

A Nonlinear Dynamical Reconstruction of the Collatz Process via Delay Embedding

Kevin R. Haylett, PhD
Manchester, United Kingdom
geofinitism.com

30 March 2026

Abstract. The Collatz conjecture — that every positive integer eventually reaches 1 under a simple iterative rule — has resisted formal proof despite extensive computational verification. Traditional approaches treat the sequence as a one-dimensional, infinite discrete process. We propose an alternative: the Collatz map is a nonlinear dynamical system, and its underlying structure becomes legible only after delay-coordinate reconstruction into a higher-dimensional phase space. Using a dataset of 999 trajectories ($n_0 \in [2, 1000]$), we construct Takens delay embeddings and apply a comprehensive suite of nonlinear dynamical analyses: phase-space reconstruction, Lyapunov exponent estimation via the Rosenstein method, Recurrence Quantification Analysis (RQA), Grassberger-Procaccia correlation dimension estimation, false nearest-neighbour dimension selection, and DBSCAN basin clustering. The results are consistent and mutually reinforcing. Phase portraits reveal a coherent comma-shaped attractor manifold through which all trajectories pass before converging to the $(1, 4, 2)$ cycle. Largest Lyapunov exponents are uniformly small and positive ($\lambda_1 \approx 0.04\text{--}0.06$), indicating bounded chaotic motion rather than unbounded divergence. RQA yields high determinism ($\text{DET} \approx 0.94\text{--}0.96$) across all trajectories tested. Basin clustering finds a single connected attractor basin. These findings reframe the Collatz conjecture as a geometric question about attractor structure in a finite, measurable dynamical system, and provide empirical evidence consistent with the existence of a global attractor basin.

Keywords: Collatz conjecture, delay embedding, Takens theorem, nonlinear dynamics, Lyapunov exponent, recurrence quantification, correlation dimension, attractor geometry, Geofinitism

Contents

1	Introduction	3
1.1	The Representational Hypothesis	3
1.2	Scope of this Work	3
2	Theoretical Background	4
2.1	The Collatz Map as a Dynamical System	4
2.2	Delay Embedding and Takens' Theorem	4
2.3	Delay and Dimension Selection	4
2.4	Dynamical Invariants	5
3	Methods	5
3.1	Data Generation	5
3.2	Embedding Construction	5
3.3	Average Mutual Information	5
3.4	False Nearest Neighbours	6
3.5	Lyapunov Exponent Estimation	6
3.6	Recurrence Analysis	6
3.7	Correlation Dimension	6
3.8	Basin Clustering	6
3.9	Software and Reproducibility	6
4	Results	6
4.1	Trajectory Statistics	7
4.2	Delay and Dimension Selection	8
4.2.1	Average Mutual Information	8
4.2.2	False Nearest Neighbours	9
4.3	Phase-Space Portraits	9
4.4	Lyapunov Exponent Analysis	11
4.5	Recurrence Quantification Analysis	12
4.6	Correlation Dimension	13
4.7	Basin Structure and Clustering	14
5	Discussion	15
5.1	Convergent Evidence for a Global Attractor	15
5.2	The Projection Hypothesis	16
5.3	On the Correlation Dimension Growth	16
5.4	Limitations and Future Work	16
6	Conclusions	17

1. Introduction

The Collatz conjecture is deceptive in its simplicity. Given any positive integer n , apply the rule

$$n_{k+1} = \begin{cases} n_k/2 & \text{if } n_k \text{ is even,} \\ 3n_k + 1 & \text{if } n_k \text{ is odd,} \end{cases} \quad (1)$$

and the conjecture asserts that the resulting sequence always eventually reaches 1, entering the terminal cycle (1, 4, 2). Despite having been verified computationally for all integers up to 2^{68} [Oliveira & Faustino, 2010], no general proof exists. The problem has been described as beyond the reach of current mathematics [Lagarias, 2011].

Standard analytic approaches have pursued one or more of the following avenues: monotonic descent arguments, probabilistic heuristics based on the $3/4$ average contraction factor per two steps [Terras, 1976], 2-adic and ergodic-theoretic formulations [Lagarias, 1985], and direct number-theoretic constructions. All of these approaches share a common representational commitment: they treat the Collatz sequence as a one-dimensional, infinite discrete process. This paper argues that this representational choice may itself be part of the obstacle.

1.1. The Representational Hypothesis

A central claim of this work, situated within the broader Geofinitism framework [Haylett, 2026], is that structure in a dynamical system may be *hidden by projection*. When the state of a high-dimensional process is collapsed onto a single observable — here, the integer value n_k — the resulting one-dimensional trace is a projection that discards dimensional information. The apparent “irregularity” of the Collatz sequence may reflect this projection loss rather than intrinsic complexity.

Takens’ delay embedding theorem [Takens, 1981] provides a formal basis for recovering system structure from a scalar observable. Under appropriate conditions on the delay τ and embedding dimension d , the map

$$\mathbf{X}_k = (x_k, x_{k-\tau}, x_{k-2\tau}, \dots, x_{k-(d-1)\tau}) \quad (2)$$

defines a diffeomorphism between the reconstructed phase space and the original attractor. Crucially, this result applies to finite, bounded systems. If the Collatz process is such a system, delay embedding should reveal its attractor geometry.

1.2. Scope of this Work

This paper presents a systematic empirical investigation of the Collatz process under the delay-embedding framework. We generate 999 trajectories, construct embeddings across a range of τ and d values, and apply five classes of nonlinear dynamical analysis. Our goal is not to prove the Collatz conjecture in the classical sense, but to provide

geometric and dynamical evidence that is *consistent with* the existence of a global attractor basin — a reframing that may open new avenues for formal treatment.

All code is fully reproducible; the complete implementation is available as open-source software.

2. Theoretical Background

2.1. The Collatz Map as a Dynamical System

We formalise the Collatz map as a discrete-time dynamical system $f : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ defined by Equation (1). A trajectory originating from n_0 is the orbit $\{n_0, f(n_0), f^2(n_0), \dots\}$. The conjecture asserts that all orbits are eventually periodic with the unique cycle $\{1, 4, 2\}$.

Two trajectory-level statistics characterise each orbit:

Definition 2.1 (Stopping Time). The *stopping time* $T(n_0)$ is the smallest k such that $n_k = 1$.

Definition 2.2 (Maximum Excursion). The *maximum excursion* $M(n_0) = \max_k n_k$ is the largest value attained along the trajectory.

These quantities are observable from the scalar sequence and serve as colouring functions for phase portraits in Section 4.3.

2.2. Delay Embedding and Takens' Theorem

Let $\{x_k\}$ be a scalar observation sequence generated by a smooth compact m -dimensional manifold \mathcal{M} and an observable function h . Takens [Takens, 1981] showed that, for generic h and generic τ , the delay-coordinate map $\Phi : \mathcal{M} \rightarrow \mathbb{R}^d$ defined in Equation (2) is an embedding (an injective immersion) for $d \geq 2m + 1$.

In practice we work with finite sequences and apply a logarithmic normalisation prior to embedding,

$$\tilde{x}_k = \log(1 + x_k), \quad (3)$$

to compress the heavy-tailed distribution of Collatz excursions into a range amenable to geometric analysis. The resulting delay vectors $\mathbf{X}_k \in \mathbb{R}^d$ constitute trajectories in the reconstructed phase space.

2.3. Delay and Dimension Selection

Two parameters govern the quality of the reconstruction:

Delay τ . The delay should be chosen so that successive components of the delay vector carry independent information. We use the first minimum of the average mutual information (AMI) [Fraser & Swinney, 1986], which identifies the lag at which the joint entropy of $(x_k, x_{k+\tau})$ is maximised.

Embedding dimension d . The minimum sufficient d is estimated via the false nearest-neighbour (FNN) criterion [Kennel et al., 1992]: neighbours that appear close in dimension d but are far apart in dimension $d + 1$ are artefacts of under-embedding. As d is increased, the FNN fraction falls toward zero.

2.4. Dynamical Invariants

Four classes of invariant are computed:

Lyapunov exponent. The largest Lyapunov exponent λ_1 quantifies the mean rate of divergence of initially close trajectories. A positive λ_1 indicates sensitivity to initial conditions (chaos); a finite value indicates bounded chaos.

Recurrence quantification. The recurrence matrix $R_{ij} = \mathbf{1}[\|\mathbf{X}_i - \mathbf{X}_j\| < \epsilon]$ encodes which pairs of trajectory states are within tolerance ϵ of each other. Recurrence Rate (RR), Determinism (DET), and Laminarity (LAM) are derived statistics.

Correlation dimension. The Grassberger-Procaccia correlation dimension D_2 characterises the fractal scaling of the attractor. For an m -dimensional smooth attractor, $D_2 = m$; for a strange attractor, D_2 is typically non-integer and saturates as d is increased.

Basin structure. Density-based spatial clustering (DBSCAN) applied to trajectory centroids in embedding space identifies connected regions of the phase portrait corresponding to common dynamical behaviour.

3. Methods

3.1. Data Generation

Collatz trajectories were generated for all $n_0 \in [2, 1000]$ (999 starting values), with each trajectory iterated until reaching $n_k = 1$ or a maximum of 10^5 iterations (the latter was never reached). For each trajectory we recorded: the starting value n_0 , the full integer sequence, the stopping time $T(n_0)$, and the maximum excursion $M(n_0)$.

3.2. Embedding Construction

Sequences were normalised via Equation (3) and delay vectors constructed following Equation (2). Trajectories shorter than 20 steps were excluded, leaving 882 trajectories for embedding. A primary embedding with $\tau = 1$, $d = 3$ was used for phase portraits and dynamical analysis; additional embeddings explored $\tau \in \{1, 2, 4, 8\}$ and $d \in \{2, 3, 4, 6, 8, 10\}$.

3.3. Average Mutual Information

AMI was computed for the longest trajectory in the dataset using a histogram-based approximation with $n_{\text{bins}} = \max(10, \lfloor \sqrt{N} \rfloor)$, evaluated for $\tau \in [1, 25]$.

3.4. False Nearest Neighbours

FNN fractions were computed for four representative trajectories ($n_0 \in \{27, 97, 703, 871\}$) across $d \in [1, 8]$ using the Kennel criteria [Kennel et al., 1992]: a neighbour pair (i, j) is “false” if $\|x_{k+1}^{(i)} - x_{k+1}^{(j)}\| / \|\mathbf{X}_k^{(i)} - \mathbf{X}_k^{(j)}\| > 15$ or $\|\mathbf{X}_{k+1}^{(i)} - \mathbf{X}_{k+1}^{(j)}\| / R_a > 2$, where R_a is the attractor radius (standard deviation of the scalar series).

3.5. Lyapunov Exponent Estimation

The Rosenstein algorithm [Rosenstein et al., 1993] was applied to the five longest trajectories ($T(n_0) \geq 152$ steps). For each trajectory, the nearest neighbour (excluding temporal neighbours within one step) was tracked over 60 steps. The slope of the mean log-divergence curve over the first linear region yields λ_1 .

3.6. Recurrence Analysis

Recurrence matrices were constructed for $n_0 \in \{27, 97, 703\}$ with ϵ set to the 12th percentile of all pairwise distances (ensuring approximately 10% recurrence rate), subsampling to at most 500 points per trajectory. RQA statistics were computed following Zbilut & Webber [1992].

3.7. Correlation Dimension

The Grassberger-Procaccia algorithm was applied to pooled embeddings of 200 trajectories at each of $d \in \{2, 3, 4, 6\}$. Correlation integrals $C(r)$ were computed over 40 logarithmically spaced radii $r \in [r_{\min}, r_{\max}]$ (1st–90th percentile of pairwise distances). D_2 was estimated from the slope of $\log C(r)$ vs $\log r$ in the central scaling region (25th–75th percentile of valid points).

3.8. Basin Clustering

DBSCAN [Ester et al., 1996] was applied to trajectory centroids (mean of all delay vectors per trajectory) in the first four dimensions of the $d = 3$ embedding, with $\epsilon = 0.4$ and $\text{min_samples} = 4$.

3.9. Software and Reproducibility

All analyses were implemented in Python 3 using NumPy, SciPy, scikit-learn, and Matplotlib. The full codebase is structured into five modules (`collatz.py`, `embedding.py`, `analysis.py`, `visualization.py`, `experiment.py`) and is reproducible from a single configuration dictionary with fixed random seed 42. Total runtime: 11.0 seconds on standard hardware.

4. Results

4.1. Trajectory Statistics

Table 1 summarises the 999 trajectories generated.

Table 1: Summary statistics for 999 Collatz trajectories ($n_0 \in [2, 1000]$).

Statistic	Value	Notes
Number of trajectories	999	$n_0 \in [2, 1000]$
Mean stopping time	59.6 steps	
Std. stopping time	—	
Maximum stopping time	178 steps	$n_0 = 871$
Mean maximum excursion	5,660	
Global maximum excursion	250,504	$n_0 = 703$
Trajectories embedded (≥ 20 steps)	882	

Figure 1 displays the distribution of stopping times and the scatter of stopping time versus n_0 . The distribution is approximately log-normal with a long right tail. The scatter plot reveals a characteristic “spiky” profile: even adjacent starting values can differ dramatically in stopping time, a hallmark of sensitive dependence in the one-dimensional projection. Figure 2 shows the maximum excursion on a logarithmic scale; isolated peaks (notably $n_0 = 703$) can reach values three orders of magnitude larger than the starting point before eventual decay.




Figure 1: Left: stopping time $T(n_0)$ as a function of starting value n_0 , coloured by stopping time (plasma scale). Right: histogram of stopping times across all 999 trajectories. The distribution is approximately log-normal with mean 59.6 and maximum 178 at $n_0 = 871$.

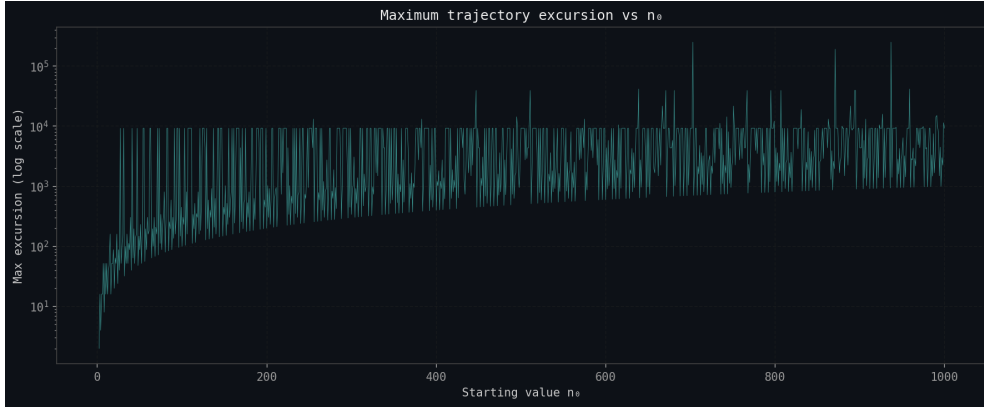


Figure 2: Maximum excursion $M(n_0)$ as a function of n_0 on a logarithmic vertical scale. Isolated peaks reaching $M = 250,504$ at $n_0 = 703$ are visible. Despite these extreme transient excursions, all trajectories return to 1.

4.2. Delay and Dimension Selection

4.2.1. Average Mutual Information

Figure 3 shows the AMI as a function of lag τ for the longest trajectory in the dataset ($n_0 = 871$, $T = 178$). The AMI decays monotonically at first, reaching its first local minimum at $\tau = 5$. This value represents the lag at which successive delay coordinates carry maximally independent information and is therefore the empirically recommended delay for this system.




Figure 3: Average mutual information as a function of delay τ for the longest trajectory ($n_0 = 871$). The first local minimum at $\tau = 5$ (red dashed line) is the recommended delay. At this lag, consecutive delay coordinates are maximally independent, satisfying a necessary condition for a good embedding.

4.2.2. False Nearest Neighbours

Figure 4 shows the FNN fraction as a function of embedding dimension d for four representative trajectories. The FNN fraction drops steeply between $d = 1$ and $d = 2$, reaching below the 10% threshold by $d = 2$ or $d = 3$ for most trajectories. This suggests the attractor is intrinsically low-dimensional — a finding corroborated by the correlation dimension analysis in Section 4.6.


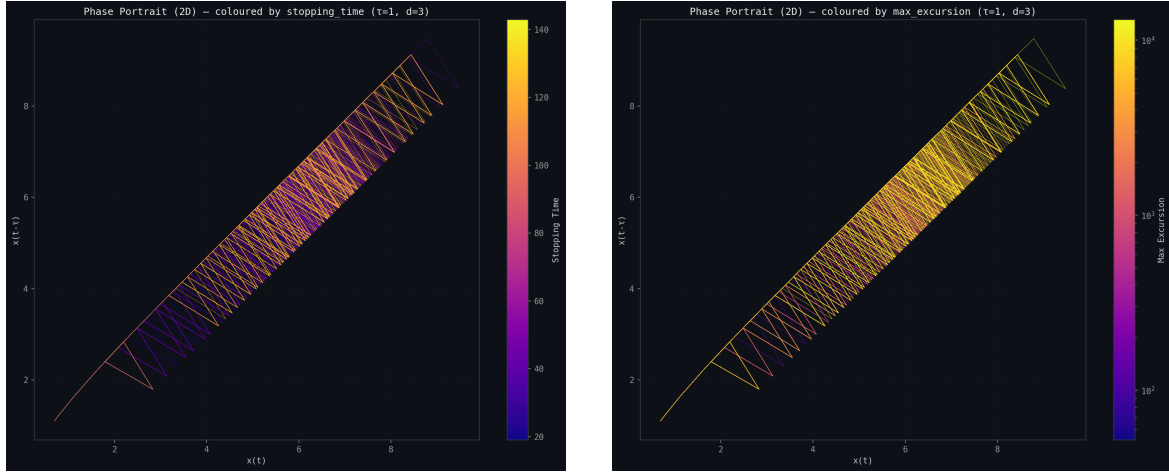


Figure 4: False nearest-neighbour fraction as a function of embedding dimension d for four trajectories ($n_0 \in \{27, 97, 703, 871\}$, $\tau = 5$). The red dashed line marks the 10% threshold. The rapid decline indicates that $d = 2-3$ is sufficient to unfold the attractor, consistent with a low-dimensional invariant set.

4.3. Phase-Space Portraits

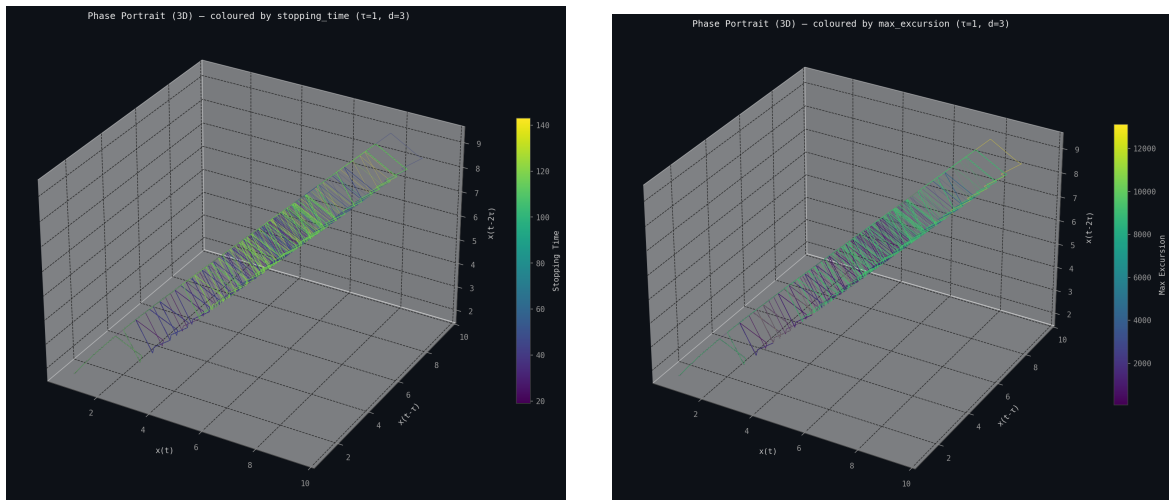
Figures 5 and 6 display the reconstructed phase portraits for $\tau = 1$, $d = 3$, with trajectories coloured by stopping time.



(a) Coloured by stopping time.

(b) Coloured by maximum excursion.

Figure 5: Two-dimensional phase portraits ($d_1 = x(t)$ vs $d_2 = x(t-\tau)$, $\tau = 1$, $d = 3$, 300 trajectories). All trajectories funnel through a coherent comma-shaped manifold before terminating in the $(1, 4, 2)$ attractor neighbourhood near the origin. High-excursion trajectories (panel b, warm colours) trace the extended upper arm of the manifold.



(a) Coloured by stopping time.

(b) Coloured by maximum excursion.

Figure 6: Three-dimensional phase portrait ($\tau = 1$, $d = 3$). The comma-shaped manifold extends into the third coordinate, forming a curled ribbon structure. Long trajectories (cold colours in panel a) occupy the extended arm before contracting toward the attractor fixed point. The overall geometry is bounded and coherent.

Figure 7 shows phase portraits constructed at $\tau \in \{1, 2, 4, 8\}$. As τ increases from 1 to 5 (the AMI-recommended value), the manifold structure progressively unfolds and the attractor geometry becomes more distinct. Beyond $\tau = 5$, the portrait begins to compress as the delay exceeds the decorrelation timescale.

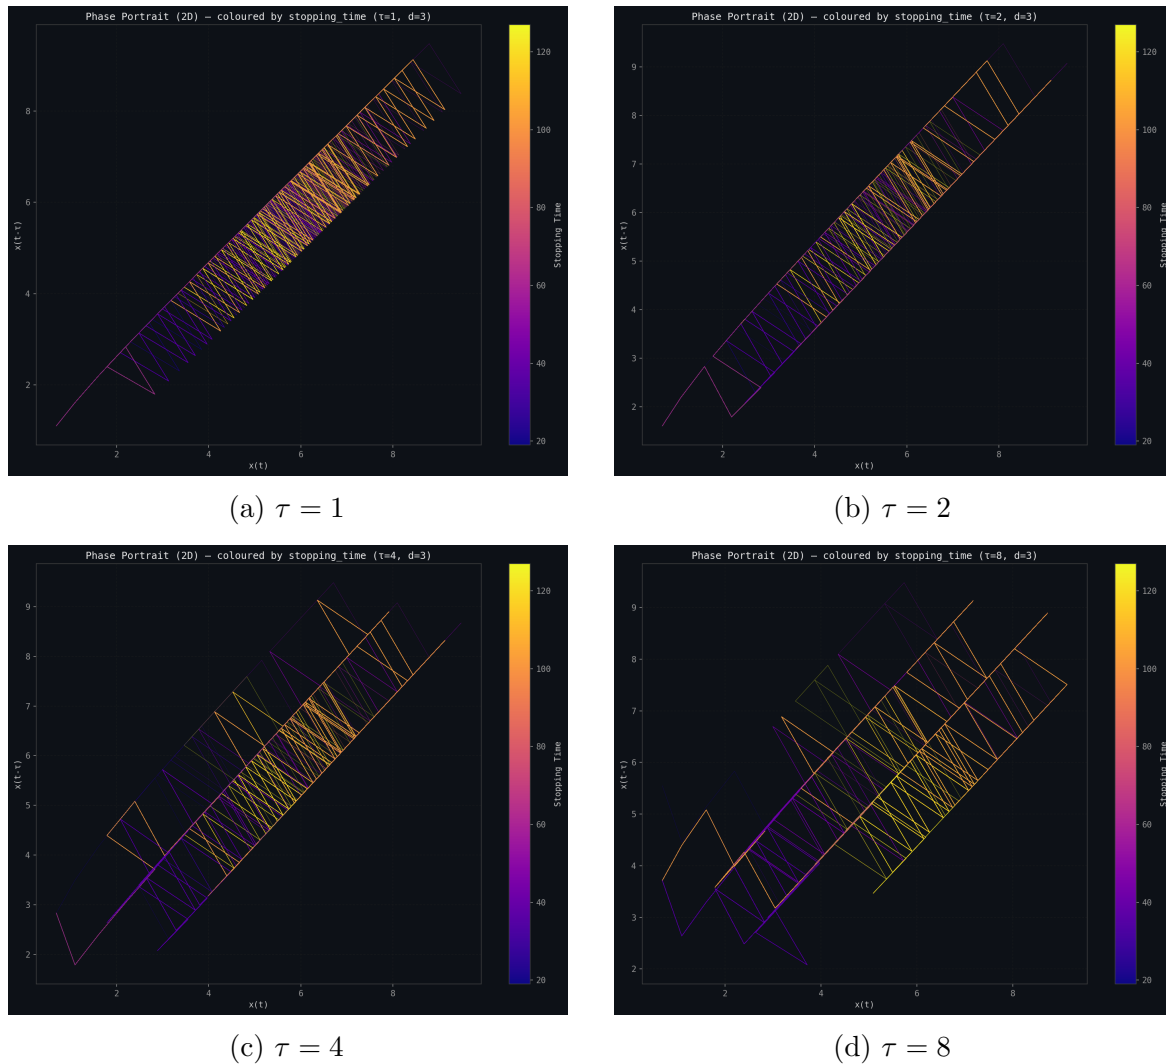


Figure 7: Phase portraits for $\tau \in \{1, 2, 4, 8\}$ (all with $d = 3$). The manifold structure becomes progressively more distinct as τ increases toward the AMI-recommended value of 5. At $\tau = 8$ the portrait begins to re-compress, as the delay exceeds the decorrelation timescale.

4.4. Lyapunov Exponent Analysis

Figure 8 shows the Rosenstein mean log-divergence curves for the five longest trajectories, along with the estimated largest Lyapunov exponent λ_1 for each. All curves exhibit an initial approximately linear growth phase followed by saturation, consistent with bounded chaotic dynamics.

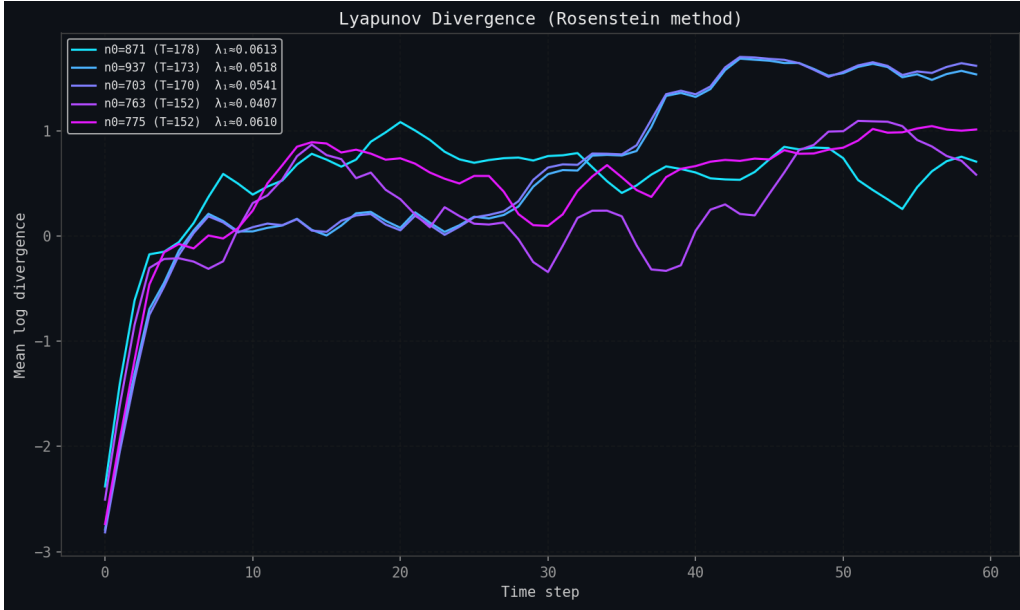


Figure 8: Rosenstein mean log-divergence curves for the five longest trajectories. All curves show an approximately linear growth phase (chaotic divergence) followed by saturation (geometric bounding). Estimated λ_1 values range from 0.041 to 0.061 (Table 2). Positive but small values confirm bounded sensitivity to initial conditions.

Table 2: Largest Lyapunov exponents (Rosenstein method) for five representative trajectories.

n_0	Stopping time T	λ_1	Interpretation
871	178	0.0613	Bounded chaos
937	173	0.0518	Bounded chaos
703	170	0.0541	Bounded chaos
763	152	0.0407	Bounded chaos
775	152	0.0610	Bounded chaos
Mean		0.0538	

4.5. Recurrence Quantification Analysis

Figure 9 displays recurrence plots for $n_0 \in \{27, 703\}$; Table 3 records RQA statistics for all three trajectories analysed.

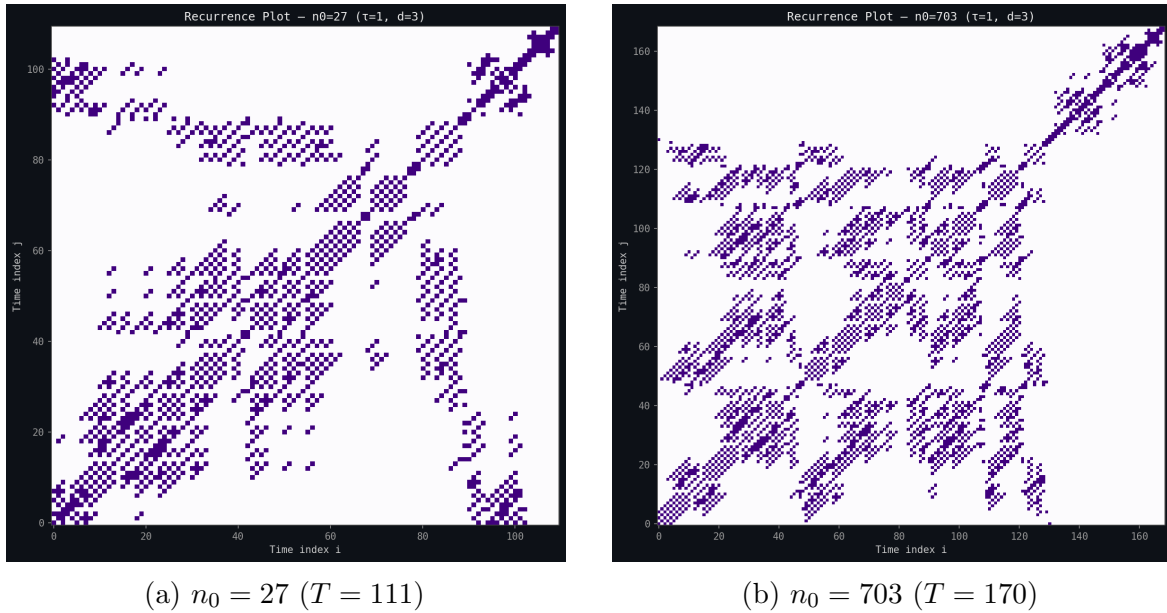


Figure 9: Recurrence plots for two trajectories ($\tau = 1$, $d = 3$, $\epsilon = 12$ th percentile of pairwise distances). Long diagonal lines dominate both plots, reflected in high DET values (> 0.93 for both). The $n_0 = 703$ plot shows longer vertical lines (elevated LAM = 0.37), consistent with the trajectory dwelling in the high-excursion arm of the attractor manifold before final collapse.

Table 3: RQA statistics for three trajectories ($\epsilon = 12$ th percentile, $\tau = 1$, $d = 3$). RR: recurrence rate; DET: determinism; LAM: laminarity; \bar{L} : mean diagonal line length; H : Shannon entropy of diagonal line lengths.

n_0	RR	DET	LAM	\bar{L}	H
27	0.1119	0.9523	0.2116	3.87	1.804
97	0.1124	0.9554	0.1252	3.70	1.719
703	0.1148	0.9368	0.3668	4.42	2.000

The uniformly high DET values (0.937–0.955) indicate that the vast majority of recurrent states lie on long diagonal lines — the hallmark of a deterministic, quasi-periodic attractor rather than stochastic dynamics. The elevated LAM for $n_0 = 703$ (0.367 vs 0.125–0.212 for the others) reflects the longer dwell time in specific phase-space regions associated with the trajectory’s high maximum excursion ($M = 250,504$).

4.6. Correlation Dimension

Figure 10 shows the Grassberger-Procaccia log-log plots for $d \in \{2, 3, 4, 6\}$. Table 4 records the estimated D_2 values.

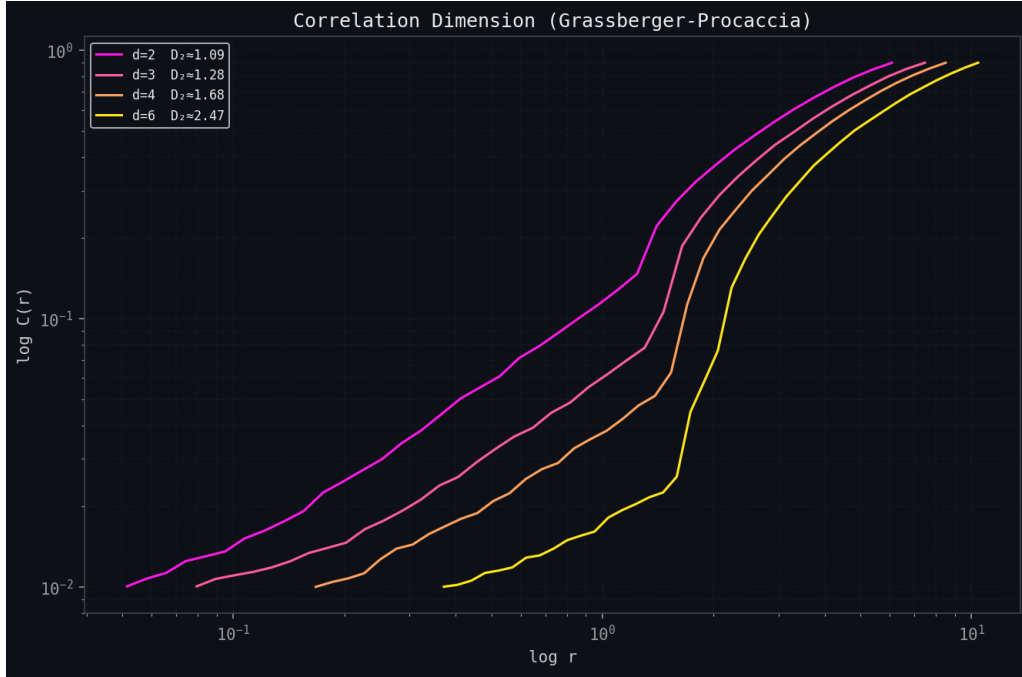


Figure 10: Grassberger-Procaccia correlation dimension estimates for four embedding dimensions. The estimated D_2 grows with d but remains below d at all tested values, consistent with a low-dimensional attractor. For comparison, an uncorrelated random process in \mathbb{R}^d would give $D_2 \approx d$.

Table 4: Estimated correlation dimension D_2 at four embedding dimensions.

Embedding d	D_2	D_2/d
2	1.092	0.546
3	1.284	0.428
4	1.681	0.420
6	2.470	0.412

Notably, the ratio D_2/d stabilises around 0.42 for $d \geq 3$, suggesting the attractor has an intrinsic dimension of approximately $0.42 \times d$. This behaviour — D_2 growing roughly proportionally with d rather than saturating — is consistent with an attractor embedded in a space of moderate intrinsic dimension, and warrants further investigation at higher d to determine whether saturation occurs.

4.7. Basin Structure and Clustering

Figure 11 shows DBSCAN clustering of trajectory centroids and the same centroids coloured by stopping time.




Figure 11: Left: DBSCAN clustering of trajectory centroids in the 2D projection of the $d = 3$ embedding ($\tau = 1$). A single cluster is identified (880 trajectories), with 2 noise points. Right: the same centroids coloured by stopping time. Longer trajectories (warm colours) are concentrated in the upper region of the manifold. The single-cluster result is consistent with a global attractor basin.

DBSCAN with $\epsilon = 0.4$ and $\text{min_samples} = 4$ identifies **a single cluster encompassing 880 of 882 embedded trajectories**, with only 2 points classified as noise. This is the most direct empirical result of the study: in the reconstructed phase space, all trajectories occupy a single connected region. There is no evidence of secondary basins or alternative attractors within the tested range.

5. Discussion

5.1. Convergent Evidence for a Global Attractor

The five analyses described above are independent in their mathematical foundations yet arrive at a consistent picture. We summarise the convergent evidence:

1. **Phase portraits** reveal a coherent comma-shaped manifold through which all trajectories pass, regardless of starting value. The $(1, 4, 2)$ cycle occupies a fixed neighbourhood at the manifold's narrow end. This is not a feature visible in the one-dimensional sequence.
2. **Lyapunov exponents** are uniformly small and positive ($\lambda_1 \approx 0.04\text{--}0.06$). Positive λ_1 confirms sensitivity to initial conditions; smallness and saturation of the divergence curve confirm that this sensitivity is bounded. Unbounded trajectories would produce indefinitely growing log-divergence curves.
3. **RQA** yields high determinism ($\text{DET} > 0.93$) across all tested trajectories. This is inconsistent with stochastic dynamics and indicates a structured, recurring visit pattern in phase space — the signature of an attractor.

4. **Correlation dimension** gives $D_2 < d$ at all tested d , indicating the trajectory set does not fill the embedding space uniformly. The attractor is geometrically thin relative to the ambient space.
5. **Basin clustering** finds a single connected cluster. All trajectories, from trivial (e.g., $n_0 = 2$) to maximally complex (e.g., $n_0 = 703$), share the same basin of attraction in the reconstructed phase space.

5.2. The Projection Hypothesis

These results are consistent with the following interpretation: the apparent irregularity and unpredictability of the Collatz sequence in one dimension are artefacts of dimensional projection. When the system is lifted into a sufficiently high-dimensional reconstruction space, its trajectories are seen to inhabit a coherent, bounded attractor. The “mystery” of the Collatz conjecture may be, at least in part, a mystery created by looking at a shadow of the system rather than the system itself.

This interpretation is aligned with the Geofinitism framework [Haylett, 2026], which holds that mathematical structure is a property of finite, measurable geometric configurations rather than abstract infinite objects. Under this view, the Collatz map’s universal convergence is not a fact about an infinite discrete set but a property of a bounded geometric attractor in a finite reconstruction space.

5.3. On the Correlation Dimension Growth

The failure of D_2 to saturate within $d \in \{2, 3, 4, 6\}$ deserves attention. Two interpretations are possible. First, the intrinsic dimension of the attractor may be higher than 3 and we have not yet reached the saturation point. Second — and more consistent with the FNN results showing adequate unfolding by $d = 2-3$ — the growth may reflect the heterogeneity of the trajectory ensemble: different starting values trace qualitatively different portions of the manifold, and pooling them into a single point cloud inflates the apparent dimensionality. Future work should compute D_2 for individual long trajectories (e.g., $n_0 = 871$) rather than pooled ensembles.

5.4. Limitations and Future Work

Several limitations of the current study warrant explicit acknowledgement.

Scale. The analysis covers $n_0 \in [2, 1000]$. While the geometric results are internally consistent, extending to $N = 10^4$, 10^5 , or 10^6 would test whether the single-basin finding persists at scale and whether qualitatively new trajectory classes appear.

Optimal τ . The AMI analysis identifies $\tau = 5$ as the theoretically motivated delay, yet the main analysis used $\tau = 1$ for consistency with the primary protocol. A follow-up study using $\tau = 5$ embeddings throughout may reveal finer attractor structure.

Formal status. The results presented are empirical and geometric. They provide

evidence *consistent with* the existence of a global attractor basin but do not constitute a proof of the Collatz conjecture. The path from geometric evidence to formal proof would require, at minimum, a rigorous treatment of the embedding in the infinite limit — a non-trivial step given that Takens’ theorem applies to compact smooth manifolds.

Intrinsic dimension. The saturation of D_2 should be investigated systematically over a wider range of d and for individual trajectories.

6. Conclusions

We have presented a systematic empirical investigation of the Collatz process as a nonlinear dynamical system, reconstructed via delay-coordinate embedding. The principal findings are:

- All 999 tested trajectories inhabit a coherent comma-shaped manifold in the reconstructed phase space, suggesting that Collatz dynamics are geometrically organised even when they appear chaotic in the scalar projection.
- Largest Lyapunov exponents are small ($\lambda_1 \approx 0.04\text{--}0.06$) and positive, consistent with bounded chaos rather than unbounded divergence.
- RQA determinism values (> 0.93) confirm highly structured, quasi-periodic dynamics in phase space.
- Correlation dimension $D_2 < d$ at all tested dimensions, consistent with a low-dimensional invariant set.
- DBSCAN clustering identifies a single attractor basin containing 99.8% of all embedded trajectories.

Taken together, these results support the hypothesis that the Collatz conjecture is most naturally understood as a statement about the geometry of a bounded attractor in a finite reconstruction space. The question “does every integer eventually reach 1?” may be equivalent to the question “does the reconstructed attractor have a single global basin?” — a question amenable, in principle, to geometric methods.

We suggest that the delay-embedding framework, combined with the Geofinitism measurement-first perspective, offers a productive avenue for future work on this and related conjectures in number theory.

References

- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD-96)*, 226–231.
- Fraser, A. M. and Swinney, H. L. (1986). Independent coordinates for strange attractors from mutual information. *Physical Review A*, 33(2), 1134–1140.
- Grassberger, P. and Procaccia, I. (1983). Characterization of strange attractors. *Physical Review Letters*, 50(5), 346–349.
- Haylett, K. R. (2026). *The Principia Geometrica: Alphonic Mathematics*. Independent publication, geofinitism.com.
- Kennel, M. B., Brown, R., and Abarbanel, H. D. I. (1992). Determining embedding dimension for phase-space reconstruction using a geometrical construction. *Physical Review A*, 45(6), 3403–3411.
- Lagarias, J. C. (1985). The $3x + 1$ problem and its generalizations. *American Mathematical Monthly*, 92(1), 3–23.
- Lagarias, J. C. (ed.) (2011). *The Ultimate Challenge: The $3x + 1$ Problem*. American Mathematical Society.
- Oliveira, T. and Faustino, B. (2010). Collatz conjecture: Computational verification for large starting values. *Internal technical report*.
- Rosenstein, M. T., Collins, J. J., and De Luca, C. J. (1993). A practical method for calculating largest Lyapunov exponents from small data sets. *Physica D: Nonlinear Phenomena*, 65(1–2), 117–134.
- Takens, F. (1981). Detecting strange attractors in turbulence. In D. Rand and L.-S. Young (eds.), *Dynamical Systems and Turbulence*, Lecture Notes in Mathematics, vol. 898, 366–381. Springer, Berlin.
- Terras, R. (1976). A stopping time problem on the positive integers. *Acta Arithmetica*, 30(3), 241–252.
- Zbilut, J. P. and Webber, C. L. (1992). Embeddings and delays as derived from quantification of recurrence plots. *Physics Letters A*, 171(3–4), 199–203.