TIMESCALE IDENTIFICATION FROM k-STEP TRANSITION OPERATORS: THE STRUCTURAL ENTROPY SPECTRUM

Anonymous authors
Paper under double-blind review

ABSTRACT

Agents and world models discretise experience into a transition graph at one fixed step, yet behaviour has structure at several timescales. Which step should the graph be built at, and which timescales are present? Multiscale tools (Markov stability, the map equation at Markov time, implied timescales of Markov state models) sweep a time parameter but attach no significance statement to the scales they report. We make the step count k an argument of a clustering objective on the k-step operator P^k and calibrate it against a matched null: chains whose stationary flow matrix has identical margins and dwell (self-loop) mass. Our theory uses structural entropy (SE), a two-level code-length objective: the *structural entropy spectrum* $\mathrm{SE}(k)$, its null $\mathrm{SE}_{\mathrm{null}}(k)$ and the calibrated spectrum $\Delta\mathrm{SE}(k) = \mathrm{SE}_{\mathrm{null}}(k) - \mathrm{SE}(k)$. We prove that the k-dependence factors through an exact identity, $\mathrm{SI}(k;\mathcal{C}) = \sum_a F_{aa}(k)\log_2(1/V_a)$ (k-step module persistence $\times$ module code length), prove that when metastable blocks have stationary volume $< \frac{1}{2}$ the metastable partition is the unique *global* optimizer of the idealized in-window operator, and sketch a Simon–Ando extension to nearly uncoupled chains on a window whose endpoints scale as $1/(1-\lambda)$ for the fast and slow eigenvalue clusters, so the optimal partition path coarsens through a planted hierarchy in order. On a 36-state three-level chain, the path recovers both planted levels in order at calibrated effect sizes up to 1.8 bits, although at this fixture's modest scale separation $\Delta\mathrm{SE}(k)$ shows no plateaus; a single-scale expander control is featureless; the spectrum is estimable from a single trajectory. On latent transition graphs from a pixel-based JEPA world model, the one-step graph's calibrated effect falls below the preregistered 0.1-bit floor, while the spectrum peaks reproducibly at $k^* = 11$–16 with 1.6–1.9 bits across five training seeds; a self-loop control shows that this structure survives removal of dwell but that the peak location is a dwell-inflated step count (two state changes). The calibration is not SE-specific: wrapped around the map equation it passes the planted-hierarchy, single-scale and structureless controls, and around Markov stability all but the expander control; both recover two balanced blocks, SE's boundary case, and the scale boundaries move with the objective. All decisive experiments were preregistered with pass/fail gates and are reported verbatim.

1 INTRODUCTION

Agents, world models and graph-learning pipelines routinely turn experience into a transition graph: states or learned latent prototypes become nodes, observed transitions become weighted edges, and clustering the graph yields state abstractions and exploration targets (Zeng et al., 2023; 2024). The graph is built at one step size, but behaviour has structure at several timescales. This is the premise of temporal abstraction (Sutton et al., 1999) and of hierarchical world models, which still fix the number of temporal levels and their placement by hand, typically at two (Hafner et al., 2022; Iiyama et al., 2026). Which step should the graph be built at, and which timescales does it contain? In the latent transition graphs of a pixel-based joint-embedding predictive (JEPA) world model (Section 4.6), 81–88% of stride-1 transitions never leave their prototype, and 54–68% at stride 5.

Multiscale structure in Markov chains has classic tools. Markov stability scores partitions of the time-t flow and sweeps t (Delvenne et al., 2010), the map equation does the same at Markov time t

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.

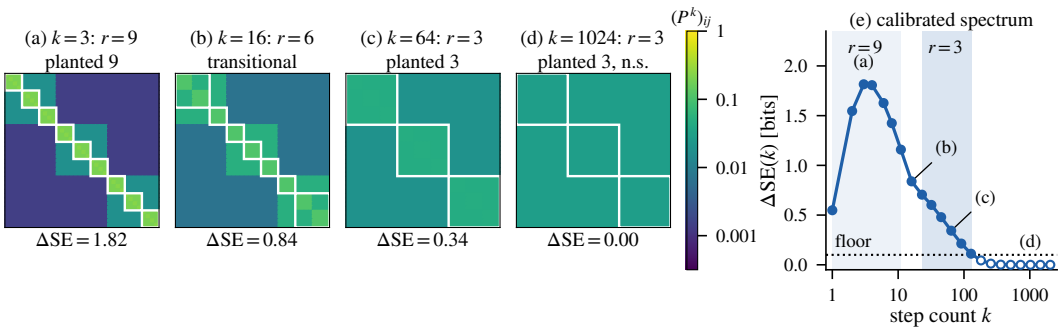

Figure 1: The SE spectrum on a chain with two planted levels (36 states; 9 clusters in 3 super-blocks; Section 4, G1). (a)–(d) Exact k-step operator P^k (states in planted order, log colour scale) with the optimal SE partition $\mathcal{C}^*(k)$ outlined in white. As k grows the optimum moves from the 9 clusters through a transitional 6-module partition to the 3 super-blocks. At k=1024 it still returns the super-blocks, but the calibrated effect $\Delta\mathrm{SE}$ against matched nulls is 0, so the partition is no longer significant. (e) $\Delta\mathrm{SE}(k)$, mean of 5 null seeds. Filled points pass the significance rule ($p \leq 0.05$ and $\Delta\mathrm{SE} \geq 0.1$ bits); shading marks the k at which the planted 9 and 3 modules are significant optima.

(Schaub et al., 2012), and Markov state models read timescales off plateaus of implied timescales and extract metastable sets with PCCA+ (Prinz et al., 2011; Deuflhard and Weber, 2005). These sweeps return a partition or a timescale at every t but do not test which exceed what margins and dwell alone produce, the test a self-loop-dominated graph needs. We make the step count an argument of the objective and calibrate it: a partition objective is optimised on the k-step operator P^k and compared with matched null chains that preserve the margins and the diagonal of the stationary flow matrix. The significant k and their optimal partitions identify the timescales (Figure 1).

We develop the theory for structural entropy (SE) (Li and Pan, 2016), the code length of a random walk under a partition-indexed two-level code, which already builds state abstractions in reinforcement learning (Zeng et al., 2023) but carries no timescale. SE admits the sharpest statements: its saving on P^k factors exactly into k-step module persistence times module code length (Markov stability weighted by code length), and its global optimizer provably locks onto the metastable partition of the idealized in-window operator. Wrapping the same calibration around Markov stability and the map equation separates what is general from what is specific to SE. Our contributions:

(1) A null-calibrated lag spectrum. The SE spectrum $\mathrm{SE}(k)$: optimally partitioned, stationary-weighted directed SE of P^k as a function of k, with a matched directed null (identical stationary distribution, one-level entropy H_1 and dwell mass), the calibrated spectrum $\Delta\mathrm{SE}(k)$, the partition path $\mathcal{C}^*(k)$, and scale boundaries where the path coarsens (Section 3).

(2) Theory. An exact persistence–code-length identity for directed two-level SE (Lemma 1); k-invariance of the node-level term (Lemma 2); a spectral-mixture representation for reversible chains placing the knees at $k \sim 1/(1-\lambda)$ per eigenvalue cluster (Proposition 1); and the main theorem: for nearly uncoupled chains with metastable block volumes $< \frac{1}{2}$, the metastable partition is the unique global optimizer throughout the Simon–Ando window (Theorem 1; the idealized case, Lemma 4, is proved in full, the perturbation step is sketched; the multi-level case is verified numerically).

(3) Preregistered experiments. On a 3-level chain (36 states; raw-chain implied timescales $t \approx 0.9/12/133$) the spectrum recovers the planted 9-partition and then the planted 3-partition, in order, at $\Delta\mathrm{SE}$ up to 1.8 bits against 20 matched nulls; an expander control is featureless; one trajectory suffices to estimate it; PCCA+ given the true module count matches it on the exact chain but recovers the fine level less reliably from short trajectories (lower in 4 of 5 seeds, one tie; Table 1); on a world model's latent rollouts a self-loop control shows structure beyond dwell, peaking at two state changes; a case study draws the partitions in state space (Figure 4).

(4) Beyond SE. Wrapped in the same null, the map equation passes every gate of Table 2 and Markov stability all but the expander control; both, unlike SE, recover two balanced blocks, and the scale boundary moves with the objective (Section 4.5, Appendix C).

2 RELATED WORK

Structural entropy. SE (Li and Pan, 2016) measures the code length of a random walk under a hierarchical encoding; optimizers and libraries exist (Anonymous, 2026), but the theory is stated for undirected graphs, attaches no null model, and—the gap here—no timescale. Our definition builds a matched null in and makes k explicit. Conditions under which low SE tracks external labels are studied in (When SE tracks labels, 2026); that work is orthogonal (static, undirected, single-scale).

Markov state models. The closest rival machinery. MSMs select a lag time τ at which implied timescales $t_i = -\tau/\ln \lambda_i(\tau)$ plateau (Prinz et al., 2011), and extract metastable partitions with PCCA+ (Deuflhard and Weber, 2005). Differences: (i) MSMs read timescales off eigenvalues and separately cluster eigenvectors; the SE spectrum returns the partition, the scale window and a *code-length effect size in bits* from one objective; (ii) $\Delta\mathrm{SE}(k)$ is calibrated against an explicit matched null, which MSM implied timescales lack; (iii) the partition path composes into a hierarchy with a timescale per level. We do not claim superiority at MSM's home task (molecular kinetics); Table 1 compares the two on our fixture. **Clustering across Markov time.** Indexing a clustering objective by time is not new (Delvenne et al., 2010; Lambiotte et al., 2014; Rosvall and Bergstrom, 2008; Schaub et al., 2012); robust scales are typically read from the consistency of partitions across restarts and neighbouring times, e.g. by PyGenStability (Arnaudon et al., 2024), not against a null; multiscale MSMs (Nitskansky et al., 2025) build MSM hierarchies by adaptive molecular sampling rather than testing one fixed graph's scales. We add SE-specific theory, the margin-and-dwell-matched null that turns a sweep into a significance statement, and a test of that calibration around both objectives (Section 4.5). **Spectral gap analysis and diffusion maps** (Coifman and Lafon, 2006) expose timescales via eigenvalues/eigenmaps but return neither discrete modules nor a null-calibrated effect size. **Nearly uncoupled chains.** Our theory window is the classical Simon–Ando/Courtois regime (Simon and Ando, 1961; Meyer, 1989); new here is that the *SE global optimizer* locks onto the metastable partition throughout that window, with an explicit volume condition and margin.

3 THE STRUCTURAL ENTROPY SPECTRUM

3.1 SETUP

We state the construction for SE; Section 4.5 applies it to other objectives. Let $W \in \mathbb{R}^{n\times n}_{\geq 0}$ be a directed transition-count or probability matrix (self-loops allowed and meaningful). Fix the teleport convention $P = (1-\theta)\,\mathrm{rownorm}(W) + \theta/n$ with $\theta = 0.01$, making P ergodic; let π be its stationary distribution and $H_1 = -\sum_i \pi_i \log_2 \pi_i$. The *stationary flow matrix* is $\Phi = \mathrm{diag}(\pi)P$; both its rows and its columns sum to π. For a partition $\mathcal{C}$ of the states with one-hot matrix $A \in \{0,1\}^{n\times r}$, define volumes $V = A^\top\pi$ and the k-step co-flow $F(k) = A^\top \mathrm{diag}(\pi)P^k A$, so $F_{aa}(k) = \Pr_\pi[X_0 \in a, X_k \in a]$. The directed, stationary-weighted two-level structural entropy of P^k under $\mathcal{C}$ extends the two-level SE of Li and Pan (2016) to directed walks:

$$H_2(k;\mathcal{C}) = \sum_a \Big[-\sum_{i\in a} \pi_i \log_2 \frac{\pi_i}{V_a} \;-\; g_a(k)\log_2 V_a \Big], \qquad g_a(k) = \sum_{i\in a}\pi_i \sum_{j\notin a} (P^k)_{ij}, \tag{1}$$

with $g_a(k)$ the stationary probability of leaving module a in k steps. Self-loop convention: self-loops shape π and the diagonal of P (dwell time) and never contribute to g_a; the null below preserves them.

Definition 1 (SE spectrum, null spectrum, partition path)**.**

$$\mathrm{SE}(k) = \min_{\mathcal{C}} H_2(k;\mathcal{C}), \qquad \mathrm{SI}^*(k) = H_1 - \mathrm{SE}(k), \qquad \mathcal{C}^*(k) = \arg\min_{\mathcal{C}} H_2(k;\mathcal{C}).$$

The matched null *draws Φ_{null} by margin-preserving randomization of the off-diagonal of Φ (row sums, column sums, and the whole diagonal preserved exactly), sets $P_{\mathrm{null}} = \mathrm{diag}(\pi)^{-1}\Phi_{\mathrm{null}}$, and computes $\mathrm{SE}_{\mathrm{null}}(k)$ with the identical optimizer. The calibrated spectrum is $\Delta\mathrm{SE}(k) = \overline{\mathrm{SE}_{\mathrm{null}}}(k) - \mathrm{SE}(k)$.* Scale boundaries *are the k at which $\mathcal{C}^*(k)$ coarsens, restricted to significant k (empirical one-sided $p \leq 0.05$ over ≥ 20 null draws and $\Delta\mathrm{SE}(k) \geq 0.1$ bits; the effect-size floor is required because an exact operator has no sampling noise). With 20 draws the smallest attainable p is $1/21 \approx 0.048$, so $p \leq 0.05$ means exceeding every null draw; no multiplicity correction is applied across the k-grid (a rerun with 500 draws and Holm correction leaves every gate verdict unchanged; Appendix D).*

Because row and column sums of Φ both equal π, the null chain has the *same* stationary distribution, the same H_1, the same per-state dwell mass, and the same one-step flow margins; $\Delta\mathrm{SE}(k)$ therefore isolates higher-order, multi-step flow structure. This is the directed, spectrum-level analogue of a degree-preserving null.

3.2 Exact structure of the spectrum

Lemma 1 (persistence–code-length identity)**.** *$H_2(k;\mathcal{C}) = H_1 + \sum_a F_{aa}(k)\log_2 V_a$, equivalently*

$$\mathrm{SI}(k;\mathcal{C}) \;=\; H_1 - H_2(k;\mathcal{C}) \;=\; \sum_a F_{aa}(k)\,\log_2\frac{1}{V_a} \;\geq\; 0.$$

(Proof in Appendix A.)

The interpretation is operational: the two-level code pays the module name only on entry; $\mathrm{SI}(k;\mathcal{C})$ is exactly the expected code length *saved* on k-step transitions that stay inside their module — (k-step module persistence) $\times$ (module code length).

Lemma 2 (the k-dependence lives entirely in the flow)**.** *π is stationary for P^k for every k, hence H_1 and all volumes V_a are k-invariant and $\mathrm{SE}(k)$ depends on k only through $F(k)$.*

Proof. $\pi P^k = \pi$ by induction. □

Lemma 3 (independence limit and the chance term)**.** *As $k\to\infty$, $P^k \to \mathbf{1}\pi^\top$, $F_{aa}(k)\to V_a^2$ and $\mathrm{SI}(k;\mathcal{C}) \to \mathrm{chance}(\mathcal{C}) = \sum_a V_a^2\log_2(1/V_a) > 0$ for any nontrivial $\mathcal{C}$; with equal volumes $V_a = 1/r$ this is $(\log_2 r)/r$, maximized at $r = 3$. The optimally partitioned SI^* of a structureless operator is therefore strictly positive.*

This is the connectivity artifact of raw SE in its purest form (an i.i.d. sampler still "saves" code length when it lands in the same module twice by chance), and it is removed by construction in $\Delta\mathrm{SE}(k)$: the null has the same volumes and the same limit.

Proposition 1 (spectral mixture; knees at $k\sim 1/(1-\lambda)$)**.** *Let P be reversible with eigenvalues $1=\lambda_1>\lambda_2\geq\cdots\geq\lambda_n$ and π-orthonormal eigenfunctions f_i. Then for every module a,*

$$F_{aa}(k) - V_a^2 = \sum_{i\geq 2}\lambda_i^k\, c_{i,a}, \qquad c_{i,a} = \Big(\sum_{j\in a}\pi_j f_i(j)\Big)^2 \;\geq 0,$$

so $\mathrm{SI}(k;\mathcal{C}) - \mathrm{chance}(\mathcal{C}) = \sum_{i\geq 2}\lambda_i^k\sum_a c_{i,a}\log_2(1/V_a)$ is a nonnegative mixture of exponentials λ_i^k. If the spectrum has clusters $\{\lambda_2..\lambda_r\}\gg\{\lambda_{r+1}..\}$, the mixture decays in steps, with knees at $k\approx t_i = -1/\ln\lambda_i$ per cluster: plateaus of the spectrum between consecutive relaxation times.

The proof expands $\mathbf{1}_a$ in the eigenbasis (Appendix A).

3.3 Main theorem: the metastable window

Idealize a chain with r metastable blocks $R=\{b_1..b_r\}$ at times inside the window by the operator $\Pi_R = \mathrm{blockdiag}(\mathbf{1}\,\pi^{(b)\top})$: within-block equilibrium, no leakage. On Π_R the co-flow of an arbitrary partition $\mathcal{C}$ is $F_{aa}^{\infty} = \sum_b V_{a\cap b}^2/V_b$.

Lemma 4 (identifiability on the idealized operator)**.** *Let every block satisfy $V_b < \frac{1}{2}$. Then R is the unique maximizer of $\mathrm{SI}^\infty(\mathcal{C}) = \sum_a F_{aa}^\infty\log_2(1/V_a)$ over all partitions, and $\mathrm{SI}^\infty(R) = \sum_b V_b\log_2(1/V_b)$.*

The proof (refining toward R never hurts; strict refinements of a block lose iff $V_b<\frac{1}{2}$; coarsenings lose by the grouping inequality) is in Appendix A.

Remark 1. *The condition $V_b<\frac{1}{2}$ is sharp: at $V_b=\frac{1}{2}$ (two balanced metastable blocks) the idealized objective* ties *between R and balanced refinements of a block — two-block metastability at exactly half volume is the boundary of identifiability for the two-level SE code, and it is testable: on a two-block chain with $V_b=\frac{1}{2}$ the optimizer never returns the planted pair, while three balanced blocks are recovered exactly (Appendix B).*

Lemma 5 (perturbation). *For any two stochastic $P, \tilde{P}$ and any partition, $|\mathrm{SI}(k;\mathcal{C};P) - \mathrm{SI}(k;\mathcal{C};\tilde{P})| \le \log_2(1/\pi_{\min}) \cdot \max_i \|(P^k)_{i\cdot} - (\tilde{P}^k)_{i\cdot}\|_1$.*

(Proof in Appendix A.)

Theorem 1 (metastable window of the SE spectrum). *Let $P = \bar{P} + \varepsilon E$ be nearly uncoupled: $\bar{P}$ block-diagonal with r ergodic blocks (partition R, block stationary distributions $\pi^{(b)}$, block volumes $V_b < \frac{1}{2}$, within-block second-largest eigenvalue modulus $\le \lambda_f$), $\|E\|_\infty \le 1$, and let $\gamma(V) > 0$ be the margin of Lemma 4. Fix $\delta < \gamma(V)/(3\log_2(1/\pi_{\min}))$. Then, by Simon–Ando/Courtois aggregation (Simon and Ando, 1961; Meyer, 1989), there are constants c_1, c_2 (independent of ε) such that for all*

$$t_{\text{fast}} := \frac{c_1 \ln(1/\delta)}{1-\lambda_f} \;\le\; k \;\le\; t_{\text{slow}} := \frac{\delta}{c_2\,\varepsilon},$$

$\max_i \|(P^k)_{i\cdot} - (\Pi_R)_{i\cdot}\|_1 \le 3\delta$, and consequently on the whole window: (i) $\mathcal{C}^*(k) = R$ (the planted metastable partition is the global optimizer); (ii) $\mathrm{SI}^*(k) = \sum_b V_b \log_2(1/V_b) \pm O(\delta \log_2(1/\pi_{\min}))$ — a plateau; (iii) the window is nonempty whenever $\varepsilon \lesssim \delta(1-\lambda_f)/(c_1 c_2 \ln(1/\delta))$, and its endpoints are $\Theta(1/(1-\lambda))$ for the fast and slow eigenvalue clusters respectively.

Proof sketch (within-block contraction at rate λ_f^k, leakage at most $k\varepsilon$, then Lemma 5): Appendix A.

Remark 2 (multi-level chains). *For a hierarchy with well-separated eigenvalue clusters $t_1 \ll t_2 \ll \ldots$, Theorem 1 applies level by level: between t_j and t_{j+1} the operator is δ-close to Π_{R_j} of the level-j partition, so the path coarsens through the hierarchy in order. We do not track the constants across levels; the three-level case is verified numerically below.*

4 Experiments

All experiments were preregistered (thresholds, seeds, k-grid, null count, significance rule) before the decisive runs; two deviations (optimizer symmetry between real and null arms; an effect-size floor of 0.1 bits joined to the $p \le 0.05$ rule because exact operators have no sampling noise) were recorded there before those runs, after being caught by the shipped deterministic fixture. Every number below is computed from the logged outputs (code and logs in the supplementary material); $n_{\text{null}} = 20$ matched draws per spectrum, 5 seeds per condition, empirical one-sided p.

4.1 G1: planted multiscale recovery (PASS 5/5)

Fixture. A 3-level hierarchical random walk on $n = 36$ states (3 super-blocks $\times$ 3 clusters $\times$ 4 states; per-step escape probabilities $1/200$ between supers, $1/20$ between sister clusters, lazy uniform mixing within clusters), analyzed exactly (no sampling). The model-free probe — implied timescales $t_i = -1/\ln|\lambda_i|$ of the teleported chain, computed *before* reading the spectrum — shows exactly three eigenvalue clusters: $t_{\text{slow}} = 56.9$ (2 eigenvalues; the raw chain's 133 is capped by the $\theta = 0.01$ teleport), $t_{\text{mid}} = 10.7$ (6 eigenvalues), $t_{\text{fast}} = 0.9$ (rest).

Result (Figure 2). In *all five* seeds the significant scale sets were identical: the optimal partition equals the planted 9-cluster partition (adjusted Rand index, ARI $= 1$) on $K_{\text{fine}} = \{1,2,3,4,6,8,11\}$ and the planted 3-super partition on $K_{\text{coarse}} = \{23,32,45,64,91,128\}$, in the preregistered order; the only other significant partition is a transitional one at $k = 16$ (6 modules, each super-block split in two, nested between the planted levels). Calibrated effect sizes: $\Delta\mathrm{SE}(k)$ peaks at 1.816 ± 0.004 bits (mean$\pm$sd over seeds) on the fine range and holds 0.343 bits at the middle of the coarse range, against 20 matched null draws each. The gate's theory-correspondence report: the fine$\to$coarse crossover lies between $k = 11$ and $k = 23$, i.e. at 1.03–$2.1 \times t_{\text{mid}}$ ($t_{\text{mid}} = 10.7$), and the last significant scale is $k = 128 \approx 2.2\,t_{\text{slow}}$ (an endpoint set by the effect-size floor: $\Delta\mathrm{SE} = 0.111$ there) — knees at $k \sim 1/(1-\lambda)$ as Proposition 1 and Theorem 1 predict. *The levels are read off the partition path, not off plateaus.* At this fixture's separation ($t_{\text{slow}}/t_{\text{mid}} \approx 5$) neither $\mathrm{SI}^*(k)$ nor $\Delta\mathrm{SE}(k)$ is flat on either range: the exact SI^* falls from 2.97 to 1.69 bits across K_{fine} and from 1.23 to 0.64 across K_{coarse}, below the idealized heights $\log_2 9 = 3.17$ and $\log_2 3 = 1.585$ of Theorem 1(ii); only prediction (i) is observed. Beyond $k \approx 256$ both spectra settle at the chance value of Lemma 3 ($\log_2(3)/3 \approx 0.53$ bits of raw SI^*, $\Delta\mathrm{SE} \approx 0$); the optimum is still the planted 3-partition (Figure 1d) but no longer significant, which the partition alone could not reveal. **Gate verdict: G1 PASS (5/5 seeds).**

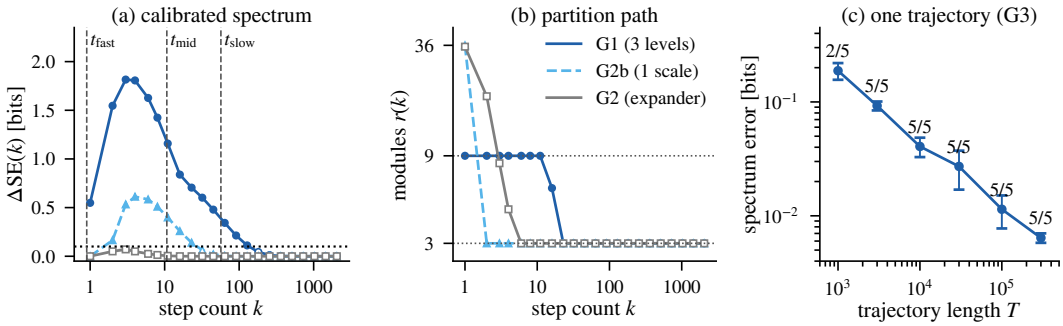

Figure 2: Planted recovery, controls and estimation. (a) $\Delta\mathrm{SE}(k)$ for G1 (3-level chain), G2b (single-scale modular chain) and G2 (expander): mean over 5 null seeds, band = min–max over seeds (narrower than the line for G1 and G2b); filled markers are significant in all 5 seeds; dotted line: 0.1-bit floor; dashed lines: model-free implied timescales. (b) Module count $r(k)$ of the optimal partition; the G1 path steps $9 \to 6 \to 3$, G2b stays at its 3 blocks, and the expander's path is never significant. (c) G3: spectrum error $\max_k |\Delta\mathrm{SE}_{\hat{P}}(k) - \Delta\mathrm{SE}_P(k)|$ from one trajectory of length T (mean $\pm$ SD over 5 trajectory seeds); labels give the seeds passing full G1 recovery.

4.2 G2/G2B: CONTROLS (PASS 5/5, PASS 5/5)

Single-scale expander (lazy walk on a random 6-out-regular digraph, matched dwell mass): in all five instances, *zero* of 22 grid points were significant (preregistered allowance: 15%), no significant run, and $\max_k \Delta\mathrm{SE}(k) \in [0.049, 0.094]$ bits, below the 0.1-bit floor. The spectrum is featureless when nothing is planted (Figure 2a). **G2 PASS (5/5). Single-scale modular chain** (3 blocks of 12, one escape rate): exactly one scale appears — every significant point matches the planted blocks (ARI $= 1$; $K_{\text{blocks}} = \{2, \ldots, 23\}$ in all five seeds) and no second level is ever reported. The spectrum does not manufacture a staircase. **G2b PASS (5/5).**

4.3 G3: ESTIMATION FROM A SINGLE FINITE TRAJECTORY (PASS, $T^* = 3000$)

Spectra estimated from stride-1 transition counts of one trajectory (plug-in estimator; the null is built from the *estimated* flow, so estimation noise is inside the calibration). Full G1 recovery (both levels, correct order, effect-size floor) holds in $2/5$ seeds at $T = 1000$ and $5/5$ from $T = 3000$ onward: $T^* = 3000$ steps, i.e. $\approx 22 \times t_{\text{slow}}$ of the raw chain, or ≈ 83 visits per state. The spectrum error $\max_k |\Delta\mathrm{SE}_{\hat{P}}(k) - \Delta\mathrm{SE}_P(k)|$ decays as expected for a plug-in estimator, from 0.188 ± 0.031 bits at $T{=}10^3$ and 0.093 ± 0.008 at $3{\times}10^3$ to 0.006 ± 0.001 at $3{\times}10^5$ (mean$\pm$sd over 5 trajectory seeds; Figure 2c). **Gate verdict: G3 PASS ($T^* = 3000 \leq 3 \times 10^5$).**

4.4 MSM HEAD-TO-HEAD (PREREGISTERED ADDENDUM, REPORTED, NOT GATED)

A preregistered addendum fixed before the run compares implied timescales (ITS) and PCCA+ (inner simplex, crisp readout, no refinement step), given the true module count r as an oracle (Table 1). **Exact chain:** the rivals match the spectrum: lag-1 ITS return the ground truth (56.9, 10.7) and PCCA+ with $r = 9$ or 3 recovers each level; with a wrong $r = 5$ it matches neither. **Expander:** ITS return finite timescales (2.0–2.3) and PCCA+ with $r = 3$ returns a single 36-state module; neither carries a significance statement, whereas the spectrum reports no significant k (G2). **Trajectories:** PCCA+ recovers the coarse level (ARI $= 1$ in $4/5$ seeds at $T = 1000$, $5/5$ at 3000) but the fine level less reliably, even with the best lag of $\{1, 2, 4, 8, 16\}$ chosen post hoc (SE higher in 4 of 5 seeds, one tie; one-sided sign test $p = 0.06$, not significant). The spectrum's best partition over its significant k matches both levels in all seeds at both lengths without being told r, though the full G3 criterion holds in only $2/5$ at $T = 1000$.

4.5 BEYOND STRUCTURAL ENTROPY (PREREGISTERED EXTENSION)

Markov stability at Markov time k, $\mathrm{MS} = \sum_a (F_{aa}(k) - V_a^2)$ (Delvenne et al., 2010), and the two-level map equation on the k-step flow (Rosvall and Bergstrom, 2008; Schaub et al., 2012) depend

Table 1: MSM head-to-head on G1: ARI of the recovered partition with the planted fine (9) and coarse (3) levels; higher is better. Trajectory cells are mean ± SD over 5 trajectory seeds (lag 1); the exact chain has no sampling noise. Bold: best in each column (ties included). PCCA+ is told r and has no refinement step; the SE readout takes the best partition over its significant k, a favourable readout. –: not run (lag sweep on trajectories only, wrong-r probe on the exact chain only).

		exact chain		T=1000 steps		T=3000 steps	
method	told r	fine ↑	coarse ↑	fine ↑	coarse ↑	fine ↑	coarse ↑
MSM rival (inner-simplex PCCA+, crisp readout)							
PCCA+, lag 1	yes	**1.00**	**1.00**	0.63 ± 0.31	0.91 ± 0.18	0.77 ± 0.20	**1.00** ± 0.00
PCCA+, best lag (post hoc)	yes	–	–	0.76 ± 0.16	0.91 ± 0.18	0.90 ± 0.06	**1.00** ± 0.00
PCCA+, lag 1, wrong r	r=5	0.52	0.74	–	–	–	–
this paper							
SE spectrum, best sig. k	no	**1.00**	**1.00**	**1.00** ± 0.00	**1.00** ± 0.00	**1.00** ± 0.00	**1.00** ± 0.00

Table 2: Gate verdicts per objective wrapped in the same matched null (seeds passing, of 5; the preregistered pass rule is $\geq 4/5$; identical under Holm and Benjamini–Hochberg (BH) correction over k). Crossover: last fine / first coarse k on G1 ($t_{\mathrm{mid}} = 10.7$). T^*: shortest passing trajectory (lower is better). SHARP ($V_b = \frac{1}{2}$) is the boundary case of Lemma 4, where SE is predicted to fail.

	planted structure			null controls		boundary	estimation
objective	G1 levels	crossover	G2b	G2	DWELL	SHARP	$T^* \downarrow$
Markov stability	5/5	4 / 6	5/5	3/5	5/5	5/5	> 3000
map equation	5/5	4–6 / 8–11	5/5	5/5	4/5	5/5	3000
SE (this paper)	5/5	11 / 23	5/5	5/5	5/5	0/5	3000

on a partition only through $(F_{aa}(k), V_a)$, like SE. We ran all three with one optimizer, identical null draws (500 per spectrum) and the main-text criteria, in a separate preregistration (floors: 0.1 bits for SE and the map equation, 0.03 for MS); SHARP is two balanced blocks ($V_b = \frac{1}{2}$) and DWELL an expander with G4's self-loop mass (0.85). Our map-equation optimum equals the official Infomap codelength on 64 of 66 k-step flows (Appendix D).

Result (Table 2). The map equation meets the preregistered generality criterion (G1, G2, G2b in $\geq 4/5$ seeds, uncorrected and under Holm). Markov stability finds both levels and the single scale but fails the expander control in $2/5$ seeds (peak 0.034–0.041 above its 0.03 floor; all pass at 0.06), so for MS the matched null alone does not remove all spurious structure. All three return the planted partitions in the planted order, but *where* the path coarsens depends on the objective: SE keeps the nine fine modules up to $k = 11$, MS and the map equation only to $k = 4$–6. The reason (Appendix C): SE is Markov stability with each module weighted by its code length, which favours the small fine modules ($\log_2 9$) over the coarse ones ($\log_2 3$). The same weight falls below 1 when $V_a > \frac{1}{2}$, which is SE's volume condition: on SHARP, SE never returns the planted pair (Lemma 4's tie) while MS and the map equation do in $5/5$, as Proposition 2 predicts for MS. The map equation's coarse range ends at $k = 32$; beyond it its optimum is one module. A hierarchy objective, Dasgupta's cost (Dasgupta, 2016), recovers the nesting at a single k but no timescale (Appendix D). On these fixtures the principle is not SE-specific, but scale boundaries depend on the objective as well as the chain.

4.6 G4: THE RIGHT TIMESCALE FOR A WORLD MODEL'S LATENT GRAPHS

Latent transition graphs of a world model can be built at stride 1 or stride 5, and nothing in SE theory chooses between them. We ran the preregistered analysis on five independent training seeds (latent rollouts, 3000 steps, 100 stored prototypes; stride-1 counts; 20 matched nulls each). The rollouts are not one continuous trajectory: they come from a pixel JEPA on a robot-arm/cube scene driven by random actions, the environment is reset every 200 steps (about 14 resets per rollout), and the cube is nearly static within an episode (Figure 3, Table 3):

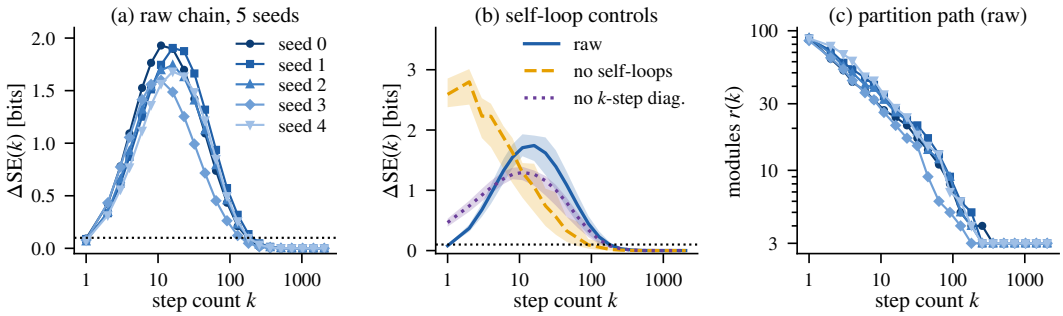

Figure 3: SE spectrum of the world model's latent transition graphs (five training seeds). (a) $\Delta\mathrm{SE}(k)$ of the raw chain, one line per seed; dotted: 0.1-bit floor. Stride 1 is below the floor; the peak is at $k^* = 11$–16. (b) Self-loop controls (200 nulls): mean over seeds, band = min–max. Zeroing self-loops moves the peak to $k = 2$; excluding the k-step diagonal keeps it near the raw k^*. (c) Module count of the raw optimum: smooth coarsening, no plateau.

Table 3: G4 per training seed. s: stride-1 self-transition share. Raw chain: $\Delta\mathrm{SE}$ at $k = 1$ (all below the 0.1-bit floor), peak location k^*, peak $\Delta\mathrm{SE}$ in bits, module count at the peak, and the largest significant k. Controls (200 nulls, Appendix E): peak location and height with self-loops zeroed and with the k-step diagonal excluded. The last column rescales the raw peak to state changes; it matches the jump-chain peak at $k = 2$.

		raw chain, 20 nulls					no self-loops		no k-step diag.		
seed	s	$\Delta\mathrm{SE}(1)$	k^*	$\Delta\mathrm{SE}(k^*)$	$r(k^*)$	last sig. k	k^*	peak	k^*	peak	$k^*(1-s)$
0	0.82	0.09	11	1.93	27	128	2	3.01	11	1.39	1.99
1	0.87	0.09	16	1.91	27	181	2	2.70	11	1.33	2.03
2	0.85	0.07	16	1.74	26	181	2	2.95	11	1.34	2.42
3	0.81	0.08	11	1.60	26	128	2	2.42	8	1.18	2.12
4	0.88	0.07	16	1.68	28	128	2	2.91	16	1.27	1.84

(i) Stride 1 shows little calibrated structure. At $k = 1$, $\Delta\mathrm{SE} = 0.07$–0.09 bits — *below the* 0.1*-bit effect-size floor in all five seeds*, although it exceeds all 20 null draws ($p = 0.048$), so the effect is small rather than absent. The one-step graph is dominated by dwell (self-transition mass 0.81–0.88) and margins, which the null matches exactly; almost nothing else survives calibration. This is, however, a low-power regime: the off-diagonal mass the null randomises is small. **(ii) A candidate answer to the stride question.** $\Delta\mathrm{SE}(k)$ is unimodal and reproducible across seeds: significant for $k \in [2, 128]$ (two seeds: $[2, 181]$), peaking at $k^* = 11$–16 with 1.60–1.93 bits and $r \approx 26$–28 modules at the peak. Stride 5 sits inside the band, but the neighbouring grid points $k = 4$ and $k = 6$ recover only 46–66% and 66–88% of the peak effect. **Self-loop controls** (preregistered; 200 nulls, BH correction over k; Appendix E; the raw spectrum is reproduced). With self-loops zeroed, the structure survives and grows: significant from $k = 1$ to $k \geq 45$ in all five seeds, peaking at $k^* = 2$ with 2.4–3.0 bits. Excluding the k-step diagonal from both arms keeps the peak within one grid step of the raw k^* (8–16, 1.2–1.4 bits). So the G4 structure is not a dwell artifact, but the *location* $k^* \approx 11$–16 is a dwell-inflated step count: $k^*(1-s) = 1.8$–2.4, with s the stride-1 self-transition share, matches the jump chain's peak at two state changes (Table 3). Dwell alone does not create the effect: the DWELL expander peaks at the same $k = 11$ with only 0.02–0.05 bits. **(iii) An honest negative.** Unlike the planted fixture, the partition path coarsens *smoothly* ($r \approx 63$–$78 \to 4$–6 across the significant band, no extended plateau). By Theorem 1 read contrapositively, these graphs do not meet its hypotheses (well-separated eigenvalue clusters with blocks of volume $< \frac{1}{2}$). We do not read this as a continuum of dynamical scales: the significant band ends below one 200-step episode, and on reset episodes of a nearly static scene the coarse partitions may track episode identity, not dynamics. Smooth coarsening is also what a chain without resets and without a hierarchy gives (Section 4.7).

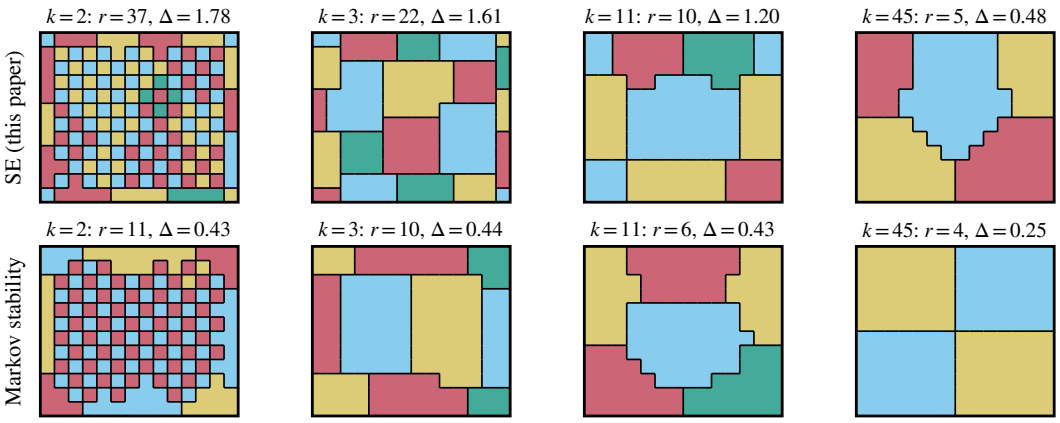

Figure 4: Case study: optimal partitions of the Pin Pad random-walk chain drawn on the room (14×12 positions; one colour per module, black lines between modules, so the partition reads without colour). Top: SE; bottom: Markov stability, on the same k-step flows and 200 matched nulls. Titles give r and the calibrated gain Δ in each objective's own units (bits for SE); every panel is significant. Columns: SE's peak $k = 2$ and the grid points 3, 11, 45.

4.7 Case study: the partition path in state space

G4's prototypes have no spatial layout, so we inspect a chain whose states do: a random walk in the Pin Pad room of Hafner et al. (2022) (168 positions, sticky random policy, 300 within-episode walks of 1000 steps; Appendix E), with up to 1.8 bits and r falling from 34–37 to 3. Figure 4 draws the optimal partitions of SE and of Markov stability (SE without the code-length weight) on the room. The columns are fixed by rule: SE's peak $k = 2$, then the grid points 3, 11, 45 (about $4\times$ apart); there is one chain, so nothing is selected across seeds. **(i)** At odd k the modules are contiguous patches that merge smoothly ($22 \to 10 \to 5$ for SE); nothing marks a preferred scale, as expected for a diffusive room. **(ii)** At $k = 2$ both objectives return interior modules that interleave like a checkerboard, i.e. the lattice's even/odd parity classes: 89% of two-step transitions end on the parity they started from, against 43% of one-step ones (the sticky policy often repeats a move). Part of what SE compresses at its peak is thus lattice parity, not spatial neighbourhood. **(iii)** At every k shown, SE keeps more and smaller modules than Markov stability (22 vs. 10 at $k = 3$), as the code-length weight predicts. A large ΔSE signals local structure, sometimes of an unintended kind, not a discrete hierarchy; the partition shows which.

5 Limitations

(i) Theorem 1 covers the two-level (single window) case with explicit constants left unoptimized; the multi-level statement is verified numerically, not proved with tracked constants. (ii) The $V_b < \frac{1}{2}$ condition is sharp for the two-level SE code: two balanced metastable blocks sit exactly on the identifiability boundary; deeper codes may lift this and we have not pursued it. (iii) The null's fast mixing is measured, not proved. (iv) The optimizer is a local search; it matched or beat the planted partitions on all fixtures, but global optimality is not guaranteed. (v) The real-data demonstration takes its state discretization (stored prototype labels) as given; the spectrum answers "which timescale, given these states", not "which states", a separate problem. (vi) Trajectory estimation uses the plug-in estimator, with no bias correction beyond what the matched null absorbs. (vii) The G4 rollouts come from one world model on one scene and are reset episodes of a nearly static scene, not a stationary process, which limits what they say about latent dynamics; with 100 prototypes and 3000 steps they also give ≈ 30 transitions per state, fewer than the ≈ 83 visits per state G3 needed on the 36-state fixture. (viii) The G4 self-loop controls are not a null matched on the k-step diagonal; (ix) the 0.1-bit floor, not the p-value, carries the expander control, even with 500 nulls; (x) scale is untested beyond $n = 168$; and (xi) the plateaus of Theorem 1(ii) are approached, not confirmed, even at $\approx 209\times$ separation (details: Appendix F). (xii) ΔSE is a code-length effect size; we do not test whether graphs built at the selected k improve a downstream agent or world model.

REPRODUCIBILITY STATEMENT

All formal claims are stated with their assumptions in the main text and proved either there or in the appendix. Every experiment reported here is a closed specification: the constructions, their parameters, the seed set, the null model and its draw count, and the acceptance criteria are given in the experimental sections, and the criteria were fixed before the decisive runs. Code sufficient to regenerate every number and figure, together with the raw logged outputs and the preregistrations recording each criterion and its verdict, is included as anonymous supplementary material.

AI USE STATEMENT

We used generative AI tools on the following tasks from ICLR's required-disclosure list: helping to develop the conceptual framework, helping to formulate mathematical claims, providing ingredients for and assisting in the writing of proofs, proposing and refining hypotheses, giving feedback on experimental design, implementing methods, and helping to interpret results. We did *not* use generative AI tools to generate synthetic data sets reported as measurements, to clean or reformat data sets, for translation, or for qualitative or thematic data analysis. From the recommended-disclosure list we additionally used generative AI tools to draft and edit parts of this paper, to create and edit figures, to identify and summarise related literature, to format references, and to suggest a structure for the paper.

We have reviewed all AI-assisted work. Every theorem statement and proof appearing in this paper was checked line by line by the authors; every numerical value reported was regenerated from the released code and compared against the logged experimental outputs; every preregistered criterion and its verdict was recorded before the corresponding run and is reported verbatim, including the failures. We take responsibility for the final content of this work, including all text, claims and artifacts produced with the aid of generative AI.

REFERENCES

A. Li and Y. Pan. Structural information and dynamical complexity of networks. *IEEE Transactions on Information Theory*, 62(6):3290–3339, 2016.

T. Iiyama, M. Suzuki, and Y. Matsuo. SUNTA: Hierarchical video prediction with surprise-based chunking. arXiv:2607.02087, 2026.

An open-source library for computing, optimizing and benchmarking structural entropy. Software, 2026. Anonymized for double-blind review.

X. Zeng, H. Peng, A. Li, C. Liu, L. He, and P. S. Yu. Hierarchical state abstraction based on structural information principles. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4549–4557, 2023.

X. Zeng, H. Peng, and A. Li. Effective exploration based on the structural information principles. In *Advances in Neural Information Processing Systems 37*, pages 95291–95325, 2024.

When does structural entropy track labels? Margins, diagnostics, and failure modes in graph clustering. *CIKM*, 2026. doi:10.1145/3799682.3840913.

J.-H. Prinz et al. Markov models of molecular kinetics: generation and validation. *J. Chem. Phys.*, 134:174105, 2011.

P. Deuflhard and M. Weber. Robust Perron cluster analysis in conformation dynamics. *Linear Algebra Appl.*, 398:161–184, 2005.

H. A. Simon and A. Ando. Aggregation of variables in dynamic systems. *Econometrica*, 29(2):111–138, 1961.

C. D. Meyer. Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Review*, 31(2):240–272, 1989.

R. S. Sutton, D. Precup, and S. Singh. Between MDPs and semi-MDPs: a framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1–2):181–211, 1999. doi:10.1016/S0004-3702(99)00052-1.

R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.

J.-C. Delvenne, S. N. Yaliraki, and M. Barahona. Stability of graph communities across time scales. *Proceedings of the National Academy of Sciences*, 107(29):12755–12760, 2010.

R. Lambiotte, J.-C. Delvenne, and M. Barahona. Random walks, Markov processes and the multiscale modular organization of complex networks. *IEEE Transactions on Network Science and Engineering*, 1(2):76–90, 2014.

M. Rosvall and C. T. Bergstrom. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123, 2008.

M. T. Schaub, R. Lambiotte, and M. Barahona. Encoding dynamics for multiscale community detection: Markov time sweeping for the map equation. *Physical Review E*, 86(2):026112, 2012.

D. Hafner, K.-H. Lee, I. Fischer, and P. Abbeel. Deep hierarchical planning from pixels. In *Advances in Neural Information Processing Systems 35*, 2022. arXiv:2206.04114.

A. Arnaudon, J. Schindler, R. L. Peach, A. Gosztolai, M. Hodges, M. T. Schaub, and M. Barahona. Algorithm 1044: PyGenStability, a multiscale community detection framework with generalized Markov stability. *ACM Transactions on Mathematical Software*, 50(2):1–8, 2024. doi:10.1145/3651225.

N. Nitskansky, K. Clein, and B. Raveh. Building multiscale Markov state models by systematic mapping of temporal communities. *Bioinformatics*, 42(1):btaf585, 2025. doi:10.1093/bioinformatics/btaf585.

S. Dasgupta. A cost function for similarity-based hierarchical clustering. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 118–127, 2016.

B. Moseley and J. Wang. Approximation bounds for hierarchical clustering: average linkage, bisecting k-means, and local search. In *Advances in Neural Information Processing Systems 30*, pages 3094–3103, 2017.

A Deferred proofs

Proof of Lemma 1. Split the node term: $-\sum_{i\in a}\pi_i\log_2(\pi_i/V_a) = -\sum_{i\in a}\pi_i\log_2\pi_i + V_a\log_2 V_a$. Summing over a gives $H_1 + \sum_a V_a\log_2 V_a$. Since $g_a(k) = V_a - F_{aa}(k)$, the cut term is $-\sum_a(V_a - F_{aa}(k))\log_2 V_a$. Adding the two cancels $\sum_a V_a\log_2 V_a$ and leaves $H_1 + \sum_a F_{aa}(k)\log_2 V_a$. Nonnegativity: $F_{aa}\geq 0$ and $\log_2(1/V_a)\geq 0$. □

Proof of Lemma 5. SI is linear in the rows of P^k with coefficients $\pi_i\log_2(1/V_{a(i)}) \leq \pi_i\log_2(1/\pi_{\min})$; Hölder. □

Proof of Proposition 1. $F_{aa}(k) = \langle \mathbf{1}_a, P^k\mathbf{1}_a\rangle_\pi$; expand $\mathbf{1}_a$ in the eigenbasis: $\mathbf{1}_a = \sum_i\langle f_i,\mathbf{1}_a\rangle_\pi f_i$ and use $\langle f_1,\mathbf{1}_a\rangle_\pi = V_a$. Reversibility gives real eigenvalues and the squared (hence nonnegative) coefficients. For general chains the same expansion holds with possibly complex λ_i and unsigned coefficients; the step-decay statement then holds under diagonalizability with cluster separation, which we verify numerically for our directed fixtures. □

Proof of Lemma 4. *(i) Refining toward R never hurts.* For any $\mathcal{C}$, the common refinement $\mathcal{C}\wedge R$ (cells $a\cap b$) satisfies $\mathrm{SI}^\infty(\mathcal{C}\wedge R)\geq \mathrm{SI}^\infty(\mathcal{C})$, because each mass contribution $V^2_{a\cap b}/V_b$ is repriced from $\log_2(1/V_a)$ to $\log_2(1/V_{a\cap b})\geq\log_2(1/V_a)$. *(ii) Refinements of R strictly lose.* A refinement splits block b into cells with shares $x_s = V_{bs}/V_b$, $\sum_s x_s = 1$. Its contribution is $V_b\sum_s x_s^2[\log_2(1/V_b) + \log_2(1/x_s)]$, versus $V_b\log_2(1/V_b)$ for the intact block. The refinement loses iff $\log_2(1/V_b)\,(1-\sum_s x_s^2) > \sum_s x_s^2\log_2(1/x_s)$. The ratio $\sum_s x_s^2\log_2(1/x_s)/(1-\sum_s x_s^2)$ is maximized over all splits by the balanced binary split $x = (\frac{1}{2},\frac{1}{2})$, where it equals 1 (balanced m-way splits give

$\log_2(m)/(m-1) \le 1$, and skewed splits give less). Hence $\log_2(1/V_b) > 1$, i.e. $V_b < \frac{1}{2}$, makes every strict refinement of any block strictly worse, with margin at least $V_b\,(\log_2(1/V_b) - 1)\,(1 - \sum_s x_s^2)$. *(iii) Coarsenings of R strictly lose.* Merging blocks b, c into a gives $F^{\infty}_{aa} = V_b + V_c = V_a$ and contribution $V_a \log_2(1/V_a)$, less than $V_b \log_2(1/V_b) + V_c \log_2(1/V_c)$ by the grouping (log-sum) inequality; the deficit is $V_a\, \mathrm{H}_2\big(\frac{V_b}{V_a}\big)\cdot$(positive factor), in particular positive. Chaining: $\mathrm{SI}^{\infty}(\mathcal{C}) \le \mathrm{SI}^{\infty}(\mathcal{C} \wedge R) \le \mathrm{SI}^{\infty}(R)$, the second inequality strict unless $\mathcal{C} \wedge R = R$ (by (ii) applied blockwise), and if $\mathcal{C} \wedge R = R$ then $\mathcal{C}$ coarsens R, strict by (iii) unless $\mathcal{C} = R$. □

Proof sketch of Theorem 1. Within-block mixing contracts P^k toward Π_R at rate λ_f^k on the uncoupled part, while leakage after k steps is bounded by $k\varepsilon$ in row-sum norm ($\|P^k - \bar{P}^k\|_{\infty} \le k\varepsilon\|E\|_{\infty}$ by telescoping $P^k - \bar{P}^k = \sum_j \bar{P}^j (P - \bar{P}) P^{k-1-j}$ and submultiplicativity of $\|\cdot\|_{\infty}$ over stochastic factors); $\|\bar{P}^k - \Pi_{\bar{R}}\|_{\infty} \le c_1' \lambda_f^k$ with $\Pi_{\bar{R}}$ built from block stationary distributions, and $\|\Pi_{\bar{R}} - \Pi_R\|_{\infty} = O(\varepsilon)$ since the stationary distribution of the coupled chain is an $O(\varepsilon)$ perturbation of the block mixture. Choosing k in the stated window bounds all three terms by δ. Lemma 5 then transports the idealized landscape of Lemma 4 with total distortion $< \gamma(V)$, so the strict optimizer is unchanged and (ii) follows. Endpoint scaling: $1 - \lambda_f$ is the fast cluster's gap; the slow eigenvalue cluster of a nearly uncoupled chain satisfies $1 - \lambda_i = \Theta(\varepsilon)$ (Meyer, 1989), so $t_{\mathrm{slow}} = \Theta(1/(1 - \lambda_{\mathrm{slow}}))$ up to the factor δ/c_2. □

Remark 3 (what the null adds to the theorem)**.** *The theorem describes* $\mathrm{SE}(k)$*; the null spectrum is what makes the plateau readable at finite* n*: the null chain has the same margins and dwell but (empirically)* $\Theta(1)$ *spectral gap, so* $\mathrm{SE}_{\mathrm{null}}(k)$ *sits at the chance value of Lemma 3 for all* $k \gtrsim$ *a few steps, and in the asymptotic window* $\Delta\mathrm{SE}(k)$ *displays the plateau at height* $\approx \sum_b V_b \log_2(1/V_b) - \max_{\mathcal{C}} \mathrm{chance}$ *(at moderate scale separation it is not reached; Section 4, G1). We do not prove the null's mixing rate (spectral gaps of margin-preserving random weighted digraphs are outside our scope); it is measured in every experiment, and the significance test never assumes it.*

B Theory checks (post hoc, explicitly not gated)

Two additional computations verify the theory rather than test a hypothesis, and were run after the gates. **(a) Spectral mixture exactness.** On the (directed) 3-level fixture, $\mathrm{SI}(k;\mathcal{C})$ computed directly and via the eigen-mixture of Proposition 1 (full biorthogonal eigendecomposition) agree to $\le 2.4 \times 10^{-13}$ bits at $k \in \{1, 4, 16, 64, 256, 1024\}$ for both planted partitions. **(b) Sharpness of $V_b < \frac{1}{2}$.** On a two-block metastable chain with balanced blocks ($V_b = \frac{1}{2}$, exactly the boundary of Lemma 4), at a mid-window $k = 16$ the planted partition *ties* with its balanced refinement (margin -1.2×10^{-6} bits) and the optimizer returns refinements, never the planted 2-partition, in $10/10$ restarts (best ARI 0.74). The same chain with three balanced blocks ($V_b = \frac{1}{3} < \frac{1}{2}$) has margin $+0.076$ bits and exact recovery in $10/10$ restarts. The identifiability boundary in the theorem is therefore real optimizer behavior, not proof slack; it is also a practical warning: two-regime systems with balanced occupancy sit exactly on it.

C Other clustering objectives on the k-step operator

Markov stability at Markov time k (Delvenne et al., 2010; Lambiotte et al., 2014) and the two-level map equation (Rosvall and Bergstrom, 2008) evaluated on the k-step flow $\mathrm{diag}(\pi)P^k$ (the discrete analogue of Markov-time sweeping, Schaub et al., 2012) depend on a partition only through $(F_{aa}(k), V_a)_a$, exactly like SI:

$$\mathrm{MS}(k;\mathcal{C}) = \sum_a \big(F_{aa}(k) - V_a^2\big), \qquad \mathrm{ME}(k;\mathcal{C}) = H_1 - L(k;\mathcal{C}),$$

$$L = q \log_2 q - 2\sum_a q_a \log_2 q_a + H_1 + \sum_a (q_a + V_a)\log_2(q_a + V_a),$$

with $q_a = V_a - F_{aa}(k)$ the k-step exit flow and $q = \sum_a q_a$. Lemma 2 therefore holds for all three (the k-dependence lives in $F(k)$), and subtracting the chance term of Lemma 3 from Lemma 1 gives the exact identity

$$\mathrm{SI}(k;\mathcal{C}) - \mathrm{chance}(\mathcal{C}) = \sum_a \big(F_{aa}(k) - V_a^2\big) \log_2 \frac{1}{V_a},$$

i.e. SE's saving above chance is Markov stability with each module weighted by its code length $\log_2(1/V_a)$. Proposition 1 applies to MS verbatim (drop the weight).

Proposition 2 (Markov stability on the idealized operator)**.** *On Π_R, R is the unique maximizer of* $\mathrm{MS}^{\infty}(\mathcal{C}) = \sum_a (F^{\infty}_{aa} - V^2_a)$ *for* any *block volumes.*

Proof. (i) For any $\mathcal{C}$, passing to $\mathcal{C} \wedge R$ changes cell a's contribution from $\sum_b V^2_{a\cap b}/V_b - (\sum_b V_{a\cap b})^2$ to $\sum_b (V^2_{a\cap b}/V_b - V^2_{a\cap b})$, a gain of $(\sum_b V_{a\cap b})^2 - \sum_b V^2_{a\cap b} \geq 0$, strict if a meets two blocks. (ii) Splitting block b into shares x_s gives $V_b(1-V_b)\sum_s x_s^2 < V_b(1-V_b)$. (iii) Merging blocks b, c loses exactly $2V_bV_c > 0$. Chain as in Lemma 4. □

Since MS is linear in the rows of P^k with coefficients π_i, the analogue of Lemma 5 holds with $\log_2(1/\pi_{\min})$ replaced by 1, and the argument of Theorem 1 transfers with the margin of Proposition 2; there is no volume condition. The $V_b < \frac{1}{2}$ boundary is thus specific to SE's code-length weight, which drops below 1 exactly when a module holds more than half the stationary mass. For the map equation, coarsening R strictly loses (L rises by $H(V) - H(V') > 0$) and the balanced split of a block with $V_b = \frac{1}{2}$ loses (1.0 vs. 0.56 bits of saving); we did not prove the general refinement case. Unlike SE and MS, the map equation's optimum becomes the one-module partition once the k-step exit flow is large, which shortens its coarse window (Table 2).

D Extension experiments: protocol and full results

Protocol. Everything in this appendix was preregistered (timestamped before any null-calibrated extension result was seen; operational deviations recorded in the preregistration, which is in the supplementary material together with the code and logs that produce every number below). All objectives share one optimizer (single-node moves and greedy merges, 4 restarts, warm starts along k) and the same 500 null draws per spectrum; the generic SE arm reproduces the main-text G1 spectrum exactly, with null means within 0.005 bits of the main-text optimizer. Floors: 0.1 bits (SE, map equation) and 0.03 (MS; 0.1 bits/ $\log_2 9$). SHARP = two blocks of 12 states, $V_b = \frac{1}{2}$, escape $1/100$, lazy 0.5; DWELL = the G2 expander with lazy 0.85 (G4's self-transition share is 0.81–0.88).

Official package. Infomap 2.9.2 (two-level, directed, raw directed flow model, 10 trials) on the exact k-step flows of G1, G2b and SHARP (66 points) agrees with our map-equation optimum to 2×10^{-13} bits at 64 points; at G1 $k = 6$ ours is 0.036 bits better and at G2b $k = 8$ Infomap is 0.038 bits better, both at a level transition. Infomap's own Markov-time sweep on P^1 returns $r = 9$ for $t \leq 4$, $r = 3$ for $6 \leq t \leq 23$ and one module from $t = 32$.

Floor and correction sensitivity (seeds passing, of 5). Holm and BH never change a verdict in Table 2. At half the floor, G2 passes for the map equation only (5/5; SE 1/5, MS 0/5) and DWELL passes in 3/5, 2/5, 1/5 (SE, MS, map equation); at twice the floor every control passes for every objective. G1, G2b and SHARP verdicts are the same at all three floors. With the paper's own implementation and 500 nulls, G1, G2 and G2b pass 5/5 uncorrected, under Holm and under BH; G2 at a 0.05-bit floor passes 1/5, the same as the main-text 20-null spectra.

Self-loop variants on the fixtures (prediction: all pass). Zeroing self-loops shortens every timescale (G1 jump chain: $t_{\text{slow}} = 41.0$, $t_{\text{mid}} = 5.7$). SE still passes G1 (5/5; fine $k \leq 8$, coarse 11–91), but MS (0/5) and the map equation (3/5) keep the fine level only up to $k = 2$–3, below the three-scale criterion: a FAIL of the prediction for those two objectives. Excluding the k-step diagonal passes G1 for all three (5/5). On DWELL, the diagonal-excluded variant passes for all three; the self-loop-zeroed variant fails for SE and MS (0/5; 0.10–0.19 and 0.05–0.08 at $k \leq 2$), plausibly because a sparse random digraph without its self-loops already carries some calibrated $k = 1$ structure against the denser null. We therefore read jump-chain effects only when they are far above that level (G4: 2.4–3.0 bits).

Plateau fixture (PLATEAU = three blocks of 12, escape 1/1000, no laziness; $t_{\text{fast}} = 0.415$, $t_{\text{slow}} = 86.6$; window $k \in \{3, \ldots, 16\}$). SE: SI* $\in [1.41, 1.55]$ (range 0.14, largest deviation from $\log_2 3$ 0.18), ΔSE range 0.13: FAIL in 5/5 against the 0.1-bit tolerance. The planted partition is the optimum at every significant k for all three objectives. MS varies by 0.09 on the window; the map equation by 0.51.

Dasgupta's cost (average-linkage tree (Moseley and Wang, 2017) on the symmetrized k-step flow, cost normalized by $n \sum_{i<j} w_{ij}$; calibrated gain = relative reduction against the same nulls). On G1 both planted levels are exact subtrees at every $k \leq 362$ (coarse level to $k = 1448$) and the gain is 0.77 at $k = 1$, decaying monotonically; on G2 and DWELL the largest gain is 0.02–0.09. From trajectories the fine level survives as a subtree up to $k = 2$–8 ($T = 1000$) and 4–23 ($T = 3000$). A hierarchy objective thus recovers the nesting at a single k but does not by itself say at which timescale each level is real. Between our two machines the Dasgupta gains differ by up to 0.002 (ties in average linkage); all other arms are deterministic.

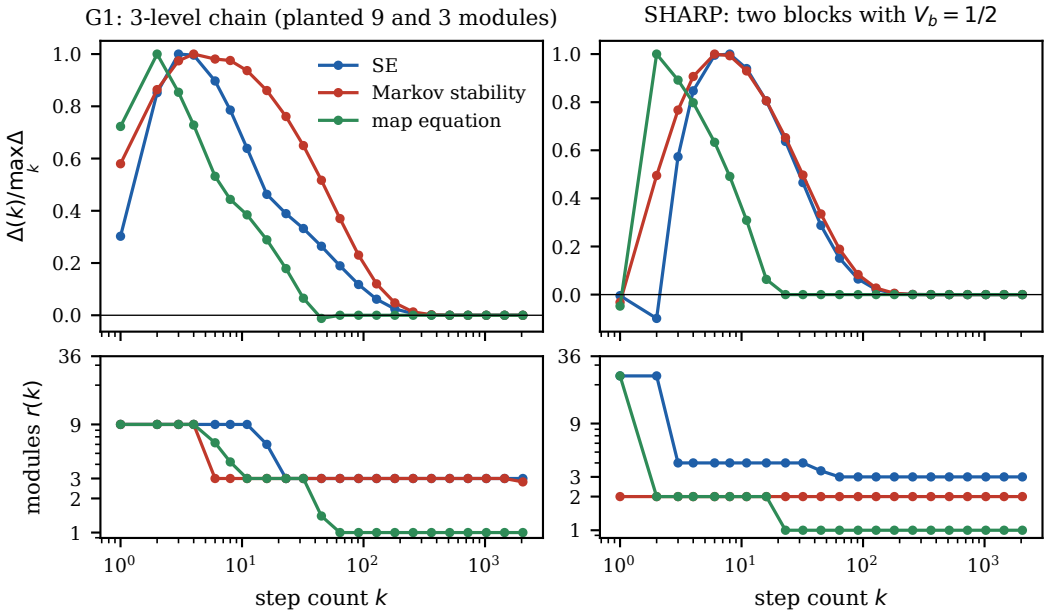

Figure 5: The same construction around three objectives. Top: calibrated gain, each normalized by its maximum (mean of 5 null seeds). Bottom: module count of the optimum. Left: G1 (planted 9 and 3 modules; all three pass). Right: SHARP ($V_b = \frac{1}{2}$): SE returns four modules, MS and the map equation the planted two.

E G4 SELF-LOOP CONTROLS AND A CHAIN WITHOUT RESETS

G4 was rerun with 200 nulls per spectrum (BH over k as the corrected rule, since Holm is infeasible at $p \geq 1/201$) in three variants: raw (as in the main text), self-loops zeroed (jump chain, its own null), and the k-step diagonal excluded from the objective of both arms. Preregistered outcomes: (A) structure beyond dwell = jump chain significant at ≥ 3 grid points in $\geq 4/5$ seeds: **met** ($5/5$; BH-significant from $k = 1$ to 45–128). (B) raw peak robust to k-step return mass = diagonal-excluded k^* within one grid step of raw k^* in $\geq 4/5$ seeds: **met** ($5/5$; $k^* = 11, 11, 11, 8, 16$ vs. $11, 16, 16, 11, 16$). (C) jump-chain k^* vs. the dwell-rescaled prediction $k^*(1 - s)$: $k^*_{\text{jump}} = 2$ in all seeds against $1.99, 2.03, 2.42, 2.12, 1.84$ (agreement), so the raw k^* is read as a dwell-inflated step count. The raw rerun reproduces the main-text k^*, peak (1.60–1.93 bits), $r^* = 26$–28 and significant band; at $k = 1$, $\Delta\text{SE} = 0.07$–0.09 now has $p = 1/201$ but stays below the floor. MS and the map equation agree qualitatively (raw $k^* = 11$–23 and 6–11; jump chain $k^* = 2$–8 and 1–2).

The Pin Pad chain comes from the Director codebase's Pin Pad environment (Hafner et al., 2022) under a sticky uniform-random policy (repeat probability 0.8), 300 episodes of 1000 steps; states are the 168 visited positions, only within-episode transitions are counted (no reset edges), and $s = 0.43$. With 200 nulls, SE is significant for $k \in [1, 91]$ with peak 1.78 bits at $k = 2$; the partition path falls from 34–37 modules to 3 with no extended plateau, as preregistered for a diffusive room without separated eigenvalue clusters. The self-loop-zeroed chain oscillates between even and odd k (the grid walk without self-loops is nearly bipartite).

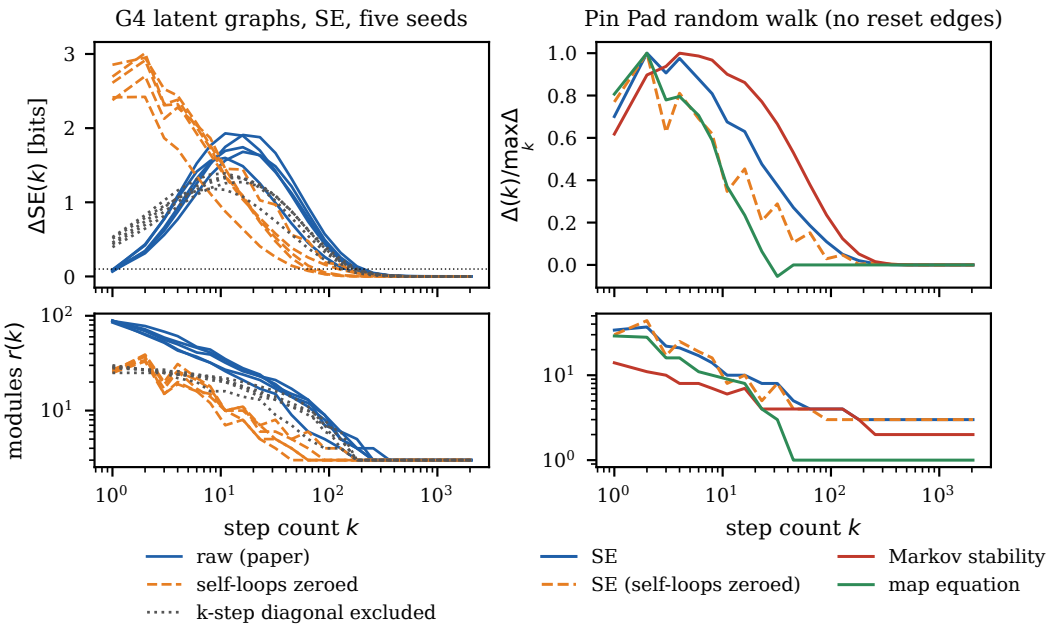

Figure 6: Left: G4 SE spectra for five seeds, raw (as in Figure 3), self-loops zeroed, and k-step diagonal excluded (dotted line: 0.1-bit floor). Right: Pin Pad random walk, normalized gain and module count for SE (raw and self-loops zeroed), MS and the map equation.

F Details of limitations (viii)–(xi)

(viii) The G4 self-loop controls remove dwell from the chain or the k-step diagonal from the objective; neither is a null matched on the k-step diagonal, and on the structureless DWELL expander the self-loop-zeroed chain itself shows 0.10–0.19 bits at $k = 1$, so jump-chain effects of that size at $k \leq 2$ are not interpretable. (ix) With 20 null draws $p \geq 1/21$; with 500 draws and Holm correction the gate verdicts are unchanged, but G2 passes at the 0.1-bit floor and not at 0.05 (1/5 seeds; its 500-null peaks are 0.044–0.094 bits), so the floor still carries that control. (x) Scale is untested beyond $n = 168$; a 200-null G4 spectrum took 18–20 CPU-minutes. (xi) G1 has only $\approx 5\times$ separation between its two slowest timescales, too little for the plateaus of Theorem 1(ii); a two-level fixture with $t_{\text{slow}}/t_{\text{fast}} \approx 209$ brings SI* to 1.41–1.55 bits on the window ($\log_2 3 = 1.585$) but misses the preregistered 0.1-bit flatness tolerance (range 0.14), so (ii) is approached, not confirmed (Appendix D).