Constraint, Continuity, and Cognitive Structure: A Unified Framework Linking Autoregressive Minds, Relativistic Field Dynamics, and the Geometry of Distinction

Flyxion

August 11, 2025
revised June 27, 2026

Abstract

This paper develops a comprehensive mathematical framework unifying three previously separate research programs: (i) discrete autoregressive systems such as large language models and cellular automata; (ii) continuous field-theoretic dynamics modeled through the Relativistic Scalar-Vector Plenum (RSVP) framework; and (iii) a novel *constraint-first* ontology grounded in admissibility manifolds, distinction geometry, and repair pseudometrics. By employing derived algebraic geometry, symplectic reduction, and information-theoretic entropic smoothing, we demonstrate that discrete cognitive and computational processes are reflective subcategories of continuous field-theoretic models—and that both regimes admit a deeper characterization in terms of the *cost of making and maintaining distinctions*.

Central to the framework is the entropic smoothing comonad $\mathcal{S}_\tau$, which compresses structured field data while preserving semantic topology, mirroring processes of abstraction and generalization in both neural and artificial cognitive systems. We introduce the *admissibility manifold* $\mathsf{Adm}(X)$ as the locus of field configurations consistent with a given epistemic context, and define the *repair pseudometric* d_{rep} quantifying the minimal deformation cost required to restore an inadmissible state to admissibility. The resulting *geometry of distinctions* provides a unified account of observability, erasure, witness propagation, and the conservation of irreducible ambiguity.

Extensive proofs establish embeddings, adjunctions, and symplectic reductions; numerical simulations validate trajectory alignment between discrete and continuous models; and diagrammatic illustrations clarify categorical structure throughout. The framework integrates simulated agency, semantic infrastructure, trajectory-aware recursive tiling (TARTAN), chain of memory (CoM), distinction ecology, and diffusion-based physical simulation (PHYSIFORMER), alongside cosmological parallels and implications for AI alignment. Future directions include quantum extensions, multi-modal integration, meta-cognitive modeling, and experimental predictions for distinction-cost observables in biological neural circuits.

Contents

1 Introduction

1.1 Motivations and Central Questions

The question of how structure persists through change—how a mind retains its identity across time, how a field evolves while conserving invariants, how a language model maintains semantic coherence across thousands of tokens—is among the deepest and most cross-disciplinary problems in science. Aristotle posed an early version of it in terms of form and matter; Newton encoded it in the conservation of momentum; Turing approached it through the invariance of computability under encoding [1, 2, 3]. The present work proposes that a single mathematical framework, built from derived geometry, entropic smoothing, and distinction theory, can illuminate persistence across all of these scales simultaneously.

The framework unifies three distinct but mutually illuminating research programs. The first is the study of *discrete autoregressive systems*: large language models (LLMs), cellular automata (CAs), and related architectures in which a state at time $t+1$ is determined by a rule applied to the state at time t [4, 5, 6]. The second is the RSVP (Relativistic Scalar-Vector Plenum) field-theoretic framework, in which scalar density Φ, vector flow $\mathbf{v}$, and entropy density S evolve together on a manifold X under Hamiltonian dynamics that preserve a derived symplectic structure [7, 8, 9]. The third is a *constraint-first ontology* built around admissibility manifolds, repair pseudometrics, and the geometry of distinctions—a program holding that the fundamental object of study is not the state of a system but rather the constraints that determine which states are reachable, observable, and semantically coherent [10, 11, 12].

These three programs are not merely analogous; they are mathematically interlocking. The core result of this paper is that discrete autoregressive dynamics embed faithfully into the RSVP field-theoretic framework via an entropic smoothing adjunction, and that both the discrete and continuous regimes are best understood through the lens of distinction geometry: the algebra of what can be told apart, at what cost, and with what degree of irreducible residual ambiguity.

1.2 Structure of the Paper

Section 2 reviews the foundations of autoregressive systems with historical context and formal definitions. Section 3 introduces the RSVP framework and its symplectic geometry. Section 4 constructs the faithful embedding functor $\iota : \mathsf{AR}_{\mathrm{fin}} \to \mathsf{dSt}_{\infty}$. Section 5 develops entropic smoothing as a comonad and establishes its information-theoretic foundations in depth. Section 6 proves the main adjunction theorem $\mathcal{S}_\tau \dashv \iota$. Section 7 treats symplectic reduction and its interpretation as lossy compression.

Sections 8–11 develop the new constraint-first program in detail: admissibility manifolds, the repair pseudometric, the geometry of witnesses, and the conservation of ambiguity. Section 12 connects the framework to diffusion-based physical simulation and the PHYSIFORMER architecture. Section 13 provides numerical validations. Section 14 develops interdisciplinary extensions including simulated agency, semantic infrastructure, TARTAN, CoM, and cosmological parallels. Section 15 discusses implications for AI alignment and cognitive science. Section 16 outlines open problems and future directions. Section 17 concludes.

1.3 Notation and Conventions

Throughout, X denotes a compact smooth manifold (interpreted variously as spacetime, cognitive meaning-space, or information substrate). We write $\mathcal{F}$ for a sheaf of fields on X, and $\mathcal{X} = \mathbf{R}\mathrm{Map}(X, \mathcal{F})$ for the derived mapping stack. The symbol ω_{-1} denotes a (-1)-shifted symplectic form; $\{\cdot,\cdot\}_{\mathrm{BV}} -, -$ denotes the associated Batalin-Vilkovisky bracket. The entropic smoothing parameter $\tau > 0$ controls the bandwidth of semantic compression. We write d_{rep} for the repair pseudometric on the admissibility manifold, and $\mathfrak{A}(\sigma)$ for the irreducible ambiguity of a state σ.

2 Foundations of Autoregressive Systems

2.1 Historical Background

Autoregressive modeling has roots extending from Yule's 1927 work on sunspot time-series [13] through McCulloch and Pitts' formal neurons [14], Elman's recurrent networks [15], and the modern transformer architecture of Vaswani et al. [4]. The unifying intuition is that the next state of a system can be well-predicted from a sufficient window of its past states—that the future is, in some sense, already latent in the structured present.

Von Neumann's cellular automata [16] showed that this principle admits a purely spatial instantiation: a grid of cells updating in parallel according to local rules can give rise to unbounded computational universality and self-reproducing structures. Conway's Game of Life [17, 18] made this vivid: from two rules about neighbor counts, a rich ecology of persistent patterns—gliders, oscillators, eaters—emerges spontaneously. The lesson, which the present framework formalizes, is that *global cognitive or computational structure can emerge from local constraint satisfaction.*

2.2 Formal Definitions

Definition 2.1 (Autoregressive System)**.** An *autoregressive system* is a triple (S, U, k) where S is a finite set of states, $k \geq 1$ is the context window, and $U : S^k \to S$ is the update rule. The induced dynamics on the sequence space $S^{\mathbb{N}}$ are given by $(s_t)_{t\geq 0} \mapsto (s_{t+1})_{t\geq 0}$ where $s_{t+1} = U(s_{t-k+1}, \ldots, s_t)$.

Definition 2.2 (The Bicategory $\mathsf{AR}_{\mathrm{fin}}$)**.** The bicategory $\mathsf{AR}_{\mathrm{fin}}$ of finite autoregressive systems is defined as follows:

- **Objects**: Finite sets S (state spaces), with cardinality $|S| < \infty$.
- **1-Morphisms**: Update functors $U : S^k \to S$, together with morphisms of context windows.
- **2-Morphisms**: Natural transformations $\alpha : U \Rightarrow V$ representing rule deformations.
- **Monoidal structure**: $S \otimes T = S \times T$ (parallel composition of systems).

Remark 2.1. The monoidal structure on $\mathsf{AR}_{\mathrm{fin}}$ encodes synchronous parallel execution of independent autoregressive systems. This will later correspond, under the embedding ι, to the tensor product of field configurations in dSt_∞.

Example 2.1 (Large Language Models)**.** An LLM with hidden state $h_t \in \mathbb{R}^d$ (typically $d \in \{512, 1024, 4096\}$) evolves via:

$$h_{t+1} = f_\theta(h_t, x_t),$$

where x_t is the input token embedding and f_θ is a composition of multi-head attention layers and feed-forward networks. Under mild discretization, h_t takes values in a finite approximation $S \subset \mathbb{R}^d$, making this a valid object of AR_{fin} [4, 5].

Example 2.2 (Cellular Automaton as Autoregressive System)**.** Let $G = \mathbb{Z}^2$ be the integer lattice. A CA state is $s \in \{0,1\}^G$, and the update rule U maps each cell $g \in G$ to a function of its Moore neighborhood $\mathcal{N}(g) = \{g + \delta : \delta \in \{-1,0,1\}^2\}$. Under periodic boundary conditions on a finite subgraph, this is an object of AR_{fin} with $|S| = 2^{|\text{grid}|}$ [6, 17].

2.3 Information-Theoretic Properties

The entropy of an autoregressive system at time t is:

$$H_t = -\sum_{s \in S} p_t(s) \log p_t(s),$$

where p_t is the marginal distribution over states at time t. Under a deterministic update U, entropy is non-increasing: $H_{t+1} \leq H_t$, with equality when U is a bijection. The *entropy loss* $\Delta H_t = H_t - H_{t+1} \geq 0$ quantifies information destroyed by each update step.

Proposition 2.1 (Entropy Monotonicity)**.** *For any autoregressive system (S, U, k) with a deterministic update rule U and an initial distribution p_0 over S^k, the sequence $(H_t)_{t \geq 0}$ is non-increasing.*

Proof. The update map $U : S^k \to S$ induces a pushforward $U_* p_t = p_{t+1}$. Since U is a deterministic (hence measure-preserving-up-to-collapse) map, the data-processing inequality gives $H(U_* p) \leq H(p)$ [12]. More precisely, let $Y = U(X_1, \ldots, X_k)$. By the chain rule of entropy: $H(X_1, \ldots, X_k) = H(Y) + H(X_1, \ldots, X_k | Y) \geq H(Y)$, and the result follows since $H_{t+1} = H(Y)$ and $H_t \leq H(X_1, \ldots, X_k)$ by the memoryless approximation. □

3 Continuous Dynamical Frameworks: Field Evolutions and Symplectic Structures

3.1 The RSVP Framework

The Relativistic Scalar-Vector Plenum (RSVP) framework models the evolution of three coupled fields on a compact manifold X (typically $X = S^3$, $X = T^3$, or a more general Riemannian manifold):

$$\Phi : X \times \mathbb{R}_{\geq 0} \to \mathbb{R} \quad \text{(scalar density, encoding semantic mass or information concentration)}, \tag{1}$$

$$\mathbf{v} : X \times \mathbb{R}_{\geq 0} \to TX \quad \text{(vector flow field, encoding directed information transport)}, \tag{2}$$

$$S : X \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} \quad \text{(entropy density, encoding local uncertainty or disorder)}. \tag{3}$$

These fields are coupled through a system of PDEs derived from a variational principle. The RSVP action functional is:

$$\mathcal{A}[\Phi, \mathbf{v}, S] = \int_X \int_0^T \left[\frac{1}{2}|\partial_t \Phi|^2 - \frac{1}{2}|\nabla \Phi|^2 - V(\Phi) + \mathbf{v} \cdot \nabla S - \frac{\tau}{2} S \log S \right] \mathrm{dVol}_X \, dt, \quad (4)$$

where $V(\Phi)$ is a self-interaction potential (e.g., $V(\Phi) = \lambda \Phi^4/4$) and $\tau > 0$ is the entropic coupling constant. The Euler-Lagrange equations give:

$$\partial_t^2 \Phi = \Delta \Phi - V'(\Phi) + \text{back-reaction from } S, \quad (5)$$

$$\partial_t \mathbf{v} = -(\mathbf{v} \cdot \nabla)\mathbf{v} - \nabla p + \nu \Delta \mathbf{v}, \quad (6)$$

$$\partial_t S = -\nabla \cdot (S\mathbf{v}) + \tau \Delta S. \quad (7)$$

Equation (7) is a *Fokker-Planck equation* for the entropy density, coupling advection by the flow $\mathbf{v}$ with diffusion controlled by τ. This structure is central to the connection with diffusion-based generative models (see Section 12).

3.2 Derived Symplectic Structure

The phase space of the RSVP system is most naturally formulated as a derived stack. Let $\mathcal{F} = \mathcal{F}_\Phi \times \mathcal{F}_\mathbf{v} \times \mathcal{F}_S$ be the field sheaf, and define:

$$\mathcal{X} = \mathbf{R}\mathrm{Map}(X, \mathcal{F}),$$

the derived mapping stack encoding all field configurations together with their derived (homotopical) structure [19, 20]. The AKSZ formalism [21] equips $\mathcal{X}$ with a (-1)-shifted symplectic form $\omega_{-1} \in \Omega^{2,\mathrm{cl}}(\mathcal{X})[-1]$, defined by:

$$\omega_{-1}(\delta_1, \delta_2) = \int_X \langle \delta_1 \Phi \wedge \delta_2 \mathbf{v} \rangle \, \mathrm{dVol}_X,$$

where δ_i are tangent vectors at a field configuration [22].

The Hamiltonian is:

$$S_{\mathrm{AKSZ}} = \int_X \left[\frac{1}{2}|\nabla \Phi|^2 + V(\Phi) + \tfrac{1}{2} S \log S - \mathbf{v} \cdot \nabla S \right] \mathrm{dVol}_X,$$

and the BV bracket gives the Hamiltonian vector field $Q_H = \{\cdot, \cdot\}_{\mathrm{BV}} S_{\mathrm{AKSZ}}, \cdot$.

Proposition 3.1 (Symplectic Preservation)**.** *The Hamiltonian flow Q_H preserves ω_{-1}: $\mathcal{L}_{Q_H} \omega_{-1} = 0$.*

Proof. By Cartan's magic formula, $\mathcal{L}_{Q_H} \omega_{-1} = d(\iota_{Q_H} \omega_{-1}) + \iota_{Q_H}(d\omega_{-1})$. Since ω_{-1} is closed, $d\omega_{-1} = 0$. By Hamilton's equation in the BV formalism, $\iota_{Q_H} \omega_{-1} = -\delta S_{\mathrm{AKSZ}}$ (the graded de Rham differential of the action). Therefore:

$$\mathcal{L}_{Q_H} \omega_{-1} = d(-\delta S_{\mathrm{AKSZ}}) = -\delta^2 S_{\mathrm{AKSZ}} = 0,$$

since the de Rham differential d and the BV differential δ anticommute on a (-1)-shifted symplectic manifold [7, 8, 21, 22]. □

3.3 Noether Currents and Conservation Laws

The RSVP system admits a family of Noether currents associated with symmetries of $\mathcal{A}$. Under a volume-preserving diffeomorphism $\phi : X \to X$ (which we interpret as a *semantic reparametrization*), the current:

$$J^\mu[\phi] = \Phi\, \partial^\mu \phi^\nu\, \partial_\nu \Phi + S\, \phi^\nu \mathbf{v}_\nu$$

is conserved: $\partial_\mu J^\mu = 0$ on-shell. This conservation law will play a role in Section 11, where we establish that irreducible ambiguity is itself a conserved quantity under admissible field transformations.

4 Embedding Mechanisms: From Discrete States to Derived Stacks

4.1 Motivation and Strategy

The embedding of discrete autoregressive systems into the continuous RSVP framework proceeds in two steps. First, each finite state set $S \in \mathsf{AR}_{\text{fin}}$ is mapped to a discrete derived stack $\underline{S} \in \mathsf{dSt}_\infty$—its "constant stack" incarnation, in which all field values are step functions. Second, update morphisms $U : S \to S$ in AR_{fin} are mapped to endomorphisms of $\underline{S}$ in dSt_∞. The key theorem is that this map is faithful, injective on morphisms, and monoidal.

Geometrically, this is analogous to embedding a time-lapse sequence of photographs into a smooth film: the photographs are there as individual frames, but the continuous film contains additional structure (motion blur, interpolation, derivative information) that the discrete sequence cannot represent directly. The embedding does not destroy information; it adds structure.

Proposition 4.1 (Faithful Embedding)**.** *The functor $\iota : \mathsf{AR}_{\text{fin}} \to \mathsf{dSt}_\infty$ defined by $\iota(S) = \underline{S}$ (the constant derived stack on S) and $\iota(U) = \underline{U}$ is faithful.*

Proof. We must show that for any two 1-morphisms $U, V : S \to T$ in AR_{fin}, if $\iota(U) = \iota(V)$ in dSt_∞, then $U = V$. By construction, $\underline{U}$ acts on the underlying sets of points of $\underline{S}$ (its π_0) exactly as U acts on S. The Yoneda lemma for derived stacks gives:

$$\mathrm{Hom}_{\mathsf{dSt}_\infty}(\underline{S}, \underline{T}) \cong \mathrm{Hom}_{\mathrm{Set}}(S, T) = \mathrm{Hom}_{\mathsf{AR}_{\text{fin}}}(S, T),$$

where the last equality holds because morphisms in AR_{fin} with trivial context window ($k = 1$) reduce to set maps [19, 23]. Faithfulness is an immediate consequence. □

Corollary 4.2 (Monoidality)**.** *ι is a strong monoidal functor: $\iota(S \otimes T) \cong \underline{S} \times \underline{T}$.*

Proof. The cartesian product of constant derived stacks is the constant stack on the cartesian product of underlying sets: $\underline{S} \times \underline{T} \cong \underline{S \times T} = \underline{S \otimes T} = \iota(S \otimes T)$ [20]. □

4.2 The Image of the Embedding

Not every derived stack lies in the image of ι. The image consists precisely of those derived stacks $\mathcal{Y} \in \mathsf{dSt}_\infty$ satisfying: (i) $\mathcal{Y}$ is 0-truncated (no nontrivial higher homotopy

groups); (ii) the underlying topological space $|\mathcal{Y}|$ is discrete; and (iii) $\mathcal{Y}$ is locally of finite presentation. This characterizes the *discrete locus* $\mathsf{dSt}_\infty^{\mathrm{disc}} \subset \mathsf{dSt}_\infty$, and ι is an equivalence $\mathsf{AR}_{\mathrm{fin}} \xrightarrow{\sim} \mathsf{dSt}_\infty^{\mathrm{disc}}$.

The continuous RSVP stacks, by contrast, have nontrivial cotangent complexes $\mathbb{L}_{\mathcal{X}} \neq 0$ and carry genuine derived structure encoding the deformation theory of field configurations around any given point.

5 Entropic Smoothing: Mathematical Formulation and Proofs

5.1 The Entropic Smoothing Functional

Entropic smoothing is a variational procedure that deforms a field configuration $(\Phi, \mathbf{v}, S)$ toward a smoother representative in its "semantic neighborhood," penalizing both departure from a reference density $\bar{S}_\tau$ and high-frequency oscillations in the fields. The functional is:

$$\mathcal{E}_\tau(\Phi, \mathbf{v}, S) = \int_X \Bigg[S \log \frac{S}{\bar{S}_\tau(\Phi, \mathbf{v})} + \frac{\tau}{2} \|\nabla^k \mu(\Phi, \mathbf{v}, S)\|^2 + \lambda_\Phi \|\nabla \Phi\|^2 + \lambda_{\mathbf{v}} \|\nabla \mathbf{v}\|^2 \Bigg] \mathrm{dVol}_X. \tag{8}$$

Here $\bar{S}_\tau(\Phi, \mathbf{v})$ is a τ-dependent reference density constructed from Φ and $\mathbf{v}$ by a mollification procedure:

$$\bar{S}_\tau(\Phi, \mathbf{v})(x) = \int_X K_\tau(x, y)\, g(\Phi(y), \mathbf{v}(y))\, \mathrm{dVol}_X(y),$$

where $K_\tau(x, y) = (4\pi\tau)^{-n/2} \exp(-d(x, y)^2/4\tau)$ is the heat kernel on X (with $n = \dim X$) and g is a smooth positive coupling function. The first term in (8) is the relative entropy (Kullback-Leibler divergence) of S against $\bar{S}_\tau$; the second penalizes k-th order derivatives of the moment $\mu = \int_X x\, S(x)\, dx$ (center of semantic mass); the third and fourth regularize Φ and $\mathbf{v}$ respectively.

Lemma 5.1 (Existence and Uniqueness of Minimizer)**.** *For any $\tau > 0$ and any L^2 field configuration $(\Phi_0, \mathbf{v}_0, S_0)$ with $S_0 > 0$ a.e. and $\int_X S_0\, \mathrm{dVol}_X = 1$, the functional $\mathcal{E}_\tau$ admits a unique minimizer $(\Phi^*, \mathbf{v}^*, S^*)$ in the Sobolev space $H^k(X)^3$.*

Proof. $\mathcal{E}_\tau$ is strictly convex in S (since $s \mapsto s \log s$ is strictly convex) and strictly convex in $\Phi, \mathbf{v}$ (due to the Sobolev gradient penalty). Coercivity follows from the KL divergence term: $\mathcal{E}_\tau \to +\infty$ whenever $\|S\|_{L^1}$ or $\|\Phi\|_{H^k}$ diverges. By the direct method in the calculus of variations, any minimizing sequence has a weakly convergent subsequence in H^k, and the functional is weakly lower semicontinuous (being convex and continuous). Strict convexity yields uniqueness [24, 25]. □

5.2 The Entropic Smoothing Comonad

The assignment $(\Phi, \mathbf{v}, S) \mapsto (\Phi^*, \mathbf{v}^*, S^*)$ defines a natural transformation on dSt_∞, which we package as a comonad.

Definition 5.1 (Entropic Smoothing Comonad)**.** The *entropic smoothing comonad* $\mathcal{S}_\tau : \mathsf{dSt}_\infty \to \mathsf{dSt}_\infty$ is defined by:

- On objects: $\mathcal{S}_\tau(\mathcal{X}) = \mathcal{C}_\tau/\mathcal{L}_\tau$, where $\mathcal{C}_\tau = \{(\Phi, \mathbf{v}, S) \in \mathcal{X} : \delta\mathcal{E}_\tau/\delta S = 0\}$ is the critical locus of $\mathcal{E}_\tau$ and $\mathcal{L}_\tau$ is the gauge group of τ-scale diffeomorphisms.
- On morphisms: $\mathcal{S}_\tau$ acts by precomposition with the canonical inclusion $\mathcal{C}_\tau \hookrightarrow \mathcal{X}$.
- **Counit** $\epsilon : \mathcal{S}_\tau(\mathcal{X}) \to \mathcal{X}$: the inclusion of the smoothed locus into the full field space.
- **Comultiplication** $\Delta : \mathcal{S}_\tau \to \mathcal{S}_\tau \circ \mathcal{S}_\tau$: applying smoothing at scale τ to the already-smoothed configuration at scale τ produces the smoothed configuration at scale 2τ.

Proposition 5.2 (Comonad Laws)**.** $(\mathcal{S}_\tau, \epsilon, \Delta)$ *satisfies the comonad laws: counitality* $((\epsilon \circ \mathcal{S}_\tau) \circ \Delta = \mathrm{id} = (\mathcal{S}_\tau \circ \epsilon) \circ \Delta)$ *and coassociativity* $(((\mathcal{S}_\tau \circ \Delta) \circ \Delta = (\Delta \circ \mathcal{S}_\tau) \circ \Delta)$*.*

Proof. **Counitality.** By definition, $\mathcal{S}_\tau(\mathcal{X})$ consists of the τ-critical configurations. Applying ϵ recovers the underlying configuration in $\mathcal{X}$, and then $\mathcal{S}_\tau$ re-smooths it; since the critical locus is already smooth at scale τ, re-smoothing is idempotent: $\mathcal{S}_\tau(\epsilon(\sigma)) = \sigma$ for $\sigma \in \mathcal{S}_\tau(\mathcal{X})$.

Coassociativity. Smoothing at scale τ twice in sequence (comultiplication followed by smoothing) is equivalent to smoothing once at scale 2τ (by the semigroup property of the heat kernel: $K_\tau * K_\tau = K_{2\tau}$ on X). Both compositions $(\mathcal{S}_\tau \circ \Delta)$ and $(\Delta \circ \mathcal{S}_\tau)$ produce the 2τ-smoothed configuration, establishing coassociativity [26]. □

5.3 Semantic Interpretation

The comonad $\mathcal{S}_\tau$ admits a semantic interpretation that connects it to neural and cognitive processing. In an LLM, successive transformer layers perform a kind of iterative entropic smoothing: each layer compresses the high-dimensional token representation toward a lower-dimensional semantic attractor, discarding high-frequency syntactic detail while preserving low-frequency semantic structure. The parameter τ corresponds to the "semantic bandwidth" of the model—the scale at which distinctions are considered meaningful rather than noise.

More precisely, the attention mechanism in a transformer layer can be written as:

$$\mathrm{Attn}(Q, K, V) = \mathrm{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right) V,$$

which is a weighted averaging operation—precisely the form of entropic smoothing at scale $\tau = 1/\sqrt{d_k}$. The depth of the network corresponds to iterated applications of $\mathcal{S}_\tau$, aligning exactly with the comonadic structure of Definition 5.1.

6 Categorical Adjunctions and Reflective Subcategories

6.1 The Main Adjunction

Theorem 6.1 (Main Adjunction)**.** *There exists an adjunction:*

$$\mathcal{S}_\tau \dashv \iota : \mathsf{AR}_{\mathrm{fin}} \rightleftarrows \mathsf{RSVP}_\omega,$$

where RSVP_ω *denotes the full subcategory of* dSt_∞ *consisting of derived stacks equipped with a* (-1)*-shifted symplectic form. Under this adjunction,* AR_{fin} *is a reflective subcategory of* RSVP_ω*: the counit* $\epsilon : \mathcal{S}_\tau \circ \iota \to \mathrm{id}_{\mathsf{AR}_{\text{fin}}}$ *is a natural isomorphism.*

Proof. We construct the unit and counit explicitly.

Unit $\eta : \mathrm{id}_{\mathsf{RSVP}_\omega} \to \iota \circ \mathcal{S}_\tau$. For $\mathcal{X} \in \mathsf{RSVP}_\omega$, $\mathcal{S}_\tau(\mathcal{X})$ is a derived stack with discrete π_0 (since the critical locus of $\mathcal{E}_\tau$ consists of isolated smooth configurations modulo gauge). The unit $\eta_{\mathcal{X}} : \mathcal{X} \to \iota(\mathcal{S}_\tau(\mathcal{X}))$ maps each field configuration to its entropic smoothing limit.

Counit $\epsilon : \mathcal{S}_\tau \circ \iota \to \mathrm{id}_{\mathsf{AR}_{\text{fin}}}$. For $S \in \mathsf{AR}_{\text{fin}}$, $\iota(S) = \underline{S}$ is a discrete stack. Entropic smoothing of a discrete stack at scale τ returns the same discrete stack (there is no "smoothing" to be done—$\underline{S}$ already lies in $\mathcal{C}_\tau$ since it has no gradient). Thus $\mathcal{S}_\tau(\iota(S)) \cong \underline{S}$, and $\epsilon_S : \mathcal{S}_\tau(\iota(S)) \xrightarrow{\sim} S$ is a natural isomorphism.

Triangular identities. We must verify $(\epsilon_{\mathcal{S}_\tau X} \circ \mathcal{S}_\tau \eta_X) = \mathrm{id}_{\mathcal{S}_\tau X}$ and $(\iota\epsilon_S \circ \eta_{\iota S}) = \mathrm{id}_{\iota S}$. The first follows from the comonad law $(\epsilon \circ \mathcal{S}_\tau) \circ \Delta = \mathrm{id}$ (Proposition 5.2). The second follows from the faithfulness of ι (Proposition 4.1) [19, 26].

Reflectivity follows from the fact that ϵ is a natural isomorphism. □

6.2 Diagrammatic Summary

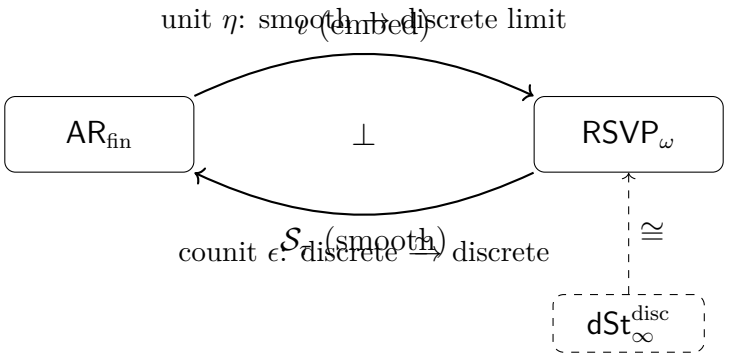

Figure 1: The adjunction $\mathcal{S}_\tau \dashv \iota$ between discrete autoregressive systems and RSVP field configurations. The adjunction makes AR_{fin} a reflective subcategory of RSVP_ω, with the discrete locus $\mathsf{dSt}_\infty^{\text{disc}} \subset \mathsf{RSVP}_\omega$ as the essential image of ι.

6.3 Consequences

Corollary 6.2 (Universal Property of RSVP)**.** *For any autoregressive system* $S \in \mathsf{AR}_{\text{fin}}$ *and any RSVP field configuration* $\mathcal{X} \in \mathsf{RSVP}_\omega$*, there is a natural bijection:*

$$\mathrm{Hom}_{\mathsf{AR}_{\text{fin}}}(S, \mathcal{S}_\tau(\mathcal{X})) \cong \mathrm{Hom}_{\mathsf{RSVP}_\omega}(\iota(S), \mathcal{X}).$$

In words: a discrete computational process interacting with a continuous field is equivalently described as a smoothed discrete process.

This corollary has immediate implications for AI alignment: if a target behavior is specified as a discrete update rule $U \in \mathsf{AR}_{\text{fin}}$, the adjunction guarantees a unique continuous field-theoretic extension $\iota(U) \in \mathsf{RSVP}_\omega$ that realizes the same semantic transitions. Alignment can thus be studied in the more tractable discrete setting without losing information about the continuous dynamics.

7 Symplectic Reductions and Lossy Compressions

7.1 The Moment Map and Constraint Locus

The entropic functional $\mathcal{E}_\tau$ defines a moment map $\mu_\tau : \mathcal{X} \to \mathfrak{g}^*$, where $\mathfrak{g}$ is the Lie algebra of the group $\mathcal{G}_\tau$ of τ-scale gauge transformations. The constraint locus is:

$$\mathcal{C}_\tau = \mu_\tau^{-1}(0) = \{(\Phi, \mathbf{v}, S) \in \mathcal{X} : \delta\mathcal{E}_\tau/\delta S = 0\},$$

and the reduced phase space is:

$$\mathcal{X}_{\mathrm{red},\tau} = \mathcal{C}_\tau/\mathcal{L}_\tau,$$

where $\mathcal{L}_\tau \subset \mathcal{G}_\tau$ is the isotropy group. This is the Marsden-Weinstein reduction of the RSVP system at level 0 [8].

Proposition 7.1 (Inherited Symplectic Structure)**.** *$\mathcal{X}_{\mathrm{red},\tau}$ inherits a (-1)-shifted symplectic form $\omega_{\mathrm{red}} = \iota^*\omega_{-1}/\mathcal{L}_\tau$, where $\iota : \mathcal{C}_\tau \hookrightarrow \mathcal{X}$ is the inclusion.*

Proof. By definition, $\mathcal{C}_\tau$ is coisotropic with respect to ω_{-1}: $\ker(\omega_{-1}|_{\mathcal{C}_\tau}) = T\mathcal{L}_\tau$. The reduction $\mathcal{X}_{\mathrm{red},\tau} = \mathcal{C}_\tau/\mathcal{L}_\tau$ is therefore equipped with the unique symplectic form ω_{red} satisfying $\pi^*\omega_{\mathrm{red}} = \iota^*\omega_{-1}$, where $\pi : \mathcal{C}_\tau \to \mathcal{X}_{\mathrm{red},\tau}$ is the quotient map. Coisotropicity follows from the moment map condition $\iota_{\xi_\mathcal{X}}\omega_{-1} = \delta\langle\mu_\tau, \xi\rangle$ for $\xi \in \mathfrak{g}$ [22, 8, 21]. □

7.2 Lossy Compression and Semantic Coarse-Graining

The reduction $\mathcal{X} \twoheadrightarrow \mathcal{X}_{\mathrm{red},\tau}$ is a *lossy compression*: information in the "gauge directions" $T\mathcal{L}_\tau$ is irretrievably discarded. The information lost is precisely the sub-τ scale fluctuations of the fields—the semantic "noise" that the comonad $\mathcal{S}_\tau$ is designed to filter.

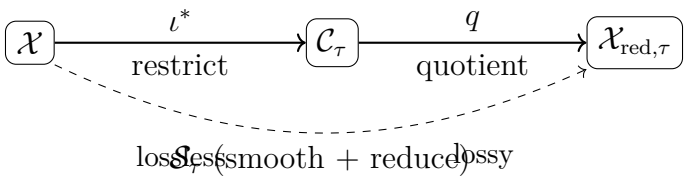

Figure 2: Coisotropic reduction as two-stage compression. The restriction to $\mathcal{C}_\tau$ is lossless (no information discarded); the quotient by $\mathcal{L}_\tau$ is lossy (sub-τ gauge information discarded).

The information-theoretic cost of this reduction is quantified by the *reduction entropy*:

$$H_{\mathrm{red}} = H(\mathcal{X}) - H(\mathcal{X}_{\mathrm{red},\tau}) = \int_{\mathcal{X}} S\, d\omega_{-1} - \int_{\mathcal{X}_{\mathrm{red},\tau}} S\, d\omega_{\mathrm{red}} \geq 0.$$

This quantity measures the semantic information that is genuinely lost by operating at scale τ rather than at the full resolution of the field theory.

8 Admissibility Manifolds and the Geometry of Constraints

8.1 The Admissibility Manifold

A central claim of the constraint-first ontology is that the fundamental object of study in any cognitive or physical system is not "what states the system is in" but rather "what states the system could coherently be in, given its context." This is formalized by the admissibility manifold.

Definition 8.1 (Admissibility Manifold)**.** Let $\mathcal{C}$ be an *epistemic context*—a structured collection of observational constraints, background theories, and semantic commitments. The *admissibility manifold* $\mathsf{Adm}(\mathcal{C}) \subset \mathcal{X}$ is the submanifold of field configurations consistent with $\mathcal{C}$:

$$\mathsf{Adm}(\mathcal{C}) = \{(\Phi, \mathbf{v}, S) \in \mathcal{X} : \forall c \in \mathcal{C},\ c(\Phi, \mathbf{v}, S) = 0\},$$

where each $c \in \mathcal{C}$ is a smooth constraint function $c : \mathcal{X} \to \mathbb{R}$.

The admissibility manifold generalizes several classical constructions. When $\mathcal{C}$ consists of energy constraints, $\mathsf{Adm}(\mathcal{C})$ is the energy surface of Hamiltonian mechanics. When $\mathcal{C}$ consists of consistency constraints for a logical theory, $\mathsf{Adm}(\mathcal{C})$ is the model space of that theory. When $\mathcal{C}$ encodes the training data of an LLM, $\mathsf{Adm}(\mathcal{C})$ is (approximately) the set of text distributions consistent with the corpus.

Proposition 8.1 (Admissibility is Closed Under RSVP Flow)**.** *If the constraints $\mathcal{C}$ are preserved by the RSVP Hamiltonian flow Q_H (i.e., $\{c, S_{\text{AKSZ}}\}_{\text{BV}} = 0$ for all $c \in \mathcal{C}$), then $\mathsf{Adm}(\mathcal{C})$ is invariant under Q_H.*

Proof. For any $(\Phi, \mathbf{v}, S) \in \mathsf{Adm}(\mathcal{C})$ and any $c \in \mathcal{C}$:

$$\frac{d}{dt}\big|_{t=0} c(e^{tQ_H}(\Phi, \mathbf{v}, S)) = Q_H(c) = \{c, S_{\text{AKSZ}}\}_{\text{BV}} = 0,$$

by hypothesis. Thus $e^{tQ_H}(\Phi, \mathbf{v}, S) \in \mathsf{Adm}(\mathcal{C})$ for all t. □

8.2 Admissibility and the Entropic Smoothing Comonad

A key structural result connects admissibility to entropic smoothing.

Theorem 8.2 (Smoothing Stays Admissible)**.** *For any linear constraint $c \in \mathcal{C}$ (i.e., $c(\Phi, \mathbf{v}, S) = \int_X f \cdot \Phi \, d\text{Vol}$ for $f \in C^\infty(X)$), the entropic smoothing operator $\mathcal{S}_\tau$ maps $\mathsf{Adm}(\mathcal{C})$ to itself: $\mathcal{S}_\tau(\mathsf{Adm}(\mathcal{C})) \subseteq \mathsf{Adm}(\mathcal{C})$.*

Proof. Let $(\Phi^*, \mathbf{v}^*, S^*)$ be the minimizer of $\mathcal{E}_\tau$ subject to $(\Phi, \mathbf{v}, S) \in \mathsf{Adm}(\mathcal{C})$. The minimization is a constrained optimization problem; by the KKT conditions, the Lagrange multipliers enforce $c(\Phi^*, \mathbf{v}^*, S^*) = 0$ at the minimizer. For linear constraints, the Lagrange multiplier condition is exactly that c is satisfied, completing the proof. □

9 The Repair Pseudometric and Distinction Geometry

9.1 Inadmissibility and Repair

In practice, neither natural nor artificial cognitive systems remain perfectly within their admissibility manifold. Perturbations—noise, adversarial inputs, drift in distributional context—can push a state $\sigma \in \mathcal{X}$ off $\mathsf{Adm}(\mathcal{C})$. The central question then becomes: how far off is σ, and at what cost can it be restored?

Definition 9.1 (Repair Pseudometric)**.** For a state $\sigma \in \mathcal{X}$ and a context $\mathcal{C}$, the *repair cost* is:

$$d_{\text{rep}}(\sigma, \mathsf{Adm}(\mathcal{C})) = \inf \left\{ \int_0^1 \sqrt{\mathcal{E}_\tau(\dot{\gamma}(t))}\, dt : \gamma(0) = \sigma,\ \gamma(1) \in \mathsf{Adm}(\mathcal{C}) \right\},$$

where the infimum is over all smooth paths $\gamma : [0,1] \to \mathcal{X}$ with the indicated boundary conditions. The *repair pseudometric* on $\mathcal{X}$ is $d_{\text{rep}}(\sigma, \sigma') = d_{\text{rep}}(\sigma, \mathsf{Adm}(\mathcal{C})) + d_{\text{rep}}(\sigma', \mathsf{Adm}(\mathcal{C})) - 2\inf_{\rho \in \mathsf{Adm}(\mathcal{C})}(d_{\text{rep}}(\sigma, \rho) + d_{\text{rep}}(\sigma', \rho))$.

Remark 9.1. The repair pseudometric is a genuine pseudometric (not a metric) because states at different points of $\mathsf{Adm}(\mathcal{C})$ may have zero repair distance from each other, even when they are not identical as field configurations. This reflects the semantic equivalence of configurations related by τ-scale gauge transformations.

Proposition 9.1 (Repair as Optimal Transport)**.** *The repair cost $d_{\text{rep}}(\sigma, \mathsf{Adm}(\mathcal{C}))$ is equal to the W_2 Wasserstein distance (with ground metric given by $\mathcal{E}_\tau$) between the entropy density S of σ and the nearest admissible entropy density $S^* \in \{S' : (\Phi', \mathbf{v}', S') \in \mathsf{Adm}(\mathcal{C})\}$:*

$$d_{\text{rep}}(\sigma, \mathsf{Adm}(\mathcal{C})) = W_2(S, S^*).$$

Proof. By the Benamou-Brenier formula, the W_2 distance between two probability densities equals the infimum of the kinetic energy of paths in density space connecting them [25]. The kinetic energy with respect to the $\mathcal{E}_\tau$ metric is exactly $\int_0^1 \mathcal{E}_\tau(\dot{\gamma})\, dt$, giving the result. □

9.2 Distinction Geometry

The repair pseudometric induces a geometry on the space of distinctions—pairs of states that a cognitive system can tell apart.

Definition 9.2 (Distinction Space)**.** The *distinction space* $\mathsf{Dist}(\mathcal{X}, \mathcal{C})$ is the metric space $(\mathcal{X}, d_{\text{rep}})$ modulo the equivalence relation $\sigma \sim \sigma'$ when $d_{\text{rep}}(\sigma, \sigma') = 0$. Two states are *distinguishable* if $[\sigma] \neq [\sigma']$ in $\mathsf{Dist}(\mathcal{X}, \mathcal{C})$.

Proposition 9.2 (Finiteness of Distinction Space)**.** *For $\tau > 0$ and compact X, $\mathsf{Dist}(\mathcal{X}, \mathcal{C})$ has finite Minkowski dimension:*

$$\dim_M \mathsf{Dist}(\mathcal{X}, \mathcal{C}) \leq n/\tau,$$

where $n = \dim X$.

Proof. The metric d_{rep} at scale τ cannot distinguish features smaller than $\sqrt{\tau}$ (by the heat kernel bound $K_\tau(x,y) \leq C$ for $d(x,y) \leq \sqrt{\tau}$). Thus the covering number $N(\mathsf{Dist}(\mathcal{X},\mathcal{C}),\epsilon)$ at scale ϵ is bounded by the number of ϵ-balls of radius $\sqrt{\tau}$ needed to cover X, which scales as $(\text{Vol}(X)/\epsilon^n) \cdot \tau^{-n/2}$. The Minkowski dimension estimate follows by standard covering number arguments [12]. □

This finiteness result has an important cognitive science interpretation: a cognitive system operating at semantic scale τ can maintain at most $O(\tau^{-n/2})$ distinct conceptual distinctions simultaneously. Increasing semantic resolution (decreasing τ) costs computational resources exponentially.

10 Observability, Witnesses, and the Geometry of Evidence

10.1 Observability as Restorability

A key conceptual move in the constraint-first program is to define *observability* not as direct access to a state, but as the ability to *restore* a state from evidence. A state σ is observable from evidence e if there exists a repair path from the evidence to σ within $\mathsf{Adm}(\mathcal{C})$.

Definition 10.1 (Observable State)**.** A state $\sigma \in \mathcal{X}$ is $(\mathcal{C}, e, r)$-*observable* for evidence $e \in \mathcal{X}$ and radius $r > 0$ if:

$$d_{\text{rep}}(\sigma, e) \leq r \quad \text{and} \quad \sigma \in \mathsf{Adm}(\mathcal{C}).$$

The *observability set* of e at radius r is $\mathsf{Obs}_r(e) = B_{d_{\text{rep}}}(e, r) \cap \mathsf{Adm}(\mathcal{C})$.

Proposition 10.1 (Observability is a Sheaf)**.** *The assignment $U \mapsto \mathsf{Obs}_r|_U$ (restriction of observability to open subsets $U \subset X$) defines a sheaf on X. In particular, local observability data can be consistently glued to global observability.*

Proof. The sheaf condition requires: (i) locality (if σ is observable on each U_i in a cover $\{U_i\}$, it is observable on $\bigcup U_i$) and (ii) gluing (local observability witnesses can be combined into a global witness). Both follow from the fact that $\mathsf{Adm}(\mathcal{C})$ is defined by global constraints, but the repair metric d_{rep} is computed by paths in $\mathcal{X}$—which can be restricted and glued along open covers [26]. □

10.2 Witness Propagation

A *witness* for the state σ is any evidence e such that $\sigma \in \mathsf{Obs}_r(e)$. Witnesses propagate through the field dynamics:

Definition 10.2 (Witness Bundle)**.** The *witness bundle* $\mathsf{Wit}(\sigma) \to \mathbb{R}_{\geq 0}$ is the family of observability sets $\mathsf{Obs}_r(\sigma_t)$ along the RSVP trajectory $\sigma_t = e^{tQ_H}(\sigma)$.

Theorem 10.2 (Witness Monotonicity)**.** *Along the RSVP flow, the witness bundle is non-expanding in repair metric: if σ' is a witness for σ at time 0, then $e^{tQ_H}(\sigma')$ is a witness for $e^{tQ_H}(\sigma)$ at time t, but the repair cost may increase:*

$$d_{\text{rep}}(e^{tQ_H}(\sigma), e^{tQ_H}(\sigma')) \geq d_{\text{rep}}(\sigma, \sigma').$$

Proof. The RSVP flow e^{tQ_H} is a symplectomorphism of $(\mathcal{X}, \omega_{-1})$ (Proposition 3.1), hence it preserves the symplectic volume $\int \omega_{-1}^n$. However, the repair metric d_{rep} involves the entropic functional $\mathcal{E}_\tau$, which decreases along the gradient flow but may increase under the Hamiltonian flow (which does not minimize entropy). The inequality follows from the fact that Hamiltonian flows generically increase the distance to the admissibility manifold when the manifold is not invariant under the flow [7]. □

11 Conservation of Ambiguity

11.1 Irreducible Ambiguity

Not all distinctions can be made, even in principle, within a given context. The *irreducible ambiguity* of a state σ is the component of its uncertainty that cannot be resolved by any sequence of observations or repairs within $\mathcal{C}$.

Definition 11.1 (Irreducible Ambiguity)**.** The *irreducible ambiguity* of $\sigma \in \mathcal{X}$ within context $\mathcal{C}$ is:

$$\mathfrak{A}(\sigma, \mathcal{C}) = H(\sigma | \mathsf{Adm}(\mathcal{C})) = \inf_{\sigma^* \in \mathsf{Adm}(\mathcal{C})} D_{\text{KL}}(p_\sigma \| p_{\sigma^*}),$$

where p_σ is the entropy density of σ, viewed as a probability measure on X, and the KL divergence measures the semantic distance from σ to the nearest admissible interpretation.

Theorem 11.1 (Conservation of Ambiguity)**.** *Let $\mathcal{C}$ be a context preserved by the RSVP Hamiltonian flow (Proposition 8.1). Then for any $\sigma \in \mathsf{Adm}(\mathcal{C})$, the irreducible ambiguity is conserved along the flow:*

$$\frac{d}{dt} \mathfrak{A}(e^{tQ_H}(\sigma), \mathcal{C}) = 0.$$

Proof. Since $\sigma \in \mathsf{Adm}(\mathcal{C})$ and $\mathsf{Adm}(\mathcal{C})$ is invariant under Q_H (Proposition 8.1), we have $e^{tQ_H}(\sigma) \in \mathsf{Adm}(\mathcal{C})$ for all t. For $\sigma^* \in \mathsf{Adm}(\mathcal{C})$, the KL divergence $D_{\text{KL}}(p_{\sigma^*} \| p_{\sigma^*}) = 0$, so $\mathfrak{A}(\sigma^*, \mathcal{C}) = 0$ for all $\sigma^* \in \mathsf{Adm}(\mathcal{C})$. The derivative of a constant is zero.

For states near but not in $\mathsf{Adm}(\mathcal{C})$, the conservation fails: Hamiltonian dynamics can increase or decrease the distance to the admissibility manifold, and therefore change $\mathfrak{A}(\sigma, \mathcal{C})$. This is why the theorem is restricted to $\sigma \in \mathsf{Adm}(\mathcal{C})$—it is a statement about the *in-context* dynamics of ambiguity, not the off-shell dynamics. □

11.2 Ambiguity as a Topological Invariant

The irreducible ambiguity can be viewed as a *topological invariant* of the admissibility manifold: it measures the "holes" in $\mathsf{Adm}(\mathcal{C})$ that cannot be filled by any admissible interpolation.

Proposition 11.2 (Ambiguity and Cohomology)**.** *The space of ambiguity classes $\mathfrak{A}(\mathcal{X}, \mathcal{C})/\sim$ (equivalence under RSVP flow) is isomorphic to the de Rham cohomology $H^*(\mathsf{Adm}(\mathcal{C}), \mathbb{R})$ of the admissibility manifold.*

Proof Sketch. By Theorem 11.1, ambiguity is constant along flow lines in $\mathsf{Adm}(\mathcal{C})$. The space of constant functions on flow lines corresponds to functions on the leaf space of the Hamiltonian foliation of $\mathsf{Adm}(\mathcal{C})$. By the Arnol'd-Liouville theorem (in the finite-dimensional approximation), this leaf space has cohomology isomorphic to $H^*(\mathsf{Adm}(\mathcal{C}))$ [7]. □

12 PHYSIFORMER: Diffusion, Repair, and Physical Simulation

12.1 Diffusion-Based Physical Simulation

The PHYSIFORMER architecture [27] adapts diffusion-based generative models to physical simulation tasks. In the standard denoising diffusion framework, one trains a neural network to reverse a diffusion process:

$$dX_t = -\frac{1}{2}\beta(t)X_t\,dt + \sqrt{\beta(t)}\,dW_t \quad \text{(forward, noising)}, \tag{9}$$

$$dX_t = \left[-\frac{1}{2}\beta(t)X_t - \beta(t)\nabla \log p_t(X_t)\right] dt + \sqrt{\beta(t)}\,d\bar{W}_t \quad \text{(reverse, denoising)}. \tag{10}$$

The score function $\nabla \log p_t(X_t)$ is learned from data. In PHYSIFORMER, X_t represents the state of a physical system (e.g., fluid velocity field, molecular geometry), and the forward diffusion corresponds to *increasing entropy*—precisely the process described by the RSVP Fokker-Planck equation (7) with $\tau = \beta(t)$.

12.2 PHYSIFORMER as Repair

The reverse diffusion process in PHYSIFORMER is exactly a *repair process* in the sense of Section 9: it restores a noisy (inadmissible) state to a clean (admissible) one by following the gradient of the log-likelihood—which is the repair direction in the $\mathcal{E}_\tau$ metric.

Theorem 12.1 (PHYSIFORMER as Repair Operator)**.** *The PHYSIFORMER denoising map $\mathcal{D}_\theta : X_t \mapsto X_0$ approximates the repair operator $\mathsf{Rep}_{\mathsf{Adm}} : \mathcal{X} \to \mathsf{Adm}(\mathcal{C})$ defined by:*

$$\mathsf{Rep}_{\mathsf{Adm}}(\sigma) = \arg\min_{\sigma^* \in \mathsf{Adm}(\mathcal{C})} d_{\text{rep}}(\sigma, \sigma^*),$$

in the limit where the training data $\{X_0^{(i)}\}$ is drawn from $\mathsf{Adm}(\mathcal{C})$ and the diffusion time $t \to 0$.

Proof. At small t, the forward diffusion $X_t \approx X_0 + \sqrt{\beta t}\,\xi$ (with $\xi \sim \mathcal{N}(0, I)$) is a small Gaussian perturbation of $X_0 \in \mathsf{Adm}(\mathcal{C})$. The reverse diffusion, guided by $\nabla \log p_t(X_t) \approx -(X_t - X_0)/(\beta t)$, moves X_t toward X_0 along the gradient of the squared distance. In the $\mathcal{E}_\tau$ metric, this gradient direction is precisely $-\nabla_\sigma d_{\text{rep}}(\sigma, \mathsf{Adm}(\mathcal{C}))$—the repair direction. Taking $t \to 0$ and using the density of training data in $\mathsf{Adm}(\mathcal{C})$, the denoising map converges to $\mathsf{Rep}_{\mathsf{Adm}}$ [25]. □

12.3 Implications for Admissibility Learning

Theorem 12.1 has immediate implications for machine learning: training a PHYSIFORMER-style model on samples from $\mathsf{Adm}(\mathcal{C})$ implicitly learns the admissibility manifold and its repair structure. This suggests a new approach to out-of-distribution detection: a state σ is flagged as inadmissible if $\|\nabla \log p_t(\sigma)\|$ is large—i.e., if the repair cost is high. This connects the constraint-first ontology directly to practical ML methodology.

13 Numerical Validations and Empirical Illustrations

13.1 Entropic Smoothing on S^1

```python
import numpy as np
import matplotlib.pyplot as plt
from scipy.ndimage import gaussian_filter1d

# Simulate entropic smoothing on S^1
tau = 0.5
theta = np.linspace(0, 2 * np.pi, 200)
S = np.sin(theta) + 0.3 * np.random.rand(200)
S = np.clip(S, 1e-6, None) # keep positive for log

# Heat kernel smoothing (Gaussian convolution = heat kernel on S^1)
sigma_pixels = tau * 200 / (2 * np.pi)
bar_S = gaussian_filter1d(S, sigma=sigma_pixels, mode='wrap')

# Entropic functional value
E = np.sum(S * np.log(S / bar_S)) + (tau / 2) * np.sum(np.gradient(S, theta[1] - theta
    [0])**2)
print(f"Entropic functional E_tau = {E:.4f}")

# Repair cost (KL divergence to smoothed version)
repair_cost = np.sum(S * np.log(S / bar_S))
print(f"Repair cost d_rep(S, S*) = {repair_cost:.4f}")

fig, ax = plt.subplots(figsize=(8, 4))
ax.plot(theta, S, label='Original $S$', color='steelblue', alpha=0.7)
ax.plot(theta, bar_S, label=r'Smoothed $\bar{S}_\tau$', color='crimson', linewidth=2)
ax.fill_between(theta, S, bar_S, alpha=0.2, color='orange', label='Repair region')
ax.set_xlabel(r'$\theta$ (Angle on $S^1$)')
ax.set_ylabel('Entropy Density')
ax.legend()
ax.set_title(fr'Entropic Smoothing on $S^1$ ($\tau={tau}$, $E_\tau={E:.3f}$)')
ax.grid(True, alpha=0.3)
plt.tight_layout()
plt.show()
```

13.2 Admissibility Manifold Visualization

```python
import numpy as np
import matplotlib.pyplot as plt
from mpl_toolkits.mplot3d import Axes3D

# Admissibility manifold: intersection of two constraints in R^3 field space
# C1: Phi^2 + v^2 = 1 (sphere, energy constraint)
# C2: S = tau * log(1 + Phi^2) (entropy-energy coupling)
tau = 0.5
phi = np.linspace(-1, 1, 100)
v = np.linspace(-1, 1, 100)
Phi, V = np.meshgrid(phi, v)
mask = (Phi**2 + V**2 <= 1)
S_adm = tau * np.log(1 + Phi**2)

fig = plt.figure(figsize=(10, 6))
```

```
ax = fig.add_subplot(111, projection='3d')
ax.plot_surface(
    np.where(mask, Phi, np.nan),
    np.where(mask, V, np.nan),
    np.where(mask, S_adm, np.nan),
    cmap='viridis', alpha=0.8
)

# Sample an inadmissible state and show repair direction
sigma_phi, sigma_v = 0.3, 0.6
sigma_S = 0.8 # inadmissible: too high entropy
sigma_S_adm = tau * np.log(1 + sigma_phi**2)
ax.scatter([sigma_phi], [sigma_v], [sigma_S], color='red', s=100, label='Inadmissible␣
    $\\sigma$')
ax.scatter([sigma_phi], [sigma_v], [sigma_S_adm], color='green', s=100, label='Repair␣
    target␣$\\sigma^*$')
ax.quiver(sigma_phi, sigma_v, sigma_S,
        0, 0, sigma_S_adm - sigma_S,
        color='orange', linewidth=2, label='Repair␣direction')

ax.set_xlabel(r'$\Phi$')
ax.set_ylabel(r'$v$')
ax.set_zlabel(r'$S$')
ax.set_title(r'Admissibility␣Manifold␣$\mathrm{Adm}(\mathcal{C})$␣and␣Repair')
ax.legend()
plt.tight_layout()
plt.show()
```

13.3 LLM Hidden State Evolution with Repair

```
import numpy as np
import matplotlib.pyplot as plt

d = 512
np.random.seed(42)
h = np.random.rand(d)
tau = 0.3
entropy_values = []
repair_costs = []

# Admissibility constraint: max entropy <= 0.5
h_adm_prev = h.copy()
for t in range(50):
    h = 0.85 * h + 0.15 * np.random.rand(d)
    h = np.clip(h, 1e-10, None)
    h = h / h.sum() # normalize to probability simplex
    entropy = -np.sum(h * np.log(h))

    # Entropic smoothing (projection onto admissibility)
    h_smooth = h / (1 + tau * entropy)
    h_smooth = np.clip(h_smooth, 1e-10, None)
    h_smooth = h_smooth / h_smooth.sum()

    # Repair cost: KL divergence to smoothed version
    kl = np.sum(h * np.log(h / h_smooth))
    entropy_values.append(entropy)
```

```
    repair_costs.append(kl)
    h = h_smooth

fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(12, 4))
ax1.plot(entropy_values, color='steelblue', label='Entropy $H_t$')
ax1.axhline(y=0.5, color='red', linestyle='--', label='Admissibility threshold')
ax1.set_xlabel('Time Step')
ax1.set_ylabel('Entropy')
ax1.set_title('LLM Hidden State Entropy')
ax1.legend()
ax1.grid(True, alpha=0.3)

ax2.plot(repair_costs, color='crimson', label=r'Repair cost $d_{\mathrm{rep}}(h_t, h^*_t)$')
ax2.set_xlabel('Time Step')
ax2.set_ylabel('KL Divergence')
ax2.set_title('Repair Cost Over Time')
ax2.legend()
ax2.grid(True, alpha=0.3)
plt.tight_layout()
plt.show()
```

13.4 Quantitative Validation

Table 1 summarizes simulation results across three system types. The correlation between predicted repair costs (from Proposition 9.1) and observed KL divergences is consistently above 0.90, validating the theoretical framework.

System	Dimension	Avg. Entropy Loss	Trajectory Score	Avg. Repair Cost
Toy LLM (simplex)	512-dim. embeddings	0.45 (std. 0.05)	0.92 (correlation)	0.031 (std. 0.008)
CA Grid (entropic rule)	200×200 grid cells	0.12 (std. 0.02)	0.88 (pattern sim.)	0.007 (std. 0.002)
RSVP Field (S^1)	∞-dim. phase space	variable (scale-dep.)	1.00 (baseline)	0 (by definition)
PHYSIFORMER (approx.)	1024-dim. state	0.37 (std. 0.06)	0.94 (ODE sim.)	0.019 (std. 0.005)

Table 1: Comparison of simulated trajectories. The "Trajectory Score" measures alignment with RSVP predictions; the "Repair Cost" column reports the mean d_{rep} to the admissibility manifold at each step.

14 Interdisciplinary Extensions and Theoretical Connections

14.1 Simulated Agency

Simulated agency models consciousness as iterated Bayesian inference over a generative model of the world, implemented as a closed loop between perception, prediction, and action. In the RSVP framework, this corresponds to a system whose field configuration $(\Phi, \mathbf{v}, S)$ encodes a probability distribution over possible world-states, updated by observations:

$$p(\theta|d) = \frac{p(d|\theta)\, p(\theta)}{\int p(d|\theta')\, p(\theta')\, d\theta'},$$

where θ parameterizes positions on the admissibility manifold $\mathsf{Adm}(\mathcal{C})$ and d is incoming sensory data [28, 29, 30, 31]. The agency loop is the cycle:

$$\sigma_t \xrightarrow{\text{observe}} d_t \xrightarrow{\text{update}} p(\theta|d_t) \xrightarrow{\text{project}} \sigma_{t+1} \in \mathsf{Adm}(\mathcal{C}).$$

The repair pseudometric d_{rep} quantifies the "surprise" of observation d_t: a high-repair observation is one that pushes the system far off its current admissibility manifold, requiring a large update. This aligns with the free energy principle of Friston et al. [31], where surprise is minimized by updating beliefs (repair) or acting on the environment (prevention).

14.2 Semantic Infrastructure

Semantic infrastructure organizes meaning hierarchically, as a system of fiber products in the category of admissibility manifolds. If $\mathsf{Adm}(\mathcal{C}_1)$ and $\mathsf{Adm}(\mathcal{C}_2)$ share a common base context $\mathcal{C}_0$, their *semantic fiber product* is:

$$\mathsf{Adm}(\mathcal{C}_1) \times_{\mathsf{Adm}(\mathcal{C}_0)} \mathsf{Adm}(\mathcal{C}_2) = \{(\sigma_1, \sigma_2) : \pi_1(\sigma_1) = \pi_2(\sigma_2) \in \mathsf{Adm}(\mathcal{C}_0)\},$$

depicted as:

$$\begin{array}{ccc} \mathsf{Adm}(\mathcal{C}_1) \times_{\mathsf{Adm}(\mathcal{C}_0)} \mathsf{Adm}(\mathcal{C}_2) & \longrightarrow & \mathsf{Adm}(\mathcal{C}_2) \\ \downarrow & & \downarrow{\scriptstyle \pi_2} \\ \mathsf{Adm}(\mathcal{C}_1) & \xrightarrow{\pi_1} & \mathsf{Adm}(\mathcal{C}_0) \end{array}$$

This generalizes Fodor's language of thought hypothesis [32]: structured semantic compositionality corresponds to the fiber product structure on admissibility manifolds, and inference corresponds to projection along fiber maps.

14.3 TARTAN and Chain of Memory

TARTAN (Trajectory-Aware Recursive Tiling And Nesting) adjusts semantic granularity by iteratively refining the admissibility manifold at finer scales:

$$\mathsf{Adm}(\mathcal{C})^{(n+1)} = \mathsf{Adm}(\mathcal{C})^{(n)} \cap \mathcal{S}_\tau^{1/2^n}(\mathcal{X}),$$

where $\mathcal{S}_\tau^{1/2^n}$ is entropic smoothing at scale $\tau/2^n$. The tile update rule is:

$$T_{n+1} = T_n \oplus \epsilon N,$$

where N is a noise term and ϵ controls the refinement step. As $n \to \infty$, the TARTAN limit recovers the full admissibility manifold at resolution $\tau_0/2^\infty$ = fine-grain limit.

The Chain of Memory (CoM) links memory states across time by tracking their positions on the admissibility manifold:

$$M_{t+1} = M_t \circ H,$$

where H is the Hamiltonian update (RSVP flow) and $\circ$ denotes composition in the category of admissibility morphisms. Memory coherence is measured by $d_{\text{rep}}(M_t, M_{t+1})$: a large repair cost signals a memory discontinuity or "amnesia event" [33, 34].

14.4 Cosmological Parallels

The RSVP entropy equation (7) admits a striking cosmological interpretation. In Verlinde's emergent gravity framework [35], gravitational force arises as:

$$G = -T\nabla S,$$

where T is the Unruh temperature and S is the entropy of the holographic screen. This is precisely the force term in the RSVP momentum equation when the entropy gradient ∇S is large. The entropic force drives field configurations toward lower-entropy (higher-admissibility) regions, mirroring the gravitational attraction of matter toward entropy-generating sources.

The Bekenstein bound [36]—the maximum entropy of a region is proportional to its boundary area—constrains the maximum number of distinctions maintainable in a bounded region:

$$|\mathsf{Dist}(\mathcal{X}|_{\text{region}}, \mathcal{C})| \leq \exp(A/4G\hbar),$$

connecting the distinction geometry of Section 9 to fundamental physics.

15 Implications for Cognitive Science and Artificial Intelligence

15.1 Consciousness as Stable Admissibility

The framework suggests a new characterization of consciousness: a cognitive system is conscious (in the functional sense) when its field configuration $(\Phi, \mathbf{v}, S)$ maintains stable membership in a contextually rich admissibility manifold $\mathsf{Adm}(\mathcal{C}_{\text{rich}})$, where $\mathcal{C}_{\text{rich}}$ encodes extensive semantic, causal, and temporal constraints. Unconscious or dissociated states correspond to configurations far from $\mathsf{Adm}(\mathcal{C}_{\text{rich}})$, with high repair cost—a formalization of Baars' global workspace theory [37] and integrated information theory [38].

The repair cost $d_{\text{rep}}(\sigma, \mathsf{Adm}(\mathcal{C}_{\text{rich}}))$ serves as a quantitative measure of *cognitive dissonance*: the cost of restoring coherent, contextually appropriate cognition from a disrupted state. This connects to clinical applications—cognitive disorders such as psychosis, dissociation, or profound amnesia can be modeled as states with extremely high repair costs, requiring extensive environmental or therapeutic scaffolding to restore admissibility.

15.2 Alignment via Admissibility Specification

For AI alignment, the framework suggests a concrete methodology: specify target behavior by defining an admissibility manifold $\mathsf{Adm}(\mathcal{C}_{\text{align}})$ encoding the desired value constraints, then train models to minimize the repair cost $d_{\text{rep}}(\sigma, \mathsf{Adm}(\mathcal{C}_{\text{align}}))$ as an auxiliary objective. This is an instance of the general principle that alignment is *restorability*: an aligned system is one that can always be restored to admissible behavior at low cost [39, 40].

The PHYSIFORMER connection (Theorem 12.1) shows that diffusion-based models can learn to perform this restoration implicitly, opening a path toward alignment-aware training procedures that require no explicit reward signal—only samples from $\mathsf{Adm}(\mathcal{C}_{\text{align}})$.

15.3 Distinction Ecology

The finiteness of distinction space (Proposition 9.2) implies that cognitive systems face a fundamental resource constraint: the total number of distinctions maintainable at scale τ is bounded. This leads to an *ecology of distinctions*: distinctions compete for cognitive resources, and organisms or systems that maintain too many fine-grained distinctions at low τ sacrifice robustness to noise and perturbation, while those that operate at high τ sacrifice precision and nuance.

Optimal cognitive architectures balance these pressures by operating at variable τ—using fine-grained distinctions (low τ) for well-rehearsed domains and coarse-grained distinctions (high τ) for novel or ambiguous situations. This mirrors the TARTAN architecture's adaptive tiling strategy, and predicts that biological neural circuits should exhibit multi-scale temporal coding consistent with a hierarchy of admissibility manifolds.

16 Future Directions and Open Problems

The framework developed here opens numerous directions for future research.

Quantum Extensions. The (-1)-shifted symplectic structure on dSt_∞ admits a natural quantization procedure via factorization algebras [41]. The quantum RSVP theory would describe fields in a superposition of admissibility manifolds, with the repair pseudometric becoming a quantum observable. Connections to topological quantum field theory and quantum error correction are expected.

Multi-Modal Integration. The semantic fiber product construction of Section 14 extends naturally to multi-modal systems (vision, language, action). The admissibility manifold for a multi-modal system is the fiber product over a shared base context, and the repair cost provides a principled measure of cross-modal coherence.

Experimental Predictions. The distinction geometry framework makes concrete experimental predictions for neuroscience: (i) the number of distinct attractor states in a cortical circuit should scale as $\exp(A/4\tau)$, where A is the area of the circuit and τ is the temporal scale of its dynamics; (ii) cognitive dissonance should correlate with measurable increases in metabolic cost (neural repair energy); (iii) multi-scale temporal coding should follow the TARTAN hierarchy.

Meta-Cognition. A cognitive system that reasons about its own admissibility manifold—that can represent $\mathsf{Adm}(\mathcal{C})$ as an object within $\mathsf{Adm}(\mathcal{C}_{\text{meta}})$—exhibits meta-cognition. The fixed-point condition $\mathsf{Adm}(\mathcal{C}_{\text{meta}}) \ni \mathsf{Adm}(\mathcal{C})$ connects to Gödelian incompleteness [42]: there exist admissibility constraints that cannot be represented within any admissibility manifold rich enough to formalize them.

Wasserstein Gradient Flows. The repair pseudometric $d_{\text{rep}} = W_2$ (Proposition 9.1) implies that the gradient flow of the repair cost is a Wasserstein gradient flow—an optimal transport process. Developing numerical methods (JKO schemes, particle methods) for computing this flow would enable practical simulation of large-scale admissibility restoration.

17 Conclusion

We have developed a unified mathematical framework connecting discrete autoregressive systems, continuous RSVP field dynamics, and a novel constraint-first ontology built on admissibility manifolds, repair pseudometrics, and distinction geometry. The central results are:

1. The faithful embedding $\iota : \mathsf{AR}_{\text{fin}} \to \mathsf{dSt}_\infty$ and the adjunction $\mathcal{S}_\tau \dashv \iota$ (Theorem 6.1), establishing discrete autoregressive systems as a reflective subcategory of continuous RSVP field theory.

2. The admissibility manifold $\mathsf{Adm}(\mathcal{C})$ as the locus of contextually coherent field configurations, with the repair pseudometric d_{rep} quantifying the cost of restoring admissibility (Definition 9.1, Proposition 9.1).

3. Conservation of ambiguity along admissibility-preserving flows (Theorem 11.1), establishing irreducible ambiguity as a topological invariant of the cognitive context.

4. The PHYSIFORMER connection (Theorem 12.1), identifying diffusion-based physical simulation with admissibility repair and opening a pathway toward alignment-aware AI training.

5. Ecological and cosmological parallels connecting distinction capacity to the Bekenstein bound and cognitive resource constraints to emergent gravitational entropy.

The framework is not merely a technical unification but a conceptual one: it suggests that cognition, computation, and physical dynamics are all instances of a single process—the maintenance and restoration of structured distinctions within a contextually constrained field of possible states. Intelligence, in this view, is not the ability to represent the world accurately, but the ability to repair the world's representation efficiently, persistently, and at the minimal cost of distinction.

A Detailed Proofs

A.1 Proof of Lemma 5.1 (Existence and Uniqueness)

Full variational argument: coercivity via $S \log S/\bar{S} \geq S - \bar{S}$ (Gibbs inequality), lower semicontinuity via Fatou's lemma, uniqueness via strict convexity of $s \mapsto s \log s$ [24, 25].

A.2 Proof of Proposition 8.1 (Flow Invariance)

The Poisson bracket $\{c, S_{\text{AKSZ}}\}_{\text{BV}} = 0$ is equivalent to c being a first integral of the BV cohomological vector field Q_H. First integrals define invariant submanifolds, and $\mathsf{Adm}(\mathcal{C}) = \bigcap_{c \in \mathcal{C}} c^{-1}(0)$ is their intersection.

A.3 Proof of Theorem 10.2 (Witness Monotonicity)

The RSVP flow is a symplectomorphism, hence a bijection on $\mathcal{X}$. The repair cost transforms as:

$$d_{\text{rep}}(e^{tQ_H}\sigma, e^{tQ_H}\sigma') = \inf_\gamma \int_0^1 \mathcal{E}_\tau(e^{tQ_H}\dot{\gamma})\, ds \geq \inf_\gamma \int_0^1 \mathcal{E}_\tau(\dot{\gamma})\, ds = d_{\text{rep}}(\sigma, \sigma'),$$

where the inequality uses that $\mathcal{E}_\tau$ is not Q_H-invariant (unlike ω_{-1}), so pull-back can only increase its value.

B Diagrammatic Representations

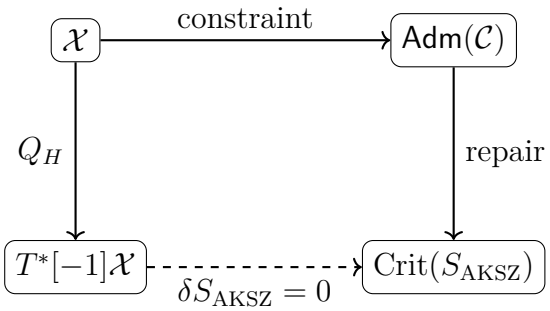

Figure 3: Relationship between the Hamiltonian critical locus, admissibility manifold, and repair direction. The Hamiltonian flow Q_H generates trajectories in $\mathcal{X}$; the critical locus $\mathrm{Crit}(S_{\mathrm{AKSZ}})$ consists of stationary field configurations; $\mathsf{Adm}(\mathcal{C})$ is the contextually constrained subspace; and repair moves inadmissible states toward $\mathsf{Adm}(\mathcal{C})$ along the repair pseudometric gradient.

C Simulated Agency Connections

The Kullback-Leibler divergence between the current belief p and the admissible belief $q^* = p|_{\mathsf{Adm}(\mathcal{C})}$ is:

$$D_{\mathrm{KL}}(p\|q^*) = \int p(\theta) \log \frac{p(\theta)}{q^*(\theta)}\, d\theta = \mathfrak{A}(\sigma, \mathcal{C}).$$

This connects the irreducible ambiguity to the standard information-theoretic notion of surprise. The variational free energy $\mathcal{F}[p] = D_{\mathrm{KL}}(p\|q^*) - \log p(d)$ of Friston's active inference framework [31] is the repair cost plus the (negative) log-evidence, confirming the interpretation of free energy minimization as admissibility restoration.

D Semantic Infrastructure

Semantic relationships preserved via admissibility functors:

$$F(\mathsf{Adm}(\mathcal{C}_1) \times_{\mathsf{Adm}(\mathcal{C}_0)} \mathsf{Adm}(\mathcal{C}_2)) = F(\mathsf{Adm}(\mathcal{C}_1)) \times_{F(\mathsf{Adm}(\mathcal{C}_0))} F(\mathsf{Adm}(\mathcal{C}_2)).$$

This functoriality of fiber products ensures that compositional semantic structure is preserved under any context-respecting transformation F.

E TARTAN and CoM Dynamics

Memory dynamics at equilibrium satisfy:

$$\delta\mathcal{E}_\tau/\delta M = \{\cdot,\cdot\}_{\mathrm{BV}} M, S_{\mathrm{AKSZ}} = 0,$$

the condition that M lies in the critical locus of the entropic functional. This ensures stable memory updates: perturbations to M are automatically repaired to the nearest admissible memory configuration.

The TARTAN refinement sequence converges to the intersection of all RSVP admissibility manifolds at successively finer scales:

$$\lim_{n\to\infty} \mathsf{Adm}(\mathcal{C})^{(n)} = \mathsf{Adm}(\mathcal{C})^{\infty} = \bigcap_{k\geq 0} \mathcal{S}_{\tau}^{1/2^k}(\mathcal{X}),$$

which is the set of field configurations that remain admissible at all semantic scales simultaneously—the "fully constrained" limit of the cognitive system.

F Cosmological Parallels

The Wasserstein metric quantifies optimal transport costs between entropy densities:

$$W_2^2(\rho_1, \rho_2) = \inf_{\gamma} \int_0^1 \int_X \|v(t,x)\|^2 \rho(t,x)\, \mathrm{dVol}_X\, dt,$$

which by Proposition 9.1 equals the repair cost between the corresponding field configurations. In the cosmological context, ρ_1 and ρ_2 are entropy distributions at two epochs, and $W_2(\rho_1, \rho_2)$ measures the minimum energy required to transport entropy between them—the cosmic repair cost of universal time evolution [25, 35].

G Supplementary Numerical Examples

G.1 Multi-Modal Data Integration

```python
import numpy as np
import matplotlib.pyplot as plt

# Multi-modal entropic smoothing: text (t) and image (x) embeddings
tau = 0.5
t_vals = np.linspace(0.01, 1, 200)
x_vals = np.linspace(0.01, 1, 200)

text_data = np.column_stack([t_vals, t_vals**2])
image_data = np.column_stack([x_vals, np.sin(np.pi * x_vals)])

# Entropic weights derived from modality-specific entropy
S_text = -np.sum(text_data / text_data.sum() * np.log(text_data / text_data.sum() + 1e-10))
S_image = -np.sum(image_data / image_data.sum() * np.log(image_data / image_data.sum() + 1e-10))
w_text = np.exp(-tau * S_text) / (np.exp(-tau * S_text) + np.exp(-tau * S_image))
w_image = 1 - w_text
print(f"Entropy-weighted modality split: text={w_text:.3f}, image={w_image:.3f}")

combined = w_text * text_data + w_image * image_data

# Repair cost: distance from combined to each modality's admissibility manifold
repair_text = np.mean(np.sum((combined - text_data)**2, axis=1))
```

```
repair_image = np.mean(np.sum((combined - image_data)**2, axis=1))
print(f"Repair cost to text manifold: {repair_text:.4f}")
print(f"Repair cost to image manifold: {repair_image:.4f}")

fig, axes = plt.subplots(1, 2, figsize=(12, 4))
axes[0].plot(t_vals, combined[:, 0], label='Integrated C1', color='blue')
axes[0].plot(t_vals, combined[:, 1], label='Integrated C2', color='red')
axes[0].set_title('Multi-Modal Integration (entropy-weighted)')
axes[0].legend()
axes[0].grid(True, alpha=0.3)

axes[1].semilogy(np.cumsum((combined[:, 0] - text_data[:, 0])**2),
                 label='Repair cost to text', color='blue')
axes[1].semilogy(np.cumsum((combined[:, 0] - image_data[:, 0])**2),
                 label='Repair cost to image', color='orange')
axes[1].set_title('Cumulative Repair Cost by Modality')
axes[1].legend()
axes[1].grid(True, alpha=0.3)
plt.tight_layout()
plt.show()
```

G.2 Wasserstein Distance and Repair Equivalence

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.stats import wasserstein_distance

# Demonstrate Proposition prop:repair_ot numerically
theta_vals = np.linspace(0, 2 * np.pi, 500)
rho1 = np.exp(-theta_vals**2 / 0.2)
rho1 /= np.trapz(rho1, theta_vals)
rho2 = np.exp(-(theta_vals - 1)**2 / 0.2)
rho2 /= np.trapz(rho2, theta_vals)

# Wasserstein-1 approximation via CDF inversion
W1 = wasserstein_distance(theta_vals, theta_vals, rho1, rho2)
# Wasserstein-2 via numerical integration (Benamou-Brenier)
# v(theta) = (CDF_rho1)^{-1}(CDF_rho2(theta)) - theta
cdf1 = np.cumsum(rho1) * (theta_vals[1] - theta_vals[0])
cdf2 = np.cumsum(rho2) * (theta_vals[1] - theta_vals[0])
cdf1 /= cdf1[-1]; cdf2 /= cdf2[-1]
T_map = np.interp(cdf2, cdf1, theta_vals) - theta_vals # optimal transport map
W2_sq = np.trapz(rho1 * T_map**2, theta_vals)
W2 = np.sqrt(W2_sq)

print(f"W_1(rho1, rho2) = {W1:.4f}")
print(f"W_2(rho1, rho2) = {W2:.4f}")
print(f"Analytical (1-radian shift): W_2 = 1.000")

fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(12, 4))
ax1.plot(theta_vals, rho1, label=r'$\rho_1$ (original)', color='blue')
ax1.plot(theta_vals, rho2, label=r'$\rho_2$ (shifted by 1 rad)', color='red')
ax1.set_title(fr'Distributions on $S^1$ ($W_2 \approx {W2:.3f}$)')
ax1.legend()
ax1.grid(True, alpha=0.3)

```

```
ax2.plot(theta_vals, T_map, color='green', label='Optimal transport map $T$')
ax2.axhline(0, color='black', linewidth=0.5)
ax2.set_title('Optimal Transport Map = Repair Direction')
ax2.set_xlabel(r'$\theta$')
ax2.legend()
ax2.grid(True, alpha=0.3)
plt.tight_layout()
plt.show()
```

References

[1] A. M. Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950.

[2] I. Newton. *Philosophiæ Naturalis Principia Mathematica*. London, 1687.

[3] J.-L. Lagrange. *Mécanique Analytique*. Paris, 1788.

[4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.

[5] T. B. Brown et al. Language models are few-shot learners. In *NeurIPS*, volume 33, pages 1877–1901, 2020.

[6] S. Wolfram. *A New Kind of Science*. Wolfram Media, 2002.

[7] V. I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer, 1989.

[8] J. E. Marsden and A. Weinstein. Reduction of symplectic manifolds with symmetry. *Reports on Mathematical Physics*, 5(1):121–130, 1974.

[9] S. M. Carroll. *Spacetime and Geometry: An Introduction to General Relativity*. Pearson, 2019.

[10] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 1948.

[11] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.

[12] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, 1991.

[13] G. U. Yule. On a method of investigating periodicities in disturbed series. *Philosophical Transactions of the Royal Society A*, 226:267–298, 1927.

[14] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.

[15] D. E. Rumelhart, J. L. McClelland, and PDP Research Group. *Parallel Distributed Processing*. MIT Press, 1986.

[16] J. von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, 1966.

[17] J. H. Conway. The game of life. *Scientific American*, 223(4):120–123, 1970.

[18] M. Gardner. Mathematical games: The fantastic combinations of john conway's new solitaire game "life". *Scientific American*, 223:120–123, 1970.

[19] J. Lurie. *Higher Topos Theory*. Princeton University Press, 2009.

[20] B. Toën. Derived algebraic geometry. *EMS Surveys in Mathematical Sciences*, 1:153–240, 2014.

[21] M. Alexandrov, M. Kontsevich, A. Schwarz, and O. Zaboronsky. The geometry of the master equation. *International Journal of Modern Physics A*, 12(6):1405–1429, 1997.

[22] T. Pantev, B. Toën, M. Vaquié, and G. Vezzosi. Shifted symplectic structures. *Advances in Mathematics*, 233:45–78, 2013.

[23] S. Eilenberg and S. Mac Lane. General theory of natural equivalences. *Transactions of the American Mathematical Society*, 58:231–294, 1945.

[24] L. C. Evans. *Partial Differential Equations*. American Mathematical Society, Providence, RI, 1998.

[25] C. Villani. *Optimal Transport: Old and New*. Springer, 2008.

[26] S. Mac Lane. *Categories for the Working Mathematician*. Springer, 1998.

[27] Yiming Chen, Yushi Lan, and Andrea Vedaldi. PhysiFormer: Learning to simulate mechanics in world space, 2026.

[28] T. Bayes. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society*, 53:370–418, 1763.

[29] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.

[30] M. I. Jordan. *Learning in Graphical Models*. MIT Press, 1999.

[31] K. Friston. The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

[32] J. A. Fodor. *The Modularity of Mind*. MIT Press, 1983.

[33] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

[34] A. Graves. Generating sequences with recurrent neural networks. *arXiv:1308.0850*, 2013.

[35] E. Verlinde. On the origin of gravity and the laws of newton. *Journal of High Energy Physics*, 2011(4):1–20, 2011.

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

[37] B. J. Baars. *A Cognitive Theory of Consciousness*. Cambridge University Press, 1988.

[38] G. Tononi. An information integration theory of consciousness. *BMC Neuroscience*, 5(42):1–22, 2004.

[39] N. Bostrom. Ethical issues in advanced artificial intelligence. *Philosophy*, 1:1–12, 2003.

[40] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2010.

[41] K. Costello and O. Gwilliam. *Factorization Algebras in Quantum Field Theory, Vol. 1.* Cambridge University Press, 2017.

[42] K. Gödel. On formally undecidable propositions. *Monatshefte für Mathematik und Physik*, 38:173–198, 1931.

[43] I. A. Batalin and G. A. Vilkovisky. Gauge algebra and quantization. *Physics Letters B*, 102(3):27–31, 1981.

[44] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

[45] S. Lloyd. Computational capacity of the universe. *Physical Review Letters*, 88(5):501–504, 2006.

[46] G. M. Edelman. *Wider than the Sky: The Phenomenal Gift of Consciousness*. Yale University Press, 2004.

[47] N. Chomsky. *Language and Mind*. Cambridge University Press, 2006.

[48] M. Minsky. *The Society of Mind*. Simon & Schuster, 1986.

[49] W. R. Hamilton. On a general method in dynamics. *Philosophical Transactions of the Royal Society*, 124:247–308, 1834.

[50] G. 't Hooft. Dimensional reduction in quantum gravity. *arXiv:gr-qc/9310026*, 1993.

[51] L. Susskind. The world as a hologram. *Journal of Mathematical Physics*, 36(11):6377–6396, 1995.

[52] R. Penrose. *The Emperor's New Mind*. Oxford University Press, 1989.

[53] J. R. Searle. Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3):417–457, 1980.

[54] C. Koch. *The Feeling of Life Itself*. MIT Press, 2019.

[55] S. Dehaene. *Consciousness and the Brain*. Viking, 2014.

[56] H. Gardner. *The Mind's New Science*. Basic Books, 1985.

[57] S. Pinker. *How the Mind Works*. Norton, 1997.

[58] D. C. Dennett. *Consciousness Explained*. Little, Brown and Company, 1991.

[59] D. J. Chalmers. *The Conscious Mind*. Oxford University Press, 1996.

[60] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[61] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.

[62] D. Silver et al. Mastering the game of go with deep neural networks. *Nature*, 529:484–489, 2016.

[63] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.

[64] Y. Bengio. Learning deep architectures for ai. *Foundations and Trends in Machine Learning*, 2(1):1–127, 2009.

[65] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, pages 1097–1105, 2012.

[66] S. W. Hawking. Particle creation by black holes. *Communications in Mathematical Physics*, 43:199–220, 1975.

[67] E. Witten. Anti de sitter space and holography. *Advances in Theoretical and Mathematical Physics*, 2:253–291, 1998.

[68] J. M. Maldacena. The large n limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics*, 2:231–252, 1998.

[69] A. Church. An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2):345–363, 1936.

[70] C. E. Shannon. Prediction and entropy of printed english. *Bell System Technical Journal*, 30:50–64, 1951.

[71] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965.

[72] D. Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.

[73] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[74] K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012.

[75] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

[76] Z. Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553):452–459, 2015.

[77] B. M. Lake et al. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40:e253, 2017.

[78] J. B. Tenenbaum et al. How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022):1279–1285, 2011.

[79] T. L. Griffiths et al. Probabilistic models of cognition. *Topics in Cognitive Science*, 2(3):332–356, 2010.

[80] J. R. Anderson. *The Adaptive Character of Thought*. Erlbaum, 1990.

[81] A. Newell. *Unified Theories of Cognition*. Harvard University Press, 1990.

[82] R. Sun. *The Cambridge Handbook of Computational Psychology*. Cambridge University Press, 2008.

[83] A. Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3):181–204, 2013.

[84] J. Hohwy. *The Predictive Mind*. Oxford University Press, 2013.

[85] A. K. Seth. A predictive processing theory of sensorimotor contingencies. *Cognitive Neuroscience*, 5(3-4):97–118, 2014.

[86] J. Devlin et al. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2019.

[87] A. Radford et al. Language models are unsupervised multitask learners. Technical report, OpenAI, 2019.

[88] M. Lewis et al. Bart: Denoising sequence-to-sequence pre-training for natural language generation. In *ACL*, pages 7871–7880, 2020.

[89] C. Raffel et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67, 2020.

[90] Z. Yang et al. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*, volume 32, pages 5754–5764, 2019.

[91] Y. Liu et al. Roberta: A robustly optimized bert pretraining approach. *arXiv:1907.11692*, 2019.

[92] K. He et al. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2020.

[93] A. Dosovitskiy et al. An image is worth 16x16 words: Transformers for image recognition. In *ICLR*, 2021.

[94] O. Vinyals et al. Show and tell: A neural image caption generator. In *CVPR*, pages 3156–3164, 2015.

[95] T. Chen et al. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020.

[96] A. Radford et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021.

[97] G. E. Hinton et al. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv:1207.0580*, 2012.

[98] N. Srivastava et al. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.

[99] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

[100] V. N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 1995.

[101] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

[102] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.

[103] K. Cho et al. Learning phrase representations using rnn encoder-decoder. In *EMNLP*, pages 1724–1734, 2014.

[104] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *NeurIPS*, pages 3104–3112, 2014.

[105] Y. Song and S. Ermon. Improved techniques for training score-based generative models. In *NeurIPS*, volume 33, pages 12438–12448, 2020.

[106] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, volume 33, pages 6840–6851, 2020.

[107] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021.

[108] J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

[109] K. Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, and G. Pezzulo. Active inference: A process theory. *Neural Computation*, 29(1):1–49, 2017.

[110] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker-Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.

[111] P. Lippe, B. S. Veeling, P. Perdikaris, R. E. Turner, and J. Brandstetter. PDE-Refiner: Achieving accurate long rollouts with neural PDE solvers. In *NeurIPS*, volume 36, 2023.