HOVER: Higher-Order Vanishing Endomorphism Recovery

Markku-Juhani O. Saarinen[0000-0002-2555-235X]*

Tampere University, Finland
markku-juhani.saarinen@tuni.fi

Abstract. Higher-Order Vanishing (HOV), introduced by Hemmert and Wiemers at CRYPTO 2026, distinguishes Goppa codes and was recently extended to McEliece key recovery. We introduce a new variant of HOV, HOVER (Higher-Order Vanishing Endomorphism Recovery), that replaces the costly low-rank direction-search phase of HOV with a much faster linear algebra step. Our attack is based on the observation that the first catalecticants (the coefficient matrices of first partial derivatives) of the public HOV kernel define a contraction tensor; HOVER computes the coefficient endomorphisms that preserve all relations of this tensor. In the clean case, in which the scalar-extended HOV kernel is spanned by exactly the m hidden Frobenius directions, this public algebra is $\mathbb{F}_{2^m}$, and the eigenspaces of a field generator are exactly those directions, already ordered by Frobenius. We give a basis-invariant construction, a structural classification of its possible matrix-algebra outputs, and explicit acceptance conditions that return $\perp$ rather than an uncertified output. During our end-to-end public-only experiments, we recovered keys for five Technology Innovation Institute (TII) McEliece challenges, including TII-252, whose label originally represented a 2^{252} brute-force enumeration model. However, our analysis does not indicate that HOVER would threaten Classic McEliece parameters in its present form. Lastly, we note that HOVER's main cryptanalytic shortcut was discovered with a large language model, underscoring the strong potential of AI cryptanalysis.

Keywords: McEliece · Goppa codes · Higher-Order Vanishing · Tensor Endomorphisms · Key Recovery · AI Cryptanalysis

1 Introduction

After being dropped from NIST PQC Standardization in 2025 [1], the Classic McEliece [7] cryptosystem has recently been standardized by ISO [13]. McEliece is often presented as a conservative choice for post-quantum key establishment since straightforward Information Set Decoding (ISD) attacks have seen relatively little progress since the scheme was originally proposed in 1978 [14].

However, there has been a surge of new algebraic or structured attacks since 2025 (see Table 1). They have been used to break many TII Challenge parameter

* Corresponding author.

sets that were thought to be well beyond grasp just a few years ago [2, 4, 10–12, 18, 20]. Algebraic attacks are now considered such a serious security threat to Classic McEliece parameters that in September 2026 the German security agency BSI started recommending against its use in new applications [6].

Table 1. Selected cryptanalytic attacks against McEliece since 2025, schematic applicability conditions, and asymptotic claims

Attack	Type	Applicability Conditions	Asymptotics
Syzygy [16]	distinguisher	Shortened dual has excess higher syzygies.	subexp $2^{o(t)}$
Rank-2 [4]	key recovery	Rank-2 relation at low solving degree.	subexp $2^{o(t)}$
HOV [12]	distinguisher	Random HOV kernel dies; Goppa not.	–
HOV [11]	key recovery	Clean Frobenius-split HOV kernel.	–
Pfaffian [2]	key recovery	Very high rate; square defect visible.	poly in (m, t)
Hold-out [10]	key recovery	Kernel survives past Hermite threshold.	qpoly $n^{O(\log n)}$

We note that several of the recent results in Table 1 only present a *distinguisher*, an algorithm that can tell if an input is a McEliece public key or a random bit string of equal length without a trapdoor. A structural distinguisher for a McEliece public code does not necessarily break a cryptosystem that uses it. More devastating attacks are either *key recovery* attacks that find a secret key for a given public key, or *message recovery* attacks that find plaintext information given ciphertexts and a public key. This work focuses on key recovery attacks. Our main contributions are:

1. **A new key-recovery algorithm.** We define the HOVER (Higher-Order Vanishing Endomorphism Recovery) construction and describe how it can be implemented. We also analyze the conditions under which our method applies effectively.
2. **Computational results.** We report the first public-only recovery of TII-129, TII-213, TII-240, TII-246, and TII-252 secret keys from their respective challenge public keys. TII-83 and TII-248 are kernel-to-key controls. An artifact binds all seven HOVER-recovered keys to the original matrices.

2 McEliece Public Keys and Secret Keys

We use the parameter names of McEliece's original presentation and the Classic McEliece specification, together with the key notation summarized in Table 2. In McEliece's generator-matrix presentation, let $q = 2^m$ and write

$$\Gamma = (g, \alpha_0, \ldots, \alpha_{n-1}), \tag{1}$$

where $g \in \mathbb{F}_q[x]$ is monic and irreducible of degree t and the $\alpha_i \in \mathbb{F}_q$ are distinct. Denote the associated binary Goppa code by $\Gamma_2 \subseteq \mathbb{F}_2^n$, and let $\mathbf{G} \in \mathbb{F}_2^{k \times n}$ be

its generator matrix. For the proper codes considered below, $\dim \Gamma_2 = k = n - mt$. McEliece originally took $n = q$; Classic McEliece also allows shortened supports with $n \leq q$. We retain this McEliece public-key view throughout. The methods apply equivalently to the Niederreiter parity-check formulation [15] used by Classic McEliece [7, 13].

Table 2. Structure of McEliece public and secret keys [14]. Our methods also apply to the Niederreiter [15] formulation used by Classic McEliece [7, 13].

$\mathbf{G} \in \mathbb{F}_2^{k \times n}$	Generator matrix for the binary Goppa code Γ_2.
$\mathbf{P} \in \mathbb{F}_2^{n \times n}$	Uniform random permutation matrix.
$\mathbf{S} \xleftarrow{\$} \mathrm{GL}_k(\mathbb{F}_2)$	Uniform random nonsingular (invertible) matrix.
$\mathcal{D}_{\Gamma_2}$	An efficient t-error-correcting decoder for Γ_2.
$pk \leftarrow \mathbf{G}' = \mathbf{SGP}$	Scrambled generator matrix $\mathbf{G}' \in \mathbb{F}_2^{k \times n}$; kn bits in raw form.
$sk \leftarrow (\mathbf{S}, \mathbf{P}, \mathcal{D}_{\Gamma_2})$	Secret descrambling and decoding information.

A key recovery must turn the structural defect detected by a distinguisher into labeled support information, or another representation that decodes the public code. A standard endpoint is Sidelnikov–Shestakov recovery: after a sufficiently large Generalized Reed–Solomon (GRS) code has been exposed, its support and multiplier can be reconstructed in polynomial time [19]. The difficult intermediate task is to select the useful GRS-related object inside a much larger public relation space.

Higher-order vanishing gives such a relation space for binary Goppa codes. Hemmert and Wiemers compute homogeneous polynomials whose derivatives vanish at the columns of a shortened public dual [12]. Hemmert's key recovery constructs special linear combinations that omit one Frobenius GRS block and then completes them to a Goppa key [11]. We call a public intermediate procedure that extracts these useful objects from an already-computed HOV relation space a *Selector Algorithm*. It identifies, and when needed orders, certified directions or linear combinations whose derivatives omit Frobenius GRS blocks; it returns $\perp$ on failure. It neither computes the HOV relation space nor completes the Goppa key. Hemmert recognizes these combinations through unusually small first-catalecticant rank, using a structured projective MinRank search over $\mathbb{F}_q$. We refer to this unnamed step of his attack as Hemmert's Selector Algorithm.

Key Observation. Our observation is that the complete family of catalecticants retains more information than the rank of any one combination. It defines a public three-way contraction tensor. Binary endomorphisms of its coefficient mode that preserve every contraction relation form an associative algebra. For a clean HOV kernel this algebra is the hidden extension field. A field generator splits into the HOV directions, and squaring supplies their cyclic order. Thus linear algebra replaces both projective search and subsequent ordering.

This construction is related to transverse-operator and tensor-centroid methods [5, 9], but fixes the derivative mode and asks only for compatible actions on HOV coefficients. It is complementary to other recent structural attacks: conductor methods cover generic high-rate alternant codes [3], low-rank polar forms give another route to binary Goppa structure [8], a Pfaffian-based polynomial attack covers the very-high-rate binary-Goppa regime [2], and a rank-two formulation gives a heuristic subexponential attack on McEliece [4]. The syzygy distinguisher provides a further public relation space but not the tensor considered here [16].

3 Higher-Order Vanishing

Let Γ be as in (1), assume that Γ_2 is proper, and put $\boldsymbol{\alpha} = (\alpha_0, \dots, \alpha_{n-1})$. Scalar extension of the public dual has the Frobenius–GRS decomposition

$$\Gamma_2^{\perp} \otimes_{\mathbb{F}_2} \mathbb{F}_q = \bigoplus_{i=0}^{m-1} \mathcal{G}_i, \qquad \mathcal{G}_i = \mathrm{GRS}_t\big(\boldsymbol{\alpha}, g(\boldsymbol{\alpha})^{-1}\big)^{[2^i]}. \tag{2}$$

Here evaluation and inversion are componentwise. Properness is exactly the dimension-mt condition needed for the displayed sum to be direct.

Fix an HOV degree $p \geq 3$ and put

$$\begin{aligned} &e = 1 + \left\lfloor \log_2 \frac{t-1}{p-1} \right\rfloor, \quad s_* = (e+2)t - (p-1)2^e - p, \\ &\kappa = mt - s_*, \qquad\qquad a = m - e - 2, \qquad b = (m-1)t - s_*. \end{aligned} \tag{3}$$

Here e is an auxiliary dyadic index used only to set the other quantities, s_* is the number of shortened coordinates, κ is the resulting number of variables, a is the window length, and b is the catalecticant-rank bound. A *consecutive HOV window* is a combination supported on a cyclically adjacent hidden directions. After a public permutation and row reduction, shorten $\Gamma_2^{\perp}$ on s_* pivot coordinates. The modified HOV map evaluates selected derivatives of square-free degree-p forms at the remaining public columns. Its binary kernel

$$L = \ker \widetilde{\varphi}^{(p)}_{H_{s_*}} \tag{4}$$

is the input to both Hemmert's Selector Algorithm and the HOVER Selector Algorithm. For a random code of the same dimensions, this kernel should be zero; we check this for every parameter profile used below.

Hemmert proves that the scalar extension of L contains determinant forms $h_0, \dots, h_{m-1}$. Each is the determinant of a $p \times p$ matrix of linear forms, and their projective lines form a Frobenius orbit. Any linear combination supported on a consecutive directions omits one of the blocks $\mathcal{G}_i$ in (2). Its first catalecticant consequently has rank at most b, and its derivatives lift to a subspace of $\bigoplus_{j \neq i} \mathcal{G}_j$ [11].

Hemmert's Selector Algorithm recognizes such combinations by the low-rank bound. Its parameter rule includes

$$ap^2 > b, \tag{5}$$

which the low-rank heuristic needs so that, among sums of a determinant directions, essentially only consecutive windows have rank at most b. HOVER uses the existence and block-omission results, but does not use (5): it first recovers the individual directions from the tensor formed by all catalecticants. The clean-kernel condition that the h_i span $L \otimes \mathbb{F}_q$ remains essential.

4 The Compatible-Endomorphism Algebra

4.1 Public Construction

For $S \subseteq [\kappa] = \{1, \ldots, \kappa\}$, write $x_S = \prod_{j \in S} x_j$. If $f = \sum_{|U|=p} f_U x_U$ is a binary square-free degree-p form, its first catalecticant is the coefficient matrix of its first partial derivatives:

$$\mathrm{Cat}(f)_{r,S} = [x_S]\frac{\partial f}{\partial x_r} = \begin{cases} f_{S\cup\{r\}}, & r \notin S, \\ 0, & r \in S, \end{cases} \qquad r \in [\kappa], \quad S \in \binom{[\kappa]}{p-1}.$$

Thus row r is the coefficient vector of $\partial f/\partial x_r$, and $\mathrm{Cat}(f)^\top z$ is the coefficient vector of the directional derivative $\sum_r z_r \partial f/\partial x_r$. Set $d = \dim_{\mathbb{F}_2} L$ and choose a binary basis $v_1, \ldots, v_d$ of L. Its first catalecticants are

$$M_i = \mathrm{Cat}(v_i) \in \mathbb{F}_2^{\kappa \times N}, \qquad N = \binom{\kappa}{p-1}. \tag{6}$$

They define the contraction

$$\beta : \mathbb{F}_2^\kappa \times L \longrightarrow \mathbb{F}_2^N, \qquad \beta(z, \lambda) = \left(\sum_{i=1}^{d} \lambda_i M_i\right)^\top z. \tag{7}$$

The raw linearization $\widetilde{T} \in \mathbb{F}_2^{N \times \kappa d}$ has entries $\widetilde{T}_{S,(r,i)} = (M_i)_{r,S}$, with its columns indexed by the basis $e_r \otimes v_i$ of $\mathbb{F}_2^\kappa \otimes L$. Let $\rho = \operatorname{rank} \widetilde{T}$ and $D = \operatorname{im} \widetilde{T} = \operatorname{span}(\operatorname{im} \beta)$. By *row compression* we mean choosing coordinates on D, equivalently row-reducing $\widetilde{T}$ and retaining its ρ independent rows. This gives a full-row-rank matrix $T \in \mathbb{F}_2^{\rho \times \kappa d}$ for the surjection $T : \mathbb{F}_2^\kappa \otimes L \twoheadrightarrow D$, with $\ker T = \ker \widetilde{T}$. Define

$$\mathcal{E}_\beta = \{B \in \mathrm{End}_{\mathbb{F}_2}(L) : (I_\kappa \otimes B)(\ker T) \subseteq \ker T\}. \tag{8}$$

Lemma 1 (Kernel criterion). *An endomorphism B lies in $\mathcal{E}_\beta$ if and only if it induces a unique $A_B \in \mathrm{End}_{\mathbb{F}_2}(D)$ satisfying*

$$A_B \beta(z, \lambda) = \beta(z, B\lambda). \tag{9}$$

Consequently $\mathcal{E}_\beta$ is a unital associative algebra. A change of basis in L conjugates it, and a coordinate change in D leaves it unchanged.

Proof. If $A_B T = T(I_\kappa \otimes B)$, preservation of $\ker T$ is immediate. Conversely, if B preserves $\ker T$, define

$$A_B(Tu) = T(I_\kappa \otimes B)u. \tag{10}$$

The definition is independent of the chosen preimage precisely because the kernel is preserved. This proves (9). The identity, sums, and scalar multiples are compatible; if B_1, B_2 induce A_1, A_2, then $A_1 A_2 T = T(I_\kappa \otimes B_1 B_2)$, proving closure under composition.

A basis change P in L sends the coefficient matrix B to PBP^{-1}. A coordinate change in D left-multiplies T by an invertible matrix and does not change its kernel. This proves the stated invariance.

If the rows of Z span the right kernel of this full-row-rank T, the kernel criterion is the binary linear system

$$T(I_\kappa \otimes B)Z^\top = 0 \tag{11}$$

in the d^2 entries of B. Written entrywise, it has $\rho(\kappa d - \rho)$ scalar equations, but still only d^2 columns. Thus one nullspace computation obtains the whole algebra. For TII-252, $(\kappa, d, \rho) = (74, 10, 90)$ gives 58,500 equations in 100 unknowns, replacing a search in $\mathbb{P}^9(\mathbb{F}_{2^{10}})$.

Proposition 1 (Public field certificate). *Suppose $d = m$, $\dim_{\mathbb{F}_2} \mathcal{E}_\beta = m$, and some $B \in \mathcal{E}_\beta$ has irreducible characteristic polynomial of degree m. Then*

$$\mathcal{E}_\beta = \mathbb{F}_2[B] \simeq \mathbb{F}_q. \tag{12}$$

Over $\mathbb{F}_q$, B has m distinct eigenlines and coefficientwise squaring permutes them in one cycle.

Proof. Irreducibility implies that the minimal and characteristic polynomials of B are equal and of degree m. Hence $\mathbb{F}_2[B]$ is a field of binary dimension m contained in $\mathcal{E}_\beta$. Equality follows from the assumed dimension. The irreducible polynomial is separable; its roots are one Frobenius orbit in $\mathbb{F}_q$. Since B has binary entries, squaring an eigenvector squares its eigenvalue, so the eigenspaces form the same cycle.

This certificate is entirely public. It says what the computed algebra is; the next result says when its eigenlines have the HOV meaning needed for key recovery.

4.2 Structural Identification

Extend scalars to $\mathbb{F}_q$. For a hidden direction h_i, write

$$\beta_i(z) = \beta(z, h_i), \qquad D_i = \operatorname{im} \beta_i, \qquad R_i = \ker \beta_i. \tag{13}$$

Assumption 1 (Clean orbit and direct images). *The kernel has $d = m$, and $L \otimes \mathbb{F}_q$ has a basis $h_0, \ldots, h_{m-1}$ with $[h_{i+1}] = [h_i^{[2]}]$. Moreover,*

$$D \otimes \mathbb{F}_q = D_0 \oplus \cdots \oplus D_{m-1}. \tag{14}$$

The direct-image condition is stronger than equality of the individual ranks. It is nevertheless an exact public rank test after the candidate eigenlines have been recovered.

Theorem 1 (Algebra classification). *Under Assumption 1, partition the indices by $i \sim j$ when $R_i = R_j$. The classes are cosets of a subgroup of $\mathbb{Z}/m\mathbb{Z}$ and have a common size c. With $\tau = m/c$,*

$$\mathcal{E}_\beta \simeq M_c(\mathbb{F}_{2^\tau}), \qquad \dim_{\mathbb{F}_2} \mathcal{E}_\beta = mc. \tag{15}$$

Proof. Work over $K = \mathbb{F}_q$ and write $Bh_i = \sum_j \mu_{ji} h_j$. Let A_B be the compatible output action and let π_j project the direct sum in (14) onto D_j. Compatibility gives

$$(\pi_j A_B|_{D_i})\beta_i(z) = \mu_{ji}\beta_j(z). \tag{16}$$

If $\mu_{ji} \neq 0$, every $z \in R_i$ therefore belongs to R_j. Frobenius transport gives all β_i the same rank, so $R_i \subseteq R_j$ implies $R_i = R_j$. Thus every compatible action is block diagonal for the equal-kernel partition.

Conversely, when $R_i = R_j$, both β_i and β_j induce isomorphisms from the same quotient $\mathbb{F}_q^\kappa / R_i$ to D_i and D_j. Consequently, every linear map between directions in one equal-kernel class has a compatible output map. Hence

$$\mathcal{E}_\beta \otimes_{\mathbb{F}_2} K \simeq \prod_{\nu=1}^{\tau} M_c(K). \tag{17}$$

Frobenius translates the index set, so an invariant equivalence relation on $\mathbb{Z}/m\mathbb{Z}$ consists of cosets of one subgroup. The $\tau = m/c$ matrix factors are permuted transitively. Galois descent gives a central simple algebra of degree c over $\mathbb{F}_{2^\tau}$; finite fields have trivial Brauer group, so it is $M_c(\mathbb{F}_{2^\tau})$. Its binary dimension is $c^2\tau = mc$.

Corollary 1 (HOVER identification). *Under Assumption 1, if the R_i are pairwise distinct, then $\mathcal{E}_\beta \simeq \mathbb{F}_q$. The eigenlines over $\mathbb{F}_q$ of any element having degree m over $\mathbb{F}_2$ are exactly $[h_0], \ldots, [h_{m-1}]$, and Frobenius gives their cyclic order.*

Proof. Pairwise separation makes $c = 1$ in Theorem 1, giving $\mathcal{E}_\beta \simeq \mathbb{F}_q$. After scalar extension, the algebra is diagonal in the hidden basis. A degree-m field element has m distinct conjugate eigenvalues, so its eigenspaces are exactly the hidden lines. Coefficientwise Frobenius supplies their cyclic order.

The classification is also a failure-mode description. If $c > 1$, the algebra reveals τ Frobenius-related blocks but leaves a residual search in $\mathbb{P}^{c-1}(\mathbb{F}_q)$. If (14) fails, reducible or nonsemisimple algebras may occur. Neither outcome proves that the original HOV MinRank problem has no solution; the present HOVER implementation refuses and falls back.

Algorithm 1: HoverSelector() : The HOVER Selector Algorithm for one shortened HOV kernel (high-level pseudocode).

Input: A binary HOV kernel L, its shortening data, and (m, t, p).
Output: A certified record of Frobenius-ordered HOV directions and their lifted block-omitting derivative spaces, or $\perp$.

Check $\dim L = m$ and replay every basis vector in the full HOV operator;
Build the first catalecticants, T, $\ker T$, and the system in (11);
Compute $\mathcal{E}_\beta$ and check identity, closure, and $\dim \mathcal{E}_\beta = m$;
Find $B \in \mathcal{E}_\beta$ with irreducible degree-m characteristic polynomial; otherwise **return** $\perp$;
Split B over the attack copy of $\mathbb{F}_q$ using an explicit field isomorphism;
Check m one-dimensional eigenspaces, one Frobenius cycle, and (14); otherwise **return** $\perp$;
Form the m cyclic windows of length $a = m - e - 2$;
Differentiate, lift through the public shortening, and retain only windows passing the exact rank and block-space gates;
return *the Frobenius-ordered directions and lifted spaces, or* $\perp$ *if a gate fails.*

4.3 A Fail-Closed Selector Algorithm

We call an algorithm *fail-closed* when every failed or ambiguous check returns $\perp$ rather than a candidate output. Algorithm 1 states the new part of the attack. A field may be constructed internally with a modulus different from that of the field used for the public matrix. Therefore, eigenvector coefficients are transported through an explicit root-defined field isomorphism; integer encodings are never treated as field maps.

An eigenvector is defined only up to a nonzero scalar. This is harmless: choosing one representative on each line spans the same consecutive window, and Hemmert's block-omission result allows arbitrary nonzero coefficients in that window. We retest all ranks and lifted row spaces after the choice.

5 From HOVER Directions to a Goppa Key

Algorithm 2 composes the Selector Algorithm with the two completion routes used in this work. The subroutine DirectionCompletion() uses the Segre–Veronese geometry of one direction to recover a partial support and then completes its labels. The subroutine WindowCompletion() aligns windows, takes Frobenius intersections and Schur products, and performs GRS recovery. Each returns $\perp$ when an exact gate fails. The predicate VerifyKey() is precisely the complete public row-space test in (22).

5.1 Frobenius Alignment and GRS Extraction

For every accepted shortening, differentiation of a consecutive window gives a public space contained in one block-omitting sum $V_i = \bigoplus_{j \neq i} \mathcal{G}_j$. The local Frobenius order is known, but independent shortenings have an unknown cyclic phase.

Algorithm 2: Hover() : HOVER key recovery (high-level pseudocode).

Input: A McEliece public generator matrix $\mathbf{G}'$, parameters (n, m, t, p), and a limit N_{try} on public shortening trials.

Output: An equivalent Goppa key Γ', or $\perp$.

Compute a basis H_{pub} of $\mathrm{RowSpace}(\mathbf{G}')^{\perp}$ and the HOV values s_*, a, b;
Initialize an empty list $\mathcal{C}$ of accepted selector records;
For at most N_{try} public pivot shortenings, construct the shortened dual and modified degree-p HOV operator, then compute its kernel L;
Apply HoverSelector$(L, \text{shortening data}, m, t, p)$ to each kernel and collect every non-$\perp$ record in $\mathcal{C}$;
If $p = 3$, apply DirectionCompletion$(H_{\mathrm{pub}}, \mathcal{R})$ to each $\mathcal{R} \in \mathcal{C}$ and return the first candidate passing VerifyKey();
$\Gamma' \leftarrow$ WindowCompletion$(H_{\mathrm{pub}}, \mathcal{C})$;
If $\Gamma' \neq \perp$ and VerifyKey$(H_{\mathrm{pub}}, \Gamma')$, **return** Γ';
return $\perp$

We fix the first phase and enumerate all m shifts of each subsequent shortening. Equal row spaces are merged; exact dimensions of accumulated sums, Frobenius intersections, and Schur products reject incompatible branches. With three shortenings and $m \leq 10$, exhaustive alignment has at most $m^2 \leq 100$ phase states.

Once the accumulated space is sufficiently large, successive Frobenius intersections isolate one or two adjacent $\mathcal{G}_i$. A componentwise product then exposes a GRS code of sufficient dimension for Sidelnikov–Shestakov recovery. This part follows Hemmert's completion, but HOVER supplies every local window at once and in cyclic order. A low accumulated dimension is not accepted. The implementation tests the complete intersection trace and GRS dimension before invoking recovery.

5.2 One-Direction Geometry

For degree three, a single recovered direction has enough geometry to bypass shortening alignment altogether. The following determinantal description also explains why.

Proposition 2 (Rational-normal-scroll containment). *Suppose a recovered HOV direction is $h = \det Q$, where Q is a $p \times p$ matrix of linear forms, and $\operatorname{rank} \mathrm{Cat}(h) = p^2$. After a row-coordinate change, every nonzero projective column of its order-$(p-1)$ derivative space has the form*

$$(1, \alpha_j, \ldots, \alpha_j^{p-1}) \otimes w_j. \tag{18}$$

Thus the columns lie on the rational normal scroll given by the Segre–Veronese variety

$$\mathrm{Seg}\big(\nu_{p-1}(\mathbb{P}^1) \times \mathbb{P}^{p-1}\big) \subseteq \mathbb{P}^{p^2-1}. \tag{19}$$

Proof. The full catalecticant-rank hypothesis identifies the order-$(p-1)$ derivative span of $\det Q$ with the evaluated entry span of Q. In the HOV determinant construction, the evaluation of Q at coordinate j has rank one, and one factor is proportional to $(1, \alpha_j, \ldots, \alpha_j^{p-1})$. Vectorizing the rank-one matrix gives (18). These are exactly the coordinates of the indicated Segre–Veronese embedding.

On every reported degree-three space, exact quadratic-ideal and linear-syzygy tests recognize the two tensor factors. Factoring each nonzero public column then returns the labeled $\mathbb{P}^1$ coordinate. This is used only after HOVER has selected a direction; the scroll equations alone were not observed to select it from the full HOV kernel.

5.3 Completing Labels and Verifying the Key

Two completion paths apply after HOVER selection. The degree-four completion is also used after MinRank selection in the tabulated runs. For the degree-four targets, choosing a second shortening on the first shortening's public pivot coordinates forces maximal accumulated growth. This completes TII-129 and contributes to TII-213. The remaining steps recover a full GRS support and then solve for a degree-t Goppa polynomial.

For TII-240, TII-246, and TII-252, one degree-three HOVER factor and Proposition 2 expose $n - 36$ labeled coordinates; the 36 shortening positions give zero columns in that factor. Restricted Goppa recovery fixes the common row operation. For each omitted public column, the ratio of its first-degree and constant reconstructed parity-check entries determines at most one compatible unused nonroot of g. Exhaustive bipartite matching restores all 36 labels. Ambiguity, a repeated support entry, a root of g, or a failed row-space check causes rejection.

In every path, the final verification reconstructs the full Frobenius parity-check matrix from a candidate Γ' of the form in (1). It compares its scalar-extended row space with the original public dual. Therefore, a returned key is an equivalent decoding key regardless of which shortening, field root, phase, or GRS normalization produced it. Random choices can cause a refusal, but cannot make an incorrect key pass this test.

6 TII Key Recoveries

6.1 Inputs and Structural Measurements

The five rows marked "new" below were absent from Hemmert's public solution archive at the cited commit.[1] Their public inputs are verbatim Track-2 matrices

[1] Tobias Hemmert, *Recovered Keys for TII McEliece Challenges*, software and data archive, `https://github.com/tobhem/key-recovery-mceliece-tii-solutions`, commit `25f5305` (2026).

from the TII challenge repository of Kirshanova and Esser.[2] TII-83 and TII-248 reproduce keys already published by Hemmert and exercise the kernel-to-key pipeline against known answers.

Table 3 reports the invariants that decide HOVER before any secret-side comparison. Here $\kappa = mt - s_*$ is the number of variables after shortening, and $\widehat{h}_i$ denotes a public representative of the ith recovered eigenline. In every row,

$$\operatorname{rank} T = \sum_{i=0}^{m-1} \operatorname{rank} \operatorname{Cat}(\widehat{h}_i), \tag{20}$$

so the recovered eigenline images form a direct sum. Every row also has $d = m$, while the algebra dimension and irreducible generator certify the simple-field case. Write ρ_{eig} for the rank of each recovered eigenline and ρ_{win} for the rank of each accepted window. The measurements obey the exact law $\rho_{\mathrm{win}} = \min\{a\rho_{\mathrm{eig}}, b\}$: the bound b is active in the first four rows, whereas the final row has $\rho_{\mathrm{eig}} = p^2 = 9$ and $a\rho_{\mathrm{eig}} = 45 < b = 63$.

Table 3. Exact public HOVER structure. Targets on one row have identical structural dimensions. Here $d = \dim L$ and $\rho = \operatorname{rank} T$; ρ_{eig} and ρ_{win} are the ranks of every eigenline and accepted window.

Target	p	κ	d	ρ	$\dim \mathcal{E}_\beta$	ρ_{eig}	ρ_{win}
TII-83	4	35	8	80	8	10	30
TII-248	3	46	9	81	9	9	39
TII-129	4	61	9	144	9	16	52
TII-213	4	66	9	144	9	16	56
TII-240/246/252	3	74	10	90	10	9	45

An exhaustive subset enumeration provides a separate check on the cyclic-window interpretation. For TII-83, exactly eight of the $\binom{8}{5} = 56$ five-line subsets have the target rank; for TII-248, TII-129, and TII-213, exactly nine of $\binom{9}{5} = 126$ do. The recorded subset identities match the Frobenius-consecutive windows returned by HOVER for TII-83, TII-248, and TII-129; for TII-213, the retained record provides only the count. The HOVER Selector Algorithm obtains its ordered windows from one field split rather than testing all subsets.

6.2 End-to-End Results

We call an execution *full-from-public* when it starts from the original public matrix, computes every HOV kernel it uses, and completes and verifies a key. The five new recoveries in Table 4 have this form. For TII-129 and TII-213, the

[2] Elena Kirshanova and André Esser, *TII McEliece Key-Recovery Challenge Data and Scripts*, software and challenge-data repository, `https://github.com/ElenaKirshanova/tii_decoding_challenge`, commit `e502cd5` (2023).

tabulated timings are from our implementation of Hemmert's Selector Algorithm (MinRank). Subsequent HOVER kernel-to-key replays recovered and verified equivalent keys from retained public-derived kernels without using the previously recovered keys. The two controls also start from saved public-derived kernels and test the selector and completion pipeline against previously published keys.

Table 4. Measured HOV/HOVER recovery records. The new rows are full-from-public; the controls are saved-kernel replays. The TII number is a historical challenge label, not a present security estimate.

Target	Status	(n, m, t)	p	Execution route	Wall time (h:min:s)	Resident set size record
TII-129	new	$(509, 9, 9)$	4	MinRank, three shortenings	4:30:28	80.30 GiB
TII-213	new	$(496, 9, 10)$	4	MinRank, two shortenings	6:44:37	$\geq$ 144 GiB
TII-240	new	$(1010, 10, 11)$	3	one-direction recovery	0:03:06	2.70 GiB†
TII-246	new	$(1009, 10, 11)$	3	one-direction recovery	0:03:11	2.70 GiB†
TII-252	new	$(1008, 10, 11)$	3	one-direction recovery	0:06:15	2.70 GiB†
TII-83	control	$(253, 8, 5)$	4	saved-kernel replay	0:00:24.3	1.70 GiB
TII-248	control	$(482, 9, 7)$	3	saved-kernel replay	0:13:38	0.40 GiB

†The 2.70-GiB figures use the largest-process reference measurement for this recovery path; the transient laptop process tree measured 9.79 and 10.71 GiB on TII-240 and TII-246. The TII-213 entry is a recorded lower bound; platform details are in Section 6.3.

For TII-240, TII-246, and TII-252, degree four would create a roughly $3.0 \cdot 10^6$ by $1.75 \cdot 10^6$ vanishing operator. HOVER instead admits degree three after the random-code kernel is checked to be zero: each shortening has a dimension-10 HOV kernel, a dimension-10 field algebra, ten rank-9 eigenlines, and ten rank-45 consecutive windows. In the notation of (3), $ap^2 = 45 < b = 63$, so the published rarity condition (5) fails exactly where the HOVER Selector Algorithm succeeds. This is a lower-degree, instance-certified optimization; it does not prove that degree three succeeds throughout a parameter family.

The five new rows report recorded full-from-public executions. The integrated path now tries one-shortening completion first for degree-three targets and can fall back to three-shortening alignment; Section 6.3 records saved-kernel replay evidence separately from the table's timings. On all nine retained degree-three kernels, a separate comparison also found that the usual consecutive-window intersection test returns exactly one direct eigenline lift, with the ten matches forming one Frobenius cycle. This supports a future shortcut but is not used to rewrite the executed recovery record.

The supplementary artifact binds the exact public matrices to canonical recovered supports and Goppa polynomials and verifies each public row space independently of the attack implementation. Section 6.3 gives the exact verification equation and platform details.

The label TII-252 originally encoded a 2^{252} enumeration model involving candidate Goppa polynomials and support sets. That historical model is not

a modern security estimate. The defensible result is therefore "recovery of the challenge TII-252," not "breaking 252-bit security."

6.3 Verification and Measurement Record

Independent Key Verification. The artifact[3] contains the exact public matrices, their upstream locations and SHA-256 digests, canonical recovered supports and Goppa polynomials, run records, and a standalone Sage verifier independent of the attack implementation. For each candidate $\Gamma = (g, \alpha_0, \ldots, \alpha_{n-1})$ it reconstructs $\mathbb{F}_q$, checks the declared parameters and field modulus, checks that the support is distinct and avoids the roots of the monic irreducible g, and forms

$$H_{\mathrm{rec}}[jt+u, i] = \left(g(\alpha_i)^{-1}\alpha_i^u\right)^{2^j}, \qquad 0 \le j < m, \quad 0 \le u < t. \tag{21}$$

It accepts only when

$$\mathrm{RowSpace}_{\mathbb{F}_q}(H_{\mathrm{rec}}) = \mathrm{RowSpace}_{\mathbb{F}_q}(H_{\mathrm{pub}}). \tag{22}$$

Structural Replay. For every shortening used by the attack, the retained record replays the HOV kernel against the complete vanishing operator, checks multiplication and the identity in $\mathcal{E}_\beta$, verifies an irreducible field generator, splits it through an explicit field isomorphism, and records every eigenline, Frobenius edge, catalecticant rank, lifted row-space digest, and completion gate. Equation (20) is evaluated from those public spaces. The degree-three parameter profile also has a zero matched random-code kernel. These checks establish the hypotheses on the reported instances; they are not used to infer a family-wide success probability.

Tabulated Run Measurements. TII-129 and TII-213 ran on a 28-vCPU large-memory server; the other rows of Table 4 ran on a 12-core AMD Ryzen AI 9 HX 370 laptop. Cross-machine wall times are measurements, not a controlled speed comparison.

Later One-Shortening Evidence. The restricted-key replays opened only the original public matrix, one saved public-derived HOVER space, and its provenance. They did not open a saved key, phase choice, or completed support. TII-240, TII-246, and TII-252 recovered 974, 973, and 972 labeled coordinates, respectively, and uniquely matched all 36 omissions from 50, 51, and 52 unused nonroots. The retained replays took 207.86, 262.57, and 250.67 s, excluding HOV kernel construction. These are post-shortening measurements and are deliberately not substituted for the full-from-public times in Table 4.

[3] Verification artifact: `https://github.com/mjosaarinen/tii-solved`

7 Cost, Applicability, and Limitations

7.1 What the Selector Algorithm Changes

Hemmert's Selector Algorithm is estimated to test approximately

$$c_{\text{search}} \simeq \frac{1}{m}\, 2^{m(e+2)} \tag{23}$$

catalecticant ranks before finding a block-omitting combination [11]. Conditional on an explicit dimension-m kernel L, HOVER replaces this by construction of the contraction and one binary system in m^2 unknowns, followed by degree-m polynomial factorization and eigenspaces over $\mathbb{F}_q$. These operations, the phase enumeration for the reported values of m, GRS recovery, and final verification are polynomial in their explicit matrix dimensions.

This comparison has two consequences. First, at a fixed HOV degree satisfying the published rarity condition, HOVER removes the projective search but may not change the dominant cost. Second, it can matter much more by removing that condition. At $(m, t) = (10, 11)$, degree four was selected to make low-rank diagnostics reliable, whereas HOVER's field algebra separates the directions already at degree three. Reducing the HOV operator from roughly $3.0 \cdot 10^6$ by $1.75 \cdot 10^6$ to roughly $6.7 \cdot 10^4$ by $6.5 \cdot 10^4$ (rows by columns) is the main practical gain in the last three rows of Table 4.

7.2 Upstream HOV Work Factor

The modified degree-p HOV operator has

$$\binom{\kappa}{p} \quad \text{columns and} \quad k \sum_{\ell=1}^{\lfloor \log_2 p \rfloor} \binom{\kappa}{p - 2^\ell} \quad \text{rows}, \qquad \kappa = mt - s_*, \quad k = n - mt. \tag{24}$$

It is sparse, so the attack uses iterative binary linear algebra rather than dense elimination, but the binomial dimensions remain binding. Calling the HOVER Selector Algorithm polynomial means polynomial in this explicit HOV-kernel representation; it does not make the complete attack polynomial in the public key length.

For the five Classic McEliece parameter sets, Hemmert's parameter rule gives HOV degrees between 22 and 52 and estimates the leading sparse-kernel term between approximately 2^{297} and 2^{690} operations [11]. Hemmert's Selector Algorithm is lower order there. HOVER therefore does not turn any Classic parameter set into an executable attack, even under the clean identification assumptions.

7.3 When HOVER Is Effective

Figure 1 locates reported recoveries by extension degree m and Goppa degree t. It is an empirical map: a point does not determine the length n, HOV degree p, or instance-dependent conditions for success.

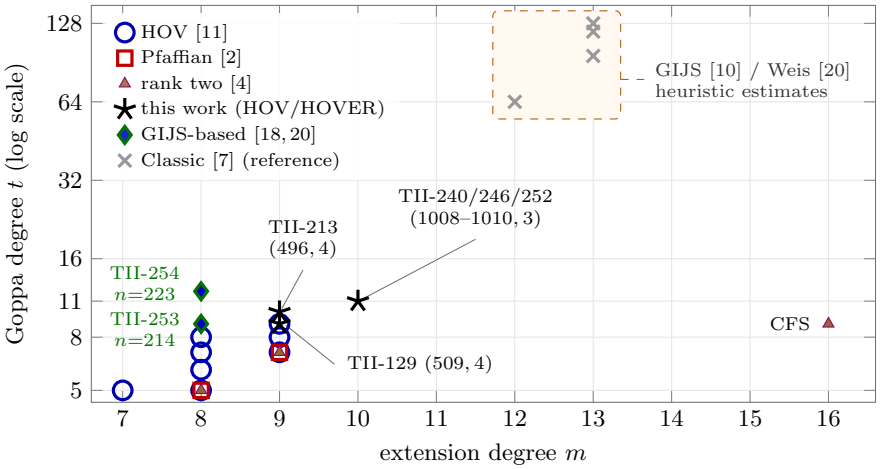

Fig. 1. Reported binary-Goppa challenge and CFS key recoveries by (m, t); coincident marks denote several methods. This work's annotations give (n, p). Gray crosses show the five Classic McEliece parameter sets (two share $(13, 128)$). *Addendum:* green diamonds are recoveries based on Ghoshal et al. (GIJS) [10]: TII-253 by Weis [20] and TII-254 by a two-anchor variant [18]. The dashed box groups Classic reference cells covered by heuristic key-recovery estimates in GIJS [10] and Weis [20], with no executed recovery or success region implied. Distinguisher-only results and generated test keys are omitted.

The numerical parameters (n, m, t) determine candidate values of p and s_*, but they do not alone determine success. The effective regime has four publicly testable features:

1. the HOV operator in (24) is computationally reachable, and the matched random-code profile has zero kernel;
2. the structured kernel has dimension m and is verified against the complete HOV operator;
3. its compatible algebra passes the field certificate, and the resulting eigenline images pass the direct-sum and separation ranks; and
4. at least one completion route reaches a GRS code or a partial support and then passes complete public-key verification.

The first and fourth items depend on resources and the instance. The second contains the clean-kernel premise inherited from HOV. The public field, Frobenius-cycle, rank, and verification tests certify all subsequent outputs but do not, by themselves, prove that every key in a parameter family has the same clean orbit. Thus the experiments establish seven exact instances, not a success probability over key generation.

If the algebra has the matrix form in (15) with $c > 1$, it still reduces the original projective problem to a residual $\mathbb{P}^{c-1}$ search. No natural HOV target

in the present data has this shape, so that branch is a theorem-level classification rather than a claimed attack implementation. Any other algebra shape, failed field map, ambiguous phase, or incomplete support makes the implementation return $\perp$ with the failed check recorded. Only the final row-space equality authorizes a key.

Addendum. Ghoshal et al. (GIJS) extend their hold-out distinguisher to heuristic key recovery, estimated at 2^{130}–2^{149} bit operations for the Classic McEliece parameter sets [10]. Weis lowers the cost of a distinguisher run by about 20 bits and estimates key recovery at 2^{94}–2^{102} in the same memory-free cost model; he also recovered TII-253 [20]. A two-anchor holdout/Hermite variant subsequently recovered TII-254 [18].

8 Conclusion

HOVER recovers the hidden directions of a clean higher-order-vanishing kernel as eigenspaces of a public, compatible endomorphism algebra. The computation is basis invariant and uses one binary nullspace rather than projective MinRank enumeration. The classification theorem explains both the observed field and the residual ambiguity when derivative kernels coincide.

HOVER recovers keys for all five newly solved TII challenges and reproduces two published controls. Removing the low-rank-rarity requirement makes degree three usable for three $t = 11$ instances. One-shortening partial-support completion also removes phase alignment for those instances; the tabulated full-from-public executions remain the timing record. A returned equivalent Goppa key is unconditionally verified against the complete public matrix.

Generic success of the clean algebra remains unproved, and the upstream HOV kernels estimated for Classic McEliece remain beyond practical reach.

Addendum. On 29 September 2026, the German Federal Office for Information Security (BSI) advised against Classic McEliece in new applications [6], reversing a position BSI had held since 2020. This work's ePrint version [17] is cited along with more recent attacks [10, 18, 20] as informing that reassessment.

On the Use of AI in Cryptanalysis

We used OpenAI GPT-5.6-Sol and Claude Opus 5 for theory formulation, proofs, analysis, and preparing and executing experiments. We also used the LLMs and Grammarly to improve the presentation.

The discovery of HOVER's main cryptanalytic shortcut can largely be credited to large language models, guided by hundreds of directed prompts and manual edits over several months under close human supervision.

Acknowledgments. The author acknowledges CSC – IT Center for Science, Finland, for computational resources.

Disclosure of Interests. The author has no competing interests to declare that are relevant to the content of this article.

References

1. Alagic, G., Bros, M., Ciadoux, P., Cooper, D., Dang, Q., Dang, T., Kelsey, J., Lichtinger, J., Liu, Y.K., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., Silberg, H., Smith-Tone, D., Waller, N.: Status report on the fourth round of the NIST post-quantum cryptography standardization process. Interagency or Internal Report NIST IR 8545, National Institute of Standards and Technology (March 2025). https://doi.org/10.6028/NIST.IR.8545
2. Bardet, M., Lemoine, A., Tillich, J.P.: An attack on the CFS scheme and on TII McEliece challenges. In: Heninger, N., Rosulek, M. (eds.) Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 16803, pp. 380–412. Springer (2026). https://doi.org/10.1007/978-3-032-35398-6_13, https://eprint.iacr.org/2026/430
3. Bardet, M., Mora, R., Tillich, J.P.: Polynomial time key-recovery attack on high rate random alternant codes. IEEE Trans. Inf. Theory **70**(6), 4492–4511 (2024). https://doi.org/10.1109/TIT.2023.3334592, https://arxiv.org/abs/2304.14757
4. Briaud, P., Lemoine, A., Randriambololona, H., Tillich, J.P.: A heuristic subexponential attack on the McEliece cryptosystem. Cryptology ePrint Archive, Paper 2026/1232 (2026), https://eprint.iacr.org/2026/1232
5. Brooksbank, P.A., Maglione, J., Wilson, J.B.: Tensor isomorphism by conjugacy of Lie algebras. Journal of Algebra **604**, 790–807 (2022). https://doi.org/10.1016/j.jalgebra.2022.04.014, https://arxiv.org/abs/2005.04046
6. BSI: Notes on recent developments concerning Classic McEliece (September 2026), https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Crypto/Notes_Classic_McEliece.pdf
7. Classic McEliece Team: Classic McEliece: conservative code-based cryptography: guide for security reviewers (October 2022), https://classic.mceliece.org/mceliece-security-20221023.pdf, document of official Classic McEliece web site.
8. Couvreur, A., Mora, R., Tillich, J.P.: A new approach based on quadratic forms to attack the McEliece cryptosystem. In: Guo, J., Steinfeld, R. (eds.) Advances in Cryptology - ASIACRYPT 2023 - 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4-8, 2023, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 14441, pp. 3–38. Springer (2023). https://doi.org/10.1007/978-981-99-8730-6_1, https://eprint.iacr.org/2023/950
9. First, U., Maglione, J., Wilson, J.B.: A spectral theory for transverse tensor operators. arXiv:1911.02518 [math.SP] (2020). https://doi.org/10.48550/arXiv.1911.02518
10. Ghoshal, A., Ishai, Y., Jain, A., Sun, N.: Quasipolynomial cryptanalysis of the McEliece cryptosystem (or: PIR meets McEliece). Cryptology ePrint Archive, Paper 2026/1630 (August 2026), https://eprint.iacr.org/2026/1630
11. Hemmert, T.: Key recovery for the McEliece cryptosystem using higher-order vanishing. Cryptology ePrint Archive, Paper 2026/1339 (2026), https://eprint.iacr.org/2026/1339
12. Hemmert, T., Wiemers, A.: Distinguishing Goppa codes using higher-order vanishing. In: Heninger, N., Rosulek, M. (eds.) Advances in Cryptology – CRYPTO 2026. 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17–20, 2026, Proceedings, Part IV. Lecture Notes in Computer

Science, vol. 16803, pp. 191–223. Springer (2026). https://doi.org/10.1007/978-3-032-35398-6_7, https://eprint.iacr.org/2025/1661
13. ISO: Information technology – security techniques – encryption algorithms – part 2: Asymmetric ciphers – amendment 2. Standard ISO/IEC 18033-2:2006/Amd 2:2026, International Organization for Standardization (2026), https://www.iso.org/standard/86890.html
14. McEliece, R.J.: A public-key cryptosystem based on algebraic coding theory. Progress Report DSN PR 42-44, Deep Space Network (1978), https://ipnpr.jpl.nasa.gov/progress_report2/42-44/44N.PDF
15. Niederreiter, H.: Knapsack-type cryptosystems and algebraic coding theory. Problems of Control and Information Theory **15**(2), 159–166 (1986), http://real-j.mtak.hu/7997/
16. Randriambololona, H.: The syzygy distinguisher. In: Fehr, S., Fouque, P.A. (eds.) Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part VI. Lecture Notes in Computer Science, vol. 15606, pp. 324–354. Springer (2025). https://doi.org/10.1007/978-3-031-91095-1_12, https://eprint.iacr.org/2024/1193
17. Saarinen, M.J.O.: HOVER: Higher-order vanishing endomorphism recovery. Cryptology ePrint Archive, Paper 2026/1778 (August 2026), https://eprint.iacr.org/2026/1778
18. Saarinen, M.J.O.: Two-anchor Holdout/Hermite: Solving the TII-254 McEliece key recovery challenge. Cryptology ePrint Archive, Paper 2026/1986 (September 2026), https://eprint.iacr.org/2026/1986
19. Sidelnikov, V.M., Shestakov, S.O.: On the insecurity of cryptosystems based on generalized Reed-Solomon codes. Discrete Mathematics and Applications **2**(4), 439–444 (1992). https://doi.org/10.1515/dma.1992.2.4.439
20. Weis, S.A.: Improving GIJS key recovery for Classic McEliece. Cryptology ePrint Archive, Paper 2026/1984 (September 2026), https://eprint.iacr.org/2026/1984