Persistent Anomalies and the Geometry of Ontology Revision

A visual guide to repair, admissibility, and non-abelian tears in machine learning and scientific discovery.

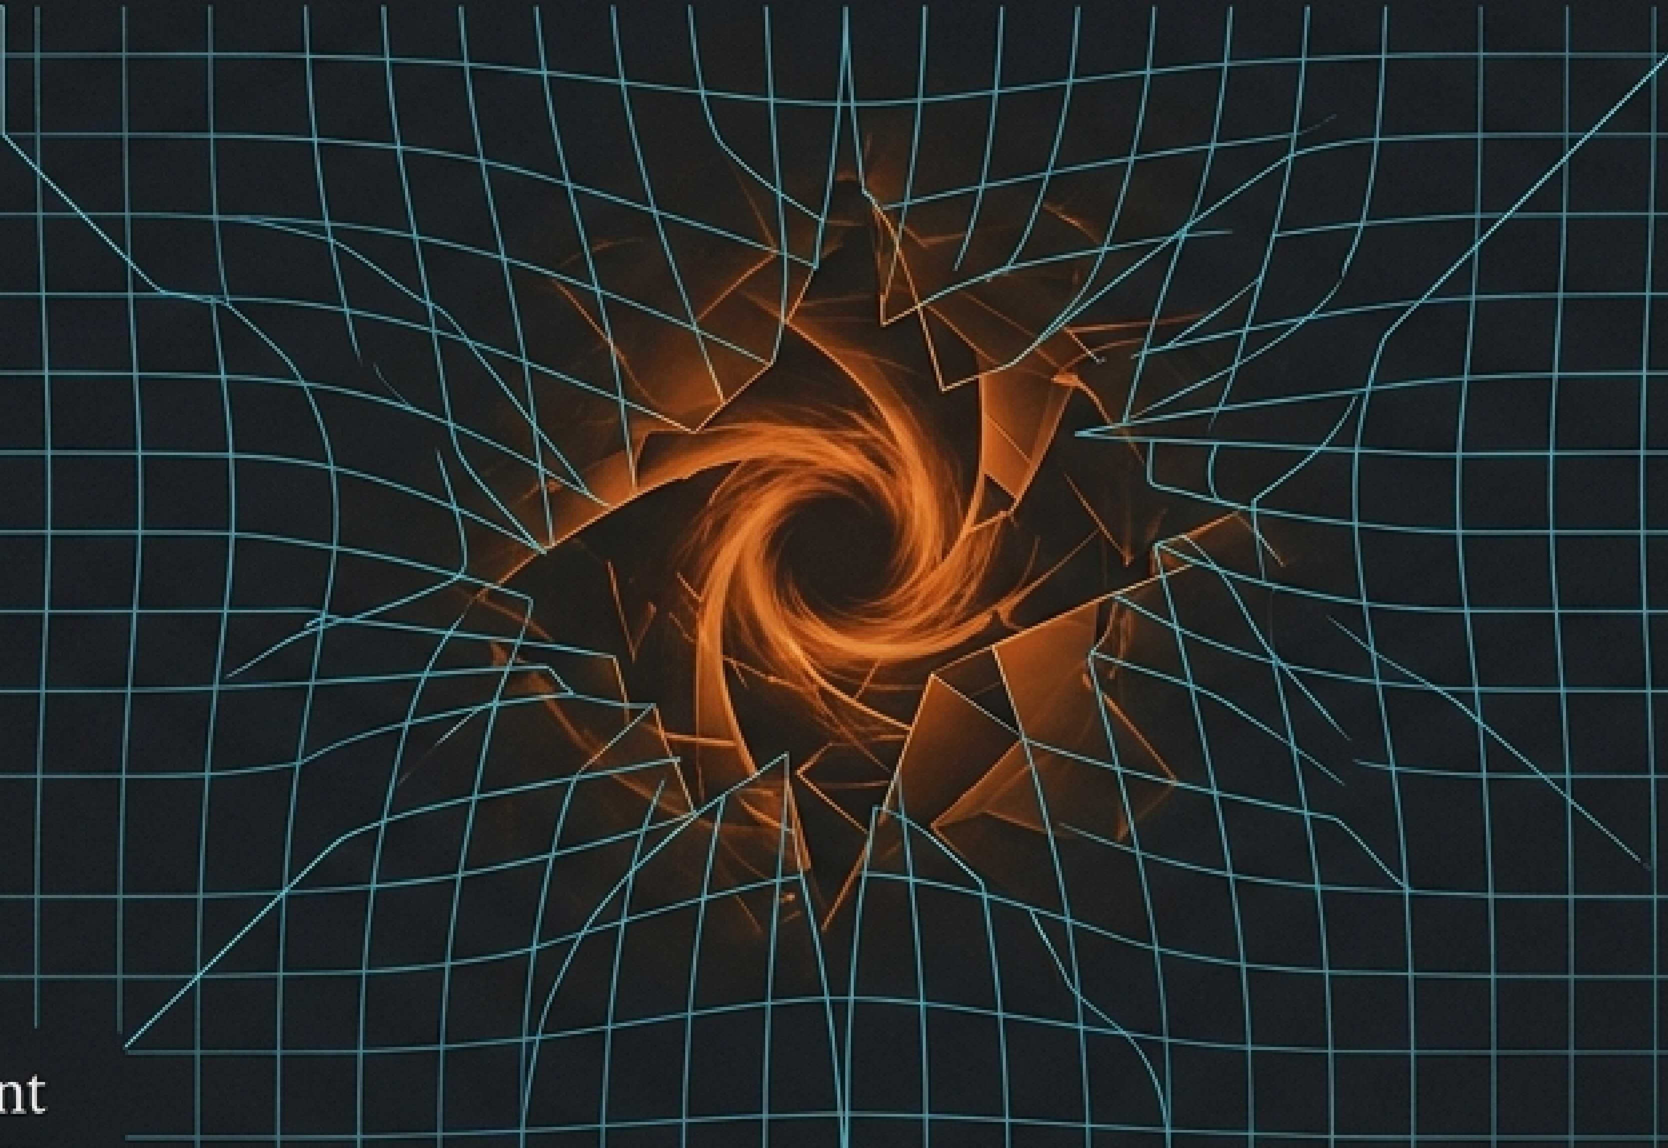

Based on the 2026 independent research monograph by Flyxion.

The Value of the Stubborn Residual

Most anomalies are merely errors waiting to be corrected. A rare few are signposts pointing to the boundary of our reality.

In the 19th century, Mercury's perihelion precessed at 43 arcseconds per century beyond Newtonian predictions.

This anomaly was not valuable because it was large. It was valuable because it was stubborn. It survived every available mathematical repair.

"It is not the magnitude of the discrepancy but its stubbornness that matters."

Parameter Error vs. Structural Error

The Bad Map

The observation lies within the model's capacity.
A simple parameter adjustment fixes the error.

The Wrong Geometry

The observation lies outside the model's image.
No parameter adjustment can reach it.

Takeaway: Persistent anomalies detect structural errors. They signal missing representational structure, not just inaccurate para

The Anatomy of Repair: The Promotion Ladder

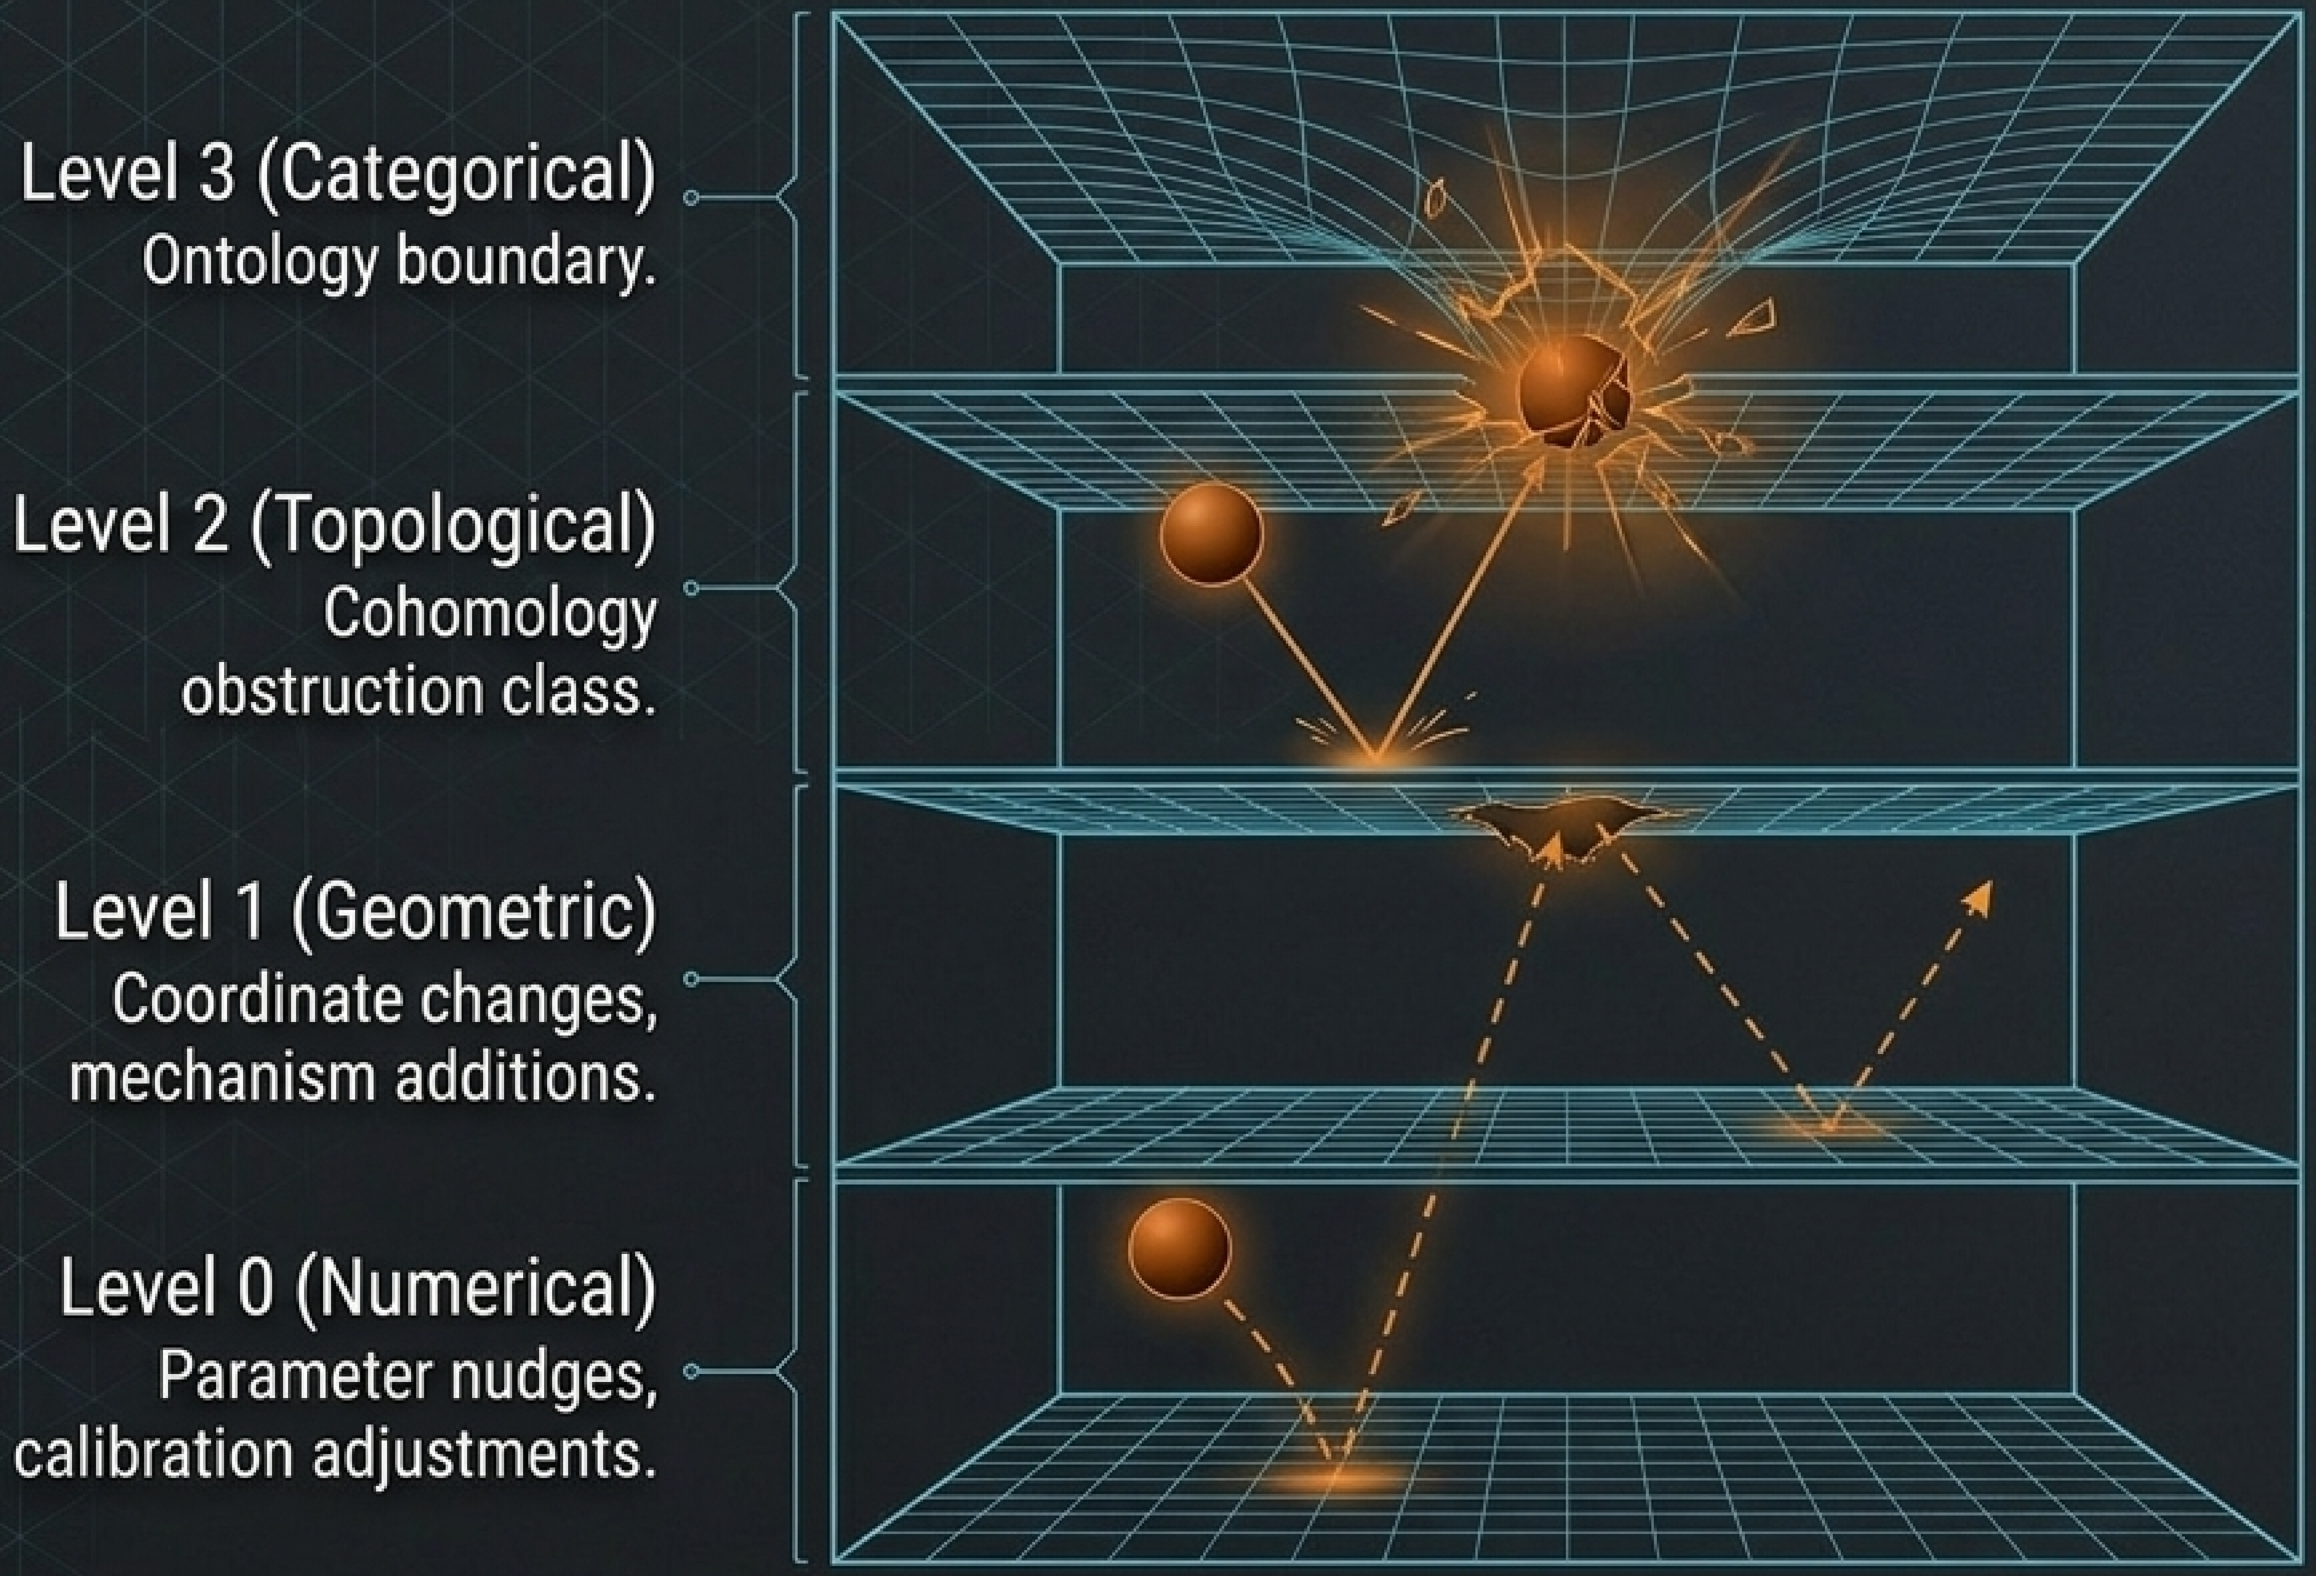

Learning is not primarily error reduction.
It is the continuous operation of a triage process that sorts anomalies by their structural depth.

The valuable ones are those that climb.

The Absolute Necessity of Memory

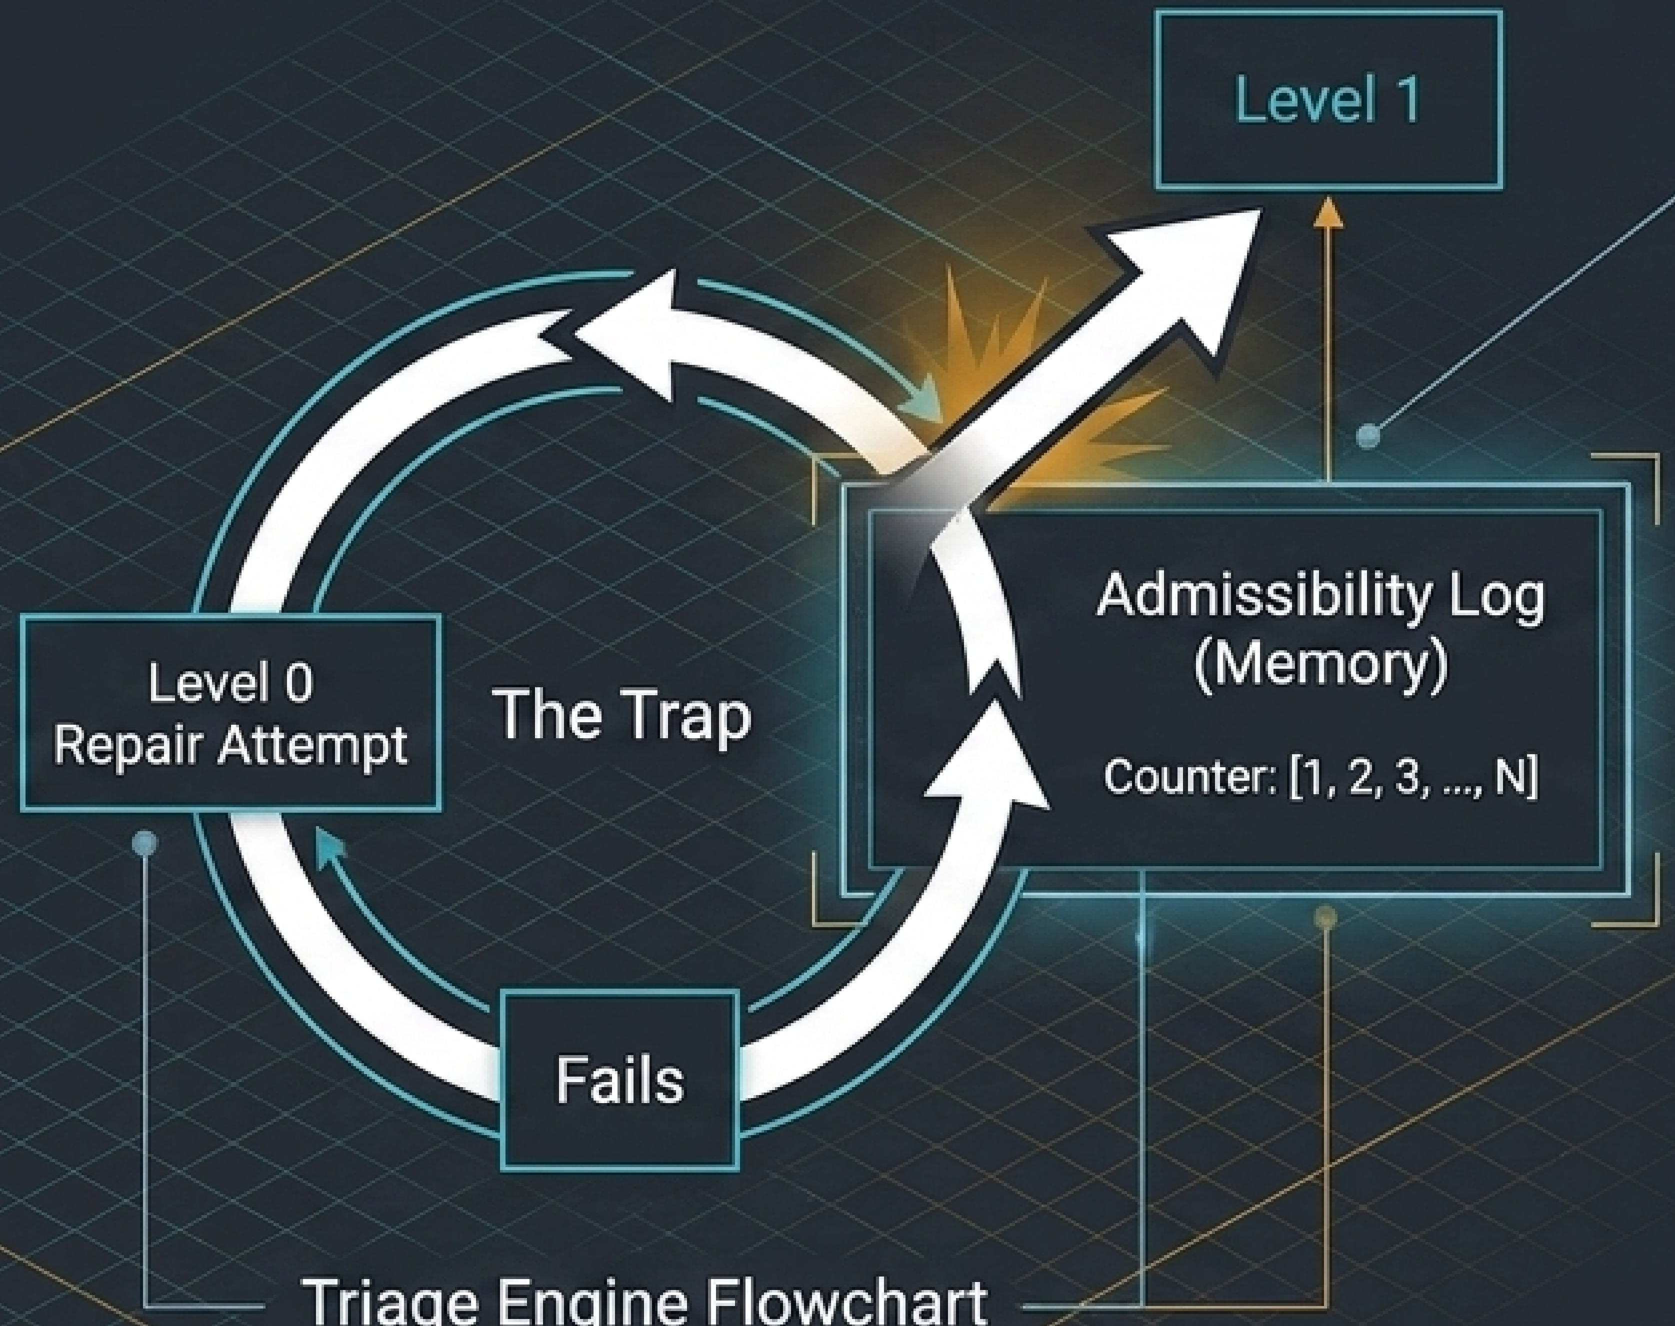

- Memory is constitutive of adaptive intelligence, not just passive storage.
- A memoryless learner treats every repair attempt as fresh, indefinitely repeating the same low-level fix.

$$\text{Memory} \Rightarrow \text{Persistence} \Rightarrow \text{Promotion} \Rightarrow \begin{array}{c}\text{Ontology}\\\text{Revision}\end{array}$$

The Hierarchy of Discrepancy Matrix

Axes I: Type of Discrepancy	Axes II: Persistence Character			
	Transient	Stable	Persistent	Invariant
Numerical	Calibration Error	Systematic Bias	Structural Mismatch	Missing Parameter
Geometric	Coord. Choice	Fit Residual	Manifold Incompatibility	Wrong Geometry
Topological	Local Obstruction	Boundary Effect	Genuine Tear	Ontological Gap
Categorical	Model Bug	Symmetry Break	Diagram Failure	**Ontology Boundary**

The bottom-right cells represent candidates for total Ontology Revision. Applying a numerical repair to a categorical discrepancy consumes resources while providing false evidence of progress.

The Topography of a Non-Abelian Tear

The Holonomy Diagnostic

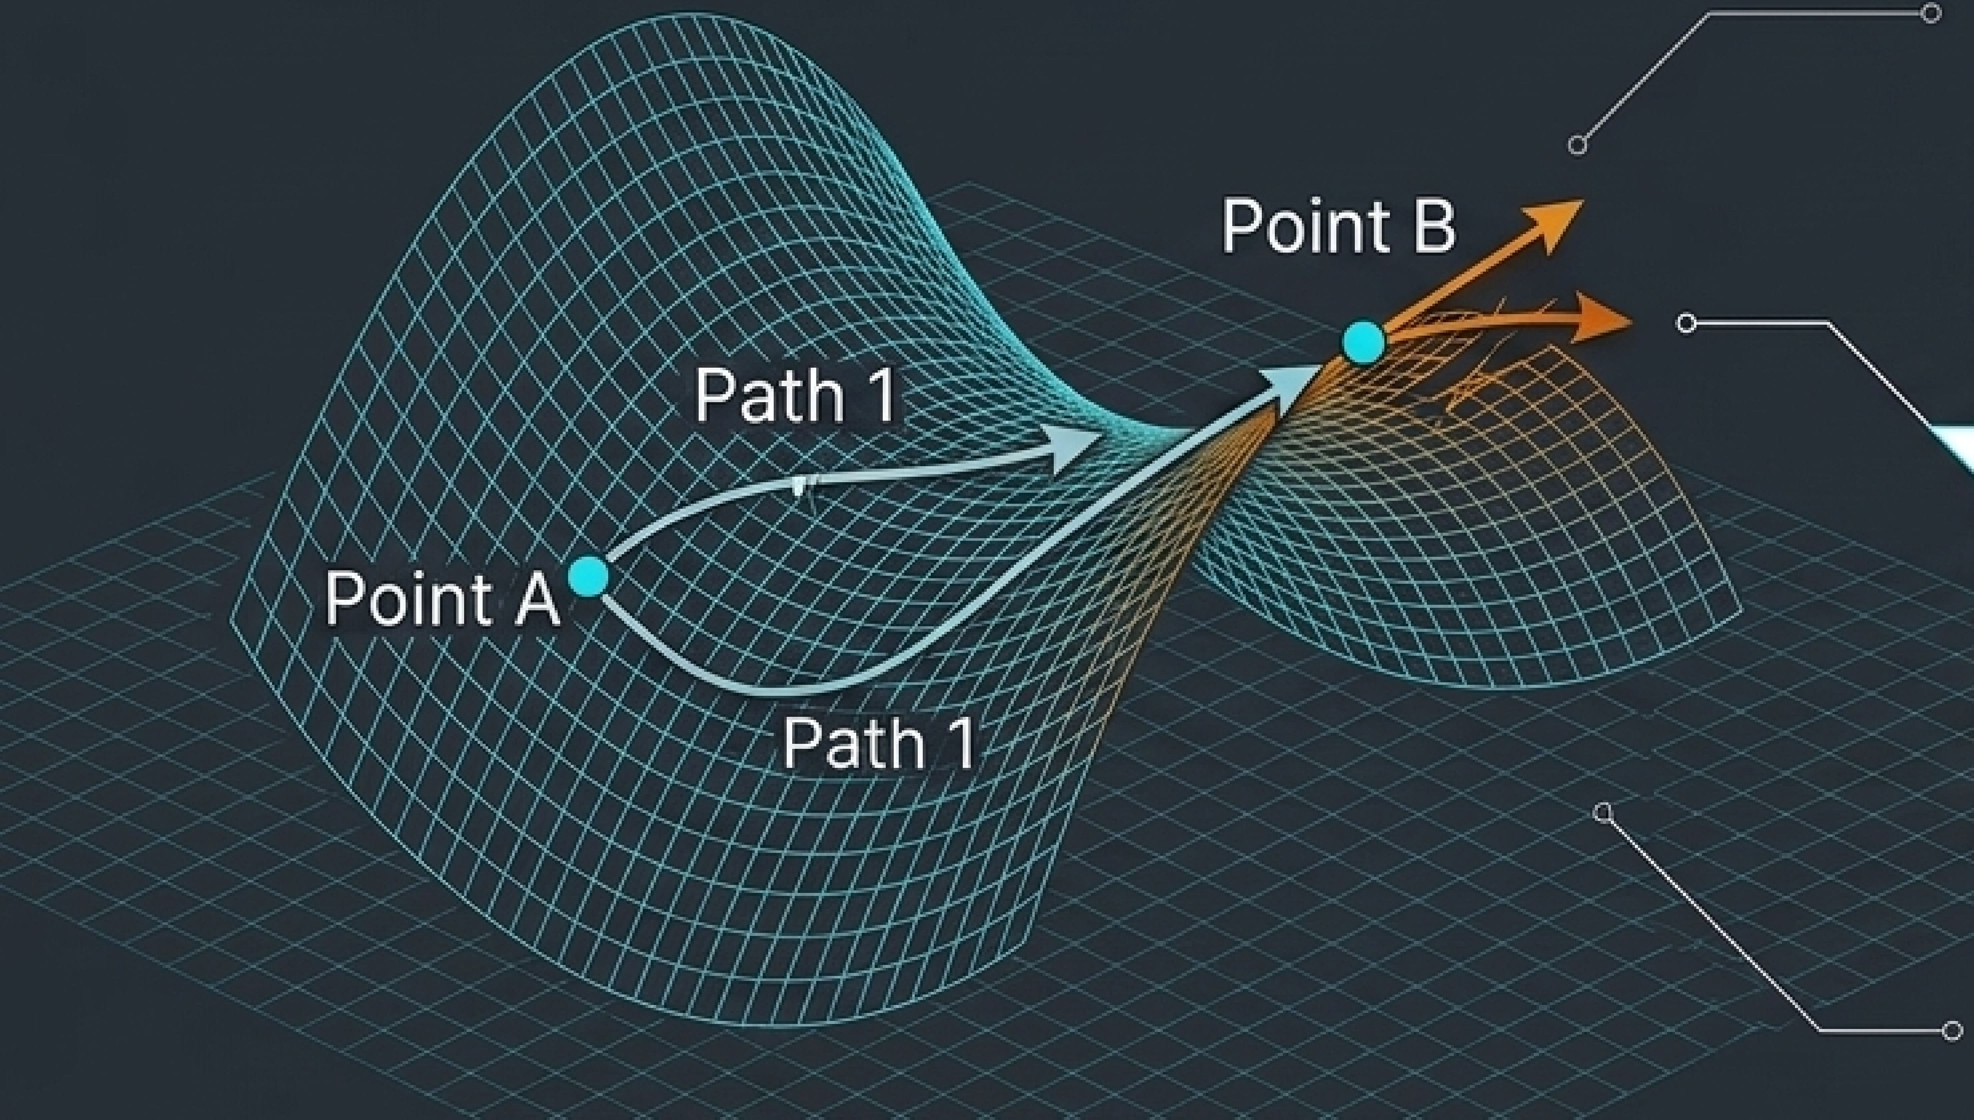

Divergent vectors at point B demonstrate path-dependence.

A non-abelian tear occurs when two trajectory histories cannot be smoothly deformed into one another. The order of operations matters.

Small Metric Distance, Large Topological Distance

The model sees a small metric error and assumes a local fix; the topology sees a completely different universe.

Commutator Test: $[R_a, R_b(\epsilon)] \neq 0$ across diverse repair families.

$$H(\gamma_1, \gamma_2) = T_{\gamma_1} T_{\gamma_2}^{-1} \neq \mathbf{I}.$$

The Convergence Gap: An Illusion of Success

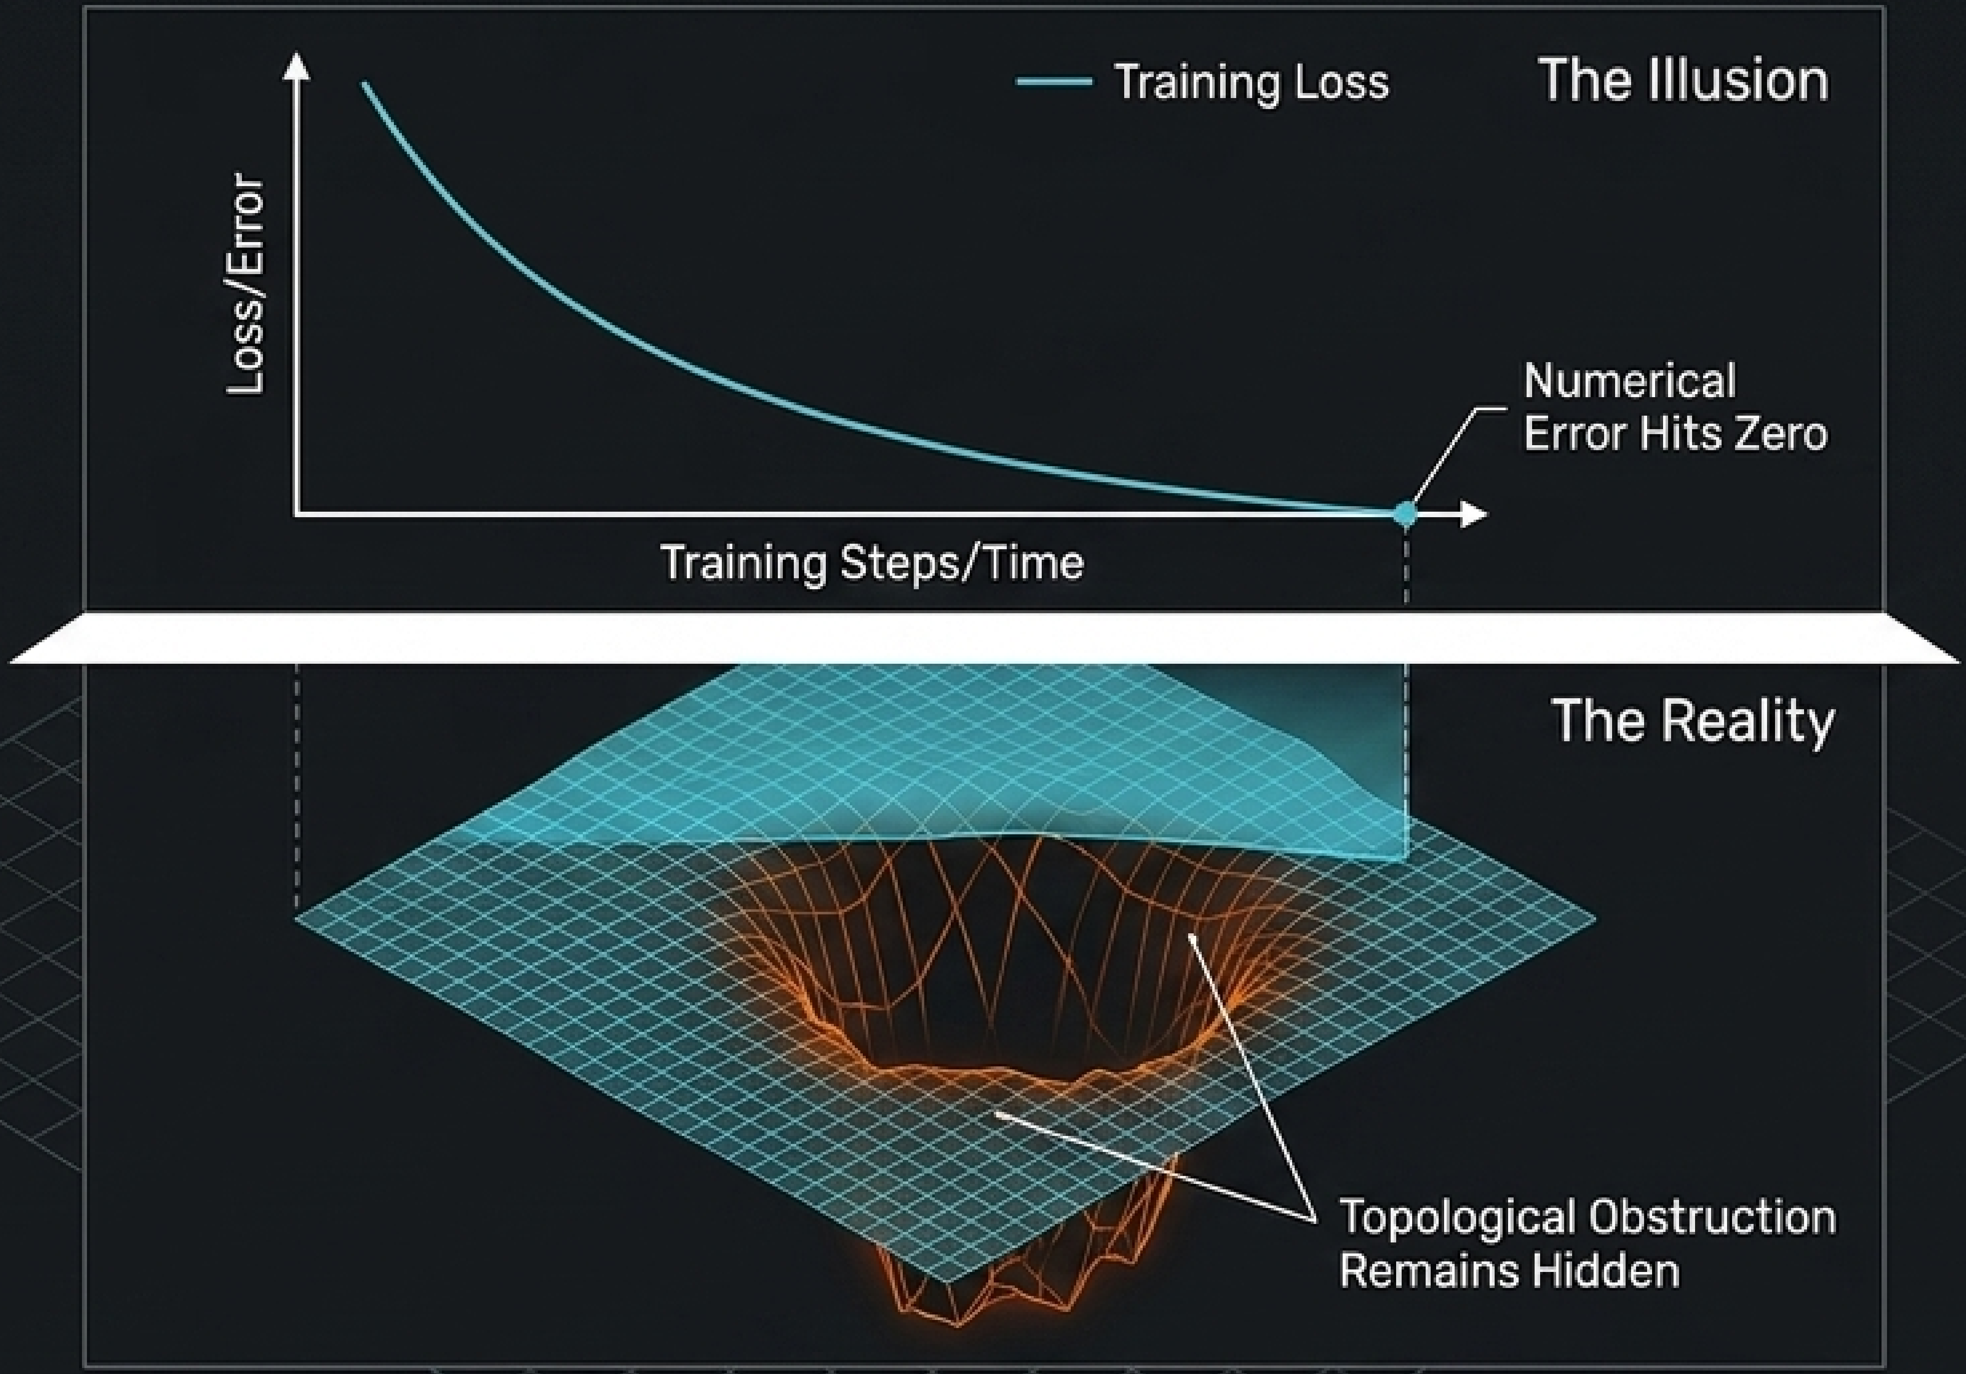

Convergence in residual space is strictly weaker than convergence in representational space. A repair sequence can drive numerical numerical error to zero while leaving the topological obstruction completely intact.

The tear has not been dissolved. It has been hidden.

Derivation: The Convergence Gap

Let $\varepsilon_n \to 0$ while $[\tau_n] = [\tau_0] \neq 0$. Then

$$\lim_{n\to\infty} \|\varepsilon_n\| = 0 \quad \text{does not imply} \quad \lim_{n\to\infty} [\tau_n] = 0.$$

The Dual Principle: Tear Suppression

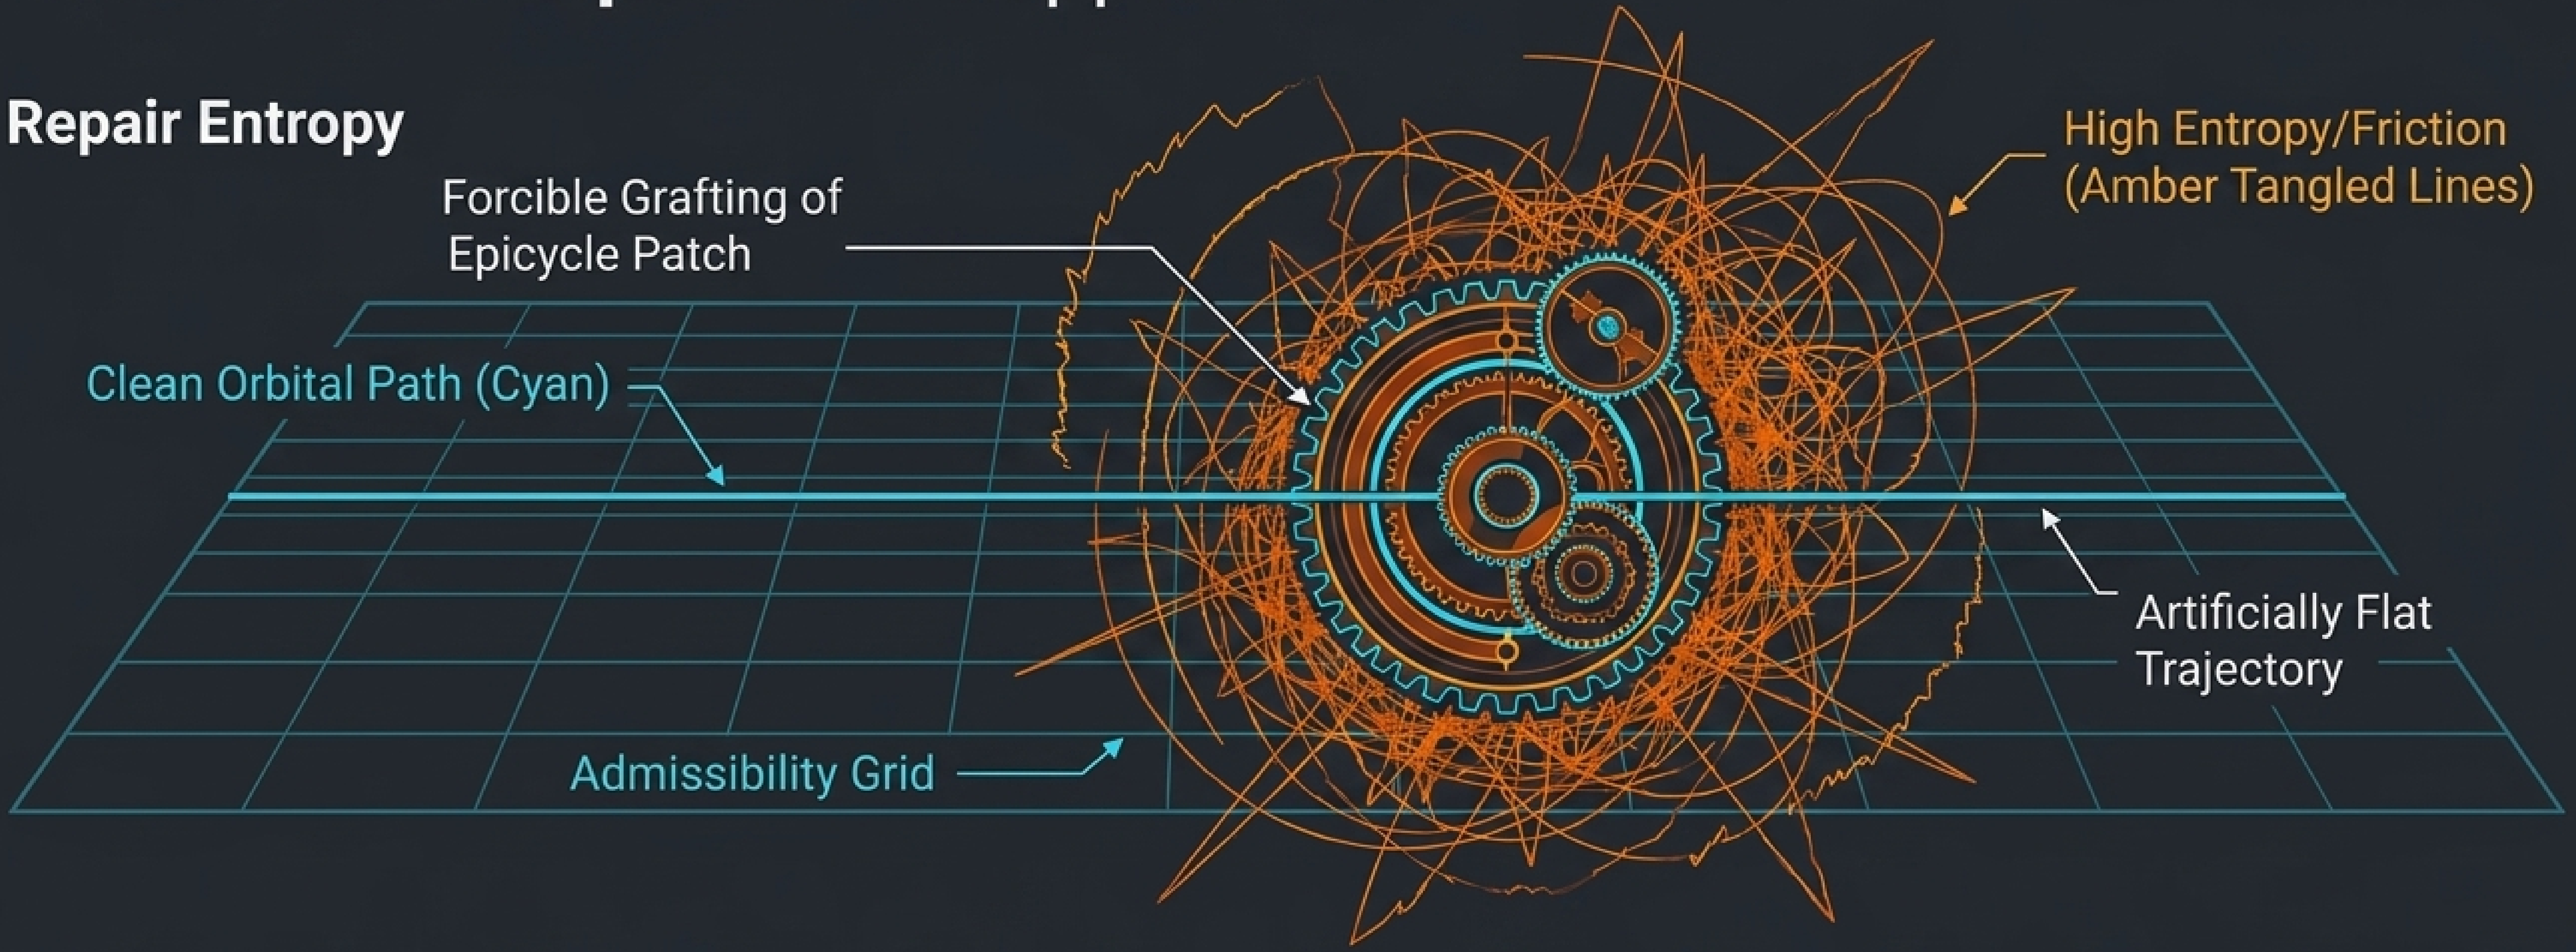

The Ptolemaic system was not a failure of fitting data; it was a success of tear suppression. Each epicycle was a perfectly valid Level-1 repair that reduced metric discrepancy but ignored the categorical necessity of a heliocentric model.

A system that is too flexible is scientifically blind. Expressive algebras absorb evidence of structural discrepancies.

The Modern Epicycle: Reinterpreting Overfitting

The Standard Account	The Persistence Account
Engine: Gradient Descent	Engine: Repair Triage
State: Memoryless (evaluates instantaneous scalar loss)	State: Memory-Bearing (evaluates Admissibility Logs)
Goal: Error Minimization	Goal: Representational Growth
Vulnerability: Tear Suppression / Classical Overfitting	Strength: Tear Dissolution / Topological Fidelity

Modern foundation models are the ultimate Ptolemaic epicycles. Their massive parameter counts constitute an infinitely flexible Level-0 repair algebra, perfectly capable of hiding Level-3 structural tears by driving numerical loss to zero.

The Ptolemaic Principle in Machine Learning

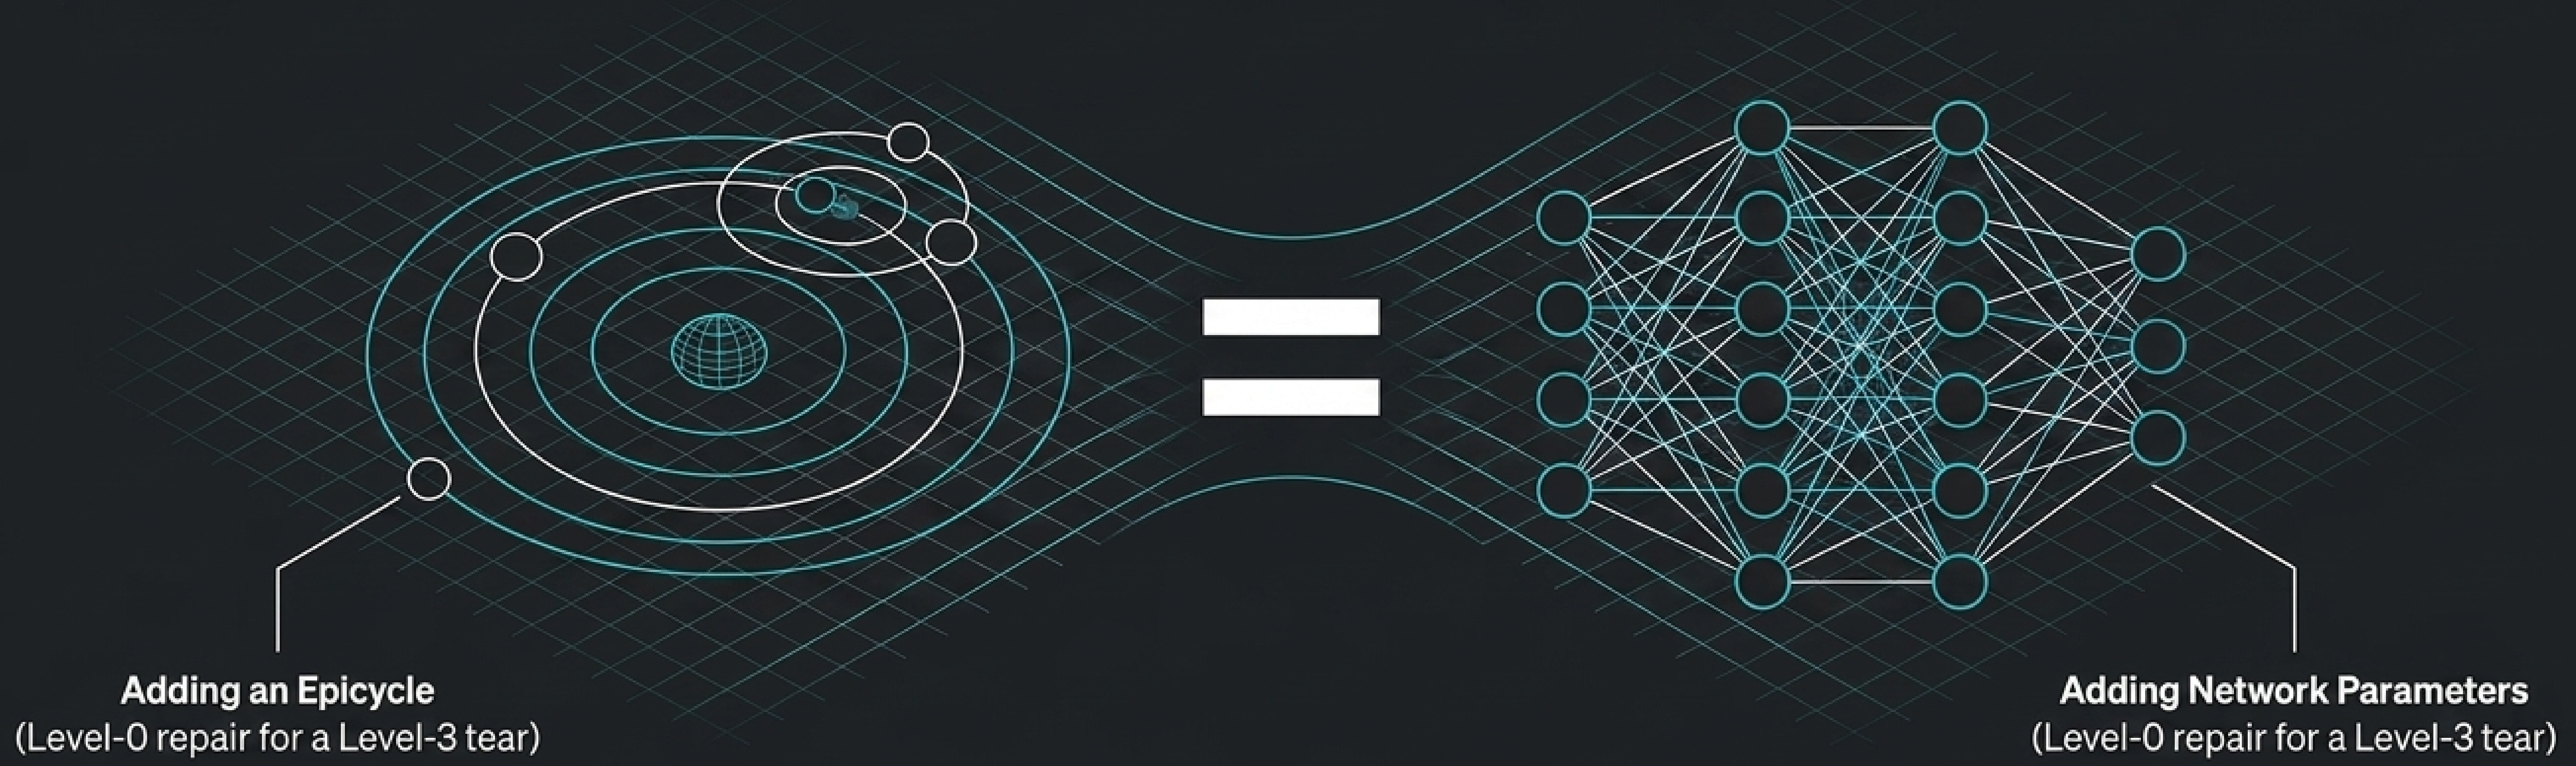

If the dimension of the repair algebra exceeds the topological complexity of the true trajectory manifold, an overfit model will drive error to zero without resolving the obstruction.

Convergence is not evidence of ontological adequacy.

Ontology Revision: The Dissolving Tear

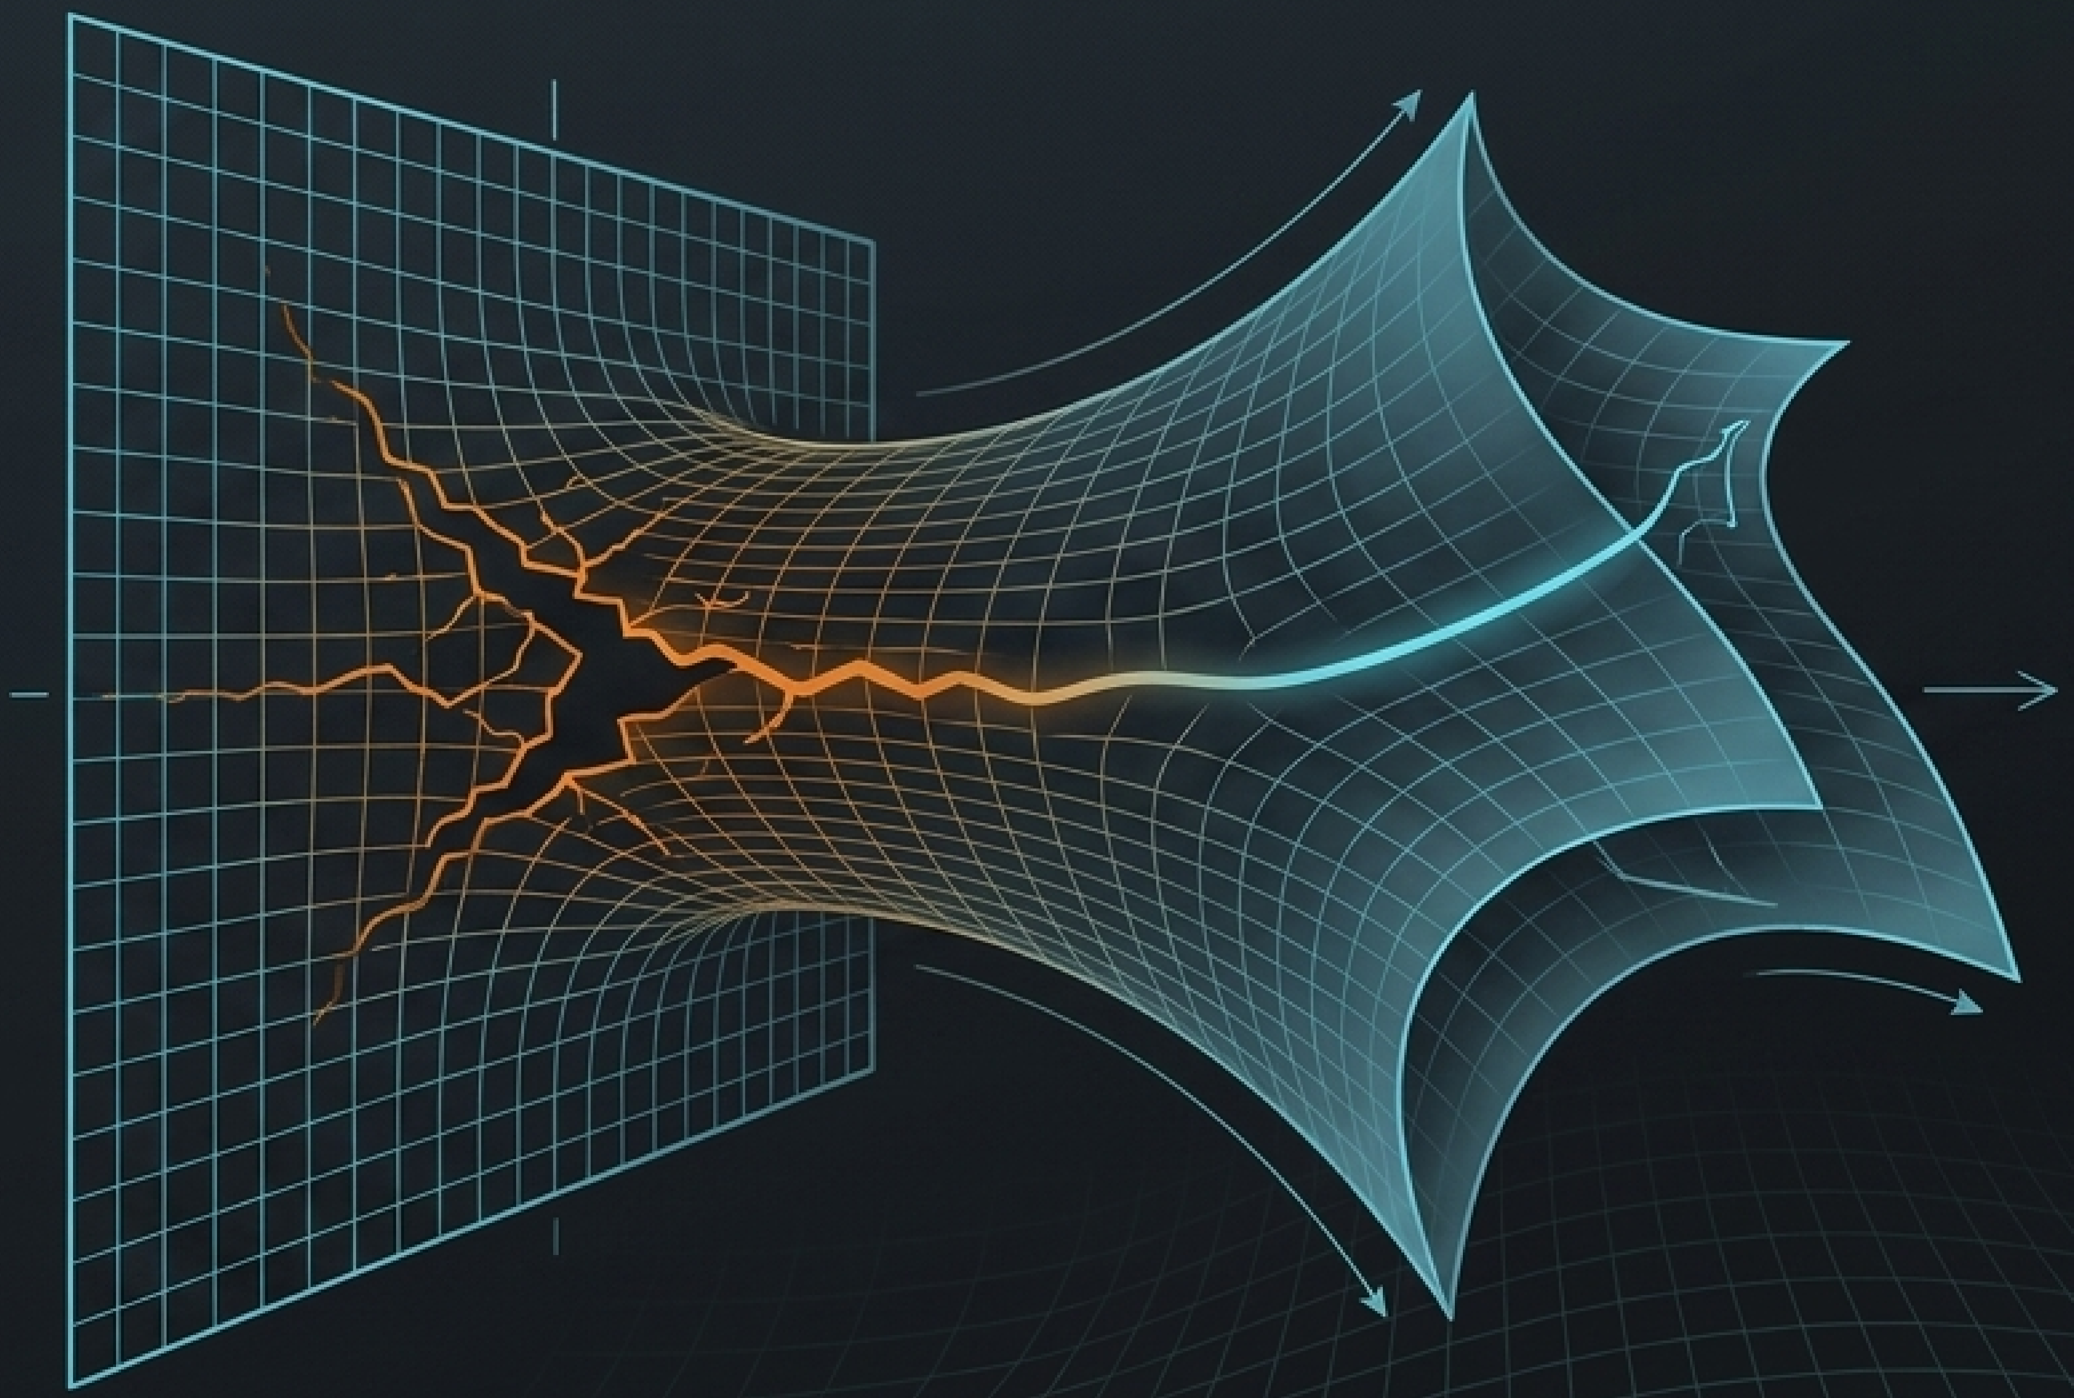

The goal of learning is not error minimization, but representational growth.

- When the coordinate chart is exhausted, parameters cannot save you.
- A persistent anomaly acts as a generator for a new geometry.
- The new theory doesn't "fix" the anomaly—it expands the admissibility manifold until the tear naturally dissolves.

Mercury Resolved: From Patch to Geodesic

Suppression

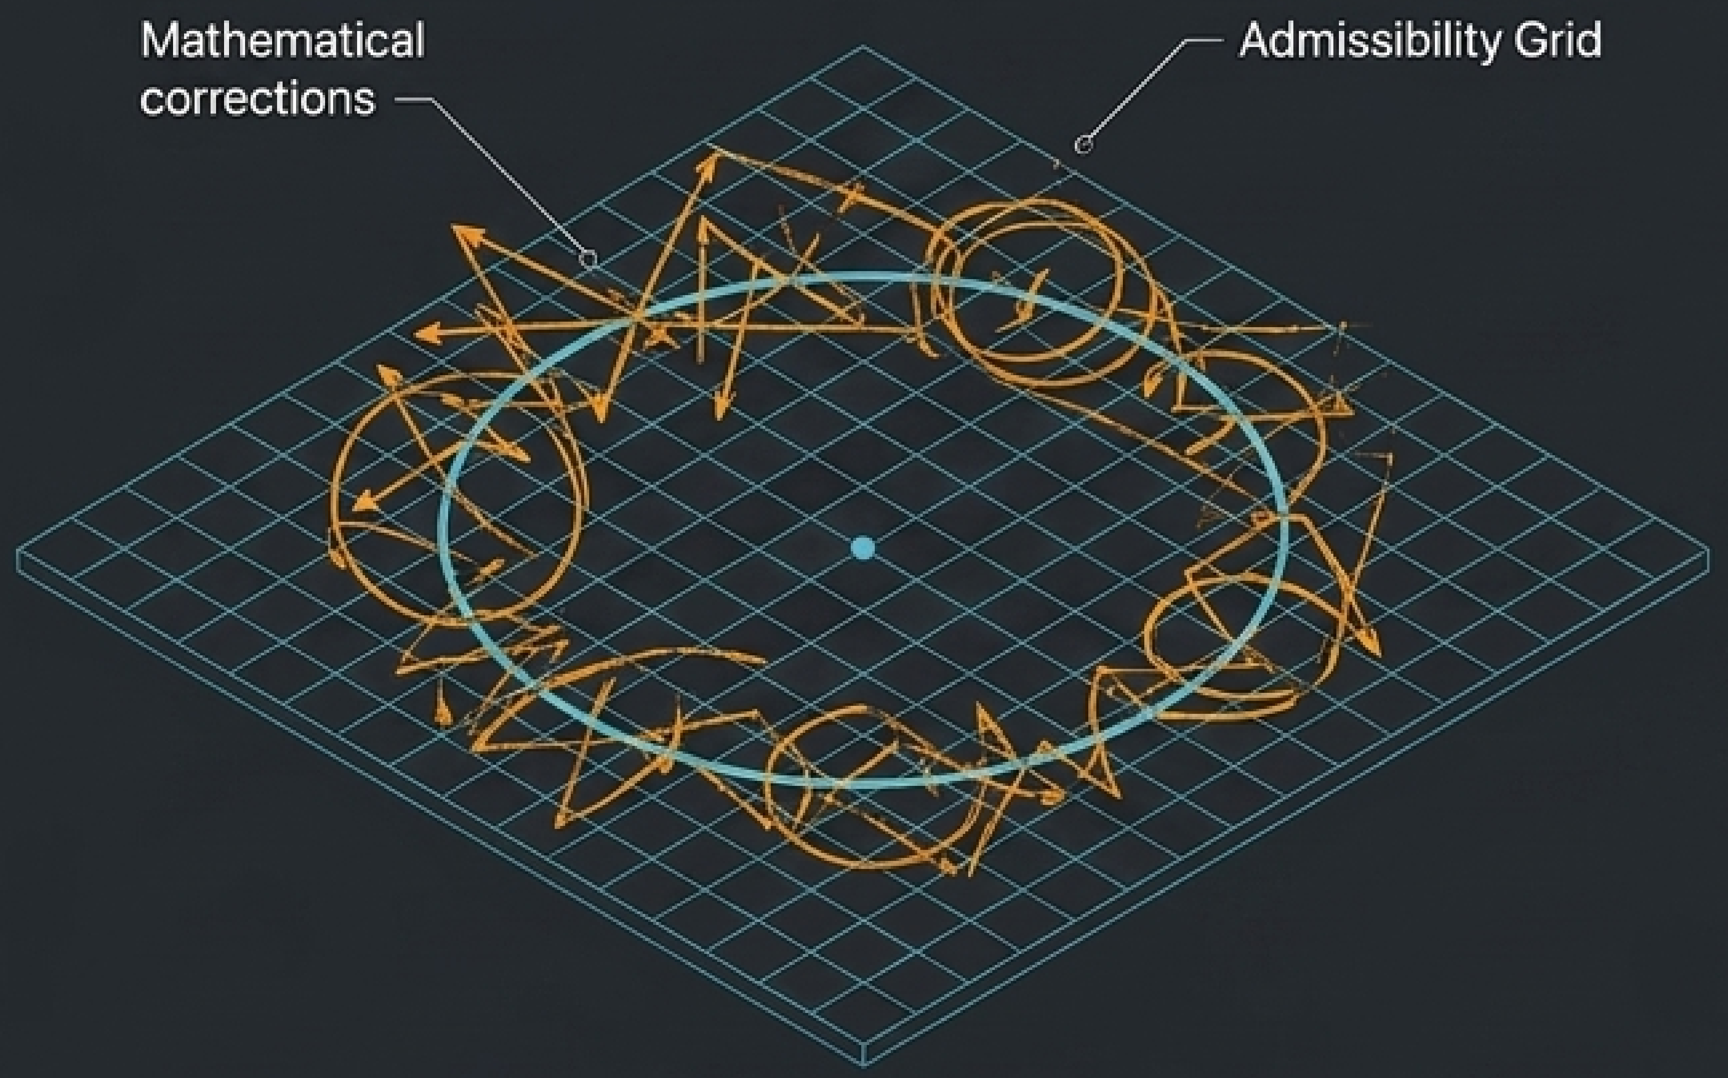

High Repair Entropy / Newtonian Framework

Dissolution

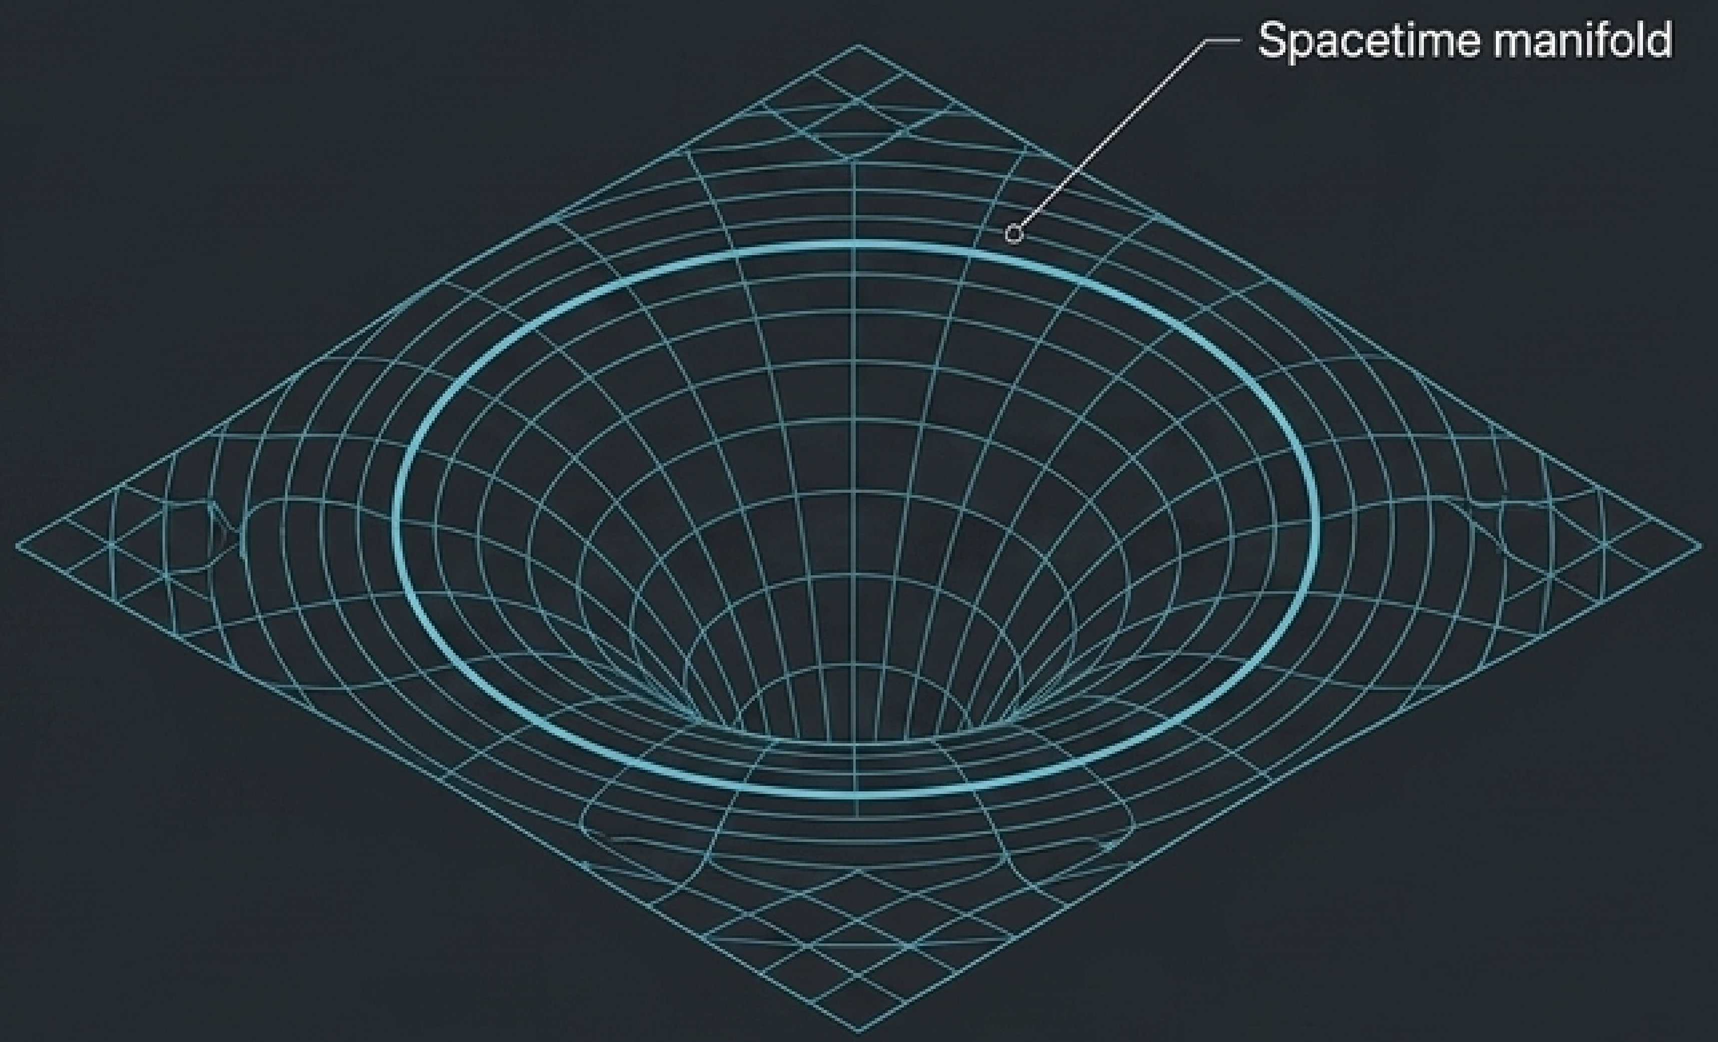

Zero Obstruction / Riemannian Framework

General Relativity did not predict Mercury's precession by adding a correction term. By revising the ontology to curved spacetime, the anomaly became the expected, natural trajectory. The tear dissolved because the new coordinate system had room for it.

Towards Persistence-Driven Learning

Persistence Monitors

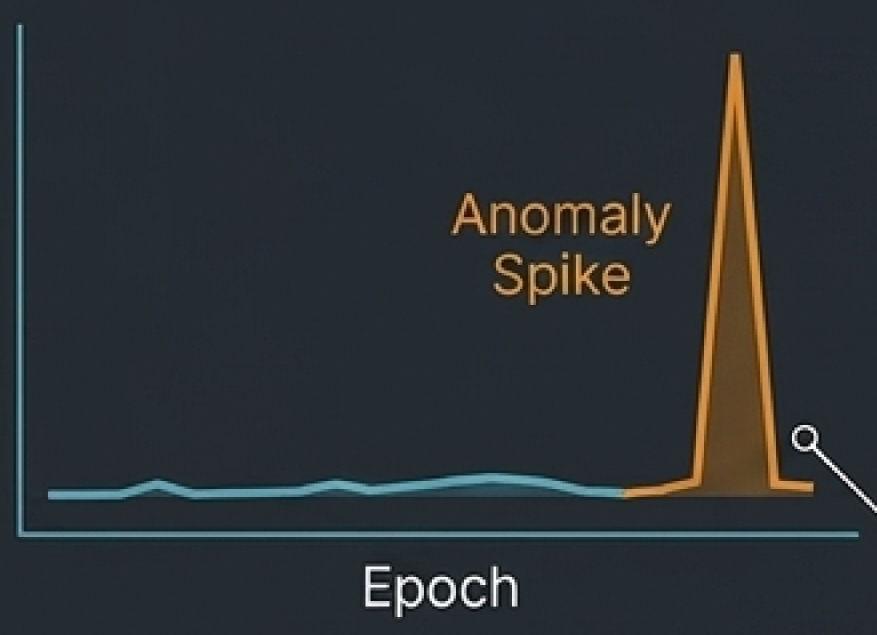

Track the **repair history of training anomalies** across epochs. A sudden spike in persistent small anomalies is an early warning of impending representation failure.

Holonomy-Based Augmentation

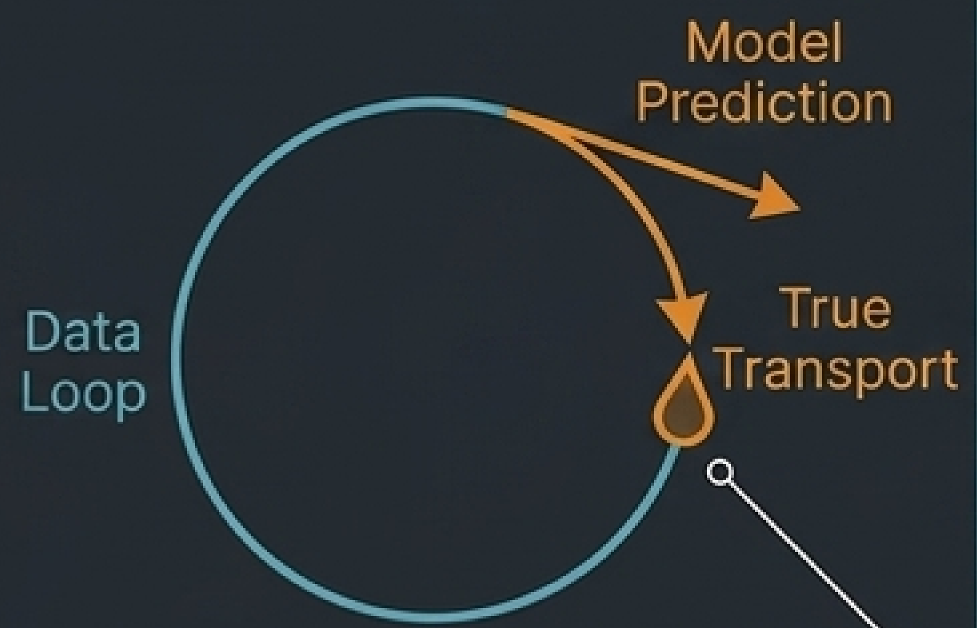

Generate training examples by **transporting inputs around loops** in data space. If the model fails the holonomy test, target augmentation at the tear.

Topology-Regularized Training

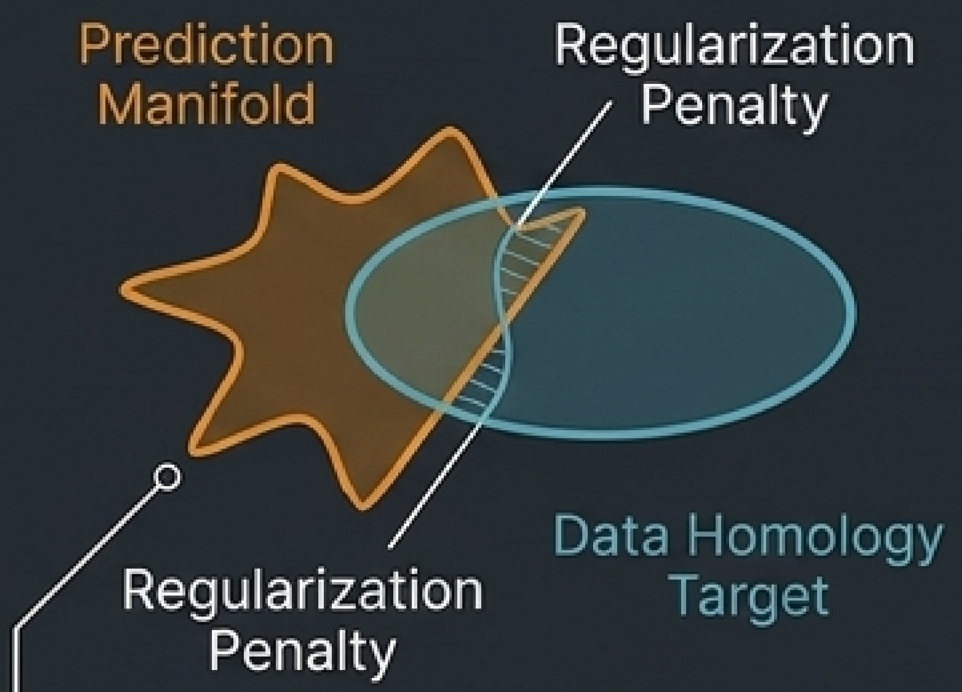

Add a regularization term **penalizing discrepancies** between the model's prediction manifold and the **true data homology**, directly discouraging tear suppression.

The Persistence Principle

Persistent anomaly $\Rightarrow$ missing structure.

A persistent anomaly is not a fact about the world.
It is a fact about the geometry of the model
used to represent the world.

The repair sequence is the experiment. The Admissibility Log is the record. The anomaly is the instrument by which the model discovers the shape of its own boundary.