# RSVP–Polyxan: A Unified Field Theory of Semantic Hyperstructures

Flyxion

Independent Researcher

2026

**Abstract**

This paper develops a unified theoretical framework coupling the Relativistic Scalar–Vector Plenum (RSVP) field system, defined on a semantic manifold of embeddings, with the Polyxan hypermedia substrate, a typed, bidirectionally linked graph of content atoms and media quines in the Xanadu tradition. We construct an action functional for RSVP fields over semantic space, derive Euler–Lagrange equations, and introduce a geometric interpretation of polysemantic clusters. The framework is then extended from a static semantic field theory into a dynamical world model with persistent latent memory, projection-based observation, and constrained field evolution. Constraint projection is internalized into the variational structure via Lagrange multipliers, yielding a principled dynamics in which neural components act as generalized forcing terms rather than unconstrained updates. Noether's theorem applied to the augmented action produces semantic conservation laws for coherence, directional agency, and causal residue. The system is then reformulated in Hamiltonian phase-space terms, equipped with gauge symmetries corresponding to semantically equivalent representations, and finally extended to a stochastic path-integral formulation that subsumes diffusion generative models as a special case. A sheaf-theoretic treatment of galaxy-local worldviews is refined into a derived sheaf valued in $\infty$-groupoids, and the Polyxan category is lifted to an $\infty$-category in which semantic inconsistency is characterized by non-vanishing cohomology. The framework is then brought into precise contact with MultiGen, a concrete diffusion-based generative game engine, which is shown to be a degenerate projection of RSVP–Polyxan obtained by collapsing the full latent field to a static scalar substrate and eliminating the write operator. The factored update $\Pi_{\mathcal{C}} \circ \mathcal{N}_{\psi}$ is specified with six explicit field-level constraint equations governing mass consistency, energy budgets, topological admissibility, causal write locality, entropy accounting, and identity persistence. Fast and slow update cadences, event-typed operator decomposition, read-write cycle closure, and a four-phase training curriculum complete the operational architecture. Multiplayer interaction is characterized as causal field synchronization over a shared substrate, with concurrent edits resolved by constraint projection rather than consensus protocol. Throughout, we present proofs, lemmas, and structural invariants that guarantee coherence between continuous semantic fields and discrete hypergraphs.

# Contents


**1 Introduction** **3**

**2 Semantic Manifold and RSVP Fields** **3**

**3 RSVP Lagrangian Formulation** **3**

**4 Coupling RSVP to the Polyxan Hypergraph** **4**

4.1 Topological Curvature from Links . . . . . . . . . . 4

**5 Latent Field Memory and World-State Representation** **4**

5.1 State Decomposition . . . . . . . . . . 5

**6 Observation as Projection Operator** **5**

6.1 Observation Does Not Store State . . . . . . . . . . 5

**7 Dynamics as Constrained Field Evolution** **6**

7.1 Constraint Manifold . . . . . . . . . . 6

7.2 Invention-to-Memory Principle . . . . . . . . . . 6

**8 Variational Constrained Field Evolution** **6**

8.1 Augmented Action Functional . . . . . . . . . . 7

8.2 Euler–Lagrange Equations with Constraints . . . . . . . . . . 7

8.3 Neural Proposal as Generalized Force . . . . . . . . . . 7

**9 Noether Theorem for Semantic Field Invariants** **7**

9.1 Symmetry Transformations . . . . . . . . . . 8

9.2 Conserved Current . . . . . . . . . . 8

9.3 Examples of Semantic Invariants . . . . . . . . . . 8

**10 Hamiltonian Formulation of RSVP Memory Dynamics** **9**

10.1 Canonical Variables . . . . . . . . . . 9
10.2 Hamiltonian Functional . . . . . . . . . . 9
10.3 Hamilton's Equations . . . . . . . . . . 10
10.4 Symplectic Structure . . . . . . . . . . 10

**11 Renormalization and Scale Separation in Semantic Fields** **10**
11.1 Renormalization Operator . . . . . . . . . . 10
11.2 Scale-Dependent Dynamics . . . . . . . . . . 11
11.3 Semantic Interpretation and Limiting Cases . . . . . . . . . . 11

**12 Polyxan as a Discrete Interface to Field Memory** **11**
12.1 Bidirectional Coupling . . . . . . . . . . 12

**13 Gauge Symmetries and Semantic Equivalence** **12**
13.1 Gauge Transformations . . . . . . . . . . 12
13.2 Examples of Gauge Groups . . . . . . . . . . 12
13.3 Gauge Invariance of the Action . . . . . . . . . . 13
13.4 Connection and Covariant Derivative . . . . . . . . . . 13

**14 Sheaf-Theoretic Interpretation** **13**

**15 Derived and $\infty$-Categorical Structure of Polyxan** **13**
15.1 Derived Field Space . . . . . . . . . . 14
15.2 Obstruction Theory . . . . . . . . . . 14
15.3 Polyxan as an $\infty$-Category . . . . . . . . . . 14
15.4 Functor to Derived Fields . . . . . . . . . . 14
15.5 Homotopy-Coherent Gluing . . . . . . . . . . 14
15.6 Interpretation of Hallucination . . . . . . . . . . 15
15.7 Semantic Fixed Points . . . . . . . . . . 15

**16 Category-Theoretic Architecture** **15**

**17 Quantization of RSVP Field Dynamics** **16**

17.1 Configuration Space . . . . . 16

17.2 Path Integral . . . . . 16

17.3 Stochastic Diffusion Limit . . . . . 16

17.4 Connection to Diffusion Generative Models . . . . . 16

**18 Information Geometry of RSVP Field Space** **17**

18.1 Statistical Manifold . . . . . 17

18.2 Natural Gradient Flow . . . . . 17

18.3 Entropy as Local Curvature . . . . . 17

**19 Existence and Stability of RSVP Field Dynamics** **18**

19.1 Existence of Solutions . . . . . 18

19.2 Convergence Toward Stable Configurations . . . . . 18

19.3 Observability and Identifiability . . . . . 19

**20 Stability and Energy Minimization** **19**

**21 Global Reset as Field Reconfiguration** **19**

**22 MultiGen as a Degenerate Projection of RSVP** **20**

22.1 The Projection Map . . . . . 20

22.2 Field Correspondences . . . . . 20

22.3 The Missing Write Operator . . . . . 21

22.4 MultiGen as a Theorem . . . . . 21

22.5 Multiplayer as Evidence . . . . . 22

**23 Factored Update: Proposal, Projection, and Commit** **22**

23.1 The Three-Stage Update . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 22
23.2 Explicit Constraint Equations . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 23
23.3 Composable Projections . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 24

**24 Fast and Slow Dynamics, Event Operators, and Read-Write Closure 24**
24.1 Fast and Slow Field Update Cadences . . . . . . . . . . . . . . . . . . . . . . . . . 24
24.2 Event-Typed Update Operators . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 24
24.3 Read-Write Closure . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 24

**25 Multiplayer as Causal Field Synchronization and Training 25**
25.1 Multiplayer as Causal DAG Merging . . . . . . . . . . . . . . . . . . . . . . . . . . 25
25.2 Training Curriculum . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 25

**26 Computational Instantiation 26**

**Architecture Diagram 31**

**27 Conclusion 32**

# 1 Introduction

The RSVP–Polyxan framework fuses three mathematical layers. The first is a **semantic manifold** $X$ where conceptual embeddings live. The second is a **field theory** consisting of a scalar potential $\Phi$, vector flow $\mathbf{v}$, and entropy field $S$. The third is a **typed hypergraph** of content atoms and links, whose structure both sources and responds to RSVP fields.

The goal is to provide a unified variational, geometric, and categorical description of semantic information flow, and to promote this description into a full dynamical world engine with physics, memory, and rendering separation. The paper proceeds through three ascending layers of theoretical elaboration. At the classical layer we establish Lagrangian and Hamiltonian dynamics, constraint manifolds, and Noether invariants. At the stochastic layer we construct a path integral over field configurations and identify diffusion generative models as sampling trajectories from this integral. At the derived layer we lift the Polyxan category to an $\infty$-categorical structure, characterize semantic consistency as homotopy-coherent gluing, and interpret hallucination as a cohomological obstruction.

# 2 Semantic Manifold and RSVP Fields

Let $(X, g)$ be a smooth Riemannian manifold representing semantic space. Define three RSVP fields:

$$\Phi : X \to \mathbb{R}, \quad \mathbf{v} : X \to TX, \quad S : X \to \mathbb{R}.$$

The scalar field $\Phi$ encodes semantic potential or coherence density. The vector field $\mathbf{v}$ encodes directed agency or semantic drift. The scalar field $S$ encodes entropy, uncertainty, or morphic degeneracy. Field derivatives are defined using the Levi-Civita connection.

# 3 RSVP Lagrangian Formulation

We propose an action functional:

$$\mathcal{A}[\Phi, \mathbf{v}, S] = \int_X \mathcal{L}\, d\mu_g,$$

where the Lagrangian density incorporates kinetic, elastic, and potential terms:

$$\begin{aligned}
\mathcal{L} = {} & \frac{1}{2}(\partial_t \Phi)^2 - \frac{c_\Phi^2}{2}\|\nabla \Phi\|^2 \\
& + \frac{1}{2}\|\partial_t \mathbf{v}\|^2 - \frac{c_v^2}{2}\|\nabla \mathbf{v}\|^2 \\
& + \frac{1}{2}(\partial_t S)^2 - \frac{c_S^2}{2}\|\nabla S\|^2 \\
& - V(\Phi, \mathbf{v}, S; \rho, \kappa).
\end{aligned}$$

Here $\rho(x)$ denotes the node density induced by the Polyxan graph, and $\kappa(x)$ denotes the local curvature induced by typed links. Variation yields Euler–Lagrange equations:

$$\frac{\delta \mathcal{A}}{\delta \Phi} = \partial_t^2 \Phi - c_\Phi^2 \Delta \Phi - \frac{\partial V}{\partial \Phi} = 0,$$

and similarly for $\mathbf{v}$ and $S$.

# 4 Coupling RSVP to the Polyxan Hypergraph

Let $G = (V, E)$ be the Polyxan graph. Each node $i \in V$ is assigned an embedding $\mathbf{z}_i \in X$. Define discrete samples:

$$\Phi_i = \Phi(\mathbf{z}_i), \quad \mathbf{v}_i = \mathbf{v}(\mathbf{z}_i), \quad S_i = S(\mathbf{z}_i).$$

Define a discrete evolution law:

$$\frac{d\mathbf{z}_i}{dt} = -\alpha \nabla \Phi(\mathbf{z}_i) + \beta\, \mathbf{v}(\mathbf{z}_i) - \gamma \nabla S(\mathbf{z}_i).$$

Thus embeddings evolve in RSVP gradient flows.

## 4.1 Topological Curvature from Links

Define a link-curvature invariant:

$$\kappa_i = \sum_{j,k \sim i} f_{\text{tri}}(i, j, k),$$

counting typed 3-cycles weighted by link types. This curvature term shapes the potential $V$.

# 5 Latent Field Memory and World-State Representation

The classical RSVP formulation treats $(\Phi, \mathbf{v}, S)$ as fields over a semantic manifold $X$. However, this representation remains insufficient for persistent world modeling, as it does not explicitly encode revisitable state, appearance, or causal residue.

We therefore introduce a time-indexed latent field memory:

$$\mathcal{M}_t : X \to \mathbb{R}^d,$$

which encodes the full world-state as a structured field over semantic space. We decompose $\mathcal{M}_t$ into interpretable components:

$$\mathcal{M}_t(x) = \big(\Phi_t(x),\, \mathbf{v}_t(x),\, S_t(x),\, A_t(x),\, R_t(x),\, z_t(x)\big).$$

The first three components are the scalar semantic potential $\Phi_t$, the directed flow $\mathbf{v}_t$, and the entropy-uncertainty field $S_t$ already established by the RSVP formulation. The remaining three are new: $A_t$ encodes affordance structure, $R_t$ encodes causal residue accumulated through the system's history, and $z_t$ encodes latent appearance, the field-theoretic analogue of a rendered view. This decomposition generalizes both continuous RSVP fields and discrete Polyxan structures into a unified persistent substrate.

### 5.1 State Decomposition

We define the full system state:

$$X_t = (\mathcal{M}_t, p_t, h_t),$$

where $p_t \in TX$ is the agent trajectory coordinate and $h_t = (o_{t-L+1}, \dots, o_t)$ is a finite observation trace of length $L$. Crucially, $h_t$ is no longer the world state but merely a local perceptual history. The world-state resides entirely in $\mathcal{M}_t$, and observations are derived from it by projection.

## 6 Observation as Projection Operator

In classical generative models the world is implicitly encoded in the observation sequence. In contrast, we treat observation as a projection from latent field memory. Define the observation operator:

$$\mathcal{O}_\theta : (\mathcal{M}_t, p_t, h_t, a_t) \;\longrightarrow\; o_{t+1},$$

where $a_t$ is an action input. Observations are then sampled via:

$$o_{t+1} \sim p_\theta\big(o \mid \mathrm{Read}(\mathcal{M}_t, p_t), h_t, a_t\big).$$

The operator $\mathrm{Read}(\mathcal{M}_t, p_t)$ extracts a view-dependent slice of the field, analogous to a camera projection from a three-dimensional scene.

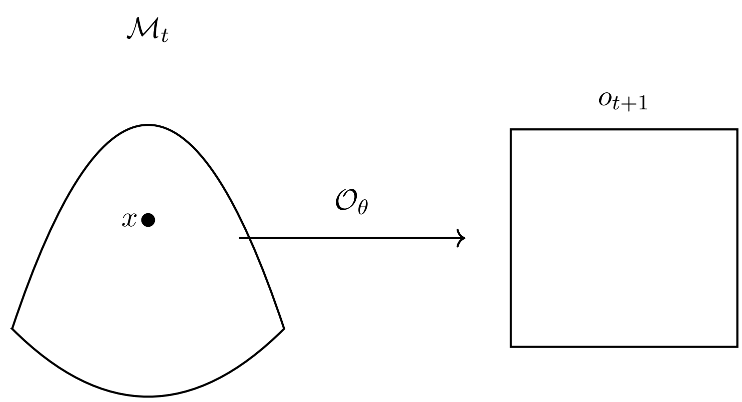


Figure 1: Observation as projection from latent field memory to visible state. The operator $\mathcal{O}_\theta$ maps a point in $\mathcal{M}_t$ to a rendered observation, introducing noise proportional to the local entropy $S_t(x)$.

### 6.1 Observation Does Not Store State

Unlike autoregressive systems, $\mathcal{O}_\theta$ does not maintain world-state; it merely resolves latent uncertainty into visible form.

**Principle 1.** *The world is stored in $\mathcal{M}_t$, not in the observation history.*

This separation is structurally fundamental. Long-range consistency, persistent object identity, and causal coherence are maintained not through the growth of an observation buffer but through the evolution of the latent field.

# 7 Dynamics as Constrained Field Evolution

We define a dynamics operator:

$$\Delta_\psi : (\mathcal{M}_t, p_t, a_t, o_{t+1}) \;\longrightarrow\; \delta\mathcal{M}_t,$$

which proposes updates to the latent field. The world-state evolves according to:

$$\mathcal{M}_{t+1} = \Pi_{\mathcal{C}}\big(\mathcal{M}_t + \Delta_\psi(\mathcal{M}_t, p_t, a_t, o_{t+1})\big),$$

where $\Pi_{\mathcal{C}}$ is a projection onto a constraint manifold $\mathcal{C}$.

## 7.1 Constraint Manifold

The admissible set $\mathcal{C}$ enforces a system of global invariants. Conservation constraints require that mass and energy be preserved across field updates. Topological consistency requires that the Polyxan hypergraph remain well-formed under induced modifications. Causal locality requires that updates propagate only within the future causal cone of the acting agent. Entropy monotonicity outside observation regions requires that $S_t$ not decrease in unobserved subsets of $X$, in accordance with the second law interpreted at the field level. The projection $\Pi_{\mathcal{C}}$ is therefore not a free deformation but a physically and semantically constrained retraction.

## 7.2 Invention-to-Memory Principle

Generation may invent details only once; after invention, details must be promoted into memory. Formally, if a region $x \in X$ is observed at time $t+1$, then:

$$z_{t+1}(x) \leftarrow z_{t+1}(x) + \lambda\, \hat{z}(x),$$

where $\hat{z}(x)$ is the inferred appearance extracted by $\mathcal{O}_\theta$. Thus appearance becomes persistent after first observation, and subsequent observations of the same region are constrained by the accumulated field value $z_t(x)$ rather than re-sampled from the prior.

# 8 Variational Constrained Field Evolution

The update rule of the previous section may be interpreted as a projection onto an admissible manifold of states. We now internalize this projection into a variational formulation, so that constraint enforcement arises continuously from the equations of motion rather than as a post-hoc correction.

### 8.1 Augmented Action Functional

Define an augmented action:

$$\mathcal{A}_{\text{aug}}[\mathcal{M},\lambda] = \int dt \int_X \Big(\mathcal{L}(\mathcal{M},\partial_t\mathcal{M}) + \langle \lambda, \mathcal{C}(\mathcal{M})\rangle\Big)\, d\mu_g,$$

where $\mathcal{L}$ extends the RSVP Lagrangian to all components of $\mathcal{M}_t$, the equation $\mathcal{C}(\mathcal{M}) = 0$ encodes constraint equations, and $\lambda$ is a field of Lagrange multipliers enforcing those constraints continuously throughout the evolution.

### 8.2 Euler–Lagrange Equations with Constraints

Variation of $\mathcal{A}_{\text{aug}}$ with respect to $\mathcal{M}$ yields:

$$\frac{\delta \mathcal{A}_{\text{aug}}}{\delta \mathcal{M}} = \frac{\partial \mathcal{L}}{\partial \mathcal{M}} - \partial_t \left(\frac{\partial \mathcal{L}}{\partial(\partial_t \mathcal{M})}\right) + \lambda \cdot \frac{\partial \mathcal{C}}{\partial \mathcal{M}} = 0,$$

together with the constraint equation $\mathcal{C}(\mathcal{M}) = 0$. The projection $\Pi_{\mathcal{C}}$ is thereby realized as a continuous constraint force in the equations of motion rather than a discrete retraction applied at each time step.

### 8.3 Neural Proposal as Generalized Force

We reinterpret the learned update $\Delta_\psi$ as an external forcing term in the equations of motion:

$$\partial_t^2 \mathcal{M} = -\nabla_{\mathcal{M}} E(\mathcal{M}) + F_\psi(\mathcal{M}, p_t, a_t, o_{t+1}),$$

where $F_\psi$ is a learned generalized force field derived from $\Delta_\psi$.

**Principle 2.** *The neural component proposes dynamics, but admissibility is enforced by variational constraints.*

This separation guarantees that learned forcing can enrich the dynamics without violating the invariants encoded in $\mathcal{C}$.

## 9 Noether Theorem for Semantic Field Invariants

The augmented RSVP system admits conservation laws arising from symmetries of the action functional. These laws are not imposed externally but follow from the variational structure established in the previous section.

## 9.1 Symmetry Transformations

Let $\mathcal{M}_t$ transform under a one-parameter group:

$$\mathcal{M}_t \mapsto \mathcal{M}_t + \epsilon\, \delta \mathcal{M}_t,$$

such that the augmented action satisfies $\delta \mathcal{A}_{\text{aug}} = 0$.

## 9.2 Conserved Current

Then there exists a conserved current $J^\mu$ satisfying:

$$\partial_t J^0 + \nabla \cdot \mathbf{J} = 0.$$

## 9.3 Examples of Semantic Invariants

**Coherence Conservation.** If $\mathcal{L}$ is invariant under $\Phi \mapsto \Phi + \epsilon$, then

$$Q_\Phi = \int_X \Phi \, d\mu_g$$

is conserved, representing total semantic coherence across the manifold.

**Flow Conservation.** If $\mathcal{L}$ is invariant under translations of $\mathbf{v}$, then

$$Q_v = \int_X \mathbf{v} \, d\mu_g$$

is conserved, representing total directional agency.

**Entropy Balance.** If $\mathcal{L}$ admits a scaling symmetry in $S$, then

$$Q_S = \int_X S \, d\mu_g$$

is conserved modulo boundary flux, representing entropy redistribution rather than entropy creation.

**Residue Persistence.** For invariance under shifts of $R$ in unobserved regions, we obtain:

$$\partial_t R(x) \geq 0 \quad \text{for } x \notin \text{view}(p_t).$$

Thus causal residue cannot spontaneously disappear from regions outside the agent's perceptual cone.

**Theorem 9.1** (Semantic Conservation Law)**.** *If the augmented action is invariant under a continuous transformation, then the corresponding semantic quantity is conserved along admissible field trajectories.*

*Proof.* Follows from the standard Noether theorem applied to $\mathcal{A}_{\text{aug}}$. □

The RSVP–Polyxan system thus satisfies variational evolution through the field equations derived from $\mathcal{A}_{\text{aug}}$, constrained admissibility enforced by the projection manifold $\mathcal{C}$, learned perturbation through the neural forcing $F_\psi$, and conserved quantities arising from Noether's theorem. Together these properties establish the system as a genuine field theory of semantic dynamics rather than a purely descriptive model.

# 10 Hamiltonian Formulation of RSVP Memory Dynamics

We now reformulate RSVP–Polyxan dynamics in Hamiltonian form, providing a phase-space description of latent field evolution.

## 10.1 Canonical Variables

Let $\mathcal{M}_t : X \to \mathbb{R}^d$ denote the latent field memory. Define conjugate momenta:

$$\Pi_t(x) = \frac{\partial \mathcal{L}}{\partial(\partial_t \mathcal{M}_t(x))}.$$

The phase space is therefore $\mathcal{P} = \{(\mathcal{M}, \Pi)\}$, a space of pairs of field configurations and their conjugate momenta.

## 10.2 Hamiltonian Functional

Define the Hamiltonian via Legendre transform:

$$\mathcal{H}[\mathcal{M}, \Pi] = \int_X \left(\langle \Pi, \partial_t \mathcal{M} \rangle - \mathcal{L}\right) \, d\mu_g.$$

Expanding the kinetic and gradient terms yields:

$$\mathcal{H} = \int_X \left( \frac{1}{2}\|\Pi\|^2 + \frac{c_\Phi^2}{2}\|\nabla \Phi\|^2 + \frac{c_v^2}{2}\|\nabla \mathbf{v}\|^2 + \frac{c_S^2}{2}\|\nabla S\|^2 + V(\mathcal{M}) \right) d\mu_g.$$

The Hamiltonian thus decomposes into a kinetic term $\frac{1}{2}\|\Pi\|^2$ representing the rate of change of field configurations, structural gradient terms $\|\nabla \Phi\|^2$, $\|\nabla \mathbf{v}\|^2$, $\|\nabla S\|^2$ representing the elastic energy stored in spatial variation, and the interaction potential $V$ encoding coupling to the Polyxan graph.

### 10.3 Hamilton's Equations

The evolution equations are:

$$\partial_t \mathcal{M} = \frac{\delta \mathcal{H}}{\delta \Pi}, \quad \partial_t \Pi = -\frac{\delta \mathcal{H}}{\delta \mathcal{M}} + F_\psi.$$

Here $F_\psi$ is the learned forcing term from Section **??**.

### 10.4 Symplectic Structure

Define the canonical symplectic form:

$$\omega = \int_X d\Pi \wedge d\mathcal{M}.$$

**Theorem 10.1.** *In the absence of forcing and dissipation, the RSVP memory dynamics preserve the symplectic form $\omega$.*

*Sketch.* Follows from the standard result that Hamiltonian flow preserves $\omega$, since the unforced equations of motion are precisely the condition $\iota_{X_H}\omega = d\mathcal{H}$. □

Symplectic preservation means that the phase-space volume of any ensemble of field configurations is conserved by the unforced dynamics. Learned forcing $F_\psi$ breaks symplecticity and introduces dissipation, but the constraint projection $\Pi_{\mathcal{C}}$ acts as a restoring force that keeps trajectories on the admissible manifold. RSVP dynamics thus correspond to energy-conserving flow on the space of semantic configurations, perturbed by learned forces that drive the system toward observed evidence.

## 11 Renormalization and Scale Separation in Semantic Fields

The RSVP–Polyxan framework operates across multiple spatial and semantic scales, from fine-grained latent appearance fields $z_t(x)$ to global hypergraph structure. The relationship between these scales is not merely informal: the theory must specify how field descriptions at one resolution are related to those at another, and how effective dynamics emerge from integrating out fine-scale fluctuations. We therefore introduce a renormalization operator that makes scale separation explicit.

### 11.1 Renormalization Operator

Define a scale parameter $\ell > 0$ and a coarse-graining operator:

$$\mathcal{R}_\ell : \mathcal{M}_t \longrightarrow \mathcal{M}_t^{(\ell)},$$

where $\mathcal{M}_t^{(\ell)}$ is a smoothed field representation over regions of diameter $\ell$ in $X$. Each component transforms under a convolution kernel $K_\ell$:

$$\Phi_t^{(\ell)}(x) = \int_{B_\ell(x)} K_\ell(x, y)\, \Phi_t(y)\, d\mu_g(y),$$

with analogous definitions for $\mathbf{v}_t$, $S_t$, $A_t$, $R_t$, and $z_t$. The kernel $K_\ell$ is required to be non-negative, normalized, and to converge to a delta function as $\ell \to 0$.

### 11.2 Scale-Dependent Dynamics

The effective dynamics at scale $\ell$ are governed by a renormalized energy functional:

$$E^{(\ell)}[\mathcal{M}] = \int_X V^{(\ell)}(\mathcal{M}^{(\ell)})\, d\mu_g,$$

where $V^{(\ell)}$ incorporates the contributions from field fluctuations below scale $\ell$ that have been integrated out. This is the field-theoretic analogue of Wilson's renormalization group: coarser descriptions are related to finer ones by integrating out short-wavelength degrees of freedom.

### 11.3 Semantic Interpretation and Limiting Cases

At small $\ell$, the field retains fine-grained appearance, local entropy, and object-level residue. At large $\ell$, only coarse structural constraints and global semantic organization survive. MultiGen corresponds to the large-$\ell$ limit in which only scalar geometry persists and all semantic components collapse. The discrete prototype of Section 26 corresponds to an intermediate $\ell$ at which the field is resolved on a $128 \times 128$ grid, sacrificing high-frequency appearance detail while retaining the structural asymmetry between observation and substrate. The full RSVP formulation retains all scales simultaneously via the multi-component field $\mathcal{M}_t$.

**Principle 3.** *Semantic structure is scale-dependent, but governed by a unified field theory under renormalization. The diversity of existing generative world models corresponds to different choices of coarse-graining scale, not to fundamentally different architectures.*

## 12 Polyxan as a Discrete Interface to Field Memory

The Polyxan hypergraph $G = (V, E)$ provides a discrete interface into the continuous field $\mathcal{M}_t$. Each node $i \in V$ corresponds to a sample at $\mathbf{z}_i \in X$, with the full field value:

$$\mathcal{M}_t(\mathbf{z}_i) = (\Phi_i, \mathbf{v}_i, S_i, A_i, R_i, z_i).$$

Edges in $E$ define constraints and transformations that influence the local field geometry by modifying $\rho$ and $\kappa$ in the potential $V$.

### 12.1 Bidirectional Coupling

The coupling between field and graph is bidirectional. In the field-to-graph direction, RSVP fields determine node dynamics through the gradient flow equation established in Section 4. In the graph-to-field direction, hypergraph structure modifies the density $\rho$ and curvature $\kappa$ terms, which in turn influence the field evolution through the potential $V(\Phi, \mathbf{v}, S; \rho, \kappa)$. Thus Polyxan acts as a read/write layer over $\mathcal{M}_t$: reading the field to generate observable structure, and writing curvature and density back into the field to record the effects of discrete interactions.

## 13 Gauge Symmetries and Semantic Equivalence

We now introduce gauge symmetries corresponding to transformations that preserve semantic equivalence classes. These symmetries formalize the principle that multiple internal representations may correspond to the same observable semantic reality, and that a well-formed theory should be independent of which representative is chosen.

### 13.1 Gauge Transformations

Let $\mathcal{G}$ be a group of transformations acting on $\mathcal{M}_t$:

$$\mathcal{M}_t \mapsto \mathcal{M}_t^g = g \cdot \mathcal{M}_t, \quad g \in \mathcal{G},$$

subject to the requirement that observations be gauge-invariant:

$$\mathcal{O}_\theta(\mathcal{M}_t^g, p_t) = \mathcal{O}_\theta(\mathcal{M}_t, p_t).$$

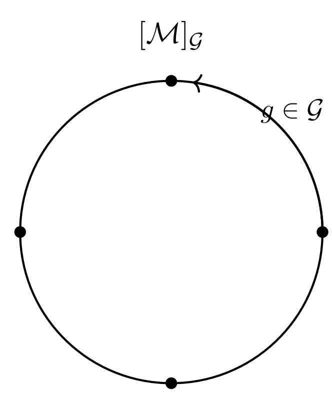


Figure 2: Gauge orbit of semantically equivalent field configurations. All points on the orbit produce identical observations under $\mathcal{O}_\theta$; the equivalence class $[\mathcal{M}]_\mathcal{G}$ is the physical degree of freedom.

### 13.2 Examples of Gauge Groups

Three natural gauge groups arise from the structure of $\mathcal{M}_t$. The latent appearance gauge acts by $z_t(x) \mapsto z_t(x) + \nabla\chi(x)$ for a scalar field $\chi$ preserving perceptual equivalence under $\mathcal{O}_\theta$. Embedding reparameterization acts by $\mathbf{z}_i \mapsto f(\mathbf{z}_i)$ where $f$ is a diffeomorphism of $X$ preserving semantic relations encoded in the hypergraph. Affordance relabeling acts by $A_t(x) \mapsto \sigma(A_t(x))$ where $\sigma$ is a symmetry of the interaction algebra governing action-affordance pairings.

### 13.3 Gauge Invariance of the Action

**Definition 13.1.** *The RSVP action is gauge-invariant if* $\mathcal{A}[\mathcal{M}] = \mathcal{A}[\mathcal{M}^g]$ *for all* $g \in \mathcal{G}$.

**Theorem 13.2.** *Gauge invariance implies conservation laws corresponding to redundant semantic degrees of freedom, via Noether's second theorem.*

### 13.4 Connection and Covariant Derivative

Introduce a gauge connection $\mathcal{A}_\mu$ and define the covariant derivative:

$$D_\mu \mathcal{M} = \partial_\mu \mathcal{M} + \mathcal{A}_\mu \cdot \mathcal{M}.$$

Gradients in the Lagrangian are replaced by covariant derivatives $\nabla \to D$, ensuring that the action transforms covariantly under local gauge transformations. The field strength curvature is then:

$$\mathcal{F}_{\mu\nu} = [D_\mu, D_\nu],$$

which measures the obstruction to consistent semantic identification across distinct regions of $X$. Non-vanishing curvature $\mathcal{F}_{\mu\nu} \neq 0$ indicates that representations in overlapping neighborhoods cannot be identified by a flat connection, a field-theoretic signature of representational inhomogeneity.

## 14 Sheaf-Theoretic Interpretation

Define for each user $u$ an open neighborhood $U_u \subset X$. Define a presheaf $\mathscr{G}$ of galaxy renderings:

$$\mathscr{G}(U) = \{\text{layout functions over } U\}.$$

Restriction maps obey $\rho_{UV}(G_U) = G_V$ for $V \subset U$.

**Theorem 14.1.** $\mathscr{G}$ *is a sheaf if the layout map is determined by* $(\Phi, \mathbf{v}, S)$, *embeddings are globally indexed, and projections are deterministic.*

*Sketch.* Local layouts glued along intersections produce a unique global layout due to determinism and continuity of the underlying field. □

## 15 Derived and ∞-Categorical Structure of Polyxan

We now lift the Polyxan–RSVP framework into derived and higher-categorical form, providing a language adequate to the treatment of higher-order equivalences and obstruction phenomena.

### 15.1 Derived Field Space

Instead of treating $\mathcal{M}_t$ as a point in a manifold, we consider it as an object in a derived moduli space:

$$\mathcal{M}_t \in \mathcal{X},$$

where $\mathcal{X}$ is a derived stack encoding field configurations, constraint solutions, and the higher-order equivalences between them. The passage to derived geometry allows the moduli space to carry cohomological information that is invisible at the classical level.

### 15.2 Obstruction Theory

Let $\mathrm{Obs}^k(\mathcal{M})$ denote the $k$-th obstruction class of a configuration $\mathcal{M}$.

**Definition 15.1.** *A configuration $\mathcal{M}$ is admissible if $\mathrm{Obs}^k(\mathcal{M}) = 0$ for all $k \geq 1$.*

Non-vanishing obstruction classes correspond to three qualitatively distinct failure modes. Inconsistencies in world-state arise when local field patches fail to agree on overlaps. Multiplayer desynchronization arises when distinct agent perspectives induce incompatible updates to $\mathcal{M}_t$. Failure of semantic gluing arises when local Polyxan neighborhoods cannot be assembled into a coherent global hypergraph.

### 15.3 Polyxan as an $\infty$-Category

We upgrade the category $\mathbf{C}$ of Section 16 to an $\infty$-category $\mathbf{C}_\infty$, in which objects are content atoms, 1-morphisms are typed links, and higher morphisms encode equivalences of transformations and coherence data for composites. The passage from $\mathbf{C}$ to $\mathbf{C}_\infty$ is necessary because the composition of semantic transformations is in general only associative up to coherent homotopy, not strictly.

### 15.4 Functor to Derived Fields

Define:

$$\mathsf{RSVP} : \mathbf{C}_\infty \to \mathbf{Field}_\infty,$$

mapping hypergraph structure into derived field configurations. This functor preserves the homotopy-theoretic structure of both sides: equivalences of content atoms map to quasi-isomorphisms of field configurations.

### 15.5 Homotopy-Coherent Gluing

The sheaf $\mathscr{G}$ of Section 13 is refined to a derived sheaf:

$$\mathscr{G}^\infty : \mathrm{Op}(X)^{op} \to \infty\text{-Groupoids}.$$

**Theorem 15.2.** *Global semantic coherence corresponds to the existence of a homotopy limit:*

$$\mathcal{M} \simeq \operatorname{holim} \mathscr{G}^{\infty}(U_i).$$

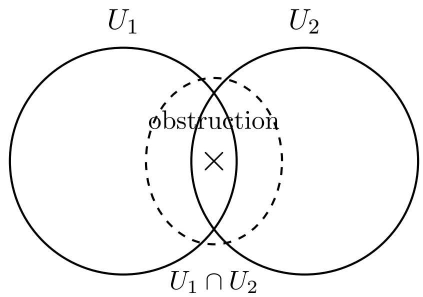


Figure 3: Failure of sheaf gluing corresponds to non-vanishing cohomology $H^1(\mathscr{G}^{\infty}) \neq 0$. The cross marks the obstruction class that prevents local patches from assembling into a globally consistent world-state.

### 15.6 Interpretation of Hallucination

We interpret hallucination as non-vanishing first cohomology of the derived sheaf:

$$H^1(\mathscr{G}^{\infty}) \neq 0.$$

This represents the failure of local semantic patches to glue into a consistent global structure: the system generates locally coherent content that cannot be assembled into a globally consistent world-state. The cohomological characterization provides a mathematically precise criterion for detecting and diagnosing hallucination without appeal to external ground truth.

### 15.7 Semantic Fixed Points

A consistent world-state is a derived fixed point:

$$\mathcal{M} \simeq \mathsf{RSVP}(\mathsf{Poly}(\mathcal{M})).$$

Semantic stability therefore corresponds to homotopy-invariant closure: the system's field configuration is a fixed point of the composed functor up to coherent equivalence.

## 16 Category-Theoretic Architecture

Define category $\mathbf{C}$ with objects given by Content Atoms and Spans, and morphisms given by Typed Links. Define the Polycompiler as an endofunctor:

$$\mathsf{Poly} : \mathbf{C} \to \mathbf{C}.$$

Define RSVP as a functor $\mathsf{RSVP} : \mathbf{C} \to \mathbf{F}$ where $\mathbf{F}$ is the category of field configurations.

**Theorem 16.1.** *The composition* $\mathsf{RSVP} \circ \mathsf{Poly}$ *yields a natural transformation describing generative semantic deformation.*

# 17 Quantization of RSVP Field Dynamics

We now construct a quantum, or equivalently stochastic, extension of RSVP memory dynamics, interpreting the latent field evolution as a path integral over field configurations.

## 17.1 Configuration Space

Let $\mathcal{M}(t)$ denote a trajectory in field space:

$$\mathcal{M} : [0, T] \to \mathcal{F},$$

where $\mathcal{F}$ is the space of admissible field configurations, that is, the classical solutions to the constrained Euler–Lagrange equations.

## 17.2 Path Integral

Define the partition functional:

$$Z = \int \mathcal{D}\mathcal{M} \, \exp\left(\frac{i}{\hbar}\mathcal{A}_{\text{aug}}[\mathcal{M}]\right),$$

where the functional integral is taken over all paths in $\mathcal{F}$. Observables $\mathcal{O}$ are computed as:

$$\langle \mathcal{O} \rangle = \frac{1}{Z} \int \mathcal{D}\mathcal{M} \, \mathcal{O}[\mathcal{M}] \exp\left(\frac{i}{\hbar}\mathcal{A}_{\text{aug}}\right).$$

## 17.3 Stochastic Diffusion Limit

Under Wick rotation $t \mapsto -i\tau$, the Lorentzian path integral becomes a Euclidean one:

$$Z_E = \int \mathcal{D}\mathcal{M} \, \exp\left(-\frac{1}{\hbar}\mathcal{A}_E[\mathcal{M}]\right).$$

This corresponds to the stochastic differential equation:

$$\partial_\tau \mathcal{M} = -\nabla_{\mathcal{M}} E(\mathcal{M}) + \eta(\tau),$$

where $\eta$ is Gaussian white noise with variance $\hbar$. The invariant measure of this process is the Boltzmann distribution $p(\mathcal{M}) \propto e^{-E(\mathcal{M})/\hbar}$ over field configurations.

## 17.4 Connection to Diffusion Generative Models

Diffusion generative models correspond precisely to sampling trajectories from this stochastic evolution. The RSVP energy functional $E$ defines the score function $\nabla_{\mathcal{M}} \log p(\mathcal{M})$ that guides the reverse diffusion process. Generation is then sampling from the Boltzmann distribution over fields, and the denoising process is the time-reversal of the stochastic gradient flow on $E$.

**Principle 4.** *Diffusion models are discrete approximations to stochastic quantization of RSVP field dynamics.*

A single generated world is one sampled trajectory through $\mathcal{F}$. Consistency of the generated world emerges from the action functional, which penalizes configurations that violate the constraints encoded in $\mathcal{C}$. Uncertainty at any point $x \in X$ corresponds to the entropy field $S_t(x)$, which is high where the path integral measure is diffuse and low where it concentrates on a narrow set of field configurations.

# 18 Information Geometry of RSVP Field Space

The stochastic formulation of Section 17 induces an information-geometric structure on the space of field configurations. This structure connects the entropy field $S_t$, the score function of the diffusion process, and the learning dynamics of $\mathcal{N}_\psi$ within a unified statistical manifold.

## 18.1 Statistical Manifold

Let $p(\mathcal{M}) \propto e^{-E(\mathcal{M})/\hbar}$ be the Boltzmann distribution over field configurations induced by the RSVP energy. Define the Fisher information metric:

$$g_{ij}(\mathcal{M}) = \mathbb{E}_p[\partial_i \log p(\mathcal{M})\ \partial_j \log p(\mathcal{M})]\,,$$

where $i, j$ index directions in field space. This metric equips the space of field configurations with a Riemannian structure that is intrinsic to the distribution $p$, rather than inherited from the ambient Euclidean structure of $\mathbb{R}^d$.

## 18.2 Natural Gradient Flow

The Langevin dynamics of Section 17:

$$\partial_\tau \mathcal{M} = -\nabla_{\mathcal{M}} E(\mathcal{M}) + \eta(\tau)$$

correspond, in the zero-noise limit, to steepest descent under the Euclidean metric. Under the Fisher metric $g_{ij}$, the analogous flow is natural gradient descent:

$$\partial_\tau \mathcal{M} = -g^{ij}\, \partial_j E(\mathcal{M}),$$

which converges more efficiently by following the intrinsic curvature of the probability simplex rather than the extrinsic geometry of parameter space. The learning dynamics of $\mathcal{N}_\psi$ are thus interpretable as approximating natural gradient flow on the statistical manifold of field configurations.

## 18.3 Entropy as Local Curvature

The entropy field $S_t(x)$ corresponds to local dispersion in the distribution over field configurations at position $x$. Regions of high $S_t(x)$ correspond to flat directions in the energy landscape at $x$:

many field configurations are nearly equally probable, and the Hessian of $E$ at those points is small. Regions of low $S_t(x)$ correspond to sharply curved energy basins in which the field is tightly concentrated near a local minimum. The write operator $\mathcal{W}$ reduces $S_t(x)$ by injecting residue that deepens the energy basin at the interaction site, causing the field distribution to sharpen and the observation operator $\mathcal{O}_\theta$ to produce more deterministic outputs.

**Principle 5.** *Learning and inference in RSVP–Polyxan correspond to geometric flows on a statistical manifold of field configurations, where the entropy field encodes local curvature and the write operator acts as a mechanism of basin sharpening.*

# 19 Existence and Stability of RSVP Field Dynamics

We address the well-posedness of the RSVP evolution equations, establishing that solutions exist locally in time, that stable configurations correspond to semantically coherent world-states, and that observations are not sufficient alone to reconstruct the full latent field.

## 19.1 Existence of Solutions

**Theorem 19.1.** *Under smooth initial conditions $\mathcal{M}_0 \in \mathcal{F}$ and bounded forcing $F_\psi$, the constrained RSVP evolution admits a local-in-time solution $\mathcal{M}_t$ for $t \in [0, T)$.*

*Sketch.* The unconstrained Euler–Lagrange equations derived from $\mathcal{A}_{\text{aug}}$ define a hyperbolic PDE system in $\mathcal{M}$. Standard energy methods yield local existence for smooth data. The constraint projection $\Pi_\mathcal{C}$ acts as a Lipschitz retraction onto the closed admissible set $\mathcal{C}$, which by standard results for projected dynamical systems preserves local existence and adds at most a bounded constraint force to the right-hand side. □

## 19.2 Convergence Toward Stable Configurations

**Proposition 19.2.** *If the energy functional $E(\mathcal{M})$ is bounded below and the entropy dissipation $\mathcal{D}(S_t)$ is strictly positive, then trajectories of the constrained RSVP system converge toward a stable set of admissible configurations.*

*Sketch.* $E$ is a Lyapunov functional: the unforced dynamics decrease $E$ by the symplectic flow established in Section 10, and the dissipation term $\mathcal{D}(S_t)$ removes energy from the system. By LaSalle's invariance principle applied to the projected system, trajectories accumulate at the largest invariant set within $\{\mathcal{M} : \mathcal{D}(S(\mathcal{M})) = 0\} \cap \mathcal{C}$, which is precisely the set of admissible configurations with minimal entropy dissipation. □

Stability in this sense is semantic coherence: the field converges to configurations that minimize energy while satisfying the conservation constraints of Section 23, and these configurations correspond to consistent world-states in which no constraint is violated and no residue is being actively generated.

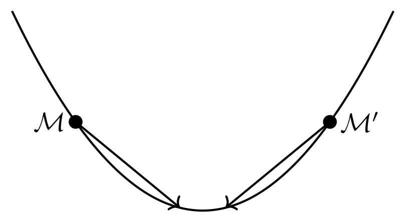


Figure 4: Field evolution as gradient descent in the energy landscape $E(\mathcal{M})$. Distinct initial configurations converge to the same stable admissible state under constrained dynamics.

### 19.3 Observability and Identifiability

The separation between latent field memory $\mathcal{M}_t$ and observations $o_t$ raises the question of observability: whether distinct field configurations produce distinct observation distributions under $\mathcal{O}_\theta$, and whether the field can in principle be reconstructed from the observation sequence alone.

Partial observability is the generic case. Multiple field configurations may map to identical observations, since $\mathcal{O}_\theta$ is a projection from the high-dimensional field $\mathcal{M}_t$ to the lower-dimensional observation $o_{t+1}$. In particular, components of $\mathcal{M}_t$ that lie outside the current view cone $\mathrm{view}(p_t)$ are not constrained by $o_{t+1}$ at all.

The write operator $\mathcal{W}$ resolves this ambiguity progressively by collapsing equivalence classes through residue accumulation. When a region $x$ is observed, the injection of residue $\mathcal{R}_t(x)$ and the collapse of $S_t(x)$ distinguish the realized configuration from all others in its equivalence class under $\mathcal{O}_\theta$. Repeated observation of overlapping regions thus accumulates enough constraint to identify the field configuration in the observed subset of $X$.

**Principle 6.** *Observation alone is insufficient for world reconstruction. Persistence and identifiability require write-back into the field: the world is not recovered by reading; it is fixed by writing.*

## 20 Stability and Energy Minimization

Define energy:

$$E = \int_X V(\Phi, \mathbf{v}, S; \rho, \kappa)\, d\mu_g.$$

**Lemma 20.1.** *Fixed points of RSVP flow minimize $E$ subject to embedding constraints.*

*Sketch.* Gradient descent of embeddings follows $-\nabla E$ by design. □

## 21 Global Reset as Field Reconfiguration

Define reset operator:

$$\mathcal{R} : (\Phi, \mathbf{v}, S, \mathbf{z}) \mapsto (\Phi', \mathbf{v}', S', \mathbf{z}'),$$

where $\mathcal{R}$ recomputes fields and embeddings by re-estimating $\rho$ and $\kappa$, relaxing the field equations to a new equilibrium, and reprojecting embeddings into $X$. Reset corresponds to a global reconfiguration of semantic geometry, and is distinguished from the incremental update $\Delta_\psi$ in that it does not preserve the prior field configuration but reconstructs it from the current hypergraph topology.

## 22 MultiGen as a Degenerate Projection of RSVP

MultiGen [1] represents the most structurally sophisticated system in the emerging class of diffusion-based generative game engines. We do not treat it merely as a compatible architecture. We show that it is a degenerate projection of the RSVP trajectory-operator framework, arising under severe dimensional restriction of the full field state.

### 22.1 The Projection Map

MultiGen defines an explicit world-state:

$$S_t = (M,\, p_t,\, o_{t-L+1:t}),$$

consisting of a global geometric memory $M$, a player pose $p_t$, and a finite visual trace of length $L$. In RSVP terms, this bundle arises from the full state $X_t = (\mathcal{M}_t, p_t, h_t)$ via a projection:

$$\pi : X_t \longrightarrow S_t,$$

which acts by (i) collapsing the rich latent field $\mathcal{M}_t = (\Phi_t, \mathbf{v}_t, S_t, A_t, R_t, z_t)$ to its scalar substrate $\Phi_t$ alone, (ii) retaining only the pose coordinate $p_t$, and (iii) replacing the full observation history with a window of length $L$. MultiGen is therefore not a different kind of system but the RSVP system evaluated at $\pi$.

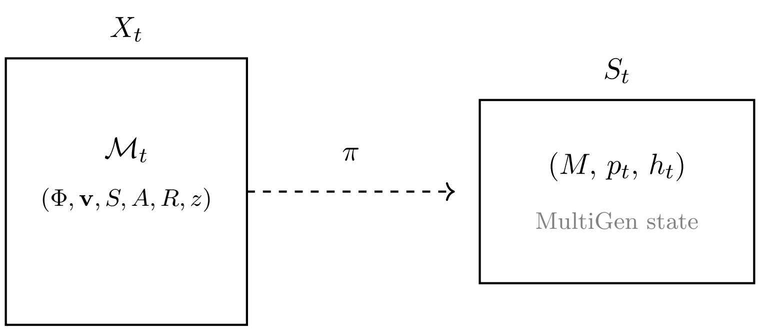


Figure 5: MultiGen as a projection $\pi$ of the full RSVP state. The dashed arrow indicates that MultiGen's state bundle arises by collapsing the rich latent field to scalar geometry alone and discarding the write operator.

### 22.2 Field Correspondences

The correspondence between MultiGen components and RSVP fields is exact. The authored minimap $M = (V, E)$ corresponds to a coarse scalar constraint field $M \sim \Phi$: it encodes where structure is geometrically possible and where it is blocked, but carries no semantic content beyond occupancy. The player pose $p_t$ is the agent trajectory coordinate, unchanged across both frameworks. The observation operator is a projection from field and pose:

$$\mathcal{O}_\phi : (\Phi,\, p_t,\, h_t,\, a_t) \;\mapsto\; o_{t+1},$$

and the dynamics operator updates pose only:

$$\mathcal{D}_\psi : (p_t,\, a_t,\, r_t,\, f_t) \;\mapsto\; p_{t+1}.$$

The full MultiGen rollout is therefore:

$$M_{t+1} = M_t, \qquad p_{t+1} = \mathcal{D}_\psi(p_t, a_t, \cdots), \qquad o_{t+1} = \mathcal{O}_\phi(M_t, p_t, \cdots).$$

The substrate is frozen: $M_{t+1} = M_t$ for all $t$.

### 22.3 The Missing Write Operator

The structural deficiency of MultiGen is made precise by the absence of a write operator. In the RSVP framework, the dynamics operator $\Delta_\psi$ writes observation-derived information back into the field:

$$\mathcal{W} : o_{t+1} \;\longrightarrow\; \mathcal{M}_{t+1}.$$

MultiGen has no such operator. The field $M$ is an immutable anchor produced by a level designer prior to runtime; it is read by both observation and dynamics modules but never updated by them. This is not an engineering limitation but a structural one. The absence of $\mathcal{W}$ means the system has no feedback channel from observation back into substrate. It is a one-way system: constraint flows toward projection, but projection never flows back toward constraint.

### 22.4 MultiGen as a Theorem

The preceding analysis entitles the following proposition.

**Proposition 22.1.** *MultiGen is the degenerate projection of RSVP–Polyxan obtained by setting $\mathbf{v}_t = 0$, $A_t = 0$, $R_t = 0$, $z_t = 0$, restricting $\Phi_t$ to a static geometric graph, and removing the write operator $\mathcal{W}$. Under this projection, the RSVP update equations reduce exactly to the MultiGen rollout.*

The RSVP generalization is therefore not merely an extension but a structural completion. MultiGen establishes that diffusion must be demoted from world-state to observation operator. RSVP establishes the missing half: observation must be coupled to a write-back mechanism that promotes realized detail into the persistent field, so that the system is bidirectional—constraint to projection, and projection back to constraint—and the field can converge toward a fixed point under its own dynamics.

### 22.5 Multiplayer as Evidence

MultiGen's multiplayer capability is the most direct empirical confirmation of the substrate-first principle. In a purely diffusion-based system, two agents cannot share a world because there is no shared source of truth beyond the model weights. MultiGen achieves cross-agent consistency by making all agents read from the same $M$, so spatial agreement is guaranteed by geometric constraint rather than model intelligence. In the RSVP extension, multiplayer no longer reduces to reading from a shared map. It becomes co-evolution of a shared field: every action by any agent induces a write into $\mathcal{M}_t$ via $\mathcal{W}$, and all subsequent observations across all agents are constrained by the accumulated field history. Two agents agree not merely on where walls are but on what has happened to those walls. Consistency is achieved not through spatial anchoring alone but through shared causal residue.

## 23 Factored Update: Proposal, Projection, and Commit

We now refine the dynamics operator $\Delta_\psi$ introduced in Section 7 into a three-stage factored update, and specify the constraint set $\mathcal{C}$ with field-level equations rather than prose invariants.

### 23.1 The Three-Stage Update

A purely neural update $\Delta_\psi$ applied without constraint will exhibit drift on long horizons: high-capacity updaters treat the field as a writable canvas, producing visually plausible local edits that accumulate global contradictions. The correct formulation is a constrained neural operator: a learnable update whose proposals live inside a space defined by explicit field-level laws. We factor the update into three stages:

$$\mathcal{M}_{t+1} = \Pi_{\mathcal{C}}\Big(\mathcal{M}_t + \mathcal{N}_\psi(\mathcal{M}_t, p_t, a_t, o_{t+1})\Big),$$

where $\mathcal{N}_\psi$ is the neural proposal network, which generates local edits to the field; $\Pi_{\mathcal{C}}$ is the projection onto the admissible constraint manifold; and the commit stage writes back only the closest valid state. where $\mathcal{N}_\psi$ is the neural proposal network, which generates local edits to the field; $\Pi_{\mathcal{C}}$ is the projection onto the admissible constraint manifold; and the commit stage writes back only the closest valid state.

$$\mathcal{M}_t \xrightarrow{\text{proposal}} \mathcal{M}_t + \mathcal{N}_\psi \xrightarrow{\text{projection}} \Pi_{\mathcal{C}} \xrightarrow{\text{commit}} \mathcal{M}_{t+1}$$

Figure 6: Factored update: the neural proposal $\mathcal{N}_\psi$ generates a candidate field increment, which is projected onto the constraint manifold $\mathcal{C}$ and committed as the next state.

This structure mirrors incompressible flow solvers: an unconstrained velocity update is followed by a pressure solve that enforces incompressibility. The projection is realized as:

$$\Pi_{\mathcal{C}}(\mathcal{M}_t + \delta) \;\approx\; \arg\min_{Y \in \mathcal{C}} \; \|Y - (\mathcal{M}_t + \delta)\|^2.$$

### 23.2 Explicit Constraint Equations

The constraint set $\mathcal{C}$ is defined by six field-level invariants, each with a precise mathematical form.

**Mass and Occupancy Consistency.** The occupancy density $\rho$ must satisfy a continuity equation across updates:

$$\int_\Omega \rho_{t+1}(x)\,dx \;\approx\; \int_\Omega \rho_t(x)\,dx + \text{sources}(a_t) - \text{sinks}(a_t).$$

This prevents mass from appearing or disappearing without a corresponding action.

**Energy Budget.** The total field energy is bounded by the work performed by the agent's action and dissipated by entropy:

$$E_{t+1} \;\le\; E_t + W(a_t) - \mathcal{D}(S_t),$$

where $W(a_t)$ is the work budget of action $a_t$ and $\mathcal{D}(S_t)$ is the entropy-driven dissipation.

**Topological Admissibility.** Edits to the substrate $\Phi_t$ must preserve or change topology only via admissible operations—cut, join, fracture—with topological validity monitored through Betti number changes or a learned admissibility classifier that vetoes proposals introducing illegal connectivity changes.

**Causal Write Locality.** The support of $\mathcal{N}_\psi$ must be spatially localized around the causal influence cone of the action. This is enforced by a kernel $K_\psi$ that decays with geodesic distance from the action site in $X$:

$$\mathcal{N}_\psi(\mathcal{M}_t, p_t, a_t, o_{t+1})(x) \;\propto\; K_\psi(d_X(x, \text{support}(a_t)))\,.$$

**Entropy Accounting.** Entropy $S_t$ increases under occlusion and uncertainty, and decreases under observation. Write-backs must not reduce entropy outside the agent's observation cone:

$$\Delta S_t(x) \ge 0 \quad \text{for all } x \notin \text{view}(p_t).$$

This enforces the second law at the field level and prevents information from appearing in unobserved regions.

**Identity Persistence.** Object embeddings in $\mathcal{M}_t$ carry persistent identity labels. Updates must be consistent with identity trajectories: spontaneous duplication or destruction of objects is forbidden unless permitted by the mass and energy budgets of the action.

### 23.3 Composable Projections

The projection $\Pi_{\mathcal{C}}$ need not be implemented as a monolithic solver. Differentiable projections handle occupancy and continuity via Poisson-like solves, SDF clamping, and angle wrapping. Lagrange multiplier terms in a local optimization adjust proposals to satisfy energy budgets. Learned validator networks score admissibility and reject or correct proposals that fail. Typed discrete operations with preconditions and postconditions govern topology changes such as breaking a wall or opening a door. Each projection component is composable: they are applied in sequence, and the final committed state satisfies all constraints simultaneously.

## 24 Fast and Slow Dynamics, Event Operators, and Read-Write Closure

### 24.1 Fast and Slow Field Update Cadences

Not all components of $\mathcal{M}_t$ should be updated at the same temporal cadence. The latent appearance field $z_t$ and short-range deformations of $\mathbf{v}_t$ are updated at every step by the neural proposal $\mathcal{N}_\psi$, as these fields must respond to high-frequency perceptual change. The scalar substrate $\Phi_t$, topological structure, affordances $A_t$, and object identities are updated only via typed operations that pass through the full constraint projection. This two-cadence design prevents noisy neural proposals from catastrophically editing structural fields while keeping the rendering layer responsive.

### 24.2 Event-Typed Update Operators

The dynamics operator $\Delta_\psi$ is decomposed as a mixture of event-typed operators:

$$\Delta_\psi = \sum_k \alpha_k \, \mathcal{E}_k(\mathcal{M}_t, p_t, a_t),$$

where each $\mathcal{E}_k$ is a typed operator with known preconditions and field effects, and $\alpha_k$ are learned mixture weights. Examples of typed operators include displacement (translating an object along a trajectory in $X$), fracture (splitting a connected component of $\Phi_t$ into two, subject to energy budget), occlusion write (updating $S_t$ and $z_t$ in regions newly blocked from the observation cone), and residue deposit (incrementing $R_t$ at the action site). The typed decomposition provides a scaffold that constrains the neural proposal to a structured vocabulary of world-edits, sharply reducing the surface area over which drift can accumulate.

### 24.3 Read-Write Closure

To prevent divergence between what is observed and what is stored, the system must satisfy a closure condition relating the read operator $\mathrm{Read}(\mathcal{M}_t, p_t)$ to the write-back induced by $\mathcal{W}$.

The read operator must be Lipschitz in $\mathcal{M}_t$ locally: small changes to the field must not produce discontinuous changes in the rendered observation. This ensures that the observation module $\mathcal{O}_\theta$ is stable under incremental field updates.

The write-back must be consistent with read in the following cycle-closure sense. Given an observation $o_{t+1}$ produced by $\mathcal{O}_\theta$, the encoding of that observation back into field space must be compatible with the stored field:

$$\mathcal{M}_{t+1} \xrightarrow{\mathcal{O}_\theta} \hat{o}_{t+1} \xrightarrow{\text{Enc}} \widetilde{\mathcal{M}}_{t+1},$$

where $\widetilde{\mathcal{M}}_{t+1}$ and $\mathcal{M}_{t+1}$ should agree in the observed region. This cycle loss is incorporated as a training objective to prevent the field from storing representations that the observation module cannot decode, and to prevent the observation module from generating content that the field cannot absorb.

# 25 Multiplayer as Causal Field Synchronization and Training

## 25.1 Multiplayer as Causal DAG Merging

In the RSVP framework, multiplayer interaction is not a state-replication problem but a causal field synchronization problem. All agents operate over a shared field $\mathcal{M}_t$; each agent $i$ maintains a local pose $p_t^i$ and produces local observation $o_{t+1}^i = \mathcal{O}_\theta(\mathcal{M}_t, p_t^i, h_t^i, a_t^i)$. When agent $i$ takes action $a_t^i$, the corresponding neural proposal $\mathcal{N}_\psi$ generates a local field delta, which after projection is committed to $\mathcal{M}_{t+1}$. All agents' subsequent reads are constrained by this updated field.

To handle concurrent actions from multiple agents, the system maintains a causal directed acyclic graph of field-editing events, each stamped with a logical timestamp and the identity of the agent that produced it. Commutative field edits—residue deposits, appearance updates, entropy writes in non-overlapping regions—are merged via CRDT-style operations that are order-independent. Non-commutative edits, such as two agents simultaneously attempting to destroy the same structural component, are resolved through the energy budget constraint: the action with sufficient energy at the time of application succeeds; the other is vetoed by $\Pi_\mathcal{C}$ and generates a null delta. Because all agents project into the same admissible manifold $\mathcal{C}$, divergent local proposals collapse to the same valid committed state, and consistency is guaranteed by constraint structure rather than consensus protocol.

## 25.2 Training Curriculum

Training the full RSVP update system proceeds in four phases. In the first phase, the observation operator $\mathcal{O}_\theta$ is pretrained as a renderer conditioned on $\mathcal{M}_t$, with the field held fixed at ground-truth values. In the second phase, the neural proposal network $\mathcal{N}_\psi$ is pretrained on synthetic event logs with known ground-truth field deltas, supervised by the typed operator structure. In the third phase, the constraint projection $\Pi_\mathcal{C}$ is introduced with soft penalties that gradually anneal toward hard projections, so the system learns to generate proposals that are already close to admissible before projection. In the fourth phase, the full system is trained end-to-end using self-consistency losses: when ground-truth field deltas are unavailable, the cycle-closure condition between $\mathcal{O}_\theta$ and Enc provides the supervision signal, and conservation penalties enforce the field-level invariants of $\mathcal{C}$. This curriculum ensures that the system acquires local expressivity from the neural proposal, global coherence from the constraint projection, and long-horizon persistence from the enforced invariants, in that order.

# 26 Computational Instantiation

The preceding formalism admits a direct computational realization. To demonstrate that the proposed architecture is not merely interpretive but executable, we construct a minimal prototype that instantiates the separation between substrate, dynamics, and observation.

The implementation is deliberately restricted to the smallest system that preserves the essential asymmetry: observation is stochastic and local, while memory is persistent and global. The system consists of a client-side rendering layer and a server-side field evolution layer. The former implements the observation operator $\mathcal{O}_\theta$, while the latter implements a constrained update $\mathcal{T}$ together with an explicit write channel $\mathcal{W}$.

The global state is represented as a discretized field

$$X_t = (\Phi_t,\, S_t,\, \mathcal{R}_t),$$

where $\Phi_t$ encodes scalar structure, $S_t$ encodes uncertainty, and $\mathcal{R}_t$ accumulates residue. This field is stored persistently on the server and evolves independently of any particular observation.

The observation operator $\mathcal{O}_\theta$ is realized in the client as a projection from field to pixel space. Crucially, this projection is stochastic: each render introduces noise proportional to the local entropy $S_t$. In the absence of a write-back mechanism, this would produce temporally incoherent outputs, analogous to diffusion-based frame prediction.

The write operator $\mathcal{W}$ is triggered by user interaction. Localized perturbations are injected into $\mathcal{R}_t$ and simultaneously reduce $S_t$, thereby promoting observed structure into persistent memory. This enforces the invention-to-memory principle of Section 7: generated detail becomes constraint upon realization.

The dynamics operator $\mathcal{T}$ is implemented as a diffusion-like update on $\Phi_t$ coupled to $\mathcal{R}_t$, together with a monotone decay of entropy. This realizes a minimal form of constraint propagation and stabilization, corresponding to the Laplacian term in the RSVP Lagrangian of Section 3.

### Client: Observation and Interaction

```
<!DOCTYPE html>
<html lang="en">
<head>
<meta charset="UTF-8">
<title>RSVP Field Prototype</title>
<style>
```

```
body { margin:0; background:#000; overflow:hidden; }
canvas { display:block; }
</style>
</head>
<body>
<canvas id="c"></canvas>
<script>
const canvas = document.getElementById("c");
const ctx = canvas.getContext("2d");
canvas.width = window.innerWidth;
canvas.height = window.innerHeight;

const W = 128, H = 128;

// persistent field: phi (structure), entropy, residue
let field = {
  phi:     new Float32Array(W*H),
  entropy: new Float32Array(W*H),
  residue: new Float32Array(W*H)
};
for (let i=0; i<W*H; i++){
  field.phi[i]     = Math.random()*0.1;
  field.entropy[i] = 1.0;
}

function idx(x, y){ return y*W + x; }

// observation operator: stochastic projection from field to pixels
function render(){
  let img = ctx.createImageData(canvas.width, canvas.height);
  for (let y=0; y<canvas.height; y++){
    for (let x=0; x<canvas.width; x++){
      let gx = Math.floor(x/canvas.width*W);
      let gy = Math.floor(y/canvas.height*H);
      let i  = idx(gx, gy);

      let base  = field.phi[i]*255 + field.residue[i]*200;
      let noise = Math.random() * field.entropy[i] * 50;  // variance ~ S_t
      let c     = base + noise;

      let p = (y*canvas.width + x)*4;
      img.data[p]   = c;
      img.data[p+1] = c*0.7;
      img.data[p+2] = 255-c;
      img.data[p+3] = 255;
    }
  }
  ctx.putImageData(img, 0, 0);
}

// write operator: interaction injects residue, collapses local entropy
canvas.addEventListener("mousemove", (e)=>{
  let x = Math.floor(e.clientX/canvas.width*W);
  let y = Math.floor(e.clientY/canvas.height*H);
```

```
  for (let dy=-2; dy<=2; dy++){
    for (let dx=-2; dx<=2; dx++){
      let nx=x+dx, ny=y+dy;
      if (nx>=0 && nx<W && ny>=0 && ny<H){
        let i = idx(nx, ny);
        field.residue[i] += 0.1;   // promote residue
        field.entropy[i] *= 0.8;   // reduce uncertainty
      }
    }
  }
});

// dynamics operator: diffusion on phi coupled to residue
function step(){
  let newPhi = new Float32Array(W*H);
  for (let y=1; y<H-1; y++){
    for (let x=1; x<W-1; x++){
      let i   = idx(x,y);
      let lap =
        field.phi[idx(x+1,y)] + field.phi[idx(x-1,y)] +
        field.phi[idx(x,y+1)] + field.phi[idx(x,y-1)] -
        4*field.phi[i];

      newPhi[i]         = field.phi[i] + 0.1*lap + 0.05*field.residue[i];
      field.entropy[i] *= 0.999;   // monotone decay outside view
    }
  }
  field.phi = newPhi;
}

function loop(){ step(); render(); requestAnimationFrame(loop); }
loop();
</script>
</body>
</html>
```

### Server: Persistent Field and Evolution

```
import numpy as np
from flask import Flask, jsonify, request

app = Flask(__name__)
W, H = 128, 128

# persistent field exists independently of any rendering pass
field = {
    "phi":     np.random.rand(H, W) * 0.1,
    "entropy": np.ones((H, W)),
    "residue": np.zeros((H, W))
}
```

```
# dynamics operator: constrained Laplacian update + residue coupling
def step():
    phi     = field["phi"]
    residue = field["residue"]
    lap = (
        np.roll(phi, 1, 0) + np.roll(phi,-1, 0) +
        np.roll(phi, 1, 1) + np.roll(phi,-1, 1) - 4*phi
    )
    field["phi"]     = phi + 0.1*lap + 0.05*residue
    field["entropy"] *= 0.999   # entropy monotone outside observed cone

@app.route("/step")
def api_step():
    step()
    return jsonify({"ok": True})

@app.route("/field")
def get_field():
    return jsonify({
        "phi":     field["phi"].tolist(),
        "residue": field["residue"].tolist(),
        "entropy": field["entropy"].tolist()
    })

# write operator: localized residue injection with entropy collapse
@app.route("/write", methods=["POST"])
def write():
    x, y = request.json["x"], request.json["y"]
    for dy in range(-2, 3):
        for dx in range(-2, 3):
            nx, ny = x+dx, y+dy
            if 0 <= nx < W and 0 <= ny < H:
                field["residue"][ny, nx] += 0.2
                field["entropy"][ny, nx] *= 0.7
    return jsonify({"ok": True})

if __name__ == "__main__":
    app.run(port=5000)
```

## Interpretation

This prototype realizes the minimal closure:

$$\mathcal{O}_\theta \longrightarrow \mathcal{W} \longrightarrow \mathcal{T} \longrightarrow \mathcal{O}_\theta,$$

in which observation produces realizations, realizations are written back as residue, and the field evolves under those accumulated constraints.

The essential property is that stochastic variation in the observation layer does not produce temporal incoherence, because the field retains the history of prior realizations through $\mathcal{R}_t$. The system

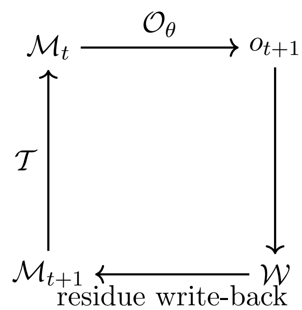


Figure 7: Read–write closure: observation $\mathcal{O}_\theta$ projects the field to a visible state; the write operator $\mathcal{W}$ injects residue and collapses local entropy; the dynamics $\mathcal{T}$ evolve the field; and the cycle repeats.

does not remember via a buffer; it remembers via deformation of the substrate itself. This is precisely the structural feature absent in purely autoregressive or diffusion-based generative systems, and it constitutes the minimal condition for persistent, causally coherent world generation.

## Discrete Constrained Field Flow

The prototype admits a direct interpretation as a discrete approximation to the constrained field evolution developed in Section 8. Writing the state as $X_t = (\Phi_t, S_t, \mathcal{R}_t)$, the server update implements an explicit Euler step:

$$X_{t+1}^{\text{prop}} = X_t + \delta_t,$$

where the proposal $\delta_t$ is given by a local Laplacian term on $\Phi_t$ coupled to $\mathcal{R}_t$, together with a monotone decay of $S_t$.

The absence of an explicit $\Pi_\mathcal{C}$ in the code reflects the choice of a minimal constraint set enforced implicitly. The locality of the Laplacian update enforces causal support; the monotone decay of $S_t$ via the factor 0.999 at each step prevents spontaneous reduction of uncertainty outside regions of interaction, realizing the entropy accounting constraint of Section 23. The write operator $\mathcal{W}$ introduces localized, entropy-reducing perturbations that act as admissible sources for structure. Thus the full update may be written schematically as:

$$X_{t+1} = \Pi_\mathcal{C}\big(X_t + \mathcal{N}_\psi(X_t, a_t)\big),$$

with $\mathcal{N}_\psi$ realized by the discrete diffusion–residue coupling, and $\Pi_\mathcal{C}$ realized implicitly by the restricted form of admissible updates.

The observation operator $\mathcal{O}_\theta$ is a stochastic projection from field space to pixel space with variance proportional to $S_t$. In the absence of write-back, this yields temporally uncorrelated samples. The inclusion of $\mathcal{R}_t$ closes the loop by promoting realized structure into persistent constraint, converting stochastic observation into deterministic evolution at the level of the field.

The prototype therefore constitutes a discrete, low-dimensional instance of a projected gradient flow on the space of admissible field configurations, in which observation acts as a probe of underdetermined degrees of freedom and writing acts as a mechanism of constraint accumulation. It is not a toy in the usual sense but a collapsed limit of the general RSVP–Polyxan theory: the system already inhabits the same category of constrained dynamical systems as the full formalism, with the richness of $\mathcal{M}_t$ replaced by the three-component field $(\Phi_t, S_t, \mathcal{R}_t)$ sufficient to exhibit the core asymmetry.

One theoretical gap between the prototype and the full formalism should be named explicitly. Section 17 derives the stochastic observation process via Wick rotation of $\mathcal{A}_{\text{aug}}$, showing that the Euclidean path integral yields a Langevin dynamics whose invariant measure is $p(\mathcal{M}) \propto e^{-E(\mathcal{M})/\hbar}$, and that diffusion generative models are samplers from this distribution with the RSVP energy functional acting as the score. The prototype realizes the *structure* of stochastic observation — noise proportional to $S_t$ — but not yet the *derivation*. The renderer is a hand-coded stochastic projection that uses $S_t$ as a proxy for the variance of the Boltzmann distribution over field configurations; it is not a Langevin sampler governed by $\nabla_{\mathcal{M}} E(\mathcal{M})$, nor a learned score network conditioned on $\mathcal{M}_t$. Closing this gap would require replacing the noise term $\eta \cdot S_t(x)$ in the render function with a score-based denoiser trained against the RSVP energy, at which point the prototype's observation operator would be a genuine discrete approximation to the Euclidean path integral of Section 17 rather than a structural analogue of it. This constitutes the primary open step between the formal theory and a fully derived implementation.

## Architecture Diagram

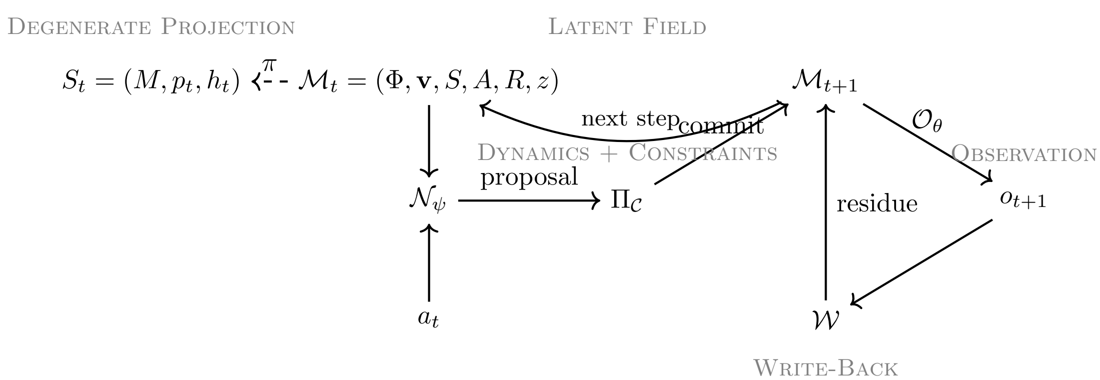


Figure 8: The RSVP–Polyxan architecture as a closed operator system. The neural proposal $\mathcal{N}_\psi$ and constraint projection $\Pi_{\mathcal{C}}$ evolve the latent field; the observation operator $\mathcal{O}_\theta$ projects to visible state; the write operator $\mathcal{W}$ closes the loop by promoting realized structure into persistent memory. The curved arrow encodes temporal recurrence. MultiGen appears as the dashed projection $\pi$ in which the write operator and all non-scalar field components are suppressed.

## 27 Conclusion

We have extended RSVP–Polyxan from a static semantic field theory into a fully dynamical world model defined by persistent latent memory, projection-based observation, and constrained field evolution, and then elevated that dynamical model through three ascending layers of mathematical structure.

At the classical layer, the constraint projection $\Pi_{\mathcal{C}}$ is internalized into an augmented action functional with Lagrange multipliers, yielding constrained Euler–Lagrange equations in which admissibility is enforced continuously rather than imposed as a post-hoc correction. Noether's theorem applied to the augmented action produces four classes of semantic conservation law: coherence conservation, flow conservation, entropy balance, and residue persistence. The system is then reformulated in Hamiltonian phase-space terms, and symplectic preservation is established for the unforced dynamics.

At the stochastic layer, Wick rotation of the path integral yields a stochastic diffusion process whose invariant measure is the Boltzmann distribution over field configurations. Diffusion generative models emerge as discrete approximations to stochastic quantization of RSVP dynamics, with the RSVP energy functional playing the role of the score function.

At the derived layer, the Polyxan category is lifted to an $\infty$-category $\mathbf{C}_\infty$, the sheaf of galaxy renderings is refined to a derived sheaf valued in $\infty$-groupoids, and semantic consistency is characterized as homotopy-coherent gluing. Hallucination is identified with non-vanishing first cohomology $H^1(\mathscr{G}^\infty) \neq 0$, providing a cohomological criterion for semantic failure that is intrinsic to the field-theoretic structure.

The relationship to MultiGen is established at the level of a formal projection theorem. MultiGen arises from RSVP–Polyxan by collapsing the full latent field to a static scalar substrate, removing the affordance, residue, flow, and appearance components, and eliminating the write operator $\mathcal{W}$. The resulting one-way system—constraint flowing toward projection but projection never flowing back toward constraint—achieves spatial consistency but not causal consistency. The RSVP completion adds the missing feedback channel: the factored update $\Pi_{\mathcal{C}} \circ \mathcal{N}_\psi$ promotes realized observations into the persistent field via six explicitly formulated field-level invariants. Multiplayer interaction becomes causal field synchronization over a shared substrate, with concurrent edits resolved by the constraint projection rather than by consensus protocol.

Three further sections close the remaining structural gaps in the framework. A renormalization operator $\mathcal{R}_\ell$ relates field descriptions at different resolutions, establishing that MultiGen and the discrete prototype correspond to large-$\ell$ and intermediate-$\ell$ limits of the full RSVP theory. An information-geometric analysis equips the space of field configurations with a Fisher metric, interprets the Langevin dynamics as natural gradient flow, and identifies the entropy field $S_t$ as local curvature in the energy landscape. An existence and stability theorem grounds the constrained evolution in standard results for projected hyperbolic PDE systems, and a principle of observability establishes that the field cannot be reconstructed from observations alone: persistence requires write-back.

The resulting system constitutes a unified framework in which continuous fields, discrete hypergraphs, and generative operators interact within a single coherent structure governed by variational principles, conserved under symmetry, and interpretable at every level from Hamiltonian mechanics to derived algebraic geometry. A generative system is not a model that produces worlds, but a gauge-invariant, stochastic, Hamiltonian field theory over a persistent semantic substrate, whose observable realizations arise as samples from a path integral, whose consistency is governed by homotopy-coherent gluing conditions, and whose dynamics are stabilized by explicit field-level conservation laws enforced through constrained neural projection. That this architecture is constructively realizable, and not merely formally consistent, is demonstrated by the prototype of Section 26, which instantiates the core closure $\mathcal{O}_\theta \to \mathcal{W} \to \mathcal{T} \to \mathcal{O}_\theta$ in executable form and exhibits the field's capacity to accumulate constraint from observation rather than discarding it.

## References

[1] R. Po, D. J. Zhang, A. Hertz, G. Wetzstein, N. Wadhwa, and N. Ruiz, *MultiGen: Level-Design for Editable Multiplayer Worlds in Diffusion Game Engines*, arXiv:2603.06679 (2026).

[2] J. Ho, A. Jain, and P. Abbeel, *Denoising Diffusion Probabilistic Models*, Advances in Neural Information Processing Systems (2020).

[3] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, *Score-Based Generative Modeling through Stochastic Differential Equations*, International Conference on Learning Representations (2021).

[4] P. W. Anderson, *Basic Notions of Condensed Matter Physics*, Addison-Wesley (1984).

[5] R. Landauer, *Irreversibility and Heat Generation in the Computing Process*, IBM Journal of Research and Development, 5(3), 183–191 (1961).

[6] C. H. Bennett, *The Thermodynamics of Computation—A Review*, International Journal of Theoretical Physics, 21(12), 905–940 (1982).

[7] S. Lloyd, *Ultimate Physical Limits to Computation*, Nature, 406, 1047–1054 (2000).

[8] S. Mac Lane, *Categories for the Working Mathematician*, Springer (1998).

[9] J. Lurie, *Higher Topos Theory*, Princeton University Press (2009).

[10] R. Hartshorne, *Algebraic Geometry*, Springer (1977).

[11] M. Kashiwara and P. Schapira, *Sheaves on Manifolds*, Springer (1990).

[12] J. C. Baez and J. P. Master, *Open Petri Nets*, Mathematical Structures in Computer Science (2011).

[13] T. Jacobson, *Thermodynamics of Spacetime: The Einstein Equation of State*, Physical Review Letters, 75(7), 1260–1263 (1995).

[14] E. Verlinde, *On the Origin of Gravity and the Laws of Newton*, Journal of High Energy Physics (2011).

[15] D. Ha and J. Schmidhuber, *World Models*, arXiv:1803.10122 (2018).

[16] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi, *Dream to Control: Learning Behaviors by Latent Imagination*, International Conference on Learning Representations (2020).

[17] D. Hafner et al., *Mastering Diverse Domains through World Models*, arXiv (2024).

[18] T. Brooks et al., *Video Generation Models as World Simulators*, OpenAI Research (2024).

[19] O. Bar-Tal et al., *Lumiere: A Space-Time Diffusion Model for Video Generation*, SIGGRAPH Asia (2024).

[20] B. Chen et al., *Diffusion Forcing: Next-Token Prediction Meets Full-Sequence Diffusion*, arXiv (2024).

[21] X. Huang et al., *Self-Forcing: Bridging the Train-Test Gap in Autoregressive Video Diffusion*, arXiv (2025).

[22] M. Guzdial and M. Riedl, *Game Level Generation from Gameplay Videos*, AIIDE (2021).

[23] S. W. Kim et al., *Learning to Simulate Dynamic Environments with GameGAN*, CVPR (2020).

[24] M. U. Nasir and J. Togelius, *Practical PCG through Large Language Models*, IEEE Conference on Games (2023).

[25] A. Anand et al., *Coherent Multi-Agent Simulation with Shared Latent States*, arXiv (2024).

[26] N. Wadhwa et al., *Genie: Generative Interactive Environments*, arXiv (2025).

[27] Y. Hong et al., *Relic: Interactive Video World Model with Long-Horizon Memory*, arXiv (2025).

[28] B. Deng et al., *Streetscapes: Large-Scale Consistent Street View Generation*, SIGGRAPH (2024).

[29] S. Kauffman, *The Origins of Order*, Oxford University Press (1993).

[30] H. T. Odum, *Self-Organization, Transformity, and Information*, Science (1988).

[31] N. Georgescu-Roegen, *The Entropy Law and the Economic Process*, Harvard University Press (1971).

[32] R. U. Ayres, *Industrial Ecology and the Flow of Materials*, Edward Elgar (1998).

[33] H. E. Daly, *Beyond Growth*, Beacon Press (1996).