Epistemic Limits of the Embedded Observer

Peter David Fagan
School of Informatics, University of Edinburgh

May 30, 2026

Abstract

The Platonic observer is replaced by a physically embedded computational subsystem. Bounded by the fundamental limits imposed by Conservation-Congruent Encoding (CCE), the observer is forced to project a massively large set of microscopic realities into a strictly smaller set of macroscopic equivalence classes. Consequently, state evolution under exact reversible dynamics becomes predictively intractable within the observer's formal system. Furthermore, a continuous state space requires infinite encoding precision and unbounded operational resources to model exactly, rendering the continuum physically unobservable and forcing epistemic reliance on discrete macroscopic equivalence classes. Expanding upon the physical limits of inference and the computational capacity bounds of the universe, this framework proves that computing the precise trajectory of the universe is physically impossible without exceeding the observer's finite conserved-resource budgets.

1 Axioms of the Embedded Observer

Let an observer $\mathcal{O}$ utilize a physical encoding substrate $\mathcal{M}$ to process and store information about a microscopic physical system $\mathcal{E}$, which possesses state space $\Omega_{\mathcal{E}}$. The first two theorems treat the case in which $\Omega_{\mathcal{E}}$ is an effectively finite discrete state space; the final theorem treats the putative continuum as a stronger counterfactual limit.

Axiom 1 (Geometric and Encoding Capacity Bounds)**.** *The observer's ability to store information is strictly bounded by its physical embedding: $\mathcal{M} \subseteq \mathcal{O} \subset \mathcal{E}$. By the Bekenstein-Hawking area law [4, 5], maximum encoding capacity is strictly monotone and finite for any bounded physical substrate: $I_{\max}(\mathcal{M}) \leq I_{\max}(\mathcal{O}) < I_{\max}(\mathcal{E}) < \infty$.*

Axiom 2 (Conservation-Congruent Processing). *Encoded states are physical distinctions whose creation, preservation, and erasure must respect conserved quantities. In accordance with Landauer's principle and the Conservation-Congruent Encoding (CCE) framework [6, 3], irreversible information processing operations (e.g., erasure or merging) of* ΔI_{irr} *nats demand a proportional export of entropy into physical resource channels. For exchanged conserved quantities* Q_a *and their intensive conjugates* λ_a*, this CCE resource cost is strictly bounded:* $\sum_a \lambda_a \Delta Q_a \geq \Delta I_{\mathrm{irr}}$*.*

2 The Physical Intractability Theorem

Theorem 1 (State-Level Epistemic Projection). *An isomorphic, complete microscopic model of* $\mathcal{E}$ *or* $\mathcal{O}$ *is physically forbidden. The observer must utilize a massive informational compression, mapping the finite microstate space* $\Omega_{\mathcal{E}}$ *into a strictly smaller finite set of discrete macroscopic encodings* Z*.*

Proof. An exact model of $\mathcal{E}$ requires an active encoding capacity $I(\mathcal{M}) \geq I_{\max}(\mathcal{E})$, violating the bounds of Axiom 1. Similarly, an exact live self-model demands perpetual recursive updating of the observer's encoded information. Correcting this state mismatch requires unphysical conserved-resource dissipation (Axiom 2) or prohibitive auxiliary encoding capacity (Axiom 1). Thus, the system is forced to use a many-to-one projection $\pi_\Phi : \Omega_{\mathcal{E}} \to Z$. Because $|Z| \ll |\Omega_{\mathcal{E}}|$ and the capacity of Z is finite, any metastable encoding basin $B_z = \pi_\Phi^{-1}(z) \subset \Omega_{\mathcal{E}}$ possesses a vast multiplicity of microstates. Therefore, there exist distinct microstates $x_1 \neq x_2 \in B_z$ that map to the exact same macroscopic encoding z, rendering them strictly indistinguishable to the observer. □

Theorem 2 (Dynamical Intractability). *For systems evolving under non-trivial discrete reversible dynamics, there exist physically realized future and past states of* $\mathcal{E}$ *that cannot be uniquely predicted or retrodicted by the observer's formal model* π_Φ*.*

Proof. Let $x_1, x_2 \in B_z$ be two distinct microstates that map to the same preserved initial macrostate z_0 at discrete time t_0. At the fundamental level, the true internal and external states evolve under an exact, deterministic, and reversible unit-step map T, so that $T_\tau = T^\tau$ for elapsed discrete time steps $\tau \in \mathbb{Z}_{>0}$. For any macrostate z, define the sets of distinct macroscopic futures and pasts by

$$\mathcal{F}_\tau(z) = \{\pi_\Phi(T_\tau x) : x \in \pi_\Phi^{-1}(z)\}, \qquad \mathcal{P}_\tau(z) = \{\pi_\Phi(T_\tau^{-1} x) : x \in \pi_\Phi^{-1}(z)\}.$$

In non-trivial systems, such measure-preserving dynamics exhibit effective mixing relative to the observer's coarse-graining, so localized microscopic differences may be dispersed across the finite system. For some $\tau \in \mathbb{Z}_{>0}$, the forward evolutions of these initially indistinguishable microstates cross different macroscopic partition boundaries, yielding diverging macroscopic futures such that $|\mathcal{F}_\tau(z_0)| > 1$.
Conversely, macroscopic dissipation and retrodictive incompleteness emerge whenever $|\mathcal{P}_\tau(z_0)| > 1$, directly as an artifact of the many-to-one projection, which irreversibly merges distinct microscopic histories into shared macroscopic attractors.
The universe realizes exactly one microscopic trajectory. To deterministically infer this true trajectory from the macroscopic degeneracy, the observer must process $I_{\text{dec}}(\tau)$ additional nats. By Axiom 2, this irreversible acquisition demands a physical resource payment $\sum_a \lambda_a \Delta Q_a^{\text{dec}} \geq I_{\text{dec}}(\tau)$. For dynamics with growing branch entropy relative to encoding capacity, this cost inevitably exceeds the finite conserved-resource capacity of $\mathcal{O}$, rendering the true trajectory physically intractable. □

3 The Epistemic Boundary of the Continuum

The preceding result establishes physical intractability even for a finite discrete microstate space. A continuous microscopic state space does not evade this limit; it strengthens it by demanding exact distinctions that no embedded physical substrate can finitely encode, preserve, or operationally resolve.

Theorem 3 (Epistemic Unobservability of the Continuum)**.** *If the microscopic state space $\Omega_\mathcal{E}$ is a continuous manifold, then an exact embedded model of a microscopic state or trajectory is physically forbidden by finite encoding capacity and finite operational resources. Consequently, continuity cannot be physically established by any finite embedded observer; only finite coarse-grained equivalence classes are observable.*

Proof. Assume, counterfactually, that $\Omega_\mathcal{E}$ is a continuous manifold and that a microstate $x \in \Omega_\mathcal{E}$ possesses at least one coordinate q in a bounded continuum chart. To specify q to finite resolution ϵ over a chart of size L requires at least $\ln(L/\epsilon)$ nats. Exact specification requires the limit

$$I_{\text{exact}}(q) = \lim_{\epsilon \to 0} \ln(L/\epsilon) = \infty.$$

Thus even a single exact continuous coordinate requires infinite encoding precision. An exact embedded model of x therefore demands $I(\mathcal{M}) = \infty$, contradicting the finite Bekenstein-Hawking capacity bound $I_{\max}(\mathcal{M}) \leq I_{\max}(\mathcal{O}) < \infty$ from Axiom 1.

A second obstruction is operational rather than representational. Since the observer can only preserve finite macroscopic encodings, any physical measurement again factors through a projection $\pi_\Phi : \Omega_\mathcal{E} \to Z$ into a finite set of records. For any nonempty macrostate $z \in Z$, the basin $B_z = \pi_\Phi^{-1}(z)$ contains uncountably many microstates. To isolate the unique true trajectory from this basin, the observer would have to resolve which candidate cell contains the realized state. At any finite refinement containing N distinguishable candidate cells, this discrimination requires at least $\ln N$ nats of additional specification; as the refinement approaches exact continuity, $N \to \infty$ and the required specification diverges.

By Axiom 2, processing this diverging informational acquisition demands a proportional physical resource payment. Under the CCE bound, any irreversible discrimination of ΔI_{irr} nats requires

$$\sum_a \lambda_a \Delta Q_a \geq \Delta I_{\text{irr}}.$$

As exact continuum discrimination sends $\Delta I_{\text{irr}} \to \infty$, the corresponding conserved-resource cost also diverges, guaranteeing that the observer exhausts its finite physical budgets before exact discrimination can be completed.

This divergence is independent of the particular resource accounting scheme. Any physical implementation of discrimination must be carried by a finite substrate, through finite operations, using finite available resources. Exact continuity demands completion of an unbounded refinement procedure, so no embedded observer with finite operational resources can complete this discrimination.

Therefore, the observer's formal system is epistemically cut off from the continuum itself. It receives only finite records $z \in Z$ and finite sequences of such records. Any continuous model fitted to those records is structurally underdetermined by the available physical data, since multiple continuous trajectories remain compatible with the same finite observational history. The continuum may serve as a mathematical representation, but it cannot be physically proven from within an embedded finite observer. □

4 Conclusion

For an embedded subsystem, physics is the science of what can be physically recorded, processed, and conserved; under finite encoding and CCE resource bounds, the observable is necessarily coarse-grained and discrete. The same constraints that make observation possible also impose absolute limits on predictability: they compress highly complex microscopic realities into stable macroscopic basins. If the underlying state space is finite and discrete, exact reversible dynamics remain dynamically intractable because indistinguishable microstates can possess distinguishable futures and pasts. If the underlying state space is continuous, the obstruction becomes stronger: exact continuum states require infinite precision and unbounded operational resources, placing them beyond physical observation.

Epistemic incompleteness is therefore not a paradox of algorithmic undecidability, but a physical constraint. Computing the precise trajectory of the universe is physically intractable without exhausting the observer's finite encoding capacity and conserved-resource budgets; proving a continuum from within the universe is likewise impossible for the same physical reason.

References

[1] David H. Wolpert. Physical limits of inference. *Physica D: Nonlinear Phenomena*, 237(9):1257–1281, 2008.

[2] Seth Lloyd. Computational capacity of the universe. *Physical Review Letters*, 88(23):237901, 2002.

[3] Peter David Fagan. Conservation-Congruent Encoding. Preprint (v2), April 24, 2026.

[4] Jacob D. Bekenstein. Black holes and entropy. *Physical Review D*, 7(8):2333–2346, 1973.

[5] Stephen W. Hawking. Particle creation by black holes. *Communications in Mathematical Physics*, 43(3):199–220, 1975.

[6] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.