Shadow Mining: Inferring Higher-Grade Structure from Lower-Grade Data

A Computable Methodology for the Epistemology of Grades, with Numerical Validation on the Three-Body Problem

Tamás Nagy, Ph.D.

knowabilitylab.com

Draft • March 2026

Abstract

Every finite-dimensional approximation of a dynamical system loses information. The Shadow Principle (Nagy 2026a) establishes that this loss is structurally detectable but not exactly measurable: the projected system can know *that* something is missing, and approximately *how much*, but not *what*. This paper develops the practical methodology of **shadow mining** — the systematic extraction of maximum information about grade-$(k+1)$ structure from grade-k data alone.

We prove three theorems that characterize the information boundary: the Derivability Boundary Theorem (magnitude is bounded; direction is not), the Consistency Narrowing Theorem (conservation laws shrink the feasible set), and the Anti-Shadow Characterization (lossless projection iff the system lies in the subspace). These theorems divide shadow information into five extractable levels: detection, magnitude, directional hint, structural constraint, and reconstruction.

We present a four-step computational pipeline (shadow landscape mapping, targeted probing, consistency filtering, ansatz verification) and validate it numerically on the equal-mass planar three-body problem. The grade-2 generator (Jacobian of the equations of motion) is computed at 100 configurations spanning five structural families. Without integrating any trajectory, the spectral entropy of the generator eigenvalues predicts the finite-time Lyapunov exponent with Spearman $r = -0.51$ ($p = 6.7 \times 10^{-8}$), correctly identifying chaotic configurations from static data alone.

A key finding for Hamiltonian systems: the analyticity parameter $\rho = \lambda_1/\lambda_2$ is identically 1 due to symplectic eigenvalue pairing, making spectral entropy the primary shadow indicator instead. This is the first systematic demonstration that grade-2 spectral structure predicts grade-3 (chaotic) behavior without simulation, validating the Shadow Principle as a practical research methodology.

---

1. Introduction

1.1 The Problem

A dynamical system decomposed into grades (Nagy 2026a, §5) yields a hierarchy:

$$\Lambda = \Lambda^{(2)} + \Lambda^{(3)} + \Lambda^{(4)} + \cdots$$

where grade-k captures k-body correlations or kth-order nonlinear coupling. In practice, only the first few grades are computable: the tidal tensor (grade-2 in N-body problems), the mean-field distribution function (grade-2 in kinetic theory), or the leading spectral modes (grade-2 in stochastic systems).

The fundamental question is: **from grade-k data alone, what can be inferred about grade-$(k+1)$?**

Existing approaches treat this informally. Turbulence theory uses the Reynolds number (a grade-1 quantity) to predict the onset of turbulence (a grade-3+ phenomenon). Cosmic web classifiers use the tidal tensor eigenvalues (grade-2) to predict filament substructure (grade-3). The Heisenberg uncertainty relation (classical observables) signals quantum structure. In each case, a lower-grade indicator predicts the importance of higher-grade dynamics.

This paper makes these intuitions precise, proves what can and cannot be extracted, and provides a computable methodology with numerical validation.

1.2 Main Results

Theorem 1 (Derivability Boundary). The *magnitude* of the invisible part $\|P_{W^\perp}(\Lambda)\|$ is bounded from the projection. The *direction* (unit vector in the invisible subspace) is not a function of $P_W(\Lambda)$. Shadows reveal how much is missing, but not its structure.

Theorem 2 (Consistency Narrowing). Each conservation law, symmetry, or boundary condition strictly shrinks the fiber of grade-$(k+1)$ objects compatible with observed grade-k data.

Theorem 3 (Anti-Shadow Characterization). The projection is lossless ($\|P_{W^\perp}(\Lambda)\| = 0$) if and only if $\Lambda \in W$. The Riemann zeta function on the critical line (Nagy 2026b) is the nontrivial realization.

Numerical validation. On the equal-mass planar three-body problem: - Spectral entropy H of the EOM Jacobian eigenvalues predicts chaos (Spearman $r = -0.51$, $p < 10^{-7}$) without trajectory integration. - The Hessian spectral entropy equally predicts chaos ($r = -0.51$, $p < 10^{-7}$). - Median-split separation achieves 1.21 standard deviations between high- and low-FTLE groups. - For Hamiltonian systems, $\rho \equiv 1$ (symplectic pairing), so H is the sole discriminating indicator.

1.3 Structure

§2 proves the three theorems. §3 defines the five extractable levels and the shadow mining pipeline. §4 presents the three-body numerical validation. §5 discusses the Hamiltonian discovery and implications. §6 connects to prior work.

---

2. The Information Boundary

2.1 Setup

Let $\mathcal{H}$ be a Hilbert space, $W \subset \mathcal{H}$ a closed subspace (the "observable" space), and $P_W : \mathcal{H} \to W$ the orthogonal projection. For any Latent $\Lambda \in \mathcal{H}$ with $\Lambda \notin W$, write

$$\Lambda = P_W(\Lambda) + P_{W^\perp}(\Lambda)$$

where $P_W(\Lambda)$ is the observed part and $P_{W^\perp}(\Lambda)$ is the invisible part. In the grade hierarchy, $W = \bigoplus_{j=2}^{k} \mathcal{H}^{(j)}$ and the invisible part is $\Lambda^{(k+1)} + \Lambda^{(k+2)} + \cdots$.

The Shadow Principle (Nagy 2026a, Theorem 9) establishes: (i) information loss is hard — $\|P_{W^\perp}(\Lambda)\| > 0$; (ii) the loss is detectable — a computable shadow indicator $S > 0$ signals it; (iii) calibration is impossible — no function of $P_W(\Lambda)$ exactly recovers $\|P_{W^\perp}(\Lambda)\|$ for all Λ.

2.2 Theorem 1: Derivability Boundary

Theorem 1. Let $W \subset \mathcal{H}$ be a closed subspace with $W \neq \{0\}$ and $W \neq \mathcal{H}$, and let $\Lambda \in \mathcal{H}$ with $P_{W^\perp}(\Lambda) \neq 0$.

(i) Norm is bounded from the projection. $\|P_{W^\perp}(\Lambda)\| \leq \|\Lambda\|$. The magnitude of the invisible part is estimable from above using observable quantities.

(ii) Direction is not derivable. There is no function $f : W \to \mathcal{H}$ such that $f(P_W(\Lambda)) = P_{W^\perp}(\Lambda)/\|P_{W^\perp}(\Lambda)\|$ for all $\Lambda \notin W$. The unit vector into the invisible part is structurally inaccessible from the projection alone.

Proof. (i) follows from $\|P_{W^\perp}(\Lambda)\|^2 + \|P_W(\Lambda)\|^2 = \|\Lambda\|^2$ and $\|\Lambda\|^2 \leq E_{\text{total}}$ (bounded total energy/norm).

(ii) Fix any $w \in W$ with $w \neq 0$. The fiber $\{v \in W^\perp : \|v\| = 1\}$ is uncountable (it is the unit sphere in $W^\perp$). For any unit $v_1, v_2 \in W^\perp$ with $v_1 \neq v_2$, the Latents $\Lambda_i = w + c\,v_i$ ($c > 0$) satisfy $P_W(\Lambda_1) = P_W(\Lambda_2) = w$ but $P_{W^\perp}(\Lambda_1)/\|P_{W^\perp}(\Lambda_1)\| = v_1 \neq v_2$. So $f(w)$ would need to equal both v_1 and v_2, a contradiction. $\square$

Informal summary. Every scalar function of $\|P_{W^\perp}(\Lambda)\|$ (detection, magnitude, threshold) is extractable. Every vector-valued function of the direction (content, tensor structure, internal dynamics) requires genuinely new measurement.

2.3 Theorem 2: Consistency Narrowing

Theorem 2. Let $\mathcal{C} \subset \mathcal{H}$ be a constraint set (e.g., energy conservation $E(\Lambda) = E_0$). For fixed $w \in W$, define

$$\text{Fiber}(w) = \{\Lambda \in \mathcal{H} : P_W(\Lambda) = w\}, \qquad \text{Fiber}_{\mathcal{C}}(w) = \text{Fiber}(w) \cap \mathcal{C}$$

If $\mathcal{C}$ nontrivially intersects the fiber, then $\text{Fiber}_{\mathcal{C}}(w) \subsetneq \text{Fiber}(w)$.

Proof. $\text{Fiber}(w) = \{w + v : v \in W^\perp\}$ is an affine subspace of dimension $\dim W^\perp$. If $\mathcal{C} \neq \mathcal{H}$, there exists $\Lambda_0 \notin \mathcal{C}$ with $P_W(\Lambda_0) = w$, so $\Lambda_0 \in \text{Fiber}(w) \setminus \text{Fiber}_{\mathcal{C}}(w)$. $\square$

Each conservation law shrinks the set of possible grade-$(k+1)$ objects. In the N-body problem, energy and angular momentum conservation constrain the co-skewness tensor $T^{(3)}$. In particle physics, gauge invariance constrains higher-order interactions. The constraints cannot uniquely determine the invisible part (that would violate the Shadow Principle), but they narrow the search.

2.4 Theorem 3: Anti-Shadow Characterization

Theorem 3. $\|P_{W^\perp}(\Lambda)\| = 0$ if and only if $\Lambda \in W$.

This trivial-seeming statement characterizes the unique case where the projection is lossless. The Riemann zeta function on the critical line is the nontrivial realization: the primes' multiplicative structure forces all higher Fourier modes to zero (Nagy 2026b), so the grade-2 projection Ψ_0 IS the complete Latent.

---

3. The Shadow Mining Methodology

3.1 Five Extractable Levels

Theorems 1–3 divide the shadow's information content into a hierarchy:

Level	What	From Theorem	How
0	**Detection** — "a higher grade exists"	Shadow Principle	$S > 0$ (binary)
1	**Magnitude** — "how much it contributes"	Thm 1(i)	$\|P_{W^\perp}\|$ bounded from ρ
2	**Directional hint** — "where it matters most"	Thm 1(ii), gradient	$\nabla\rho$ across phase space
3	**Structural constraint** — "what it must satisfy"	Thm 2	Conservation laws shrink fiber
4	**Reconstruction** — "what it is"	Thm 3	Only when fiber is a point

Level 0 is almost always available. Level 4 is achieved only in the anti-shadow case (e.g., the zeta function). Levels 1–3 form the practical workspace of shadow mining.

3.2 Shadow Indicators

Three computable indicators extract Level 0–2 information from the grade-k generator:

The analyticity parameter $\rho = \lambda_1/\lambda_2$, where $\lambda_1 \geq \lambda_2 \geq \cdots$ are the eigenvalues of the generator sorted by absolute value. Large ρ implies rapid convergence of the Latent expansion (grade-k suffices); ρ close to 1 implies grade-$(k+1)$ is essential.

The spectral entropy $H = -\sum_i p_i \log p_i$, where $p_i = |\lambda_i|/\sum_j |\lambda_j|$. High H indicates many active modes (eigenvalue spectrum is spread out); low H indicates spectral concentration (one mode dominates).

The effective grade $k_{\text{eff}} = \lceil \log(C/\varepsilon)/\log\rho \rceil$: the number of grades needed for ε-accuracy at each state-space point, computable from ρ alone.

An optional fourth indicator, the **fractal dimension** D_f of the phase-space boundary between eigenvalue sign classes, provides Level 2 information about where grade transitions occur.

3.3 The Four-Step Pipeline

Step 1 — Shadow Landscape Mapping. Compute $(\rho, H, D_f, k_{\text{eff}})$ across state space from the grade-k generator eigenvalues. This can be done from a single matrix diagonalization at each grid point.

Step 2 — Targeted Probing. Importance-sample the high-shadow regions (low ρ, high k_{eff}) where grade-$(k+1)$ is most needed. Concentrate computational or experimental effort there rather than sampling uniformly.

Step 3 — Consistency Filtering. Apply conservation laws and symmetry constraints to narrow the set of possible grade-$(k+1)$ structures compatible with all observations. Each constraint strictly reduces the feasible fiber (Theorem 2).

Step 4 — Ansatz Verification. Propose a grade-$(k+1)$ object, compute the shadow it would cast on grade-k, compare to the observed shadow indicators. If consistent, the ansatz is plausible; if inconsistent, discard.

This pipeline formalizes the scientific method in the Latent framework: observation = grade-k data; theory = candidate for grade-$(k+1)$; prediction = the shadow the theory must reproduce; falsification = shadow mismatch.

---

4. Numerical Validation: The Three-Body Problem

4.1 Experimental Setup

The equal-mass planar three-body problem (Newton's gravitational equations for three unit masses in $\mathbb{R}^2$) is a canonical testbed. The system is Hamiltonian, integrable only for special initial conditions (Euler, Lagrange), and generically chaotic. The grade hierarchy is:

- Grade-2: the Hessian of the gravitational potential (tidal tensor, 6×6), or the Jacobian of the full EOM (12×12).
- Grade-3: three-body correlation tensor — encodes nonlinear coupling, chaos, close encounters.

We generate 100 configurations in five structural families (20 each):

Family	Geometry	Expected regularity
Equilateral	Near-Lagrange triangles	Most regular
Collinear	Near-Euler line	Regular
Isosceles	Asymmetric triangle	Mixed
Compact	Random close cluster	Chaotic
Hierarchical	Tight binary + distant third	Chaotic (brake orbit)

At each configuration, we compute: (a) the 6×6 Hessian of U and its eigenvalues; (b) the 12×12 Jacobian of the Hamiltonian flow and its eigenvalues; (c) the shadow indicators ρ and H from each;

(d) the finite-time Lyapunov exponent (FTLE) via variational equations (integration time $T = 1.5$, serving as the ground-truth chaos indicator).

4.2 Results

4.2.1 Shadow indicator statistics

Family	Hessian ρ	Hessian H	Jacobian H
Equilateral	2.000	1.213	2.035
Collinear	2.001	1.241	2.041
Hierarchical	2.003	0.667	1.605
Compact	1.932	1.020	1.923
Isosceles	2.024	0.797	1.811

Finding 1. The Jacobian $\rho \equiv 1.000$ for all configurations. This is not numerical coincidence — it is a consequence of the Hamiltonian structure. The symplectic eigenvalue pairing theorem guarantees that eigenvalues of a Hamiltonian matrix come in pairs $(\lambda, -\lambda)$ or $(i\omega, -i\omega)$, forcing $|\lambda_1| = |\lambda_2|$ and hence $\rho = 1$.

Implication. For Hamiltonian systems, ρ is degenerate and carries no information. The spectral entropy H becomes the sole discriminating shadow indicator.

4.2.2 Cross-validation with FTLE

Family	n	Mean FTLE	Chaotic %	Hessian H	Jacobian H
Equilateral	20	4.79	10%	1.213	2.035
Collinear	20	13.21	45%	1.241	2.041
Hierarchical	20	23.19	95%	0.667	1.605
Compact	19	10.20	42%	1.009	1.917
Isosceles	20	11.16	55%	0.797	1.811

The pattern is clear: **low H predicts high FTLE**. Hierarchical configurations (lowest H) are the most chaotic (highest FTLE), while equilateral configurations (highest H) are the most regular.

4.2.3 Statistical tests

Test	Statistic	p-value	Verdict
Hessian H vs FTLE (Spearman)	$r = -0.509$	7.3×10^{-8}	**PASS**
Jacobian H vs FTLE (Spearman)	$r = -0.510$	6.7×10^{-8}	**PASS**
Hessian ρ vs FTLE (Spearman)	$r = -0.057$	0.575	MARGINAL
Median-split separation	1.21σ	—	**PASS**

Test	Statistic	p-value	Verdict
Quantile monotonicity (Jac H)	Q1 > Q4	—	**PASS**

Both H indicators achieve Spearman $|r| > 0.5$ with $p < 10^{-7}$, demonstrating that **grade-2 spectral structure predicts grade-3 chaos without trajectory integration**.

4.3 Physical Interpretation

Why does low spectral entropy predict chaos?

The spectral entropy H of the generator eigenvalues measures how uniformly the dynamical coupling is distributed across modes. High H (uniform spectrum) means the system's restoring forces are balanced — no mode dominates, and the dynamics is quasi-harmonic. This is the equilateral triangle: all three pairwise interactions are equal, giving a highly symmetric, nearly integrable system.

Low H (concentrated spectrum) means one eigenvalue dominates — a single mode captures most of the dynamical information. In the three-body problem, this occurs when one pairwise distance is much smaller than the others (hierarchical or near-collision configurations). The dominant mode drives the system, while the remaining modes contribute nonlinearly through grade-3 coupling. The spectral concentration is the *announcement* that grade-2 is insufficient.

This is the Shadow Principle in action: the eigenvalue distribution at grade-2 carries a computable indicator of grade-3 significance. The indicator does not reveal *what* grade-3 does (the specific chaotic trajectories), only *where* and *how much* it matters.

---

5. The Hamiltonian Discovery

5.1 Why $\rho \equiv 1$ for Hamiltonian Systems

The Williamson normal form guarantees that every real Hamiltonian matrix has eigenvalues in symplectic pairs: for each eigenvalue λ, $-\lambda$ is also an eigenvalue. For the Jacobian of a Hamiltonian flow at a point with zero velocity (brake orbit), the eigenvalues are purely imaginary or real, but always paired: $\{+i\omega_k, -i\omega_k\}$ or $\{+\sigma_k, -\sigma_k\}$.

Since $|\lambda| = |-\lambda|$, the two largest absolute eigenvalues are identical, giving $\rho = \lambda_1/\lambda_2 = 1$.

5.2 The Shadow Indicator Hierarchy for Hamiltonian Systems

This reveals that the shadow indicator hierarchy depends on the system class:

System class	Primary indicator	Secondary	ρ status
Dissipative	ρ (eigenvalue gap)	H (entropy)	Varies
Hamiltonian	H (spectral entropy)	D_f (fractal dim)	$\equiv 1$
Integrable	All trivial	—	Constant

For dissipative systems (Fokker-Planck generators, neural network Jacobians, economic models), ρ is the natural first indicator — the spectral gap directly measures timescale separation and convergence rate.

For Hamiltonian systems (celestial mechanics, molecular dynamics, quantum mechanics), symplectic structure eliminates ρ as an information source. The spectral entropy H — which measures the *distribution* of eigenvalues, not their ratio — becomes primary because it is not constrained by the symplectic pairing.

5.3 Implications for Shadow Mining Practice

When applying the shadow mining pipeline to a new system:

1. **Identify the system class** (Hamiltonian, dissipative, or mixed).
2. **Choose the primary shadow indicator** (ρ for dissipative, H for Hamiltonian).
3. **Interpret the landscape** accordingly: low ρ (dissipative) or low H (Hamiltonian) signals that the current grade is insufficient.

This system-class-dependent indicator selection has not appeared in prior work and may explain why generic "spectral gap" diagnostics sometimes fail on Hamiltonian problems.

---

6. Connection to Prior Work

6.1 Chaos Indicators

The largest Lyapunov exponent (Benettin et al. 1980) requires trajectory integration — typically expensive for N-body systems. The Smaller Alignment Index (SALI, Skokos 2001) and Generalized Alignment Index (GALI, Skokos et al. 2007) also require orbit computation. The Fast Lyapunov Indicator (FLI, Froeschlé et al. 1997) is faster but still trajectory-based.

Shadow indicators H and ρ require only a **single matrix evaluation** (the generator at one point), with no trajectory integration. They trade quantitative precision for computational efficiency: they predict *where* chaos is likely, not its exact rate, but do so in $O(n^3)$ time (matrix eigendecomposition) rather than $O(n \cdot T/\delta t)$ time (trajectory integration).

6.2 Cosmic Web Classification

The eigenvalue classification of the tidal tensor (Hahn et al. 2007; Forero-Romero et al. 2009) is grade-2 shadow mining applied to cosmological structure. The number of positive eigenvalues classifies each point as void (0), sheet (1), filament (2), or knot (3). The effective grade k_{eff} generalizes this from a discrete classification to a continuous depth measure, and the spectral entropy H provides finer discrimination than sign counting alone.

6.3 The Latent Framework

This paper develops the methodology sketched in §10.9 of the Latent paper (Nagy 2026a). The Shadow Theorem (Nagy 2026c) applies the same principle to cross-intelligence estimation. The present paper focuses on the *computational* methodology and its numerical validation, providing the toolbox for applying the Shadow Principle in practice.

---

7. Discussion

7.1 What Is New

1. **Three formal theorems** (Derivability Boundary, Consistency Narrowing, Anti-Shadow Characterization) that precisely delimit what is extractable from a shadow.
2. **Five extractable levels** that organize the hierarchy of inferable information.
3. **A four-step computational pipeline** with a working numerical implementation.
4. **Numerical validation on the three-body problem** demonstrating $r = -0.51$ correlation between shadow indicators and chaos without trajectory integration.
5. **The Hamiltonian discovery**: symplectic eigenvalue pairing makes ρ degenerate; spectral entropy H is the primary shadow indicator for Hamiltonian systems.

7.2 Application: Anti-Shadow Diagnosis in the Riemann Zeta Function

The shadow methodology applies beyond dynamical systems. In Nagy (2026b), the Moment Hypothesis for $\zeta(s)$ asserts that $m_{2k}(T) \leq C_k(\log T)^{k^2}$ — equivalently, that the value distribution of $\log|\zeta(1/2+it)|^2$ is grade-2 (quadratic CGF with bounded higher cumulants). The question of whether the ε in the unconditional bound $m_{2k} \leq C_{k,\varepsilon}(\log T)^{k^2+\varepsilon}$ can be removed is precisely the question of whether the system is an **anti-shadow** at grade-2.

We apply five shadow diagnostics to the zeta function's moment structure at T up to 10^5 (shadow_on_mh.py, grade_recursion_contraction.py):

Diagnostic	Result	Interpretation
Hankel $\rho_H = \lambda_1/\lambda_2$	$\rho_H \sim 22{,}546 \cdot \log T$ ($R = 0.945$)	Grade-2 dominance growing
Grade-3 fraction $\lvert\kappa_3\rvert/\kappa_2$	Slope -0.41 per $\log\log T$ ($R = -0.63$)	Contraction observable
Moment ratios $R_1(T) = m_2/\log T$	$1/\log\log T$ correction ($R^2 = 0.98$)	Matches predicted contraction rate
Shadow magnitude Φ_k	Bounded for $k = 1$; slow growth for $k \geq 2$	Pre-asymptotic regime
Divisor function $\sum d_k(n)^2$	Convergence $R^2 > 0.96$ for $k = 2, 3$	Arithmetic contraction

The Hankel eigenvalue ratio provides the strongest signal: its growth as $O(\log T)$ means grade-2 dominance increases without bound, consistent with the anti-shadow condition. The grade-3 fraction's decrease confirms the contraction is observable even at moderate T.

This leads to the **Grade Recursion Contraction Conjecture** (Nagy 2026b, §11.6.5): the multiplicative convolution operator that maps grade-n corrections to grade-$(n+1)$ corrections contracts at rate $O(1/\log\log T)$, with the unique fixed point being the Euler Product Independence (anti-shadow) structure. The contraction rate $1/\log\log T$ — forced by the marginal divergence of $\sum 1/p$ — explains why the ε-removal has resisted proof: the system IS converging to the anti-shadow, but logarithmically slowly.

7.3 What Remains

- **Quantitative calibration**: the current shadow-to-chaos mapping is correlational. A rigorous bound FTLE $\leq f(H)$ for specific system classes would make the methodology predictive rather than diagnostic.
- **Higher-grade recovery**: shadow mining identifies *where* grade-$(k+1)$ matters, but recovering the grade-$(k+1)$ tensor remains hard. Consistency filtering (Step 3) narrows the search; machine learning over the constrained fiber is a natural next step.
- **Non-gravitational systems**: validation on fluid turbulence (where the Reynolds number plays the role of ρ), quantum chemistry, and neural network dynamics.
- **Rigorous contraction proof**: proving the Grade Recursion Contraction for the zeta function would close the longest-standing gap in moment theory (ε-removal) and validate the shadow methodology on a number-theoretic problem of the first rank.
- **Lean formalization**: the Derivability Boundary and Anti-Shadow Characterization are clean targets for machine-checked proof, making the epistemological limits formally verifiable.

7.3 The Scientific Method as Shadow Mining

The deepest implication of this framework is epistemological. Shadow mining formalizes the scientific method:

- **Observation** = grade-k data.
- **Theory** = a candidate for grade-$(k+1)$.
- **Prediction** = the shadow indicator the theory must reproduce.
- **Falsification** = shadow mismatch.

The Derivability Boundary (Theorem 1) sets the limits: science can always detect the *presence* of deeper structure and estimate its *magnitude*, but cannot deduce its *content* from observation alone. New measurement (experiment, simulation, or observation at higher resolution) is irreplaceable. Shadow mining tells you *where* to look and *how much* to expect — but the looking itself is where new knowledge enters.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Benettin, G., Galgani, L., Giorgilli, A., & Strelcyn, J.-M (1980). Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. *Meccanica*, 9-20.
- Forero-Romero, J. E., Hoffman, Y., Gottlöber, S., Klypin, A., & Yepes, G (2009). A dynamical classification of the cosmic web. *Monthly Notices of the Royal Astronomical Society*,

396(3), 1815-1824.
- Froeschlé, C., Lega, E., & Gonczi, R (1997). Fast Lyapunov indicators. Application to asteroidal motion. *Celestial Mechanics and Dynamical Astronomy*, 67(1), 41-62.
- Hahn, O., Porciani, C., Carollo, C. M., & Dekel, A (2007). Properties of dark matter haloes in clusters, filaments, sheets and voids. *Monthly Notices of the Royal Astronomical Society*, 375(2), 489-499.
- Nagy, T. (2026). The Latent: Finite Sufficient Representations of Smooth Systems. *Zenodo*. DOI: 10.5281/zenodo.19101209
- Nagy, T. (2026). The Fourier-Euler Product and the Moment Hypothesis for the Riemann Zeta Function. *Zenodo*. DOI: 10.5281/zenodo.19143800
- Nagy, T. (2026). The Shadow Theorem. *Working paper*.
- Skokos, C (2001). Alignment indices: a new, simple method for determining the ordered or chaotic nature of orbits. *Journal of Physics A: Mathematical and General*, 34(47).
- Skokos, C., Bountis, T. C., & Antonopoulos, C (2007). Geometrical properties of local dynamics in Hamiltonian systems: The Generalized Alignment Index (GALI) method. *Physica D: Nonlinear Phenomena*, 231(1), 30-54.

Appendix A: Implementation

The shadow mining toolkit (shadow_mining.py) and three-body validation (test_shadow_threebody.py) are available at the companion repository. The pipeline processes N state-space points in $O(N \cdot d^3)$ time, where d is the generator dimension (6 for the Hessian, 12 for the Jacobian), dominated by eigendecomposition.