AUDITING COEFFICIENT IDENTIFIABILITY IN PHYSICS-CONSTRAINED DICTIONARY MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Low data-fitting error and a small physics residual do not ensure that the coefficients of a dictionary-based PDE solution are identifiable. For a fixed linear operator and sampling design, we develop a pre-training structural audit based on the classical joint null space of the data and residual maps. Without observed response values, the audit identifies coefficient perturbations that leave both predictions and residuals unchanged, returns explicit witnesses, and reports which directions the physics penalty constrains. Tolerance scans, higher-precision calculations, and symbolic checks distinguish functional identities from sampling-induced dependencies and numerical near-dependencies. Across nineteen dictionaries, the audit detects multi-term dependencies missed by pairwise and triple screens and identifies four dictionaries with identically zero residual penalties. In a public PySINDy constrained-regression example, steady-state observations leave seven coefficient directions unresolved by the data and constraint together. On these data, the published optimizer fits the sampled derivatives with a relative ℓ_2 error of 1.7×10^{-5} while omitting the growth and coupling terms. Quadrupling the steady-state samples leaves all seven directions unresolved. At the original sample budget, transient observations remove these blind directions. The audit identifies which individual coefficients and coefficient combinations are determined under the specified observations and physics constraints. It does not validate the imposed operator or guarantee improved reconstruction accuracy.

Keywords: Coefficient identifiability, Physics-informed learning, Structural auditing, Dictionary-based models, Null-space analysis.

1 INTRODUCTION

A fitted scientific model often invites a claim about its coefficients: a term is present or absent, a growth rate has a particular value, or a coupling is weak. Such claims require more than accurate predictions. Different coefficients can produce the same measured field, and an imposed equation can leave that ambiguity intact. An optimizer must still return one vector, making an arbitrary representative look like a measured mechanism. This matters in both fixed dictionaries for PDE solutions, where coefficients weight candidate fields, and constrained equation regression, where they weight candidate terms (Brunton et al., 2016; Rudy et al., 2017; Kaptanoglu et al., 2022). In each setting the useful question is which individual coefficients or combinations the stated observations and constraint actually determine.

The ambiguity is visible even with exact data and a correct equation. Let $\varphi_1 = \sin(kx)\cos(kt)$, $\varphi_2 = \sin(k(x{+}t))$, and $\varphi_3 = \sin(k(x{-}t))$. All three solve the one-dimensional wave equation, and $\varphi_1 = \frac{1}{2}\varphi_2 + \frac{1}{2}\varphi_3$. Thus $(1, 0, 0)$ and $(0, \frac{1}{2}, \frac{1}{2})$ give the same field at every observation point and the same zero residual. Their equivalence persists at zero data error and at any physics-penalty weight. Dropping φ_1 removes the coefficient ambiguity when observations separate the two travelling waves; the residual remains identically zero. The data and physics term therefore supply distinct information. In larger dictionaries, the analogous dependence can involve more than two or three atoms and be hidden by finite precision or a restricted sampling design.

The underlying algebra is established. Uniqueness of stacked data–penalty systems is governed by intersecting kernels in general-form regularization (Van Loan, 1976; Hansen, 1989); Scheipl & Greven (2016) use this overlap for a pre-fit warning about non-identifiable coefficients in penalized

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

regression. Constrained system identification also has full-rank conditions (Loiseau & Brunton, 2018), while equation-discovery work studies conditioning and identifiability of the recovered law (Rudy et al., 2017; Scholl et al., 2026; Hauger et al., 2024). These established criteria provide the algebraic starting point. Our audit reports the directions invisible to data, those closed by the imposed physics, and the coefficient readings that survive the remaining ambiguity. It separates exact functional identities from tolerance-dependent near-dependencies (Appendix C).

The distinction affects the kind of scientific statement one can make. If a direction is blind to both maps, reporting one coordinate from a fitted vector conceals a whole family of equally supported values. A combination orthogonal to that direction may still be reportable. If the direction is blind only to data, its value follows from the chosen constraint and should be labelled as equation-implied. If a direction is merely near-zero, the conclusion depends on the numerical scale and the accuracy demanded. These cases can yield similar fitted losses yet require different interpretations and different remedies: changing the dictionary, collecting informative observations, or stating a precision-qualified result. Our report makes those choices visible for the model and design actually used.

We therefore examine two maps on the same coefficient vector: one maps coefficients to sampled observations and the other to imposed residuals or equality constraints. Before training and without response values, the audit returns the directions blind to both maps with explicit witnesses, the data-blind directions constrained by physics, and a test for whether a proposed linear reading of the coefficients is determined. It also checks whether the result survives changes in tolerance and precision. Exact linear theory supports global statements; for nonlinear residuals the construction reports local first-order sensitivity. Scientific use of the report assumes a justified imposed equation; predictive accuracy requires a separate evaluation.

The main public case makes the reporting consequence concrete. In a pre-registered audit of a constrained-regression pipeline, steady-state observations leave a published coupling coefficient A and growth coefficient g individually undetermined, although the combination $g+mA$ is determined on the observed cycle. Four times more steady-state data leave the ambiguity; transient observations with the original sample count remove it (Section 4.1). An exploratory matched-sample-size refit returns all nine true terms on transient data at the optimizer's iteration limit. Our contributions are:

- *Determined readings.* We characterize coefficient freedoms left by data and the imposed constraint, distinguish data-blind directions that physics resolves, and test whether any specified linear combination is uniquely determined (Section 2).
- *An auditable report.* Witness vectors, residual-term support, and tolerance and precision checks separate exact identities from near-dependencies. Across nineteen dictionaries the audit finds multi-term relations missed by pairwise and triple screens and four identically zero residual penalties (Figure 3, Tables 5–9).
- *Consequences in public pipelines.* The registered main case shows which coefficients and combinations can be reported under distinct observation designs and what a matched-sample-size refit returns; two further rule-selected studies test the report on official example suites, mostly with well-excited designs.

2 Structural Criteria

2.1 Two maps on one coefficient vector

Both maps act on one coefficient vector α under a fixed sampling design. For a solution dictionary $u = \Phi^\top \alpha$, the physics map comes from the imposed operator $\mathcal{N}$. For constrained equation regression $u_t = \Theta(u)\xi$, the regression and equality-constraint matrices are the two maps on ξ (Section 4.1).

A dictionary $\{\varphi_i\}_{i=1}^K$ of real functions on a domain Ω is evaluated at observation points $z_1, \ldots, z_{N_\mathrm{d}}$ and residual points $y_1, \ldots, y_{N_\mathrm{c}}$. Writing $\Phi \in \mathbb{R}^{K\times N_\mathrm{d}}$, $\Phi_{in} = \varphi_i(z_n)$, and $R \in \mathbb{R}^{K\times N_\mathrm{c}}$, $R_{im} = (\mathcal{N}[\varphi_i])(y_m)$, the representation is $u = \Phi^\top \alpha$, the observations are $u_\mathrm{data} = \Phi^\top \alpha_\star + \varepsilon$ with $\mathbb{E}[\varepsilon] = 0$, $\mathrm{Cov}[\varepsilon] = \sigma^2 I$, and the training objective is

$$J_\lambda(\alpha) = \tfrac{1}{N_\mathrm{d}} \left\| \Phi^\top \alpha - u_\mathrm{data} \right\|^2 + \tfrac{\lambda}{N_\mathrm{c}} \left\| R^\top \alpha \right\|^2, \qquad \lambda \geq 0. \tag{1}$$

$\mathcal{N}$ is the **imposed** operator chosen by the practitioner. The audit assesses this specified equation; agreement with the data is a separate empirical question. For nonlinear operators we use the Fréchet Jacobian at a specified u_0 (Appendix B.3); this is a local first-order sensitivity diagnostic only. All required derivatives are computed symbolically. Rescale each atom to unit root-mean-square on the observation points, a reparametrization of α applied identically to the rows of R, and set

$$A_{\mathrm{d}} = \frac{\Phi^\top}{\sqrt{N_{\mathrm{d}}}}, \qquad A_{\mathrm{p}} = \frac{R^\top}{\sqrt{N_{\mathrm{c}}}}, \qquad \tilde{y} = \frac{u_{\mathrm{data}}}{\sqrt{N_{\mathrm{d}}}}, \qquad J_\lambda(\alpha) = \|A_{\mathrm{d}}\alpha - \tilde{y}\|^2 + \lambda\|A_{\mathrm{p}}\alpha\|^2. \tag{2}$$

For a coefficient perturbation v, $A_{\mathrm{d}}v$ is how much the predictions at the observation points change and $A_{\mathrm{p}}v$ how much the residual changes. The two Gram matrices are

$$G_{\mathrm{d}} = A_{\mathrm{d}}^\top A_{\mathrm{d}} = \tfrac{1}{N_{\mathrm{d}}}\Phi\Phi^\top, \qquad G_{\mathrm{p}} = A_{\mathrm{p}}^\top A_{\mathrm{p}} = \tfrac{1}{N_{\mathrm{c}}}RR^\top, \qquad H_\lambda = G_{\mathrm{d}} + \lambda\, G_{\mathrm{p}} = \tfrac{1}{2}\nabla^2 J_\lambda, \tag{3}$$

so $v \in \ker G_{\mathrm{d}}$ means the predictions do not move and $v \in \ker G_{\mathrm{p}}$ means the residual does not move, both exactly, since $\|A_{\mathrm{d}}v\|^2 = v^\top G_{\mathrm{d}} v$. For the wave example, symbolic checks give $\operatorname{rank} G_{\mathrm{d}} = 2$ and $G_{\mathrm{p}} = 0$; dropping φ_1 makes G_{d} full rank while G_{p} remains zero. This is the classical stacked-system condition (Appendix B.1.3).

2.2 Directions neither term constrains

The central object is

$$V_0 \;=\; \ker A_{\mathrm{d}} \cap \ker A_{\mathrm{p}} \;=\; \ker G_{\mathrm{d}} \cap \ker G_{\mathrm{p}}. \tag{4}$$

For $v \in V_0$, moving the coefficients along $\alpha + tv$ leaves both terms of equation 1 unchanged at every λ: relative to the current data and physics terms it is a jointly blind direction, which we also call a *gauge* direction (the word is overloaded; Appendix B.1.4). On the running example V_0 is spanned by $(1, -\frac{1}{2}, -\frac{1}{2})$, the difference of the two representations.

Proposition 1 (uniqueness)**.** *For fixed $\lambda > 0$ the minimizer of the bi-quadratic objective equation 1 is unique iff $V_0 = \{0\}$.*

The proposition applies to the quadratic objective in equation 1; networks, sparsity terms and additional priors change the objective. For readings it gives a test: with the columns of B_0 a basis of V_0, a linear reading $\ell^\top\alpha$ takes one value on the whole solution set iff $B_0^\top \ell = 0$. The minimum-norm solution selects the representative with zero component in V_0, which does not resolve the non-identifiability, and prediction-based criteria, cross-validation and GCV included, cannot distinguish the representations on the given design, since every functional of fitted values is invariant along $\ker G_{\mathrm{d}}$ (Proposition 4, Appendix B.1.2).

2.3 What the physics term closes

A direction in $\ker A_{\mathrm{d}}$ is unresolved by the observations. The physics term is sensitive to it iff it lies outside $\ker A_{\mathrm{p}}$, and we call such a direction *physics-visible*; the number of independent data-blind directions the residual block removes is the rank it adds to the data block,

$$\Delta r \;=\; \operatorname{rank}\begin{bmatrix} A_{\mathrm{d}} \\ A_{\mathrm{p}} \end{bmatrix} - \operatorname{rank} A_{\mathrm{d}} \;=\; \dim\ker A_{\mathrm{d}} - \dim V_0, \tag{5}$$

by rank–nullity applied to $\ker[A_{\mathrm{d}}; A_{\mathrm{p}}] = \ker A_{\mathrm{d}} \cap \ker A_{\mathrm{p}}$. In finite precision each rank is numerical; we use the common singular-value cut $\sqrt{\tau s_{\mathrm{ref}}}$ for the data and balanced stacked matrices (Section 3.2). On every corpus library $G_{\mathrm{d}} \neq 0$, so $s_{\mathrm{ref}} = \|G_{\mathrm{d}}\|_2$, and the rank identity is checked directly. At $\Delta r = 0$ the physics block adds no identifiable dimension, although statistical benefit remains possible. At $\Delta r > 0$ it closes data-blind directions under the imposed operator; an incorrect operator can close them to incorrect values. What the equation implies on a data-blind direction is a solution set, not a number:

Proposition 2 (equation-implied completion)**.** *Let the columns of B be an orthonormal basis of $\ker G_{\mathrm{d}}$, write $\alpha = \alpha_{\mathrm{vis}} + Bc$ with $\alpha_{\mathrm{vis}} \perp \ker G_{\mathrm{d}}$, and set $M = B^\top G_{\mathrm{p}} B$. The coefficient vectors that share the visible component and minimize the residual form the set $\alpha_{\mathrm{vis}} + Bc^\star + V_0$ with $c^\star = -M^+ B^\top G_{\mathrm{p}}\,\alpha_{\mathrm{vis}}$ and $B\ker M = V_0$; a linear reading $\ell^\top\alpha$ is constant on it iff $\ell \perp V_0$.*

The completion needs α_{vis} and is therefore computed after a data fit. The audit itself uses no observed responses, and the data alone cannot establish the correctness of the equation-implied completion (Theorem 1, Appendix B.3.1). Diagonalizing the pair $(G_{\mathrm{d}}, G_{\mathrm{p}})$ refines the count into the exact four classes behind Table 1. With numerical truncation, separately thresholded counts are diagnostics rather than an exact decomposition, and their arithmetic need not reproduce $\operatorname{rank} G_{\mathrm{p}}$. $\nu_j \in [0,1]$ is the physics share of direction j, the eigenvalues of the $G_{\mathrm{d}}+\lambda_{\mathrm{ref}}G_{\mathrm{p}}$-whitened $\lambda_{\mathrm{ref}}G_{\mathrm{p}}$ on the complement of V_0, with λ_{ref} the balanced weight of Section 3. The classes answer distinct questions: the running example has one gauge direction and an inert penalty, its two-function variant has no gauge direction and the same inert penalty, and three corpus libraries have $\dim V_0 = 0$ and $\operatorname{rank} G_{\mathrm{p}} = 0$ at once (Section 4.2). For pruning, only the part of $\ker G_{\mathrm{d}}$ inside $\ker G_{\mathrm{p}}$ is redundant (Appendix B.1.5).

Table 1: Exact four-way decomposition of coefficient space: the common null space plus three classes on its complement. Here ν_j is the generalized physics share, with $\nu_j = 0$ or 1 interpreted exactly. Numerical Figure 2 uses stated endpoint cuts; its bins are not these exact subspace dimensions.

class	definition	who can determine it	consequence
gauge	$v \in V_0$	neither term, at any λ	quotient or redesign
physics-only	$\nu_j = 1$	the equation, modulo V_0	equation-implied value
data-only	$\nu_j = 0$	the data only	penalty inert
jointly seen	$0 < \nu_j < 1$	both	normal regime

2.4 Where jointly blind directions come from

Lemma 1 (structural existence, lower bound)**.** *If the dictionary contains p atoms lying in some d-dimensional subspace S of functions with $p > d$, then $\dim V_0 \geq p - d$, for every choice of observation and residual points.*

The hypothesis is checkable on a dictionary before any data exist. It covers the first of two sources, which the audit keeps apart because their remedies differ. *Functional identity:* the atoms are dependent as functions, no sampling separates them, and a linear operator cannot either, since $\mathcal{N}[\sum_i c_i\varphi_i] = 0$; the running example is one instance, and the fourteen-atom wave library of Section 4 is the same instance three times over, six identities in all, invisible until expanded. *Sampling coincidence:* the atoms are distinct functions that agree at the observation points (grid aliasing); here $\ker A_{\mathrm{d}}$ is a property of the design, redesigning the sampling set can remove it, and the operator, evaluated at its own points, may resolve it meanwhile. Every fixed-design statement concerns inference from the stated observation design; new sampling points can change its null space.

Proposition 3 (detection completeness)**.** *In exact arithmetic $V_0 = \ker[A_{\mathrm{d}}; A_{\mathrm{p}}]$ and a rank-revealing decomposition of the stacked matrix returns all of it, whatever the supports of its elements. An exhaustive enumeration of all relations with support at most three, tested against* both *given maps, returns $V_{\leq 3} = \operatorname{span}\{v \in V_0 : \|v\|_0 \leq 3\}$ and misses a direction iff $V_{\leq 3} \subsetneq V_0$; a nonexhaustive triple screen can miss more.*

The gap is real: on one corpus library V_0 is spanned by an identity on five atoms, $V_{\leq 3} = \{0\}$, and the triple enumeration returns nothing (Section 4.3).

3 A Numerically Reliable Structural Audit

3.1 Procedure and outputs

The audit takes a dictionary $\{\varphi_i\}_{i=1}^K$, the imposed operator $\mathcal{N}$, a sampling design and a tolerance τ. It needs no observed response values, no training and no user-selected λ; if the dictionary was itself selected from data, every conclusion is conditional on that selection. Algorithm 1 (Appendix B.1) states the audit line by line. At a fixed tolerance it uses up to four symmetric $K \times K$ eigendecompositions (joint sum, two separate Grams, and whitened physics share), $O(K^2(N_{\mathrm{d}} + N_{\mathrm{c}}) + K^3)$ for the

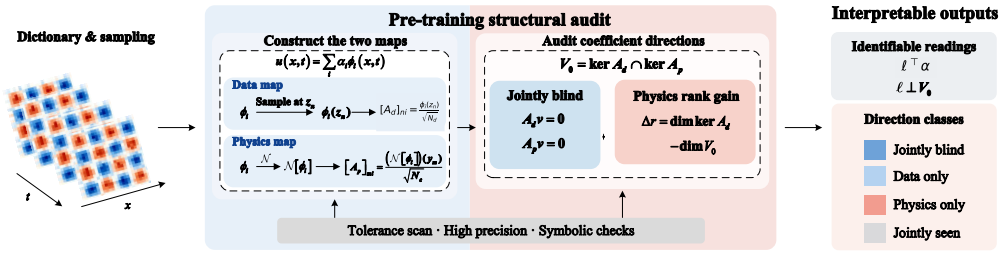


Figure 1: Pre-training structural audit of a prescribed solution dictionary. Given the dictionary, imposed operator, and sampling design, the audit constructs the normalized data and physics maps A_{d} and A_{p}. Their joint null space identifies coefficient directions invisible to both terms; the physics rank gain counts how many data-blind directions the operator constrains. The outputs are identifiable linear readings and the four direction classes. No observed response values or trained model are needed.

core linear algebra, seconds for K in the hundreds, against $O(K^3)$ least-squares problems for the triple enumeration it replaces, which by Proposition 3 is also strictly less complete. Its four outputs (Table 7) are joint near-zero directions with witnesses, the support and rank of the residual term, the numerical rank gain Δr_τ, and the tolerance and precision diagnostics. Four corpus libraries have a zero physics block, so the procedure handles that branch explicitly. When both blocks are nonzero, $\lambda_{\mathrm{ref}} = \|G_{\mathrm{d}}\|_2/\|G_{\mathrm{p}}\|_2$ balances their norms; when only G_{p} is nonzero, we normalize it by its own norm. These conventions preserve the counts under $\mathcal{N} \to c\mathcal{N}$ (Section 4.4).

The support of the residual term is the most practitioner-legible item: the gradient $2\lambda G_{\mathrm{p}}\alpha$ lies in a subspace of dimension $\operatorname{rank} G_{\mathrm{p}}$, and the atoms with $\mathcal{N}[\varphi_i] \neq 0$ index the only coefficients the penalty depends on. On the fourteen-atom wave library the penalty reduces, to the digits shown, to $\lambda(4781.12\,\alpha_{12}^2 - 594.80\,\alpha_{12}\alpha_{13} + 4781.12\,\alpha_{13}^2)$, every other entry of G_{p} being 0.000: a two-coordinate ridge on the two spurious atoms and no function of the twelve carrying the wave content. A noise-calibrated companion count is defined in Appendix B.1.6.

3.2 Finite precision and tolerance

A zero in the theory is exact; a zero in the computation is an eigenvalue below a threshold. With H_{ref} and s_{ref} defined by Algorithm 1, let $V_\tau = \operatorname{span}\{v_j : \lambda_j(H_{\mathrm{ref}}) \leq \tau s_{\mathrm{ref}}\}$ and $d_\tau = \dim V_\tau$. When $G_{\mathrm{d}} \neq 0$, $s_{\mathrm{ref}} = \|G_{\mathrm{d}}\|_2$; when $G_{\mathrm{d}} = 0$ the physics block provides the reference scale. Always $V_0 \subseteq V_\tau$, and the audit's job is to say how much of V_τ is V_0. It distinguishes three tiers: *strict* dependencies, zero in exact arithmetic and confirmed by an identity, all that Lemma 1 lower-bounds; *near-zero* directions below the working cut, poorly resolved by the Gram eigendecomposition but not necessarily by a direct SVD of the stacked maps, and not identities; and *near* dependencies whose count moves with the threshold, the sloppy regime of Chis et al. (2016), where an integer reports the threshold rather than the dictionary. The procedure scans nine logarithmically spaced cuts $\tau \in [10^{-14}, 10^{-6}]$ and reports the interval on which each count is constant. A plateau establishes numerical stability alone: a small nonzero eigenvalue below the scanned range produces the same plateau, so strict dependencies need the separate 40-digit argument of Section 4.4. The rule is: **At the operating cut a local plateau requires agreement at 10^{-9}, 10^{-10} and 10^{-11}.** A constant run elsewhere is reported as an interval, not as operating-cut stability. Thus heat1d has a $d_\tau = 12$ run from 10^{-10} to 10^{-13} while the mean-field late window's seven directions end at 10^{-10} and have no such plateau. This rule checks dimension only: the orientation of V_τ and the projection of a requested reading onto it also require inspection before a numerical reading is called stable. Accordingly we distinguish three report states: a strict identity confirmed analytically, a near-zero reading stable only under a stated numerical cut and scale, and a threshold-sensitive case with no reading verdict. The exact uniqueness test using V_0 cannot be applied to V_τ without the stated approximation. We use $\tau = 10^{-10}$ throughout, fixed before evaluation.

4 EXPERIMENTS

Section 4.1 asks what the audit finds and compares refits on two registered sampling designs in a public pipeline; later sections ask whether it separates the phenomena it should, what it detects beyond simpler checks, and how stable its counts are. Every number is script-generated and test-pinned in the accompanying code supplement.

4.1 A PUBLIC EXAMPLE: FROM AUDIT TO REFITTING

Table 2: Audit of PySINDy's mean-field example (Kaptanoglu et al., 2021; 2022): degree-2 library (30 coefficients), energy-preserving constraint, and registered sampling schemes. A is the first-equation x_0x_2 coupling. †Threshold-sensitive count with no tolerance plateau (Appendix B.2.2).

observations	samples	data-blind	closed	jointly blind	reading of A
published recipe (long transient)	49999	0	0	0	determined
steady state (on the limit cycle)	4998	15	8	**7**	**not determined**
steady state, four times the data	19998	15	8	**7**	**not determined**
transient, same number of samples	4998	0	0	0	determined
two initial conditions, same number	4996	0	0	0	determined
late window of the published trajectory	4998	12	5	7†	no strict claim

On steady-state data, 15 of 30 directions are data-blind, the constraint closes 8, and 7 remain jointly blind across all nine cuts (Table 2). Growth g, coupling A, rotation rate and slaving are individually undetermined. With cycle height $m = 10^{-2}$, $x_2 = m$ makes $gx_0 + Ax_0x_2 = (g + mA)x_0$; this combination is determined. Exact rational elimination on the full 30-coefficient library confirms that $g + mA$ annihilates all seven joint null directions; g and A each vary along them (Appendix B.2.2). Two explicit witnesses with $A = -1.5$ and $A = -0.5$ preserve sampled derivatives and the constraint to numerical precision, yet differ from the true trajectory by 24% and 50% over $0 \le t < 50$. TrappingSR3 on the steady-state data achieves a 1.7×10^{-5} relative derivative error while omitting growth, coupling and slaving (exploratory; Appendix B.2.2). Four times more steady-state data leaves seven blind directions; equal-sample-size transient or two-initial-condition data remove all. The late-window count has no tolerance plateau and supports no strict claim. That transients help is known for this model (Brunton et al., 2016); the audit identifies the specific coefficient readings and constraint contributions behind that observation.

An *exploratory* matched-sample-size refit compares the two designs (Table 3; protocol in Appendix B.2.2). The returned transient fit recovers all nine true terms and lowers held-out trajectory error at the optimizer's iteration limit. The audit resolves the structural question; solver convergence is assessed separately.

Table 3: Matched-sample-size TrappingSR3 refit (exploratory). Coefficient columns are absolute errors; support is recovered true terms out of nine, then false positives, at cut 0.005. Trajectory error averages three held-out initial conditions. Full protocol in Appendix B.2.2.

design	$\lvert\hat{g} - g\rvert$	$\lvert\hat{A} - A\rvert$	$\lvert\widehat{g + mA} - (g + mA)\rvert$	support	trajectory	converged
steady, 4998	1.00×10^{-2}	1.00	3.79×10^{-7}	$2/9, +2$	27.5%	yes
transient, 4998	3.90×10^{-6}	5.68×10^{-6}	3.84×10^{-6}	$9/9, +0$	0.023%	no

Two further registered studies cover nine auditable PySINDy and DeePyMoD designs (Appendix B.2.3). Eight have no blind direction, including the previously known C3 case (seven of eight after excluding it). In the inviscid MHD exception, the energy constraint closes all 24 data-blind directions among 168 coefficients. The audit locates their equation-implied readings, beyond the warning from a 10^{11} condition number.

4.2 SEPARATING THE THREE PHENOMENA

Nineteen dictionaries in three pools, of sizes $K = 2$ to $K = 43$: twelve *design*, five *held-out* written after every constant was frozen, and two built by PySINDy's generator rules (Kaptanoglu et al., 2022) (construction and full audit in Appendix D, Table 9). Its purpose-built probes measure detection on selected designs, with no population-frequency estimate.

Table 4: Five dictionaries, each a different answer to the two questions of Section 1, at $\tau = 10^{-10}$; the first two are the running example and its two-function variant. *blind* counts eigenvalues of G_{d} at most $\tau\|G_{\mathrm{d}}\|_2$ and d_τ is the near-zero count of Section 3.2, both numerical (Table 9).

dictionary	K	blind	d_τ	Δr_τ	$\operatorname{rank} G_{\mathrm{p}}$	reading
$\varphi_1, \varphi_2, \varphi_3$	3	1	1	0	0	equivalent representations; penalty inert
φ_2, φ_3	2	0	0	0	0	coefficients pinned by data; penalty still inert
wave1d_minimal	2	0	0	0	0	identifiable, penalty inert
aliased_heat	6	1	0	1	2	data-blind direction resolved by the operator
wave1d_full	14	8	8	0	2	eight near-zero, six strict; penalty touches two atoms

Table 4 separates unidentifiability, an inert penalty and a penalty that completes the observations: removing the redundant function eliminates both the data-null and joint-null directions while the residual block remains zero, whereas the aliased probe is the mirror image, one blind direction and no gauge direction, so $\Delta r_\tau = 1$. Each finding has its own remedy. For the exact equivalence the audit returns the witness $(1, -\frac{1}{2}, -\frac{1}{2})$, and the practitioner reports the equivalence class or fixes a representative and says so. To make the aliased probe *data-identifiable* without relying on the imposed operator, the observation design must change: on the aliased probe the blind direction survives refining the residual grid from 4096 to 65536 points and disappears when the observation grid is refined to 16384 points or shifted off the alias. Across the corpus $\Delta r_\tau > 0$ on exactly three libraries and $\operatorname{rank} G_{\mathrm{p}} = 0$ on four (Table 9).

Near-zero directions are common. Seven of nineteen dictionaries have a nonzero near-null space V_τ at the operating tolerance, the share running to 8 of 14 on the wave and Klein–Gordon libraries (Figure 2); how much of each is exact functional dependency and how much numerical near-dependency is the question of Section 4.4.

Four dictionaries carry a residual term that is the zero function. On wave1d_minimal, wave2d, wave2d_quadruple and the held-out euler_bernoulli_null every atom solves the imposed equation exactly, so $R = 0$, $\operatorname{rank} G_{\mathrm{p}} = 0$, and the residual term is identically zero on the whole representation space: no gradient at any λ, and on wave1d_full and klein_gordon gradient on two coordinates only. A dictionary built from exact solutions has already encoded the equation, as Trefftz methods do by design (Hiptmair et al., 2016); the fault is a diagnostic that reports it as constraint. Two of the four come from the companion submission's benchmark suite (Anonymous, 2026), whose triple-enumeration diagnostic returns 0.000 on all four, a score documented to mean that the residual loss supplies non-redundant constraint; the audit reports, and symbolic differentiation confirms, that it supplies none. Three of the four also have $\dim V_0 = 0$, the case Section 2.3 distinguished from a gauge direction.

4.3 DETECTION POWER AND THE CLASSICAL COMPUTATION

Coherence recovers 15 of 35 working-precision dimensions and triple enumeration 30; the latter misses the order-5 identity of Proposition 3 and four heat1d directions. On heat1d, its 28 relations have projected rank 12 but only eight principal angles within tolerance, so projected rank would over-count. Section 4.4 confirms only 19 of the 35 as strict identities. A relation with nonzero *out* does not certify lossless atom deletion; the stacked SVD agrees with the audit at the same cut.

4.4 THRESHOLD, PRECISION AND NORMALIZATION

A tolerance-dependent count requires a stated cut; three checks quantify its stability and interpretation.

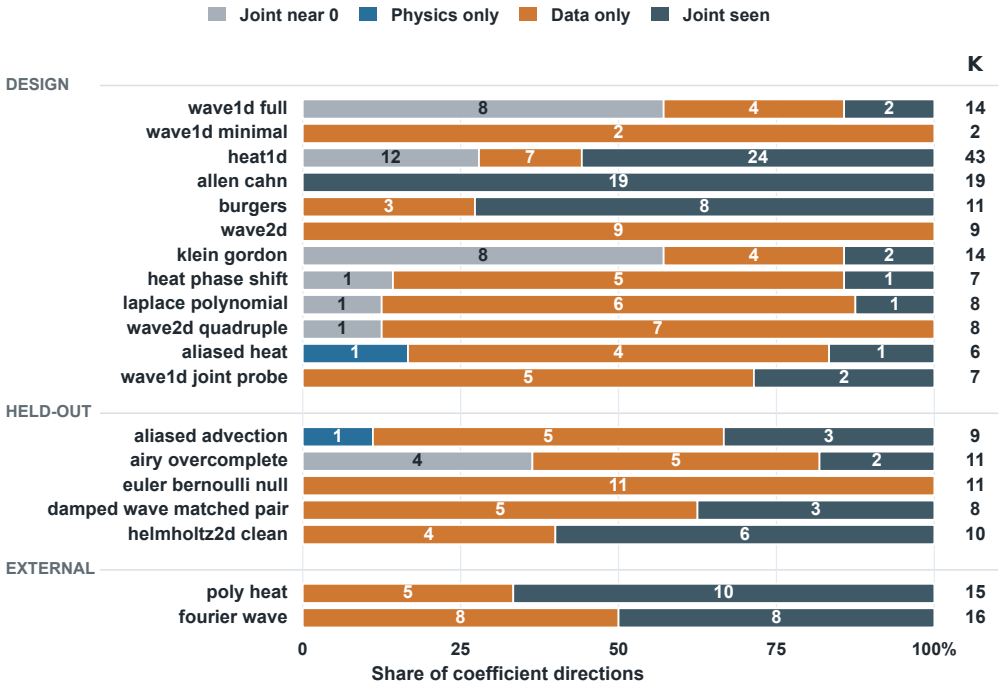


Figure 2: Composition of the operational direction bins for all nineteen dictionaries. Bar widths show fractions of K; numbers inside segments are direction counts, and the right column gives K. The joint near-zero bin has d_τ directions at $\tau = 10^{-10}$. On its complement, the physics-only bin counts $\nu_j \geq 1 - \tau_\nu$, the data-only bin counts $\nu_j \leq \tau_\nu$, and the jointly seen bin counts $\tau_\nu < \nu_j < 1 - \tau_\nu$, with $\tau_\nu = 10^{-8}$. The bins sum to K by construction. The separate numerical rank gain Δr_τ is reported in Table 9, not inferred from the physics-only bin. For nonlinear operators the physics block is the Fréchet Jacobian at the anchored solution, so these are local first-order bins. Their physics-visible sum need not equal the independently thresholded $\operatorname{rank} G_{\mathrm{p}}$ in Table 9; heat1d is the example. Figure 3 shows tolerance changes and Table 6 which near-zero directions are strict.

Table 5: Three detectors on the seven dictionaries with $V_\tau \neq 0$. *rec*: principal angles between returned relations and V_τ with $\sin\theta_i \leq 10^{-3}$ (a different criterion from the triple screen's equally valued residual cut); *out* and *out*$_0$: largest share outside V_τ and the confirmed identities, respectively. The final column is a plain stacked-SVD check at the matching cut, agreeing on all nineteen libraries (Appendix B.1.3).

			pairwise (coherence)			triple enumeration			audit	stacked SVD
library	K	d_τ	rec	out	out$_0$	rec	out	out$_0$	dim	dim
wave1d_full	14	8	7	0.63	0.71	8	0.63	0.71	**8**	8
klein_gordon	14	8	8	0.71	0.71	8	0.71	0.71	**8**	8
heat1d	43	12	0	–	–	**8**	1.00	1.00	**12**	12
airy_overcomplete	11	4	0	–	–	4	0.00	0.00	**4**	4
heat_phase_shift	7	1	0	–	–	1	0.00	0.00	**1**	1
laplace_polynomial	8	1	0	–	–	1	0.00	0.00	**1**	1
wave2d_quadruple	8	1	0	–	–	**0**	–	–	**1**	1
corpus total (19)		35	15			30			**35**	35

Precision. Table 6 separates near-zero from strict identities. Four libraries have only confirmed identities; in wave1d_full and klein_gordon, $d_\tau = 8$ includes six identities and two directions at the float64 floor. Heat1d has no strict identity; its $d_\tau = 12$ is a numerical count.

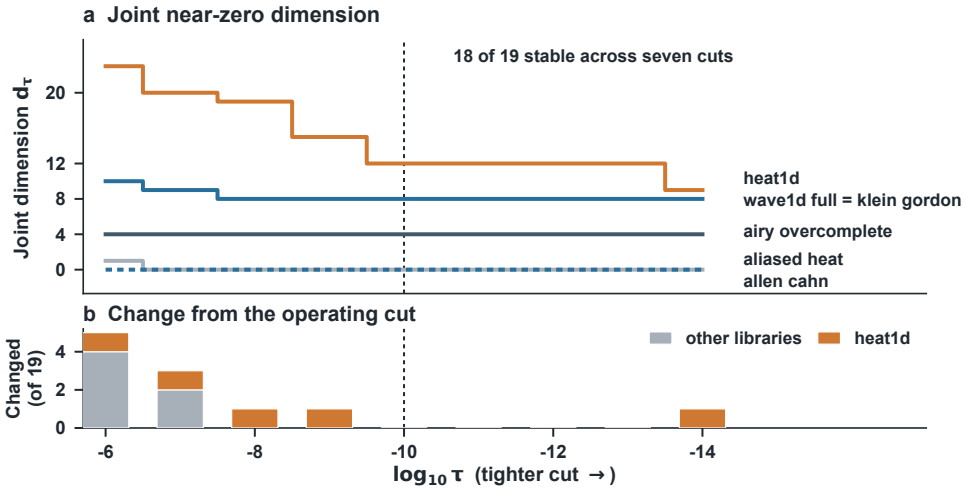


Figure 3: We scan the eigenvalue cut of H_{ref} over nine log-spaced values (eight orders of magnitude) on all nineteen dictionaries: (a) the joint near-zero dimension d_τ each reports, the quantity the audit outputs; (b) how many dictionaries change their count at each cut. Dotted line: operating point.

Table 6: Near-zero against strict, on every dictionary with a near-zero subspace. d_τ is the audit's count at the operating cut; *identities* the number of independent relations in the 40-digit null space confirmed *symbolically*, $\sum_i c_i \varphi_i$ simplifying to exactly 0 with the wavenumbers as symbols; *next* the smallest nonzero eigenvalue at 40 digits.

library	d_τ	identities	next	reading
heat_phase_shift	1	1	3.9×10^{-2}	strict
laplace_polynomial	1	1	5.4×10^{-3}	strict
wave2d_quadruple	1	1	1.1×10^{-1}	strict
airy_overcomplete	4	4	4.4×10^{-2}	strict
wave1d_full	8	6	1.5×10^{-16}	6 strict, 2 near-zero at the float64 floor
klein_gordon	8	6	1.1×10^{-15}	6 strict, 2 near-zero at the float64 floor
heat1d	12	0	1.8×10^{-26}	near: no strict relation

Threshold. All but heat1d retain their count over seven cuts in $[10^{-14}, 10^{-8}]$ (Figure 3); heat1d ranges from 23 to 9 across the full scan. At 10^{-6}, wave1d_joint_probe and locally linearized Allen–Cahn each acquire one spurious direction; both wave-library counts rise from 8 to 10. **Normalization.** Rescaling $(\mathcal{N}, \lambda)$ to $(c\mathcal{N}, \lambda/c^2)$ preserves counts; singular values change by at most 3×10^{-15} relative to the largest for $c \in [10^{-4}, 10^{4}]$. Fixing $\lambda = 1$ instead changes the heat count. Noise-calibrated counts use the raw Gram to retain magnitudes (Appendix B.2.1).

5 Conclusion

The audit uses the classical joint null space to return witnesses and test coefficient readings. Across nineteen dictionaries, it separates exact identities from numerical near-dependencies and finds relations missed by low-order screens. In the public case, seven jointly blind directions leave g and A individually undetermined while preserving $g + mA$. Equal-budget transient observations remove these directions; an exploratory refit recovers all nine true terms at the optimizer's iteration limit.

These results connect structural diagnosis to sampling choices and coefficient interpretation. Conclusions depend on the dictionary, sampling design and imposed operator (Theorem 1). Nonlinear and hybrid cases receive local first-order diagnoses; the selected corpus provides case-specific evidence (Appendix B.3). Predictive accuracy and operator validity require separate evaluation.

AI USE STATEMENT

Generative AI tools were used in three parts of this work, and their role differs by part. *Research artifacts.* The five held-out libraries (Appendix D) were constructed by LLM agents that were given only the four-class taxonomy, with no access to the diagnostic, to its thresholds or to any of its outputs; that use is part of the frozen evaluation protocol and is stated where the pool is introduced. *Code.* The implementation, the experiment drivers and the tests were written with AI coding assistance. *Writing.* AI tools were used for drafting and editing the manuscript and for figure layout.

We have reviewed all AI-assisted work. No agent output is used unverified: for every library, including the agent-designed ones, the residuals and the identities this paper reports are recomputed by symbolic differentiation in the pipeline (Appendix D), and the numerical claims are regenerated by repository scripts and guarded by tests that the authors ran and inspected (code supplement). The problem formulation, the proofs, the design constants and the order in which they were frozen, and every reported value were checked by the authors. We take responsibility for the final content of this work, including the text, the claims and the artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This is a methodological study on synthetic dictionaries and analytically specified sampling designs. It involves no human subjects, no personal data and no new data collection, and it releases no model that could be deployed on people. The audit's scope and limitations are stated with the method (Section 3) and in the conclusion (Section 5).

REPRODUCIBILITY STATEMENT

Every reported number is produced by scripts in the accompanying repository. For a linear operator all audited quantities are closed form, so the whole corpus regenerates in minutes on a single CPU core with no training, and the reproduction guards are run as tests (code supplement). The audit is stated line by line in Appendix B.1, the dictionaries are audited in Appendix D and defined in full in the code supplement, and the order in which the design constants and held-out pool were frozen is recorded in the accompanying protocols. Because the dimension estimates depend on choices that are easy to leave implicit, the sensitivity of the reported counts to tolerance, working precision and column normalization is reported explicitly in Section 4.4, and Appendix B.1.1 separates what float64 spectra can support from what requires exact arithmetic.

REFERENCES

Anonymous. Symbolic where possible, neural where necessary: Dual-level interpretability for PDE reconstruction from noisy data, 2026. Companion submission, under double-blind review.

Gert-Jan Both, Subham Choudhury, Pierre Sens, and Remy Kusters. DeepMoD: Deep learning for model discovery in noisy data. *Journal of Computational Physics*, 428:109985, 2021.

Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.

Emmanuel J. Candès, Yonina C. Eldar, Deanna Needell, and Paige Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1): 59–73, 2011.

Cecilia Casolo, Sören Becker, and Niki Kilbertus. Identifiability challenges in sparse linear ordinary differential equations. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2506.09816.

Zhao Chen, Yang Liu, and Hao Sun. Physics-informed learning of governing equations from scarce data. *Nature Communications*, 12:6136, 2021.

Oana-Teodora Chis, Alejandro F. Villaverde, Julio R. Banga, and Eva Balsa-Canto. On the relationship between sloppiness and identifiability. *Mathematical Biosciences*, 2016.

Fabrício Olivetti de França et al. Improving genetic programming for symbolic regression with equality graphs. *Philosophical Transactions of the Royal Society A*, 2025.

Charles B. Delahunt and J. Nathan Kutz. A toolkit for data-driven discovery of governing equations in high-noise regimes. *IEEE Access*, 10:31210–31234, 2022.

David L. Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization. *Proceedings of the National Academy of Sciences*, 100(5): 2197–2202, 2003.

Nathan Doumèche, Francis Bach, Gérard Biau, and Claire Boyer. Physics-informed machine learning as a kernel method. In *Proceedings of the 37th Conference on Learning Theory (COLT)*, volume 247 of *PMLR*, pp. 1399–1450, 2024.

Nathan Doumèche, Francis Bach, Gérard Biau, and Claire Boyer. Physics-informed kernel learning. *Journal of Machine Learning Research*, 26(124), 2025.

Urban Fasel, J. Nathan Kutz, Bingni W. Brunton, and Steven L. Brunton. Ensemble-SINDy: Robust sparse model discovery in the low-data, high-noise limit, with active learning and control. *Proceedings of the Royal Society A*, 478:20210904, 2022.

Matteo Gallo, Fabio Anselmi, and Paolo Lazzari. Attractor geometry determines the identifiability limits of system discovery. arXiv:2607.18490, 2026.

Per Christian Hansen. Regularization, GSVD and truncated GSVD. *BIT Numerical Mathematics*, 29:491–504, 1989.

Per Christian Hansen. Regularization Tools version 4.0 for Matlab 7.3. *Numerical Algorithms*, 46: 189–194, 2007.

Koji Hashimoto, Yuji Hirono, and Akiyoshi Sannai. Unification of symmetries inside neural networks: Transformer, feedforward and neural ODE. arXiv:2402.02362, 2024.

Hillary Hauger, Philipp Scholl, and Gitta Kutyniok. Robust identifiability for symbolic recovery of differential equations. arXiv:2410.09938, 2024.

Ralf Hiptmair, Andrea Moiola, and Ilaria Perugia. A survey of Trefftz methods for the Helmholtz equation. In *Building Bridges: Connections and Challenges in Modern Approaches to Numerical Partial Differential Equations*, volume 114 of *Lecture Notes in Computational Science and Engineering*, pp. 237–279. Springer, 2016.

Carl Jidling, Niklas Wahlström, Adrian Wills, and Thomas B. Schön. Linearly constrained Gaussian processes. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Kadierdan Kaheman, J. Nathan Kutz, and Steven L. Brunton. SINDy-PI: a robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *Proceedings of the Royal Society A*, 476:20200279, 2020.

Alan A. Kaptanoglu, Jared L. Callaham, Christopher J. Hansen, Aleksandr Aravkin, and Steven L. Brunton. Promoting global stability in data-driven models of quadratic nonlinear dynamics. *Physical Review Fluids*, 6:094401, 2021.

Alan A. Kaptanoglu, Brian M. de Silva, Urban Fasel, Kadierdan Kaheman, Andy J. Goldschmidt, Jared Callaham, Charles B. Delahunt, Zachary G. Nicolaou, Kathleen Champion, Jean-Christophe Loiseau, J. Nathan Kutz, and Steven L. Brunton. PySINDy: A comprehensive Python package for robust sparse system identification. *Journal of Open Source Software*, 7(69):3994, 2022.

Ronald Katende. Counterfactual operator relevance for PDE discovery: Screening, pruning, and identifiability. arXiv:2506.20181, 2025.

Takeshi Koshizuka and Issei Sato. Understanding generalization in physics informed models through affine variety dimensions. arXiv:2501.18879, 2025.

Aditi S. Krishnapriyan, Amir Gholami, Shandian Zhe, Robert M. Kirby, and Michael W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, 2021.

Raphael Leiteritz and Dirk Pflüger. How to avoid trivial solutions in physics-informed neural networks. arXiv:2112.05620, 2021.

Jean-Christophe Loiseau and Steven L. Brunton. Constrained sparse Galerkin regression. *Journal of Fluid Mechanics*, 838:42–67, 2018.

Torkel E. Loman, Alexander P. Browning, and Ruth E. Baker. Structural functional identifiability and model discovery in differential equation models. arXiv:2606.30289, 2026.

Niall M. Mangan, Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Inferring biological networks by sparse identification of nonlinear dynamics. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, 2(1):52–63, 2016.

Daniel A. Messenger and David M. Bortz. Weak SINDy for partial differential equations. *Journal of Computational Physics*, 443:110525, 2021.

Erion Morina, Philipp Scholl, and Martin Holler. Symbolic recovery of PDEs from measurement data. arXiv:2602.15603, 2026.

Sangnam Nam, Mike E. Davies, Michael Elad, and Rémi Gribonval. The cosparse analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30–56, 2013.

M. Z. Naser. Fundamental flaws of physics-informed neural networks and explainability methods in engineering systems. *Computers & Industrial Engineering*, 212:111704, 2026.

Cornelius Otchere and Michael Shields. Effective dimensionality as an operator invariant for physics-preserving constraint adaptation in physics-informed neural networks. arXiv:2606.06171, 2026.

C. C. Paige and M. A. Saunders. Towards a generalized singular value decomposition. *SIAM Journal on Numerical Analysis*, 18(3):398–405, 1981.

Yang Pan and Helmut Bölcskei. Recovering governing equations from solution data: Identifiability bounds for linear and nonlinear ODEs. arXiv:2606.27285, 2026.

Samuel H. Rudy, Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Data-driven discovery of partial differential equations. *Science Advances*, 3(4):e1602614, 2017.

Fabian Scheipl and Sonja Greven. Identifiability in penalized function-on-function regression models. *Electronic Journal of Statistics*, 10(1):495–526, 2016.

Philipp Scholl, Aras Bacho, Holger Boche, and Gitta Kutyniok. Symbolic recovery of differential equations: The identifiability problem. *Machine Learning*, 115:139, 2026. doi:10.1007/s10994-026-07068-0; first preprint arXiv:2210.08342, 2022.

Tejas Pradeep Shirodkar. Dead directions: Geometric singular learning. arXiv:2606.05957, 2026.

Zakhar Shumaylov, Peter Zaika, Philipp Scholl, Gitta Kutyniok, Lior Horesh, and Carola-Bibiane Schönlieb. When is a system discoverable from data? Discovery requires chaos. arXiv:2511.08860, 2025.

Joel A. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, 2004.

Charles F. Van Loan. Generalizing the singular value decomposition. *SIAM Journal on Numerical Analysis*, 13(1):76–83, 1976.

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021.

Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022.

Zongren Zou, Xuhui Meng, and George Em Karniadakis. Correcting model misspecification in physics-informed neural networks, 2024.

A PROOFS

Notation as in Section 2; G_{d}, G_{p} and H_λ are symmetric positive semidefinite. Numerical checks are included in the released code supplement.

Lemma 2 (consistency)**.** $b := \Phi u_{\mathrm{data}}/N_{\mathrm{d}} \in \mathrm{range}(G_{\mathrm{d}}) \subseteq \mathrm{range}(H_\lambda)$ *for every realization of* u_{data}*; hence* J_λ *attains its minimum and the normal equations* $H_\lambda \alpha = b$ *are consistent.*

Proof. $b \in \mathrm{range}(\Phi) = \mathrm{range}(\Phi\Phi^\top) = \mathrm{range}(G_{\mathrm{d}})$. Both summands of H_λ are positive semidefinite, so $\ker H_\lambda \subseteq \ker G_{\mathrm{d}}$ for every $\lambda \geq 0$, and taking orthogonal complements gives $\mathrm{range}(G_{\mathrm{d}}) \subseteq \mathrm{range}(H_\lambda)$. □

Proof of Proposition 1. J_λ is a convex quadratic with Hessian $2H_\lambda$; by Lemma 2 a minimizer exists and the minimizer set is $\hat{\alpha} + \ker H_\lambda$. For a positive semidefinite G, $v^\top G v = 0$ iff $Gv = 0$. Applied to both summands with $\lambda > 0$, $H_\lambda v = 0$ iff $v^\top G_{\mathrm{d}} v = 0$ and $v^\top G_{\mathrm{p}} v = 0$, i.e. iff $v \in V_0$. □

Proof of Lemma 1. The p atoms lie in a d-dimensional space, so at least $p - d$ linearly independent coefficient vectors c satisfy $\sum_i c_i \varphi_i = 0$ *as a function*. For each: $\Phi^\top c = 0$ at any data points, and by linearity $\sum_i c_i \mathcal{N}[\varphi_i] = \mathcal{N}[\sum_i c_i \varphi_i] = \mathcal{N}[0] = 0$, so $R^\top c = 0$ at any collocation points; hence $c \in V_0$. □

Proof of Proposition 2. Write $G_{\mathrm{p}} = LL^\top$. $\ker M$: $Mc = 0$ iff $\|L^\top Bc\| = 0$ iff $Bc \in \ker G_{\mathrm{p}}$; and $Bc \in \ker G_{\mathrm{d}}$ always, so $Bc \in V_0$. Conversely every $v \in V_0 \subseteq \ker G_{\mathrm{d}}$ is $v = Bc$ with $c = B^\top v$ and $Mc = 0$, so $B \ker M = V_0$. Range: $B^\top G_{\mathrm{p}} \alpha_{\mathrm{vis}} = B^\top L (L^\top \alpha_{\mathrm{vis}}) \in \mathrm{range}(B^\top L) = \mathrm{range}(M)$, so the normal equations $Mc = -B^\top G_{\mathrm{p}} \alpha_{\mathrm{vis}}$ are consistent, $c^\star = -M^+ B^\top G_{\mathrm{p}} \alpha_{\mathrm{vis}}$ solves them, and the quadratic $c \mapsto \|A_{\mathrm{p}}(\alpha_{\mathrm{vis}} + Bc)\|^2 = (\alpha_{\mathrm{vis}} + Bc)^\top G_{\mathrm{p}} (\alpha_{\mathrm{vis}} + Bc)$ has minimizer set $c^\star + \ker M$; its image under $c \mapsto \alpha_{\mathrm{vis}} + Bc$ is $\alpha_{\mathrm{vis}} + Bc^\star + V_0$, on which $\ell^\top \alpha$ is constant iff $\ell^\top v = 0$ for every $v \in V_0$. □

Implementation note. The code additionally caps $\|Bc\|$ at $\|\alpha_{\mathrm{vis}}\|$ as a guard against ill-conditioning. If the cap activates, the returned vector is a *truncated approximate completion*; scaling $c^\star$ need not minimize the residual and Proposition 2 does not describe that returned vector. The cap is inactive on both aliasing probes of the corpus (aliased_heat and aliased_advection).

Proof of Proposition 3. $[A_{\mathrm{d}}; A_{\mathrm{p}}]v = 0$ iff $A_{\mathrm{d}} v = 0$ and $A_{\mathrm{p}} v = 0$, so $\ker[A_{\mathrm{d}}; A_{\mathrm{p}}] = \ker A_{\mathrm{d}} \cap \ker A_{\mathrm{p}} = V_0$, and a rank-revealing decomposition returns a basis of the kernel of the matrix it is given, with no condition on the supports of the basis vectors. A relation with support at most three that annihilates both maps is a vector of V_0. Exhaustively enumerating and spanning all such relations gives $V_{\leq 3}$ by definition, so some direction of V_0 is missed exactly when $V_{\leq 3} \neq V_0$. A partial enumeration returns only a subspace of $V_{\leq 3}$. □

Proof of Theorem 1. (a) $\Phi^\top \alpha_B = \Phi^\top \alpha_A + t\, \Phi^\top v = \Phi^\top \alpha_A$ exactly, and the noise law does not depend on the truth, so the laws of $u = \Phi^\top \alpha + \varepsilon$ coincide. (b) The law of $g(u)$ is a functional of the law of u; with $q := P(g{=}1)$ common to both and opposite correct answers, the success probability under the uniform prior is $\frac{1}{2}q + \frac{1}{2}(1 - q) = \frac{1}{2}$. □

Proof of Proposition 4. $\Phi^\top v = 0$, so fitted values $\Phi^\top \hat{\alpha}$ do not depend on data-blind components of $\hat{\alpha}$, nor $\Phi^\top \alpha_\star$ on those of $\alpha_\star$; every listed criterion is a functional of fitted values (and u). For GCV specifically, procedures differing only in data-blind components induce the identical fitted-value map $u \mapsto \Phi^\top \hat{\alpha}(u)$, hence identical hat matrix and identical effective degrees of freedom. □

B ALGORITHM, EXPERIMENTAL DETAILS AND SCOPE

This appendix states the audit line by line with its numerical conventions (Appendix B.1), gives the details of the experiments the main text summarizes (Appendix B.2), and states the scope of the claims (Appendix B.3). It also states Proposition 4 and Theorem 1.

B.1 ALGORITHM AND NUMERICAL CONVENTIONS

Algorithm 1 The structural audit. Input: dictionary $\{\varphi_i\}_{i=1}^{K}$, imposed operator $\mathcal{N}$ (or a coefficient constraint), observation and residual points, and cuts τ, τ_ν. A nonlinear operator also requires a linearization point u_0. No observed response values, training, or user-selected penalty weight are required in the linear case.

1: $\Phi_{in} \leftarrow \varphi_i(z_n)$; for linear $\mathcal{N}$, set $R_{im} \leftarrow \mathcal{N}[\varphi_i](y_m)$; for nonlinear $\mathcal{N}$, set $R_{im} \leftarrow D\mathcal{N}[u_0](\varphi_i)(y_m)$. Rescale both rows by the observation rms where nonzero.
2: $G_{\mathrm{d}} \leftarrow \Phi\Phi^{\top}/N_{\mathrm{d}}$, $G_{\mathrm{p}} \leftarrow RR^{\top}/N_{\mathrm{c}}$
3: **if** $G_{\mathrm{d}} = G_{\mathrm{p}} = 0$: set $H_{\mathrm{ref}} \leftarrow 0$, $s_{\mathrm{ref}} \leftarrow 1$; **else if** $G_{\mathrm{p}} = 0$: set $H_{\mathrm{ref}} \leftarrow G_{\mathrm{d}}$, $s_{\mathrm{ref}} \leftarrow \|G_{\mathrm{d}}\|_2$; **else if** $G_{\mathrm{d}} = 0$: set $\lambda_{\mathrm{ref}} \leftarrow 1/\|G_{\mathrm{p}}\|_2$, $H_{\mathrm{ref}} \leftarrow \lambda_{\mathrm{ref}} G_{\mathrm{p}}$, $s_{\mathrm{ref}} \leftarrow 1$; **else**: set $\lambda_{\mathrm{ref}} \leftarrow \|G_{\mathrm{d}}\|_2/\|G_{\mathrm{p}}\|_2$, $H_{\mathrm{ref}} \leftarrow G_{\mathrm{d}} + \lambda_{\mathrm{ref}} G_{\mathrm{p}}$, $s_{\mathrm{ref}} \leftarrow \|G_{\mathrm{d}}\|_2$
4: eigendecompose H_{ref}; eigenvectors with eigenvalue $\leq \tau s_{\mathrm{ref}}$ span the near-zero subspace V_τ (Section 3.2), with witnesses
5: separately threshold G_{d} at $\tau\|G_{\mathrm{d}}\|_2$ and G_{p} at $\tau\|G_{\mathrm{p}}\|_2$ for the data-blind count and physics rank; set $\Delta r_\tau = \mathrm{blind} - d_\tau$ and obtain physics support from nonzero residual rows
6: on the complement of V_τ, whiten by H_{ref} and eigendecompose the whitened physics block for shares ν_j. Count $\nu_j \leq \tau_\nu$ as data-only, $\nu_j \geq 1 - \tau_\nu$ as physics-only, and the remaining j as jointly seen; keep Δr_τ as a separate rank count
7: **repeat** lines 4–6 over logarithmically spaced τ, including when $G_{\mathrm{p}} = 0$; report the constant-count intervals and the operating-cut stability verdict (Section 3.2)

Three conventions in Algorithm 1 decide what the counts mean. Line 1 uses symbolic derivatives, so R carries no stencil error. For nonlinear operators the resulting matrix is a Jacobian at u_0, so its null directions are only first-order sensitivity statements. The unit-rms rescaling makes numerical counts invariant to nonzero rescalings of atoms with nonzero observation rms. An atom zero on the observation design is left unscaled; its numerical counts are *not* unit-invariant without an additional scale convention. The weight λ_{ref} makes them invariant to a rescaling $\mathcal{N} \to c\mathcal{N}$ of the whole operator. Line 4 uses a symmetric eigensolver on $G_{\mathrm{d}} + \lambda_{\mathrm{ref}} G_{\mathrm{p}}$ rather than a generalized solver on the pair, so that $G_{\mathrm{p}} = 0$ (four of our libraries) needs no special treatment beyond the branch of line 3; only after whitening the complement of V_τ by H_{ref} does the whitened physics block have eigenvalues that are shares in $[0, 1]$. At a fixed tolerance the general branch uses up to four eigendecompositions; the zero-block branch needs only the data Gram decomposition. The cost is $O(K^2(N_{\mathrm{d}} + N_{\mathrm{c}}) + K^3)$ as in Section 3; the tolerance scan repeats lines 4–6 once per cut, and the symbolic differentiation, the 40-digit recomputation and the symbolic confirmation of identities are outside that count. Table 7 lists what the audit reports.

Table 7: What the audit reports, and what a reader learns from each item.

output	what it tells the practitioner
joint near-zero directions, with witnesses	coefficient combinations the two linearized maps weakly distinguish at the stated cut; only symbolic identities certify strict blindness
support and rank of the residual term	the coefficients the physics term depends on at all
numerical rank gain Δr_τ	the increase in numerical rank after stacking the physics block, at the stated normalization and cut; the exact direction interpretation is Eq. (5). An equation-implied completion requires a data-determined visible component (Proposition 2)
tolerance and precision diagnostics	intervals of constant dimension; individual witness directions and requested readings also require stability checks

B.1.1 EXACT AGAINST FLOAT64 SPECTRA

Figure 4 shows the three tiers of Section 3.2 on three libraries by recomputing the atom Gram at 40 decimal digits on the same random points as the float64 computation.

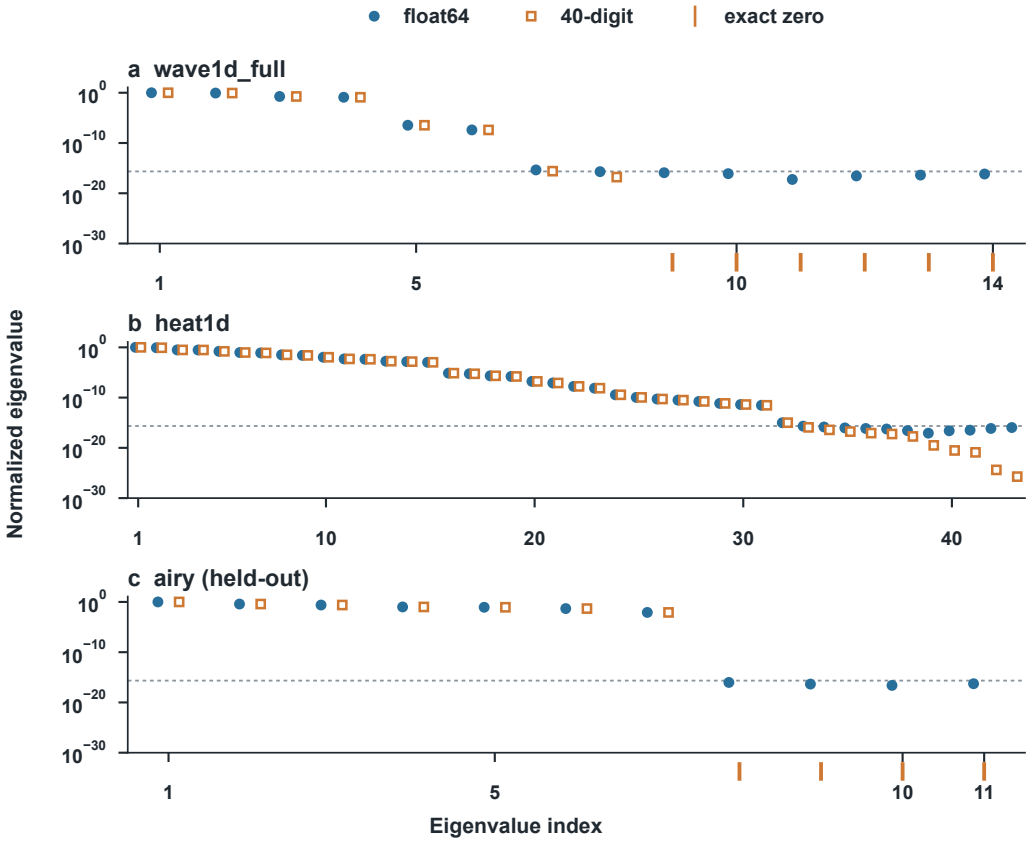


Figure 4: Normalized atom-Gram spectra $\lambda_i/\lambda_{\max}$ at float64 (filled circles) and 40-digit precision (open squares), computed on the same random points. Markers are offset slightly to show coincident values. Orange ticks below each axis mark symbolically confirmed zeros, which have no position on a logarithmic axis; the dotted line marks float64 resolution. The wave1d_full plateau of eight contains six strict identities, heat1d's plateau of twelve contains none, and all four Airy plateau modes are exact.

B.1.2 PREDICTION-ERROR CRITERIA

Proposition 4 (prediction-error criteria cannot see data-blind components)**.** *For $v \in \ker G_{\mathrm{d}}$, every functional of fitted values on the design (prediction error, design-based cross-validation, and GCV, whose degrees of freedom depend only on the fitted-value map) is invariant to $\langle v, \hat{\alpha}-\alpha_\star \rangle$. A criterion selecting λ by estimated prediction error therefore cannot register any benefit the penalty provides along data-blind directions.*

The audit tests this degeneracy separately from prediction-based tuning. The analogous invariance for smoothness penalties was observed by Scheipl & Greven (2016).

B.1.3 WHAT IS CLASSICAL

H_λ is the normal-equation matrix of general-form Tikhonov regularization; rewriting equation 1 as a stacked least squares on $[\Phi \mid \sqrt{\lambda}R]$ is immediate; and $\ker \Phi^\top \cap \ker R^\top = \{0\}$ guarantees

uniqueness. The last is the full-column-rank hypothesis of Van Loan's GSVD theorem (Van Loan, 1976), as presented by Hansen (1989); the general GSVD of Paige & Saunders (1981), the form LAPACK implements, exists for any pair, and a common null space appears in it as a zero block rather than as an obstruction. The condition characterizes uniqueness of the penalized minimizer: any $v \in \ker \Phi^\top \cap \ker R^\top$ leaves both terms invariant. We use this classical identification and its truncated-GSVD tolerance apparatus (Hansen, 1989).

The audit attaches witness directions, source attribution and reading tests to that dimension. The stacked-SVD column of Table 5 validates the implementation: at a matched cut a rank-revealing decomposition of $[A_\mathrm{d}; \sqrt{\lambda_\mathrm{ref}} A_\mathrm{p}]$ returns the same dimension on all nineteen libraries, as it must, since the eigenvalues of H_ref are the squared singular values of the stacked matrix; on the seven libraries with a near-zero subspace the two null spaces coincide as subspaces, with largest principal-angle sine 6.0×10^{-6} (heat1d, the library without a strict relation; at most 3×10^{-8} on the others) and residuals $\|A_\mathrm{d} v\|$, $\sqrt{\lambda_\mathrm{ref}}\|A_\mathrm{p} v\|$ of at most 9.1×10^{-7} on every returned basis vector. Whether a direction is an exact identity or an artifact of this sampling design is invisible in a float64 spectrum and takes the forty-digit and symbolic passes of Table 6, which confirm 19 of the 35 near-zero directions found across the corpus. The count is well defined only once the cut is anchored and the columns normalized: a detector that fixes $\lambda = 1$ reports 20, 14 or 8 directions for the same heat library as the operator is rescaled (Section 4.4), and a loose cut reports a spurious direction on a library that has none. And a dimension says nothing about which coefficients the residual penalty reaches, which separates the four classes of Table 1: on four libraries the penalty touches no coefficient at all. The comparison of screens in Table 5 is at the level of returned subspaces, through principal angles and the share of each witness outside V_τ, not at the level of counts.

B.1.4 Terminology

We call the directions in V_0 *gauge directions*, after the physical intuition of a gauge freedom: a change of coefficients that leaves both loss terms unchanged on the given design. The word is overloaded in the neural-network literature, where "gauge" and "dead" directions denote parametrization symmetries and Fisher-degenerate directions of a network (Hashimoto et al., 2024; Shirodkar, 2026); here the directions belong to a specified library–operator pair. "Joint-kernel directions" is an equivalent neutral name.

B.1.5 The existence hypothesis in practice

The hypothesis of Lemma 1 does not require the d-dimensional subspace S to consist of solutions; linearity is used only to pass $\mathcal{N}$ through the vanishing combination. The motivating case is a solution slice the library oversamples: each wavenumber group of the 1D-wave library places $p = 4$ atoms in a $d = 2$ slice (three groups, 6 identities), and the Airy library places $p = 9$ exact solutions in a $d = 5$ slice (4 identities; the bound is attained, all four confirmed symbolically). The general form also covers identities among atoms that are not all solutions (laplace_polynomial, $p = 3$, $d = 2$: x^2 and y^2 are not harmonic, $x^2 - y^2$ is). Its hypothesis can be checked on a library before any data exist, and the two generator-built libraries of our corpus happen not to meet it.

The exact classes and numerical bins of Table 1 have different interpretations. In exact arithmetic their subspaces form a direct-sum decomposition of the K-dimensional coefficient space. Figure 2 instead partitions the *retained generalized shares* into endpoint and interior bins after a separate joint near-zero cut. Its four integers sum to K by construction, but are not, in general, dimensions of the four exact subspaces, and its physics-only bin need not equal the separately computed numerical rank gain Δr_τ. Neither report by itself removes an atom. Exact class dimensions concern directions, not counts of atoms: exact redundancy can be eliminated by selecting a linearly independent subset of the columns of $[A_\mathrm{d}; A_\mathrm{p}]$, which keeps every output the two maps can produce on this design, but that selection is a choice of representative and is in general not unique (on the fourteen-atom wave library each of the six strict identities relates three atoms, any of which may be dropped, and the two further directions of V_τ are not identities at all, Table 6). When the dependency holds only at the sampling points, the reduced dictionary spans the same values on the design but a smaller function class away from it.

B.1.6 THE NOISE-CALIBRATED DATA-BLIND COUNT

The algebraic count ignores noise. Given a noise level σ and a target coefficient accuracy δ, both supplied by the user, a companion count is $\#\{j : \gamma_j < \sigma^2/(N_\mathrm{d}\delta^2)\}$, with γ_j the eigenvalues of the *unnormalized* data Gram $\Phi\Phi^\top/N_\mathrm{d}$ in the user's units: under the linear observation model the data-only least-squares estimate has standard deviation $\sigma/\sqrt{N_\mathrm{d}\gamma_j}$ along the jth eigenvector, so this is the number of directions the observations alone cannot resolve to δ. Exact null directions ($\gamma_j = 0$) are reported separately, and the fixed representative a pseudo-inverse assigns there, whose sampling variance is zero, is not a resolved component. The count does not bound what a trusted constraint can achieve, and it is not an error bound for biased estimators; it is reported alongside the algebraic one because "can I trust the coefficient of φ_i to three digits from these data" is the noise-calibrated question.

B.2 EXPERIMENTAL DETAILS

B.2.1 TRANSFER TO CONSTRAINED SPARSE REGRESSION

What the audit requires is a coefficient vector and two quadratic forms in it, and a second block appears in sparse regression $u_t = \Theta\xi$ as soon as linear constraints are imposed on ξ, as constrained and symmetry-informed sparse regression routinely do. For a quadratic Galerkin model $\dot{a}_k = c_k + \sum_i L_{ki}a_i + \sum_{i\le j} Q_{kij}a_ia_j$, energy preservation of the quadratic term is one linear condition per cubic monomial: with the two lower indices of Q sorted, $Q_{kij} + Q_{ijk} + Q_{jki} = 0$ for distinct i, j, k, $Q_{kii} + Q_{iik} = 0$ for $i \neq k$, and $Q_{iii} = 0$; constrained sparse regression imposes it as a hard constraint (Kaptanoglu et al., 2022), and the Lorenz coefficients satisfy it exactly. Stacking the three equations over a degree-2 polynomial library gives 30 coefficients and 10 constraint rows. For the data block we take the lack-of-excitation failure of system identification: the z-axis is an invariant set of Lorenz, so a trajectory started on it has $x = y = 0$ and every atom involving x or y vanishes along the data. There $\kappa(\Theta) = \infty$, and that is all a condition number reports. The audit reports 21 data-blind directions of 30, 9 of them closed by the constraint and 12 gauge: all six blind linear coefficients, since an energy constraint acts on quadratic coefficients only, and six of the fifteen blind quadratic ones. Starting at distance 10^{-6} from the axis and sweeping the window T, the trajectory escapes between $T = 1$ and $T = 2$ and the design is then well conditioned; before the escape the Gram is numerically singular (κ between 10^{13} and 10^{19}), the algebraic count is not monotone in T (9, 6, 9 at $T = 0.25$, 0.5, 1) and the scan finds no plateau, the third tier of Section 3.2; at $T = 0.25$ the noise-calibrated count of Appendix B.1.6 is 21 of 30 at $\sigma = 10^{-3}$, $\delta = 10^{-2}$, against an algebraic count of 9. A smaller variant imposing the Lorenz reflection symmetry on one equation gives 7 data-blind of 10, 4 determined by the symmetry, 3 gauge, with witnesses on x^2, xy and y^2. This transfer check shows how the same eigendecomposition separates directions hidden by a condition number; the corpus findings remain specific to their tested libraries.

B.2.2 THE MEAN-FIELD CASE

The case of Section 4.1 was registered before any of its quantities was computed; the protocol, with its commit and deviation log, ships with the code. It pins the source (PySINDy 2.1.0; the mean-field cell of the Trapping SINDy example at a fixed commit, by file hash), the observation schemes, the readings, the numeric criteria and the null outcomes, and discloses what was known beforehand: a hand analysis predicted that steady-state data determine $g + mA$ but not g and A separately. That growth and saturation are not separately identifiable from steady-state data is known in reduced-order modelling, and the transient-versus-attractor point was made on this very model (Brunton et al., 2016). This registered case demonstrates the audit's coefficient-level report on a public example. The protocol was followed without deviation.

The model is $\dot{x}_0 = \mu x_0 - x_1 - x_0x_2$, $\dot{x}_1 = \mu x_1 + x_0 - x_1x_2$, $\dot{x}_2 = -x_2 + x_0^2 + x_1^2$ with $\mu = 10^{-2}$, integrated with the example's settings; the library is PySINDy's degree-2 polynomial library with bias, ten terms per equation; the constraint is the 10×30 energy-preserving matrix returned by PySINDy itself, which the published coefficients satisfy with residual exactly 0. The audit concerns these two blocks only; the relaxation, the stability term, thresholds and the solver's truncation in the full TrappingSR3 objective choose a representative among the coefficient vectors the two blocks allow. One synthetic case says nothing about how often the situation arises.

Readings. A linear reading $\ell^\top \alpha$ in raw coefficient units is determined by data and constraint iff $B_0^\top \ell = 0$ for a basis B_0 of the jointly blind subspace (Proposition 2); B_0 is the audit's basis mapped back to raw coefficient coordinates and re-orthonormalized there. We call a reading not determined when $\|B_0^\top \ell\|/\|\ell\| > 10^{-6}$ and determined below 10^{-10}. To keep the estimated coefficient apart from the fixed weight, write g for the coefficient of x_0 in the first equation (true value μ) and $m = \mu = 10^{-2}$ for the height of the cycle. On the steady-state scheme the shares are 0.50 for the coupling A and for the slaving coefficient (the coefficient of x_0^2 in the third equation), 5.0×10^{-3} for g, 7.1×10^{-3} for the rotation rate, and 9.9×10^{-11} for $g + mA$; only these five readings were registered, and every ℓ with $B_0^\top \ell = 0$ is determined too. With four times the data the last share is 2.2×10^{-10}, between the two thresholds, and receives no *numerical* verdict. The strict conclusion for $g + mA$ comes instead from exact arithmetic: take ten distinct rational points on the cycle, parameterized by $x_0 = (1 - t^2)/(10(1 + t^2))$, $x_1 = 2t/(10(1 + t^2))$, $x_2 = m = 1/100$, with $t = 0, \ldots, 9$. Let D be the block-diagonal degree-2 library matrix at these points and let C be the imposed 10×30 constraint. Exact rational elimination gives $\operatorname{rank} D = 15$ and $\operatorname{rank}[D; C] = 23$. For the row $\ell^\top$ with entries $\ell_g = 1$, $\ell_A = m$ and all others zero, $\operatorname{rank}[D; C; \ell^\top] = 23$. Hence ℓ lies in the row space of $[D; C]$ and annihilates all seven joint null directions; this is an exact statement for the cycle's quadratic-function span, independent of the two float64 shares above. On the published recipe, the transient and the two-initial-condition scheme every share is 0.

Witness. The jointly blind direction v along which A moves by one unit has eight nonzero entries: $+1$ on x_0x_2 in the first equation and on x_1x_2 in the second, $-m$ on x_0 in the first and on x_1 in the second, and in the third equation -1 on x_0^2 and on x_1^2, $\Delta c = m/(1 + m^2)$ on the constant and $\Delta L = m^2/(1 + m^2)$ on x_2. On the cycle, where $x_0^2 + x_1^2 = x_2 = m$, the first two equations change by $x_0x_2 - mx_0 = 0$ and $x_1x_2 - mx_1 = 0$ and the third by $\Delta c + m\,\Delta L - m = 0$; $(\Delta c, \Delta L)$ is the minimum-norm solution of that last equation, which is why it is close to, and not equal to, (m, m^2). The constraint ties the coupling to the slaving coefficient, so its residual is unchanged too. The two witnesses of Section 4.1 are the true model $\pm\frac{1}{2}v$. Every quantity below is computed in raw physical units, with the library evaluated on the sampled states and no column normalization. The prediction change is $e_{\mathrm{pred}} = \max_{n,k} |\dot{X}_{\pm,nk} - \dot{X}_{\mathrm{true},nk}| / \operatorname{rms}(\dot{X}_{\mathrm{true}})$, where $\dot{X}_\pm$ and $\dot{X}_{\mathrm{true}}$ are the derivatives the two coefficient vectors predict at the n-th sampled state for equation k, and the rms is over all samples and equations; the constraint residual is $\|C\alpha\|_2$ with C the constraint matrix in the same units; and the off-sample difference is $\|X_{\mathrm{w}} - X_{\mathrm{true}}\|_F / \|X_{\mathrm{true}}\|_F$ for state trajectories integrated from the published initial condition on $0 \le t < 50$ with step 0.01 and the example's integrator settings. The closed form was checked against the stored witness after the freeze and is labelled exploratory in the protocol's log. "Removed" for another scheme means no jointly blind direction at the operating cut, inside a plateau of the scan, and $\rho = \lambda_{\min}(Q^\top (G_{\mathrm{d}} + \lambda_{\mathrm{ref}} G_{\mathrm{p}}) Q)/\|G_{\mathrm{d}}\|_2 > 10^{-8}$ with Q the witness subspace in that scheme's coordinates: ρ is -3×10^{-17} with four times the steady-state data and 0.077, 0.080 and 0.071 for the transient, the two initial conditions and the published recipe.

The late window. The last 5000 samples of the published trajectory, what a user who discards the transient would fit, are near the cycle and not on it, and the audit says so: the jointly blind count is 7 at cuts down to 10^{-10}, 1 at 10^{-11} and 10^{-12} and 0 below, with no plateau at the operating cut, while the noise-calibrated count is 15 as on the cycle. The explicit witness does not meet the registered thresholds there (prediction change 3×10^{-6}), so no strict claim is made; this is the third tier of Section 3.2.

The optimizer's reading. Run with the example's hyperparameters, TrappingSR3 on the steady-state data converges to $\dot{x}_0 = -0.9999x_1 - 0.0100x_1x_2$, $\dot{x}_1 = 0.9999x_0 + 0.0100x_0x_2$, $\dot{x}_2 = 0.0001x_0x_1$, and on the published recipe to the published model to four decimals (there it stops at its iteration limit, as recorded). The first is a legitimate representative of the solution set: its derivatives match the exact ones on the steady-state samples to a relative ℓ_2 error of 1.7×10^{-5} (exploratory, as above), and nothing in that fit reveals that three of its mechanisms are choices of the solver.

Exploratory matched-sample-size refit. Following the audit, we trained the same TrappingSR3 call separately on the first 5000 raw states of the steady-state and published transient trajectories. Finite differences drop two endpoints, leaving 4998 derivative rows in each fit. Both use the same degree-2 library, energy-preserving constraint, and unchanged settings ($\mathrm{reg_weight_lam} = 0$, $\eta = 10^5$, $\gamma = -1$, $\mathrm{max_iter} = 5000$); no tuning uses held-out trajectories. Support means $|\hat{\alpha}_i| > 0.005$ in raw units, fixed before comparing fits. We integrate each returned model over $0 \le t < 25$ with step

0.01 from three initial conditions absent from both training sets: $(0.01, 0.01, 0)$, $(-0.20, 0.15, 0.03)$, and $(0.15, -0.12, 0.02)$. The steady fit's relative trajectory errors are 13.7%, 43.2%, and 25.6%; the transient fit's are 0.0260%, 0.0208%, and 0.0224%. Table 3 reports their mean and coefficient errors. The transient fit returns an estimate at the iteration limit; convergence remains unverified. This post-registration comparison is listed in the protocol's deviation log and leaves the registered audit outcomes unchanged.

B.2.3 Two further registered studies on public example suites

The case above was chosen by us. To see what the audit reports when the cases are chosen by a rule, two further protocols were frozen, each in a commit of its own, before any of their quantities was computed: one over the official PySINDy examples that impose an equality constraint, one over the official DeePyMoD examples with the default library. Both fix the selection rule, the mapping to $(A_{\mathrm{d}}, A_{\mathrm{p}})$, the thresholds, the wording of null outcomes and what may not be changed afterwards; both ship with the code, with deviation logs that record the launches that stopped before any audit number existed. Most of what they return is null, and we report it as registered.

Table 8: The ten PySINDy cases selected by rule. n: number of coefficients; "joint": jointly blind; counts at the operating cut. Tags: "known before", the result was known when the protocol was frozen; "data floor", below the registered data floor, no claim. The three registered initial-condition variants of the MHD case agree.

case	system, source	n	rank A_{p}	data-blind	closed	joint	tag
C1a, C1b	Lorenz, feature overview	30	2	0	0	0	
C2	Lorenz, lecture notebook	27	10	0	0	0	
C3	mean-field, trapping	30	10	0	0	0	known before
C4	oscillator, trapping	30	10	0	0	0	
C5a, C5b	Lorenz, trapping	30	10	0	0	0	
C6	triad MHD, trapping	168	56	24	24	0	
C7	plasma, 7 modes	245	28	0	0	0	data floor
C8	plasma, 16 modes	2432	952	0	0	0	data floor

PySINDy. The protocol pins PySINDy 2.1.0 and the external example repositories its documentation fetches, 27 notebooks. Its rule keeps a fit when the optimizer imposes a hard linear equality constraint on the coefficients, a stored output prints the fitted coefficients, the states come from the pinned code or files alone, and the library consists of monomials; 13 notebooks name such an optimizer, and the rule yields ten cases (Table 8). The data block is the library on the published recipe's rows, noise-free, and the physics block is the constraint matrix as the optimizer applies it. The mean-field recipe (C3) was already known when the protocol was frozen; it is included in the descriptive count but not counted as new evidence. Seven PySINDy cases are null: on the Lorenz, mean-field and oscillator recipes the published trajectory determines every library coefficient at every one of the nine cuts, so the equality constraints are not needed for identifiability. The two plasma cases read a file stored to six decimals; the registered data floor, $\varepsilon^2_{\mathrm{data}} \approx 2.5 \times 10^{-9}$ and 2.7×10^{-9}, lies above the operating cut, so their counts (0 in both) support no claim. The inviscid triad MHD example is the one non-null case, and its outcome was registered in advance from a hand analysis: the vector field has four exact quadratic invariants, each gives one exact relation among the 28 library terms in each of the six equations, and the energy constraint sees all of them. The audit returns 24 data-blind directions among the 168 coefficients, all 24 closed by the constraint, none jointly blind, constant over all nine cuts; the 24 registered relation vectors span the audited subspace (largest principal-angle sine 8×10^{-8}), a symbolic check confirms the four invariants, and an exact rational computation gives the constraint rank 24 on that subspace. Three registered initial-condition offsets agree on every number. Moving along a witness changes the fit to the exact derivatives by at most 4×10^{-19} while coefficients of pure squares and of the products $x_i x_{i+3}$ change by order one; the published model prints two members of that set, $-0.001\, x_2 x_5$ and $0.001\, x_1^2$, and no mechanism coefficient is affected. Pooling two trajectories of the notebook's own generator leaves 18 data-blind directions, and pooling five leaves none (ρ between 1.4×10^{-7} and 2.3×10^{-4} over the three offsets, against the registered 10^{-8}). The condition number of the library matrix, about 10^{11}, signals that something is wrong; it does not say that there are four directions per equation, which terms they involve, or that the constraint already in the pipeline closes every one of them.

DeePyMoD. The protocol pins DeePyMoD 2.2.0. Its rule keeps the examples that use the one-dimensional default library with one network output on fields produced by the package or stored in it: six examples, three designs, each with the twelve-term library $u^p \partial_x^d u$, $p \le 2$, $d \le 3$. No example imposes a linear condition on the coefficients, so A_p is empty and the audit reduces to the data block (Algorithm 1); the design audited is the noise-free field the network is trained to reproduce. The result is null: none of the three published designs has a blind direction at any of the nine cuts, nor on any of 20 registered subsamples of the training size, and every published reading is determined. Only the Burgers design built from the package's generator supports that statement at the operating cut (normalized condition number 54; 40-digit null space empty); the two stored fields need numerical derivatives, and the registered derivative floor lies above the cut for both, slightly for the stored Burgers field (smallest eigenvalue 3.3×10^{-3}, three orders of magnitude above its floor) and grossly for the stored KdV field, for which we claim nothing. A second arm was designed, not selected, and registered as the confirmation of a known fact (Rudy et al., 2017): a single KdV soliton also solves one-way advection, so the library cannot identify KdV from it, and two solitons of different speeds can. The protocol fixed the fields (the package's own one- and two-soliton generators on the stored example's grid), a second parameter set, and the prediction: exactly three blind directions on one soliton, spanned by three named relations that follow from $f'' = cf - 3f^2$, and none on two. Every numerical count agrees on both parameter sets: three blind directions on each single-soliton field, constant over the nine cuts, a 40-digit null space of dimension three that coincides with the span of the three relations, and none on the two-soliton fields. One pass of the pipeline's own selection and refit returns advection, $u_t = -c\,u_x$, on the single solitons and $u_t = -6uu_x - u_{xxx}$ on two; the two readings fit the single-soliton design equally (relative residual below 10^{-15}), differ by c in the coefficient of u_x, and agree on every identifiable functional, for instance $\xi_{u_x} + c\,\xi_{u_{xxx}} = -c$. The two-soliton design removes the single-soliton witnesses (ρ between 9×10^{-6} and 10^{-3}). The registered verdict of this arm is nevertheless negative: the protocol required every relation to be confirmed symbolically within 600 seconds on both parameter sets, and on one robustness field one relation was recognised with rational coefficients but no simplification route finished in time, so by the protocol's letter the arm is "parameter-dependent" and neither prediction is claimed as a registered result. A check added after the freeze, and marked so, shows the unconfirmed relation to be an exact rational combination of two of the three registered ones.

What the two studies show. Rule-selected published examples are mostly well excited: eight of the nine auditable case-designs listed in Section 4.1 have no blind direction (seven of eight after excluding the previously known C3), and there the audit's report is that the published readings are determined by the data. The one exception was predicted before it was computed, and there the constraint already in the pipeline does the closing; nothing in either suite is jointly blind. A jointly blind direction appeared only under the observation-design intervention of Section 4.1. These rule-selected cases support a descriptive result for the tested suites; population prevalence would require a representative sample.

B.3 Scope and limitations

B.3.1 Impossibility on a fixed design

Theorem 1 (impossibility on a fixed design)**.** *Fix the design (Φ fixed; noise law independent of the truth). For $v \in \ker G_\mathrm{d}$, $t \neq 0$, the observations have identical distribution under $\alpha_A = \alpha_\mathrm{vis}$ and $\alpha_B = \alpha_\mathrm{vis} + tv$, at every noise level and for any number of repeated observations on that design. Consequently any measurable binary rule $g(u) \in \{0, 1\}$ meant to tell the two settings apart satisfies $P_A(g{=}1) = P_B(g{=}1)$; if the correct answers differ, g is right with probability exactly $\frac{1}{2}$ under the uniform prior on $\{A, B\}$.*

Here impossibility means statistical indistinguishability on a fixed design: two parameter values induce the same law of the observations, so no estimator, however expensive, separates them, and whether to trust an equation-implied value on such a direction cannot be settled from the data. The hypothesis can hold exactly: on the aliased heat probe the data-blind dependence is a trigonometric identity at the sample points, and the relative difference of 2.0×10^{-14} between the two noiseless data vectors is float64 rounding error. The scope is a fixed design: re-designing the sampling set can break $\ker G_\mathrm{d}$. The theorem applies to the fixed observation design.

Where the audit adds value. For a small, hand-curated library under a trusted operator, the audit offers a cheap confirmation. Its value is concentrated where libraries are overcomplete, automatically generated, or aliased against the sampling design, and where the coefficients themselves are the deliverable.

Nonlinear operators. Fréchet linearization $\mathcal{N}'[u_0]$ reproduces R on linear operators to below 10^{-8} relative (the correctness check) and supplies the Gauss–Newton curvature at u_0 for nonlinear ones. The resulting Jacobian null directions describe local first-order sensitivity; establishing a curve of exact solutions or a second minimizer requires further analysis. The corpus linearizes at the true solution. Application to noisy data would linearize at a fitted solution and assess error away from that anchor.

Model-class scope. The audit needs two quadratic forms in one coefficient vector. In PDE-FIND-style regression, $u_t = \Theta(u)\xi$, a linear constraint on ξ supplies the second block (Appendix B.2.1, Section 4.1); without it, the analysis reduces to the conditioning of Θ. In PINN-SR-style hybrids the residual acts on the network output, giving both network and library coefficients residual gradients (Chen et al., 2021). The zero-gradient result of Section 4.2 applies to the fixed dictionary parameterization. The linear-algebra criterion also covers $N_\mathrm{d} < K$ designs; experiments on that regime and a noise-calibrated choice of τ remain open.

C Extended related work

This appendix details how the coefficient audit relates to existing uniqueness and identifiability results.

The joint null space is classical. For a quadratic objective with a data block and a second block, the set of coefficient changes that leave both unchanged is $\ker A_\mathrm{d} \cap \ker A_\mathrm{p} = \ker[A_\mathrm{d}; A_\mathrm{p}]$; the minimizer set is one solution plus that space, and a linear reading is unique on it iff it is orthogonal to that space. This is the uniqueness condition of general-form Tikhonov regularization, written as a stacked least-squares problem in Hansen (2007) (§2.7.1) after Van Loan (1976) and Hansen (1989); the cosupport condition of cosparse recovery (Nam et al., 2013), where the stack of measurement and analysis rows is the uniqueness condition for a signal under one known operator; and, in physics-constrained system identification, the full-rank condition on the stacked regression and equality-constraint matrices under which Loiseau & Brunton (2018) obtain a unique constrained solution (their §2.2; a hard constraint and a finite-weight penalty are different estimators with the same obstruction). Coherence and spark (Donoho & Elad, 2003), ERC (Tropp, 2004) and D-RIP (Candès et al., 2011) characterize a single matrix or its composition. Building on the joint-kernel criterion, this paper delivers the outputs listed in Section 1: the direction-by-direction report with witnesses, the separation of functional identities from sampling and numerical degeneracy by tolerance scan, 40-digit spectra and symbolic confirmation, the reading test applied to published coefficients, and the registered cases.

Identifiability of penalized estimators. Scheipl & Greven (2016) prove the kernel-overlap criterion as an *iff* for penalized function-on-function regression (their Props. 3.3/3.4), ship a pre-fit diagnostic that measures the overlap of the two kernels by subspace angles and warns the user, note that the penalty then decides which representation is reported, observe that validation error cannot reveal the problem, and discuss penalties derived from partial differential operators. Our analysis predicts overlap from the solution space of an arbitrary linear operator (Lemma 1) and reports its consequences for specific coefficient readings. Recent identifiability theory for equation learning is complementary: solution-set metrics and sample complexity (Pan & Bölcskei, 2026), differential-algebraic functional identifiability (Loman et al., 2026), and the sparse-ODE phase transition of Casolo et al. (2026), who show that generic ℓ_1/ℓ_2 regularization does not resolve unidentifiability from the observations. We study a penalty that shrinks toward an imposed operator's solution set. Conditional on that operator, it constrains a data-blind direction modulo V_0 exactly when the direction lies outside $\ker G_\mathrm{p}$ (Proposition 2). The observations-alone identifiability result of Casolo et al. (2026) concerns a different constraint class.

Symbolic recovery and its identifiability. Scholl et al. (2026) ask whether the differential equation itself is unique given the solution; their S-FRanCo procedure builds a candidate feature library, polynomial features included, and reads the singular values of its evaluation matrix, and Hauger et al. (2024) extend the criterion to noisy data and estimated derivatives with SVD and Weyl-type thresholds. These works recover an unknown equation from an evaluated feature library. We audit the joint information supplied by a fixed representation and an imposed operator, using symbolic derivatives of known atoms. Morina et al. (2026) combine a data term with a PDE-residual penalty in an all-at-once objective and study reconstruction and the conditions under which the true law is recovered. Our analysis concerns the joint kernel of two fixed maps. Katende (2025) (v2) proves an aliasing theorem on the null space of the operator-evaluation design and adds conditioning checks before screening, multi-experiment designs and constraint manifolds. Our audit evaluates the data and imposed-constraint blocks together.

SINDy and constrained regression. Brunton et al. (2016) show on the cylinder-wake mean-field model that transients off the slow manifold are needed to recover the corresponding terms (their SI §4.3.2); Rudy et al. (2017) build $\kappa(\Theta)$ into their selection criterion and show that a single KdV soliton cannot be told from linear advection while trajectories of different amplitude can; Delahunt & Kutz (2022) cull dependent candidates by a leave-one-out R^2 and, for evaluation, map a discovered model to the equivalent representation closest to the truth; Loiseau & Brunton (2018) and Trapping SINDy (Kaptanoglu et al., 2021) impose physical equality constraints on the coefficients, and the software consolidates all of this (Kaptanoglu et al., 2022). Section 4.1 runs on their example and reports the directions closed and left by its constraint, explicit witnesses, and the resulting coefficient readings. Other responses address statistical or algorithmic aspects: ensemble bagging (Fasel et al., 2022), weak forms (Messenger & Bortz, 2021), thresholded sparse regression (Both et al., 2021; Chen et al., 2021); the triple heuristic of Anonymous (2026) is an order-3 relative of the leave-one-out cull. The joint audit additionally checks each dependency against the residual block. Implicit-SINDy and SINDy-PI compute or bypass null spaces of derivative-augmented libraries (Mangan et al., 2016; Kaheman et al., 2020); equality saturation removes syntactic redundancy by comparing candidate expressions (de França et al., 2025). Our sampled Gram matrices detect linear dependencies on the observation design, including grid aliasing; nonlinear relations among atoms require a different check.

Attractor geometry and excitation. Gallo et al. (2026) read the identifiability ceiling of SINDy and symbolic regression off the smallest eigenvalue of the invariant-measure moment matrix, a single-block quantity computable before discovery. Shumaylov et al. (2025) characterize discoverability in classes of continuous or analytic vector fields through trajectory geometry and, in their §6, show how first integrals create ambiguity and under which conditions added physical knowledge restores uniqueness. Their result concerns function classes richer than a finite known library. Our cases give a finite-dictionary, fixed-design decomposition of the attractor's blind directions into those a stated conservation constraint closes and those that remain, with the readings that follow.

Representations that encode the physics exactly. When every atom satisfies the imposed equation, the residual term is identically zero on the representation space. Trefftz methods build such spaces on purpose (Hiptmair et al., 2016), and linearly constrained Gaussian processes transform a representation so that a linear constraint holds by construction (Jidling et al., 2017). A zero residual block is then a design choice, and what remains to check is whether the observations separate the combinations of those solutions, which is the data block alone; the audit reports the zero block as such (Section 4.2).

Physics-informed learning theory and failure modes. Doumèche et al. (2024; 2025) analyse a data term plus a PDE-residual term as a kernel method; their positive-weight Sobolev-norm regularizer makes the function estimate unique by construction. They study convergence of that estimate; we study coefficients of a redundant representation. Koshizuka & Sato (2025) relate generalization to the dimension of the affine variety cut out by the physics constraints on a linear hypothesis class, nonlinear equation constraints included; for a linear constraint that dimension is the dimension of the operator kernel, the special case our second block describes. Otchere & Shields (2026) promote a Fisher-based effective dimensionality of network-parameter directions as an a-priori diagnostic. Our output-invariant directions live in coefficient space. The well-studied PINN failures are opti-

mization pathologies (Krishnapriyan et al., 2021; Wang et al., 2021; 2022); that residual losses alone constrain only up to the operator kernel is folklore there (Leiteritz & Pflüger, 2021), and misspecification has been corrected during training (Zou et al., 2024). Our audit finds jointly blind coefficient directions before training. Position-paper critiques document the symptom, near-identical residuals from distinct parameters (Naser, 2026); our two-map criterion identifies the responsible directions. On the words "gauge" and "dead" directions (Hashimoto et al., 2024; Shirodkar, 2026), see Appendix B.1.4.

D CORPUS

Nineteen libraries in three pools. Twelve *design* libraries: six taken from the benchmark suite of Anonymous (2026), four constructed to instantiate distinct identity types (a Klein–Gordon family, a phase-shift pair, a polynomial Laplace family, an order-5 identity), and two purpose-built probes of the joint condition (an aliased heat pair and a wave probe). Five *held-out* libraries were designed after every constant was frozen, by LLM agents given only the class taxonomy and no access to the diagnostic, its thresholds or its outputs. Two *external-recipe* libraries were generated by PySINDy's PolynomialLibrary and FourierLibrary construction rules (Kaptanoglu et al., 2022), as an anchor against false alarms on healthy dictionaries. Derivatives are computed symbolically, never by finite differences. For linear operators the physics block contains $R_{im} = \mathcal{N}[\varphi_i](y_m)$; for nonlinear operators it contains $R_{im} = D\mathcal{N}[u_0](\varphi_i)(y_m)$ at the anchored solution. The solv column counts atoms whose analytic residual vanishes by `sympy` simplification; for nonlinear operators it uses the original residual, not the Jacobian. Table 9 audits every library.

Anchored truths. Wherever a ground-truth coefficient vector is needed (the noise-calibrated count and the tests), it is anchored at the first atom whose analytic residual vanishes, so that the truth satisfies the imposed equation exactly.

Table 9: Corpus audit at $\tau = 10^{-10}$ in unit-rms coordinates. *solv*: atoms with zero analytic residual; *blind*: numerical nullity of G_{d}; d_τ: joint near-zero count; $\Delta r_\tau = \mathrm{blind} - d_\tau$: numerical rank gain (checked by stacked SVD); *rank*: eigenvalues of G_{p} above $\tau\|G_{\mathrm{p}}\|_2$. Nonlinear rows use the local Jacobian Gram. Figure 2 uses a separate share cut $\tau_\nu = 10^{-8}$; its physics-only plus jointly seen bins need not equal this rank. † marks nonlinear operators.

library	K	solv	blind	d_τ	Δr_τ	rank	library	K	solv	blind	d_τ	Δr_τ	rank
wave1d_full	14	12	8	8	0	2	aliased_heat	6	4	1	0	1	2
wave1d_minimal	2	2	0	0	0	0	wave1d_joint_probe	7	5	0	0	0	2
heat1d	43	6	16	12	4	22	*held-out*						
allen_cahn†	19	2	0	0	0	19	aliased_advection	9	4	1	0	1	4
burgers†	11	1	0	0	0	8	airy_overcomplete	11	9	4	4	0	2
wave2d	9	9	0	0	0	0	euler_bernoulli_null	11	11	0	0	0	0
klein_gordon	14	12	8	8	0	2	damped_wave_matched_pair	8	5	0	0	0	3
heat_phase_shift	7	6	1	1	0	1	helmholtz2d_clean	10	4	0	0	0	6
laplace_polynomial	8	6	1	1	0	1	*external-recipe*						
wave2d_quadruple	8	8	1	1	0	0	poly_heat	15	2	0	0	0	10
							fourier_wave	16	8	0	0	0	8