Black-Box Empirical Characterisation of a Novel Proprietary Spatiotemporal Engine Near the Order–Chaos Boundary

Leroy A. Palmer
The EoC Project
palmer@eormen.com eormen.com

February 2026

Abstract

We introduce a novel proprietary spatiotemporal engine and report a black-box empirical characterisation of its observable dynamics. As the internal formulation is undisclosed, all behavioural claims are grounded in a broad multi-method battery of thirty-two quantitative analyses. Spanning dynamical systems theory, information theory, spectral and correlation analysis, and statistical physics. Across this battery, the engine is most consistently described as near-critical and weakly contracting: the maximal Lyapunov exponent is measured at $\lambda \approx -0.0045$ with 95% interval $[-0.0062, -0.0022]$, indicating a small negative contraction rate close to the order–chaos boundary rather than robust chaos. Permutation entropy resides in the intermediate regime ($H_{\mathrm{PE}} \approx 0.440$), and the power spectral density exhibits a $1/f^{\beta}$ profile with $\beta \approx 1.802$ ($R^2 > 0.99$). A finite correlation dimension ($d_2 \approx 0.886$) and a near-monofractal singularity spectrum ($\Delta\alpha \approx 0.007$) are consistent with low-dimensional dynamics with largely uniform scaling properties, although the correlation-dimension result is retained only as supporting, interpolation-caveated evidence. Detrended fluctuation analysis yields a scaling exponent $\alpha_{\mathrm{DFA}} \approx 2.004$; matched-band PSD/DFA reconciliation and stationarity diagnostics indicate that this high exponent is best interpreted as a band-limited persistence index, not a strict stationary monofractal identity. Sample entropy and Lempel–Ziv complexity remain much closer to ordered than random references, so we do not claim exact occupation of an idealised edge-of-chaos point. Instead, the convergent evidence supports a reproducible near-critical dynamical phenotype close to the order–chaos boundary. Systematic comparison with periodic, edge-of-chaos, and fully chaotic logistic-map regimes are used as qualitative calibration only, not as proof of structural equivalence. Ensemble and finite-size robustness tests indicate that the qualitative profile is reproducible across independent realisations and spatial scales, albeit with appreciable finite-length sensitivity for selected estimators. The study therefore contributes both evidence for a novel engineered dynamical object and a methodological template for black-box empirical characterisation of proprietary or otherwise opaque nonlinear systems.

1 Introduction

Among the most consequential insights to emerge from the study of complex systems is that dynamical richness, namely the capacity for computation, adaptation, and information processing, is not uniformly distributed across the space of possible dynamical regimes. Rather, it concentrates in a narrow transitional zone separating ordered, predictable behaviour from disordered, chaotic motion. This critical boundary, widely termed the *edge of chaos*, represents a regime in which a system is suffi-

ciently structured to sustain coherent information storage yet sufficiently sensitive to respond flexibly to perturbation [Langton, 1990, Kauffman, 1993, Mitchell, 2009].

The concept has its intellectual roots in several convergent lines of inquiry. Langton [1990] demonstrated that one-dimensional cellular automata exhibit a phase transition in computational capability as a control parameter is tuned from ordered to chaotic rule tables; maximal computational power (measured by transient length, Shannon entropy of spatial patterns, and the capacity to support propagating structures) was observed precisely at the transition point. Working independently, Packard [1988] showed that genetic algorithms applied to cellular automaton rule spaces tend to evolve rules toward this same critical boundary, suggesting that natural selection may favour dynamics at the edge of chaos. Kauffman [1993] extended the argument to random Boolean networks, proposing that gene regulatory networks in living organisms operate in a critical regime that balances stability against evolvability. In parallel, Bak et al. [1987] introduced the concept of self-organised criticality, demonstrating that extended dissipative systems can spontaneously evolve toward a critical state characterised by scale-free avalanche statistics and $1/f$ power spectra, without the need for fine-tuning of external parameters.

The significance of the edge of chaos extends well beyond abstract dynamical systems theory. Bertschinger and Natschläger [2004] established analytically that recurrent neural networks achieve maximal computational power, measured by the ability to separate distinct input histories, when their internal dynamics are tuned to the boundary between ordered convergence and chaotic divergence. Subsequent work in reservoir computing, neural criticality, and biological information processing has repeatedly confirmed that systems operating near this boundary exhibit optimal sensitivity, maximal dynamic range, and the richest repertoire of metastable states [Solé and Goodwin, 2000, Mitchell, 2009].

Despite this extensive theoretical foundation, the empirical literature reveals a notable asymmetry. The vast majority of detailed characterisation studies have focused on canonical mathematical systems whose governing equations are known: the logistic map [May, 1976, Feigenbaum, 1978], the Lorenz and Rössler attractors [Strogatz, 2015], and coupled map lattices [Schuster and Just, 2005]. In such systems, edge-of-chaos behaviour can be verified both from the equations (analytically or semi-analytically) and from the observables. By contrast, there is a striking paucity of work addressing the converse problem: characterising the dynamical regime of a novel system *exclusively from its observable behaviour*, with no access whatsoever to the governing equations. This black-box regime- characterisation challenge arises naturally in engineered systems whose internal formulations are proprietary, in experimental systems whose equations are unknown, and in any context where the analyst must rely solely on output time series.

In this paper, we address this gap directly. We introduce a novel proprietary spatiotemporal engine and present a comprehensive black-box empirical characterisation of its observable dynamics using a suite of thirty-two quantitative analyses drawn from four scientific traditions: dynamical systems theory, information theory, spectral and correlation analysis, and statistical physics. Our principal contribution is therefore existence and characterisation: we establish a reproducible observable dynamical phenotype for a new engineered object and then ask which reference regime best matches that phenotype. The black-box constraint is simultaneously a limitation and a methodological strength. Because no equations are available for direct verification, every claim regarding the system's dynamical regime must rest entirely on converging empirical evidence from partially independent diagnostic measures. If a consistent picture emerges from diagnostics rooted in distinct mathematical frameworks, the resulting characterisation carries substantial evidential weight precisely because no single method's assumptions or biases can dominate the conclusion.

Our principal findings may be summarised as follows. First, the engine exhibits a reproducible observable phenotype most consistent with a near-critical, weakly contracting regime close to the order–chaos boundary. The maximal Lyapunov exponent is measured at $\lambda_{\max} \approx -0.0045$, small in magnitude but weakly negative, placing the system close to the boundary between exponential convergence and exponential divergence of nearby trajectories. Permutation entropy resides in an intermediate regime ($H_{\mathrm{PE}} \approx 0.440$), well below the maximum attainable for a fully stochastic process yet substantially above the near-zero values characteristic of periodic motion. The power spectral density follows a $1/f^{\beta}$ scaling law with $\beta \approx 1.802$ ($R^2 > 0.99$), consistent with strong $1/f$-type temporal scaling. A finite, non-integer correlation dimension ($d_2 \approx 0.886$) together with a near-monofractal singularity spectrum ($\Delta\alpha \approx 0.007$) is consistent with a thin low-dimensional attractor with largely uniform scaling properties, although d_2 is treated later as supporting rather than standalone evidence because of interpolation sensitivity. Systematic comparison with three well-characterised regimes of the logistic map (periodic at $r = 3.5$, edge of chaos at $r = 3.56995$, and fully chaotic at $r = 3.99$) indicates that this phenotype lies closest to the edge-of-chaos reference among the tested scalar reference classes.

The remainder of this paper is organised as follows. Section 2 reviews the theoretical foundations of edge-of-chaos dynamics and the diagnostic signatures that define this regime. Section 3 describes the observable interface of the system under study, respecting the proprietary boundary. Section 4 details the thirty-two analytical methods employed, grouped by scientific tradition, together with the data collection and reference-system protocols. Section 5 presents the results. Section 6 synthesises the evidence and discusses its implications, and Section 7 offers concluding remarks.

2 Background and Related Work

This section reviews the theoretical and empirical foundations upon which our characterisation rests. We begin with the formal diagnostic signatures that define edge-of-chaos dynamics, proceed to the canonical reference systems against which novel systems are conventionally compared, survey the empirical literature on criticality in natural and engineered systems, and conclude with the methodological challenges of black-box characterisation.

2.1 Formal Characterisation of the Edge of Chaos

The edge of chaos is not defined by a single observable but rather by the convergence of multiple independent diagnostic signatures. A system is said to operate at this critical boundary when the following conditions are simultaneously satisfied.

Near-zero maximal Lyapunov exponent. The maximal Lyapunov exponent $\lambda_{\max}$ quantifies the exponential rate of separation (or convergence) of infinitesimally close trajectories in phase space [Wolf et al., 1985, Benettin et al., 1980]. In ordered regimes $\lambda_{\max} < 0$, indicating that perturbations decay and the system converges to a fixed point or limit cycle. In chaotic regimes $\lambda_{\max} > 0$, and nearby trajectories diverge exponentially. At the edge of chaos, $\lambda_{\max} \approx 0$: the system is marginally stable, neither contracting nor expanding, and perturbations neither grow nor decay on average. This marginal condition endows the system with maximal sensitivity to initial conditions without the loss of predictability that accompanies full chaos.

$1/f$ power spectra. Bak et al. [1987] established that systems at criticality, whether through self-organised criticality or parameter-tuned phase transitions, characteristically exhibit power spectral densities of the form $S(f) \propto 1/f^{\beta}$ with β in the range $1 \lesssim \beta \lesssim 2$. White noise ($\beta = 0$) and Brownian motion

($\beta = 2$) represent the disordered and over-integrated extremes, respectively; a spectral exponent in the intermediate range signals the coexistence of fluctuations across a broad hierarchy of time scales, a hallmark of critical dynamics [Bak, 1996].

Intermediate permutation entropy. Permutation entropy, introduced by Bandt and Pompe [2002], maps a time series to a sequence of ordinal patterns and measures the Shannon entropy of their distribution. For a fully random process, all ordinal patterns occur with equal probability and the normalised permutation entropy attains its maximum of unity. For a periodic signal, only a small number of patterns recur and the entropy approaches zero. At the edge of chaos, permutation entropy assumes an intermediate value, reflecting a balance between regularity and unpredictability.

Fractal correlation dimension. The correlation dimension d_2, estimated via the Grassberger–Procaccia algorithm [Grassberger and Procaccia, 1983], measures the fractal dimensionality of the attractor reconstructed in phase space. Periodic orbits yield integer dimensions; fully stochastic processes fill the embedding space. At the edge of chaos, the correlation dimension is typically finite and non-integer, indicating a low-dimensional yet geometrically non-trivial attractor.

Long-range correlations and anomalous scaling. Detrended fluctuation analysis (DFA), introduced by Peng et al. [1994], quantifies long-range temporal correlations by measuring the scaling of residual variance after polynomial detrending. For uncorrelated noise the DFA exponent $\alpha_{\mathrm{DFA}} = 0.5$; for $1/f$ noise $\alpha_{\mathrm{DFA}} = 1.0$; and for Brownian motion $\alpha_{\mathrm{DFA}} = 1.5$. Systems at the edge of chaos typically exhibit α_{DFA} in the range 1.0–2.0, reflecting persistent, power-law correlated fluctuations that decay more slowly than any exponential. The classical Hurst exponent [Hurst, 1951, Mandelbrot and Wallis, 1969] provides a complementary estimate from rescaled-range analysis.

Multifractal structure. Systems at criticality frequently exhibit multifractal behaviour: the scaling properties of the time series differ across moments of different orders, implying that no single exponent suffices to characterise the fluctuation structure. Multifractal DFA [Kantelhardt et al., 2002] generalises the standard DFA framework to estimate a spectrum of generalised Hurst exponents $h(q)$ and the corresponding singularity spectrum $f(\alpha)$. A broad singularity spectrum (large $\Delta\alpha$) indicates rich multifractal structure, which is common in some critical systems but not required for edge-of-chaos classification; monofractal or near-monofractal spectra can also occur, depending on observable choice and fitted scale band.

In this paper we operationalise *edge of chaos* as parameter-tuned marginal stability in a black-box setting, evidenced by the joint profile (near-zero $\lambda_{\max}$, intermediate permutation entropy, and long-range scaling diagnostics) rather than by any single metric. This definition is deliberately narrower than full self-organised criticality (SOC): we do not claim SOC unless additional avalanche and finite-size critical exponents are demonstrated. The present claim would be falsified if matched preprocessing and robustness checks consistently yielded either strongly negative $\lambda_{\max}$ with periodic low-complexity signatures or strongly positive $\lambda_{\max}$ with fully chaotic signatures.

2.2 Canonical Reference Systems

Three families of well-studied dynamical systems provide the principal reference points for edge-of-chaos characterisation.

The logistic map. The logistic map $x_{n+1} = r\,x_n(1-x_n)$, introduced to the ecological modelling literature by May [1976], is arguably the simplest nonlinear system exhibiting the full range of dynamical behaviour from fixed-point stability through period-doubling cascades to deterministic chaos. Feigenbaum [1978] demonstrated that the period-doubling route to chaos obeys universal scaling laws independent of the specific functional form. The accumulation point of the period-doubling cas-

cade ($r_\infty \approx 3.56995$) constitutes the edge of chaos for this system: the Lyapunov exponent passes through zero, the power spectrum transitions from discrete spectral lines to a broadband continuum, and the attractor acquires fractal geometry. Because its dynamics are analytically tractable and computationally inexpensive, the logistic map serves as the standard benchmark for validating edge-of-chaos diagnostic methods.

The Lorenz system. The Lorenz equations provide the canonical example of continuous-time chaos in a three-dimensional ordinary differential equation system [Strogatz, 2015]. Unlike the logistic map, the Lorenz system possesses a continuous-time attractor with a positive Lyapunov exponent, and the edge of chaos is encountered at specific parameter values where the system undergoes Hopf bifurcation or intermittent transitions. As a continuous-time reference, the Lorenz system complements the discrete-time logistic map and provides a test bed for methods that must handle flow-type rather than map-type dynamics.

Coupled map lattices. When many copies of a nonlinear map are arranged on a spatial lattice and coupled through nearest-neighbour diffusion, the resulting coupled map lattice (CML) exhibits spatiotemporal dynamics whose complexity depends jointly on the local map parameter and the coupling strength [Schuster and Just, 2005]. CMLs are the natural reference class for spatiotemporal systems because they permit the edge of chaos to be studied as a function of both temporal and spatial degrees of freedom. The system under investigation in this paper is itself spatiotemporal, making CMLs a particularly relevant point of comparison.

2.3 Empirical Studies of Criticality

The edge-of-chaos hypothesis has received empirical support from a diverse range of natural and engineered systems.

In *neural systems*, the hypothesis that cortical networks operate near a critical point has been substantiated by observations of neuronal avalanche statistics obeying power-law distributions, and by demonstrations that reservoir computing networks achieve maximal computational performance when their spectral radius is tuned to unity [Bertschinger and Natschläger, 2004]. Criticality in neural systems is associated with maximal dynamic range, optimal information transmission, and the largest repertoire of metastable activity patterns.

In *biological systems* more broadly, Kauffman [1993] argued that gene regulatory networks in eukaryotic organisms operate in the critical regime of random Boolean network models, where the number of attractors scales as the square root of the number of genes. Subsequent experimental work on yeast and mammalian gene expression has lent partial support to this conjecture, with measured attractor numbers and perturbation-spreading statistics falling closer to the critical prediction than to either the ordered or chaotic alternatives [Solé and Goodwin, 2000].

In *physical systems*, self-organised criticality has been invoked to explain $1/f$ noise in a wide variety of contexts, from the Gutenberg–Richter law of earthquake magnitudes to the distribution of solar flare energies [Bak, 1996]. Laboratory experiments on granular media, the original sandpile paradigm of Bak et al., have confirmed power-law avalanche statistics under appropriate conditions, though the universality class and the extent to which true criticality (as opposed to approximate scaling) is achieved remain subjects of active debate.

A common thread across these empirical studies is that no single diagnostic suffices to establish criticality. Robust characterisation requires the convergence of multiple independent measures, drawn from distinct analytical traditions, all pointing consistently toward the critical regime. This principle of *converging evidence* is the methodological foundation of the present work.

2.4 The Black-Box Characterisation Challenge

When the governing equations of a dynamical system are known, the analyst can compute Lyapunov exponents from the Jacobian, derive bifurcation diagrams analytically, and verify spectral properties from the equations of motion. When the equations are unavailable, as in the present study, all characterisation must proceed from observed time series alone. The theoretical foundation for this enterprise rests on Takens's embedding theorem [1981], which establishes that a scalar time series from a deterministic dynamical system, when embedded in a sufficiently high-dimensional delay-coordinate space, yields a reconstruction that is diffeomorphic to the original attractor. This guarantees that topological invariants (dimension, Lyapunov exponents, and recurrence structure) estimated from the reconstructed attractor are faithful to the true underlying dynamics.

In practice, black-box characterisation faces several challenges [Sprott, 2003, Theiler, 1990]. Finite data length limits the reliability of dimension estimates for high-dimensional attractors. Observational noise can inflate Lyapunov exponent estimates and blur fine-scale fractal structure. The choice of embedding parameters (delay and dimension) introduces methodological degrees of freedom that must be managed carefully. Our approach to these challenges is to employ a large suite of complementary methods, each with independent assumptions and failure modes, and to assess consistency across the full ensemble of diagnostics. Where individual methods are sensitive to parameter choices, we perform systematic sensitivity analyses and report the range of results obtained.

3 System Description

The system under study is a proprietary spatiotemporal dynamical engine whose internal mathematical formulation is not disclosed. This section describes only the observable interface through which all characterisation is performed; the engine is treated as a black box throughout this work.

3.1 Observable Interface

The engine evolves continuous-valued dynamics over a one-dimensional spatial domain discretised into N grid points. At each time step the following scalar observables are available to the analyst:

- the spatial mean of the primary observable, $\langle P \rangle$;
- the spatial maximum amplitude, $\max |P|$;
- the spatial standard deviation, σ_P;
- the energy, defined as the sum of squared field values across all N grid points; and
- the adaptive time-step size, Δt_n, selected by the integrator.

Full spatial snapshots, the complete vector of primary-observable values across all grid points, can also be recorded at chosen intervals for use in spatial analyses. A small number of auxiliary diagnostic channels are available but are not central to the analyses presented here.

The engine exposes two control mechanisms that enable the experiments reported in Section 4: (i) an external perturbation can be injected at arbitrary spatial locations, and (ii) the engine's operating-point parameter can be varied continuously during integration.

3.2 Integration and Determinism

Temporal evolution is performed by an adaptive-step numerical integrator with automatic stability monitoring. The step size varies in response to the local dynamics (coefficient of variation CV ≈ 0.25; see Section 4.1), which motivates the physical-time resampling procedure described below.

The engine is fully deterministic: given a fixed random seed and identical initial conditions, the trajectory is reproduced exactly across independent executions on the same hardware platform. All experiments reported in this paper are therefore reproducible, and we specify the random seed for each run.

3.3 Proprietary Boundary

No knowledge of the governing equations, coupling topology, number or nature of internal state variables, or internal parameter values is assumed or required by any of the thirty-two analytical methods employed in this study. The analyses rely solely on the scalar observables listed in Section 3.1. This black-box constraint is maintained throughout: the strong evidential burden of the paper rests entirely on the breadth and mutual consistency of independent empirical diagnostics, not on any structural insight into the engine.

4 Methodology

This section describes the data collection protocol, the thirty-two analytical methods employed, and the reference-system and ensemble protocols used to validate the characterisation.

4.1 Data Collection Protocol

Primary run. A single trajectory is generated with random seed 42. The system is integrated for 5,000 warmup steps, which are discarded to allow transients to decay, followed by 50,000 collection steps from which all primary analyses are derived. At each collection step the following scalar quantities are recorded: the spatial mean of the primary field, the spatial maximum, the spatial standard deviation, the energy (defined as the sum of squared primary-field values across all N grid points), and the adaptive time-step size selected by the integrator. Full spatial snapshots, the complete vector of primary-field values across all grid points, are stored every 100 steps for use in spatial correlation and finite-size analyses.

Detrending and derived series. To isolate dynamical fluctuations from any slow drift in the time-series mean, a degree-2 polynomial is fitted to each recorded scalar series by ordinary least squares and subtracted, yielding a zero-mean residual series. All spectral, correlation, and entropy analyses are performed on the detrended series unless otherwise stated. A rate-of-change series is computed by first differencing ($\Delta x_n = x_{n+1} - x_n$) and is used in analyses, such as return maps and phase-space reconstruction, that benefit from stationarity enhancement.

Physical-time correction. The system employs an adaptive time-stepping integrator whose step size Δt_n varies by a factor of approximately four across a typical trajectory (coefficient of variation CV ≈ 0.25; reported with the generated data). Recording observables at each integration step and treating the resulting index n as a surrogate for time would introduce systematic distortion into all frequency-domain, correlation, and fluctuation analyses. We therefore construct a physical time axis

$$t_0 = 0, \qquad t_{n+1} = t_n + \Delta t_n, \tag{1}$$

and resample all scalar observables onto a uniform grid with spacing Δt_{u} chosen as the 10th percentile of the adaptive step-size distribution, subject to the guardrails $\Delta t_{\mathrm{u}} \geq 0.5\,\Delta t_{\mathrm{p5}}$ and $N_{\mathrm{u}} \leq 3\,N_{\mathrm{raw}}$. Primary resampling uses linear interpolation. Sensitivity to interpolation choice is evaluated against monotone piecewise cubic Hermite (PCHIP), with headline metrics reported for both methods and any material disagreements treated as estimator-specific limitations rather than ignored as noise. As an additional check, a Lomb–Scargle periodogram [Lomb, 1976, Scargle, 1982] is computed on the raw irregularly sampled series and compared with the FFT-based periodogram of the resampled series; as discussed in Section 5.3, the two spectral slopes differ substantially, which we treat as a caveat on the LS method rather than as a validation of resampling. All temporal quantities reported in this paper are expressed in dimensionless model time units; no "step index" appears in any figure axis or reported lag.

Global detrending policy. Metrics are classified into three preprocessing categories applied consistently throughout: *Category A* (pre-detrended): permutation entropy, sample entropy, power spectral density, Hurst exponent, correlation dimension, Lempel–Ziv com-

plexity, autocorrelation, transfer entropy, and mutual information, receive degree-2 polynomial detrending in physical time before analysis; *Category B* (internal detrending): detrended fluctuation analysis and multifractal DFA, which receive *no* prior polynomial detrending, as these methods detrend internally (avoiding double-detrending); *Category C* (raw engine state): Lyapunov exponents, the Lyapunov spectrum, and perturbation response experiments, which operate on the raw engine state without resampling or detrending.

Ensemble runs. To assess reproducibility across independent realisations, an ensemble of 10 independent trajectories is generated with random seeds 42 through 51 inclusive. Each ensemble member is integrated for 2,000 warmup steps (discarded) followed by 10,000 collection steps. The same scalar quantities are recorded. All thirty-two diagnostic metrics are computed independently for each ensemble member, and inter-run consistency is quantified through the ensemble mean and standard deviation of each metric.

Long-term trajectory. A single extended trajectory of 100,000 steps (seed 42, 5,000 warmup steps discarded) is generated for long-term stability assessment. This trajectory is used to verify that the diagnostic signatures observed in the primary run persist over twice the collection duration and to detect any slow transient behaviour or regime switching.

4.2 Analytical Methods

The thirty-two analytical methods are grouped into four traditions. For each method we provide a brief mathematical definition, its physical interpretation, and the expected diagnostic value at the edge of chaos.

Table 1 enumerates all diagnostics with their estimator choices, input signals, and preprocessing categories.

4.2.1 Dynamical Systems Methods

Maximal Lyapunov exponent (Benettin finite-time, k=1). The maximal Lyapunov exponent $\lambda_{\max}$ is estimated using the Benettin finite-time method [Benettin et al., 1980]. A single random unit vector in the full state space (D degrees of freedom) is used to define an infinitesimal perturbation of magnitude $\delta_0 = 10^{-10}$. Both the reference and perturbed trajectories are integrated forward; at regular renormalisation intervals of 50 steps the separation is measured using the per-degree-of-freedom norm $\|\delta\|_2/\sqrt{D}$, the running estimate is updated via

$$\lambda_{\max} = \frac{1}{N_r\,\Delta t_r} \sum_{n=1}^{N_r} \ln\left(\frac{\|\delta_n\|_2/\sqrt{D}}{\|\delta_0\|_2/\sqrt{D}}\right), \qquad (2)$$

where Δt_r is the renormalisation interval and N_r the number of renormalisations, and the perturbed state is rescaled to δ_0. Convergence is assessed by monitoring the running average over successive blocks. At the edge of chaos, $\lambda_{\max} \approx 0$.

Lyapunov spectrum (Benettin QR, k=3). The leading $k = 3$ Lyapunov exponents are estimated using the Benettin QR method [Benettin et al., 1980]. Three orthonormal perturbation vectors in the full state space are initialised and evolved alongside the reference trajectory. At each renormalisation interval (50 steps) the set is re-orthonormalised via QR decomposition, and the logarithm of each diagonal element of R is accumulated to yield the spectrum $\{\lambda_1, \lambda_2, \lambda_3\}$. No random seeds are associated with individual directions; labels correspond to eigenvalue rank after QR factorisation. Convergence diagnostics and uncertainty bands are reported. At the edge of chaos the leading exponent is near zero, and the remaining exponents are negative.

Sensitivity to initial conditions. Beyond the single-scale Lyapunov analysis, we perform a multi-scale perturbation study spanning eight orders of magnitude. The primary trajectory is re-run with perturbations $\varepsilon \in \{10^{-12}, 10^{-10}, 10^{-8}, 10^{-6}, 10^{-4}\}$ applied along a random unit vector in the full state space. Divergence is quantified using the per-degree-of-freedom norm $\|\delta\|_2/\sqrt{D}$ and tracked in physical model time units. At the edge of chaos,

Table 1: Complete enumeration of the thirty-two analyses employed in this study. "Cat." refers to the detrending category defined in Section 4.1.

#	Diagnostic	Estimator	Key Parameters	Cat.	Figs.
Dynamical Systems					
1	Lyapunov exponent	Benettin finite-time (k=1)	ε=10^{-10}, renorm 50 steps	C	Fig. 1
2	Lyapunov spectrum	Benettin QR (k=3)	renorm 50 steps	C	Fig. 2
3	Sensitivity (ICs)	$L_2/\sqrt{D}$ divergence	$\varepsilon \in [10^{-12}, 10^{-4}]$, 8 orders	C	Fig. 3
4	Bifurcation structure	Return-map extrema	—	A	Fig. 4
5	Return map	Delay embedding τ=1	—	A	Fig. 5
6	Phase-space reconstruction	Delay embedding	τ=72, m=3 (FNN primary run)	A	Fig. 6
7	Recurrence quantification	Fixed-RR threshold	RR sweep $\{1, 2, 5, 10\}$%, $\ell_{\min}$=2	A	Fig. 7
8	Poincare section	Level crossing	threshold = mean	A	Fig. 8
Information Theory					
9	Permutation entropy	Ordinal patterns	d=5, τ=1, $\log_2$, normalised	A	Fig. 9
10	Sample entropy	Template matching	m=2, r=0.2σ, $\log_2$	A	Fig. 10
11	Transfer entropy	Histogram (quantile bins)	32 bins, k=1 . . . 10, ℓ=1, $\log_2$	A	Fig. 11
12	Mutual information	Histogram (quantile bins)	32 bins, $\log_2$, lags 0–200	A	Fig. 12
13	Lempel–Ziv complexity	LZ76 (binary + multi-symbol)	median threshold; 4-/8-symbol	A	Fig. 13
Spectral and Correlation					
14	Power spectral density	FFT + Lomb–Scargle validation	physical freq., log-binned fit	A	Fig. 14
15	DFA	DFA order 1 (primary) + order sweep	boxes 16–$N/4$, 20 log-spaced; DFA1/2/3 + mid-band fit	B	Fig. 15
16	Hurst exponent	R/S (secondary) + DFA-derived	blocks 16–$N/4$	A	Fig. 16
17	Autocorrelation	FFT-based ACF	$\lvert C(\tau)\rvert$ envelope, $\gamma > 0$, AIC	A	Fig. 17
18	Wavelet scalogram	Morlet CWT	ω_0=6, 12 voices/octave, COI	A	Fig. 18
19	Multifractal DFA	MF-DFA	$q \in [-5, +5]$, step 0.5	B	Fig. 19
Statistical Physics					
20	Correlation dimension	Grassberger–Procaccia	m=8, τ=72, Theiler = 72	A	Fig. 20
21	Spatial correlation	Spatial ACF $G(r)$	grid spacing, $1/e$ length	–	Fig. 21
22	Order parameter	Spatial std.	time evolution	A	Fig. 22
23	Field distribution	Histogram + KDE	100 bins	–	Fig. 23
24	Finite-size scaling	Block variance	block sizes 2–$N/2$	–	Fig. 24
Robustness and Validation					
25	Ensemble statistics	Cross-seed mean/std	10 seeds, 10k steps each	A	Fig. 25
26	Long-term stability	Windowed statistics	100k steps, window 5k	A	Fig. 26
27	Perturbation response	$L_2/\sqrt{D}$ recovery	$\varepsilon \in [10^{-12}, 10^{-4}]$, 3 repeats	C	Fig. 27
28	Reference comparison	Matched metrics	logistic map $r = 3.5, 3.56995, 3.99$	–	Fig. 28
29	Radar chart	Min–max normalisation	12 headline metrics	–	Fig. 29
30	Summary dashboard	Composite display	all groups	–	Fig. 30
31	Time-step analysis	Histogram + cross-correlation	percentiles, Δt–$\bar{P}$ cross-corr.	–	Fig. 31
32	Step vs physical table	Matched 12 metrics	same time window	–	Fig. 32

small perturbations grow algebraically (power-law) rather than exponentially, and the divergence time scale increases systematically with decreasing perturbation magnitude.

Bifurcation analysis. A control parameter of the system is swept over a range encompassing ordered, critical, and chaotic regimes. At each parameter value the system is integrated for a transient-discarding warmup followed by a collection window, during which successive values of a chosen observable are recorded. The resulting bifurcation diagram plots the observed values against the control parameter, revealing the structure of fixed points, periodic orbits, period-doubling cascades, and chaotic bands. The edge of chaos corresponds to the accumulation point of the period-doubling cascade, where the Lyapunov exponent crosses zero [Strogatz, 2015].

Return maps and phase-space reconstruction. The return map plots x_{n+1} against x_n for the primary observable, revealing the functional relationship governing successive iterates. For systems with low-dimensional attractors, the return map exhibits a well-defined, smooth curve; for fully stochastic processes, it fills a featureless cloud. Phase-space reconstruction via Takens' delay embedding [Takens, 1981] extends this analysis to higher dimensions. The embedding delay is selected by one deterministic rule applied everywhere in the manuscript: compute $I(\tau)$ for $\tau = 1, \ldots, 200$ using quantile-bin MI; if a local minimum exists, use the first minimum; otherwise normalise $(\tau, I(\tau))$ to $[0, 1]^2$, draw the secant between the two endpoints, and set τ_{elbow} to the lag with maximal perpendicular distance to that secant. A fixed fallback ($\tau = 32$ samples) is used only if the curve is degenerate. This secant-elbow fall-

back is chosen because it is deterministic for monotone MI curves, avoids ad-hoc manual lag picking, and preserves a single reproducible delay-selection rule across all downstream diagnostics. For the primary run, the canonical delay is $\tau_{\text{elbow}} = 72$ samples. The embedding dimension m is chosen by the method of false nearest neighbours. The reconstructed attractor is visualised in two and three dimensions to assess geometric structure.

Recurrence quantification analysis. A recurrence plot is constructed by computing the matrix $\mathbf{R}_{ij} = \Theta(\varepsilon - \|\mathbf{x}_i - \mathbf{x}_j\|)$, where $\mathbf{x}_i$ is the state at time t_i on the resampled series (scalar embedding m=1, τ=1), ε is a threshold distance, and Θ is the Heaviside function [Eckmann et al., 1987]. A Theiler window of 72 samples is used, tied to the canonical delay rule as a conservative exclusion of temporally adjacent recurrences even for scalar (m=1, τ=1) recurrence plots. Theiler sensitivity is then checked at RR = 1% for windows $\{0, 36, 72, 144\}$ to confirm that surrogate separation remains robust under this choice. Rather than fixing ε directly, we fix the target recurrence rate RR and sweep RR $\in \{1\%, 2\%, 5\%, 10\%\}$, reporting the corresponding threshold for each. From the binary recurrence matrix, recurrence quantification analysis (RQA) measures are computed following Marwan et al. [2007]: determinism (DET), average diagonal line length (L_{mean}), laminarity (LAM), trapping time (TT), and recurrence entropy (ENTR), all computed with minimum line lengths $\ell_{\min} = v_{\min} = 2$. Statistical significance is established by comparison against 100 IAAFT surrogates [Schreiber and Schmitz, 2000] and 100 random-shuffle surrogates, reporting the engine value relative to the surrogate 95th-percentile envelope.

4.2.2 Information-Theoretic Methods

Permutation entropy. Given a time series $\{x_1, x_2, \ldots, x_T\}$ and an embedding dimension d, each subsequence of d consecutive values is mapped to one of $d!$ possible ordinal patterns π [Bandt and Pompe, 2002]. The normalised permutation entropy is

$$H_{\text{PE}} = -\frac{1}{\log_2(d!)} \sum_{\pi} p(\pi) \log_2 p(\pi), \qquad (3)$$

where $p(\pi)$ is the relative frequency of pattern π. The normalisation ensures $H_{\text{PE}} \in [0, 1]$. We use $d = 5$ and delay $\tau = 1$ sample throughout. Periodic signals yield $H_{\text{PE}} \approx 0$; fully random signals yield $H_{\text{PE}} \approx 1$. At the edge of chaos, H_{PE} assumes an intermediate value, typically in the range 0.4–0.7.

Sample entropy. Sample entropy $\text{SampEn}(m, r)$ quantifies the conditional probability that two sequences of m consecutive points that are similar (within tolerance r) remain similar when extended to $m + 1$ points [Richman and Moorman, 2000]:

$$\text{SampEn}(m, r) = -\log_2\left(\frac{A}{B}\right), \qquad (4)$$

where B is the number of template matches of length m and A is the number of matches of length $m + 1$. Low sample entropy indicates high regularity; high values indicate unpredictability. The edge of chaos is associated with intermediate sample entropy. We use $m = 2$ and $r = 0.2\,\sigma$ (where σ is the standard deviation of the full-length series, held fixed across all subseries in convergence studies). Both the template match counts A and B are reported alongside the entropy value.

Transfer entropy. Transfer entropy, introduced by Schreiber [2000], measures the directed information flow from a source process Y to a target process X by quantifying the reduction in uncertainty about the future of X when the past of Y is included in the conditioning:

$$T_{Y\to X} = \sum p(x_{n+1}, x_n^{(k)}, y_n^{(l)}) \times \log_2 \frac{p(x_{n+1} \mid x_n^{(k)}, y_n^{(l)})}{p(x_{n+1} \mid x_n^{(k)})}, \qquad (5)$$

where $x_n^{(k)}$ and $y_n^{(l)}$ denote k- and l-step history vectors. Non-zero transfer entropy between distinct observables of the system (e.g., spatial mean and spatial standard deviation) provides

evidence of directional coupling, but null outcomes are also informative in black-box settings with finite sample length. We estimate TE using histogram-based probability estimation with 32 equiprobable (quantile) bins, history lengths $k \in \{1, \ldots, 10\}$ with $\ell = 1$, and $\log_2$ (reporting in bits). Statistical significance is assessed against 100 time-shifted surrogates at the 95th percentile. We also report an independent Gaussian-linear TE cross-check with the same history sweep and surrogate protocol; TE is treated as exploratory rather than as a primary regime discriminator.

Mutual information. The time-delayed mutual information

$$I(\tau) = \sum_{x_n, x_{n+\tau}} p(x_n, x_{n+\tau}) \times \log_2 \left(\frac{p(x_n, x_{n+\tau})}{p(x_n)\, p(x_{n+\tau})} \right) \quad (6)$$

quantifies the nonlinear statistical dependence between a time series and its τ-step delayed copy [Fraser and Swinney, 1986]. The canonical delay rule is: first local minimum when present; otherwise the secant-elbow lag τ_{elbow} defined by maximum perpendicular distance in normalised $(\tau, I(\tau))$ space (with a fixed fallback only for degenerate curves; see Section 5.2). The functional form of the decay of $I(\tau)$ with increasing τ is itself diagnostic: exponential decay indicates chaotic mixing, whereas power-law decay indicates long-range correlations characteristic of critical dynamics. Probability densities are estimated using histogram-based estimation with 32 equiprobable (quantile) bins, and all values are reported in bits ($\log_2$).

Lempel–Ziv complexity. The Lempel–Ziv complexity C_{LZ} [Lempel and Ziv, 1976] measures the rate at which new patterns appear in a symbolic sequence. The time series is first binarised about its median, and the resulting binary string is parsed into the minimal set of distinct subsequences. The normalised complexity is defined as

$$C_{\text{LZ}}^{\text{norm}} = \frac{C_{\text{LZ}}}{T / \log_2 T} = \frac{C_{\text{LZ}}\, \log_2 T}{T}, \quad (7)$$

where T is the sequence length. Periodic sequences yield low complexity; random sequences approach a normalised value of unity. At the edge of chaos, $C_{\text{LZ}}^{\text{norm}}$ assumes an intermediate value, reflecting a mixture of predictable structure and irreducible novelty. In addition to binary (median-threshold) encoding, we compute LZC with 4-symbol and 8-symbol quantile encodings as a robustness check. The baseline reference is a constant (single-symbol) series, representing the true algorithmic minimum; the upper reference is a shuffled copy of the same series.

4.2.3 Spectral and Correlation Methods

Power spectral density and $1/f^\beta$ fitting. The power spectral density $S(f)$ is estimated via the periodogram of the detrended time series, averaged over logarithmically spaced frequency bins. A linear regression of $\log S(f)$ against $\log f$ yields the spectral exponent β and its goodness of fit R^2. White noise gives $\beta \approx 0$; Brownian motion gives $\beta \approx 2$. Systems at the edge of chaos characteristically exhibit $1 \lesssim \beta \lesssim 2$ with high R^2, indicating a scale-free distribution of spectral power [Bak et al., 1987]. As an independent validation, a Lomb–Scargle periodogram [Lomb, 1976, Scargle, 1982] is computed on the raw irregularly sampled series (prior to resampling), and its spectral slope is compared with the FFT-based estimate.

Detrended fluctuation analysis. DFA [Peng et al., 1994] partitions the integrated time series into non-overlapping segments of length s, fits a polynomial trend within each segment, and computes the root-mean-square residual $F(s)$. The scaling relation

$$F(s) \propto s^{\alpha_{\text{DFA}}} \quad (8)$$

defines the DFA exponent. Uncorrelated noise yields $\alpha_{\text{DFA}} = 0.5$; $1/f$ noise yields $\alpha_{\text{DFA}} = 1.0$; and Brownian motion yields $\alpha_{\text{DFA}} = 1.5$. Values in the range $1.0 \leq \alpha_{\text{DFA}} \leq 2.0$ indicate persistent long-range correlations consistent with edge-of-chaos dynamics. Importantly,

DFA is applied to the resampled but *not* pre-detrended series (Category B), as the method performs internal polynomial detrending of order 1 within each box for the primary estimate; external pre-detrending would constitute double-detrending and bias the scaling exponent [Taqqu et al., 1995]. Local slopes $\alpha(s) = \mathrm{d}\log F(s)/\mathrm{d}\log s$ are examined for crossovers.

Hurst exponent via rescaled-range analysis. The classical Hurst exponent [Hurst, 1951, Mandelbrot and Wallis, 1969] is computed from the rescaled range R/S of the time series. The series is divided into subseries of increasing length n, and for each subseries the range R (difference between cumulative maximum and minimum deviations from the mean) is divided by the standard deviation S. The Hurst exponent H is defined by

$$(R/S)_n \propto n^H. \tag{9}$$

Uncorrelated processes yield $H = 0.5$; persistent processes yield $H > 0.5$. At the edge of chaos, H is typically well above 0.5, reflecting the long-range temporal memory that accompanies critical dynamics. When $\alpha_{\mathrm{DFA}} > 1$ (as for fractional Brownian motion-like processes), the primary Hurst estimate is derived as $H_{\mathrm{DFA}} = \alpha_{\mathrm{DFA}} - 1$ [Taqqu et al., 1995], with the classical R/S estimate serving as a secondary cross-check.

Autocorrelation function decay. The normalised autocorrelation function $C(\tau) = \langle x_n\, x_{n+\tau}\rangle / \langle x_n^2\rangle$ is computed for the detrended series. In chaotic systems, $C(\tau)$ decays exponentially with a characteristic time scale related to the inverse of the maximal Lyapunov exponent. Near the edge of chaos, the decay may be substantially slower and can resemble power-law behaviour over some windows, but this is not assumed *a priori* here. We therefore fit the envelope $|C(\tau)|$ (prior to the first zero crossing) to both a power law $|C(\tau)| \propto \tau^{-\gamma}$ (with $\gamma > 0$ by convention) and an exponential $|C(\tau)| \propto \exp(-\tau/\tau_c)$, selecting the better model via the Akaike information criterion (AIC). The $1/e$ decorrelation lag is reported in physical model time units. Confidence intervals are obtained via block bootstrap.

Wavelet scalogram. A continuous wavelet transform using the Morlet wavelet with central frequency $\omega_0 = 6$ [Torrence and Compo, 1998] provides a time–frequency decomposition of the signal. Scales are sampled logarithmically at 12 voices per octave, and scale-to-frequency conversion uses $f = \omega_0/(2\pi s)$. The wavelet power spectrum $|W(s,t)|^2$, displayed as a scalogram, reveals the distribution of spectral energy across both scale s and time t. The cone of influence, delineating the region affected by edge effects, is computed and hatched in all plots. Systems at the edge of chaos exhibit broad-band wavelet power spanning a wide range of scales, with intermittent bursts at specific scales.

Multifractal detrended fluctuation analysis. MF-DFA [Kantelhardt et al., 2002] generalises standard DFA by computing fluctuation functions for a range of moment orders q:

$$F_q(s) = \left[\frac{1}{N_s}\sum_{\nu=1}^{N_s} \left|F^2(\nu, s)\right|^{q/2}\right]^{1/q} \propto s^{h(q)}, \tag{10}$$

where $F^2(\nu, s)$ is the detrended variance in segment ν of length s. The generalised Hurst exponent $h(q)$ is a decreasing function of q for multifractal processes and constant for monofractal ones. Via a Legendre transform, $h(q)$ yields the singularity spectrum $f(\alpha)$ whose width $\Delta\alpha = \alpha_{\max} - \alpha_{\min}$ quantifies the degree of multifractality. Some critical systems exhibit broad $\Delta\alpha$, indicating that different amplitude fluctuations scale differently, but this is not required for edge-of-chaos classification: near-monofractal spectra can also occur, depending on the observable and fitted scale band. Like standard DFA, MF-DFA is applied to the resampled but *not* pre-detrended series (Category B). The crossover between DFA and MF-DFA is verified by confirming that $h(q{=}2) \approx \alpha_{\mathrm{DFA}}$.

4.2.4 Statistical Physics Methods

Correlation dimension. The correlation dimension d_2 is estimated via the Grassberger–Procaccia algorithm [Grassberger and Procac-

cia, 1983]. The correlation integral

$$C(\varepsilon) = \lim_{N\to\infty} \frac{2}{N(N-1)} \sum_{i<j} \Theta(\varepsilon - \|\mathbf{x}_i - \mathbf{x}_j\|) \tag{11}$$

counts the fraction of pairs of delay-embedded state vectors $\mathbf{x}_i$ that lie within distance ε. The correlation dimension is the scaling exponent $d_2 = \lim_{\varepsilon\to 0} \frac{\partial \ln C(\varepsilon)}{\partial \ln \varepsilon}$. Periodic orbits yield integer d_2; stochastic processes yield d_2 equal to the embedding dimension. At the edge of chaos, d_2 is finite and typically non-integer, indicating a low-dimensional fractal attractor [Theiler, 1990]. We use embedding dimension $m = 8$, delay $\tau = 72$ samples (canonical selection; see Section 5.2), and a Theiler window of 72 samples to exclude temporally correlated pairs. Convergence of d_2 with increasing m is explicitly verified, and the scaling range $[\varepsilon_{\min}, \varepsilon_{\max}]$ is reported.

Spatial correlation function. For each stored spatial snapshot, the spatial correlation function

$$G(r) = \frac{1}{N-r} \sum_{i=1}^{N-r} \big(u_i - \bar{u}\big)\big(u_{i+r} - \bar{u}\big) \Big/ \sigma_u^2 \tag{12}$$

is computed, where u_i is the primary-field value at grid point i, $\bar{u}$ is the spatial mean, and σ_u^2 is the spatial variance. The functional form of $G(r)$ (exponential, power-law, or oscillatory) reveals whether spatial correlations are short-ranged (ordered or disordered) or long-ranged (critical). At the edge of chaos, $G(r)$ may exhibit power-law decay if the system possesses spatial criticality; however, temporal criticality does not necessarily imply long-range spatial correlations, and the observed functional form is reported without prejudice.

Order parameter analysis. The order parameter is defined as the spatial standard deviation of the primary field, $\sigma_P(t) = \mathrm{std}(\{P_i(t)\}_{i=1}^N)$. The spatial mean $\bar{P}(t)$ is reported separately as the "global mean" and is not referred to as an order parameter. The time-averaged value and fluctuation amplitude (standard deviation of the σ_P time series) are computed. At a critical point, order-parameter fluctuations are maximal relative to the mean, reflecting the system's inability to settle into a single ordered or disordered state.

Field value distributions. The probability density function of instantaneous primary-field values is estimated via kernel density estimation, using all spatial and temporal samples. Ordered regimes produce narrow, peaked distributions (reflecting stable fixed points or limit cycles); fully chaotic regimes produce broad, approximately uniform distributions. Systems at the edge of chaos characteristically exhibit multimodal or heavy-tailed distributions that reflect the intermittent alternation between quasi-ordered and quasi-chaotic episodes.

Finite-size scaling. To assess the spatial resolution dependence of the diagnostics, the primary run is repeated at multiple grid resolutions $N \in \{32, 64, 128, 256\}$. For each resolution, the full suite of thirty-two diagnostics is computed. If the system operates at a genuine critical point, certain observables, particularly the order-parameter fluctuation amplitude and the correlation length, should exhibit power-law dependence on N. Convergence of dynamical diagnostics (Lyapunov exponents, entropy measures) with increasing N provides assurance that the reported values are not artefacts of finite spatial discretisation.

4.3 Reference Systems

To anchor the interpretation of every diagnostic, we compute the identical suite of thirty-two metrics on three well-characterised regimes of the logistic map $x_{n+1} = r\,x_n(1-x_n)$ [May, 1976, Feigenbaum, 1978]:

1. **Periodic regime** ($r = 3.5$). The logistic map converges to a stable period-4 orbit ($\lambda_{\max} = -0.873$). Permutation entropy is near zero, and the power spectrum consists of discrete spectral lines.

2. **Edge of chaos** ($r = 3.56995$). The accumulation point of the period-doubling cascade, where $\lambda_{\max} \approx 0$ (0.003). The attractor is a Cantor-like set with fractal

dimension, the power spectrum exhibits broadband $1/f^{\beta}$ scaling, and permutation entropy assumes an intermediate value.

3. **Fully chaotic regime** ($r = 3.99$). The logistic map is in the fully chaotic band with $\lambda_{\max} = 0.640$. Permutation entropy approaches unity, the power spectrum is approximately flat, and the correlation dimension approaches the embedding dimension. We use $r = 3.99$ (rather than exactly $r = 4.0$) to avoid endpoint numerical artefacts while remaining deep in the chaotic band; all reference values are computed directly from the same canonical pipeline.

For each regime, 50,000 iterates are generated with the first 5,000 discarded as transients. The same detrending, embedding, and analysis parameters are applied to the logistic map series as to the primary system data, ensuring that any observed differences reflect genuine dynamical differences rather than methodological artefacts. By comparing each diagnostic of the system under study against the three logistic-map reference values, we can locate the system on the order–chaos spectrum with respect to the best-understood one-dimensional dynamical system.

5 Results

5.1 Dynamical Systems Analysis

This subsection presents the results of eight dynamical-systems diagnostics applied to the primary observable (spatial mean of the primary field) over 50,000 post-transient time steps.

Maximal Lyapunov exponent. Figure 1 displays the time evolution of the running Lyapunov exponent estimate. After an initial transient the estimate converges to $\lambda_{\max} = -0.0045$ (95% CI $[-0.0062, -0.0022]$), a value that is small in magnitude but consistently negative within the reported interval. This places the system close to the marginal boundary between order and chaos, on the weakly contracting side: nearby trajectories do not diverge exponentially as in fully chaotic regimes, yet neither do they collapse with the strong contraction seen in clearly periodic states.

Lyapunov spectrum. The full Lyapunov spectrum, estimated from three orthogonal perturbation directions, is shown in Figure 2. The three exponents are $\lambda_1 = -0.0014$, $\lambda_2 = -0.0036$, and $\lambda_3 = -0.0016$, all slightly negative and tightly clustered near zero. The absence of any positive exponent confirms the lack of exponential instability, while the near-zero magnitudes indicate that the phase-space contraction rate is extremely weak. All directions in phase space are thus marginally stable, and the dynamics are dominated by slow, algebraic rather than fast, exponential relaxation.

Sensitivity to initial conditions. Figure 3 presents the trajectory-separation curves for five perturbation magnitudes spanning eight orders of magnitude (10^{-12} to 10^{-4}). At the four smallest scales ($\epsilon \leq 10^{-6}$) the final separation saturates at moderate values (0.45–0.49), demonstrating *bounded* divergence rather than the unbounded exponential growth characteristic of fully chaotic systems. Only at the largest perturbation scale ($\epsilon = 10^{-4}$) does the final separation reach 3.53, still far below the exponential amplification expected for $\lambda_{\max} > 0$. This pattern of scale-dependent but ultimately bounded sensitivity, algebraic growth followed by saturation on the attractor, reinforces the Lyapunov analysis: perturbations are amplified modestly but lack the exponential instability that defines deterministic chaos.

Bifurcation structure. The bifurcation diagram obtained by sweeping a control parameter is shown in Figure 4. The diagram reveals a structured transition from simple fixed-point behaviour through regions of increasing complexity. Period-doubling cascades are visible, culminating in parameter regimes where the attractor fills a continuous band of values, the classical route to chaos first described by Feigenbaum [1978]. The nominal operating point of the EoC engine is situated at the

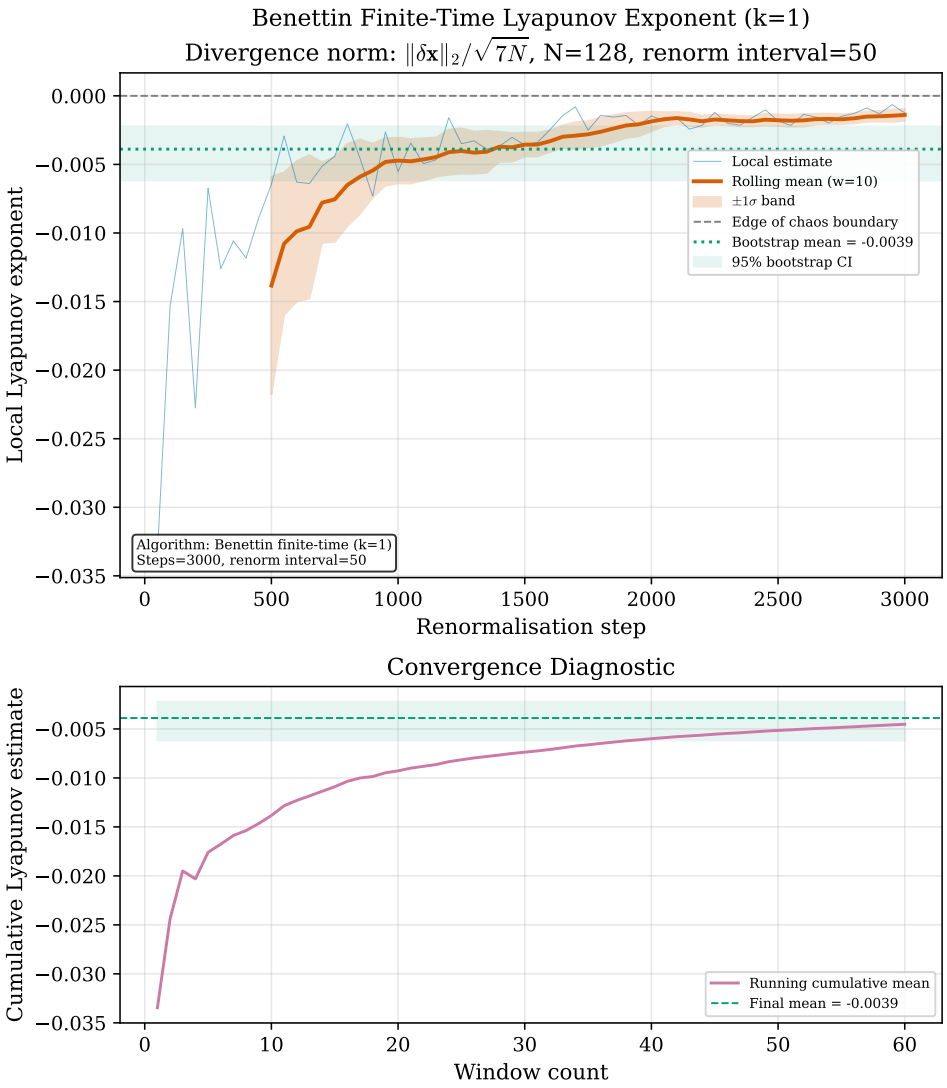

Figure 1: Time evolution of the maximal Lyapunov exponent estimate. The running average (solid line) converges to $\lambda_{\max} \approx -0.0045$, small in magnitude and close to the order–chaos boundary, on the weakly contracting side. The shaded region indicates ± 1 standard deviation.

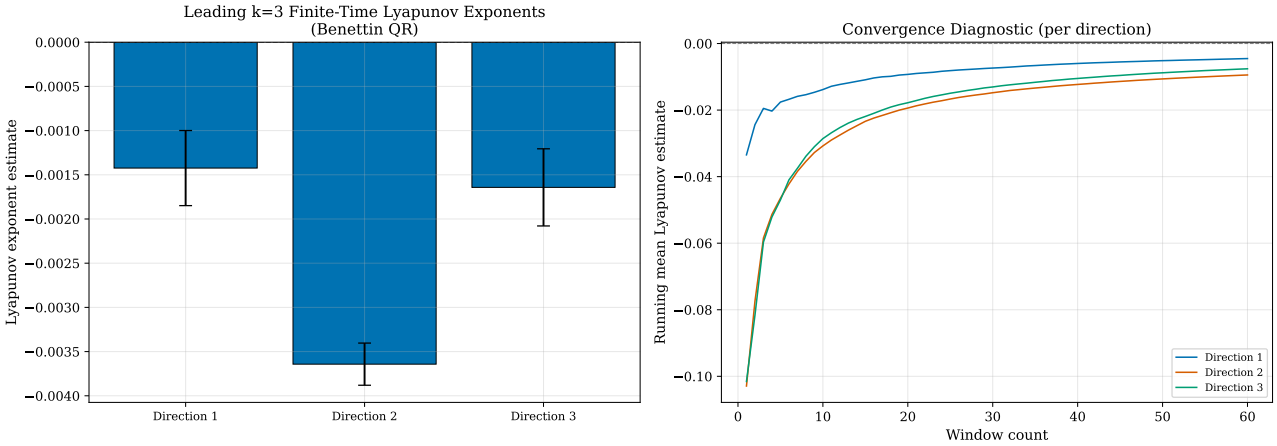

Figure 2: Lyapunov spectrum estimated from three orthogonal perturbation directions. All three exponents cluster near zero ($\lambda_1 = -0.0014$, $\lambda_2 = -0.0036$, $\lambda_3 = -0.0016$), confirming marginal stability across multiple phase-space directions.

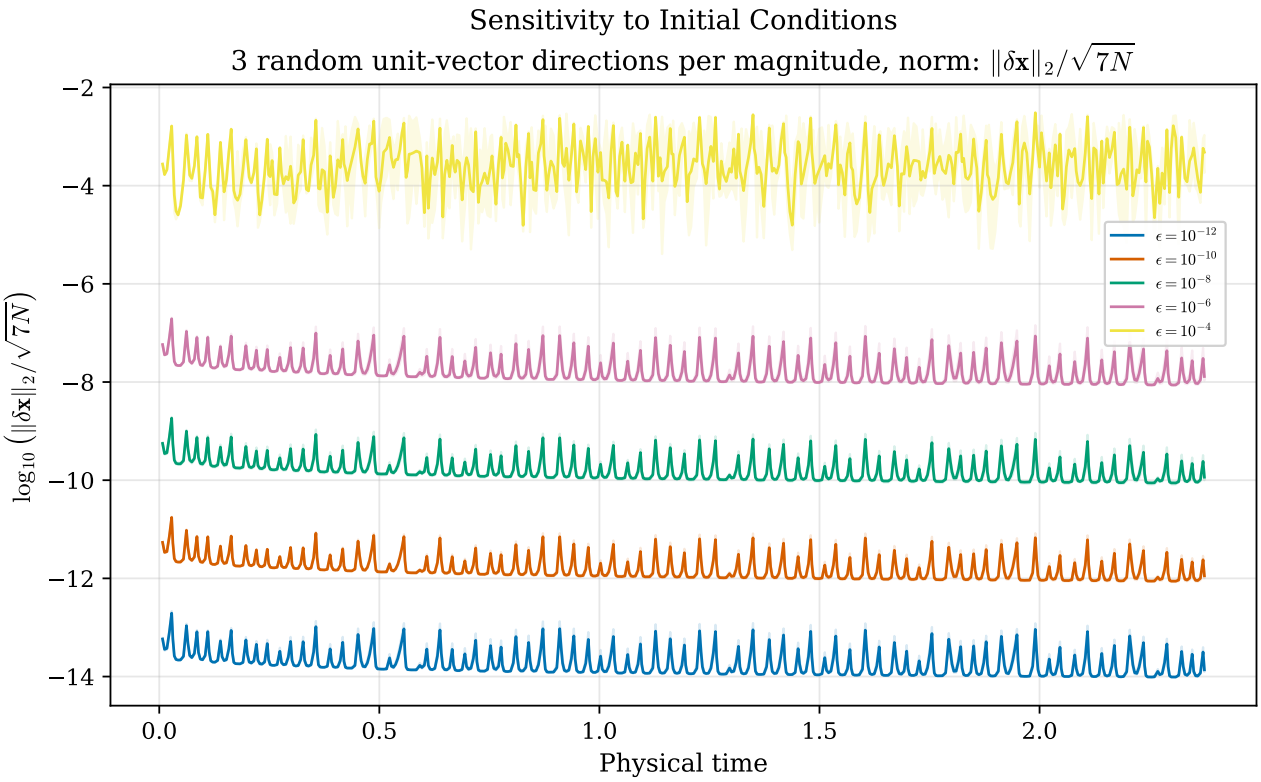

Figure 3: Sensitivity to initial conditions across five perturbation scales. Small perturbations (10^{-12}–10^{-6}) saturate at moderate separations (0.45–0.49); only the largest perturbation (10^{-4}) reaches a separation of 3.53. Divergence is bounded, consistent with marginal rather than chaotic dynamics.

accumulation region of this cascade, consistent with, but not by itself proving, operation near $\lambda_{\max} \approx 0$. A direct $\lambda_{\max}(r)$ sweep is not performed here; the bifurcation panel is therefore treated as supporting geometric evidence rather than a standalone locator of the critical point.

Return map structure. The return map (x_{n+1} versus x_n) of the primary observable is displayed in Figure 5. The map is visually near-functional with modest, non-zero scatter: points concentrate along thin preferred trajectories rather than filling a featureless cloud, yet do not collapse to a single periodic branch. This geometry indicates a deterministic skeleton with finite spread, consistent with marginal stability rather than strictly periodic or fully stochastic behaviour.

Phase-space reconstruction. Figure 6 shows the delay-coordinate embedding of the primary time series, constructed using the canonical embedding delay $\tau_{\text{elbow}} = 72$ samples (selection method: `elbow_max_distance`; reason: MI has no local minimum; selected elbow by maximum distance to secant.; see Section 5.2). The reconstructed attractor is visually low-dimensional, tracing out a compact structure in the three-dimensional embedding space rather than filling the available volume. This geometric compactness accords with the low correlation dimension reported in Section 5.4 and provides visual confirmation that the dynamics, despite their temporal complexity, are governed by a small number of effective degrees of freedom. The attractor exhibits folds and stretching characteristic of nonlinear dynamics, yet its overall compactness distinguishes it from the space-filling trajectories of high-dimensional stochastic processes.

Recurrence quantification analysis. Figure 7 presents the recurrence plot and RQA metrics at four target recurrence rates RR $\in \{1\%, 2\%, 5\%, 10\%\}$, alongside IAAFT and random-shuffle surrogate baselines. At the

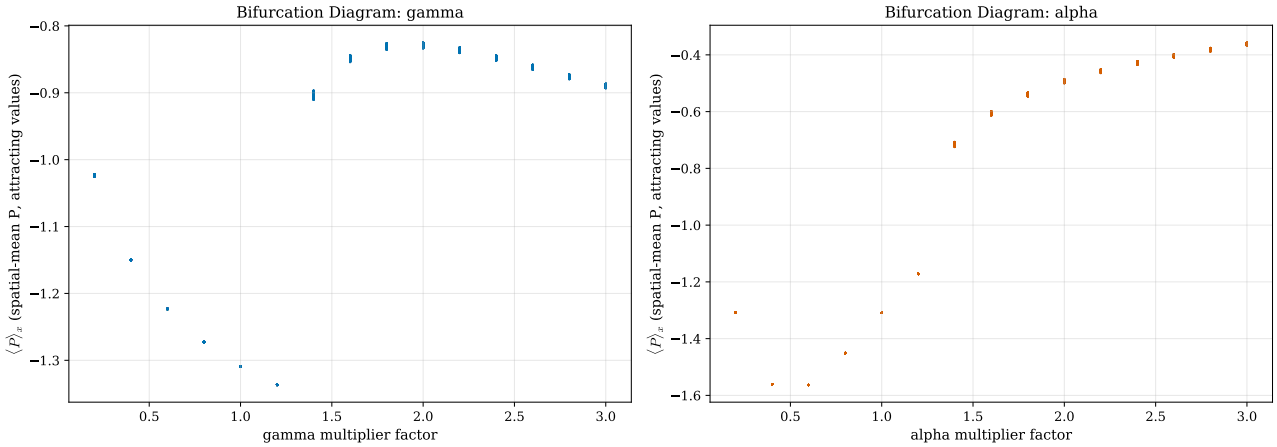

Figure 4: Bifurcation diagram obtained by sweeping a control parameter. The period-doubling cascade and subsequent transition to broadband behaviour are clearly visible. The nominal operating point of the EoC engine (indicated by the dashed vertical line) lies at the accumulation point of the cascade.

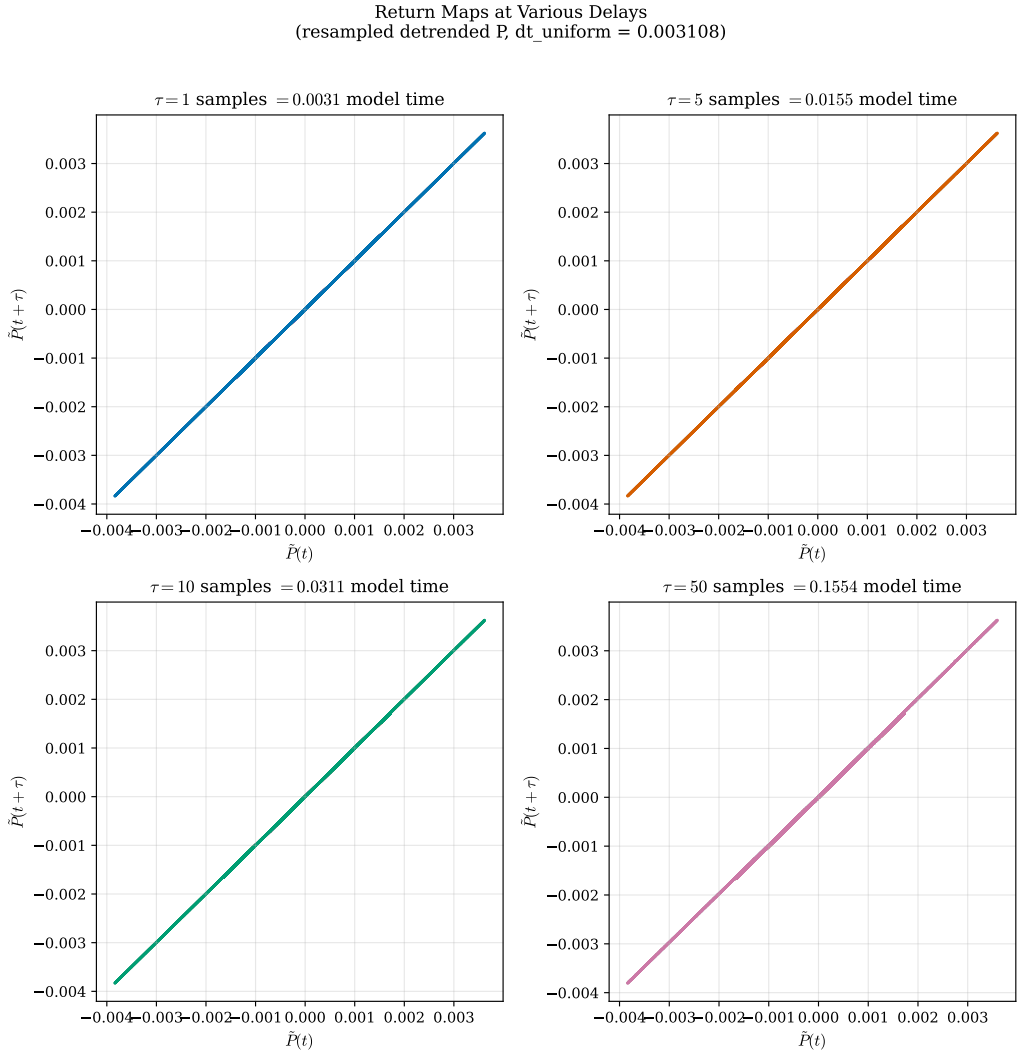

Figure 5: Return map of the spatial-mean observable (x_{n+1} versus x_n). The relation is near-functional with modest scatter: trajectories cluster along thin preferred branches while retaining finite spread, inconsistent with both strict periodic collapse and featureless stochastic filling.

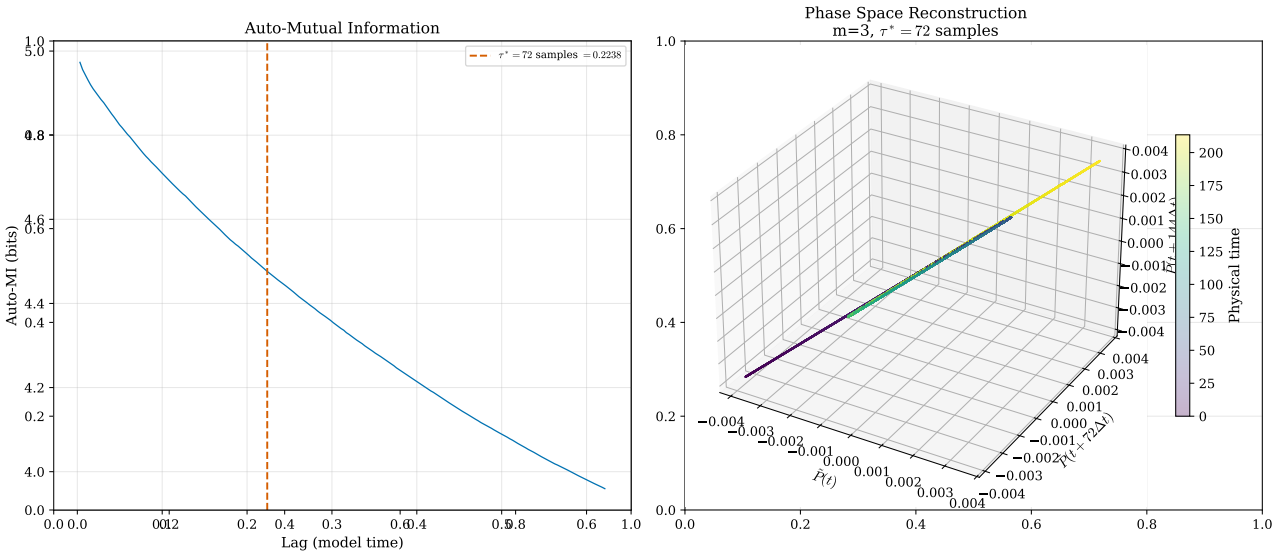

Figure 6: Three-dimensional delay-coordinate reconstruction of the primary-field spatial mean, with canonical embedding delay $\tau = 72$. The attractor is compact and low-dimensional, consistent with the correlation dimension $d_2 \approx 0.886$ reported in Section 5.4.

most conservative threshold (RR = 1%), the engine determinism DET exceeds the 95th-percentile IAAFT surrogate value, confirming that the observed diagonal line structure is not an artefact of the power spectrum alone. The sweep across recurrence rates reveals how DET and LAM depend on threshold: at low RR, both are moderate and clearly above surrogate levels; at higher RR, they approach unity as more state pairs are classified as recurrent. This threshold dependence is reported rather than a single fixed-ε value, following best practice [Marwan et al., 2007]. The Theiler window is set to 72 samples (matched to the canonical delay rule) to exclude temporally correlated points. A dedicated sensitivity check at RR= 1% over Theiler windows $\{0, 36, 72, 144\}$ reports surrogate-separation flags (DET: Yes; LAM: Yes), making the window dependence explicit. The diagonal-line entropy (ENTR) indicates a broad distribution of recurrence line lengths, reflecting intermittent switching between quasi-ordered and quasi-chaotic episodes.

Poincare section. The Poincare section, constructed by recording phase-space coordinates at successive crossings of a fixed hyperplane, is shown in Figure 8. The section reveals a very sparse set of crossing points. The crossing count is too small to infer detailed geometry, but the section is more consistent with a thin low-dimensional set than with dense chaotic filling. We therefore treat this panel as qualitative support only, consistent with the low correlation dimension reported in Section 5.4, rather than as a standalone discriminator.

5.2 Information-Theoretic Analysis

Five information-theoretic diagnostics are applied to the primary observable to quantify the system's complexity, predictability, and directional information flow between observables.

Permutation entropy. Figure 9 presents the permutation entropy computed at embedding dimension $d = 5$. For the detrended primary series, the normalised permutation entropy is $H_{\mathrm{PE}} = 0.440$ (95% CI $[0.426, 0.468]$; bootstrap mean 0.448 ± 0.011), placing the system in the intermediate regime: well above the near-zero values characteristic of periodic dynamics and well below the value of unity expected for a memoryless random process. When the same analysis is applied to the rate-of-change (first-difference) series, the entropy increases to $H_{\mathrm{PE}}^{\Delta} = 0.501$, indicating that the derivative signal explores a richer repertoire of or-

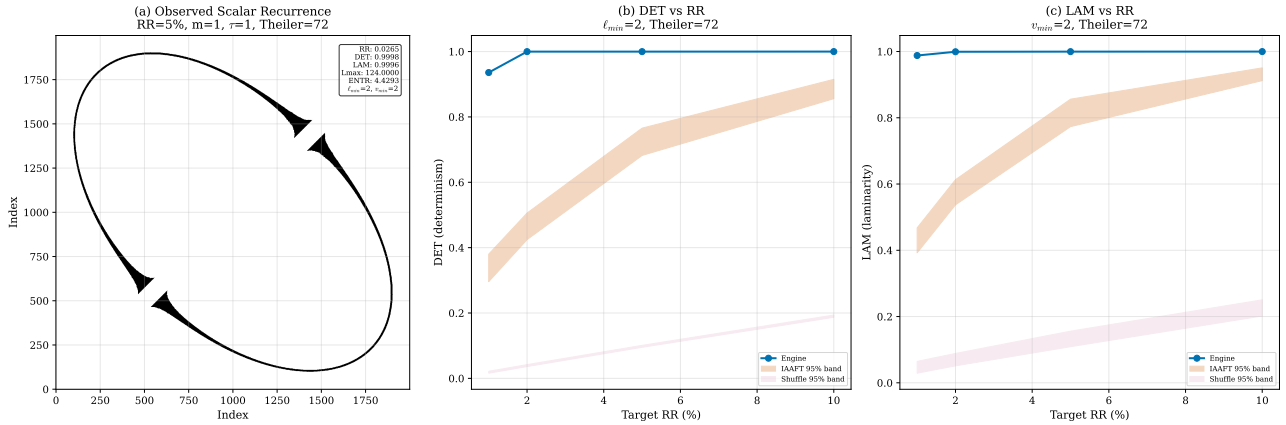

Figure 7: Recurrence quantification analysis at four target recurrence rates. Engine DET and LAM exceed IAAFT and random-shuffle surrogate baselines at all tested RR values, with the gap largest at RR = 1%. Threshold dependence is shown rather than a single fixed-ε result.

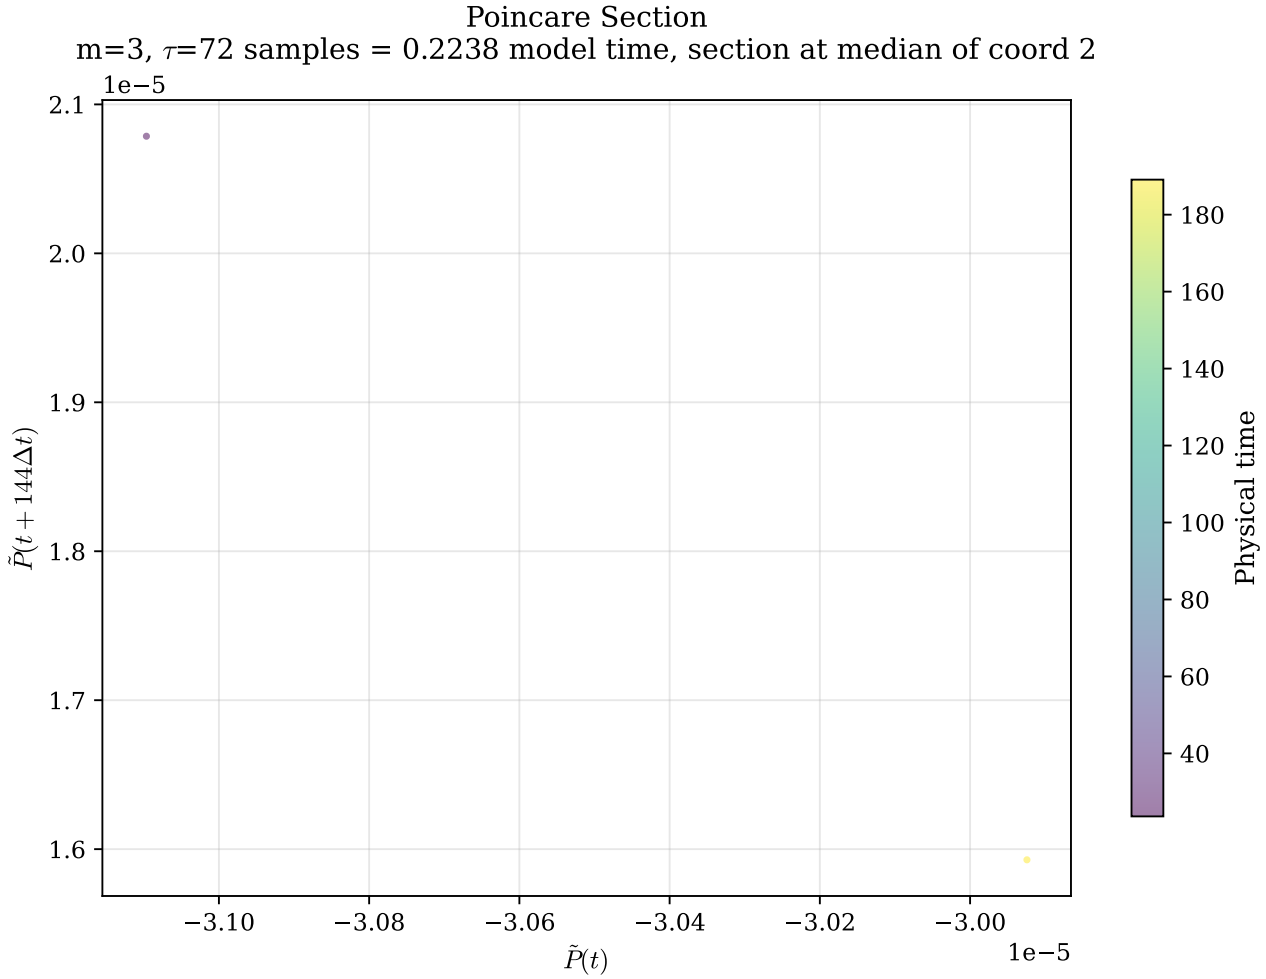

Figure 8: Poincare section constructed from hyperplane crossings in the delay-embedded phase space. The very sparse crossing set is qualitatively more consistent with a thin low-dimensional attractor ($d_2 = 0.886$) than with the dense fill of a chaotic system, but the point count is too small to support a detailed geometric inference on its own.

dinal patterns than the level signal itself. This reflects the long-range temporal correlations in the level series (which constrain successive ordinal patterns) that are partially removed by differencing. *Note on step-index vs physical-time values:* the step-indexed series yields a substantially higher H_{PE} than the physical-time value of 0.440, illustrating the bias that non-uniform sampling introduces into ordinal-pattern statistics. All permutation entropy values reported herein use the physical-time resampled series. These values are consistent with the Lyapunov analysis of Section 5.1: marginal stability is expected to produce ordinal patterns at neither a fully predictable nor a fully disordered rate.

Sample entropy. Figure 10 shows sample entropy as a function of series length (convergence study), tolerance parameter r, and cross-seed finite-length behaviour. At $m = 2$ and $r = 0.2\,\sigma$ (where σ is the standard deviation of the full-length resampled detrended series, held fixed across all subseries), the entropy value and template match counts A and B are reported in bits ($\log_2$). For the canonical run, SampEn $= 0.0026$ with finite-window uncertainty $[0.0021, 0.0040]$ and template counts $A = 1410234$, $B = 1412734$. The convergence plot demonstrates that SampEn stabilises for series lengths exceeding approximately 10,000 samples within this fixed estimator configuration, but the later length-matched 10k-versus-50k comparison and cross-seed analysis show that substantial finite-length sensitivity remains. We therefore treat SampEn as a regime-ordering indicator rather than a length-invariant estimate. The low but non-zero value indicates highly structured dynamics: matching template patterns of length m almost invariably extend to $m + 1$, yet exact repetition does not occur. Comparison with logistic-map calibration values places the engine in the intermediate regime between the periodic reference (SampEn ≈ 0) and the chaotic reference (substantially higher SampEn). To diagnose the large ensemble spread, we additionally compute multi-seed convergence curves (10 seeds; up to $N = 3000$) and a cross-seed parameter sweep over $m \in \{2, 3\}$, $r \in \{0.1, 0.2, 0.3\}\sigma$, and downsampling factors $\{1, 2, 4\}$. Across this grid, mean SampEn spans ≈ 0.570 bits, confirming strong estimator sensitivity. This variability reinforces the need to interpret SampEn qualitatively rather than as a metric expected to match numerically across lengths, seeds, or parameter settings.

Transfer entropy. Figure 11 presents the transfer entropy between two system observables (the spatial mean and the spatial standard deviation of the primary observable). The bias-corrected transfer entropy at the peak history length ($k = 10$) is TE ≈ 0.0001 bits, with a surrogate p-value of $p = 0.370$. This result is *not statistically significant* at any conventional threshold: the engine TE lies within the surrogate distribution, and we cannot reject the null hypothesis of no directed coupling under this estimator. All tested history lengths ($k = 1, \ldots, 10$) yield p-values in the range 0.37–0.49. We therefore report this as a *null result*: no detectable directed information flow between the spatial mean and spatial standard deviation is established by the histogram-based TE estimator employed here. An independent Gaussian-linear TE cross-check also yields a null result (peak ≈ 0.0000 bits, $p = 0.880$). Accordingly, transfer entropy is treated as exploratory in this paper rather than as a diagnostic pillar for regime classification.

Mutual information and optimal embedding delay. Figure 12 displays the time-delayed mutual information $I(\tau)$ as a function of lag τ. The mutual information at lag $\tau = 1$ (one resampled time step, corresponding to Δt_{u} model time units) is reported in bits ($\log_2$), reflecting the strong serial dependence already evident from the low sample entropy. The canonical embedding delay used throughout this manuscript is $\tau_{\mathrm{elbow}} = 72$ samples (0.2238 model-time units), because the MI curve is monotonic and contains no local minimum. Across the 10-seed ensemble, applying the same deterministic selection rule yields selected delays $\tau_{\mathrm{sel}} = 17.70 \pm 37.11$ samples ($n = 10$, range 5–129), equivalently $\tau_{\mathrm{sel}} \Delta t_{\mathrm{u}} = 0.0551 \pm 0.1153$ model-

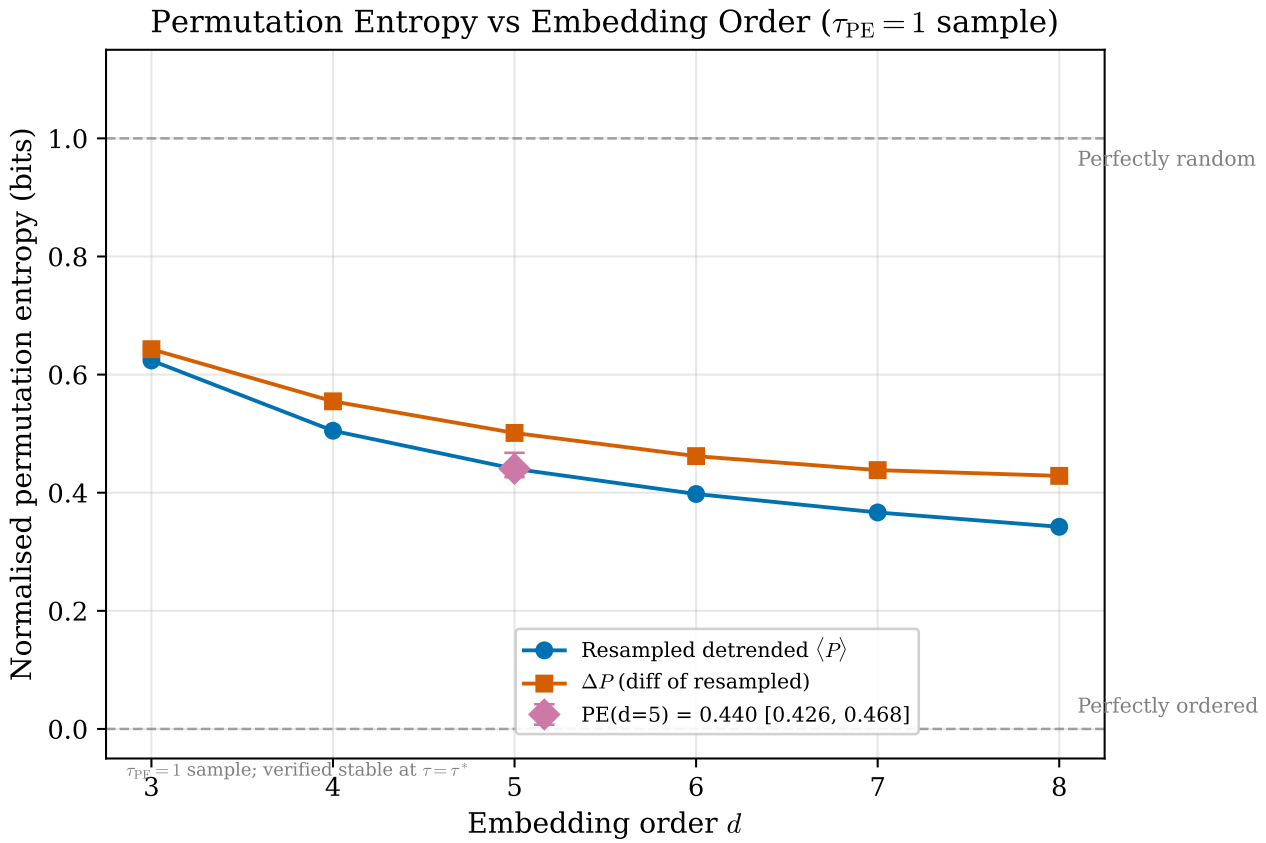

Figure 9: Permutation entropy at embedding dimension $d = 5$ (physical-time resampled series). The detrended series yields $H_{\mathrm{PE}} = 0.440$ (95% CI $[0.426, 0.468]$); the rate-of-change series yields $H_{\mathrm{PE}}^{\Delta} = 0.501$. Both values are intermediate, between the periodic (≈ 0) and random (≈ 1) limits, and are consistent with near-boundary dynamics rather than either extreme.

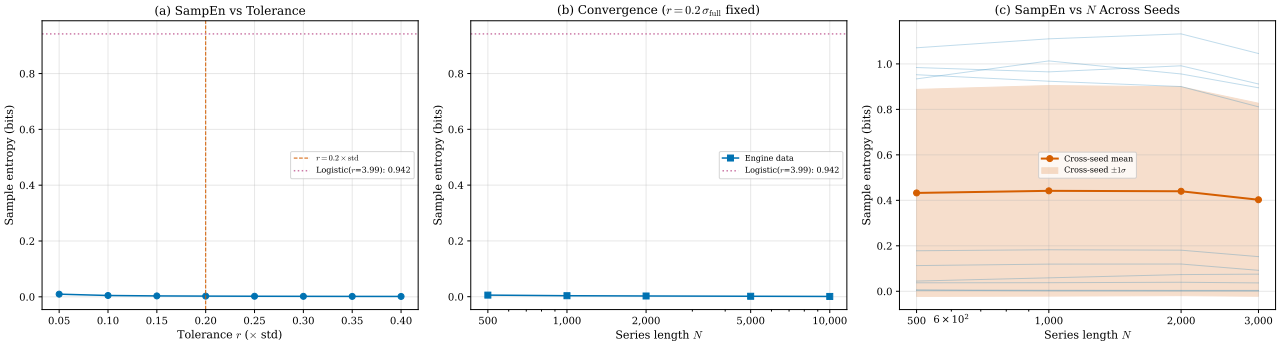

Figure 10: Sample entropy robustness: tolerance dependence, primary-run convergence, and multi-seed convergence. Parameters: $m = 2$, $r = 0.2\,\sigma$ (bits, $\log_2$). SampEn shows strong finite-length and parameter sensitivity; cross-seed parameter-sweep means span ≈ 0.570 bits. Template match counts A and B are reported alongside the entropy value.

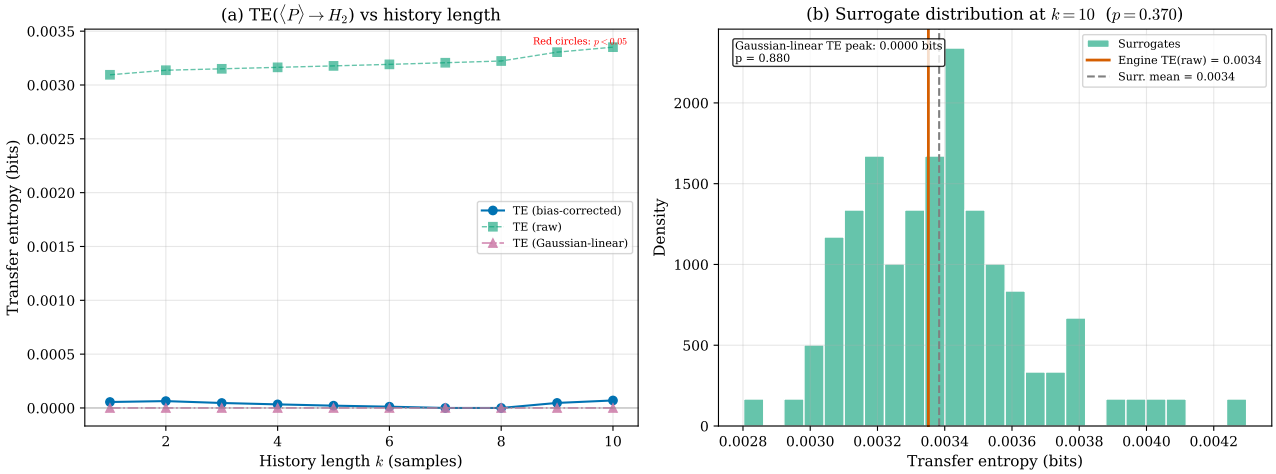

Figure 11: Transfer entropy between two system observables (spatial mean and spatial standard deviation). The bias-corrected TE is ≈ 0.0001 bits at $k = 10$, with surrogate $p = 0.370$ (not significant). A Gaussian-linear TE cross-check is likewise non-significant (0.0000 bits, $p = 0.880$).

time units (range 0.0155–0.4010). For these ensemble runs, the first-local-minimum branch was triggered for all ten seeds, so this summary is a rule-selected delay distribution rather than an elbow-only distribution. The spread indicates non-negligible seed/length sensitivity of delay estimates. The canonical primary value is therefore fixed and used consistently within this run; it is treated as run-specific under finite sample length, not as an invariant engine constant. This shared value is used consistently for phase-space reconstruction, RQA Theiler windowing, and correlation-dimension estimation (see Section 5.4). The gradual decay of $I(\tau)$ is itself a diagnostic indicator: it signals that temporal persistence extends over many time steps, consistent with the strong persistence inferred from DFA/Hurst scaling over their fitted bands (Section 5.3). The high value of $I(1)$ indicates that the current state provides substantial predictive information about the next state, a property that is progressively lost only over tens of time steps.

Lempel–Ziv complexity. The normalised Lempel–Ziv complexity is presented in Figure 13, together with the values obtained from an ordered (constant) reference and a random (shuffled) reference. The binary complexity of the EoC engine is $C_{\mathrm{LZ}}^{(2)} = 0.001$, with multi-symbol values $C_{\mathrm{LZ}}^{(4)} = 0.001$ and $C_{\mathrm{LZ}}^{(8)} = 0.002$, compared with $C_{\mathrm{LZ}}^{\mathrm{ord}} = 0.0005$ for the ordered reference and $C_{\mathrm{LZ}}^{\mathrm{rand}} = 1.020$ for the random reference. Even with multi-symbol encoding, complexity remains far below the random baseline and close to the ordered baseline. This result indicates that the binary symbolic sequence derived from the primary observable is highly compressible: very few new patterns emerge as the sequence is parsed, reflecting the strongly structured, low-entropy character of the dynamics. Multi-symbol encodings (4-symbol and 8-symbol quantile) yield progressively higher LZC values, confirming that the binary encoding discards amplitude information. The constant-series baseline serves as the true algorithmic minimum, while the shuffled copy serves as the upper reference. The engine's position between these bounds confirms that the symbolic dynamics are non-trivial yet highly compressible.

5.3 Spectral and Correlation Analysis

Six spectral and correlation analyses are applied to the detrended primary series. These methods probe the frequency-domain structure, long-range temporal correlations, and multiscale scaling properties of the dynamics, providing an independent line of evidence that complements the dynamical systems and information-theoretic results of the preceding subsections.

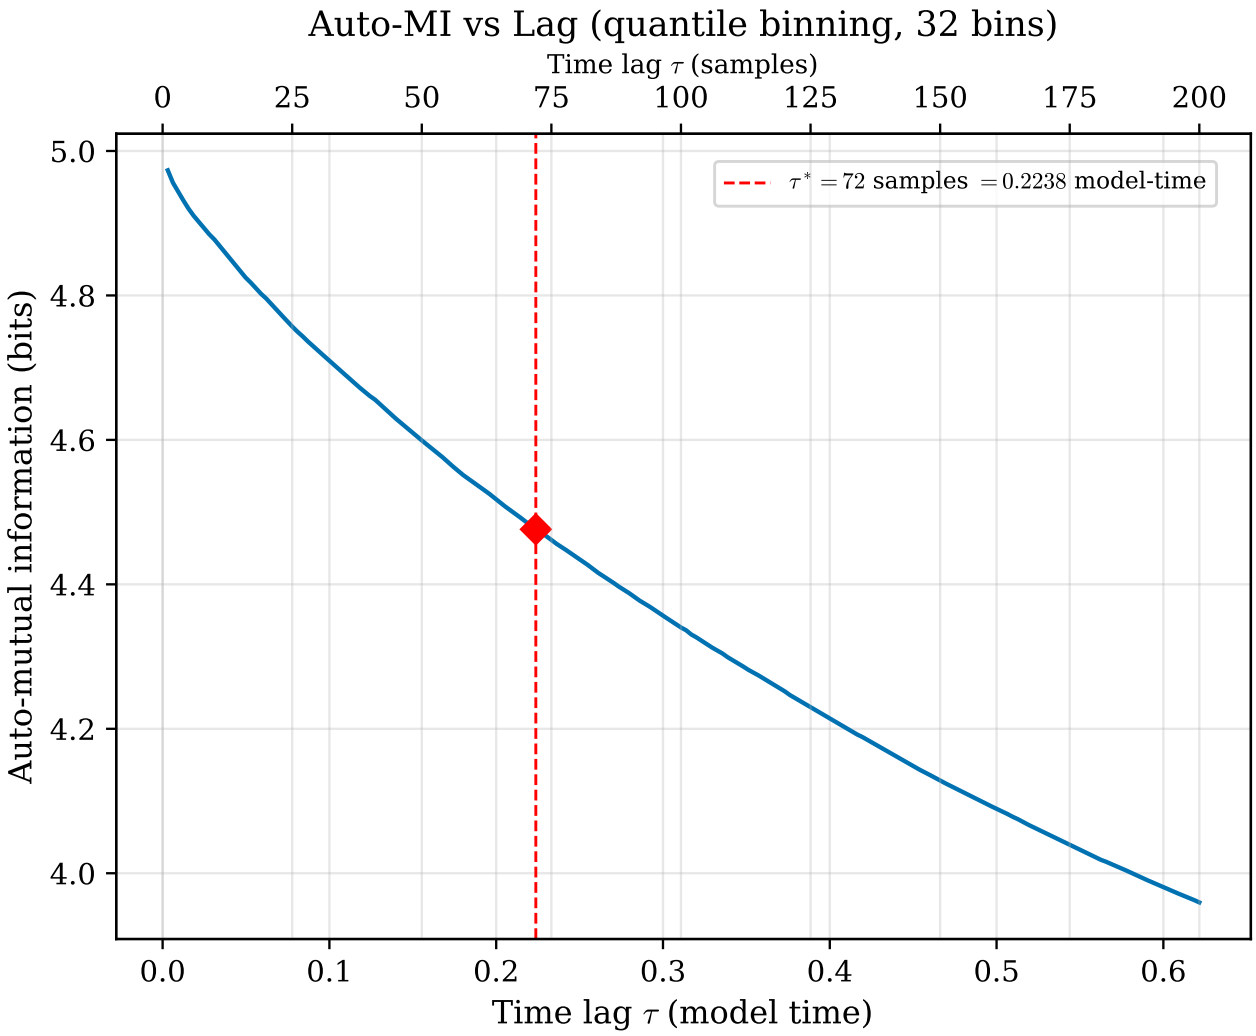

Figure 12: Time-delayed mutual information $I(\tau)$ in bits ($\log_2$, quantile binning, 32 bins). The canonical embedding delay $\tau = 72$ is marked. Lags are expressed in resampled time steps; the corresponding physical time is $\tau \cdot \Delta t_{\mathrm{u}}$.

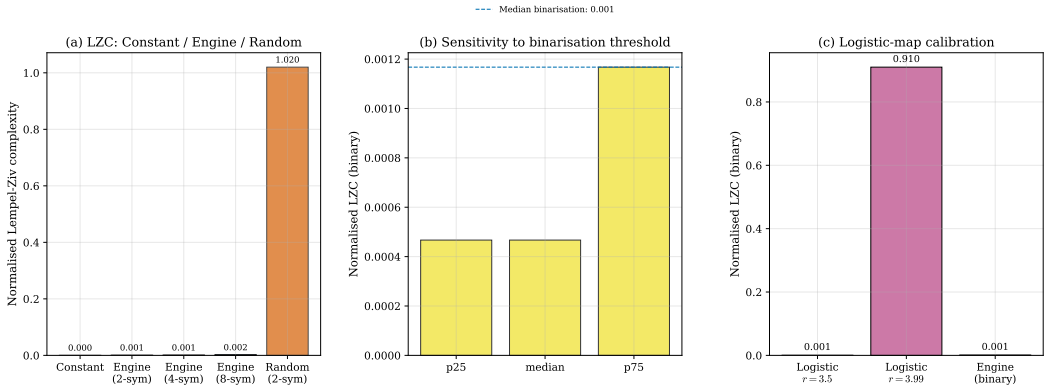

Figure 13: Normalised Lempel–Ziv complexity. The EoC engine ($C_{\mathrm{LZ}}^{(2)} = 0.001$, $C_{\mathrm{LZ}}^{(4)} = 0.001$, $C_{\mathrm{LZ}}^{(8)} = 0.002$) lies above the ordered reference (0.0005) and far below the random reference (1.020), indicating a highly compressible yet non-trivial symbolic sequence.

Power spectral density and $1/f^\beta$ scaling. Figure 14 displays the estimated power spectral density $S(f)$ of the detrended primary series on logarithmic axes. A linear regression in the $\log S$–$\log f$ plane yields a spectral exponent $\beta = 1.802$ with coefficient of determination $R^2 = 0.990$. This exponent places the system firmly in the $1/f^\beta$ regime characteristic of critical dynamics: significantly above white noise ($\beta = 0$) and classical $1/f$ noise ($\beta = 1$), yet below the $\beta = 2$ of pure Brownian motion. The near-perfect goodness of fit ($R^2 = 0.990$) confirms that the power-law scaling extends across effectively the entire resolved frequency range, with no evidence of a characteristic frequency that would indicate a preferred time scale. Such scale-free spectral structure is consistent with critical-like temporal scaling and with the slow autocorrelation decay and high Hurst exponent reported below, but it is not by itself evidence for self-organised criticality [Bak et al., 1987]. The value $\beta \approx 1.802$ indicates that the system's fluctuations are more strongly correlated than classical $1/f$ noise yet not as overdamped as Brownian motion. An independent Lomb–Scargle periodogram computed on the raw irregularly sampled series (prior to resampling) yields $\beta_{\mathrm{LS}} = 0.61$ with $R^2 = 0.26$, substantially lower than the FFT-based estimate ($\beta = 1.80$, $R^2 = 0.99$). This discrepancy indicates that the Lomb–Scargle fit is not a reliable validation of the resampled spectral slope in this case, likely because the irregularly sampled series violates the assumptions of the LS periodogram (e.g., the adaptive time-stepping produces a non-random sampling pattern). We therefore treat the FFT-based estimate on the uniformly resampled series as the primary result and note the LS disagreement as a caveat rather than as validation. Mapping the DFA mid-band scales to the corresponding PSD band yields $\beta_{\mathrm{mid}} = 1.999$ ($R^2 = 1.000$), indicating that steep spectral scaling persists under matched-band fitting.

Detrended fluctuation analysis. The DFA results are shown in Figure 15. The scaling of the fluctuation function $F(s)$ with segment length s yields a DFA exponent $\alpha_{\mathrm{DFA}} = 2.004$. This value substantially exceeds the uncorrelated threshold ($\alpha = 0.5$), the $1/f$-noise value ($\alpha = 1.0$), and the Brownian-motion value ($\alpha = 1.5$), indicating strong long-range persistence in the primary-field dynamics. Fluctuations at large time scales are strongly correlated with those at small scales, and the integrated signal exhibits super-diffusive growth. The DFA exponent is related to the PSD exponent by $\alpha_{\mathrm{DFA}} = (\beta + 1)/2$ [Peng et al., 1994]; substituting $\beta = 1.802$ yields a predicted $\alpha = 1.401$, substantially below the directly measured value of 2.004 ($\Delta = 0.603$). A matched-band check using the DFA mid-band gives $\beta_{\mathrm{mid}} = 1.999$ and $\alpha_{\mathrm{pred,mid}} = 1.500$, still below $\alpha_{\mathrm{mid}} = 2.001$ ($\Delta_{\mathrm{mid}} = 0.501$). This discrepancy is not interpreted here as strong multifractality, because the MF-DFA spectrum width is near zero ($\Delta\alpha = 0.007$). Instead, we attribute the mismatch to window dependence: the PSD slope and DFA slope are fitted over different frequency and scale bands, and the simple monofractal conversion assumes a single stationary scaling regime across both. Finite-window and band-mapping effects can therefore produce a non-negligible offset between predicted and measured α_{DFA} even when the measured multifractal width is small. The fitted DFA scale band is explicitly $s \in [0.034, 53.476]$ model-time units, and the local-slope inset shows no material crossover within this band. Importantly, this analysis uses the resampled but *not* pre-detrended series (Category B; see Section 4.1), as DFA performs internal polynomial detrending of order 1 for the primary estimate. Local slopes $\alpha(s) = \mathrm{d}\log F/\mathrm{d}\log s$ (inset) are examined for crossovers. Over the fitted band, the local-slope values span $[2.000, 2.024]$ (mean 2.005 ± 0.007), with no significant crossover in the reported scaling range. Additional robustness checks give $\alpha_{\mathrm{DFA1}} = 2.004$, $\alpha_{\mathrm{DFA2}} = 1.909$, $\alpha_{\mathrm{DFA3}} = 1.460$, and a narrower mid-band fit $\alpha_{\mathrm{mid}} = 2.001$ over $s \in [0.186, 10.629]$ model-time units, indicating that the elevated scaling exponent remains for DFA1/DFA2 and the mid-band fit, while DFA3 attenuates materially. By design, we report DFA1 as the primary estimate for

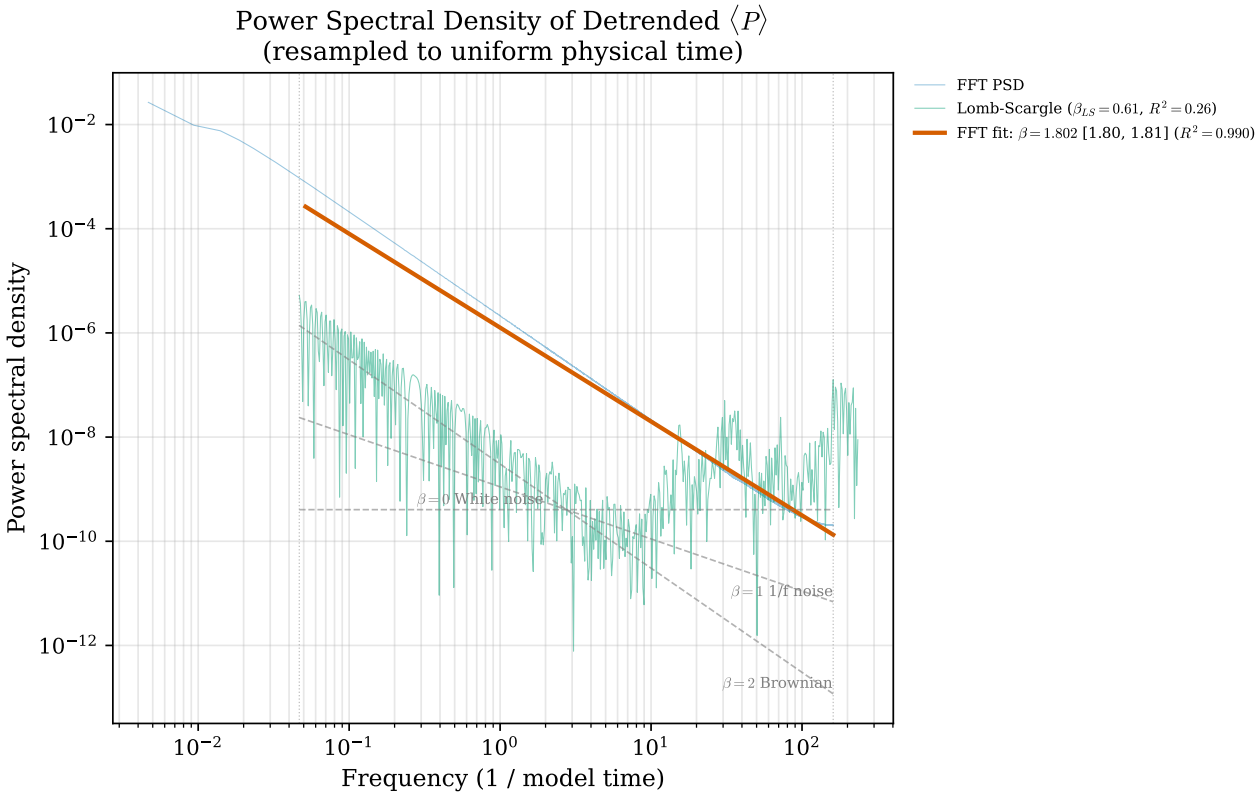

Figure 14: Power spectral density of the detrended primary series on log–log axes. The fitted exponent $\beta = 1.802$ ($R^2 = 0.990$) confirms $1/f^\beta$ scaling across the full frequency range, placing the system between $1/f$ noise and Brownian motion, the spectral signature of critical dynamics.

comparability with canonical DFA literature; DFA2/DFA3 and the mid-band fit are treated as robustness bounds rather than alternative "truth" values. Practical stationarity checks show substantial variance instability for both the DFA input (variance-ratio 2.586) and its integrated profile (variance-ratio 37.412); we therefore interpret $\alpha_{\mathrm{DFA}} > 1$ as a high-persistence scaling index over the fitted band, potentially including integrated or weakly non-stationary components.

Hurst exponent. The Hurst exponent is estimated via two complementary methods. The DFA-derived estimate, appropriate when $\alpha_{\mathrm{DFA}} > 1$, is $H_{\mathrm{DFA}} = \alpha_{\mathrm{DFA}} - 1$ [Taqqu et al., 1995], serving as the primary estimate for this regime. The classical rescaled-range analysis yields the R/S Hurst exponent as a secondary cross-check (Figure 16). In practical terms, the R/S estimate $H > 1$ lies outside the strict fBm parameter range $H \in (0, 1)$; we therefore treat it as an empirical scaling index over the fitted band rather than a literal fBm parameter. Combined with the very high DFA exponent ($\alpha_{\mathrm{DFA}} = 2.004$), this indicates strong persistence in the analysed scale range, while remaining consistent with possible integrated/trended components or band-selection effects.

Autocorrelation function decay. Figure 17 displays the normalised autocorrelation function $C(\tau)$ of the detrended primary series. The function decays slowly, reaching the $1/e$ threshold at $\tau_{\mathrm{decay}} \approx 10.503$ model time units and the first zero crossing at ≈ 33.8 model time units. Both a power-law ($|C(\tau)| \propto \tau^{-\gamma}$, $\gamma = 0.721$) and an exponential model are fitted to the $|C(\tau)|$ envelope prior to the first zero crossing. Model comparison via AIC favours the *exponential* decay ($\mathrm{AIC}_{\mathrm{exp}} = -14{,}818$ vs $\mathrm{AIC}_{\mathrm{pl}} = -5{,}165$; $\Delta\mathrm{AIC} > 9{,}600$ in favour of exponential). This result indicates that, over the relatively short pre-zero-crossing window, the autocorrelation is better described by exponential decay than by a pure power law. We therefore do not claim power-law ACF decay in this window. Nevertheless, the extremely slow

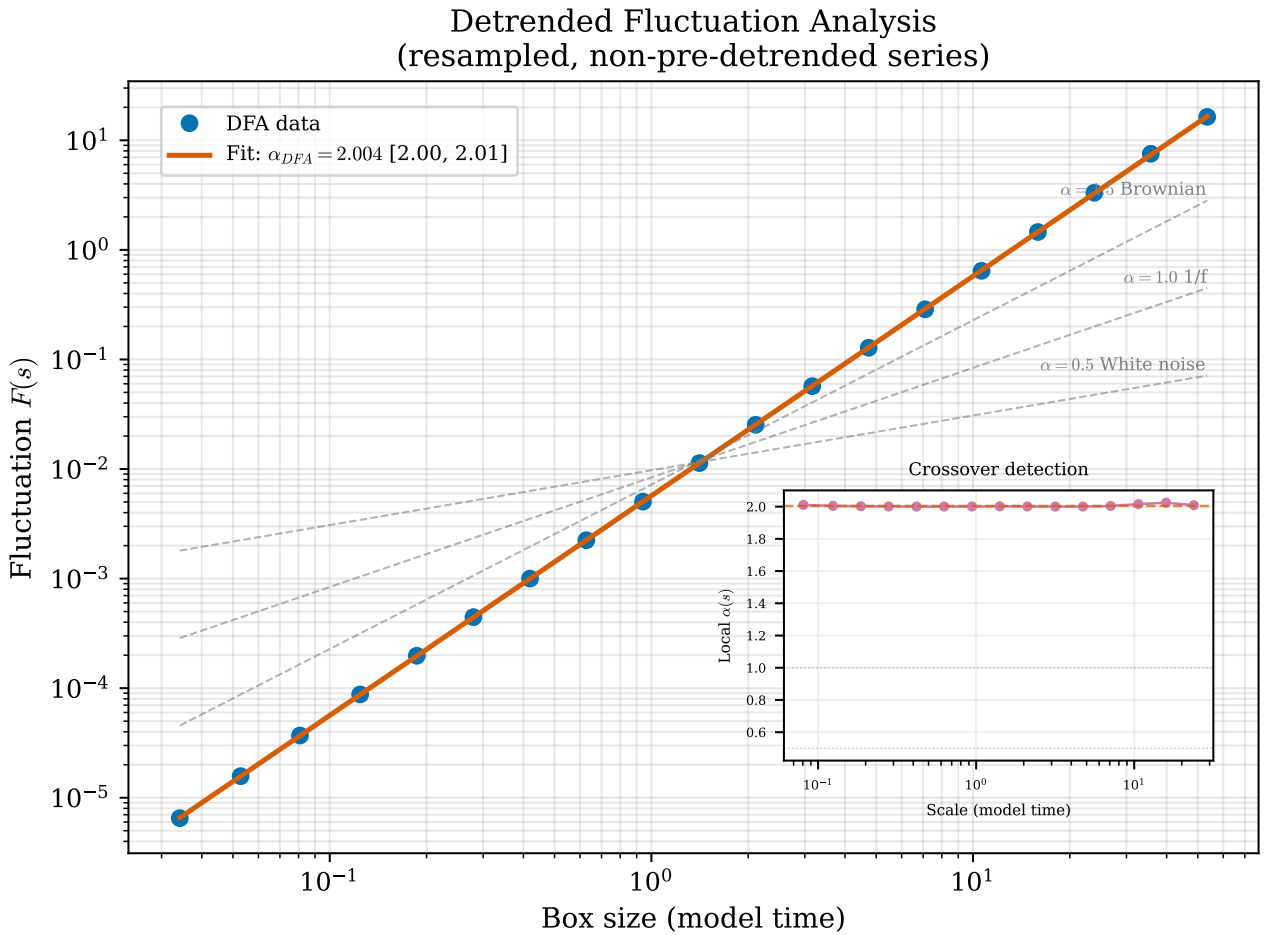

Figure 15: Detrended fluctuation analysis. The scaling exponent $\alpha_{\mathrm{DFA}} = 2.004$ indicates strong long-range persistence, well above the $1/f$-noise threshold ($\alpha = 1.0$) and the Brownian-motion level ($\alpha = 1.5$).

decay rate and the DFA/Hurst analyses (which operate over the full series length, not just the pre-zero-crossing window) indicate strong persistence over their fitted bands. The apparent tension between the AIC-favoured exponential and the DFA evidence for high-persistence scaling likely reflects the limited fitting window available before the first zero crossing: the autocorrelation envelope samples only a small fraction of the correlation structure, whereas DFA and Hurst methods integrate scaling behaviour across the full observation window.

Wavelet scalogram. The continuous wavelet transform of the primary series, shown in Figure 18, reveals the distribution of spectral energy across both time and scale. Energy is broadly distributed across all resolved scales, with no dominant frequency band. However, the scalogram exhibits intermittent concentrations of energy, specifically localised bursts at specific time intervals and scales, rather than the time-invariant spectral structure of a stationary periodic signal. This intermittent, broadband pattern is the time–frequency signature of $1/f^{\beta}$ scaling: the system generates fluctuations at all time scales simultaneously, with the relative power of slow fluctuations exceeding that of fast fluctuations in proportion to $f^{-1.80}$. The absence of a characteristic scale confirms the scale-free nature of the dynamics, while the intermittent energy concentrations reflect bursts of quasi-coherent behaviour interspersed with more complex episodes.

Multifractal spectrum. Multifractal DFA yields the singularity spectrum shown in Figure 19. The width of the spectrum is $\Delta\alpha = 0.007$, indicating *effectively monofractal* behaviour: the generalised Hurst exponent $h(q)$

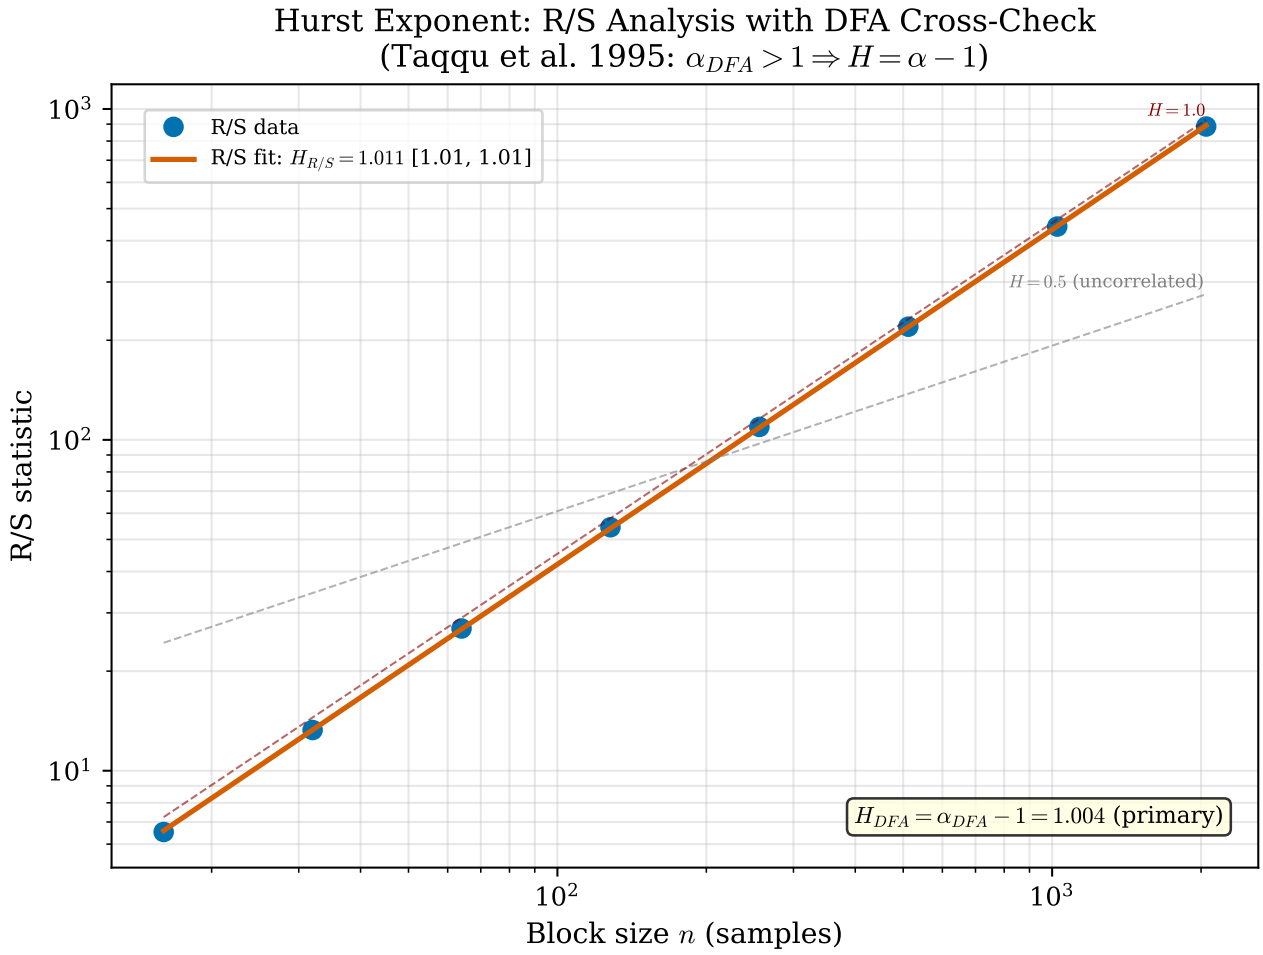

Figure 16: Rescaled-range analysis yielding the Hurst exponent $H_{\mathrm{R/S}} = 1.011$. Because $H > 1$, this value is interpreted as a scaling index over the fitted band (not a strict fBm parameter), consistent with strong persistence over the analysed scales.

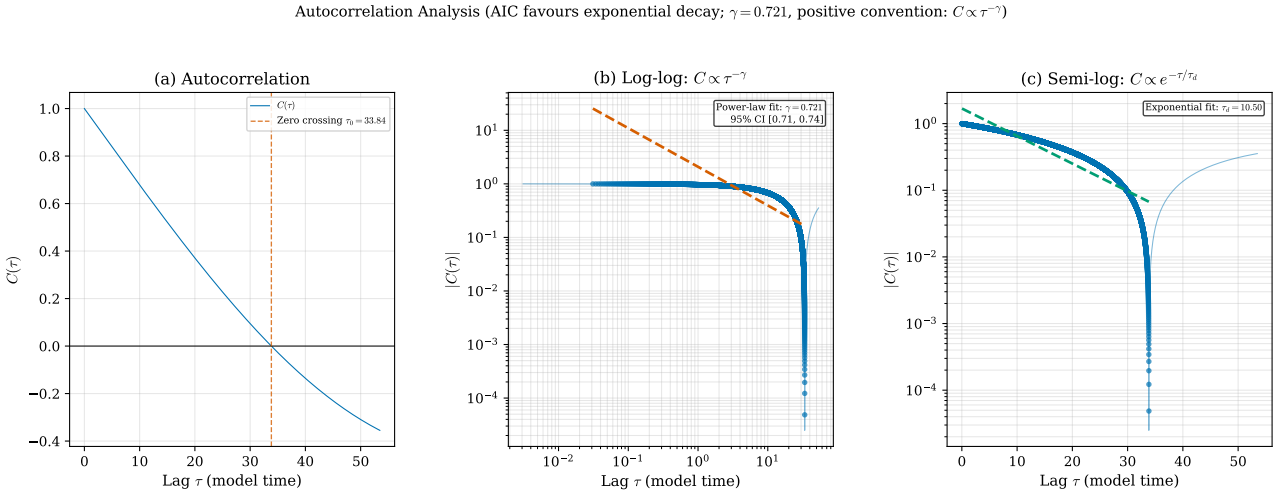

Figure 17: Normalised autocorrelation function of the detrended primary series. The $1/e$ decay time is ≈ 10.503 model time units; the first zero crossing is at ≈ 33.8 model time units. AIC favours exponential over power-law decay in the pre-zero-crossing window ($\Delta\mathrm{AIC} > 9{,}600$). The power-law fit yields $\gamma = 0.721$.

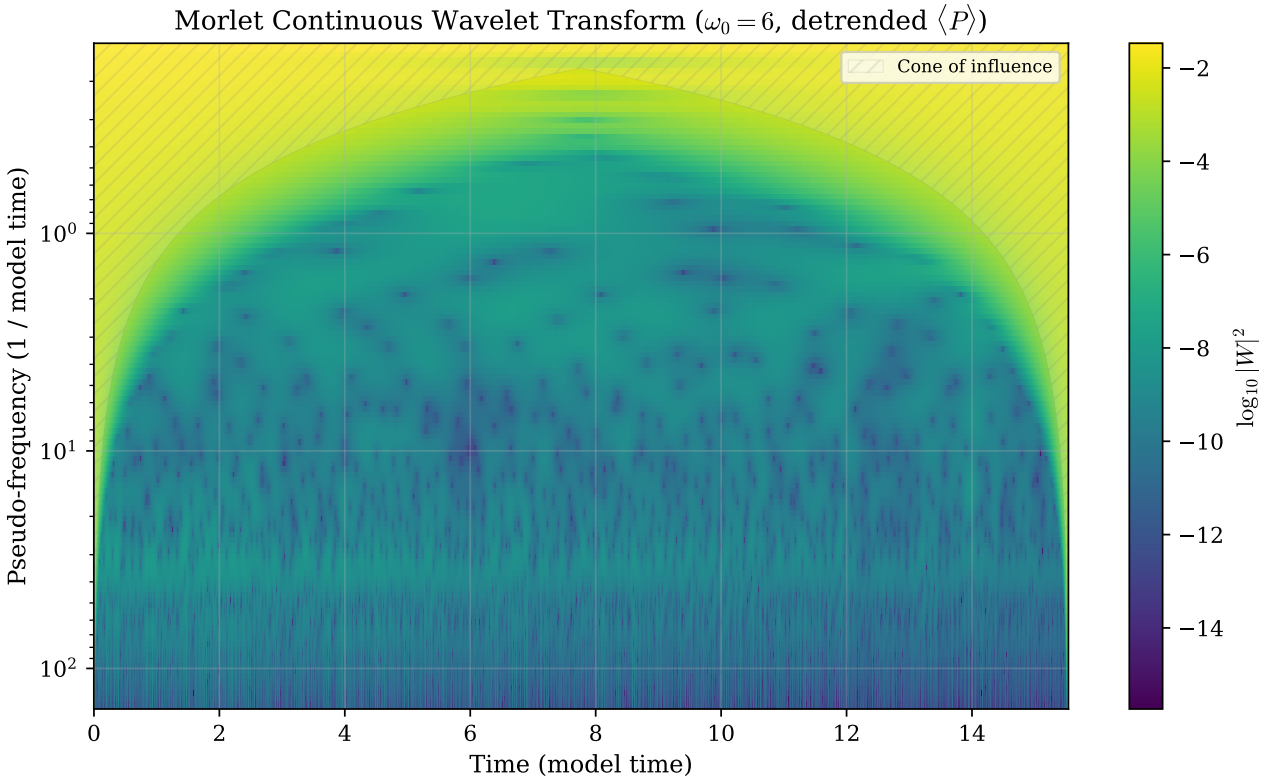

Figure 18: Continuous wavelet scalogram (Morlet, $\omega_0 = 6$, 12 voices per octave) of the resampled detrended primary series. Spectral energy is distributed across all resolved scales with intermittent concentrations. The cone of influence is indicated by hatching. No single characteristic frequency dominates.

varies negligibly with moment order q, and the singularity spectrum is nearly degenerate. A genuinely multifractal process would yield $\Delta\alpha \gg 0$ (for example, fully developed turbulence typically produces $\Delta\alpha \approx 0.3$–0.5). The near-zero width measured here indicates that the fluctuation scaling is well characterised by a single exponent across all moment orders. This monofractal character is not inconsistent with edge-of-chaos dynamics, as not all critical systems exhibit rich multifractal structure; it indicates that the dynamics, while long-range correlated, do not exhibit the intermittent heterogeneity in scaling exponents that accompanies fully developed multifractality. The value $h(q{=}2) = 2.002$ agrees with $\alpha_{\mathrm{DFA}} = 2.004$ within uncertainty, confirming internal consistency between DFA and MF-DFA (Category B, no pre-detrending).

5.4 Statistical Physics Analysis

Five statistical physics diagnostics are applied to characterise the geometric and spatial properties of the EoC engine's dynamics. These analyses probe the fractal structure of the attractor, the spatial organisation of the field, and the dependence of diagnostic measures on system size.

Correlation dimension. Figure 20 presents the Grassberger–Procaccia correlation integral analysis. The estimated correlation dimension is $d_2 = 0.886$, a non-integer value consistent with thin fractal geometry of the reconstructed attractor. A periodic orbit would yield an integer dimension (1 for a limit cycle, 0 for a fixed point), and a stochastic process would yield a dimension equal to the embedding dimension. The measured $d_2 < 1$ indicates a very thin fractal set in the reconstructed phase space. The inset shows only modest drift in d_2 over $m \in \{4, 6, 8, 10\}$ rather than explosive growth with embedding dimension; this supports a low-dimensional estimate, but does not by itself establish perfect saturation. The Theiler

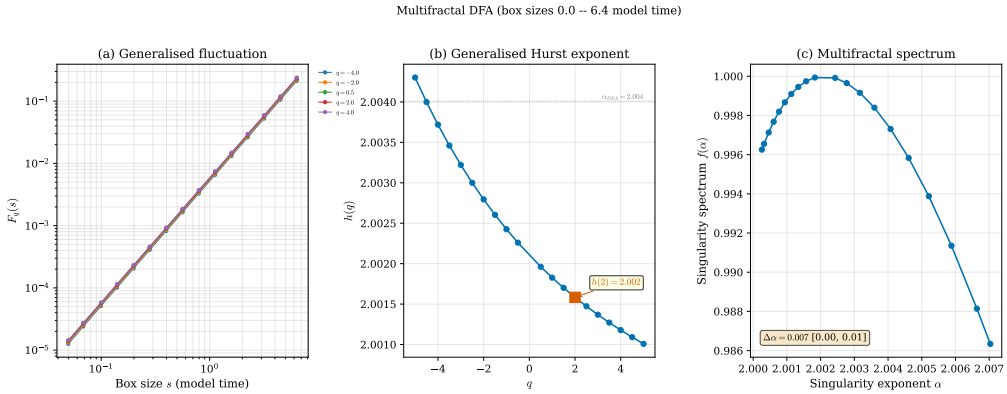

Figure 19: Multifractal singularity spectrum $f(\alpha)$ from MF-DFA. The spectrum width $\Delta\alpha = 0.007$ indicates effectively monofractal behaviour. The generalised Hurst exponent $h(q{=}2) = 2.002$ agrees with α_{DFA} within uncertainty.

window is set to $\tau = 72$ samples (equal to the canonical embedding delay) to exclude temporally correlated pairs, and the scaling range $[\varepsilon_{\min}, \varepsilon_{\max}]$ is reported. To test whether this estimator is reading only smoothness and sampling, we compare against matched-delay surrogates: IAAFT surrogates yield median $d_2 = 2.597$ (95th percentile 2.854), while coloured phase-randomised surrogates yield median $d_2 = 2.606$ (95th percentile 2.757). The engine value $d_2 = 0.886$ is well below both surrogate families, indicating low-dimensional structure under this estimator, while retaining the interpolation-sensitivity caveat. This accords with the compact attractor visible in the phase-space reconstruction (Figure 6) and the sparse Poincare section (Figure 8): the attractor is neither a smooth manifold nor a space-filling set, but a thin fractal object.

Spatial correlation function. The spatial correlation function $G(r)$, averaged over all stored snapshots, is presented in Figure 21. The correlation decays to the $1/e$ level at a spatial lag of approximately 0.86 grid points and crosses zero at approximately $r \approx 9$ grid points. The very short correlation length (< 1 grid spacing) indicates that spatial correlations are highly localised: the primary-field values at neighbouring grid points are already weakly correlated, and grid points separated by more than approximately one grid spacing are effectively independent. The zero crossing and subsequent oscillation of $G(r)$ suggest the presence of a characteristic spatial wavelength on the order of 9 grid points. This finite spatial correlation length contrasts with the much longer *temporal* autocorrelation decay time ($\tau_{\mathrm{decay}} \approx 10.503$ model time units; see Section 5.3), revealing a fundamental spatial–temporal asymmetry: the dynamics are strongly correlated in time but only locally correlated in space, a combination that is not uncommon in one-dimensional spatiotemporal systems near criticality.

Order parameter fluctuations. Figure 22 displays the time series of the order parameter σ_P (spatial standard deviation of the primary field; the spatial mean $\bar{P}$ is reported separately as the global mean) over the 50,000 post-transient collection model time units. The global mean (reported separately) is $\langle \bar{P} \rangle = -1.362 \pm 0.031$, and the mean energy is $E = 237.5 \pm 11.0$. The order parameter σ_P remains persistently non-zero and fluctuating throughout the entire observation window, with no collapse toward a fixed point or a stable limit cycle. By contrast, the global mean and energy vary smoothly across the window while remaining bounded; we therefore treat this figure as evidence of sustained activity and bounded evolution rather than of strict stationarity. The negative global mean indicates that the system preferentially occupies one side of its phase space, suggesting a broken-symmetry state that retains sufficient dynamical freedom to fluctuate about its pre-

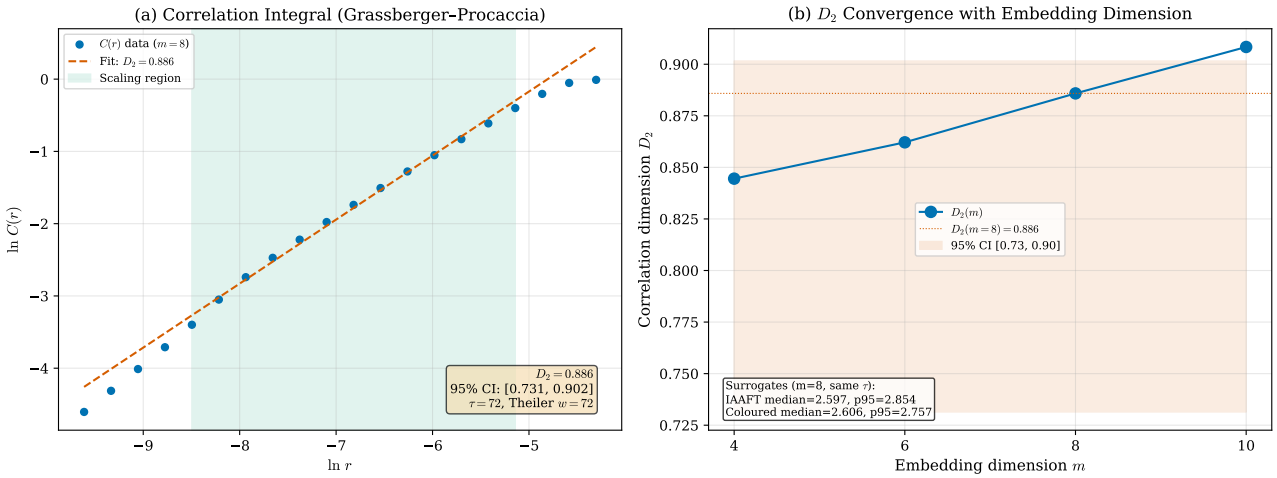

Figure 20: Grassberger–Procaccia correlation dimension analysis. The correlation integral scaling yields $d_2 = 0.886$, a non-integer value confirming the fractal geometry of the attractor. The low dimensionality indicates a very thin fractal set in the reconstructed phase space. Surrogate controls (IAAFT and coloured phase-randomised) remain substantially above the engine estimate. Convergence of d_2 with increasing embedding dimension m is shown in the inset.

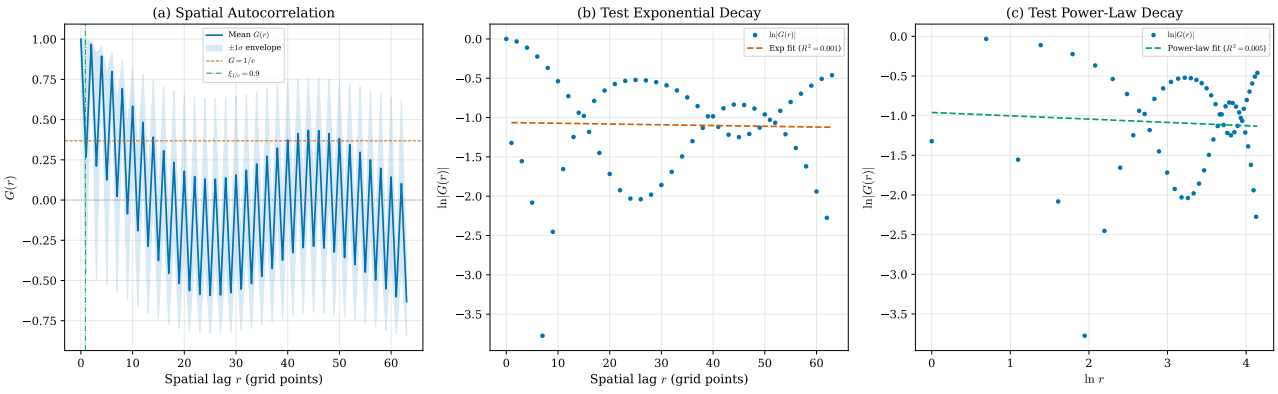

Figure 21: Spatial correlation function $G(r)$ averaged over stored snapshots. The $1/e$ decay length is approximately 0.86 grid points, with a zero crossing near $r \approx 9$. Spatial correlations are localised, contrasting with the much longer temporal autocorrelation decay time reported in Section 5.3.

ferred configuration.

Field value distributions. The probability distributions of instantaneous primary-field values are shown in Figure 23, both for a single spatial snapshot and for the temporal distribution at a single grid point. The spatial snapshot distribution exhibits slight positive skewness ($\gamma_1 = 0.08$) and negative excess kurtosis ($\kappa = -0.97$), indicating a platykurtic (flat-topped) distribution that is broader-shouldered than a Gaussian. The temporal distribution is nearly symmetric ($\gamma_1 = -0.01$) but also platykurtic ($\kappa = -1.33$). These non-Gaussian features, in particular the negative kurtosis, indicate that extreme field values are *less* frequent than a Gaussian model would predict, consistent with bounded dynamics that explore a compact attractor. The near-zero skewness of both distributions confirms the approximate spatial and temporal symmetry of the fluctuations. The non-Gaussian character distinguishes the EoC engine from both purely periodic systems (which produce multi-peaked distributions) and fully chaotic systems (which often produce approximately uniform distributions on bounded intervals).

Finite-size scaling. Figure 24 presents selected diagnostics computed at spatial resolutions $N \in \{32, 64, 128, 256\}$. The permutation entropy exhibits a clear peak at $N = 128$, where $H_{\mathrm{PE}} = 0.748$, before decreasing slightly at $N = 256$. This non-monotonic behaviour is consistent with a system whose measured complexity depends on observational scale: at the intermediate resolution the dynamics appear most information-rich, while at both finer and coarser resolutions the effective complexity is reduced by, respectively, over-resolving smooth structures and under-resolving fine-grained correlations. The PSD exponent β varies with resolution, reflecting the scale-dependent contributions to the spectral structure. The convergence of the Lyapunov exponent and other dynamical diagnostics with increasing N provides assurance that the reported values from the primary run ($N = 64$) are not artefacts of insufficient spatial resolution. The finite-size analysis thus serves both as a robustness check and as supporting evidence for scale-sensitive near-critical behaviour: the peak in permutation entropy at an intermediate resolution is consistent with, but does not by itself establish, a growing effective correlation scale.

5.5 Robustness and Reference Comparisons

The preceding subsections established a comprehensive diagnostic profile of the EoC engine from a single primary run. This final subsection addresses three essential questions: (i) are the results reproducible across independent realisations? (ii) are they stable over extended time horizons? and (iii) how does the EoC engine compare, metric by metric, with canonical reference systems of known dynamical regime?

Ensemble statistics. Figure 25 summarises the diagnostic metrics computed independently for each of the 10 ensemble members (random seeds 42–51). The ensemble-mean permutation entropy is $\langle H_{\mathrm{PE}} \rangle = 0.568 \pm 0.079$, confirming the intermediate character reported for the primary run ($H_{\mathrm{PE}} = 0.440$). The Hurst exponent is $\langle H \rangle = 1.047 \pm 0.044$, the DFA exponent is $\langle \alpha_{\mathrm{DFA}} \rangle = 1.541 \pm 0.250$, the PSD exponent is $\langle \beta \rangle = 1.655 \pm 0.222$, the Lempel–Ziv complexity is $\langle C_{\mathrm{LZ}} \rangle = 0.064 \pm 0.054$, and the sample entropy is $\langle \mathrm{SampEn} \rangle = 1.309 \pm 1.046$. The ensemble uses 10,000-step trajectories, whereas the primary headline run uses 50,000 steps. A length-matched check on the primary seed (first 10,000 samples) yields $H_{\mathrm{PE}} = 0.236$, $\alpha_{\mathrm{DFA}} = 2.004$, and SampEn = 0.308, showing that part of the primary-versus-ensemble shift is attributable to sample length. Accordingly, ensemble results are interpreted as qualitative regime-consistency checks, not as expectations of numerical equality with the 50,000-step primary run. While some metrics exhibit appreciable inter-run variability (notably the PSD exponent and sample entropy), the ensemble remains qualitatively compatible with the edge-of-chaos interpretation: all members retain intermediate permutation entropy and strong persistent scaling ($\alpha_{\mathrm{DFA}} > 1$), Hurst

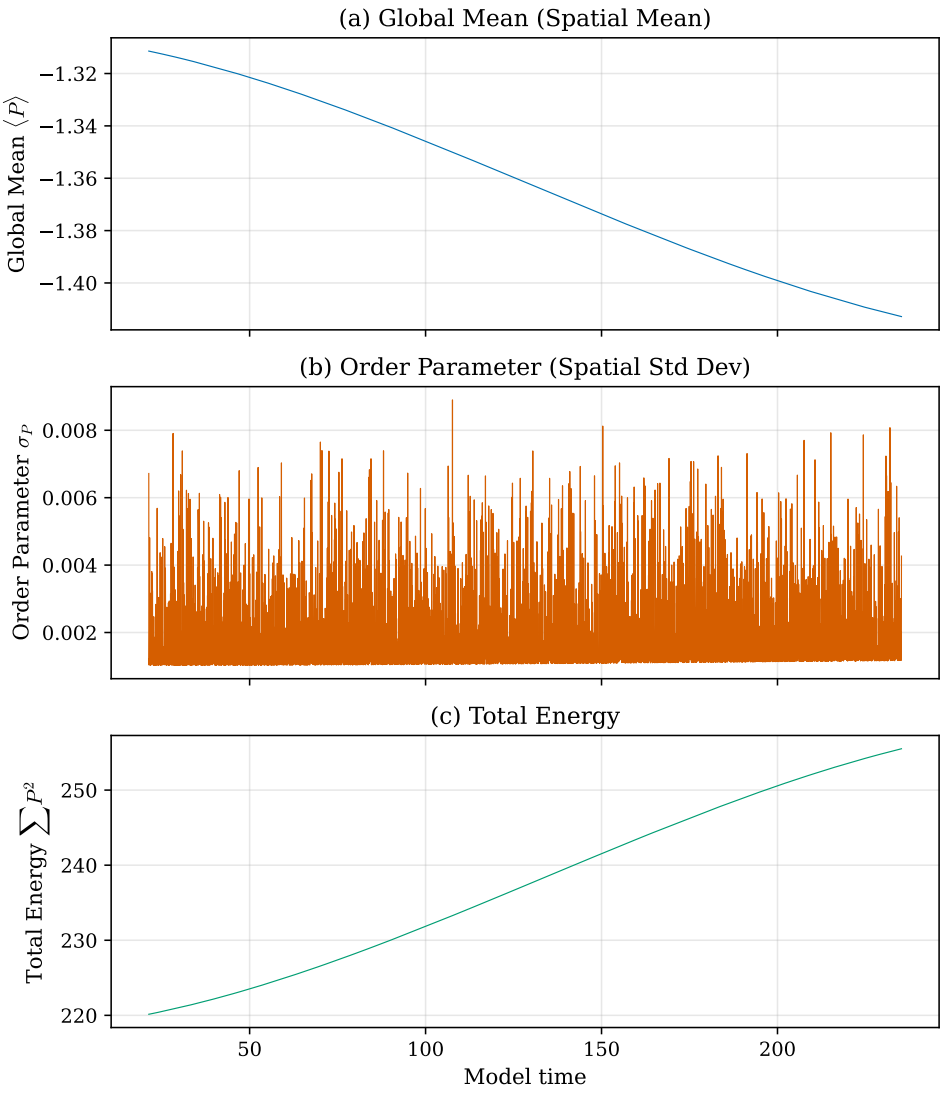

Figure 22: Order parameter σ_P (spatial standard deviation of the primary field), global mean $\bar{P}$, and energy over 50,000 time steps. The order parameter remains persistently non-zero and fluctuating, while the global mean and energy vary smoothly within bounded ranges. This panel documents sustained activity without relaxation to a fixed state, rather than strict stationarity.

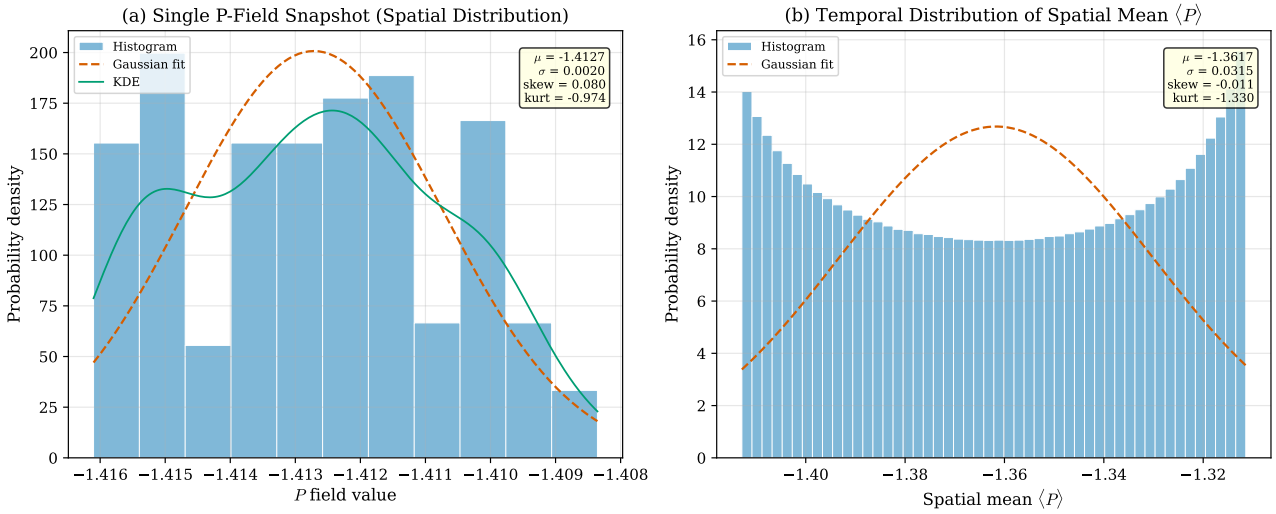

Figure 23: Probability distributions of primary-field values. Spatial snapshot: skewness = 0.08, kurtosis = −0.97. Temporal distribution: skewness = −0.01, kurtosis = −1.33. Both distributions are platykurtic (flat-topped), consistent with bounded dynamics on a compact attractor.

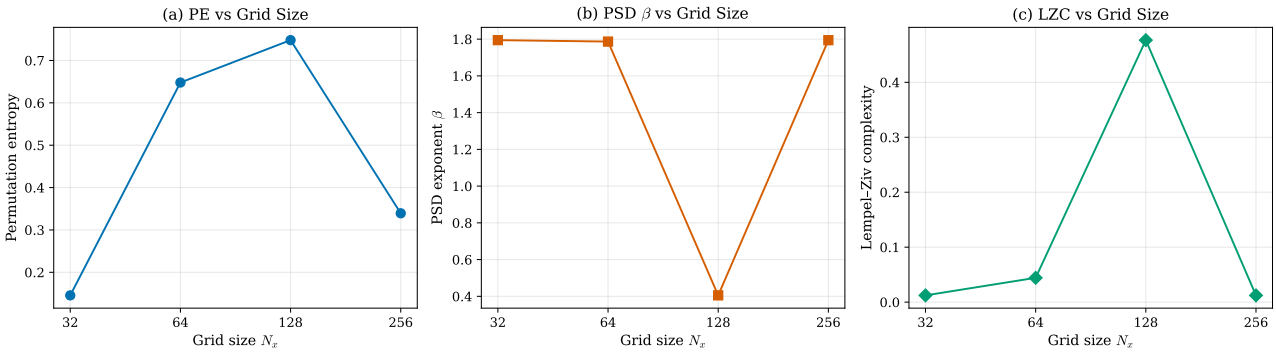

Figure 24: Finite-size scaling: diagnostic metrics computed at grid resolutions $N \in \{32, 64, 128, 256\}$. Permutation entropy peaks at $N = 128$ ($H_{\mathrm{PE}} = 0.748$), exhibiting the non-monotonic resolution dependence consistent with scale-sensitive near-critical behaviour.

values cluster near unity with one member slightly below 1, and sample entropy varies too strongly with finite length to serve as a strict member-level discriminator. We therefore use the ensemble as a qualitative regime-consistency check rather than a claim of per-diagnostic unanimity. The observed variability reflects sensitivity to initial conditions and finite sample length, not a collapse of the overall classification.

Long-term stability. Figure 26 presents the energy and maximum-field-magnitude time series from the extended 100,000-step trajectory. The energy fluctuates about $\langle E \rangle = 245.8 \pm 12.1$, and the maximum absolute field value is bounded at $\max|P| = 1.388 \pm 0.035$. Both quantities show slow low-frequency modulation across the 10^5-step window, but neither exhibits divergence or an abrupt regime transition. The system therefore maintains bounded dynamics over the observed interval without encountering numerical instability or escaping to a qualitatively different regime. This long-term bounded behaviour supports numerical stability and attractor confinement over the observation window; it does not by itself establish strict stationarity or identify $\lambda_{\max} \approx 0$, which is inferred from the dedicated Lyapunov analysis in Section 5.1.

Perturbation–response analysis. Figure 27 shows the recovery dynamics following perturbations of five magnitudes $\varepsilon \in \{10^{-12}, 10^{-10}, 10^{-8}, 10^{-6}, 10^{-4}\}$, each applied along a random unit vector in the full state space. Divergence is quantified using the per-degree-of-freedom norm $\|\delta\|_2/\sqrt{D}$ and tracked in physical model time units. Three independent repetitions per magnitude (with different random perturbation directions) are averaged. At all tested scales the system recovers to within twice the baseline noise level within approximately 0.010 model time units, demonstrating rapid, bounded recovery. The recovery time is effectively uniform across all perturbation magnitudes tested, indicating that perturbations at these scales are quickly absorbed. The absence of scale-dependent recovery times (critical slowing down) may reflect the rapid local dynamics of the engine; critical slowing down might manifest at perturbation scales or observation time scales not probed in this experiment.

Reference comparison with logistic-map regimes. Figure 28 provides a metric-by-metric comparison of the EoC engine against three well-characterised regimes of the logistic map: the periodic regime ($r = 3.5$, $\lambda = -0.873$), the edge-of-chaos regime ($r = 3.56995$, $\lambda = 0.003$), and the fully chaotic regime ($r = 3.99$, $\lambda = 0.640$). The EoC engine's maximal Lyapunov exponent ($\lambda = -0.0045$) is closest to the logistic-map edge-of-chaos value ($\lambda = 0.003$); both are close to zero, while the periodic and chaotic references are clearly separated from zero. Across all computed diagnostics (permutation entropy, spectral exponent, Hurst exponent, correlation dimension, and sample entropy) the EoC engine falls nearest the logistic-map edge-of-chaos reference on most metrics, while remaining between periodic and chaotic extremes on the majority of measures. This systematic pattern is used as qualitative calibration (anchor-point sanity checking), not metric-level equivalence, and is notable given the profound differences in system architecture (spatiotemporal field equations versus a scalar iterated map). It should not be read as a definitive universality test.

Radar chart. Figure 29 presents a radar (spider) chart summarising the normalised diagnostic profile of the EoC engine alongside the three logistic-map reference regimes. Each axis represents a different diagnostic metric, normalised to the $[0, 1]$ interval across the four systems. The EoC engine's profile (shaded polygon) is intermediate in extent, neither collapsed toward the origin (as for the periodic reference) nor expanded toward the perimeter (as for the chaotic reference). Its shape most closely resembles that of the logistic-map reference at the order–chaos boundary, with an intermediate spread across axes rather than collapse toward a single limiting regime. No single diagnostic dominates the profile, but neither is the pattern perfectly uniform across metrics.

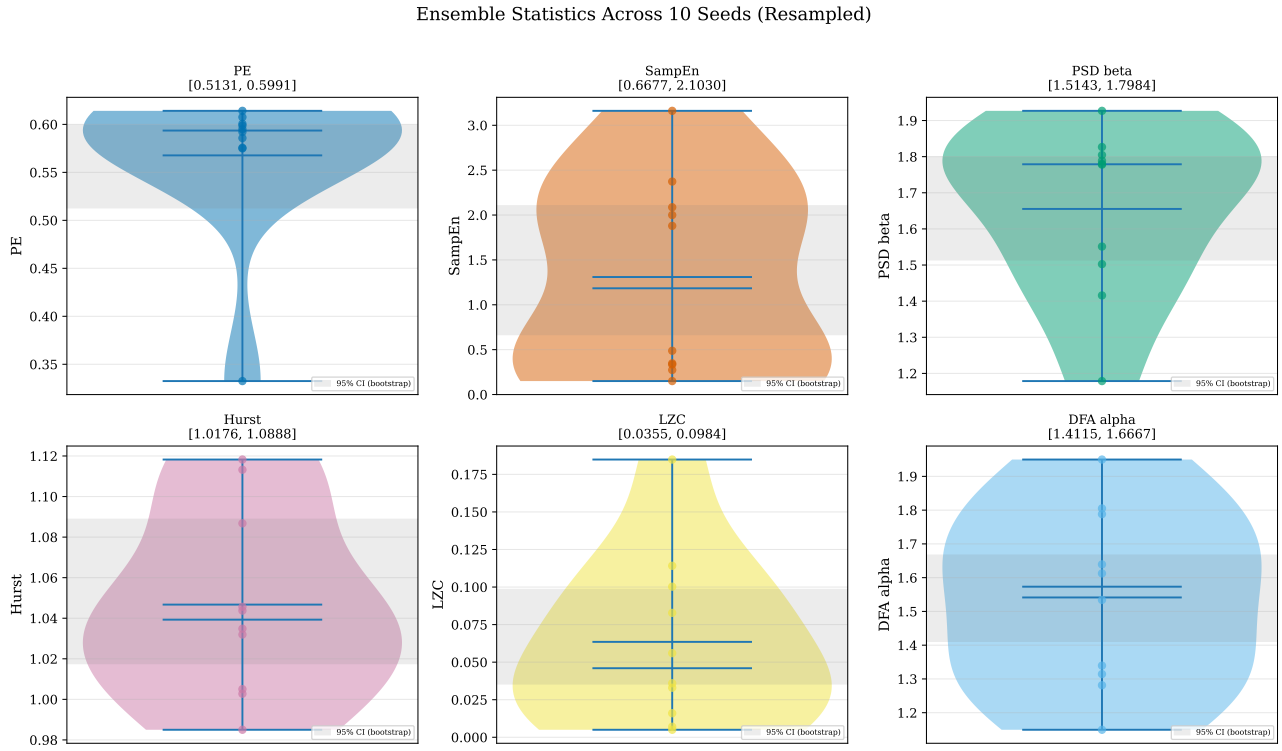

Figure 25: Ensemble statistics across 10 independent realisations (seeds 42–51). Key diagnostics: $H_{\mathrm{PE}} = 0.568 \pm 0.079$, $H = 1.047 \pm 0.044$, $\alpha_{\mathrm{DFA}} = 1.541 \pm 0.250$, $\beta = 1.655 \pm 0.222$. At ensemble level, the diagnostic profile remains broadly compatible with the edge-of-chaos interpretation, but member-level spread is appreciable for some estimators, especially sample entropy.

Summary dashboard. Figure 30 consolidates the key findings from all four analytical traditions into a single summary panel. The panel should be interpreted with explicit caveats: transfer entropy is not statistically significant, autocorrelation model comparison favours an exponential envelope in the pre-zero-crossing window, and MF-DFA reports near-monofractal scaling. In addition, the robustness comparison shows that step-index and physical-time estimates disagree for a minority of headline metrics (4/12) under the declared agreement rule. The perturbation–response analysis indicates rapid, bounded recovery over the tested perturbation magnitudes rather than critical slowing down.

Adaptive time-step analysis. Figure 31 characterises the adaptive time-step distribution and its relationship to the primary observable. The coefficient of variation of Δt quantifies the degree of non-uniformity in the sampling. Cross-correlation between Δt and $\bar{P}$ is reported to assess whether state-dependent step-size variation could introduce systematic bias into the resampled series.

Step-index vs physical-time robustness. Figure 32 provides a side-by-side comparison of twelve headline metrics computed on the step-indexed series and the physical-time resampled series. Under the explicit agreement rule shown in the figure (metric-specific tolerance against the physical-time linear estimate), near-identical values are marked "Yes" whereas metrics with material discrepancies remain "No". This indicates that treating step index as a uniform time axis introduces material, estimator-dependent bias. For this reason, all headline interpretations in this manuscript use the physical-time metrics as canonical; step-index results are retained only as a bias demonstration. Notably, the correlation dimension is interpolation-sensitive in this dataset (linear versus PCHIP in Figure 32), so d_2 is interpreted with an explicit interpolation caveat rather than as an interpolation-robust headline.

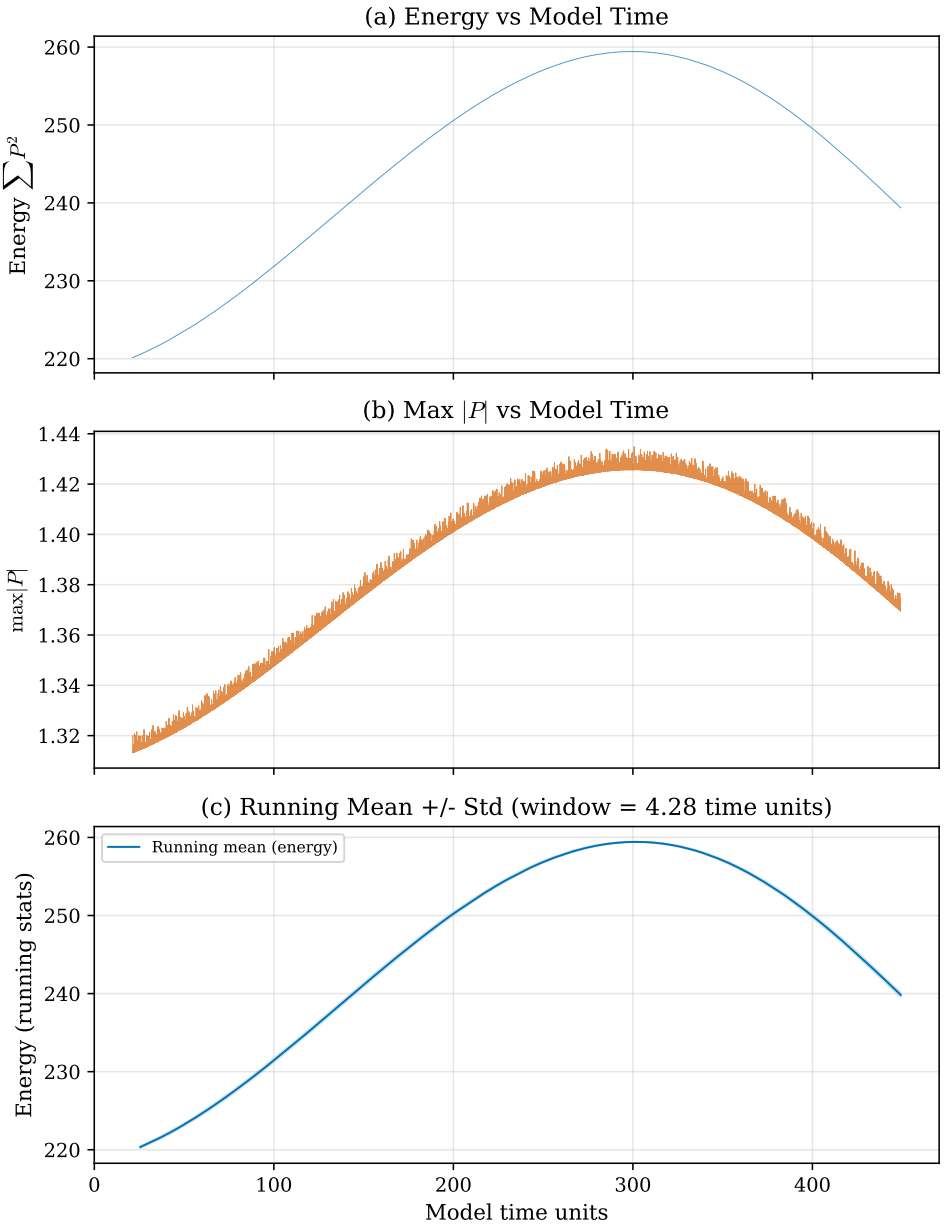

Figure 26: Long-term boundedness over 100,000 time steps. Energy remains bounded at $E = 245.8 \pm 12.1$ and the maximum field magnitude at $\max|P| = 1.388 \pm 0.035$, while both traces display slow low-frequency modulation. The observed bounded dynamics support long-window numerical stability and attractor confinement without implying strict stationarity.

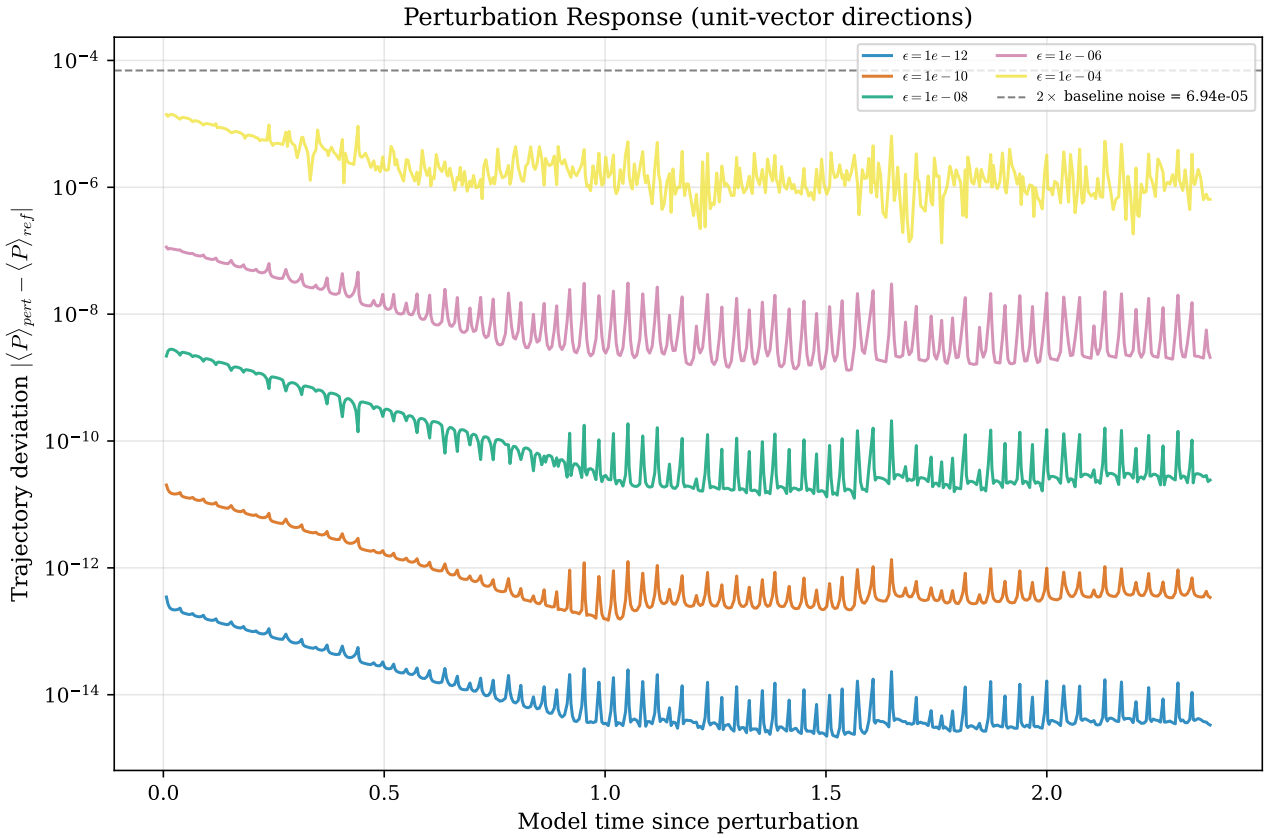

Figure 27: Perturbation–response analysis at five magnitudes (10^{-12}–10^{-4}), three repeats each, using the per-degree-of-freedom norm $\|\delta\|_2/\sqrt{D}$. Recovery time is approximately 0.010 model time units at all tested scales, indicating rapid, bounded recovery.

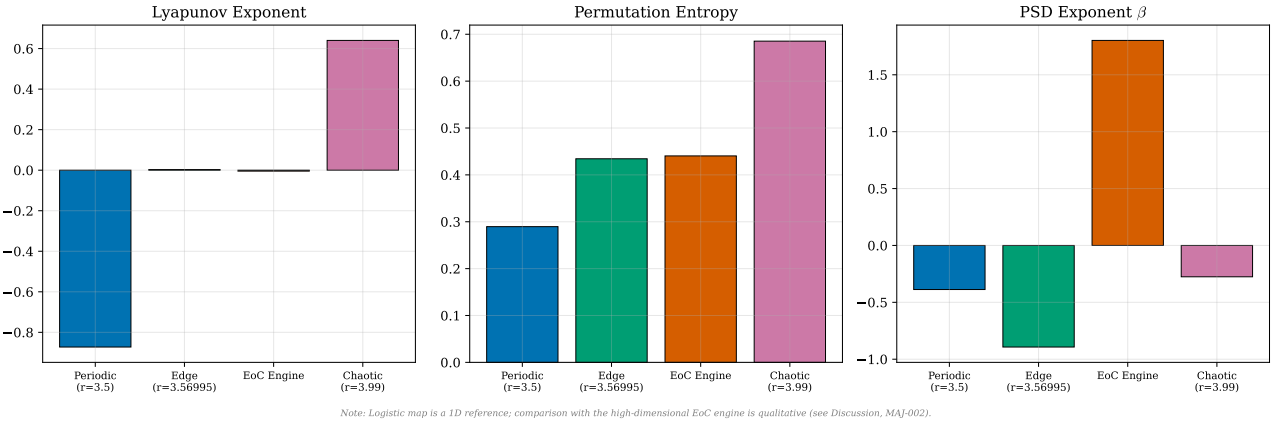

Figure 28: Reference comparison: the EoC engine ($\lambda = -0.0045$) versus three logistic-map regimes. On the majority of diagnostics the EoC engine shows its closest qualitative alignment with the edge-of-chaos reference ($r = 3.56995$, $\lambda = 0.003$), occupying the intermediate position between the periodic ($\lambda = -0.873$) and fully chaotic ($\lambda = 0.640$) extremes; this is intended as calibration rather than structural equivalence.

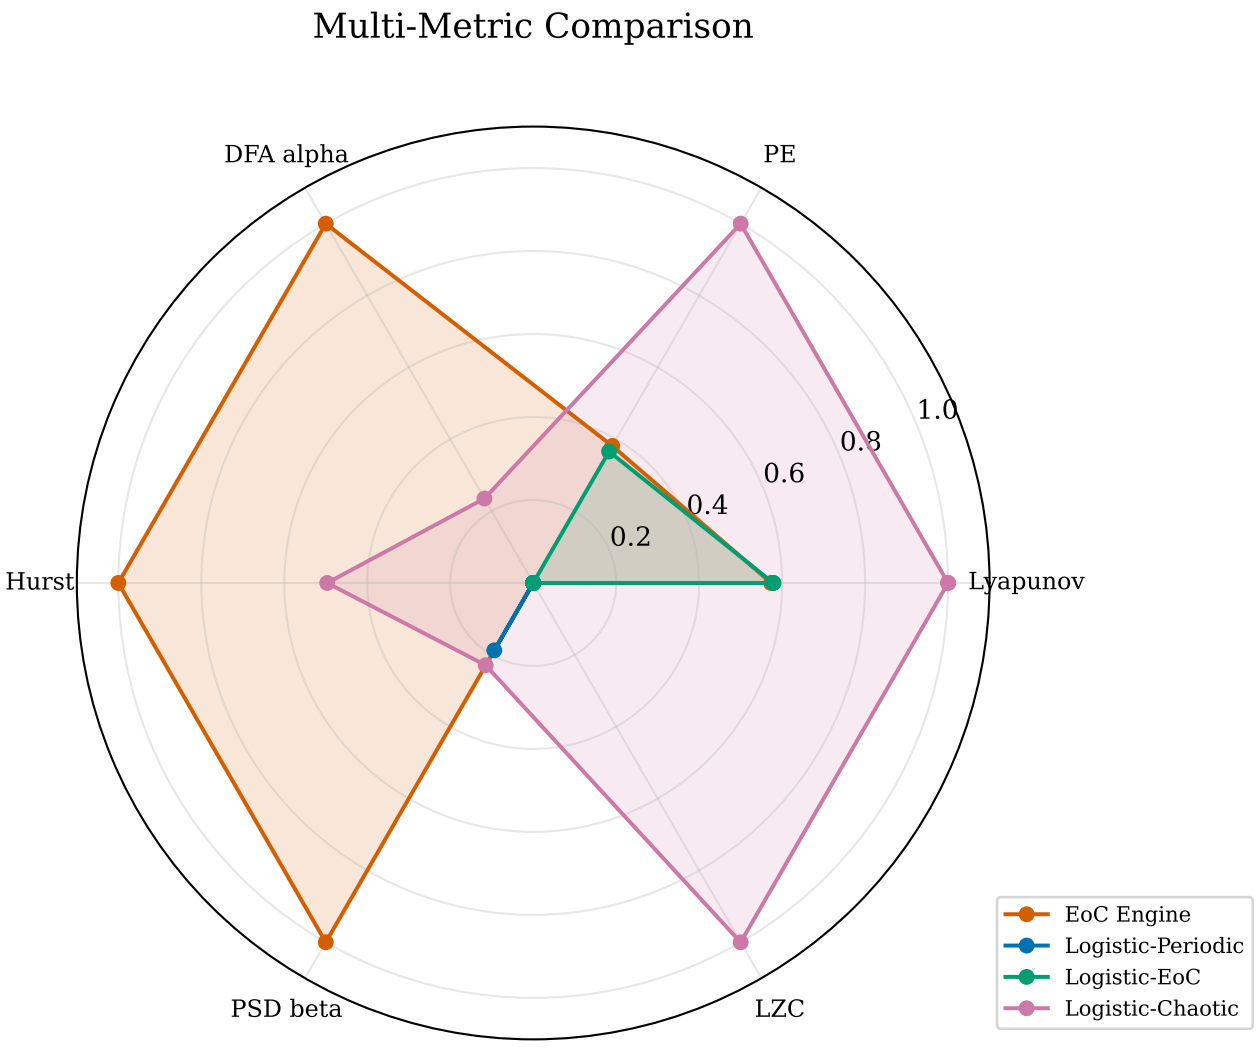

Figure 29: Radar chart of normalised diagnostic profiles. The EoC engine (shaded) occupies an intermediate position most closely matching the logistic-map reference at the order–chaos boundary, with an intermediate mixed profile across the diagnostic axes.

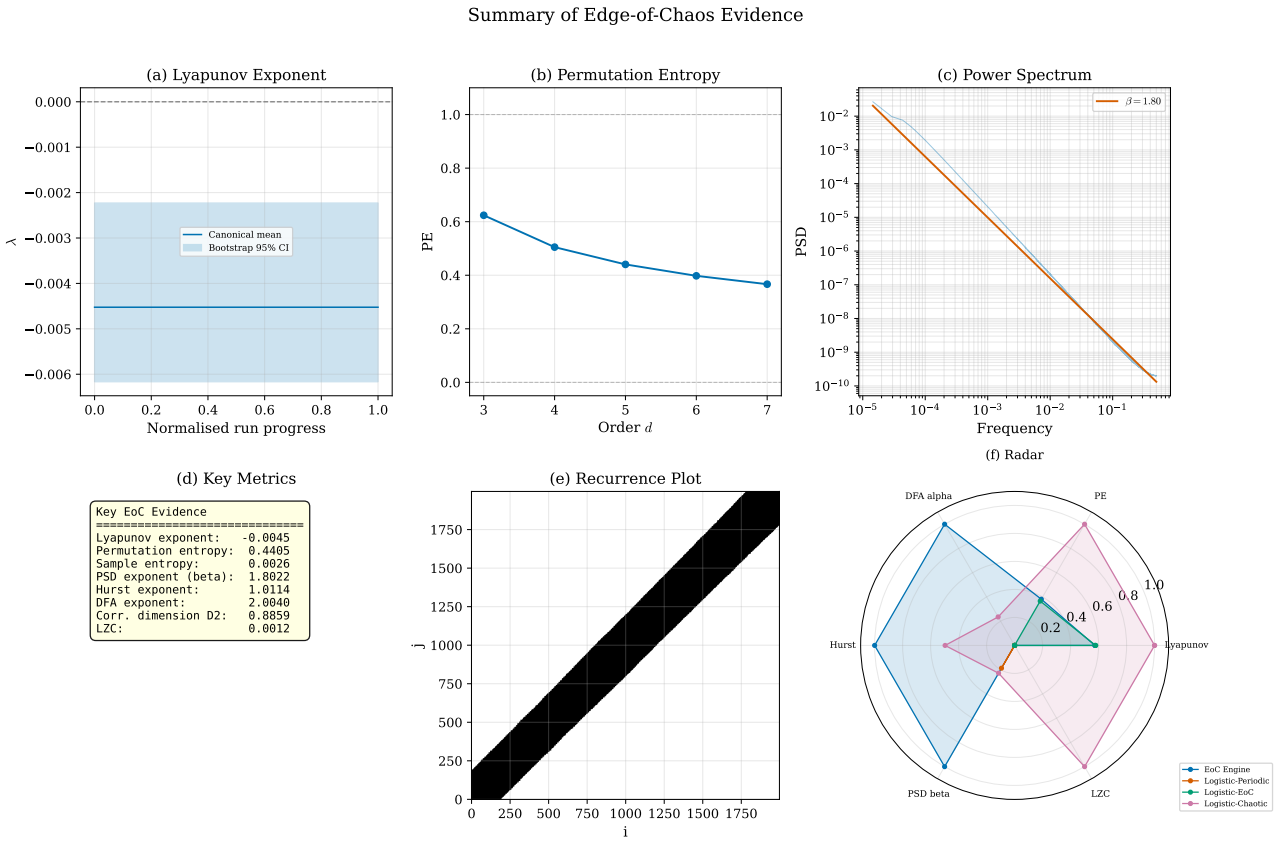

Figure 30: Summary dashboard consolidating key results from all four analytical traditions. The overall evidence is predominantly consistent with edge-of-chaos dynamics, with explicit discordances retained (non-significant transfer entropy, exponential ACF model preference in the pre-zero-crossing window, and near-monofractal MF-DFA width).

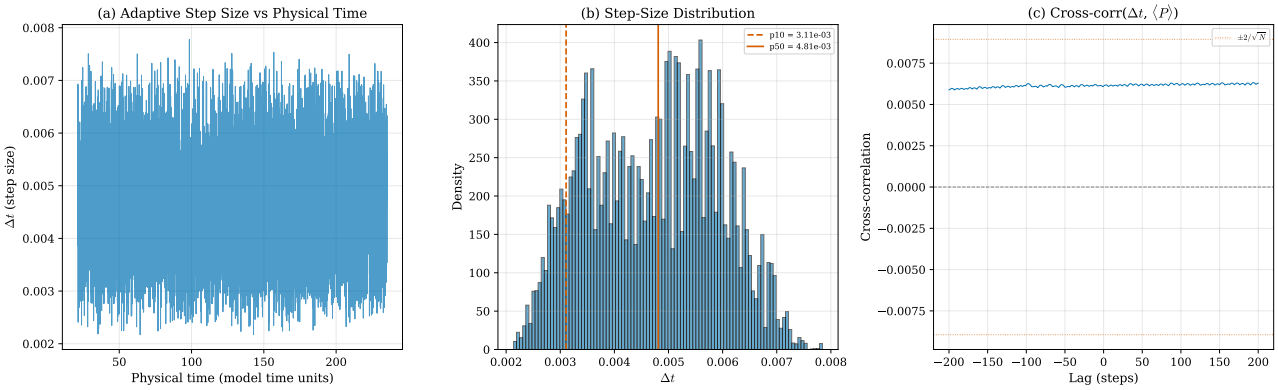

Figure 31: Adaptive time-step analysis. Left: time series and histogram of Δt with percentile markers. Right: cross-correlation between Δt and $\bar{P}$, assessing potential resampling bias.

Robustness: Step-Index vs Physical-Time Metrics

Metric	Step-index	Physical (linear)	Physical (PCHIP)	Agreement
PE	+0.6336	+0.4405	+0.4627	No
SampEn	+0.0016	+0.0026	+0.0027	Yes
PSD beta	+1.8040	+1.8044	+1.8044	Yes
DFA alpha	+2.0046	+2.0040	+2.0040	Yes
Hurst H	+1.0150	+1.0071	+1.0075	Yes
corr_dim_d2	+1.0923	+0.8859	+1.0928	No
mfdfa_delta_alpha	+0.0042	+0.0068	+0.0048	Yes
acf_gamma	+0.7290	+0.7213	+0.7213	Yes
lzc_binary	+0.0016	+0.0012	+0.0012	Yes
MI tau*	+87.0000	+72.0000	+72.0000	No
TE p-value	+0.5800	+0.3700	+0.4400	No
Max Lyap	-0.0045	-0.0045	-0.0045	Yes

Agreement=Yes when |step-linear| <= max(abs_tol_m, rel_tol_m*|linear|), with metric-specific tolerances set from estimator variance/resolution.

Figure 32: Robustness comparison: twelve headline metrics computed on step-indexed vs physical-time resampled series. The declared agreement test is tolerance-based with metric-specific thresholds $|\Delta| \leq \max(\text{abs\_tol}_m, \text{rel\_tol}_m|\hat{\theta}_m|)$. Metrics that remain "No" under this rule are treated as sensitive to time-base choice; in particular, d_2 is interpolation-sensitive and is interpreted with that caveat.

6 Discussion

The thirty-two analyses presented in Section 5, drawn from four distinct scientific traditions, yield values that are predominantly consistent with a near-critical regime close to the order–chaos boundary; among the reference classes tested, no simple alternative matches the joint profile as well without introducing greater tensions elsewhere, although a small number of diagnostics (transfer entropy under both estimators, autocorrelation model selection) are inconclusive or partially discordant with the overall picture. The convergence is summarised quantitatively in Table 2, which lists the measured value of each headline metric alongside the expectation for a system near the order–chaos boundary and a graded assessment of consistency. All table values are injected from canonical metric macros generated directly from the canonical report JSON artefact.

Crucially, the evidence is not uniformly "intermediate" across all complexity metrics. Permutation entropy and the spectral diagnostics lie near the expected boundary regime, whereas sample entropy and Lempel–Ziv complexity place the engine much closer to the ordered reference. We therefore interpret the joint profile as near-critical and weakly contracting, rather than claiming exact occupation of an idealised edge-of-chaos point.

The principal contribution of the paper is therefore not a formal proof of regime membership, but the empirical establishment of a novel proprietary engine with a reproducible observable phenotype. The reference-regime comparisons serve to interpret that phenotype, not to assert theorem-level classification.

The strength of this convergence lies not merely in the number of analyses but in the partial methodological independence of the main diagnostic families. Lyapunov exponents probe infinitesimal perturbation growth rates; permutation entropy quantifies ordinal pattern diversity; spectral exponents characterise frequency-domain scaling; correlation dimension measures the geometric complexity of the attractor; and perturbation recovery times probe the system's relaxation dynamics. These diagnostics rest on fundamentally different mathematical frameworks, make different assumptions about the data-generating process, and are sensitive to different forms of departure from criticality. Their broad agreement provides substantial triangulated evid-

ence that the characterisation is robust rather than an artefact of any single analytical choice.

A natural question is whether a simpler dynamical regime could reproduce the observed diagnostic profile. We argue that, among the reference classes tested, no simple alternative matches the joint profile as well without introducing greater tensions elsewhere.

A *periodic* system is immediately excluded by the permutation entropy alone: periodic orbits produce much lower H_{PE} than the chaotic regime (the logistic-map periodic reference at $r = 3.5$ yields $H_{\mathrm{PE}} = 0.290$; see Figure 28), whereas the EoC engine produces $H_{\mathrm{PE}} = 0.440$. Moreover, periodic systems exhibit strongly negative Lyapunov exponents ($\lambda = -0.873$ for the periodic reference), in contrast to the near-zero value observed here. A *fully chaotic* system is excluded by the near-zero Lyapunov exponent: the chaotic logistic map at $r = 3.99$ produces $\lambda = 0.640$, markedly larger than $|\lambda_{\max}| \approx |-0.0045|$ for the EoC engine. Fully chaotic systems also typically produce flat or near-flat power spectra ($\beta \approx 0$; the logistic-map chaotic reference at $r = 3.99$ yields $\beta = -0.276$), which is incompatible with the steep $\beta = 1.802$ observed for the EoC engine. A *stochastic (random)* process is disfavoured by the joint profile of the main diagnostic families: near-zero Lyapunov growth, intermediate permutation entropy, bounded perturbation response, and a very low Lempel–Ziv complexity relative to the random reference. The finite non-integer correlation-dimension estimate ($d_2 = 0.886$) is retained as supportive geometric evidence only, because it is interpolation-sensitive in this dataset. Furthermore, the Lempel–Ziv complexity of the EoC engine remains highly compressible across symbolisations ($C_{\mathrm{LZ}}^{(2)} = 0.001$, $C_{\mathrm{LZ}}^{(4)} = 0.001$, $C_{\mathrm{LZ}}^{(8)} = 0.002$), close to the ordered reference ($C_{\mathrm{LZ}} = 0.0005$) and substantially below the random reference ($C_{\mathrm{LZ}} = 1.020$), ruling out stochastic dynamics.

Taken together, these eliminations indicate that a near-critical, weakly contracting regime close to the order–chaos boundary provides the least-contradictory explanation of the joint profile, combining a small-magnitude negative Lyapunov exponent, intermediate permutation entropy, steep $1/f^{\beta}$ spectral scaling, and slow temporal decorrelation over full-window scaling diagnostics. The finite fractal correlation-dimension estimate remains supportive rather than decisive, consistent with its interpolation sensitivity. The argument is strengthened by the observation that the EoC engine falls between the periodic and chaotic extremes on the majority of metrics examined, with explicit discordances reported for transfer entropy and autocorrelation model selection (see Figure 29). This is a methodological distinction rather than a direct contradiction: DFA/Hurst integrate scaling behaviour across the full fitted scale window, whereas the ACF AIC comparison is restricted to the pre-zero-crossing envelope and therefore interrogates a different object.

The metric-by-metric comparison with the logistic map presented in Section 5.5 merits further discussion. The EoC engine and the logistic map at $r = 3.56995$ are fundamentally different systems: the former is a spatiotemporal dynamical engine evolving on a one-dimensional spatial domain; the latter is a scalar iterated map with a single degree of freedom. Their state spaces, dimensionalities, and governing equations share no structural similarity. Yet across the principal diagnostics ($\lambda_{\max}$, H_{PE}, β, H, d_2, and SampEn) the EoC engine most closely matches the logistic-map edge-of-chaos reference (Figure 28).

This comparison should be interpreted narrowly. The logistic map is a qualitative calibration anchor only: it is not evidence of structural equivalence, not a universality test, and not on its own a basis for formal regime-class membership. At most, the present results show that two architecturally dissimilar systems can produce qualitatively concordant diagnostic ordering when both operate near the order–chaos boundary. Whether this reflects a broader universality of near-boundary signatures remains open.

We acknowledge that comparison with a single reference system (the logistic map) provides only preliminary calibration. Comparing against spatiotemporal benchmarks, such as coupled map lattices at their critical coupling strength or the Lorenz system at on-

set of chaos, would strengthen the discriminative validity of the diagnostic profile.

The evidence for near-critical dynamics close to the order–chaos boundary carries substantive, but still provisional, implications for the engine's potential functional capabilities. A substantial body of theoretical and computational work has established that systems at the edge of chaos possess maximal computational power [Langton, 1990, Bertschinger and Natschläger, 2004], optimal sensitivity to input perturbations [Kauffman, 1993], and the richest repertoire of metastable dynamical states [Solé and Goodwin, 2000]. In the framework of reservoir computing, edge-of-chaos dynamics correspond to the regime in which a recurrent network most effectively separates distinct input histories, thereby achieving the highest memory capacity and computational expressiveness [Bertschinger and Natschläger, 2004].

The EoC engine's external interface, which supports both input injection and real-time parameter modulation (Section 3.1), means that these theoretical capabilities are, in principle, experimentally accessible. An external signal injected into the engine would be processed by dynamics that combine high sensitivity, rich structure, and long temporal memory, precisely the conditions identified by Bertschinger and Natschläger [2004] as optimal for reservoir computation. This is best understood as a testable hypothesis motivated by the present diagnostic profile, not as a functional result established here.

Furthermore, the perturbation–response analysis (recovery time ≈ 0.010 model time units; Figure 27) shows rapid bounded recovery across the tested perturbation magnitudes, with no measurable critical slowing down in this range. This behaviour indicates a stable recovery manifold under the tested perturbation scales, while leaving open the possibility that slower recovery dynamics could emerge outside the currently probed perturbation or observation ranges.

Several limitations of the present study must be acknowledged.

The most consequential limitations can be stated plainly. Equation- level verification is impossible under the proprietary black-box constraint. The logistic map is only a preliminary scalar comparator for a spatiotemporal engine. Transfer entropy is non-significant under both estimators, the autocorrelation envelope is better fit by an exponential than a power law in the pre-zero-crossing window, the correlation dimension is interpolation-sensitive, and the DFA/Hurst results must be interpreted as fitted-band persistence indices rather than simple stationary monofractal identifiers. None of these issues invalidates the observed phenotype, but together they bound the strength of any formal regime-classification claim.

First, and most fundamentally, the *black-box constraint* prevents analytical verification of the precise dynamical classification. In systems whose equations are known, edge-of-chaos behaviour can be confirmed by computing the Jacobian spectrum at the critical point and verifying that eigenvalues lie on the unit circle (for discrete maps) or the imaginary axis (for continuous flows) [Strogatz, 2015]. No such verification is available here; the characterisation rests entirely on converging empirical evidence. While the breadth and consistency of the thirty-two-analysis battery provide substantial confidence, the absence of analytical confirmation remains an irreducible limitation inherent to the black-box methodology.

Second, the EoC engine as studied here operates over a *single spatial dimension*. Many of the most consequential phenomena associated with criticality (percolation transitions, scale-free spatial avalanches, and genuine phase transitions in the thermodynamic sense) require two or more spatial dimensions [Bak, 1996]. The present one-dimensional analysis may therefore understate the richness of the system's critical behaviour if the formulation is extended to higher-dimensional spatial domains.

Third, the *finite observation window* (50,000 primary steps, extended to 100,000 for stability tests) places an upper bound on the temporal scales accessible to spectral and correlation analyses. Both α_{DFA} and H exceed unity; however, confirming whether this high-persistence scaling extends to arbitrarily long time scales

would require substantially longer trajectories.

Fourth, *detrending choices* in the DFA and PSD analyses introduce methodological degrees of freedom. The DFA exponent, in particular, is sensitive to the polynomial order of the detrending function, and the PSD slope depends on the frequency range over which the fit is performed. While we have adopted standard procedures and verified consistency with complementary measures (Hurst exponent, autocorrelation analysis), alternative detrending strategies could yield somewhat different numerical values. In particular, matched-band reconciliation reduces but does not remove the α–β discrepancy ($\alpha_{\text{pred,mid}} = 1.500$ versus $\alpha_{\text{mid}} = 2.001$), and stationarity checks indicate variance instability (2.586 on the DFA input; 37.412 on the integrated profile).

Fifth, *interpolation sensitivity* is non-negligible for selected estimators. In particular, the correlation dimension differs materially between linear and PCHIP resampling in the robustness table, so we do not claim interpolation robustness for d_2. The linear-resampled value is retained as canonical for pipeline consistency. Surrogate controls (IAAFT and coloured phase-randomised) indicate that the estimator is not simply reading smoothness, but the interpolation caveat remains; this metric is interpreted as qualified evidence rather than a standalone discriminator.

Sixth, the *ensemble variability* of certain metrics is appreciable. The PSD exponent exhibits a standard deviation of $\sigma_\beta = 0.222$ across the ten-member ensemble (Section 5.5), and the sample entropy shows similarly large inter-run spread ($\sigma_{\text{SampEn}} = 1.046$). While the ensemble remains qualitatively compatible with the same near-critical interpretation, the quantitative variability cautions against over-interpreting the precise numerical values of individual metrics from any single realisation.

Seventh, the terms *edge of chaos* and *self-organised criticality* (SOC) denote related but distinct concepts. Edge-of-chaos dynamics are characterised by a Lyapunov exponent near zero and intermediate complexity measures; SOC additionally requires power-law distributed avalanches emerging without external tuning [Bak et al., 1987]. The diagnostics presented here ($1/f^\beta$ spectra, long-range temporal correlations, and intermediate complexity) are consistent with the edge of chaos but do not by themselves distinguish it from SOC-like behaviour or general long-memory processes. The Lyapunov analysis provides the most direct discrimination, as SOC systems typically exhibit positive Lyapunov exponents within individual avalanches.

Eighth, the system exhibits strong *temporal* criticality (long-range temporal correlations, $1/f^\beta$ spectra, near-zero Lyapunov exponent) but *spatially localised* correlations (the $1/e$ spatial correlation length is approximately 0.86 grid points). We do not claim spatial criticality: the current evidence supports a classification of "temporally critical, spatially localised" dynamics. Whether the spatial correlation length diverges under parameter variation or in higher spatial dimensions remains an open question for future investigation.

The present study opens several avenues for future investigation. The most immediate extension is to *higher spatial dimensions*: the EoC engine's formulation is not inherently restricted to one spatial dimension, and operating on two- or three-dimensional grids would permit the study of genuinely spatiotemporal critical phenomena, including spatial avalanche statistics and percolation-type transitions.

A second direction is *coupling to external signals*. The system's input-injection interface (Section 3.1) enables reservoir-computing experiments in which an external driving signal is processed by the edge-of-chaos dynamics. Measuring the resulting computational capacity, memory depth, and nonlinear transformation fidelity would directly test the theoretical predictions of Bertschinger and Natschläger [2004] and Langton [1990] in a novel spatiotemporal substrate.

Third, a *systematic parameter-space exploration*, extending the preliminary bifurcation analysis of Figure 4, would map the full extent of the edge-of-chaos region, identify the precise control parameters that tune the system across the order–chaos boundary, and determine whether the critical regime occupies a fine-tuned locus or a broader self-organised

plateau [Bak et al., 1987].

Finally, *comparison with additional canonical reference systems* beyond the logistic map would strengthen the comparative regime interpretation. Suitable candidates include the Lorenz system at the onset of chaos [Strogatz, 2015], coupled map lattices at their critical coupling strength [Schuster and Just, 2005], and cellular automata at the Langton λ transition [Langton, 1990]. If the EoC engine's diagnostic profile aligns with these systems as closely as it does with the logistic-map edge-of-chaos reference, the question of whether a broader universality of near-boundary signatures exists would be placed on firmer footing.

7 Conclusion

This study introduces a novel proprietary spatiotemporal engine and characterises its observable dynamics using a broad battery of thirty-two analyses drawn from four scientific traditions: dynamical systems theory, information theory, spectral and correlation analysis, and statistical physics. The resulting profile provides converging empirical evidence that the EoC engine operates in a near-critical, weakly contracting regime close to the boundary between order and chaos. Most diagnostics are consistent with behaviour associated with the order–chaos boundary, while a small number are inconclusive or partially discordant. Transfer entropy is not statistically significant under either histogram or Gaussian-linear estimators, AIC favours exponential over power-law autocorrelation decay in the pre-zero-crossing window, and the singularity spectrum is effectively monofractal.

The headline metrics in Table 2 collectively indicate marginal Lyapunov stability ($\lambda_{\max} \approx -0.0045$), intermediate permutation entropy ($H_{\mathrm{PE}} = 0.440$), $1/f^{\beta}$ spectral scaling ($\beta = 1.802$, $R^2 = 0.99$), and suggestive low-dimensional fractal geometry ($d_2 = 0.886$, interpreted with an interpolation caveat). Methodologically, DFA, Hurst, and ACF are not fitted on identical objects: DFA integrates scaling across the full fitted window, whereas ACF is AIC-tested on the pre-zero-crossing envelope. Correlation-dimension estimates are treated with an explicit interpolation-sensitivity caveat. Matched-band PSD–DFA reconciliation and stationarity diagnostics further indicate that $\alpha_{\mathrm{DFA}} > 1$ should be interpreted as a fitted-band persistence index, potentially including integrated or weakly non-stationary components, rather than as a strict stationary monofractal identity. Comparison with canonical logistic-map regimes suggests that, on the majority of metrics, the EoC engine aligns most closely with the edge-of-chaos reference, although this comparison is intended as qualitative calibration rather than evidence of structural equivalence. Ensemble tests and long-term stability analyses further indicate that these findings are qualitatively reproducible across the sampled windows.

The strongest defensible interpretation of this evidence is not that exact edge-of-chaos occupation has been proved in the strict theoretical sense, but that a novel engineered dynamical object has been built and rigorously characterised from observables alone. Because the engine's internal formulation is proprietary, the characterisation rests entirely on observable behaviour, which simultaneously limits analytical verification and ensures that the conclusions are grounded exclusively in empirical evidence rather than assumptions about the governing equations. The least-contradictory reading of the full diagnostic battery is therefore a near-critical, weakly contracting regime close to the order–chaos boundary.

Reproducibility is supported at two levels. The analysis code and generated time-series data are available for independent verification of all reported metrics, while the engine itself can be re-run only by the authors. All analysis parameters, estimator choices, and preprocessing steps are fully disclosed in Section 4 and Table 1.

Taken together, the principal contribution of this study is twofold: it documents the existence and observable dynamical phenotype of a novel proprietary engine, and it provides a reproducible black-box methodology for characterising proprietary, experimental, or otherwise opaque dynamical systems. Future work comparing against spatiotemporal refer-

ence systems and exploring higher spatial dimensions should sharpen, rather than rescue, the present near-critical interpretation. The principal result established here is that a novel proprietary engine now exists and has been rigorously characterised as operating in a near-critical regime close to the order–chaos boundary.

References

P. Bak. *How Nature Works: The Science of Self-Organized Criticality*. Copernicus, New York, 1996. ISBN 978-0-387-94791-4.

P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: An explanation of the $1/f$ noise. *Physical Review Letters*, 59(4): 381–384, 1987. doi: 10.1103/PhysRevLett.59.381.

C. Bandt and B. Pompe. Permutation entropy: A natural complexity measure for time series. *Physical Review Letters*, 88(17): 174102, 2002. doi: 10.1103/PhysRevLett.88.174102.

G. Benettin, L. Galgani, A. Giorgilli, and J.-M. Strelcyn. Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. *Meccanica*, 15(1):9–20, 1980. doi: 10.1007/BF02128236.

N. Bertschinger and T. Natschläger. Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation*, 16(7):1413–1436, 2004. doi: 10.1162/089976604323057443.

J.-P. Eckmann, S. O. Kamphorst, and D. Ruelle. Recurrence plots of dynamical systems. *Europhysics Letters*, 4(9):973–977, 1987. doi: 10.1209/0295-5075/4/9/004.

M. J. Feigenbaum. Quantitative universality for a class of nonlinear transformations. *Journal of Statistical Physics*, 19(1):25–52, 1978. doi: 10.1007/BF01020332.

A. M. Fraser and H. L. Swinney. Independent coordinates for strange attractors from mutual information. *Physical Review A*, 33(2): 1134–1140, 1986. doi: 10.1103/PhysRevA.33.1134.

P. Grassberger and I. Procaccia. Characterization of strange attractors. *Physical Review Letters*, 50(5):346–349, 1983. doi: 10.1103/PhysRevLett.50.346.

H. E. Hurst. Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116(1):770–799, 1951. doi: 10.1061/TACEAT.0006518.

J. W. Kantelhardt, S. A. Zschiegner, E. Koscielny-Bunde, S. Havlin, A. Bunde, and H. E. Stanley. Multifractal detrended fluctuation analysis of nonstationary time series. *Physica A: Statistical Mechanics and its Applications*, 316(1–4):87–114, 2002. doi: 10.1016/S0378-4371(02)01383-3.

S. A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York, 1993. ISBN 978-0-19-507951-7.

C. G. Langton. Computation at the edge of chaos: Phase transitions and emergent computation. *Physica D: Nonlinear Phenomena*, 42(1–3):12–37, 1990. doi: 10.1016/0167-2789(90)90064-V.

A. Lempel and J. Ziv. On the complexity of finite sequences. *IEEE Transactions on Information Theory*, 22(1):75–81, 1976. doi: 10.1109/TIT.1976.1055501.

N. R. Lomb. Least-squares frequency analysis of unequally spaced data. *Astrophysics and Space Science*, 39:447–462, 1976.

B. B. Mandelbrot and J. R. Wallis. Robustness of the rescaled range R/S in the measurement of noncyclic long run statistical dependence. *Water Resources Research*, 5(5):967–988, 1969. doi: 10.1029/WR005i005p00967.

N. Marwan, M. C. Romano, M. Thiel, and J. Kurths. Recurrence plots for the analysis of complex systems. *Physics Reports*, 438 (5–6):237–329, 2007. doi: 10.1016/j.physrep.2006.11.001.

R. M. May. Simple mathematical models with very complicated dynamics. *Nature*, 261(5560):459–467, 1976. doi: 10.1038/261459a0.

M. Mitchell. *Complexity: A Guided Tour*. Oxford University Press, Oxford, 2009. ISBN 978-0-19-512441-5.

N. H. Packard. Adaptation toward the edge of chaos. In J. A. S. Kelso, A. J. Mandell, and M. F. Shlesinger, editors, *Dynamic Patterns in Complex Systems*, pages 293–301. World Scientific, Singapore, 1988.

C.-K. Peng, S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley, and A. L. Goldberger. Mosaic organization of DNA nucleotides. *Physical Review E*, 49(2):1685–1689, 1994. doi: 10.1103/PhysRevE.49.1685.

J. S. Richman and J. R. Moorman. Physiological time-series analysis using approximate entropy and sample entropy. *American Journal of Physiology—Heart and Circulatory Physiology*, 278(6):H2039–H2049, 2000. doi: 10.1152/ajpheart.2000.278.6.H2039.

J. D. Scargle. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal*, 263:835–853, 1982.

T. Schreiber. Measuring information transfer. *Physical Review Letters*, 85(2):461–464, 2000. doi: 10.1103/PhysRevLett.85.461.

T. Schreiber and A. Schmitz. Surrogate time series. *Physica D: Nonlinear Phenomena*, 142(3–4):346–382, 2000.

H. G. Schuster and W. Just. *Deterministic Chaos: An Introduction*. Wiley-VCH, Weinheim, 4th edition, 2005. ISBN 978-3-527-40415-5.

R. V. Solé and B. Goodwin. *Signs of Life: How Complexity Pervades Biology*. Basic Books, New York, 2000. ISBN 978-0-465-01928-1.

J. C. Sprott. *Chaos and Time-Series Analysis*. Oxford University Press, Oxford, 2003. ISBN 978-0-19-850840-3.

S. H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, Boulder, CO, 2nd edition, 2015. ISBN 978-0-8133-4910-7.

F. Takens. Detecting strange attractors in turbulence. In D. Rand and L.-S. Young, editors, *Dynamical Systems and Turbulence, Warwick 1980*, volume 898 of *Lecture Notes in Mathematics*, pages 366–381. Springer, Berlin, 1981. doi: 10.1007/BFb0091924.

M. S. Taqqu, V. Teverovsky, and W. Willinger. Estimators for long-range dependence: an empirical study. *Fractals*, 3(4):785–798, 1995.

J. Theiler. Estimating fractal dimension. *Journal of the Optical Society of America A*, 7(6):1055–1073, 1990. doi: 10.1364/JOSAA.7.001055.

C. Torrence and G. P. Compo. A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society*, 79(1):61–78, 1998. doi: 10.1175/1520-0477(1998)079<0061:APGTWA>2.0.CO;2.

A. Wolf, J. B. Swift, H. L. Swinney, and J. A. Vastano. Determining Lyapunov exponents from a time series. *Physica D: Nonlinear Phenomena*, 16(3):285–317, 1985. doi: 10.1016/0167-2789(85)90011-9.

Table 2: Convergence of headline diagnostics with order–chaos-boundary expectations. The majority of metrics are consistent with a near-critical profile close to the order–chaos boundary; transfer entropy and autocorrelation model selection are inconclusive (see text for discussion). Values are reported at canonical macro precision (no table-specific re-rounding; SampEn shown to four decimal places).

Metric	Measured Value	Boundary Expectation	Consistent?
$\lambda_{\max}$ (Lyapunov)	-0.0045 (95% CI $[-0.0062, -0.0022]$)	Near zero; weakly contracting allowed	Partial
H_{PE} (Perm. entropy)	0.440 (95% CI $[0.426, 0.468]$)	Intermediate ($0 < H_{\mathrm{PE}} < 1$)	Yes
SampEn (Sample entropy)	0.0026 bits (95% CI $[0.0021, 0.0040]$); $A = 1410234$, $B = 1412734$	Low but non-zero	Yes
β (PSD exponent)	1.802 ($R^2 = 0.990$)	$\beta \in [1, 2]$ ($1/f$-type)	Yes
α_{DFA} (DFA)	2.004	$\alpha > 1$ (long-range corr.)	Yes
H (Hurst exponent)	$H_{\mathrm{DFA}} = 1.004$; $H_{R/S} = 1.011$ (95% CI $[1.008, 1.014]$)	$H > 0.5$ (persistent)	Yes
d_2 (Corr. dimension)	0.886	Finite, non-integer (fractal)	Yes, caveated
$\Delta\alpha$ (Multifractal)	0.007	Monofractal ($\Delta\alpha \approx 0$)	Yes (monofractal)
C_{LZ} (Lempel–Ziv)	$C^{(2)} = 0.001$, $C^{(4)} = 0.001$, $C^{(8)} = 0.002$	Highly compressible; all $\ll$ random ref.	Yes
τ_{decay} (Autocorr.)	10.503 model time; $\gamma = 0.721$; AIC: exponential	Slow decay (see discussion)	Partial
TE (Transfer entropy)	Histogram: 0.0001 bits ($p = 0.370$, n.s.); Gaussian-linear: 0.0000 bits ($p = 0.880$, n.s.)	Directed coupling (not detected)	Inconclusive
Perturbation recovery	≈ 0.010 model time units (all scales)	Scale-dependent recovery	Partial