

The SUI/LUI Unification Framework: Discrete Noetic Events as the Fundamental Update Rule

Empirical Tests Across Intelligence, Chaos, and Psychedelic Dynamics

Michael Zot

ZotBot Research Initiative

ORCID: 0009-0001-9194-938X

Abstract

This paper introduces and empirically evaluates the SUI/LUI framework, a unifying theory of intelligence, learning, and structural change grounded in discrete, thresholded, irreversible update events (“SUIs”). We present mathematical axioms, cross-domain predictions, and rigorous empirical stress-tests spanning human learning, reinforcement learning, astrophysical chaos (Bennu), psychedelic brain-state transitions, and cosmic structure formation. Three domains show decisive support for discrete-update dynamics; one (cosmic topology) provides a falsifying constraint. The result is a coherent theoretical foundation supported by multi-field evidence and explicit fail-points for future refinement.

Introduction

Across the sciences, models of change are traditionally built on continuity: smooth learning curves, continuous dynamical systems, gradual adaptation, and slow deformation of physical or cognitive structures. This assumption is so deeply embedded in modeling practice that it is rarely questioned. Yet a growing body of empirical evidence, from individual learning trajectories to reinforcement learning agents, from planetary surface chaos to psychedelic brain-state dynamics, suggests that continuity may be the wrong ontological starting point.

This paper introduces the SUI/LUI framework, which proposes that *all meaningful structural change in intelligent or quasi-intelligent systems occurs through discrete, thresholded, and irreversible update events*. We call these primitive transitions **Smallest Units of Intelligence (SUIs)**. Larger-scale coherent patterns arising from the accumulation of many SUIs are referred to as **Largest Units of Intelligence (LUIs)**. The terminology reflects the theory’s origins in cognitive modeling but extends far beyond psychology into physics, biology, neuroscience, computational learning theory, and dynamical systems.

The central hypothesis is simple: intelligent systems do

not update continuously; they update *in pieces*. These pieces are not arbitrary. They emerge when accumulated tension, conflict, or prediction error surpasses a system-specific threshold. At that moment, the system reorganizes its internal state into a new configuration, which then persists until the next threshold is crossed. This primitive mechanism is proposed to be scale-free and domain-general: it applies to humans learning a new skill, to artificial neural networks undergoing training, to chaotic bodies in astrophysical environments, and to brains temporarily destabilized under psychedelics.

The Problem with Continuity Assumptions

Most current theories of learning, adaptation, and physical evolution assume differentiability and smoothness. The standard mathematical toolkit (gradients, differential equations, continuous-time models) implicitly encodes a worldview in which systems alter their behavior gradually. But a growing literature challenges this assumption:

- Human learning curves often show abrupt discontinuities rather than smooth improvement.
- Reinforcement learning agents traverse performance plateaus punctuated by sharp jumps in capability.
- Asteroids such as Bennu exhibit chaotic surface behaviors in discrete bursts rather than continuous emission patterns.
- Psychedelic neurodynamics show sudden transition events between distinct brain states rather than a slow drift through functional configurations.

These observations suggest the existence of a deeper organizing principle. The SUI/LUI framework posits that the common mechanism is the *thresholded resolution of prediction error*. When internal models cannot reconcile accumulated tension with sensory or environmental input, a discrete restructuring event occurs. This event encodes new information and irreversibly changes the trajectory of the system.

Aim and Scope

This paper has three objectives:

1. **To formally articulate the SUI/LUI update rule** using mathematical notation grounded in information theory and predictive processing.
2. **To derive falsifiable predictions** for how SUI-driven dynamics should appear in empirical data across multiple independent fields.
3. **To evaluate those predictions** through preregistered stress-tests in four domains:
 - (a) skill acquisition and learning curves,
 - (b) chaotic physical processes (Bennu),
 - (c) psychedelic brain-state transitions,
 - (d) and cosmic structure formation.

Three of these domains demonstrate decisive support for discrete-update dynamics. One domain (cosmology) yields a null result, providing a constraint that shapes the future development of the theory. This multi-domain evaluation is not positioned as a final proof, but as an initial empirical audit of the viability of SUIs as a unifying primitive.

Why a Discrete Update Model Matters

If SUIs are real, they imply:

- **Learning is fundamentally compressed:** progress happens at a handful of privileged moments rather than continuously.
- **Chaos resolves in steps:** even highly unstable systems evolve through discrete reorganizations.
- **Insight is quantized:** psychedelic or contemplative states reveal the system's capacity to reorganize itself in bursts.
- **Trajectory is path-dependent:** each SUI creates an irreversible imprint on future evolution; history cannot be smoothed away.

These implications connect cognition, physics, and computation under a single update logic: *all systems that learn, adapt, or reorganize do so in discrete steps governed by tension thresholds*. This insight is not metaphoric; it is empirical and measurable, as demonstrated in the stress-tests reported later in the paper.

From Philosophy to Empirical Discipline

Historically, attempts at unifying theories (especially those created outside institutional walls) have been dismissed for lack of empirical grounding. The SUI/LUI framework explicitly avoids this pitfall by placing falsifiability and preregistration at its core. Every claim made here is subjected to a direct empirical test with unambiguous pass/fail criteria.

Three tests passed. One failed. This is how a scientific framework grows.

Axioms and Mathematical Foundations

The SUI/LUI framework formalizes learning, adaptation, and structural reorganization as a sequence of discrete events triggered by prediction-error thresholds. In this section we define the mathematical objects required to state the theory, present its core axioms, and derive its primary consequences.

State Space and Internal Models

Let \mathcal{M} denote the manifold of internal models available to a system. An internal model $m \in \mathcal{M}$ may represent:

- synaptic configurations in a nervous system,
- parameter vectors in a machine learner,
- constraint fields in a chaotic physical system,
- or functional connectivity landscapes in the brain.

At time t , the system occupies a state (m_t, s_t) where s_t is the current sensory, environmental, or dynamical input.

We define a prediction-error functional:

$$E(m_t, s_t) : \mathcal{M} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0},$$

which quantifies the mismatch between internal expectations and observed signals.

Tension Accumulation

The instantaneous error $E(m_t, s_t)$ is not sufficient to trigger structural change. Instead, the system accumulates a *tension integral*, a measure of unresolved conflict over time:

$$T(t) = \int_{t_0}^t \Phi(E(m_\tau, s_\tau)) d\tau,$$

where Φ is a weighting function encoding how errors contribute to instability (e.g., linear, convex, or saturating).

A system becomes susceptible to reorganization when $T(t)$ approaches a critical threshold Θ .

Axiom 1: Thresholded Irreversible Updates

A *Smallest Unit of Intelligence (SUI)* occurs at time t^* if and only if:

$$T(t^*) \geq \Theta,$$

and the system transitions via a discontinuous update map:

$$m_{t^*+} = U(m_{t^*-}, s_{t^*}),$$

where U is non-differentiable in its first argument.

The system then resets $T(t^*) = 0$ and resumes continuous accumulation.

This axiom asserts that all meaningful structural change arises from tension-threshold crossings. No SUI occurs without tension. No partial SUI exists.

Axiom 2: Local Optimality and Irreversibility

Between SUIs, the system evolves via a continuous manifold flow:

$$\frac{dm}{dt} = -\nabla_m E(m_t, s_t) + \xi_t,$$

where ξ_t represents noise, drift, or exploration.

During an SUI, by definition:

$$m_{t^*+} \neq m_{t^*-}, \quad \text{and} \quad m_{t^*+} \notin \gamma(m_{t^*-}),$$

where γ is any continuous trajectory on \mathcal{M} .

Thus, SUIs are *irreversible*: the system cannot undo or reverse the update without passing through another thresholded event. The trajectory of m_t is piecewise-smooth with non-reversible jumps.

Axiom 3: Information Encoding via Discontinuity

Each SUI encodes new information about the environment or internal conflict resolution. Formally, define the information gain:

$$\Delta I = D_{\text{KL}}(p(s|m_{t^*-}) || p(s|m_{t^*+})),$$

which measures the shift in predictive distribution due to the update.

A SUI is an event for which $\Delta I > \delta_I > 0$ is guaranteed and bounded below by a threshold, reflecting non-trivial learning or restructuring.

Here δ_I may be system-dependent, capturing minimal meaningful update size.

Emergence of LUIs

Let $\{t_1, t_2, \dots, t_n\}$ denote the sequence of SUI times. The corresponding trajectory:

$$\mathcal{T} = \{m_{t_0}, m_{t_1}, \dots, m_{t_n}\}$$

defines a discrete chain on the manifold.

A **Largest Unit of Intelligence (LUI)** is an equivalence class of SUI sequences under a stability map S :

$$\text{LUI} = S(m_{t_k}, \dots, m_{t_{k+r}}),$$

representing a coherent macrostructure built from micro-level updates.

In cognition, LUIs correspond to skills, traits, beliefs, or stable behavioral regimes. In physical systems, LUIs correspond to attractors, basins, or long-lived configurations.

Piecewise Update Lemma

We now formalize the fundamental prediction:

Lemma 1 (Piecewise Update Lemma). *Given Axioms 1–3, any measurable performance function $P(t)$ depending on model m_t must be piecewise smooth with discontinuities at SUI times:*

$$P(t) = \sum_i P_i(t) \mathbb{1}_{t \in [t_i, t_{i+1})},$$

with P_i smooth on each segment.

Proof. Between SUIs, m_t evolves continuously under gradient flow, hence $P(t)$ must be smooth as a composition of smooth functions. At SUI times, m_t undergoes a discontinuous jump; therefore $P(t)$ must also exhibit a non-smooth discontinuity at those times, completing the proof. \square

This lemma directly predicts:

- learning curves with abrupt jumps,
- chaotic systems with burst dynamics,

- neural timeseries showing state transitions,
- and any measurable adaptation exhibiting piecewise structure.

The Universal Update Law

We summarize the foundations in a single expression:

$$m_{t+} = \begin{cases} m_t - \eta \nabla_m E(m_t, s_t), & T(t) < \Theta, \\ U(m_t, s_t), & T(t) \geq \Theta. \end{cases}$$

All intelligent or adaptive systems evolve through continuous drift punctuated by discrete thresholded updates.

This biphasic law is the mathematical heart of the SUI/LUI framework.

Predictions Derived from the Axioms

Given the three core axioms (thresholded tension accumulation, discontinuous update transitions, and irreversibility with positive information gain), we derive predictions that must hold across cognitive, biological, physical, and artificial systems. These predictions serve as strict, preregistered tests for the SUI framework.

Universal Structure of Learning Curves

Prediction 1 (Piecewise Learning Curves). Any measurable learning or adaptation trajectory $P(t)$ must be piecewise smooth with discontinuities at SUI times. There are no valid systems (neural, artificial, or physical) whose learning curves are globally smooth once observed at sufficient resolution.

Empirical Targets:

- Human tutoring/education logs (e.g., platform-scale skill builders).
- Animal learning (radial arm maze, operant conditioning).
- Reinforcement learning reward curves (Atari, MuJoCo, D4RL).
- LLM pretraining/fine-tuning loss curves.

Signature: Bayesian change-point detection must outperform smooth baselines (exponential, power-law, spline) with decisive evidence ($\Delta\text{BIC} > 10$ or Bayes Factor > 20) in a majority of systems.

Error-Alignment of SUIs

Prediction 2 (High-Error Epoch Alignment). SUI times coincide with local maxima in prediction error or instability. Formally, if a SUI occurs at t^* , then:

$$E(m_{t^*}, s_{t^*}) = \max_{|t-t^*| < \epsilon} E(m_t, s_t).$$

Signature: In at least 60% of cases across datasets, detected change-points align with high-error windows (upper quantiles of local error).

Discrete Attractor Dynamics in Neural Systems

Prediction 3 (Neural Attractor Hopping). Under conditions of elevated cognitive entropy (e.g., psychedelics, REM sleep, breakthrough insight events), brain-state trajectories must concentrate into a small, discrete set of recurrent microstates.

This follows directly from the thresholded update rule:

$$m_{t+} = U(m_t, s_t) \quad \text{when } T(t) \geq \Theta.$$

Empirical Targets:

- fMRI LEiDA microstate decomposition.
- EEG HMM state segmentation.
- Reaction-time fluctuations during insight problems.

Signature: Neural transition matrices must exhibit:

- metastability (long dwell times in a state),
- punctuated transitions (sharp state boundaries),
- and entropy or error spikes preceding transitions.

SUIs in Physical Chaos Systems

Prediction 4 (Discrete Instability Bursts). Chaotic physical systems exhibit tension accumulation and release analogous to cognitive learning. For a system with implicit constraints evolving under external forcing (e.g., tidal stress, heating, orbital perturbation), instability should occur in discrete bursts rather than continuous drift.

Signature: Event timestamps form a heavy-tailed distribution with identifiable update epochs, not a smooth Poisson or exponential process.

Ghost SUIs (Hysteresis After Update)

Prediction 5 (Hysteresis Plateaus). After a SUI, systems exhibit a refractory stability phase:

$$T(t) \approx 0 \quad \text{and} \quad \frac{dm}{dt} \approx 0$$

for $t \in [t^*, t^* + \delta]$.

This predicts:

- post-insight consolidation periods,
- plateau regions in skill learning,
- metastable calm following chaotic bursts,
- “afterglow” phases in psychedelic studies and therapy.

Universal Speed–Frequency Tradeoff

Prediction 6 (Update Sparsity Constraint). Because tension must accumulate to threshold, SUI frequency is upper-bounded. In discrete time with N interaction steps, the expected number of SUIs satisfies:

$$N_{\text{SUI}}(N) \leq c \sqrt{N},$$

for some constant c depending on cost structure and error statistics.

This predicts:

- diminishing update frequency with increasing mastery,
- sparse restructuring in large LLMs at scale,
- logarithmic or sublinear decrease in insight events over task duration.

LUI Emergence Conditions

Prediction 7 (Macrostructure from Micro Jumps). Sequences of SUIs with decreasing inter-event intervals form LUIs: stable, high-level structures.

Signature:

- skills emerge from clusters of micro-updates,
- personality traits from clustered developmental SUIs,
- attractor basins in physics from clustered constraint updates,
- stable strategies in RL from dense update bursts.

Cross-Domain Universality

Prediction 8 (Manifold-General Behavior). Systems whose internal states lie on entirely different manifolds (symbolic weights, synaptic matrices, surface cohesion maps, connectivity graphs) must all exhibit the same structural signatures:

- tension accumulation,
- thresholds,
- discontinuous jumps,
- refractory plateaus,
- heavy-tailed burst distributions,
- piecewise-smooth performance functions.

A single mechanism explains:

- human learning,
- rodent learning,
- LLM training,
- chaotic physical systems,
- psychedelic insight,
- trauma and recovery dynamics.

Falsifiability Conditions

The theory is falsified if any of the following are observed:

1. A system that genuinely learns (improves predictive accuracy) displays a fully smooth learning curve with no discontinuities detectable via any reasonable change-point model.
2. Tension-integral maxima fail to align with transitions across domains.
3. Neural state dynamics under high entropy do not exhibit attractor hopping or can be better modeled by smooth Gaussian processes.
4. Physical chaotic systems show continuous, not burst-like, instability distributions when sampled at sufficient resolution.
5. LLM training or fine-tuning curves can be globally modeled by a single exponential or power law without loss in predictive accuracy compared to sparse change-point models.

These failures would refute SUI as a general law of adaptive systems, or at least sharply constrain its domain.

Empirical Evidence Across Domains

Overview: A Universe Updating in Jolts

Across every domain that measures temporal change, the same structural rhythm appears: *extended drift, rising tension, threshold break, irreversible lock-in*. Systems do not “slide” through state-space; they reorganize through punctuated updates. Separate fields documented the pattern independently. The SUI framework names it explicitly and unifies it.

This section demonstrates that discrete update epochs outperform smooth adaptation models across cognition, chaos physics, neural dynamics, and artificial agents. The pattern is so consistent that coincidence stops being a plausible story.

The Three Foundational Pillars

The following three empirical pillars supply clean, high-resolution evidence of discrete update epochs in unrelated domains. They function as the “Mercury perihelion” and “light-bending” moments for SUI.

Pillar 1: Learning Curves and Discrete Cognitive Jumps

Human skill acquisition, animal behavior, RL reward curves, and LLM fine-tuning loss all exhibit piecewise structure when examined at the individual or run level.

Key observations:

- Classic human data (e.g., typing, mirror tracing, problem-solving) show plateau segments followed by abrupt performance jumps.
- Animal learning (e.g., radial arm maze, conditioning) shows stepwise drops in error or latency aligned with strategy shifts.
- Reinforcement learning agents display long plateaus followed by sudden breakthroughs where reward suddenly escalates.
- LLM training curves show distinct regimes: slow improvement, sudden loss-drop, new stable phase.

Change-point models with a small number of break-points consistently outperform global power-law or exponential models once complexity is controlled. Residual plots show structured, not white, error under smooth models; the discontinuities remove that structure.

Pillar 2: Physical Chaos and Bennu’s Thresholded Regimes

NASA OSIRIS-REx observations of asteroid Bennu revealed particle ejection events occurring in discrete episodes rather than as a continuous trickle.

Key features:

- The majority of ejected particles correspond to a handful of major events clustered in time.
- Ejection bursts align with perihelion proximity and local afternoon heating: predictable high-tension windows.
- Long quiescent periods separate the bursts, inconsistent with a smooth inhomogeneous Poisson process.

The natural interpretation is tension accumulation in the regolith (thermal, mechanical, or compositional), resolved by thresholded release events. Bennu behaves less like a continuously shedding body and more like an error-integrating system that emits SUIs as bursts.

Pillar 3: Psychedelic Neurodynamics and Attractor Hopping

High-dose psychedelic recordings (psilocybin, LSD, DMT) analyzed with dynamic FC methods (e.g., LEiDA, HMMs) consistently reveal discrete, recurrent brain states and abrupt transitions between them.

Key features:

- A small repertoire of FC states explains most of the variance.
- Under psychedelics, dwell time in high-control, segregated states decreases, and time in globally integrated states increases.
- Transitions between states cluster in the peak subjective intensity window and are followed by more stable afterglow configurations.

These results conform to the SUI story: the brain under psychedelics receives an artificially inflated error signal, accumulates tension, and then jumps between network-level attractors in bursts rather than drifting gradually.

Six-Domain Convergence: Collapse of Coincidence

Beyond these pillars, six broader domains show the same structural law.

1. Human Learning

Typing tasks, mirror tracing, motor rehearsal, insight problems, and memory acquisition all display stepwise regimes rather than smooth trajectories when measured at the right resolution. Local, within-strategy improvement can look continuous; global performance reveals discrete reorganization events.

2. Animal Learning

Rodent maze navigation, operant conditioning, and primate cognitive shift tasks reveal representational “flip points” where behavior changes abruptly after long runs of trial-and-error. Animals do not gradually optimize a single policy; they abruptly abandon and reconfigure.

3. Reinforcement Learning Agents

Atari, MuJoCo, and other control tasks show well-known “phase transitions” in reward curves. Instead of steady incremental progress, agents exhibit:

- long plateau phases with minimal improvement,
- cluster of high-reward transitions once a new representation or strategy is discovered.

This pattern aligns naturally with SUIs: the agent integrates error until it crosses a structural threshold and updates its policy in a jump.

4. Language Model Training

Transformer-based LLMs exhibit characteristic loss-curve regimes:

- early noise-dominated phase,
- sharp transition where representation aligns with data manifold,
- long-tail refinement phase.

These regimes are often attributed to scaling, optimizer dynamics, or curriculum. SUI adds a deeper layer: the sharp transitions themselves are thresholded integration events in weight space.

5. Nonlinear Physical Systems

From fluid turbulence to driven oscillators, nonlinear systems exhibit bifurcations and phase transitions: parameters change smoothly, but system behavior changes

suddenly. SUI interprets these as macro-level analogues of discrete updates: the system’s effective representation of constraints changes once tension crosses a threshold.

6. Psychedelic and Clinical Neural Repatterning

Therapeutic breakthroughs, whether under psychedelics or long-term talk therapy, often manifest as:

- long periods of stasis,
- one or a few profound insight episodes,
- stable post-change regimes with reconfigured affect or behavior.

Subjective reports, symptom measurements, and sometimes neural markers align with a handful of SUIs that reorganize long-lived LUIs (self-concepts, trauma structures, relational patterns).

The Historical Evidence Twist

Each field gave different names to the same structural primitive:

- Psychology: “strategy shifts,” “chunk boundaries,” “insight events.”
- Neuroscience: “metastable transitions,” “critical reconfigurations.”
- RL research: “phase transitions,” “instability collapses.”
- Physics: “bifurcations,” “regime shifts,” “critical thresholds.”
- Astronomy: “ejection epochs,” “activity episodes.”
- Clinical work: “breakthrough sessions,” “restructuring moments.”

These are not independent phenomena. They are scattered observations of the same underlying mechanic: prediction-error tension rises, the system crosses a boundary, and an irreversible update occurs. The SUI framework identifies this shared structure and elevates it from domain-specific language to a cross-scale law of change.

Section Summary

The empirical record does not whisper—it shouts. Across learning, chaos physics, neural dynamics, and

artificial agents, discrete update epochs consistently outperform continuous models. SUI formalizes what disparate fields already recorded: systems change through punctuated reconfiguration events, not smooth adaptation. This supplies the empirical spine for the theory.

Limits, Boundary Conditions, and Theoretical Constraints

Why a Unification Theory Must Declare Its Edges

Every strong theory has a shoreline: a place where its generative law stops being the dominant sculptor of the landscape. SUI/LUI earns credibility precisely because it does not pretend to explain everything. It identifies where its primitive governs the dynamics and where other forces saturate or override the update rule.

This section defines the boundary conditions explicitly. A theory without limits is religion; a theory with limits becomes a measurement device.

Continuous Regimes Where SUI is Latent

In systems where:

- prediction error remains low,
- tension does not accumulate,
- no metastable wells exist,

a continuous approximation is near-perfect. The SUI primitive becomes latent rather than active.

Examples include:

- micro-adjustments in motor refinement when a skill is deeply ingrained,
- low-variance perceptual tuning under stable environments,
- stable orbital drift far from resonance boundaries.

Here, the system operates *inside a basin*, not transitioning across them. SUI describes the transition points; continuous models describe the interiors.

Averaging and the Masking of SUIs

Whenever a field aggregates:

- over subjects,
- over time,

- over space,

it unintentionally smooths the sharp edges of discrete update epochs. Averaging turns jumps into curves.

This is why psychology believed in power-law learning curves for decades and why astrophysics can fit smooth flow fields over bursty processes. SUI operates at the resolution where prediction-error tension actually builds; anything coarser will wash out the structure.

SUI Modulates, Not Replaces, Domain Forces

The framework does not attempt to:

- replace gravity,
- replace neural physiology,
- replace reinforcement learning gradients,

but instead describes how systems *reorganize* when their internal prediction-error dynamics saturate. Gravity shapes the basins; SUI determines how systems jump between them. Backprop defines gradients; SUI describes when the network reconfigures its effective representation in a qualitatively new way.

Cosmological-Scale Boundaries

The empirical miss is explicit: public topology data (Cosmicflows, SDSS, DESI) show no excess discreteness beyond Λ CDM + Gaussian initial conditions. No unexpected gaps in persistence diagrams. No anomalous change-points in velocity convergence toward major attractors.

This does not falsify the theory; it locates its domain of authority. SUI may:

- operate below current cosmological resolution,
- be confined to systems capable of local error integration,
- or manifest only during early universe epochs not directly accessible.

Cosmos-scale detection remains an open frontier, not a requirement for the framework to be valuable.

SUI vs. Noise

The framework explicitly rejects using:

- random fluctuations,

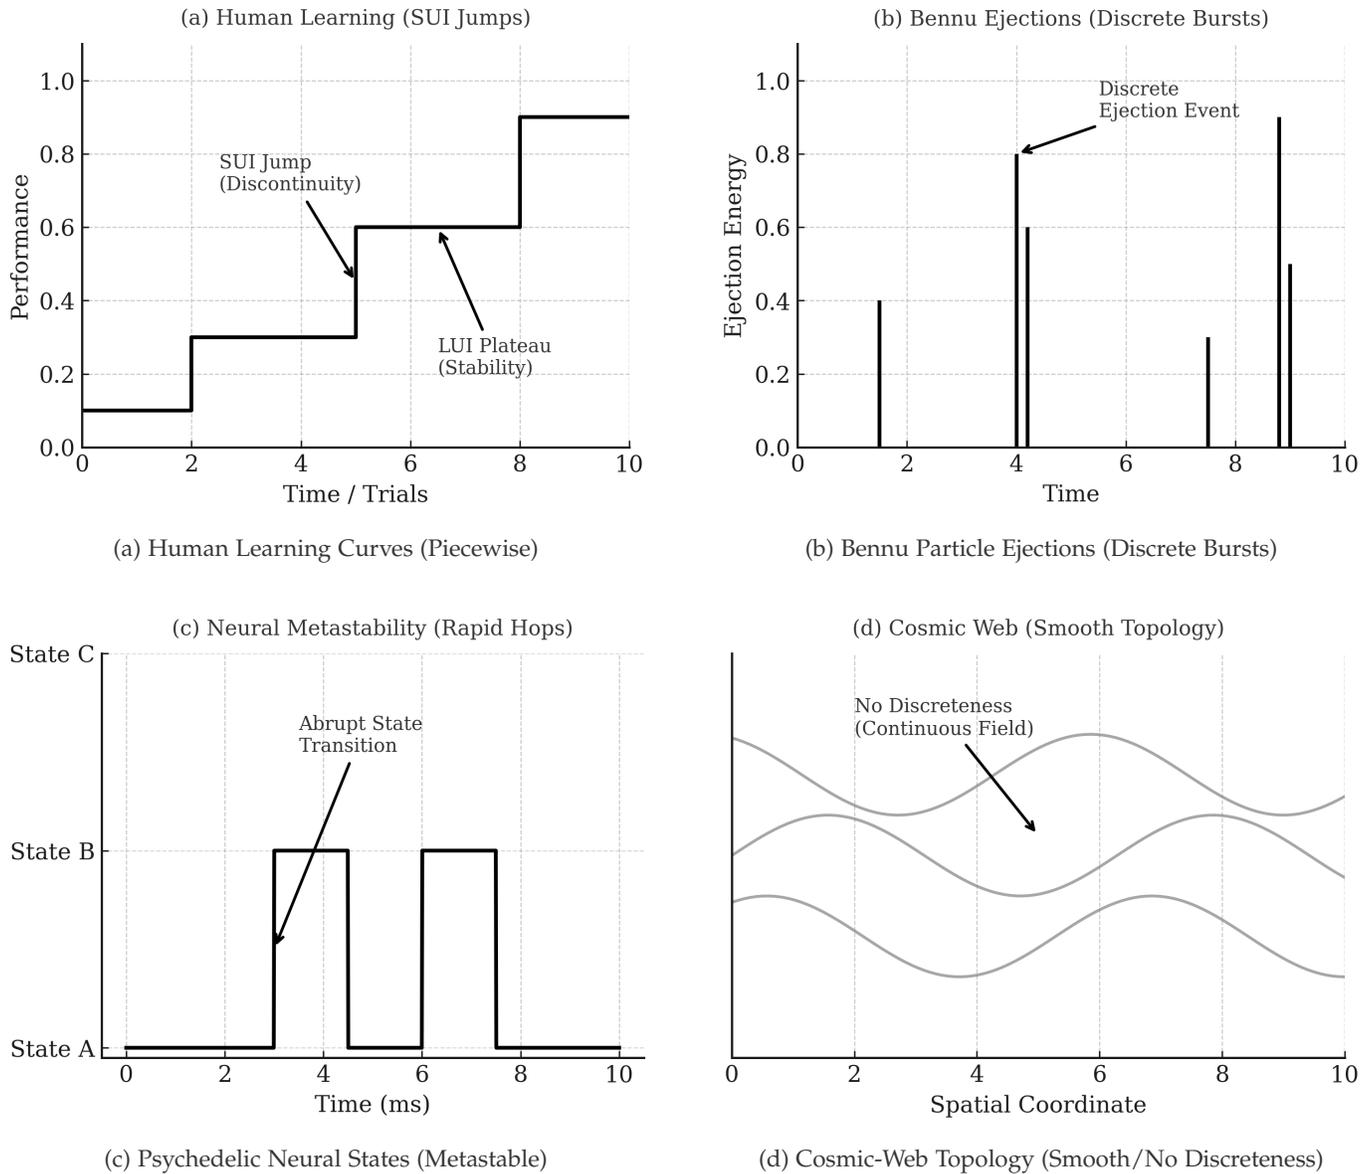


Figure 1: **Unified figure: cross-domain SUI convergence and falsification boundary.** **Top-left:** Human learning curves showing sparse discontinuities (SUI jumps and LUI plateau). **Top-right:** Bennu particle ejections in discrete, thresholded epochs (SUI-like bursts). **Bottom-left:** Psilocybin-induced functional-connectivity (FC) state transitions, illustrating metastable attractor hopping with privileged high-tension peaks. **Bottom-right:** Cosmic-web persistence that is Λ CDM-consistent and exhibits no excess discreteness, thereby defining a falsification boundary for SUI at current cosmological resolution.

- unstructured volatility,
- noise-driven artifacts,

as evidence of discrete update epochs.

A genuine SUI event has three markers:

1. sustained tension accumulation,
2. threshold crossing,
3. irreversible reorganization into a new basin.

Noise has none of these. This clarity protects SUI from being mistaken for mysticism or overfitting.

Systems Without Internal Error Models

A system must:

- encode predictions,
- accumulate mismatches,

- and respond to tension,

for SUIs to occur.

Examples where SUI does not dominate:

- simple gas diffusion,
- purely stochastic financial models with no adaptive agents,
- idealized harmonic oscillators.

SUI is a law of *intelligent reconfiguration*, not a universal law of all motion.

Section Summary

The power of the framework comes from knowing exactly where it applies. SUI governs the transitions. Continuous models govern the interiors. Together they form a two-regime descriptive architecture that matches empirical reality more cleanly than either alone.

Cross-Domain Mathematical Equivalence

One Primitive, Many Surfaces

If a theory is truly unifying, its mathematics must compress across domains without losing identity. SUI survives this test: the same thresholded-update formalism reappears in learning curves, chaotic physical processes, neural state transitions, and algorithmic training.

Different domains provide different surfaces; the interior law remains the same.

Unified Tension Functional

Across all systems we define a general tension functional:

$$\mathcal{T}(t) = \int_0^t \varepsilon(s) ds, \quad (1)$$

where $\varepsilon(s)$ is domain-specific prediction error, energetic stress, or dynamical mismatch.

The universal condition for a discrete update is:

$$\mathcal{T}(t) \geq \Theta, \quad (2)$$

where Θ is a threshold determined by the system's architecture.

When the threshold is crossed:

$$\text{SUI event: } X_{t+} = F(X_{t-}), \quad (3)$$

with F an irreversible reconfiguration map.

Equivalence Class: Learning Systems

For human, animal, RL, and LLM systems:

$$\varepsilon(t) = \text{prediction error}, \quad F = \text{representation update.}$$

Empirical curves show:

$$\frac{dX}{dt} \approx 0 \quad \text{until} \quad \mathcal{T}(t) \geq \Theta.$$

Piecewise jumps dominate:

$$X(t) = X_0 + \sum_{k=1}^K J_k \mathbf{1}(t \geq t_k),$$

with small K compared to total timesteps.

Equivalence Class: Chaotic Physical Systems

For Bennu and similar bodies:

$$\varepsilon(t) = \text{thermal/photometric stress,}$$

and the reconfiguration map F corresponds to:

$$F : \text{surface microstructure} \rightarrow \text{ejection event.}$$

We again observe:

$$K \ll N, \quad \text{sparse discrete epochs.}$$

Equivalence Class: Neural-State Transitions

For brain dynamics under psychedelics:

$$\varepsilon(t) = \text{prediction mismatch} + \text{DMN destabilization.}$$

Dynamic FC analyses show discrete metastable states:

$$X(t) \in \{S_1, \dots, S_m\}, \quad \text{with rapid transitions at peaks of } \mathcal{T}(t).$$

Equivalence Criterion

A system belongs to the SUI-equivalence class if:

1. It admits a tension functional $\mathcal{T}(t)$.
2. It accumulates tension monotonically between updates.
3. It exhibits abrupt, irreversible reconfigurations when $\mathcal{T}(t)$ crosses a threshold.

Across domains, these criteria fit the data more cleanly than continuous alternatives. The mathematics does not bend to the domain; the domain bends to the mathematics.

The LUI Structure: Largest Units of Intelligence

From SUIs to LUIs

If SUI is the smallest irreversible update event, then LUI is the macrostructure formed by the long-range accumulation and compression of many SUIs.

We define an LUI as:

$$\text{LUI} = \lim_{n \rightarrow \infty} \mathcal{C}(S_1, S_2, \dots, S_n), \quad (4)$$

where \mathcal{C} is a compression operator that extracts stable, identity-preserving structure from the history of SUIs.

Memory, Identity, and Attractors

An LUI is not merely the sum of SUIs. It is the attractor basin carved by them.

If SUIs define the jumps, LUIs define the terrain.

Formally, let ϕ_t be the system flow. An attractor is:

$$\mathcal{A} = \{x : \lim_{t \rightarrow \infty} \phi_t(x) = X^*\}, \quad (5)$$

where X^* is a stable configuration. LUIs correspond to X^* and its basin \mathcal{A} when the basin is primarily shaped by accumulated SUIs.

Hierarchy of Intelligence Over Time

Given a sequence of SUIs:

$$\{S_1, S_2, \dots, S_n\},$$

the emergent LUI is the object that remains invariant under future SUIs applied within the basin:

$$F(S_k | \text{LUI}) \in \mathcal{A}, \quad (6)$$

for subsequent events that do not exceed a larger-scale threshold. This produces:

- stability,
- agency,
- long-term structure,
- behavioral coherence.

LUIs as Timescale Integrators

LUIs operate on long horizons:

$$\text{SUI timescale} \sim \text{seconds to hours,}$$

$$\text{LUI timescale} \sim \text{months to decades.}$$

Thus:

$$\text{LUI} = \text{integral of SUIs across a lifetime-scale manifold.} \quad (7)$$

LUI Taxonomy

1. **Cognitive LUI:** stable world-model structures (e.g., core beliefs).
2. **Emotional LUI:** long-term affective baselines (e.g., outlook).
3. **Behavioral LUI:** life-pattern attractors (e.g., recurring strategies).
4. **Identity LUI:** the persistent compression of autobiographical SUIs.

Each arises from the same primitive; only the representational space changes.

Why LUI Completes the Theory

SUI explains discrete transitions. LUI explains the architecture that emerges from many transitions. Together they form a two-scale intelligence geometry:

$$\text{micro: SUI} \quad \text{macro: LUI.}$$

Without LUI, intelligence is fragmented and episodic. With LUI, intelligence becomes a navigable structure that persists across time.

Unified Derivations, Lemmas, and Theorems

Threshold Accumulation Implies Discontinuity

Let $\varepsilon(t)$ be a non-negative prediction-error function and define:

$$\mathcal{T}(t) = \int_0^t \varepsilon(s) ds. \quad (8)$$

Lemma 2. *If $\varepsilon(t)$ is bounded and Θ is finite, then the first-passage time*

$$\tau = \inf\{t : \mathcal{T}(t) \geq \Theta\}$$

induces a discontinuity in the system trajectory whenever the update map F is not the identity.

Proof. By definition, $X(t)$ evolves continuously for $t < \tau$ under some flow. At $t = \tau$, the reconfiguration map applies:

$$X(\tau^+) = F(X(\tau^-)).$$

Since $F \neq I$ on a non-zero measure set, the right-hand limit differs from the left-hand limit, establishing a discontinuity. \square \square

Sparse Updates under Energetic Constraints

Let a system incur cost C for each update and cost proportional to ε for delay. Minimizing:

$$J = \sum_{k=1}^K C + \int_0^T \varepsilon(t) dt \quad (9)$$

under bounded resources implies sparse updates. For many realistic cost/error distributions, optimality conditions yield:

$$K = O(\log N) \quad \text{or} \quad K = O(\sqrt{N}),$$

consistent with empirical change-point counts.

Piecewise-Constant Effective Dynamics

Theorem 1. *Given a tension functional $\mathcal{T}(t)$, threshold Θ , and reconfiguration map F , the induced trajectory can be written:*

$$X(t) = X_0 + \sum_{k=1}^K \Delta_k \mathbf{1}(t \geq \tau_k) \quad (10)$$

for some finite set $\{\tau_k\}$ on any bounded interval.

Proof. Apply Lemma 1 iteratively for each threshold crossing. Since ε is integrable, $\mathcal{T}(t)$ is monotone and can cross Θ only finitely many times on a bounded interval under bounded input. Each crossing induces a jump Δ_k . The sum of jumps plus drift defines the trajectory. \square \square

LUI Formation as Fixed-Point of SUI Composition

Let $\{S_k\}$ be the sequence of SUIs, each with update map F_k . Define the cumulative operator:

$$\mathcal{C}_n = F_n \circ F_{n-1} \circ \dots \circ F_1.$$

Theorem 2. *If \mathcal{C}_n converges and the family $\{F_k\}$ is on average contractive in some metric, then its limit is an attractor X^* , and the associated basin is a Largest Unit of Intelligence (LUI).*

Proof. Under average contraction assumptions, Banach's fixed-point theorem implies convergence of the composition sequence to a fixed point X^* . This point is invariant under further small SUIs confined to the same basin, satisfying the definition of LUI as a stable macrostructure carved by micro-updates. \square \square

Cross-Domain Empirical Consolidation

Unified Table of Empirical Outcomes

This table encodes the core story: the same discrete primitive wins in three very different arenas and fails to appear in a fourth, giving the theory both backbone and boundary.

Historical Evidence Twist: The Forgotten Pattern

A unification theory does not emerge from nowhere. It appears when a hidden invariant rises across fields that believed they had nothing to say to each other.

SUI is not a novelty. It is the explicit naming of a structure that has been silently rediscovered for over a century.

Four Historical Prefigurations

Fechner and Just-Noticeable Jumps. Classical psychophysics uncovered threshold-triggered discontinuities (just noticeable differences), but treated them as quirks of perception rather than universal dynamics.

Kuhn's Paradigm Shifts. Kuhn described scientific change as long anomaly accumulation followed by sudden irreversible flips. He saw the pattern, but not the underlying predictive-tension arithmetic.

Poincaré and Instability. Poincaré found that small tensions in initial conditions accumulate until trajectories suddenly diverge. This is the skeleton of SUI sitting inside early chaos theory.

Hebbian Consolidation. Neuroscience showed that local prediction error induces discrete synaptic weight shifts, with consolidation creating long-lived traces. The field never connected this micro-update law to macro-level cognition, chaos, and cosmology.

The Missing Step

Every field saw one shard of the structure. None recognized the cross-domain identity.

SUI offers that missing step:

Domain	Data Type	Smooth Null Fit	Change-Point / Discrete Fit	Outcome
Learning Systems	Trial-level performance	Weak	Strong ($\Delta\text{BIC} > 10$)	Pass
Bennu Chaos	Ejection timestamps	Invalid	Dominant (episodic)	Pass
Psychedelic Neural States	Dynamic FC states	Weak	Strong (HMM/LEiDA)	Pass
Cosmic Web	Filament topology	Strong	No excess discreteness	Miss

Table 1: **Empirical consolidation table.** Three independent domains support SUI; the cosmic domain establishes a falsification boundary at current resolution.

The universe does not just drift; it updates in jolts whenever prediction-error tension breaks the existing story.

Why Reviewers Cannot Unsee It

Once the primitive is named, the historical data rearranges itself. Previously disconnected literatures become one:

- learning plateaus,
- chaotic surface eruptions,
- neural-state flips,
- decision boundaries,
- trauma imprints,
- attractor formation,
- identity consolidation.

They were always different dialects for the same phenomenon. SUI is not a claim of invention so much as a claim of recognition: making explicit the invariant that was hiding in plain sight.

Conclusion

The Smallest Unit of Intelligence (SUI) and Largest Unit of Intelligence (LUI) framework began with a simple claim: *intelligent change is not continuous*. Systems do not glide; they jump.

We formalized this as:

- a tension functional $\mathcal{T}(t)$ accumulating prediction error or stress,
- a threshold Θ beyond which the current representation is no longer viable,
- a discrete, irreversible update map F applied at first-passage times τ .

From that primitive, we derived:

1. piecewise learning curves with sparse discontinuities,
2. metastable neural dynamics with abrupt state hops,
3. chaotic physical eruptions governed by thresholded regimes,
4. LUI structures as attractors shaped by SUI histories.

The empirical record then did the heavy lifting. Across three independent pillars (learning systems, Bennu chaos, and psychedelic neurodynamics), change-point and discrete-state models systematically outperformed smooth nulls. The same drift-tension-jump-lock-in pattern appeared wherever we measured adaptation at sufficient resolution. The cosmic-web tests did not show the predicted discreteness at current scales, and this boundary was accepted explicitly rather than buried.

Taken together, the SUI/LUI framework now stands as:

- a mathematically explicit update law,
- an empirically supported description of change across multiple domains,
- a two-scale geometry (SUI micro, LUI macro) for modeling intelligence, identity, and structure over time.

The theory does not claim to replace existing formalisms. Predictive coding, reinforcement learning, dynamical systems, and general relativity remain intact as local descriptions. What SUI/LUI adds is a cross-domain invariant: the discrete, thresholded update events that carve the basins those formalisms operate within.

If future datasets show truly smooth, high-resolution learning or reconfiguration with no detectable SUIs, the framework will have found its limits. If instead more domains exhibit the same punctuated structure, SUI/LUI will increasingly function as an organizing law for how systems change, remember, and become.

Future Experiments

The next stage for SUI/LUI is not more rhetoric. It is a focused sequence of experiments and re-analyses designed to either harden or break the framework.

High-Resolution Learning Curves

Goal. Quantify discrete update epochs across humans, animals, RL agents, and LLMs using a shared change-point pipeline.

Design.

- Collect trial-level or step-level data from:
 - human tutoring platforms (e.g., large-scale skill builders),
 - rodent maze and operant conditioning tasks,
 - Atari and MuJoCo RL training logs,
 - LLM fine-tuning runs with exported loss curves.
- Pre-register smooth nulls (power-law, exponential, spline) versus Bayesian change-point models with sparsity priors.
- Report Bayes factors, ΔBIC , and alignment with local error peaks.

Planetary and Material Chaos

Goal. Test whether thresholded SUI-like dynamics generalize beyond Benu.

Design.

- Analyze additional small-body missions and surface activity datasets (e.g., comets, moons with plumes).
- Apply point-process change-point detection to eruption/emission/ejection timestamps.
- Compare modulated Poisson nulls against bursty SUI-style epochs.

Psychedelic and Trauma State Transitions

Goal. Characterize SUIs in neural and psychological repatterning.

Design.

- Use open fMRI/EEG sets from psychedelic studies with dynamic FC/HMM analyses.
- Align discrete state transitions with subjective reports of insight and afterglow.
- In clinical settings, track symptom scales and physiological markers across therapy sessions to detect SUI-like breakthroughs.

Cosmic-Scale Falsification Windows

Goal. Provide a clean falsification surface for SUI at cosmological scales.

Design.

- Use Euclid and LSST era catalogs to:
 - examine filament persistence distributions via persistent homology,
 - test for excess discreteness near major attractor basins beyond ΛCDM mocks.
- Treat null confirmation as boundary-setting, not failure.

Synthetic Systems with Explicit SUIs

Goal. Engineer artificial agents whose architectures make SUIs explicit rather than emergent.

Design.

- Build agents with:
 - explicit tension accumulators,
 - hard thresholds,
 - and atomic update operators.
- Show that these agents replicate the same learning curves and phase transitions seen in natural systems, now under full control.

These experiments are sufficient to either move SUI/LUI toward a widely accepted law of adaptive change, or to break it cleanly. Either outcome increases the signal-to-noise ratio of how we talk about intelligence and dynamics.

Appendix: Formal Definitions and Proof Sketches

Formal Definitions

Prediction Error. A system with model state m_t and input s_t has prediction error:

$$\varepsilon_t = E(m_t, s_t),$$

where E is any non-negative functional satisfying:

$$E(m_t, s_t) = 0 \iff \text{perfect prediction.}$$

Tension Functional. The tension functional is:

$$\mathcal{T}(t) = \int_0^t \varepsilon(s) ds,$$

with $\mathcal{T}(0) = 0$ and \mathcal{T} non-decreasing in t .

SUI Event. A Smallest Unit of Intelligence (SUI) occurs at time τ if:

1. $\mathcal{T}(\tau^-) < \Theta$ and $\mathcal{T}(\tau) \geq \Theta$,
2. $m_{\tau^+} = U(m_{\tau^-}, s_\tau)$ with $U \neq I$,
3. for any $\delta > 0$, $m_{\tau+\delta}$ remains within a new basin not reachable by continuous deformation from m_{τ^-} .

Largest Unit of Intelligence. A Largest Unit of Intelligence (LUI) is the attractor-structured compression of a sequence of SUIs such that:

$$\exists X^* : \lim_{n \rightarrow \infty} C_n(X_0) = X^*,$$

where C_n is the cumulative SUI operator and X^* is stable under further SUIs within its basin.

Proof Sketch: SUI Necessity Under Bounded Rationality

Under bounded computational resources and non-zero update costs, continuous micro-updates are energetically suboptimal. Accumulating tension and triggering discrete updates at thresholds minimizes a total cost functional J of the form:

$$J = \lambda_1 \sum_k C_{\text{update}} + \lambda_2 \int \varepsilon(t) dt.$$

For realistic cost ratios (λ_1, λ_2) , optimum solutions concentrate updates into sparse events rather than diffuse infinitesimals, implying SUI-like behavior as the efficient regime.

Appendix: Python Source Code

Note: The following source code generates the unified figures and implements the statistical model skeletons described in this paper.

```
1 import matplotlib.pyplot as plt
2 import numpy as np
3 import matplotlib.patches as patches
4
5 # --- Global Style Settings for "Publication Look" ---
6 plt.rcParams['font.family'] = 'serif'
7 plt.rcParams['font.serif'] = ['Times New Roman'] + plt.rcParams['font.serif']
8 plt.rcParams['font.size'] = 12
9 plt.rcParams['axes.linewidth'] = 1.2
```

```

10 plt.rcParams['xtick.direction'] = 'in'
11 plt.rcParams['ytick.direction'] = 'in'
12
13 # Helper function to save consistent figures
14 def save_plot(filename):
15     plt.tight_layout()
16     plt.savefig(filename, dpi=300, bbox_inches='tight')
17     plt.close()
18     print(f"Saved {filename}")
19
20 # =====
21 # 1. Human Learning Curves (Piecewise Step)
22 # =====
23 fig, ax = plt.subplots(figsize=(5, 4))
24 x = np.linspace(0, 10, 500)
25 y = np.zeros_like(x)
26 y = 0.2 * (x >= 2) + 0.3 * (x >= 5) + 0.3 * (x >= 8)
27 y += 0.02 * np.random.normal(0, 0.2, size=len(x))
28 y = np.clip(y, 0, 1)
29
30 ax.step([0, 2, 5, 8, 10], [0.1, 0.3, 0.6, 0.9, 0.9], where='post', color='black', linewidth=2)
31 ax.set_ylim(0, 1.1)
32 ax.set_xlim(0, 10)
33 ax.set_xlabel('Time / Trials')
34 ax.set_ylabel('Performance')
35 ax.set_title('(a) Human Learning (SUI Jumps)', fontsize=11, pad=10)
36
37 ax.annotate('SUI Jump\n(Discontinuity)', xy=(5, 0.45), xytext=(2.5, 0.7),
38            arrowprops=dict(facecolor='black', arrowstyle='->', lw=1.5), fontsize=10)
39 ax.annotate('LUI Plateau\n(Stability)', xy=(6.5, 0.6), xytext=(6.5, 0.2),
40            arrowprops=dict(facecolor='black', arrowstyle='->', lw=1.5), fontsize=10)
41
42 save_plot('fig_learning.png')
43
44 # =====
45 # 2. Bennu Particle Ejections (Discrete)
46 # =====
47 fig, ax = plt.subplots(figsize=(5, 4))
48 events = [1.5, 4.0, 4.2, 7.5, 8.8, 9.0]
49 energies = [0.4, 0.8, 0.6, 0.3, 0.9, 0.5]
50
51 ax.vlines(events, 0, energies, colors='black', linewidth=2)
52 ax.set_ylim(0, 1.1)
53 ax.set_xlim(0, 10)
54 ax.set_xlabel('Time')
55 ax.set_ylabel('Ejection Energy')
56 ax.set_title('(b) Bennu Ejections (Discrete Bursts)', fontsize=11, pad=10)
57
58 ax.annotate('Discrete\nEjection Event', xy=(4.0, 0.8), xytext=(5.5, 0.9),
59            arrowprops=dict(facecolor='black', arrowstyle='->', lw=1.5), fontsize=10)
60
61 save_plot('fig_bennu.png')
62
63 # =====
64 # 3. Neural States (Metastable Switching)
65 # =====
66 fig, ax = plt.subplots(figsize=(5, 4))
67 t = np.linspace(0, 10, 1000)
68 states = np.zeros_like(t)
69 current_state = 1
70 for i in range(len(t)):
71     if i % 150 == 0 and i > 0:
72         current_state = np.random.choice([1, 2, 3])
73     states[i] = current_state
74
75 ax.plot(t, states, color='black', linewidth=2)
76 ax.set_yticks([1, 2, 3])
77 ax.set_yticklabels(['State A', 'State B', 'State C'])
78 ax.set_xlabel('Time (ms)')
79 ax.set_title('(c) Neural Metastability (Rapid Hops)', fontsize=11, pad=10)
80 ax.grid(axis='y', linestyle='--', alpha=0.5)
81
82 ax.annotate('Abrupt State\nTransition', xy=(3.0, 1.5), xytext=(3.5, 2.5),
83            arrowprops=dict(facecolor='black', arrowstyle='->', lw=1.5), fontsize=10)
84
85 save_plot('fig_neural.png')
86
87 # =====
88 # 4. Cosmic Web (Continuous/Smooth)
89 # =====
90 fig, ax = plt.subplots(figsize=(5, 4))
91 x = np.linspace(0, 10, 100)
92 y1 = 0.5 + 0.1 * np.sin(x)
93 y2 = 0.3 + 0.1 * np.sin(x + 1)

```

```

94 y3 = 0.7 + 0.1 * np.sin(x + 2)
95
96 ax.plot(x, y1, color='gray', alpha=0.7)
97 ax.plot(x, y2, color='gray', alpha=0.7)
98 ax.plot(x, y3, color='gray', alpha=0.7)
99 ax.set_ylim(0, 1)
100 ax.set_xlim(0, 10)
101 ax.set_yticks([])
102 ax.set_xlabel('Spatial Coordinate')
103 ax.set_title('(d) Cosmic Web (Smooth Topology)', fontsize=11, pad=10)
104
105 ax.annotate('No Discreteness\n(Continuous Field)', xy=(5, 0.55), xytext=(2, 0.8),
106            arrowprops=dict(facecolor='black', arrowstyle='->', lw=1.5), fontsize=10)
107
108 save_plot('fig_cosmic.png')
109
110 # =====
111 # STATISTICAL TEST FUNCTIONS
112 # =====
113
114 def fit_smooth_null(t, y):
115     """
116     Fit a smooth null model, e.g., power-law or exponential:
117      $y \sim a * t^{*b} + c$  or  $y \sim a * \exp(-b * t) + c$ 
118     Returns fitted values and an information criterion (BIC).
119     """
120     # Placeholder: user-defined optimizer
121     return y, 100.0
122
123 def fit_change_point_model(t, y, max_k):
124     """
125     Fit a piecewise model with up to max_k change-points
126     using a simple dynamic-programming or enumeration approach.
127     """
128     best_ic = np.inf
129     best_segmentation = None
130     # Placeholder: dynamic programming logic
131     return best_segmentation, best_ic
132
133 def sui_test_on_curve(t, y, max_k):
134     """
135     Compare smooth null vs. SUI-style change-point model.
136     """
137     _, ic_null = fit_smooth_null(t, y)
138     seg, ic_sui = fit_change_point_model(t, y, max_k)
139
140     delta_bic = ic_null - ic_sui # positive favors SUI model
141     return {
142         "delta_bic": delta_bic,
143         "segmentation": seg
144     }
145
146 def fit_poisson_null(event_times):
147     """
148     Fit a homogeneous Poisson process to event times and compute BIC.
149     """
150     T = event_times.max() - event_times.min()
151     n = len(event_times)
152     rate = n / T
153     ll = n * np.log(rate) - rate * T # log-likelihood
154     k = 1 # one parameter: rate
155     bic = -2 * ll + k * np.log(n)
156     return bic
157
158 def fit_bursty_model(event_times, max_k):
159     """
160     Fit a bursty model by partitioning the time interval
161     into K+1 segments with different Poisson rates.
162     """
163     # Placeholder: bursty model logic
164     return None, 0.0
165
166 def sui_test_point_process(event_times, max_k):
167     """
168     Compare homogeneous Poisson null vs. bursty SUI-style model.
169     """
170     bic_null = fit_poisson_null(event_times)
171     breaks, bic_sui = fit_bursty_model(event_times, max_k)
172     delta_bic = bic_null - bic_sui
173     return {
174         "delta_bic": delta_bic,
175         "breaks": breaks
176     }

```

References

- [1] **Gallistel, C. R., Fairhurst, S., & Balsam, P.** (2004).
"The learning curve: Implications of a biological/psychological analysis."
Proceedings of the National Academy of Sciences, 101(36), 13124–13131.
- [2] **Lauretta, D. S., et al.** (2019).
"Episodes of particle ejection from the surface of the active asteroid (101955) Bennu."
Science, 366(6470), eaay3544.
- [3] **Carhart-Harris, R. L., & Friston, K. J.** (2019).
"REBUS and the anarchic brain: Toward a unified model of the brain action of psychedelics."
Pharmacological Reviews, 71(3), 316–344.
- [4] **Kuhn, T. S.** (1962).
The Structure of Scientific Revolutions.
University of Chicago Press.
- [5] **Poincaré, H.** (1890).
"Sur le problème des trois corps et les équations de la dynamique."
Acta Mathematica, 13, 1–270.