A SHARP SIX-NEIGHBOR EXPRESSIVITY THRESHOLD FOR THIRD-ORDER GEOMETRIC INVARIANTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Many geometric learners compress a local point cloud into pooled low-order invariants before regression. When does that compression first destroy information? We determine the exact answer for the complete reflection-even third-order correlation in three dimensions. Every single-type cloud with at most five distinct neighbors is reconstructed, up to permutation and $O(3)$, whereas six neighbors admit non-congruent clouds with identical complete third-order information. Equivalently, rank-at-most-three positive-semidefinite Gram matrices of order at most five are reconstructible from their unlabelled three-vertex principal-submatrix decks, and order six is the first failure. The five-neighbor result is computer assisted but exact: all $B_{10} = 115{,}975$ edge-equality patterns are reduced to five maximal ambiguity templates, resolved by parallel geometric/algebraic arguments and thirteen rational ideal-membership certificates. Every exact six-neighbor collision has repeated radii and lies in a measure-zero discriminant; the known four-parameter family is locally a codimension-eleven collision branch. The threshold has direct learning consequences. Below it, every continuous invariant target on a compact domain factors continuously through the representation; at six neighbors, a smooth four-vertex target imposes an analytic Bayes error floor. On 11,000 exact collision pairs, third-order models converge to that floor regardless of data or capacity. Fourth-order cards and a sufficiently deep edge-incidence MPNN remove the error, while depth-0/1 controls with matched parameter counts do not. Perturbation experiments show that the repair persists near the collision manifold. The result separates a representation-level obstruction from basis truncation, optimization, and model size.

1 INTRODUCTION

Geometric learning systems must respect rigid-motion and permutation symmetries without discarding information needed for prediction. This tension is central to molecular and materials learning, where local environments are encoded by symmetry-adapted descriptors or equivariant messages before an energy, force, or other observable is predicted (Bartók et al., 2013; Thompson et al., 2015; Drautz, 2019; Batatia et al., 2022; Batzner et al., 2022). It is also a general question about point-cloud representations: when is an invariant feature map injective, and when can no downstream learner repair its collisions?

Atom-centred density correlations provide a particularly clean hierarchy. Second-order correlations give the SOAP power spectrum, third-order correlations the bispectrum used by SNAP, and higher products form scalar ACE bases. The untruncated hierarchy is systematically improvable (Dusson et al., 2022), but every fixed order pools small motifs without retaining which motifs share points. Counterexamples are known at several orders (Pozdnyakov et al., 2020; Nigam et al., 2024; Domina & Ceriotti, 2026), and recent numerical work studies inversion from compressed correlations (Nigam et al., 2026). These works establish that incompleteness exists, but they do not identify the smallest cloud on which the complete reflection-even third-order invariant fails.

We close this finite-cardinality gap and translate it into an expressivity threshold. Our contributions are:

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1. We prove that the complete reflection-even third-order correlation is globally injective for at most five distinct points in $\mathbb{R}^3$, including repeated distances and singular configurations.
2. We combine this reconstruction theorem with a known six-point family to prove that six is the exact first failure. The key mechanism is that rank-three positive-semidefinite (PSD) realizability eliminates every abstract five-vertex triangle-deck ambiguity.
3. We derive continuous factorization below the threshold and, at the threshold, endpoint, Bayes-risk, and integrated gradient-error lower bounds for a smooth invariant target.
4. We isolate the missing information experimentally. Increasing data or capacity leaves a D_3 learner at the analytic floor; D_4 and an edge-incidence MPNN repair the task. A parameter-matched depth ablation locates the repair at two propagation rounds rather than at model size.

The claim is deliberately narrower than a statement about all equivariant networks. Higher-order features, parity-odd channels, message passing, or neighbor-indexed equivariants can retain information removed by third-order scalar pooling. Our theorem identifies exactly where that particular information bottleneck first becomes possible.

2 Representation and reconstruction problem

Let $X = \{x_1, \dots, x_N\} \subset \mathbb{R}^3$ be a single-type point cloud relative to a fixed origin, with distinct x_i. We quotient by $O(3) \times S_N$. Writing $\mu_X = \sum_i \delta_{x_i}$, define the complete reflection-even third-order correlation

$$C_3(X) = \int_{O(3)} (Q_{\#}\mu_X)^{\otimes 3} \, \mathrm{d}H(Q), \tag{1}$$

where H is normalized Haar measure. This is an ideal infinite-resolution object; finite radial and angular implementations are projections of it.

The Gram matrix $G_X = (x_i \cdot x_j)$ is PSD with rank at most three and determines X up to $O(3)$. For $I \subset [N]$, let $G_X[I]$ be a principal submatrix. Define the unlabelled Gram three-deck

$$\mathcal{D}_3(X) = \left\{\!\!\left\{ [G_X[I]]_{S_3} : I \in \binom{[N]}{3} \right\}\!\!\right\}, \tag{2}$$

together with repeated-index cards, equivalently the one- and two-index decks. Explicitly, $\mathcal{D}_3 = (D_1, D_2, D_3^{\neq})$, with D_k empty for $k > N$; thus $N = 1, 2$ retain the full Gram matrix as a lower-order card.

Lemma 1 (Correlation–deck equivalence)**.** *For fixed N, $C_3(X) = C_3(Y)$ if and only if $\mathcal{D}_3(X) = \mathcal{D}_3(Y)$.*

Proof sketch. The $O(3)$ orbit of an ordered triple is classified by its 3×3 Gram matrix. Pushing equation equation 1 through the Gram map gives the atomic measure of all ordered triple Gram matrices. Conversely, each Gram atom specifies a unique compact orbit and hence its Haar orbit measure. Quotienting the tensor slots gives equation equation 2 and the repeated-index cards. The full argument is in Supplementary Section S1. □

Thus the learning representation is exactly a multiset of local Gram cards. The use of $O(3)$ matters: parity-odd $SO(3)$ pseudoscalars are not functions of Gram matrices and are outside our theorem.

3 Sharp expressivity threshold

Theorem 1 (Sharp third-order threshold)**.** *If $N \le 5$ and $\mathcal{D}_3(X) = \mathcal{D}_3(Y)$, then $Y = QX_\pi$ for some $Q \in O(3)$ and $\pi \in S_N$. For $N = 6$, there exist distinct-point clouds with equal three-decks that are not related by $O(3) \times S_6$.*

The six-point pair can be chosen with integer coordinates. Let

$$u = (-2, -1, 0), \qquad v = (-2, 0, -1), \qquad w = (-1, 0, 2),$$
$$s = (0, -1, -2), \qquad t = (0, -1, 2),$$

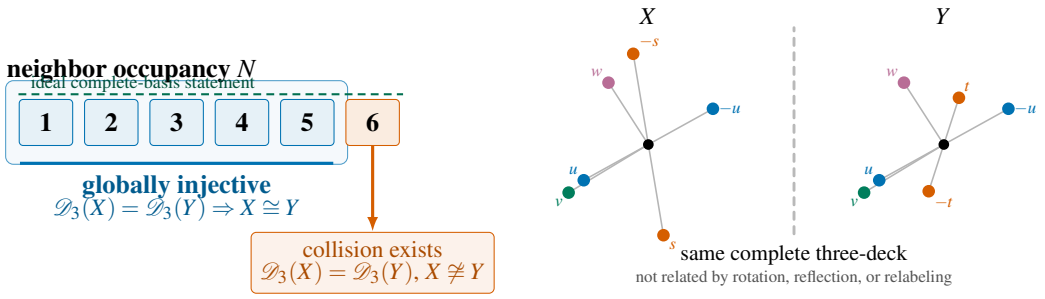


Figure 1: **Left:** the complete third-order invariant is injective through five points and first fails at six. **Right:** an exact integer collision. The two clouds share the same complete Gram three-deck but have non-permutation-congruent full Gram matrices.

Table 1: Five ambiguity templates, in edge order $01, 02, 03, 04, 12, 13, 14, 23, 24, 34$. M1–M4 have equal radii; R41 has radius multiplicity $4 + 1$. Full branch conditions are in S4–S5.

	First pattern	Second pattern	Resolution
M1	`0011234562`	`0011235462`	projected planar frames
M2	`0011023240`	`0013014220`	equilateral-anchor Schur constraints
M3	`0123234401`	`0123241340`	principal minors; 7 certificates
M4	`0011203212`	`0012113202`	adjugate Schur equations; 6 certificates
R41	`0123245653`	`0125243635`	kernel/range and planar frames

and set

$$X = \{u, v, w, s, -s, -u\}, \qquad Y = \{u, v, w, t, -t, -u\}. \tag{3}$$

Their order-one, two, and three decks agree exactly, but none of the 720 simultaneous row–column permutations maps G_X to G_Y. This pair is a member of the continuous six-neighbor family in Domina & Ceriotti (2026); our new contribution is the universal lower side of the threshold.

For the continuous family, $h = \|u\|^2$, $e = \|v\|^2 = \|w\|^2$, $a = u \cdot v$, $b = u \cdot w$, and $v \cdot w = 0$, with $h, e > 0$ and $he > a^2 + b^2$. The integer pair has $(h, e, a, b) = (5, 5, 4, 2)$; the rational audit point $(5, 3, 1, 2)$ below is a different member.

Proposition 1 (Size of the six-neighbor obstruction)**.** *Every noncongruent exact C_3 collision satisfies $\Delta_{\mathrm{rad}}(G_X) = \Delta_{\mathrm{rad}}(G_Y) = 0$, where $\Delta_{\mathrm{rad}}(G) = \prod_{i<j}(G_{ii} - G_{jj})$. Consequently the collision set has zero 15-dimensional measure on the regular six-neighbor Gram quotient (rank-three PSD matrices modulo permutations, locally with trivial stabilizer). At $(h, e, a, b) = (5, 3, 1, 2)$, the family in Supplementary Section S2 is a smooth four-dimensional paired-collision branch, whose projection to either environment quotient has local codimension eleven.*

Distinct diagonal entries label every vertex and recover every off-diagonal entry from the cards. Supplementary Section S2 supplies the measure argument and exact rank-38 audit of all 96 compatible card matchings. The dimension claim is local to this branch, not a global classification.

Why five is nontrivial. General node- and edge-weighted complete graphs on five vertices are not reconstructible from weighted subtriangle histograms (Boutin & Kemper, 2007). We enumerate all $B_{10} = 115{,}975$ equality patterns of the ten edges of K_5 for every radius multiplicity. Abstract collisions occur only for equal radii and radius type $4 + 1$: respectively 80 and 43 value-aligned nonisomorphic pairs. Every candidate is a color coarsening of one of five maximal templates. Three templates are eliminated by geometric arguments; the other two have parallel Schur-complement and principal-minor proofs and thirteen exact rational ideal-membership identities. In every branch, PSD rank at most three forces a vertex relabeling or coincident points.

Geometric intuition: gluing cards to an anchor triangle. The generic mechanism is simpler than the exceptional-case algebra suggests. Fix a nonsingular three-point card with Gram block B. A fourth point is described relative to this anchor by a dot-product vector $p \in \mathbb{R}^3$. PSD rank three

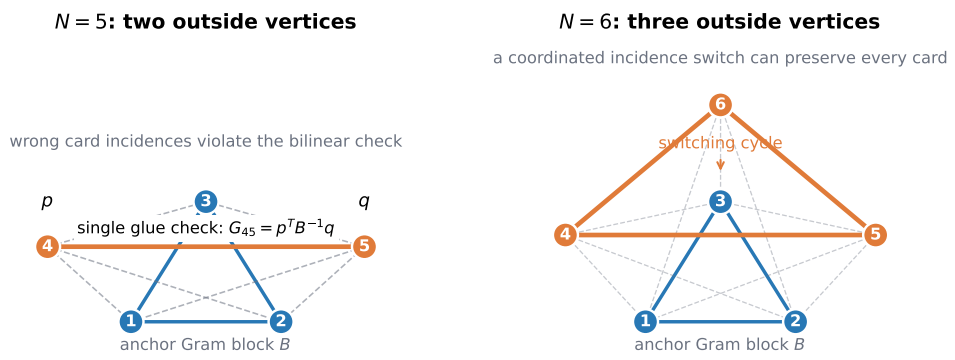


Figure 2: Anchor-triangle intuition. At five vertices only one bilinear edge tests how the two outside cards are glued. At six, three outside vertices form the first switching cycle. The drawing explains the generic mechanism; the template proof covers singular anchors and repeated Gram entries.

leaves no unobserved height: its squared radius and coordinate are forced by

$$G_{44} = p^T B^{-1} p, \qquad x_4 = X_B^{-1} p, \tag{4}$$

where $X_B X_B^T = B$. With five points there are only two outside vectors p, q. The vanishing Schur complement is the concrete rank-three constraint

$$\begin{pmatrix} G_{44} & G_{45} \\ G_{45} & G_{55} \end{pmatrix} = \begin{pmatrix} p^T \\ q^T \end{pmatrix} B^{-1} \begin{pmatrix} p & q \end{pmatrix}. \tag{5}$$

Its diagonal entries validate each outside point separately; its single off-diagonal entry is the global glue check $G_{45} = p^T B^{-1} q$. Abstract triangle-deck ambiguities try to attach entries of p and q to different anchor vertices; the diagonal constraints and the single cross check above generically reject such attachments. The complete template analysis shows that every exceptional attachment either relabels or degenerates. Supplementary Sections S3–S5 give the complete enumeration, template arguments, certificate targets, and trust boundary. For a singular anchor, PSD requires $p, q \in \operatorname{im} B$, but leaves normal components: $G_{44} = p^T B^{+} p + \|x_4^{\perp}\|^2$. The residual Schur block must be PSD of rank at most $3 - \operatorname{rank} B$. Thus range conditions alone do not reconstruct the outside points; the global template/minor analysis handles singular anchors and repeated entries.

At six points there are three outside vectors and three outside–outside edges. A coordinated switch around this incidence cycle can preserve the multiset of all triangle cards while changing their global gluing; the family in Supplementary Section S2 realizes exactly such a switch. Thus the threshold is not a naive parameter count: rank three fixes every point once its incidences are known, while six vertices are the first number that permits two different globally consistent incidence assignments.

4 Consequences for Learning

Theorem 1 yields both a positive and a negative statement about downstream learning. For $R, \eta > 0$, let $\mathcal{X}_N(R, \eta)$ be the quotient of clouds satisfying $\|x_i\| \le R$, $\|x_i\| \ge \eta$, and $\|x_i - x_j\| \ge \eta$. This is compact. Use the coordinate orbit metric $d_{\mathrm{coord}}(X, Y) = \min_{Q,\pi}(N^{-1} \sum_i \|x_i - Q y_{\pi(i)}\|^2)^{1/2}$. Normalize $\bar{C}_3 = N^{-3} C_3$ and give measures the weak topology, metrized by W_1 on the compact support with average product ground metric. Coupling matched neighbors under the same Haar rotation gives $W_1(\bar{C}_3(X), \bar{C}_3(Y)) \le d_{\mathrm{coord}}(X, Y)$ (S1).

Corollary 1 (Continuous factorization below the threshold)**.** *For $N \le 5$, every continuous $O(3) \times S_N$-invariant target $V : \mathcal{X}_N(R, \eta) \to \mathbb{R}$ factors as*

$$V = f \circ C_3 \tag{6}$$

for a continuous f on $C_3(\mathcal{X}_N(R, \eta))$. Consequently, any universal learner on the compact descriptor image can approximate every such V.

Proof. The continuous map C_3 is injective by Theorem 1. A continuous bijection from a compact space to a Hausdorff image is a homeomorphism. Take $f = V \circ C_3^{-1}$. □

Injectivity implies uniform continuity of the inverse on each such compact domain, but not a Lipschitz inverse. Our five-neighbor transverse planar test has $d_{\mathrm{coord}} = \Theta(\varepsilon)$ and card distance $\Theta(\varepsilon^2)$ (measured slope 2.000000; S13), distinguishing coordinate stability from Gram stability, as motivated by inverse reconstruction studies (Nigam et al., 2026).

Proposition 2 (Conditioning at a collision)**.** *For a noncongruent exact collision (X_0, Y_0) and C^1 coordinate perturbations of size σ, the optimal-assignment card distance obeys $d_{D_3}(X_\sigma, Y_\sigma) = O(\sigma)$, while $d_{\mathrm{coord}}(X_\sigma, Y_\sigma) \geq d_{\mathrm{coord}}(X_0, Y_0) - O(\sigma)$. Outside a finite union of proper linear subspaces of perturbation directions, $d_{D_3} = \Theta(\sigma)$; hence the inverse secant ratio is $\Theta(\sigma^{-1})$.*

The finite matching expansion and its exceptional directions are proved in S2.

At six points, the failure is witnessed by a low-degree target, not only by an abstract orbit-separating construction. Define

$$T(G) = \sum_{i,j,k,\ell \text{ all distinct}} G_{ij} G_{jk} G_{k\ell}. \tag{7}$$

This injective path-homomorphism sum is a smooth $O(3) \times S_6$-invariant polynomial. For equation equation 3, exact integer arithmetic gives $T(G_Y) - T(G_X) = 192$.

Proposition 3 (Irreducible energy and gradient error)**.** *Let a prediction factor as $\widehat{V} = f \circ C_3$. If $C_3(X) = C_3(Y)$ and $\Delta = |V(X) - V(Y)|$, then*

$$\max\{|\widehat{V}(X) - V(X)|, |\widehat{V}(Y) - V(Y)|\} \geq \Delta/2. \tag{8}$$

For conditional probabilities p and $1 - p$, the minimum conditional MSE is $p(1-p)\Delta^2$. For the gradient claim, additionally assume $V, \widehat{V}$ are C^1 on an open neighborhood of an absolutely continuous coordinate path γ from X to Y, of length $L(\gamma) > 0$. Then

$$\operatorname*{ess\,sup}_t \|\nabla \widehat{V} - \nabla V\| \geq \Delta / L(\gamma). \tag{9}$$

The endpoint bound is the triangle inequality. The gradient bound follows by integrating $(\nabla \widehat{V} - \nabla V) \cdot \gamma'$ and using equality of the predicted endpoint energies. Neither bound depends on data size, network capacity, or optimization.

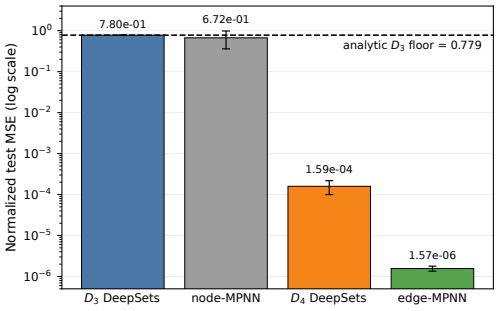
5 Controlled collision-family benchmark

Exact paired data and target. We sample the four-parameter six-point family described in Supplementary Section S2. Each parameter tuple gives two noncongruent clouds with bitwise-identical canonical D_3 decks. We use 8,000/1,000/2,000 pairs for training/validation/test and always split pairs, not individual branches. The target is the path polynomial in equation equation 7. On test pairs, the largest paired D_3 difference is zero, the smallest flattened D_4 distance is 1.922, and the mean target gap is 2.107. The analytic D_3 Bayes floor is 0.7789 after normalization by the training-target variance 1.616023. We draw $h, e \sim U[0.8, 1.2]$, $a/\sqrt{he} \sim U[0.65, 0.85]$, $b/\sqrt{he} \sim U[0.25, 0.45]$, rejecting $a^2 + b^2 \geq 0.95he$. This sampling region differs from the local-dimension audit point. Models minimize standardized MSE with AdamW; the original three baselines use batch 512, learning rate 10^{-3}, and at most 400 epochs, while the stable edge model uses batch 256, 3×10^{-4}, at most 800 epochs, gradient clipping and a plateau scheduler (full settings in S13).

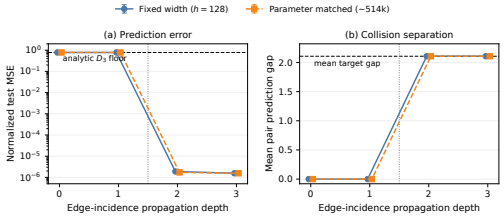
Representation repair, not a larger regressor. We compare DeepSets over complete D_3 or D_4 cards with complete-Gram-graph MPNNs. All models are invariant to neighbor order; five seeds are used. The stable edge MPNN updates edge states through incident edges, thereby retaining which pairwise relations share a vertex. Table 2 and Figure 3 show that D_3 saturates its exact floor, while D_4 and edge-incidence propagation recover the target. A conventional node-state MPNN is highly variable on this benchmark and is reported as a diagnostic, not as the main repair. It uses three

Table 2: Exact-collision regression, mean ± population standard deviation over five seeds. NMSE is raw test MSE divided by the training-target variance. The true mean pair gap is 2.107.

Representation	Parameters	Test NMSE	Predicted gap
D_3 DeepSets	67,073	0.78021 ± 0.00030	0
Node MPNN	396,289	0.672 ± 0.312	0.516 ± 0.759
D_4 DeepSets	67,585	$(1.59 \pm 0.59) \times 10^{-4}$	2.10462 ± 0.00146
Edge MPNN	513,793	$(1.565 \pm 0.217) \times 10^{-6}$	2.10685 ± 0.00018
D_3 analytic floor	–	0.77887	0



(a) Changing what information is retained.



(b) Fixed-width and parameter-matched depth controls.

Figure 3: **The barrier is architectural, not parametric. (a)** Complete D_3 pooling remains at the analytic floor; D_4 and edge-incidence propagation separate the pair. **(b)** Depth 0 and 1 remain near the floor in these runs, whereas depth 2 reduces NMSE by 4.1×10^5 (fixed width) and 4.5×10^5 (parameter matched). Error bars show five seeds.

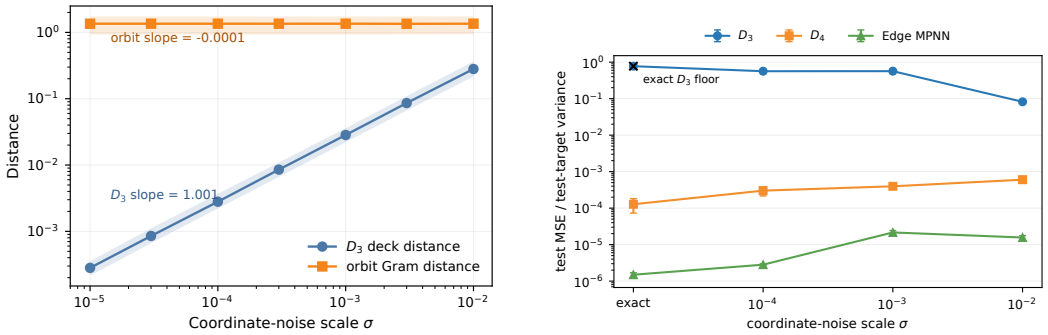
residual node-update layers with mean neighbor aggregation; seed variability alone does not diagnose an impossibility theorem for node MPNNs. The implemented directed-edge update is

$$\bar{h}_{ij}^t = \frac{\sum_{k \neq i,j} h_{ki}^t}{N-2}, \quad h_{ij}^{t+1} = \mathrm{LN}\left[\mathrm{GRU}_t\left(\mathrm{LN}(\phi_t[h_{ij}^t, \bar{h}_{ij}^t]), h_{ij}^t\right)\right]. \tag{10}$$

Initialization encodes (G_{ii}, G_{ij}, G_{jj}); readout applies an MLP to concatenated mean/max edge pools. Two rounds expose non-backtracking three-edge walks, including four-vertex paths. The observed depth transition is specific to this implementation and task, not a universal depth lower bound.

The depth transition. At fixed width, depths $0, 1, 2, 3$ obtain mean NMSE $0.7810, 0.7789, 1.90 \times 10^{-6}, 1.57 \times 10^{-6}$. Matching every depth to approximately 514k parameters changes these to $0.7805, 0.7789, 1.72 \times 10^{-6}, 1.57 \times 10^{-6}$; the maximum parameter mismatch is 0.083%. Predicted pair gaps simultaneously jump from below 4×10^{-7} to 2.107. Capacity and data sweeps for the pooled D_3 model remain pinned to the analytic floor (Supplementary Section S8), ruling out underfitting as the explanation.

Conditioning and noisy learning. We perturb both branches independently in coordinates; the noisy regression runs reuse random directions across noise levels, not between branches. Card-set distance is computed by optimal assignment with minimization over internal card permutations; orbit distance minimizes the full Gram discrepancy over all 720 vertex permutations. At zero noise the median D_3 distance is 8.17×10^{-15} while the median orbit distance is 1.352. The former grows linearly with coordinate noise (log–log slope 1.001), but the latter stays constant (slope -1.2×10^{-4}). In separately trained noisy-data experiments, D_4 and the edge MPNN stay below 6.1×10^{-4} and 2.2×10^{-5} NMSE, respectively, through noise $\sigma = 10^{-2}$. The D_3 error eventually falls because noise makes the formerly identical inputs distinguishable; that does not contradict the exact floor.



(a) Metric conditioning around the collision manifold.

(b) Independent branch noise, coupled across noise levels.

Figure 4: **Near-collision behavior.** The two physical orbits remain far apart while D_3 separates only proportionally to perturbation size. Higher order or incidence-aware models preserve their repair under the same noise. Shading/error bars show empirical quantiles or standard deviations over three seeds for the noise sweep (five for the exact benchmark).

6 Relation to modern geometric learning

Completeness and expressivity of invariant or equivariant geometric networks are active topics (Hordan et al., 2024; Li et al., 2025; Sverdlov & Dym, 2025). Those results ask whether particular graph architectures can separate general point clouds, often by retaining vertex incidence or iteratively propagating information. Our question is complementary: after a fixed-order motif histogram has already erased incidence, at what cardinality does that loss first become irreversible? The exact answer for third-order reflection-even pooling in three dimensions is six.

For atomistic models, conventional SOAP power spectra are lower order and may fail earlier (Pozdnyakov et al., 2020). Complete scalar bispectrum/SNAP features, and scalar ACE models restricted to the same correlation information, inherit the six-neighbor obstruction. Finite implementations may fail earlier. Higher-order ACE, MACE-style higher-order messages, parity-odd channels, neighbor-centred message passing, and species information may separate the displayed pair; Theorem 1 does not claim otherwise.

How common modelling choices change the obstruction. The theorem is best read as a representation diagnostic, not as a claim that every modern atomistic network fails at six neighbors. Fourth-order cards separate the benchmark directly; directed edge-incidence message passing retains which edges share vertices and can learn the same repair; parity-odd channels may add information absent from an $O(3)$ Gram representation. Species channels also refine the deck, but do not provide a universal guarantee: any homogeneous six-neighbor environment remains in the single-type setting.

To quantify this scope on the integer pair, we exhaustively assign each of its six vertices to $q = 1, \dots, 4$ species channels and compare the complete colored decks of orders one through three. Under a uniformly sampled shared assignment, exactly q^2 of the q^6 colorings retain the collision, giving survival probability q^{-4}: 1, $1/16$, $1/81$, and $1/256$. The surviving patterns have form $(\alpha, \beta, \beta, \alpha, \alpha, \alpha)$ in the displayed vertex order, including the homogeneous case $\alpha = \beta$. Thus generic chemical labels rapidly break this particular obstruction, while structured or elemental environments do not. This finite ablation limits the scope of the counterexample without weakening the single-type threshold theorem. Conditioning instead on fixed stoichiometry, an exact enumeration of all 11 integer partitions of six gives survival $1/15$ for A_4B_2, zero for A_3B_3, A_5B, and every composition using three or more species, and one for A_6 (S13). These probabilities assume uniform shared site assignments within the specified composition, not equilibrium chemistry.

7 LIMITATIONS AND CONCLUSION

The theorem concerns unit-weight, single-type, distinct-neighbor local clouds modulo $O(3)$. It does not establish a whole-structure collision, a failure in a named material, or an error in a specific electronic-structure method. The benchmark target is a controlled smooth invariant, chosen to isolate representation error. Real-data prevalence and chemically weighted variants are separate questions; the displayed family has local codimension eleven and all exact collisions lie in the radial discriminant. This measure-zero statement does not bound the prevalence of finite-tolerance near collisions. A validated appendix-only scan of 20,000 untyped six-neighbor environments from an ethanol trajectory reranked 7,000 candidate pairs and found no pair with relative D_3 distance below 10^{-2} and relative orbit distance above 0.1. This negative audit is evidence that the strict event is not common in that sample, not evidence that the obstruction is chemically impossible. An additional exhaustive scan of 160 environments from each of the existing ethanol, aspirin, and uracil files checks 38,149 eligible pairs with the same thresholds and finds zero events. Three synthetic stress sets (128 environments each: two perturbed octahedral sets and random equal-radius clouds) similarly give zero events among 24,384 pairs (S13). Radial degeneracy is necessary, not sufficient; heterogeneous bond lengths alone do not guarantee a near-collision exclusion. These finite geometric audits are not molecular dynamics simulations. Force gaps are confounded by element and site identity, and frame energies are not local labels; we therefore make no real-material failure claim. The proof is computer assisted: finite enumeration and rational identities are machine checked, while the reduction and three geometric template arguments remain explicit mathematical lemmas requiring ordinary peer review. We do not benchmark physical energy/force fitting or MACE/GemNet here; the repair models are controlled representation tests, not proposed state-of-the-art force fields.

Within this scope, the conclusion is sharp. Complete third-order reflection-even pooling is sufficient for all continuous invariant learning problems on compact clouds through five points, and first loses that guarantee at six. The threshold converts a qualitative incompleteness observation into an exact boundary and provides a controlled test case for distinguishing information limits from optimization and data limitations.

Reproducibility statement. The supplementary package contains the standard-library exact checker, archived C++17 enumeration sources, thirteen rational certificates, collision data generator, CPU reference experiment, PyTorch DeepSets and invariant-MPNN implementations, exact species-channel ablation, a robust rMD17-style trajectory scanner, configurations, unit tests, and plotting scripts. The exact audit has one command; a source-to-output mode regenerates every five-vertex enumeration. Section S13 provides the new exhaustive scans, fixed-composition audit, stability test, raw results and reproduction commands.

AI use statement. AI-assisted tools were used for symbolic-search orchestration, code drafting, and language editing. Every mathematical certificate reported in the paper is verified by independent exact arithmetic, all generated experimental code is released, and the authors take full responsibility for the claims and text.

REFERENCES

Albert P. Bartók, Risi Kondor, and Gábor Csányi. On representing chemical environments. *Phys. Rev. B*, 87:184115, 2013. doi: 10.1103/PhysRevB.87.184115.

Ilyes Batatia, Dávid Péter Kovács, Gregor N. C. Simm, Christoph Ortner, and Gábor Csányi. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. In *Advances in Neural Information Processing Systems*, volume 35, pp. 11423–11436, 2022.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13:2453, 2022. doi: 10.1038/s41467-022-29939-5.

Mireille Boutin and Gregor Kemper. Lossless representation of graphs using distributions. *arXiv:0710.1870*, 2007.

Michelangelo Domina and Michele Ceriotti. Using large language models to probe the limits of atom-centered structural descriptors. *arXiv:2607.26984*, 2026.

Ralf Drautz. Atomic cluster expansion for accurate and transferable interatomic potentials. *Phys. Rev. B*, 99:014104, 2019. doi: 10.1103/PhysRevB.99.014104.

Geneviève Dusson, Markus Bachmayr, Gábor Csányi, Ralf Drautz, Simon Etter, Cas van der Oord, and Christoph Ortner. Atomic cluster expansion: Completeness, efficiency and stability. *J. Comput. Phys.*, 454:110946, 2022. doi: 10.1016/j.jcp.2022.110946.

Snir Hordan, Tal Amir, Steven J. Gortler, and Nadav Dym. Complete neural networks for complete euclidean graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 12482–12490, 2024. doi: 10.1609/aaai.v38i11.29141.

Zian Li, Xiyuan Wang, Shijia Kang, and Muhan Zhang. On the completeness of invariant geometric deep learning models. In *International Conference on Learning Representations*, 2025.

Jigyasa Nigam, Sergey N. Pozdnyakov, Kevin K. Huguenin-Dumittan, and Michele Ceriotti. Completeness of atomic structure representations. *APL Mach. Learn.*, 2:016110, 2024. doi: 10.1063/5.0160740.

Jigyasa Nigam, Tuong Phung, Ameya Daigavane, Aria Mansouri Tehrani, and Tess Smidt. Reconstructing local environments from concise atomistic representations. *arXiv:2607.20338*, 2026.

Sergey N. Pozdnyakov, Michael J. Willatt, Albert P. Bartók, Christoph Ortner, Gábor Csányi, and Michele Ceriotti. Incompleteness of atomic structure representations. *Phys. Rev. Lett.*, 125:166001, 2020. doi: 10.1103/PhysRevLett.125.166001.

Yonatan Sverdlov and Nadav Dym. On the expressive power of sparse geometric MPNNs. In *International Conference on Learning Representations*, 2025.

A. P. Thompson, L. P. Swiler, C. R. Trott, S. M. Foiles, and G. J. Tucker. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *J. Comput. Phys.*, 285:316–330, 2015. doi: 10.1016/j.jcp.2014.12.018.

Supplementary Material for "A Sharp Six-Neighbor Expressivity Threshold for Third-Order Geometric Invariants"

Anonymous authors

This supplement supplies the complete definitions, the five-neighbor classification, the paper proofs for the geometric templates, exact certificate statements for the algebraic templates, all controlled-learning experiments, the species audit, the real-trajectory diagnostic, and the reproducibility boundary. The result is computer assisted: finite enumerations and all thirteen rational polynomial identities are checked by the supplied code and retained as a complete formal evidence chain. In parallel, the radius-partition reconstruction, Schur-complement eliminations, and geometric proofs are written below for independent human review. All empirical values are tied to explicit JSON summaries; mean error bars are population standard deviations over seeds.

S1 Setting and the correlation–deck equivalence

Let $X = \{x_1, \dots, x_N\} \subset \mathbb{R}^3$ be a set of distinct neighbor vectors relative to a central atom. The environment is single species and unit weight. Two environments are equivalent when an orthogonal map and a permutation relate them. Put

$$\mu_X = \sum_{i=1}^{N} \delta_{x_i}, \qquad C_3(X) = \int_{O(3)} (Q_{\#}\mu_X)^{\otimes 3} \, \mathrm{d}H(Q). \tag{S1}$$

The Gram matrix $G_X = (x_i \cdot x_j)$ is PSD and has rank at most three. Conversely every PSD matrix of rank at most three is the Gram matrix of vectors in $\mathbb{R}^3$, and two such vector lists are related by $O(3)$ exactly when their Gram matrices agree. Permuting the vectors simultaneously permutes the rows and columns.

For $I \subseteq [N]$, write $G_X[I]$ for the principal submatrix. The distinct-index three-deck is

$$\mathscr{D}_3^{\neq}(X) = \left\{ \left\{ [G_X[I]]_{S_3} : I \in \binom{[N]}{3} \right\} \right\}. \tag{S2}$$

The complete deck $\mathscr{D}_3(X)$ also contains repeated-index cards, equivalently the one- and two-index decks. Multiplicities are always retained.

Lemma 1 (Correlation–deck equivalence)**.** *For fixed occupancy, $C_3(X) = C_3(Y)$ if and only if $\mathscr{D}_3(X) = \mathscr{D}_3(Y)$.*

Proof. For an ordered triple $z = (z_1, z_2, z_3)$, let $\Phi(z) = (z_a \cdot z_b)_{a,b=1}^3$. Two triples lie on the same $O(3)$-orbit if and only if their Gram matrices are equal. The pushforward of the correlation is the finite atomic measure

$$\Phi_{\#} C_3(X) = \sum_{(i,j,k) \in [N]^3} \delta_{\mathrm{Gram}(x_i, x_j, x_k)}. \tag{S3}$$

Equality of correlations therefore gives equality of ordered Gram multisets. Conversely, each Gram atom specifies a unique compact $O(3)$-orbit, and Haar averaging gives the unique invariant probability measure on that orbit. Equality of the atomic Gram measures reconstructs the same sum of orbit measures and hence the same correlation. Finally, quotienting by the permutation of tensor slots turns the ordered multiset into the principal three-cards plus the repeated-index lower decks. □

S1.1 Weak topology, metric continuity, and factorization

Fix $R, \eta > 0$ and define the compact labelled configuration space

$$\widetilde{\mathscr{X}_N}(R, \eta) = \left\{ (x_i)_{i=1}^N : \ \eta \le \|x_i\| \le R, \quad \|x_i - x_j\| \ge \eta \ (i \ne j) \right\}. \tag{S4}$$

Let $\mathscr{X}_N(R,\eta)$ be its quotient by the compact group $O(3)\times S_N$. On the quotient put

$$d_{\mathrm{orb}}([X],[Y]) = \min_{Q\in O(3),\,\pi\in S_N} \left(\frac{1}{N}\sum_{i=1}^{N} \|x_i - Qy_{\pi(i)}\|^2\right)^{1/2}. \tag{S5}$$

The minimum is attained. The usual compact-group quotient argument shows that Eq. (S5) is a metric and induces the quotient topology.

Normalize the empirical measure and the correlation by

$$\nu_X = \frac{1}{N}\sum_{i=1}^{N}\delta_{x_i}, \qquad \overline{C}_3(X) = \int_{O(3)} (Q_{\#}\nu_X)^{\otimes 3}\, dH(Q) = N^{-3}C_3(X). \tag{S6}$$

Thus $\overline{C}_3(X)$ is a probability measure on $K_R = (\overline{B}_R)^3$. Give K_R the ground metric

$$\rho_3(z,z') = \frac{1}{3}\sum_{a=1}^{3}\|z_a - z'_a\|. \tag{S7}$$

On $\mathscr{P}(K_R)$ use the weak topology. Because K_R is compact, this topology is metrized both by the 1-Wasserstein distance W_1 and by the bounded-Lipschitz distance

$$d_{\mathrm{BL}}(\mu,\nu) = \sup_{\|f\|_\infty\le 1,\ \mathrm{Lip}(f)\le 1} \left|\int f\,d\mu - \int f\,d\nu\right|. \tag{S8}$$

Proposition 1 (Lipschitz continuity of the normalized correlation)**.** *For all $[X],[Y]\in\mathscr{X}_N(R,\eta)$,*

$$W_1(\overline{C}_3(X),\overline{C}_3(Y)) \le d_{\mathrm{orb}}([X],[Y]). \tag{S9}$$

In particular, $\overline{C}_3 : \mathscr{X}_N(R,\eta)\to\mathscr{P}(K_R)$ is continuous for the weak, W_1, and bounded-Lipschitz topologies.

Proof. Fix Q and π. Sample i,j,k independently and uniformly from $[N]$ and sample U from normalized Haar measure. Couple the two correlations by

$$(Ux_i, Ux_j, Ux_k) \quad\text{and}\quad (UQy_{\pi(i)}, UQy_{\pi(j)}, UQy_{\pi(k)}). \tag{S10}$$

The expected ρ_3 cost equals $N^{-1}\sum_i \|x_i - Qy_{\pi(i)}\|$, which is bounded by the corresponding root-mean-square cost. Minimize over Q,π to obtain Eq. (S9). The final statement follows because W_1 and d_{BL} metrize weak convergence on compact support. □

Corollary 1 (Continuous factorization and universality)**.** *Suppose $N\le 5$. Every continuous invariant $V : \mathscr{X}_N(R,\eta)\to\mathbb{R}$ has a unique factorization*

$$V = f\circ\overline{C}_3, \qquad f\in C\big(\overline{C}_3(\mathscr{X}_N(R,\eta))\big). \tag{S11}$$

If $\mathscr{A}$ is uniformly dense in the continuous functions on the compact descriptor image, then $\{a\circ\overline{C}_3 : a\in\mathscr{A}\}$ is uniformly dense in $C(\mathscr{X}_N(R,\eta))$.

Proof. The reconstruction result for $N\le 5$ proved in Sections S3–S5 makes $\overline{C}_3$ injective. Proposition 1 makes it continuous, and a continuous bijection from a compact space to its Hausdorff image is a homeomorphism. Hence $f = V\circ\overline{C}_3^{-1}$ is the unique continuous factor on the image. Uniform density there gives the last assertion after composition. □

The use of $O(3)$ is essential. Parity-odd $SO(3)$ channels are not functions of Gram matrices and are not considered here.

S2 Exact six-neighbor collision

Define

$$u=(-2,-1,0), \qquad v=(-2,0,-1), \qquad w=(-1,0,2), s \qquad =(0,-1,-2), t \qquad =(0,-1,2),$$

and

$$X=\{u,v,w,s,-s,-u\}, \qquad Y=\{u,v,w,t,-t,-u\}. \tag{S12}$$

In the displayed order their Gram matrices are

$$G_X = \begin{pmatrix} 5 & 4 & 2 & 1 & -1 & -5 \\ 4 & 5 & 0 & 2 & -2 & -4 \\ 2 & 0 & 5 & -4 & 4 & -2 \\ 1 & 2 & -4 & 5 & -5 & -1 \\ -1 & -2 & 4 & -5 & 5 & 1 \\ -5 & -4 & -2 & -1 & 1 & 5 \end{pmatrix}, \tag{S13}$$

$$G_Y = \begin{pmatrix} 5 & 4 & 2 & 1 & -1 & -5 \\ 4 & 5 & 0 & -2 & 2 & -4 \\ 2 & 0 & 5 & 4 & -4 & -2 \\ 1 & -2 & 4 & 5 & -5 & -1 \\ -1 & 2 & -4 & -5 & 5 & 1 \\ -5 & -4 & -2 & -1 & 1 & 5 \end{pmatrix}. \tag{S14}$$

The checker constructs every one-, two-, and three-index Gram card using integer arithmetic, canonically sorts within a card and across the multiset, and finds equality including multiplicities. It then tests all 720 simultaneous row-column permutations and finds none mapping G_X to G_Y. This is an exact proof of deck equality and noncongruence for this pair.

For context, the pair is a member of a parametric family from Ref. [1]. Choose independent u, v, w satisfying

$$u^2 = h, \quad v^2 = w^2 = e, \quad u \cdot v = a, \quad u \cdot w = b, \quad v \cdot w = 0, \tag{S15}$$

where $h, e > 0$ and $\Delta_0 = he - a^2 - b^2 > 0$. Set $c = \Delta_0/e$, and define s, t by

$$\begin{aligned} u \cdot s &= c, \quad v \cdot s = b, \quad w \cdot s = -a, \\ u \cdot t &= c, \quad v \cdot t = -b, \quad w \cdot t = a. \end{aligned} \tag{S16}$$

Then $s^2 = t^2 = h$, and the two sets $\{u, v, w, s, -s, -u\}$ and $\{u, v, w, t, -t, -u\}$ have the same symbolic three-deck. The integer representative proves that noncongruence holds away from a proper algebraic exceptional set.

S2.1 The radial discriminant and ambient measure zero

For a labelled Gram matrix define

$$\Delta_{\mathrm{rad}}(G) = \prod_{1 \le i < j \le 6} (G_{ii} - G_{jj}). \tag{S17}$$

Although Δ_{rad} changes sign under some vertex permutations, its square is a well-defined polynomial on the permutation quotient.

Proposition 2 (Every exact collision is radially singular)**.** *If two six-neighbor clouds have the same complete three-deck and are not equivalent under $O(3) \times S_6$, then*

$$\Delta_{\mathrm{rad}}(G_X) = \Delta_{\mathrm{rad}}(G_Y) = 0. \tag{S18}$$

Consequently, exact collisions form a set of zero 15*-dimensional Hausdorff measure in the regular rank-three six-neighbor Gram quotient.*

Proof. Suppose first that the six diagonal entries of G_X are distinct. Deck equality gives the same multiset of diagonal values for G_Y, hence a unique global matching of vertices by their squared radii. Within every three-card, the diagonal values then label its vertices without ambiguity. For every pair i, j, choose $k \notin \{i, j\}$; the card on $\{i, j, k\}$ identifies and recovers G_{ij}. Thus all entries of the two Gram matrices agree after the radius matching, contradicting noncongruence. The same argument with X, Y interchanged proves the first assertion.

The manifold of rank-three PSD Gram matrices of order six has dimension

$$6 \cdot 3 - \dim O(3) = 18 - 3 = 15. \tag{S19}$$

Quotienting by the finite group S_6 does not change this local dimension. In Gram-factor or Cholesky gauge charts, Δ_{rad}^2 is a nonzero real analytic function: independent small rescalings of the six nonzero vectors make their squared radii distinct while preserving rank three and, in the interior of the separated-cloud domain, all strict inequalities. Its zero set therefore has zero 15-dimensional measure in every chart. The lower-rank stratum has dimension at most $6 \cdot 2 - \dim O(2) = 11$ and also has zero 15-dimensional measure. A finite quotient preserves null sets, completing the proof. □

This is an ambient measure statement, not a claim that the full collision set is a pure-dimensional algebraic variety. In particular, it does not by itself determine the volume of an ε-near-collision tube.

S2.2 Exact local Jacobian rank of the four-parameter family

We now compute the local dimension of the displayed family at the rational audit point

$$(h,e,a,b) = (5,3,1,2), \qquad c = \frac{10}{3}. \tag{S20}$$

The 3×3 principal block on u,v,w has determinant $30 > 0$, the minimum squared inter-neighbor distance on both sides is $10/3$, the stabilizer of G_X in S_6 is trivial, and none of the 720 permutations conjugates G_X to G_Y. All statements use exact rational arithmetic.

Let $\mathscr{M}_3$ denote the smooth manifold of rank-three PSD 6×6 matrices near this point. If K is a 6×3 kernel matrix for G, its tangent space is

$$T_G\mathscr{M}_3 = \{H \in \mathrm{Sym}_6 : K^T H K = 0\}. \tag{S21}$$

At Eq. (S20), exact kernel bases are

$$K_X = \begin{pmatrix} -1 & 1 & 1 \\ -1/3 & 1/3 & 0 \\ 1 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad K_Y = \begin{pmatrix} -1 & 1 & 1 \\ 1 & -1 & 0 \\ 1/3 & -1/3 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{S22}$$

Thus the two rank constraints contribute twelve independent linear equations in the 42 upper-triangular entries of a pair (H_X, H_Y).

There are 15 distinct unlabelled three-card values at the audit point. Their nonunit multiplicities are $4,2,2$, and every compatible pair of cards has a unique internal vertex matching. Hence there are exactly

$$4!\,2!\,2! = 96 \tag{S23}$$

compatible local deck matchings. For each matching m, let $E_m(H_X,H_Y) = 0$ be the 120 linearized card-entry equalities. Define the combined linear map

$$L_m(H_X,H_Y) = \big(K_X^T H_X K_X,\ K_Y^T H_Y K_Y,\ E_m(H_X,H_Y)\big). \tag{S24}$$

Proposition 3 (Local four-dimensional paired-collision branch)**.** *For every one of the 96 matchings,*

$$\operatorname{rank} L_m = 38. \tag{S25}$$

Consequently, the deck-equality map restricted to $\mathscr{M}_3\times\mathscr{M}_3$ has differential rank $38-12=26$, and its local zero set has dimension $30-26=4$. The four-parameter family is exactly this local paired-collision germ. Its projection to either regular 15-dimensional environment quotient is a smooth local branch of codimension eleven.

Proof. The script `verify_local_jacobian.py` uses only `fractions.Fraction`. It constructs all 96 matchings, the twelve rows in Eq. (S21), and the card equations, and performs exact Gaussian elimination. It obtains rank 38 for every matching. Under the row and column convention printed by the checker, one representative 38×38 minor equals

$$-\frac{1280}{81} \neq 0. \tag{S26}$$

This is a rational rank certificate, not a floating-point estimate.

The rank-three manifold has dimension 15, so the paired manifold has dimension 30. Linear algebra gives

$$\operatorname{rank}\big(E_m|_{T_{G_X}\mathscr{M}_3\times T_{G_Y}\mathscr{M}_3}\big) = 38-12 = 26. \tag{S27}$$

The implicit-function theorem therefore gives a smooth four-dimensional zero set for each compatible matching. The symbolic deck checker verifies that the displayed family satisfies every matching among the persistent multiplicity classes. Its parameter Jacobian has rank four because the entries $(G_{00},G_{11},G_{01},G_{02})$ are exactly (h,e,a,b). A same-dimensional embedded submanifold contained in each local zero set is open in it, so all 96 germs coincide with the family germ. Incompatible matchings have positive card discrepancy at the audit point and cannot enter a sufficiently small neighborhood.

Finally, noncongruence persists because the minimum over the finite set of permutation-conjugacy discrepancies is positive at the audit point. The trivial stabilizer makes the quotient locally a manifold, and projection to G_X retains rank four through the entries (h,e,a,b). This proves the codimension statement. □

The conclusion is deliberately local at the paired point (G_X,G_Y). It does not exclude other components elsewhere or assert that the full collision locus has pure codimension eleven.

S2.3 Perturbations and inverse conditioning

For unlabelled three-cards define

$$d_{\mathrm{card}}([A],[B]) = \min_{P\in S_3} \|A - PBP^T\|_F. \tag{S28}$$

Writing $M = \binom{N}{3}$, define the root-mean-square optimal-assignment distance

$$d_{D_3}(X,Y) = \min_{\tau\in S_M} \left(\frac{1}{M} \sum_{I\in\binom{[N]}{3}} d_{\mathrm{card}}\big([G_X[I]],[G_Y[\tau(I)]]\big)^2 \right)^{1/2}. \tag{S29}$$

This is the idealized card-set metric used in the conditioning experiment. The analogous ordered-atom W_1 metric follows by replacing the finite root-mean-square assignment by the uniform empirical Gram measure; the proof below is unchanged.

Proposition 4 (First-order splitting of an exact collision)**.** *Let (X_0,Y_0) be a noncongruent exact collision with $\delta_0 = d_{\mathrm{orb}}(X_0,Y_0) > 0$. Suppose*

$$x_i(\sigma) = x_i + \sigma\xi_i + o(\sigma), \qquad y_i(\sigma) = y_i + \sigma\zeta_i + o(\sigma). \tag{S30}$$

Then

$$d_{D_3}(X_\sigma,Y_\sigma) = O(\sigma), \tag{S31}$$

$$d_{\mathrm{orb}}(X_\sigma,Y_\sigma) \geq \delta_0 - \sigma(\|\xi\|_N + \|\zeta\|_N) - o(\sigma), \tag{S32}$$

where $\|\xi\|_N = (N^{-1}\sum_i \|\xi_i\|^2)^{1/2}$. Hence the inverse secant conditioning d_{orb}/d_{D_3} is $\Omega(\sigma^{-1})$.

Moreover, outside a finite union of proper linear subspaces of directions (ξ,ζ),

$$d_{D_3}(X_\sigma,Y_\sigma) = \Theta(\sigma), \qquad \frac{d_{\mathrm{orb}}(X_\sigma,Y_\sigma)}{d_{D_3}(X_\sigma,Y_\sigma)} = \Theta(\sigma^{-1}). \tag{S33}$$

Proof. Choose any card matching and internal permutations that realize the exact deck equality at $\sigma = 0$. Each Gram entry is a smooth quadratic function of the coordinates, so this fixed matching has cost $O(\sigma)$, proving Eq. (S31). The triangle inequality for d_{orb} gives

$$d_{\mathrm{orb}}(X_\sigma,Y_\sigma) \geq d_{\mathrm{orb}}(X_0,Y_0) - d_{\mathrm{orb}}(X_\sigma,X_0) - d_{\mathrm{orb}}(Y_\sigma,Y_0), \tag{S34}$$

and the identity matchings on the last two terms give Eq. (S32).

Let $\mathfrak{M}$ be the finite set of card and internal matchings compatible at the base point. Matchings outside $\mathfrak{M}$ have a positive zeroth-order cost and cannot be optimal for small σ. If $\dot\Gamma^X_I, \dot\Gamma^Y_J$ are the first derivatives of the matched Gram cards, the finite minimum expansion is

$$d_{D_3}(X_\sigma,Y_\sigma) = \sigma c(\xi,\zeta) + o(\sigma), \tag{S35}$$

where

$$c(\xi,\zeta) = \min_{(\tau,(P_I))\in\mathfrak{M}} \left(\frac{1}{M}\sum_I \|\dot\Gamma^X_I - P_I \dot\Gamma^Y_{\tau(I)} P_I^T\|_F^2 \right)^{1/2}. \tag{S36}$$

Thus $c = 0$ is a finite union of linear kernels. Each kernel is proper: take $\zeta_i = 0$ and $\xi_i = x_i$. Every X card then has a nonzero diagonal Gram derivative because the neighbors are nonzero, while every Y derivative vanishes. Hence no compatible matching annihilates this direction. The complement of the exceptional union is open, dense, and of full Lebesgue measure in direction space; there $c > 0$, proving Eq. (S33). □

S3 Combinatorial reduction for five neighbors

Fix the edge order

$$01, 02, 03, 04, 12, 13, 14, 23, 24, 34. \tag{S37}$$

A restricted-growth string of length ten represents a set partition of these edges, hence an equality pattern of the off-diagonal Gram entries. Diagonal colors represent equality of squared radii.

Equality of numerical three-decks implies equality of the corresponding colored three-decks. In the enumeration we also quotient a deck by one global renaming of its edge colors. This deliberately enlarges the candidate set: a numerical equality must use the same actual inner-product values in the two matrices, whereas the abstract search initially forgets their names. Thus eliminating every enumerated abstract candidate is sufficient.

Every set partition of the ten edges is generated exactly once as a restricted-growth string, so all

$$B_{10} = 115{,}975 \tag{S38}$$

patterns are covered for each radius partition. A graph is canonicalized under the subgroup of S_5 preserving its diagonal colors. A three-card is canonicalized under its S_3 vertex action, and a deck under card permutation and a single global edge-color renaming.

Table S1: Complete five-vertex equality-pattern enumeration.

Radius multiplicities	Canonical graphs	Collision groups	Pairs
(5)	1,299	46	80
$4+1$	5,525	43	43
$3+2$	10,636	0	0
$3+1+1$	20,413	0	0
$2+2+1$	30,303	0	0
$2+1+1+1$	59,047	0	0
$1+1+1+1+1$	direct reconstruction	0	0

We retain this exhaustive enumeration as one complete proof chain. We now give, in parallel, a direct reconstruction proof for every radius partition other than (5) and $4+1$. The following elementary gluing lemmas explicitly include repeated edge values.

Lemma 2 (Two-record gluing)**.** *Let $r_1, r_2 \in S_1 \times \cdots \times S_m$ with $m \geq 2$. The unordered pair $\{r_1, r_2\}$ is determined by all of its two-coordinate projections.*

Proof. If the two records agree in every coordinate there is nothing to prove. Otherwise choose a coordinate q on which they differ. In the projection to coordinates (q,k), the two distinct q-values uniquely attach the two k-values to the corresponding records. This recovers every coordinate; equal values cause no ambiguity because exchanging equal entries changes nothing. □

Lemma 3 (Marked 2×2 reconstruction)**.** *A 2×2 array is determined up to row and column permutations by the multiset of its two unordered rows and the multiset of its two unordered columns. The conclusion remains true when rows and columns carry additional marks, with permutations restricted to preserve those marks.*

Proof. Fix the order of the first row. Modulo a global column swap, the second row has only two possible orientations. They give column-pair multisets

$$\{\{\{a,c\},\{b,d\}\}\} \quad \text{and} \quad \{\{\{a,d\},\{b,c\}\}\}. \tag{S39}$$

If these are equal, direct comparison of the two possible pair matchings gives $a=b$ or $c=d$ (or an equality making the same row/column swap explicit), in which case the two arrays are already equivalent. Marks only remove allowed permutations and simultaneously identify the corresponding row or column; they cannot create a new ambiguity. □

Lemma 4 (Two anchors on a decorated triangle)**.** *Let three indistinguishable vertices have signatures $s_i = (u_i, v_i)$ and internal edge labels w_{ij}. The data*

$$\{\{(u_i,v_i) : i=1,2,3\}\}, \qquad \{\{(w_{ij},\{u_i,u_j\}) : i<j\}\}, \qquad \{\{(w_{ij},\{v_i,v_j\}) : i<j\}\} \tag{S40}$$

determine the decorated triangle up to a permutation of its vertices.

Proof. If all signatures coincide, S_3 induces every permutation of the three edges, so only the multiset of w_{ij} matters. If the signature multiplicities are $2+1$, one of the coordinates separates the singleton; it distinguishes the edge within the doublet from the two cross edges, which are exchanged by the residual transposition. If all signatures are distinct and one coordinate has three distinct values, its endpoint pairs identify every edge. If a coordinate has type $2+1$, it identifies the edge inside the doublet; the other coordinate distinguishes the two cross edges because the two vertices in the doublet have distinct full signatures. If one coordinate is constant, distinctness of the full signatures makes the other coordinate distinct. These cases exhaust all equality patterns and prove the claim. □

Lemma 5 (Decorated two-sheeted triangle)**.** *Let $U = (u_{i\alpha})$ be a 3×2 array, considered modulo row permutations and one global column swap, and let w_{ij} label the three row edges. The following data determine (U, w) up to that action:*

$$\{\{\{u_{i1}, u_{i2}\} : i = 1,2,3\}\}, \tag{S41}$$

$$\big\{\big\{(w_{ij}, \{u_{i\alpha}, u_{j\alpha}\}) : i < j,\ \alpha = 1,2\big\}\big\}. \tag{S42}$$

Proof. First align the three unordered rows and the edge-decorated base triangle. For three distinct edge labels, Eq. (S42) groups the two sheets edge by edge. For edge-label multiplicity $2+1$, the unique edge is fixed and the other two are exchanged by the stabilizer of the edge-colored triangle. When all edge labels agree, every permutation of the three base edges is induced by a row permutation. In each case, comparing the union of the two cards on an edge with the two endpoint row-multisets fixes the remaining attachment; if two such unions coincide, the corresponding row exchange is already an allowed stabilizer. Thus any residual ambiguity is only an orientation choice within the three aligned unordered rows.

Modulo a global column swap, two lifts of the same unordered rows differ by flipping one row, say (x, y). Cancel from Eq. (S42) the two unchanged cards on the edge not incident to that row. If the two incident edge labels differ, the remaining cards can be compared edge by edge; Lemma 3 says that equality requires $x = y$, or both of the other rows to be constant. If the incident labels are equal, direct matching of the four remaining unordered pairs gives exactly three possibilities:

(i) $x = y$;
(ii) both other rows are constant;
(iii) the other two ordered rows are reverses of one another.

In (i) the flip changes nothing; in (ii) it is a global column swap; in (iii) it is a global column swap followed by exchanging the other two rows. Hence all lifts with the same data are equivalent. □

Proposition 5 (Hand reconstruction of the nonambiguous radius splits)**.** *A numerically edge-labelled K_5 is reconstructed from its radius-coloured three-deck whenever its radius multiplicities are*

$$1+1+1+1+1, \quad 2+1+1+1, \quad 3+1+1, \quad 2+2+1, \quad 3+2. \tag{S43}$$

Thus only (5) and $4+1$ can reach the five maximal ambiguity templates.

Proof. With five distinct radii, the diagonal values name the vertices inside every card and every edge is read directly. For $2+1+1+1$, the three singleton radii are anchors: cards containing the repeated-radius vertex and two anchors give all two-coordinate projections of two three-entry incidence records, so Lemma 2 glues them; cards containing both repeated-radius vertices recover their internal edge.

For $3+1+1$, cards containing one repeated-radius vertex and both singleton anchors give its signature (u_i, v_i). The two families of cards containing two repeated-radius vertices and one anchor give precisely Eq. (S40), so Lemma 4 reconstructs the internal triangle. For $2+2+1$, the four cards containing one vertex from each doublet and the singleton give a marked 2×2 cross array; the $2+1$ cards give its row and column pairs. Lemma 3 reconstructs it, and the two within-doublet edges are read directly.

Finally, for $3+2$, the cards containing one vertex of the triple and both vertices of the doublet give the three unordered rows in Eq. (S41). Cards containing two vertices of the triple and one vertex of the doublet give Eq. (S42), with the internal triple edge as decoration. Lemma 5 reconstructs the cross array and the decorated internal triangle; the doublet edge is common to and read from the first family of cards. This covers all listed partitions. □

The independent checker `verify_radius_reconstruction_lemmas.py` enumerates all equality patterns in the finite cores of these gluing lemmas: 15 marked 2×2 patterns and 21,147 nine-position patterns for each of Lemmas 4 and 5. The original B_{10} enumeration remains unchanged and independently proves the same radius-split classification.

Every one of the 80 aligned equal-radius pairs is a color coarsening of one of four maximal pairs, and every one of the 43 radius-$4+1$ pairs is a coarsening of a fifth. In the edge order (S37) these are Here a coarsening means applying one shared map from template colors to output colors, after allowed vertex permutations on the two sides. The coverage counts overlap: M1, M2, M3, and M4 cover 51, 14, 19, and 8 aligned pairs, respectively; R41 covers all 43 of its pairs.

The count 80 retains the common numerical color labels across a pair. If a pair is further quotiented by a simultaneous global renaming of these common values, the same candidates form 52 abstract pair orbits. The $4+1$ count is 43 under both conventions. The audit checks the larger aligned collection.

For $N = 4$, the analogous exhaustive calculation covers all $B_6 = 203$ edge patterns for radius partitions $4, 3+1, 2+2, 2+1+1, 1+1+1+1$ and finds no collision. The cases $N \le 3$ are immediate because the full Gram matrix is already a card. It remains only to prove that M1–M4 and R41 have no nonisomorphic PSD rank-three realization.

Table S2: Five maximal ambiguous colored-deck templates.

Template	First graph	Second graph
M1	0011234562	0011235462
M2	0011023240	0013014220
M3	0123234401	0123241340
M4	0011203212	0012113202
R41	0123245653	0125243635

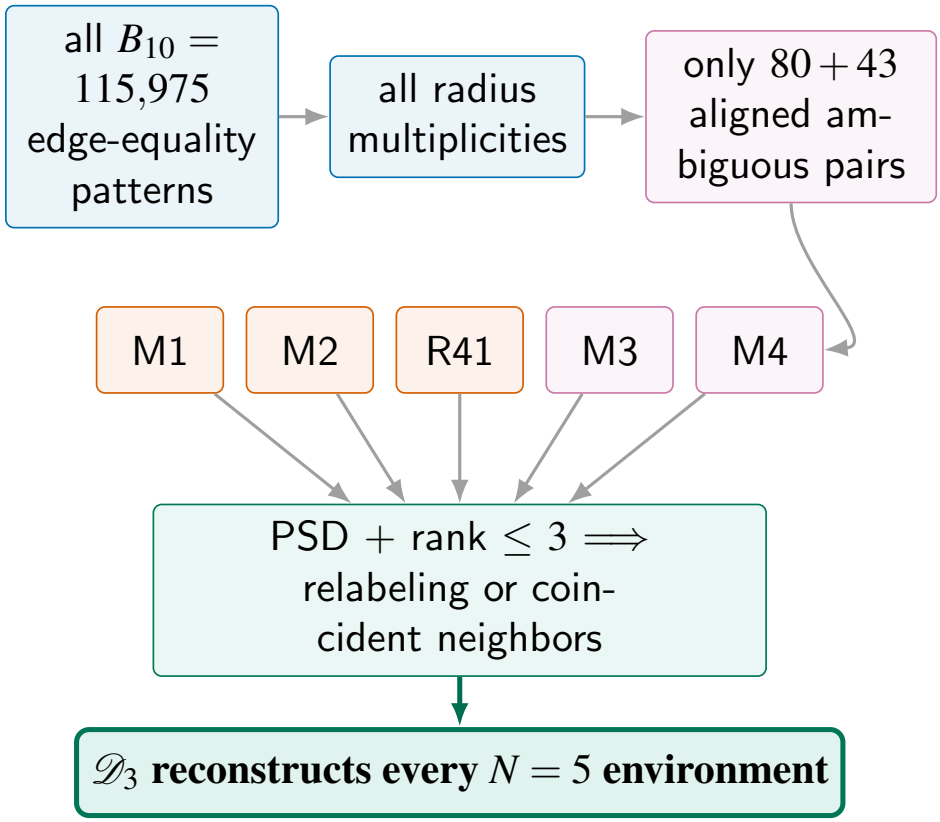


Figure S1: Auditable proof architecture. The finite combinatorial search and rational identities remain machine checked. In parallel, the correlation–deck reduction, radius-split reconstruction, five geometric templates, and the added M3/M4 hand eliminations are written for ordinary mathematical review.

S4 Geometric elimination of M1, M2, and R41

For equal-radius templates, scale the common diagonal to one. None of the arguments assumes that different template colors take different numerical values, so all coarsenings are covered.

S4.1 Template M1

Write its seven edge values as a,b,c,d,e,f,g. Relative to the common unit vector v_0, vertices $1,2$ have dot product a with v_0 and mutual dot product c; vertices $3,4$ have dot product b with v_0 and the same mutual dot product c. Projection onto $v_0^{\perp}$ gives within-pair Gram matrices

$$U = \begin{pmatrix} 1-a^2 & c-a^2 \\ c-a^2 & 1-a^2 \end{pmatrix}, \qquad V = \begin{pmatrix} 1-b^2 & c-b^2 \\ c-b^2 & 1-b^2 \end{pmatrix}. \tag{S44}$$

The residual cross matrix for the first graph is

$$M = \begin{pmatrix} d-ab & e-ab \\ f-ab & g-ab \end{pmatrix}, \tag{S45}$$

while the second graph uses M^T. Let $S = 2^{-1/2}\left(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right)$. The Hadamard basis diagonalizes both U and V. If $\widetilde{M} = S^T M S$, then

$$\widetilde{M}_{12} = \tfrac{1}{2}[(d-g)+(f-e)], \qquad \widetilde{M}_{21} = \tfrac{1}{2}[(d-g)-(f-e)]. \tag{S46}$$

When U,V are nonsingular, rank at most two of the two projected four-vector Gram matrices gives

$$V = M^T U^{-1} M = M U^{-1} M^T. \tag{S47}$$

Comparison of the first diagonal component in the Hadamard basis yields $\widetilde{M}_{12}^2 = \widetilde{M}_{21}^2$, hence

$$(d-g)(e-f) = 0. \tag{S48}$$

Either factor provides an explicit vertex relabeling.

If U is singular, its antisymmetric eigenvalue is $1-c$. Its vanishing makes two unit neighbors coincide. Otherwise the symmetric eigenvalue vanishes and $\mathbf{1} \in \ker U$. PSD block structure then implies both $M^T\mathbf{1} = 0$ and $M\mathbf{1} = 0$, giving $e = f$ and $d = g$. The same argument holds for singular V. Thus M1 relabels or contains repeated neighbors.

S4.2 Template M2

Both graphs contain the equilateral base (0,1,2), with Gram matrix

$$H = (1-a)I_3 + a\mathbf{1}\mathbf{1}^T. \tag{S49}$$

The dot vectors of the remaining vertices are

$$p = (b,c,c), \quad q = (b,d,e), \qquad p' = (b,b,c), \quad q' = (d,e,c), \tag{S50}$$

and their mutual dot product is a. For $-1/2 < a < 1$,

$$H^{-1} = \frac{I_3}{1-a} - \frac{a}{(1-a)(1+2a)}\mathbf{1}\mathbf{1}^T. \tag{S51}$$

Subtracting the unit-norm equations for p, p' gives

$$\frac{c^2-b^2}{1+2a} = 0. \tag{S52}$$

The branch $c = b$ relabels. In the branch $c = -b$, subtracting the two mutual-dot equations gives $d+e = 0$. The norm and mutual-dot equations reduce to

$$b^2(3+5a) = 1+a-2a^2, \qquad b^2(1+3a) = a(1+a-2a^2). \tag{S53}$$

The positive-definite interval excludes the vanishing common factor and the inconsistent denominators. Eliminating b^2 gives $5a^2 = 1$, then $b^2 = d^2 = 1/5$. Hence all five edge values are $\pm 1/\sqrt{5}$, with $c = -b$ and $e = -d$. All eight sign choices are permutation-equivalent, which the exact relabeling checker confirms.

At $a = 1$, the three base vertices coincide. At $a = -1/2$, their sum is zero; range conditions for both p, p' force $b = c = 0$, an isomorphic coarsening. M2 gives no counterexample.

S4.3 The $4+1$ template R41

The common four-point base is

$$H = \begin{pmatrix} 1 & a & b & c \\ a & 1 & c & e \\ b & c & 1 & g \\ c & e & g & 1 \end{pmatrix}, \tag{S54}$$

and the unique-radius fifth point has dot vectors

$$p = (d, f, f, d), \qquad p' = (f, d, d, f). \tag{S55}$$

Put $\mathbf{1} = (1,1,1,1)$, $z = (1,-1,-1,1)$, $s = d+f$, and $t = d-f$. Then

$$p = \tfrac{s}{2}\mathbf{1} + \tfrac{t}{2}z, \qquad p' = \tfrac{s}{2}\mathbf{1} - \tfrac{t}{2}z. \tag{S56}$$

If $t = 0$, the matrices agree. If $s, t \neq 0$, PSD block structure puts $\mathbf{1}, z \in \operatorname{Range} H$. Since $\operatorname{rank} H \leq 3$, a kernel vector has the form $(x, y, -y, -x)$. Unless $c = 1$, which duplicates base points, both coordinates are nonzero. The kernel equation gives $a = g$ and $b = e$, each of which relabels the graphs.

It remains to take $s = 0$, so $p' = -p$ and $z \in \operatorname{Range} H$. Choose $w \neq 0$ whose dots with the four base vectors u_i are $(1,-1,-1,1)$, and define

$$x = u_0 - u_3, \quad y = u_1 - u_2, \quad r = u_0 + u_3, \quad q = u_1 + u_2. \tag{S57}$$

Then $w \cdot x = w \cdot y = 0$, $w \cdot r = 2$, $w \cdot q = -2$, and

$$\|x\|^2 = \|y\|^2 = 2(1-c), \qquad \|r\|^2 = \|q\|^2 = 2(1+c), \tag{S58}$$

with $x \perp r$ and $y \perp q$. After removing the equal-magnitude components parallel to w, the two orthogonal frames in the plane $w^\perp$ have equal corresponding lengths. They therefore have equal or opposite orientation, so $x \cdot q = \pm r \cdot y$. Since

$$a - g = \tfrac{1}{2}(r \cdot y + x \cdot q), \qquad b - e = \tfrac{1}{2}(x \cdot q - r \cdot y), \tag{S59}$$

we obtain $(a-g)(b-e) = 0$, and either factor relabels. R41 is harmless.

S5 Exact algebraic elimination of M3 and M4

This section states the certificate logic and every factorized target. A certificate file serializes an identity

$$T = \sum_{i,\alpha} q_{i,\alpha} x^\alpha F_i, \qquad q_{i,\alpha} \in \mathbb{Q}, \tag{S60}$$

where F_i are the stated rank equations. The independent verifier parses all coefficients with `fractions.Fraction`, expands both sides, and requires every coefficient of the residual to be exactly zero. Therefore every common zero of the generators over $\mathbb{C}$, and in particular every real PSD realization, satisfies $T = 0$.

S5.1 Template M3

With variables A, B, C, D, E, the two matrices are

$$G_A = \begin{pmatrix} 1 & A & B & C & D \\ A & 1 & C & D & E \\ B & C & 1 & E & A \\ C & D & E & 1 & B \\ D & E & A & B & 1 \end{pmatrix}, \tag{S61}$$

$$G_B = \begin{pmatrix} 1 & A & B & C & D \\ A & 1 & C & E & B \\ B & C & 1 & D & E \\ C & E & D & 1 & A \\ D & B & E & A & 1 \end{pmatrix}. \tag{S62}$$

If both have rank at most three, every principal four-minor vanishes. The nine distinct such minors generate an ideal I. The first certificate proves

$$(A-B)(A-D)(3A+1)(B-D)(C-E)\in I. \tag{S63}$$

The factors $A=B$ and $C=E$ directly relabel the matrices. On $A=D$ and $B=D$, additional certificates prove

$$(A-B)(A-E)(C-E)(C-A)=0, \tag{S64}$$
$$(B-A)(B-C)(B-E)(C-E)=0. \tag{S65}$$

These two branch relations also admit short hand-checkable principal-minor proofs. Let M_i^X and M_i^Y be the principal 4×4 determinants obtained by deleting vertex i from G_A and G_B, respectively. Direct determinant expansion gives, after setting $D=A$,

$$(A-B)(A-E)(C-E)(C-A)=\frac{1}{2}\left(M_0^X+M_3^X+M_4^X-M_0^Y-M_2^Y-M_4^Y\right), \tag{S66}$$

and, after setting $D=B$,

$$(B-A)(B-C)(B-E)(C-E)=\frac{1}{2}\left(M_0^X+M_1^X+M_4^X-M_0^Y-M_3^Y-M_4^Y\right). \tag{S67}$$

Thus Eqs. (S64)–(S65) follow immediately from the rank-three principal-minor conditions. The six-term rational certificates for the same two identities are nevertheless retained and checked as an independent proof route. All terminal factors are either previously handled or one of the explicit relabelings checked by `verify_branch_identities.py`.

Only $A=-1/3$ remains. Take the common principal block

$$H=\begin{pmatrix}1 & A & B\\ A & 1 & C\\ B & C & 1\end{pmatrix} \tag{S68}$$

and adjoining dot vectors

$$p=(C,D,E),\quad q=(D,E,A),\qquad p'=(C,E,D),\quad q'=(D,B,E). \tag{S69}$$

Let $h=\det H$ and $J=\operatorname{adj} H$. The six adjugate Schur equations are

$$\begin{aligned} &p^TJp-h=0, && q^TJq-h=0, && p^TJq-Bh=0,\\ &p'^TJp'-h=0, && q'^TJq'-h=0, && p'^TJq'-Ah=0. \end{aligned} \tag{S70}$$

Their exact difference contains the split

$$F_4-F_1=-\frac{3B+1}{9}(D-E)\left[(1-3B)(D+E)+6C(C+1)\right]. \tag{S71}$$

The factor $3B+1$ is $B=A$. On $D=E$, a certificate gives

$$(B-A)(D-A)(B-D)(C-D)=0. \tag{S72}$$

On

$$R=(1-3B)(D+E)+6C(C+1)=0, \tag{S73}$$

another certificate gives

$$\left[(B-A)(D-A)(B-D)(C-E)\right]^2=0. \tag{S74}$$

The $D=A$ and $B=D$ subbranches close by the specialized forms of Eqs. (S64) and (S65); the remaining factors directly relabel. Hence M3 has no nonisomorphic PSD rank-three realization.

No invertibility assumption on H is hidden here. For a PSD block matrix, an adjoining dot vector lies in $\operatorname{Range} H$. If H is singular of rank two, $\operatorname{adj} H$ maps into $\ker H$ and annihilates $\operatorname{Range} H$; at lower rank $\operatorname{adj} H=0$. Since also $h=0$, all equations (S70) remain necessary on the boundary.

S5.2 Template M4

Use the common base

$$H = \begin{pmatrix} 1 & A & B \\ A & 1 & C \\ B & C & 1 \end{pmatrix} \tag{S75}$$

and adjoining vectors

$$p = (A,C,A), \quad q = (B,B,C), \qquad p' = (A,B,B), \quad q' = (C,A,C), \tag{S76}$$

whose common mutual dot product is D. The six adjugate Schur equations, defined exactly as in Eq. (S70), generate the relevant ideal. The elimination certificate gives

$$(A-C)(A-1)^4(A+1)(3A+1)(5A^2-1) = 0. \tag{S77}$$

The factor $A = C$ relabels. At $A = 1$, unit neighbors coincide; at $A = -1$, two unit vectors equal the antipode of a third and hence coincide. At $A = -1/3$, two certificates force

$$B = A, \qquad C = A. \tag{S78}$$

Finally, after adjoining the generator $5A^2 - 1$, three certificates give

$$\begin{aligned} (A-C)(A+C) &= 0, \\ (A-C)(5B^2-1) &= 0, \\ (A-C)(D-B) &= 0. \end{aligned} \tag{S79}$$

Outside the already relabeled branch $A = C$, these imply $C = -A$, $B = \pm A$, and $D = B$. Thus $B = A$ or $B = C$, and both cases have explicit vertex permutations. M4 gives no counterexample.

S5.2.1 Parallel hand elimination for M4

The certificate chain above remains a complete exact proof. We now give an independent elementary elimination from the same six necessary Schur equations. Exact expansion first gives

$$F_5 - F_1 = (A-C)(A+C)\left((A-C)^2 - (1-B)^2\right). \tag{S80}$$

The branch $A = C$ relabels. On $C = -A$,

$$F_2 - F_4 = 4AB(A^2 - B^2). \tag{S81}$$

If $A = 0$, then $A = C$. If $B = 0$, the two necessary equations reduce to

$$F_1 = (1-A^2)(4A^2-1), \qquad F_2 = -1 + 3A^2 - A^4, \tag{S82}$$

which have no common root. Hence $B = sA$ for $s \in \{1,-1\}$. Substitution in $F_1 = 0$ gives

$$(1-A^2)(5A^2-1) = 0. \tag{S83}$$

The root $A^2 = 1$ makes two unit neighbors coincide: for $A = 1$ the edge with value A joins equal vectors, while for $A = -1$ the edge with $C = 1$ does so. Thus $A^2 = 1/5$. Modulo $5A^2 - 1$, the third equation is

$$F_3 = \frac{2}{5}(sA-1)(D-sA). \tag{S84}$$

Since $sA \neq 1$, this forces $D = sA = B$. The cases $B = A$ and $B = C$ have the explicit relabelings printed by `verify_branch_identities.py`.

It remains to consider

$$(A-C)^2 = (1-B)^2. \tag{S85}$$

Distinct unit neighbors imply $A < 1$ and $B < 1$. Put

$$R = 2(1-A) > 0, \qquad S = 2(1-B) > 0, \qquad t = R/S > 0. \tag{S86}$$

If $A-C=1-B$, then $C=A+B-1$ and exact substitution gives

$$F_1 = \frac{S^3}{16}\left[4-4t-12t^2+S(-1+4t^2+4t^3)\right], \tag{S87}$$

$$F_6 - F_3 = \frac{S^3}{16}\left[-4+4t+4t^2+S(1-2t^2-2t^3)\right]. \tag{S88}$$

Eliminating S between the bracketed linear equations yields $8t^5=0$, contrary to $t>0$.

If $A-C=-(1-B)$, then $C=A-B+1$ and

$$F_1 = \frac{S^3}{16}\left[S(4t^3-8t^2+4t-1)-12t^2+20t-4\right], \tag{S89}$$

$$F_6 - F_3 = \frac{S^3}{16}\left[S(2t^3-4t^2+2t-1)-4t^2+4t+4\right]. \tag{S90}$$

Eliminating S now gives $8(t-1)^5=0$. Hence $t=1$, and either displayed equation gives $S=4$. Therefore $A=B=-1$ and $C=1$, so two base neighbors coincide. All branches of Eq. (S80) have now been eliminated or explicitly relabeled. The identities in this hand proof are separately checked coefficient by coefficient in `verify_branch_identities.py`; none of the existing M4 rational certificates is removed.

S5.3 Certificate inventory

Table S3 lists every serialized certificate. "Terms" is the number of nonzero monomial–generator multipliers in Eq. (S60), not the number of terms after expansion.

Table S3: Exact rational ideal-membership certificates.

File	Target or branch	Terms
`m3_general_split_certificate.json`	Eq. (S63)	853
`m3_general_A_eq_D_certificate.json`	Eq. (S64) on $A=D$	6
`m3_general_B_eq_D_certificate.json`	Eq. (S65) on $B=D$	6
`m3_deq_certificate.json`	Eq. (S72) on $D=E$	119
`m3_d_eq_a_certificate.json`	Eq. (S64) on $D=A$	78
`m3_b_eq_d_certificate.json`	Eq. (S65) on $B=D$	87
`m3_relation_certificate.json`	Eq. (S74) with $R=0$	581
`m4_elimination_certificate.json`	Eq. (S77)	332
`m4_a_neg_third_B_certificate.json`	$B-A=0$ at $A=-1/3$	27
`m4_a_neg_third_C_certificate.json`	$C-A=0$ at $A=-1/3$	27
`m4_root_C_certificate.json`	$(A-C)(A+C)=0,\ 5A^2=1$	84
`m4_root_B_certificate.json`	$(A-C)(5B^2-1)=0,\ 5A^2=1$	82
`m4_root_D_certificate.json`	$(A-C)(D-B)=0,\ 5A^2=1$	195

Sections S3–S5 prove reconstruction for $N\le 5$. Section S2 proves failure at $N=6$, completing the sharp-threshold theorem.

S6 Consequences for descriptor-based models

Let D be any descriptor and let every prediction factor as $\widehat{V}=f\circ D$.

Proposition 6 (Energy error floor)**.** *If $D(X)=D(Y)$ and $\Delta=|V(X)-V(Y)|$, then*

$$\max\{|\widehat{V}(X)-V(X)|,|\widehat{V}(Y)-V(Y)|\}\ge\Delta/2. \tag{S91}$$

If the two targets have conditional probabilities p and $1-p$ at their shared descriptor value, the minimum conditional mean-squared error is $p(1-p)\Delta^2$.

Proof. Both predictions equal a common c. The pointwise bound is $\Delta \le |V(X)-c| + |c-V(Y)|$. Minimizing $p(c-V(X))^2 + (1-p)(c-V(Y))^2$ gives the conditional statement. □

For the exact pair in Section S2, let P_π be the permutation matrix of $\pi \in S_6$ and define

$$V_X(Z) = \prod_{\pi \in S_6} \|G_Z - P_\pi G_X P_\pi^T\|_F^2. \tag{S92}$$

This is a smooth $O(3) \times S_6$-invariant polynomial in the coordinates. It vanishes on the orbit of X and is positive on Y. Hence a model factoring through $\mathscr{D}_3$ is not universal for smooth invariant six-neighbor site energies.

Proposition 7 (Integrated force error)**.** *Let V and $\widehat{V} = f \circ D$ be differentiable along an absolutely continuous path $\gamma : [0,1] \to (\mathbb{R}^3)^N$ from X to Y. With $F = -\nabla V$ and $\widehat{F} = -\nabla \widehat{V}$,*

$$\int_0^1 \|\widehat{F} - F\| \, \|\gamma'\| \, \mathrm{d}t \ge \Delta. \tag{S93}$$

If $L(\gamma) > 0$, then

$$\operatorname*{ess\,sup}_{t \in [0,1]} \|\widehat{F} - F\| \ge \Delta / L(\gamma). \tag{S94}$$

Proof. The predicted endpoint energies agree because $D(X) = D(Y)$. The fundamental theorem for line integrals gives

$$\Delta = \left| \int_0^1 (\widehat{F} - F)(\gamma(t)) \cdot \gamma'(t) \, \mathrm{d}t \right|, \tag{S95}$$

and the two estimates follow from the triangle inequality and the definition of path length. □

These propositions are conditional on a target with $\Delta > 0$. They prove a function-class limitation and an unavoidable error when the target separates the pair. They do not establish that a named material, electronic-structure method, or phase transition realizes such a gap.

S7 Descriptor interpretation and scope

The ideal object in the theorem has complete radial and angular resolution. Conventional SOAP power spectra are lower-order correlations and may fail at four neighbors [2]. Complete reflection-even bispectrum/SNAP features and scalar ACE models restricted to the same correlation information inherit the six-neighbor obstruction [3–5]. A nonlinear model downstream of a fixed descriptor cannot undo an exact collision.

The theorem does not apply unchanged to every model called ACE, SOAP, SNAP, or an equivariant neural network. Finite basis truncation may cause earlier information loss. Higher-order ACE terms, parity-odd $SO(3)$ channels, equivariant features, message passing, species labels, and the collection of environments at other centers can add information that separates the pair. The result is local and does not automatically create a collision of a whole periodic structure or of a sum of site energies.

Pozdnyakov *et al.* demonstrated a practical error plateau for the lower-order version of this mechanism on a methane data set [2]. Parsaeifard and Goedecker found quasi-constant SOAP and ACSF manifolds that impair four-body torsional learning [6]; the subsequent comment distinguished near-degeneracy, which can be mitigated by resolution and targeted data, from exact collisions, which cannot [7]. These are relevant precedents, not direct validation of the six-neighbor pair.

S8 Controlled benchmark and experimental protocol

S8.1 Data generation and exact integrity checks

For each sampled parameter tuple (h, e, a, b) in the family of Section S2, we factor the two PSD Gram matrices into coordinates and compute canonical Gram decks. A k-card is canonicalized by enumerating all $k!$ internal vertex permutations and selecting the lexicographically minimal upper triangle; the card multiset is then sorted. The resulting D_3 tensor has shape 20×6 and D_4 has shape 15×10. Pair equality and noncongruence are checked before examples are serialized. Dataset splits are made at the pair level with independent seeds.

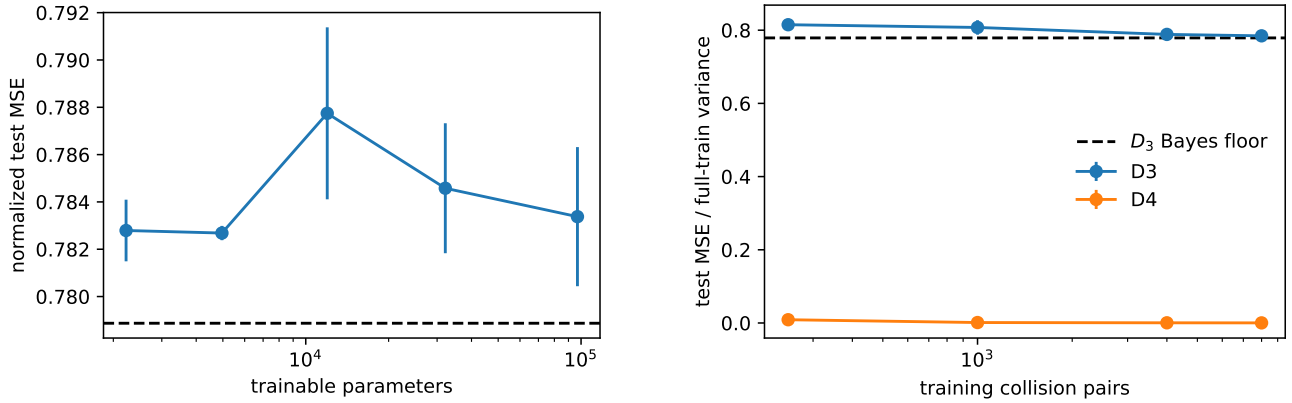
Targets are standardized with training statistics. The train mean and standard deviation are -0.87561 and 1.27123. On the test split the mean paired target gap is 2.10678 and the exact normalized midpoint/Bayes floor for an equiprobable paired sample is 0.778871. All reported regression experiments use five seeds unless a table explicitly states three.

Table S4: Exact data audit for the 11,000-pair controlled benchmark.

Split	Collision pairs	Max paired D_3 gap	Min paired D_4 distance
Train	8,000	0	1.854
Validation	1,000	0	1.904
Test	2,000	0	1.922

S8.2 Independent CPU reference, data, and capacity sweeps

The independent CPU reference flattens the canonical deck and uses a two-hidden-layer ReLU MLP, batch size 512, Adam learning rate 10^{-3}, at most 300 epochs, and early stopping with patience 20. With 8,000 training pairs, the width-128 D_3 model obtains 0.78458 ± 0.00275 NMSE and predicts paired branches equally to within 7×10^{-11} on average. Replacing only the input by D_4 yields $(2.793 \pm 0.603) \times 10^{-4}$ NMSE and mean predicted pair gap 2.10775.



(a) D_3 capacity sweep.

(b) Data sweep in one common variance unit.

Figure S2: Neither model capacity nor additional collision pairs remove the third-order information floor. Widths 16–256 span 2,225–97,025 parameters. Data-sweep points use the 8,000-pair training-target variance so that the analytic floor is a horizontal line.

Each data-subset optimization standardizes with that subset's own target statistics. Consequently, in native subset-normalized units the analytic floors are 0.78902, 0.76271, 0.76735, and 0.77887 for 250, 1,000, 4,000, and 8,000 training pairs. Figure S3 reports this audit directly; Figure S2(b) converts raw errors and floors to the common 8,000-pair variance unit before plotting. Thus the horizontal line in the main data-sweep figure is intentional, not a reuse of the wrong normalization.

S8.3 Higher order and incidence-aware repairs

The PyTorch suite compares DeepSets over canonical cards, a conventional node-state MPNN, and a stable edge-state MPNN on the complete weighted Gram graph. The edge architecture updates an edge from incident edge states, preserving the incidence information erased by motif histogramming. It uses one global scalar normalization statistic; columnwise normalization would spuriously depend on vertex order. CUDA runs use AdamW, batch size 512, gradient clipping for the MPNNs, and record device, software versions, wall time, and peak allocated memory.

S8.4 Propagation-depth ablation

The preregistered test compares edge-incidence depths $0, 1, 2, 3$ in two groups: fixed hidden width 128, and widths chosen to match the depth-3 model's 513,793 parameters within 0.2%. Every one of the 35 fits matches its preregistered separating/nonseparating behavior.

The depth-1 to depth-2 NMSE ratio is 4.11×10^5 at fixed width and 4.53×10^5 after parameter matching. This locates the repair in the edge-incidence receptive field rather than parameter count.

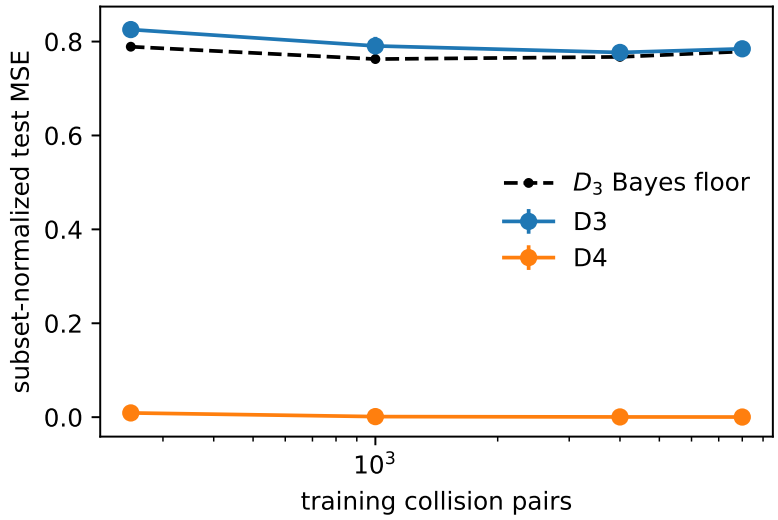


Figure S3: Native subset-normalized data sweep. The dashed analytic floor is a curve because each training subset defines its own target scale.

Table S5: Full-data model comparison (five seeds).

Model	Parameters	Test NMSE	Mean pair gap
D_3 DeepSets	67,073	0.780211 ± 0.000305	0
Node MPNN	396,289	0.6721 ± 0.3119	0.516 ± 0.759
D_4 DeepSets	67,585	$(1.592 \pm 0.592) \times 10^{-4}$	2.10462
Edge MPNN	513,793	$(1.565 \pm 0.217) \times 10^{-6}$	2.10685
Analytic D_3 floor	–	0.778871	0

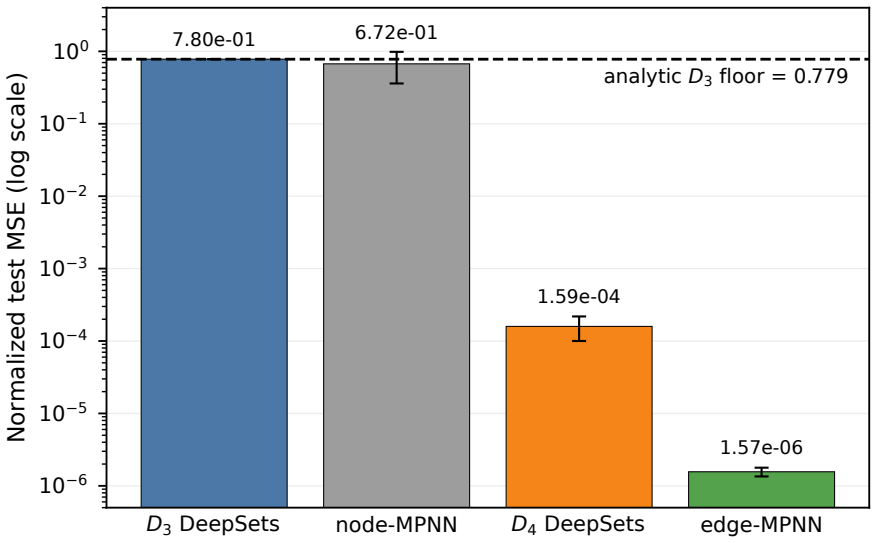


Figure S4: The stable edge MPNN and D_4 cards repair the collision task. The large variance of the node MPNN motivates the mechanistic depth audit rather than a claim based on a single successful seed.

Table S6: Edge-incidence depth ablation. Entries are mean ± population standard deviation over five seeds.

Group / depth	Hidden width	Parameters	Test NMSE / mean gap
Fixed / 0	128	66,817	0.78104 ± 0.00009 / 1.4×10^{-7}
Fixed / 1	128	215,809	0.778908 ± 0.000012 / 3.4×10^{-7}
Fixed / 2	128	364,801	$(1.896 \pm 0.130) \times 10^{-6}$ / 2.10690
Fixed / 3	128	513,793	$(1.565 \pm 0.217) \times 10^{-6}$ / 2.10685
Matched / 0	357	513,367	0.78054 ± 0.00013 / 1.0×10^{-7}
Matched / 1	198	514,009	0.778896 ± 0.000006 / 2.9×10^{-7}
Matched / 2	152	513,457	$(1.721 \pm 0.212) \times 10^{-6}$ / 2.10683
Matched / 3	128	513,793	$(1.565 \pm 0.217) \times 10^{-6}$ / 2.10685

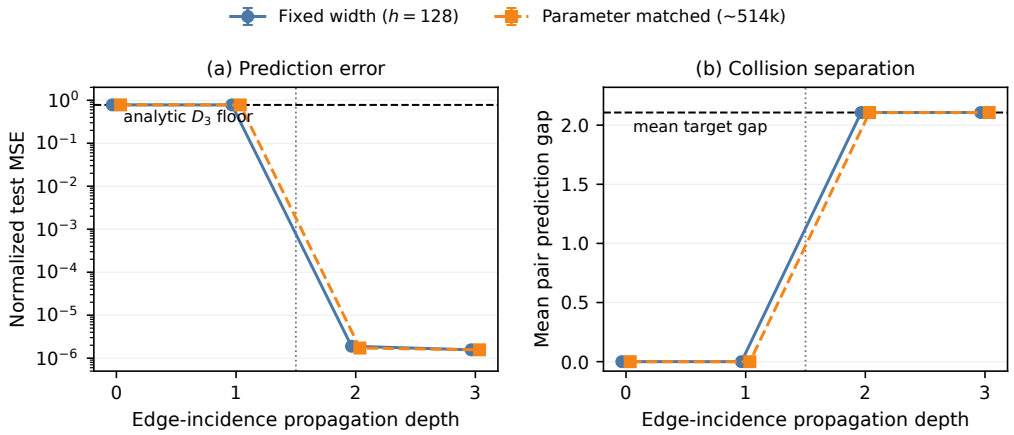


Figure S5: Depth ablation. One propagation round is still nonseparating; two rounds recover both the target and the paired gap in fixed-width and parameter-matched controls.

S9 Conditioning and noise robustness

S9.1 Descriptor versus orbit distance

For 1,000 independent exact pairs at each of eight noise scales, both branches receive independent isotropic coordinate perturbations. Card distance uses optimal bipartite matching after minimizing each card over all internal permutations. Orbit distance is the minimum full-Gram Frobenius distance over all 720 vertex permutations. This avoids the discontinuity of subtracting two lexicographic representatives near ties.

At $\sigma = 0$, the maximum D_3 distance is 2.16×10^{-14} and the median orbit distance is 1.35193. For nonzero scales, median D_3 distance is fitted by $28.48\sigma^{1.00083}$, while the orbit-distance log–log slope is -1.23×10^{-4}. Thus an arbitrarily small descriptor displacement can coexist with order-one orbit separation near the exact collision set.

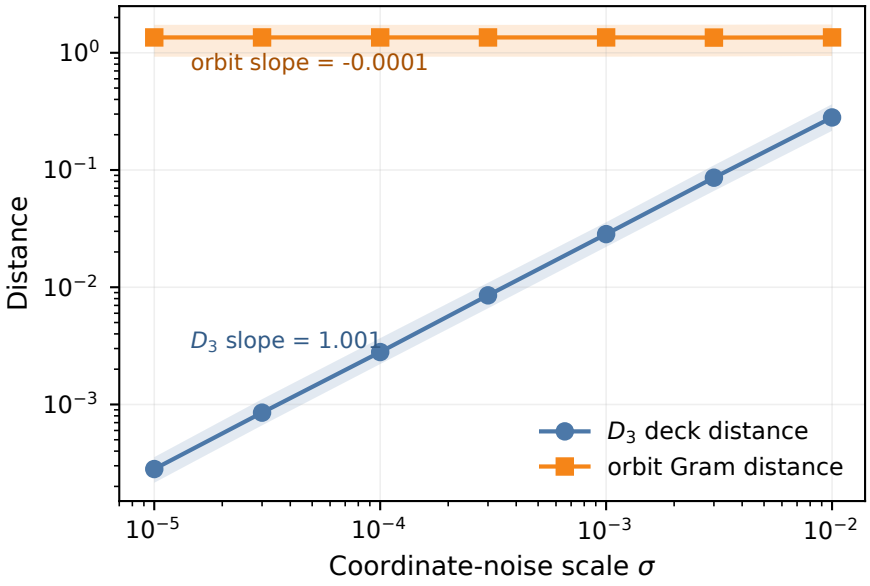


Figure S6: Verified conditioning experiment. Lines show medians and shaded regions show the 10–90% empirical quantiles across 1,000 pairs per scale.

S9.2 Regression under coupled coordinate noise

For $\sigma \in \{0, 10^{-4}, 10^{-3}, 10^{-2}\}$, we regenerate train, validation, and test sets, train three seeds for each model, and normalize raw test MSE by the target variance of that test split. The exact-collision floor in these plot units is 0.773912 and is shown only at $\sigma = 0$.

Table S7: Noise ablation: mean test NMSE over three seeds.

σ	D_3	D_4	Edge MPNN
0	$7.752\,10^{-1}$	$1.278\,10^{-4}$	$1.497\,10^{-6}$
10^{-4}	$5.650\,10^{-1}$	$3.016\,10^{-4}$	$2.821\,10^{-6}$
10^{-3}	$5.667\,10^{-1}$	$3.966\,10^{-4}$	$2.157\,10^{-5}$
10^{-2}	$8.205\,10^{-2}$	$6.018\,10^{-4}$	$1.567\,10^{-5}$

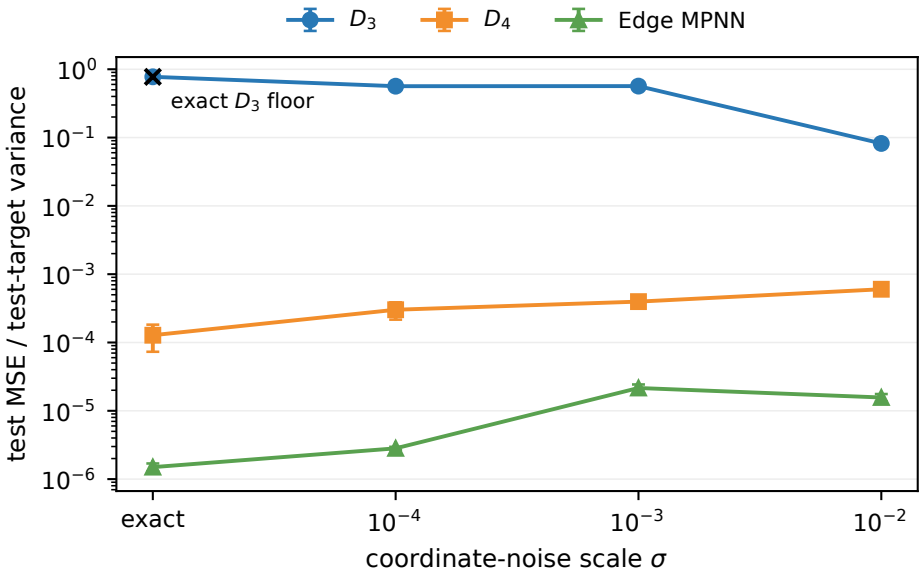


Figure S7: Noise ablation. D_3 improves at larger noise because the input collision is broken, while D_4 and edge incidence solve both the exact and perturbed tasks. Error bars are population standard deviations over seeds.

S10 Species channels and empirical prevalence

S10.1 Exact species-scope audit

We enumerate all $B_6 = 203$ canonical partitions of the six vertices into species classes, using the same vertex-to-species assignment on both clouds. Only the homogeneous partition and the two-species pattern (a,b,b,a,a,a) (up to species renaming) retain equal colored decks of orders one through three. Conditional on using exactly two species, the surviving fraction is $1/31$; every assignment using at least three distinct species separates this integer pair. With q available labels sampled uniformly, the unconditional survival probability is q^{-4} because exactly q^2 of q^6 assignments survive. Species therefore reduces prevalence for this witness but cannot yield universal completeness: homogeneous environments remain valid counterexamples.

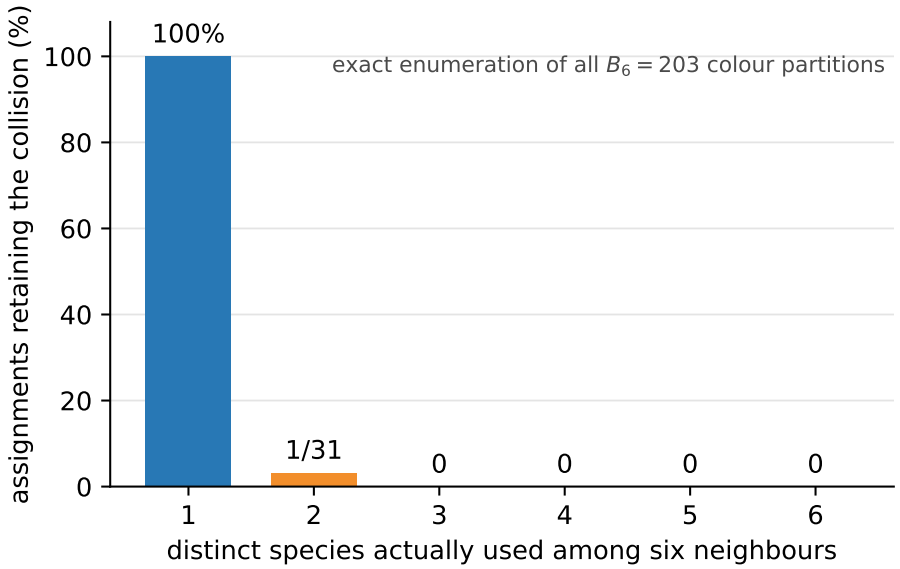


Figure S8: Exact species audit over all 203 canonical color partitions.

S10.2 Validated ethanol trajectory diagnostic

We use the trajectory scan only as a negative prevalence audit, not as a materials-failure experiment. From an rMD17-style ethanol trajectory, the scanner samples 20,000 center environments, retains the six nearest neighbors, ignores element labels, retrieves close canonical D_3 decks, and reranks 7,000 candidate pairs by optimal within-card/between-card matching and full Gram-orbit distance. A strict candidate is preregistered as

$$d_{D_3}^{\mathrm{rel}} < 10^{-2}, \qquad d_{\mathrm{orbit}}^{\mathrm{rel}} > 0.1. \tag{S96}$$

No strict candidate is found. On an exhaustive 200-environment audit subset, candidate and post-rerank recall@100 are both 1.0. Synthetic positive controls give maximum D_3 distance zero and minimum orbit separation 0.11608.

Table S8: Final validated real-trajectory diagnostic.

Quantity	Result
Sampled environments	20,000
Reranked candidate pairs	7,000
Audit candidate recall@100	1.0
Audit post-rerank recall@100	1.0
Strict separated near-collisions	0
Spearman $\rho(D_3, \mathrm{orbit})$	0.96657
Synthetic control: max D_3 / min orbit	0 / 0.11608

The force-label gaps from this untyped scan are diagnostic only because species and chemical-site identity are confounded. Frame energy is global and cannot be attributed to one six-neighbor environment. The correct conclusion is therefore limited: no strict separated near-collision was found among the 7,000 validated checked pairs. The experiment neither estimates the measure of the exact algebraic locus nor proves absence from ethanol trajectories.

S11 Reproducibility and trust boundary

From the verification-package root, run

```
python3 verify_all.py
```

to check, with the Python standard library only:

1. all four-neighbor equality patterns for every radius partition;
2. all thirteen rational certificates in Table S3;
3. the low-degree M3/M4 hand identities and every terminal M1–M4/R41 relabeling;
4. the finite cores of the hand radius-reconstruction lemmas;
5. coverage of all 80 equal-radius and 43 radius-$4+1$ collision pairs by the five maximal templates;
6. the six-neighbor coordinates, lower decks, symbolic family deck, and all 720 full-matrix relabelings;
7. all 96 compatible local deck matchings and the exact rank-38 Jacobian calculation in Proposition 3.

The expected final line is

```
ALL EXACT CHECKS PASSED
```

For a source-to-output regeneration of the five-vertex enumeration, run

```
python3 verify_all.py --full-enumeration
```

This additionally compiles the archived C++17 enumerators in a temporary directory and regenerates all ten edge-color cases for radius partitions $5,4+1,3+2,3+1+1,2+2+1,2+1+1+1$, including every zero-collision output. The all-distinct case is reconstructed directly.

The machine audit establishes the serialized polynomial identities, finite enumerations, local Jacobian ranks, template coverage, and explicit relabelings. It is retained as a complete evidence chain rather than replaced by the added hand proofs. Conversely, the correlation–deck equivalence, topological factorization, radius-split gluing lemmas, and all five template arguments are intentionally written out for direct mathematical review. Certificate-generation scripts may use NumPy to discover convenient Macaulay minors, but neither the archived identities nor any verifier depends on floating-point linear algebra.

S12 Relation to prior reconstruction work

Boutin and Kemper studied reconstruction of point configurations from distance or area distributions [8] and later node- and edge-weighted complete graphs reconstructed from weighted subtriangle distributions [9]. General weighted K_5 already has ambiguities. The present result is not a claim that triangle-deck reconstruction is new; it identifies a sharp geometric exception in which PSD rank-three realizability removes every five-vertex ambiguity.

Earlier atomistic work established general incompleteness of low-order density correlations [2], proposed alternative complete representations [10], and exhibited degeneracies persisting to higher correlation order [1]. To our knowledge, those works do not provide the universal $N \leq 5$ reconstruction theorem and hence do not determine the sharp five/six occupancy threshold for the reflection-even bispectrum. The six-neighbor family itself is credited to Ref. [1]; the new contribution is the matching lower bound and exact classification.

References

[1] Michelangelo Domina and Michele Ceriotti. Using large language models to probe the limits of atom-centered structural descriptors. *arXiv:2607.26984*, 2026.

[2] Sergey N. Pozdnyakov, Michael J. Willatt, Albert P. Bartók, Christoph Ortner, Gábor Csányi, and Michele Ceriotti. Incompleteness of atomic structure representations. *Phys. Rev. Lett.*, 125:166001, 2020.

[3] A. P. Thompson, L. P. Swiler, C. R. Trott, S. M. Foiles, and G. J. Tucker. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *J. Comput. Phys.*, 285:316–330, 2015.

[4] Ralf Drautz. Atomic cluster expansion for accurate and transferable interatomic potentials. *Phys. Rev. B*, 99:014104, 2019.

[5] Geneviève Dusson, Markus Bachmayr, Gábor Csányi, Ralf Drautz, Simon Etter, Cas van der Oord, and Christoph Ortner. Atomic cluster expansion: Completeness, efficiency and stability. *J. Comput. Phys.*, 454:110946, 2022.

[6] Behnam Parsaeifard and Stefan Goedecker. Manifolds of quasi-constant SOAP and ACSF fingerprints and the resulting failure to machine learn four-body interactions. *J. Chem. Phys.*, 156:034302, 2022.

[7] Sergey N. Pozdnyakov, Michael J. Willatt, Albert P. Bartók, Christoph Ortner, Gábor Csányi, and Michele Ceriotti. Comment on "manifolds of quasi-constant SOAP and ACSF fingerprints and the resulting failure to machine learn four-body interactions". *J. Chem. Phys.*, 157:177101, 2022.

[8] Mireille Boutin and Gregor Kemper. On reconstructing n-point configurations from the distribution of distances or areas. *Adv. Appl. Math.*, 32:709–735, 2004.

[9] Mireille Boutin and Gregor Kemper. Lossless representation of graphs using distributions. *arXiv:0710.1870*, 2007.

[10] Jigyasa Nigam, Sergey N. Pozdnyakov, Kevin K. Huguenin-Dumittan, and Michele Ceriotti. Completeness of atomic structure representations. *APL Mach. Learn.*, 2:016110, 2024.

S13 Revision: clarifications and additional geometric audits

This section supplements the preceding S1–S12, whose complete proof pages are retained. Where wording or normalization differs, the clarifications below supersede the earlier wording. No molecular dynamics trajectories were generated for this revision, and no physical energy/force model was trained. New computations use seed 20260922. Raw pair arrays and summaries are released alongside the scripts.

S13.1 Definitions, regularity, and normalization

The complete deck includes orders one and two, including when $N < 3$; the distinct-index order-k deck is empty for $k > N$. For fixed N, this is equivalent to keeping the repeated-index cards in the ordered correlation. The integer example uses $(h,e,a,b) = (5,5,4,2)$, while $(5,3,1,2)$ is a separate rational point for the local-dimension calculation. Here $h = \|u\|^2$, $e = \|v\|^2 = \|w\|^2$, $a = u \cdot v$, $b = u \cdot w$. The ambient rank-three Gram manifold has dimension $18 - 3 = 15$; at a point with trivial permutation stabilizer its finite quotient is locally a manifold. The 96-matching rank calculation is local, not a global pure-dimension assertion about the collision set.

For the integrated gradient/force bound in S6, assume V and $\widehat{V}$ are C^1 on an open neighborhood of the absolutely continuous coordinate path γ. Then their compositions with γ are absolutely continuous and obey the chain rule almost everywhere. This is a sufficient condition for the line-integral argument; the endpoint and Bayes bounds require no differentiability. Merely writing "differentiable along a path" does not by itself supply all these hypotheses.

A singular anchor requires both the range condition and residual Schur constraints. With $P = (p,q)$ and outside Gram block C, $C - P^T B^+ P \succeq 0$ has rank at most $3 - \operatorname{rank} B$. Range membership alone is not as restrictive as invertible-anchor reconstruction: the normal components remain. The global template proofs cover this case.

To distinguish metrics, write $d_G = \min_\pi \|G_X - P_\pi G_Y P_\pi^T\|_F$ and $d_{\mathrm{xyz}} = \min_{Q,\pi} \|X - QY_\pi\|_F$ for unnormalized distances. The coordinate metric in S1 and the revised main text is $d_{\mathrm{coord}} = d_{\mathrm{xyz}}/\sqrt{N}$. The numerical code reports the unnormalized card-assignment distance

$$d_{\mathrm{cards}}^2 = \min_\tau \sum_I \min_{P \in S_3} \|G_X[I] - P G_Y[\tau(I)] P^T\|_F^2.$$

It equals $\sqrt{\binom{N}{3}}$ times the RMS distance in S2, Eq. (S29). In particular the original S9 conditioning numbers use the unnormalized version, and its "orbit distance" is d_G, not d_{coord}. These constant factors do not change the slopes, but coordinate and Gram distances should not be interchanged near rank-deficient clouds.

S13.2 Five-neighbor stability below the threshold

Four fixed clouds are tested: a generic rank-three cloud, its planar projection, five distinct collinear points, and a regular unit pentagon. For each cloud, 64 independent Gaussian directions in $\mathbb{R}^{5\times 3}$ are normalized to unit Frobenius norm. The same directions are used at $\varepsilon = 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}$, yielding 1,280 perturbation comparisons. The script records all distance quantiles. Across these random-direction tests, the median ratio d_G/d_{cards} stays approximately 0.50; this is a finite directional diagnostic, not a uniform stability bound.

There is also an explicit transverse planar test. Let the two nonzero coordinate columns of X be

$$X = \begin{pmatrix} .2 & .3 & 0 \\ -.8 & .4 & 0 \\ .6 & -.9 & 0 \\ 1.1 & .3 & 0 \\ -.4 & -.6 & 0 \end{pmatrix}.$$

Project $z_0 = (.3, -.7, .2, .9, -.4)^T$ orthogonally off $\mathrm{col}(X)$ and normalize to obtain $\|z\| = 1$. Set $Y_\varepsilon = X + \varepsilon z e_3^T$. Then

$$G_Y - G_X = \varepsilon^2 z z^T, \qquad d_{\mathrm{xyz}}(X, Y_\varepsilon) = \varepsilon \quad \text{for sufficiently small } \varepsilon.$$

The latter follows from $X^T z = 0$, orthogonal Procrustes alignment and the positive gap to every nonidentity permutation of this distinct-radius base. Its three-cards also have a positive matching gap, so the identity card matching is optimal for small enough ε. Accordingly $d_{\mathrm{cards}} = \Theta(\varepsilon^2)$ whereas $d_{\mathrm{coord}} = \varepsilon/\sqrt{5}$. The measured log–log slope is 2.00000006, and the unnormalized ratio $d_{\mathrm{xyz}}/d_{\mathrm{cards}}$ grows from 4.9824 to 4.9824×10^4 over the five noise scales. Thus continuous reconstruction does not imply a locally Lipschitz inverse from Gram-card distance to coordinate-orbit distance at this planar cloud. It remains consistent with a weaker modulus of continuity.

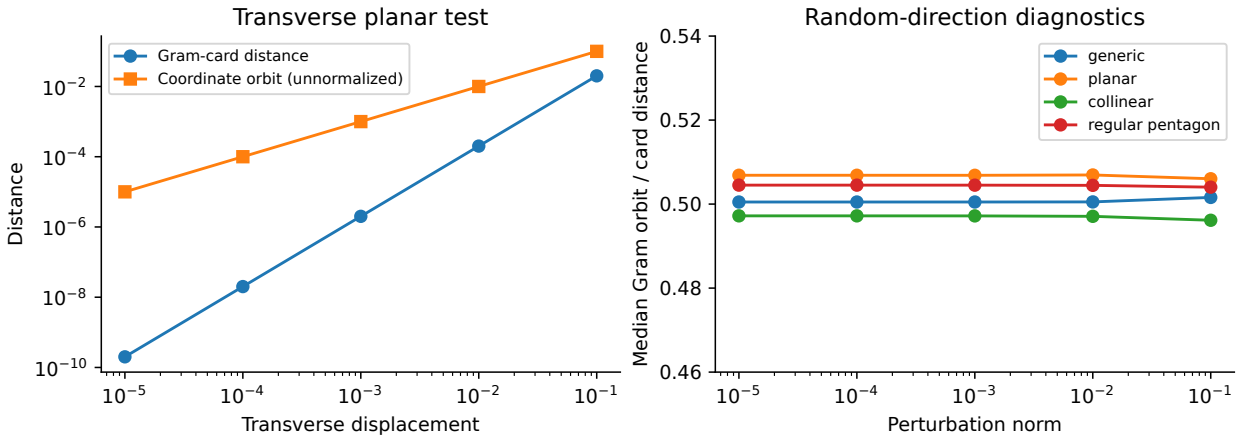


Figure S13.1: New five-neighbor stability diagnostic. Left: a transverse planar perturbation changes coordinates linearly and Gram cards quadratically. Right: median Gram-orbit/card ratios for 64 random directions per fixed cloud. These are different metrics and different directional tests.

S13.3 Exact fixed-stoichiometry enumeration

For each integer partition of six, fix a set of named species with those counts and exhaust all distinct site assignments. The same assignment is used on the displayed vertex ordering of both collision branches. Colored decks of orders one, two and three are canonicalized with integer arithmetic, including multiplicities. The results are:

Composition	Assignments	Surviving	Probability
A_6	1	1	1
A_5B	6	0	0
A_4B_2	15	1	$1/15$
A_3B_3	20	0	0
A_4BC	30	0	0
A_3B_2C	60	0	0
$A_2B_2C_2$	90	0	0
A_3BCD	120	0	0
A_2B_2CD	180	0	0
A_2BCDE	360	0	0
$ABCDEF$	720	0	0

The surviving patterns remain $(\alpha,\beta,\beta,\alpha,\alpha,\alpha)$. This conditions the earlier q^{-4} iid-label result on composition. Uniform assignments under fixed counts are a combinatorial model; the probabilities are not chemical equilibrium frequencies and do not cover independently assigned labels on the two branches.

S13.4 Exhaustive finite geometric stress tests

We generate 128 environments in each of three synthetic sets: unit octahedral vertices $\{\pm e_1,\pm e_2,\pm e_3\}$ with independent Gaussian coordinate noise of standard deviation 0.01 or 0.05, and six independently sampled unit vectors (normalized isotropic Gaussians). The last set has equal radii by construction, up to floating-point rounding. These are geometric stress tests, not sampled equilibrium states of a crystal or liquid.

For each of the existing rMD17-format ethanol, aspirin and uracil files, we sample 160 distinct frames without replacement, choose one uniformly random center per frame, and extract its six nearest neighbors. Atom types are ignored for this representation diagnostic. Every within-set pair is compared, excluding trajectory pairs separated by fewer than ten frame indices. Thus there is no approximate candidate retrieval or reranking truncation in the new scans. The frame/center indices and input-file SHA256 hashes are archived. Existing trajectory coordinates are reused as provided; the original 20,000-environment ethanol audit remains a separate experiment.

Relative distances are $d_{\text{cards}}/\max(\|\mathscr{C}_X\|_F,\|\mathscr{C}_Y\|_F)$, where $\mathscr{C}_X$ stacks all full 3×3 cards, and $d_G/\max(\|G_X\|_F,\|G_Y\|_F)$. All 720 vertex permutations and all six internal card permutations are considered, with optimal bipartite card assignment. The vectorized implementation agrees with the independent original scalar implementation on random clouds of sizes 3–6 and on the exact integer collision control.

Table S13.1: New exhaustive within-set scans. $r_{\min}$ is the median over environments of $\min_{i<j}|G_{ii}-G_{jj}|/\max_i G_{ii}$.

Set	Environments	Eligible pairs	Strict events	$r_{\min}$
Octahedron, noise 0.01	128	8,128	0	0.00122
Octahedron, noise 0.05	128	8,128	0	0.00597
Random equal radii	128	8,128	0	0
Ethanol	160	12,715	0	0.0245
Aspirin	160	12,717	0	0.0144
Uracil	160	12,717	0	0.0179

No pair meets both relative card distance < 0.01 and relative Gram-orbit distance > 0.1. The full grid uses card thresholds $0.0025, 0.005, 0.01, 0.02, 0.05, 0.1$ and orbit thresholds $0.02, 0.05, 0.1$; all counts, distances and configurations are archived. These negative results are limited to the enumerated sets. Pair observations share environments and trajectory frames may be correlated, so pair counts are not independent Bernoulli trials and no population prevalence bound is inferred. Repeated radii are necessary for an exact collision, not sufficient. Conversely, heterogeneous bond lengths do not alone guarantee a finite lower bound on near-collision distances.

S13.5 Existing benchmark re-audit and reproducibility

Recomputing statistics from the saved five-seed metric files gives the predicted-gap mean and population standard deviation 2.1046243 ± 0.0014559 for D_4, 2.1068493 ± 0.0001822 for the edge MPNN, and 0.5155240 ± 0.7591451 for the node MPNN. The D_3 gap is exactly zero in all five archived runs. These replace the missing gap standard deviations in the earlier S8 table. The raw training-target variance is 1.6160230 (float64 re-evaluation of the stored targets); the recomputed normalized paired Bayes floor is 0.77887119. This is a re-analysis of existing runs, not a claim that those models were retrained for this revision.

The base generator samples h, e independently uniformly in $[0.8, 1.2]$, $a/\sqrt{he}$ in $[0.65, 0.85]$, $b/\sqrt{he}$ in $[0.25, 0.45]$, rejecting $a^2 + b^2 \geq 0.95he$. The dataset seed is 2027. Whole pairs stay together in the 8,000/1,000/2,000 training/validation/test split. The original D_3, D_4, node-MPNN configuration uses AdamW, standardized MSE, batch size 512, learning rate 10^{-3}, weight decay 10^{-6}, at most 400 epochs and validation patience 40. The stable edge configuration uses batch size 256, learning rate 3×10^{-4}, weight decay 10^{-6}, at most 800 epochs, patience 100, gradient clipping at 1, full precision, and ReduceLROnPlateau with patience 20, factor 0.5, and minimum rate 10^{-6}. Thus the earlier blanket statement that all CUDA runs use batch 512 is superseded. The archived configurations remain the reproducibility source; settings should not be inferred from the model name alone.

Noise is independent between branches; resetting the noise seed at each σ couples directions only across noise levels. Each perturbed target is recomputed. The noisy learning sweep has three seeds, whereas the exact benchmark and depth comparison have five. These are separately budgeted experiments, and their error bars should not be interpreted as having equal precision. The directed-edge formula in the revised main text matches `src/models.py`; its mean is over predecessor states before the message MLP. Its two-round receptive field includes four-vertex paths but also shorter closed walks. The measured depth transition does not establish a universal minimal depth across architectures. No cause for the node-model seed variance is established by these statistics alone.

Exact proof checks. The refreshed package retains all thirteen certificates and both additional radius-reconstruction and local-Jacobian checkers. The rational audit obtains rank 38 for all 96 compatible matchings and representative determinant $-1280/81$. A separate source-to-output C++17 audit reproduces all archived K5 outputs across the six nontrivial radius multiplicities. The original geometric arguments remain mathematical proofs requiring human review.

Commands from the supplementary-package root.

```
python proof/verify_all.py
python proof/verify_all.py --full-enumeration
cd experiments
python run_supplementary.py --data-root /path/to/rmd17
```

The new experiment requires NumPy and SciPy; figure generation additionally uses matplotlib. The real-data files are not redistributed; their hashes and sampled indices identify the exact inputs. Synthetic configurations and all evaluated geometric pairs are included. Physical energy/force fitting and comparisons to MACE or GemNet remain outside this controlled study.