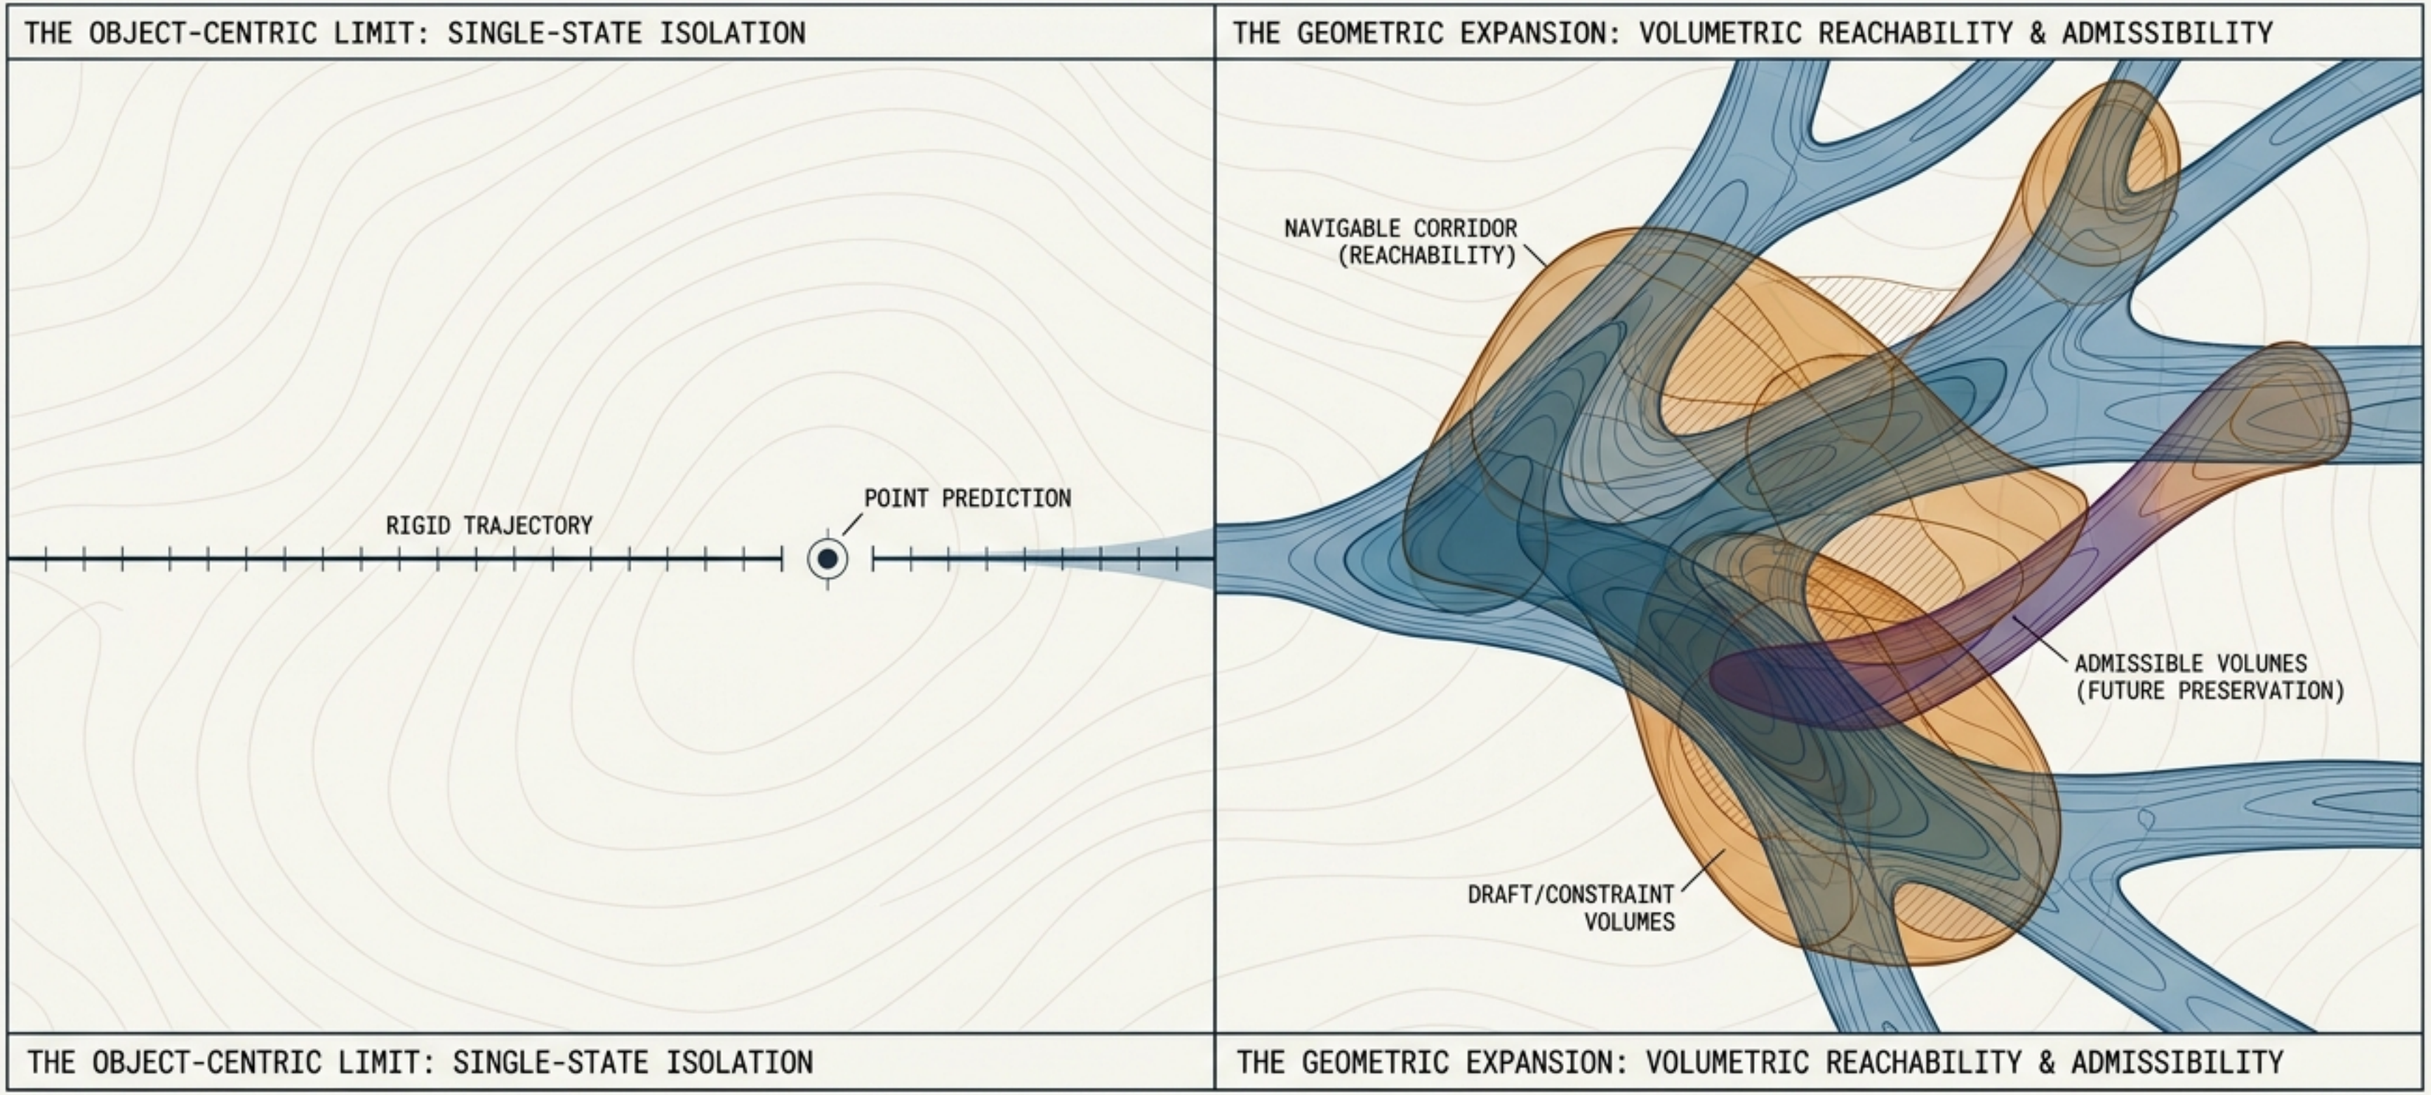

BEYOND PREDICTION: THE GEOMETRY OF FUTURE PRESERVATION

REACHABILITY, ADMISSIBILITY, AND THE LIMITS OF OBJECT-CENTRIC AI

Intelligence is Navigation, Not Search

	Search (Object-Centric)	**Navigation** (Process-Centric)
		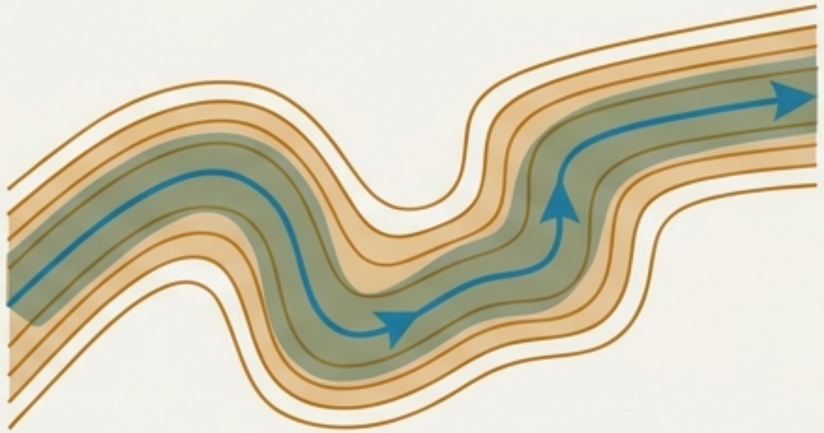
	$\mathbf{x}^* = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{argmax}} f(\mathbf{x})$	$\mathbf{x}(t+1) = \mathbf{x}(t) + \mathbf{v}(\mathbf{x}, t), \gamma(t) \subset \mathcal{A}(\mathbf{x}(t))$
Memory	Retrieval of a stored item	Reconstruction under constraint
Identity (Ship of Theseus)	Preservation of same physical parts	Preservation of same admissible future set ($\mathcal{A}(\mathbf{x})$)
Software	Tracking literal file differences	Preserving computational reachability
AI Objective	Finding the unique optimal state	Staying inside the navigable corridor

The Empirical Anomaly in Speculative Decoding

The Assumption

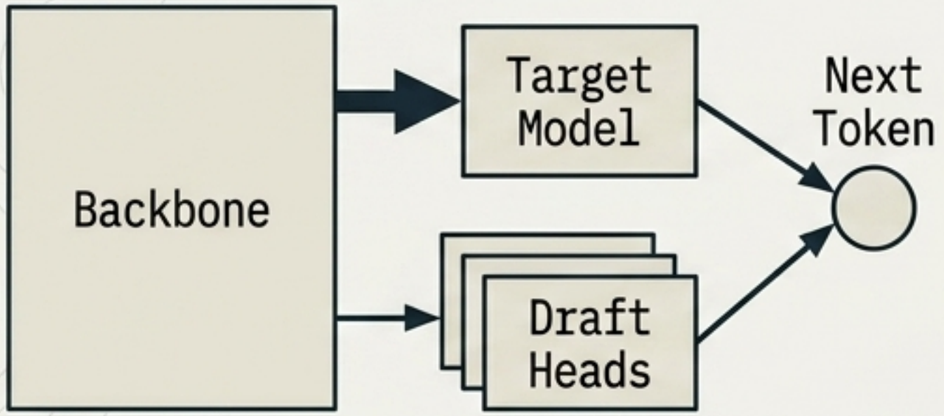

The standard assumption: Acceptance drops during RL because draft heads lose predictive accuracy and become stale.

The Reality (The Bebop Result)

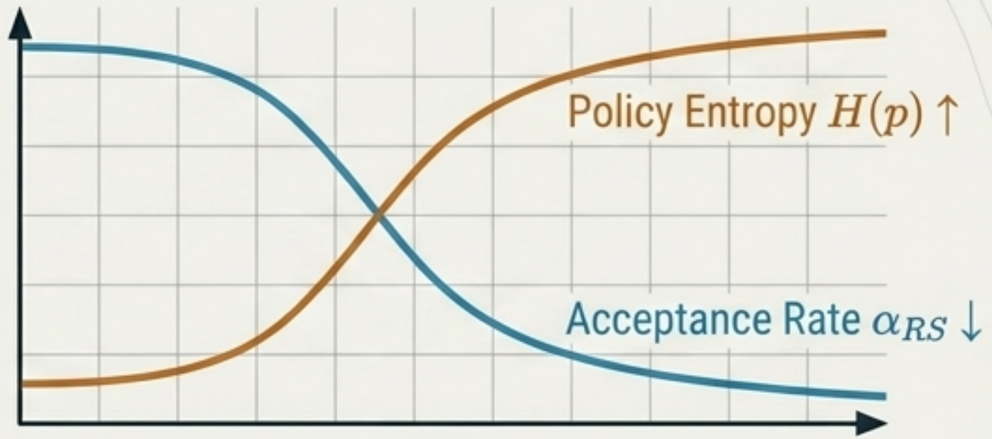

$$\Delta\alpha_{mismatch} \approx 0$$

The Bebop Result proves the acceptance drop is driven almost entirely by policy entropy, not model staleness. The backbone moves substantially during RL, but the admissible manifold's topology remains stable. The relevant metric is not predictive accuracy, but shared probability mass.

Geometric Overlap Governs System Acceptance

CE/KL Loss (Pointwise)

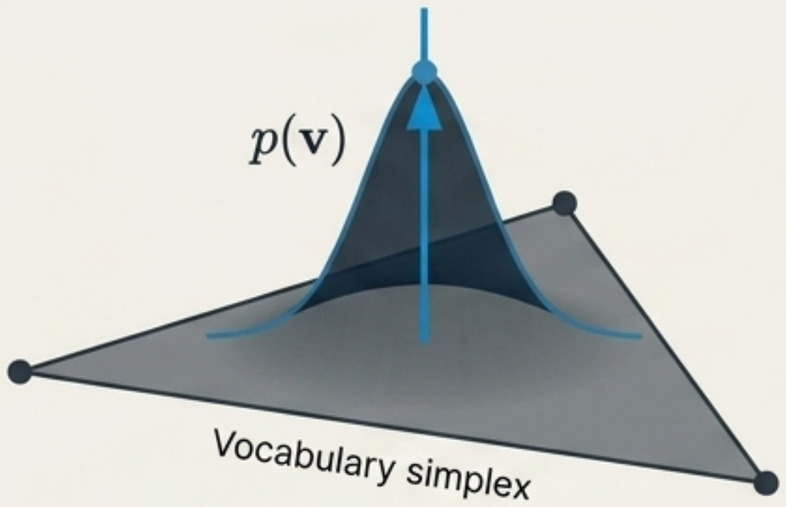

Optimizing for a single point

TV Loss (Geometric)

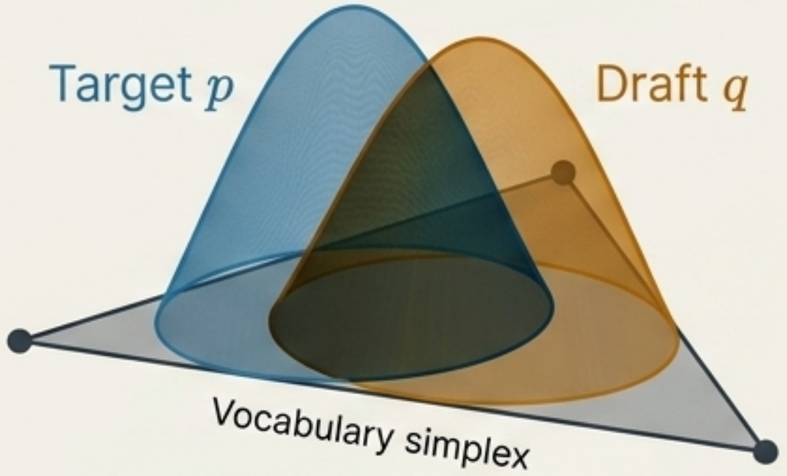

Optimizing for shared region

The Rejection Sampling Math:

$$\alpha_{RS} = \sum_{v} \min(p(v), q(v)) = 1 - d_{TV}(p, q)$$

The Finding: Optimizing Total Variation (TV) overlap directly outperforms optimizing CE or KL. The draft learns to approximate the shape of the **admissible region** rather than the location of its mode. Access to a shared region matters more than identification of a shared point.

The Five Levels of Future Orientation

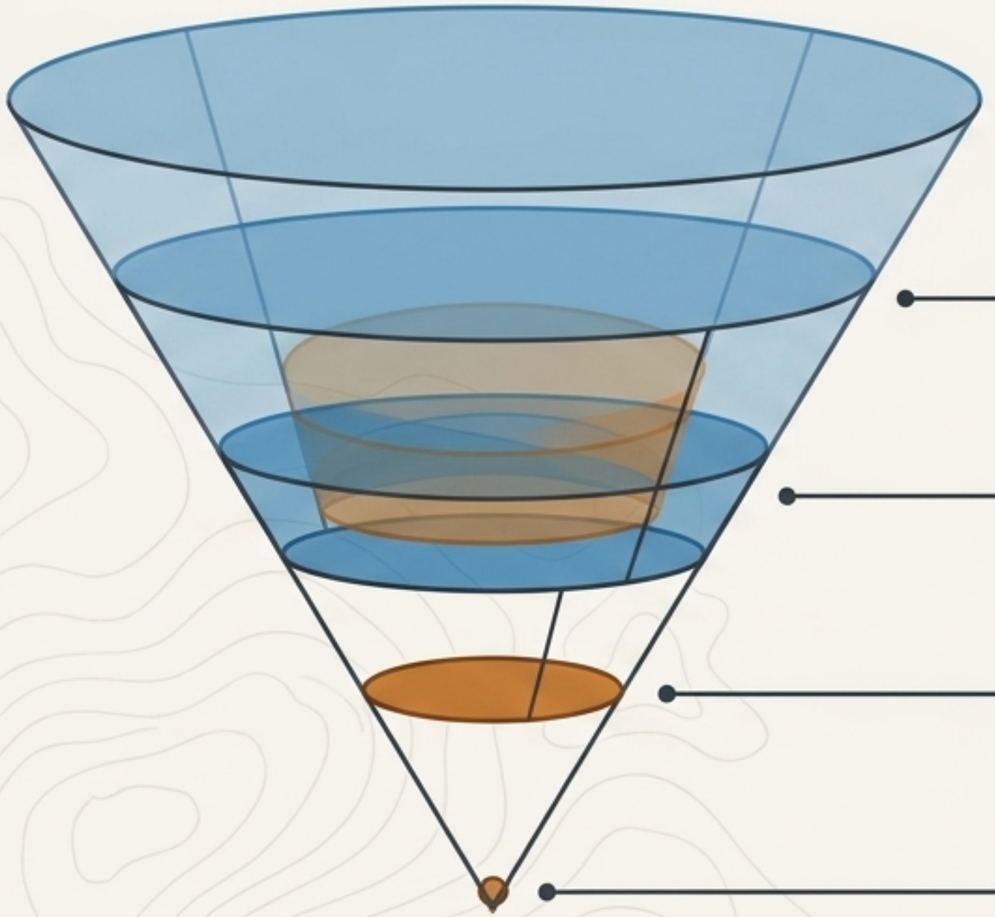

Level 4: Ontological Reachability
Which description spaces are reachable?
Metric: $\Omega(T_1, T_2)$

Level 3: Admissibility Overlap
Which coherent futures remain?
Metric: $\mathcal{O}(p, q) = \text{Vol}(\mathcal{A}_p \cap \mathcal{A}_q) \ / \ \text{Vol}(\mathcal{A}_p)$

Level 2: Reachable Trajectories
What futures overlap?
$\text{Vol}(R_p(T) \cap R_q(T))$

Level 1: Distributional Overlap
What mass is shared?
$1 - d_{\text{TV}}(p, q)$

Level 0: Token Prediction
What is y_t?
$p(y_t \mid y_{<t})$

Proposition 6.1: TV overlap is not a metaphor. It is trajectory overlap restricted to horizon 1, projected onto a flat vocabulary simplex. Each level subsumes the previous. Prediction is exact only in degenerate, deterministic settings.

Entropy is a Geometrically Blind Proxy

$$H(p_1) = H(p_2) = \log 100$$

Isotropic Diffusion

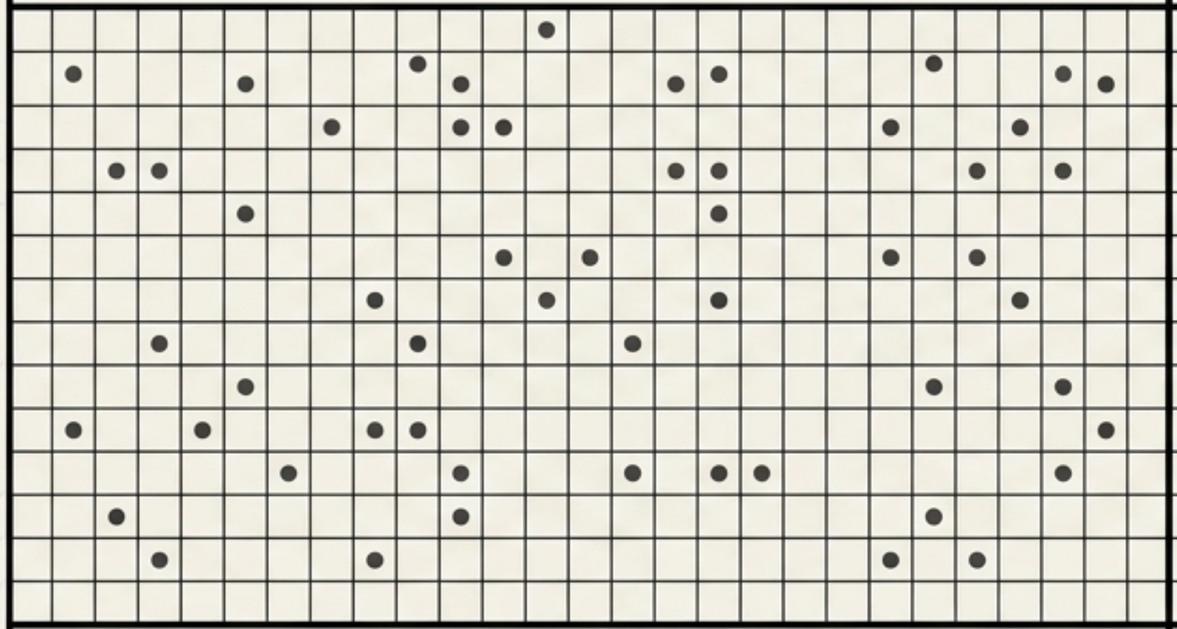

Entropy measures uncertainty—how spread out probability mass is across a space.

Admissible Volume $\mathcal{V}_{\mathcal{A}}(p)$ measures possibility—how much of the future remains coherently reachable.

Anisotropic Diffusion

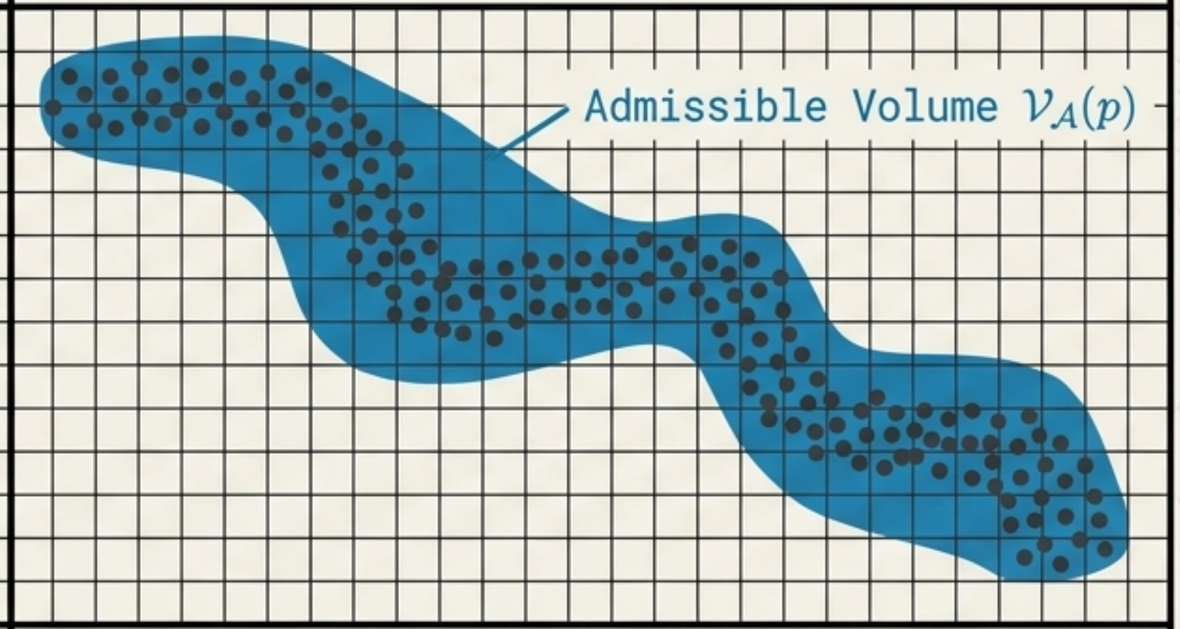

Support vs Admissibility: Support measures probabilistically reachable states. Admissibility measures coherent possibility. $\mathcal{A}(p) \subseteq \text{supp}(p)$.

Standard entropy cannot distinguish a random scatter from a navigable semantic corridor.

Speculative Decoding as Field Transport

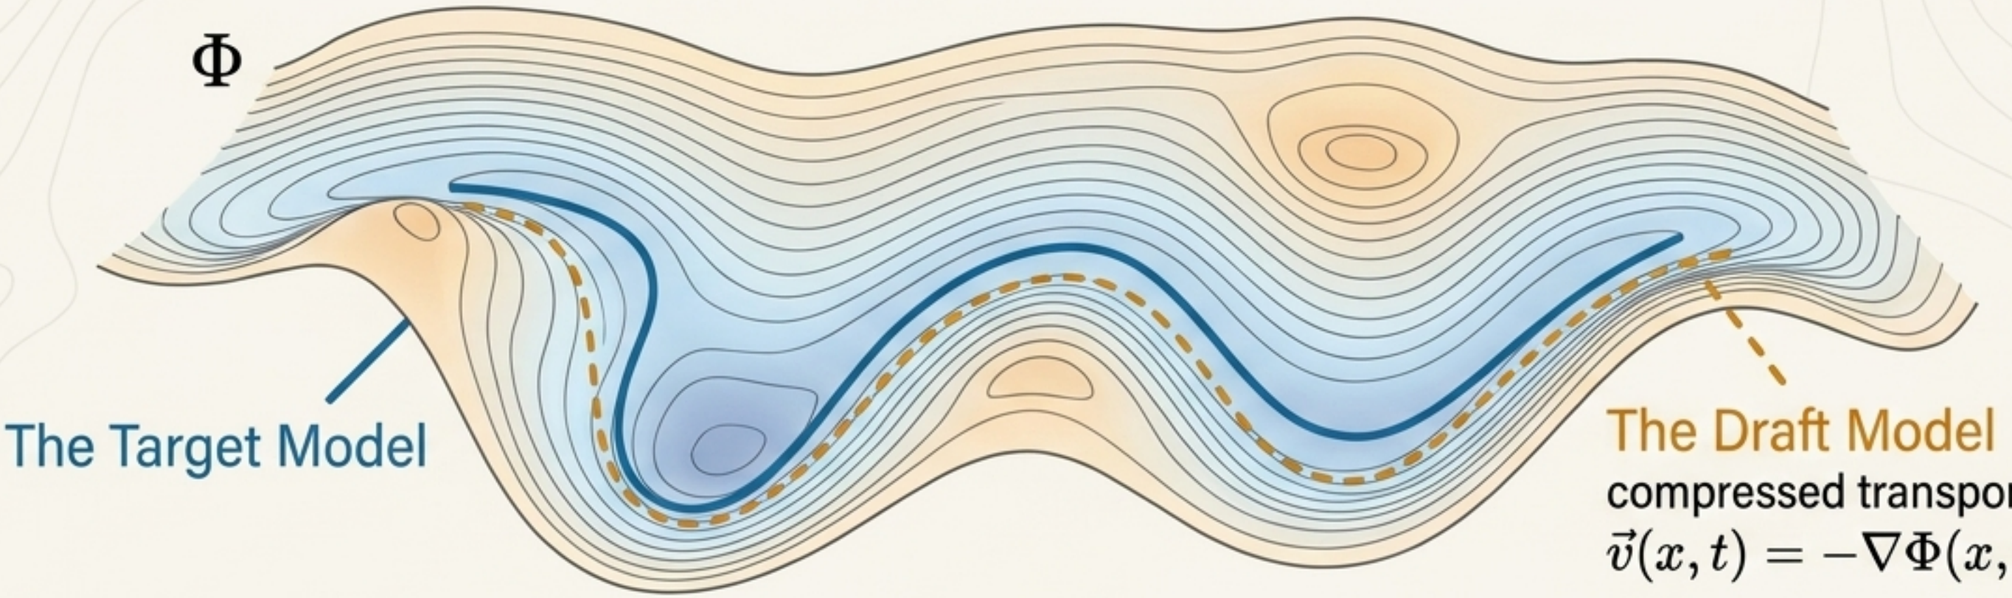

CLIO Projection

Draft heads are projection operators $\pi_k: \mathcal{H} \to \Delta_{|V|}$. When policy entropy rises, the admissible region expands. The projection must adaptively widen to preserve access to the future without needing to evaluate the full constraint field.

RSVP Metric Deformation

Reinforcement Learning (RL) acts primarily as a metric deformation ($g_t \neq g_{t+\Delta t}$), reweighting trajectories rather than changing fundamental paths. It preserves the topology of the admissible basin ($\mathcal{A}_t \cong \mathcal{A}_{t+\Delta t}$), which is why draft-model mismatch remains negligible.

The Industrial Convergence on Reachability

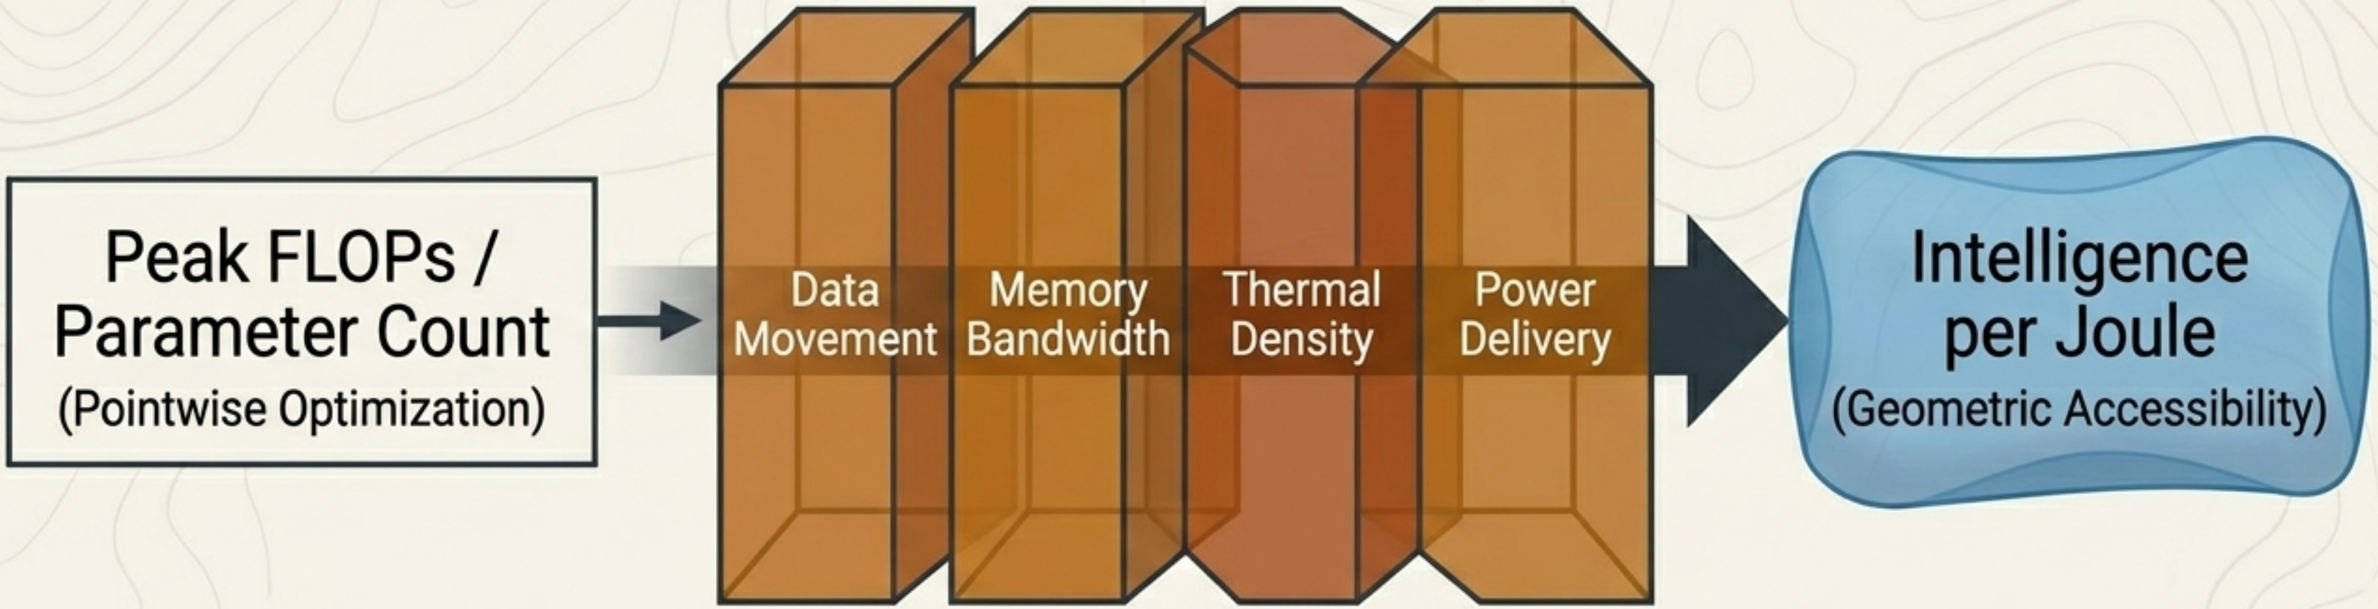

The AI+HW 2035 roadmap implicitly adopts the **Reachability Principle**. The target is not improving predictive accuracy, but preserving access to capabilities under severe constraints.

Proposition 4.1: If System A uses 100x less energy than System B for identical tasks, its admissible future set strictly contains B's:

$$\mathcal{A}(A) \supset \mathcal{A}(B)$$

Efficiency is fundamentally an admissibility-restoration program.

The Rematching Archive Inverts the Thin Pipe

Shannon's Pipe (Classical)

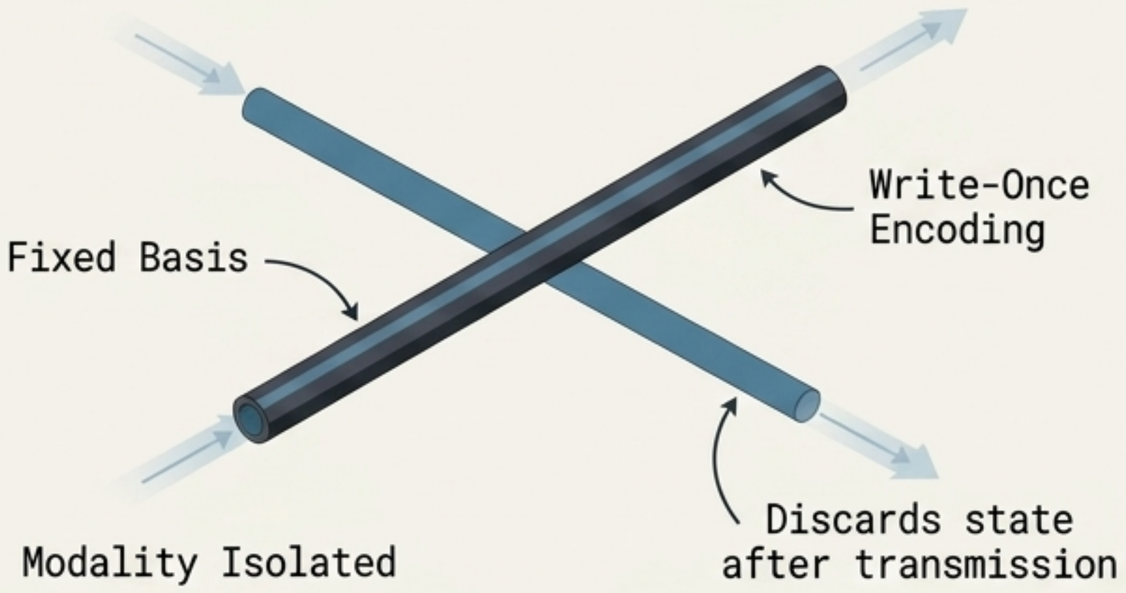

The Rematching Archive (New)

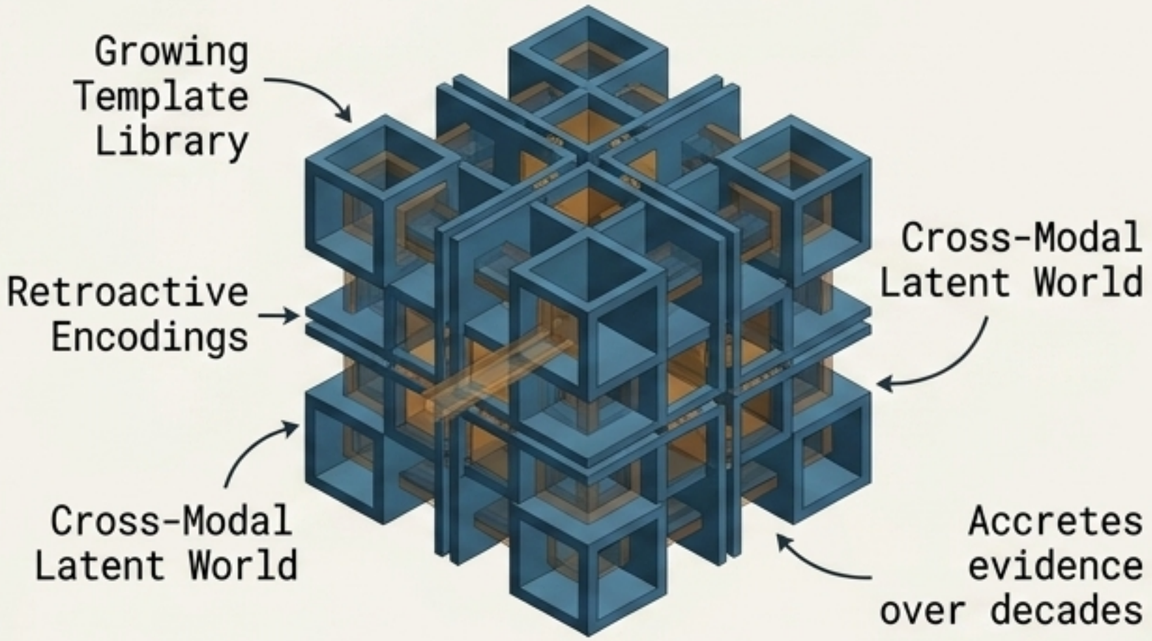

Shannon built a pipe for transient transmission. An archive is an epistemic agent. In a rematching archive, file size is a relational property, not an intrinsic one. The Marginal Cost is:

$$\text{cost}(o) = \Lambda(D) - \Lambda(D \setminus \{o\})$$

The 1,000th photo of a cat costs near zero bytes because the geometry is already reachable. Storage cost tracks reachability, not raw data volume.

The Devil’s Staircase of Ontological Expansion

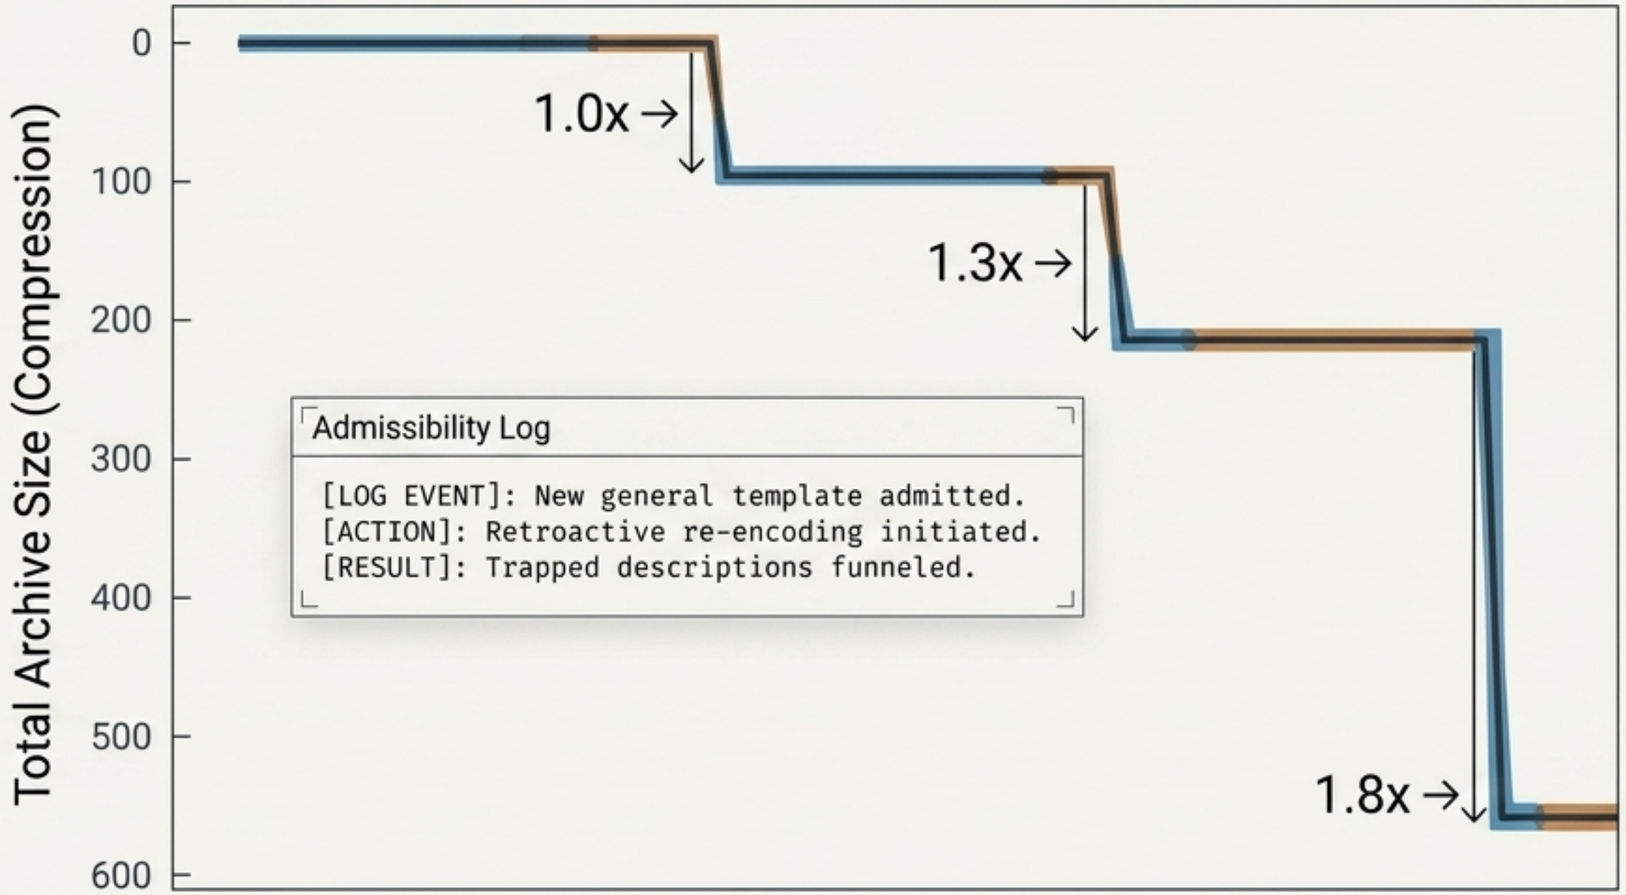

Compression does not improve smoothly. It improves in waves triggered by the admission of highly general templates that retroactively rewrite the archive.

The Math of Collapse:

$$C(\tau) = \Delta E(\tau) = \delta_T - \delta_{T \cup \{\tau\}}$$

Each plateau is a stratum of description space. The jumps represent phase transitions where trapped descriptions funnel through a newly opened ontological channel.

Persistent Anomalies as Gradients Toward Discovery

The Ontological Deficit Gradient

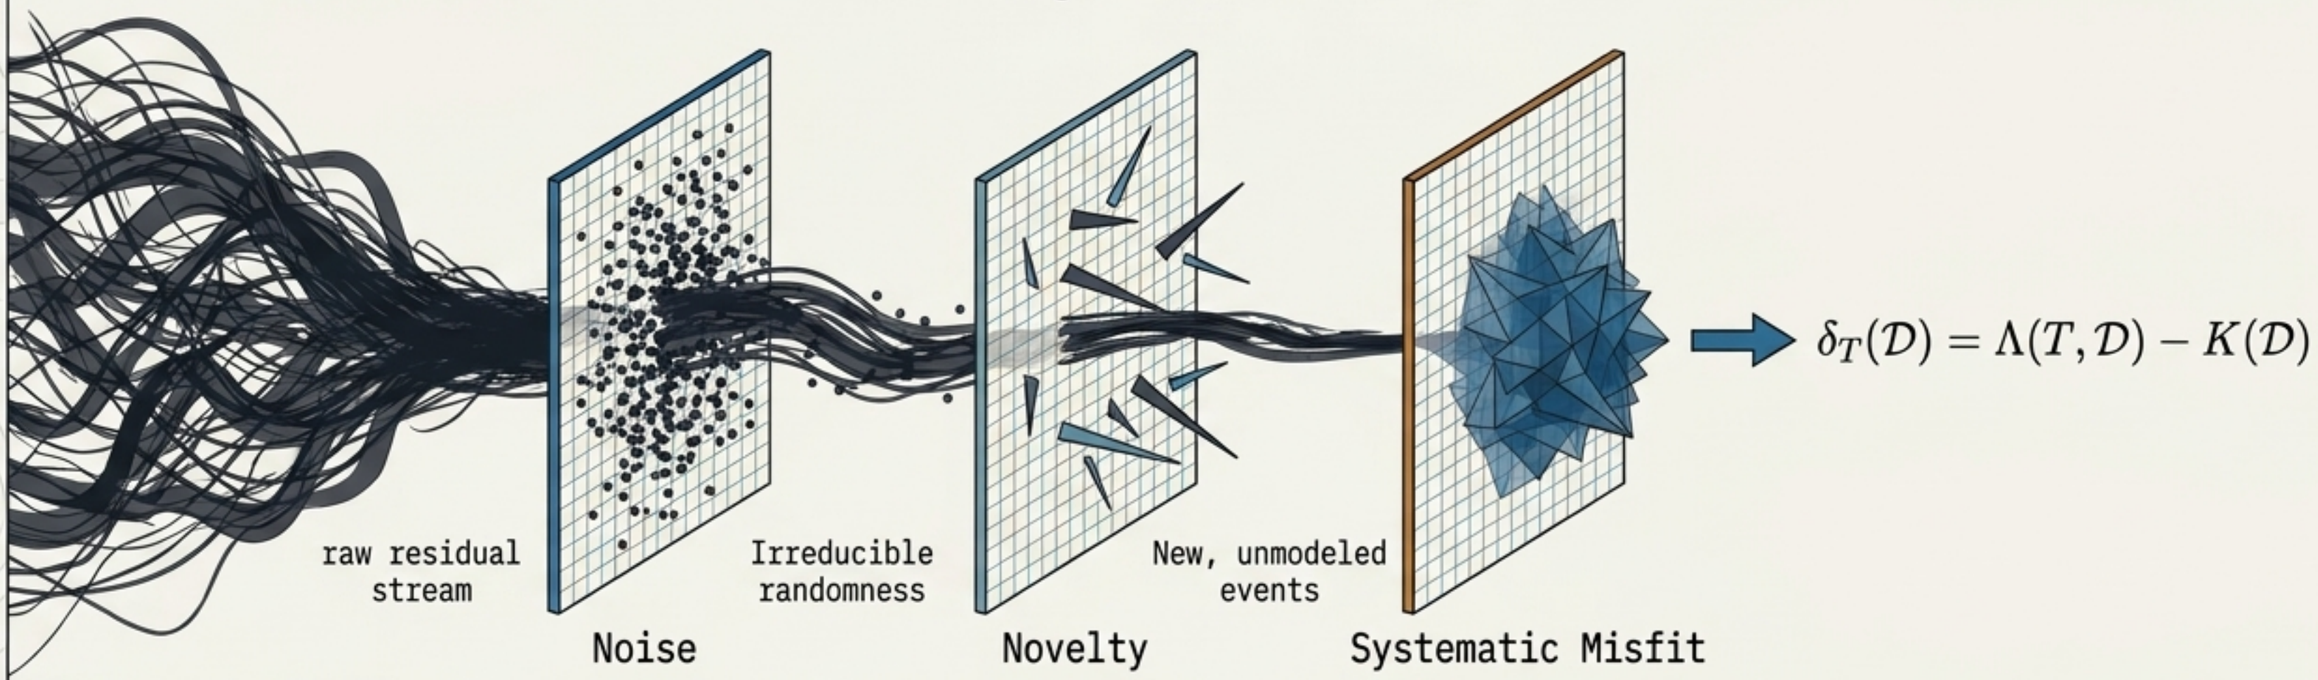

The **Ontological Deficit** $\delta_T(\mathcal{D})$ measures what the current ontology cannot say. A persistent anomaly is not a failure of the theory; it is a mathematical gradient pointing out of the current stratum toward a new discovery.

Anomaly Pressure (P_A): When trapped anomaly mass builds up without a valid template, it forces an ontology revision. The archive's next discovery is indicated directly by the geometry of the residual.

The World-Model Crossover Threshold (n^*)

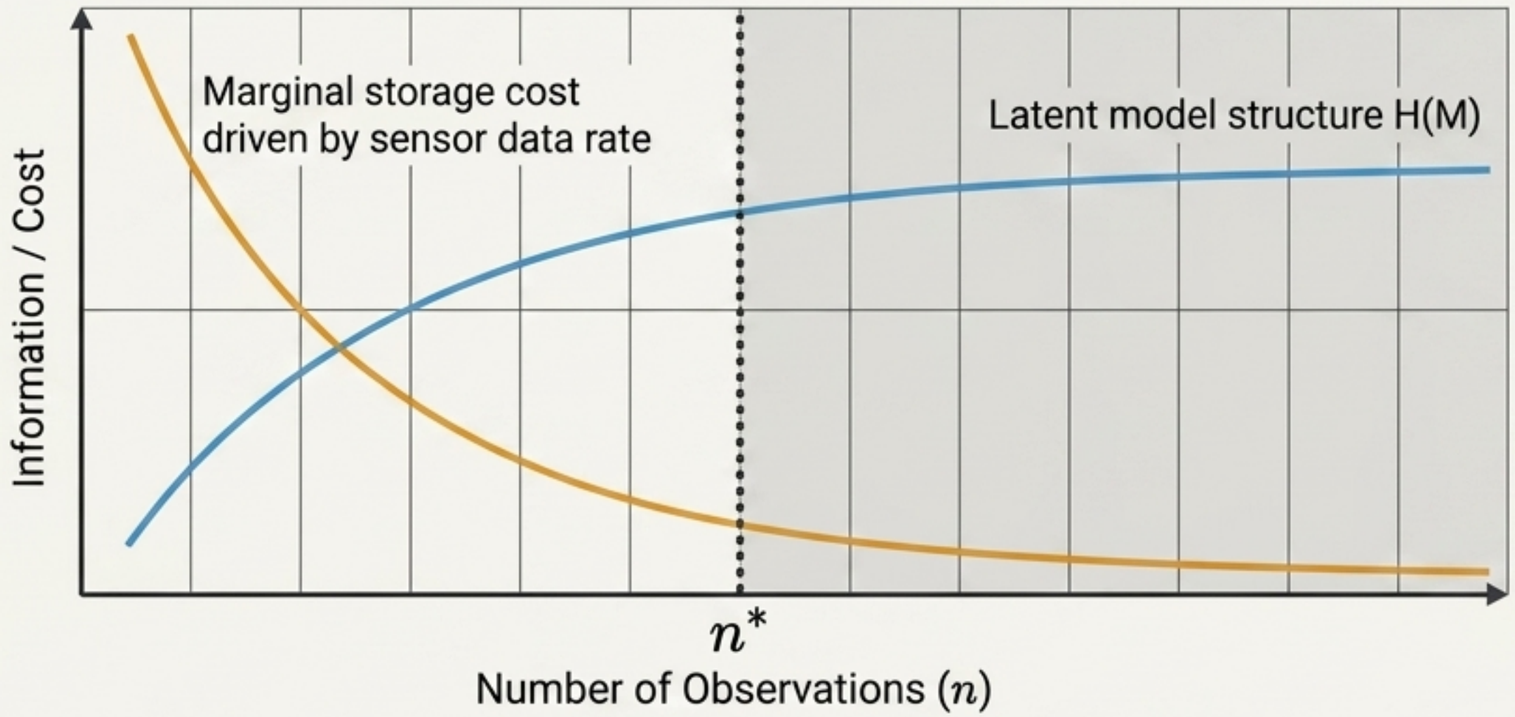

The Theorem

There exists a finite horizon n^* past which the latent model is the dominant information-bearing object.

$$I(M; \{O_1, \ldots, O_n\}) \to H(M)$$

The Consequence

Before n^*, observations are primary. After n^*, the theory is primary, and new observations are just residuals. Storage cost tracks admissible volume, not signal volume. Observations become projections of a shared latent world.

A Unified Characterization of Theory

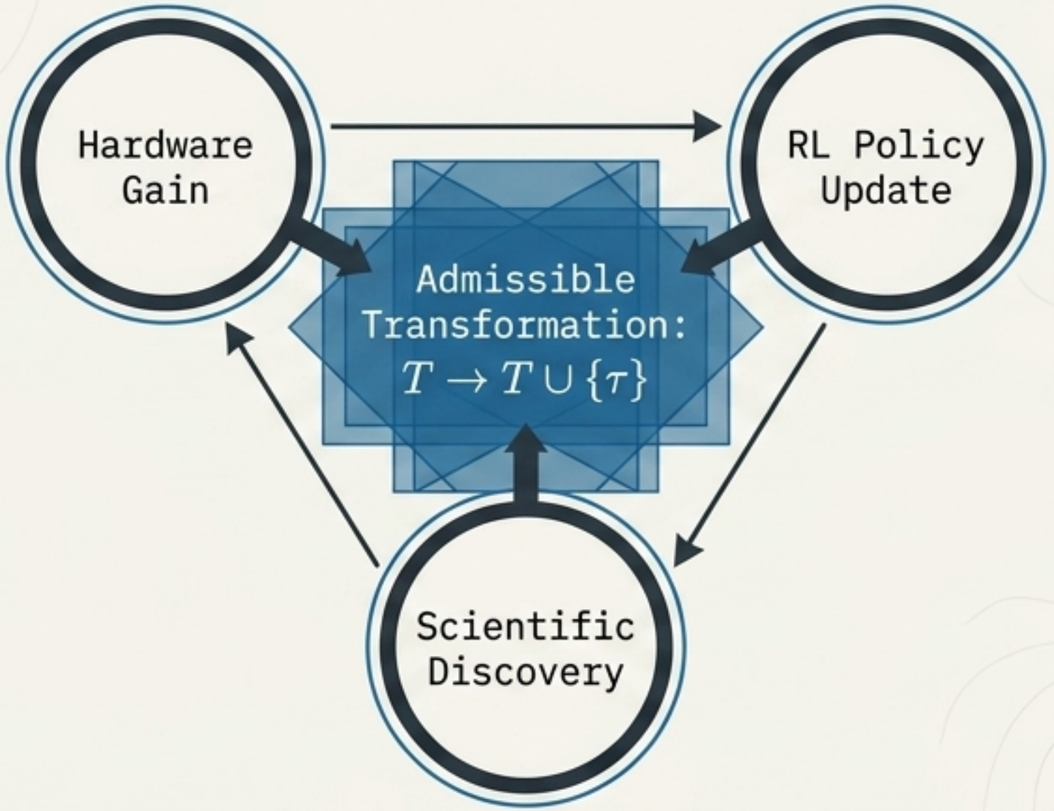

Principle 2 (Explanation as Ontological Enlargement): A successful theory must do three things:

1. Reduce Ontological Deficit:	2. Expand Reachable Description Space:	3. Increase Future Compression:
$\delta_{T \cup \{\tau\}} < \delta_T$	$\mathcal{R}_T(\mathcal{D}) \subset \mathcal{R}_{T \cup \{\tau\}}(\mathcal{D})$	$\Lambda(T \cup \{\tau\}, \mathcal{D}') < \Lambda(T, \mathcal{D}')$

Scientific theories, compression templates, RL updates, and chip design are structurally identical: they construct new admissible routes through a constrained landscape. Explanation creates futures.

Toward an Admissibility-Based Learning Objective

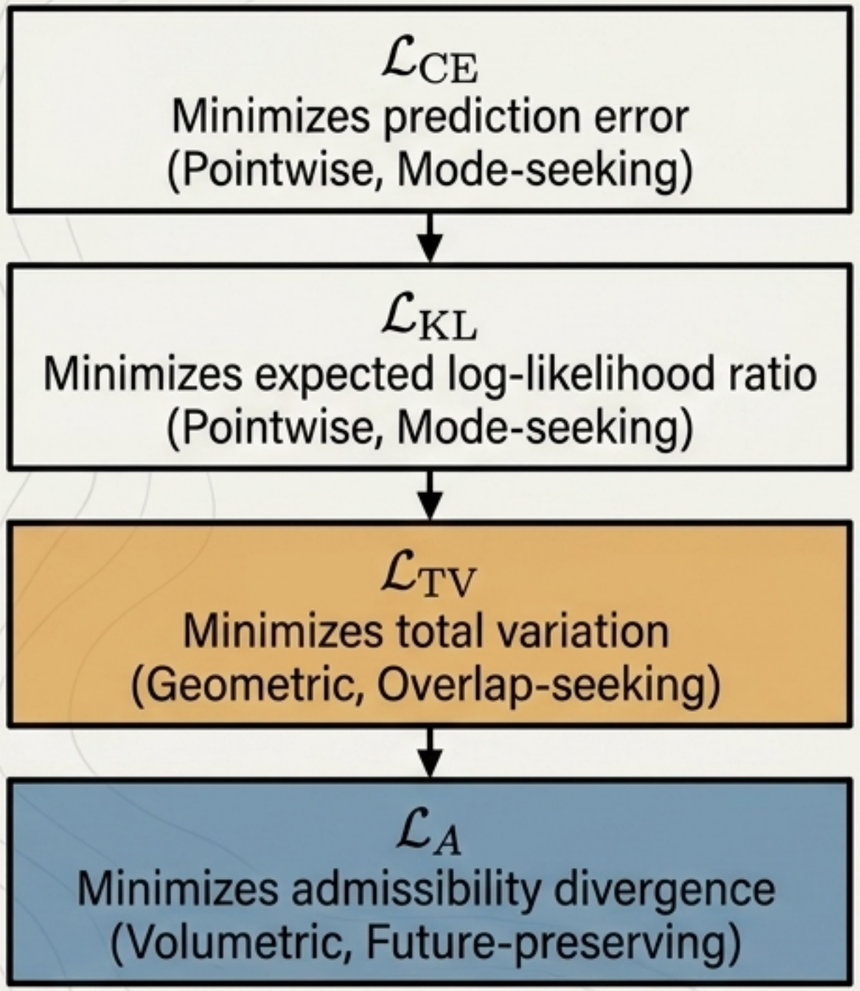

The Admissibility Objective

$$\mathcal{L}_A(p, q) = 1 - \frac{\mathrm{Vol}(\mathcal{A}_p \cap \mathcal{A}_q)}{\mathrm{Vol}(\mathcal{A}_p)}$$

This objective directly measures the fraction of the target's admissible futures that remain admissible under the draft. $\mathcal{L}_{\mathrm{TV}}$ is strictly a degenerate, flat-simplex case of $\mathcal{L}_A$.

Future RL

Instead of rewarding high-reward trajectories, we reward policies for maintaining high admissibility overlap with a reference policy: $r_A(p, p_{ref})$.

The Paradigm Inversion

Standard Framing	Reachability Framing
Intelligence ≈ prediction	Intelligence ≈ future preservation
Next token is the objective	Access to continuations is the objective
Search: find the optimal state	Navigation: stay in the admissible corridor
Entropy bounds performance	Admissible volume governs tracking difficulty
CE/KL optimize faithfulness	TV/**admissibility** optimize coverage
State preservation ⇒ identity	Future preservation ⇒ identity
Repair ≈ state restoration	Repair ≈ **admissibility** restoration
Discovery preserves options	Discovery creates new options

The Reachability Principle: Identity is an invariant of future accessibility, not of present state. The highest form of performance is **enlarging the space of futures that can be conceived, described, and reached.**