Deterministic Partial Algorithms for Matrix-Space Conjugacy over Finite Fields

Jingchuan Ma[a,*], Yanhua Liu[a], Qiaoyun Huang[a]

[a] *Fuzhou University Zhicheng College, Fuzhou, Fujian, China*

Abstract

Given two ordered families of square matrices over a finite field, we seek an ambient conjugation and an invertible mixing of the family that carry one input to the other. The matrices need not be linearly independent. We give deterministic partial algorithms in an explicit finite-field model: source-side recognition certifies completeness against every target, and every returned transporter is verified. For a matrix space with nondegenerate trace pairing, two successive orthogonal projections produce an intrinsic matrix pair. A characteristic-wise nonzero polynomial certificate gives source density at least $1-(6n^2-4)s/q-2/q^3$, where $n \geq 3$ is the matrix size, $2 \leq m \leq n^2-2$ is the number of slices, and $s = \min\{m, n^2-m\}$. Thus the density tends to one whenever $n^2 s = o(q)$. A complementary one-dimensional-hull construction, together with low-dimensional and endpoint solvers, gives a recognized source family of iid-uniform mass greater than $9/(1216q)$ for every prime power q and all positive n, m. Both algorithms have bit complexity polynomial in $n, m, \log q$. An image-and-kernel lift recovers the declared coefficient action, including dependent presentations. We separate the exceptional odd-characteristic two-by-two case and the arithmetic limitations of higher-moment constructions from the universal partial guarantee.

Keywords: matrix spaces, conjugacy, finite fields, deterministic algorithms, average-case analysis, trace geometry

1. Introduction

Let $A = (A_1, \ldots, A_m)$ and $B = (B_1, \ldots, B_m)$ be ordered families in $\mathrm{M}_n(\mathbb{F}_q)$. We consider the constructive problem of finding $(S, T) \in \mathrm{GL}_n(\mathbb{F}_q) \times \mathrm{GL}_m(\mathbb{F}_q)$ such that

$$B_i = \sum_{j=1}^{m} t_{ij} S A_j S^{-1} \qquad (1 \leq i \leq m),$$

or deciding that no such pair exists. At the level of spans, this is conjugacy of matrix spaces; the term *matrix code* is customary when a matrix space is regarded as a linear code in the rank metric. The input representation matters: m is the declared number of matrices, whereas the span may have smaller dimension. Our output includes both transformations, not just an isomorphism invariant or a decision about the spans.

Unlike simultaneous similarity of ordered tuples, this problem permits arbitrary $\mathrm{GL}_m(\mathbb{F}_q)$ mixing: the corresponding matrices are not given. An ordered-tuple canonical form alone does not remove that mixing.

*Corresponding author. ORCID identifiers: Jingchuan Ma, 0009-0001-3703-471X; Yanhua Liu, 0000-0002-6076-9968; Qiaoyun Huang, 0009-0008-2122-7307.

Email addresses: kiciot@qq.com (Jingchuan Ma), Lyhwa@fzu.edu.cn (Yanhua Liu), huangqy@fdzcxy.edu.cn (Qiaoyun Huang)

Prior trace geometry and the missing step. The conjugation-invariant trace hull $H(C) = C \cap C^{\perp_b}$, for $b(X, Y) = \mathrm{Tr}(XY)$, is established in the matrix-code algorithms of Couvreur and Levrat (2025), including the broader left–right equivalence problem. Li et al. (2026) give a randomized bit-polynomial algorithm on a $1/\Theta(q)$ fraction of iid sources when $m = n$, with constructive coset and candidate-verification steps, and ask for arbitrary coefficient dimension at the same scale. We build on their trace geometry and witness recovery; our additional step is the construction of intrinsic matrix covariants.

These covariants transform only by ambient conjugation. If C has zero hull, write P_C for its trace-orthogonal projection and set

$$A_C = P_C(I_n), \qquad M_C = P_C(A_C^2).$$

This pair is intrinsic, with no scalar ambiguity. If the hull is a line $\langle A_0 \rangle$, a transverse trace moment selects a nondegenerate complement of that line. Trace duality then produces weighted intrinsic matrices in the complement. In either case, scalar common centralizer reduces ambient recovery to a one-dimensional intertwiner space. The main probability arguments analyze these constructions under the actual source law rather than substituting independent random matrices.

1.1. Principal guarantees

The field is supplied as $\mathbb{F}_q = \mathbb{F}_p[z]/(g)$, with an irreducible polynomial g included in the input. We do not construct a field representation from the integer q. A partial algorithm first recognizes a family of sources; once a source is recognized, it is complete against every target and never returns `FAILURE`. All probabilities below concern iid-uniform source entries. The algorithms themselves are deterministic.

The zero-hull construction gives the simplest entry point and the high-density guarantee.

Theorem A (High-density matrix-space recovery)**.** *Let q be any prime power, $n \geq 3$, and $2 \leq m \leq n^2 - 2$. Put $s = \min\{m, n^2 - m\}$. There is a deterministic partial algorithm, with bit complexity* $\mathrm{poly}(n, m, \log q)$*, whose recognized source event $\mathsf{Rec}_{n,m}$ satisfies*

$$\Pr[\mathsf{Rec}_{n,m}] \geq 1 - \frac{(6n^2 - 4)s}{q} - \frac{2}{q^3}.$$

For every recognized source and arbitrary target, it constructs and verifies both transformations or returns an exact `NOT_CONJUGATE` answer.

A negative right-hand side is only a vacuous lower bound, not evidence of small-field success. The estimate gives density tending to one when $n^2 \min\{m, n^2 - m\} = o(q)$; in particular, $n^4 = o(q)$ suffices uniformly over the interior range. The proof constructs a nonzero polynomial in the actual iid slice entries separately in every characteristic. Its degree controls the source density, but computing that polynomial is not part of recovery.

A second construction supplies an absolute guarantee even when this large-field estimate is uninformative.

Theorem B (Uniform all-parameter partial guarantee)**.** *For every prime power $q \geq 2$ and $n, m \geq 1$, there is a deterministic* $\mathrm{poly}(n, m, \log q)$*-time partial algorithm and an efficiently recognizable source family $\mathcal{G}_{q,n,m}$ such that*

$$\Pr[\mathcal{G}_{q,n,m}] > \frac{9}{1216q} > \frac{1}{136q}.$$

On each recognized source it is complete against every arbitrary target, returns an invertible ordered witness precisely when one exists, and otherwise returns `NOT_CONJUGATE`.

This includes the binary field, $n = 1$, dependent full-image presentations, and all declared slice counts. A source-only thinning $\mathcal{G}^\star \subseteq \mathcal{G}$ has mass between $1/(136q)$ and $4/q$ (Appendix F). This certificate compares the result with the two-sided source scale in the prior problem statement; it does not increase the recognized family or strengthen the algorithm. The two algorithms can be combined by source-side dispatch, so the high-density and uniform guarantees need not be traded against one another.

Endpoint procedures yield density $1 - O_{n,m}(1/q)$ for fixed $n \geq 3, m \geq 1$, and for fixed $n = 2, m \neq 2$. In characteristic two, a separate projected multiplication procedure recognizes exactly $(1 - q^{-1})(1 - q^{-2})^2$ of iid sources at $n = m = 2$. We do not assert a high-density theorem for odd-characteristic $(2, 2)$. The detailed statements and their dependencies are given in Sections 4, 7, and 8.

1.2. Computational scope

We compute a transporter (S, T), not a canonical form. Recognition and recovery have polynomial worst-case cost, with completeness restricted to recognized sources; no all-source solver or average-running-time theorem for a complete solver is asserted. The characteristic-two dispatch uses a deterministic degree-at-most-three splitter, not a general factoring oracle.

Matrix-space conjugacy also arises in algebraic-isomorphism complexity and matrix-code cryptography (Grochow and Qiao, 2023; Grochow et al., 2022; Couvreur and Levrat, 2025). Our iid-source guarantees do not automatically transfer to reduction outputs or cryptosystem key distributions.

Organization. Sections 2–3 fix the input model, trace geometry, and coefficient lift. Section 4 presents the zero-hull algorithm and its high-density analysis first. Sections 5–7 develop the complementary one-hull construction and uniform dispatch. Section 8 completes the small matrix size. Exact examples and a coarse operation budget are given in Section 9; Section 10 isolates arithmetic boundaries. Appendices retain the geometric witness, formed-space counts, source-scale certificate, and secondary arithmetic proofs.

2. Preliminaries and problem formulation

2.1. Explicit finite fields

Throughout, $q = p^e$ and $\mathbb{F}_q = \mathbb{F}_p[z]/(g)$, where the irreducible polynomial g of degree e is supplied as part of the input. Residue-class coordinates give deterministic field arithmetic and inversion in $\mathrm{poly}(\log q)$ bit time. We count field operations separately; constructing a field representation from the integer q is not part of the problem.

2.2. Matrix spaces and trace geometry

Put $W = \mathrm{M}_n(\mathbb{F}_q)$ and $N = n^2$, with trace pairing

$$b(X, Y) = \mathrm{Tr}(XY). \tag{1}$$

Definition 2.1 (Matrix code, orthogonal, and hull)**.** A *matrix code* is a linear subspace $\mathcal{C} \leq W$. Its orthogonal and trace hull are

$$\mathcal{C}^{\perp_b} = \{Y \in W : b(X, Y) = 0 \text{ for every } X \in \mathcal{C}\}, \qquad H(\mathcal{C}) = \mathcal{C} \cap \mathcal{C}^{\perp_b}.$$

Two codes are conjugate if $\mathcal{D} = S\mathcal{C}S^{-1}$ for some $S \in \mathrm{GL}_n(\mathbb{F}_q)$.

Lemma 2.2 (Trace geometry under conjugation)**.** *Over every finite field, b is symmetric, nondegenerate, and invariant under simultaneous conjugation. If $\mathcal{D} = S\mathcal{C}S^{-1}$, then*

$$\mathcal{D}^{\perp_b} = S\mathcal{C}^{\perp_b}S^{-1}, \qquad H(\mathcal{D}) = SH(\mathcal{C})S^{-1}.$$

If $H(\mathcal{C}) = \langle A_0 \rangle$, then $\mathrm{Tr}(A_0^2) = 0$; every nonzero $B_0 \in H(\mathcal{D})$ has the form $B_0 = \alpha S A_0 S^{-1}$ for a unique $\alpha \in \mathbb{F}_q^\times$.

Proof. The identities $\mathrm{Tr}(XY) = \mathrm{Tr}(YX)$ and $\mathrm{Tr}(SXYS^{-1}) = \mathrm{Tr}(XY)$ give symmetry and invariance. The units E_{ij}, E_{ji} are dual, so the form is nondegenerate. For $Y = SZS^{-1}$, the identity $b(SXS^{-1}, Y) = b(X, Z)$ proves orthogonal covariance; intersection with $\mathcal{D}$ gives hull covariance. Finally $A_0 \in \mathcal{C} \cap \mathcal{C}^{\perp_b}$ gives $b(A_0, A_0) = 0$, and two generators of the conjugate hull line differ by a unique nonzero scalar. □

Proposition 2.3 (Trace-orthogonal projection)**.** *If $H(\mathcal{C}) = 0$, the equations*

$$b(P_{\mathcal{C}}(X), Y) = b(X, Y) \qquad (Y \in \mathcal{C}) \tag{2}$$

define a unique linear projection $P_{\mathcal{C}} : W \to \mathcal{C}$ with $\ker P_{\mathcal{C}} = \mathcal{C}^{\perp_b}$ and $P_{\mathcal{C}}|_{\mathcal{C}} = \mathrm{id}$. Thus $W = \mathcal{C} \oplus \mathcal{C}^{\perp_b}$. For $\mathcal{D} = S\mathcal{C}S^{-1}$,

$$P_{\mathcal{D}}(SXS^{-1}) = SP_{\mathcal{C}}(X)S^{-1}. \tag{3}$$

In any basis $C_1, \ldots, C_d$ of $\mathcal{C}$, write $G = (\mathrm{Tr}(C_i C_j))_{i,j}$ and $h = (\mathrm{Tr}(XC_j))_j$. Then $P_{\mathcal{C}}(X) = \sum_i a_i C_i$, where $Ga = h$.

Proof. The map $\mathcal{C} \to \mathcal{C}^*$ induced by b is an isomorphism because its radical is $H(\mathcal{C}) = 0$. Its inverse gives existence, uniqueness and linearity in (2), as well as the coordinate formula. The same equation identifies the restriction and kernel of $P_{\mathcal{C}}$. Trace invariance shows that $SP_{\mathcal{C}}(X)S^{-1}$ satisfies the defining equations on $\mathcal{D}$; uniqueness proves covariance. □

Spans, kernels, orthogonals, hulls and projections are consequently computed by Gaussian elimination. The one-hull construction will retain the homogeneous weights caused by the scalar ambiguity in A_0.

2.3. Ordered presentations and the source model

An ordered tuple $A = (A_1, \ldots, A_m) \in W^m$ defines

$$\phi_A : \mathbb{F}_q^m \to W, \quad \phi_A(x) = \sum_i x_i A_i, \qquad \mathcal{C}_A = \mathrm{im}\, \phi_A, \quad r_A = \mathrm{rank}\, \phi_A. \tag{4}$$

The declared length m is not replaced by the image rank r_A when the slices are dependent. Our action convention is

$$B_i = \sum_{j=1}^{m} t_{ij} S A_j S^{-1} \qquad (1 \le i \le m), \quad S \in \mathrm{GL}_n(\mathbb{F}_q),\ T \in \mathrm{GL}_m(\mathbb{F}_q). \tag{5}$$

For column coefficient vectors this reads

$$\phi_B = \mathrm{Ad}_S \circ \phi_A \circ T^{\mathsf{T}}, \qquad \mathrm{Ad}_S(X) = SXS^{-1}, \tag{6}$$

since the coefficient of SA_jS^{-1} in $\phi_B(x)$ is $\sum_i x_i t_{ij} = (T^{\mathsf{T}}x)_j$.

The source tuple is iid-uniform in W^m; all mn^2 entries are independent and uniform. The target $B \in W^m$ is arbitrary.

Definition 2.4 (Source-recognizable partial search and decision)**.** A deterministic MCC algorithm is *source-recognizable* if it first evaluates a deterministic predicate depending only on A. If the predicate accepts, then for every target B the algorithm is complete: it returns an invertible pair (S, T) satisfying (5) if and only if such a pair exists, and otherwise returns `NOT_CONJUGATE`. Every returned pair is checked against all m slice equations. The outcome `FAILURE` is permitted only when the source predicate rejects; target-side tests do not convert a recognized source into a failure.

3. Ordered presentations and transporter lifting

We lift an image-code conjugator to the declared coefficient action and then count the iid presentations of each image. This treats dependent and surjective presentations without identifying an ordered generating tuple with a basis.

Put $N = n^2$ and $V = \mathbb{F}_q^m$, retaining the maps and action from Section 2. In column vectorization, set

$$P_A = [\mathrm{vec}(A_1) \ \cdots \ \mathrm{vec}(A_m)] \in \mathbb{F}_q^{N\times m}, \qquad P_{S,A} = [\mathrm{vec}(SA_1S^{-1}) \ \cdots \ \mathrm{vec}(SA_mS^{-1})].$$

The coefficient convention (6) becomes

$$P_B = P_{S,A}T^{\mathsf{T}}. \tag{7}$$

3.1. The ordered-tensor/image-code orbit bridge

Lemma 3.1 (Ordered-tensor/image-code orbit lemma)**.** *For every prime power q, every $n, m \geq 1$, and arbitrary $A, B \in \mathrm{M}_n(\mathbb{F}_q)^m$, the following conditions are equivalent.*

1. *There are $S \in \mathrm{GL}_n(\mathbb{F}_q)$ and $T \in \mathrm{GL}_m(\mathbb{F}_q)$ satisfying (5).*
2. *The ranks satisfy $r_A = r_B$, and there is an $S \in \mathrm{GL}_n(\mathbb{F}_q)$ such that*

$$\mathcal{C}_B = S\mathcal{C}_AS^{-1}.$$

For a fixed S in the second condition, a corresponding T can be constructed by a deterministic algorithm. In the explicit-field model, its bit complexity is polynomial in n, m, and $\log q$. Image conjugacy already implies the displayed rank equality, which records the first rejection test.

Proof. Equation (7) gives necessity because both actions are invertible. Conversely fix S with $\mathcal{C}_B = S\mathcal{C}_AS^{-1}$, and put $P = P_{S,A}$, $Q = P_B$, $r = \operatorname{rank} P = \operatorname{rank} Q$. Gaussian elimination supplies preimage columns $U_A, U_B \in \mathbb{F}_q^{m\times r}$ and kernel bases $K_A, K_B \in \mathbb{F}_q^{m\times(m-r)}$ such that PU_A, QU_B are bases of the common image. Then

$$E_A = [U_A \ K_A], \qquad E_B = [U_B \ K_B] \tag{8}$$

are invertible: applying P or Q to a column relation first kills the preimage coefficients, and the kernel coefficients then vanish. The unique image-basis change $H \in \mathrm{GL}_r(\mathbb{F}_q)$ satisfies

$$QU_B = (PU_A)H. \tag{9}$$

With empty blocks allowed at $r = 0, m$, set

$$R = E_A \begin{pmatrix} H & 0 \\ 0 & I_{m-r} \end{pmatrix} E_B^{-1}. \tag{10}$$

All factors are invertible and $PRE_B = [PU_AH \ 0] = [QU_B \ 0] = QE_B$, so $PR = Q$. Moreover,

$$RK_B = K_A, \qquad R(\ker \phi_B) = \ker \phi_A. \tag{11}$$

Here $\ker P = \ker \phi_A$ since Ad_S is invertible. Therefore

$$T = R^{\mathsf{T}} \tag{12}$$

satisfies every ordered slice equation. All bases and matrices above are computed by deterministic linear algebra. □

Remark 3.2 (Ranks and rectangular regimes). The lift always acts on $\mathbb{F}_q^m$. At rank zero take $T = I_m$; at rank m the kernel blocks are empty and the lift for a fixed S is unique. For $0 < r < m$, replacing I_{m-r} in (10) by any $L \in \mathrm{GL}_{m-r}(\mathbb{F}_q)$ gives $RK_B = K_AL$. No step assumes $m \leq N$. If $m = N = r$, then $R = P^{-1}Q$; if $m > N$ the kernel blocks remain essential. When $r = N$, both images are W, so $S = I_n$ suffices and all surjective presentations lie in one coefficient-mode orbit.

Algorithm 1 Lift a verified ambient conjugator

Require: Ordered tuples A, B with conjugate image codes under a given S
Ensure: An invertible coefficient witness T satisfying all slice equations

1: Form $P = P_{S,A}$ and $Q = P_B$; let their common rank be r
2: Choose $E_A = [U_A\ K_A]$, $E_B = [U_B\ K_B] \in \mathrm{GL}_m(\mathbb{F}_q)$ by elimination
K_A, K_B span $\ker P, \ker Q$; PU_A, QU_B are image bases
3: Solve $PU_AH = QU_B$ for $H \in \mathrm{GL}_r(\mathbb{F}_q)$
4: $R \leftarrow E_A \operatorname{diag}(H, I_{m-r})E_B^{-1}$; $T \leftarrow R^{\mathsf{T}}$
5: Verify $PR = Q$, $\det S \det T \neq 0$, and every ordered slice equation
6: **return** (S, T)

3.2. Deterministic lifting and verification

Proposition 3.3 (Algorithmic coefficient lift)**.** *Given A, B and $S \in \mathrm{GL}_n(\mathbb{F}_q)$, one can either reject the claim $\mathcal{C}_B = S\mathcal{C}_AS^{-1}$ or construct the matrix T from Lemma 3.1. With classical arithmetic, the construction and its verification use*

$$O(mn^3 + n^2m^2 + m^3)$$

field operations.

Proof. Compare the ranks and reduced column spaces of $P_{S,A}$ and P_B. If they agree, apply the construction in Lemma 3.1 and verify invertibility and

$$B_i = \sum_{j=1}^{m} t_{ij} S A_j S^{-1} \qquad (1 \le i \le m). \tag{13}$$

Conjugating m slices costs $O(mn^3)$ field operations; rectangular elimination and slice verification cost $O(n^2m^2)$; adapted-basis inversion and multiplication cost $O(m^3)$. Explicit-field arithmetic then gives the stated bit-polynomial bound. This procedure presupposes an ambient candidate S and does not solve image-code conjugacy. □

3.3. The iid rank law and conditional image distribution

We now return to the ordered iid source model. Since vectorization is a bijection, independent uniform slices make P_A uniform in $\mathbb{F}_q^{N\times m}$.

Proposition 3.4 (Rectangular rank law and uniform conditional image)**.** *Let A be iid-uniform in $\mathrm{M}_n(\mathbb{F}_q)^m$. For every $0 \le r \le \min\{N, m\}$,*

$$\#\{P \in \mathbb{F}_q^{N\times m} : \operatorname{rank}(P) = r\} = \prod_{i=0}^{r-1} \frac{(q^N - q^i)(q^m - q^i)}{q^r - q^i}, \tag{14}$$

and hence

$$\Pr[r_A = r] = q^{-(N-r)(m-r)} \prod_{i=0}^{r-1} \frac{(1 - q^{i-N})(1 - q^{i-m})}{1 - q^{i-r}}. \tag{15}$$

Conditional on $r_A = r$, the image code $\mathcal{C}_A$ is uniform on $\mathrm{Gr}(r, W)$.

Proof. For each fixed $C \in \mathrm{Gr}(r, W)$, the maps with image C are the surjections $\mathbb{F}_q^m \twoheadrightarrow C$, numbering

$$s_{m,r} = \prod_{i=0}^{r-1} (q^m - q^i). \tag{16}$$

Multiplication by the Gaussian coefficient $\begin{bmatrix} N \\ r \end{bmatrix}_q = \prod_{i=0}^{r-1}(q^N - q^i)/(q^r - q^i)$ gives (14); division by q^{Nm} gives (15). The fibre size $s_{m,r}$ is independent of C, hence

$$\Pr[\mathcal{C}_A = C \mid r_A = r] = \frac{1}{\begin{bmatrix} N \\ r \end{bmatrix}_q}. \tag{17}$$

These are generating m-tuples, and are bases only when $m = r$. At $r = 0$ all products are empty and the unique tuple is zero, with probability q^{-Nm}. □

The injective and surjective specializations are, respectively, for $m \le N$,

$$\Pr[r_A = m] = \prod_{j=N-m+1}^{N} (1 - q^{-j}), \tag{18}$$

whereas for $m \ge N$,

$$\Pr[r_A = N] = \prod_{j=m-N+1}^{m} (1 - q^{-j}). \tag{19}$$

At $m = N$ the two formulas coincide.

Corollary 3.5 (Exact probability transfer)**.** *For any $0 \le r \le \min\{N, m\}$ and any family $\mathcal{E}_r \subseteq \mathrm{Gr}(r, W)$,*

$$\Pr[r_A = r,\ \mathcal{C}_A \in \mathcal{E}_r] = \Pr[r_A = r]\, \frac{|\mathcal{E}_r|}{\begin{bmatrix} N \\ r \end{bmatrix}_q}. \tag{20}$$

Proof. Apply the probability chain rule and then (17). □

Remark 3.6. Equation (20) is an exact conditional calculation, not an assertion that rank and image are independent. It also remains a statement about the original ordered iid tuple: the coefficient dimension is m on both sides, even when $r < m$.

4. The zero-hull projection algorithm and high-density sources

A nondegenerate trace form on the input space permits a particularly simple construction: project the identity, square the result, and project again. We first specify and recover from this pair, then bound the density of its recognition event. The proof separates a geometric witness from its polynomial certificate and from the actual iid source law.

Throughout the section, put

$$W = \mathrm{M}_n(\mathbb{F}_q), \qquad N = n^2, \qquad b(X, Y) = \mathrm{Tr}(XY).$$

4.1. The intrinsic projection pair

Let $E \le W$ have zero trace hull. Proposition 2.3 then gives the trace-orthogonal projection $P_E : W \to E$. Define

$$A_E = P_E(I_n), \qquad M_E = P_E(A_E^2). \tag{21}$$

Definition 4.1 (Projection-pair good code)**.** An r-dimensional code $E \le W$ is *projection-pair good* if

$$H(E) = 0 \qquad \text{and} \qquad \mathrm{Cent}(A_E, M_E) = \mathbb{F}_q I_n. \tag{22}$$

For brevity, we call such a code ZP-good.

Both matrices in (21), including their scales, are intrinsic. In particular, no basis of E, projective generator, eigenvalue, or polynomial factor is selected. This gives literal covariance rather than covariance up to an unknown scalar.

Proposition 4.2 (Covariance, complementation, and recognition)**.** *Let $E \le W$ be ZP-good and let $S \in \mathrm{GL}_n(\mathbb{F}_q)$. For $F = SES^{-1}$ one has*

$$A_F = SA_ES^{-1}, \qquad M_F = SM_ES^{-1}. \tag{23}$$

Moreover, $E^{\perp_b}$ is ZP-good and

$$A_{E^{\perp_b}} = I_n - A_E, \tag{24}$$

$$M_{E^{\perp_b}} = I_n - A_E + A_E^2 - M_E, \tag{25}$$

$$\mathrm{Cent}(A_{E^{\perp_b}}, M_{E^{\perp_b}}) = \mathrm{Cent}(A_E, M_E). \tag{26}$$

Given a basis of E, ZP-goodness and the pair (21) *are computable by deterministic finite-field linear algebra.*

Proof. Equation (23) follows from Proposition 2.3, first with I_n and then with A_E^2. Since $H(E) = 0$, the ambient space is the orthogonal direct sum $W = E \oplus E^{\perp_b}$ and $P_{E^{\perp_b}} = \mathrm{id}_W - P_E$. This gives $A_{E^{\perp_b}} = I_n - A_E$. Applying the same identity to $(I_n - A_E)^2$ and using $P_E(I_n) = P_E(A_E) = A_E$ yields

$$P_{E^{\perp_b}}((I_n - A_E)^2) = I_n - A_E + A_E^2 - M_E,$$

in every characteristic. Commuting with $I_n - A_E$ is equivalent to commuting with A_E; under that condition, commuting with the second matrix in (25) is equivalent to commuting with M_E. This proves (26).

For recognition, form the trace Gram matrix of any basis of E. Its invertibility is equivalent to $H(E) = 0$. Two solves with that matrix compute $P_E(I_n)$ and $P_E(A_E^2)$. The stacked equations

$$XA_E = A_EX, \qquad XM_E = M_EX$$

then compute the common centralizer, whose dimension is one exactly when it is $\mathbb{F}_qI_n$. All tests are basis independent by the uniqueness of the orthogonal projection. □

For a full-rank m-slice image code C, the interior recognizer uses the smaller of C and its orthogonal complement:

$$s = \min\{m, N - m\}, \qquad E(C) = \begin{cases} C, & m \le N/2, \\ C^{\perp_b}, & m > N/2. \end{cases} \tag{27}$$

The second branch is computed from the same source code C. It is not replaced, in either the algorithm or the probability proof, by an independent tuple of s random matrices.

For an ordered m-slice source A with image code $C = \mathcal{C}_A$, define

$$\mathsf{Rec}_{n,m} = \{r_A = m \text{ and } E(C) \text{ is ZP-good}\}, \tag{28}$$

Here r_A is the coefficient-map rank. This recognition condition is separate from the probability estimate proved below.

4.2. Deterministic recovery against an arbitrary target

The polynomial certificate is used only for counting. The algorithm tests the full Gram determinant and the full common-centralizer dimension on the actual source. Recovery then consists of one simultaneous linear transporter solve.

Proposition 4.3 (Projection-pair recovery and ordered lift)**.** *Let $n \ge 3$, $2 \le m \le N - 2$, and let A satisfy $\mathsf{Rec}_{n,m}$. There is a deterministic $\mathrm{poly}(n, m, \log q)$-time algorithm satisfying Definition 2.4 against every ordered target $B \in \mathrm{M}_n(\mathbb{F}_q)^m$.*

Algorithm 2 Zero-hull matrix-space recovery

Require: Bases of $C, D \le \mathrm{M}_n(\mathbb{F}_q)$ of equal dimension d
Ensure: A verified conjugator, `NOT_CONJUGATE`, or source-side `FAILURE`

Let $E = C$ if $d \le n^2/2$; otherwise let $E = C^{\perp_b}$
Compute the trace Gram matrix of E; if it is singular, return `FAILURE`
Compute $A_E = P_E(I_n)$ and $M_E = P_E(A_E^2)$ by Gram solves
If $\dim \mathrm{Cent}(A_E, M_E) \neq 1$, return `FAILURE`
Form $F = D$ or $D^{\perp_b}$ using the same dimension branch
If its Gram matrix is singular, return `NOT_CONJUGATE`
Compute A_F, M_F and $J = \{Z : ZA_E = A_F Z,\ ZM_E = M_F Z\}$
If $\dim J \neq 1$, return `NOT_CONJUGATE`
Let S be its deterministic nonzero basis matrix
If S is singular or $D \neq SCS^{-1}$, return `NOT_CONJUGATE`
return S

Proof. Compute the target coefficient rank and reject if it is not m. Put $C = \mathcal{C}_A$, $D = \mathcal{C}_B$, and form

$$E = E(C), \qquad F = \begin{cases} D, & m \le N/2, \\ D^{\perp_b}, & m > N/2. \end{cases}$$

Reject if $H(F) \neq 0$, compute (A_E, M_E) and (A_F, M_F), and solve

$$XA_E = A_F X, \qquad XM_E = M_F X. \tag{29}$$

Return `NOT_CONJUGATE` if this solution space is not a line or if a nonzero generator of the line is singular. If $D = S_0 C S_0^{-1}$, covariance and Lemma 2.2 give $F = S_0 E S_0^{-1}$, and

$$\mathrm{Sol}\,(29) = S_0\,\mathrm{Cent}(A_E, M_E) = S_0(\mathbb{F}_q I_n). \tag{30}$$

Thus a genuine conjugator makes the solution space exactly one-dimensional and every nonzero element invertible. Conversely, if the solution space has dimension one with invertible generator S, verify the complete identity $D = SCS^{-1}$. Failure of this identity rejects a collision of the marked pairs and cannot discard a different pair transporter, because every solution of (29) is a scalar multiple of S. Orthogonal complementation is involutive, so conjugacy of E and F is equivalent to conjugacy of the original codes C and D.

After the exact code check, apply Proposition 3.3. Its kernel-aware construction returns R and $T = R^{\mathsf{T}}$ and verifies (5). Here the source has full column rank; the common lift keeps the same action convention as all other branches.

All operations are ranks, orthogonal-complement and Gram solves, commutator and intertwiner linear systems, matrix inversion, row-space comparison, and the coefficient lift. In the explicit-field model of Section 2, their bit complexity is $\mathrm{poly}(n, m, \log q)$. No root extraction, finite-field enumeration, or polynomial factorization is invoked. □

The code-level procedure in Algorithm 2 uses the complement only to reduce matrix dimensions in the Gram solves. The ordered-source recognizer $\mathsf{Rec}_{n,m}$ additionally requires full coefficient rank; this requirement is part of the probability model, not of code-level covariance.

Example 4.4 (Two intrinsic projections over $\mathbb{F}_5$). Let $E = \langle X, Y \rangle \le \mathrm{M}_3(\mathbb{F}_5)$, where

$$X = \begin{pmatrix} 1 & 1 & 3 \\ 1 & 4 & 1 \\ 3 & 2 & 4 \end{pmatrix}, \qquad Y = \begin{pmatrix} 0 & 1 & 1 \\ 2 & 4 & 2 \\ 4 & 2 & 0 \end{pmatrix}.$$

In this basis the trace Gram matrix is $\mathrm{diag}(2,1)$ and $(\operatorname{Tr} X, \operatorname{Tr} Y) = (4,4)$. Hence the first solve gives $A_E = 2X + 4Y$. The second gives $M_E = 3X + 3Y$; explicitly,

$$A_E = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 4 & 0 \\ 2 & 2 & 3 \end{pmatrix}, \qquad M_E = \begin{pmatrix} 3 & 1 & 2 \\ 4 & 4 & 4 \\ 1 & 2 & 2 \end{pmatrix}.$$

The stacked commutator system has rank eight in nine unknowns, so $\mathrm{Cent}(A_E, M_E) = \mathbb{F}_5 I_3$. Thus the code is recognized even though the large-field lower bound below is vacuous at $q = 5$. This is a verified example, not a lower-bound estimate from sampling. Section 9 continues it through a dependent ordered presentation and full recovery.

4.3. Geometric nonemptiness in every characteristic

We next prove that the predicate in Definition 4.1 is a genuine nonempty open condition for every interior code dimension when $n \geq 3$. The statement is geometric: the witness may live over an extension of the prime field.

Proposition 4.5 (All-characteristic geometric witnesses)**.** *Fix $n \geq 3$ and $2 \leq r \leq N - 2$. In every characteristic, the Grassmannian of r-subspaces of M_n contains, after extension of the prime field, a ZP-good code. No assertion is made that every small $\mathbb{F}_q$ contains a rational witness.*

Proof. The full characteristic-wise construction is given in Appendix A. It constructs an intrinsic pair with scalar centralizer and then completes its span to a nondegenerate code of every permitted dimension. □

4.4. An actual-coordinate bad-locus certificate

We now convert Proposition 4.5 into a polynomial bound in the entries sampled by the iid source model. Fix $n \geq 3$ and $2 \leq r \leq N - 2$, and let $X_1, \ldots, X_r$ be formal $n \times n$ matrices. These Nr entries are the actual affine coordinates of an ordered r-tuple. Set

$$G_{ij} = \operatorname{Tr}(X_i X_j), \qquad g = \det G, \qquad t_i = \operatorname{Tr}(X_i), \tag{31}$$

$$\widetilde{a} = \operatorname{adj}(G)t, \qquad \widetilde{A} = \sum_i \widetilde{a}_i X_i, \tag{32}$$

$$\widetilde{s}_i = \operatorname{Tr}(\widetilde{A}^2 X_i), \qquad \widetilde{d} = \operatorname{adj}(G)\widetilde{s}, \qquad \widetilde{M} = \sum_i \widetilde{d}_i X_i. \tag{33}$$

On $g \neq 0$, the code $E = \langle X_1, \ldots, X_r \rangle$ has zero hull, and the two Gram solves give exactly

$$A_E = \frac{\widetilde{A}}{g}, \qquad M_E = \frac{\widetilde{M}}{g^3}. \tag{34}$$

The homogeneous total degrees in the actual slice-entry variables are

$$\deg g = 2r, \qquad \deg \widetilde{A} = 2r, \qquad \deg \widetilde{M} = 6r. \tag{35}$$

Indeed, $\operatorname{adj}(G)t$ has degree $2r - 1$, while $\widetilde{s}$, $\operatorname{adj}(G)\widetilde{s}$, and $\widetilde{M}$ have degrees $4r + 1$, $6r - 1$, and $6r$, respectively.

To test the centralizer after clearing denominators, use the fixed complement

$$V_0 = \{Z \in \mathrm{M}_n : z_{nn} = 0\}, \qquad \mathrm{M}_n = \mathbb{F}_q I_n \oplus V_0,$$

which remains valid even when the characteristic divides n. For matrices A, M, define

$$\Phi_{A,M} : V_0 \longrightarrow W \oplus W, \qquad Z \longmapsto ([Z, A], [Z, M]). \tag{36}$$

Then $\mathrm{Cent}(A, M) = \mathbb{F}_q I_n$ if and only if $\operatorname{rank} \Phi_{A,M} = N - 1$. Let Δ_J be an $(N-1)$-row maximal minor of the matrix of $\Phi_{\widetilde{A},\widetilde{M}}$. If it uses a rows from the $\widetilde{A}$ block, then

$$\deg \Delta_J \leq 2ra + 6r(N - 1 - a) \leq 6r(N - 1). \tag{37}$$

Proposition 4.6 (Characteristic-wise hypersurface certificate)**.** *For every prime characteristic p, every $n \geq 3$, and every $2 \leq r \leq N-2$, one can select a maximal minor $\Delta_J = \Delta_{J(p,n,r)}$ such that*

$$F_{p,n,r} = g\Delta_J \tag{38}$$

is a nonzero polynomial over $\mathbb{F}_p$ of degree at most

$$\deg F_{p,n,r} \leq (6N-4)r = (6n^2-4)r. \tag{39}$$

The failure of ZP-goodness for the span of the ordered tuple is contained in $\{F_{p,n,r} = 0\}$; equality need not hold. The selected minor may depend on the characteristic, but the recognizer in Definition 4.1 does not.

Proof. Choose an ordered basis of a geometric witness from Proposition 4.5. At that point $g \neq 0$ and the reduced commutator map has rank $N-1$. By (34), the cleared pair specializes to (gA_E, g^3M_E); nonzero rescaling of its two entries does not change the common centralizer. Some maximal minor Δ_J is therefore nonzero at the witness. It follows that $g\Delta_J$ is a nonzero formal polynomial over $\mathbb{F}_p$. Equations (35) and (37) give

$$\deg(g\Delta_J) \leq 2r + 6r(N-1) = (6N-4)r.$$

If $g = 0$, the certificate vanishes. If $g \neq 0$ but the common centralizer is larger than the scalars, every maximal minor vanishes. Thus every failure is in the zero set of the selected certificate. Conversely, the selected minor can vanish while another maximal minor is nonzero, so only containment, not equality, is justified. □

For an iid-uniform ordered r-tuple, Schwartz–Zippel applied to the actual Nr slice entries gives the explicitly truncated estimate

$$\Pr[F_{p,n,r} \neq 0] \geq \left[1 - \frac{(6N-4)r}{q}\right]_+, \qquad [x]_+ = \max\{x, 0\}. \tag{40}$$

For small fields the displayed lower bound can be zero. The exact recognizer remains valid there, but neither geometric nonemptiness nor (40) asserts a rational witness in every small field.

4.5. Transfer to the original ordered iid source

For $0 \leq k \leq N$, write

$$R_{N,k}(q) = \prod_{i=0}^{k-1}(1-q^{i-N}) = \prod_{j=N-k+1}^{N}(1-q^{-j}). \tag{41}$$

This is the exact probability that k iid vectors in the N-space W are linearly independent. Let $\eta_{n,r}(q)$ denote the fraction of r-spaces in $\mathrm{Gr}(r, W)$ that are ZP-good. Since every r-space has the same number of ordered bases, Proposition 3.4 and (40) imply

$$R_{N,r}(q)\eta_{n,r}(q) \geq \left[1 - \frac{(6N-4)r}{q}\right]_+, \tag{42}$$

and hence the same right-hand side divided by $R_{N,r}(q)$ is a lower bound for $\eta_{n,r}(q)$.

Proposition 4.7 (Actual ordered-source probability)**.** *Let $n \geq 3$ and $2 \leq m \leq N-2$, and put $s = \min\{m, N-m\}$. For an iid-uniform ordered source $A = (A_1, \ldots, A_m)$, let $C = \mathcal{C}_A$ and use the event* (28)*. Then*

$$\Pr[\mathsf{Rec}_{n,m}] \geq \frac{R_{N,m}(q)}{R_{N,s}(q)}\left[1 - \frac{(6N-4)s}{q}\right]_+. \tag{43}$$

In particular, for fixed n and m,

$$\Pr[\mathsf{Rec}_{n,m}] \geq 1 - \frac{(6n^2-4)s}{q} - \frac{q^{-(N-m+1)}}{1-q^{-1}} = 1 - O_{n,m}(q^{-1}) \tag{44}$$

for all sufficiently large q.

Proof. Conditional on $r_A = m$, the image C is exactly uniform on $\mathrm{Gr}(m, W)$ by Proposition 3.4. If $m \le N/2$, then $E(C) = C$. If $m > N/2$, orthogonal complementation is a bijection

$$\mathrm{Gr}(m, W) \longrightarrow \mathrm{Gr}(N - m, W),$$

so $E(C) = C^{\perp_b}$ is exactly uniform on $\mathrm{Gr}(s, W)$. Consequently, in both cases,

$$\Pr[\mathsf{Rec}_{n,m}] = R_{N,m}(q)\eta_{n,s}(q). \tag{45}$$

Combining this identity with (42) at dimension s proves (43). This argument transfers a Grassmannian fraction through a bijection; it never represents $C^{\perp_b}$ by independent random slices.

For the asymptotic form, the elementary product inequality gives

$$1 - R_{N,m}(q) \le \sum_{j=N-m+1}^{N} q^{-j} < \frac{q^{-(N-m+1)}}{1 - q^{-1}}.$$

Since $R_{N,s}(q) \le 1$, the right-hand side of (43) is at least $R_{N,m}(q)[1 - (6N - 4)s/q]_+$. Expanding this product for sufficiently large q proves (44). The notation $O_{n,m}$ fixes both n and m; the explicit estimate also yields the following consequence. □

Corollary 4.8 (High density with growing parameters)**.** *For every prime power $q \ge 2$, every $n \ge 3$, and every $2 \le m \le n^2 - 2$, the event of* (28) *satisfies*

$$\Pr[\mathsf{Rec}_{n,m}] \ge 1 - \frac{(6n^2 - 4)\min\{m, n^2 - m\}}{q} - \frac{2}{q^3}. \tag{46}$$

Consequently, let $n = n(q)$ and $m = m(q)$ vary with q in this interior range. If

$$n(q)^2 \min\{m(q), n(q)^2 - m(q)\} = o(q), \tag{47}$$

then the recognized iid source proportion tends to one. In particular, $n(q)^4 = o(q)$ is sufficient uniformly over all interior $m(q)$. The same deterministic algorithm and arbitrary-target output guarantees apply, with bit complexity $\mathrm{poly}(n, m, \log q)$.

Proof. Put $N = n^2$, $s = \min\{m, N - m\}$, and $a = (6N - 4)s$. The rank-product estimate and $N - m + 1 \ge 3$ give

$$1 - R_{N,m}(q) < \frac{q^{-(N-m+1)}}{1 - q^{-1}} \le \frac{2}{q^3}.$$

Since $R_{N,s}(q) \le 1$, (43) implies $\Pr[\mathsf{Rec}_{n,m}] \ge R_{N,m}(q)[1 - a/q]_+$. If $a/q \ge 1$, the right-hand side of (46) is nonpositive, so the claimed bound is automatic. Otherwise,

$$R_{N,m}(q)(1 - a/q) = 1 - a/q - (1 - R_{N,m}(q))(1 - a/q) \ge 1 - a/q - (1 - R_{N,m}(q)),$$

which proves (46). Under (47), both error terms tend to zero. The uniform sufficient condition follows from $s \le n^2/2$, hence $a \le 3n^4$. Recovery is the existing procedure of Proposition 4.3; no new recognizer or field operation is introduced. □

For example, if λ is a parameter with $\log q = \Theta(\lambda)$ and n bounded by a polynomial in λ, this sufficient condition holds for every interior m. In the supplied-field model, the same algorithm then also has bit complexity polynomial in λ.

For comparison with the interior bound, the source predicates already used at the three endpoints have exact iid masses

$$m = 1: \qquad \Pr[\mathrm{Tr}(A_1) \neq 0] = 1 - q^{-1}, \tag{48}$$

$$m = N-1: \quad \Pr[r_A = N-1 \text{ and the normal line is trace-good}] = \prod_{j=1}^{N-1}(1-q^{-j}), \tag{49}$$

$$m \geq N: \qquad \Pr[r_A = N] = \prod_{j=m-N+1}^{m}(1-q^{-j}). \tag{50}$$

For (49), the full-rank probability is $R_{N,N-1}(q)$ and a uniform hyperplane has a trace-good normal with fraction $q^{N-1}(q-1)/(q^N-1)$; multiplying cancels the factor $1-q^{-N}$. Each mass in (48)–(50) is $1 - O_{n,m}(q^{-1})$ for fixed n, m.

5. One-dimensional-hull covariants and recovery

This section develops the intrinsic data used when the trace hull is a line. The construction has three parts. A transverse trace functional first splits the code into its hull and a nondegenerate complement. A Riesz representative then records a second trace functional on that complement. Finally, a simultaneous intertwiner of the hull generator and its Riesz representative recovers the ambient conjugator whenever their common centralizer is scalar. All statements below track the projective choice of hull generator explicitly. We write $d = \dim \mathcal{C}$ and reserve r, s for moment indices.

Recall that a matrix $Z \in \mathrm{M}_n(\mathbb{F}_q)$ is *cyclic* when its minimal polynomial has degree n. Equivalently, the matrices $I_n, Z, \ldots, Z^{n-1}$ are linearly independent. Vectorizing these powers therefore tests cyclicity by Gaussian elimination, without polynomial factorization.

5.1. Transverse complements and Riesz representatives

Assume that

$$H(\mathcal{C}) = \langle A_0 \rangle$$

for a nonzero $A_0 \in \mathcal{C}$. For every integer $r \geq 0$, define the linear functional and its value on the hull generator by

$$\ell_r^{A_0}(X) = \mathrm{Tr}(A_0^r X), \qquad \mu_r(A_0) = \ell_r^{A_0}(A_0) = \mathrm{Tr}(A_0^{r+1}). \tag{51}$$

Here and below $A_0^0 = I_n$. We call s *transverse* if $\mu_s(A_0) \neq 0$, and then put

$$\mathcal{C}_0 = \ker(\ell_s^{A_0}|_{\mathcal{C}}). \tag{52}$$

When at least one such index lies in $\{0, \ldots, n-1\}$, we write

$$s(A_0) = \min\{s \in \{0, \ldots, n-1\} : \mu_s(A_0) \neq 0\}$$

for the least transverse index.

Lemma 5.1 (Transverse decomposition)**.** *If s is transverse, then*

$$\mathcal{C} = \langle A_0 \rangle \oplus \mathcal{C}_0,$$

and the restriction of $b(X,Y) = \mathrm{Tr}(XY)$ to $\mathcal{C}_0$ is nondegenerate.

Proof. Equation (51) gives $\ell_s^{A_0}(A_0) = \mu_s(A_0) \neq 0$. Thus $\ell_s^{A_0}|_{\mathcal{C}}$ is nonzero, $\mathcal{C}_0$ has codimension one in $\mathcal{C}$, and $A_0 \notin \mathcal{C}_0$. This proves the direct-sum assertion.

Let $Y \in \mathcal{C}_0$ be orthogonal to $\mathcal{C}_0$. Since $A_0 \in H(\mathcal{C}) \subseteq \mathcal{C}^{\perp_b}$, it is also true that $b(Y, A_0) = 0$. The direct sum then implies $b(Y, \mathcal{C}) = 0$, so $Y \in H(\mathcal{C}) = \langle A_0 \rangle$. Its membership in $\mathcal{C}_0$ forces $Y = 0$. Hence $b|_{\mathcal{C}_0}$ has zero radical. □

Nondegeneracy gives a basis-independent Riesz representative for every further moment functional.

Definition 5.2 (Riesz representative)**.** For a transverse s and any $r \geq 0$, let $M_r = M_r(\mathcal{C}, A_0; s)$ be the unique element of $\mathcal{C}_0$ satisfying

$$\operatorname{Tr}(M_r X) = \operatorname{Tr}(A_0^r X) \qquad \text{for every } X \in \mathcal{C}_0. \tag{53}$$

The definition uses the subspace $\mathcal{C}$, not an ordered basis of it. For an explicit computation, choose a basis $E_1, \ldots, E_{d-1}$ of $\mathcal{C}_0$, form

$$G_{ij} = \operatorname{Tr}(E_i E_j), \qquad h_i = \operatorname{Tr}(A_0^r E_i),$$

and solve $Gc = h$. Lemma 5.1 makes G nonsingular, and $M_r = \sum_j c_j E_j$. A change of basis changes this linear system but not its unique solution in $\mathcal{C}_0$.

5.2. The quotient projection identity

The probability argument uses a projection in a nondegenerate quotient, not an independently sampled auxiliary matrix. Since $A_0 \in \mathcal{C}^{\perp_b}$, one has $\mathcal{C} \leq A_0^{\perp_b}$; since A_0 is isotropic, define

$$Q_{A_0} = A_0^{\perp_b}/\langle A_0 \rangle, \qquad U = \mathcal{C}/\langle A_0 \rangle \leq Q_{A_0}, \tag{54}$$

and let $\pi_{A_0} : A_0^{\perp_b} \to Q_{A_0}$ be the quotient map. The ambient nondegeneracy of b gives $\operatorname{rad}(b|_{A_0^{\perp_b}}) = \langle A_0 \rangle$, so the induced form $b_{Q_{A_0}}$ on Q_{A_0} is nondegenerate. The induced form on U is also nondegenerate because $H(\mathcal{C}) = \langle A_0 \rangle$. Consequently the orthogonal projection $P_U : Q_{A_0} \to U$ is well defined.

For transverse s, set

$$v_{r,s}(A_0) = A_0^r - \frac{\mu_r(A_0)}{\mu_s(A_0)} A_0^s. \tag{55}$$

The vector $v_{r,s}(A_0)$ lies in $A_0^{\perp_b}$, since its pairing with A_0 is zero.

Lemma 5.3 (Exact quotient identity)**.** *For every transverse s and every $r \geq 0$,*

$$\pi_{A_0}(M_r) = P_U(\pi_{A_0}(v_{r,s}(A_0))). \tag{56}$$

Proof. The direct sum in Lemma 5.1 identifies $\mathcal{C}_0$ with U. Given $u \in U$, let $X_0 \in \mathcal{C}_0$ be its unique representative. Equations (52) and (53) give

$$\begin{aligned} b_{Q_{A_0}}(\pi_{A_0}(M_r), u) &= \operatorname{Tr}(A_0^r X_0) \\ &= \operatorname{Tr}(A_0^r X_0) - \frac{\mu_r(A_0)}{\mu_s(A_0)} \operatorname{Tr}(A_0^s X_0) \\ &= b_{Q_{A_0}}(\pi_{A_0}(v_{r,s}(A_0)), u). \end{aligned}$$

Since $\pi_{A_0}(M_r)$ belongs to U, uniqueness of orthogonal projection in the nondegenerate subspace U proves (56). □

The nonzero object needed to apply the marked-projection argument is the class of $v_{r,s}(A_0)$, rather than its projection onto U.

Lemma 5.4 (Nonzero marked class)**.** *Suppose that A_0 is cyclic, $0 \leq r, s < n$, and s is transverse. Then*

$$\pi_{A_0}(v_{r,s}(A_0)) = 0 \quad \Longleftrightarrow \quad r \in \{1, s\}. \tag{57}$$

In particular, for $n \geq 3$ the choices

$$(s, r) = (0, 2), \qquad \text{or} \qquad r = 0 \text{ when } s > 0,$$

give a nonzero marked class.

Proof. The hull condition gives $\mu_1(A_0) = \mathrm{Tr}(A_0^2) = 0$, so a transverse index cannot equal one. If $r = 1$, then $v_{r,s}(A_0) = A_0$, while if $r = s$ the marked vector is zero. Both classes therefore vanish.

Conversely, a vanishing class gives a relation

$$A_0^r - \frac{\mu_r(A_0)}{\mu_s(A_0)} A_0^s - \lambda A_0 = 0$$

for some $\lambda \in \mathbb{F}_q$. If r, s, and 1 are distinct, the corresponding polynomial is nonzero and has degree less than n. This contradicts cyclicity, which makes the minimal polynomial of A_0 have degree n. This proves (57). When $s > 0$ is the least transverse index, $\mu_0(A_0) = \mathrm{Tr}(A_0) = 0$, so $v_{0,s}(A_0) = I_n$; its class could vanish only if A_0 were scalar. □

Lemma 5.4 does not say that $\pi_{A_0}(M_r)$ is always nonzero. The projection in (56) can vanish for some U, and such subspaces belong to the exceptional event controlled in the source-density analysis.

5.3. Scaling and conjugacy covariance

The following identities account for the projective ambiguity in A_0.

Proposition 5.5 (Generator weights and combined covariance)**.** *Let $\beta \in \mathbb{F}_q^\times$. Rescaling the source hull generator gives*

$$\ell_r^{\beta A_0} = \beta^r \ell_r^{A_0}, \qquad \mu_r(\beta A_0) = \beta^{r+1} \mu_r(A_0), \tag{58}$$

$$\ker(\ell_s^{\beta A_0}|_{\mathcal{C}}) = \mathcal{C}_0, \qquad M_r(\mathcal{C}, \beta A_0; s) = \beta^r M_r(\mathcal{C}, A_0; s), \tag{59}$$

$$v_{r,s}(\beta A_0) = \beta^r v_{r,s}(A_0). \tag{60}$$

The least transverse index, when it exists, is unchanged. Moreover $Q_{\beta A_0} = Q_{A_0}$, U and P_U are unchanged, and

$$\mathrm{Cent}(\beta A_0, \beta^r M_r) = \mathrm{Cent}(A_0, M_r).$$

Now let $\mathcal{D} = S\mathcal{C}S^{-1}$ and choose a target hull generator

$$B_0 = \alpha S A_0 S^{-1}, \qquad \alpha \in \mathbb{F}_q^\times. \tag{61}$$

Let $\mathcal{D}_0 = \ker(\ell_s^{B_0}|_{\mathcal{D}})$ and let $N_r = M_r(\mathcal{D}, B_0; s)$. Then

$$\ell_r^{B_0}(SXS^{-1}) = \alpha^r \ell_r^{A_0}(X), \qquad \mu_r(B_0) = \alpha^{r+1} \mu_r(A_0), \tag{62}$$

$$\mathcal{D}_0 = S\mathcal{C}_0 S^{-1}, \qquad N_r = \alpha^r S M_r S^{-1}, \tag{63}$$

$$v_{r,s}(B_0) = \alpha^r S v_{r,s}(A_0) S^{-1}. \tag{64}$$

In particular, $s(B_0) = s(A_0)$ whenever the least transverse index is defined. Let $\overline{\mathrm{Ad}}_S : Q_{A_0} \to Q_{B_0}$ be the isometry induced by Ad_S, write π_{B_0} for the target quotient map, and put $V = \mathcal{D}/\langle B_0 \rangle$. Then

$$\pi_{B_0}(v_{r,s}(B_0)) = \alpha^r \overline{\mathrm{Ad}}_S(\pi_{A_0}(v_{r,s}(A_0))), \tag{65}$$

$$P_V \circ \overline{\mathrm{Ad}}_S = \overline{\mathrm{Ad}}_S \circ P_U. \tag{66}$$

Proof. Equations (58) follow by expanding powers of βA_0. Multiplication of a functional by the nonzero scalar β^s does not change its kernel. Uniqueness in Definition 5.2 then gives (59). The ratio μ_r/μ_s has weight $r - s$, which proves (60). The hull line and its orthogonal are unchanged by rescaling, and multiplying either member of a centralizer pair by a nonzero scalar does not change the common centralizer.

For the combined identities, trace invariance gives

$$\mathrm{Tr}(B_0^r S X S^{-1}) = \alpha^r \mathrm{Tr}(A_0^r X).$$

This proves (62) and maps the two transverse kernels to one another. Applying the Riesz defining equation after this change of variables proves (63). Substitution in (55) proves (64). Finally, Ad_S preserves the trace form, so the induced quotient map is an isometry. Orthogonal-projection covariance follows from uniqueness. □

If A_0 is replaced by βA_0 while B_0 is fixed, the scale in (61) becomes α/β. Hence

$$(\alpha/\beta)^r M_r(\mathcal{C}, \beta A_0; s) = \alpha^r M_r(\mathcal{C}, A_0; s),$$

so the pair-intertwiner equations do not depend on the chosen source generator. Similarly, replacing B_0 by γB_0 multiplies N_r by γ^r, sends every candidate scale α to $\gamma\alpha$, and leaves the corresponding pair-intertwiner spaces unchanged.

The pair (A_0, M_r) is therefore intrinsic only with its displayed homogeneous weights. It is not a canonical representative of the code: Proposition 5.5 compares two supplied projective hull frames but does not select a distinguished frame from an orbit.

5.4. Recovery from a scalar common centralizer

For a selected transverse pair (s, r), the source recognizer explicitly tests

$$\mathrm{Cent}(A_0, M_r) := \{Z \in \mathrm{M}_n(\mathbb{F}_q) : ZA_0 = A_0Z,\ ZM_r = M_rZ\} = \mathbb{F}_q I_n. \tag{67}$$

Cyclicity alone does not imply (67); it only identifies $\mathrm{Cent}(A_0)$ with $\mathbb{F}_q[A_0]$. The scalar-pair condition is a separate, deterministic linear-algebra test.

Proposition 5.6 (Recovery with a complete scale list)**.** *Suppose that a source code $\mathcal{C}$ has one-dimensional hull generated by a cyclic A_0, that $0 \le r, s < n$, that s is transverse, and that* (67) *holds for the selected r. Let $\mathcal{D} \le \mathrm{M}_n(\mathbb{F}_q)$ be arbitrary.*

After rejecting necessary invariant mismatches, suppose that $H(\mathcal{D}) = \langle B_0 \rangle$ and define the target complement and Riesz representative N_r from $(\mathcal{D}, B_0; s)$. Put

$$\rho = \frac{\mu_s(B_0)}{\mu_s(A_0)}, \qquad \Lambda = \{\alpha \in \mathbb{F}_q^\times : \alpha^{s+1} = \rho\}. \tag{68}$$

If an exact procedure supplies the complete list Λ, then conjugacy of $\mathcal{C}$ and $\mathcal{D}$ can be decided and a conjugator can be recovered using at most $|\Lambda| \le s + 1 \le n$ simultaneous linear systems.

Proof. The target checks include equality of code dimensions and hull dimensions, cyclicity of B_0, the moment-zero pattern used to select s, transversality of s, and the scalar common-centralizer condition for (B_0, N_r). Each is necessary under conjugacy by Proposition 5.5. A mismatch therefore gives exact `NOT_CONJUGATE`.

For every $\alpha \in \Lambda$, solve the linear space

$$\mathcal{I}_\alpha = \{X :\ X(\alpha A_0) = B_0X, \quad X(\alpha^r M_r) = N_rX\}. \tag{69}$$

If $\mathcal{D} = S\mathcal{C}S^{-1}$, there is a unique scalar α satisfying (61). Equation (62) places it in the complete list Λ, and (63) places S in $\mathcal{I}_\alpha$. For this scale,

$$\mathcal{I}_\alpha = S\,\mathrm{Cent}(A_0, M_r) = S(\mathbb{F}_q I_n). \tag{70}$$

Thus the correct intertwiner space is one-dimensional and every nonzero element is invertible.

The algorithm rejects a zero-dimensional intertwiner space, a space of dimension greater than one, or a one-dimensional space with a singular generator. Indeed, if any invertible X belongs to $\mathcal{I}_\alpha$, then the same calculation as in (70) gives $\mathcal{I}_\alpha = X\,\mathrm{Cent}(A_0, M_r)$, so that space must be a line. For each remaining generator X, it checks the exact code identity

$$\mathcal{D} = X\mathcal{C}X^{-1}. \tag{71}$$

This final check prevents an intertwiner of the derived pair from being accepted unless it conjugates the full code. Conversely, a genuine conjugator occurs in the complete scale list and passes (71). Exhausting Λ therefore makes a negative answer sound.

For ordered presentations, Proposition 3.3 constructs and verifies $T \in \mathrm{GL}_m(\mathbb{F}_q)$ once X is known, including for rank-deficient coefficient maps. Together with the complete scale list, this gives the output semantics of Definition 2.4. □

The scalar common centralizer is used only in (70). It yields a projectively unique transporter between two supplied intrinsic pairs, not a canonical form for a single matrix code. The exact code and slice checks remain necessary.

5.5. Selected deterministic scale branches

The general geometry above does not by itself provide a deterministic root algorithm for every transverse index. The all-parameter algorithm uses only the following scale branches when it invokes the one-hull Riesz construction.

Lemma 5.7 (Deterministic splitting in degree at most three)**.** *Let $\mathbb{F}_q = \mathbb{F}_2[z]/(g)$ be supplied explicitly, with $q = 2^e$, and let $0 \neq f \in \mathbb{F}_q[X]$ have degree at most three. All its distinct $\mathbb{F}_q$-roots can be listed deterministically in* $\mathrm{poly}(\log q)$ *bit time. The algorithm uses $O(e)$ constant-degree gcd splitting attempts and no finite-field enumeration.*

Proof. A nonzero constant has no roots. Otherwise make f monic, compute $X^q \bmod f$ by e repeated squarings, and form

$$h = \gcd(f, X^q - X).$$

Since $(X^q - X)' = -1$, h is the squarefree product of exactly the linear factors of f over $\mathbb{F}_q$, even if f has repeated roots. If $h = 1$, return the empty list.

Let $\beta_1, \dots, \beta_e$ be the supplied polynomial basis of $\mathbb{F}_q/\mathbb{F}_2$. In $\mathbb{F}_q[X]/(h)$ compute, for $i = 1, \dots, e$,

$$\tau_i(X) = \sum_{j=0}^{e-1} (\beta_i X)^{2^j} \pmod{h}.$$

Each value at a root a equals $\mathrm{Tr}_{\mathbb{F}_q/\mathbb{F}_2}(\beta_i a) \in \{0, 1\}$. Start with the factor list $[h]$. For each i and each unresolved factor u, replace u by the nonconstant factors among $\gcd(u, \tau_i)$ and $u/\gcd(u, \tau_i)$. The first factor contains exactly its trace-zero roots, and the second exactly its trace-one roots. Linear factors may be left unchanged.

For completeness, the field trace pairing is nondegenerate. The polynomial $T(Y) = Y + Y^2 + \cdots + Y^{2^{e-1}}$ is nonzero of degree $q/2 < q$, so it is not zero on every element of $\mathbb{F}_q$. Thus the trace map is nonzero. For $a \neq b$, multiplication by $a - b$ is bijective, so some $\beta \in \mathbb{F}_q$ has $\mathrm{Tr}(\beta(a - b)) = 1$. Writing β in the basis shows that at least one basis functional separates a from b. After all e stages no factor contains two different roots. Every remaining factor is therefore linear, and its root is read from its coefficients and verified in f.

There are at most three factors at any stage, hence at most $3e$ splitting gcds. Constructing the e trace polynomials takes $O(e^2)$ constant-degree field operations; the initial rational-root gcd takes $O(e)$ such operations. Arithmetic, inversion, and constant-degree gcds in the explicit field all have bit cost polynomial in e. This proves the stated deterministic complexity. □

Corollary 5.8 (Deterministic selected branches)**.** *Assume $n \geq 3$ and that the source satisfies the one-hull, cyclicity, transversality, and scalar-common-centralizer tests above. The recovery procedure in Proposition 5.6 is deterministic in* $\mathrm{poly}(n, m, \log q)$ *time in each of the following cases.*

1. *If q is odd, the selected branch has $s = 0$ and $r = 2$. The unique possible scale is*

$$\alpha = \frac{\mathrm{Tr}(B_0)}{\mathrm{Tr}(A_0)}.$$

Algorithm 3 Selected one-hull recovery

Require: Equal-dimensional codes $C, D \le \mathrm{M}_n(\mathbb{F}_q)$, $n \ge 3$
Ensure: A verified conjugator, `NOT_CONJUGATE`, or source-side `FAILURE`

Compute $H(C)$; if it is not a line, return `FAILURE`
Let A_0 be its deterministic normalized generator
Set $(s, r) = (0, 2)$ in odd characteristic and $(s, r) = (2, 0)$ otherwise
If A_0 is noncyclic or $\mathrm{Tr}(A_0^{s+1}) = 0$, return `FAILURE`
Construct C_0 and M_r; if $\dim \mathrm{Cent}(A_0, M_r) \neq 1$, return `FAILURE`
Reject by `NOT_CONJUGATE` any mismatch of these necessary source properties in D
Compute the target hull generator B_0, complement D_0, and representative N_r
$\rho \leftarrow \mathrm{Tr}(B_0^{s+1}) / \mathrm{Tr}(A_0^{s+1})$
List all $\alpha \in \mathbb{F}_q^\times$ with $\alpha^{s+1} = \rho$ using the selected splitter
for each listed α **do**
    $J \leftarrow \{Z : Z(\alpha A_0) = B_0 Z,\ Z(\alpha^r M_r) = N_r Z\}$
    if $\dim J = 1$ and its generator S is invertible and $D = SCS^{-1}$ **then**
        return S
    end if
end for
return `NOT_CONJUGATE`

2. *If $q = 2^e \ge 4$, the selected branch has $s = 2$ and $r = 0$. All possible scales are the at most three roots of*
$$Z^3 - \frac{\mathrm{Tr}(B_0^3)}{\mathrm{Tr}(A_0^3)}.$$
They are listed deterministically by Lemma 5.7.
3. *If $q = 2$, the selected binary one-hull branch also has $s = 2$ and $r = 0$, and the only nonzero scale is $\alpha = 1$.*

Proof. In the first case, $\mu_0(A_0) = \mathrm{Tr}(A_0) \neq 0$, and (62) has weight one. It therefore determines α by a single division. The choice $r = 2$ also makes the marked class nonzero by Lemma 5.4.

In characteristic two, self-orthogonality gives $\mathrm{Tr}(A_0)^2 = \mathrm{Tr}(A_0^2) = 0$. The selected nonzero cubic moment therefore has least transverse index $s = 2$, and the Riesz representative M_0 has weight zero. For $q = 2^e \ge 4$, apply Lemma 5.7 to $f(Z) = Z^3 - \rho$, where $\rho = \mathrm{Tr}(B_0^3)/\mathrm{Tr}(A_0^3)$. For $q = 2$ the nonzero scale is one.

All remaining operations are power computation, cyclicity tests, Gram solves, centralizer and intertwiner linear systems, matrix inversion, image/kernel basis construction, and exact equality tests. Their dimensions are polynomial in n and m, and arithmetic in the explicitly supplied field is polynomial in $\log q$. □

For a later transverse index s, Proposition 5.6 remains an exact conditional recovery statement, but constructing its complete root list may require general finite-field root extraction. We assume no deterministic polynomial-time bound for that general operation.

Example 5.9 (A one-hull covariant and its scalar weight). Over $\mathbb{F}_5$, take $C = \langle A_0, B_0 \rangle$, where
$$A_0 = \begin{pmatrix} 0 & 1 & 1 \\ 0 & 4 & 4 \\ 2 & 0 & 0 \end{pmatrix}, \qquad B_0 = \begin{pmatrix} 2 & 2 & 1 \\ 2 & 0 & 3 \\ 4 & 3 & 0 \end{pmatrix}.$$
The trace Gram matrix is $\begin{pmatrix} 0 & 0 \\ 0 & 3 \end{pmatrix}$, $\mathrm{Tr}(A_0) = 4$, and I_3, A_0, A_0^2 are independent. The trace-zero

complement is the line spanned by $Z = \begin{pmatrix} 2 & 4 & 3 \\ 2 & 3 & 1 \\ 3 & 3 & 0 \end{pmatrix}$. The Gram solve gives $M_2 = 3Z$ and $\mathrm{Cent}(A_0, M_2) = \mathbb{F}_5 I_3$. For the ambient conjugator S in Example 9.1, choose the target hull generator $2SA_0S^{-1}$. Recomputing the target Gram solve gives $N_2 = 4SM_2S^{-1}$, not SM_2S^{-1}. The trace ratio is 2 and the weighted intertwiner system in Algorithm 3 recovers the same projective line of conjugators. This displays the distinction between basis independence and invariance under rescaling a hull generator.

6. Source predicates and uniform probability bounds

We now quantify the one-hull construction for every characteristic and every interior dimension. Put $N = n^2$. Code-level probabilities refer to a uniform Grassmannian; ordered-source probabilities are obtained only through the exact rank law in Proposition 3.4. The narrow source-scale certificate $\mathcal{G}^\star$ is defined separately in Appendix F.

6.1. The code predicates

We first name the two code-level predicates that occur in the interior rows. For $X, Y \in \mathrm{M}_2(\mathbb{F}_q)$, let

$$b_{\det}(X, Y) = \det(X + Y) - \det(X) - \det(Y). \tag{72}$$

Definition 6.1 (The predicate Good_{22})**.** Let $C \leq \mathrm{M}_2(\mathbb{F}_q)$ be two-dimensional. We write $\mathsf{Good}_{22}(C)$ if the following conditions hold.

1. The restriction of Tr to C is nonzero.
2. The line $K = C \cap \ker(\mathrm{Tr})$ has a nonzero generator U with $\det(U) = 0$.
3. For a vector $X \in C$ with $\mathrm{Tr}(X) = 1$, one has $b_{\det}(X, U) \neq 0$.

This predicate is independent of the displayed choices. Rescaling U rescales the value in the third condition by the same nonzero scalar. If X is replaced by $X + aU$, then $b_{\det}(U, U) = 2\det(U) = 0$, so the third condition is unchanged. Since $\mathrm{Tr}(U) = \det(U) = 0$, Cayley–Hamilton also gives $U^2 = 0$; thus the second condition is equivalently the nilpotent trace-kernel condition.

Definition 6.2 (The predicate $\mathsf{Good}_{\mathrm{int}}$)**.** Let $n \geq 3$, let $2 \leq d \leq N-2$, and let $C \in \mathrm{Gr}(d, \mathrm{M}_n(\mathbb{F}_q))$. We write $\mathsf{Good}_{\mathrm{int}}(C)$ if $H(C) = \langle A_0 \rangle$ is a line, A_0 is cyclic, and the following field-appropriate conditions hold.

1. If q is odd, then $\mathrm{Tr}(A_0) \neq 0$ and
$$\mathrm{Cent}\big(A_0, M_2(C, A_0; 0)\big) = \mathbb{F}_q I_n.$$
2. If q is even, then $\mathrm{Tr}(A_0^3) \neq 0$ and
$$\mathrm{Cent}\big(A_0, M_0(C, A_0; 2)\big) = \mathbb{F}_q I_n.$$

Here the Riesz representatives are those of Definition 5.2.

The definition does not depend on the nonzero generator of the hull line. Indeed, Proposition 5.5 shows that M_2 has weight two in the odd branch and that M_0 has weight zero in the even branch. Cyclicity, nonvanishing of the selected moment, and scalarity of the common centralizer are therefore projective conditions. In characteristic two, the hull identity $\mathrm{Tr}(A_0^2) = 0$ also gives $\mathrm{Tr}(A_0) = 0$, so a nonzero cubic trace makes $s = 2$ the least transverse index.

6.2. Rank and hull estimates

We record two estimates used in the probability calculation. Both hold uniformly at $q = 2$.

Lemma 6.3 (Rank products)**.** *For $m \le N$, put*

$$p_{\mathrm{col}}(N, m, q) = \prod_{j=N-m+1}^{N} (1 - q^{-j}).$$

This is the probability that the coefficient map has rank m. If $2 \le m \le N - 2$, then

$$p_{\mathrm{col}}(N, m, q) > \frac{3}{4}. \tag{73}$$

Moreover, for every $k \ge 3$,

$$\prod_{j=1}^{k} (1 - q^{-j}) > \frac{1}{4}. \tag{74}$$

Proof. The exact product is Proposition 3.4. For the interior range its first exponent is at least three, and therefore

$$p_{\mathrm{col}}(N, m, q) \ge \prod_{j=3}^{N} (1 - q^{-j}) > 1 - \sum_{j=3}^{\infty} q^{-j} = 1 - \frac{1}{q^2(q-1)} \ge \frac{3}{4}.$$

For the second assertion, split off the first factor and apply the same finite-product bound to the tail:

$$\prod_{j=1}^{k} (1 - q^{-j}) > (1 - q^{-1}) \left(1 - \sum_{j=2}^{\infty} q^{-j}\right) = (1 - q^{-1}) \left(1 - \frac{1}{q(q-1)}\right) \ge \frac{1}{4}.$$

The strict inequalities come from replacing a finite product by an infinite tail bound. □

Lemma 6.4 (Uniform upper bound for a one-dimensional hull)**.** *Let $n \ge 3$, $2 \le d \le N - 2$, and let C be uniform in $\mathrm{Gr}(d, \mathrm{M}_n(\mathbb{F}_q))$. Then*

$$\Pr[\dim H(C) = 1] \le \frac{4}{q}. \tag{75}$$

Proof. If $H(C) = L$ is a line, then L is isotropic and $L \subseteq C \subseteq L^{\perp_b}$. For a fixed isotropic line, omitting the additional condition that C/L be nondegenerate only increases the count. The resulting ratio to all d-subspaces is

$$R_{N,d} = \frac{\begin{bmatrix} N-2 \\ d-1 \end{bmatrix}_q}{\begin{bmatrix} N \\ d \end{bmatrix}_q} = \frac{(q^d - 1)(q^{N-d} - 1)}{(q^N - 1)(q^{N-1} - 1)}. \tag{76}$$

In characteristic two, the isotropic vectors are exactly the trace-zero matrices, because $\mathrm{Tr}(X^2) = \mathrm{Tr}(X)^2$. There are therefore $(q^{N-1} - 1)/(q - 1)$ isotropic lines, and the union bound gives

$$\frac{(q^d - 1)(q^{N-d} - 1)}{(q - 1)(q^N - 1)} \le \frac{2}{q}.$$

For $q = 2$ the last inequality is immediate; for $q \ge 4$ it follows after dividing by q^N and using $N \ge 9$.

Suppose now that q is odd. After all entries except x_{11} are fixed, the equation $\mathrm{Tr}(X^2) = 0$ is a nonzero quadratic equation in x_{11}. It has at most two solutions. Hence the number of isotropic lines is less than $2q^{N-1}/(q - 1)$. Since $q \ge 3$ and $N \ge 9$,

$$(1 - q^{-N})(1 - q^{-(N-1)}) > \frac{3}{4}.$$

Multiplying the line bound by (76) gives

$$\Pr[\dim H(C) = 1] < \frac{8}{3(q-1)} \le \frac{4}{q}.$$

This proves (75) in every characteristic. □

The lower estimate requires the selected cyclic moment and Riesz-pair conditions, rather than the hull dimension alone. The following proposition states the exact sequential payments used for that intersection.

Proposition 6.5 (Interior good-code probability)**.** *Let $n \ge 3$, $2 \le d \le N-2$, and let C be uniform in $\mathrm{Gr}(d, \mathrm{M}_n(\mathbb{F}_q))$. Then*

$$\Pr[\mathsf{Good}_{\mathrm{int}}(C)] > \frac{3}{304q}. \tag{77}$$

Consequently, for an iid ordered source in the fourth row of Definition F.1,

$$\frac{9}{1216q} < \Pr[A \in \mathcal{G}^{\star}_{q,n,m}] \le \frac{4}{q}. \tag{78}$$

Proof. Put $D = N-2$ and $k = d-1$. For an isotropic line L, taking inverse image under $L^{\perp} \to L^{\perp}/L$ bijects the nondegenerate k-subspaces of the quotient with the codes having hull exactly L. Thus the hull line is counted first, and the correlated projection is counted in that same quotient. The exact fibre calculation and the centralizer-witness union bound are given in Appendix E.1.

The resulting sequential lower bounds are as follows. In the first seven rows, p_1 is the mass of a one-dimensional hull with the specified cyclic, transverse generator, and p_2 is the conditional scalar-centralizer probability. In the binary rows the calculation restricts further to oo quotients as defined in Lemma D.1; the last column is the resulting unconditional lower bound.

field and size	p_1	p_2	$\Pr[\mathsf{Good}_{\mathrm{int}}]$
q odd, $n = 3$	$> 3/(16q)$	$> 1/19$	$> 3/(304q)$
q odd, $n = 4$	$> 1/(8q)$	$> 1/3$	$> 1/(24q)$
q odd, $n \ge 5$ odd	$\ge 1/(5q)$	$\ge 1/7$	$\ge 1/(35q)$
q odd, $n \ge 6$ even	$> 1/(10q)$	$> 1/3$	$> 1/(30q)$
$q \ge 4$ even, $n = 3$	$> 1/(4q)$	$> 1/3$	$> 1/(12q)$
$q \ge 4$ even, $n = 4$	$> 1/(4q)$	$> 1/2$	$> 1/(8q)$
$q \ge 4$ even, $n \ge 5$	$> 1/(8q)$	$> 1/7$	$> 1/(56q)$
$q = 2$, $n = 3$	oo quotient calculation		$> 3/128$
$q = 2$, $n = 4$	oo quotient calculation		$> 1/64$
$q = 2$, $n \ge 5$	oo quotient calculation		$> 1/128$

The odd-characteristic factors use Lemmas B.1–B.4 and C.1. For characteristic two with $q \ge 4$, the factors follow from the direct trace-zero and incidence calculation in Appendix E.1, culminating in (E.43). For $q = 2$, the factors use Lemma C.1, Corollary C.2, and Lemmas D.1–D.3. The endpoint maximization and all constant comparisons are in Appendix E. The weakest row is q odd, $n = 3$, proving the code-level bound.

Conditional on full column rank, the image code is uniform by Proposition 3.4. The rank probability is strictly greater than $3/4$ by Lemma 6.3. Multiplying these conditional factors gives $9/(1216q)$. The upper bound follows from $\mathsf{Good}_{\mathrm{int}} \Rightarrow \dim H(C) = 1$ and Lemma 6.4. □

6.3. A small-dimensional source count

Lemma 6.6 (The 2-by-2 good-plane count)**.** *For a uniform two-dimensional code $C \le \mathrm{M}_2(\mathbb{F}_q)$,*

$$\Pr[\mathsf{Good}_{22}(C)] = g_2(q) := \frac{q(q^2-1)}{(q^2+1)(q^2+q+1)}, \qquad \frac{1}{3q} < g_2(q) < \frac{1}{q}. \tag{79}$$

Proof. There are $q+1$ projective nonzero nilpotent lines in $\mathrm{M}_2(\mathbb{F}_q)$: the image and kernel of a nonzero nilpotent matrix are the same line in $\mathbb{F}_q^2$, and each such line determines one projective matrix line. Fix a nilpotent line and conjugate it to $L = \langle E_{12} \rangle$. A trace-one matrix has the form

$$E = \begin{pmatrix} a & b \\ c & 1-a \end{pmatrix}, \qquad b_{\det}(E, E_{12}) = -c.$$

Thus the nonzero determinant-polarization condition in Definition 6.1 is $c \neq 0$, and

$$\det E = 0 \quad \Longleftrightarrow \quad b = \frac{a(1-a)}{c}.$$

Each pair $(a, c) \in \mathbb{F}_q \times \mathbb{F}_q^\times$ therefore gives such a matrix E and a good plane $\langle E_{12}, E \rangle$ with trace-kernel L. Conversely, in any good plane with trace-kernel L, the trace-one matrices are $X + tE_{12}$, where the lower-left entry c of X is nonzero. Since

$$\det(X + tE_{12}) = \det X - tc,$$

exactly one of them has determinant zero. Hence each good plane contributes exactly one of the displayed matrices E, with no repeated planes in the (a, c) count. There are consequently $q(q-1)$ good planes for each fixed L. Each good plane has a unique trace-kernel line, so the total is $(q+1)q(q-1) = q(q^2-1)$. Division by

$$\begin{bmatrix} 4 \\ 2 \end{bmatrix}_q = (q^2+1)(q^2+q+1)$$

gives the displayed fraction. The lower inequality is equivalent to

$$2q^4 - q^3 - 5q^2 - q - 1 > 0,$$

which after writing $q = x + 2$ has positive coefficients and positive constant term. The upper inequality follows from

$$(q^2+1)(q^2+q+1) - q^2(q^2-1) = q^3 + 3q^2 + q + 1 > 0.$$

□

7. Assembly of the all-parameter deterministic algorithm

We now assemble the local recovery procedures into one algorithm for every matrix size and every declared number of slices. The assembly has two distinct source sets. Appendix F defines the certified subevent $\mathcal{G}^\star_{q,n,m}$, whose probability is bounded on both sides. This section defines the larger recognized family $\mathcal{G}_{q,n,m}$ on which the same algorithm is complete. We retain both sets because the upper bound for $\mathcal{G}^\star_{q,n,m}$ does not impose an upper bound on all sources that the algorithm recognizes.

Throughout this section, put $N = n^2$. A nonzero vector in a matrix code or an orthogonal line is chosen deterministically by scaling its first nonzero vectorized coordinate to one. All predicates below are independent of this normalization.

7.1. The recognized source family

Definition 7.1 (The larger all-parameter source family)**.** For $A = (A_1, \dots, A_m) \in \mathrm{M}_n(\mathbb{F}_q)^m$, define $A \in \mathcal{G}_{q,n,m}$ by the first applicable row below.

1. If $n = 1$, every source belongs to $\mathcal{G}_{q,1,m}$.
2. If $n \geq 2$ and $m = 1$, require $\mathrm{Tr}(A_1) \neq 0$.
3. If $(n, m) = (2, 2)$, require $r_A = 2$ and $\mathsf{Good}_{22}(\mathcal{C}_A)$, with Good_{22} as in Definition 6.1.
4. If $n \geq 3$ and $2 \leq m \leq N - 2$, require $r_A = m$ and $\mathsf{Good}_{\mathrm{int}}(\mathcal{C}_A)$, with the field-dependent one-hull predicate from Definition 6.2.
5. If $n \geq 2$ and $m = N - 1$, require $r_A = N - 1$ and, writing $\mathcal{C}_A^{\perp_b} = \langle A_0^{\perp} \rangle$, require $\mathrm{Tr}(A_0^{\perp}) \neq 0$.
6. If $n \geq 2$ and $m \geq N$, require $r_A = N$.

These six rows are mutually exclusive and cover every $n, m \geq 1$. They are also source-only conditions. Ranks, image codes, trace hulls, orthogonal lines, the two-space predicate, and the one-hull rigid-pair predicate are all computed by the deterministic operations described in Sections 2–5. The event of Definition F.1 was obtained by source-only thinning of these rows. Hence

$$\mathcal{G}^{\star}_{q,n,m} \subseteq \mathcal{G}_{q,n,m}. \tag{80}$$

7.2. A complete solver for trace-good lines

The line solver is used directly when $m = 1$ and, after orthogonal complementation, when $m = N - 1$.

Proposition 7.2 (Trace-normalized line solver)**.** *Let $\mathcal{C} = \langle A_0 \rangle$ and $\mathcal{D} \leq \mathrm{M}_n(\mathbb{F}_q)$, where $\mathrm{Tr}(A_0) \neq 0$. There is a deterministic $\mathrm{poly}(n, \log q)$ algorithm which returns $S \in \mathrm{GL}_n(\mathbb{F}_q)$ exactly when $\mathcal{D} = S\mathcal{C}S^{-1}$, and otherwise returns* `NOT_CONJUGATE`*.*

Proof. First check that $\mathcal{D}$ is a line. Choose its normalized generator B_0. If $\mathrm{Tr}(B_0) = 0$, the codes cannot be conjugate. Otherwise set

$$\alpha = \frac{\mathrm{Tr}(B_0)}{\mathrm{Tr}(A_0)}. \tag{81}$$

If $\mathcal{D} = S\mathcal{C}S^{-1}$, then $B_0 = \beta S A_0 S^{-1}$ for a unique $\beta \in \mathbb{F}_q^{\times}$, and trace invariance forces $\beta = \alpha$. Thus code conjugacy is equivalent to ordinary similarity of αA_0 and B_0.

Compute rational canonical forms of these two matrices, together with transforming bases. Unequal forms give `NOT_CONJUGATE`; equal forms give a matrix S satisfying $B_0 = S(\alpha A_0)S^{-1}$. The latter identity is checked before S is returned. Rational canonical form and a transforming basis can be computed by polynomial Euclidean and Smith-form arithmetic over $\mathbb{F}_q[x]$. This requires neither irreducible factorization nor root extraction. □

The normalization in Proposition 7.2 is only an implementation convention. Replacing A_0, B_0 by uA_0, vB_0 changes α to $(v/u)\alpha$, so the matrix tested for similarity changes by the same target scaling. The resulting code-level decision is unchanged.

7.3. Scalar, hyperplane, and full-image cases

Three rows of Definition 7.1 reduce to complete linear classifications.

Proposition 7.3 (Boundary and full-image solvers)**.** *For a recognized source, image-code conjugacy is decided and an ambient conjugator is obtained as follows.*

1. *If $n = 1$, compare $r_A, r_B \in \{0, 1\}$. Equal ranks give conjugate image codes with $S = 1$; unequal ranks give* `NOT_CONJUGATE`*.*
2. *If $n \geq 2$ and $m = N - 1$, compare the two coefficient ranks. For equal full ranks, compute the normal lines $\mathcal{C}_A^{\perp_b}$ and $\mathcal{C}_B^{\perp_b}$ and apply Proposition 7.2.*

3. *If $n \geq 2$ and $m \geq N$, compare the two coefficient ranks. If both are N, then $\mathcal{C}_A = \mathcal{C}_B = \mathrm{M}_n(\mathbb{F}_q)$ and one may take $S = I_n$.*

Every negative answer in these procedures is exact.

Proof. For $n = 1$, conjugation is trivial and the only image codes are 0 and $\mathbb{F}_q$. This proves the first claim.

For the second claim, Lemma 2.2 gives the more general identity

$$(S\mathcal{C}S^{-1})^{\perp_b} = S\mathcal{C}^{\perp_b}S^{-1}.$$

Ambient nondegeneracy also gives $(\mathcal{C}^{\perp_b})^{\perp_b} = \mathcal{C}$. Conjugacy of two full-rank hyperplanes is therefore equivalent to conjugacy of their normal lines. A conjugate target normal has nonzero trace, so every rejection by the line solver is necessary.

In the last case, surjectivity makes both images equal to the whole matrix space. Conversely, coefficient rank is invariant under the ordered action, so a target of smaller rank cannot be conjugate to a surjective source. □

7.4. The exceptional two-dimensional matrix space

The only interior value for $n = 2$ is $m = 2$. Its solver uses the determinant-polarization pair from the predicate Good_{22} rather than a one-hull Riesz pair. For $X, Y \in \mathrm{M}_2(\mathbb{F}_q)$, put

$$\Delta(X, Y) = \det(X + Y) - \det(X) - \det(Y) = \mathrm{Tr}(X)\,\mathrm{Tr}(Y) - \mathrm{Tr}(XY). \tag{82}$$

This polynomial identity remains valid in characteristic two.

Let $\mathcal{C}$ be a good two-space. The restriction of Tr to $\mathcal{C}$ is nonzero, and

$$\mathcal{C} \cap \ker(\mathrm{Tr}) = \langle N \rangle.$$

Choose $X \in \mathcal{C}$ with $\mathrm{Tr}(X) = 1$, set

$$\lambda = \Delta(X, N), \qquad N_{\mathcal{C}} = \lambda^{-1}N, \qquad E_{\mathcal{C}} = X - \det(X)N_{\mathcal{C}}, \tag{83}$$

and recall that goodness requires $\det(N) = 0$ and $\lambda \neq 0$.

Lemma 7.4 (Canonical pair for a good two-space)**.** *The pair $(E_{\mathcal{C}}, N_{\mathcal{C}})$ in* (83) *is independent of the choices of N and X, is conjugacy covariant, and satisfies*

$$E_{\mathcal{C}}^2 = E_{\mathcal{C}}, \qquad N_{\mathcal{C}}^2 = 0, \qquad \mathrm{Cent}(E_{\mathcal{C}}, N_{\mathcal{C}}) = \mathbb{F}_q I_2.$$

Proof. Rescaling N rescales λ by the same factor, so $N_{\mathcal{C}}$ is unchanged. Replacing X by $X + tN$ leaves λ unchanged because $\Delta(N, N) = 2\det(N) = 0$. Since $\Delta(X, N_{\mathcal{C}}) = 1$ and $\det(N_{\mathcal{C}}) = 0$, one has

$$\det(X + aN_{\mathcal{C}}) = \det(X) + a.$$

This identity shows that $E_{\mathcal{C}}$ is also unchanged. Trace, determinant, and Δ are conjugacy invariant, which proves covariance. The two power identities follow from the 2×2 Cayley–Hamilton identity.

Conjugate $E_{\mathcal{C}}$ to $\mathrm{diag}(1, 0)$ and write

$$N_{\mathcal{C}} = \begin{pmatrix} a & b \\ c & -a \end{pmatrix}.$$

The equations $\Delta(E_{\mathcal{C}}, N_{\mathcal{C}}) = 1$ and $\det(N_{\mathcal{C}}) = 0$ give $a = -1$ and $bc = -1$. Thus $b, c \neq 0$. A matrix commuting with $E_{\mathcal{C}}$ is diagonal, and commuting also with $N_{\mathcal{C}}$ forces its two diagonal entries to agree. This proves the centralizer identity. □

Proposition 7.5 (Complete $(2,2)$ code solver)**.** *Let $\mathcal{C}$ satisfy $\mathsf{Good}_{22}(\mathcal{C})$ and let $\mathcal{D} \le \mathrm{M}_2(\mathbb{F}_q)$ be arbitrary. Deterministic linear algebra returns $S \in \mathrm{GL}_2(\mathbb{F}_q)$ exactly when $\mathcal{D} = S\mathcal{C}S^{-1}$, and otherwise returns* `NOT_CONJUGATE`*.*

Proof. Reject unless $\dim \mathcal{D} = 2$ and $\mathsf{Good}_{22}(\mathcal{D})$ holds. These are necessary conjugacy invariants. Construct the two canonical pairs and solve

$$\mathcal{J}(\mathcal{C}, \mathcal{D}) = \{Z : ZE_{\mathcal{C}} = E_{\mathcal{D}}Z,\ ZN_{\mathcal{C}} = N_{\mathcal{D}}Z\}. \tag{84}$$

If a conjugator S exists, covariance and Lemma 7.4 give

$$\mathcal{J}(\mathcal{C}, \mathcal{D}) = S\,\mathrm{Cent}(E_{\mathcal{C}}, N_{\mathcal{C}}) = S(\mathbb{F}_q I_2).$$

Consequently a space of dimension other than one contains no invertible conjugator. In dimension one, test the determinant of a nonzero generator Z and then check the exact identity $\mathcal{D} = Z\mathcal{C}Z^{-1}$. These tests accept every genuine conjugator and reject every false candidate. □

7.5. Interior one-hull solvers in every characteristic

Suppose now that $n \ge 3$ and $2 \le m \le N-2$. A source in $\mathcal{G}_{q,n,m}$ has $r_A = m$, so its image is an m-space. The predicate $\mathsf{Good}_{\mathrm{int}}$ selects one of the following projectively well-defined pairs.

1. For odd q, the hull generator A_0 is cyclic and $\mathrm{Tr}(A_0) \neq 0$. Use $(s, r) = (0, 2)$ and the pair (A_0, M_2).
2. For $q = 2^e \ge 4$, the hull generator A_0 is cyclic and $\mathrm{Tr}(A_0^3) \neq 0$. Use $(s, r) = (2, 0)$ and the pair (A_0, M_0).
3. For $q = 2$, use the same indices $(s, r) = (2, 0)$. The nonzero generator of the hull line is unique and satisfies $\mathrm{Tr}(A_0^3) = 1$.

In every row the source predicate also requires $\mathrm{Cent}(A_0, M_r) = \mathbb{F}_q I_n$.

Proposition 7.6 (Complete interior code solver)**.** *For every source code selected above and every arbitrary target code $\mathcal{D} \le \mathrm{M}_n(\mathbb{F}_q)$, there is a deterministic $\mathrm{poly}(n, m, \log q)$ procedure that returns $S \in \mathrm{GL}_n(\mathbb{F}_q)$ exactly when $\mathcal{D} = S\mathcal{C}S^{-1}$, and otherwise returns* `NOT_CONJUGATE`*.*

Proof. The selected source satisfies cyclicity, transversality and scalar common centralizer, the hypotheses of Proposition 5.6. By Corollary 5.8, its complete scale list is the singleton $\mathrm{Tr}(B_0)/\mathrm{Tr}(A_0)$ in odd characteristic, the roots of $Z^3 - \mathrm{Tr}(B_0^3)/\mathrm{Tr}(A_0^3)$ for $q = 2^e \ge 4$, and $\{1\}$ for $q = 2$. The deterministic splitter of Lemma 5.7 constructs the only nonlinear list. Each list has at most three elements. The cited recovery proposition supplies the complete intertwiner search, invertibility and code checks, and exact negative answer for every target. □

The separate binary analysis in Appendix D concerns source mass; recovery still uses the weight-zero Riesz equations.

7.6. The master search-and-decision procedure

We can now state the algorithm independently of the probability proof.

Definition 7.7 (All-parameter partial MCC algorithm)**.** First evaluate the six-row predicate $A \in \mathcal{G}_{q,n,m}$ of Definition 7.1; reject the source with `FAILURE` if it fails. Otherwise compare the source and target ranks and invoke the complete code solver for that row. A rank mismatch or an exact negative answer from that solver returns `NOT_CONJUGATE`. For every ambient candidate S, verify $\mathcal{C}_B = S\mathcal{C}_A S^{-1}$ and apply Algorithm 1, whose matrix R satisfies $P_B = P_{S,A}R$ and $R(\ker \phi_B) = \ker \phi_A$. Return

$$T = R^{\mathsf{T}} \tag{85}$$

together with S only after checking invertibility and all ordered slice equations. Exhausting the candidates returns `NOT_CONJUGATE`.

Proposition 7.8 (Exact output semantics)**.** *On every source $A \in \mathcal{G}_{q,n,m}$, the algorithm of Definition 7.7 satisfies the full output semantics of Definition 2.4.*

Proof. The scalar, hyperplane and full-image rows use Proposition 7.3; the other rows use Propositions 7.2, 7.5 and 7.6. Each code solver is complete on its source predicate. If an ordered witness exists, Lemma 3.1 ensures equal ranks and image conjugacy, so the selected solver returns an ambient conjugator and the same lemma lifts it on the original m-dimensional coefficient space. Final verification gives soundness; completeness makes candidate exhaustion an exact negative answer. Only source rejection returns `FAILURE`.

This reasoning applies to any source-certified complete branch: if its code solver rejects a target, another branch cannot rescue that target. It is the justification for choosing one branch in the union dispatcher of Section 9, not an assertion that the source families are identical. □

7.7. Deterministic complexity

Proposition 7.9 (Encoded-input running time)**.** *In the explicit-field model of Section 2, the algorithm in Definition 7.7 has deterministic bit complexity* $\mathrm{poly}(n, m, \log q)$.

Proof. All recognition and recovery steps are polynomial-size linear systems, except the deterministic rational-canonical-form computation in the line/hyperplane solver and the selected characteristic-two scale list. The former is justified in Proposition 7.2; the latter uses Lemma 5.7, with at most three candidates and $O(\log q)$ trace-coordinate separators. Cyclicity is tested by the rank of $I_n, A_0, \dots, A_0^{n-1}$. Each ambient candidate incurs the $O(mn^3 + n^2m^2 + m^3)$ lift of Proposition 3.3. All costs are deterministic and polynomial in the explicit input; Section 9 gives a coarse operation budget. Source-density counts are not executed by the algorithm, and no general factoring or root-extraction routine is assumed. □

7.8. Probability assembly and the universal theorem

All probabilities below refer to the original ordered source tuple with independent uniform entries. Conditional on $r_A = r$, the image $\mathcal{C}_A$ is uniform on $\mathrm{Gr}(r, \mathrm{M}_n(\mathbb{F}_q))$ by Proposition 3.4. Thus the rank factors and the code-level factors combine by the chain rule, not by an independence assumption.

Proposition 7.10 (Uniform mass of the larger source family)**.** *For every prime power $q \geq 2$ and all $n, m \geq 1$,*

$$\Pr[A \in \mathcal{G}_{q,n,m}] > \frac{9}{1216q} > \frac{1}{136q}. \tag{86}$$

Proof. For $n = 1$ every source is recognized. The trace-good line row has exact mass $1 - q^{-1}$. At $(n, m) = (2, 2)$, (F.5) gives $(q-1)^3(q+1)^2/q^6 > 1/(4q)$. The hyperplane row has mass $\prod_{j=1}^{N-1}(1 - q^{-j}) > 1/4$ by (F.7); the full-image row has mass $\prod_{j=m-N+1}^{m}(1 - q^{-j}) > 1/4$ by (19) and Lemma 6.3.

For the interior row, the authoritative field/size table in Proposition 6.5 gives code probability $> 3/(304q)$, including its separately derived binary cases. Conditional image uniformity and the full-rank factor $> 3/4$ therefore give

$$\Pr[A \in \mathcal{G}_{q,n,m}] > \frac{3}{4}\,\frac{3}{304q} = \frac{9}{1216q}.$$

The weakest code-level row is odd q, $n = 3$; every other row is larger. Finally $9 \cdot 136 = 1224 > 1216$. □

The following theorem is the detailed form of Theorem B; in particular, its proof proves the headline statement from the Introduction.

Theorem 7.11 (All-parameter deterministic MCC)**.** *Let $q \geq 2$ be any prime power, let $n, m \geq 1$, and suppose that $\mathbb{F}_q = \mathbb{F}_p[z]/(g)$ is supplied explicitly. There are efficiently recognizable source events*

$$\mathcal{G}^{\star}_{q,n,m} \subseteq \mathcal{G}_{q,n,m}$$

and a deterministic partial MCC search-and-decision algorithm of bit complexity $\mathrm{poly}(n, m, \log q)$ *with the following properties. For an iid-uniform ordered source $A \in \mathrm{M}_n(\mathbb{F}_q)^m$,*

$$\frac{1}{136q} \leq \Pr[A \in \mathcal{G}^{\star}_{q,n,m}] \leq \frac{4}{q}, \tag{87}$$

$$\Pr[A \in \mathcal{G}_{q,n,m}] > \frac{9}{1216q} > \frac{1}{136q}. \tag{88}$$

In particular, $\Pr[A \in \mathcal{G}^{\star}_{q,n,m}] = \Theta(1/q)$ with absolute constants. No upper bound is asserted for $\mathcal{G}_{q,n,m}$.

For every $A \in \mathcal{G}_{q,n,m}$ and every arbitrary target $B \in \mathrm{M}_n(\mathbb{F}_q)^m$, the algorithm returns `YES` *together with $(S, T) \in \mathrm{GL}_n(\mathbb{F}_q) \times \mathrm{GL}_m(\mathbb{F}_q)$ satisfying* (5) *if and only if such a pair exists. Every returned pair is verified against all m slice equations. If no such pair exists, the algorithm returns* `NOT_CONJUGATE`*. Once the source recognizer accepts, the algorithm never returns* `FAILURE`*.*

Proof. The two-sided estimate (87) and efficient recognizability of $\mathcal{G}^{\star}_{q,n,m}$ are established in Theorem F.2. The row-by-row comparison following Definition 7.1 proves the inclusion (80). Proposition 7.10 proves (88). Proposition 7.8 proves the search, decision, and arbitrary-target statements, while Proposition 7.9 proves deterministic encoded-input polynomial time. These statements cover all six mutually exclusive rows of Definition 7.1, so no parameter regime remains outside the dispatch. □

8. The two-by-two matrix size

The two-dimensional subspaces of $\mathrm{M}_2(\mathbb{F}_q)$ form the only interior parameter regime in which the zero-hull pair $(P_{\mathcal{C}}(I_2), P_{\mathcal{C}}(P_{\mathcal{C}}(I_2)^2))$ cannot have scalar common centralizer. This section first proves that collapse directly from Cayley–Hamilton. It then separates two consequences that must not be conflated. In characteristic two, the full projected multiplication tensor still gives a deterministic high-density recovery procedure. Appendix G treats the distinct odd-characteristic obstruction; it is not a complexity lower bound.

Throughout the section, let

$$W = \mathrm{M}_2(k), \qquad b(X, Y) = \mathrm{Tr}(XY),$$

where the ground field k will be specified in each subsection. If a two-plane $\mathcal{C} \leq W$ has $H(\mathcal{C}) = 0$, write

$$P_{\mathcal{C}} : W \longrightarrow \mathcal{C}$$

for the b-orthogonal projection, and put

$$A_{\mathcal{C}} = P_{\mathcal{C}}(I_2), \qquad M_{\mathcal{C}} = P_{\mathcal{C}}(A_{\mathcal{C}}^2). \tag{89}$$

The defining identity

$$\mathrm{Tr}(P_{\mathcal{C}}(Y)X) = \mathrm{Tr}(YX) \qquad (X \in \mathcal{C}) \tag{90}$$

also defines the full projected multiplication

$$\mu_{\mathcal{C}}(X, Y) = P_{\mathcal{C}}(XY) \qquad (X, Y \in \mathcal{C}). \tag{91}$$

Both objects are basis free. If $\mathcal{D} = S\mathcal{C}S^{-1}$, trace invariance and uniqueness in (90) give

$$P_{\mathcal{D}}(SYS^{-1}) = SP_{\mathcal{C}}(Y)S^{-1}, \tag{92}$$

$$A_{\mathcal{D}} = SA_{\mathcal{C}}S^{-1}, \qquad \mu_{\mathcal{D}}(SXS^{-1}, SYS^{-1}) = S\mu_{\mathcal{C}}(X, Y)S^{-1}. \tag{93}$$

There is no projective scaling in these covariance laws.

8.1. Cayley–Hamilton collapse of a single anchor

The obstruction to the pair in (89) is characteristic independent.

Lemma 8.1 (Single-anchor projected-power collapse)**.** *Let k be any field and let $\mathcal{C} \leq \mathrm{M}_2(k)$ be a zero-hull two-plane. Then*

$$A_{\mathcal{C}}^2 = \mathrm{Tr}(A_{\mathcal{C}})A_{\mathcal{C}} - \det(A_{\mathcal{C}})I_2, \tag{94}$$

$$M_{\mathcal{C}} = (\mathrm{Tr}(A_{\mathcal{C}}) - \det(A_{\mathcal{C}}))A_{\mathcal{C}}. \tag{95}$$

More generally, $P_{\mathcal{C}}(f(A_{\mathcal{C}}))$ belongs to $kA_{\mathcal{C}}$ for every $f \in k[z]$. Consequently no tuple obtained only from projected powers of $A_{\mathcal{C}}$ has scalar common centralizer.

Proof. Cayley–Hamilton for a 2×2 matrix is

$$X^2 - \mathrm{Tr}(X)X + \det(X)I_2 = 0,$$

which gives (94). Since $A_{\mathcal{C}} \in \mathcal{C}$, the projection fixes $A_{\mathcal{C}}$; by definition it also sends I_2 to $A_{\mathcal{C}}$. Applying $P_{\mathcal{C}}$ to (94) proves (95).

Euclidean reduction modulo the characteristic polynomial writes $f(A_{\mathcal{C}}) = uI_2 + vA_{\mathcal{C}}$ for some $u, v \in k$. Projection therefore gives $(u+v)A_{\mathcal{C}}$. If $A_{\mathcal{C}}$ is nonscalar, then $kI_2 + kA_{\mathcal{C}} \leq \mathrm{Cent}(A_{\mathcal{C}})$; if it is scalar, its centralizer is all of $\mathrm{M}_2(k)$. In either case the common centralizer of any finite family of scalar multiples of $A_{\mathcal{C}}$ is strictly larger than kI_2. □

The same proof covers any defined rational function of the single anchor: the inverse of an invertible element of $kI_2 + kA_{\mathcal{C}}$ remains in that two-dimensional algebra. Lemma 8.1 is not a statement about genuinely mixed products. The characteristic-two result below retains the entire bilinear map (91).

8.2. Characteristic two: complete projected multiplication

Assume in this subsection that $k = \mathbb{F}_q$ with $q = 2^e$. For $X, Y \in W$, define the determinant polarization

$$\beta(X, Y) = \det(X + Y) + \det(X) + \det(Y). \tag{96}$$

Lemma 8.2 (Trace geometry in characteristic two)**.** *The form β is nondegenerate and alternating on the four-dimensional space W, and*

$$\beta(X, Y) = \mathrm{Tr}(XY) + \mathrm{Tr}(X)\,\mathrm{Tr}(Y). \tag{97}$$

Let (X, Y) be an ordered basis of a two-plane $\mathcal{C}$, and set

$$u = \mathrm{Tr}(X), \qquad v = \mathrm{Tr}(Y), \qquad w = \mathrm{Tr}(XY), \qquad D = uv + w.$$

Then

$$\det \begin{pmatrix} \mathrm{Tr}(X^2) & \mathrm{Tr}(XY) \\ \mathrm{Tr}(XY) & \mathrm{Tr}(Y^2) \end{pmatrix} = D^2. \tag{98}$$

In particular, $H(\mathcal{C}) = 0$ if and only if $D \neq 0$. In that case

$$A_{\mathcal{C}} = \frac{vX + uY}{D}, \qquad \mathrm{Tr}(A_{\mathcal{C}}) = 0. \tag{99}$$

Moreover, $A_{\mathcal{C}} = 0$ if and only if $\mathrm{Tr}\,|_{\mathcal{C}} = 0$, while $A_{\mathcal{C}} = I_2$ if and only if $I_2 \in \mathcal{C}$.

Proof. Write $X = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ and $Y = \left(\begin{smallmatrix} a' & b' \\ c' & d' \end{smallmatrix}\right)$. Expanding (96) gives

$$\beta(X, Y) = ad' + a'd + bc' + b'c,$$

which proves alternation and nondegeneracy by inspection. Expanding the right side of (97) gives the same expression.

In characteristic two, $\mathrm{Tr}(X^2) = u^2$ and $\mathrm{Tr}(Y^2) = v^2$. Hence the Gram determinant is $u^2v^2 + w^2 = (uv + w)^2$, proving (98). The radical of $b|_{\mathcal{C}}$ is zero exactly when this determinant is nonzero.

The vector on the right of (99) pairs with X and Y as u and v, respectively. It is therefore the Riesz representative in (90). Its trace is $(vu + uv)/D = 0$. The same characterization shows that it vanishes exactly when the trace functional vanishes on $\mathcal{C}$. Finally, an orthogonal projection fixes every vector in its image, so $I_2 \in \mathcal{C}$ implies $A_{\mathcal{C}} = I_2$; the converse is immediate from $A_{\mathcal{C}} \in \mathcal{C}$. □

Definition 8.3 (Characteristic-two good plane)**.** A two-plane $\mathcal{C} \leq \mathrm{M}_2(\mathbb{F}_q)$ is *CT-good* if

$$H(\mathcal{C}) = 0, \qquad A_{\mathcal{C}} \notin \{0, I_2\}. \tag{100}$$

This predicate is intrinsic and deterministically recognizable by constant size linear algebra. For a CT-good plane, choose any $B_{\mathcal{C}} \in \mathcal{C}$ with $\mathrm{Tr}(B_{\mathcal{C}}) = 1$. Such a vector exists because $A_{\mathcal{C}} \neq 0$ is the Riesz representative of the restricted trace functional. Lemma 8.2 gives

$$\mathrm{Gram}_b(A_{\mathcal{C}}, B_{\mathcal{C}}) = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}. \tag{101}$$

Put

$$\delta_{\mathcal{C}} = \det(A_{\mathcal{C}}), \qquad \epsilon_{\mathcal{C}} = \det(B_{\mathcal{C}}). \tag{102}$$

Lemma 8.4 (Normalized multiplication table)**.** *For a CT-good plane, the full multiplication* $X \star_{\mathcal{C}} Y = \mu_{\mathcal{C}}(X, Y)$ *satisfies*

$$A_{\mathcal{C}} \star_{\mathcal{C}} A_{\mathcal{C}} = \delta_{\mathcal{C}} A_{\mathcal{C}}, \tag{103}$$

$$A_{\mathcal{C}} \star_{\mathcal{C}} B_{\mathcal{C}} = B_{\mathcal{C}} \star_{\mathcal{C}} A_{\mathcal{C}} = (1 + \delta_{\mathcal{C}}) A_{\mathcal{C}} + \delta_{\mathcal{C}} B_{\mathcal{C}}, \tag{104}$$

$$B_{\mathcal{C}} \star_{\mathcal{C}} B_{\mathcal{C}} = B_{\mathcal{C}} + \epsilon_{\mathcal{C}} A_{\mathcal{C}}. \tag{105}$$

Furthermore,

$$[A_{\mathcal{C}}, B_{\mathcal{C}}] = I_2 + A_{\mathcal{C}}, \qquad \mathrm{Cent}(A_{\mathcal{C}}, B_{\mathcal{C}}) = \mathbb{F}_q I_2. \tag{106}$$

Proof. Cayley–Hamilton gives $A_{\mathcal{C}}^2 = \delta_{\mathcal{C}} I_2$ and $B_{\mathcal{C}}^2 = B_{\mathcal{C}} + \epsilon_{\mathcal{C}} I_2$ in characteristic two. Projection proves (103) and (105).

Write $P_{\mathcal{C}}(A_{\mathcal{C}} B_{\mathcal{C}}) = xA_{\mathcal{C}} + yB_{\mathcal{C}}$. Pairing first with $A_{\mathcal{C}}$ and then with $B_{\mathcal{C}}$, using (90) and (101), gives $y = \delta_{\mathcal{C}}$ and $x + y = 1$. This proves the displayed formula for $A_{\mathcal{C}} \star_{\mathcal{C}} B_{\mathcal{C}}$. Cyclicity of trace gives the same two pairings for $B_{\mathcal{C}} A_{\mathcal{C}}$, so the projected product is commutative.

Polarizing Cayley–Hamilton yields

$$XY + YX = \mathrm{Tr}(X)Y + \mathrm{Tr}(Y)X + \big(\mathrm{Tr}(XY) + \mathrm{Tr}(X)\,\mathrm{Tr}(Y)\big) I_2.$$

Substitution of the normalized Gram data proves the commutator identity in (106). The matrix $A_{\mathcal{C}}$ is nonscalar. Indeed, if $A_{\mathcal{C}} = \lambda I_2$ with $\lambda \neq 0$, then $I_2 \in \mathcal{C}$, so the projection fixes I_2 and forces $A_{\mathcal{C}} = I_2$. Definition 8.3 excludes this case as well as $A_{\mathcal{C}} = 0$. Thus $\mathrm{Cent}(A_{\mathcal{C}}) = \mathbb{F}_q I_2 + \mathbb{F}_q A_{\mathcal{C}}$. If $xI_2 + yA_{\mathcal{C}}$ also commutes with $B_{\mathcal{C}}$, then $y(I_2 + A_{\mathcal{C}}) = 0$. Since $A_{\mathcal{C}} \neq I_2$, this forces $y = 0$ and proves the centralizer identity. □

The normalized vector $B_{\mathcal{C}}$ is not intrinsic, but its ambiguity is exactly one scalar. This produces a complete constant-size candidate list.

Lemma 8.5 (Complete characteristic-two candidate equation)**.** *Let $\mathcal{C}, \mathcal{D}$ be CT-good planes, and choose normalized bases $(A_{\mathcal{C}}, B_{\mathcal{C}})$ and $(A_{\mathcal{D}}, B_{\mathcal{D}})$ as above. Every linear isomorphism preserving the trace form and the anchors has the form*

$$\varphi_c(A_{\mathcal{C}}) = A_{\mathcal{D}}, \qquad \varphi_c(B_{\mathcal{C}}) = B_{\mathcal{D}} + cA_{\mathcal{D}} \tag{107}$$

for a unique $c \in \mathbb{F}_q$. It preserves the full projected multiplication if and only if

$$\delta_{\mathcal{C}} = \delta_{\mathcal{D}}, \qquad \delta_{\mathcal{C}} c^2 + c = \epsilon_{\mathcal{C}} + \epsilon_{\mathcal{D}}. \tag{108}$$

If the common value δ is zero, the second equation has the unique solution $c = \epsilon_{\mathcal{C}} + \epsilon_{\mathcal{D}}$. If $\delta \neq 0$, the substitution $z = \delta c$ gives

$$z^2 + z = \delta(\epsilon_{\mathcal{C}} + \epsilon_{\mathcal{D}}), \tag{109}$$

whose solution set has cardinality zero or two and is found by $\mathbb{F}_2$-linear Gaussian elimination.

Proof. The trace-one vectors in $\mathcal{D}$ are exactly $B_{\mathcal{D}} + cA_{\mathcal{D}}$, because $\mathrm{Tr}(A_{\mathcal{D}}) = 0$. Equation (101) also shows that every map in (107) is an isometry. Comparing the images of $A_{\mathcal{C}} \star A_{\mathcal{C}}$ in the two multiplication tables forces $\delta_{\mathcal{C}} = \delta_{\mathcal{D}}$. Under this equality, comparing the images of $B_{\mathcal{C}} \star B_{\mathcal{C}}$ gives exactly the second equation in (108); the mixed-product identity then follows from (104). Conversely, the same three table comparisons show that (108) is sufficient.

When $\delta \neq 0$, multiplying the equation by δ proves (109). The map $z \mapsto z^2 + z$ is $\mathbb{F}_2$-linear on the e-dimensional bit-vector space underlying the supplied representation of $\mathbb{F}_q$, and its kernel is $\{0, 1\}$. Constructing its binary matrix and row reducing therefore returns either no solution or one affine coset of this two-element kernel. This includes $q = 2$ and takes $\mathrm{poly}(\log q)$ bit operations. □

Equation (109) is an additive prime-field linear system. It neither enumerates $\mathbb{F}_q$ nor invokes a multiplicative root-list routine.

The associated partial algorithm tests only the source code for CT-goodness. If that source-only test fails, it returns `FAILURE`; all subsequent tests against the target are exact decision steps.

Theorem 8.6 (Characteristic-two high-density recovery)**.** *Let $q = 2^e$, and let $\mathbf{A} = (A_1, A_2) \in \mathrm{M}_2(\mathbb{F}_q)^2$ be an ordered source whose image code $\mathcal{C}_{\mathbf{A}}$ is CT-good. Against every ordered target $\mathbf{B} = (B_1, B_2) \in \mathrm{M}_2(\mathbb{F}_q)^2$, a deterministic $\mathrm{poly}(\log q)$ algorithm returns `YES` with $(S, T) \in \mathrm{GL}_2(\mathbb{F}_q) \times \mathrm{GL}_2(\mathbb{F}_q)$ satisfying*

$$B_i = \sum_{j=1}^{2} t_{ij} S A_j S^{-1} \qquad (i = 1, 2) \tag{110}$$

if and only if the two ordered tensors are MCC-isomorphic. Otherwise it returns `NOT_CONJUGATE`. Once the source recognizer accepts, the procedure never returns `FAILURE`.

Proof. Compute the target image code $\mathcal{D} = \mathcal{C}_{\mathbf{B}}$. A rank other than two, a nonzero trace hull, or a target anchor in $\{0, I_2\}$ is an exact conjugacy-invariant mismatch and gives `NOT_CONJUGATE`. Otherwise choose trace-one vectors $B_{\mathcal{C}}$ and $B_{\mathcal{D}}$ by deterministic linear algebra, compute the four constants in (102), and construct the complete zero-, one-, or two-element list from Lemma 8.5.

For every listed c, solve the simultaneous linear system

$$SA_{\mathcal{C}} = A_{\mathcal{D}} S, \qquad SB_{\mathcal{C}} = (B_{\mathcal{D}} + cA_{\mathcal{D}})S. \tag{111}$$

If an invertible solution S_0 exists, the full solution space is

$$S_0 \operatorname{Cent}(A_{\mathcal{C}}, B_{\mathcal{C}}) = S_0(\mathbb{F}_q I_2)$$

by Lemma 8.4. Hence it is enough to retain exactly the one-dimensional solution spaces whose nonzero generator is invertible. For each retained generator, verify $\mathcal{D} = S\mathcal{C}_{\mathbf{A}}S^{-1}$ exactly.

Apply the coefficient lift of Proposition 3.3. It constructs the right multiplier R, sets $T = R^{\mathsf{T}}$ with the action convention of Section 3, and verifies both identities in (110). In the present full-rank two-slice regime both coefficient kernels are zero, but using the general lift fixes the transpose orientation and retains the manuscript's common verification layer.

Soundness follows from the final slice checks. For completeness, suppose a true ordered-tensor isomorphism exists. Covariance in (93) sends the source anchor to the target anchor and the chosen source trace-one vector to $B_{\mathcal{D}} + cA_{\mathcal{D}}$ for some c. The induced map preserves the complete multiplication table, so Lemma 8.5 places this c in the enumerated list. Its ambient conjugator solves (111) and survives every exact check. Thus exhaustion of the candidate list proves an exact negative answer rather than a heuristic rejection. All matrix systems have constant size, and the only growing system is the e-dimensional binary system in (109). □

Proposition 8.7 (Exact characteristic-two source mass)**.** *For every $q = 2^e$, the number of CT-good two-planes in $\mathrm{M}_2(\mathbb{F}_q)$ is*

$$q^2(q^2 - 1). \tag{112}$$

Consequently, for two iid uniform source slices, the exact probability that their image is CT-good is

$$\Pr[\mathcal{C}_{\mathbf{A}} \text{ is CT-good}] = (1 - q^{-1})(1 - q^{-2})^2. \tag{113}$$

Proof. By Lemma 8.2, zero-hull planes are the nondegenerate two-planes for the symplectic form β. Count ordered pairs (X, Y) with $\beta(X, Y) \neq 0$; each nondegenerate two-plane contributes exactly $|\mathrm{GL}_2(q)|$ such ordered bases. There are $q^4 - 1$ choices for a nonzero first vector, and then $q^4 - q^3$ choices for a second vector with nonzero β-pairing. Dividing by $|\mathrm{GL}_2(q)| = (q^2 - 1)(q^2 - q)$ gives

$$\#\{\mathcal{C} : H(\mathcal{C}) = 0\} = q^2(q^2 + 1). \tag{114}$$

It remains to remove the two excluded anchor loci. The condition $A_{\mathcal{C}} = 0$ says that $\mathcal{C}$ lies in the trace-zero hyperplane $L = \ker(\mathrm{Tr})$. The radical of $\beta|_L$ is $\mathbb{F}_q I_2$, and $L/\mathbb{F}_q I_2$ is a nondegenerate symplectic two-space. A nondegenerate two-plane in L is precisely a complement to this radical. After fixing one complement, all complements are graphs of linear maps from a two-space to a line, giving exactly q^2 such planes.

The condition $A_{\mathcal{C}} = I_2$ says $I_2 \in \mathcal{C}$. A plane of the form $\langle I_2, Y \rangle$ is nondegenerate exactly when $\beta(I_2, Y) = \mathrm{Tr}(Y) \neq 0$. Normalize Y to trace one. There are q^3 trace-one matrices, and two normalized choices determine the same plane exactly when they differ by a multiple of I_2. This gives another q^2 planes. The two excluded families are disjoint, so subtracting them from (114) proves (112).

Every two-plane has $|\mathrm{GL}_2(q)|$ ordered bases. The iid sample space has q^8 ordered pairs, including dependent pairs. Therefore

$$\frac{q^2(q^2 - 1)|\mathrm{GL}_2(q)|}{q^8} = (1 - q^{-1})(1 - q^{-2})^2,$$

which proves (113) under the original ordered-source distribution. □

Combining Theorem 8.6 and Proposition 8.7 gives the fixed-parameter high-density statement in characteristic two, including the smallest field $\mathbb{F}_2$. No target distribution was used.

8.3. The fixed-parameter theorem

The following theorem assembles the fixed-parameter consequences stated in the Introduction.

Theorem 8.8 (Fixed-parameter high-density MCC)**.** *Let the finite field be supplied in the explicit model of Section 2, and let the first ordered tensor be iid uniform while the second is arbitrary.*

1. *For every fixed $n \geq 3$ and $m \geq 1$, there is an efficiently recognizable source family of mass*
$$1 - O_{n,m}(q^{-1})$$
as q tends to infinity through prime powers.
2. *For $n = 2$ and every fixed $m \neq 2$, the same conclusion holds.*
3. *For $(n, m) = (2, 2)$ in characteristic two, a separate intrinsic branch has exact recognized iid mass*
$$(1 - q^{-1})(1 - q^{-2})^2. \tag{115}$$
Its recovery and exact count are Theorem 8.6 and Proposition 8.7, respectively.

On every source recognized in these three cases, a deterministic $\mathrm{poly}(n, m, \log q)$ *partial algorithm returns a pair $(S, T) \in \mathrm{GL}_n(\mathbb{F}_q) \times \mathrm{GL}_m(\mathbb{F}_q)$ satisfying (5) if and only if such a pair exists. Every positive answer is verified against all ordered slice equations; otherwise the algorithm returns the exact answer* `NOT_CONJUGATE`*. Once its source recognizer accepts, it never returns* `FAILURE`*.*

The present high-density theorem excludes $(n, m) = (2, 2)$ in odd characteristic. This exclusion concerns the intrinsic high-density route analyzed here; it is not a hardness statement and does not weaken the all-parameter theorem of Section 7. When $n = 1$, every source is tractable by coefficient-space linear algebra.

Proof. For $n \geq 3$ and $2 \leq m \leq N - 2$, combine Propositions 4.7 and 4.3. The exact endpoint masses (48)–(50) and the complete endpoint solvers of Propositions 7.2 and 7.3 cover $m = 1$, $m = N - 1$, and $m \geq N$. These ranges exhaust every fixed m when $n \geq 3$.

When $n = 2$, one has $N = 4$. The same endpoint branches cover $m = 1$, $m = 3$, and $m \geq 4$, proving the second assertion. The characteristic-two $(2, 2)$ construction and deterministic arbitrary-target recovery are proved in Theorem 8.6, and its exact count is Proposition 8.7; the resulting probability is (115). This also makes clear why that clause is separate from the literal projection-pair argument above. □

Theorem 8.8 is the fixed-parameter formulation. For growing dimensions, Corollary 4.8 supplies the explicit interior bound and its sufficient growth condition. The truncation in (40) remains necessary for small fields.

9. Combined dispatch, exact examples, and cost

9.1. One source recognizer for the complementary algorithms

The all-parameter family $\mathcal{G}$ need not contain the high-density zero-hull family. To combine guarantees without changing either proof, define $\mathcal{R}$ to be the union of $\mathcal{G}$, the interior event $\mathsf{Rec}_{n,m}$, and the full-rank CT-good event at $n = m = 2$ in characteristic two. The trace-good endpoint branches already belong to $\mathcal{G}$. Evaluate the applicable source predicates and cache their intrinsic data before inspecting the target. At least one complete local solver is available precisely on this union.

Proposition 7.8's source-certified branch argument applies unchanged: a negative answer from the selected complete solver is conclusive. Thus Algorithm 4 preserves both headline guarantees on their respective events without any independence assumption. The thinning $\mathcal{G}^\star$ is not used by this dispatch.

Algorithm 4 Combined partial solver for ordered matrix-space conjugacy

Require: An explicit $\mathbb{F}_q$ and tuples $A, B \in \mathrm{M}_n(\mathbb{F}_q)^m$
Ensure: Verified (S, T), exact `NOT_CONJUGATE`, or source-only `FAILURE`

Compute the source rank, image code, and the applicable predicates
Build the list of applicable source-certified branches:
    zero-hull, characteristic-two two-plane, and the six-row $\mathcal{G}$ dispatch
if this list is empty **then**
    return `FAILURE`
end if
Compute the target rank; if it differs from the source rank, return `NOT_CONJUGATE`
Choose the first available branch in the displayed deterministic order
Run its complete code solver on the arbitrary target
if the solver returns `NOT_CONJUGATE` **then**
    return `NOT_CONJUGATE`
end if
for each ambient candidate S retained by that branch **do**
    Verify image-code equality; then apply Algorithm 1
    if the ordered equations hold **then**
        return the verified pair (S, T)
    end if
end for
return `NOT_CONJUGATE`

9.2. The image-and-kernel lift on a dependent presentation

Example 9.1 (Recovering the declared coefficient action). Continue Example 4.4 and present its code by $\mathbf{A} = (X, Y, X + Y)$. Thus $m = 3$ but $r_A = 2$. Let

$$S = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \qquad T_0 = \begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix},$$

and form $B_i = \sum_j (T_0)_{ij} S A_j S^{-1}$. The actual target is

$$B_1 = \begin{pmatrix} 1 & 4 & 1 \\ 1 & 2 & 2 \\ 1 & 0 & 4 \end{pmatrix}, \quad B_2 = \begin{pmatrix} 1 & 4 & 1 \\ 1 & 2 & 2 \\ 1 & 0 & 4 \end{pmatrix}, \quad B_3 = \begin{pmatrix} 1 & 4 & 2 \\ 4 & 4 & 3 \\ 0 & 1 & 2 \end{pmatrix}.$$

The target code has the projected pair $(SA_E S^{-1}, SM_E S^{-1})$. Its simultaneous intertwiner space with the source pair is $\mathbb{F}_5 S$. Algorithm 1, using the first pivot columns of the coefficient matrices, returns

$$T = \begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & 1 \\ 2 & 1 & 0 \end{pmatrix}.$$

Direct multiplication gives $P_B = P_{S,A} T^{\mathsf{T}}$ and verifies all 27 target matrix entries. Here T need not equal T_0: dependent presentations have coefficient-kernel freedom. The returned T is invertible, so it witnesses the original declared three-dimensional coefficient action, rather than a reduced two-slice problem.

This presentation does not lie in the full-rank interior event $\mathsf{Rec}_{3,3}$; it illustrates the code-level recovery and ordered lift, not that event's probability. Replacing the target by three zero matrices is rejected immediately by the invariant rank comparison. More generally, a tuple intertwiner is accepted only after verifying the entire image code and all original slices.

9.3. A coarse operation budget

Write $N = n^2$ and $d = \dim C \le \min\{m, N\}$. Ordinary dense elimination, without a fast-multiplication assumption, gives the following budget for the projection and one-hull branches, including ordered lifting.

Proposition 9.2 (Explicit linear-algebra budget)**.** *Apart from the line/hyperplane rational-canonical-form routine, the selected algorithms use*

$$O(mN \min\{m, N\} + m^3 + N^3 + e^2)$$

operations in $\mathbb{F}_q$ when $q = 2^e$, together with $O(e^3)$ binary operations for a characteristic-two Artin–Schreier solve when that branch is used. For odd q, the e^2 and binary-solve terms are absent. These are coarse upper bounds, not claims of fast practical performance.

Proof. An $N \times m$ coefficient matrix is reduced in $O(mN \min\{m, N\})$ operations. A trace Gram matrix on d basis matrices costs $O(d^2N)$ to assemble and $O(d^3)$ to reduce. Since $d \le N$, these costs are bounded by $O(N^3)$. Orthogonal complements and their Gram matrices have the same bound. Matrix powers and the cyclicity test use at most n dense $n \times n$ multiplications, followed by elimination on $N \times n$ vectorizations; both costs are absorbed by $O(N^3)$. Centralizers and intertwiners have at most N unknowns and a constant multiple of N equations. Riesz representatives require at most two Gram solves. In the selected characteristic-two scale routine, the polynomial degrees are bounded by three; Lemma 5.7 uses $O(e^2)$ field operations and at most three scale candidates. Proposition 3.3 adds $O(mn^3 + n^2m^2 + m^3)$ operations. The inequalities $mn^3 \le m^3 + N^3$ and $Nm^2 \le m^3 + N^3$ absorb these terms. The two-plane additive solve is elimination on an e-dimensional binary space. This proves the displayed budget. □

For the line and hyperplane branches, add the cost of deterministic rational canonical form with a transforming basis on $n \times n$ matrices. Polynomial Euclidean and Smith-form arithmetic supplies a polynomial bound, as used in Proposition 7.2; no claim about a specific fast exponent is needed here. Thus the end-to-end bit cost remains $\mathrm{poly}(n, m, \log q)$ through the supplied field representation. The probability estimates do not belong to this operation count.

Computational material. The computational scripts and recorded outputs supporting the worked examples, the characteristic-two degree-at-most-three splitter checks, the retained finite regression checks, and the eight shifted-polynomial certificates in Lemma E.3 are archived on Zenodo as *Computational Checks for Deterministic Partial Algorithms for Matrix-Space Conjugacy over Finite Fields*, version 1.0.0 (Ma, 2026). The archive was prepared by Jingchuan Ma. The archive includes reproduction instructions and a mapping from each check to the corresponding statement in the manuscript. The scripts use Python's standard library. These computations provide auxiliary validation only; they are not a complete implementation of the combined dispatcher, do not constitute benchmark evidence, and are not dependencies of the universal mathematical proofs.

10. Arithmetic boundaries of the constructions

The selected scale routines do not provide a general root-extraction primitive. The following boundaries explain the restriction to those routines.

For an arbitrary matrix, all positive power traces vanish exactly when each nonzero eigenvalue has algebraic multiplicity divisible by the characteristic. Equivalently, $\chi_A(x) = x^{e_0}h(x)^p$ with $h(0) \neq 0$. Proposition H.1 proves this using a Vandermonde system. Outside that family, the least transverse moment has index at most $n - 1$, but its projective scale can require an equation $\alpha^k = \rho$ of degree growing with n. The geometric Riesz construction therefore does not, by itself, provide a deterministic algorithm for all adaptive moments.

Table 1: The closest constructive comparison. Each density concerns the stated source law; no target distribution is assumed for our guarantees.

	LLQTW	Universal construction	Zero-hull construction
Parameter scope	$m = n$	All $q \geq 2, n, m \geq 1$	$n \geq 3$, interior m
Source guarantee	$1/\Theta(q)$	$> 9/(1216q)$	$\geq 1-(6n^2-4)s/q-2/q^3$
Algorithm	Randomized bit-polynomial	Deterministic bit-polynomial	Deterministic bit-polynomial
Intrinsic data	Hull and spectral/coefficient data	Hull moments and small-dimensional branches	Two projections from a zero-hull code
Coverage caveat	Dimension-matched construction	Partial, not all-source	Useful density only in a large-field regime

Proposition H.3 gives a fixed five-dimensional source family with scalar pair centralizer whose ambient transporters reveal fifth roots. This pointwise reduction does not force a partial random-source solver to recognize that family. Likewise, Proposition H.4 identifies an r-th-root interface in worst-case companion-line and trace-dual hyperplane search. Projective-similarity decision can nevertheless use invariant factors and exponentiation without first constructing the root (Proposition H.5). Full statements and proofs are in Appendix H.

For odd-characteristic two-planes in M_2, a specified generic projective-transporter cover has degree two and a nonsplit quadratic function-field extension (Theorem G.2). Appendix G proves the precise geometric statement, including a dominant chart. This no-section statement concerns the specified cover, not finite-field complexity; the universal good-two-plane branch does not depend on it.

These distinctions matter for comparison with deterministic factoring algorithms: for example, the general framework of Guo (2020) explicitly uses GRH and hypotheses on a lifted polynomial. None of those hypotheses, and no general factoring oracle, enters the selected MCC procedures.

11. Comparison with algebraic-isomorphism algorithms

Matrix spaces versus ordered tuples. The broad tensor-isomorphism relationships established by Grochow and Qiao (2023) explain why matrix-space problems recur in algebraic isomorphism. The earlier algorithms of Grochow et al. (2022) study other source-density and dimension-exponential regimes. Such worst-case reductions do not identify their random input distributions with ours. For ordered matrix tuples, Qiao and Sun (2024) give randomized Las Vegas polynomial-time canonical-form algorithms. Bongartz and Friedland (2026) study complete invariants for simultaneous similarity. Neither problem includes the unknown coefficient mixing of an unordered matrix space. Our scalar-centralizer certificates allow elementary intertwiners instead of a canonical-form subroutine.

The direct trace-hull lineage. Table 1 compares the two source guarantees with Li et al. (2026). Their Sections 2.3 and 3 already include constructive conjugacy-coset and candidate checking. In Couvreur and Levrat (2025, Section 2.3), the conjugacy specialization has a heuristic polynomial-time interpretation, which is distinct from a proved uniform iid-source guarantee. The new covariants and conditional counts here supply the latter without changing the original coefficient dimension.

The broader left–right matrix-code equivalence problem is treated by Narayanan et al. (2024); Reijnders et al. (2024); Couvreur and Levrat (2025). Tensor monomial isomorphism and matrix-subcode equivalence supply neighboring code-equivalence settings (D'Alconzo, 2024; Bardet et al., 2026), but their acting groups are not the inverse-diagonal conjugation subgroup studied

here. In particular, Highway's relevant multiplier outputs lie in the trace-zero hyperplane; the odd-characteristic nonzero-trace branch cannot simply be transferred to that output distribution. We make no direct cryptosystem-break claim.

Trace duality and finite formed spaces. Complementary-dual codes and code hulls predate the present algorithms (Massey, 1992; Sendrier, 1997, 2000). Nondegenerate-subspace and formed-space enumeration is likewise established mathematics (Klopsch and Voll, 2009; Glasby et al., 2022, 2023; Carlet et al., 2019). Here it is applied to marked projections. The cyclic-matrix and separability results of Neumann and Praeger (1995); Wall (1999) are external inputs, not new contributions. The characteristic-two and small-size counts for the selected law are proved in the appendices.

Neighboring symbolic-computation tasks. Murray and Roney-Dougal (2011) construct classical-group homomorphisms and isometries of forms, with explicit Las Vegas and, for some outputs, discrete-logarithm qualifications. A form isometry need not be induced by matrix conjugation; those algorithms do not solve our coefficient-mixed matrix-space problem. Kreuzer (2025) relates monomial linear-code equivalence to projective point sets and polynomial isomorphism, with regularity-dependent complexity. This is useful algebraic-isomorphism context, not an identification of its code action with MCC. We rely on none of these neighboring tasks as an oracle in the deterministic solver.

12. Conclusion

Trace geometry supplies intrinsic matrices without first matching the input generators. On zero-hull spaces, two projections yield a characteristic-wise polynomial certificate and an explicit high-density bound. On a complementary family, one-hull moments and exact quotient counts yield a uniform positive source fraction for every finite field and all matrix and coefficient dimensions. A deterministic image-and-kernel lift recovers the ordered coefficient action as well as the ambient conjugator.

The result remains a partial-algorithm theorem. Geometric nonemptiness is not a small-field witness guarantee, source probability is not algorithmic randomness, and construction-specific arithmetic obstructions are not complexity lower bounds. Removing the odd-characteristic $(2, 2)$ high-density exception or expanding the recognized families requires additional arguments beyond those given here.

Appendix A. Geometric witnesses for the projection algorithm

This appendix proves Proposition 4.5. All fields used for the geometric construction may be extensions of the prescribed prime field.

Proof. We first reduce the projection constraints to four trace identities. Work over a field k of the prescribed characteristic and choose matrices A, M. Put

$$U = I_n - A, \qquad V = A^2 - M, \qquad L = \langle A, M \rangle, \qquad R = \langle U, V \rangle. \tag{A.1}$$

The condition $L \perp R$ is equivalent to

$$\mathrm{Tr}(A^2) = \mathrm{Tr}(A), \qquad \mathrm{Tr}(AM) = \mathrm{Tr}(M) = \mathrm{Tr}(A^3), \qquad \mathrm{Tr}(A^2M) = \mathrm{Tr}(M^2). \tag{A.2}$$

If a nondegenerate r-space E satisfies $L \le E \le R^{\perp_b}$, then $U, V \in E^{\perp_b}$ and hence

$$P_E(I_n) = A, \qquad P_E(A^2) = M.$$

It therefore suffices to construct a pair satisfying (A.2), with scalar common centralizer, and then to complete L to such a nondegenerate space E.

Odd characteristic. Let $x_1, \dots, x_{n-1}$ be algebraically independent over the prime field. Put

$$c = \sum_{i<n} (x_i^2 - x_i),$$

and choose, in an algebraic extension of the rational function field, a root a_n of

$$z^2 - z + c = 0.$$

Set $a_i = x_i$ for $i < n$, $A = \mathrm{diag}(a_1, \dots, a_n)$, and

$$p_j = \sum_{i=1}^{n} a_i^j, \qquad t = p_1, \qquad s_3 = p_3.$$

The defining quadratic gives $p_2 = p_1 = t$. The a_i are pairwise distinct, and t and $n - t$ are nonzero. Indeed, any equality $a_n = a_j$, $t = 0$, or $t = n$ would turn the defining quadratic into a nonzero polynomial in the independent x_i; in the last two cases the cross monomial $2x_1x_2$ already prevents vanishing. Thus these are nonempty open conditions.

Adjoin a further transcendental u. Solve the invertible Vandermonde system

$$\begin{pmatrix} 1 & 1 & 1 \\ a_1 & a_2 & a_3 \\ a_1^2 & a_2^2 & a_3^2 \end{pmatrix} \begin{pmatrix} d_1 \\ d_2 \\ d_3 \end{pmatrix} = \begin{pmatrix} s_3 \\ s_3 \\ u \end{pmatrix}, \qquad d_i = 0 \quad (i > 3). \tag{A.3}$$

Let $D = \mathrm{diag}(d_1, \dots, d_n)$ and let

$$Q_0 = E_{12} + E_{23} + \cdots + E_{n-1,n} + E_{n1}.$$

With

$$w = \frac{u - \sum_i d_i^2}{2}, \qquad Q(w) = Q_0 + wE_{21}, \qquad M = D + Q(w), \tag{A.4}$$

the only length-two closed walk in $Q(w)$ uses the two edges E_{12} and wE_{21}, so $\mathrm{Tr}(Q(w)^2) = 2w$. Equations (A.3)–(A.4) give

$$\mathrm{Tr}(M) = \mathrm{Tr}(AM) = s_3 = \mathrm{Tr}(A^3), \qquad \mathrm{Tr}(A^2M) = u = \mathrm{Tr}(M^2).$$

Together with $p_2 = p_1$, these are precisely (A.2).

Since A has simple spectrum, every matrix commuting with A is diagonal. The coefficient-one Hamiltonian cycle in Q_0 forces all diagonal entries of a matrix commuting also with M to agree. Hence $\mathrm{Cent}(A, M) = kI_n$. The Gram determinants on L and R are

$$\det(b|_L) = tu - s_3^2, \tag{A.5}$$

$$\det(b|_R) = (n - t)(p_4 - u) - (t - s_3)^2. \tag{A.6}$$

They are nonzero because u is transcendental and both t and $n-t$ are nonzero. Finally, A, M, U, V are independent: comparison of off-diagonal parts in a relation reduces its diagonal part to $\alpha A^2 + \beta A + \gamma I_n = 0$, and a quadratic polynomial cannot vanish at the $n \geq 3$ distinct eigenvalues of A unless all coefficients vanish.

Thus $K = R^{\perp_b}$ is a nondegenerate $(N-2)$-space containing the nondegenerate two-space L. Its orthogonal complement to L inside K is nondegenerate of dimension $N - 4$. Over the algebraic closure it contains a nondegenerate subspace of every dimension from 0 to $N - 4$. Adjoining one of dimension $r - 2$ orthogonally to L gives the required E.

Characteristic two. Here the trace form satisfies

$$b(X, X) = \mathrm{Tr}(X^2) = \mathrm{Tr}(X)^2. \tag{A.7}$$

For a nondegenerate r-space E and $A = P_E(I_n)$, the vector representing the square-root functional in (A.7) has norm r mod 2. Equivalently,

$$\mathrm{Tr}(A) = r \bmod 2. \tag{A.8}$$

To see this directly, split the restricted symmetric form orthogonally into anisotropic lines and alternating hyperbolic planes. Each anisotropic line contributes one to the norm of the representing vector, while each alternating plane contributes zero; the number of lines has the parity of r.

Let $\epsilon \in \{0, 1\} \subset k$ be the image of r mod 2. For algebraically independent $x_1, \ldots, x_{n-1}$ set

$$a_i = x_i \ (i < n), \qquad a_n = \epsilon + \sum_{i<n} x_i, \qquad A = \mathrm{diag}(a_1, \ldots, a_n), \qquad s_3 = \sum_i a_i^3. \tag{A.9}$$

The eigenvalues are pairwise distinct and $\mathrm{Tr}(A) = \mathrm{Tr}(A^2) = \epsilon$. Both s_3 and $s_3 + \epsilon$ are nonzero: after the substitution in (A.9), the mixed monomial $x_1^2 x_2$ occurs with coefficient one, and hence the cubic sum cannot be either constant. This is where $n \geq 3$ is used.

Use the same three-column Vandermonde matrix as in (A.3), now with right-hand side $(s_3, s_3, s_3^2)^{\mathsf{T}}$, and again put $d_i = 0$ for $i > 3$. With

$$D = \mathrm{diag}(d_1, \ldots, d_n), \qquad M = D + Q_0, \tag{A.10}$$

one obtains

$$\mathrm{Tr}(M) = \mathrm{Tr}(AM) = s_3 = \mathrm{Tr}(A^3), \qquad \mathrm{Tr}(A^2 M) = s_3^2 = \mathrm{Tr}(M^2),$$

where the last equality uses (A.7). Thus (A.2) holds. Simple spectrum and the Hamiltonian cycle again give $\mathrm{Cent}(A, M) = kI_n$, and the quadratic-polynomial argument again proves that A, M, U, V are independent.

Put $d = N - r$ and let $\bar{d} \in \{0, 1\}$ be its parity. Direct substitution in the trace identities gives

$$G_L = \begin{pmatrix} \epsilon & s_3 \\ s_3 & s_3^2 \end{pmatrix}, \qquad G_R = \begin{pmatrix} \bar{d} & \epsilon + s_3 \\ \epsilon + s_3 & (\epsilon + s_3)^2 \end{pmatrix}. \tag{A.11}$$

Since s_3 and $s_3 + \epsilon$ are nonzero,

$$\dim \mathrm{rad}(L) = \epsilon, \qquad \dim \mathrm{rad}(R) = \bar{d}. \tag{A.12}$$

Let $K = R^{\perp_b}$. Then $\mathrm{rad}(K) = \mathrm{rad}(R)$, and the independence of A, M, U, V gives $L \cap R = 0$.

If $\epsilon = 0$, set $Q = L$. If $\epsilon = 1$, the vector $\rho = s_3 A + M$ spans $\mathrm{rad}(L)$. It does not lie in $K^{\perp_b} = R$, so choose $x \in K$ with $b(x, \rho) = 1$ and replace it by

$$x' = x + b(x, A)A.$$

Then $b(x', A) = 0$ and $b(x', \rho) = 1$. Moreover, x' is orthogonal to both A and $U = I_n - A$, so $\mathrm{Tr}(x') = 0$ and $b(x', x') = 0$ by (A.7). Consequently $Q = \langle A, \rho, x' \rangle$ is nondegenerate and contains L. In both parities,

$$\dim Q = 2 + \epsilon. \tag{A.13}$$

The orthogonal remainder $J = K \cap Q^{\perp_b}$ is alternating: every $Y \in J$ is orthogonal to A and U, hence to I_n, so $\mathrm{Tr}(Y) = 0$ and (A.7) gives $b(Y, Y) = 0$. Furthermore,

$$\dim(J/\operatorname{rad} J) = N - 4 - \epsilon - \bar{d}.$$

The desired residual dimension

$$h = r - 2 - \epsilon$$

is nonnegative and even, and

$$(N - 4 - \epsilon - \bar{d}) - h = d - 2 - \bar{d} \geq 0.$$

Choose $h/2$ hyperbolic pairs in the symplectic quotient and lift them to a nondegenerate h-space $J_h \leq J$. Then $E = Q \perp J_h$ is nondegenerate, has dimension r, contains L, and lies in $R^{\perp_b}$. The initial reduction therefore gives $A_E = A$, $M_E = M$, and $\mathrm{Cent}(A_E, M_E) = kI_n$.

The constructions in both characteristics are symbolic. Algebraically independent parameters and extension fields prove geometric nonemptiness; they are not being specialized to claim a rational point over each small finite field. □

Appendix B. Odd-characteristic formed-space counts

This appendix fixes the type convention and derives the counts used in Proposition 6.5. All forms in this appendix are nondegenerate symmetric bilinear forms over an odd finite field.

Lemma B.1 (Norm counts, types, and orthogonal orders)**.** *For an even-dimensional space V of dimension $2a$, define its type by $\epsilon(V) = \chi_q((-1)^a \det b_V)$. Then*

$$|O(2a+1,q)| = 2q^{a^2}\prod_{i=1}^{a}(q^{2i}-1), \qquad |O^{\epsilon}(2a,q)| = 2q^{a(a-1)}(q^a-\epsilon)\prod_{i=1}^{a-1}(q^{2i}-1).$$

The number of isotropic lines is $(q^{2a}-1)/(q-1)$ in dimension $2a+1$, and $(q^{a-1}+\epsilon)(q^a-\epsilon)/(q-1)$ in dimension $2a$. For the trace form on $\mathrm{M}_n(\mathbb{F}_q)$ with even n,

$$\epsilon(\mathrm{M}_n) = \chi_q((-1)^{n/2}).$$

An isotropic-line quotient $L^{\perp}/L$ has the same even-dimensional type as its ambient even-dimensional space.

Proof. For a diagonal form $Q(x) = \sum_{i=1}^{s} a_i x_i^2$, additive-character orthogonality counts a norm fibre by

$$\#\{Q(x) = t\} = q^{-1}\sum_{u\in\mathbb{F}_q}\psi(-ut)\prod_i\sum_{x\in\mathbb{F}_q}\psi(ua_ix^2).$$

For $u \neq 0$, the inner sum is $\chi_q(ua_i)G$, where $G = \sum_x \psi(x^2)$ and $G^2 = \chi_q(-1)q$. The latter identity follows by writing $G = \sum_{a\neq 0}\chi_q(a)\psi(a)$, multiplying the two sums, and substituting $b = at$; the inner sum over $a \neq 0$ is $q-1$ when $1+t=0$ and -1 otherwise. Thus, in dimension $2a$, the zero fibre has size $q^{2a-1} + \epsilon(q-1)q^{a-1}$ and each nonzero fibre has size $q^{2a-1} - \epsilon q^{a-1}$. In dimension $2a+1$, the zero fibre has size q^{2a} and the fibre at $t \neq 0$ has size $q^{2a} + \chi_q((-1)^a \det b_V)\chi_q(t)q^a$. Subtracting the zero vector and dividing by $q-1$ gives the line counts.

These counts also show that every form of dimension at least three contains a hyperbolic plane. Splitting such planes reduces classification to a line or a binary plane, where determinant square class distinguishes the two possibilities. Given a nonzero isotropic vector v, there are q^{s-2} isotropic w satisfying $b(v,w) = 1$: after one such w_0 is chosen, write $w = w_0 + z - \frac{1}{2}b(z,z)v$, with $z \in \langle v, w_0\rangle^{\perp}$. Isometries are transitive on these ordered hyperbolic pairs because an isometry of the pair extends by an isometry of their complements, which have equal dimension and determinant class. Therefore

$$|O(V)| = (\#\{v \neq 0 : Q(v) = 0\})q^{s-2}|O(H^{\perp})|.$$

The base orders are $|O(1,q)| = 2$, $|O^+(2,q)| = 2(q-1)$, and $|O^-(2,q)| = 2(q+1)$. For the last two, use a split isotropic basis, or identify the anisotropic plane with the norm form of $\mathbb{F}_{q^2}/\mathbb{F}_q$;

its isometries are the norm-one multiplications and their compositions with field conjugation. Induction gives the displayed orders.

For the matrix-unit basis of M_n, each diagonal unit has Gram entry one, and each pair E_{ij}, E_{ji} contributes $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$. Consequently $\det b_W = (-1)^{n(n-1)/2}$. When n is even, multiplying by $(-1)^{n^2/2}$ gives the square class $(-1)^{n/2}$. Finally, choose a hyperbolic partner to a generator of L. Then $L^\perp/L \cong H^\perp$; the determinant of H is -1, which cancels the change in the exponent in the definition of type. □

Lemma B.2 (Nondegenerate subspace counts)**.** *Write the orders in Lemma B.1 as o_{2a+1} and o^τ_{2a}. For $1 \le k \le D-1$, the four formulas* (E.13)–(E.16) *count all nondegenerate k-subspaces of a D-space, with the indicated ambient type.*

Proof. A nondegenerate U gives an orthogonal decomposition $V = U \perp U^\perp$. For prescribed isometry classes of the two summands, $O(V)$ acts transitively on such U: take any isometries between corresponding summands and form their direct sum. Its stabilizer is exactly $O(U) \times O(U^\perp)$, not a quotient or a subgroup exchanging the summands, since U itself is marked.

If D is odd, exactly one summand is even. Its two possible types give the two terms $1/o^+_{2a} + 1/o^-_{2a}$, or the analogous complement terms. The other summand's determinant class is forced, but its odd orthogonal order is independent of that class. If D and k are even, the two even types multiply to the ambient type ϵ, giving the sum over $(\tau, \epsilon\tau)$. If D is even and k is odd, there are two determinant classes for the odd summand, the complement class is forced, and both choices have the same order; this gives the factor 2. All summand dimensions in these uses are positive, so no zero-dimensional type is being counted twice. This derives every parity case and every denominator in the four formulas. □

Lemma B.3 (Trace slices of the isotropic cone)**.** *Let $W = \mathrm{M}_n(\mathbb{F}_q)$ with q odd, $N = n^2$, and $A \neq 0$ uniform subject to $\mathrm{Tr}(A^2) = 0$. If n is even, $N = 2s$, then* (E.28) *gives the exact probability $\Pr[\mathrm{Tr}(A) \neq 0]$. If n is odd, the lower bound* (E.21) *holds.*

Proof. The trace functional is $b(I_n, \cdot)$ and $b(I_n, I_n) = n$ in the field. If this norm is nonzero, the slice $\mathrm{Tr}(A) = a$ is $A = (a/n)I_n + Y$, with $Y \in I_n^\perp$ and $b(Y, Y) = -a^2/n$. Apply the nonzero norm-fibre formula from the proof of Lemma B.1 to this $(N-1)$-space. When N is even, each $a \neq 0$ contributes $q^{N-2} + \epsilon q^{s-1}$. If $n = 0$ in the field, complete I_n to a hyperbolic pair. For fixed $a = b(I_n, A) \neq 0$ and arbitrary coordinates in the orthogonal remainder, the zero-norm equation determines the coefficient of I_n uniquely. Each such slice then has q^{N-2} points. Dividing these counts by $T_\epsilon = q^{N-1} - 1 + \epsilon(q-1)q^{s-1}$ proves the even formula.

For odd N, put $u = q^{(N-1)/2}$. When $n \neq 0$, each nonzero trace slice contains at least $q^{N-2} - q^{(N-3)/2}$ points. Dividing their sum by $q^{N-1} - 1 = u^2 - 1$ gives $(1 - q^{-1})u/(u+1) = (1 - q^{-1})/(1 + q^{-(N-1)/2})$. When $n = 0$, the hyperbolic-pair calculation gives the larger value $(q-1)q^{N-2}/(q^{N-1} - 1)$. □

Lemma B.4 (Noncyclic subtraction with the actual conditioning)**.** *For $n \ge 2$, the all-matrix noncyclic estimate is*

$$\Pr_{A \in \mathrm{M}_n(\mathbb{F}_q)}[A \text{ is noncyclic}] < \frac{1}{(q^2-1)(q-1)}.$$

This is the Neumann–Praeger bound (Neumann and Praeger, 1995), also stated explicitly in Wall (1999, Eq. (1.1), p. 254). For any nonzero cone or slice population of size T, its conditional noncyclic fraction is less than $q^N/[T(q^2-1)(q-1)]$.

Proof. The unconditional bound is the cited result. At most the entire unconditional exceptional population can lie in the conditioned set; division by T proves the assertion. In odd dimension take $T = q^{N-1} - 1$. For even dimension, $T = T_\epsilon \ge q^{N-1}(1 - \delta)$, where $\delta = (q-1)q^{-N/2} + q^{1-N}$. In characteristic two take $T = q^{N-1} - 1$. These substitutions are exactly (E.22), (E.29), and (E.40), respectively. □

Appendix C. Krylov-pair lower counting

Lemma C.1 (Krylov-pair lower counting)**.** *Let q be a prime power, $n \geq 1$, and $\mathcal{F}$ a set of monic polynomials of degree n over $\mathbb{F}_q$. Fix $0 \neq v \in \mathbb{F}_q^n$ and put $B_n = |\mathrm{GL}_n(\mathbb{F}_q)|/(q^n - 1)$. The number of matrices A for which v is cyclic and $\chi_A \in \mathcal{F}$ is exactly $B_n|\mathcal{F}|$. The number of pairs (A, w) with w cyclic for A and $\chi_A \in \mathcal{F}$ is exactly $|\mathrm{GL}_n(\mathbb{F}_q)||\mathcal{F}|$. Consequently the number of cyclic matrices with characteristic polynomial in $\mathcal{F}$ is at least $B_n|\mathcal{F}|$.*

For the projective conclusion, assume additionally that $\varnothing \neq \mathcal{Z} \subseteq \mathrm{M}_n(\mathbb{F}_q) \setminus \{0\}$, $\lambda\mathcal{Z} = \mathcal{Z}$ for every $\lambda \in \mathbb{F}_q^\times$, and

$$f_\lambda(t) := \lambda^n f(t/\lambda) \in \mathcal{F} \qquad (f \in \mathcal{F},\ \lambda \in \mathbb{F}_q^\times).$$

Let $\mathcal{A}_\mathcal{F}$ be the set of cyclic matrices A with $\chi_A \in \mathcal{F}$, and assume $\mathcal{A}_\mathcal{F} \subseteq \mathcal{Z}$. Then the fraction of lines in $\mathcal{Z}/\mathbb{F}_q^\times$ represented by $\mathcal{A}_\mathcal{F}$ is at least $B_n|\mathcal{F}|/|\mathcal{Z}|$. This last expression is the exact fraction for the subset on which the specified v is cyclic, not in general the exact fraction of all good lines.

Proof. Choose an ordered basis $E = (v, v_2, \ldots, v_n)$ and $f \in \mathcal{F}$. There is a unique matrix $A = E\,\mathrm{Comp}(f)E^{-1}$ satisfying $Av_i = v_{i+1}$ for $i < n$ and the last-column relation prescribed by f. Thus $E = (v, Av, \ldots, A^{n-1}v)$. Conversely this Krylov basis and f are uniquely determined by A when v is cyclic. The number of bases with first vector v is

$$(q^n - q)(q^n - q^2)\cdots(q^n - q^{n-1}) = B_n.$$

This proves the fixed-vector assertion; summing over $q^n - 1$ nonzero vectors proves the pair count.

For a cyclic matrix with characteristic polynomial f, its module is $\mathbb{F}_q[x]/(f)$. A vector is cyclic exactly when its residue class is a unit, so its number of cyclic vectors is

$$c(f) = q^n \prod_{h \mid f \text{ irreducible}} (1 - q^{-\deg h}) \leq q^n - 1.$$

The pair count is therefore a weighted sum of matrices, with weights $c(f)$ that need not be constant. Dividing by $q^n - 1$ gives the asserted lower bound, not an equality for the whole population. Equivalently, the fixed-v matrices are an explicitly counted subset. For the projective conclusion, $\chi_{\lambda A}(t) = \lambda^n \chi_A(t/\lambda)$, and the Krylov columns for λA are those for A multiplied by $1, \lambda, \ldots, \lambda^{n-1}$. Thus both $\mathcal{A}_\mathcal{F}$ and its fixed-v subset are scaling invariant under the stated closure condition. Each nonzero projective matrix line has exactly $q - 1$ generators. Dividing the ambient and subset cardinalities by $q - 1$ proves the line statement and the asserted exact subset fraction. □

Corollary C.2 (Trace conditions used in the interior count)**.** *Write $f = x^n + c_1x^{n-1} + c_2x^{n-2} + \cdots + c_n$, with $n \geq 3$. In odd characteristic the conditions $\mathrm{Tr}(A^2) = 0$ and $\mathrm{Tr}(A) \neq 0$ allow exactly $(q-1)q^{n-2}$ characteristic polynomials. In characteristic two the conditions $\mathrm{Tr}(A) = 0$ and $\mathrm{Tr}(A^3) \neq 0$ allow the same number. Over $\mathbb{F}_2$, the latter number is 2^{n-2}.*

Proof. Newton identities give $\mathrm{Tr}(A) = -c_1$ and $\mathrm{Tr}(A^2) = c_1^2 - 2c_2$. In odd characteristic choose $c_1 \neq 0$, set $c_2 = c_1^2/2$, and choose the remaining $n - 2$ coefficients freely. In characteristic two set $c_1 = 0$; the third Newton identity becomes $\mathrm{Tr}(A^3) = c_3$. Choose $c_3 \neq 0$ and all other coefficients freely. For $q = 2$ the choice is $c_3 = 1$. Both polynomial families satisfy the scaling-closure assumption of Lemma C.1, since $\mathrm{Tr}((\lambda A)^j) = \lambda^j \mathrm{Tr}(A^j)$ for $\lambda \neq 0$. □

For odd $n = 3$, the nonzero isotropic cone has size $q^8 - 1$, and $B_3(q-1)q = q^4(q-1)^3(q+1)$. For $n = 4$ the trace form is of plus type, so the cone size is $q^{15} - 1 + (q-1)q^7$. These give exactly the displayed subset ratios δ_3^{odd} and δ_4^{odd}. In characteristic two the nonzero trace-zero set has size $q^{n^2-1} - 1$. Finally

$$B_n 2^{n-2} = 2^{n^2-2} \prod_{j=1}^{n-1} (1 - 2^{-j}) = G_n.$$

Thus the equalities defining B_n, the displayed δ expressions, and G_n are exact algebraic counts or definitions. Their uses as fractions of the entire cyclic population are lower bounds. The inequalities needed for the probability constants are established in Appendix E.

Appendix D. Binary formed spaces and correlated projection

In this appendix b is a nondegenerate symmetric bilinear form over $\mathbb{F}_2$. Its characteristic vector is the unique c such that $b(x,x) = b(c,x)$ for every x.

Lemma D.1 (Binary orthogonal orders and marked decompositions)**.** *A nonalternating space of dimension D is isometric to $\mathbb{F}_2^D$ with Gram matrix I_D. Its characteristic vector is nonzero and has norm $D \bmod 2$. Put*

$$o_{2a} = 2^{a^2} \prod_{i=1}^{a-1} (2^{2i} - 1), \qquad o_{2a+1} = 2^{a^2} \prod_{i=1}^{a} (2^{2i} - 1).$$

For $1 \le k \le D-1$, let $\mathrm{ND}_{oo}(D,k)$ count nondegenerate k-subspaces U such that both U and $U^\perp$ are nonalternating. Then

$$\mathrm{ND}_{oo}(D,k) = \frac{o_D}{o_k o_{D-k}}.$$

In particular this gives all four parity cases of (E.46)*.*

Proof. Split off anisotropic lines until the remainder is alternating. A line $\langle e \rangle$ of norm one together with a hyperbolic pair u, v can be replaced by the orthonormal basis $e+u, e+v, e+u+v$. This absorbs every alternating plane and proves the classification. In an orthonormal basis c is the sum of the basis vectors, so $b(c,c) = D \bmod 2$.

If $D = 2a+1$, the decomposition $\langle c \rangle \perp c^\perp$ has symplectic second summand. Restriction gives $O(D,2) \cong \mathrm{Sp}(2a,2)$. If $D = 2a$, choose d with $b(c,d) = 1$. Necessarily $b(d,d) = 1$, and $\langle c, d \rangle$ has Gram matrix $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right)$; its complement is symplectic. The image of d has 2^{D-1} choices. For each, an isometry of the symplectic complement extends uniquely while fixing c. Thus $|O(D,2)| = 2^{D-1} |\mathrm{Sp}(D-2,2)|$. Counting an ordered symplectic pair gives the recurrence $|\mathrm{Sp}(2a,2)| = (2^{2a}-1) 2^{2a-1} |\mathrm{Sp}(2a-2,2)|$, which proves the two orders.

There is an oo decomposition in every stated dimension, by splitting an orthonormal basis into two nonempty parts. Given two such subspaces, isometries between the corresponding nonalternating summands extend by direct sum, so $O(D,2)$ is transitive. The stabilizer of the marked subspace is $O(k,2) \times O(D-k,2)$. Orbit–stabilizer proves the ratio. Inserting the two order formulas gives the four parity cases, including $k = 1$ and $k = D-1$ with $o_1 = 1$. □

Lemma D.2 (Binary correlated projection law)**.** *Let V be a nonalternating nondegenerate D-space over $\mathbb{F}_2$, with characteristic vector c, and let U be uniform among the $\mathrm{ND}_{oo}(D,k)$ subspaces. Then $P_U(c)$ is uniform on*

$$X_{D,k} = \begin{cases} c^\perp \setminus \{0\}, & D \text{ odd}, \ k \text{ even}, \\ c + (c^\perp \setminus \{0\}), & D \text{ odd}, \ k \text{ odd}, \\ \{x : b(x,x) = 1\}, & D \text{ even}, \ k \text{ odd}, \\ c^\perp \setminus \{0, c\}, & D \text{ even}, \ k \text{ even}. \end{cases}$$

The first two supports each have size $2^{D-1} - 1$, the third has size 2^{D-1}, and the fourth has size $2^{D-1} - 2$.

Proof. Write $x = P_U(c)$. This is the characteristic vector of U, while $c - x$ is the characteristic vector of $U^\perp$. Nonalternation makes both nonzero, and Lemma D.1 gives $b(x,x) = k \bmod 2$. Hence x belongs to the stated support.

The acting group is $O(V, b)$, and every element fixes c because the characteristic vector is unique. For odd D, its restriction to $c^\perp$ is the full symplectic group. This group is transitive on nonzero vectors: extend either chosen vector to a symplectic basis and send one such basis to the other. It is also transitive on the corresponding translates by c.

For even D, any norm-one x forms with c a nondegenerate plane of Gram matrix $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right)$. The complement is symplectic, so a map sending (c, x) to (c, x') extends by an isometry of the complements. For $x \in c^\perp \setminus \{0, c\}$, the radical of $c^\perp$ is $\langle c \rangle$. Choose $y \in c^\perp$ with $b(x, y) = 1$. Then x, y form a hyperbolic pair, and their orthogonal complement is nonalternating with characteristic vector c. Do the same for x'. The map between the hyperbolic pairs extends by an isometry of the complements; that isometry necessarily sends their characteristic vectors c to c. This proves transitivity on the fourth support too.

Finally, orthogonal projection is equivariant: $P_{gU}(c) = gP_U(c)$. The oo family is a nonempty transitive orbit by Lemma D.1. The image is a nonempty invariant subset of the appropriate transitive support, and hence the whole support. For any two support points, an element carrying one to the other bijects their projection fibres. All fibres have equal size, which proves the asserted uniform distribution and the cardinalities. □

Lemma D.3 (Intersections with a centralizer witness space)**.** *If E contains c and $\dim E \le \ell$, its intersection with the four supports above has size at most, respectively,*

$$2^{\ell-1} - 1, \quad 2^{\ell-1} - 1, \quad 2^{\ell-1}, \quad 2^\ell - 2.$$

Proof. For odd D, $b(c, c) = 1$, so $b(c, \cdot)|_E$ is nonzero and its two fibres each have size $2^{\dim E-1}$. Excluding 0 or c gives the first two bounds. For even D, the norm-one fibre in E is either empty or has size $2^{\dim E-1}$, giving the third. The last support excludes two distinct vectors already in E, giving the fourth. □

Appendix E. Incidence bounds and exact constant derivations

Lemma E.1 (Centralizer witness dimensions)**.** *For nonscalar $Z \in \mathrm{M}_n(\mathbb{F}_q)$, $\dim \mathrm{Cent}(Z) \le (n-1)^2+1$. If A_0 is cyclic, $L = \langle A_0 \rangle$, and $Z \in \mathbb{F}_q[A_0]$ is nonscalar, then $E_Z = (\mathrm{Cent}(Z) \cap A_0^\perp)/L$ has dimension at most $(n-1)^2$. There are $(q^{n-1}-1)/(q-1)$ scalar-translate-and-rescale witness classes in $\mathbb{F}_q[A_0] \setminus \mathbb{F}_q I_n$.*

Proof. The dimension of a homogeneous linear system is unchanged by extension to an algebraic closure. For Jordan blocks of lengths a, b with the same eigenvalue, the intertwiner dimension is $\min(a, b)$; for distinct eigenvalues it is zero. Summing these dimensions gives $\dim \mathrm{Cent}(Z) = \sum_{\lambda,j} (\lambda'_{\lambda,j})^2$. The positive column lengths in this sum total n. Unless Z is scalar, there are at least two such positive lengths. Among partitions of n into at least two positive parts, the sum of squares is at most $(n-1)^2+1$: transferring mass from any smaller part to a largest part increases the sum until only the parts $n-1, 1$ remain. The quotient by L, which is contained in the intersection defining E_Z, subtracts one dimension. This proves the two bounds without assuming that restriction of the trace functional decreases dimension. Finally $\mathbb{F}_q[A_0]$ has dimension n, and the witness directions are exactly the lines in $\mathbb{F}_q[A_0]/\mathbb{F}_q I_n$. Replacing Z by $aZ + bI_n$ with $a \neq 0$ preserves its centralizer, proving the stated count. □

Lemma E.2 (Endpoint maximization)**.** *Let I be the number of isotropic lines in a nondegenerate D-space, and let J denote the chosen number of witnesses times q^ℓ. If $I + J < q^{D-1} - 1$, the upper bound $(R_0 + R_1 + JR_*)/(1 - IR_*)$ is largest at $k = 1, D-1$ and there equals*

$$\frac{q^{D-1} + q - 2 + (q-1)J}{q^D - 1 - (q-1)I}. \tag{E.1}$$

Proof. Set $Q = q^D$, $a = q^{D-1} - 1$, and $x = q^k + q^{D-k}$. The bound is

$$F(x) = \frac{a(x-2) + J(Q+1-x)}{a(Q-1) - I(Q+1-x)}.$$

The denominator is positive since $IR_* < I/a < 1$, and

$$F'(x) = \frac{a(Q-1)(a-I-J)}{[a(Q-1) - I(Q+1-x)]^2} > 0.$$

The sequence $q^k + q^{D-k}$ has its maximum at the two endpoints; substitution gives (E.1). □

Appendix E.1. Complete interior calculation

We give the full proof of Proposition 6.5. The formed-space formulas are proved in Appendix B, the cyclic-line subset counts in Appendix C, and the binary projection law in Appendix D. The remaining constant comparisons are certified in Lemma E.3 below. In particular, numerical regression tests are not premises of this proof.

Proof. Put $D = N - 2$ and $k = d - 1$. Thus

$$1 \le k \le D - 1, \tag{E.2}$$

including $k = 1, D-1$, which correspond to $d = 2, N-2$.

The exact conditioned space. Let $L = \langle A_0 \rangle$ be an isotropic line and put $Q_L = L^{\perp_b}/L$. Inverse image under $L^{\perp_b} \to Q_L$ is a bijection

$$\left\{ \begin{array}{c} U \le Q_L : \ \dim U = k, \\ U \cap U^{\perp_b} = 0 \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{c} C \le \mathrm{M}_n(\mathbb{F}_q) : \ \dim C = d, \\ H(C) = L \end{array} \right\}. \tag{E.3}$$

Indeed, the radical of the inverse image projects onto $U \cap U^{\perp_b}$ and already contains L. Write $\mathrm{ND}(Q_L, k)$ for the number on the left. Hence a fixed line contributes

$$\frac{\mathrm{ND}(Q_L, k)}{\begin{bmatrix} N \\ d \end{bmatrix}_q} \tag{E.4}$$

to the uniform-Grassmannian probability. Conditioned on $H(C) = L$, the quotient $U = C/L$ is therefore exactly uniform over these nondegenerate k-subspaces. Lemma 5.3 applies to this same U; it does not introduce an independently sampled projection.

We use the following elementary incidence ratios in a D-space:

$$R_0 = \frac{q^{D-k} - 1}{q^D - 1}, \qquad R_1 = \frac{q^k - 1}{q^D - 1}, \tag{E.5}$$

$$R_* = \frac{(q^k - 1)(q^{D-k} - 1)}{(q^D - 1)(q^{D-1} - 1)}. \tag{E.6}$$

They count, respectively, all k-subspaces contained in $v^{\perp_b}$, all those containing v, and all those satisfying $\langle x \rangle \le U \le (v-x)^{\perp_b}$ for a fixed nonendpoint projection value x. These are Gaussian-binomial ratios. If x and $v - x$ were nonzero and dependent, their common line would lie in $U \cap U^{\perp_b}$, so that case cannot occur for a nondegenerate U. If $I(Q_L)$ is the number of isotropic lines in Q_L, the radical-line union bound gives

$$\rho(Q_L, k) := \frac{\mathrm{ND}(Q_L, k)}{\begin{bmatrix} D \\ k \end{bmatrix}_q} \ge 1 - I(Q_L) R_*. \tag{E.7}$$

For later use, suppose that A_0 is cyclic and has the selected nonzero moment. In the odd branch put

$$\bar{v} = \pi_L\left(A_0^2 - \frac{\mathrm{Tr}(A_0^3)}{\mathrm{Tr}(A_0)} I_n\right),$$

and in characteristic two put $\bar{v} = \pi_L(I_n)$. These classes are nonzero by Lemma 5.4. Cyclicity gives $\mathrm{Cent}(A_0) = \mathbb{F}_q[A_0]$. For a nonscalar $Z \in \mathbb{F}_q[A_0]$, set

$$E_Z = (\mathrm{Cent}(Z) \cap A_0^{\perp_b})/L \le Q_L. \tag{E.8}$$

By Lemma E.1, the rational-canonical centralizer formula

$$\dim \mathrm{Cent}(Z) = \sum_f \deg(f) \sum_j (\lambda'_{f,j})^2$$

yields

$$\dim \mathrm{Cent}(Z) \le N - 2n + 2, \qquad \dim E_Z \le (n-1)^2 =: \ell. \tag{E.9}$$

The bound follows by maximizing the displayed partition sum away from the scalar partition: the maximum is $(n-1)^2 + 1$. Moreover $\bar{v} \in E_Z$, and the quotient projection identity gives

$$[Z, M_r] = 0 \quad \Longleftrightarrow \quad P_U(\bar{v}) \in E_Z, \qquad r = 2 \text{ in odd characteristic}, \quad r = 0 \text{ otherwise}. \tag{E.10}$$

The two endpoint values $0, \bar{v}$ are common to all witnesses and are counted once. Intermediate values may be unioned over at most $q^n - q$ nonscalar elements of $\mathbb{F}_q[A_0]$. Combining (E.5)–(E.7) therefore gives the explicit safe bound

$$\Pr_U[\mathrm{Cent}(A_0, M_r) \ne \mathbb{F}_q I_n \mid H(C) = L] \le \frac{R_0 + R_1 + (q^n - q) q^\ell R_*}{1 - I(Q_L) R_*}. \tag{E.11}$$

For $n = 3, 4$ we use the sharper projectivization: since $\mathrm{Cent}(aZ + bI_n) = \mathrm{Cent}(Z)$ for $a \ne 0$, the number of witness directions is only

$$P_n = \frac{q^{n-1} - 1}{q - 1}, \tag{E.12}$$

so $q^n - q$ in (E.11) can be replaced by P_n. All formulas above hold over the full range (E.2).

Odd characteristic. Lemmas B.1 and B.2 prove the orders, type convention, and all four subspace counts used below. For q odd, put

$$o_{2a+1} = 2q^{a^2} \prod_{i=1}^{a} (q^{2i} - 1),$$

$$o_{2a}^{\tau} = 2q^{a(a-1)}(q^a - \tau) \prod_{i=1}^{a-1} (q^{2i} - 1), \qquad \tau \in \{+1, -1\}.$$

If Q_L has odd dimension $2t + 1$, orbit–stabilizer gives

$$\mathrm{ND}(Q_L, 2a) = \frac{o_{2t+1}}{o_{2(t-a)+1}} \left(\frac{1}{o_{2a}^{+}} + \frac{1}{o_{2a}^{-}}\right), \tag{E.13}$$

$$\mathrm{ND}(Q_L, 2a+1) = \frac{o_{2t+1}}{o_{2a+1}} \left(\frac{1}{o_{2(t-a)}^{+}} + \frac{1}{o_{2(t-a)}^{-}}\right). \tag{E.14}$$

If Q_L has dimension $2t$ and type $\epsilon \in \{+1, -1\}$, then

$$\mathrm{ND}^{\epsilon}(Q_L, 2a) = \sum_{\tau=\pm 1} \frac{o^{\epsilon}_{2t}}{o^{\tau}_{2a} o^{\epsilon\tau}_{2(t-a)}}, \tag{E.15}$$

$$\mathrm{ND}^{\epsilon}(Q_L, 2a+1) = \frac{2o^{\epsilon}_{2t}}{o_{2a+1} o_{2(t-a-1)+1}}. \tag{E.16}$$

The numbers of isotropic lines in dimensions $2t+1$ and $2t$ are

$$S_{2t+1} = \frac{q^{2t}-1}{q-1}, \qquad S^{\epsilon}_{2t} = \frac{(q^{t-1}+\epsilon)(q^t-\epsilon)}{q-1}. \tag{E.17}$$

Thus (E.4), summed over the applicable isotropic lines, is an exact formula, not an asymptotic estimate.

When n is odd, N and D are odd and (E.13)–(E.17) give

$$\Pr[\dim H(C) = 1] = \rho(Q_L, k) \frac{(q^d-1)(q^{N-d}-1)}{(q-1)(q^N-1)} \geq \frac{1}{2q}. \tag{E.18}$$

Here (E.7) simplifies to $\rho(Q_L, k) > (q-2)/(q-1)$; substitution in the displayed product proves the last inequality for every $1 \leq k \leq D-1$. For $n = 3$, a direct cyclic Krylov count gives the sharper information needed after this conditioning. For fixed $0 \neq v \in \mathbb{F}_q^n$, the number of possible initial Krylov bases is

$$B_n = \frac{|\mathrm{GL}_n(\mathbb{F}_q)|}{q^n-1} = \prod_{i=1}^{n-1} (q^n - q^i). \tag{E.19}$$

By Lemma C.1, the fraction of isotropic lines whose generator is cyclic with nonzero trace is at least the following fixed-vector subset fraction:

$$\delta_3^{\mathrm{odd}}(q) = \frac{q^4(q-1)^3(q+1)}{q^8-1} > \frac{3}{8}. \tag{E.20}$$

For odd $n \geq 5$, zero-cone slicing and subtraction of the entire noncyclic population give

$$\Pr[\mathrm{Tr}(A_0) \neq 0 \mid \mathrm{Tr}(A_0^2) = 0, A_0 \neq 0] \geq \frac{1-q^{-1}}{1+q^{-(N-1)/2}}, \tag{E.21}$$

$$\Pr[A_0 \text{ noncyclic} \mid \mathrm{Tr}(A_0^2) = 0, A_0 \neq 0] < \frac{q}{(q^2-1)(q-1)(1-q^{-(N-1)})}. \tag{E.22}$$

Lemma B.4 proves the stated conditional use of the all-matrix noncyclic bound. The estimates in Lemma E.3 show that their difference is at least $2/5$. This is subtraction inside the isotropic-line population, not an independence argument.

When n is even, write $N = 2s$. The trace form has type $\epsilon = \chi_q((-1)^{n/2}) \in \{+1, -1\}$, where χ_q is the quadratic character; the matrix-unit determinant computation in Lemma B.1 shows that Q_L has dimension D and the same type. Equations (E.15)– (E.17) give the exact identity

$$\Pr[\dim H(C) = 1] = \frac{S^{\epsilon}_N \mathrm{ND}^{\epsilon}(Q_L, k)}{\begin{bmatrix} N \\ d \end{bmatrix}_q}. \tag{E.23}$$

For $n = 4$ one obtains, uniformly in $1 \leq k \leq 13$,

$$\rho(Q_L, k) > 1 - \frac{(q^6+1)(q^7-1)/(q-1)}{q^{13}-1} > \frac{4}{9}, \tag{E.24}$$

$$\Pr[\dim H(C) = 1] > \frac{256}{729q} > \frac{1}{3q}. \tag{E.25}$$

The same fixed-vector count, by Lemma C.1, lower-bounds the cyclic, nonzero-trace line fraction by

$$\delta_4^{\mathrm{odd}}(q) = \frac{B_4(q-1)q^2}{q^{15}-1+(q-1)q^7} > \frac{208}{547} > \frac{3}{8}. \tag{E.26}$$

For even $n \geq 6$, the radical-line estimate in the type-ϵ quotient gives $\rho(Q_L, k) > 3/8$; moreover

$$\frac{S_N^\epsilon}{q^{N-1}-1} > \frac{1}{q}, \qquad \frac{(q^d-1)(q^{N-d}-1)}{q^N-1} > (1-q^{-2})^2.$$

Substitution in (E.23) yields

$$\Pr[\dim H(C) = 1] > \frac{8}{27q} > \frac{1}{4q}. \tag{E.27}$$

For the next conditional, let $T_\epsilon = q^{N-1} - 1 + \epsilon(q-1)q^{s-1}$. Lemma B.3 gives the exact slicing identity

$$\Pr[\mathrm{Tr}(A_0) \neq 0 \mid A_0 \text{ isotropic}, A_0 \neq 0] = \begin{cases} \dfrac{(q-1)(q^{N-2}+\epsilon q^{s-1})}{T_\epsilon}, & \mathrm{char}(\mathbb{F}_q) \nmid n, \\ \dfrac{(q-1)q^{N-2}}{T_\epsilon}, & \mathrm{char}(\mathbb{F}_q) \mid n. \end{cases} \tag{E.28}$$

Both values exceed $3/5$. With $\delta = (q-1)q^{-s} + q^{1-N} < 1/16$, subtraction of the whole all-matrix exceptional population gives

$$\Pr[A_0 \text{ noncyclic} \mid A_0 \text{ isotropic}, A_0 \neq 0] < \frac{q}{(q^2-1)(q-1)(1-\delta)} < \frac{1}{5}. \tag{E.29}$$

Thus the cyclic, nonzero-trace fraction is greater than $2/5$.

It remains to insert the common-centralizer probability. In (E.11), the right-hand side is fractional-linear in $q^k + q^{D-k}$. Its derivative has the sign of $q^{D-1} - 1 - J - I(Q_L)$, where $J = (q^n - q)q^\ell$ (or $J = P_n q^\ell$ after projectivization). The sign is positive in every case above, so the maximum occurs at $k = 1, D-1$. For odd $n \geq 5$, division by $q^{D-1}-1$ gives $I(Q_L)/(q^{D-1}-1) = 1/(q-1)$ and $J/(q^{D-1}-1) < (q-2)/(q-1)$. For even $n \geq 6$, the corresponding bounds are $I(Q_L)/(q^{D-1}-1) < 5/8$ and $J/(q^{D-1}-1) < 1/8$. Thus the asserted derivative sign also holds at the endpoint values in (E.2). For $n = 3$, using (E.12), the endpoint escape is

$$\frac{q^7 - 3q^6 + q^4 - q + 2}{q^6(q-1)} > \frac{1}{19}. \tag{E.30}$$

For $n = 4$, the endpoint failure probability is at most

$$\frac{q^{13} + q^{12} - q^9 + q - 2}{q^{14} - q^{13} - q^7 + q^6} < \frac{2}{3}. \tag{E.31}$$

For odd $n \geq 5$, the endpoint form of the coarser bound in (E.11) is at most

$$\frac{1}{q-1} + \frac{q-2}{q^{D-1}(q-1)} + q^{4-n}(1-q^{1-n}); \tag{E.32}$$

it decreases in q, n, and its value at $(q, n) = (3, 5)$ is less than $6/7$. For even $n \geq 6$, put $I_\epsilon = S_D^\epsilon$ and $J = (q^n - q)q^\ell$. Retaining the quotient type gives the endpoint failure bound

$$\frac{q^{D-1} + q - 2 + (q-1)J}{q^D - 1 - (q-1)I_\epsilon} \leq \frac{153}{242} < \frac{2}{3}. \tag{E.33}$$

Let p_H denote the hull probability, p_{line} the conditional cyclic nonzero-trace line fraction, and p_{esc} the subsequent scalar-centralizer escape probability. We have proved the sequential table

n	p_H	p_{line}	p_{esc}	$\Pr[\mathsf{Good}_{\mathrm{int}}]$
3	$\geq 1/(2q)$	$> 3/8$	$> 1/19$	$> 3/(304q)$
4	$> 1/(3q)$	$> 3/8$	$> 1/3$	$> 1/(24q)$
odd $n \geq 5$	$\geq 1/(2q)$	$\geq 2/5$	$\geq 1/7$	$\geq 1/(35q)$
even $n \geq 6$	$> 1/(4q)$	$> 2/5$	$> 1/3$	$> 1/(30q)$.

(E.34)

Characteristic two with $q \geq 4$. Here the isotropic vectors are precisely the trace-zero matrices, so their projective lines number

$$S = \frac{q^{N-1} - 1}{q - 1}. \tag{E.35}$$

Every cyclic line with $\mathrm{Tr}(A_0^3) \neq 0$ is nonscalar. Since $\mathrm{Tr}(A_0) = 0$, we have $I_n \in L^{\perp_b}$, and the class $c = \pi_L(I_n)$ is nonzero because $I_n \notin L$. The quotient Q_L is nondegenerate, so $x \mapsto b_{Q_L}(c, x)$ is a nonzero linear functional. The trace identity gives $b_{Q_L}(x, x) = b_{Q_L}(c, x)^2$. Hence Q_L is nonalternating, and its isotropic vectors form exactly the hyperplane $\ker b_{Q_L}(c, \cdot)$ of dimension $D - 1$. Consequently

$$I(Q_L) = \frac{q^{D-1} - 1}{q - 1}, \qquad \rho(Q_L, k) \geq 1 - I(Q_L)R_* > \frac{q-2}{q-1} \geq \frac{2}{3}. \tag{E.36}$$

The line-incidence factor obtained from (E.4) is

$$\frac{S\begin{bmatrix} D \\ k \end{bmatrix}_q}{\begin{bmatrix} N \\ d \end{bmatrix}_q} = \frac{(q^d - 1)(q^{N-d} - 1)}{(q-1)(q^N - 1)} > \frac{3}{4q}, \tag{E.37}$$

where $d, N - d \geq 2$; hence the endpoint values of d are included.

For $n = 3, 4$, Lemma C.1 lower-bounds the cyclic, nonzero-cubic-trace line fraction by the fixed-vector subset ratio

$$\delta_n^{\mathrm{even}}(q) = \frac{B_n(q-1)q^{n-2}}{q^{N-1} - 1} > (1 - q^{-1})^2 \prod_{j=2}^{n-1} (1 - q^{-j}) > \frac{1}{2}. \tag{E.38}$$

For $n \geq 5$, fix all trace-zero matrix entries except a_{12}, a_{23}, a_{31}. The polynomial $\mathrm{Tr}(A_0^3)$ is multilinear in these variables with nonzero cubic coefficient. Thus

$$\Pr[\mathrm{Tr}(A_0^3) \neq 0 \mid \mathrm{Tr}(A_0) = 0, A_0 \neq 0] \geq \frac{(1 - q^{-1})^3}{1 - q^{-(N-1)}} > \frac{27}{64}, \tag{E.39}$$

$$\Pr[A_0 \text{ noncyclic} \mid \mathrm{Tr}(A_0) = 0, A_0 \neq 0] < \frac{q}{(q^2 - 1)(q - 1)(1 - q^{1-N})} < \frac{1}{8}. \tag{E.40}$$

Their difference is greater than $1/4$. Multiplying the appropriate good-line fraction, (E.36), and (E.37), directly over the good lines, gives

$$\Pr[\dim H(C) = 1, A_0 \text{ cyclic}, \mathrm{Tr}(A_0^3) \neq 0] > \begin{cases} 1/(4q), & n = 3, 4, \\ 1/(8q), & n \geq 5. \end{cases} \tag{E.41}$$

This direct sum over good lines avoids assigning a fictitious uniform marginal to all characteristic-two hull lines.

For $n = 3$, the endpoint expression in (E.30) is greater than $1/3$ for $q \geq 4$. For $n = 4$, projectivization and the actual nonalternating denominator give the failure bound

$$\frac{q^{13} + q^{12} - q^9 + q - 2}{q^{13}(q - 1)} < \frac{1}{2}. \tag{E.42}$$

For $n \geq 5$, substitution of (E.36) in (E.11), followed by the same endpoint maximization, is bounded by (E.32) and hence gives escape probability greater than $1/7$; here $I(Q_L)/(q^{D-1}-1) = 1/(q-1)$ and $J/(q^{D-1}-1) < (q-2)/(q-1)$ verify the derivative sign. Hence

n	Pr[cyclic one– hull]	Pr[scalar cent. | previous]	Pr[$\mathsf{Good}_{\mathrm{int}}$]
3	$> 1/(4q)$	$> 1/3$	$> 1/(12q)$
4	$> 1/(4q)$	$> 1/2$	$> 1/(8q)$
$n \geq 5$	$> 1/(8q)$	$> 1/7$	$> 1/(56q)$.

(E.43)

The binary field. By Lemma C.1, over $\mathbb{F}_2$ a fixed nonzero vector is cyclic for exactly

$$G_n = 2^{N-2} \prod_{j=1}^{n-1} (1 - 2^{-j}) \tag{E.44}$$

matrices satisfying $\mathrm{Tr}(A_0) = 0$ and $\mathrm{Tr}(A_0^3) = 1$. Since there are $2^{N-1} - 1$ nonzero trace-zero matrices, the fraction of isotropic lines satisfying the two intrinsic predicates is

$$\delta_{\mathrm{cyc}} \geq \frac{G_n}{2^{N-1}-1} > \frac{1}{2} \prod_{j=1}^{\infty} (1 - 2^{-j}) > \frac{9}{64}. \tag{E.45}$$

The last comparison follows from

$$\prod_{j=1}^{\infty} (1 - 2^{-j}) = \frac{1}{2}\frac{3}{4} \prod_{j=3}^{\infty} (1 - 2^{-j}) > \frac{3}{8} \left(1 - \sum_{j=3}^{\infty} 2^{-j} \right) = \frac{9}{32}.$$

For $n = 3, 4$ the first ratio in (E.45) is, respectively, $16/85$ and $5376/32767$. Every such line again has a nonalternating quotient, with characteristic vector $c = \pi_L(I_n)$. Let $\mathrm{ND}_{oo}(D, k)$ count the nondegenerate k-subspaces U for which both U and $U^{\perp_b}$ are nonalternating. For reference, put

$$o_{2a} = 2^{a^2} \prod_{i=1}^{a-1} (2^{2i} - 1), \qquad o_{2a+1} = 2^{a^2} \prod_{i=1}^{a} (2^{2i} - 1).$$

Lemma D.1 specifies the oo orbit and its product stabilizer. Writing D, k in the four possible parity forms gives

$$\mathrm{ND}_{oo}(D,k) = \begin{cases} \dfrac{o_{2a}}{o_{2b+1} o_{2(a-b-1)+1}}, & D = 2a,\ k = 2b+1, \\ \dfrac{o_{2a+1}}{o_{2b+1} o_{2(a-b)}}, & D = 2a+1,\ k = 2b+1, \\ \dfrac{o_{2a+1}}{o_{2b} o_{2(a-b)+1}}, & D = 2a+1,\ k = 2b, \\ \dfrac{o_{2a}}{o_{2b} o_{2(a-b)}}, & D = 2a,\ k = 2b. \end{cases} \tag{E.46}$$

Direct simplification of (E.46) gives, throughout (E.2),

range	$\inf_{1 \leq k \leq D-1} \dfrac{\mathrm{ND}_{oo}(D,k)}{\begin{bmatrix} D \\ k \end{bmatrix}_2}$
$n = 3,\ D = 7$	$48/127 > 3/8$
$n = 4,\ D = 14$	$16773120/44731051 > 1/3$
$n \geq 5$	$> 297/2048 > 1/8$.

(E.47)

The first two minima occur by evaluating (E.46) at $k = 1, \ldots, D-1$; complementation makes the lists symmetric. The last inequality follows uniformly by writing the same ratios as products

of factors $1-4^{-i}$. More explicitly, if $\gamma_Q(t)=\prod_{i=1}^{t}(1-Q^{-i})$, the four parity cases in (E.46), after division by $2^{k(D-k)}$, are

$$\begin{cases} \dfrac{\gamma_4(D/2-1)}{\gamma_4((k-1)/2)\gamma_4((D-k-1)/2)}, & D \text{ even, } k \text{ odd,} \\ \dfrac{(1-2^{-(D-k)})\gamma_4((D-1)/2)}{\gamma_4((k-1)/2)\gamma_4((D-k)/2)}, & D \text{ odd, } k \text{ odd,} \\ \dfrac{(1-2^{-k})\gamma_4((D-1)/2)}{\gamma_4(k/2)\gamma_4((D-k-1)/2)}, & D \text{ odd, } k \text{ even,} \\ \dfrac{\gamma_4(D/2-1)}{\gamma_4(k/2-1)\gamma_4((D-k)/2-1)}, & D \text{ even, } k \text{ even.} \end{cases} \tag{E.48}$$

Since

$$\gamma_4(\infty) > \frac{3}{4}\left(1-\sum_{i=2}^{\infty}4^{-i}\right)=\frac{11}{16},$$

these expressions exceed $33/64$. Furthermore

$$\frac{\begin{bmatrix}D\\k\end{bmatrix}_2}{2^{k(D-k)}} < \frac{1}{\gamma_2(\infty)}, \qquad \gamma_2(\infty) > \frac{9}{32},$$

which proves the last line of (E.47).

Lemma D.2 proves that for a uniform oo subspace the correlated projection $P_U(c)$ is uniform on

$$\mathcal{X}_{D,k} = \begin{cases} c^{\perp_b}\setminus\{0\}, & D \text{ odd, } k \text{ even,} \\ c+(c^{\perp_b}\setminus\{0\}), & D \text{ odd, } k \text{ odd,} \\ \{x : b(x,x)=1\}, & D \text{ even, } k \text{ odd,} \\ c^{\perp_b}\setminus\{0,c\}, & D \text{ even, } k \text{ even.} \end{cases} \tag{E.49}$$

Indeed, the orthogonal group fixing c is transitive both on the *oo* orbit and on each displayed support, while $U \mapsto P_U(c)$ is equivariant; its fibres therefore have equal size. This is the actual conditional projection law. The support sizes are $2^{D-1}-1$, 2^{D-1}, and $2^{D-1}-2$ in the corresponding cases. If $E \ni c$ and $\dim E \le \ell$, its intersection with this support has size at most

$$2^{\ell-1}-1, \quad 2^{\ell-1}, \quad 2^{\ell}-2, \tag{E.50}$$

respectively. There are exactly $2^{n-1}-1$ projective centralizer witness directions. Applying the union bound to (E.50) gives

$$n=3: \quad \Pr[\text{nonscalar cent.} \mid oo] \le \frac{3(2^{4-1}-1)}{63}=\frac{1}{3}, \tag{E.51}$$

$$n=4: \quad \Pr[\text{nonscalar cent.} \mid oo] \le \begin{cases} 7/32, & k \text{ odd,} \\ 3570/8190, & k \text{ even,} \end{cases} < \frac{1}{2}. \tag{E.52}$$

For $n \ge 5$, with $\ell=(n-1)^2$, the three corresponding bounds are

$$\frac{(2^{n-1}-1)(2^{\ell-1}-1)}{2^{D-1}-1}, \quad \frac{(2^{n-1}-1)2^{\ell-1}}{2^{D-1}}, \quad \frac{(2^{n-1}-1)(2^{\ell}-2)}{2^{D-1}-2}. \tag{E.53}$$

In their respective parity ranges these bounds are less than $1/8$: the boundary checks are $8\cdot 15(2^{15}-1)<2^{22}-1$ for odd $n=5$ and $8\cdot 31(2^{25}-2)<2^{33}-2$ for even $n=6$; increasing n only decreases the ratios. Thus the centralizer escape conditioned on *oo* is greater than $7/8$, and (E.47) gives the combined estimate

$$\frac{\#\{U:\ U \text{ is } oo \text{ and has scalar-centralizer escape}\}}{\begin{bmatrix}D\\k\end{bmatrix}_2} > \frac{1}{8}\frac{7}{8}=\frac{7}{64}. \tag{E.54}$$

The strict inequality uses the strict *oo*-density bound.

Finally, summing the fixed-line fibres over all $2^{N-1}-1$ isotropic lines gives the exact basic incidence factor

$$\frac{(2^{N-1}-1)\begin{bmatrix}D\\k\end{bmatrix}_2}{\begin{bmatrix}N\\d\end{bmatrix}_2} = \frac{(2^d-1)(2^{N-d}-1)}{2^N-1} > (1-2^{-d})(1-2^{-(N-d)}) \geq \frac{9}{16}. \tag{E.55}$$

For $n=3,4$ its minimum at $d=2,N-2$ is, respectively, $381/511$ and $49149/65535$, both greater than $2/3$. Combining the factors in their conditional order gives

$$n=3: \quad \Pr[\mathsf{Good}_{\mathrm{int}}(C)] > \underbrace{\frac{9}{64}}_{\text{good hull line}} \underbrace{\frac{3}{8}}_{oo\text{ quotient}} \underbrace{\frac{2}{3}}_{\text{basic incidence}} \underbrace{\frac{2}{3}}_{\text{centralizer escape}} = \frac{3}{128}, \tag{E.56}$$

$$n=4: \quad \Pr[\mathsf{Good}_{\mathrm{int}}(C)] > \underbrace{\frac{9}{64}}_{\text{good hull line}} \underbrace{\frac{1}{3}}_{oo\text{ quotient}} \underbrace{\frac{2}{3}}_{\text{basic incidence}} \underbrace{\frac{1}{2}}_{\text{centralizer escape}} = \frac{1}{64}, \tag{E.57}$$

$$n\geq 5: \quad \Pr[\mathsf{Good}_{\mathrm{int}}(C)] > \underbrace{\frac{9}{64}}_{\text{good hull line}} \underbrace{\frac{7}{64}}_{oo\text{ quotient and escape}} \underbrace{\frac{9}{16}}_{\text{basic incidence}} = \frac{567}{65536} > \frac{1}{128}. \tag{E.58}$$

In particular, the factor $7/64$ in the last line already combines the nondegenerate formed-space density with centralizer escape; the factor $9/16$ is the separate Gaussian incidence factor.

The weakest row above is the odd-characteristic $n=3$ row, so all cases give

$$\Pr[\mathsf{Good}_{\mathrm{int}}(C)] > \frac{3}{304q}. \tag{E.59}$$

Conditional on $\operatorname{rank}(P_A)=m$, Proposition 3.4 makes C_A exactly uniform on the m-Grassmannian. The chain rule and Lemma 6.3 therefore give

$$\Pr[A\in\mathcal{G}^{\star}_{q,n,m}] = p_{\mathrm{col}}(N,m,q)\Pr_C[\mathsf{Good}_{\mathrm{int}}(C)] > \frac{3}{4}\frac{3}{304q} = \frac{9}{1216q}.$$

For the upper bound, $\mathsf{Good}_{\mathrm{int}}(C)$ implies $\dim H(C)=1$. Lemma 6.4, with the rank probability bounded by one, gives the other side of (78). □

Appendix E.2. Constant comparisons

Lemma E.3 (Exact constants used above)**.** *All constant inequalities and derivative signs used in the preceding calculation hold throughout their displayed ranges. In particular the weakest odd row gives* $3/(304q)$*, and the binary rows are strictly larger than* $3/(304\cdot 2)$*.*

Proof. We record algebraic certificates and uniform estimates to make the numerical margins checkable. For a rational expression $R(q)>0$, cancel common factors, take its numerator with positive denominator on the stated range, and substitute $q=x+q_0$. A polynomial with nonnegative coefficients and positive constant is positive for all $x\geq 0$. Here the following rows use, respectively,

$$\frac{q(q-2)}{(q-1)^2}(1-q^{-2})^2-\frac{1}{2}, \quad \delta_3^{\mathrm{odd}}-\frac{3}{8}, \quad \delta_4^{\mathrm{odd}}-\frac{208}{547}, \quad \frac{5}{9}-\frac{(q^6+1)(q^7-1)}{(q-1)(q^{13}-1)},$$
$$\frac{q^7-3q^6+q^4-q+2}{q^6(q-1)}-\frac{1}{19}, \quad \frac{2}{3}-\frac{q^{13}+q^{12}-q^9+q-2}{q^{14}-q^{13}-q^7+q^6},$$
$$\frac{q^7-3q^6+q^4-q+2}{q^6(q-1)}-\frac{1}{3}, \quad \frac{1}{2}-\frac{q^{13}+q^{12}-q^9+q-2}{q^{13}(q-1)}.$$

The table reports the shifted numerator degree, its least coefficient, and its constant coefficient. Every numerator is specified by the corresponding rational expression above, so these certificates can be verified by polynomial expansion without sampling field orders.

	q_0	deg	min coefficient	constant
odd hull constant	3	3	1	5
n=3 odd Krylov	3	6	5	132
n=4 odd Krylov	3	13	339	56888
n=4 formed density	3	13	5	790600
n=3 odd escape	3	6	18	31
n=4 odd escape	3	13	2	28065
n=3 even escape	4	6	2	254
n=4 even escape	4	13	1	11359572

For example, the shifted numerator in the $n = 3$ odd escape row is

$$18x^6 + 286x^5 + 1822x^4 + 5825x^3 + 9368x^2 + 6104x + 31.$$

This is the tightest centralizer margin in the proof. Also $208/547 > 3/8$ since $1664 > 1641$.

For the hull bounds, $R_* < 1/(q^{D-1} - 1)$. In odd dimension, $I/(q^{D-1} - 1) = 1/(q - 1)$, so $\rho > (q - 2)/(q - 1)$; the first polynomial certificate gives $p_H > 1/(2q)$. In the $n = 4$ plus-type quotient the fourth certificate gives $\rho > 4/9$. Moreover $S_{16}^{+}/(q^{15} - 1) > 1/(q - 1) > 1/q$ and the remaining incidence factor exceeds $(1 - q^{-2})^2 \geq 64/81$, yielding $256/(729q)$. For even $n \geq 6$, $D \geq 34$ and

$$\frac{I}{q^{D-1} - 1} \leq \frac{1}{q - 1} \frac{1 + q^{1-D/2}}{1 - q^{-(D-1)}} < \frac{5}{8}.$$

Thus $\rho > 3/8$; the same incidence estimate gives $8/(27q)$. The inequality $S_N^{\epsilon}/(q^{N-1} - 1) > 1/q$ follows directly on multiplying by the positive denominators, with the minus type giving the smaller numerator.

For odd $n \geq 5$, the trace fraction in Lemma B.3 exceeds $32/51$ because $q^{-(N-1)/2} < 1/16$. The conditional noncyclic fraction is less than $(3/16)/(1 - 1/16) = 1/5$. Their difference is greater than $2/5$. For even $n \geq 6$, both $(q - 1)q^{-N/2}$ and $\delta = (q - 1)q^{-N/2} + q^{1-N}$ are less than $1/64$. The trace fraction is therefore greater than $(2/3 - 1/64)/(1 + 1/64) = 25/39 > 3/5$, and the noncyclic fraction is less than $(3/16)/(1 - 1/64) = 4/21 < 1/5$. These estimates prove the line factors without assuming independence.

To check the derivative sign in Lemma E.2, write $a = q^{D-1} - 1$. For $n = 3$, $I/a \leq 1/2$ and $J/a = (q + 1)q^4/(q^6 - 1) < 1/2$. For $n = 4$, $I/a < 5/9$ and $J/a = (q^2 + q + 1)q^9/(q^{13} - 1) < 1/3$. For odd $n \geq 5$, or nonalternating quotients with $n \geq 5, q \geq 4$, $I/a = 1/(q - 1)$ and

$$J/a < \frac{q^{4-n}}{1 - q^{-(D-1)}} < \frac{q - 2}{q - 1}.$$

For even $n \geq 6$ in odd characteristic, the preceding bound gives $I/a < 5/8$ and $J/a < 1/8$. Hence $I + J < a$ in every case.

For the large odd or nonalternating quotient, the endpoint expression is

$$\frac{1}{q - 1} + \frac{q - 2}{(q - 1)q^{D-1}} + q^{4-n}(1 - q^{1-n}).$$

Each term decreases as $q \geq 3$ or $n \geq 5$ increases. For the second term the logarithmic derivative in q is $1/((q - 2)(q - 1)) - (D - 1)/q < 0$. For the third, differentiating gives $q^{3-n}[(4 - n) + (2n - 5)q^{1-n}] < 0$; increasing n decreases it by comparing adjacent exponents. At $(q, n) = (3, 5)$ the sum is $1/2 + 1/(2 \cdot 3^{22}) + 80/243 < 6/7$. In an even-dimensional odd-characteristic quotient

the worst type is plus and its endpoint denominator is $q^D(1-q^{-1})(1-q^{-D/2})$. Consequently the failure bound is less than

$$\frac{1/(q-1)+(q-2)q^{1-D}/(q-1)+q^{4-n}}{1-q^{-D/2}} < \frac{1/2+1/81+1/9}{1-1/81} = \frac{101}{160} < \frac{153}{242} < \frac{2}{3}.$$

This justifies the retained constant in (E.33).

For $q \geq 4$ even, the line-incidence factor is greater than $(1-q^{-2})^2/(q-1) > 3/(4q)$. For $n = 3, 4$, the Krylov lower bound exceeds $(3/4)^2(15/16)(63/64) = 8505/16384 > 1/2$. For larger n, fixing the other trace-zero entries leaves a multilinear polynomial in a_{12}, a_{23}, a_{31} with nonzero top coefficient. A polynomial with a nonzero full multilinear monomial in t variables is nonzero on at least $(q-1)^t$ tuples: apply induction to the coefficient of its last variable, then exclude at most one root in that variable. This proves the cubic support estimate. Its difference from the noncyclic bound is greater than $27/64 - 1/8 > 1/4$.

For binary spaces the four product identities (E.48) follow by cancellation in $o_D/(o_k o_{D-k})$. The finite exceptional ratios for $1 \leq k \leq \lfloor D/2 \rfloor$ are

D	$\mathrm{ND}_{oo}(D,k)/\begin{bmatrix} D \\ k \end{bmatrix}_2$
7	63/127, 48/127, 1680/3937
14	8192/16383, 16773120/44731051, 58720256/134193153, 37544263680/91564461397, 116500987904/274693384191, 2785101742080/6684205681981, 66571993088/157894622409.

Complementation gives the other values, proving the two stated minima. For all remaining dimensions the already displayed estimates $\gamma_4(\infty) > 11/16$ and $\gamma_2(\infty) > 9/32$ give the uniform bound $297/2048$. For odd $n \geq 5$ the first centralizer ratio in (E.53) is strictly below $2^{2-n} \leq 1/8$; the factors removed from its numerator dominate the factor removed from its denominator. For even $n \geq 6$, the second ratio is below 2^{2-n} and the third below $2^{3-n} \leq 1/8$ by the same comparison. Only these respective parity ranges are used. The basic incidence has its minimum at $d = 2, N-2$, since $2^d + 2^{N-d}$ is largest there. This proves the binary margins and completes all comparisons. □

Appendix F. A two-sided source-scale certificate

This appendix thins the recognized family of Definition 7.1; none of its additional scalar equalities is required by the recovery algorithm or imposed on the target.

Appendix F.1. A six-row source event

For the hyperplane row, let $\mathcal{E}_{\mathrm{hyp}}(A)$ denote the joint condition

$$\mathrm{rank}(P_A) = N-1 \quad \text{and} \quad C_A^{\perp_b} = \langle Z \rangle \text{ with } \mathrm{Tr}(Z) \neq 0. \tag{F.1}$$

The trace condition is independent of the chosen nonzero generator Z. It implies $I_n \notin C_A$. Consequently, on $\mathcal{E}_{\mathrm{hyp}}(A)$ the determinant

$$\Delta_A = \det[\mathrm{vec}(A_1) \ \cdots \ \mathrm{vec}(A_{N-1}) \ \mathrm{vec}(I_n)] \tag{F.2}$$

is nonzero.

Definition F.1 (The certified event)**.** For every prime power $q \geq 2$ and every $n, m \geq 1$, define $\mathcal{G}^{\star}_{q,n,m} \subseteq \mathrm{M}_n(\mathbb{F}_q)^m$ by the following six mutually exclusive parameter rows.

1. If $n = 1$, then $A \in \mathcal{G}^{\star}_{q,1,m}$ if and only if $A_1 = 0$. The remaining $m-1$ entries are unrestricted.
2. If $n \geq 2$ and $m = 1$, then $A \in \mathcal{G}^{\star}_{q,n,1}$ if and only if $\mathrm{Tr}(A_1) = 1$.
3. If $(n, m) = (2, 2)$, then $A \in \mathcal{G}^{\star}_{q,2,2}$ if and only if $\mathrm{rank}(P_A) = 2$ and $\mathsf{Good}_{22}(C_A)$ holds.

4. If $n \geq 3$ and $2 \leq m \leq N-2$, then $A \in \mathcal{G}^{\star}_{q,n,m}$ if and only if $\operatorname{rank}(P_A) = m$ and $\mathsf{Good}_{\mathrm{int}}(C_A)$ holds.
5. If $n \geq 2$ and $m = N-1$, then $A \in \mathcal{G}^{\star}_{q,n,N-1}$ if and only if $\mathcal{E}_{\mathrm{hyp}}(A)$ holds and $\Delta_A = 1$.
6. If $n \geq 2$ and $m \geq N$, then $A \in \mathcal{G}^{\star}_{q,n,m}$ if and only if

$$\det\big[\operatorname{vec}(A_1) \ \cdots \ \operatorname{vec}(A_N)\big] = 1. \tag{F.3}$$

The remaining $m - N$ slices are unrestricted.

The six rows cover every positive pair (n, m). The prescribed scalar equalities in the first, second, fifth, and sixth rows are deliberate source-only thinnings. They are not orbit invariants and will never be imposed on an arbitrary target.

Theorem F.2 (Certified two-sided source probability)**.** *For every prime power $q \geq 2$ and every $n, m \geq 1$, the event in Definition F.1 is deterministically recognizable in* $\mathrm{poly}(n, m, \log q)$ *bit time and satisfies*

$$\frac{1}{136q} \leq \Pr[A \in \mathcal{G}^{\star}_{q,n,m}] \leq \frac{4}{q}. \tag{F.4}$$

On this event, the selected MCC algorithm satisfies the arbitrary-target search, verification, and exact-decision requirements of Definition 2.4.

Proof. We calculate the iid probability in each of the six disjoint rows.

If $n = 1$, the single constraint $A_1 = 0$ has probability $1/q$, independently of the unrestricted later coordinates. If $n \geq 2$ and $m = 1$, the matrix trace is a surjective linear functional because $\operatorname{Tr}(E_{11}) = 1$; hence $\Pr[\operatorname{Tr}(A_1) = 1] = 1/q$.

For $(n, m) = (2, 2)$, the full-column-rank probability is $(1 - q^{-4})(1 - q^{-3})$. Conditional image uniformity and Lemma 6.6 give the exact value

$$\Pr[A \in \mathcal{G}^{\star}_{q,2,2}] = (1 - q^{-4})(1 - q^{-3})g_2(q). \tag{F.5}$$

The rank product is greater than $3/4$, so this probability is greater than $1/(4q)$; it is less than $1/q$. At $q = 2$, its exact value is $9/64$.

The interior row is covered by Proposition 6.5. Its strict lower coefficient satisfies

$$\frac{9}{1216} > \frac{1}{136}, \tag{F.6}$$

because $9 \cdot 136 = 1224 > 1216$.

Consider next $n \geq 2$ and $m = N - 1$. Conditional on full column rank, C_A is a uniform hyperplane, and orthogonal complementation makes $C_A^{\perp_b}$ a uniform line. The fraction of lines generated by a matrix of nonzero trace is

$$\frac{1 - q^{-1}}{1 - q^{-N}}.$$

Multiplication by the full-rank probability cancels the $j = N$ factor and gives the exact joint probability

$$p_{\mathrm{hyp}}(N, q) = \Pr[\mathcal{E}_{\mathrm{hyp}}(A)] = \prod_{j=1}^{N-1} (1 - q^{-j}) > \frac{1}{4}. \tag{F.7}$$

For a fixed good hyperplane, choose one ordered basis matrix B_C. Every other ordered basis is uniquely $B_C U$ with $U \in \mathrm{GL}_{N-1}(\mathbb{F}_q)$, and

$$\det[B_C U, \operatorname{vec}(I_n)] = \det(U) \det[B_C, \operatorname{vec}(I_n)].$$

The determinant map on $\mathrm{GL}_{N-1}(\mathbb{F}_q)$ is onto $\mathbb{F}_q^\times$ with equal fibres. Exactly $1/(q-1)$ of these presentations have $\Delta_A = 1$. Therefore

$$\Pr[A \in \mathcal{G}^\star_{q,n,N-1}] = \frac{p_{\mathrm{hyp}}(N,q)}{q-1}, \qquad \frac{1}{4q} < \Pr[A \in \mathcal{G}^\star_{q,n,N-1}] \le \frac{2}{q}. \tag{F.8}$$

Finally, let $n \ge 2$ and $m \ge N$. The first N vectorized slices form a uniform square matrix, and the remaining slices are free. Equal nonzero determinant fibres give

$$\Pr[A \in \mathcal{G}^\star_{q,n,m}] = \frac{|\mathrm{SL}_N(\mathbb{F}_q)|}{q^{N^2}} = \frac{1}{q-1}\prod_{j=1}^{N}(1-q^{-j}). \tag{F.9}$$

Lemma 6.3 again yields

$$\frac{1}{4q} < \Pr[A \in \mathcal{G}^\star_{q,n,m}] \le \frac{2}{q}. \tag{F.10}$$

This probability is independent of $m - N$.

These calculations prove (F.4). The predicates are deterministic rank, trace, determinant, hull and centralizer tests. Each row thins the corresponding row of Definition 7.1; hence completeness follows from Proposition 7.8. The additional source equalities are never tested on the target, and coefficient lifting still uses the declared m. □

The constants in (F.4) are absolute, so $\Pr[\mathcal{G}^\star_{q,n,m}] = \Theta(1/q)$ uniformly in the stated parameters. This is the two-sided source certificate corresponding to the $1/\Theta(q)$ formulation in LLQTW Problem 1.5 (Li et al., 2026). The larger recognized event $\mathcal{G}_{q,n,m}$ of Definition 7.1 contains $\mathcal{G}^\star_{q,n,m}$ and has probability strictly greater than $9/(1216q)$, but no $O(1/q)$ upper bound is asserted for that larger event.

Appendix G. The specified odd-characteristic transporter cover

Appendix G.1. Generic projective-transporter cover

We now isolate the exact geometric statement available in odd characteristic. Let k_0 be a field of characteristic different from two, and base-change to an algebraic closure k. A rational section over the smaller field would remain a section after this base change. Put

$$G = \mathrm{PGL}_2(k), \qquad h(\mathcal{C}) = \mathrm{Tr}(A_{\mathcal{C}}),$$

and consider the geometric open locus

$$X = \{\mathcal{C} \in \mathrm{Gr}(2,W) : H(\mathcal{C}) = 0,\ h(\mathcal{C})(h(\mathcal{C})-2) \neq 0\}. \tag{G.1}$$

The projection is regular on the zero-hull chart because it is obtained by inverting the trace Gram matrix. Thus (G.1) is open in the geometrically integral fourfold $\mathrm{Gr}(2,4)$.

Proposition G.1 (Geometric orbits and stabilizers)**.** *The locus X is nonempty and dense. For every $\mathcal{C} \in X$, put*

$$e = A_{\mathcal{C}}, \qquad e_0 = e - \frac{h(\mathcal{C})}{2} I_2, \tag{G.2}$$

and choose a nonzero $u \in \mathcal{C} \cap \ker(\mathrm{Tr})$. Then

$$e_0^2 = \frac{h(\mathcal{C})(2-h(\mathcal{C}))}{4} I_2, \qquad u^2 = \frac{\mathrm{Tr}(u^2)}{2} I_2, \qquad e_0 u = -u e_0, \tag{G.3}$$

and both displayed squares are nonzero. The projective stabilizer scheme is reduced and étale of length two, with

$$N_G(\mathcal{C}) = \{1, [e_0]\} \cong C_2. \tag{G.4}$$

Moreover, h is a complete geometric orbit parameter on X: two points of X are G-conjugate if and only if they have the same value of h.

Proof. The projection identity gives $\mathrm{Tr}(e^2) = b(e,e) = b(I_2,e) = h(\mathcal{C})$. Since $h(\mathcal{C}) \neq 0$, the trace-zero line in $\mathcal{C}$ is spanned by u, and the nondegeneracy of $b|_{\mathcal{C}}$ gives $\mathrm{Tr}(u^2) \neq 0$. The identity $\mathrm{Tr}(e^2) = \mathrm{Tr}(e)^2 - 2\det(e)$ gives $\det(e) = (h(\mathcal{C})^2 - h(\mathcal{C}))/2$. Substitution into Cayley–Hamilton yields the displayed formula for e_0^2. Since u is traceless, Cayley–Hamilton similarly gives $u^2 = \mathrm{Tr}(u^2)I_2/2$. Finally, polarized Cayley–Hamilton gives

$$e_0u + ue_0 = \mathrm{Tr}(e_0u)I_2 = 0,$$

because $b(e,u) = \mathrm{Tr}(u) = 0$. This proves (G.3). The condition $h(\mathcal{C}) \neq 2$ makes e_0 invertible, while $\mathrm{Tr}(u^2) \neq 0$ makes u invertible.

Any ambient normalizer of $\mathcal{C}$ fixes the intrinsic anchor e and preserves the intrinsic line $\mathcal{C} \cap \ker(\mathrm{Tr})$. A representative commuting with e_0 has the form $S = \alpha I_2 + \beta e_0$, and $SuS^{-1} = \lambda u$ is equivalent, using $e_0u = -ue_0$, to

$$\alpha(1-\lambda) = 0, \qquad \beta(1+\lambda) = 0. \tag{G.5}$$

Invertibility and odd characteristic leave exactly the scalar component and the scalar-e_0 component, proving the set-theoretic identity (G.4). For reducedness, an element of the Lie stabilizer in $\mathfrak{pgl}_2$ first commutes with e_0, so it is γe_0 modulo scalars. Preservation of the line ku then requires $2\gamma e_0u \in ku$. The vectors u and e_0u are independent, since both u and e_0 are invertible and e_0 is nonscalar. Hence $\gamma = 0$. The stabilizer is therefore reduced étale of length two.

It remains to identify the orbits. For

$$t(1+t^2) \neq 0, \qquad E_t = \begin{pmatrix} \dfrac{1+t}{1+t^2} & 0 \\ 0 & \dfrac{1-t}{1+t^2} \end{pmatrix}, \qquad U_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

put $\mathcal{C}_t = \langle E_t, U_1 \rangle$. Direct calculation gives

$$\mathrm{Gram}_b(E_t, U_1) = \mathrm{diag}\left(\frac{2}{1+t^2}, 2\right),$$

so this plane has zero hull, and

$$P_{\mathcal{C}_t}(I_2) = E_t, \qquad h(\mathcal{C}_t) = \frac{2}{1+t^2}, \qquad E_t - \frac{h(\mathcal{C}_t)}{2}I_2 = \frac{t}{1+t^2}\,\mathrm{diag}(1,-1). \tag{G.6}$$

These examples show that X is nonempty; since it is a nonempty open in an integral Grassmannian, it is dense.

For an arbitrary $\mathcal{C} \in X$, choose $t \in k$ with $t^2 = (2 - h(\mathcal{C}))/h(\mathcal{C})$. The two traceless anchors in (G.3) and (G.6) have the same nonzero square, so they are conjugate. In an eigenbasis for this anchor, the anticommutation identity makes u off diagonal with both entries nonzero. A diagonal conjugation, together with rescaling the generator of its line, sends ku to kU_1. Thus $\mathcal{C}$ is conjugate to $\mathcal{C}_t$. This proves that equal h-values give conjugate planes; the converse follows from conjugacy invariance of h. □

Let $M = \mathbb{A}^1_k \setminus \{0,2\}$ and regard h as a morphism $X \to M$. Define the reduced equal-h incidence

$$\mathcal{I} = (X \times_M X)_{\mathrm{red}} = \{(\mathcal{C},\mathcal{D}) \in X \times X : h(\mathcal{C}) = h(\mathcal{D})\}_{\mathrm{red}}. \tag{G.7}$$

Proposition G.1 says that its geometric points are exactly the geometrically conjugate source-target pairs.

Theorem G.2 (Generic odd-characteristic transporter boundary)**.** *The projective-transporter morphism*

$$\mu : G \times X \longrightarrow \mathcal{I}, \qquad \mu(g,\mathcal{C}) = (\mathcal{C}, g\mathcal{C}g^{-1}), \tag{G.8}$$

is dominant and generically finite separable of degree two. After shrinking the target to a dense open, it is finite étale of degree two. If

$$L = k(G \times X), \qquad K = k(\mathcal{I}),$$

then L/K is a nontrivial quadratic Kummer extension. More precisely, there are $\rho \in L$ and $\Delta \in K^\times$ such that

$$L = K(\rho), \qquad \rho^2 = \Delta, \qquad \Delta \notin K^{\times 2}. \tag{G.9}$$

Consequently, this specified projective-transporter cover admits no rational section over the generic point of $\mathcal{I}$, equivalently over the function field K.

Proof. Proposition G.1 makes (G.8) surjective on geometric points. For a pair $(\mathcal{C}, \mathcal{D})$ and one transporter g_0, the scheme-theoretic fibre is the translate of the projective stabilizer,

$$g_0 N_G(\mathcal{C}) = \{g_0, g_0[e_0(\mathcal{C})]\}.$$

It is therefore reduced and étale of length two by (G.4). Since $G \times X$ is irreducible of dimension $3 + 4 = 7$, its image $\mathcal{I}$ is irreducible of dimension seven, and μ is generically finite of degree two. The reduced generic fibre makes the extension separable. Standard generic shrinking then makes the morphism finite étale of degree two.

The nontrivial deck involution is

$$\sigma(g, \mathcal{C}) = (g[e_0(\mathcal{C})], \mathcal{C}). \tag{G.10}$$

It has no fixed point, and its two-element orbits are exactly the fibres of μ. Hence $[L : K] = 2$ and $K = L^\sigma$. Choose $r \in L$ moved by σ and set

$$\rho = \frac{r - \sigma(r)}{2}, \qquad \Delta = \rho^2.$$

Then $\sigma(\rho) = -\rho$, so $\Delta \in K$, while $\rho \notin K$. Thus $L = K(\rho)$ and (G.9) holds except possibly for the nonsquare assertion. If $\Delta = b^2$ for some $b \in K$, then $(\rho/b)^2 = 1$. Odd characteristic forces $\rho/b \in \{1, -1\}$, contradicting $\rho \notin K$. Therefore the Kummer class is nontrivial.

A rational section on a dense open would induce a K-algebra retraction $L \to K$. In particular, it would send ρ to a square root of Δ in K, contradicting $\Delta \notin K^{\times 2}$. This proves the no-section statement. □

Appendix G.2 gives an explicit seven-dimensional dominant chart, identifies the fixed field, and verifies nonsplitting by a valuation calculation.

This no-section theorem concerns the specified cover. It neither proves finite-field hardness nor restricts the determinant-polarization branch of the universal partial algorithm.

Appendix G.2. An explicit dominant chart for the Kummer cover

Use G, X, $\mathcal{I}$ and $\mathcal{C}_t$ from Appendix G.1. The intrinsic equation is $\rho^2 = \Delta$; the following full-dimensional base change gives the explicit equation $\rho^2 = a$.

Let

$$T_0 = \{t : t(1 + t^2) \neq 0\}, \qquad J = \mathrm{diag}(1, -1), \qquad H = \{1, [J]\} \leq G, \qquad Q = G/H.$$

The normal-form calculation in Proposition G.1 shows that

$$T_0 \times G_x \times Q \longrightarrow \mathcal{I}, \quad (t, x, gH) \longmapsto (x\mathcal{C}_t x^{-1}, xg\mathcal{C}_t g^{-1} x^{-1}) \tag{G.11}$$

is dominant. Its dimension is $1 + 3 + 3 = 7$, so it retains both the full source-conjugation and relative-target parameters. The pullback of the transporter cover is

$$T_0 \times G_x \times G_g \longrightarrow T_0 \times G_x \times (G_g/H).$$

On a dense Gauss chart, represent the relative element by

$$g(\rho,u,v)=\begin{pmatrix}\rho & \rho u\\ v\rho & 1+v\rho u\end{pmatrix},\qquad \det(g)=\rho\neq 0. \tag{G.12}$$

Right multiplication by J, followed by a projectively irrelevant scalar -1, acts as

$$(\rho,u,v)\longmapsto(-\rho,-u,v).$$

With $s=u/\rho$, $a=\rho^2$, and $F=k(T_0\times G_x)$, the exact fixed-field calculation is

$$F(\rho,u,v)=F(\rho,s,v),\qquad F(\rho,s,v)^H=F(a,s,v),\qquad \rho^2=a. \tag{G.13}$$

The a-adic valuation of a is one, whereas the valuation of every square in $F(a,s,v)^\times$ is even. Thus a is nonsquare and the dominant pullback is nonsplit. This seven-dimensional chart, rather than a lower-dimensional diagonal specialization, supports the explicit equation (G.13).

Appendix H. Higher moments and arithmetic interfaces

The selected algorithm uses trace ratios and the bounded-degree splitter. Here we analyze the different arithmetic interfaces of adaptive moments and worst-case projective similarity. The conclusions concern the stated constructions and source families, not general complexity lower bounds.

Appendix H.1. Higher trace moments and projective scale

We first identify exactly when power traces can provide a transverse functional. The characterization is independent of cyclicity.

Proposition H.1 (Vanishing of all positive power traces)**.** *Let $A\in \mathrm{M}_n(\mathbb{F}_q)$, and let $p=\mathrm{char}(\mathbb{F}_q)$. Over an algebraic closure, let $\lambda_1,\dots,\lambda_t$ be the distinct nonzero eigenvalues of A, with algebraic multiplicities $e_1,\dots,e_t$, and let e_0 be the algebraic multiplicity of zero. Then the following conditions are equivalent:*

1. $\mathrm{Tr}(A^k)=0$ *for every* $k\geq 1$*;*
2. $p\mid e_i$ *for every* $1\leq i\leq t$*;*
3. $$\chi_A(x)=x^{e_0}h(x)^p\qquad \text{for some } h\in\mathbb{F}_q[x] \text{ with } h(0)\neq 0. \tag{H.1}$$

If these conditions fail, then $\mathrm{Tr}(A^k)\neq 0$ for some $1\leq k\leq n$. Consequently

$$s(A)=\min\{s\in\{0,\dots,n-1\}:\mathrm{Tr}(A^{s+1})\neq 0\} \tag{H.2}$$

is defined outside the family in (H.1)*. For every $\beta\in\mathbb{F}_q^\times$ and $S\in\mathrm{GL}_n(\mathbb{F}_q)$,*

$$s(\beta SAS^{-1})=s(A). \tag{H.3}$$

Proof. For every $k\geq 1$, the zero eigenvalues make no contribution and

$$\mathrm{Tr}(A^k)=\sum_{i=1}^{t}e_i\lambda_i^k. \tag{H.4}$$

The $t\times t$ matrix $(\lambda_i^k)_{1\leq k,i\leq t}$ is the usual Vandermonde matrix with exponents $0,\dots,t-1$, followed by multiplication by $\mathrm{diag}(\lambda_1,\dots,\lambda_t)$. It is therefore nonsingular. Vanishing of the first t power traces is equivalent to the vanishing of every e_i in the ground characteristic, which is equivalent to $p\mid e_i$ for every $1\leq i\leq t$. This condition also makes (H.4) vanish for every positive k.

Finite fields are perfect. Hence the multiplicity of every nonzero root is divisible by p precisely when the exponent of every irreducible factor other than x is divisible by p. This proves the equivalence with (H.1). If the conditions fail, one of the first $t\leq n$ traces is nonzero, proving the last assertion. Finally, $\mathrm{Tr}((\beta SAS^{-1})^k)=\beta^k\,\mathrm{Tr}(A^k)$ for every $k\geq 1$, so the zero pattern of the power traces, and hence their least nonzero index, is unchanged. This proves (H.3). □

The upper bound $s(A) \le n-1$ cannot be replaced by a bound independent of n. If $p \nmid n$ and $a \neq 0$, the companion matrix of $x^n - a$ is cyclic and satisfies

$$\operatorname{Tr}(A^k) = 0 \quad (1 \le k < n), \qquad \operatorname{Tr}(A^n) = na \neq 0.$$

Moreover, cyclicity alone does not ensure that (H.2) exists. A cyclic nilpotent Jordan block has all positive power traces zero. In characteristic p, the invertible cyclic block $J_p(a)$ with $a \neq 0$ also has $\operatorname{Tr}(J_p(a)^k) = pa^k = 0$ for every $k \ge 1$. These examples lie in the exceptional family of Proposition H.1.

Outside that family, the geometric construction from Section 5 applies at the canonical transverse index, but it does not by itself determine the projective scale.

Proposition H.2 (Adaptive geometry and its scale interface)**.** *Suppose that $n \ge 3$ and $H(\mathcal{C}) = \langle A_0 \rangle$, where A_0 is cyclic, and suppose that A_0 is outside the exceptional family in* (H.1)*. Let $s = s(A_0)$ and put*

$$\mathcal{C}_0 = \ker\big(\operatorname{Tr}(A_0^s \cdot)|_{\mathcal{C}}\big).$$

Then $\mathcal{C} = \langle A_0 \rangle \oplus \mathcal{C}_0$, the trace form on $\mathcal{C}_0$ is nondegenerate, and the Riesz representatives of Section 5 are defined. With

$$r = \begin{cases} 2, & s = 0, \\ 0, & s > 0, \end{cases} \qquad v_{r,s} = A_0^r - \frac{\operatorname{Tr}(A_0^{r+1})}{\operatorname{Tr}(A_0^{s+1})} A_0^s,$$

the marked class of $v_{r,s}$ in $A_0^{\perp_b}/\langle A_0 \rangle$ is nonzero. If $U = \mathcal{C}/\langle A_0 \rangle$ and $M_r \in \mathcal{C}_0$ is the corresponding Riesz representative, then

$$\pi_{A_0}(M_r) = P_U(\pi_{A_0}(v_{r,s})). \tag{H.5}$$

If $\mathcal{D} = S\mathcal{C}S^{-1}$ and a target hull generator is chosen as $B_0 = \alpha S A_0 S^{-1}$, then the target representative N_r satisfies

$$N_r = \alpha^r S M_r S^{-1}, \qquad \frac{\operatorname{Tr}(B_0^{s+1})}{\operatorname{Tr}(A_0^{s+1})} = \alpha^{s+1}. \tag{H.6}$$

Thus the one-hull decomposition, Riesz construction, and correlated projection law extend to the adaptive index; the remaining datum in (H.6) *is the projective scale α.*

Proof. Lemma 5.1 gives the complement and its nondegeneracy. Since $\operatorname{Tr}(A_0^2) = 0$, the least transverse index is not one. For $s = 0$, the selected exponents are $0, 1, 2 < n$; for $s > 0$, they are $0, 1, s < n$. Cyclicity therefore gives the nonzero marked class by Lemma 5.4. The quotient identity and covariance are precisely Lemma 5.3 and Proposition 5.5 at these indices. □

The scale equation can be sharpened by combining its nonzero homogeneous weights. Let

$$K(A_0) = \{k \in \{1, \dots, n\} : \operatorname{Tr}(A_0^k) \neq 0\}, \qquad d = \gcd K(A_0).$$

For a conjugate target, every $k \in K(A_0)$ gives

$$\rho_k = \frac{\operatorname{Tr}(B_0^k)}{\operatorname{Tr}(A_0^k)} = \alpha^k. \tag{H.7}$$

Bézout coefficients for the integers in $K(A_0)$ compute α^d as a product of powers of the ρ_k. If $\gcd(d, q-1) = 1$, inversion of d modulo $q-1$ recovers α by one exponentiation. Otherwise, complete recovery within this interface proceeds from a complete root list. For example, taking the least $k = s + 1 \le n$, list every root of

$$X^k - \rho_k.$$

There are at most $k \le n$ roots. Filter them by all equations (H.7), and apply the exact intertwiner, code, and ordered-slice checks of Proposition 5.6. Listing only one root is not sufficient: different roots can induce different pair-intertwiner systems, and only another root may conjugate the full code. Exhausting the list makes a negative answer exact. No complete root-list procedure for arbitrary (k, q) is assumed by the all-parameter algorithm.

Appendix H.2. A fifth-root family

The scale ambiguity persists even when the adaptive Riesz pair has scalar common centralizer.

Proposition H.3 (A route-specific fifth-root reduction)**.** *Let p be prime with $p \equiv 1 \pmod{25}$. In $\mathrm{M}_5(\mathbb{F}_p)$, index the standard basis by $0, \dots, 4$ and define*

$$Ae_i = e_{i+1} \ (0 \le i < 4), \qquad Ae_4 = e_0, \qquad M = \frac{1}{3}\operatorname{diag}(0,1,2,3,4).$$

Put $\mathcal{C}_p = \langle A, M\rangle$. For $\rho \in \mathbb{F}_p^\times$, let

$$B_\rho e_i = e_{i+1} \ (0 \le i < 4), \qquad B_\rho e_4 = \rho e_0, \qquad \mathcal{D}_{p,\rho} = \langle B_\rho, M\rangle.$$

Then $H(\mathcal{C}_p) = \langle A\rangle$, its canonical transverse index is $s = 4$, the selected Riesz index is $r = 0$ with $M_0 = M$, and

$$\operatorname{Cent}(A, M) = \operatorname{Cent}(B_\rho, M) = \mathbb{F}_p I_5.$$

Moreover,

$$\mathcal{C}_p \text{ and } \mathcal{D}_{p,\rho} \text{ are conjugate} \iff \rho \in (\mathbb{F}_p^\times)^5. \tag{H.8}$$

Whenever the codes are conjugate, any ambient transporter determines a fifth root of ρ using field and 5×5 matrix operations.

Proof. The trace Gram matrices of the displayed bases of $\mathcal{C}_p$ and $\mathcal{D}_{p,\rho}$ are both

$$\begin{pmatrix} 0 & 0 \\ 0 & 10/3 \end{pmatrix}.$$

Indeed, A^2, B_ρ^2, AM, and $B_\rho M$ have zero diagonal, while $\operatorname{Tr}(M^2) = 10/3 \neq 0$. Hence

$$H(\mathcal{C}_p) = \langle A\rangle, \qquad H(\mathcal{D}_{p,\rho}) = \langle B_\rho\rangle.$$

Furthermore,

$$\operatorname{Tr}(A^k) = 0 \quad (1 \le k < 5), \qquad \operatorname{Tr}(A^5) = 5,$$

so $s = 4$. The transverse complement is $\langle M\rangle$, and $\operatorname{Tr}(M) = \operatorname{Tr}(M^2) = 10/3$ makes its weight-zero Riesz representative equal to M. A matrix commuting with M is diagonal because the five diagonal entries of M are distinct. Commutation with either connected weighted five-cycle A or B_ρ then makes all diagonal entries equal, proving the two common-centralizer identities.

Suppose first that $\rho = \alpha^5$. For

$$S_\alpha = \operatorname{diag}(1, \alpha^{-1}, \alpha^{-2}, \alpha^{-3}, \alpha^{-4}),$$

direct calculation gives

$$B_\rho = \alpha S_\alpha A S_\alpha^{-1}, \qquad S_\alpha M S_\alpha^{-1} = M.$$

Thus $S_\alpha \mathcal{C}_p S_\alpha^{-1} = \mathcal{D}_{p,\rho}$.

Conversely, let $Q\mathcal{C}_pQ^{-1} = \mathcal{D}_{p,\rho}$. Hull covariance gives a unique $\delta \in \mathbb{F}_p^\times$ such that

$$QAQ^{-1} = \delta B_\rho.$$

Taking fifth powers and using $A^5 = I_5$ and $B_\rho^5 = \rho I_5$ yields

$$1 = \delta^5 \rho, \qquad (\delta^{-1})^5 = \rho.$$

The scalar δ is read from any nonzero corresponding matrix entry and the last identity is checked directly. This proves (H.8) and the extraction assertion. □

The source $\mathcal{C}_p$ is independent of ρ and satisfies the pointwise adaptive predicate. Proposition H.3 is therefore a pointwise, family-restricted reduction from promised fifth-root extraction to MCC search: the search procedure must handle every member of this source–target family. A partial algorithm guaranteed only on a c/q fraction of iid sources may return `FAILURE` on this fixed source. Its average-source guarantee alone does not activate the reduction. The congruence $p \equiv 1 \pmod{25}$ ensures that $25 \mid p - 1$; it separates this family from the elementary exponent inversion available when fifth powering is invertible on the relevant subgroup.

Appendix H.3. Worst-case line and hyperplane endpoints

For this subsection only, r denotes the companion size and root exponent, rather than the moment index used above. For $u \in \mathbb{F}_q^\times$, let

$$J_r(u) = \mathrm{Comp}(x^r - u), \qquad L_r = \langle J_r(1) \rangle, \qquad L_{r,c} = \langle J_r(c) \rangle.$$

Proposition H.4 (Companion lines and their trace duals)**.** *For every finite field $\mathbb{F}_q$, every $r \geq 2$, and every $c \in \mathbb{F}_q^\times$,*

$$L_r \text{ and } L_{r,c} \text{ are conjugate} \iff c \in (\mathbb{F}_q^\times)^r. \tag{H.9}$$

This remains valid when $\mathrm{char}(\mathbb{F}_q)$ *divides r. The same equivalence holds for the hyperplanes $L_r^{\perp_b}$ and $L_{r,c}^{\perp_b}$.*

Proof. If the lines are conjugate, then some $S \in \mathrm{GL}_r(\mathbb{F}_q)$ and $\gamma \in \mathbb{F}_q^\times$ satisfy

$$S J_r(1) S^{-1} = \gamma J_r(c).$$

Scalar multiplication changes the characteristic polynomial by

$$\chi_{\gamma J_r(c)}(x) = \gamma^r \chi_{J_r(c)}(x/\gamma) = x^r - \gamma^r c.$$

Comparison with $\chi_{J_r(1)}(x) = x^r - 1$ gives $(\gamma^{-1})^r = c$. This argument uses neither eigenvalues nor separability.

Conversely, if $x^r = c$, put

$$S_x = \mathrm{diag}(1, x^{-1}, \ldots, x^{-(r-1)}).$$

A basis-vector calculation gives

$$J_r(c) = x S_x J_r(1) S_x^{-1},$$

which proves line conjugacy. Finally, conjugation is an isometry of $b(X, Y) = \mathrm{Tr}(XY)$, so $(SLS^{-1})^{\perp_b} = SL^{\perp_b} S^{-1}$. Applying orthogonal complementation twice proves the hyperplane assertion. □

Existence of the projective scale can nevertheless be decided without first constructing it.

Proposition H.5 (Decision and search for projective similarity)**.** *Let $A, B \in \mathrm{M}_n(\mathbb{F}_q)$ be nonzero. Whether B is similar to αA for some $\alpha \in \mathbb{F}_q^\times$ can be decided deterministically using invariant factors, integer gcds, and exponentiation in $\mathbb{F}_q$. If the answer is positive, constructing α and a transporter requires at most one h-th-root operation, where $h \leq n$.*

Proof. Write the monic invariant factors of A as

$$f_i^A(x) = x^{d_i} + \sum_{j=1}^{d_i} a_{ij} x^{d_i - j}.$$

The corresponding invariant factor of αA is

$$f_i^{\alpha A}(x) = \alpha^{d_i} f_i^A(x/\alpha) = x^{d_i} + \sum_{j=1}^{d_i} a_{ij} \alpha^j x^{d_i - j}. \tag{H.10}$$

Unequal degree lists or unequal zero patterns give an exact negative answer. For every remaining pair of nonzero corresponding coefficients, write $\rho_\ell = b_{ij}/a_{ij}$ when the coefficient of $x^{d_i - j}$ in the target invariant factor is b_{ij}, and put $j_\ell = j$. The necessary scale equation is

$$\alpha^{j_\ell} = \rho_\ell. \tag{H.11}$$

Put $Q = q - 1$ and

$$h = \gcd(Q, j_1, \dots, j_t).$$

Extended Euclid gives integers $u_1, \dots, u_t, v$ such that $\sum_\ell u_\ell j_\ell + vQ = h$. Define

$$\eta = \prod_{\ell=1}^{t} \rho_\ell^{u_\ell}.$$

The simultaneous system (H.11) has a solution if and only if

$$\eta^{Q/h} = 1, \qquad \rho_\ell = \eta^{j_\ell/h} \quad (1 \le \ell \le t). \tag{H.12}$$

Necessity follows from $\eta = \alpha^h$. Conversely, the first test says that η lies in the image of $x \mapsto x^h$ on the cyclic group $\mathbb{F}_q^\times$, and the second test makes any solution of

$$\alpha^h = \eta$$

satisfy every original weight equation. Thus (H.12) decides existence by gcds and binary exponentiation. If there are no nonzero coefficient equations, all invariant factors are matching monomials and $\alpha = 1$ works.

When equations are present, $h \le \min_\ell j_\ell \le n$. One root of $\alpha^h = \eta$, followed by deterministic rational-canonical-form bases, constructs S with $B = S(\alpha A)S^{-1}$. Direct verification of this identity completes the search procedure. □

The trace-good endpoint branches of Section 7 fix their scales by a nonzero trace and do not require the root operations above. Their positive source mass does not provide a transporter on every companion-line or hyperplane instance.

CRediT authorship contribution statement

Jingchuan Ma: Conceptualization, Methodology, Formal analysis, Validation, Software, Writing – original draft, Writing – review and editing.

Yanhua Liu: Supervision, Writing – review and editing.

Qiaoyun Huang: Validation, Writing – review and editing.

Funding

No funding was received for this work.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data and code availability

No external empirical dataset was used in this study. The mathematical proofs are contained in the article. The computational scripts, recorded outputs, worked-example checks, finite regression checks, and reproduction instructions used for auxiliary validation are archived on Zenodo as *Computational Checks for Deterministic Partial Algorithms for Matrix-Space Conjugacy over Finite Fields*, version 1.0.0, DOI: 10.5281/zenodo.22643362. These computations provide auxiliary validation and are not dependencies of the universal mathematical proofs.

Preprint version history

A preprint version of this work and its revision history are available as IACR Cryptology ePrint Archive Report 2026/1664, https://eprint.iacr.org/2026/1664. The ePrint record concerns the same work as the present manuscript and is not an independent prior algorithm.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used OpenAI's ChatGPT and Codex to assist with language polishing and with the clarification and formal presentation of mathematical proofs. The authors reviewed and edited all AI-assisted material and take full responsibility for the content of the article.

References

Bardet, M., Brion, C., Gaborit, P., Haiech, M., Neveu, R., 2026. The matrix subcode equivalence problem and its application to signature with MPC-in-the-head. Designs, Codes and Cryptography 94. doi:https://doi.org/10.1007/s10623-026-01836-8, arXiv:2507.15377. article 76.

Bongartz, K., Friedland, S., 2026. Complete invariants for simultaneous similarity. Advances in Mathematics 499, 111040. doi:https://doi.org/10.1016/j.aim.2026.111040, arXiv:2601.00379.

Carlet, C., Mesnager, S., Tang, C., Qi, Y., 2019. New characterization and parametrization of LCD codes. IEEE Transactions on Information Theory 65, 39–49. doi:https://doi.org/10.1109/TIT.2018.2829873, arXiv:1709.03217.

Couvreur, A., Levrat, C., 2025. Highway to hull: An algorithm for solving the general matrix code equivalence problem, in: Kalai, Y.T., Kamara, S.F. (Eds.), Advances in Cryptology – CRYPTO 2025, Springer Nature Switzerland, Cham. pp. 253–283. doi:https://doi.org/10.1007/978-3-032-01855-7_9, arXiv:2504.01230. IACR Cryptology ePrint Archive, Paper 2025/596.

D'Alconzo, G., 2024. Monomial isomorphism for tensors and applications to code equivalence problems. Designs, Codes and Cryptography 92, 1961–1982. doi:https://doi.org/10.1007/s10623-024-01375-0. IACR Cryptology ePrint Archive, Paper 2023/396.

Glasby, S.P., Ihringer, F., Mattheus, S., 2023. The proportion of non-degenerate complementary subspaces in classical spaces. Designs, Codes and Cryptography 91, 2879–2891. doi:https://doi.org/10.1007/s10623-023-01235-3, arXiv:2207.04678.

Glasby, S.P., Niemeyer, A.C., Praeger, C.E., 2022. Random generation of direct sums of finite non-degenerate subspaces. Linear Algebra and its Applications 649, 408–432. doi:https://doi.org/10.1016/j.laa.2022.05.016, arXiv:2111.02198.

Grochow, J.A., Qiao, Y., 2023. On the complexity of isomorphism problems for tensors, groups, and polynomials I: Tensor isomorphism-completeness. SIAM Journal on Computing 52, 568–617. doi:https://doi.org/10.1137/21M1441110, arXiv:1907.00309.

Grochow, J.A., Qiao, Y., Tang, G., 2022. Average-case algorithms for testing isomorphism of polynomials, algebras, and multilinear forms. Journal of Groups, Complexity, Cryptology 14. doi:https://doi.org/10.46298/jgcc.2022.14.1.9431, arXiv:2012.01085.

Guo, Z., 2020. Deterministic polynomial factoring over finite fields: A uniform approach via P-schemes. Journal of Symbolic Computation 96, 22–61. doi:https://doi.org/10.1016/j.jsc.2019.02.011.

Klopsch, B., Voll, C., 2009. Igusa-type functions associated to finite formed spaces and their functional equations. Transactions of the American Mathematical Society 361, 4405–4436. doi:https://doi.org/10.1090/S0002-9947-09-04671-6, arXiv:math/0603565.

Kreuzer, M., 2025. Code equivalence, point set equivalence, and polynomial isomorphism. URL: https://arxiv.org/abs/2511.06843v2. arXiv:2511.06843v2 (2026 revision).

Li, T., Li, Y., Qiao, Y., Tao, D., Wang, Y., 2026. On the average-case complexity landscape for Tensor-Isomorphism-complete problems over finite fields. URL: https://arxiv.org/abs/2604.00591v1, arXiv:2604.00591v1.

Ma, J., 2026. Computational checks for deterministic partial algorithms for matrix-space conjugacy over finite fields. Zenodo. URL: https://doi.org/10.5281/zenodo.22643362, doi:https://doi.org/10.5281/zenodo.22643362. version 1.0.0.

Massey, J.L., 1992. Linear codes with complementary duals. Discrete Mathematics 106–107, 337–342. doi:https://doi.org/10.1016/0012-365X(92)90563-U.

Murray, S.H., Roney-Dougal, C.M., 2011. Constructive homomorphisms for classical groups. Journal of Symbolic Computation 46, 371–384. doi:https://doi.org/10.1016/j.jsc.2010.09.003.

Narayanan, A.K., Qiao, Y., Tang, G., 2024. Algorithms for matrix code and alternating trilinear form equivalences via new isomorphism invariants, in: Joye, M., Leander, G. (Eds.), Advances in Cryptology – EUROCRYPT 2024, Springer Nature Switzerland, Cham. pp. 160–187. doi:https://doi.org/10.1007/978-3-031-58734-4_6. IACR Cryptology ePrint Archive, Paper 2024/368.

Neumann, P.M., Praeger, C.E., 1995. Cyclic matrices over finite fields. Journal of the London Mathematical Society 52, 263–284. doi:https://doi.org/10.1112/jlms/52.2.263.

Qiao, Y., Sun, X., 2024. Canonical forms for matrix tuples in polynomial time, in: 2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS), IEEE. pp. 780–789. doi:https://doi.org/10.1109/FOCS61266.2024.00054, arXiv:2409.12457.

Reijnders, K., Samardjiska, S., Trimoska, M., 2024. Hardness estimates of the code equivalence problem in the rank metric. Designs, Codes and Cryptography 92, 833–862. doi:https://doi.org/10.1007/s10623-023-01338-x. IACR Cryptology ePrint Archive, Paper 2022/276.

Sendrier, N., 1997. On the dimension of the hull. SIAM Journal on Discrete Mathematics 10, 282–293. doi:https://doi.org/10.1137/S0895480195294027.

Sendrier, N., 2000. Finding the permutation between equivalent linear codes: The support splitting algorithm. IEEE Transactions on Information Theory 46, 1193–1203. doi:https://doi.org/10.1109/18.850662.

Wall, G.E., 1999. Counting cyclic and separable matrices over a finite field. Bulletin of the Australian Mathematical Society 60, 253–284. doi:https://doi.org/10.1017/S000497270003639X.