41}
5	0.73	64.52	579K	2^{65}
6	0.67	57.00	115K	2^{97}
7	0.62	54.42	411K	✗ 2^{139}
8	0.56	52.38	20K	✗ 2^{193}
9–17	0.51–0.07	≤ 51	600K	✗ $> 2^{264}$

of dataset generation, model size, and training compute are needed to figure out the scaling laws of our approach.

A notable pattern during training as shown in Figure 7, is that the loss often almost stagnates for an extended period without the gradients vanishing before abruptly decreasing at a specific training step, denoted as $T_{q,m,t}$. This drop in loss tends to consistently occur much later for larger values of t, though the exact nature of the dependency between t and $T_{q,m,t}$ remains unclear. This raises a question about the applicability of gradient-based optimization on such tasks.

Larger codes with Code Puncturing. We applied the strategy discussed in Section 4.3 to evaluate the model trained on binary Goppa codes with $m = 7, n = 128$ on codes of parameters $n = 1024, m = 10, t = 2$ using algorithm 2 with 1000 trials. This experiment yields a 70% accuracy suggesting that there

Table 4: Classification Accuracy (%) on balanced $10k$ eval datasets of binary irreducible Goppa codes of length $n = 256$. The model used is the same throughout the paper: a tokenizer-free encoder only transformer with 4 layers and $d = 1024$ embedding dimension.

Degree (t)	2	3	4	5	6
$n = 256, m = 8$	98.06	98.38	60.36	54.74	51.66

are probably unknown relationships between families of binary Goppa codes over different field extensions.

5.3 Distinguishing MDPC and QC-MDPC Codes

We adopt the same framework outlined in `BIKE` [2]. Specifically, we take $n_0 = 2, n = 2k$, implying that QC-MDPC code has rate $\frac{1}{2}$ and the corresponding parity-check matrix is composed of two circulant blocks and the codes have a rate of . We train the model on codes of length $n = 158$ and vary the row weight w. Taking a block size $k = 79$ prime and odd values of $w/2$ ensures that the circulants are invertible in $\mathbb{F}_q$, which explains the values of w considered in Table 5. This table shows that we can distinguish MDPCcodes up to $w = 14$ while for QC-MDPC, we could only distinguish codes with row weight $w = 6$.

This outcome is somewhat surprising, as one might expect the additional structure introduced in the Quasi-Cyclic case to make classification easier rather than harder. However, the circulant structure seems to introduce constraints that makes it more challenging for the model to extract distinguishing features. An avenue of improvement is to elaborate an effective representation of this structure in a way that helps the learning of the model.

6 Hidden Goppa Code Problem

We introduce a new problem related to Goppa codes stronger than distinguishing but weaker than the key-recovery problem.

*Problem 4 (Hidden Goppa Code (*HGC*) problem). Given a parameter $\zeta > 0$, and matrix $\widetilde{\mathbf{G}} \in \mathcal{M}_{k\times n}(\mathbb{F}_q \cup \{*\})$ with at most ζ placeholder symbols $*$, the* HGC *problem asks to find – if it exists – a completion $\widehat{\mathbf{G}} \in \mathcal{M}_{k\times n}(\mathbb{F}_q)$ of $\widetilde{\mathbf{G}}$ (i.e. obtained by replacing all placeholder symbols $*$ by field elements) $\widehat{G}$ is a valid generator matrix for a Goppa code $\mathcal{G}(\boldsymbol{\alpha}, g)$ with $\boldsymbol{\alpha} \in \mathbb{F}_{q^m}^n$ a support and $g(x) \in \mathbb{F}_{q^m}[x]$ a Goppa polynomial of degree t.*

This problem is trivial in the case of random linear codes, since any solution makes a valid code. However, for Goppa codes, the structure imposed by the algebraic properties of the code constrains the space of possible solutions. This structure is defined by the Goppa polynomial and the support set, which are

Table 5: Maximum evaluation accuracy on Moderate-Density Parity-Check (MDPC) codes versus Quasi-Cyclic MDPC codes on the distinguishing task. The number of training steps needed to achieve these accuracies is also reported.

Code length & dim (n, k)	Code	Row weight (w)	Eval Accuracy	Train. Steps
158, 79	MDPC	10	97.14	445K
		11	74.36	265K
		12	65.39	200K
		13	58.28	145K
		14	54.90	220K
		16	51.73	335K
		18	51.05	220K
	QC-MDPC	6	98.02	78K
		10	51.31	738K
		14	51.36	1.08M
		18	51.21	905K

not directly visible in the generator matrix. Therefore, any solution to the HGC problem must implicitly respect these hidden parameters, making it a non-trivial task.

`DeepRecover`. Using the same transformer architecture as for the `DeepDistinguisher` with several hidden elements ranging between $\zeta = 1$ and $\zeta = 80$ out of $mt(n-mt)$ entries (e.g. 624 for $m = 6, t = 2, n = 64$), we were able to successfully train the model on this task achieving a component-wise accuracy as high as 80% for binary [64-52]-Goppa codes over extension degree $m = 6$ as shown in Table 6.

We find that larger values of ζ accelerate the model training but converge to a less optimal solution. Further investigations are required to understand better the limits of the feasibility of this problem. It is evident that there is a theoretical upper bound of ζ beyond which the number of possible solutions explodes and we think that our model works partially because, for the values we chose of ζ, the solution is either unique or there are not many solutions, allowing the model to recover the solution that we used to generate the given sample (generate a Goppa code, hide some entries, then ask the model to recover that exact solution instead of recovering any valid solution). It's worth noting that thanks to our `DeepDistinguisher`, we could also train the `DeepRecover` model to recover any valid solution since we can test in a gradient-friendly way whether a matrix is Goppa or not with high accuracy.

7 Conclusion

In this research, we looked at how well machine learning can tell apart different types of codes, specifically Goppa and Alternant codes. Our results show that Goppa codes are generally easier to distinguish, especially when the degree of

Degree (t)	Best Accuracy
2	0.80
3	0.76
4	0.64
5	0.58
6	0.50

Table 6: Best component-wise accuracy for each polynomial degree t after training on Goppa codes with parameters $q = 2, n = 64, m = 6$. An accuracy of 50% is as good as random guessing.

the Goppa polynomial t and the extension degree m are lower. However, as these parameters increase, it becomes harder to achieve high accuracy. We noticed that the training process often hits a plateau before improving, which seems to be related to the degree t. We also show that trained models on small code lengths are generalizable to larger codes without any further training.

Alternant codes, on the other hand, are more difficult to distinguish. We only managed to get better-than-random accuracy for small values of t. This suggests that Alternant codes don't have enough characteristics which makes them harder to differentiate using our current methods.

We also explored MDPC and QC-MDPC codes. We found that MDPC codes can be distinguished up to a certain row weight, but QC-MDPC codes are tougher, with successful differentiation only at very low row weights.

This work is a first step in applying machine learning to code-based cryptography, opening up new possibilities for research. Future work could focus on improving these models, trying to distill a classical approach or algorithm that our model may be approximating, and figuring out the recurring behavior of the gradient descent when training on mathematical problems.

References

1. Albrecht, M.R., Bernstein, D.J., Chou, T., Cid, C., Gilcher, J., Lange, T., Maram, V., Maurich, I.V., Misoczki, R., Niederhagen, R., Paterson, K.G., Persichetti, E., Peters, C., Schwabe, P., Sendrier, N., Szefer, J., Tjhai, C.J., Tomlinson, M., Wang, W.: Classic McEliece: conservative code-based cryptography (Oct 2022), `https://inria.hal.science/hal-04288769`, round 4 submission to the NIST call for postquantum cryptographic primitives
2. Aragon, N., Barreto, P., Bettaieb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Gaborit, P., Ghosh, S., Gueron, S., Güneysu, T., et al.: Bike: bit flipping key encapsulation (2022)
3. Ba, J.L., Kiros, J., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
4. Bardet, M., Faugère, J.C., Salvy, B., Yang, B.Y.: Asymptotic behaviour of the degree of regularity of semi-regular polynomial systems. In: The Effective Methods in Algebraic Geometry Conference – MEGA 2005. pp. 1–14 (2005)

5. Bardet, M., Mora, R., Tillich, J.P.: Polynomial time key-recovery attack on high rate random alternant codes. IEEE Transactions on Information Theory (2023)
6. Berlekamp, E., McEliece, R., van Tilborg, H.: On the inherent intractability of certain coding problems. IEEE Transactions on Information Theory **24**(3), 384–386 (May 1978)
7. Bernstein, D.J.: Grover vs. Mceliece. In: Sendrier, N. (ed.) Post-Quantum Cryptography, Third International Workshop, PQCrypto 2010, Darmstadt, Germany, May 25-28, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6061, pp. 73–80. Springer (2010). https://doi.org/10.1007/978-3-642-12929-2_6, http://dx.doi.org/10.1007/978-3-642-12929-2_6
8. Bernstein, D.J., Buchmann, J., Dahmen, E. (eds.): Post-quantum cryptography. Mathematics and Statistics Springer-11649; ZDB-2-SMA, Springer Berlin Heidelberg, Berlin, Heidelberg (2009), http://opac.inria.fr/record=b1128738
9. Buchberger, B.: Bruno Buchberger's PhD thesis 1965: An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal. Journal of Symbolic Computation **41**(3-4), 475–511 (2006)
10. Buchberger, B., Collins, G.E., Loos, R.G.K., Albrecht, R.: Computer algebra symbolic and algebraic computation. SIGSAM Bull. **16**(4), 5–5 (1982)
11. Chen, L., Jordan, S., Liu, Y.K., Moody, D., Peralta, R., Perlner, R., Smith-Tone, D.: Report on post-quantum cryptography. Reasearch report NISTIR 8105, NIST (2003), http://csrc.nist.gov/publications/drafts/nistir-8105/nistir_8105_draft.pdf
12. Cluzeau, M., Tillich, J.: On the code reverse engineering problem. In: Kschischang, F.R., Yang, E. (eds.) 2008 IEEE International Symposium on Information Theory, ISIT 2008, Toronto, ON, Canada, July 6-11, 2008. pp. 634–638. IEEE (2008). https://doi.org/10.1109/ISIT.2008.4595063, https://doi.org/10.1109/ISIT.2008.4595063
13. Courtois, N.T., Finiasz, M., Sendrier, N.: How to achieve a mceliece-based digital signature scheme. In: Boyd, C. (ed.) Advances in Cryptology - ASIACRYPT 2001, 7th International Conference on the Theory and Application of Cryptology and Information Security, Gold Coast, Australia, December 9-13, 2001, Proceedings. Lecture Notes in Computer Science, vol. 2248, pp. 157–174. Springer (2001). https://doi.org/10.1007/3-540-45682-1_10, https://doi.org/10.1007/3-540-45682-1_10
14. Couvreur, A., Mora, R., Tillich, J.P.: A new approach based on quadratic forms to attack the mceliece cryptosystem. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 3–38. Springer (2023)
15. Couvreur, A., Otmani, A., Tillich, J.P.: Polynomial time attack on wild mceliece over quadratic extensions. IEEE Transactions on Information Theory **63**(1), 404–427 (2016)
16. Cox, D.A., Little, J.B., O'Shea, D.: Ideals, Varieties and Algorithms. Springer Verlag (2005)
17. Cybenko, G.: Approximation by superpositions of a sigmoidal function. Mathematics of Control, Signals and Systems **2**(4), 303–314 (1989)
18. Dallot, L.: Sécurité de protocoles cryptographiques fondés sur les codes correcteurs d'erreurs. (Security of cryptographic protocols based on error correcting codes). Ph.D. thesis, University of Caen Normandy, France (2010), https://tel.archives-ouvertes.fr/tel-01102440
19. Debris-Alazard, T.: Code-based cryptography: Lecture notes. CoRR **abs/2304.03541** (2023). https://doi.org/10.48550/ARXIV.2304.03541, https://doi.org/10.48550/arXiv.2304.03541

20. Faugère, J., Din, M.S.E., Spaenlehauer, P.: On the complexity of the generalized minrank problem. J. Symb. Comput. **55**, 30–58 (2013). https://doi.org/10.1016/J.JSC.2013.03.004, https://doi.org/10.1016/j.jsc.2013.03.004
21. Faugere, J.C., Gauthier-Umana, V., Otmani, A., Perret, L., Tillich, J.P.: A distinguisher for high-rate mceliece cryptosystems. IEEE Transactions on Information Theory **59**(10), 6830–6844 (2013)
22. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge, MA (2016)
23. Goppa, V.D.: A new class of linear correcting codes. Problemy Peredachi Informatsii **6**(3), 24–30 (1970)
24. Gryak, J., Haralick, R.M., Kahrobaei, D.: Solving the conjugacy decision problem via machine learning. Exp. Math. **29**(1), 66–78 (2020). https://doi.org/10.1080/10586458.2018.1434704, https://doi.org/10.1080/10586458.2018.1434704
25. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)
26. Hornik, K.: Approximation capabilities of multilayer feedforward networks. Neural Networks **4**(2), 251–257 (1991)
27. Kim, B.D., Vasudevan, V.A., D'Oliveira, R.G., Cohen, A., Stahlbuhk, T., Médard, M.: Cryptanalysis via machine learning based information theoretic metrics. arXiv preprint arXiv:2501.15076 (2025)
28. Kingma, D.P.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
29. Li, C., Wenger, E., Allen-Zhu, Z., Charton, F., Lauter, K.: SALSA VERDE: a machine learning attack on Learning With Errors with sparse small secrets. In: Proc. of NeurIPS (2023)
30. Li, C.Y., Sotáková, J., Wenger, E., Malhou, M., Garcelon, E., Charton, F., Lauter, K.: Salsa Picante: A Machine Learning Attack on LWE with Binary Secrets. In: Proc. of ACM CCS (2023)
31. MacWilliams, F.J., Sloane, N.J.A.: The theory of error-correcting codes (1977), https://api.semanticscholar.org/CorpusID:118260868
32. McEliece, R.J.: A public-key cryptosystem based on algebraic. Coding Thv **4244**, 114–116 (1978)
33. Miller, E., Sturmfels, B.: Combinatorial Commutative Algebra, Graduate Texts in Mathematics, vol. 227. Springer-Verlag, New York (2005)
34. Misoczki, R., Tillich, J.P., Sendrier, N., Barreto, P.S.: MDPC-McEliece: New McEliece variants from moderate density parity-check codes. In: 2013 IEEE international symposium on information theory. pp. 2069–2073. IEEE (2013)
35. National Institute of Standards and Technology: Module-Lattice-Based Digital Signature Standard. Federal Information Processing Standards Publication 204, U.S. Department of Commerce (Aug 2024), https://csrc.nist.gov/pubs/fips/204/final
36. National Institute of Standards and Technology: Module-Lattice-Based Key-Encapsulation Mechanism Standard. Federal Information Processing Standards Publication 203, U.S. Department of Commerce (Aug 2024), https://csrc.nist.gov/pubs/fips/203/final
37. National Institute of Standards and Technology: Stateless Hash-Based Digital Signature Standard. Federal Information Processing Standards Publication 205, U.S. Department of Commerce (Aug 2024), https://csrc.nist.gov/pubs/fips/205/final

38. National Institute of Standards and Technology (NIST): Post-Quantum Cryptography: Round 4 Submissions. https://csrc.nist.gov/projects/post-quantum-cryptography/round-4-submissions (2024), accessed: 2024-12-31
39. Overbeck, R., Sendrier, N.: Code-based cryptography, pp. 95–145. Springer Berlin Heidelberg, Berlin, Heidelberg (2009). https://doi.org/10.1007/978-3-540-88702-7_4, http://dx.doi.org/10.1007/978-3-540-88702-7_4
40. Randriambololona, H.: The syzygy distinguisher. arXiv preprint arXiv:2407.15740 (2024)
41. Reed, I.S., Solomon, G.: Polynomial codes over certain finite fields. Journal of the society for industrial and applied mathematics **8**(2), 300–304 (1960)
42. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. Journal of the ACM (JACM) **56**(6), 1–40 (2009)
43. Stevens, S., Wenger, E., Li, C., Nolte, N., Saxena, E., Charton, F., Lauter, K.: Salsa fresca: Angular embeddings and pre-training for ml attacks on learning with errors. arXiv preprint arXiv:2402.01082 (2024)
44. Torres, R.C., Sendrier, N.: Analysis of information set decoding for a sub-linear error weight. In: Takagi, T. (ed.) Post-Quantum Cryptography - 7th International Workshop, PQCrypto 2016, Fukuoka, Japan, February 24-26, 2016, Proceedings. Lecture Notes in Computer Science, vol. 9606, pp. 144–161. Springer (2016). https://doi.org/10.1007/978-3-319-29360-8_10, http://dx.doi.org/10.1007/978-3-319-29360-8_10
45. Valembois, A.: Detection and recognition of a binary linear code. Discret. Appl. Math. **111**(1-2), 199–218 (2001). https://doi.org/10.1016/S0166-218X(00)00353-X, https://doi.org/10.1016/S0166-218X(00)00353-X
46. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems (NIPS). pp. 5998–6008 (2017)
47. Wenger, E., Chen, M., Charton, F., Lauter, K.E.: Salsa: Attacking lattice cryptography with transformers. Proc. of NeurIPS (2022)
48. Wenger, E., Saxena, E., Malhou, M., Thieu, E., Lauter, K.: Benchmarking attacks on learning with errors. Cryptology ePrint Archive, Paper 2024/1229 (2024), https://eprint.iacr.org/2024/1229

A Figures

B Implementation details

Datasets are generated using SageMath and saved in files for training. One noticeable artifact that occurs when standardizing non standard form binary codes in SageMath is that the resulting distribution of the generator matrix entries are not uniform due to the algorithm used to swap columns. Plotting the probabilities that an entry is 1 shows that some cells are less likely to be one. This occurs regardless of the code family. One way to avoid this is to simply discard non standard codes meaning we keep only about 29% of the codes. We use a small transformer of size $\approx 50m$ trainable parameters with 4 layers and embedding dimension $d = 1024$. Our implementation is using torch and is based on the public implementation of SALSA[47].

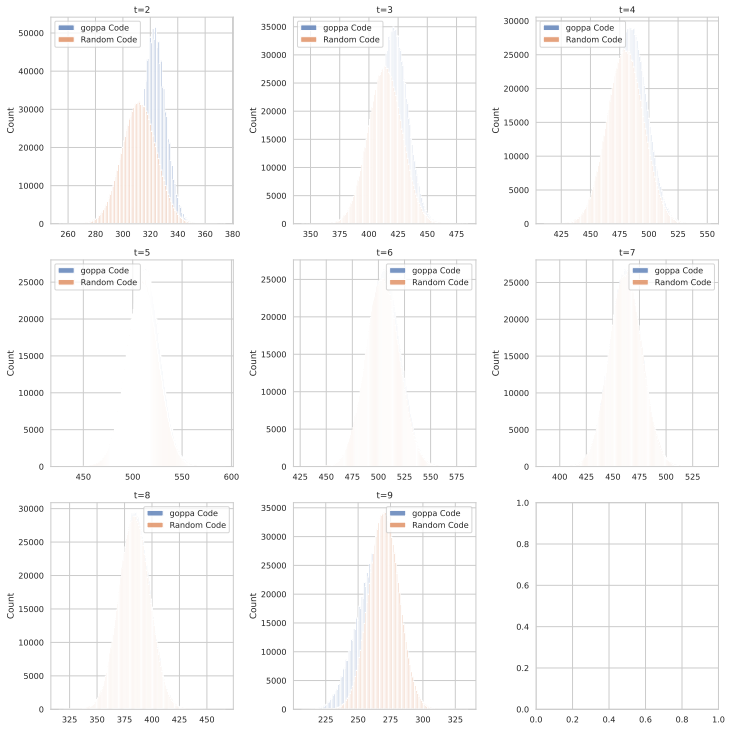

Fig. 6: Histograms of the total hamming weight of the generator matrix of Goppa codes versus random linear codes: $w_{\mathscr{F}}(\mathbf{G}) = \sum_{i=1}^{k} w_H(g_i)$ for binary codes with parameters $n = 64, m = 6$ while varying polynomial degree t.

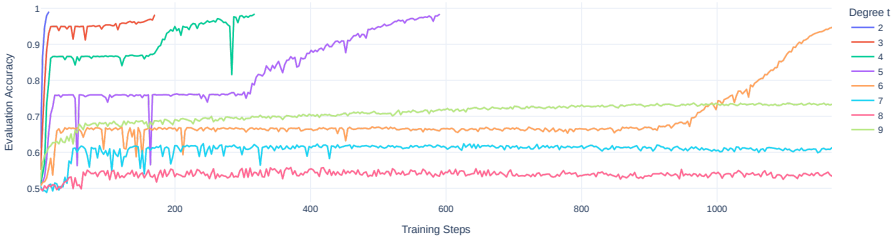

Fig. 7: Evolution of evaluation accuracy during training of the classifier over Goppa codes vs random codes with parameters $n = 64, q = 2, m = 6$. The color represents different values of the polynomial degree t.