Inference as Constraint Closure: A Variational, Spectral, and Sheaf-Theoretic Framework for Reconstruction

Flyxion

April 7, 2026

Abstract

We present a unified theoretical framework for inference grounded in the principle of constraint closure over structured state spaces. Moving beyond the traditional division between discriminative and generative methods, we formulate inference as the search for a configuration that satisfies a family of projection constraints while remaining admissible under underlying dynamics. This perspective is developed through three complementary formalisms: a variational formulation in which inference corresponds to minimizing a residual functional, a spectral formulation in which reconstruction emerges through modal compression and entropy redistribution, and a sheaf-theoretic formulation in which global consistency arises from the gluing of local sections.

We prove that reconstruction yields a unique solution when the induced projection map is injective over the feasible set, the residual functional is strictly convex, and cohomological obstruction vanishes. Under these conditions, inference corresponds to convergence toward a fixed point of a contractive consistency operator. Failure modes arise when degeneracy or flatness prevents unique closure.

This framework subsumes a wide range of systems, including model-based vision, variational inference, and multiscale data assimilation, and provides a unified account of perception as entropy-respecting constraint satisfaction. Rather than treating computation as symbolic manipulation or prediction as direct mapping, we interpret inference as a physical and mathematical process of reconstructing coherent structure from incomplete projections. This shift yields a general framework for inference that extends across physical, computational, and semantic domains.

1 Introduction: From Pose Estimation to Constraint Closure

The problem of 3D human pose estimation has traditionally been framed as either a regression problem over image features or an optimization problem over parametric models. In the hybrid formulation proposed by Ammar Qammaz, these two paradigms are combined: discriminative methods provide coarse localization, while generative methods refine estimates through model-based optimization. While this framing is operationally effective, it obscures a deeper structural interpretation.

The present essay advances the claim that the hybrid method is more accurately understood as a constraint closure system. Rather than estimating pose as a direct output, the system iteratively enforces consistency between multiple projections of an underlying latent configuration. The observed RGBD data, the neural network joint estimates, and the personalized 3D model together define a set of constraints whose joint satisfaction determines the admissible pose.

In this view, the algorithm does not "predict" pose. It reconstructs it as the unique configuration that satisfies a collection of observational and structural constraints. This reinterpretation aligns naturally with the Relativistic Scalar-Vector Plenum (RSVP) framework and the Yarncrawler formalism, in which inference is framed as the search for a fixed point within a constrained equivalence class of field configurations.

2 Related Work

The framework developed in this work intersects multiple established traditions across physics, statistics, and computer science. We briefly situate the present approach relative to these domains, emphasizing both continuity and distinction.

2.1 Variational Inference and Energy-Based Models

Variational formulations of inference have a long history in statistical mechanics and machine learning. The use of energy functionals to define probability distributions and optimization objectives is central to both maximum entropy methods [2] and modern variational inference [9, 10]. Energy-based models similarly interpret inference as the minimization of a scalar potential over configuration space.

The present framework adopts this variational perspective but extends it in two directions. First, the residual functional $\mathcal{R}(X)$ is not treated purely as a statistical objective but as a physically grounded quantity incorporating admissibility constraints derived from underlying dynamics. Second, inference is not framed solely as optimization or sampling, but as *constraint closure* under a family of projection operators. This shifts emphasis from likelihood maximization to consistency enforcement across heterogeneous observations.

2.2 Statistical Mechanics and Thermodynamic Inference

The interpretation of inference in thermodynamic terms has been explored extensively, particularly in the context of free energy minimization and probabilistic inference [25]. The correspondence between Gibbs distributions and variational objectives provides a bridge between statistical physics and learning systems.

Our formulation aligns with this tradition by introducing a partition function over admissible configurations and interpreting reconstruction as ground state selection in the low-temperature limit. However, we emphasize a non-equilibrium perspective: inference is modeled as a driven dissipative process subject to external constraints, rather than relaxation toward unconstrained equilibrium. This places the framework closer to non-equilibrium statistical mechanics than to classical equilibrium formulations.

2.3 Partial Differential Equations and Field Theory

The use of variational principles and Euler–Lagrange equations to describe physical systems is foundational in field theory [6]. Inverse problems and data assimilation similarly employ PDE-constrained optimization to reconstruct hidden states from partial observations.

The present work extends this paradigm by treating observational constraints as source terms within a field-theoretic formulation. The resulting system can be interpreted as a forced field equation, where reconstruction corresponds to finding a configuration satisfying both intrinsic dynamics and extrinsic constraints. This unifies inverse problems in physics with inference in computational systems.

2.4 Multiscale Methods and Renormalization

Multiscale structure plays a central role in both physics and signal processing. Renormalization group methods describe how systems evolve across scales, while wavelet and spectral decompositions provide practical tools for representing multiscale phenomena [20].

Our spectral formulation connects reconstruction to modal decomposition, while a TARTAN-inspired tiling structure introduces a complementary spatial decomposition. Together, these define dual multiscale representations—frequency and spatial—that govern the flow of information and entropy during inference. This dual perspective extends classical renormalization ideas to constraint-based reconstruction systems.

2.5 Sheaf Theory and Topological Consistency

Sheaf-theoretic methods have been increasingly applied to problems involving local-to-global consistency, particularly in topology and data analysis [16, 18]. These approaches formalize the conditions under which local data can be glued into a globally consistent structure.

The present framework adopts this perspective by interpreting reconstruction as the existence of a global section satisfying compatibility conditions across overlapping domains. The introduction of cohomological obstruction provides a precise characterization of failure modes, linking topological inconsistency to degeneracy in inference.

2.6 Computer Vision and Inverse Problems

In computer vision, model-based reconstruction methods have long combined observational data with geometric priors [22, 23]. Hybrid approaches integrating discriminative and generative components have been shown to improve robustness and accuracy in tasks such as pose estimation.

Our reinterpretation of such systems places them within a broader theoretical context. Rather than viewing hybridization as a heuristic combination of techniques, we show that it arises naturally from the requirement to enforce multiple constraints simultaneously. This recasts vision systems as instances of a general constraint closure architecture.

2.7 Summary of Contributions

While each of the above traditions addresses aspects of inference, they are typically developed in isolation. Variational methods emphasize optimization, statistical mechanics emphasizes distributions, field theory emphasizes dynamics, and sheaf theory emphasizes consistency. The contribution of this work is to unify these perspectives within a single framework, in which variational optimization corresponds to energy minimization, statistical inference corresponds to sampling over admissible configurations, field dynamics enforce admissibility, spectral and multiscale methods regulate complexity, and sheaf-theoretic conditions govern global consistency. Rather than replacing these frameworks, the present work provides a unifying perspective that makes explicit their underlying structural commonalities, yielding a general framework for reconstruction that extends across physical, computational, and semantic domains.

3 Field-Theoretic Representation

Definition 1 (RSVP State)**.** *A system state is defined as a triple*

$$X = (\Phi, \mathbf{v}, S)$$

where Φ is a scalar field, $\mathbf{v}$ is a vector field, and S is an entropy field over a domain Ω.

Within the RSVP framework, a physical or perceptual system is represented as such a triple, where scalar potential, vector flow, and entropy define the state of a structured plenum. Perception corresponds to the reconstruction of such states under partial observations.

The pose estimation problem can be embedded into this framework by identifying the human body configuration h with a corresponding RSVP state X_h. Observations such as RGB and depth images define projection operators

$$\Pi_i : X_h \to \mathcal{Y}_i$$

mapping latent configurations into observable spaces. Neural network outputs and model constraints define additional projections, producing a family of observational conditions

$$\Pi_i(X_h) \approx y_i.$$

The objective function described in the original thesis—combining depth consistency and joint alignment—can then be interpreted as a discretized action functional measuring deviation from projection consistency. Minimizing this functional corresponds to driving the system toward a low-entropy configuration consistent with all available observations. Thus, pose estimation becomes equivalent to finding a field configuration X^* such that:

$$\Pi_i(X^*) \approx y_i \quad \forall i.$$

4 Constraint Sets and Feasible Configurations

Definition 2 (Feasible Set). *The feasible set $\mathcal{F}$ is defined as*

$$\mathcal{F} = \{[X] \in \mathcal{A} \mid \|\Pi_i(X) - y_i\| \leq \epsilon_i \; \forall i\},$$

where $\mathcal{A}$ is the admissible configuration space modulo equivalence.

This formulation corresponds directly to the reconstruction paradigm: observations do not determine a unique configuration directly but constrain the solution space. This corresponds exactly to a two-stage projection process: a coarse projection from raw observations into an approximate pose manifold, followed by a refinement operator that iteratively reduces constraint violations.

In Yarncrawler terms, the refinement operator is a consistency operator acting on admissible configurations. The fixed point of this operator represents the reconstructed pose. The necessity of hybridization—explicitly argued in the original thesis—is thus revealed as a structural requirement: no single projection operator is sufficient to enforce all constraints simultaneously.

5 Hybrid Inference as Projection Composition

Definition 3 (Consistency Operator). *A consistency operator $\mathcal{C}$ is a mapping*

$$\mathcal{C} : \mathcal{A} \to \mathcal{A}$$

that reduces projection residuals:

$$\|\Pi_i(\mathcal{C}(X)) - y_i\| \leq \|\Pi_i(X) - y_i\|.$$

The reconstructed pose corresponds to a fixed point of $\mathcal{C}$.

Theorem 1 (Fixed Point Reconstruction). *If $\mathcal{C}$ is contractive over $\mathcal{F}$, then there exists a unique fixed point $[X^*]$ such that*

$$\mathcal{C}(X^*) = X^*,$$

and X^* satisfies all projection constraints.

This establishes pose estimation as a convergence process rather than a direct prediction.

6 Identifiability and Personalized Models

Generic body models introduce degeneracy in the inverse mapping from observations to configurations. Multiple latent states may produce indistinguishable projections.

Definition 4 (Projection Injectivity)**.** *The induced projection map*

$$\tilde{\Pi} : X/{\sim} \to \prod_i \mathcal{Y}_i$$

is injective over $\mathcal{F}$ if distinct configurations produce distinct observations.

Personalized models reduce this degeneracy by embedding subject-specific structure into the constraint set, increasing the injectivity of the induced projection map.

Proposition 1. *Personalized geometry increases the injectivity of $\tilde{\Pi}$ over $\mathcal{F}$, improving identifiability and stability of reconstruction.*

Thus, personalization functions as symmetry breaking in the inverse problem. A central empirical result of the original thesis is that personalized body models significantly improve tracking accuracy. Within the present framework, this improvement is not empirically contingent but structurally necessary: it corresponds to satisfying an identifiability condition.

7 Timescale Separation

The thesis introduces a two-stage architecture separating 3D scanning from tracking. While motivated pragmatically, this design can be interpreted as a separation of timescales in the RSVP dynamics.

Scanning corresponds to the slow evolution of the underlying field geometry Φ, while tracking corresponds to fast updates in configuration under fixed geometry. By decoupling these processes, the system prevents error propagation from high-frequency fluctuations into the structural substrate.

Proposition 2. *Decoupling slow structural evolution from fast inference dynamics prevents error accumulation and stabilizes reconstruction.*

In field-theoretic terms, this enforces a quasi-static approximation on the manifold of admissible configurations, stabilizing the inference process. The model becomes a slowly varying constraint surface, while pose estimation is a rapid flow constrained to that surface.

8 Optimization as Entropy Descent

The use of stochastic optimization methods such as Particle Swarm Optimization can be reinterpreted as entropy descent in the configuration space. Each candidate pose represents a microstate, and the objective function defines an energy landscape over these states.

Definition 5 (Entropy-Reducing Flow)**.** *A trajectory* $X(t)$ *is entropy-reducing if*

$$\frac{d}{dt}S(X(t)) \leq 0$$

under the dynamics induced by $-\nabla E$*.*

The iterative refinement process reduces entropy by concentrating probability mass around configurations that satisfy the constraints. GPU acceleration enables the evaluation of large ensembles of hypotheses, effectively sampling the space of possible configurations and driving convergence toward the global minimum.

This aligns with the RSVP view of inference as a thermodynamic process: perception is not instantaneous computation, but a relaxation toward a low-entropy state consistent with environmental constraints.

9 Statistical Mechanics Interpretation

The reconstruction problem admits a natural interpretation in statistical mechanics. Define a Gibbs distribution over configurations:

$$p(X) = \frac{1}{Z} \exp\left(-\beta \mathcal{R}(X)\right),$$

where β is an inverse temperature parameter and Z is the partition function:

$$Z = \int_{\mathcal{A}} \exp\left(-\beta \mathcal{R}(X)\right) \, dX.$$

Definition 6 (Free Energy Functional)**.** *The free energy is defined as:*

$$F = -\frac{1}{\beta} \log Z.$$

Proposition 3. *In the zero-temperature limit* $\beta \to \infty$*, the Gibbs distribution concentrates on minimizers of* $\mathcal{R}(X)$*, recovering the deterministic reconstruction problem.*

Thus, inference can be interpreted as sampling from a thermodynamic ensemble, with reconstruction corresponding to ground state selection.

9.1 Entropy–Energy Decomposition

The residual functional $\mathcal{R}(X)$ plays the role of an effective energy, while entropy arises from the volume of admissible configurations.

Proposition 4. *The variational problem is equivalent to minimizing the free energy:*

$$F = \mathbb{E}_p[\mathcal{R}(X)] - \frac{1}{\beta} H(p),$$

where $H(p)$ *is the Shannon entropy.*

This establishes a direct equivalence with variational inference and maximum entropy principles.

10 RSVP PDE Embedding and Variational Structure

We now formalize the reconstruction process within the full RSVP field dynamics. Let the system state be

$$X = (\Phi, \mathbf{v}, S)$$

defined over a compact domain $\Omega \subset \mathbb{R}^3$, with

$$\Phi \in H^1(\Omega), \quad \mathbf{v} \in H^1(\Omega, \mathbb{R}^3), \quad S \in L^2(\Omega).$$

We define a residual functional capturing both observational and dynamical constraints:

$$\mathcal{R}(X) = \sum_i \|\Pi_i(X) - y_i\|^2 + \lambda_\Phi \|\mathcal{L}_\Phi(\Phi, \mathbf{v}, S)\|^2 + \lambda_v \|\mathcal{L}_v(\Phi, \mathbf{v}, S)\|^2 + \lambda_S \|\mathcal{L}_S(\Phi, \mathbf{v}, S)\|^2,$$

where $\mathcal{L}_\Phi, \mathcal{L}_v, \mathcal{L}_S$ are RSVP field operators encoding admissibility conditions such as flow consistency, entropy balance, and coupling constraints.

The reconstruction problem becomes:

$$X^* = \arg\min_{X \in \mathcal{A}} \mathcal{R}(X).$$

This formulation embeds pose estimation into a variational field theory, where admissible configurations correspond to approximate solutions of the RSVP PDE system.

11 Action Principle and Euler–Lagrange Equations

The residual functional $\mathcal{R}(X)$ defines an action:

$$\mathcal{S}[X] = \int_\Omega \mathcal{L}(X, \nabla X)\, dx + \sum_i \|\Pi_i(X) - y_i\|^2.$$

Definition 7 (Lagrangian Density)**.** *The Lagrangian $\mathcal{L}$ encodes intrinsic RSVP dynamics:*

$$\mathcal{L} = \mathcal{L}_\Phi + \mathcal{L}_v + \mathcal{L}_S + \mathcal{L}_{\text{coupling}}.$$

Proposition 5. *Stationary points of $\mathcal{S}$ satisfy Euler–Lagrange equations:*

$$\frac{\delta \mathcal{S}}{\delta X} = 0.$$

Thus, reconstruction corresponds to solving a constrained field theory with external forcing terms induced by observations.

11.1 Interpretation as Forced Field Dynamics

Observational constraints act as source terms:

$$\mathcal{L} \;\rightarrow\; \mathcal{L} + J(X),$$

where J encodes projection mismatch.

Proposition 6. *Reconstruction is equivalent to finding a field configuration satisfying:*

$$\mathcal{E}(X) = J,$$

where $\mathcal{E}$ is the Euler–Lagrange operator.

This establishes a direct analogy with classical field theory under external forcing.

12 Gradient Flow and Lyapunov Functional

We define a gradient flow over the configuration space:

$$\frac{dX}{dt} = -\nabla \mathcal{R}(X).$$

Definition 8 (Lyapunov Functional)**.** *The functional $\mathcal{R}(X)$ is a Lyapunov functional if along trajectories $X(t)$:*

$$\frac{d}{dt}\mathcal{R}(X(t)) \leq 0.$$

Proposition 7. *Under the gradient flow dynamics, $\mathcal{R}(X)$ is non-increasing:*

$$\frac{d}{dt}\mathcal{R}(X(t)) = -\|\nabla \mathcal{R}(X)\|^2 \leq 0.$$

Thus, reconstruction corresponds to a dissipative process driving the system toward a stable minimum of $\mathcal{R}$.

Proposition 8. *If $\mathcal{R}$ is strictly convex over $\mathcal{F}$, then the gradient flow converges to a unique global minimizer X^*.*

This provides a formal stability guarantee for the reconstruction process.

13 Stochastic Dynamics and Fokker–Planck Evolution

The deterministic gradient flow can be generalized to stochastic dynamics via a Langevin equation:

$$dX_t = -\nabla \mathcal{R}(X_t)\, dt + \sqrt{2D}\, dW_t,$$

where W_t is a Wiener process and D is a diffusion coefficient.

Proposition 9. *The probability density $p(X, t)$ evolves according to the Fokker–Planck equation:*

$$\frac{\partial p}{\partial t} = \nabla \cdot (p \nabla \mathcal{R}) + D\Delta p.$$

Proposition 10. *The stationary distribution of this dynamics is:*

$$p^*(X) \propto \exp\left(-\frac{\mathcal{R}(X)}{D}\right).$$

Thus, stochastic inference corresponds to diffusion in configuration space biased by constraint violations.

13.1 Ergodicity and Convergence

Proposition 11. *If the dynamics is ergodic and $\mathcal{R}$ is confining, time averages converge to ensemble averages under p^*.*

This provides a statistical guarantee for reconstruction via sampling-based methods.

14 Consistency Operator and Fixed Point Structure

Definition 9 (Gradient Consistency Operator)**.** *Define the operator*

$$\mathcal{C}(X) = X - \eta \nabla \mathcal{R}(X)$$

for step size $\eta > 0$.

This operator represents a discrete-time approximation of the gradient flow.

Proposition 12. *A configuration X^* is a fixed point of $\mathcal{C}$ if and only if*

$$\nabla \mathcal{R}(X^*) = 0,$$

i.e., X^* is a stationary point of the residual functional.

Thus, fixed points of $\mathcal{C}$ correspond exactly to reconstructed configurations.

15 Identifiability Theorem in RSVP Reconstruction

Theorem 2 (RSVP Identifiability Theorem)**.** *Let $\mathcal{F}$ be the feasible set defined by observational constraints, and let $\mathcal{R}$ be strictly convex over $\mathcal{F}$. Assume the induced projection map*

$$\tilde{\Pi} : X/\sim \to \prod_i \mathcal{Y}_i$$

is injective on $\mathcal{F}$.

Then:

1. *There exists a unique minimizer $[X^*] \in \mathcal{F}$,*
2. *$[X^*]$ is the unique fixed point of the consistency operator $\mathcal{C}$,*
3. *The gradient flow converges to $[X^*]$ from any initial condition in $\mathcal{F}$.*

Proof. Strict convexity ensures uniqueness of the minimizer. Injectivity of $\tilde{\Pi}$ guarantees that distinct equivalence classes cannot satisfy the same observational constraints, eliminating degeneracy. Convergence follows from standard results on gradient flows in Hilbert spaces: if $\mathcal{R}$ is coercive, continuously differentiable, and has Lipschitz-continuous gradient, then trajectories converge to the unique minimizer. □

16 Failure Modes and Degeneracy

Two fundamental failure modes arise when the conditions of the theorem are violated.

Definition 10 (Projection Degeneracy)**.** *Projection degeneracy occurs when $\tilde{\Pi}$ is not injective over $\mathcal{F}$, leading to multiple indistinguishable configurations.*

Definition 11 (Regularizer Flatness)**.** *Regularizer flatness occurs when $\mathcal{R}$ is not strictly convex, allowing multiple local minima.*

Proposition 13. *In the presence of either projection degeneracy or regularizer flatness, reconstruction may fail to produce a unique solution.*

Personalized models mitigate projection degeneracy by increasing injectivity, while dynamical regularization mitigates flatness by enforcing curvature in $\mathcal{R}$.

17 Spectral Formulation and Modal Decomposition

We now reformulate the RSVP reconstruction problem in the spectral domain. Let Ω be a compact domain with appropriate boundary conditions. We expand each field in an orthonormal basis $\{\psi_k\}_{k=0}^{\infty}$:

$$\Phi(x) = \sum_k \hat{\Phi}_k \psi_k(x), \quad \mathbf{v}(x) = \sum_k \hat{\mathbf{v}}_k \psi_k(x), \quad S(x) = \sum_k \hat{S}_k \psi_k(x).$$

The state X is represented by spectral coefficients $\hat{X} = (\hat{\Phi}_k, \hat{\mathbf{v}}_k, \hat{S}_k)_{k=0}^{\infty}$, and the residual functional becomes:

$$\mathcal{R}(\hat{X}) = \sum_i \|\Pi_i(\hat{X}) - y_i\|^2 + \sum_k \left(\lambda_\Phi \|\hat{\mathcal{L}}_\Phi(\hat{X})_k\|^2 + \lambda_v \|\hat{\mathcal{L}}_v(\hat{X})_k\|^2 + \lambda_S \|\hat{\mathcal{L}}_S(\hat{X})_k\|^2 \right).$$

Proposition 14. *Under suitable regularity conditions, the spectral representation diagonalizes linear components of the RSVP operators, allowing independent or weakly coupled evolution of modes.*

This yields a modal interpretation of reconstruction: low-frequency modes encode global structure, while high-frequency modes encode fine detail.

Definition 12 (Spectral Entropy). *Define the spectral entropy:*

$$S_{\text{spec}} = \sum_k w_k \|\hat{X}_k\|^2,$$

where weights w_k increase with frequency.

Proposition 15. *Gradient flow preferentially suppresses high-frequency modes, yielding an implicit regularization toward smooth, low-complexity configurations.*

In practice, we approximate $\hat{X} \approx (\hat{X}_0, \dots, \hat{X}_N)$, yielding a finite-dimensional system suitable for GPU implementation.

18 Renormalization and Multiscale Structure

The spectral decomposition induces a natural hierarchy of scales. Let $X^{(k)}$ denote the contribution from modes up to frequency k.

Definition 13 (Coarse-Graining Operator). *Define a coarse-graining operator:*

$$\mathcal{G}_\Lambda(X) = \sum_{k \le \Lambda} \hat{X}_k \psi_k.$$

Proposition 16. *Coarse-graining reduces high-frequency entropy while preserving low-frequency structure.*

18.1 Renormalized Residual

Define the effective residual at scale Λ:

$$\mathcal{R}_\Lambda(X) = \mathcal{R}(\mathcal{G}_\Lambda(X)).$$

Proposition 17. *As Λ increases, $\mathcal{R}_\Lambda$ converges to $\mathcal{R}$, yielding a multiscale refinement process.*

This provides a renormalization-group interpretation of reconstruction, where inference proceeds from coarse global structure to fine detail.

18.2 Connection to TARTAN

The TARTAN framework can be interpreted as a spatial analogue of renormalization, decomposing the domain into tiles rather than frequencies.

Proposition 18. *Spectral and spatial decompositions form dual renormalization schemes for reconstruction.*

19 Sheaf-Theoretic Formulation and Gluing

We now reinterpret the reconstruction process in categorical terms. Let Ω be covered by open sets $\{U_i\}$. Define a presheaf $\mathcal{S}$ assigning to each U_i the set of admissible local configurations:

$$\mathcal{S}(U_i) = \{X|_{U_i} \text{ consistent with local observations}\}.$$

Restriction maps $\rho_{ij} : \mathcal{S}(U_i) \to \mathcal{S}(U_j)$ encode compatibility between overlapping regions.

Definition 14 (Gluing Condition)**.** *A family $\{s_i\}$ satisfies the gluing condition if:*

$$\rho_{ij}(s_i) = \rho_{ji}(s_j) \quad \forall i, j.$$

Proposition 19. *A global section $s \in \mathcal{S}(\Omega)$ exists if and only if all local sections satisfy pairwise compatibility constraints.*

This corresponds exactly to projection consistency across observations. We measure failure of global consistency via the Čech 1-cocycle:

$$\delta_{ij} = \rho_{ij}(s_i) - \rho_{ji}(s_j).$$

Definition 15 (Obstruction Class)**.** *The obstruction class $[\delta] \in H^1(\Omega, \mathcal{S})$ measures failure of global consistency.*

Theorem 3. *A global reconstruction exists if and only if $[\delta] = 0$.*

Thus, reconstruction corresponds to annihilating cohomological obstruction.

Proposition 20. *Repeated application of $\mathcal{C}$ corresponds to iterative cohomological repair, converging to a global section when obstruction vanishes.*

This establishes a direct equivalence:

$$\text{Constraint closure} \quad \longleftrightarrow \quad \text{Sheaf gluing}.$$

20 Categorical Interpretation

Let **Obs** be the category of observational spaces and **Field** the category of RSVP configurations modulo equivalence. Projection operators define a functor:

$$\Pi : \mathbf{Field} \to \mathbf{Obs}.$$

The reconstruction problem seeks a section of this functor consistent with observed data.

Theorem 4 (Functorial Reconstruction)**.** *If the projection functor is faithful and the constraint system is consistent, then reconstruction corresponds to selecting a unique object in the fiber over y.*

21 Synthesis: Spectral-Sheaf Duality

The spectral formulation decomposes reconstruction into modes, while the sheaf formulation decomposes it into regions. These represent dual factorizations:

$$\text{Global structure} \;\leftrightarrow\; \text{spectral modes}, \quad \text{local consistency} \;\leftrightarrow\; \text{sheaf sections.}$$

Proposition 21. *Reconstruction succeeds when both spectral compression and sheaf gluing conditions are satisfied.*

This establishes a duality between frequency-domain regularization and spatial consistency, unifying two fundamental perspectives on inference.

22 A One-Dimensional Toy Model of Constraint Closure

To make the framework concrete, we construct a minimal one-dimensional example in which all components of the theory can be analyzed explicitly.

22.1 Setup

Let $\Omega = [0,1]$ and consider a scalar field $\Phi(x) \in H^1(\Omega)$. We ignore vector and entropy fields for simplicity and treat entropy implicitly through the functional.

Define two projection operators:

$$\Pi_1(\Phi) = \int_0^1 \Phi(x)\,dx, \quad \Pi_2(\Phi) = \Phi(x_0),$$

where $x_0 \in (0,1)$ is fixed. Let the observed values be:

$$\Pi_1(\Phi) = m, \quad \Pi_2(\Phi) = c.$$

These correspond to a global constraint (mean value) and a local constraint (pointwise observation).

22.2 Residual Functional

Define the residual:

$$\mathcal{R}(\Phi) = \left(\int_0^1 \Phi(x)\,dx - m\right)^2 + (\Phi(x_0) - c)^2 + \lambda \int_0^1 |\nabla\Phi(x)|^2\,dx.$$

The final term enforces smoothness and plays the role of an RSVP admissibility condition.

22.3 Euler–Lagrange Equation

Taking the functional derivative yields:

$$-\lambda\Delta\Phi(x)+2\left(\int_0^1 \Phi(x')\,dx'-m\right)+2\delta(x-x_0)(\Phi(x_0)-c)=0.$$

Proposition 22. *The solution Φ^* satisfies a Poisson-type equation with a global forcing term and a localized source at x_0, explicitly realizing the claim that observations act as forcing terms in a field equation.*

22.4 Closed-Form Solution Structure

Let $\bar{\Phi}=\int_0^1 \Phi(x)\,dx$. Then the equation reduces to:

$$-\lambda\Delta\Phi(x)=2(m-\bar{\Phi})+2(c-\Phi(x_0))\delta(x-x_0).$$

The solution can be expressed using the Green's function $G(x,x')$:

$$\Phi(x)=A+B\,G(x,x_0),$$

where constants A, B are determined by the constraints.

Proposition 23. *There exists a unique solution Φ^* provided $\lambda>0$, ensuring strict convexity of $\mathcal{R}$.*

22.5 Gradient Flow Dynamics

The gradient flow is:

$$\frac{\partial\Phi}{\partial t}=-\frac{\delta\mathcal{R}}{\delta\Phi}.$$

Proposition 24. *The functional $\mathcal{R}(\Phi(t))$ decreases monotonically along trajectories, so the system converges to Φ^*.*

22.6 Spectral Interpretation

Expand Φ in Fourier modes:

$$\Phi(x)=\sum_{k=0}^{\infty}\hat{\Phi}_k\sin(k\pi x).$$

Then:

$$\int_0^1 |\nabla\Phi|^2\,dx=\sum_k k^2\hat{\Phi}_k^2.$$

Proposition 25. *Higher-frequency modes are penalized more strongly, leading to spectral smoothing. Reconstruction corresponds to suppressing high-frequency entropy while satisfying constraints.*

22.7 Statistical Mechanics Interpretation

Define:

$$p(\Phi) \propto \exp(-\beta \mathcal{R}(\Phi)).$$

Proposition 26. *As $\beta \to \infty$, $p(\Phi)$ concentrates on Φ^*. For finite β, fluctuations around Φ^* are Gaussian with covariance determined by the inverse Hessian of $\mathcal{R}$.*

22.8 Identifiability and Degeneracy

If the point constraint Π_2 is removed, the system becomes underdetermined.

Proposition 27. *Without Π_2, infinitely many functions satisfy the mean constraint, leading to degeneracy. This explicitly demonstrates the necessity of multiple projections for identifiability.*

22.9 Summary

This toy model instantiates all components of the general framework. The residual functional defines an energy landscape; gradient flow implements entropy descent; observations act as forcing terms in a field equation; spectral decomposition enforces smoothness; and multiple projections ensure identifiability. Even in one dimension, reconstruction emerges as constraint closure under variational dynamics. This example demonstrates that the abstract framework admits concrete realizations with explicitly solvable structure, bridging formal theory and implementable systems.

23 Discussion: From Toy Model to Full Systems

The simplicity of the one-dimensional model obscures none of the essential structure. Each component scales naturally: scalar fields generalize to coupled RSVP fields; point constraints generalize to multimodal observations; the Laplacian generalizes to full RSVP operators; and Fourier modes generalize to high-dimensional spectral bases. The full framework can thus be understood as a high-dimensional generalization of this minimal system, with all essential mechanisms already present in the one-dimensional case.

24 Grand Unification: Inference as Entropy-Respecting Constraint Closure

We now synthesize the preceding constructions into a unified theoretical framework. The central claim is that perception, cognition, and computation can all be described as instances of a single principle: *entropy-respecting constraint closure over structured fields*.

24.1 Unified Object of Study

Let Ω be a domain equipped with an RSVP field configuration $X = (\Phi, \mathbf{v}, S)$, and let $\{\Pi_i\}$ be a family of projection operators encoding observational, semantic, or internal constraints. We define the global reconstruction problem as:

$$X^* = \arg\min_{X \in \mathcal{A}} \mathcal{R}(X),$$

subject to admissibility under RSVP dynamics and projection consistency. This formulation is agnostic to domain: the same structure applies whether X encodes a human body configuration, a physical field, or a semantic state.

24.2 The Yarncrawler Closure Principle

Within the Yarncrawler framework, reconstruction is governed by a consistency operator $\mathcal{C} : \mathcal{A} \to \mathcal{A}$.

Theorem 5 (Yarncrawler Closure Principle)**.** *A system achieves reconstruction if and only if there exists a unique fixed point $[X^*]$ of $\mathcal{C}$ within the feasible set $\mathcal{F}$.*

This principle provides the core mechanism: inference is the iterative elimination of inconsistency.

24.3 CLIO as Recursive Inference Functor

The Cognitive Loop via In-Situ Optimization (CLIO) can be interpreted as a recursive refinement of the consistency operator. Let C_t denote the consistency operator at iteration t. Then CLIO defines a sequence $X_{t+1} = C_t(X_t)$, where C_t itself is updated based on prior residuals.

Definition 16 (CLIO Operator)**.** *CLIO is a higher-order operator:*

$$\mathcal{C}_{\text{CLIO}} : (X_t, C_t) \mapsto (X_{t+1}, C_{t+1}),$$

where both state and inference mechanism co-evolve.

Proposition 28. *CLIO induces a meta-stable trajectory in the space of operators, converging toward increasingly efficient constraint closure.*

Thus, CLIO lifts inference from a fixed operator to a self-optimizing process.

24.4 TARTAN as Multiscale Constraint Tiling

The TARTAN framework introduces a multiscale decomposition of the domain Ω into tiles $\{U_i\}$ with annotated noise and trajectory buffers. Each tile carries local constraints $\mathcal{F}_i = \{X|_{U_i} \mid \text{local consistency}\}$.

Proposition 29. *TARTAN implements a hierarchical sheaf structure, where global reconstruction emerges from recursive gluing of local constraint systems.*

Annotated noise encodes uncertainty structure, while trajectory buffers preserve temporal consistency across updates.

24.5 Simulated Agency as Sparse Projection

Simulated Agency arises when the system does not directly observe the full constraint set but instead samples sparse projections.

Definition 17 (Sparse Projection Engine)**.** *A system exhibits simulated agency if it reconstructs* X^* *from a subset of projections by leveraging structural priors and iterative refinement.*

Proposition 30. *Simulated agency corresponds to reconstruction under underdetermined constraints, where the feasible set is shaped by prior-dominant attractors.*

24.6 Entropy and Thermodynamic Consistency

Definition 18 (Entropy-Respecting Computation)**.** *A computational process is entropy-respecting if it reduces local inconsistency while preserving global thermodynamic constraints.*

Reconstruction is an entropy-reducing process constrained by admissibility conditions, and entropy is not eliminated globally but redistributed across scales and representations. This aligns inference with physical law: computation is not abstract symbol manipulation, but structured dissipation.

24.7 Unified Theorem of Inference

Theorem 6 (Unified Inference Theorem)**.** *Let* $\mathcal{A}$ *be a space of admissible RSVP configurations,* $\{\Pi_i\}$ *a family of projections, and* $\mathcal{R}$ *a strictly convex residual functional.*

Then reconstruction yields a unique configuration $[X^*]$ *if and only if:*

1. *The induced projection map is injective over* $\mathcal{F}$*,*
2. *The consistency operator* $\mathcal{C}$ *is contractive,*
3. *The system admits a Lyapunov functional* $\mathcal{R}$*,*
4. *Sheaf obstruction vanishes (*$H^1 = 0$*),*
5. *Spectral energy is bounded and convergent.*

Proof. Uniqueness follows from injectivity and convexity. Convergence follows from contractivity and standard Hilbert-space gradient flow results under coercivity and Lipschitz-continuous gradient. Existence of a global section follows from vanishing cohomology. Spectral boundedness ensures finite-energy solutions. □

24.8 Philosophical Implication

This theorem unifies all prior constructions: RSVP provides the field substrate and dynamics; Yarncrawler provides the fixed-point structure; CLIO provides recursive refinement; TARTAN provides multiscale decomposition and memory; Simulated Agency provides sparse inference under incomplete constraints. Together, they define a single architecture:

a system that reconstructs coherent structure from partial projections through entropy-respecting constraint closure.

The consequence is a shift in ontology. Systems do not "contain" representations; they *stabilize* them. Reality, perception, and cognition are not separate domains, but different regimes of the same process: the emergence of consistent structure under constraint. To perceive is to solve a reconstruction problem; to act is to modify constraints; to think is to explore the space of admissible closures.

25 Non-Equilibrium Thermodynamics of Inference

Reconstruction is inherently a non-equilibrium process. The system evolves toward a constrained steady state rather than thermodynamic equilibrium.

Definition 19 (Effective Entropy Production Rate)**.** *Define the effective entropy production:*

$$\sigma = -\frac{d}{dt}\mathcal{R}(X(t)).$$

Proposition 31. *Inference corresponds to non-negative effective entropy production,*

$$\sigma \geq 0,$$

until a steady state is reached.

25.1 Detailed Balance and Constraint Violation

Proposition 32. *At equilibrium, detailed balance holds if and only if all projection constraints are satisfied.*

Thus, constraint closure corresponds to a steady state of zero net inconsistency flow.

25.2 Physical Interpretation

The system exchanges entropy with its environment through observations. Constraints act as boundary conditions that drive the system away from equilibrium, while reconstruction restores consistency.

Proposition 33. *Inference is a driven dissipative process that minimizes constraint violation subject to admissibility.*

26 Conclusion

We have developed a unified framework for inference grounded in entropy-respecting constraint closure over structured field manifolds. Beginning from the observation that hybrid pose estimation implicitly implements a projection-consistent inference operator, we have shown that this structure generalizes to a complete theory of reconstruction encompassing variational, statistical mechanical, spectral, categorical, and thermodynamic perspectives.

The Unified Inference Theorem establishes the conditions under which reconstruction yields a unique, stable solution. The failure mode analysis clarifies when and why these conditions break down. The one-dimensional toy model confirms that all essential mechanisms are already present in the simplest non-trivial case.

This framework dissolves traditional boundaries between perception, computation, and physical law. All become instances of a single principle: the search for stable configurations under constraint. The resulting picture is not one of computation as abstract manipulation, but of computation as a physical process of entropy-guided reconstruction, operating across scales and domains.

In this sense, perception is not the passive reception of information, but the active reconstruction of coherent structure from incomplete projections under constraint.

References

[1] C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, 27(3):379–423, 1948.

[2] E. T. Jaynes, "Information Theory and Statistical Mechanics," *Physical Review*, 106(4):620–630, 1957.

[3] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Wiley, 2006.

[4] A. N. Kolmogorov, "Three Approaches to the Quantitative Definition of Information," *Problems of Information Transmission*, 1(1):1–7, 1965.

[5] H. Risken, *The Fokker–Planck Equation: Methods of Solution and Applications*. Springer, 1989.

[6] L. C. Evans, *Partial Differential Equations*, 2nd ed. American Mathematical Society, 2010.

[7] C. Villani, *Optimal Transport: Old and New*. Springer, 2009.

[8] S. Amari, *Information Geometry and Its Applications*. Springer, 2016.

[9] M. J. Wainwright and M. I. Jordan, "Graphical Models, Exponential Families, and Variational Inference," *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.

[10] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational Inference: A Review for Statisticians," *Journal of the American Statistical Association*, 112(518):859–877, 2017.

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

[12] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge University Press, 2009.

[13] W. Rudin, *Functional Analysis*, 2nd ed. McGraw–Hill, 1991.

[14] M. Reed and B. Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*. Academic Press, 1980.

[15] R. Bott and L. W. Tu, *Differential Forms in Algebraic Topology*. Springer, 1982.

[16] M. Kashiwara and P. Schapira, *Categories and Sheaves*. Springer, 2006.

[17] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed. Springer, 1998.

[18] R. Ghrist, *Elementary Applied Topology*. Createspace, 2014.

[19] R. N. Bracewell, *The Fourier Transform and Its Applications*, 3rd ed. McGraw–Hill, 2000.

[20] S. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, 2009.

[21] B. K. P. Horn, *Robot Vision*. MIT Press, 1986.

[22] R. Szeliski, *Computer Vision: Algorithms and Applications*. Springer, 2010.

[23] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2004.

[24] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.

[25] M. Mézard and A. Montanari, *Information, Physics, and Computation*. Oxford University Press, 2009.