Same Dynamics, Different Structures in Energy-Preserving Operator Fits

Sushaan Kandukoori and Aarya Patel

***Abstract*—Models $\dot{z} = J(z)\nabla E(z)$ with skew-symmetric J affine in the state and quadratic E are fitted to trajectories because they conserve energy by construction. The fit determines the dynamics without determining the coefficients: in dimension d, for any quadratic energy with invertible Hessian, the coefficient changes leaving the vector field untouched at every state form a subspace of dimension $\binom{d}{3}$, known in closed form before any data are collected: 56 of the 224 bracket coefficients at $d = 8$. Its dimension recurs at two reduced orders of a model-reduction benchmark, and on five of thirteen algebras it contains a non-isomorphic Lie algebra reproducing the field, trajectories and Jacobi identity to machine precision, so no residual or conservation diagnostic detects the substitution. A certificate computed from a hypothesized bracket and the energy alone decides in advance which of two levers restores identifiability; over thirteen algebras, its routing matched the measured outcome on all nine where both levers were run. Where it permits, a Levenberg–Marquardt retraction along the invisible directions recovers the coefficients in 8–9 steps without moving what the data determined, at two orders of magnitude less work than the penalty and constrained-solver arms that also recover them; where it does not, a second experiment with pairwise-distinct curvature ratios raises the fitted-to-true coefficient cosine from 0.336 to 1.0000, while a proportional second energy at the same cost gains nothing.**

***Index Terms*—Computational modeling, least squares methods, Lie algebras, nonlinear dynamical systems, parameter estimation, structural identifiability, system identification.**

I. INTRODUCTION

A Rigid body tumbling in free space is the smallest instance of the model class treated here. Write $z \in \mathbb{R}^3$ for the angular momentum in body coordinates and $E(z) = \frac{1}{2} z^\top A z$ for the kinetic energy, with A the inverse of the inertia tensor. Euler's equations read $\dot{z} = z \times Az$. The cross product with z is a matrix acting on Az, so the equations take the form $\dot{z} = J(z)\nabla E(z)$, where the state-dependent matrix $J(z)$ has entries $J(z)_{ij} = \sum_k z_k c_{ij}^k$ with $c_{ij}^k = \varepsilon_{ikj}$ the Levi-Civita symbol; $J(z)$ is skew-symmetric for every z. Two properties of this form are why engineers fit it rather than a black box. First, skewness of J makes $\dot{E} = \nabla E^\top J \nabla E = 0$ for *any* coefficients c, so conservation holds without drift under fitting error and over long rollouts. Second, the coefficients are a small array with physical content: they encode which quantity generates motion in which direction, and are comparable against a hypothesis, transferable to a new operating point, readable as parameters.

Preprint.

S. Kandukoori is an independent researcher (e-mail: sushaankandukoori@gmail.com).

A. Patel is with De Anza College, Cupertino, CA 95014 USA.

The same template covers systems far larger than a rigid body. Let $z \in \mathbb{R}^d$, $E(z) = \frac{1}{2} z^\top A z$ with A symmetric, and

$$J(z) = J_0 + \sum_{k=1}^{d} z_k \, c^k, \qquad J_0^\top = -J_0, \quad (c^k)^\top = -c^k. \tag{1}$$

Fitting (1) means estimating $\frac{d(d-1)}{2} + \frac{d^2(d-1)}{2}$ real numbers (J_0, c) from observed states and time derivatives, a linear least-squares problem since the field is linear in the parameters. The array c may additionally satisfy the Jacobi identity $\sum_{\circlearrowleft} c(c(x,y),z) = 0$, summed over cyclic permutations of (x, y, z), a quadratic polynomial system in its entries. When it holds, c is the table of structure constants of a d-dimensional Lie algebra, and (1) is Lie–Poisson when $J_0 = 0$, the case of every algebra experiment here; a nonzero constant block needs one further linear condition for the pair to be Poisson, stated as (11). Section II develops the vocabulary from the $d = 3$ case.

The fit does not determine the coefficients. On $\mathfrak{su}(3)$ ($d = 8$), a blind least-squares solve with the energy given, at three energy draws, drives the held-out rollout error from 0.14–0.28 to 2–$6 \cdot 10^{-3}$ as the orbit budget (the number of independent initial conditions) grows from 8 to 512. Over the same draws, the cosine between estimated and true structure constants rises to 0.235–0.309 and stops at a value computable before any data are collected (Section II-D). The solve also reports why. Of the 252 bracket parameters at $d = 8$ the design sees 196; the remaining 56, or 22%, leave the vector field pointwise unchanged. They span the kernel of the parameters-to-dynamics map, so no optimizer, no regularizer, and no additional trajectory removes them, and that kernel depends only on the state dimension and the energy Hessian, available in closed form before a single measurement is taken.

Motion along the kernel also changes the bracket's isomorphism class, the algebra's identity up to a change of basis. On five algebras of the thirteen-algebra battery of Section V-D, one indecomposable and nilpotent at $d = 4$, a Lie algebra *not isomorphic to the truth* reproduces the true vector field to $7.7 \cdot 10^{-15}$ over 3000 states and its RK4 trajectories to $7.8 \cdot 10^{-16}$ over 300 steps, at Jacobi residual at most $1.2 \cdot 10^{-16}$. Such a model passes every check a practitioner would run, conservation included: nothing in the data distinguishes it from the truth, and the two disagree on the algebra's invariants.

For a bracket fitted to measured data, this article determines what can be recovered, what cannot, how that is known in advance, and what changes the answer.

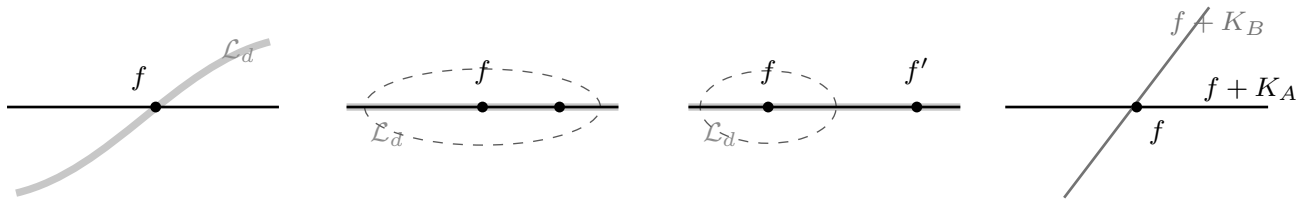

(a) $r_{\mathrm{Jac}} = 0$ *retract along* K_A (b) $r_{\mathrm{Jac}} > 0$, $\iota_H = 0$ *report the basis ambiguity* (c) $r_{\mathrm{Jac}} > 0$, $\iota_H > 0$ *run a second energy* (d) two energies $K_A \cap K_B = 0$

Fig. 1. The three cases the certificate distinguishes, and what a second experiment does. Read the panels before the notation: f is the true coefficient array, the heavy horizontal line is the set of arrays that generate the same vector field as f and are therefore indistinguishable from it in any quantity of trajectory data, the pale band is the set of arrays satisfying the Jacobi identity, and the dashed ellipse is the set reachable from f by a change of coordinates. (a) The two sets meet only at f, so imposing the Jacobi identity along the line returns the truth. (b) The line lies inside the band, so the Jacobi identity cannot choose among its points, but they all lie in the ellipse and differ from f only by a relabeling. (c) The line lies inside the band and leaves the ellipse, so f' is a structurally different algebra reproducing the same trajectories, and only a second experiment separates it from f. (d) A second energy shrinks the indistinguishable set to the intersection of two such lines, which Theorem 2 makes trivial when the curvature ratios are pairwise distinct. The panel labels r_{Jac} and ι_H are the two integers Section III-G defines; r_{Jac} counts the indistinguishable directions that keep the Jacobi identity, and ι_H how many of those are not relabelings. Schematic; not to scale.

A. Contributions

- *What trajectory data cannot see, and the ceiling it forces* (Section III, Theorem 1, Corollary 1). The unidentifiable *gauge* is, for every invertible energy Hessian A, the totally antisymmetric part of the coefficient array twisted by A in one index, of dimension $\binom{d}{3}$; the constant part J_0 is identifiable. Measured nullities match the counts at eight dimensions between $d = 3$ and $d = 12$ and on the 2D Burgers benchmark of the operator-inference literature. Because minimum-norm least squares returns the projection of the truth off the gauge, the cosine any gauge-blind estimator can reach is fixed before data collection.
- *An intersection law that says which second experiment helps* (Theorem 2). For two diagonal energies the surviving ambiguity has dimension $\sum_t \binom{|L_t|}{3}$ over the level sets of the curvature ratio vector, so pairwise-distinct ratios annihilate it. A four-arm ladder at matched total budget confirms the predicted dimensions, with a proportional-energy arm ruling out data volume as the variable.
- *An a priori certificate* (Section III-G). Exact linear algebra on the pair (hypothesized bracket, energy), computed without trajectory data, decides which repair applies: r_{Jac} routes between the two levers, and its image in cohomology, ι_H, separates fits that merely re-label the truth from fits that land on a different algebra. On the nine algebras where both levers were run the routing matched the measured outcome on every one, and the pair is strictly finer than the classical deformation invariant H^2, which misroutes three rows of the battery.
- *One procedure for the two repairs, with a predicted budget* (Section IV). Fitting inside the quotient makes the design full rank, its smallest singular value a usable margin rather than a numerical zero. A Levenberg–Marquardt retraction then moves only along directions the data cannot see, enforcing the Jacobi identity without disturbing what the data determined and at two orders of magnitude less work than the penalty and constrained-solver arms that also recover the coefficients. Because the quotient margin follows from a pilot orbit, the orbit budget meeting a target accuracy is predicted rather than tuned, and its refusals are informative. The Jacobi residual's null distribution is dimension-dependent, so the implemented gate is too.

B. Related Work

a) Structure-preserving learning of dynamics: Imposing geometry on a learned model yields conservation and long-horizon stability that a generic regressor lacks: Hamiltonian and Lagrangian networks [1], [2], symplectic integrator networks [3], port-Hamiltonian formulations [4], and the geometric-integration literature behind them [5], [6]; sparse regression [7] and operator inference [8], [9] fit interpretable coefficient arrays directly, with structure-preserving variants constraining those arrays to be energy-preserving [10]. That work asks whether the fitted model predicts and conserves; the question here is whether the fitted coefficients are the true ones.

b) Non-uniqueness in bracket and operator identification: The parameter count of the ambiguity is prior art. Gkimisis et al. [11] prove that every energy-preserving quadratic operator admits a skew-block representative with $n^3/6 - n^2/2 + n/3$ freely selectable entries, then fix one by a bookkeeping rule and set the freedom aside; their energy is $\frac{1}{2}\|x\|^2$, their linear operator sits outside the analysis, and the free directions are counted by a rank argument rather than identified. The count is theirs; what we add is that this space is the kernel of the parameters-to-dynamics map, the totally antisymmetric summand twisted by the energy, of dimension unchanged for every invertible Hessian, with J_0 splitting off as identifiable and the complementary summand making the design full rank. Hu, Ortega and Yin [12] treat the mirror problem, the bracket fixed and the Hamiltonian learned up to Casimirs, which is complementary to ours rather than competing. Latent Lie–Poisson networks [13] adopt the bracket as known. MLSD [14] recovers $\mathfrak{so}(4)$ and $\mathfrak{su}(3)$

"up to basis transformations" and repairs the ambiguity post hoc; PHAST [15] names the gauge freedom and pins it with knowledge anchors; neither characterizes what was quotiented. Exact Jacobi enforcement appears in [16] at $d = 3$.

c) What is classical at d=3: At $d = 3$ the substitution phenomenon is not new. Every unimodular linear Poisson structure on $\mathbb{R}^3$ is $\nabla C \times \cdot$ for a quadratic Casimir C, so the one-dimensional ambiguity rescales C and stays on the variety automatically, and the resulting one-parameter families sit inside the classical Bianchi classification of three-dimensional algebras [17], with equivalence classes already cataloged [18], [19]. The $\mathfrak{se}(2)$ and $\mathfrak{h}_3$ rows of Table III reproduce it and no credit is claimed for them. The claim of this article rests on $d \geq 4$, where the filiform $\mathfrak{n}_4$ (nilpotent with a one-dimensional center, hence indecomposable) cannot be a $d = 3$ example carrying an abelian summand.

d) Structural identifiability: Which parameters a model class determines from ideal, noise-free data is the classical subject of structural identifiability [20], [21], developed for systems-biology models [22] and repaired there by reparameterization [23]. The closest control-theoretic anchor is Sontag et al. [24], characterizing the input classes under which bilinear systems are identifiable. The intersection law of Section III is the energy-side analogue: where that work varies the input, a second energy with a different curvature spectrum resolves the ambiguity here, and the law says how much a given choice leaves.

e) Null-space methods and gauge fixing: Moving along a data null space to satisfy a second criterion has earlier names: the null-space shuttle in seismic tomography [25], the free-network datum in bundle adjustment [26], identifiable reparameterization [23], and task-priority null-space projection in redundant manipulators [27]. Here the second criterion is an algebraic variety, the solution set of the polynomial Jacobi equations, rather than a linear constraint or a scalar objective, with a retraction onto it exact in general d and certified well-posed before it runs.

f) Symmetry and Lie algebra discovery: Learning a symmetry group or its Lie algebra from data is an active line [28]–[30], whose methods recover generators of a group acting on data; the object here is the bracket generating the dynamics, and the ambiguity is a property of the pair (bracket, energy) rather than of the algebra alone. That distinction separates the certificate from deformation theory. Richardson [31] showed that the second cohomology H^2, which counts the structure-constant perturbations that keep an algebra a Lie algebra modulo those that merely re-label it, fails to characterize rigidity, the absence of nontrivial deformations; Nijenhuis and Richardson [32], [33] set the deformation theory used in Section III-G. The failure here is a second and independent one: H^2 is computed without reference to the energy through which the data probe the bracket, and the same algebra has different invisible cocycles under different energies.

II. BACKGROUND AND PROBLEM STATEMENT

A. The free rigid body

Let $z \in \mathbb{R}^3$ be the angular momentum of a torque-free rigid body in body-fixed principal-axis coordinates, $I = \mathrm{diag}(I_1, I_2, I_3)$ its inertia matrix, $A = I^{-1}$, and $E(z) = \frac{1}{2} z^\top A z$ its kinetic energy, with gradient $\nabla E(z) = Az$ the angular velocity. Euler's equations $\dot{z} = z \times (Az)$ use the cross product only through its matrix form

$$J(z) \;=\; \begin{bmatrix} 0 & -z_3 & z_2 \\ z_3 & 0 & -z_1 \\ -z_2 & z_1 & 0 \end{bmatrix}, \qquad J(z)^\top = -J(z), \quad (2)$$

and read

$$\dot{z} \;=\; J(z)\,\nabla E(z), \quad (3)$$

with J antisymmetric and linear in the state,

$$J(z) = \sum_{k=1}^{3} z_k\, c^k, \qquad c^k_{ij} = -\varepsilon_{ijk}, \quad (4)$$

for ε the permutation symbol. The nine free entries of the three constant matrices c^k hold everything about the body but its inertia.

Antisymmetry conserves energy: $\dot{E} = \nabla E^\top J(z) \nabla E = 0$ for *every* antisymmetric J and every E, so a fitted model conserving its own energy is no evidence its c is correct. Conservation of $\|z\|^2$ is not automatic: it holds because $z^\top J(z) = 0$ for this particular c, confining the trajectory to a momentum sphere.

B. Affine brackets

The rigid body is one of a family: fix a state dimension d and consider

$$\begin{aligned} \dot{z} &= J(z)\nabla E(z), \qquad E(z) = \tfrac{1}{2} z^\top A z, \\ J(z) &= J_0 + \sum_{k=1}^{d} z_k\, c^k, \end{aligned} \quad (5)$$

with $J_0^\top = -J_0$, $(c^k)^\top = -c^k$ for each k, and $A = A^\top$ invertible. Call J an *affine bracket*, $c = (c^1, \ldots, c^d)$ its *structure constants*. The unknowns are J_0, with $\frac{d(d-1)}{2}$ free entries, and c, with $\frac{d^2(d-1)}{2}$: twelve parameters at $d = 3$ and 252 at $d = 8$.

Substituting $\nabla E = Az$ into (5),

$$\dot{z}_i = \sum_j (J_0)_{ij} (Az)_j + \sum_{j,k} c^k_{ij}\, z_k\, (Az)_j, \quad (6)$$

so c^k_{ij} is the gain with which state k times the generalized velocity $(Az)_j = \partial_j E$ drives state i. Antisymmetry in (i, j) makes every such term an *exchange*: what coordinate i gains through channel (i, j), coordinate j loses. For the rigid body, (4) trades angular momentum cyclically among the three axes, the rate for any two set by the third component. In general, c is what an engineer wants back from data, the part of the model unchanged by a change of energy, operating point or excitation.

Two special cases bracket the family. With $c = 0$, (5) is a constant-bracket Hamiltonian system with structure matrix J_0, classical mechanics in position–momentum coordinates once J_0 is nondegenerate and in canonical form. With $J_0 = 0$ it is a *Lie–Poisson* system, where the rigid body sits. Under the Jacobi conditions of Section II-E the general member is an

affine Poisson structure in Weinstein's sense [34]; we drop "Poisson" when Jacobi is not assumed, as Theorem 1 never uses it. The affine form also covers a Lie–Poisson system in coordinates centered on an operating point: substituting $z \mapsto z + z^0$ into $\sum_k z_k c^k$ produces a constant $J_0 = \sum_k z_k^0 c^k$, so keeping J_0 covers a sensor origin away from equilibrium.

Systems of the form (5) are common in engineering: spacecraft attitude and dual-spin dynamics, Lie–Poisson on $\mathfrak{so}(3)$ [6], [35]; the heavy top and a rigid body in a fluid, on the six-dimensional $\mathfrak{e}(3)$ [6], the neutrally buoyant underwater vehicle its $\mathfrak{se}(3)$ case, analyzed for stability in these coordinates [36]; ideal fluids, magnetohydrodynamics and the Vlasov–Maxwell system, whose noncanonical Poisson brackets keep this shape in infinite dimensions [37]–[39]; port-Hamiltonian networks, whose interconnection is a state-dependent antisymmetric structure matrix beside a dissipation term [4]; and projection-based reduced-order models of quadratic fluid systems, whose energy-preserving quadratic operator is the c-block of (5) [11], [40]. The bracket is usually known from first principles; the identification problem arises where it is not: reduced-order models fitted to simulation or experiment, latent coordinates from a learned encoder, and systems whose coupling topology is itself under study.

C. Identification is a linear least-squares problem

Suppose we observe S pairs $(z^{(s)}, y^{(s)})$, $y^{(s)}$ a measured derivative or a finite difference along a sampled trajectory, with the energy model A known or estimated separately. Then (5) is linear in the unknowns. Collect the free entries of (J_0, c) into a vector $u \in \mathbb{R}^p$ with

$$p = \tfrac{d(d-1)}{2} + \tfrac{d^2(d-1)}{2}, \tag{7}$$

build the design matrix $M \in \mathbb{R}^{Sd \times p}$ whose columns are the fields of the elementary antisymmetric basis elements at the sampled states, and stack the observations into $y \in \mathbb{R}^{Sd}$. Recovery is then

$$\hat{u} = \arg\min_u \|Mu - y\|_2^2, \tag{8}$$

one linear least-squares problem, solvable by a single singular value decomposition, with no local minima and textbook error analysis [41], the move behind sparse regression of dynamics [7] and operator inference [11], [40]. The usual diagnostics endorse it: the residual falls with more data and the model rolls out accurately. The obstruction sits upstream, in (8) itself: for every $d \geq 3$, M is column-rank deficient, the deficiency independent of the noise level, the solver, and S once the states are numerous and generic (Lemma 1).

D. The failure, measured

We ran (8) blind on two ground-truth brackets, the rigid body at $d = 3$ and the eight-dimensional $\mathfrak{su}(3)$ (Section II-E), blind meaning the solve gets the states, fields and energy model and nothing about the truth. RK4 generated trajectories in segments of 12 states from n random initial conditions, $n \in \{8, 32, 128, 512\}$, over three seeds. Two quantities are reported: the conserving-rollout relative error a practitioner would monitor, from held-out initial conditions under a discrete-gradient integrator conserving the fitted energy at every step by construction; and the cosine between fitted and true structure constants, what the practitioner wants.

On $\mathfrak{su}(3)$ the rollout error falls from 0.14–0.28 at eight orbits to 2–$6 \cdot 10^{-3}$ at 512, while from 128 orbits on the cosine stops at 0.2355–0.3094, within 10^{-4} of the ceilings computed from each draw's energy. The $64\times$ budget increase improves prediction by a factor of 29 to 67 across the three draws and moves identification only to its precomputed ceiling. The rigid body behaves the same: its cosine stays within $1.2 \cdot 10^{-4}$ of its per-draw ceilings (0.2620/0.2523/0.3347) at every budget. The cap is a property of the estimator's reach, not of the truth: on these three draws the 56-dimensional invisible subspace counted below holds between 0.951 and 0.972 of the true $\mathfrak{su}(3)$ structure-constant norm, the complement of those ceilings (Section III-E), and the solve returns the one coset representative with no component in it. The fraction is a property of the draw, and other draws give other values.

When A is unknown, the energy must be co-estimated with the bracket; we alternate least squares over (J_0, c) and the energy parameters θ, at ridge $3 \cdot 10^{-3}$. Co-estimation leaves the split of overall scale between c and θ undetermined, which the cosine is blind to: it normalizes both arguments and is reported in absolute value. On the same grid it scatters from 0.565 to 1.000 at $d = 3$ and 0.001 to 0.102 at $d = 8$, with no trend in budget: no longer confined to a fixed complement of the gauge, the fit lands on an arbitrary representative of a larger solution set. A cosine of 0.003 in $\mathbb{R}^{224}$ is a factor of roughly 18 below the 0.053 two independent random unit vectors average in absolute value. A cosine of 1.00 at $d = 3$ is the same arbitrariness with better luck; neither number measures recovery (Section III-E).

The design reports the reason: at $n = 512$ initial conditions M has rank 196 of 252 columns at relative tolerance 10^{-6}, the remaining 56 directions, 22% of the parameter space, changing (J_0, c) without changing the sampled field. At $d = 3$ the same measurement gives rank 11 of 12: one invisible direction, 8%, already capping recovery. Under the inertia $I = \mathrm{diag}(1, 2, 3)$ and noiseless field data, minimum-norm least squares returns structure constants at cosine 0.420560 to the truth, agreeing to six digits with the closed form Corollary 1 computes from (d, A) and the hypothesized bracket, before any data are collected.

The same run rules out three explanations. *More data* does not lift the ceiling: the rank deficiency of M is independent of S, and the blind ladder above moves the cosine to the ceiling and no further. *A better optimizer* does not help: (8) is an exact linear solve with no landscape to get stuck in. *Noise or conditioning* is not the cause: these data are noiseless, and in the dimension sweep of Section V-B, with the same design built at generic states, fifteen orders of magnitude separate the smallest retained and largest discarded singular values of M, so the deficiency is an integer and not a tolerance decision. Regularization is a remedy rather than an explanation, and it does not help either: ridge or minimum-norm least squares still returns a member of the solution set, chosen by a norm criterion unrelated to the physics, so the choice is arbitrary in the directions the data do not constrain (Corollary 1).

What is left is a kernel in the parameters-to-dynamics map, computed in closed form in Section III.

E. The vocabulary: Lie algebras, Jacobi, Killing

1) Lie algebras and structure constants: A *Lie algebra* is a vector space $\mathfrak{g}$ with a bilinear map $[\cdot,\cdot] : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$ antisymmetric, $[x, y] = -[y, x]$, and satisfying the Jacobi identity below. A basis $e_1, \dots, e_d$ turns the bracket into a table of numbers, the *structure constants*,

$$[e_i, e_j] = \sum_k c_{ij}^k \, e_k. \tag{9}$$

The canonical example is $\mathbb{R}^3$ with the cross product, structure constants $c_{ij}^k = \varepsilon_{ijk}$ in the standard basis; this algebra is $\mathfrak{so}(3)$, the brackets of infinitesimal rotations, and (4) carries the opposite sign, the usual body-frame convention, absorbed into c. The bracket J of (5) induces $\{f, g\} = \nabla f^\top J \nabla g$ on functions of the state, giving on the coordinate functions

$$\{z_i, z_j\} = J(z)_{ij} = (J_0)_{ij} + \sum_k c_{ij}^k z_k. \tag{10}$$

The model's array c *is* a table of structure constants in the sense of (9). The algebra $\mathfrak{su}(3)$ is the 3×3 traceless anti-Hermitian matrices under $[X, Y] = XY - YX$; its $8 \times 8 \times 8$ table is standard, generated in the accompanying implementation from the Gell-Mann matrices rather than transcribed. Nothing in (5) forces c to satisfy Jacobi.

Two conditions, not one, make (5) a Poisson system. Expanding the cyclic sum $\sum_{\circlearrowleft} \{z_i, \{z_j, z_k\}\}$ by degree in z, the degree-one part is the Jacobi identity on c studied throughout this article, and the degree-zero part couples the constant block to it,

$$\sum_\ell \big((J_0)_{i\ell} c_{jk}^\ell + (J_0)_{j\ell} c_{ki}^\ell + (J_0)_{k\ell} c_{ij}^\ell \big) = 0, \tag{11}$$

which says J_0 is a scalar two-cocycle of the algebra c. An arbitrary skew J_0 does not satisfy (11): paired with the $\mathfrak{su}(3)$ constants, a generic one leaves a residual of 27.8 where the degree-one part sits at $5 \cdot 10^{-16}$, and only 8 of the 28 skew parameters are compatible, the coboundaries produced by recentering, as Whitehead's lemma requires for a semisimple algebra. The identifiability results below do not depend on (11), since Theorem 1 never uses Jacobi and J_0 is identifiable in any case, but it does bound the scope of the word Poisson and of the retraction, which moves c alone. Every Lie-algebra experiment here takes $J_0 = 0$, where (11) holds trivially; a deployment estimating a nonzero J_0 and wanting the pair to be Poisson must also enforce it, linear in J_0 at fixed c and, unlike the retraction, not free of the data (Section VI-D). There (11) constrains the pair further than $\mathcal{S}$ and the certificate account for, so both over-count the surviving invisible directions and the routing errs toward the second energy. The supplement carries the certificate to the pair, one further linear block that never loosens the test: vacuous at $d = 3$, it cuts the count at $\mathfrak{su}(3)$, $A = I$, from 21 to 3 for generic compatible J_0, rising to 5 and 7 exactly on the isospin and hypercharge strata.

2) The Jacobi identity and what it guarantees: Define, for an antisymmetric bilinear c, the trilinear map

$$\mathrm{Jac}(c)(x, y, z) = c\big(c(x, y), z\big) + c\big(c(y, z), x\big) + c\big(c(z, x), y\big), \tag{12}$$

and call c a Lie bracket when $\mathrm{Jac}(c) = 0$. For the cross product this is the vector identity $(x \times y) \times z + (y \times z) \times x + (z \times x) \times y = 0$, so the rigid body satisfies it by inspection. Equation (12) is *quadratic* in c, so the set $\mathcal{L}_d = \mathrm{Jac}^{-1}(0)$ of Lie brackets is an algebraic variety inside $\mathbb{R}^{d^2(d-1)/2}$ rather than a linear subspace: Jacobi cannot be appended to (8) as a linear constraint, and the repair in Section IV is a retraction rather than a projection.

Jacobi supplies what antisymmetry does not (Section II-A): it organizes the functions C with $J(z)\nabla C(z) = 0$, conserved under *every* energy, so the state space splits into invariant leaves on which the motion is symplectic and these Casimirs are constant [42]; for the rigid body, $\|z\|^2$ is the Casimir and the momentum spheres the leaves. A fitted bracket violating Jacobi has no such foliation, so its long-horizon rollouts have no reason to stay on the manifold the real system never leaves and geometric integrators lose their motivating guarantees [5], [6]. The isomorphism class of $\mathfrak{g}$ is defined only when Jacobi holds.

Jacobi is checkable from a fitted $\hat{c}$ alone. Every test in this article compares one statistic against a tolerance,

$$\mathcal{J}(c) \;=\; \frac{\mathrm{RMS}\big(\mathrm{Jac}(c)\big)}{\|c\|_F^2}, \tag{13}$$

the root mean square taken over all d^4 entries of (12), redundant permutations included, against the squared Frobenius norm of c. The root mean square rather than a sum makes the null scaling of Section V-H a statement about d rather than about how many entries were counted.

3) The Killing form as a basis-independent fingerprint: Three quantities computable from any fitted $\hat{c}$ do not change under a linear change of state coordinates: the signature of the *Killing form* $\kappa(x, y) = \mathrm{tr}(\mathrm{ad}_x \mathrm{ad}_y)$, which transforms by congruence, and the dimensions of the derived algebra $[\mathfrak{g}, \mathfrak{g}]$ and of the center. Two fitted brackets whose values differ cannot be the same system in different coordinates; equal values do not prove they are, so the triple is a necessary condition and what Section V reports. The same form supplies the words the battery is built from, *nilpotent*, *solvable*, *semisimple* and *compact*, which the supplement defines and computes.

4) What "up to a change of basis" means: Coordinates are a modeling choice. Substituting $z = Sz'$ with S invertible turns (5) into a system of the same form with

$$\begin{aligned} J_0' &= S^{-1} J_0 S^{-\top}, \qquad A' = S^\top A S, \\ c'^{\,m} &= S^{-1} \Big(\sum_k S_{km} c^k \Big) S^{-\top}. \end{aligned} \tag{14}$$

Two parameter triples related by (14) are the same physical system in different state coordinates: no experiment distinguishes them and every prediction transforms correctly. The honest recovery target is therefore the orbit of the true c^* under (14), the set of all tables a change of basis carries it to, the isomorphism class of the algebra and what the invariants

above probe. If A is held fixed because the energy model is not refitted, the admissible S are those with $S^\top A S = A$, a smaller but still nontrivial group.

This gives the two kinds of non-uniqueness to keep apart for the rest of the article. *Reparameterization* is the benign kind: the fit returns a different basis for the same algebra, and the invariants agree. *Impersonation* is the other kind: the fit returns a bracket lying on the Lie variety, reproducing the observed dynamics to machine precision, and *not* isomorphic to the truth, its invariants having moved. Both appear in our thirteen-algebra battery, and no residual, conservation check or Jacobi certificate separates them after the fact. The certificate of Section III-G decides which regime a given problem is in, from the candidate bracket and the energy alone and without trajectory data.

F. Problem statement

Problem 1 (Certified structure identification)**.** Given a state dimension d, a quadratic energy model $E(z) = \frac{1}{2} z^\top A z$ with A symmetric and invertible, and S samples $(z^{(s)}, y^{(s)})$ with $y^{(s)} = J^*(z^{(s)}) A z^{(s)} + n^{(s)}$ generated by an unknown affine bracket $J^* = J_0^* + \sum_k z_k c^{*k}$ with $c^* \in \mathcal{L}_d$ and $n^{(s)}$ zero-mean noise of scale σ, produce estimates $(\hat{J}_0, \hat{c})$ with a certificate, such that

- (P1) the fitted field $\hat{J}(z) A z$ agrees with the observed field to the accuracy the noise permits;
- (P2) the estimate $\hat{c}$ lies on the variety $\mathcal{L}_d$ itself, satisfying the Jacobi identity outright and not within a penalty weight of it;
- (P3) the estimate $\hat{c}$ lies in the orbit (14) of the true c^*, so the recovered system is the true system in possibly different coordinates;
- (P4) the certificate states whether (P3) is attainable at all, judged from (d, A) and the fitted candidate alone and without reference to c^*.

Property (P1) is what (8) already delivers and the blind ladder shows to be insufficient. Properties (P2)–(P4) are the content of this article, developed in turn in Sections III, III-F, III-G and IV.

III. IDENTIFIABILITY THEORY

For the model class of Section II under a quadratic energy, the parameter directions no amount of trajectory data can resolve form a subspace computed from the state dimension and the energy Hessian alone, before an experiment is designed. Table I collects the objects this section and the experiments use.

A. The evaluation map

Recall the affine bracket on $\mathbb{R}^d$,

$$J(z) \;=\; J_0 + \sum_{k=1}^{d} z_k \, c^k, \qquad J_0^\top = -J_0, \quad (c^k)^\top = -c^k, \tag{15}$$

with dynamics $\dot{z} = J(z) \nabla E(z)$. The unknowns are the constant block J_0 and the state-dependent block $c = (c_{ij}^k)$, carrying $\frac{d(d-1)}{2}$ and $\frac{d^2(d-1)}{2}$ parameters, in total $P(d) = \frac{d(d-1)}{2} + \frac{d^2(d-1)}{2}$, the p of (7); at $d = 8$, $P(8) = 28 + 224 = 252$.

TABLE I
THE RECURRING OBJECTS, EACH WITH ITS PLAIN READING. THE COHOMOLOGICAL ROWS ARE STANDARD; THE GAUGE K_A, THE HOOK PART, r_{Jac} AND ι_H ARE THIS ARTICLE'S, DEFINED IN SECTIONS II AND III WHERE EACH FIRST MATTERS.

symbol	plain reading
$z \in \mathbb{R}^d$	the measured state (angular momentum, modes)
$E = \frac{1}{2} z^\top A z$	the quadratic energy; A its Hessian
J_0, c	constant and state-dependent bracket parts
$\mathcal{L}_d$	arrays satisfying the Jacobi identity
K_A (gauge)	coefficient changes invisible to data, given A
$\Lambda^3 \mathbb{R}^d$	totally antisymmetric three-index arrays
hook part	complement of Λ^3; seen by data only at $A = I$
Z^2 (cocycles)	directions keeping Jacobi to first order
B^2 (coboundaries)	directions that only re-label the basis
$H^2 = Z^2 / B^2$	deformations with basis changes removed
$r_{\mathrm{Jac}} = \dim(Z^2 \cap K_A)$	invisible Jacobi-compatible directions
ι_H	those that are not basis changes; grades the cost
$\mathcal{S}$	brackets data cannot tell from the truth

The rigid body of Section II-A, the worked example throughout, is the $d = 3$ instance: $J_0 = 0$ and $c_{ij}^k = -\varepsilon_{kij}$, the body-frame sign of (4), turn (15) into $\dot{z} = z \times \mathbb{I}^{-1} z$ with inertia matrix $\mathbb{I}$.

For $E(z) = \frac{1}{2} z^\top A z$ with A symmetric, $\nabla E = A z$, and the *evaluation map* from parameters to their field,

$$\Phi_A(J_0, c)(z) \;=\; \Big(J_0 + \sum_k z_k c^k \Big) A z \tag{16}$$

is *linear* in the unknowns, the reason a complete answer exists. Two parameter vectors give the same field at every state, hence the same trajectory from every initial condition, if and only if their difference lies in the nullspace $\ker \Phi_A$, the *gauge* of E. Independent of the observed states, it cannot be shrunk by more initial conditions, longer rollouts or finer time steps.

Lemma 1 (Sampling)**.** *Let* $\Phi_A^{(S)}$ *be* (16) *evaluated at states* $z^{(1)}, \ldots, z^{(S)}$*. Then* $\ker \Phi_A^{(S)} \supseteq \ker \Phi_A$ *for every sample, with equality for all* $S \geq S_0(d)$ *outside a proper Zariski-closed set of state configurations.*

Proof sketch. Appending states shrinks the kernel monotonically, and whenever the sampled kernel strictly contains $\ker \Phi_A$ some state cuts it further, so equality follows after finitely many states; entries of $\Phi_A^{(S)}$ are polynomial in the states, so the rank-drop locus is the zero set of maximal minors. The supplement gives the full argument. □

This licenses reading an integer off a singular value decomposition: the supplement's sweep has a spectral gap of fifteen orders of magnitude, so reported nullities are integers, not thresholds.

B. The gauge at $d = 3$

At $d = 3$ the totally antisymmetric 3-tensors are the one-dimensional span of ε. Under the isotropic energy $A = I$, setting $c_{ij}^k = t \, \varepsilon_{kij}$ makes component i of the field

$\sum_{k,j} t\,\varepsilon_{kij} z_k z_j$ vanish, an array antisymmetric in (k,j) contracted against the symmetric $z_k z_j$: the whole line $t \mapsto t\varepsilon$ is invisible, the *overall scale of the cross product* cannot be recovered, and the true bracket lies on it. This is the spherical top: $I_1 = I_2 = I_3$ leaves the angular velocity parallel to the angular momentum and $\dot{z} = z \times Az$ vanishes.

For $A = \mathbb{I}^{-1} = \mathrm{diag}(a)$, $a_k = 1/I_k$, the same cancellation holds for the weighted tensor $u^k_{ij} = a_k \varepsilon_{kij}$, since $\sum_{k,j} a_k \varepsilon_{kij} z_k\, a_j z_j$ contracts ε against the symmetric $(a_k z_k)(a_j z_j)$, so the invisible direction rotates from ε to $a \odot \varepsilon$, still one-dimensional. Adding $t\,(a \odot \varepsilon)$ sends $\varepsilon^k \mapsto (1 + t\,a_k)\varepsilon^k$, a different rigid body with identical trajectories from every initial condition; at $d = 3$ this line is classical (Section I-B).

C. Theorem 1: the gauge for every invertible Hessian

The $d = 3$ pattern is the general one: for any invertible Hessian the invisible directions are a twisted copy of the totally antisymmetric 3-tensors.

Theorem 1 (The gauge)**.** *Let A be symmetric and invertible. Then*

$$\begin{aligned} \ker \Phi_A \;&=\; 0 \oplus K_A, \\ K_A &= \big\{ c :\; c^k_{ij} = \textstyle\sum_m A_{km}\, g_{mij},\; g \in \Lambda^3 \mathbb{R}^d \big\}, \end{aligned} \tag{17}$$

where $\Lambda^3\mathbb{R}^d$ is the space of 3-tensors changing sign under every transposition of their slots. Consequently J_0 is identifiable and, for every *invertible A,* $\dim K_A = \binom{d}{3}$ *and*

$$\mathrm{rank}\, \Phi_A = \frac{d(d-1)}{2} + \frac{d(d^2-1)}{3}. \tag{18}$$

Proof. Write $C(z) = \sum_k z_k c^k$, so $\Phi_A(J_0, c)(z) = J_0 Az + C(z)Az$, two terms homogeneous of degree 1 and 2 in z. A polynomial map vanishes identically if and only if each homogeneous component does, decoupling the constant and state-dependent conditions.

Step 1 (degree one). The constant block is fully visible: $J_0 Az \equiv 0$ gives $J_0 A = 0$ at $z = e_1, \ldots, e_d$, and A is invertible, so $J_0 = 0$.

Step 2 (degree two, isotropic core). At $A = I$, invisibility amounts to one extra sign symmetry. Write $t_{kij} := c^k_{ij}$. Component i of the quadratic field $C(z)z$ is the quadratic form of $M^{(i)}_{kj} = t_{kij}$, so the field vanishes at every state if and only if all d forms do, and a real quadratic form vanishes identically if and only if its matrix is antisymmetric, $M^{(i)} + M^{(i)\top} = 0$, that is

$$t_{kij} + t_{jik} = 0 \qquad \text{for all } i, j, k, \tag{19}$$

sign reversal under the transposition $(1\,3)$ of its slots; the model class already gives $(2\,3)$, since $c^k_{ij} = -c^k_{ji}$.

Step 3 (two transpositions generate S_3). Applying $(1\,3)$ and then $(2\,3)$,

$$t_{kij} \;=\; -\,t_{jik} \;=\; +\,t_{jki}, \tag{20}$$

so t is invariant under the cyclic shift $(k,i,j) \mapsto (j,k,i)$. Using (20) twice and then $(2\,3)$,

$$t_{ikj} \;=\; t_{jik} \;=\; t_{kji} \;=\; -\,t_{kij}, \tag{21}$$

which is sign reversal under the remaining transposition $(1\,2)$. All six elements of S_3 are then fixed: the identity trivially; the three transpositions with -1, one each by (19), the model class and (21); the two 3-cycles with $+1$ by (20). Hence $t \in \Lambda^3\mathbb{R}^d$. Conversely, for $t \in \Lambda^3\mathbb{R}^d$ and each i the array t_{kij} is antisymmetric in (k,j), contracted against the symmetric $z_k z_j$, so the field vanishes and $\ker \Phi_I = 0 \oplus \Lambda^3\mathbb{R}^d$.

Step 4 (general invertible A). Substitute $w = Az$, a bijection of $\mathbb{R}^d$ reducing the general case to the isotropic one. Since $C(A^{-1}w) = \sum_m w_m \sum_k (A^{-1})_{km} c^k$,

$$C(z)Az \;=\; C(A^{-1}w)\, w \;=\; \sum_m w_m\, g^m\, w, \tag{22}$$

where $g^m := \sum_k (A^{-1})_{km} c^k$. Each g^m combines the antisymmetric c^k, hence is antisymmetric, so g lies in the same model class and $\Phi_A(0,c)(z) = \Phi_I(0,g)(Az)$. As $z \mapsto Az$ is onto, $\Phi_A(0,c) \equiv 0$ if and only if $\Phi_I(0,g) \equiv 0$, which by Steps 2–3 holds if and only if $g \in \Lambda^3\mathbb{R}^d$. Inverting the slice-index map with A symmetric gives (17), $c^k = \sum_m A_{km} g^m$. The map $c \mapsto g$ is a linear isomorphism of the c-block, so $\dim K_A = \binom{d}{3}$.

Step 5 (rank). By rank–nullity $\mathrm{rank}\, \Phi_A = P(d) - \binom{d}{3}$, and the c-block difference simplifies: $\frac{d^2(d-1)}{2} - \frac{d(d-1)(d-2)}{6} = \frac{d(d-1)}{6}(3d - (d-2)) = \frac{d(d^2-1)}{3}$. □

Remark 1. The proof uses only invertibility of $z \mapsto Az$, so for invertible non-symmetric A it returns the twisted gauge with A_{km} replaced by A_{mk}, of the same dimension. Symmetry is imposed only because $\nabla E = Az$ requires a Hessian; the non-symmetric case is verified numerically at $d \in \{6, 8\}$ (nullity 20 and 56). Invertibility can be relaxed by one: at corank one the theorem holds verbatim, because $J_0 A = 0$ would force a skew matrix of rank one and skew rank is even. From corank two the kernel *grows*, by an amount measured in Section V-B rather than proved for general rank. Finally $K_A = \Lambda^3\mathbb{R}^d$ if and only if A is a multiple of the identity: anisotropy re-aims the gauge and never reduces it.

D. The dimension arithmetic

Of the $P(d)$ parameters, $\binom{d}{3}$ are invisible and $\frac{d(d-1)}{2} + \frac{d(d^2-1)}{3}$ identifiable: at $d = 8$, 56 against 196, with the c-block splitting as $224 = 168 + 56$. The blind spot does not wash out with scale, the invisible fraction of the c-block being $(d-2)/(3d)$, increasing to $1/3$. The supplement checks both formulas against the measured nullity and rank at generic Gaussian states. Under simultaneous change of basis the c-block splits into two inequivalent irreducible pieces, the totally antisymmetric part and its orthogonal complement, the *hook* part. Inequivalence makes them orthogonal for every invariant inner product, so at $A = I$ the projection onto the gauge is plain antisymmetrization and the identifiable component has an explicit orthonormal basis, the one Section IV fits in; the supplement gives that decomposition and its dimension count.

E. The forced ceiling

Theorem 1 fixes in advance how close any gauge-respecting estimator comes to the truth, in the cosine between estimated and true structure constants.

Corollary 1 (Spherical-top degeneracy and the forced ceiling)**.** *For a compact semisimple Lie algebra* expressed in a basis orthonormal for the negative Killing form*, the isotropic energy* $E = \frac{1}{2}\|z\|^2$ *is the quadratic Casimir and the Lie–Poisson field vanishes identically (the spherical top in* d *dimensions), since ad-invariance,* $\kappa([x,y],z) = -\kappa(y,[x,z])$ *[43], makes the structure constants totally antisymmetric, hence elements of* K_I*. Let noiseless field data come from* (J_0^*, c^*) *under* $E = \frac{1}{2}z^\top A z$ *with* A *invertible, at* $S \geq S_0(d)$ *states in general position so that Lemma 1 applies, and* $(\hat{J}_0, \hat{c})$ *the minimum-norm least-squares estimate in the Frobenius inner product. Then*

$$\hat{J}_0 = J_0^*, \quad \hat{c} = P_{K_A^\perp} c^*, \quad \cos(\hat{c}, c^*) = \frac{\|P_{K_A^\perp} c^*\|}{\|c^*\|}, \tag{23}$$

a number computable before data collection, which also upper-bounds $\cos(\hat{c}, c^*)$ *for* any *estimator returning an element of* $K_A^\perp$*. If* c^* *are the constants of a compact semisimple algebra in a Killing-orthonormal basis and* $A = I$*, the spherical top above, that ratio is* 0*: the estimate itself vanishes, so no such estimator recovers any component of* c^**, and the cosine is undefined there rather than small. The qualifier is on the basis: total antisymmetry is a property of the presentation, not of the algebra.*

Proof sketch. The least-squares solution set is the coset $u^* + \ker M$ with $\ker M = 0 \oplus K_A$, whose minimum-norm element is the orthogonal projection onto $(\ker M)^\perp$; since $\ker M$ sits inside the c-block this gives (23). Here M is the full design of (8), not the reduced design of Section IV-E, which has trivial kernel by construction. Cauchy–Schwarz extends the bound to any $\hat{c} \in K_A^\perp$, and noise does not lift it, since $M^+ n \perp \ker M$ keeps the estimate there. The supplement gives the computation. □

For the rigid body (23) has a closed form in the moments of inertia alone, vanishing for the spherical top and growing with asymmetry: at $\mathbb{I} = \mathrm{diag}(1,2,3)$ it predicts 0.420560 against a measured 0.420560, and two further cells agree to six digits. The supplement carries the derivation.

A structure-constant cosine carries no recovery information on its own, in either direction: a low value may be the forced ceiling of a perfectly consistent fit, a high value the work of an estimator returning a lucky coset representative. The ceiling is also per cell rather than per algebra. In the prespecified experiments of Section V-D, the fraction of true $\mathfrak{su}(3)$ constants inside the twisted gauge, $\|P_{K_A} c^*\|/\|c^*\|$, ranges from 0.928 to 0.949 across the certificate battery's three energy draws; for $\mathfrak{so}(3)$, 0.917 to 0.942. Those draws differ from the recovery arms', so these fractions do not pair with the ceilings above.

F. Theorem 2: which second experiment removes the gauge

Since the gauge depends on the energy, a second energy has a second gauge, and only directions invisible under both stay invisible; one spectrum from the two Hessians fixes that intersection.

Theorem 2 (Intersection law)**.** *Let* $A = \mathrm{diag}(a)$ *and* $B = \mathrm{diag}(b)$ *be invertible, set* $r_k = a_k/b_k$*, and let* $L_1, \ldots, L_\ell$ *be the level sets of* r *(the groups of coordinates sharing a ratio). Then* $K_A \cap K_B$ *is spanned by the* $a \odot e^\tau$ *with* τ *a* 3*-element subset of a single level set, where* e^τ *is the elementary totally antisymmetric tensor on* τ *and* $(a \odot e)^k_{ij} = a_k e_{kij}$*, so*

$$\dim\big(K_A \cap K_B\big) \;=\; \sum_{t=1}^{\ell} \binom{|L_t|}{3}. \tag{24}$$

In particular $b \propto a$ *leaves the gauge at* $\binom{d}{3}$*, and pairwise distinct ratios annihilate it.*

Proof sketch. A direction in both gauges satisfies $c^k_{ij} = a_k u_{kij} = b_k v_{kij}$ with u, v totally antisymmetric, so $v = r \odot u$; antisymmetry of v under two transpositions then forces $(r_k - r_i)u_{kij} = 0$ and $(r_k - r_j)u_{kij} = 0$, so a nonzero entry needs $\{k,i,j\}$ inside one level set, and each such triple contributes one direction. The supplement gives the computation and, for non-diagonal pairs with A positive definite, the reduction to Theorem 2 with level sets over the generalized eigenvalues of the pencil (B, A); that reduction is proved and not exercised numerically, since the measured cells are diagonal pairs only. □

Equation (24) is a design rule: repeat the experiment under a second quadratic energy and inspect the ratio spectrum $r_k = a_k/b_k$. Ties of size two cost nothing, since $\binom{2}{3} = 0$; the hazard is three or more coordinates sharing a ratio, leaving $\binom{|L|}{3}$ invisible directions. On the rigid body this is a second run with a different inertia tensor, its principal ratios against the first not all equal; since $d = 3$ carries only $\binom{3}{3} = 1$ gauge direction, one broken tie removes it entirely. The law is tested at $d = 6$ and, on $\mathfrak{su}(3)$, at $d = 8$ in Section V-E: a proportional second energy at matched budget leaves the cosine unchanged to twelve digits, and pairwise distinct ratios raise it to 1.0000.

G. The certificate

A second source of information costs no data at all: the Jacobi identity. Moving along an invisible direction resolves the ambiguity only if the moved bracket stops being a Lie algebra, so the question is how much of K_A stays compatible with Jacobi. With $\mathrm{Jac}(c)(x,y,z) = \sum_{\circlearrowleft} c(c(x,y),z)$ the cyclic sum, the Lie algebra structures form the variety $\mathcal{L}_d = \mathrm{Jac}^{-1}(0)$: quadratic equations in c, not a subspace. Fix $f \in \mathcal{L}_d$ and write $\mathfrak{g} = (\mathbb{R}^d, f)$. Since Jac is homogeneous quadratic it has an exact polarization, linear in the perturbation h,

$$D_f h \;:=\; \mathrm{Jac}(f+h) - \mathrm{Jac}(f) - \mathrm{Jac}(h), \tag{25}$$

whose kernel $Z^2 := \ker D_f$, the *cocycles*, are the directions along which $f + th$ satisfies Jacobi to first order. Inside them sit the *coboundaries* $B^2 := \{X \cdot f : X \in \mathfrak{gl}(d)\}$, with $(X \cdot f)^k_{ij} = \sum_l (X_{kl} f^l_{ij} - f^k_{lj} X_{li} - f^k_{il} X_{lj})$, the first-order motions produced by a change of basis of $\mathbb{R}^d$; the quotient $H^2(\mathfrak{g};\mathfrak{g}) := Z^2/B^2$ is what remains after dividing those out. Up to sign D_f is the Chevalley–Eilenberg differential with adjoint coefficients [44] and H^2 the Nijenhuis–Richardson deformation space [32], [33]; nothing from that theory beyond

these definitions is used. All three are dense linear algebra, the matrix of D_f being 4096×224 at $d = 8$. As a check, $\dim B^2 = d^2 - \dim \mathrm{Der}(\mathfrak{g})$, so $\mathfrak{su}(3)$, whose derivations are all inner and eight-dimensional, has $64-8 = 56$, the measured value.

Definition 1 (Certificate)**.** The *gauge certificate* of the pair (bracket, energy) counts the invisible directions that keep the bracket a Lie algebra to first order and are not changes of basis: with K_A from Theorem 1,

$$\iota_H \;:=\; \dim\big((Z^2 \cap K_A) + B^2\big) - \dim B^2. \tag{26}$$

Equivalently, ι_H is the dimension of the image of the invisible cocycles $Z^2 \cap K_A$ in H^2. It depends on (f, A) alone; no data enter.

Remark 2 (Which equivalence the certificate uses)*.* B^2 is built from all of $\mathfrak{gl}(d)$, so ι_H answers (P3): is the invisible motion a re-labeling of the abstract algebra? A reader delivering the coefficient array in fixed sensor coordinates, with the energy held at a known A, has the smaller symmetry group $\mathfrak{o}(A) = \{X : X^\top A + AX = 0\}$, of dimension $\frac{d(d-1)}{2}$ rather than d^2, and the two readings do not always agree. Substituting it leaves five of the eight battery rows with a nonzero intersection unchanged and raises the other three, $\mathfrak{so}(3)$, $\mathfrak{sl}(2,\mathbb{R})$ and $\mathfrak{so}(4)$, from $\iota_H = 0$ to 1, 1 and 2, which are the rows Section V-D calls reparameterization, whose invisible motion is a change of basis but not one preserving A. This article reports the $\mathfrak{gl}(d)$ certificate throughout, matching (P3); a deployment holding the stricter target should read the $\mathfrak{o}(A)$ column.

Since ι_H can exceed neither $\dim(Z^2 \cap K_A)$ nor $\dim H^2$, an algebra with $H^2 = 0$ has $\iota_H = 0$; every semisimple algebra is such, by Whitehead's lemma [45, Ch. III]. And $\iota_H = 0$ says $Z^2 \cap K_A \subseteq B^2$: every invisible first-order deformation preserving Jacobi is a change of basis. The informative case is the one that implication does not cover, algebras where H^2 is large and ι_H still zero.

The *invisible slice* $\mathcal{S} = (f + K_A) \cap \mathcal{L}_d$ collects every bracket satisfying the Jacobi identity *and* reproducing, without error, the dynamics of f under $E = \frac{1}{2} z^\top A z$, and its *tangent cone* $C_f(\mathcal{S})$ collects the directions along which other members approach f, the limits of rescaled secants $t_n(f_n - f)$ with $f_n \in \mathcal{S}$ and $f_n \to f$.

Proposition 1 (Isolation: sufficient, not necessary)**.** *Let $f \in \mathcal{L}_d$ and let A be invertible.*

(i) *$C_f(\mathcal{S}) \subseteq Z^2 \cap K_A$: every approach direction is an invisible cocycle.*

(ii) *If $Z^2 \cap K_A = 0$, then f is isolated in $\mathcal{S}$: under exact enforcement of the Jacobi identity, no other bracket near f reproduces the dynamics.*

(iii) *If $\iota_H = 0$, then $C_f(\mathcal{S}) \subseteq B^2 \cap K_A$: every invisible first-order motion along the variety is a change of basis.*

Each hypothesis is sufficient for the corresponding form of local identifiability and none is necessary: identifiability can also come from a second energy (Theorem 2), and $\iota_H > 0$ does not by itself produce a competing algebra.

Proof sketch. (i) An approach direction inherits both memberships: each secant lies in the closed subspace K_A, and $\mathrm{Jac}(f) = \mathrm{Jac}(f_n) = 0$ turns (25) into $D_f h_n = -\mathrm{Jac}(h_n)$, whose rescaled limit is annihilated by D_f. (ii) Contrapositive: a non-isolated f produces unit secants converging to a nonzero element of $Z^2 \cap K_A$, which the hypothesis forbids; $\mathcal{S}$ is the right object by Theorem 1. (iii) $\iota_H = 0$ if and only if $Z^2 \cap K_A \subseteq B^2$; combine with (i). The supplement gives all three in full. □

The three parts are not equally strong. When $Z^2 \cap K_A = 0$, part (ii) is local isolation in the full nonlinear sense, the case on every row the retraction identifies. Part (iii) is first order, and it is neither claimed nor proved that a whole neighborhood of f in $\mathcal{S}$ consists of algebras isomorphic to f. Where $\iota_H = 0$ with $Z^2 \cap K_A \neq 0$, finite-displacement behavior is measured instead (Section V-D), and is decisive: the invisible direction h satisfies $\mathrm{Jac}(h) = 0$ to machine precision, so by (25) the whole line $f + th$ lies in $\mathcal{S}$ and the question reduces to that line; at larger displacements those lines do cross isomorphism classes (Section V-D), routing such rows to the energy lever.

Section V-D shows numerically (26) is strictly finer than cohomology, the heavy top $\mathfrak{e}(3)$ recovered by the retraction despite carrying the classical Lévy–Nahas deformation class [46]. It is also not a function of the algebra alone: the energy enters through K_A, which H^2 ignores. Measured, $Z^2 \cap K_A$ is 21-dimensional for $\mathfrak{su}(3)$ at the isotropic energy ($A = I$) and 0 at every anisotropic draw, with $H^2 = 0$ throughout. In the experiments ι_H is evaluated at the true bracket. Deployment evaluates only r_{Jac}, and at the retracted candidate of Section IV: off the Lie variety $B^2 \subseteq Z^2$ can fail, so (26) is arithmetic there rather than a dimension in H^2, and the substitution of candidate for truth is tested only end to end.

IV. THE IDENTIFICATION PROCEDURE

Sections III and III-G become a procedure that certifies, fits in the identifiable quotient, then recovers the remainder along directions the data cannot see. Only the fitting stage needs data.

A. The algorithm

The procedure runs in two regimes, differing in what can be computed before data exist. Under *hypothesis testing* a candidate bracket f is available, from a first-principles model or from structure the practitioner wants to confirm; the certificate, a function of $(f, \mathcal{A})$ alone, is computable before a single trajectory is recorded and says whether the planned experiment could identify that structure if the system has it. It does not test whether the hypothesis is right; a wrong candidate is certified as readily as a right one, and what exposes it is the fit residual (Section V-G). Under *structure discovery* no candidate exists, the certificate has nothing to evaluate before the fit, and the procedure retracts first, evaluating it at the recovered bracket afterwards. Pre-data decisions claimed here belong to the first regime; the second is a post-fit check, its substitution of the recovered candidate for the truth tested only end to end (Section VI-D).

The supplement states the procedure as pseudocode and annotates it line by line. Two of its steps carry decisions worth

naming here. A multi-energy design's kernel is the intersection of the individual kernels, so the quotient basis depends on $(d, \mathcal{A})$ alone and is computed once per energy plan; and if no budget in the pilot ladder meets the target error, the procedure refuses rather than returning an unsupportable estimate.

B. The design matrix

The design matrix M maps the unknowns (J_0, c) to the predicted field, one row per (sample s, component i) pair, giving Sd rows. Order the columns as the $\binom{d}{2}$ pairs $(i<j)$ for J_0 then the $d\binom{d}{2}$ pairs $(k,(i<j))$ for c; write $g^{(s)} = Az^{(s)}$ for the energy gradient at s. The J_0 column (i,j) carries $g_j^{(s)}$ in row (s,i) and $-g_i^{(s)}$ in row (s,j); the c column $(k,(i,j))$, the same two entries scaled by $z_k^{(s)}$; all other entries vanish. The one convention to get wrong is the column order, which must match that for Λ^3 below. M has two nonzeros per column per sample, so at large S assemble it column-blockwise. The samples must also come from several initial conditions, a geometric requirement: one trajectory's states lie on a single invariant leaf (a coadjoint orbit) intersected with one energy level set. The kernel can survive that confinement, as it does on the $\mathfrak{su}(3)$ orbits of Section V at every budget, though the quotient margin drops to the numerical-rank floor (Section IV-J); on some systems it inflates (Section V-C).

C. The invisible subspace

Theorem 1 gives the invisible subspace in closed form, $K_A = \{c : c_{ij}^k = \sum_m A_{km} g_{mij},\ g \in \Lambda^3\mathbb{R}^d\}$. For a basis, enumerate the $\binom{d}{3}$ index subsets $\{p<q<r\}$, forming the array e^{pqr} for each: $+1$ on even permutations of (p,q,r), -1 on odd ones, zero elsewhere, flattened in the coefficient order above into a column of $L \in \mathbb{R}^{n\times\binom{d}{3}}$ with $n = d\binom{d}{2}$. Then

$$G_{\mathrm{raw}} = (A \otimes I_{\binom{d}{2}})\, L, \qquad G = \mathrm{orth}(G_{\mathrm{raw}}). \tag{27}$$

At $d=3$ with $A = \mathrm{diag}(a)$ a single such column results, $c_{ij}^k = a_k \varepsilon_{kij}$, and adding any multiple of it to the rigid body's structure constants changes nothing about the motion, for any inertia tensor. The rigid body is therefore a poor guide to the general case, where invisible directions grow as $d^3/6$.

G_{raw} is tall, since $n/\binom{d}{3} > 3$ for every d, and has full column rank, $A \otimes I_{\binom{d}{2}}$ being invertible with L's columns independent; the rank, rather than the tallness, legitimizes a reduced QR here, as the warning below shows. The factorization doubles as a rank test: a Hessian of corank two or more drops R's diagonal below tolerance and the routine should raise an error, the kernel growing rather than shrinking there. Cross-check the closed form once per implementation against the empirical kernel: build M on generic Gaussian states with Sd comfortably above the parameter count, take its SVD, count singular values below $10^{-8}\sigma_{\max}$. The two agree to the integer, a gap reading rather than a tolerance decision: the spectrum has a gap of some fifteen orders of magnitude at the cut, smallest retained near $2\cdot10^{-1}\sigma_{\max}$ and largest discarded near $10^{-15}\sigma_{\max}$ across $d = 3,\dots,12$. The empirical route is $O(S)$ and the closed form is not, so the latter is the default.

For the intersection step, stack the projectors onto the complements, $P = [\, I - G_1G_1^\top; \dots; I - G_NG_N^\top \,] \in \mathbb{R}^{Nn\times n}$, and take its nullspace at τ_{null}: a vector lies in every $\mathrm{span}\, G_i$ if and only if every block annihilates it. For an intersection *dimension* rather than a basis, principal angles are cheaper: the singular values of $G_A^\top G_B$ are the subspace-angle cosines; count those above $1-\tau_{\mathrm{ang}}$. They cluster at 1 or well below, so the count is insensitive to τ_{ang} over several decades.

Unpivoted reduced QR is not a rank-revealing basis: Where the input may be rank deficient, unpivoted reduced QR should not supply an orthonormal basis of its column space: `numpy.linalg.qr` returns $\min(m,n)$ columns regardless of rank, so the basis can span more than that space and every projection off it silently evaluates to zero. The implementation uses the SVD throughout. The supplement gives the two measured cases, one square and one tall, where this turned a correct reported fraction into 0.000.

D. The certificate

Definition 1 needs three objects: the polarized Jacobi operator D_f, with kernel the cocycle space Z^2; the coboundary space B^2, tangent at f to its orbit under change of basis; and the intersection $Z^2 \cap \mathrm{span}\, G$ with the invisible subspace. All three are dense linear algebra on matrices of size $\Theta(d^4)\times\Theta(d^3)$, with $\iota_H = \mathrm{rank}[\, C \mid B^2 \,] - \mathrm{rank}\, B^2$ for C a basis of that intersection. Nothing is differentiated numerically and no data enter. The supplement gives each operator's index conventions and elementary-array construction.

E. Least squares on the quotient

With Q the quotient basis, the reduced design is $D = [\, M_{J_0} \mid M_cQ \,]$ and $\hat{c} = Q\hat{x}_c$, a parameterization preferable on three counts to a full-space solve discarding small singular values. The design has full column rank, so its smallest singular value is the *quotient margin*, a meaningful conditioning number rather than an `rcond` artifact. The problem carries $\binom{d}{3}$ fewer unknowns, asymptotically one third of the c parameters. And noise cannot enter the invisible directions, absent from the model: $\hat{c}$ lies in $K_A^\perp$ at every noise level, so the ceiling of Corollary 1 is a fixed comparison rather than something a regularizer has moved (Section V-F reports how closely the reduced and full solves agree).

F. The Levenberg–Marquardt retraction

When the certificate permits, the remainder is recovered by sliding along invisible directions until the Jacobi identity holds. The search set is the invisible coset, $\hat{c}$ shifted by arbitrary elements of K_A, parameterized by one coefficient per invisible direction, $c(g) = \hat{c} + \sum_{t=1}^m g_t G_t$; minimize

$$\Phi(g) = \big\|\mathrm{Jac}\big(c(g)\big)\big\|_F^2, \qquad g\in\mathbb{R}^m, \quad g^{(0)} = 0. \tag{28}$$

Every $c(g)$ generates the same field as $\hat{c}$, so (28) cannot degrade the fit wherever it lands. The Jacobian is closed-form, $\partial\mathrm{Jac}(c(g))/\partial g_t = B(c, G_t) + B(G_t, c)$, needing no finite

differences; stacked as rows, those m arrays give $J \in \mathbb{R}^{m\times d^4}$; the step solves

$$\big(JJ^\top + \lambda I_m\big)\delta = -Jr, \qquad r = \mathrm{Jac}\big(c(g)\big) \text{ flattened.} \quad (29)$$

The implemented schedule starts at $\lambda = 10^{-3}$, accepting any step that strictly decreases Φ and dividing λ by three, otherwise tripling it and recomputing; the supplement gives the stopping tests. On $\mathfrak{su}(3)$ from a single anisotropic energy it stops at 8, 8 and 9 iterations on three pre-specified configurations, at Jacobi residual 1.6–$2.5 \cdot 10^{-16}$.

The coset parameterization excludes the radial direction c itself, keeping the retraction off the homogeneous cone down which full-space Newton steps slide toward the abelian bracket $c = 0$ (Section V-E measures that failure).

Acceptance tests the recovered $\hat{c}_J$ twice: the scale-normalized residual $\mathcal{J}(\hat{c}_J)$ must fall below τ_d and $\dim(\ker \mathrm{DJac}(\hat{c}_J) \cap K_A)$ must vanish, the second the certificate re-evaluated at the recovered bracket rather than at a truth that a deployment lacks. The threshold must be dimension-aware: the same normalized residual has median approximately $\sqrt{3}\, d^{-5/2}$ on random antisymmetric arrays, so a fixed 10^{-3} stops being a test between $d = 16$ and $d = 20$ (Section V). It must also be noise-aware: on noiseless data the retraction reaches 10^{-16} and a fixed 10^{-8} is the right test, but observation noise displaces the fitted quotient component off the variety by an amount no motion inside K_A undoes, so the noise rather than the optimizer sets the attainable residual. Measured on $\mathfrak{su}(3)$ at a fixed design over four decades of σ, it grows with slope 1.05 in σ ($R^2 = 0.996$) and stays under $4.1 \cdot 10^{-3}$ of the predicted relative quotient error on every cell, giving the calibrated test

$$\mathcal{J}(\hat{c}_J) \;\le\; \alpha\, \sigma \sqrt{\textstyle\sum_i s_i^{-2}}\,/\,\|\hat{u}\|, \quad (30)$$

with $\alpha = 10^{-2}$ leaving a factor of about 2.5 over the measurement.

The second test needs the same treatment, and needing it is easy to miss. Both integers it reads, the dimension of $\ker \mathrm{DJac}(\hat{c}_J)$ and how much of that kernel lies inside K_A, come from tolerance cuts: a relative singular value below 10^{-10} counts as zero, a principal-angle cosine above $1 - 10^{-8}$ counts as an intersection. Those numbers are right at an exact Lie bracket and wrong at a candidate fitted to noisy data, the only place the test is ever applied. On $\mathfrak{su}(3)$ at $\sigma = 10^{-3}$ the whole kernel lifts above the singular-value cut and its reported dimension falls from 56 to 0; on $\mathfrak{so}(3)$ the dimension survives but the kernel has rotated, the cosine falls below the angle cut, and the reported intersection falls from 1 to 0. Both failures report $r_{\mathrm{Jac}} = 0$, and $r_{\mathrm{Jac}} = 0$ is what licenses the Jacobi lever, so the error runs the unsafe way: across a grid of four algebras, five noise levels and three seeds, the exact tolerances give the wrong count on 22 of 60 cells and all 22 read zero where the truth is positive. Tying both cuts to the predicted relative quotient error $q = \sigma\sqrt{\sum_i s_i^{-2}}/\|\hat{u}\|$, already computed for (30), removes the systematic part. A tolerance cut can still be indecisive, so the implementation reports an integer only when the cut separates the two clusters by a factor of 10^2, the smallest decade at which no wrong count survives on that grid, and otherwise refuses and routes to the second energy. Refusing is not a weaker answer: the second energy does not depend on the candidate reading.

The implementation therefore uses four thresholds: (30) when a noise scale is supplied and a fixed 10^{-8} otherwise, both for retraction acceptance; the q-scaled cuts above, with their separation requirement, for the certificate re-read at the candidate; plus the dimension-aware gate of Section V-H for the different question whether an arbitrary tensor is a bracket at all.

G. Prescribing a second energy

When the certificate is positive the invisible directions are cocycles, so to first order they do not leave the Lie variety and a retraction has nothing to select on. On every battery row the displacement stays on the variety at finite size too, which the checks of Section V-D verify directly rather than prove, and no quantity of data changes that; a second experiment under a different energy does. For diagonal Hessians $A = \mathrm{diag}(a)$ and $B = \mathrm{diag}(b)$, Theorem 2 gives the intersection dimension $\sum_t \binom{|L_t|}{3}$ over the level sets L_t of the ratio vector $r_k = a_k/b_k$. To annihilate it, take $b = a \odot \rho$ with the ρ_k pairwise distinct and positive, say evenly spaced over a decade. Changing the energy is an intervention on the experiment, not a reanalysis of the same data, and it assumes the bracket persists across the two runs, the states stay the same quantities in the same coordinates, and the Hessian actually changes. A structure whose energy is set by a tunable load or a controlled stiffness admits that directly. The rigid body does not: changing the inertia tensor means changing the body, so the two runs identify a family rather than one object, a legitimate experiment but a different one. Where the energy is fixed by the representation, as with the orthonormal POD basis that pins $A = I$ in Section V-C, the lever is unavailable and the ambiguity must be quotiented rather than removed.

Proportionality is the failure case and the natural thing to try: $b = 2a$ has a constant ratio vector, one level set, and leaves the invisible subspace at its full $\binom{d}{3}$; the second experiment adds nothing. For non-diagonal pairs with A positive definite, diagonalize simultaneously by congruence; the level sets are those of the generalized eigenvalues of the pencil (B, A), which for an indefinite A need not be real. The dimension is zero as soon as the ratios are distinct, but near-ties degrade conditioning rather than dimension, so the ratio gaps should be comparable to the anisotropy of A itself. That last point is guidance; we did not sweep the gap.

H. Predicting the orbit budget

Because quotienting removed the structural rank deficiency, the quotient estimation error is the textbook one: with white field noise n and full-rank D,

$$\hat{u} - u^* = D^{+} n, \qquad \mathbb{E}\|\hat{u} - u^*\|^2 = \sigma^2 \sum_i s_i^{-2}, \quad (31)$$

with s_i the singular values of D: the weakly measured directions, the small s_i, dominate the expected error. The identity is standard; what it gives here is a right-hand side computable

from a *pilot* orbit set before the fitting data exist, so the budget $K^*(\varepsilon, \sigma)$ meeting a target follows from a formula rather than a sweep. The target is relative, $\varepsilon\|u^*\|$; a deployment, lacking u^*, substitutes the pilot estimate's norm, an argument to the budget routine. The substitute is a quotient norm, and on a noiseless pilot at the $\mathfrak{su}(3)$ gauge fractions measured here it is roughly a third of $\|u^*\|$, tightening the target rather than loosening it. Neither the fraction nor the direction of that inequality is guaranteed at another gauge fraction or on a noisy pilot. The prediction is a budget rule and not a law, in the sense made precise in Section V-F, where its calibration and the finite-difference variant are measured.

I. Computational cost

Two stages dominate, the certificate nullspace and the Levenberg–Marquardt normal matrix, both $\Theta(d^{10})$ in time and $\Theta(d^7)$ in memory, each touching a matrix with d^4 rows; everything else is at most $\Theta(d^9)$ and linear in the sample count. In float64 on one CPU core the certificate nullspace costs $0.091\,\mathrm{s}$ at $d = 8$ and $33.3\,\mathrm{s}$ at $d = 16$; end-to-end recovery on $\mathfrak{su}(3)$ with 16 orbits returns in $0.21\,\mathrm{s}$. The binding limit is memory: the certificate matrix occupies $7.3\,\mathrm{MB}$ at $d = 8$, $1.01\,\mathrm{GB}$ at $d = 16$ and $17.6\,\mathrm{GB}$ at $d = 24$, so dense factorization is comfortable to $d \approx 16$ and impractical beyond $d \approx 20$. The supplement gives the per-stage table and three reductions that would raise that ceiling, none implemented here.

J. Numerical considerations

Three practice notes govern every integer this method reports, each of them measured in the supplement. Double precision is required, for the reason Section V-I measures as F1. A reported dimension should be read off the spectral gap rather than the tolerance: where the ratio at the proposed cut exceeds about 10^4 the integer is stable over ten decades of tolerance, and where it does not, report the spectrum instead. And the quotient margin exposes a bad experiment before the fit does, sitting at the numerical-rank floor when the states come from one or two trajectories however many rows they supply, so the budget rule refuses degenerate configurations without a separate genericity test.

K. Software

The procedure is a single NumPy module with two entry points. From d, an energy model and an optional candidate bracket, `certify` returns the gauge dimension and basis, the quotient basis, ι_H, the cosine ceiling of Corollary 1, the lever, and a callable returning K^* and its predicted error or a refusal. `discover` returns $\hat{c}$, $\hat{J}_0$, the lever used, the Jacobi residual, the quotient cosine bound, the pre-retraction estimate and the retraction trace, routing by default from the certificate at the retracted candidate so no ground truth is required. Three guards are deliberate and tested: sub-`float64` input raises an error, an unrecognized mode raises one rather than defaulting silently, and the lever field is always present, even when undecidable.

TABLE II
THE FIVE QUANTITIES THIS SECTION REPORTS, EASY TO CONFUSE BECAUSE ALL FIVE ARE SMALL AND ALL FIVE ARE CALLED ERRORS SOMEWHERE IN THE LITERATURE. $\hat{u}$ IS THE FIT IN THE QUOTIENT, u^* ITS TRUTH, $\hat{c}$ THE FULL COEFFICIENT ARRAY.

quantity	what it measures
field residual $\|D\hat{u} - y\|/\|y\|$	how well the fit reproduces the data
quotient error $\|\hat{u} - u^*\|/\|u^*\|$	error in what the data determine
coefficient error $\|\hat{c} - c^*\|/\|c^*\|$	error in the array, gauge included
c-cosine $\cos(\hat{c}, c^*)$	agreement of the array up to scale
$\mathcal{J}(\hat{c})$	distance from the Jacobi variety

The test suite is eighteen cases in a few seconds, pinning the certificate integers on $\mathfrak{su}(3)$ and $\mathfrak{so}(3)$, the intersection law on proportional and ratio-distinct energy lists, end-to-end recovery with automatic routing, the negative control that routes $\mathfrak{so}(3)$ to the energy lever and identifies it once a second energy is stacked, the guards above, and two safety invariants: that the certificate never certifies $r_{\mathrm{Jac}} = 0$ where it is positive, and that the implemented acceptance tolerance is the one (30) prints. The supplement lists them. The thirteen-algebra battery reproduces bit for bit on rerun.

V. EXPERIMENTAL VALIDATION

Sections V-A–V-H measure the counting theorems as *integers*, the predictive value of ι_H, and the two repairs on noisy data at a budget chosen in advance.

A. Experimental Setup

1) The thirteen algebras: The battery is a covering design over the structural properties plausibly controlling identifiability, broad enough that some row falsifies any shortcut rule: $\mathfrak{so}(3)$, $\mathfrak{sl}(2,\mathbb{R})$, $\mathfrak{se}(2)$ and $\mathfrak{h}_3$ at $d = 3$; $\mathfrak{se}(2) \oplus \mathbb{R}$, $\mathfrak{h}_3 \oplus \mathbb{R}$, the filiform $\mathfrak{n}_4$ and $\mathfrak{aff}(1) \oplus \mathfrak{aff}(1)$ at $d = 4$; the filiform $\mathfrak{n}_5$ and $\mathfrak{h}_5$ at $d = 5$; $\mathfrak{so}(4)$ and the heavy top $\mathfrak{e}(3)$ at $d = 6$; and $\mathfrak{su}(3)$ at $d = 8$. The supplement gives every bracket. Compactness is crossed against semisimplicity, where the classical rigidity results apply and every prior bracket-learning paper reports; nilpotency appears five times as the obvious cheap predictive proxy, two of the five refuting it, with three solvable non-nilpotent rows between. The heavy top $\mathfrak{e}(3)$, the standard non-semisimple Lie–Poisson system of rigid-body mechanics, is one of three rows with $H^2 \neq 0$. Decomposability is crossed against all of it, and that cell is decisive: at $d = 3$ the phenomenon is classical (Section I-B), so a $d \geq 4$ row that is a $d = 3$ algebra with an abelian summand would prove nothing. The indecomposable $\mathfrak{n}_4$ carries the claim alone, being nilpotent with a one-dimensional center where a decomposable nilpotent algebra needs at least two.

2) Protocol: Every experiment uses a quadratic energy with A symmetric and invertible, in three families: anisotropic diagonal, the physically ordinary default; dense symmetric positive definite, where an intersection dimension's genericity is at issue; and isotropic $A = I$, a boundary case because for a compact semisimple algebra in a Killing-orthonormal basis it makes E the quadratic Casimir and the field vanishes. Trajectories come from RK4 from Gaussian initial conditions;

the unit of data is the *orbit*, one initial condition integrated for a segment of states; wherever budgets are compared the total row count is fixed. Observation noise is i.i.d. Gaussian on the observed fields, with states and ∇E exact, the regression setting the budget rule is derived for; Section V-G relaxes that. Three seeds are used throughout. Because a dimension count alone is not evidence of invisibility (Section V-I, F8–F9), every row with $Z^2 \cap K_A \neq 0$ carries five executed checks, C1–C5, on a concrete invisible direction, defined and tabulated in the supplement; on the five rows with $Z^2 \cap K_A = 0$ they are vacuous and the intersection itself runs instead. All linear algebra is `numpy float64`, with ranks and orthonormal bases read from SVD, never from unpivoted QR. Everything runs on one laptop CPU core with no GPU and no training loop, the full battery in 67 s. The supplement gives the tolerances, seeds, step size and timing ladder.

B. Verification of the Gauge and Intersection Theorems

1) Nullity and rank: The supplement tabulates the two integers of Theorem 1, measured by SVD of the design matrix at generic $\mathcal{N}(0, I)$ states for $d \in \{3, 5, 6, 7, 8, 9, 10, 12\}$. Predicted and measured agree on every row, with fifteen orders of magnitude between the smallest retained and the largest discarded singular value, so the rank decision is not a tolerance choice, and the J_0-block contributes nullity 0 throughout. At $d = 8$ the rank 196 of 252 reproduces, from an independent code path, the rank of the blind $\mathfrak{su}(3)$ solve of Section II-D: 56 of 252 bracket parameters are invisible to the data. The sweep omits $d = 4$ and $d = 11$, though the four $d = 4$ battery rows of Table III verify the gauge dimension $\binom{4}{3} = 4$.

2) Anisotropy does not shrink the gauge: The number of invisible directions is $\binom{d}{3}$ for *every* invertible A. Measured at $d = 6$ and $d = 8$ for diagonal SPD, dense SPD and generic invertible non-symmetric A, the nullities are $20, 20, 20$ and $56, 56, 56$: anisotropy re-aims the invisible subspace without resizing it, and no single quadratic energy reduces it.

3) The two boundary cases: Invertibility of A can be relaxed by one. For symmetric A of rank $r < d$ the measured kernel dimension matches the predicted $\binom{d-r}{2} + \binom{r}{3} + (d - r)\binom{r}{2} + d\binom{d-r}{2}$ at $(d, r) = (6, 5)$, $(8, 7)$, $(6, 4)$ and $(8, 5)$: at corank one it collapses to $\binom{d}{3}$, the invertible value 20 and 56, and the kernel *grows* only from corank two, to 23 and 67. Conversely, for $E = \frac{1}{2}\|z\|^2 + \frac{1}{3}\sum_l g_l z_l^3$ with all $g_l \neq 0$ the measured nullity is 0 at $d \in \{6, 8\}$, for a diagonal and for a generic symmetric cubic: a nonlinear energy removes the gauge outright at this bracket degree (the polynomial degree of $J(z)$ in z, one for the affine class). That is not a general escape, and Section V-I, F6, shows what the cubic energy costs in orbit diversity.

4) The forced ceiling: Corollary 1 predicts, before any data are collected, the cosine a gauge-blind least-squares fit reports against the true structure constants. Predicted and measured agree to six digits in three cells: 0.598200 for a mixed truth under $A = I$, 0.156128 for $\mathfrak{su}(3)$ under a diagonal A at $d = 8$, and 0.420560 for the $\mathfrak{so}(3)$ rigid body with inertia $(1, 2, 3)$. In the recovery experiments the same agreement holds cell by cell to $|\Delta| \leq 5.6 \cdot 10^{-17}$ over five energy draws. Behind the isotropic boundary case, the $\mathfrak{so}(3)$ and $\mathfrak{su}(3)$ constants have hook component zero and the $\mathfrak{su}(3)$ Lie–Poisson field under $E = \frac{1}{2}\|z\|^2$ is $7.8 \cdot 10^{-16}$ over 134 generic states, i.e. identically zero.

5) The intersection law: Theorem 2 predicts how much of the gauge a second energy removes, $\dim(K_A \cap K_B) = \sum_t \binom{|L_t|}{3}$, from the level sets L_t of the curvature ratios $r_k = a_k/b_k$ alone. Measured at $d = 6$, the predictions 20, 1 and 0, for $B = 2A$, for level sets $\{1, 2, 3\}, \{4, 5\}, \{6\}$ and for pairwise distinct ratios, are each matched to the integer; the same three at $d = 8$ give $56/56$, $1/1$ and $0/0$.

C. The Gauge on Community Benchmark Data

Every system above is hand-built, but the same computation runs on data we did not generate: the 2D Burgers snapshots published with Gkimisis et al. [11], to whom Section I-B attributes the count of freely selectable entries. The supplement gives the simulation parameters. A POD basis is orthonormal, so the reduced energy is isotropic, $A = I$, and the reduced quadratic operator is the c-block of (5).

The supplement tabulates the nullity of the sampled evaluation map against $\binom{r}{3}$ at five reduced orders. At $r = 5$ and $r = 8$ the measured nullity equals $\binom{r}{3}$; at $r \geq 10$ the count exceeds the prediction and the gap at the cut falls to 1.02, so by the rule of Section IV-J no tolerance turns it into an integer. Two controls localize the excess: at matched sample count, generic Gaussian states return $\binom{r}{3}$ at every order, as do Gaussian states whose per-coordinate scales match the decaying POD amplitudes, ruling out amplitude anisotropy. What remains is the configuration of the trajectory's states, permitted by Lemma 1, which guarantees kernel equality only outside an exceptional set; these snapshots, a single decaying solution, sit in it at $r \geq 10$. Being a single trajectory is not itself the cause; the $\mathfrak{su}(3)$ designs of Section V-I are also single-trajectory per orbit and hold the nullity at its predicted value at every budget.

At the orders a practitioner would choose here, r between 10 and 20, the basis already capturing 99.98% of the snapshot energy, one simulation does not expose the gauge as an integer. The published data do not vary the initial condition; whether that diversity restores the count here is untested.

The closed-form gauge basis lies in the kernel of the sampled design to $9 \cdot 10^{-18}$ relative, so arbitrary gauge displacements separate admissible fits: adding to the minimum-norm fit a gauge element of norm $0.24\|\hat{c}\|$ at $r = 5$ and $0.68\|\hat{c}\|$ at $r = 8$ (magnitudes chosen for demonstration) changes the generated field by $1.2 \cdot 10^{-16}$ and $6.2 \cdot 10^{-14}$. Two codes fitting these snapshots under different conventions for the free entries can disagree about two-thirds of the operator and remain indistinguishable from the data.

The fitted operator is not a Lie bracket: its scale-normalized Jacobi residual is 0.51 at $r = 5$ and 0.41 at $r = 8$, against the dimension-aware gate of Section V-H at $3 \cdot 10^{-5}$ and 10^{-5}. Theorem 1 nowhere uses the Jacobi identity and applies to this operator unchanged; the certificate of Definition 1, the retraction and the impersonation rows are defined on the Lie variety and say nothing about it. These snapshots

are a high-fidelity simulation, not a laboratory measurement (Section VI-B). Only the quadratic block of the reduced model is analyzed here; its linear part, the reduced viscous term, is symmetric and dissipative and lies outside (5).

D. The Thirteen-Algebra Battery

Which of the thirteen rows can impersonate, and does the certificate say so in advance? Table III evaluates it on all thirteen. The gauge dimension equals $\binom{d}{3}$ on every row, and r_{Jac} and ι_H are *identical across all three anisotropic energy draws*, licensing a reading of them as a property of the pair, not of a lucky draw.

TABLE III

THE CERTIFICATE ON THE THIRTEEN-ALGEBRA BATTERY. $Z^2 = \ker D_f$, $B^2 = \mathrm{im}(\mathfrak{gl}(d)\cdot f)$, $H^2 = Z^2/B^2$; GAUGE $= \dim K_A$, $r_{\mathrm{Jac}} = \dim(Z^2 \cap K_A)$, $\iota_H = \dim((Z^2 \cap K_A) + B^2) - \dim B^2$. GAUGE EQUALS $\binom{d}{3}$ ON EVERY ROW. r_{Jac} AND ι_H ARE IDENTICAL ACROSS THREE INDEPENDENT ANISOTROPIC ENERGY DRAWS ON EVERY ROW. EVERY IMPERSONATION ROW IS VERIFIED DIRECTLY BY CHECKS C1–C5 (SUPPLEMENT), NEVER INFERRED FROM A DIMENSION COUNT. THE LEVER IS NOT TABULATED BECAUSE IT IS A FUNCTION OF r_{Jac}: THE JACOBI RETRACTION WHERE $r_{\mathrm{Jac}} = 0$, A SECOND ENERGY OTHERWISE. ALL 13 ROWS ARE COMPUTED AT THREE ENERGY DRAWS EACH, WITH RANKS CUT AT AN ORTHONORMAL-BASIS TOLERANCE OF 10^{-9}.

Algebra	d	Z^2	B^2	H^2	gauge	r_{Jac}	ι_H
the retraction identifies these							
$\mathfrak{su}(3)$	8	56	56	0	56	0	0
$\mathfrak{e}(3)$ (heavy top)	6	30	29	1	20	0	0
$\mathfrak{h}_5$	5	30	10	20	10	0	0
$\mathfrak{n}_5$ (filiform)	5	24	16	8	10	0	0
$\mathfrak{aff}(1) \oplus \mathfrak{aff}(1)$	4	12	12	0	4	0	0
reparameterization: the invisible motion is a change of basis							
$\mathfrak{so}(3)$	3	6	6	0	1	1	0
$\mathfrak{so}(4)$	6	30	30	0	20	2	0
$\mathfrak{sl}(2,\mathbb{R})$	3	6	6	0	1	1	0
impersonation: a different algebra fits the same data							
$\mathfrak{se}(2)$	3	7	5	2	1	1	1
$\mathfrak{h}_3$	3	8	3	5	1	1	1
$\mathfrak{n}_4$ (filiform)	4	15	9	6	4	1	1
$\mathfrak{h}_3 \oplus \mathbb{R}$	4	19	6	13	4	2	2
$\mathfrak{se}(2) \oplus \mathbb{R}$	4	15	10	5	4	1	1

1) Finding 1: the certificate is finer than cohomology: Could the deformation count H^2 stand in for the certificate? The heavy top $\mathfrak{e}(3)$ says no, with $H^2 = 1$ along the classical Lévy–Nahas direction, so a reader working from cohomology alone would call it non-rigid and expect trouble; its invisible subspace nevertheless meets Z^2 only at the origin, the row routes to the Jacobi lever, and the retraction recovers it. The reason is that H^2 is an invariant of the bracket alone, whereas K_A depends on the energy: for $\mathfrak{su}(3)$, $Z^2 \cap K_A$ is 21-dimensional under the isotropic energy and 0 at every anisotropic draw with $H^2 = 0$ throughout, and for $\mathfrak{so}(4)$ it is 2 at diagonal energies and 0 at dense SPD ones.

2) Finding 2: impersonation is not a $d = 3$ artifact: Could impersonation be a low-dimensional accident? At $d = 3$ every unimodular linear Poisson structure on $\mathbb{R}^3$ is $\nabla C \times \cdot$, the one-dimensional gauge merely rescales the Casimir, and the resulting family moves through the classical Bianchi types (Section I-B). The claim therefore rests on the $d = 4$ rows, specifically the indecomposable $\mathfrak{n}_4$, whose supplement row carries all five checks: adding the invisible direction leaves the field, the Jacobi residual and the 300-step trajectories agreeing to 10^{-15} or better while moving both the Killing signature and the derived-algebra dimension, and between 65% and 88% of the direction lies outside B^2 across the three energy draws, so the motion is no change of basis in disguise. Across all five impersonation rows and all three draws at $t = 0.1$, the worst field difference is $7.7 \cdot 10^{-15}$, the worst Jacobi residual $1.2 \cdot 10^{-16}$, the worst trajectory deviation $7.8 \cdot 10^{-16}$. A structurally different Lie algebra reproduces the data to machine precision, and no residual, conservation or Jacobi diagnostic reports the substitution.

Two further measurements confirm the checks. Invisibility and variety membership degrade with t as the polarization identity $\mathrm{Jac}(f + th) = t^2\mathrm{Jac}(h)$ predicts: on the four rows where $\mathrm{Jac}(h)$ sits above floating-point noise the measured residuals at $t = 0.5$ and $t = 1.0$ are 25 and 100 times those at $t = 0.1$, all remaining at floating-point noise. The residual columns C1 stay below $4.1 \cdot 10^{-15}$ in K_A and $8.8 \cdot 10^{-16}$ in Z^2 across all eight rows and all three draws, so the tested direction never drifts out of the intersection. The two $d = 3$ impersonators land on the invariants of $\mathfrak{sl}(2,\mathbb{R})$ or $\mathfrak{so}(3)$, both classical Bianchi types, the three $d = 4$ impersonators on $(0,1,3)/3/1$ or $(2,1,1)/3/1$, depending on the energy draw. We report the change of isomorphism class and do not identify the target class, which would require more than these three invariants.

On the three reparameterization rows the verdict is a demonstrated isomorphism rather than an absence of detected change. On $\mathfrak{so}(3)$'s invisible line the motion rescales the slices, $\varepsilon^k \mapsto \mu_k \varepsilon^k$ with $\mu_k = 1 + ta_k$, and a diagonal change of basis $S = \mathrm{diag}(\lambda)$ acts by $c_{ij}^k \mapsto (\lambda_i\lambda_j/\lambda_k)c_{ij}^k$; matching the three cyclic triples and multiplying gives $\lambda_k = \sqrt{\mu_i\mu_j}$ over the complementary pair, defined whenever every μ is positive. That map reproduces $f + th$ to $2 \cdot 10^{-16}$ at $t = 0.1$, 0.5 and 1.0 without preserving the energy: $\|S^\top AS - A\|/\|A\|$ is 0.21, 1.26 and 3.00 at those displacements, Remark 2 made concrete and the reason the $\mathfrak{o}(A)$ certificate reads 1 on this row where the $\mathfrak{gl}(d)$ certificate reads 0.

Those verdicts are statements about the tested displacements and do not extend along the whole invisible line. At larger negative t the slice scales $\mu_k = 1 + ta_k$ change sign one at a time, at $t = -1/a_k$. Between the first and last crossing the algebra on $\mathfrak{so}(3)$'s line carries the Killing signature $(2,0,1)$ of $\mathfrak{sl}(2,\mathbb{R})$ while still generating the identical field with the Jacobi identity intact: at $t = -0.55$ under one fixed energy draw the field difference is $1.4 \cdot 10^{-15}$ with the signature moved. One scaled $\mathfrak{so}(3)$ ideal of $\mathfrak{so}(4)$ behaves the same way: between that ideal's own crossings the signature is $(2,0,4)$, that of $\mathfrak{sl}(2,\mathbb{R}) \oplus \mathfrak{so}(3)$, at field difference $2.0 \cdot 10^{-15}$, the Jacobi identity again intact. This is the Bianchi crossing of Section I-B on the battery's own rows, and why every row with nonzero intersection routes to the energy lever regardless of ι_H: motion along the gauge is a change of basis near the truth and not in the large.

3) Finding 3: nilpotency does not predict impersonation: The cheapest proxy rule, "nilpotent algebras impersonate", is false. $\mathfrak{h}_5$ and $\mathfrak{n}_5$ are nilpotent at $d = 5$ with $H^2 = 20$ and

TABLE IV
THE INTERSECTION LADDER ON $\mathfrak{su}(3)$ ($d = 8$), MATCHED TOTAL BUDGET OF 16 ORBITS $\times$ 24 STATES AT $\Delta t = 10^{-2}$, MEDIANS OVER THREE SEEDS. RECOVERY IS THE COSINE OF THE BLIND LEAST-SQUARES ESTIMATE TO THE TRUE STRUCTURE CONSTANTS. ROW 2 IS THE CONTROL: TWO ENERGIES, IDENTICAL RATIOS, ZERO GAIN.

Arm	Ratios $r_k = a_k/b_k$	$\dim(K_A \cap K_B)$ pred. / meas.	Recovery (cosine)
one energy	—	56 / 56	0.3358
$B = 2A$ (control)	all equal	56 / 56	0.3358
partial ties	$\{1,2,3\},\{4,5\},\ldots$	1 / 1	0.8839
distinct	pairwise distinct	0 / 0	1.0000

$H^2 = 8$ respectively, and both have $r_{\mathrm{Jac}} = \iota_H = 0$: the retraction identifies them. Nor does H^2 order the rows, since $\mathfrak{h}_3 \oplus \mathbb{R}$ has $H^2 = 13$ and impersonates while $\mathfrak{h}_5$ has $H^2 = 20$ and does not. In Table III, ι_H is the only column separating the recoverable rows from the impersonating ones: zero on all eight of the former, positive on all five of the latter.

E. Recovery: the Two Levers

1) The energy lever, with its control: Table IV tests the intersection law on $\mathfrak{su}(3)$ at $d = 8$ under a matched total budget of 16 RK4 orbits of 24 states, split evenly across each arm's energies, the structure constants solved blind by minimum-norm least squares. One anisotropic energy returns cosine 0.335838763882721, the ceiling $\|P_{K_A^{\perp}} f\|/\|f\| = 0.3358387633$ predicted for that curvature vector, agreeing to eight significant digits. Partial ties and pairwise distinct ratios lift the recovery as the kernel dimension falls.

The second row is the control: $B = 2A$ gives two energy functions and two trajectory sets at the same total number of states but identical curvature ratios, hence a predicted intersection of 56, the full gauge, and a measured recovery matching the single-energy arm to twelve significant digits. Data volume and energy count are thereby excluded, leaving the ratio spectrum, as Theorem 2 says.

2) The Jacobi lever against a soft-penalty baseline: Where a second energy is unavailable and the certificate reports $r_{\mathrm{Jac}} = 0$, the repair is a Levenberg–Marquardt retraction moving only inside the invisible coset $\hat{c} + K_A$, minimizing the Jacobi residual. Because every step lies in K_A, the retraction cannot change the fit to the data. On noiseless data from a single anisotropic energy and 16 orbits of 24 states, gauge-fixed least squares returns the ceiling of Section II-D, each of the three energy draws matching its prediction to $|\Delta| \le 5.6 \cdot 10^{-17}$. The retraction then reaches cosine 1.0000 in 8–9 steps at Jacobi residual 1.6–$2.5 \cdot 10^{-16}$. The standard alternative, a soft Jacobi penalty added to the loss with its weight tuned over $\lambda \in \{0.01, 0.1, 1, 10\}$, also reaches 1.0000 but needs 1506–2440 steps, stops at residual $1.1 \cdot 10^{-9}$–$4.6 \cdot 10^{-8}$, and leaves a rollout residual of 1.3 to $9.7 \cdot 10^{-5}$. On noiseless data the retraction therefore gains two to three orders of magnitude in iteration count and seven in certificate quality, and no accuracy; an accuracy difference does open under noise, in the penalty's favor (Section V-F).

The obvious alternative fails structurally, and no weight in the swept range repairs it. Minimum-norm Newton projection onto $\{\mathrm{Jac} = 0\}$ always has the radial direction available, since $\mathrm{DJac}(c)[c] = 2\,\mathrm{Jac}(c)$, so its steps flow down the homogeneous cone toward the abelian bracket $c = 0$, halving the scale at each step. Measured, it terminates at a Jacobi residual of $1.0 \cdot 10^{-10}$, a diagnostic that looks conclusive, with the cosine stuck at 0.3454 and the data residual at 0.863: the fit to the observations is destroyed. Enforcing the constraint along the invisible coset avoids this and is available only because the coset is known in closed form.

A soft penalty is the weakest way to impose a constraint, so the comparison is repeated against two solvers built for the job, an augmented Lagrangian and sequential quadratic programming, each started from the same quotient fit the retraction starts from and given the constraint in its independent form, the $d\binom{d}{3}$ components the Jacobiator has once its total antisymmetry is used. One of them succeeds: on $\mathfrak{su}(3)$ the augmented Lagrangian also reaches cosine 1.0000. What separates the methods is what they do to the component the data determined. The retraction is confined to K_A, so it cannot move that component and does not, displacing it by $7 \cdot 10^{-15}$ relative; the augmented Lagrangian moves it by $5.9 \cdot 10^{-6}$ and the soft penalty by $5.4 \cdot 10^{-6}$, both paying seven hundred to two thousand function evaluations against the retraction's eight. The two SQP arms show why the problem is hard for a general solver. At a Lie point the constraint Jacobian is rank deficient by construction, its kernel being the coboundary space, and scipy's trust-region method reports that degeneracy before driving the quotient error to 1.0, returning a bracket that satisfies the Jacobi identity and has no relation to the data; SLSQP terminates at its starting point. On $\mathfrak{so}(3)$ none of this arises: the certificate reports $r_{\mathrm{Jac}} = 1$ before any data are collected, the quotient fit is already on the variety, and every solver returns it unchanged.

F. The Certified Pipeline

The three components run as one procedure: certify, predict the orbit budget K^* from the quotient margin, collect that many orbits, fit on the quotient, retract, and report the quotient error and the certificate. Nothing in the loop is tuned against the answer.

1) End-to-end grid: On $\mathfrak{su}(3)$ over $\varepsilon \in \{0.1, 0.03, 0.01\}$ crossed with $\sigma_{\mathrm{rel}} \in \{10^{-3}, 10^{-2}\}$ and three seeds, all 18 cells land at or under the target ε at the predicted budget, with $K^* \in \{8, 16, 32\}$, and no cell reaches the refusal path. A representative cell: $\varepsilon = 0.1$, $\sigma_{\mathrm{rel}} = 10^{-3}$, seed 0, $K^* = 8$, predicted relative quotient error 0.01745, achieved 0.01425.

Both sides of that comparison come from the same predictor on the same realized design, so what the grid measures is calibration, whether the predicted error is honest about the design it was computed on, and not agreement with an independent oracle. The supplement makes the same comparison for the rule proper, empirical against predicted K^*: those agree in 32 of 32 cells, two by agreeing to refuse, since at $\sigma_{\mathrm{rel}} = 0.1$ and $\varepsilon \in \{0.03, 0.01\}$ no K in the grid suffices either way.

When K^* is chosen from pilot orbits and the fit is then made on an independent draw of initial conditions, over 45 cells on $\mathfrak{su}(3)$ ($\varepsilon \in \{0.3, 0.1, 0.03\}$, $\sigma_{\mathrm{rel}} \in \{10^{-3}, 10^{-2}, 10^{-1}\}$, five

seeds), one cell refuses and 43 of the remaining 44 meet their target, a transfer rate of 0.977, with the achieved error at 0.32 of target in the median and the single miss overshooting by 5%. The rule transfers, with a failure rate no fixed-design table can show.

2) Component ablation at matched budget: The supplement separates the contributions at the budget the rule chose.

The quotient parameterization costs nothing and improves conditioning. Quotient and full-space least squares return the same quotient error to eleven significant digits (worst relative difference $6.8 \cdot 10^{-12}$ over the 18 cells), a free consistency check: the two solvers share no code path beyond the design assembly and can agree only if the closed-form gauge basis is the actual kernel. What it improves is conditioning, the full design sitting at $\sigma_{\min}/\sigma_{\max} = 3.5 \cdot 10^{-17}$, a numerical zero, against the reduced design's $1.1 \cdot 10^{-3}$, the quotient margin itself.

The retraction moves only where the data are silent. It lifts the c-cosine from the ceiling 0.2625 to 1.0000 (worst cell 0.99992) and leaves the quotient error at 0.00705, unchanged to every digit reported. The step is by construction a combination of the gauge basis vectors, so it lies in K_A, and the unchanged quotient error is the measurement that confirms it. The Jacobi residual $\mathcal{J}$ improves from $5.3 \cdot 10^{-3}$ to $1.4 \cdot 10^{-5}$ without reaching the 10^{-16} of the noiseless case, and that gap is the point of (30): the residual a retraction can reach under noise is set by the displacement of the fitted quotient off the variety, which motion inside K_A cannot undo. The arm reported here is the forced Jacobi lever. Under the fixed 10^{-8} acceptance test the automatic mode would decline the lever on all 18 cells and route them to a second energy, as on all 15 cells of the calibration sweep; under the noise-calibrated form (30) every one of those cells is accepted.

Under noise a strongly weighted soft penalty does better. Swept over six decades, the penalty selects $\lambda \in [10^2, 10^3]$ and reaches a c-cosine of 1.0000 on every cell, a certificate of $8.0 \cdot 10^{-7}$, and a quotient error of 0.00110, six times *smaller* than the 0.00705 the retraction inherits from least squares. Jacobi constrains the identifiable component as well as the invisible one, so enforcing it in the full parameter space corrects part of the noise in what the data measured. The retraction cannot and is not meant to: moving only inside K_A is what makes its guarantee unconditional and its cost ten steps against the penalty's median 609 iterations. Stopping the weight grid at $\lambda = 1$ instead gives 0.9959 with a worst cell at 0.844, an under-tuned arm rather than an intrinsically worse one.

3) The budget rule itself: Quotienting the gauge leaves the reduced design at full column rank, so the Gauss–Markov identity (31) applies as written. The noise amplification $\sum_i s_i^{-2}$, over the singular values s_i of that design, is computable from a pilot orbit before the fitting data are collected, so K^* is predicted rather than searched. Over the 108 noisy cells on $\mathfrak{so}(3)$ and $\mathfrak{su}(3)$, 100 of them unsaturated and entering the fit, a pooled log-log fit of the error against $\sigma/s_{\min}$ has slope 1.019 with $R^2 = 0.9885$. The fitted constant has median 1.78, IQR $[1.45, 2.09]$, and range $[0.87, 3.26]$; the sharp form's prediction deviates from the measured error by 3.2% in the median and at most 24%. We call this a budget rule and not a law deliberately: it is conditioned on the realized design and claims nothing over the randomness of that design. Finite-difference truncation, substituted for observation noise, obeys the same scaling with slope 0.755, $R^2 = 0.961$, but a design-correlated constant of median 7.0, IQR $[5.0, 10.5]$. Truncation error is deterministic and correlates with the ill-conditioned directions, so the white-noise form under-predicts by a factor of about four in the median. Use the finite-difference row as a scaling with an $O(10)$ constant, never as a prediction.

G. Where the Procedure Stops Being Trustworthy

Everything above holds the modeling assumptions fixed. Three measurements relax them one at a time on the two systems either side of the certificate; the supplement carries them in full.

Noisy states move the design matrix itself, so the least-squares estimate is biased. What cannot move is structural, the invisible subspace being a function of the dimension and the Hessian alone: across 27 cells the measured nullity of the design built from noisy states is the predicted $\binom{8}{3}$ at every noise level. What degrades is the budget rule's constant, in the optimistic direction, by a median factor of 1.6 to 1.7 when the two noise levels are matched up to 10^{-2}, by 5.2 when matched at 10^{-1}, and by 533 when state noise dominates the fields; treat the rule as a lower bound there. Total least squares does not repair it.

Handing the whole procedure a wrong Hessian $A_1 = A_0 + \delta \|A_0\|_2 \operatorname{sym}(G)$, and never correcting it, costs accuracy but not safety, and the safety is almost entirely refusal: over 108 cells spanning δ up to 10^{-1} the certificate accepts three, all at $\delta = 0$. Nothing below cosine 0.99 is certified because under a misspecified energy nearly nothing is certified at all. The residual says the rest without ground truth: for $\delta \geq 10^{-3}$ it is proportional to δ with a constant between 1.2 and 1.6, the two systems agreeing within a factor of 1.25, rising to about 6 at $\delta = 10^{-4}$ as the mismatch nears the noise floor. A residual an order above that floor is a misspecified energy, not a missing lever.

The second energy is exercised on noiseless fields in Table IV, which settles a question about kernels and leaves one about conditioning. Repeating the three arms over five noise levels at a fixed total row budget separates them. The kernel statement is untouched, the stacked gauge holding at 56, 56 and 0 at every noise level; but the distinct-ratio arm gains its empty kernel at a quotient margin one to two orders of magnitude smaller, so its advantage narrows with noise and reverses at $\sigma = 10^{-1}$. It still leads on every seed at $\sigma = 10^{-2}$, though by then the cosine has fallen to a median 0.74 across a seed range of 0.40 to 0.91, against 0.26 for a single energy. The budget rule prices that trade in advance from the stacked design.

H. The Dimension-Aware Jacobi Gate

How small must a Jacobi residual be before it counts as evidence? The certificate's null distribution, its value on antisymmetric arrays that are not brackets, shrinks with d, so

a tolerance fixed at 10^{-3} rejects every random tensor through $d = 16$, where the null median is $1.58{\cdot}10^{-3}$, and accepts every one from $d = 20$, where it is $9.15 \cdot 10^{-4}$: it has stopped being a test. The implemented gate is therefore dimension-aware, $\mathcal{J}(c) < \alpha\sqrt{3}\, d^{-5/2}$ with $\alpha = 10^{-3}$, and on a 28-row battery it classifies correctly all 22 rows carrying a ground-truth label, including the nine at $d \geq 20$ the fixed gate wrongly accepts. Passing it does not mean the algebra is right: in a companion battery two other algebras and a random tensor driven onto the variety by gradient descent all pass, and only the Killing signature separates them. A Jacobi certificate is a necessary condition and should be reported alongside an isomorphism-class invariant, never alone. The supplement gives the null fit, the battery and the separation distances.

I. Failure-Mode Catalog

The method, and the diagnostics one would naturally use to check it, can report success while being wrong. The supplement catalogs nine such modes, F1–F9, each with a measured symptom and the mitigation the implementation ships, and four of them change what a practitioner should do. Under F1, `float32` structure constants inflate the numerical rank, turning $\dim Z^2 = 56$ into 19 and the intersection 21 into 9, both self-consistent and unflagged, so the module refuses sub-`float64` input. Under F2, a fixed Jacobi tolerance stops being a test at moderate dimension (Section V-H). F3 is certified Jacobi with the wrong algebra, the five impersonation rows, which no post-hoc diagnostic detects. F5 is saturation, the estimate mostly noise at relative error above 0.7, on 8 of 108 noisy cells, 5 of them worse than returning zero, where the rule should refuse rather than report.

1) F3: the failure mode with no diagnostic: On the five impersonation rows the only signals separating the two brackets are the $GL(d)$-invariants computed from the estimate, so they say which algebra you got, not whether it is the right one. What says so in advance is ι_H, computed from the hypothesized pair before a single trajectory is recorded.

2) F5 and F6: two different walls: F5 is statistical: the design is full rank and the budget rule stays accurate while the estimate is useless, so the response is to refuse, and more data at the same noise level fixes it. F6 is structural: putting every initial condition on one sphere under a cubic energy gives the sampled design integer rank deficiencies that no noise reduction touches, and only orbit diversity removes them. The mode is specific to the gauge-breaking nonlinear energy, which gains identifiability at the price of more independent orbits; the supplement gives the decay with budget and the measured counterpoint, excess nullity from the single dissipative Burgers trajectory of Section V-C at $r \geq 10$ under a quadratic reduced energy.

VI. DISCUSSION

A. A Decision Procedure

These results compose into a data-free procedure of closed-form linear algebra on the model class, run before an experiment is designed.

1) *Compute the ambiguity first.* From the state dimension alone, $\binom{d}{3}$ coefficients are invisible under a single quadratic energy; for several planned energies, Theorem 2 counts what survives all of them.
2) *Decide whether the coefficients are the deliverable.* For prediction, rollout or control synthesis the ambiguity is harmless, every member of the quotient class generating the same trajectories. What follows applies when the constants themselves are the object.
3) *Certify before collecting.* For a hypothesized bracket and the planned energy, compute $r_{\mathrm{Jac}} = \dim(Z^2 \cap K_A)$ and ι_H: one nullspace, one rank and one subspace intersection, seconds of CPU time at $d = 8$.
4) *If $r_{\mathrm{Jac}} = 0$, use the Jacobi lever.* The truth is isolated in its invisible coset (Proposition 1(ii)), so fitting in the quotient and retracting recovers it without disturbing what the data determined. Without ground truth the test is r_{Jac} re-read at the retracted candidate, and the procedure refuses where that reading is not decisive (Section IV-F).
5) *If $r_{\mathrm{Jac}} > 0$, change the experiment, not the sample size.* The surviving directions stay on the Lie variety, so Jacobi enforcement cannot choose among them, and ι_H says what the ambiguity costs: a change of basis when $\iota_H = 0$, a non-isomorphic algebra when $\iota_H > 0$. The lever is a second energy with pairwise-distinct curvature ratios.
6) *Predict the budget.* Quotienting makes the design full rank, so the parameter error obeys the Gauss–Markov trace $\sigma^2 \sum_i s_i^{-2}$ with equality and one pilot orbit fixes the orbit count for a target accuracy; its refusal branch is informative.

B. Where This Applies

Of the engineering settings listed in Section II-B, projection-based model reduction is the one where the ambiguity is already counted in the literature: 1140 of 3800 coefficients at reduced order $d = 20$ [11]. It matters wherever reduced operators from different runs are compared, averaged or interpolated [8], [9], and Section V-C finds admissible fits on those benchmark snapshots differing by two-thirds of the operator norm while generating the same field. Because the orthonormal POD basis fixes the reduced energy at $A = I$, the energy lever is unavailable there and the ambiguity must be quotiented rather than removed. In spacecraft attitude and rigid-body identification [47] the state is angular momentum, A the inverse inertia tensor, $d = 3$, and the one-dimensional ambiguity the classical Casimir rescaling. This locates the model class rather than reporting results: every number here comes from synthetic systems whose algebras are known in closed form, and the method has not been run on laboratory or field data.

C. Reproducibility

Beside the module and suite of Section IV-K, the archive carries one short producer script per experiment. Each writes a machine-readable results file, and every number printed here

is transcribed from one, its source recorded at the point of use. The archive also carries the thirteen algebra definitions in exact arithmetic, the seeds and energy draws behind every reported cell, the tolerance settings, and the interpreter and library versions. The only third-party input, the 2D Burgers snapshot set published with [11], is fetched from its public repository rather than redistributed. Every table and figure regenerates from a clean checkout with no manual steps.

D. Limitations

The scope is affine brackets with quadratic energies.

Confinement to K_A makes the retraction's guarantee unconditional and, under noise, costs accuracy: the Jacobi identity constrains the identifiable component too, so a method free to move that component can correct part of the noise in it. A soft penalty at $\lambda \in [10^2, 10^3]$ does exactly that, matching the retraction's cosine at a quotient error six times smaller (Section V-F). Where that component is wanted as accurately as possible and no unconditional guarantee is needed, it is the better arm.

The certificate is one-sided: $\iota_H > 0$ does not by itself produce a competing algebra, which is why the five impersonation rows are exhibited directly. Proposition 1 is first-order, about tangent cones: where the intersection is positive the finite-displacement conclusion is verified numerically along an explicit line rather than proved, and three energy draws leave a degenerate-energy stratum unexcluded; where it is zero the stratum is excluded, the intersection dimension being upper semicontinuous in the energy.

The budget rule is derived for Gaussian noise on observed fields with exact states and gradients, and finite-difference truncation transfers it in shape but not in constant (Section V-F). With the states noisy too the rule survives as an order of magnitude rather than an identity, optimistic by a median factor of 1.6 to 1.7 at matched noise up to 10^{-2} and by 5.2 at 10^{-1} (Section V-G); it should be read as a lower bound there. The budget grid measures calibration on the realized design, not agreement with an independent oracle.

Throughout, the energy Hessian A is treated as known without error: the gauge, the certificate and the retraction direction all condition on it. As a first sensitivity measurement we replace A by $A + \delta\|A\|_2 \operatorname{sym}(G)$ with G Gaussian, δ up to 10^{-1}, over three energy draws and three perturbation seeds per system. This leaves the three integers (gauge dimension, r_{Jac}, ι_H) unchanged in all 72 cells on $\mathfrak{so}(3)$ and $\mathfrak{su}(3)$, while the ceiling moves continuously, by at most $1.7\,\delta$ across the sweep. Isolation, where it holds, is local at the truth, while the retraction starts a full coset displacement away. No basin guarantee is given, and re-evaluating the certificate at the retracted candidate is only a check after the fact.

Finally, the certificate carries the complexity limit: its cost is dominated by the nullspace computation for the polarized Jacobi operator, whose matrix has d^4 rows and $d^2(d-1)/2$ columns, 4096×224 at $d = 8$, computed in seconds, but $1.6 \cdot 10^5 \times 3800$ at $d = 20$, about five gigabytes in double precision. Everything else scales with the $O(d^3)$ parameter count, and beyond $d \approx 20$ the certificate needs a structured or randomized nullspace method, not developed here.

VII. CONCLUSION

Fitting an affine bracket to trajectory data determines the dynamics and does not determine the coefficients. The undetermined set is a linear subspace of dimension $\binom{d}{3}$, present for every invertible energy Hessian and computable before any measurement, and at moderate dimension it is large enough that a structurally different Lie algebra can reproduce the true field and its trajectories to machine precision while satisfying the Jacobi identity. Fit residual, energy conservation and held-out rollout error do not detect the substitution, so structure recovery must be certified before the data are collected.

Two levers restore identifiability, and exact linear algebra on the pair (bracket, energy) decides in advance which applies. When the invisible directions leave the variety of Lie algebras, a retraction along them recovers the truth in a handful of steps without disturbing what the data determined, at two orders of magnitude less work than the penalty and constrained-solver arms that also recover it. When they remain on the variety, the remedy is a second experiment with a different energy curvature spectrum, and the intersection law says which spectra work. Because the fit is posed in the quotient, its design is full rank and its smallest singular value is a usable conditioning margin, from which a pilot design predicts the sample budget for a target accuracy.

What these results leave open is the behavior at higher bracket degree. A cubic energy removed the kernel at bracket degree one in the cases measured ($d \in \{6, 8\}$), but for a bracket of higher degree q the ambiguity is expected to migrate to degree $q - 1$, where neither the dimension count nor the certificate has been established.

REFERENCES

[1] S. Greydanus, M. Dzamba, and J. Yosinski, "Hamiltonian neural networks," in *Advances in Neural Information Processing Systems 32*, 2019, pp. 15 353–15 363.

[2] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, "Lagrangian neural networks," in *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020, arXiv:2003.04630.

[3] P. Jin, Z. Zhang, A. Zhu, Y. Tang, and G. E. Karniadakis, "SympNets: intrinsic structure-preserving symplectic networks for identifying Hamiltonian systems," *Neural Networks*, vol. 132, pp. 166–179, 2020.

[4] A. van der Schaft and D. Jeltsema, "Port-Hamiltonian systems theory: An introductory overview," *Foundations and Trends in Systems and Control*, vol. 1, no. 2–3, pp. 173–378, 2014.

[5] E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, 2nd ed., ser. Springer Series in Computational Mathematics. Berlin: Springer, 2006, vol. 31.

[6] J. E. Marsden and T. S. Ratiu, *Introduction to Mechanics and Symmetry: A Basic Exposition of Classical Mechanical Systems*, 2nd ed., ser. Texts in Applied Mathematics. New York: Springer, 1999, vol. 17.

[7] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proceedings of the National Academy of Sciences*, vol. 113, no. 15, pp. 3932–3937, 2016.

[8] B. Peherstorfer and K. Willcox, "Data-driven operator inference for nonintrusive projection-based model reduction," *Computer Methods in Applied Mechanics and Engineering*, vol. 306, pp. 196–215, 2016.

[9] B. Kramer, B. Peherstorfer, and K. E. Willcox, "Learning nonlinear reduced models from data with operator inference," *Annual Review of Fluid Mechanics*, vol. 56, no. 1, pp. 521–548, 2024.

[10] H. Sharma, Z. Wang, and B. Kramer, "Hamiltonian operator inference: physics-preserving learning of reduced-order models for canonical Hamiltonian systems," *Physica D: Nonlinear Phenomena*, vol. 431, p. 133122, 2022.

[11] L. Gkimisis, I. Pontes Duff, P. Goyal, and P. Benner, "On the representation of energy-preserving quadratic operators with application to Operator Inference," *Applied Mathematics Letters*, vol. 173, p. 109761, 2026, thm.: $n^3/6 - n^2/2 + n/3 = \binom{n}{3}$ entries freely selectable.

[12] J. Hu, J.-P. Ortega, and D. Yin, "A global structure-preserving kernel method for the learning of Poisson systems," *Journal of Nonlinear Science*, vol. 35, no. 4, p. 79, 2025, arXiv:2504.13396.

[13] V. Putkaradze, "Latent Lie–Poisson neural networks (LLPNNs): discovering the motion of Lie–Poisson systems through observable data and latent dynamics," *arXiv:2607.28939*, 2026.

[14] W. Hou, M. Li, and Y.-Z. You, "Machine learning symmetry discovery for integrable Hamiltonian dynamics," *Physical Review E*, vol. 113, p. 025302, 2026.

[15] S. Bhardwaj and C. Bajaj, "PHAST: port-Hamiltonian architecture for structured temporal dynamics forecasting," *arXiv:2602.17998*, 2026.

[16] M. Šípka, M. Pavelka, O. Esen, and M. Grmela, "Direct Poisson neural networks: learning non-symplectic mechanical systems," *arXiv:2305.05540*, 2023.

[17] G. F. R. Ellis and M. A. H. MacCallum, "A class of homogeneous cosmological models," *Communications in Mathematical Physics*, vol. 12, no. 2, pp. 108–141, 1969.

[18] R. Biggs and C. C. Remsing, "A classification of quadratic Hamilton–Poisson systems in three dimensions," in *Proceedings of the Fifteenth International Conference on Geometry, Integrability and Quantization*, ser. Geometry, Integrability and Quantization, vol. 15. Sofia: Avangard Prima, 2014, pp. 67–78.

[19] R. M. Tudoran, "The free rigid body dynamics: generalized versus classic," *Journal of Mathematical Physics*, vol. 54, no. 7, p. 072704, 2013, arXiv:1102.2725.

[20] R. Bellman and K. J. Åström, "On structural identifiability," *Mathematical Biosciences*, vol. 7, no. 3–4, pp. 329–339, 1970.

[21] L. Ljung and T. Glad, "On global identifiability for arbitrary model parametrizations," *Automatica*, vol. 30, no. 2, pp. 265–276, 1994.

[22] A. F. Villaverde, A. Barreiro, and A. Papachristodoulou, "Structural identifiability of dynamic systems biology models," *PLoS Computational Biology*, vol. 12, no. 10, p. e1005153, 2016.

[23] N. Meshkat and S. Sullivant, "Identifiable reparametrizations of linear compartment models," *Journal of Symbolic Computation*, vol. 63, pp. 46–67, 2014.

[24] E. D. Sontag, Y. Wang, and A. Megretski, "Input classes for identifiability of bilinear systems," *IEEE Transactions on Automatic Control*, vol. 54, no. 2, pp. 195–207, 2009.

[25] M. M. Deal and G. Nolet, "Nullspace shuttles," *Geophysical Journal International*, vol. 124, no. 2, pp. 372–380, 1996.

[26] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle adjustment: a modern synthesis," in *Vision Algorithms: Theory and Practice*. Springer, 2000, pp. 298–372.

[27] Y. Nakamura, H. Hanafusa, and T. Yoshikawa, "Task-priority based redundancy control of robot manipulators," *The International Journal of Robotics Research*, vol. 6, no. 2, pp. 3–15, 1987.

[28] N. Dehmamy, R. Walters, Y. Liu, D. Wang, and R. Yu, "Automatic symmetry discovery with Lie algebra convolutional network," in *Advances in Neural Information Processing Systems 34*, 2021, pp. 2503–2515.

[29] J. Yang, R. Walters, N. Dehmamy, and R. Yu, "Generative adversarial symmetry discovery," in *Proceedings of the 40th International Conference on Machine Learning*, ser. PMLR, vol. 202, 2023, pp. 39 488–39 508.

[30] R. T. Forestano, K. T. Matchev, K. Matcheva, A. Roman, E. B. Unlu, and S. Verner, "Deep learning symmetries and their Lie groups, algebras, and subalgebras from first principles," *Machine Learning: Science and Technology*, vol. 4, no. 2, p. 025027, 2023.

[31] R. W. Richardson, Jr., "On the rigidity of semi-direct products of Lie algebras," *Pacific Journal of Mathematics*, vol. 22, no. 2, pp. 339–344, 1967.

[32] A. Nijenhuis and R. W. Richardson, Jr., "Cohomology and deformations in graded Lie algebras," *Bulletin of the American Mathematical Society*, vol. 72, pp. 1–29, 1966.

[33] ——, "Deformations of Lie algebra structures," *Journal of Mathematics and Mechanics*, vol. 17, pp. 89–105, 1967.

[34] A. Weinstein, "Affine Poisson structures," *International Journal of Mathematics*, vol. 1, pp. 343–360, 1990.

[35] P. S. Krishnaprasad, "Lie–Poisson structures, dual-spin spacecraft and asymptotic stability," *Nonlinear Analysis: Theory, Methods & Applications*, vol. 9, no. 10, pp. 1011–1035, 1985.

[36] N. E. Leonard, "Stability of a bottom-heavy underwater vehicle," *Automatica*, vol. 33, no. 3, pp. 331–346, 1997.

[37] P. J. Morrison and J. M. Greene, "Noncanonical Hamiltonian density formulation of hydrodynamics and ideal magnetohydrodynamics," *Physical Review Letters*, vol. 45, no. 10, pp. 790–794, 1980.

[38] D. D. Holm, J. E. Marsden, T. Ratiu, and A. Weinstein, "Nonlinear stability of fluid and plasma equilibria," *Physics Reports*, vol. 123, no. 1–2, pp. 1–116, 1985.

[39] P. J. Morrison, "Hamiltonian description of the ideal fluid," *Reviews of Modern Physics*, vol. 70, no. 2, pp. 467–521, 1998.

[40] P. Benner, S. Gugercin, and K. Willcox, "A survey of projection-based model reduction methods for parametric dynamical systems," *SIAM Review*, vol. 57, no. 4, pp. 483–531, 2015.

[41] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed. Baltimore, MD: Johns Hopkins University Press, 2013.

[42] A. Weinstein, "The local structure of Poisson manifolds," *Journal of Differential Geometry*, vol. 18, no. 3, pp. 523–557, 1983.

[43] J. E. Humphreys, *Introduction to Lie Algebras and Representation Theory*, ser. Graduate Texts in Mathematics. New York: Springer, 1972, vol. 9.

[44] C. Chevalley and S. Eilenberg, "Cohomology theory of Lie groups and Lie algebras," *Transactions of the American Mathematical Society*, vol. 63, pp. 85–124, 1948.

[45] N. Jacobson, *Lie Algebras*. New York: Interscience, 1962, whitehead's lemma, Ch. III.

[46] M. Lévy-Nahas, "Deformation and contraction of Lie algebras," *Journal of Mathematical Physics*, vol. 8, no. 6, pp. 1211–1222, 1967.

[47] F. L. Markley and J. L. Crassidis, *Fundamentals of Spacecraft Attitude Determination and Control*, ser. Space Technology Library. New York: Springer, 2014, vol. 33.