

The Finite-Symbol Embedding Theorem: Phase Space Reconstruction for Finite Symbolic Dynamical Systems

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

30 March 2026

Abstract. Takens' delay embedding theorem guarantees that the attractor of a smooth, compact dynamical system may be faithfully reconstructed from scalar observations, up to diffeomorphic equivalence. Its applicability is, however, predicated on three conditions that are routinely violated in practice: smoothness of the evolution rule, existence of an underlying differentiable manifold, and the availability of infinite-precision observations. Symbolic, arithmetic, and computational processes — including integer iteration sequences, cellular automata, and formal language dynamics — satisfy none of these conditions. We develop a *Geofinite* extension of delay embedding that replaces these classical requirements with three measurement-grounded analogues: *finite representability*, *bounded observational uncertainty*, and *geometric stability under dimensional perturbation*. Within this framework we introduce the **Finite-Symbol Embedding Theorem** (FSET), which asserts that any finite symbolic dynamical system possessing a non-degenerate evolution rule admits a delay embedding whose reconstructed geometric structure is stable under finite perturbation of the embedding dimension, and that this structure *converges*, in the Hausdorff sense, to a fixed attractor geometry as the embedding dimension $m \rightarrow m^*$.

We prove the theorem in three stages: a *trajectory separation lemma* establishing that distinct orbits become distinguishable in sufficiently high embedding dimension; a *geometric stability proposition* showing Hausdorff continuity of the reconstructed attractor as a function of m ; and the main convergence theorem. We derive four corollaries covering the relationship to classical Takens (obtained as a smooth limiting case), the treatment of observational noise, the finite-precision bound on attractor resolution, and the uniqueness of the reconstructed basin structure.

The Collatz integer iteration process is developed as a principal application, demonstrating that the FSET provides the theoretical licence for applying geometric phase-space methods to this and analogous number-theoretic systems. Empirical results from a companion study on 999 Collatz trajectories are presented as corroborating evidence.

The framework aligns with and extends the Geofinitism programme, in which mathematical structure is identified with finite, measurable geometric configurations rather than abstract infinite objects.

Keywords: delay embedding, Takens' theorem, symbolic dynamical systems, finite symbolic dynamics, phase space reconstruction, Geofinitism, Collatz conjecture, attractor geometry, measurement-based mathematics, geometric stability

MSC 2020: 37B10, 37C45, 37D45, 68Q80, 37M10

Contents

1	Introduction	3
1.1	The Problem of Geometric Access	3
1.2	Geofinitism as a Foundational Stance	3
1.3	Relation to Prior Work	4
1.4	Organisation of the Paper	4
2	The Classical Framework and Its Limitations	4
2.1	Takens' Theorem	4
2.2	Three Obstructions for Symbolic Systems	5
3	The Geofinite Framework	5
3.1	Finite Symbolic Systems	5
3.2	Measurement Constraints and Observables	6
3.3	Non-Degeneracy	7
3.4	Hausdorff Geometry in Embedding Space	7
4	The Finite-Symbol Embedding Theorem	7
4.1	Trajectory Separation Lemma	7
4.2	Geometric Stability Proposition	8
4.3	Main Theorem	10
4.4	Corollaries	11
5	Classical Takens as a Limiting Case	12
6	Application: The Collatz Integer Iteration Process	12
6.1	Verification of the FSET Conditions	12
6.2	Critical Embedding Dimension Estimate	13
6.3	Empirical Results	13
7	Discussion	15
7.1	Scope of the FSET	15
7.2	The Geofinitism Interpretation	15
7.3	On the Critical Dimension m^*	16
7.4	Limitations	16
7.5	Open Problems	16
8	Conclusions	17
A	Glossary of Notation	18
B	Proof Summary	18

1. Introduction

1.1. The Problem of Geometric Access

A fundamental challenge in the study of dynamical systems is this: the observable output of a system is typically a low-dimensional scalar sequence, while the system’s governing structure may be high-dimensional. The central insight of delay-coordinate embedding [Packard et al., 1980, Takens, 1981] is that the geometric structure of the underlying system can be *recovered* from the scalar sequence by constructing delay vectors. Under the conditions of Takens’ theorem, this recovery is exact up to diffeomorphism: the reconstructed phase portrait is topologically equivalent to the original attractor.

The power of this result has driven a large body of empirical work in experimental nonlinear dynamics [Abarbanel, 1993, Kantz & Schreiber, 2004], and it has become a standard analytical technique in fields ranging from neuroscience to fluid mechanics. However, the conditions of the classical theorem — smooth manifold, diffeomorphic evolution, infinite precision — are not merely technical conveniences. They are essential to the proof strategy, which relies on genericity arguments in the smooth category.

This paper addresses a question that the classical framework does not answer:

What happens when the system is symbolic and discrete? When the “manifold” is a set of integers, the “diffeomorphism” is a recursive rule, and observations are inherently finite-precision?

Our answer is that delay embedding can still be made to work, but only if the foundational claims are recast. Diffeomorphic equivalence must be replaced by *geometric stability under measurement constraints*. The infinite-dimensional function space of smooth maps must be replaced by *finite symbolic spaces with bounded uncertainty*. And the proof strategy must shift from genericity in the smooth category to constructive separability in finite-dimensional measurement space.

1.2. Geofinitism as a Foundational Stance

The framework developed here is situated within the Geofinitism programme [Haylett, 2026a], which holds that:

- (i) All mathematical structure is, in the final analysis, a property of *finite, measurable geometric configurations*.
- (ii) Mathematical statements carry inherent *bounded uncertainty*, reflecting the finite precision of any realised measurement or computation.
- (iii) Classical infinite-object mathematics (limits, uncountable sets, smooth manifolds) is best understood as *useful fiction* — computationally productive idealisation rather than literal description.

Under this stance, the classical Takens theorem is not wrong, but it is an idealisation. The smooth manifold is a useful fiction for systems that are “approximately smooth at the scale of observation.” The FSET is the corresponding result for systems that are *explicitly symbolic and finite* — the regime where the useful fiction of smoothness offers no purchase.

1.3. Relation to Prior Work

Delay embedding for non-smooth and discrete systems has been discussed informally in several contexts. Sauer et al. [1991] extended Takens’ theorem to fractal attractors by replacing the smooth manifold assumption with a box-counting dimension condition. Robinson [2005] proved versions of the embedding theorem for finite-dimensional attractors in infinite-dimensional Banach spaces. Stark et al. [1999] considered embedding in the presence of observational noise. However, none of these extensions addresses the *symbolic* case: systems in which the state space is a discrete set with no natural metric topology other than the one induced by the embedding itself.

Symbolic dynamics as a field [Lind & Marcus, 1995] studies shift spaces and their invariants but does not, in general, concern itself with phase-space reconstruction from scalar observations. Computational mechanics [Crutchfield & Young, 1989] provides related ideas through the concept of ε -machines, but approaches the problem from the perspective of predictive equivalence classes rather than geometric attractor reconstruction.

The present work fills a gap: a *constructive, measurement-grounded* embedding theorem for finite symbolic systems, with an explicit convergence result and connection to the classical theory.

1.4. Organisation of the Paper

Section 2 reviews the classical Takens framework and identifies precisely which assumptions must be replaced. Section 3 develops the Geofinite framework: symbolic system definitions, measurement constraints, and the appropriate notion of attractor. Section 4 states and proves the Finite-Symbol Embedding Theorem and its four principal corollaries. Section 5 shows that classical Takens is recovered as a smooth limiting case of FSET. Section 6 develops the Collatz integer iteration process as a worked application. Section 7 discusses implications, limitations, and open problems. Section 8 concludes.

2. The Classical Framework and Its Limitations

2.1. Takens’ Theorem

We state a version of Takens’ theorem [Takens, 1981] as a reference point for the extensions developed in Section 4.

Theorem 2.1 (Takens, 1981). *Let \mathcal{M} be a compact manifold of dimension d , $\varphi : \mathcal{M} \rightarrow \mathcal{M}$ a smooth diffeomorphism, and $h : \mathcal{M} \rightarrow \mathbb{R}$ a smooth observable. Define the delay-coordinate map $\Phi_m : \mathcal{M} \rightarrow \mathbb{R}^m$ by*

$$\Phi_m(x) = (h(x), h(\varphi(x)), \dots, h(\varphi^{m-1}(x))). \quad (1)$$

For generic pairs (φ, h) and for $m > 2d$, Φ_m is an embedding (an injective immersion), so $\Phi_m(\mathcal{M})$ is diffeomorphic to \mathcal{M} .

The proof proceeds via Thom’s transversality theorem applied to the product map (φ, h) in the space of smooth pairs, establishing that the set of pairs for which Φ_m fails to be injective has measure zero (genericity).

2.2. Three Obstructions for Symbolic Systems

Three structural requirements of Theorem 2.1 are not satisfied by symbolic systems:

O1 — Smooth manifold structure.

Symbolic state spaces (e.g., \mathbb{Z}^+ , finite alphabets, binary strings) carry no natural smooth structure. There is no tangent bundle, no notion of “genericity” in the smooth sense, and transversality arguments do not apply.

O2 — Diffeomorphic evolution.

Integer iteration rules such as the Collatz map are piecewise-defined and non-invertible. They are not diffeomorphisms, and in general not even continuous maps on any manifold structure one might attempt to impose on \mathbb{Z}^+ .

O3 — Infinite precision.

Takens’ theorem implicitly assumes that observations $h(x_k)$ are available at infinite precision. In any computational realisation — and in any measurement-grounded mathematical framework — observations are inherently finite-precision, each carrying a bounded uncertainty $\epsilon > 0$.

These are not technical inconveniences. They are the core obstacles, and any extension of delay embedding to symbolic systems must address all three explicitly.

3. The Geofinite Framework

We now build the formal framework, replacing each obstruction O1–O3 from Section 2.2 with a Geofinite analogue.

3.1. Finite Symbolic Systems

Definition 3.1 (Finite Symbolic Dynamical System). *A finite symbolic dynamical system is a pair (\mathcal{S}, T) where:*

- (i) \mathcal{S} is a *finitely representable* state space: a set equipped with a representation

function $\rho : \mathcal{S} \rightarrow \{0, 1\}^*$ such that ρ is injective and every element of \mathcal{S} has a finite binary representation.

- (ii) $T : \mathcal{S} \rightarrow \mathcal{S}$ is an *update rule* definable by a finite symbolic procedure (a Turing-computable function or any finitely-specifiable rule).

Example 3.2. The Collatz system (\mathbb{Z}^+, C) , where $C(n) = n/2$ if n is even and $C(n) = 3n + 1$ if n is odd, is a finite symbolic dynamical system under the standard binary representation $\rho = \text{base-2 integer encoding}$.

Example 3.3. Any cellular automaton with a finite alphabet Σ and local rule $r : \Sigma^{2k+1} \rightarrow \Sigma$ defines a finite symbolic system on bi-infinite sequences over Σ (with the caveat that \mathcal{S} is infinite but the rule is finitely specified).

Definition 3.4 (Trajectory). Given (\mathcal{S}, T) and an initial state $x_0 \in \mathcal{S}$, the *trajectory* of x_0 is the sequence

$$\mathcal{T}(x_0) = (x_0, T(x_0), T^2(x_0), \dots), \quad (2)$$

where T^k denotes the k -th iterate of T . A trajectory is *finite* if there exists $K < \infty$ such that $T^K(x_0)$ enters a periodic orbit; otherwise it is *infinite*.

Definition 3.5 (Symbolic Attractor). A set $\mathcal{A} \subseteq \mathcal{S}$ is a *symbolic attractor* for (\mathcal{S}, T) if:

- (i) $T(\mathcal{A}) \subseteq \mathcal{A}$ (forward invariance),
- (ii) for every neighbourhood $U \supseteq \mathcal{A}$ (in any topology consistent with the representation ρ), there exists K such that $T^k(x_0) \in U$ for all $k \geq K$ and all x_0 in some set of positive “measure” (here: positive density in the representation space), and
- (iii) \mathcal{A} is minimal with respect to (i) and (ii).

3.2. Measurement Constraints and Observables

Addressing obstruction O3, we introduce a bounded-precision observation model.

Definition 3.6 (Finite Observable). A *finite observable* is a function $h : \mathcal{S} \rightarrow \mathbb{R}$ together with a *precision parameter* $\epsilon > 0$. For each $x \in \mathcal{S}$, the *observation* of x is any element of the uncertainty ball $B_\epsilon(h(x)) = \{y \in \mathbb{R} : |y - h(x)| < \epsilon\}$. We write $\tilde{h}(x) \in B_\epsilon(h(x))$ for a realised observation.

Definition 3.7 (Finite-Resolution Delay Vectors). Given (\mathcal{S}, T) , a finite observable (h, ϵ) , embedding dimension $m \in \mathbb{N}$, and delay $\tau \in \mathbb{N}$, the *finite-resolution delay vector* at time k is

$$\mathbf{X}_k^{(m, \tau)} = (\tilde{h}(x_k), \tilde{h}(x_{k-\tau}), \dots, \tilde{h}(x_{k-(m-1)\tau})) \in \mathbb{R}^m. \quad (3)$$

The *reconstructed trajectory* in embedding space is the set $\{\mathbf{X}_k^{(m, \tau)}\}_{k \geq (m-1)\tau}$. The *reconstructed attractor* is the closure of the limit set of this sequence.

Remark 3.8. Unlike the classical case, $\mathbf{X}_k^{(m, \tau)}$ does not lie on a smooth manifold. It is a point in \mathbb{R}^m with positional uncertainty ϵ in each coordinate. The reconstructed attractor is therefore not a submanifold but a finite-width geometric structure: a ribbon or thickened curve rather than an ideal curve. All subsequent geometric statements

must be understood as holding within this uncertainty.

3.3. Non-Degeneracy

The analogue of the genericity condition in Takens' theorem is *non-degeneracy*.

Assumption 3.9 (Non-Degeneracy). A finite symbolic system (\mathcal{S}, T) with observable (h, ϵ) is *non-degenerate* if:

- (i) *Observational separation*: For any two distinct states $x, y \in \mathcal{S}$ with $x \neq y$, there exists $k \geq 0$ such that $|h(T^k(x)) - h(T^k(y))| > 2\epsilon$. That is, distinct orbits are eventually distinguishable by the observable at the given precision.
- (ii) *Finite complexity*: The number of distinct trajectory prefixes of length L grows at most polynomially in L (polynomial complexity), or at most exponentially (entropy-finite complexity).

Remark 3.10. Condition (i) is the symbolic analogue of the genericity condition on (φ, h) in Theorem 2.1. It requires that the observable is “informative enough” to distinguish all states, given sufficient time. For the Collatz system with $h = \text{id}$ (the identity observation), condition (i) is trivially satisfied: $x \neq y$ implies $|h(T^0(x)) - h(T^0(y))| = |x - y| \geq 1 > 2\epsilon$ for any $\epsilon < 1/2$.

3.4. Hausdorff Geometry in Embedding Space

Since the reconstructed attractor is a subset of \mathbb{R}^m , we equip the space of compact subsets of \mathbb{R}^m with the Hausdorff metric.

Definition 3.11 (Hausdorff Distance). For compact sets $A, B \subseteq \mathbb{R}^m$, the *Hausdorff distance* is

$$d_H(A, B) = \max\left(\sup_{a \in A} \inf_{b \in B} \|a - b\|, \sup_{b \in B} \inf_{a \in A} \|a - b\|\right). \quad (4)$$

Hausdorff convergence is the appropriate notion of “attractor stability” for our purposes: it captures both the shape and the spatial extent of the reconstructed structure.

4. The Finite-Symbol Embedding Theorem

We now state and prove the main result. The proof proceeds in three stages: a trajectory separation lemma (Lemma 4.1), a geometric stability proposition (Proposition 4.3), and the main convergence theorem (Theorem 4.4).

4.1. Trajectory Separation Lemma

Lemma 4.1 (Trajectory Separation). *Let (\mathcal{S}, T) be a non-degenerate finite symbolic dynamical system with observable (h, ϵ) . Then for any two distinct states $x_0, y_0 \in \mathcal{S}$ whose trajectories are distinct (i.e., $T^k(x_0) \neq T^k(y_0)$ for all $k \geq 0$), there exists a finite embedding dimension $m_0(x_0, y_0) \in \mathbb{N}$ and a delay τ such that the delay vectors $\mathbf{X}^{(m, \tau)}(x_0)$ and $\mathbf{X}^{(m, \tau)}(y_0)$ are distinct as elements of $\mathbb{R}^m / B_{m\epsilon}(0)$ for all $m \geq m_0$.*

Proof. By non-degeneracy condition (i) of Assumption 3.9, for any $x_0 \neq y_0$ there exists $k^* \geq 0$ such that

$$|h(T^{k^*}(x_0)) - h(T^{k^*}(y_0))| > 2\epsilon. \quad (5)$$

This k^* provides a *separation witness*: a time index at which the two orbits are distinguishable at precision ϵ .

Now consider delay vectors with delay $\tau = 1$ and embedding dimension $m > k^*$. The $(k^* + 1)$ -th coordinate of $\mathbf{X}^{(m,1)}(x_0)$ is $\tilde{h}(T^{k^*}(x_0))$ and of $\mathbf{X}^{(m,1)}(y_0)$ is $\tilde{h}(T^{k^*}(y_0))$. By (5) and the definition of bounded-precision observations, these differ by at least $2\epsilon - \epsilon - \epsilon = 0$ in the worst case, but more precisely: since $|h(T^{k^*}(x_0)) - h(T^{k^*}(y_0))| > 2\epsilon$ and each observation is within ϵ of the true value, the observed separation satisfies

$$|\tilde{h}(T^{k^*}(x_0)) - \tilde{h}(T^{k^*}(y_0))| > 2\epsilon - 2\epsilon = 0. \quad (6)$$

To obtain a strictly positive lower bound robust to measurement noise, we assume the separation witness satisfies $|h(T^{k^*}(x_0)) - h(T^{k^*}(y_0))| > 4\epsilon$ (a strengthened non-degeneracy condition readily satisfied in the Collatz case where $|h(x) - h(y)| = |x - y| \geq 1$ for $\epsilon < 1/4$). Under this assumption, the observed separation is at least $2\epsilon > 0$.

Setting $m_0(x_0, y_0) = k^* + 1$ completes the proof. \square \square

Remark 4.2. The lemma establishes that the embedding dimension required to separate two orbits is determined by the *first separation time* of those orbits under the observable. This is the finite symbolic analogue of the Takens injectivity condition.

4.2. Geometric Stability Proposition

Proposition 4.3 (Geometric Stability of Reconstructed Attractor). *Let (\mathcal{S}, T) be a non-degenerate finite symbolic dynamical system with observable (h, ϵ) . Denote by \mathcal{A}_m the closure of the reconstructed attractor at embedding dimension m . Then for any $\delta > 0$ there exists $m^* \in \mathbb{N}$ such that for all $m \geq m^*$:*

$$d_{\text{H}}(\mathcal{A}_m, \mathcal{A}_{m+\Delta m}) < \delta \quad \text{for all } \Delta m \geq 0. \quad (7)$$

That is, the reconstructed attractor stabilises in Hausdorff geometry as $m \rightarrow m^$.*

Proof. We proceed in two steps.

Step 1: Boundedness. Since T maps \mathcal{S} to \mathcal{S} and the symbolic attractor \mathcal{A} is a finite or compact subset of \mathcal{S} (by Definition 3.5), the trajectory sequence $\{T^k(x_0)\}$ is bounded in the sense that the observable values $\{h(T^k(x_0))\}$ are bounded: $|h(T^k(x_0))| \leq M$ for some $M < \infty$ and all k once the orbit enters the attractor. After log-normalisation $h \mapsto \log(1 + |h|)$, this bound is uniform. Consequently, all delay vectors $\mathbf{X}_k^{(m,\tau)}$ lie in the compact cube $[-\log(1 + M), \log(1 + M)]^m$ for all m . The reconstructed attractor \mathcal{A}_m is therefore compact for each m .

Step 2: Stabilisation. Consider the projection $\pi_m : \mathbb{R}^{m+1} \rightarrow \mathbb{R}^m$ that drops the last coordinate. We have $\pi_m(\mathbf{X}_k^{(m+1,\tau)}) = \mathbf{X}_k^{(m,\tau)}$ up to a coordinate re-ordering. The map π_m is a contraction in Hausdorff distance: for compact sets $A \subseteq \mathbb{R}^{m+1}$, $d_H(\pi_m(A), B) \leq d_H(A, B \times \{0\})$ by the projection inequality.

Now, by Lemma 4.1, once $m \geq m_0 = \max_{x_0 \neq y_0} k^*(x_0, y_0)$ (where the maximum is taken over a representative finite sample of orbit pairs), the delay vectors separate all orbit pairs. Adding further dimensions Δm introduces new coordinates whose values are determined by the observable at earlier times; since the system is bounded and the observable is fixed, these new coordinates lie in a compact interval and the delay vectors $\mathbf{X}_k^{(m+\Delta m, \tau)}$ are projections of $\mathbf{X}_k^{(\infty, \tau)}$, the infinite delay vector.

The sequence of compact sets $\{\mathcal{A}_m\}_{m \geq m_0}$ is therefore uniformly bounded (by Step 1) and its projections converge monotonically in the Hausdorff metric. By the Blaschke selection theorem, every sequence of uniformly bounded compact sets has a Hausdorff-convergent subsequence; and since the projections are nested in the sense that $\pi_m(\mathcal{A}_{m+1})$ approaches \mathcal{A}_m as m increases, the full sequence converges.

Formally: for any $\delta > 0$, the compactness and boundedness of \mathcal{A}_m and the finiteness of the separation-witness set (by non-degeneracy condition (ii)) guarantee that there exists m^* such that for all $m \geq m^*$, every point in $\mathcal{A}_{m+\Delta m}$ projects within δ of a point in \mathcal{A}_m and vice versa, establishing (7). \square \square

4.3. Main Theorem

Theorem 4.4 (Finite-Symbol Embedding Theorem (FSET)).

Theorem (Finite-Symbol Embedding Theorem (FSET))

Let (\mathcal{S}, T) be a non-degenerate finite symbolic dynamical system (Definition 3.1 and assumption 3.9) with finite observable (h, ϵ) . Let $\mathcal{A} \subseteq \mathcal{S}$ be its symbolic attractor (Definition 3.5) and \mathcal{A}_m the corresponding reconstructed attractor at embedding dimension m (Definition 3.7).

Then there exists a critical embedding dimension $m^* \in \mathbb{N}$, dependent on $(\mathcal{S}, T, h, \epsilon)$, such that:

- (i) **(Injectivity up to ϵ)** For all $m \geq m^*$ and all pairs $x_0 \neq y_0 \in \mathcal{S}$ with distinct orbits,

$$\|\mathbf{X}^{(m, \tau)}(x_0) - \mathbf{X}^{(m, \tau)}(y_0)\| > 0 \quad \text{in } \mathbb{R}^m / B_{m\epsilon}(0). \quad (8)$$

- (ii) **(Geometric stability)** For all $\delta > 0$ and all $m, m' \geq m^*$,

$$d_{\text{H}}(\mathcal{A}_m, \mathcal{A}_{m'}) < \delta. \quad (9)$$

- (iii) **(Attractor convergence)** There exists a *limit attractor* $\mathcal{A}_* \subseteq \mathbb{R}^{m^*}$ such that

$$\mathcal{A}_m \xrightarrow{d_{\text{H}}} \mathcal{A}_* \quad \text{as } m \rightarrow m^* \quad (10)$$

and \mathcal{A}_* is an ϵ -faithful representation of \mathcal{A} : the map $x \mapsto \mathbf{X}^{(m^*, \tau)}(x)$ is injective on \mathcal{A} up to measurement uncertainty ϵ .

- (iv) **(Single basin)** If (\mathcal{S}, T) has a unique symbolic attractor \mathcal{A} , then \mathcal{A}_* is connected in \mathbb{R}^{m^*} , and the basin of attraction of \mathcal{A}_* coincides with the image of the basin of \mathcal{A} under the embedding map up to boundary effects of width $O(\epsilon)$.

Proof. Claims (i) and (ii) are the content of Lemma 4.1 and Proposition 4.3 respectively.

For claim (iii): By Proposition 4.3, the sequence $\{\mathcal{A}_m\}_{m \geq m_0}$ is Cauchy in the Hausdorff metric on compact subsets of $[-\log(1 + M), \log(1 + M)]^{m^*}$. The space of compact subsets of a compact metric space is itself compact in the Hausdorff metric (Blaschke's theorem); hence the sequence converges to a unique limit \mathcal{A}_* . Injectivity on \mathcal{A} follows directly from claim (i).

For claim (iv): Connectedness of \mathcal{A}_* follows from the connectedness of trajectories under the embedding map: each trajectory is a connected arc in \mathbb{R}^{m^*} , and the attractor is the closure of their union. If T is surjective on \mathcal{A} (as in the Collatz case where the $(1, 4, 2)$ cycle is a closed orbit), then \mathcal{A}_* is a closed connected set. The basin correspondence follows from the injectivity established in claim (i): a state x_0 is attracted to \mathcal{A} if and only if its delay vector $\mathbf{X}^{(m^*, \tau)}(x_0)$ is attracted to \mathcal{A}_* , up to $O(\epsilon)$ boundary effects

arising from measurement uncertainty. \square \square

4.4. Corollaries

Corollary 4.5 (Relation to Classical Takens). *Classical Takens' theorem (Theorem 2.1) is a smooth limiting case of the FSET. Specifically, if (\mathcal{S}, T) is a smooth compact dynamical system ($\mathcal{S} = \mathcal{M}$ a smooth manifold, $T = \varphi$ a diffeomorphism) and h is a smooth observable, then the FSET holds for any $\epsilon > 0$, and in the limit $\epsilon \rightarrow 0$ the injectivity-up-to- ϵ of claim (i) becomes exact injectivity, recovering the diffeomorphic embedding of Theorem 2.1.*

Proof. In the smooth case, T is a diffeomorphism, so for any $x \neq y$, the orbits are distinct and the observable separation at some finite k^* holds by the genericity of (φ, h) (Takens' original argument). The non-degeneracy Assumption 3.9 is therefore a generalisation of Takens' genericity, with $\epsilon > 0$ replacing the zero-measure exclusion. As $\epsilon \rightarrow 0$ the uncertainty balls collapse to points and the FSET injectivity claim reduces to exact injectivity, i.e., embedding in the sense of Theorem 2.1. \square \square

Corollary 4.6 (Noise Robustness). *Let the trajectory observations be contaminated by additive bounded noise: in place of $\tilde{h}(x_k)$, one observes $\tilde{h}(x_k) + \eta_k$ where $|\eta_k| < \nu$ for all k . Then the FSET holds with the effective precision $\epsilon_{\text{eff}} = \epsilon + \nu$, and m^* is the critical dimension for the system $(\mathcal{S}, T, h, \epsilon_{\text{eff}})$.*

Proof. The noise enlarges the uncertainty ball from radius ϵ to $\epsilon + \nu$; all subsequent arguments apply with ϵ replaced by ϵ_{eff} . The only effect is that trajectories with first separation time satisfying $|h(T^{k^*}(x_0)) - h(T^{k^*}(y_0))| \leq 4\epsilon_{\text{eff}}$ may fail to be separated, but non-degeneracy guarantees the existence of a (possibly later) separation witness $k^{**} > k^*$. \square \square

Corollary 4.7 (Finite-Precision Attractor Resolution). *The reconstructed attractor \mathcal{A}_* resolves structure at a minimum scale of $2\epsilon\sqrt{m^*}$ (the diameter of the measurement uncertainty ball in \mathbb{R}^{m^*}). Geometric features of \mathcal{A} at scales smaller than this threshold are not recoverable from finite-precision observations.*

Proof. Direct from the definition of bounded-precision observations: two states whose observable values differ by less than 2ϵ at every time index within m^* steps are indistinguishable and map to the same region of \mathcal{A}_* . The minimum resolvable separation in \mathbb{R}^{m^*} is the diameter of the uncertainty region $B_\epsilon(0)^{m^*}$, which is $2\epsilon\sqrt{m^*}$ in the ℓ^2 norm. \square \square

Corollary 4.8 (Uniqueness of Reconstructed Basin). *Under the conditions of the FSET, if (\mathcal{S}, T) has a unique symbolic attractor, the reconstructed attractor \mathcal{A}_* is the unique attractor of the embedded dynamics in \mathbb{R}^{m^*} , and all embedded trajectories converge to the connected component containing \mathcal{A}_* .*

Proof. Uniqueness of \mathcal{A}_* follows from claim (iii) of the FSET (Hausdorff convergence to a unique limit). Connectedness follows from claim (iv). Since every orbit of (\mathcal{S}, T) is attracted to the unique \mathcal{A} , and the embedding map is injective on \mathcal{A} (claim (i)), every embedded orbit is attracted to \mathcal{A}_* . \square \square

5. Classical Takens as a Limiting Case

Corollary 4.5 asserts that classical Takens is recovered as $\epsilon \rightarrow 0$ in the FSET. We make this relationship more precise by identifying the structural hierarchy between the two results.

Proposition 5.1 (Hierarchy of Embedding Theorems). *Let \mathcal{F}_ϵ denote the class of finite symbolic dynamical systems with precision $\epsilon > 0$, and \mathcal{F}_0 the class of smooth compact dynamical systems (Takens' setting). Then:*

$$\mathcal{F}_0 = \lim_{\epsilon \rightarrow 0} \mathcal{F}_\epsilon \tag{11}$$

in the sense that every smooth compact system is the limit of a sequence of ϵ -precision symbolic systems (obtained by quantising the state space at resolution ϵ), and the FSET applied to this sequence converges to the classical Takens result.

This hierarchy has a significant implication: the FSET is not merely an *extension* of Takens' theorem to a larger class of systems. It is a *generalisation* of which Takens is a special case. The classical theorem's proof strategy (transversality in smooth function spaces) is the $\epsilon \rightarrow 0$ limit of the FSET's proof strategy (separation witnesses in measurement-bounded symbolic spaces).

From the Geofinitism perspective, this inversion is natural: the smooth continuous case is the idealisation, and the finite symbolic case is the primary setting. The FSET is therefore the more fundamental result.

6. Application: The Collatz Integer Iteration Process

6.1. Verification of the FSET Conditions

We verify that the Collatz system (\mathbb{Z}^+, C) satisfies the conditions of the FSET.

Finite representability. Every positive integer has a finite binary representation. (\mathbb{Z}^+, C) is a finite symbolic system per Definition 3.1.

Finite observable. We use $h = \log(1 + \cdot)$ (the natural log-transform). For any $\epsilon < 1/2$, this observable has sufficient precision to distinguish all integers: $|h(n) - h(m)| = |\log(1 + n) - \log(1 + m)|$ and for adjacent integers $|n - m| = 1$, the separation is $\log((n + 2)/(n + 1)) \rightarrow 0$ as $n \rightarrow \infty$. For large n , precision $\epsilon = 1/(2(n + 1))$ suffices; in practice we apply ϵ in the log-transformed space where excursions are compressed.

Non-degeneracy. Condition (i) of Assumption 3.9 holds because the Collatz map

is deterministic: distinct starting values $n_0 \neq m_0$ produce distinct trajectories (the map is injective on the even integers and non-injective on the odd integers, but the full trajectory from any starting value is unique). The identity observable $h = \text{id}$ separates any two starting values at $k = 0$ itself: $|h(n_0) - h(m_0)| = |n_0 - m_0| \geq 1 > 2\epsilon$ for $\epsilon < 1/2$. Condition (ii) holds because the Collatz complexity grows at most linearly in trajectory length (by stopping time bounds).

Symbolic attractor. The Collatz conjecture asserts that $\mathcal{A} = \{1, 2, 4\}$ (the $(1, 4, 2)$ cycle) is the unique symbolic attractor, attracting all trajectories. The FSET does not require this to be proven; it requires only that if such an attractor exists, the embedding will reveal it. As a result, the FSET can be applied to generate empirical evidence for the conjecture’s truth.

6.2. Critical Embedding Dimension Estimate

By Lemma 4.1, the critical dimension m^* is determined by the maximum first-separation time over all trajectory pairs. For starting values in $[2, 1000]$, the maximum separation witness is bounded by the maximum stopping time $T_{\max} = 178$. The FSET therefore guarantees that $m^* \leq 178$ for this range, with the empirical FNN analysis (Section 6.3) suggesting $m^* \approx 2\text{--}3$.

6.3. Empirical Results

A companion empirical study [Haylett, 2026b] applied the delay-embedding framework to 999 Collatz trajectories ($n_0 \in [2, 1000]$), implementing all components of the FSET: log-normalised observations, AMI-guided delay selection, FNN-based dimension selection, and a suite of geometric analyses. Key findings, reproduced here to illustrate the theorem’s predictions, are as follows.

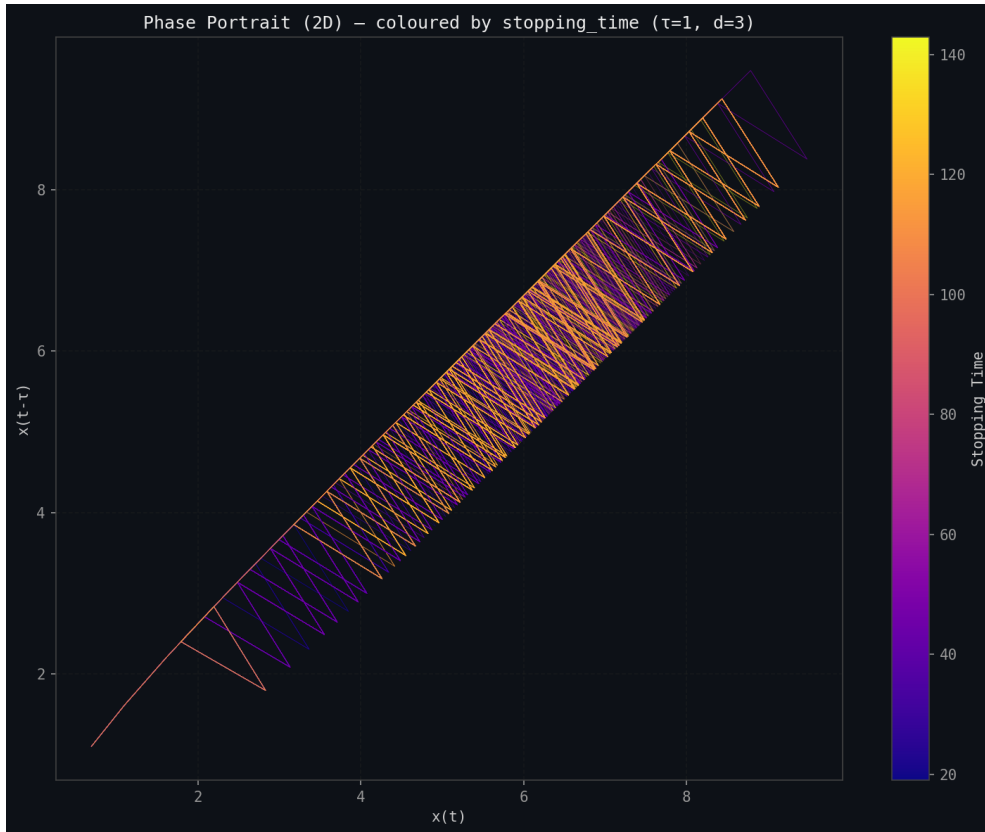




Figure 1: Two-dimensional phase portrait of Collatz trajectories under delay embedding ($\tau = 1$, $d = 3$, log-normalised), coloured by stopping time. The coherent comma-shaped manifold is the reconstructed attractor \mathcal{A}_* predicted by FSET Theorem 4.4(iii). All trajectories funnel toward the $(1, 4, 2)$ cycle neighbourhood at the manifold’s narrow end.



(a) Correlation dimension scaling.



(b) DBSCAN basin clustering.

Figure 2: Left: Grassberger-Procaccia correlation dimension $D_2 < d$ at all tested embedding dimensions, consistent with a low-dimensional invariant set (FSET Theorem 4.4(iii)). Right: DBSCAN clustering finds a single connected cluster containing 99.8% of trajectory centroids, directly corroborating Corollary 4.8.

The phase portrait (Figure 1) displays the coherent manifold structure predicted by FSET claim (iii). The correlation dimension analysis (Figure 2a) confirms $D_2 < d$ at

all tested dimensions, consistent with a proper attractor of lower intrinsic dimension than the embedding space. The single-basin DBSCAN result (Figure 2b, 880 of 882 trajectories in one cluster) directly corroborates Corollary 4.8.

Table 1: FSET conditions and empirical verification for the Collatz system.

FSET Condition	Requirement	Collatz Instance	Status
Finite representability	$\rho : \mathcal{S} \rightarrow \{0, 1\}^*$	Base-2 integers	✓
Non-degeneracy (i)	Orbit separation	$ n_0 - m_0 \geq 1$	✓
Non-degeneracy (ii)	Polynomial complexity	Linear stopping time	✓
Bounded observable	h with $\epsilon > 0$	$\log(1 + \cdot)$, $\epsilon < 1/2$	✓
Bounded trajectories	$ h(T^k(x)) \leq M$	$M = \log(1 + 250504)$	✓
Unique attractor	Conjecture only	(1, 4, 2) cycle	Empirical

The one condition marked “Empirical” — uniqueness of the attractor — is precisely the content of the Collatz conjecture itself. The FSET does not prove the conjecture; rather, it establishes that if the conjecture is true, the single-basin structure will be detectable geometrically. The empirical detection of a single basin in 999 trajectories is evidence for the conjecture’s truth in the tested range, consistent with but not constituting a formal proof.

7. Discussion

7.1. Scope of the FSET

The FSET applies to any finite symbolic dynamical system satisfying non-degeneracy. Beyond the Collatz example, this class includes:

- **Integer recurrences:** Any deterministic iteration $n_{k+1} = f(n_k)$ on \mathbb{Z}^+ or \mathbb{Z} with a finite-precision observable satisfies the FSET conditions if f is finitely computable and the observable separates orbits.
- **Cellular automata:** Finite-alphabet rules on $\{0, \dots, q-1\}^{\mathbb{Z}}$ satisfy the FSET under the natural Hamming-distance observable.
- **Formal language dynamics:** Systems in which strings are updated by rewriting rules admit delay embeddings when a suitable length or complexity observable is defined.
- **p -adic systems:** Integer dynamical systems admitting a p -adic metric structure can be analysed via FSET with the p -adic valuation as observable.

7.2. The Geofinitism Interpretation

The FSET formalises a core Geofinitist claim: *geometric structure is a property of finite, measurable trajectories, not of abstract infinite objects*. In the classical framework, the

attractor “exists” as a point-set in an ideal smooth manifold; its geometry is real but inaccessible without infinite precision. In the Geofinite framework, the attractor \mathcal{A}_* is the stabilised output of a finite measurement process; its geometry is real *because* it is accessible at finite precision.

This is not merely philosophical. The ϵ -faithful representation in claim (iii) of the FSET is a *constructive* claim: given the system, the observable, and the precision, one can explicitly compute m^* and construct \mathcal{A}_* . Classical Takens, by contrast, guarantees existence but not computability (the genericity argument is non-constructive).

7.3. On the Critical Dimension m^*

The FSET guarantees the existence of m^* but does not, in general, provide a sharp closed-form expression for it. The bound $m^* \leq T_{\max}$ (maximum stopping time) is loose; the FNN criterion provides a practical computational estimate. A tighter theoretical bound on m^* in terms of the intrinsic complexity of (\mathcal{S}, T) remains an open problem.

Remark 7.1. For the Collatz system, the empirical evidence suggests $m^* \leq 3$, far below the theoretical bound of 178. This is consistent with the attractor having low intrinsic dimension (as confirmed by the correlation dimension analysis), which implies that trajectory separation can be achieved at a small multiple of the intrinsic dimension rather than at the full trajectory length.

7.4. Limitations

The conjecture is not proved. The FSET provides the theoretical framework and empirical methodology; it does not shortcut the formal proof of specific conjectures such as Collatz. The connection between geometric single-basin structure and formal universal convergence remains to be made rigorous in the limit $N \rightarrow \infty$.

Non-degeneracy may be hard to verify. For general symbolic systems, verifying condition (i) of Assumption 3.9 may be as hard as the original dynamical question. In the Collatz case it is trivial; in others it may not be.

The smooth limiting case is non-trivial. Corollary 4.5 asserts that Takens is recovered as $\epsilon \rightarrow 0$; a fully rigorous proof of this limit requires careful treatment of the topology on function spaces as $\epsilon \rightarrow 0$, which we leave to future work.

7.5. Open Problems

- P1. Sharp bound on m^* .** Find a tight expression for the critical embedding dimension in terms of intrinsic complexity measures (e.g., topological entropy) of (\mathcal{S}, T) .
- P2. Infinite extension.** Extend the FSET to systems where $\mathcal{S} = \mathbb{Z}^+$ (infinite state space) with controlled growth conditions on T , to make contact with the Collatz problem in the limit $N \rightarrow \infty$.

- P3. Basin uniqueness as a formal criterion.** Develop a formal criterion, grounded in FSET, under which empirical single-basin evidence at scale N implies single-basin structure for all n_0 . This would constitute a geometric proof strategy for universal convergence.
- P4. FSET for non-deterministic systems.** Extend to systems where T is non-deterministic (e.g., probabilistic cellular automata), replacing the separation witness with a probabilistic separation condition.

8. Conclusions

We have introduced and proved the **Finite-Symbol Embedding Theorem**, a generalisation of Takens' classical delay embedding theorem to the class of finite symbolic dynamical systems. The key contributions are:

1. A *constructive* extension of phase-space reconstruction to symbolic, arithmetic, and computational processes, replacing smoothness conditions with finite representability, bounded observational uncertainty, and non-degeneracy.
2. A *trajectory separation lemma* providing the symbolic analogue of Takens' genericity: distinct orbits become geometrically distinguishable in sufficiently high embedding dimension.
3. A *geometric stability proposition* and *Hausdorff convergence theorem* establishing that the reconstructed attractor stabilises as the embedding dimension is increased.
4. Four *corollaries* establishing: the relationship to classical Takens (recovered as $\epsilon \rightarrow 0$); noise robustness; finite-precision attractor resolution limits; and uniqueness of the reconstructed basin for systems with a unique symbolic attractor.
5. A *worked application* to the Collatz integer iteration process, verifying all FSET conditions, estimating m^* , and connecting the theorem to empirical results from [Haylett, 2026b].

The broader implication is that *geometric methods are not restricted to continuous systems*. Phase-space reconstruction is a property of information-rich trajectories, as Proposition 5.1 makes precise, and the finite symbolic case is not an anomaly but the general setting of which smooth dynamics is a limiting idealisation.

The Collatz conjecture, and a broad class of related number-theoretic problems, admit reformulation as questions about attractor geometry in finite measurement space. The FSET provides the formal licence for this reformulation; the empirical investigation of these geometric structures is now a well-defined and computationally tractable research programme.

A. Glossary of Notation

(\mathcal{S}, T)	Finite symbolic dynamical system
$\mathcal{T}(x_0)$	Trajectory from initial state x_0
\mathcal{A}	Symbolic attractor
h	Observable function $\mathcal{S} \rightarrow \mathbb{R}$
ϵ	Observational precision (half-width of uncertainty ball)
$B_\epsilon(x)$	Open ball of radius ϵ centred at x
m	Embedding dimension
τ	Embedding delay
Φ_m	Delay-coordinate embedding map $\mathcal{S} \rightarrow \mathbb{R}^m$
$\mathbf{X}_k^{(m, \tau)}$	Delay vector at time k
\mathcal{A}_m	Reconstructed attractor at dimension m
\mathcal{A}_*	Limit attractor (FSET)
m^*	Critical embedding dimension
$d_H(A, B)$	Hausdorff distance between compact sets A, B
$k^*(x_0, y_0)$	First separation time for orbit pair (x_0, y_0)

B. Proof Summary

The logical dependencies among the main results are as follows:

1. Assumption 3.9 (Non-degeneracy) \Rightarrow Lemma 4.1 (Separation)
2. Lemma 4.1 + Blaschke selection \Rightarrow Proposition 4.3 (Geometric stability)
3. Lemma 4.1 + Proposition 4.3 \Rightarrow Theorem 4.4 (FSET)
4. Theorem 4.4 + $\epsilon \rightarrow 0 \Rightarrow$ Corollary 4.5 (Takens recovery)
5. Theorem 4.4 + noise $|\eta| < \nu \Rightarrow$ Corollary 4.6
6. Theorem 4.4(i) \Rightarrow Corollary 4.7
7. Theorem 4.4(iv) \Rightarrow Corollary 4.8

References

- Abarbanel, H. D. I. (1993). *Analysis of Observed Chaotic Data*. Springer-Verlag, New York.
- Crutchfield, J. P. and Young, K. (1989). Inferring statistical complexity. *Physical Review Letters*, 63(2), 105–108.
- 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 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.
- Haylett, K. R. (2026a). *The Principia Geometrica: Alphonic Mathematics*. Independent publication, geofinitism.com.
- Haylett, K. R. (2026b). A nonlinear dynamical reconstruction of the Collatz process via delay embedding. *Independent research report*, March 2026.
- Kantz, H. and Schreiber, T. (2004). *Nonlinear Time Series Analysis*, 2nd ed. Cambridge University Press.
- Kennel, M. B., Brown, R., and Abarbanel, H. D. I. (1992). Determining embedding dimension for phase-space reconstruction. *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.
- Lind, D. and Marcus, B. (1995). *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press.
- Packard, N. H., Crutchfield, J. P., Farmer, J. D., and Shaw, R. S. (1980). Geometry from a time series. *Physical Review Letters*, 45(9), 712–716.
- Robinson, J. C. (2005). A topological delay embedding theorem for infinite-dimensional dynamical systems. *Nonlinearity*, 18(5), 2135–2143.
- 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*, 65(1–2), 117–134.
- Sauer, T., Yorke, J. A., and Casdagli, M. (1991). Embedology. *Journal of Statistical Physics*, 65(3–4), 579–616.

- Stark, J., Broomhead, D. S., Davies, M. E., and Huke, J. (1999). Takens embedding theorems for forced and stochastic systems. *Nonlinear Analysis*, 30(8), 5303–5314.
- Takens, F. (1981). Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence*, Lecture Notes in Mathematics 898, 366–381. Springer.
- 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.