

Yarncrawler: World-State Reconstruction as Constraint Closure over a Generative Field

Flyxion

March 31, 2026

Abstract

This paper introduces Yarncrawler as a framework for reconstructing physical world-states from multi-modal sensor observations under continuous dynamical constraint. The central object is a consistency operator \mathcal{C} acting on the quotient configuration space \mathcal{X}/\sim of coupled Relativistic Scalar-Vector Plenum field triples (Φ, \mathbf{v}, S) . Each sensor modality defines a projection operator $\Pi_i : \mathcal{X} \rightarrow \mathcal{Y}_i$; each observation y_i constrains the set of field configurations consistent with what the sensor reports. The consistency operator searches for the unique admissible configuration whose forward projections are compatible with all observations simultaneously.

The main result—the Yarncrawler Identifiability Theorem—establishes that \mathcal{C} has a unique fixed point if and only if the induced map $\tilde{\Pi} : \mathcal{X}/\sim \rightarrow \prod_i \mathcal{Y}_i$ is injective on the feasible set $\mathcal{F} = \Omega_{\text{obs}} \cap \Gamma$. The Dynamical Disambiguation Theorem extends this to realistic sensing regimes where static observations underdetermine the world-state: trajectory fragments constrain the tangent bundle of the admissible manifold, removing degrees of freedom that no static projection can touch.

The paper further establishes that the variational consistency framework and the sheaf-theoretic repair framework—in which sensor modalities define local charts, observations define local sections, and closure corresponds to the existence of a unique global section—are formally equivalent. There exists a functor $\Psi : \mathbf{Proj}(\mathcal{X}) \rightarrow \mathbf{Sh}(\mathcal{X}/\sim)$ under which the consistency operator is the variational realization of the sheaf repair operator, projection degeneracy corresponds to non-trivial H^1 , and entropy collapse corresponds to cohomology trivialization. The RSVP entropy field S is identified as the local density of cohomological obstruction; its minimization under \mathcal{C} is the physical expression of sheaf repair.

Deployment at planetary, regional, and site scales is unified by treating each scale as a different geometry of the feasible set under varying noise profiles. The framework is implemented as a sensor fusion architecture combining learned trajectory operators with physics-constrained forward models, applied to coupled ecosystem and civilization dynamics calibrated from before-and-after observational pairs across modalities including infrared, LiDAR, RGB, audio, and photonic strip sensors.

Contents

1	Introduction	7
1.1	The Problem	7
1.2	The Central Result	7
1.3	The Duality	7
1.4	The Consequence	8
1.5	Structure of the Paper	8
2	The Sheaf-Variational Equivalence	9
2.1	Two Descriptions of the Same Object	9
2.2	The Sheaf-Theoretic Account	9
2.3	The Variational Account	10
2.4	The Dictionary	10
2.5	The Sheaf-Variational Equivalence	10
2.6	Consequences of the Equivalence	11
2.7	The Entropy Identification	12
3	The World-State Manifold and Its Symmetry Group	12
3.1	Configuration Space	12
3.2	Symmetry Group	13
3.3	The Admissible Set	13
3.4	The RSVP Field Equations	13
3.5	The Regularizer	14
4	Sensor Modalities as Projection Operators	14
4.1	Purpose of This Section	14
4.2	The Configuration Space and Its Topology	14
4.3	Infrared Projection	15
4.4	LiDAR Projection	15
4.5	RGB/CCD Projection	15
4.6	Audio Projection	16
4.7	Photonic Strip Projection	16
4.8	The Completeness Proposition	16
4.9	Degrees of Freedom Table	17
4.10	Each Modality as a Chart	17
5	The Feasible Set and Its Geometry	17
5.1	Quotient Structure and Identifiability Domain	17
5.2	Geometry of the Observation-Consistent Set	17
5.3	Dimensionality, Degeneracy, and Cohomology	18
5.4	Information Geometry of the Feasible Set	18
5.5	Information Volume and Closure	18

5.6	Dynamics and Non-Convex Structure	19
5.7	Interpretation	19
5.8	Optimal Sensor Selection as Cohomological Descent	19
5.9	The Morse-Variational Structure of Yarncrawler Dynamics	20
5.10	Distributional Lift: Wasserstein Gradient Flow	21
5.11	RSVP Realization of the Variational Flow	21
6	The Consistency Operator and the Identifiability Theorem	22
6.1	Setup and Standing Assumptions	22
6.2	The Feasible Set	23
6.3	The Consistency Operator	23
6.4	The Identifiability Theorem	23
6.5	The Corollary on Diagnosable Failure Modes	24
6.6	Why the Quotient Structure Is Not Optional	24
7	Trajectory Learning and the Dynamical Admissibility Set	25
7.1	The Central Obstruction	25
7.2	The Path Space and the Stochastic Kernel	25
7.3	The Deterministic Limit and Tangent Bundle Constraint	25
7.4	Trajectory Learning from Before/After Pairs	26
7.5	The Dynamical Disambiguation Theorem	26
7.6	Joint Identifiability	26
8	The Inversion Layer	27
8.1	Role Specification	27
8.2	The Encoder as Approximate Posterior	27
8.3	The Refinement Loop	27
8.4	Categorical Structure of Multi-Modal Fusion	28
8.5	The Forward Model	28
8.6	Why Not End-to-End	28
9	Closure: Definition, Criteria, and Detection	28
9.1	The Unified Definition	28
9.2	The Three Conditions	29
9.3	Equivalence of the Three Conditions	29
9.4	Operational Detection	29
9.5	Entropy Collapse at Closure	29
9.6	Failure Modes Revisited	30
10	Deployment Scales as Feasible Set Geometry	30
10.1	The Unified Framework	30
10.2	Planetary Scale	30
10.3	Regional Scale	30
10.4	Site Scale	30

10.5	Multi-Scale Fusion and Scale Consistency	31
10.6	Sensor Selection Across Scales	31
11	Relation to Prior Work	31
11.1	Data Assimilation and Ensemble Methods	31
11.2	Inverse Rendering and Neural Radiance Fields	32
11.3	Digital Twin Systems	32
11.4	Categorical and Sheaf-Theoretic Semantics	32
12	Discussion: Closure as Epistemology	32
12.1	What Closure Means	32
12.2	Constraint Primacy	33
12.3	The RSVP Fields as Constraint Carriers	33
12.4	Stigmergic Accumulation and the Refinement Loop	33
12.5	Limits of the Framework	33
12.6	Forward and Inverse as Dual Inference	34
13	Conclusion	34
A	Fréchet Manifold Structure of \mathcal{X}	35
A.1	Fréchet Manifolds	35
A.2	The Quotient Topology	35
A.3	Tangent Bundle of \mathcal{X}/\sim	35
A.4	Geodesic Convexity on \mathcal{X}/\sim	35
B	Full Proof of the Identifiability Theorem	35
B.1	Functional-Analytic Setting	35
B.2	Existence of a Minimizer	36
B.3	Uniqueness of the Minimizer	36
B.4	Passage to the Quotient	36
B.5	Full Proof of Theorem 6.11	36
C	Full Proof of the Dynamical Disambiguation Theorem	37
C.1	Measure-Theoretic Framework	37
C.2	Convergence to the Deterministic Generator	37
C.3	Stability of Trajectory Separation Under Noise	37
C.4	Full Proof of Theorem 7.10	37
D	Fisher Information Geometry of the Feasible Set	38
D.1	Purpose and Position	38
D.2	Observation Likelihood on the Quotient Space	38
D.3	The Fisher Information Metric	38
D.4	Metric Monotonicity Under Sensor Addition	39
D.5	Restriction to the Feasible Set	39
D.6	Dynamical Lifting of Null Directions	39

D.7	Closure as Vanishing Information Volume	40
D.8	Reconstruction Efficiency	40
D.9	The Effective Metric Near Closure	40
D.10	Entropy Identification	41
E	Operator Class Library	41
E.1	Structure of the Library	41
E.2	Ecological Process Operators	41
E.3	Civilizational Process Operators	42
E.4	Learning Protocol	42
E.5	Library Extensibility	42
F	Implementation File Structure	42
F.1	Structural Overview	43
F.2	Core Infrastructure	43
F.3	State Layer	43
F.4	Log Layer	43
F.5	Morphism Layer	44
F.6	Controller Layer	44
F.7	Engine Layer	45
F.8	Correspondence to the Theoretical Framework	46
G	The Constellation Construction as Iterated Constraint Collapse	46
G.1	The Initial Object	46
G.2	Stage-wise Refinement as Endofunctors	46
G.3	Sheaf Enforcement (Stages 1–4)	47
G.4	Emergence of the Morphism Algebra (Stages 5–7)	47
G.5	Variational Selection and Trajectory Space (Stages 8–10)	47
G.6	Typed Category and Compositional Structure (Stage 11)	48
G.7	Functorial Transfer (Stage 12)	48
G.8	Controller Quotient (Stage 13)	49
G.9	The Fixed Point of the Total Operator	49
G.10	Correspondence to the Main Theorems	49
G.11	Stage 14: Homeostatic Constraint (Internal Sensory Monitoring)	50
G.12	Stage 15: Presentational Projection (Observer-Accessible Interface)	51
G.13	Stage 16: The Thought System and Terminal Closure	51
H	The Constellation Architecture as Operational Sheaf	52
H.1	From Deployment Tool to Sheaf Realization	52
H.2	The Eight Strata as Factorization of \mathcal{C}	53
H.3	The Repair Loop as Coboundary Construction	54

I	The Morphism Kernel as RSVP Operator Algebra	55
I.1	Operators as RSVP Field Transformations	55
I.2	Sensor-Coupled Operator Selection	55
I.3	Context-Sensitive Memory and Operator Discovery	56
J	The Controller Quotient and Meta-Level Closure	57
J.1	Controllers as Functors	57
J.2	Natural Transformations as Controller Evolution	58
J.3	The Quotient Structure over Controller Space	59
J.4	Closure at the Meta-Level	59

1 Introduction

1.1 The Problem

Given a family of projection operators $\{\Pi_i\}_{i \in \mathcal{I}}$ mapping an unknown world-state $[X] \in \mathcal{X}/\sim$ to observable quantities $y_i \in \mathcal{Y}_i$, and given a dynamical law governing how world-states evolve, the central question of this paper is:

Under what conditions does there exist a unique trajectory in the quotient configuration space \mathcal{X}/\sim consistent with all projections simultaneously and admissible under the dynamics?

This is not a modeling question. It is an identifiability question. The paper does not propose a method for approximating world-states from sensor data—it establishes the conditions under which a unique world-state is *demanded* by the data, and proves that when those conditions hold, a specific operator converges to it.

The setting is the following. The world-state is a coupled field triple $X = (\Phi, \mathbf{v}, S)$ —scalar density, vector flow, and entropy production—evolving over a compact spatial domain under the Relativistic Scalar-Vector Plenum (RSVP) field equations. The projections Π_i correspond to sensor modalities: infrared, LiDAR, RGB, audio, photonic strips. Each observation y_i constrains which field configurations are consistent with what the sensor reports. The consistency operator \mathcal{C} searches for the configuration that minimizes total projection residual subject to physical admissibility. The dynamical admissibility set Γ restricts attention to configurations that lie on trajectories the world could have taken. The feasible set $\mathcal{F} = \Omega_{\text{obs}} \cap \Gamma$ is the intersection of all constraints. Closure—the condition that $|\mathcal{F}| = 1$ —is the condition that the world-state is uniquely determined.

1.2 The Central Result

The paper proves the following.

The problem of multi-modal consistency—finding a world-state whose forward projections are compatible with all available observations—is equivalent to the existence of a unique global section of a sheaf \mathcal{S}_Π induced by the projection family over a cover of \mathcal{X}/\sim .

More precisely: the consistency operator \mathcal{C} is the variational realization of the sheaf repair operator. Its fixed points are global sections. Its failure modes—projection degeneracy and regularizer flatness—are cohomological obstructions $H^1 \neq 0$ and degenerate H^0 . Adding sensors refines the cover and kills H^1 classes. Strengthening the dynamical constraint extends the cover into the temporal direction and kills residual ambiguity that spatial sensing cannot reach. Closure is the condition $H^0 = 1, H^1 = 0$: the sheaf has exactly one global section.

This is the Yarncrawler Identifiability Theorem and its dynamical extension, proved in Sections 6 and 7 and their appendices.

1.3 The Duality

The framework has two faces, both present in this paper.

The *forward* direction takes a world-state as given and asks how a system maintains coherence under perturbation: how it detects tears in its semantic fabric, applies repair morphisms that trivialize Čech cocycles, exports entropy through stigmergic reinforcement, and sustains homeorhetic flows rather than static equilibria.

The *inverse* direction takes sensor observations as given and asks what world-state they structurally demand. Given what the sensors report, what field configuration is forced?

These are not complementary perspectives on the same phenomenon. They are dual realizations of the same mathematical object. Both directions study the sheaf \mathcal{S}_Π over the quotient configuration space \mathcal{X}/\sim . The forward direction studies the internal dynamics of sections within \mathcal{S}_Π —how they evolve, heal, and accumulate into global coherence. The inverse direction studies the conditions under which a unique global section exists and the operator that converges to it. The RSVP field triple (Φ, \mathbf{v}, S) parametrizes the stalks in both directions. The Yarncrawler framework is the unified theory of that sheaf and its dynamics.

The object common to both is neither the operator nor the sheaf taken separately. It is the equivalence between them—the functor $\Psi : \mathbf{Proj}(\mathcal{X}) \rightarrow \mathbf{Sh}(\mathcal{X}/\sim)$ proved in Section 2, which shows that a projection family and a sheaf over a configuration space are the same structure viewed from outside and inside respectively.

1.4 The Consequence

When closure is achieved—when $|\mathcal{F}| = 1$ —three things happen simultaneously that were previously three separate conditions.

The projection residuals collapse to within noise tolerance: $\|\Pi_i(X^*) - y_i\| \leq \epsilon_i$ for all i . The cohomology of the induced sheaf trivializes: $H^1(\mathcal{U}, \mathcal{S}_\Pi) = 0$ and $\dim H^0 = 1$. The RSVP entropy field S is minimized subject to admissibility: no further constraint can be satisfied that is not already satisfied.

These are not three criteria that must be checked independently. Under the Sheaf-Variational Equivalence of Section 2, they are one criterion expressed in three languages—variational, cohomological, and field-theoretic. The entropy of the RSVP field at the fixed point is the local density of cohomological obstruction, and its minimization is the physical expression of sheaf repair.

The closed state $[X^*]$ is therefore not a best estimate. It is the unique trajectory consistent with all projections and admissible dynamics—the world-state that no alternative history can explain away.

1.5 Structure of the Paper

Section 2 establishes the Sheaf-Variational Equivalence. Sections 3 and 4 develop the prerequisite machinery. Section 5 defines the feasible set geometrically. Sections 6 and 7 contain the core theorems. Section 8 develops the inversion layer architecture. Section 9 defines closure operationally and proves equivalence of the three closure criteria. Section 10 unifies planetary, regional, and site-scale deployment. Section 11 positions the framework

against prior work. Section 12 draws the epistemological consequence. Appendices A–E provide the functional-analytic foundations.

2 The Sheaf-Variational Equivalence

2.1 Two Descriptions of the Same Object

The Yarncrawler framework has two faces. The first describes a self-refactoring system that maintains coherence by repairing tears in its semantic fabric: local sections are parsed, inconsistencies are detected as Čech cocycles, and repair morphisms trivialize them until a global section exists. The second describes a consistency operator that searches for a field configuration whose forward projections are compatible with all available sensor observations simultaneously. These descriptions appear to operate at different levels of abstraction—one categorical and topological, the other variational and analytic. This section proves they are the same description.

The identification is not a metaphor. There is a precise functor between the category of sheaves of semantic sections over a cover and the category of projection families on the quotient configuration space \mathcal{X}/\sim . Under this functor, every concept in the sheaf-theoretic account has an exact counterpart in the variational account, and vice versa.

2.2 The Sheaf-Theoretic Account

Let \mathcal{M} be a topological space equipped with an open cover $\mathcal{U} = \{U_i\}_{i \in \mathcal{I}}$. A presheaf \mathcal{S} over \mathcal{U} assigns to each U_i a category $\mathcal{S}(U_i)$ of local sections, and to each inclusion $U_j \subseteq U_i$ a restriction functor $\rho_{ij} : \mathcal{S}(U_i) \rightarrow \mathcal{S}(U_j)$.

Definition 2.1 (Gluing Condition). *A family of local sections $\{s_i \in \mathcal{S}(U_i)\}_{i \in \mathcal{I}}$ satisfies the gluing condition if for all $i, j \in \mathcal{I}$:*

$$\rho_{U_i, U_i \cap U_j}(s_i) = \rho_{U_j, U_i \cap U_j}(s_j) \quad \text{on } U_i \cap U_j.$$

A presheaf in which every compatible family of local sections admits a unique global section $s \in \mathcal{S}(\mathcal{M})$ with $s|_{U_i} = s_i$ is called a sheaf.

Definition 2.2 (Semantic Tear and Repair). *A semantic tear at overlap $U_i \cap U_j$ is a failure of the gluing condition. The discrepancy defines a Čech 1-cochain $c_{ij} = \rho_{ij}(s_j) - \rho_{ji}(s_i)$. A repair is a morphism in $\mathcal{S}(U_i \cap U_j)$ that trivializes c_{ij} . The space of unresolved tears is measured by $H^1(\mathcal{U}, \mathcal{S})$.*

Definition 2.3 (Sheaf Repair Operator). *The sheaf repair operator $\mathcal{R}_{\text{sheaf}}$ maps a family of inconsistent local sections to a minimally modified family satisfying the gluing condition, where minimality is measured by a seam cost functional $\mathcal{L}_{\text{seam}}$:*

$$\mathcal{R}_{\text{sheaf}}(\{s_i\}) = \arg \min_{\{s'_i\}} \sum_{i < j} \mathcal{L}_{\text{seam}}(s'_i|_{U_{ij}}, s'_j|_{U_{ij}}) + \sum_i \mathcal{L}_{\text{local}}(s'_i, s_i).$$

Global section existence corresponds to $\mathcal{L}_{\text{seam}} = 0$ at the optimum.

2.3 The Variational Account

The variational account centers on the consistency operator:

$$\mathcal{C}([X]) = \arg \min_{[X'] \in \mathcal{A}} \left[\sum_{i \in \mathcal{I}} \|\Pi_i(X') - y_i\|_{\mathcal{Y}_i}^2 + \mathcal{R}(X') \right]$$

where $\Pi_i : \mathcal{X} \rightarrow \mathcal{Y}_i$ are projection operators, y_i are observations, and \mathcal{R} is the RSVP-consistent regularizer encoding admissibility. A fixed point $[X^*]$ satisfies $\Pi_i(X^*) \approx y_i$ for all i within tolerance ϵ_i .

2.4 The Dictionary

The following correspondence identifies every object in the sheaf-theoretic account with its exact counterpart in the variational account.

Sheaf-theoretic concept	Variational counterpart
Open set U_i in cover \mathcal{U}	Sensor modality i with space \mathcal{Y}_i
Local section $s_i \in \mathcal{S}(U_i)$	Observation $y_i \in \mathcal{Y}_i$
Restriction map ρ_{ij}	Projection operator Π_i
Gluing condition	Projection consistency $\Pi_i(X) = y_i$
Čech 1-cochain c_{ij}	Projection residual $\ \Pi_i(X) - y_i\ $
Coboundary (trivialized cocycle)	Residual driven to zero by \mathcal{C}
Global section $s \in \mathcal{S}(\mathcal{M})$	Fixed point $[X^*] \in \mathcal{X}/\sim$
Cohomological obstruction $H^1 \neq 0$	Projection degeneracy
Flat direction in $\mathcal{L}_{\text{local}}$	Regularizer flatness
Repair morphism trivializing c_{ij}	One step of the refinement loop
Sheaf repair operator $\mathcal{R}_{\text{sheaf}}$	Consistency operator \mathcal{C}
Stigmergic threshold	Contractivity condition on \mathcal{C}
Self-maintaining growth	Convergence of refinement loop to $[X^*]$

2.5 The Sheaf-Variational Equivalence

Definition 2.4 (Category of Projection Families). *Let $\mathbf{Proj}(\mathcal{X})$ be the category whose objects are pairs $(\mathcal{I}, \{\Pi_i\}_{i \in \mathcal{I}})$ and whose morphisms are inclusions of index sets compatible with the projection structure.*

Definition 2.5 (Category of Covering Sheaves). *Let $\mathbf{Sh}(\mathcal{M})$ be the category whose objects are sheaves \mathcal{S} over covers \mathcal{U} of \mathcal{M} , and whose morphisms are morphisms of sheaves over refinements of covers.*

Theorem 2.6 (Sheaf-Variational Equivalence). *There exists a functor*

$$\Psi : \mathbf{Proj}(\mathcal{X}) \longrightarrow \mathbf{Sh}(\mathcal{X}/\sim)$$

such that:

- (i) Ψ maps each projection family $\{\Pi_i\}_{i \in \mathcal{I}}$ to a sheaf \mathcal{S}_Π over a cover of \mathcal{X}/\sim whose local sections are observation-consistent field configurations.
- (ii) The gluing condition of \mathcal{S}_Π is equivalent to the projection consistency condition $\Pi_i(X) = y_i$ for all i .
- (iii) The consistency operator \mathcal{C} is the variational realization of the sheaf repair operator $\mathcal{R}_{\text{sheaf}}$ under Ψ .
- (iv) The fixed point $[X^*]$ of \mathcal{C} corresponds under Ψ to the unique global section of \mathcal{S}_Π .
- (v) The cohomological obstruction $H^1(\mathcal{U}, \mathcal{S}_\Pi) = 0$ if and only if the projection family is identifying on \mathcal{F} .

Proof. Construction of Ψ . Given a projection family $(\mathcal{I}, \{\Pi_i\})$, construct the cover $\mathcal{U} = \{U_i\}_{i \in \mathcal{I}}$ of \mathcal{X}/\sim where

$$U_i = \{[X] \in \mathcal{A} \mid \|\Pi_i(X) - y_i\|_{\mathcal{Y}_i} \leq \epsilon_i\}.$$

Define the sheaf \mathcal{S}_Π by assigning to each U_i :

$$\mathcal{S}_\Pi(U_i) = U_i \cap \mathcal{A}.$$

For inclusions $U_j \subseteq U_i$, the restriction map ρ_{ij} is the inclusion $U_j \cap \mathcal{A} \hookrightarrow U_i \cap \mathcal{A}$.

Part (i). By construction, $\mathcal{S}_\Pi(U_i)$ consists of admissible configurations whose i -th projection is within tolerance of y_i . ✓

Part (ii). A family $\{[X_i] \in \mathcal{S}_\Pi(U_i)\}$ satisfies the gluing condition iff $[X_i] = [X_j]$ on $U_i \cap U_j$ for all i, j , which is precisely the condition that a single $[X] \in \mathcal{A}$ satisfies $\Pi_i(X) = y_i$ for all i simultaneously. ✓

Part (iii). The objective of $\mathcal{R}_{\text{sheaf}}$ translates as:

$$\sum_{i < j} \mathcal{L}_{\text{seam}} + \mathcal{L}_{\text{local}} \longleftrightarrow \sum_i \|\Pi_i(X) - y_i\|^2 + \mathcal{R}(X),$$

which is exactly the objective of \mathcal{C} . ✓

Part (iv). A fixed point $[X^*]$ of \mathcal{C} satisfies $[X^*] \in U_i$ for all i , giving a local section $s_i = [X^*]|_{U_i}$ that trivially satisfies the gluing condition. Uniqueness follows from Theorem 6.11. ✓

Part (v). $H^1(\mathcal{U}, \mathcal{S}_\Pi)$ measures failure of compatible local sections to glue into a global section, which is exactly projection degeneracy by part (ii). ✓

Functoriality. Adding sensors corresponds to refining the cover; the refinement map defines a natural transformation establishing functoriality. □

2.6 Consequences of the Equivalence

Consequence 1: The two failure modes are cohomological. Projection degeneracy is $H^1 \neq 0$; regularizer flatness is degenerate H^0 . Adding sensors can kill H^1 classes; strengthening the regularizer thins the stalks and kills degenerate H^0 .

Consequence 2: Dynamical disambiguation is cohomology in time. The dynamical admissibility set Γ extends the cover into the temporal direction. Trajectory separation kills residual H^1 classes that spatial sensors cannot. The full feasible set $\mathcal{F} = \Omega_{\text{obs}} \cap \Gamma$ is the global section of a sheaf over a spatio-temporal cover.

Consequence 3: Forward and inverse are dual. The forward direction studies the repair dynamics of \mathcal{S}_{Π} from inside—how sections evolve under stigmergic reinforcement and homeorhetic flow. The inverse direction studies the existence and uniqueness of the global section from outside. Both study the same sheaf.

2.7 The Entropy Identification

The RSVP entropy field S appears in two roles throughout this framework: as the entropy component of the field triple (Φ, \mathbf{v}, S) , and as a measure of unresolved obstruction. Under the Sheaf-Variational Equivalence these become one.

The observation-consistent set Ω_{obs} carries a natural measure induced by the Fisher information metric of the noise models $\{\mathcal{N}_i\}$ (developed in Appendix D). The volume of Ω_{obs} under this metric measures residual uncertainty. Under the equivalence:

$$\log \text{Vol}_{\text{Fisher}}(\Omega_{\text{obs}}) \longleftrightarrow \log \dim H^0(\mathcal{U}, \mathcal{S}_{\Pi}) \longleftrightarrow S_{\text{RSVP}}(X^*).$$

All three measure residual degrees of freedom after all constraints have been applied. Closure corresponds to each collapsing simultaneously. The RSVP entropy field is not merely a bookkeeping device—it is the local density of cohomological obstruction, and its minimization under \mathcal{C} is the physical expression of sheaf repair.

3 The World-State Manifold and Its Symmetry Group

3.1 Configuration Space

The configuration space is

$$\mathcal{X} = H^1(\Omega) \times H^1(\Omega; \mathbb{R}^3) \times L^2(\Omega)$$

over a compact domain $\Omega \subset \mathbb{R}^3$ with smooth boundary, equipped with the product Sobolev norm. The three components have physical interpretations that ground the abstract framework. The scalar field $\Phi \in H^1(\Omega)$ encodes ecosystem density and civilizational legitimacy potential: high Φ at a location corresponds to dense biomass, concentrated infrastructure, or high resource availability. The vector field $\mathbf{v} \in H^1(\Omega; \mathbb{R}^3)$ encodes resource flows and infrastructure transport: its divergence marks sources and sinks, its curl marks circulation patterns. The entropy field $S \in L^2(\Omega)$ encodes complexity budget and irreversibility: high S corresponds to high local disorder, biological activity, or unresolved ambiguity in the system’s reconstruction of the world-state.

As a product of Hilbert spaces, \mathcal{X} is itself a separable Hilbert space and hence a Fréchet manifold. The full Fréchet manifold structure, including the treatment of extensions to

more general field configurations, is developed in Appendix A.

3.2 Symmetry Group

A symmetry group G acts on \mathcal{X} by three classes of transformations. First, *gauge redundancies* of the field representation: rephasing and rescaling transformations that leave all physical observables invariant but change the coordinate representation of (Φ, \mathbf{v}, S) . Second, *coordinate reparametrizations* of Ω : smooth diffeomorphisms that relabel spatial positions without altering field content. Third, *observationally equivalent reconfigurations* of the civilizational layer: distinct infrastructure configurations or demographic distributions that produce identical outputs under all available sensor projections.

The action of G is smooth, proper, and isometric with respect to the product Sobolev metric, so the quotient \mathcal{X}/G is Hausdorff and inherits a natural metric.

Two distinct equivalence relations are in play. G -equivalence ($X \sim_G X'$ when $X' = g \cdot X$ for some $g \in G$) is intrinsic to the field representation. Observational equivalence ($X \sim X'$ when $\Pi_i(X) = \Pi_i(X')$ for all i) is extrinsic and depends on the projection family. In general \sim is coarser than \sim_G : two configurations may agree on all projections without being related by any element of G . When the projection family is complete in the sense of Proposition 4.10, the two equivalences coincide.

3.3 The Admissible Set

Define the admissible set

$$\mathcal{A} = \{[X] \in \mathcal{X}/\sim \mid \mathcal{R}(X) < \infty\},$$

where \mathcal{R} is the regularizer defined below. The admissible set consists of equivalence classes satisfying the RSVP field equations, ecological plausibility, and civilizational feasibility simultaneously. It is a closed convex subset of the Banach realization $\mathcal{B} = \mathcal{X}$ (Assumption 6.5).

Geodesic convexity of \mathcal{A} ensures that admissible interpolations between two configurations remain admissible. Lower semicontinuity of \mathcal{R} ensures admissibility is stable under limits. Coercivity ensures that physically implausible configurations — those with unbounded field energy or violating conservation laws — are excluded rather than merely penalized.

3.4 The RSVP Field Equations

The time evolution of (Φ, \mathbf{v}, S) is governed by the RSVP plenum equations:

$$\begin{aligned} \partial_t \Phi + \nabla \cdot (\Phi \mathbf{v}) &= 0 && \text{(scalar continuity),} \\ \partial_t \mathbf{v} + (\mathbf{v} \cdot \nabla) \mathbf{v} + \nabla \Phi &= \mathbf{f} && \text{(vector momentum),} \\ \partial_t S + \nabla \cdot (S \mathbf{v}) &= \sigma(X) && \text{(entropy production),} \end{aligned}$$

where \mathbf{f} encodes external forcing (ecological or civilizational) and $\sigma(X) \geq 0$ is the entropy production rate satisfying the second law. Solutions to these equations form the dynamical admissibility set Γ in the zero-noise limit (Section 7).

3.5 The Regularizer

The regularizer encodes coupled closure across all three constraint classes:

$$\mathcal{R}(X) = \mathcal{R}_{\text{RSVP}}(X) + \lambda_{\text{eco}} \mathcal{R}_{\text{eco}}(X) + \lambda_{\text{civ}} \mathcal{R}_{\text{civ}}(X).$$

$\mathcal{R}_{\text{RSVP}}$ is the squared L^2 norm of residuals of the three field equations above. \mathcal{R}_{eco} is zero on ecologically admissible configurations (biomass conservation, succession ordering, watershed feasibility) and strictly positive off them. \mathcal{R}_{civ} is zero on civilizationally admissible configurations (demographic feasibility, infrastructure stress bounds, resource flow conservation) and strictly positive off them.

The decomposition is notational: \mathcal{R} is a single functional encoding coupled closure. A configuration satisfying the RSVP equations but violating ecological admissibility is not merely penalized; it lies outside \mathcal{A} entirely. The weights $\lambda_{\text{eco}}, \lambda_{\text{civ}} > 0$ are fixed hyperparameters determining the relative stringency of the three constraint classes.

4 Sensor Modalities as Projection Operators

4.1 Purpose of This Section

Section 2 established that a projection family $\{\Pi_i\}_{i \in \mathcal{I}}$ and a sheaf \mathcal{S}_Π over \mathcal{X}/\sim are the same structure viewed from outside and inside respectively. What remains is to specify what the projection operators actually are—not as abstract bounded linear maps, but as concrete operators whose domains, codomains, and kernels are determined by the physics of each sensor modality.

This section defines each projection operator precisely enough that Assumption A1 can be verified directly. It establishes the RSVP correspondence: the identification of each modality with a specific component or functional of the field triple (Φ, \mathbf{v}, S) . And it proves the completeness proposition: the condition under which the projection family collectively spans the full RSVP triple.

Nothing in this section is motivated by sensor engineering. The operators are defined as they are because those definitions are what the theorems in Sections 6 and 7 require.

4.2 The Configuration Space and Its Topology

Let $\Omega \subset \mathbb{R}^3$ be a compact domain with smooth boundary. The configuration space is

$$\mathcal{X} = H^1(\Omega) \times H^1(\Omega; \mathbb{R}^3) \times L^2(\Omega)$$

with product Sobolev norm $\|X\|_{\mathcal{X}}^2 = \|\Phi\|_{H^1}^2 + \|\mathbf{v}\|_{H^1}^2 + \|S\|_{L^2}^2$. All projection operators defined below are required to descend to the quotient \mathcal{X}/\sim : they must be constant on

G -orbits.

4.3 Infrared Projection

Definition 4.1 (Infrared Projection). *The infrared projection $\Pi_{\text{IR}} : \mathcal{X} \rightarrow \mathcal{Y}_{\text{IR}}$ is defined by*

$$\Pi_{\text{IR}}(X) = (\Phi|_{\partial\Omega}, \nabla\Phi \cdot \hat{n}|_{\partial\Omega})$$

where \hat{n} is the outward unit normal to $\partial\Omega$ and $\mathcal{Y}_{\text{IR}} = L^2(\partial\Omega) \times L^2(\partial\Omega)$.

Lemma 4.2. Π_{IR} is a bounded linear operator from \mathcal{X} to \mathcal{Y}_{IR} .

Proof. The trace operator $\Phi \mapsto \Phi|_{\partial\Omega}$ is bounded from $H^1(\Omega)$ to $L^2(\partial\Omega)$ by the Sobolev trace theorem. The normal derivative is bounded from $H^1(\Omega)$ to $H^{-1/2}(\partial\Omega) \hookrightarrow L^2(\partial\Omega)$ for smooth $\partial\Omega$. \square

RSVP correspondence. Π_{IR} constrains boundary values and outward flux of Φ : energy density at the surface and rate of energy exchange. It does not directly constrain \mathbf{v} or S in the interior.

4.4 LiDAR Projection

Definition 4.3 (LiDAR Projection). *Fix a threshold $\phi_0 > 0$. The LiDAR projection $\Pi_{\text{LiDAR}} : \mathcal{X} \rightarrow \mathcal{Y}_{\text{LiDAR}}$ is defined by*

$$\Pi_{\text{LiDAR}}(X) = \mathbf{1}_{\{\Phi \geq \phi_0\}} \quad \text{together with} \quad \Pi_{\text{LiDAR}}^\perp(X) = \frac{\nabla\Phi}{|\nabla\Phi|} \Big|_{\{\Phi=\phi_0\}}.$$

The observation space is $\mathcal{Y}_{\text{LiDAR}} = L^2(\Omega; \{0, 1\}) \times L^2(\{\Phi = \phi_0\}; S^2)$.

Lemma 4.4. Π_{LiDAR} is bounded for ϕ_0 a regular value of Φ .

Proof. For ϕ_0 a regular value, $\{\Phi = \phi_0\}$ is smooth by the implicit function theorem in H^1 . The indicator is bounded in $L^2(\Omega)$ by domain volume. The normal field is bounded by the Sobolev embedding $H^1 \hookrightarrow C^0$ in three dimensions. \square

RSVP correspondence. Π_{LiDAR} constrains the geometry of level sets of Φ and normal directions to those surfaces, which simultaneously constrain boundary conditions on \mathbf{v} : flows cannot cross material boundaries.

4.5 RGB/CCD Projection

Definition 4.5 (RGB Projection). *Let $\mathbf{c} \in S^2$ be the camera viewing direction and $f : \mathbb{R} \rightarrow \mathbb{R}^3$ a material reflectance function. The RGB projection $\Pi_{\text{RGB}} : \mathcal{X} \rightarrow \mathcal{Y}_{\text{RGB}}$ is defined by*

$$\Pi_{\text{RGB}}(X) = \int_0^\infty f(\Phi(\mathbf{x}(t))) e^{-\int_0^t \Phi(\mathbf{x}(s)) ds} dt$$

where $\mathbf{x}(t) = \mathbf{x}_0 + t\mathbf{c}$ and $\mathcal{Y}_{\text{RGB}} = L^2(\mathcal{V}; \mathbb{R}^3)$ over camera pixels \mathcal{V} .

Lemma 4.6. Π_{RGB} is bounded when f is Lipschitz and $\Phi \geq 0$.

Proof. The exponential attenuation is bounded by 1. The integrand is bounded by $\|f\|_\infty$, which is finite by Lipschitz continuity and $H^1 \hookrightarrow L^\infty$ in three dimensions. \square \square

RSVP correspondence. Π_{RGB} constrains Φ along rays—a projective complement to Π_{LiDAR} 's level-set constraint. Their kernels overlap partially but not completely.

4.6 Audio Projection

Definition 4.7 (Audio Projection). Let $\mathbf{x}_m \in \Omega$ for $m = 1, \dots, M$ be microphone positions. The audio projection $\Pi_{\text{audio}} : \mathcal{X} \times [0, T] \rightarrow \mathcal{Y}_{\text{audio}}$ is

$$\Pi_{\text{audio}}(X)(t) = (\partial_t S(\mathbf{x}_m, t) + \nabla \cdot \mathbf{v}(\mathbf{x}_m, t))_{m=1}^M$$

with $\mathcal{Y}_{\text{audio}} = L^2([0, T]; \mathbb{R}^M)$.

RSVP correspondence. Π_{audio} is the only modality that directly constrains $\partial_t S$ —the rate of entropy production—and $\nabla \cdot \mathbf{v}$, encoding sources and sinks of flow. Hidden processes invisible to all spatial sensors are audible as entropy production events.

4.7 Photonic Strip Projection

Definition 4.8 (Photonic Strip Projection). Let $\gamma : [0, L] \rightarrow \Omega$ parametrize the strip by arc length. The photonic strip projection $\Pi_{\text{strip}} : \mathcal{X} \times [0, T] \rightarrow \mathcal{Y}_{\text{strip}}$ is

$$\Pi_{\text{strip}}(X)(s, t) = (\Phi(\gamma(s), t), \mathbf{v}(\gamma(s), t) \cdot \gamma'(s), \partial_t \Phi(\gamma(s), t))$$

with $\mathcal{Y}_{\text{strip}} = L^2([0, L] \times [0, T]; \mathbb{R}^3)$.

RSVP correspondence. The photonic strip provides continuous temporal records of Φ , tangential \mathbf{v} , and $\partial_t \Phi$ along a fixed transect: the only modality directly constraining temporal evolution of Φ at interior points.

4.8 The Completeness Proposition

Definition 4.9 (Spanning Condition). A projection family $\{\Pi_i\}_{i \in \mathcal{I}}$ is field-spanning if

$$\ker \tilde{\Pi} = \{[X] \in \mathcal{X}/\sim \mid \Pi_i(X) = 0 \forall i \in \mathcal{I}\} = \{[0]\}.$$

Proposition 4.10 (Completeness of the Five-Modality Family). The projection family $\{\Pi_{\text{IR}}, \Pi_{\text{LiDAR}}, \Pi_{\text{RGB}}, \Pi_{\text{audio}}, \Pi_{\text{strip}}\}$ is field-spanning on \mathcal{X}/\sim .

Proof. Suppose $[X] \in \ker \tilde{\Pi}$. $\Pi_{\text{IR}}(X) = 0$ gives $\Phi|_{\partial\Omega} = 0$. $\Pi_{\text{LiDAR}}(X) = 0$ gives $\Phi < \phi_0$ a.e. $\Pi_{\text{strip}}(X) = 0$ along strips dense in Ω gives $\Phi = 0$ on a dense subset; by H^1 continuity, $\Phi = 0$ on Ω . With $\Phi = 0$, strip vanishing also gives $\mathbf{v} \cdot \gamma' = 0$ along all strips; with three non-parallel strip directions, $\mathbf{v} = 0$. Then $\Pi_{\text{audio}}(X) = 0$ with $\mathbf{v} = 0$ gives $\partial_t S = 0$ at dense microphone positions, so S is time-constant; RSVP entropy constraints force $S = 0$. Hence $[X] = [0]$. \square \square

4.9 Degrees of Freedom Table

Modality	Field constrained	Where	New DOF closed
Infrared	$\Phi, \nabla\Phi \cdot \hat{n}$	$\partial\Omega$	Boundary energy, flux
LiDAR	Level sets of Φ , normals	Interior surfaces	Geometry; \mathbf{v} BCs
RGB	Φ along rays	Lines of sight	Interior density
Audio	$\partial_t S, \nabla \cdot \mathbf{v}$	Point measurements	Temporal entropy, flow sources
Photonic strip	$\Phi, \mathbf{v} \cdot \boldsymbol{\gamma}', \partial_t \Phi$	Transects over time	Temporal Φ ; tangential flow

4.10 Each Modality as a Chart

Under the Sheaf-Variational Equivalence, each Π_i defines a chart $U_i = \{[X] \in \mathcal{A} \mid \|\Pi_i(X) - y_i\| \leq \epsilon_i\}$ on \mathcal{X}/\sim . The five modalities define five charts. Completeness ensures the cover is fine enough to support a non-degenerate sheaf. Sensor design is sheaf design: choosing modalities is choosing the cover of \mathcal{X}/\sim .

5 The Feasible Set and Its Geometry

This section identifies the feasible set \mathcal{F} as the primary geometric and cohomological object of the Yarncrawler framework. Given the projection family $\{\Pi_i\}_{i \in \mathcal{I}}$ and observations $\{y_i\}$, define observation-consistent sets

$$U_i = \{[X] \in \mathcal{A} \mid \|\Pi_i(X) - y_i\|_{\mathcal{Y}_i} \leq \epsilon_i\}, \quad \Omega_{\text{obs}} = \bigcap_{i \in \mathcal{I}} U_i.$$

The dynamically admissible set Γ , whose full construction appears in Section 7, consists of configurations lying on admissible trajectories of the world dynamics. The feasible set is $\mathcal{F} = \Omega_{\text{obs}} \cap \Gamma$. Under the Sheaf-Variational Equivalence (Section 2), each U_i is the support of a local section of \mathcal{S}_{Π} , and \mathcal{F} is the space of admissible global sections. The identifiability and disambiguation theorems are statements about when this fiber collapses to a singleton.

5.1 Quotient Structure and Identifiability Domain

The relevant ambient space is the quotient $\mathcal{A} = \mathcal{X}/\sim$, where $X \sim X'$ if $\Pi_i(X) = \Pi_i(X')$ for all i . Multiple representatives in \mathcal{X} agreeing on all projections correspond to the same equivalence class in \mathcal{A} and hence to the same global section of \mathcal{S}_{Π} . All identifiability statements in Section 6 concern the restriction of the induced map $\tilde{\Pi} : \mathcal{A} \rightarrow \prod_i \mathcal{Y}_i$ to \mathcal{F} .

5.2 Geometry of the Observation-Consistent Set

Each U_i is closed and convex under Assumption 6.1. However, their intersection Ω_{obs} is generically anisotropic and nontrivial. Three structural features determine its geometry. *Anisotropy*: each projection constrains different directions in \mathcal{A} . *Non-aligned nullspaces*: the kernels of $\{\Pi_i\}$ need not coincide, leaving residual directions that survive all individual

constraints. *Curvature under pullback*: the metric induced by the projection family is not Euclidean in \mathcal{X} , even when each Π_i is linear. For these reasons Ω_{obs} is best understood as a stratified admissible region whose local structure reflects the overlap pattern of the sensor family; smooth manifold structure holds on regular strata under generic regularity assumptions.

5.3 Dimensionality, Degeneracy, and Cohomology

Define the effective dimension $\dim_{\text{eff}}(\mathcal{F}) = \dim(\Omega_{\text{obs}} \cap \Gamma)$. When $\dim_{\text{eff}}(\mathcal{F}) > 0$, multiple admissible configurations remain; under the Sheaf-Variational Equivalence this corresponds to multiple global sections, equivalently to non-trivial $H^1(\mathcal{U}, \mathcal{S}_{\Pi})$. Closure is the condition $\dim_{\text{eff}}(\mathcal{F}) = 0$, which under the equivalence is exactly

$$\dim H^0(\mathcal{U}, \mathcal{S}_{\Pi}) = 1, \quad H^1(\mathcal{U}, \mathcal{S}_{\Pi}) = 0.$$

Ambiguity, degeneracy, and non-trivial cohomology are the same phenomenon expressed in geometric, variational, and categorical language.

Proposition 5.1 (Feasible Set Contraction). *Under monotone addition of sensors with independent noise models, $\dim_{\text{eff}}(\mathcal{F})$ is non-increasing. Strict decrease holds when the new sensor Π_j satisfies $\ker \Pi_j \not\supseteq \bigcap_{i < j} \ker \Pi_i$.*

Proof. Adding U_j to the intersection defining Ω_{obs} cannot increase dimension. Strict decrease holds iff Π_j constrains a direction not previously closed. \square \square

5.4 Information Geometry of the Feasible Set

The projection family and noise models induce a Fisher information metric on \mathcal{A} :

$$g_F(X) = \sum_{i \in \mathcal{I}} D\Pi_i(X)^* \Sigma_i^{-1} D\Pi_i(X).$$

Directions in $\ker g_F(X)$ are observationally invisible. Two structural properties follow. *Metric monotonicity*: adding sensor j gives $g_F^{\text{new}} \succeq g_F$. *Dynamical lifting*: a null direction v at $t = 0$ may satisfy $g_F^{(\tau)}(X)(v, v) > 0$ under the time-extended metric, because the flow Φ_F^τ transports null directions into visible ones. The formal development of both properties is in Appendix D.

5.5 Information Volume and Closure

The Fisher metric induces a volume form on the regular strata of \mathcal{F} . Define the information volume $\text{Vol}_F(\mathcal{F}) = \int_{\mathcal{F}_{\text{reg}}} d\text{Vol}_{g_F}$. At closure the following collapse simultaneously:

$$\text{Vol}_F(\mathcal{F}) \rightarrow 0, \quad \dim H^0 \rightarrow 1, \quad H^1 \rightarrow 0, \quad S_{\text{RSVP}}(X^*) \text{ minimized.}$$

These are not independent criteria but one condition in four languages. Full equivalence is established in Appendix D.

5.6 Dynamics and Non-Convex Structure

While Ω_{obs} is often convex, Γ is generically non-convex. Their intersection \mathcal{F} may exhibit disconnected components, curved boundaries, and trajectory-induced foliations, corresponding under the equivalence to distinct trajectory classes and non-trivial H^1 classes. Section 7 shows that trajectory constraints eliminate these components by enforcing temporal separation.

5.7 Interpretation

The feasible set is the space of all worlds compatible with the evidence and the laws of motion. The task of Yarncrawler is not statistical prediction but geometric contraction: reduce $\dim_{\text{eff}}(\mathcal{F})$ until $|\mathcal{F}| = 1$. Under the Sheaf-Variational Equivalence, this is exactly the process of repairing \mathcal{S}_Π until a unique global section exists.

5.8 Optimal Sensor Selection as Cohomological Descent

The preceding subsections establish that residual ambiguity in \mathcal{F} corresponds, under the Sheaf-Variational Equivalence, to nontrivial cohomology in $H^1(\mathcal{U}, \mathcal{S}_\Pi)$. Each unresolved degree of freedom in \mathcal{F} is a generator of this cohomology group. Closure is therefore equivalent to the annihilation of all such generators.

This allows the sensing problem to be reformulated as a descent problem on cohomological rank. Let $r(\mathcal{U}) := \dim H^1(\mathcal{U}, \mathcal{S}_\Pi)$ and let ϕ_j denote the map on cohomology induced by adding sensor j : $\phi_j : H^1(\mathcal{U}, \mathcal{S}_\Pi) \rightarrow H^1(\mathcal{U} \cup \{U_j\}, \mathcal{S}_\Pi)$. The new sensor sends some cocycles to coboundaries, reducing the rank to $r_{n+1} = \dim \text{Im}(\phi_j)$.

For a candidate sensor j , define its marginal information gain along the residual nullspace $N_X = \ker \tilde{\Pi}|_{T_X \mathcal{F}}$ by

$$\Delta_j(X) = \inf_{\substack{v \in N_X \\ \|v\|=1}} \left\langle D\Pi_j(X)v, \Sigma_j^{-1} D\Pi_j(X)v \right\rangle.$$

Theorem 5.2 (Greedy Cohomology Descent). *Suppose that each sensor j induces a linear action on cohomology generators via $D\Pi_j$, the Fisher metric provides a nondegenerate quadratic form on the residual nullspace, and sensor effects on generators are independent up to second order. Then*

$$j^* \in \arg \max_j \int_{\mathcal{F}_{\text{reg}}} \Delta_j(X) d\text{Vol}_{g_{\mathcal{F}}}(X) \iff j^* \in \arg \min_j \dim H^1(\mathcal{U} \cup \{U_j\}, \mathcal{S}_\Pi).$$

Proof. Each generator of H^1 corresponds to a direction $v \in N_X$. Sensor j acts on this generator via $D\Pi_j(X)v$. If $D\Pi_j(X)v \neq 0$, the generator is detected and eliminated: the cocycle becomes a coboundary under the refined cover. The Fisher contribution $\langle D\Pi_j(X)v, \Sigma_j^{-1} D\Pi_j(X)v \rangle$ measures the energy of elimination. Maximizing Δ_j integrated over \mathcal{F} maximizes aggregate elimination energy. Under independence, each generator contributes additively, so the sensor maximizing this sum achieves maximal rank reduction in H^1 . \square

The criterion has a direct geometric interpretation: it selects the sensor whose projections are most sensitive along directions that remain unconstrained by all existing modalities. The objective is not to “see more” but to eliminate the specific degrees of freedom that still admit alternative world-states.

5.9 The Morse-Variational Structure of Yarncrawler Dynamics

The refinement process admits a continuous limit in which the evolution of $[X]$ is governed by a gradient flow on the consistency functional.

Definition 5.3 (Yarncrawler Energy Functional).

$$\mathcal{E}(X) = \sum_{i \in \mathcal{I}} \|\Pi_i(X) - y_i\|_{\mathcal{Y}_i}^2 + \mathcal{R}(X).$$

The gradient flow under the Fisher metric g_F is

$$\frac{dX}{dt} = -\nabla_{g_F} \mathcal{E}(X),$$

which decreases \mathcal{E} monotonically. The critical structure of \mathcal{E} recovers the cohomological content of the framework.

Theorem 5.4 (Morse Structure of Residual Ambiguity). *Under generic conditions on $\{\Pi_i\}$ and \mathcal{R} , the functional \mathcal{E} is a Morse function on each regular stratum of \mathcal{F} , and:*

- (i) *Critical points of \mathcal{E} correspond to equivalence classes locally indistinguishable under all projections.*
- (ii) *The index of a critical point equals $\dim N_X$, the number of independent ambiguity directions.*
- (iii) *Nontrivial cohomology classes in $H^1(\mathcal{U}, \mathcal{S}_{\Pi})$ correspond to non-minimal critical points of \mathcal{E} .*
- (iv) *The unique global minimum corresponds to the closure state $[X^*]$ with $\mathcal{E}(X^*) = 0$.*

Proof sketch. At a critical point X_c , $\nabla \mathcal{E}(X_c) = 0$, which holds when all projection residuals vanish or the gradient is in equilibrium with $\nabla \mathcal{R}$. The Hessian $\nabla^2 \mathcal{E}(X_c) = g_F(X_c) + \nabla^2 \mathcal{R}(X_c)$ has nullspace equal to directions invisible to all projections and unpenalized by \mathcal{R} —exactly the generators of H^1 . Non-minimal critical points have positive-index nullspace corresponding to these generators. \square \square

Corollary 5.5 (Closure as Unique Morse Minimum). *If the conditions of Theorems 6.11 and 7.10 hold, then \mathcal{E} has a unique global minimum $[X^*]$ and all other critical points are unstable. Equivalently:*

$$\text{all cohomology classes trivialized} \iff \mathcal{E} \text{ has a unique minimum.}$$

5.10 Distributional Lift: Wasserstein Gradient Flow

The pointwise framework admits a distributional generalization. Treat the state not as a point estimate but as a probability measure $\mu \in \mathcal{P}(\mathcal{A})$ supported on the current feasible set.

Definition 5.6 (Measure-Valued Energy Functional).

$$\mathcal{E}[\mu] = \int_{\mathcal{A}} \left(\sum_i \|\Pi_i(X) - y_i\|^2 + \mathcal{R}(X) \right) d\mu(X) + \tau \text{Ent}(\mu),$$

where $\text{Ent}(\mu) = \int \mu \log \mu$ is the Boltzmann entropy and $\tau > 0$ is a temperature parameter encoding exploration-exploitation balance.

Endowing $\mathcal{P}(\mathcal{A})$ with the 2-Wasserstein metric W_2 , the gradient flow of $\mathcal{E}[\mu]$ is the continuity equation

$$\partial_t \mu_t = \nabla \cdot \left(\mu_t \nabla \frac{\delta \mathcal{E}}{\delta \mu} \right), \quad \frac{\delta \mathcal{E}}{\delta \mu}(X) = \sum_i \|\Pi_i(X) - y_i\|^2 + \mathcal{R}(X) + \tau \log \mu_t(X).$$

Under the conditions of Theorems 6.11 and 7.10, as $t \rightarrow \infty$ and $\tau \rightarrow 0$:

$$\mu_t \longrightarrow \delta_{[X^*]}, \quad \text{supp}(\mu_t) \rightarrow \{[X^*]\}.$$

The pointwise consistency operator is recovered as the zero-temperature limit: $\tau \rightarrow 0 \Rightarrow \mu_t \rightarrow \delta_{X(t)}$, recovering $\frac{dX}{dt} = -\nabla \mathcal{E}(X)$.

In the distributional formulation, cohomological obstruction manifests as multi-modality of μ_t : distinct connected components of $\text{supp}(\mu_t)$ correspond to distinct global sections. The elimination of cohomology is the transition from multi-modal to unimodal concentration.

5.11 RSVP Realization of the Variational Flow

The Wasserstein gradient flow of Section 5.8 is identical to RSVP field evolution under a constraint-induced potential. This establishes that the Yarncrawler refinement process is not merely analogous to RSVP dynamics; it is a specific instantiation of RSVP dynamics under constraint-driven forcing.

Represent the measure μ_t as a density field: $\mu_t(X) \leftrightarrow \Phi(x, t)$. The continuity equation $\partial_t \mu_t + \nabla \cdot (\mu_t u_t) = 0$ becomes the RSVP scalar conservation law $\partial_t \Phi + \nabla \cdot (\Phi \mathbf{v}) = 0$, with velocity field

$$\mathbf{v} = -\nabla V_{\text{proj}} - \nabla V_{\text{adm}} + \tau \nabla S,$$

where $V_{\text{proj}}(X) = \sum_i \|\Pi_i(X) - y_i\|^2$ is the projection potential, $V_{\text{adm}}(X) = \mathcal{R}(X)$ is the admissibility potential, and $S = -\log \Phi$ is the RSVP entropy field (so $-\tau \nabla \log \Phi = \tau \nabla S$).

The resulting field equation is

$$\partial_t \Phi + \nabla \cdot [\Phi(-\nabla V_{\text{proj}} - \nabla V_{\text{adm}} + \tau \nabla S)] = 0,$$

a forced advection-diffusion equation in which projection and admissibility terms drive directed collapse toward constraint-satisfying configurations, while the entropy term induces diffusive spreading preventing premature concentration.

Closure occurs when: (i) projection mismatch vanishes, $\Pi_i(X^*) = y_i$; (ii) admissibility is satisfied, $\mathcal{R}(X^*) = \min$; and (iii) entropy collapses, $S \rightarrow -\infty$ ($\Phi \rightarrow \delta_{[X^*]}$), giving $\partial_t \Phi \rightarrow 0$.

The full equivalence is:

$$\text{Yarncrawler refinement} \equiv \text{RSVP field evolution} \equiv \text{Wasserstein gradient flow.}$$

These are not three descriptions. They are one dynamical system written in categorical form (sheaf repair), geometric form (feasible set collapse), variational form (energy minimization), and physical form (RSVP fields). Every layer of the framework is the same object: sheaf theory eliminates H^1 , geometry collapses $\dim \mathcal{F}$, information theory collapses $\text{Vol}_F(\mathcal{F})$, dynamics execute the gradient flow of \mathcal{E} , and optimization drives the convergence of \mathcal{C} . These are different coordinate systems on one process.

6 The Consistency Operator and the Identifiability Theorem

6.1 Setup and Standing Assumptions

Let $\mathcal{X} = H^1(\Omega) \times H^1(\Omega; \mathbb{R}^3) \times L^2(\Omega)$ with the product Sobolev topology (Appendix A). A symmetry group G acts smoothly, properly, and isometrically on \mathcal{X} ; the quotient \mathcal{X}/\sim is Hausdorff. We work in \mathcal{X}/\sim throughout.

Assumption 6.1 (Projection Regularity). *Each $\Pi_i : \mathcal{X} \rightarrow \mathcal{Y}_i$ is a bounded linear map to a separable Hilbert space \mathcal{Y}_i .*

Assumption 6.2 (Noise Model). *Each observation satisfies $y_i = \Pi_i(X^*) + \eta_i$ where $\|\eta_i\|_{\mathcal{Y}_i} \leq \epsilon_i$ almost surely.*

Assumption 6.3 (Regularizer Structure). *The regularizer takes the form*

$$\mathcal{R}(X) = \mathcal{R}_{\text{RSVP}}(X) + \lambda_{\text{eco}} \mathcal{R}_{\text{eco}}(X) + \lambda_{\text{civ}} \mathcal{R}_{\text{civ}}(X)$$

where $\mathcal{R}_{\text{RSVP}}$ measures violation of the RSVP field equations in L^2 , \mathcal{R}_{eco} is zero on ecologically admissible configurations and strictly positive off them, and \mathcal{R}_{civ} is zero on civilizationaly admissible configurations and strictly positive off them.

Assumption 6.4 (Admissibility Region). *The admissible set $\mathcal{A} = \{[X] \in \mathcal{X}/\sim \mid \mathcal{R}(X) < \infty\}$ is a closed convex subset of \mathcal{X}/\sim , and \mathcal{R} restricted to \mathcal{A} is strictly convex along geodesics.*

Assumption 6.5 (Banach Realization). *There exists a Banach space \mathcal{B} and continuous embedding $\iota : \mathcal{X} \hookrightarrow \mathcal{B}$ such that \mathcal{A} is closed and convex in \mathcal{B} , each Π_i extends to a bounded linear operator on \mathcal{B} , and \mathcal{R} extends to a proper lower semicontinuous functional on \mathcal{B} .*

In practice one takes $\mathcal{B} = \mathcal{X}$ itself, which is already a Hilbert space. The Banach realization is stated separately to make the compactness and weak convergence arguments in Appendix B self-contained.

6.2 The Feasible Set

Definition 6.6 (Observation-Consistent Set).

$$\Omega_{\text{obs}} = \{[X] \in \mathcal{A} \mid \|\Pi_i(X) - y_i\|_{\mathcal{Y}_i} \leq \epsilon_i \quad \forall i \in \mathcal{I}\}.$$

Lemma 6.7 (Non-emptiness). *Under Assumptions 6.1 and 6.2, $X^* \in \Omega_{\text{obs}}$ almost surely.*

Proof. $\|y_i - \Pi_i(X^*)\| = \|\eta_i\| \leq \epsilon_i$ almost surely by Assumption 6.2. \square \square

Lemma 6.8 (Closedness and Convexity). *Under Assumptions 6.1 and 6.4, Ω_{obs} is a closed convex subset of \mathcal{A} .*

Proof. Each constraint set is closed (continuity of Π_i) and convex (linearity of Π_i , convexity of norm ball). Intersection of closed convex sets is closed and convex. \square \square

The dynamically admissible set Γ is introduced as a named object here; its full definition and construction are given in Section 7. The *feasible set* is $\mathcal{F} = \Omega_{\text{obs}} \cap \Gamma$.

6.3 The Consistency Operator

Definition 6.9 (Consistency Operator).

$$\mathcal{C}(X) = \arg \min_{[X'] \in \mathcal{A}} \left[\sum_{i \in \mathcal{I}} \|\Pi_i(X') - y_i\|_{\mathcal{Y}_i}^2 + \mathcal{R}(X') \right].$$

\mathcal{C} maps \mathcal{A} to \mathcal{A} : the optimization is constrained to the admissible region. A configuration minimizing projection error but violating RSVP equations or ecological plausibility is not a candidate world-state. The regularizer defines the domain of the operator, not a preference over an unconstrained space.

Lemma 6.10 (Well-definedness). *Under Assumptions 6.1–6.5, $\mathcal{C}(X)$ exists and is unique for every $[X] \in \mathcal{A}$.*

Proof. The objective $f(X') = \sum_i \|\Pi_i(X') - y_i\|^2 + \mathcal{R}(X')$ is continuous, strictly convex (sum of convex functions plus strictly convex \mathcal{R}), and coercive on \mathcal{A} . A continuous strictly convex coercive function on a closed convex set attains its minimum at a unique point. Full proof in Appendix B. \square \square

6.4 The Identifiability Theorem

The induced map of the projection family on the quotient is:

$$\tilde{\Pi} : \mathcal{X}/\sim \longrightarrow \prod_{i \in \mathcal{I}} \mathcal{Y}_i, \quad [X] \mapsto (\Pi_i(X))_{i \in \mathcal{I}}.$$

By definition of \sim , $\tilde{\Pi}$ is well-defined.

Theorem 6.11 (Yarncrawler Identifiability). *Under Assumptions 6.1–6.5, the consistency operator \mathcal{C} has a unique fixed point $[X^*] \in \mathcal{F}$ if and only if $\tilde{\Pi}$ is injective on \mathcal{F} .*

Proof. (\Rightarrow) Suppose \mathcal{C} has a unique fixed point $[X^*] \in \mathcal{F}$. If $\tilde{\Pi}$ were not injective on \mathcal{F} , there exist distinct $[X], [X'] \in \mathcal{F}$ with $\Pi_i(X) = \Pi_i(X')$ for all i . Both achieve zero projection residual in the objective. The objective reduces to $\mathcal{R}(X)$ vs. $\mathcal{R}(X')$. Either $\mathcal{R}(X) = \mathcal{R}(X')$ —violating strict convexity of \mathcal{R} on \mathcal{A} (Assumption 6.4)—or one has strictly smaller regularizer value, making the other not a minimizer. Either way contradicts uniqueness.

(\Leftarrow) If $\tilde{\Pi}$ is injective on \mathcal{F} , then $\Omega_{\text{obs}} \cap \mathcal{F}$ contains at most one point. By Lemma 5.1, $[X^*] \in \Omega_{\text{obs}}$; by definition $[X^*] \in \Gamma$; hence $[X^*] \in \mathcal{F}$. Injectivity forces $|\mathcal{F}| = 1$. At $[X^*]$ all projection residuals are within tolerance and $\mathcal{R}(X^*)$ is minimized over the singleton \mathcal{F} . Hence $\mathcal{C}([X^*]) = [X^*]$. \square

6.5 The Corollary on Diagnosable Failure Modes

Corollary 6.12. *Identifiability fails in exactly two geometrically distinguishable ways:*

(i) Projection degeneracy: $\tilde{\Pi}$ is not injective on \mathcal{F} . Two distinct admissible world-states produce identical observations. This failure is intrinsic to the sensor family and cannot be resolved by the regularizer.

(ii) Regularizer flatness: \mathcal{R} is not strictly convex on \mathcal{F} . The regularizer assigns equal admissibility to configurations the sensors cannot distinguish. This failure is intrinsic to the model class.

These failure modes are geometrically distinguishable: projection degeneracy corresponds to a non-trivial kernel of $\tilde{\Pi}|_{\mathcal{F}}$; regularizer flatness corresponds to a flat direction of \mathcal{R} within \mathcal{F} . They can coexist.

6.6 Why the Quotient Structure Is Not Optional

In a standard inverse problem one works in \mathcal{X} directly. But in the Yarncrawler setting, two configurations related by a gauge transformation or coordinate reparametrization represent the same physical situation. Working in \mathcal{X} rather than \mathcal{X}/\sim would make uniqueness false by construction—there would always be a continuum of configurations achieving the same objective, related by the action of G .

An L^2 uniqueness statement conflates field-theoretic gauge redundancy with measurement noise. A measure-theoretic statement collapses the geometry of \mathcal{X}/\sim into observational statistics and loses the regularizer’s role as an admissibility condition. The symmetry-group formulation is not a choice among equivalent formalisms; it is the only formulation that preserves the semantic commitments of the framework.

7 Trajectory Learning and the Dynamical Admissibility Set

7.1 The Central Obstruction

Section 6 established that \mathcal{C} has a unique fixed point iff $\tilde{\Pi}$ is injective on \mathcal{F} . This condition is clean but generically unmet. Under realistic sensing, Ω_{obs} is high-dimensional. Projection degeneracy is the default condition, not a pathological case.

The resolution cannot come from within the static framework. What is required is a constraint of a different kind—one operating not on the configuration X at a single time but on the trajectory through which X was reached. The dynamical admissibility set Γ is not an additional observation; it is a restriction on which configurations are reachable under the dynamics of the world. Two configurations in Ω_{obs} that are observationally indistinguishable at one moment may evolve under those dynamics into configurations that are distinguishable at a later moment.

7.2 The Path Space and the Stochastic Kernel

Definition 7.1 (Discrete Path Space). $\mathcal{P}_\Delta = \{\mathbf{X} = ([X_0], [X_1], \dots, [X_N]) \mid [X_k] \in \mathcal{A}, N = \lfloor T/\Delta \rfloor\}$.

Definition 7.2 (Stochastic Transition Kernel). $T_\Delta : \mathcal{A} \rightarrow \mathcal{P}(\mathcal{A})$, where for each $[X] \in \mathcal{A}$, $T_\Delta([X], \cdot)$ specifies the distribution over one-step successors under the world dynamics at resolution Δ .

Definition 7.3 (Dynamically Admissible Set).

$$\Gamma = \left\{ [X] \in \mathcal{A} \mid \exists [X_0] \in \mathcal{A}, \exists k \geq 0 \text{ s.t. } \mu_{X_0}^{T_\Delta}([X_k] = [X]) > 0 \right\}.$$

Γ is the support of the union of all admissible path measures. It is a geometric object—a subset of \mathcal{A} —not a probabilistic one.

Lemma 7.4 (Invariance of Γ). *If T_Δ is time-homogeneous and \mathcal{A} is forward-invariant under T_Δ , then $\Gamma \subseteq \mathcal{A}$ and Γ is forward-invariant.*

7.3 The Deterministic Limit and Tangent Bundle Constraint

Assumption 7.5 (Generator Regularity). *There exists a Lipschitz vector field $F(\cdot; \theta) : \mathcal{A} \rightarrow T\mathcal{A}$ such that $T_\Delta([X], \cdot) \Rightarrow \delta_{[X] + \Delta F([X]; \theta)}$ weakly as $\Delta \rightarrow 0$.*

In the zero-noise limit, the dynamically admissible set reduces to:

$$\Gamma^0 = \{ [X] \in \mathcal{A} \mid \exists [X_0] \in \mathcal{A}, \exists t \in [0, T] \text{ s.t. } \Phi_F^t([X_0]) = [X] \}$$

where Φ_F^t is the flow of F at time t .

Definition 7.6 (Admissible Tangent Cone). *For $[X] \in \Gamma^0$: $\mathcal{T}_{[X]}\Gamma^0 = \{v \in T_{[X]}\mathcal{A} \mid v = F([X]; \theta) \text{ for some } \theta \in \Theta\}$.*

Proposition 7.7 (Tangent Dimension Reduction). *If $\Theta \subset \mathbb{R}^p$, then $\dim \mathcal{T}_{[X]}\Gamma^0 \leq p$. In particular, if $p < \dim T_{[X]}\mathcal{A}$, the dynamical constraint reduces the effective dimensionality of admissible directions from $\dim T_{[X]}\mathcal{A}$ to at most p .*

Proof. The map $\theta \mapsto F([X]; \theta)$ is smooth from $\Theta \subset \mathbb{R}^p$ to $T_{[X]}\mathcal{A}$; its image has dimension at most p by the rank theorem. \square

7.4 Trajectory Learning from Before/After Pairs

Definition 7.8 (Fragment-Consistent Parameter Set). *Given fragment $(y_{\text{before}}, y_{\text{after}}, \tau)$:*

$$\Theta_{(y_b, y_a, \tau)} = \{\theta \in \Theta \mid \exists [X_b] \in \Omega_{\text{obs}}(y_b), \Pi_i(\Phi_F^\tau([X_b])) \in B_{\epsilon_i}(y_{a,i}) \forall i\}.$$

Given training set \mathcal{D} , the learned parameter set is $\hat{\Theta} = \bigcap_n \Theta_{(y_b^{(n)}, y_a^{(n)}, \tau^{(n)})}$.

The operator class library. In practice $F(\cdot; \theta)$ is decomposed as:

$$F([X]; \theta) = \sum_{k=1}^K \alpha_k([X]) \cdot F_k([X]; \theta_k)$$

where each F_k corresponds to a named ecological or civilizational process—forest succession, urban expansion, agricultural cycling, infrastructure decay, hydrological change, demographic transition. Each F_k is learned from before/after pairs in which that process is dominant. The activation weights $\alpha_k([X]) \geq 0$ satisfy $\sum_k \alpha_k = 1$.

7.5 The Dynamical Disambiguation Theorem

Definition 7.9 (Trajectory Separation). *Two configurations $[X], [X'] \in \Omega_{\text{obs}}$ are trajectory-separated by $F(\cdot; \hat{\theta})$ at time $\tau > 0$ if*

$$\left\| \tilde{\Pi}(\Phi_F^\tau([X]; \hat{\theta})) - \tilde{\Pi}(\Phi_F^\tau([X']; \hat{\theta})) \right\|_{\Pi_i \mathcal{Y}_i} > \epsilon$$

for some $\epsilon > 0$ depending only on the noise levels $\{\epsilon_i\}$.

Theorem 7.10 (Dynamical Disambiguation). *Under Assumptions 6.1–7.5, if $|\Omega_{\text{obs}}| > 1$ and every pair $[X], [X'] \in \Omega_{\text{obs}}$ with $[X] \neq [X']$ is trajectory-separated by $F(\cdot; \hat{\theta})$ at some $\tau \in (0, T]$, then $|\mathcal{F}| = |\Omega_{\text{obs}} \cap \Gamma^0| = 1$.*

Proof. Let $[X], [X'] \in \Omega_{\text{obs}} \cap \Gamma^0$ be distinct. Since both lie in Γ^0 , both lie on integral curves of $F(\cdot; \hat{\theta})$. By trajectory separation, there exists $\tau \in (0, T]$ such that the forward-evolved projections differ by more than ϵ . At time τ , observed data $y(\tau)$ must lie within ϵ_i of the true trajectory; therefore at most one of $[X(\tau)], [X'(\tau)]$ lies in $\Omega_{\text{obs}}(y(\tau))$. The true state $[X^*(\tau)] \in \Omega_{\text{obs}}(y(\tau))$ by Lemma 5.1 at time τ . Applying to all pairs eliminates all but $[X^*]$. Hence $|\mathcal{F}| = 1$. \square

7.6 Joint Identifiability

Corollary 7.11 (Joint Identifiability). *The consistency operator \mathcal{C} restricted to \mathcal{F} has a unique fixed point if either: (i) $\tilde{\Pi}$ is injective on \mathcal{F} ; (ii) every pair in $\Omega_{\text{obs}} \cap \Gamma^0$ is*

trajectory-separated; or (iii) both conditions hold partially. The sufficient condition for (iii) is that $\dim(\Omega_{\text{obs}}/\ker \tilde{\Pi}) + \dim(\Gamma^0/\Gamma_{\text{indist}}^0) \geq \dim \Omega_{\text{obs}}$.

Case (iii) is the generic realistic scenario: sensors eliminate most ambiguity, trajectory learning eliminates the residual. The two mechanisms are additive in their dimensional reduction, not competing.

8 The Inversion Layer

8.1 Role Specification

The inversion layer is not the model. It is a proposal mechanism: its role is to place the iterative optimization in the basin of the correct fixed point. The consistency operator \mathcal{C} is the model; the encoder is a high-entropy initializer; the physics-constrained forward model is the contraction operator. Conflating encoder with solution is the failure mode of standard encoder-decoder stacks: doing so makes the two failure modes of Corollary 6.12 undiagnosable. Separating proposal from decision preserves the falsifiability structure.

8.2 The Encoder as Approximate Posterior

The encoder implements an approximate posterior:

$$q_\phi(X | \{y_i\}) \approx p(X | \{y_i\}) \propto \exp\left(-\sum_i d_i(\Pi_i(X), y_i) - \mathcal{R}(X)\right).$$

The encoder approximates this distribution; it does not solve it. The forward-consistency loop is what contracts that distribution. Because \mathcal{R} encodes physical admissibility, the approximate posterior is already biased toward physically plausible configurations: the encoder does not need to learn the RSVP dynamics from scratch, only to propose configurations that are approximately consistent with them. The encoder minimizes KL divergence from the true posterior, subject to its architectural constraints, treating \mathcal{R} as an implicit prior.

8.3 The Refinement Loop

The five-step loop is:

1. Encoder samples $X^{(0)} \sim q_\phi(\cdot | \{y_i\})$
2. Forward model computes $\{\hat{y}_i\} = \{\Pi_i(X^{(0)})\}$ using the operators of Section 4
3. Residuals $r_i = d_i(\hat{y}_i, y_i)$ drive a projected gradient update
4. $X^{(k+1)} = \mathcal{C}(X^{(k)})$, one step of projected gradient descent on the consistency functional
5. Iterate until $\text{Vol}_F(\mathcal{F}) < \varepsilon$ (closure criterion of Section 9)

Theorem 6.11 guarantees a unique fixed point when the identifiability conditions hold. The refinement loop approximates the path to that fixed point. Convergence rate is governed by the condition number of $\nabla^2\mathcal{E}$ at $[X^*]$, equal to $g_F(X^*) + \nabla^2\mathcal{R}(X^*)$, as established in Appendix D.

8.4 Categorical Structure of Multi-Modal Fusion

The encoder decomposes into modality-specific proposal maps $q_\phi^{(i)} : \mathcal{Y}_i \rightarrow \mathcal{P}(\mathcal{A})$ composed via a lax monoidal structure over $\prod_i \mathcal{Y}_i$. Each modality-specific encoder produces a marginal proposal; the fusion map combines them before the consistency operator acts. This describes how multi-modal proposals are combined, but the explanatory weight rests on the posterior formulation, not on the categorical structure. The categorical framing is reserved for structural analysis of how proposals compose across deployment scales (Section 10).

8.5 The Forward Model

Given a proposed $[X]$, the forward model produces $\{\Pi_i(X)\}$ using the operators of Section 4. For infrared and LiDAR, this is direct evaluation of the trace and level-set operators. For RGB, it is a volume rendering pass using the NeRF integral. For audio, it requires integrating the RSVP field equations forward in time from $[X]$ to produce $\partial_t S$ and $\nabla \cdot \mathbf{v}$ at microphone positions. For photonic strips, it evaluates the transect trace operator over $[0, T]$. The computational cost is highest for audio and strip projections; learned surrogate models replace the full forward pass in those cases during iterative refinement, with the exact operators used only for final validation.

8.6 Why Not End-to-End

An end-to-end system collapses proposal and decision into a single learned map $\{y_i\} \mapsto [X]$. This replaces \mathcal{C} with the encoder and discards the forward model. The result cannot diagnose its own failure modes (Corollary 6.12 becomes inaccessible), cannot incorporate new sensor modalities without retraining, and cannot be falsified by physical constraints. The separation of encoder from consistency operator is not an engineering preference. It is required by the identifiability structure: semantics are not learned directly; they are selected via constraint closure.

9 Closure: Definition, Criteria, and Detection

9.1 The Unified Definition

A state $[X^*] \in \mathcal{X}/\sim$ is *closed* if it is the unique fixed point of \mathcal{C} and lies on a trajectory of T_Δ consistent with all observed fragments. This definition decomposes into three operationally distinct conditions that together constitute closure in a deployed system.

9.2 The Three Conditions

Projection parity: $\|\Pi_i(X^*) - y_i\|_{\mathcal{Y}_i} \leq \epsilon_i$ for all $i \in \mathcal{I}$. The fixed point is consistent with all current observations. This is verifiable in real time by evaluating the forward model.

Dynamical stability: $[X^*] \in \Gamma^0$, meaning $[X^*]$ lies on an integral curve of $F(\cdot; \hat{\theta})$. The fixed point is reachable under the learned dynamics, not merely consistent with instantaneous observations. This rules out configurations that fit current data but could not have been produced by any admissible trajectory.

Counterfactual consistency: if $[X^*]$ is simulated forward under F for time τ and reprojected, the result must lie within ϵ_i of actual observations at time τ . This is the strongest condition: it requires the fixed point's forward trajectory to be consistent with future observations, not just its current projection.

9.3 Equivalence of the Three Conditions

Theorem 9.1 (Equivalence of Closure Criteria). *Under the assumptions of Theorems 6.11 and 7.10, the three conditions of Section 8.2 are equivalent. Each implies $\dim_{\text{eff}}(\mathcal{F}) = 0$, and $\dim_{\text{eff}}(\mathcal{F}) = 0$ implies all three.*

Proof. Projection parity tightens Ω_{obs} until $[X^*]$ is its only element. Dynamical stability ensures $[X^*] \in \Gamma^0$, so $\mathcal{F} = \{[X^*]\}$. Under the Sheaf-Variational Equivalence, $\dim_{\text{eff}}(\mathcal{F}) = 0$ is equivalent to $H^0 = 1$ and $H^1 = 0$. Counterfactual consistency is the temporal extension of projection parity: the sheaf condition holds over the spatio-temporal cover, not just the spatial cover. Each condition therefore implies $|\mathcal{F}| = 1$, and $|\mathcal{F}| = 1$ implies all three. \square \square

The three conditions are not independent criteria to be checked separately. They are three faces of the single condition $|\mathcal{F}| = 1$.

9.4 Operational Detection

Closure is operationally detectable as stability of $[X^*]$ under perturbation of ϵ_i : if tightening any constraint by $\delta\epsilon_i$ does not move the fixed point, closure has been achieved. This provides the stopping criterion for the refinement loop of Section 8. The rate of convergence to closure, as a function of sensor density and trajectory fragment coverage, is governed by the eigenspectrum of $g_F(X^*) + \nabla^2\mathcal{R}(X^*)$ established in Appendix D.

9.5 Entropy Collapse at Closure

At closure, the following hold simultaneously: $\text{Vol}_F(\mathcal{F}) = 0$, $H^1(\mathcal{U}, \mathcal{S}_{\Pi}) = 0$, $\dim H^0 = 1$, and $S_{\text{RSVP}}(X^*)$ is minimized subject to admissibility. By Theorem D.6 and Corollary D.8, these are equivalent under nondegeneracy of $g_F^{\mathcal{F}}$. The RSVP entropy field at closure is the local density of eliminated cohomological obstruction. The closed state is not a best estimate; it is the unique trajectory consistent with all projections and admissible dynamics.

9.6 Failure Modes Revisited

Corollary 6.12 identified two failure modes. Projection degeneracy ($H^1 \neq 0$) means closure cannot be achieved by additional observation at the current moment: it requires new sensor modalities (refining the cover) or temporal extension (Section 7’s disambiguation mechanism). Regularizer flatness means closure cannot be achieved by any sensor: it requires revision of the model class \mathcal{A} itself. These are the only two ways the refinement loop can fail to converge, and both are diagnosable by the criteria of this section.

10 Deployment Scales as Feasible Set Geometry

10.1 The Unified Framework

All three deployment scales are instances of the same consistency operator \mathcal{C} with different $(\Pi_i, \mathcal{N}_i, \epsilon_i)$ profiles. The scale is not a separate architectural choice; it is a parameter of the noise model and the available projection family. The identifiability and disambiguation theorems apply uniformly across scales.

10.2 Planetary Scale

At planetary scale, sensors are satellite multispectral imagery (RGB and infrared) with occasional LiDAR passes; audio and photonic strips are absent. Temporal resolution is weeks to months; spatial resolution 10–100 meters per pixel; noise levels ϵ_i are large. The feasible set Ω_{obs} is high-dimensional and weakly constrained. The system operates in the high-entropy proposal regime, relying heavily on \mathcal{R} to restrict \mathcal{F} . Trajectory learning carries the disambiguation burden: seasonal cycles, annual land cover change, and multi-year succession patterns provide the operator class library entries that contract Γ . The projection family is incomplete (Proposition 4.10 fails due to missing audio and strip modalities), but Corollary ?? applies: trajectory learning compensates for the missing static constraints.

10.3 Regional Scale

At regional scale, sensors include aerial multispectral and LiDAR surveys, distributed acoustic sensors, and periodic photonic strip transects. Temporal resolution is days to weeks; spatial resolution 1–10 meters. Ω_{obs} contracts along spatial dimensions but remains temporally undersampled. Both static projection richness and trajectory learning are required. The transition from planetary to regional scale is a cover refinement: the planetary-scale sheaf \mathcal{S}_{Π} restricts to the regional-scale sheaf $\mathcal{S}_{\Pi'}$ on the finer cover, with H^1 classes killed by the additional modalities.

10.4 Site Scale

At site scale, fixed arrays include all five modalities at high spatial and temporal density: seconds to minutes temporal resolution, centimeters to meters spatial resolution, small

ϵ_i . Ω_{obs} collapses rapidly toward a singleton. The system achieves genuine closure. Re-projection residuals (forward model outputs against actual sensor readings) provide a continuous closure detection signal implementing the operational criterion of Section 9.

10.5 Multi-Scale Fusion and Scale Consistency

Theorem 10.1 (Scale Consistency). *If the noise models $\{\mathcal{N}_i\}$ are correctly specified at each scale, the fixed point $[X^*]$ is invariant to which scale drives the final refinement step.*

Proof. The functor Ψ of Theorem 2.6 maps cover refinements to sheaf restrictions. The restriction map from the finer (site-scale) sheaf to the coarser (planetary-scale) sheaf preserves the global section, since the global section at the finer scale restricts to a global section at the coarser scale. Uniqueness at each scale (under the identifiability conditions) forces these to agree. \square

In practice: planetary constraints provide the prior on \mathcal{R} , regional constraints initialize q_ϕ , and site constraints drive the refinement loop to closure. The three scales are nested constraint systems, not separate deployments.

10.6 Sensor Selection Across Scales

The optimal sensor selection criterion of Theorem 5.2 applies uniformly across scales. At planetary scale, the most valuable next sensor is the one that closes the largest-dimensional direction in $\ker \tilde{\Pi}|_{\mathcal{F}}$: typically audio (constrains $\partial_t S$ and $\nabla \cdot \mathbf{v}$, both unconstrained by satellite imagery) or high-frequency temporal sampling (constrains $\partial_t \Phi$). At site scale, where most static degrees of freedom are already closed, the criterion favors temporal extension of the observation window over additional simultaneous modalities.

11 Relation to Prior Work

11.1 Data Assimilation and Ensemble Methods

Data assimilation methods, including ensemble Kalman filtering (EnKF) and variational 4D-Var, address the same broad problem of hidden state estimation from observations over time. Three structural differences distinguish the Yarncrawler framework. First, data assimilation works in \mathcal{X} rather than \mathcal{X}/\sim , treating gauge redundancies as noise rather than structure. Second, the regularizer in data assimilation is a prior distribution expressing preference, not an admissibility condition expressing exclusion. Third, data assimilation has no sheaf structure: there is no notion of local charts, restriction maps, or cohomological obstruction, and no counterpart to the Identifiability Theorem. In the linear Gaussian regime with trivial symmetry group, \mathcal{C} reduces to the Kalman update, making EnKF a degenerate special case.

11.2 Inverse Rendering and Neural Radiance Fields

Neural radiance fields (NeRF) and related inverse rendering methods reconstruct 3D scenes from 2D images. The RGB projection operator Π_{RGB} of Section 4 is the NeRF volume rendering integral. Three differences are structural. NeRF has no dynamics: it is a static reconstruction with no dynamical admissibility set Γ . NeRF has no RSVP field structure: the scene representation is a learned radiance function, not a physical field triple. NeRF has no regularizer encoding physical admissibility. The Yarncrawler framework strictly generalizes NeRF: RGB is one chart of the cover, and NeRF is the degenerate case where only that chart is available and dynamics are absent.

11.3 Digital Twin Systems

Digital twin literature proposes maintaining synchronized virtual replicas of physical systems. The standard approach is object-ontological: the twin represents a collection of objects whose properties are updated from sensor data. Yarncrawler departs in three ways. The primary ontology is trajectory-based: objects emerge as persistent patterns in field trajectories, not as primary entities. The twin is not updated by recognized objects but by the consistency operator acting on the full quotient configuration space. Most importantly, digital twins have no identifiability theory: there is no formal account of when a twin uniquely determines its physical counterpart. The Yarncrawler Identifiability Theorem provides exactly this.

11.4 Categorical and Sheaf-Theoretic Semantics

Sheaf theory has been applied to contextuality in quantum mechanics, to distributed knowledge in multi-agent systems, and to compositional natural language semantics. The Yarncrawler framework draws on these applications but departs in one crucial direction: the sheaf here is over a physical configuration space with RSVP field structure, not over an abstract semantic space. The stalks are Sobolev-class field configurations satisfying physical equations. The repair morphisms are steps of a variational optimization. The entropy identified with H^1 is a physical field, not an information-theoretic measure. This grounding allows the framework to make contact with sensor physics and deployment engineering while retaining the full generality of the categorical structure.

12 Discussion: Closure as Epistemology

12.1 What Closure Means

A closed Yarncrawler state is not a probabilistic estimate. It is not the most likely world-state given the observations. It is the unique world-state that is structurally demanded by the observations under the model class \mathcal{A} . The distinction has operational consequences: a probabilistic estimate can always be revised by new observations, while a closed state cannot be revised without either adding new sensor modalities (killing new H^1 classes) or

revising the model class. The closed state is, in a precise sense, as certain as the model class allows.

12.2 Constraint Primacy

The framework operationalizes a structural commitment: constraints are prior to completions. The world-state is not recovered by filling in a model template with observations; it is the unique configuration surviving after all constraints have been applied. The Sheaf-Variational Equivalence makes this precise: a world-state has meaning (is uniquely identified) exactly when the sheaf induced by its projections has a unique global section. This is constraint primacy as a mathematical theorem rather than a philosophical position.

12.3 The RSVP Fields as Constraint Carriers

The RSVP field triple (Φ, \mathbf{v}, S) is not merely a convenient parametrization of the configuration space. It is the structure that makes the regularizer \mathcal{R} non-arbitrary: the admissible set \mathcal{A} is defined by the RSVP field equations, which encode physical conservation laws, ecological succession principles, and civilizational feasibility constraints. Without the RSVP structure, the regularizer would be an arbitrary penalty and the closed state would be an artifact of the choice of penalty. With it, the closed state is the unique configuration consistent with the governing laws of the domain.

12.4 Stigmergic Accumulation and the Refinement Loop

Each step of the refinement loop of Section 8 is a repair morphism in the sheaf \mathcal{S}_{Π} : it reduces a projection residual (trivializes a cocycle) and updates the field configuration toward closure. The accumulation of constraints over time — as more sensor observations arrive — is the variational analogue of stigmergic reinforcement: each new observation narrows the feasible set. This is the physical realization of the self-refactoring dynamics that characterize the Yarncrawler framework in its forward direction. The refinement loop converges when no more constraints can be satisfied — when the global section is unique and the yarn has been fully crawled.

12.5 Limits of the Framework

Three structural limits warrant acknowledgment.

Model class dependence: closure is unique within \mathcal{A} , not in \mathcal{X} . If \mathcal{A} is misspecified — if the true world-state lies outside the admissible set — the consistency operator converges to the closest admissible state, not the true one. Non-existence of a global section is the falsification signal requiring model revision.

Trajectory learning quality: dynamical disambiguation depends on the quality of $F(\cdot; \hat{\theta})$. If the operator class library is missing a dominant process type, trajectory separation fails along the missing direction, detectable as residual $|\mathcal{F}| > 1$ after full sensor coverage.

Computational tractability: \mathcal{C} is defined as a global minimizer of a non-convex functional. The refinement loop approximates with local gradient steps. Convergence to the global

minimizer is guaranteed only up to initialization quality; multiple encoder samples may be required.

12.6 Forward and Inverse as Dual Inference

The forward direction maintains coherence by repairing tears as they arise: it is the epistemic activity of a system embedded in a world it partially controls. The inverse direction reconstructs state from observations: it is the epistemic activity of a system observing a world it does not control. The Yarncrawler framework is the unified theory of both activities. The sheaf \mathcal{S}_{Π} is the common object; the two directions are two ways of relating to it. A fully realized system alternates between them: reconstructing from observations, projecting forward, detecting model failure at divergence, and repairing when failure is detected. This is the complete epistemic loop.

13 Conclusion

The Yarncrawler Identifiability Theorem and the Dynamical Disambiguation Theorem characterize the conditions under which multi-modal sensor observations uniquely determine a physical world-state. The Sheaf-Variational Equivalence shows that this characterization is not merely a technical result about inverse problems: it is the variational realization of a categorical structure in which sensor modalities are charts, observations are local sections, and closure is the existence of a unique global section. The RSVP field triple is the substrate in which these structures are physically grounded. The Morse structure of the consistency functional, the Wasserstein gradient flow over world-state measures, and the RSVP field evolution under constraint forcing are not three descriptions of the framework — they are one dynamical system in three coordinate systems.

A world-state is uniquely recoverable from multi-modal observations if and only if the induced map from the quotient configuration space to the product of observation spaces is injective on the feasible set, or if the dynamics of the world separate configurations that static sensors cannot. In either case, closure corresponds simultaneously to cohomology trivialization, entropy collapse, Fisher volume collapse, and convergence of the Wasserstein gradient flow to a Dirac measure. These are not independent conditions; they are one condition in four languages.

Three open problems remain. First, the extension of the identifiability theorem to the stochastic case without the low-noise limit: does a version hold when σ^2 is not small? Second, the learning theory of $F(\cdot; \hat{\theta})$: what sample complexity is needed for the operator class library to ensure trajectory separation with high probability? Third, the connection between the Fisher information geometry of the feasible set and the computational complexity of the refinement loop: can convergence rate be bounded in terms of the eigenspectrum of the Fisher metric at $[X^*]$?

The Yarncrawler system is not a sensing architecture. It is a theory of when the world becomes legible — when the observations accumulated by a sensor network, constrained by physical law and dynamical history, leave no alternative configuration standing. That

condition is not an approximation target. It is a threshold: on one side, multiple histories remain consistent with the evidence; on the other, only one does. The framework identifies that threshold precisely and provides the operator that reaches it.

A Fréchet Manifold Structure of \mathcal{X}

A.1 Fréchet Manifolds

A Fréchet manifold is a topological space locally homeomorphic to a Fréchet space (a complete metrizable locally convex topological vector space) with smooth transition maps. The configuration space $\mathcal{X} = H^1(\Omega) \times H^1(\Omega; \mathbb{R}^3) \times L^2(\Omega)$ is a Hilbert space, hence also a Banach space and a Fréchet space, with its product inner product. The Fréchet framing is introduced because extensions to more general field configurations (distributional fields, fields on non-compact domains) may require genuine Fréchet rather than Hilbert structure.

A.2 The Quotient Topology

The symmetry group G acts on \mathcal{X} smoothly, properly, and isometrically. Under these conditions the quotient \mathcal{X}/G is Hausdorff. The metric on \mathcal{X}/\sim is induced from \mathcal{X} via $d([X], [X']) = \inf_{g \in G} \|X - gX'\|_{\mathcal{X}}$.

The observational quotient \mathcal{X}/\sim (by \sim) is in general coarser than \mathcal{X}/G . They coincide when the projection family is complete (Proposition 4.10), because then every sensor-invisible direction is also G -equivalent.

A.3 Tangent Bundle of \mathcal{X}/\sim

The tangent bundle $T\mathcal{X}/\sim$ is the quotient of $T\mathcal{X}$ by the action of TG . When G acts freely, $T\mathcal{X}/\sim \cong T\mathcal{X}/TG$ is a vector bundle. The admissible tangent cone $\mathcal{T}_{[X]}\Gamma^0$ (Definition 6.5 in the main text) is a linear subspace of $T_{[X]}\mathcal{X}/\sim$ of dimension at most $p = \dim \Theta$ (Proposition 6.6).

A.4 Geodesic Convexity on \mathcal{X}/\sim

Geodesics on \mathcal{X}/\sim are the images of straight lines in \mathcal{X} under the quotient map. The admissible set \mathcal{A} is geodesically convex when its preimage in \mathcal{X} is convex, which holds by definition of \mathcal{A} as the sublevel set $\{\mathcal{R}(X) < \infty\}$ of the convex functional \mathcal{R} . This justifies Assumption 6.4.

B Full Proof of the Identifiability Theorem

B.1 Functional-Analytic Setting

The proof of Theorem 6.11 relies on existence and uniqueness of minimizers of the consistency functional on \mathcal{X}/\sim . We introduce additional assumptions for the infinite-dimensional setting.

Assumption B.1 (Banach Realization, restated). *There exists a Banach space \mathcal{B} and continuous embedding $\iota : \mathcal{X} \hookrightarrow \mathcal{B}$ such that \mathcal{A} is closed and convex in \mathcal{B} , each Π_i extends to a bounded linear operator on \mathcal{B} , and \mathcal{R} extends to a proper, lower semicontinuous functional on \mathcal{B} . In practice, $\mathcal{B} = \mathcal{X}$ itself.*

Assumption B.2 (Geodesic Convexity, restated). *\mathcal{A} is geodesically convex with respect to the metric induced from \mathcal{B} , and \mathcal{R} is strictly convex along geodesics in \mathcal{A} :*

$$\mathcal{R}(\gamma(t)) < (1-t)\mathcal{R}([X]) + t\mathcal{R}([X']) \quad \forall [X] \neq [X'] \in \mathcal{A}, t \in (0,1).$$

Assumption B.3 (Coercivity). *$f([X]) = \sum_i \|\Pi_i(X) - y_i\|^2 + \mathcal{R}(X)$ satisfies $\|X_n\|_{\mathcal{B}} \rightarrow \infty \Rightarrow f([X_n]) \rightarrow \infty$.*

B.2 Existence of a Minimizer

Proposition B.4. *Under the standing assumptions, f attains a minimum on \mathcal{A} .*

Proof. Let $\{[X_n]\}$ be a minimizing sequence; coercivity gives boundedness in \mathcal{B} . Reflexivity of \mathcal{B} (if Hilbert) yields a weakly convergent subsequence $X_{n_k} \rightharpoonup X_*$. Closedness of \mathcal{A} gives $[X_*] \in \mathcal{A}$. Weak continuity of Π_i (bounded linear) and lower semicontinuity of \mathcal{R} give $f([X_*]) \leq \liminf f([X_{n_k}])$. Hence $[X_*]$ is a minimizer. \square

B.3 Uniqueness of the Minimizer

Proposition B.5. *Under the standing assumptions, the minimizer is unique.*

Proof. Suppose $[X], [X']$ are distinct minimizers. By geodesic convexity, there exists a geodesic γ in \mathcal{A} connecting them. The projection residual term is convex along γ (by linearity of Π_i). The regularizer is strictly convex along γ by assumption. Hence $f(\gamma(t)) < (1-t)f([X]) + tf([X']) = \inf f$ for $t \in (0,1)$, contradicting minimality. \square

B.4 Passage to the Quotient

Proposition B.6. *If $[X] = [X']$ in \mathcal{X}/\sim , then $f([X]) = f([X'])$. Conversely, distinct equivalence classes yield distinct objective values whenever $\tilde{\Pi}$ is injective.*

Proof. Observational equivalence means $\Pi_i(X) = \Pi_i(X')$ for all i , so residual terms are equal; \mathcal{R} is defined on equivalence classes, so regularizer terms are equal. Injectivity of $\tilde{\Pi}$ means distinct classes differ in at least one projection, so residual terms differ. \square

B.5 Full Proof of Theorem 6.11

Full proof. Well-definedness: by the existence and uniqueness propositions, \mathcal{C} is single-valued on \mathcal{A} .

Fixed points = minimizers: $[X^*] = \mathcal{C}([X^*])$ iff $[X^*]$ minimizes f over \mathcal{A} .

Injectivity implies uniqueness: if $\tilde{\Pi}$ is injective on \mathcal{F} , \mathcal{F} contains at most one class. By non-emptiness, $[X^*] \in \mathcal{F}$, so $|\mathcal{F}| = 1$. At $[X^*]$, all residuals are within tolerance; strict convexity prevents alternative minimizers. Hence $[X^*]$ is the unique fixed point.

Non-injectivity implies non-uniqueness: if $\tilde{\Pi}$ is not injective on \mathcal{F} , distinct $[X], [X'] \in \mathcal{F}$ have identical projections. Both are candidates for the minimizer. If $\mathcal{R}([X]) = \mathcal{R}([X'])$, strict convexity fails (contradiction). If $\mathcal{R}([X]) \neq \mathcal{R}([X'])$, the larger-regularizer point is not a minimizer—but then it is not in \mathcal{F} as defined, contradicting the assumption. Either way we reach a contradiction with uniqueness. \square

C Full Proof of the Dynamical Disambiguation Theorem

C.1 Measure-Theoretic Framework

Assumption C.1 (Polish Structure). *\mathcal{A} with the metric induced from \mathcal{B} is a Polish space (complete, separable metric).*

Definition C.2 (Markov Kernel, precise). *$T_\Delta : \mathcal{A} \rightarrow \mathcal{P}(\mathcal{A})$ is a Markov kernel: for each $[X]$, $T_\Delta([X], \cdot)$ is a Borel probability measure on \mathcal{A} ; for each Borel $B \subseteq \mathcal{A}$, $[X] \mapsto T_\Delta([X], B)$ is measurable.*

C.2 Convergence to the Deterministic Generator

Assumption C.3 (Generator Convergence, precise). *For all bounded continuous $\varphi : \mathcal{A} \rightarrow \mathbb{R}$:*

$$\int \varphi(X') T_\Delta([X], dX') = \varphi([X] + \Delta F([X]; \theta)) + O(\Delta^{3/2}).$$

This is weak convergence $T_\Delta([X], \cdot) \Rightarrow \delta_{[X] + \Delta F([X]; \theta)}$ at rate $\Delta^{3/2}$, corresponding to a diffusion with noise variance $O(\Delta)$.

C.3 Stability of Trajectory Separation Under Noise

Definition C.4 (Uniform Separation Margin). *A pair $[X], [X'] \in \Omega_{\text{obs}}$ is uniformly trajectory-separated if there exist $\tau > 0$ and $\delta > 0$ such that*

$$\left\| \tilde{\Pi}(\Phi_F^\tau([X])) - \tilde{\Pi}(\Phi_F^\tau([X'])) \right\| \geq \epsilon + \delta.$$

Lemma C.5. *Under generator convergence, uniformly trajectory-separated pairs remain distinguishable under stochastic evolution with probability $\geq 1 - C\Delta/\delta^2$.*

Proof. By Chebyshev applied to the $O(\Delta)$ noise variance, stochastic trajectories deviate from deterministic by $O(\sqrt{\Delta})$ with high probability. For deviation $\delta/2$: $\mathbb{P}(\|X_\tau^\Delta - X_\tau\| > \delta/2) \leq C\Delta/\delta^2$. With margin δ , the two stochastic trajectories remain separated by $\geq \epsilon$ with probability $\geq 1 - 2C\Delta/\delta^2$. \square

C.4 Full Proof of Theorem 7.10

Full proof. Let $[X], [X'] \in \Omega_{\text{obs}} \cap \Gamma^0$ be distinct and uniformly trajectory-separated at time τ .

Step 1: Both lie on integral curves of $F(\cdot; \hat{\theta})$ by definition of Γ^0 .

Step 2: At time τ , their projections differ by $\geq \epsilon + \delta > \epsilon$.

Step 3: By the noise lemma, stochastic evolutions remain separated with high probability for small Δ .

Step 4: At time τ , the true observation $y(\tau)$ has $\|y_i(\tau) - \Pi_i(X^*(\tau))\| \leq \epsilon_i$. At most one of the two separated candidates lies within ϵ_i of $y_i(\tau)$.

Step 5: The true state $[X^*(\tau)] \in \Omega_{\text{obs}}(y(\tau))$ by non-emptiness applied at time τ . Hence the false candidate is eliminated.

Step 6: Applying to all pairs eliminates all but $[X^*]$. Hence $|\mathcal{F}| = 1$. \square \square

D Fisher Information Geometry of the Feasible Set

D.1 Purpose and Position

Section 9 states that closure corresponds to vanishing information volume of the feasible set. This appendix provides the formal construction underlying that claim. The central point is that the projection family $\{\Pi_i\}$, together with noise models $\{\Sigma_i\}$, induces a Riemannian metric on the admissible quotient space. The metric measures distinguishability of nearby world-states under the available sensor suite. Closure occurs when all remaining admissible directions are eliminated.

D.2 Observation Likelihood on the Quotient Space

For each modality i , assume noise $\eta_i \sim \mathcal{N}(0, \Sigma_i)$ so the conditional density is

$$p_i(y_i | X) \propto \exp\left(-\frac{1}{2}\|y_i - \Pi_i(X)\|_{\Sigma_i^{-1}}^2\right).$$

Assuming conditional independence across modalities, the full observation log-likelihood is

$$-\ell(X; y) = \frac{1}{2} \sum_{i \in \mathcal{I}} \|y_i - \Pi_i(X)\|_{\Sigma_i^{-1}}^2.$$

This is well-defined on $\mathcal{A} = \mathcal{X}/\sim$ because each Π_i is constant on \sim -equivalence classes. The negative log-likelihood is precisely the residual sum in the consistency operator objective.

D.3 The Fisher Information Metric

Definition D.1 (Fisher Information Metric). *The Fisher information metric g_F on \mathcal{A} is the symmetric bilinear form*

$$g_F(X)(u, v) = \sum_{i \in \mathcal{I}} \langle D\Pi_i(X)u, \Sigma_i^{-1}D\Pi_i(X)v \rangle_{\mathcal{Y}_i},$$

where $D\Pi_i(X) : T_X\mathcal{A} \rightarrow \mathcal{Y}_i$ is the derivative of Π_i at $[X]$.

This is a positive semidefinite tensor. It is positive definite exactly when the local model is identifiable: no nonzero tangent direction leaves the induced observation law

unchanged to first order. The null space of g_F is therefore the infinitesimal version of projection degeneracy:

$$g_F(X)(v, v) = 0 \iff D\Pi_i(X)v = 0 \forall i.$$

The negative log-likelihood Hessian equals g_F in the Gaussian case, linking the Fisher metric directly to the curvature of the consistency operator objective.

D.4 Metric Monotonicity Under Sensor Addition

Proposition D.2 (Metric Monotonicity). *Adding a new sensor modality j yields $g_F^{\text{new}}(X) \succeq g_F(X)$ at every $[X] \in \mathcal{A}$.*

Proof. $g_F^{\text{new}}(X)(u, u) = g_F(X)(u, u) + \|D\Pi_j(X)u\|_{\Sigma_j^{-1}}^2 \geq g_F(X)(u, u)$. \square

This is the infinitesimal counterpart of Proposition 5.1: new sensors cannot reduce local distinguishability. They fail to improve it only when their derivative contribution lies entirely in the directions already constrained.

D.5 Restriction to the Feasible Set

Definition D.3 (Restricted Fisher Metric). *The restricted Fisher metric on \mathcal{F} is*

$$g_F^{\mathcal{F}}(X) = g_F(X)|_{T_X\mathcal{F} \times T_X\mathcal{F}}.$$

When \mathcal{F} is locally a smooth embedded submanifold of \mathcal{A} (which holds on regular strata under generic regularity assumptions), $g_F^{\mathcal{F}}$ defines a Riemannian metric on those strata.

Definition D.4 (Information Volume).

$$\text{Vol}_F(\mathcal{F}) = \int_{\mathcal{F}_{\text{reg}}} d\text{Vol}_{g_F^{\mathcal{F}}}.$$

Large information volume indicates many locally distinguishable admissible world-states. Zero volume indicates closure, subject to the nondegeneracy condition below.

D.6 Dynamical Lifting of Null Directions

Define the time-extended projection $\Pi_i^{(\tau)}(X) = \Pi_i(\Phi_F^\tau(X))$ and its induced metric

$$g_F^{(\tau)}(X)(u, v) = \sum_i \langle D(\Pi_i \circ \Phi_F^\tau)(X)u, \Sigma_i^{-1} D(\Pi_i \circ \Phi_F^\tau)(X)v \rangle.$$

By the chain rule, $D(\Pi_i \circ \Phi_F^\tau)(X) = D\Pi_i(\Phi_F^\tau(X)) \circ D\Phi_F^\tau(X)$.

Proposition D.5 (Metric Form of Dynamical Disambiguation). *Suppose $g_F(X)(v, v) = 0$ but $g_F^{(\tau)}(X)(v, v) > 0$ for some $\tau > 0$. Then v is a static null direction that is lifted into the observable range by the dynamics Φ_F^τ . This is the differential-geometric form of Theorem 7.10: two configurations indistinguishable at $t = 0$ become distinguishable at $t = \tau$.*

Time is therefore not an auxiliary feature—it is the mechanism by which null directions of the static Fisher metric are eliminated.

D.7 Closure as Vanishing Information Volume

Theorem D.6 (Closure Implies Vanishing Information Volume). *If $\mathcal{F} = \{[X^*]\}$, then $\text{Vol}_F(\mathcal{F}) = 0$.*

Proof. A singleton has zero-dimensional regular part; the Riemannian volume form integrates to zero. \square \square

Proposition D.7. *$\text{Vol}_F(\mathcal{F}) = 0$ with $g_F^{\mathcal{F}}$ nondegenerate on the regular part of \mathcal{F} implies $\dim_{\text{eff}}(\mathcal{F}) = 0$.*

Proof. Under nondegeneracy, positive dimension implies positive volume. \square \square

Corollary D.8 (Closure Criterion via Information Geometry). *If $g_F^{\mathcal{F}}$ is nondegenerate on the regular part of \mathcal{F} , then*

$$\mathcal{F} = \{[X^*]\} \iff \dim_{\text{eff}}(\mathcal{F}) = 0 \iff \text{Vol}_F(\mathcal{F}) = 0.$$

The nondegeneracy assumption excludes the case where \mathcal{F} has positive dimension but the metric degenerates on it (unresolved indistinguishability rather than genuine uniqueness). This is the exact condition Section 9 relies on.

D.8 Reconstruction Efficiency

Let \mathcal{F}_n denote the feasible set after n sensor modalities or constraint updates. Define:

Definition D.9 (Reconstruction Efficiency).

$$\mathcal{E}_n = 1 - \frac{\text{Vol}_F(\mathcal{F}_{n+1})}{\text{Vol}_F(\mathcal{F}_n)} \in [0, 1].$$

Values near 0 indicate the new sensor or update contributes little additional discrimination. Values near 1 indicate near-total collapse of residual admissible uncertainty. A continuous-time analogue is $\mathcal{E}(t) = -\frac{d}{dt} \log \text{Vol}_F(\mathcal{F}(t))$, which behaves as an entropy-contraction rate and is often more stable than raw volume differences.

The optimal sensor placement criterion follows: the next sensor j should be placed where the directional derivative of \mathcal{E} is maximized, which is determined by the eigenspectrum of g_F in the remaining unclosed directions.

D.9 The Effective Metric Near Closure

Near a nondegenerate minimizer $[X^*]$, the Hessian of the consistency objective $J(X) = \sum_i \|y_i - \Pi_i(X)\|_{\Sigma_i^{-1}}^2 + \mathcal{R}(X)$ is

$$\nabla^2 J(X^*) = g_F(X^*) + \nabla^2 \mathcal{R}(X^*).$$

This effective metric combines statistical distinguishability (Fisher geometry) with admissibility curvature (regularizer). The convergence rate of the refinement loop of Section 8 is governed by the condition number $\kappa = \lambda_{\max}/\lambda_{\min}$ of $\nabla^2 J(X^*)$. Directions with small Fisher information and small regularizer curvature correspond to slow convergence; dense multi-modal sensor networks increase λ_{\min} and hence accelerate convergence to closure.

D.10 Entropy Identification

The identification stated in Section 2,

$$\log \text{Vol}_F(\mathcal{F}) \longleftrightarrow \log \dim H^0(\mathcal{U}, \mathcal{S}_\Pi) \longleftrightarrow S_{\text{RSVP}}(X^*),$$

is now fully grounded. The first identification holds in the Gaussian regime because $\dim H^0$ counts linearly independent global sections, which scales with the information volume of the feasible set. The second uses the RSVP entropy field equation: at the fixed point $[X^*]$, S_{RSVP} is the integral of local entropy production over Ω , which under the consistency operator objective equals the log-volume of remaining uncertainty in the observation-consistent set. The RSVP entropy field is the local density of cohomological obstruction; its minimization under \mathcal{C} is the physical expression of sheaf repair.

E Operator Class Library

E.1 Structure of the Library

The library consists of K named process generators $\{F_k\}_{k=1}^K$, each a learned vector field on \mathcal{A} parametrized by $\theta_k \in \Theta_k$. The full dynamics decompose as

$$F([X]; \theta) = \sum_{k=1}^K \alpha_k([X]) F_k([X]; \theta_k),$$

with activation weights $\alpha_k([X]) \geq 0$ satisfying $\sum_k \alpha_k = 1$. Each F_k is learned from before/after pairs in which process k is dominant, identified by its characteristic sensor signature.

E.2 Ecological Process Operators

Forest succession (F_{succ}): slow increase in LiDAR canopy height, gradual RGB reflectance shift from grass-like to canopy-like, decrease in audio spectral entropy as bird diversity changes with succession stage. Time scale: years to decades.

Disturbance and recovery (F_{dist}): rapid decrease in LiDAR height and RGB NDVI followed by slow recovery; audio signature is absence of characteristic species. Time scale: event-driven (days), recovery (years).

Hydrological dynamics (F_{hydro}): correlated changes in infrared (soil moisture affects surface temperature), LiDAR (inundation changes geometry), and audio (water flow sounds). Time scale: seasonal to event-driven.

Phenological cycling (F_{pheno}): regular seasonal oscillations in RGB (green-up/senescence), infrared (latent heat flux), and audio (bird song seasonality). Time scale: annual.

E.3 Civilizational Process Operators

Urban expansion (F_{urban}): progressive increase in LiDAR building density and height, RGB shift to impervious surface reflectance, audio increase (traffic, machinery). Time scale: years to decades.

Infrastructure decay (F_{decay}): gradual LiDAR surface roughness change, RGB spectral shift, audio entropy increase (structural noise). Time scale: years.

Agricultural cycling (F_{agri}): regular RGB oscillations (crop growth stages), infrared ET changes, LiDAR canopy height variation, audio machinery activity. Time scale: seasonal.

Demographic change (F_{demo}): correlated changes in audio activity levels, RGB night-light proxies, and infrastructure use patterns. Time scale: years.

E.4 Learning Protocol

Each F_k is learned from before/after pairs as follows. First, identify pairs where process k is dominant by classifying the observation delta against the existing library. Second, fit θ_k by minimizing the forward-projection discrepancy:

$$\hat{\theta}_k = \arg \min_{\theta_k} \sum_{n \in \mathcal{D}_k} \sum_i \left\| \Pi_i(\Phi_{F_k}^\tau([X_b^{(n)}]; \theta_k)) - y_a^{(n)} \right\|^2.$$

Third, validate on held-out pairs by checking that $\Theta_{(y_b, y_a, \tau)} \cap \Theta_k \neq \emptyset$ for pairs labeled as process k .

E.5 Library Extensibility

The library is open: new process types F_{K+1} are added by collecting before/after pairs exhibiting the new process, fitting the generator as above, and appending to the library. The identifiability and disambiguation theorems are unaffected by library extension: they hold for any F with Lipschitz generator. The effect of extension is to reduce Γ_{indist}^0 — the set of configurations not separated by trajectory learning — which by Corollary ?? increases the rate of closure.

F Implementation File Structure

This appendix documents the reference implementation of the Yarncrawler consistency operator as a stratified filesystem. The architecture maps directly onto the theoretical structure: state definition corresponds to the configuration space \mathcal{X} , morphism algebra corresponds to the operator class library, controller evolution corresponds to iterative application of \mathcal{C} , and quotient construction corresponds to the identification of observationally equivalent world-states.

F.1 Structural Overview

The implementation is organized into seven layers, each corresponding to a structural role in the framework.

Layer	Theoretical Role
State layer	Configuration space \mathcal{X} , field types, domains
Morphism layer	Operator class library $\{F_k\}$, composition algebra
Transfer layer	Domain transfer, functoriality checking
Controller layer	Iterative application of \mathcal{C} , decision policy
Transformation layer	Natural transformations between controllers
Fixed-point layer	Closure detection, stability tracking
Quotient layer	Equivalence classes of world-states

F.2 Core Infrastructure

Shared utilities used by all layers:

```
common/  
  common_io.py          # I/O primitives, serialization, logging
```

F.3 State Layer

Defines the plenum field configuration and admissible types.

```
state/  
  initial_state.json    # Initial field configuration (Phi, v, S)  
  topology.json        # Domain geometry, boundary conditions  
  typing_rules.yaml     # Type constraints on field components  
  
  domain_A.json        # Source domain specification  
  domain_B.json        # Target domain specification  
  
  transfer_A_to_B.json # Cross-domain transfer morphism
```

F.4 Log Layer

Accumulates the trajectory history that the disambiguation theorem requires.

```
logs/  
  history.jsonl        # Full state history (before/after pairs)  
  controller_trace.jsonl # Controller decision trace  
  class_transforms.jsonl # Equivalence class transitions
```

F.5 Morphism Layer

Implements the operator class library: storage, generation, composition, transfer.

```
morphisms/
  registry/
    index.json          # Global operator index
    entries/
      morph_*.json     # Individual operator specifications

  transferred/
    *.json             # Operators transferred across domains

python/
  morphism_apply.py    # Apply operator to field configuration
  compose_morphisms.py # Categorical composition of operators
  discover_compositions.py # Find composable operator pairs
  abstract_morphism.py # Abstract operator from trajectory data

  admit_morphism.py    # Admit new operator to registry
  update_registry_stats.py # Track operator usage statistics
  prune_registry.py    # Remove underused operators
  query_registry.py    # Retrieve operators by type/signature

  transfer_morphism.py # Transfer operator across domains
  transfer_registry_slice.py # Transfer operator subset
  check_functoriality.py # Verify naturality conditions

  check_composable.py # Type-check operator composition
  type_path_resolver.py # Resolve type paths in composition
  propose_type_bridge.py # Propose bridging morphisms

  score_rollout.py     # Score operator on trajectory rollout
```

F.6 Controller Layer

Implements the iterative consistency operator \mathcal{C} and decision policies.

```
controllers/
  registry/
    beam.json          # Beam search controller
    dp.json            # Dynamic programming controller
    learned.json       # Learned neural controller
    spatial.json       # Spatially-indexed controller
    repair.json        # Sheaf repair controller
```

```

state/
  controller_state.json # Current controller parameters

transformations/
  beam_to_learned.json # Natural transformation between controllers

quotient/
  classes.json          # Equivalence classes of controllers
  class_0001.json      # Class representative and members
  class_0002.json
  representatives.json  # Selected representatives per class

python/
  apply_controller_transform.py # Apply natural transformation
  learn_controller_transform.py # Learn transformation from data

  iterate_controller.py      # One step of consistency operator
  controller_distance.py     # Distance between controller states
  check_fixed_point.py      # Test for closure (fixed point)
  update_stability.py       # Track convergence to closure

  controller_equivalence.py # Identify equivalent controllers
  build_controller_classes.py # Construct equivalence classes
  select_representatives.py # Choose class representatives

  build_controller_quotient.py # Construct quotient structure
  split_classes.py            # Refine equivalence classes
  select_controller_class.py  # Choose active class
  log_class_transform.py     # Record class transitions
  canonicalize_controller.py  # Map to canonical representative

  meta_controller.py        # Orchestrate controller selection

```

F.7 Engine Layer

Executable entry points that orchestrate the full refinement loop.

```

engine/
  apply_morphism.sh          # Apply operator to current state
  run_transfer.sh           # Execute domain transfer
  run_controller_fixed_point.sh # Run until closure detected
  run_class_controller.sh    # Run with class-level controller

```

run.sh

Top-level execution entry point

F.8 Correspondence to the Theoretical Framework

The implementation instantiates the theoretical structure as follows. The `state/` layer defines \mathcal{X} and \mathcal{A} . The `morphisms/` layer is the operator class library $\{F_k\}$ of Appendix E. The `logs/history.jsonl` file accumulates the before/after trajectory pairs used to learn $F(\cdot; \hat{\theta})$. The `iterate_controller.py` script implements one step of \mathcal{C} . The `check_fixed_point.py` script implements the closure detection criterion of Section 9: stability of $[X^*]$ under perturbation of ϵ_i . The `controllers/quotient/` directory implements the identification of observationally equivalent world-states, corresponding to the passage to \mathcal{X}/\sim . The `check_functoriality.py` script verifies that domain transfers respect the naturality conditions of the functor Ψ established in Theorem 2.6.

G The Constellation Construction as Iterated Constraint Collapse

G.1 The Initial Object

Let \mathcal{M}_0 denote the space of all syntactically valid modules:

$$\mathcal{M}_0 = \{\text{module.yaml satisfying schema}\}$$

At this stage no global compatibility is enforced. The space is combinatorially large and admits arbitrary configurations. Each subsequent stage reduces admissible degrees of freedom by introducing new compatibility conditions. The resulting object is a tower of constraint refinements converging to a unique global section.

G.2 Stage-wise Refinement as Endofunctors

Each stage k defines a restriction operator:

$$\mathcal{R}_k : \mathcal{M}_{k-1} \rightarrow \mathcal{M}_k, \quad \mathcal{M}_k = \{X \in \mathcal{M}_{k-1} \mid X \text{ satisfies constraints of stage } k\}$$

The full system is the intersection:

$$\mathcal{M}_* = \bigcap_{k=1}^{13} \mathcal{M}_k$$

The 13 stages are not a pipeline. They are a descending filtration:

$$\mathcal{M}_0 \supset \mathcal{M}_1 \supset \cdots \supset \mathcal{M}_{13}$$

and the development history documented in Appendix F through Appendix J is the construction of this filtration one level at a time.

G.3 Sheaf Enforcement (Stages 1–4)

At Stage 1, the module space is lifted into a sheaf:

$$\mathcal{S} : \mathcal{U}^{\text{op}} \rightarrow \mathbf{Set}$$

where modules are local sections and the constellation’s global section corresponds to the deployed system. The cover \mathcal{U} is the stratum partition: each stratum (Sensory, Intellect, Understanding, Decision, Reason, Will, Presentation, Thought) defines one chart.

Stages 2–4 enforce the sheaf condition by eliminating non-gluing configurations:

Stage 2 detects cocycles: `check_commutativity.sh` identifies pairs of modules whose interface contracts are incompatible, emitting structured obstruction objects corresponding to $c_{ij} \in H^1(\mathcal{U}, \mathcal{S})$.

Stage 3 expands the restriction maps by introducing a type DAG over field spaces ($H^1(\Omega) \hookrightarrow L^2(\Omega) \hookrightarrow \text{Scalar}$), replacing one-step compatibility checks with multi-hop path resolution.

Stage 4 constructs coboundaries: `propose_bridge.sh` and `apply_repair.sh` insert bridge modules that trivialize detected cocycles.

After Stage 4, the system has driven $H^1(\mathcal{U}, \mathcal{S}) \rightarrow 0$ or has produced a structured report identifying which obstructions are irreducible.

G.4 Emergence of the Morphism Algebra (Stages 5–7)

Stages 5–7 lift the system from object space to morphism space. Define:

$$\mathcal{O} = \{f : X \rightarrow X \mid f(\Phi, \mathbf{v}, S) = (a_\Phi \Phi + b_\Phi, a_v \mathbf{v} + \mathbf{b}_v, a_S S + b_S)\}$$

the space of affine RSVP operators. Stage 5 introduces the morphism kernel and normalizes all operators into this common form. Stage 6 introduces the candidate generation loop: mutation, composition, evaluation, and the persistent registry. Stage 7 moves the vector field \mathbf{v} from node-based coordinates to edge-based flux, making divergence a proper discrete conservation law and enabling geometry-respecting operator semantics.

These stages impose closure under composition, compatibility with RSVP field structure, and conservation laws on graph topology. The system transitions from a static constraint set to a dynamical algebra: \mathcal{M}_7 is not just a set of valid configurations but a set of valid configuration-evolution pairs.

G.5 Variational Selection and Trajectory Space (Stages 8–10)

Stages 8–10 introduce trajectory dependence. Define the rollout functional:

$$\mathcal{J}(f; X_0) = \sum_{t=0}^T \mathcal{L}(f_t(X_t))$$

where $X_{t+1} = F(f_t(X_t), X_t)$ is the coupled field dynamics and \mathcal{L} integrates sensor residuals, divergence penalties, boundary leakage, and semantic property verification. The admissible

operator set becomes:

$$\mathcal{O}_{\text{adm}} = \{f \mid \mathcal{J}(f) \text{ minimized under constraints}\}$$

Stage 8 implements this functional as a gated rollout scorer. Stage 9 spatializes the controller: instead of a single operator applying to the whole field, a weight field $w : \Omega \rightarrow \Delta(\mathcal{O})$ assigns a distribution over operators to each node, smoothed over the topology. Stage 10 introduces local morphism generation conditioned on regional obstruction signatures: each region of the field can generate operators targeted at its specific failure mode, and successful local repairs enter the global registry.

The effective dimension of \mathcal{M}_{10} is significantly lower than \mathcal{M}_7 : not only must configurations satisfy static compatibility conditions, they must also lie on admissible trajectories under the selected operators.

G.6 Typed Category and Compositional Structure (Stage 11)

Stage 11 introduces higher-order structure. The operator space becomes a typed category:

$$\mathcal{O} : \mathcal{T} \rightarrow \mathcal{T}$$

where objects are field types and morphisms are RSVP operators. Composition is constrained by:

$$f \circ g \text{ defined iff } \text{cod}(g) \rightsquigarrow \text{dom}(f)$$

via a path in the type DAG. The type path resolver finds chains of admissible embeddings, and the composition engine (`python/compose_morphisms.py`) verifies compatibility before any chain is formed.

Composition discovery (`python/discover_compositions.py`) identifies empirically successful operator chains from the history log, and abstraction (`python/abstract_morphism.py`) compresses these into single operators, growing the registry with reusable primitives. At this stage the operator space is no longer flat: it has algebraic structure, and \mathcal{M}_{11} is the intersection of configuration space with this structured operator algebra.

G.7 Functorial Transfer (Stage 12)

Stage 12 introduces inter-domain structure. Let $\mathcal{D}_1, \mathcal{D}_2$ be domain categories defined by their topology, sensor bundle, and field type system. A transfer map defines a functor:

$$\mathcal{F} : \mathcal{D}_1 \rightarrow \mathcal{D}_2, \quad \mathcal{F}(f \circ g) \approx \mathcal{F}(f) \circ \mathcal{F}(g)$$

Deviation from functoriality defines a transfer error, scored by `python/check_functoriality.py`. Transferred operators enter the target registry with a score discount proportional to that error.

This stage does not reduce the effective dimension of \mathcal{M}_{12} within a single domain. Its role is different: it introduces a relation of equivalence between admissible configurations

across domains, so that the system’s learned structure can be reused rather than rebuilt from scratch in each new environment.

G.8 Controller Quotient (Stage 13)

Stage 13 lifts the system to the highest level of abstraction. Controllers are functors:

$$\mathcal{K} : \mathcal{X} \rightarrow \mathcal{O}$$

mapping field states to operator sequences. Natural transformations $\eta : \mathcal{K}_1 \Rightarrow \mathcal{K}_2$ define equivalences between controllers. The quotient \mathcal{K}/\sim collapses redundant controllers into equivalence classes, and the system operates on $[\mathcal{K}]$ rather than individual controllers.

Fixed-point detection (`python/check_fixed_point.py`) identifies when a controller transformation leaves the controller unchanged within tolerance. Stability accumulation tracks how consistently each class achieves fixed-point status. The canonical representative of each class is the controller with highest stability score.

G.9 The Fixed Point of the Total Operator

The full system defines a composite restriction operator:

$$\mathcal{C}_{\text{total}} = \mathcal{R}_{13} \circ \cdots \circ \mathcal{R}_1$$

Closure is achieved when:

$$\mathcal{C}_{\text{total}}(X^*) = X^*$$

At this point all cocycles in $H^1(\mathcal{U}, \mathcal{S})$ are trivial, all type mismatches are resolved by admissible bridge morphisms, all operator gaps are filled by generated or retrieved morphisms, and all controller redundancies are collapsed in the quotient. In terms of Section 5:

$$\dim_{\text{eff}}(\mathcal{M}_{13}) = 0$$

The system is not a builder, not a learner, not a simulator. It is a constraint-saturating operator whose sole function is to eliminate all inconsistent configurations.

G.10 Correspondence to the Main Theorems

The filtration $\mathcal{M}_0 \supset \cdots \supset \mathcal{M}_{13}$ is the implementation-level instantiation of the feasible set contraction established in Sections 6 and 7. The main paper proves that the consistency operator \mathcal{C} has a unique fixed point when the projection family is injective on \mathcal{F} and dynamics separate residual ambiguity. The constellation construction proves the same result one filtration level at a time:

Stages 1–4 correspond to the static identifiability condition (Ω_{obs} contraction by sensor coverage). Stages 5–10 correspond to the dynamical disambiguation condition (Γ refinement by trajectory constraints). Stages 11–12 correspond to the type-theoretic and transfer conditions that ensure the operator algebra is internally consistent and portable. Stage

13 corresponds to the controller-level analogue of closure in \mathcal{X}/\sim : the decision procedure itself is fixed, not just the field state.

The fixed point $\mathcal{C}_{\text{total}}(X^*) = X^*$ is therefore the operational realization of the Yarn-crawler Identifiability Theorem in a system of hundreds of interconnected repositories. The constellation achieves closure when no alternative module configuration, operator sequence, or controller policy remains consistent with all imposed constraints simultaneously. Stages 14–16, which follow, extend the filtration from logical closure to viability closure, from internal closure to external legibility, and from task closure to mission closure.

G.11 Stage 14: Homeostatic Constraint (Internal Sensory Monitoring)

De Gyurky’s Sensory System is responsible not only for external environmental awareness but for constant monitoring of the system’s internal health. In the filtration, this stage introduces a viability requirement: the system may only occupy configurations that do not threaten its own capacity for continued operation.

Let $\mathcal{S}_{\text{infra}}$ denote the state space of the constellation’s infrastructure — memory consumption, build-time latency, dependency graph depth, repository count, and inter-service coupling density. Define a health functional

$$\mathcal{H} : \mathcal{M}_{13} \rightarrow \mathbb{R}$$

measuring departure from operational viability. The filtration step is:

$$\mathcal{M}_{14} = \{X \in \mathcal{M}_{13} \mid \mathcal{H}(X) \leq \mathcal{H}_{\text{max}}\}$$

where \mathcal{H}_{max} is the viability threshold. This constraint prunes from the feasible set any configuration that satisfies logical and type consistency but exceeds resource bounds: a proposed repair that creates a dependency cycle exceeding memory limits, or a composition chain that degrades build latency below operational thresholds, is rejected even if it would satisfy all earlier filtration conditions.

The effect of Stage 14 is to make the Will stratum subordinate to the Sensory stratum. Deployment is not only gated on logical closure; it is gated on metabolic stability. A constellation that achieves $H^1 = 0$ but collapses under its own resource load has not achieved closure in the full sense. \mathcal{M}_{14} is the submanifold of \mathcal{M}_{13} on which the system can sustain itself.

In the language of the paper’s main framework, Stage 14 adds an additional component to the regularizer:

$$\mathcal{R}_{14}(X) = \mathcal{R}(X) + \lambda_{\text{health}} \mathcal{H}(X)$$

so that configurations near the viability boundary are penalized even when logically admissible. The admissible set \mathcal{A} is refined to $\mathcal{A}_{14} = \mathcal{A} \cap \mathcal{M}_{14}$.

G.12 Stage 15: Presentational Projection (Observer-Accessible Interface)

De Gyurky’s Presentation System governs the transformation of internal experience into projectable form: viewing, rendering, and making legible the system’s internal state to external observers. In the filtration, this stage introduces the requirement that the system’s closed internal state must admit a coherent external representation.

Define the interface space \mathcal{I} as the space of observer-accessible artifacts: documentation, visualizations, API specifications, structured logs, and public repository metadata. Define a projection operator:

$$\mathcal{P} : \mathcal{A}_{14} \rightarrow \mathcal{I}$$

The filtration step is:

$$\mathcal{M}_{15} = \{X \in \mathcal{M}_{14} \mid \mathcal{P}(X) \approx y_{\text{obs}}\}$$

where y_{obs} is the target observation — the mission requirement expressed in the interface space. This is the presentational analogue of the sensor projection condition $\|\Pi_i(X) - y_i\| \leq \epsilon_i$: the system must not only be internally consistent but must project into the interface space in a way that satisfies the requirements of its external observers.

This stage formalizes the distinction between *internal completion* (the collapsed manifold \mathcal{M}_{14}) and *external completion* (the legible projection $\mathcal{P}(X^*)$). A constellation that has achieved internal closure but cannot explain its own state — cannot produce documentation from its registry, cannot render its dependency graph, cannot expose its sensor projections as human-readable artifacts — has not achieved full closure. \mathcal{M}_{15} is the submanifold of configurations that are both internally solved and externally legible.

The Presentation System in the implementation corresponds to the `docs/generators/` layer: scripts that generate FRD, FDD, SISD, and STIP documents automatically from the registry, README stubs from module metadata, and interface specifications from type signatures. These are not decorative outputs; they are constraints in \mathcal{M}_{15} that must be satisfied before Will executes.

G.13 Stage 16: The Thought System and Terminal Closure

The Thought System — the *Nexus Cogitationis* — is the hub of all communication and the locus of self-awareness in de Gyurky’s architecture. It is responsible for detecting not merely that a specific task is closed, but that the entire constellation has reached a state of structural closure across all strata simultaneously. In the filtration, this is the terminal stage.

Define an equivalence relation \sim_{mission} on \mathcal{M}_{15} : two configurations are mission-equivalent if they produce the same strategic outcome under the constellation’s governing objectives. Formally:

$$X \sim_{\text{mission}} X' \iff \mathcal{P}(X) = \mathcal{P}(X') \text{ and } \mathcal{H}(X) \approx \mathcal{H}(X') \text{ and } [\mathcal{K}(X)] = [\mathcal{K}(X')]$$

where $[\mathcal{K}]$ denotes the controller equivalence class from Stage 13. The terminal object is

the quotient:

$$\mathcal{M}_{16} = \mathcal{M}_{15} / \sim_{\text{mission}}$$

At this stage the system achieves what the main paper calls closure and what de Gyurky calls genuine autonomy: the generator and the trajectory are indistinguishable, the internal constraint structure and the external mission requirement are unified, and the system operates at the fixed point of its own constraints rather than searching for it. The Thought System acknowledges the global section’s existence across all sixteen strata and stabilizes the manifold.

The full filtration is:

$$\mathcal{M}_0 \supset \mathcal{M}_1 \supset \cdots \supset \mathcal{M}_{13} \supset \mathcal{M}_{14} \supset \mathcal{M}_{15} \supset \mathcal{M}_{16}$$

and the composite restriction operator is:

$$\mathcal{C}_{\text{total}} = \mathcal{R}_{16} \circ \cdots \circ \mathcal{R}_1$$

Closure is:

$$\mathcal{C}_{\text{total}}(X^*) = X^*, \quad \dim_{\text{eff}}(\mathcal{M}_{16}) = 0$$

At the fixed point: all cohomological obstructions are trivialized ($H^1 = 0$), all type mismatches are resolved, all operator gaps are filled, all controller redundancies are collapsed, the system is metabolically viable ($\mathcal{H}(X^*) \leq \mathcal{H}_{\text{max}}$), the internal state is externally legible ($\mathcal{P}(X^*) \approx y_{\text{obs}}$), and the mission equivalence class is a singleton ($|\mathcal{M}_{16}| = 1$).

This is the mathematical definition of a done autonomous state. The constellation is not working toward a goal; it is the realized solution, operating at the fixed point of its own internal and external constraints across all sixteen strata of the filtration.

H The Constellation Architecture as Operational Sheaf

H.1 From Deployment Tool to Sheaf Realization

The implementation described in Appendix F is not a deployment pipeline that happens to use sheaf language. It is a direct computational realization of the sheaf \mathcal{S}_{Π} over the quotient configuration space \mathcal{X}/\sim . The correspondence is exact, not analogical.

Each software module is a local section $s_i \in \mathcal{S}_{\Pi}(U_i)$: it has a declared type signature (the `inputs` and `outputs` fields of `module.yaml`), a declared stratum (one of the eight de Gyurky functional roles), and an explicit dependency list (the `depends_on` field). The dependency list defines the restriction maps ρ_{ij} : module j depends on module i means i ’s output must be compatible with j ’s input, i.e., $\rho_{ij}(s_i) = s_j|_{U_i \cap U_j}$.

The `check_commutativity.sh` script detects violations of the gluing condition. When module j requires field Φ in space $H^1(\Omega)$ but module i provides Φ in space $L^2(\Omega)$, the script emits a structured object of the form

{

```

"type": "cohomology_obstruction",
"from": "module_i", "to": "module_j",
"field": "Phi"
}

```

This is not an error message in the conventional sense. It is a representation of a non-trivial Čech 1-cocycle $c_{ij} \in H^1(\mathcal{U}, \mathcal{S}_\Pi)$. The system has computed that the local sections s_i and s_j do not agree on the overlap $U_i \cap U_j$.

Successful deployment — the condition under which `run.sh` completes without error and all repositories are materialized via `gh repo create` — corresponds exactly to the existence of a unique global section: $H^1(\mathcal{U}, \mathcal{S}_\Pi) = 0$ and $\dim H^0 = 1$. GitHub becomes the observable projection of a globally consistent semantic object, not the ground truth.

H.2 The Eight Strata as Factorization of \mathcal{C}

The eight de Gyurky systems — Sensory, Intellect, Understanding, Decision, Reason, Will, Presentation, Thought — are not implemented as separate components. They are implemented as functional strata through which every module is classified, and together they constitute a factorization of the consistency operator \mathcal{C} .

The *Sensory* stratum defines the projection operators $\{\Pi_i\}$: it comprises the modules that produce observations from field configurations. Every module in this stratum has typed outputs corresponding to sensor channels (infrared, LiDAR, RGB, audio, photonic strip) and its `module.yaml` declares the field space from which it projects.

The *Intellect* stratum corresponds to the encoder q_ϕ : it comprises the modules that generate candidate field configurations from observations, operating in the high-entropy regime without enforcing admissibility. This is the proposal mechanism of Section 8.

The *Understanding* stratum normalizes and packages candidate configurations for the consistency loop. In sheaf terms, it constructs the overlap data: given two locally consistent configurations from different sensors, it prepares the restriction maps for comparison.

The *Reason* stratum is the admissibility filter \mathcal{R} : modules in this stratum enforce that field configurations satisfy the RSVP equations, ecological plausibility constraints, and civilizational feasibility. A configuration that fails Reason is not merely penalized — it is excluded from \mathcal{A} entirely.

The *Decision* stratum implements operator selection. At each iteration of the refinement loop, Decision modules select among candidate morphisms by evaluating them against the current state and observations. This is the variational minimization in Definition 5.3, made concrete as a selection procedure over a ranked candidate pool.

The *Will* stratum executes the selected morphism: it applies the chosen operator to the current field configuration and, when closure is certified, materializes the result into the external world (repository creation, push, tagging). Critically, Will modules are never invoked until the closure check passes. Deployment is not a routine step in the pipeline; it is the consequence of $|\mathcal{F}| = 1$.

The *Presentation* stratum projects the internal field state back into observable form: documentation generation, README synthesis, visual output. It corresponds to the

forward model in the refinement loop (Section 8), applied not for validation but for communication.

The *Thought* stratum is the consistency operator itself: the orchestration loop in `run.sh` that iterates over all strata in the correct order, checks closure, triggers repair when it fails, and halts when the global section is achieved. Thought does not have its own modules; it is the loop that runs them.

The complete factorization is:

$$\mathcal{C} = \text{Will} \circ \text{Decision} \circ \text{Understanding} \circ \text{Intellect} \circ \text{Sensory} \circ \text{Reason}$$

with Presentation as an output projection and Thought as the iteration of this composition until fixed-point.

H.3 The Repair Loop as Coboundary Construction

The three-script repair loop — `propose_bridge.sh`, `evaluate_bridge.sh`, `apply_repair.sh` — implements the sheaf repair operator $\mathcal{R}_{\text{sheaf}}$ of Section 2 in executable form.

When `check_commutativity.sh` detects a cohomological obstruction between modules i and j over field Φ in incompatible spaces, `propose_bridge.sh` reads the structured obstruction object and queries the type graph defined in `schema/typing_rules.yaml`. The type graph is a directed acyclic graph over field spaces: $H^1(\Omega) \hookrightarrow L^2(\Omega) \hookrightarrow \text{Scalar}$. The path resolver (`engine/type_path_resolver.sh`) performs a breadth-first search over this graph to find an admissible sequence of embeddings from the provided space to the required space.

Each step of the resolved path generates one bridge module. If the path is $H^1(\Omega) \rightarrow L^2(\Omega) \rightarrow \text{Scalar}$, then on the first repair pass the system generates and inserts:

$$\text{phi_bridge_H1_to_L2} : \Phi_{H^1} \mapsto \Phi_{L^2}$$

On the second pass, module j still lacks the Scalar-type Φ it requires, but now a provider exists in $L^2(\Omega)$, and the system generates:

$$\text{phi_bridge_L2_to_Scalar} : \Phi_{L^2} \mapsto \Phi_{\text{Scalar}}$$

This incremental coboundary construction is precisely the multi-hop path completion established in Section 2: each bridge module trivializes one edge of the Čech cocycle, and the loop continues until $H^1 = 0$.

The loop is bounded by `MAX_ITERS` in `run.sh`. If closure is not achieved within that bound, the system emits a structured failure report identifying which cocycles remain non-trivial and whether the failure is projection degeneracy (the type graph contains no path between the required spaces) or regularizer flatness (multiple valid bridge paths exist and the system cannot select among them without model revision).

I The Morphism Kernel as RSVP Operator Algebra

I.1 Operators as RSVP Field Transformations

The morphism kernel implements the operator class library of Appendix E as an executable algebra over the RSVP field triple. Every morphism in the system — whether symbolic, learned, hybrid, composed, or generated — is normalized into a common affine RSVP form:

$$(\Phi, \mathbf{v}, S) \mapsto (a_\Phi \Phi + b_\Phi, a_v \mathbf{v} + \mathbf{b}_v, a_S S + b_S)$$

where the six parameters $(a_\Phi, b_\Phi, a_v, \mathbf{b}_v, a_S, b_S)$ constitute the operator’s normal form. This normalization, implemented in `python/morphism_normalize.py`, is what makes composition exact rather than approximate: the normal form of a composition $M_2 \circ M_1$ is computed analytically by substitution, so no numerical error accumulates through chains of operators.

The normalization also makes blending in operator space well-defined. A context-weighted blend of operators $\{M_k\}$ with weights $\{w_k\}$ is the operator with parameters:

$$a_\Phi^{\text{blend}} = \sum_k w_k a_\Phi^{(k)}, \quad b_\Phi^{\text{blend}} = \sum_k w_k b_\Phi^{(k)}, \quad \dots$$

This is implemented in `python/blend_morphisms.py` and used both for context-neighbor retrieval (Section I.3 below) and for multi-controller blending (Appendix J).

The RSVP field triple enters the implementation not as three independent channels but as a coupled physical object. The vector field \mathbf{v} is implemented as edge-based flux on an explicit graph topology (`state/topology.json`): each edge carries a scalar flux value, and divergence at node i is computed as the exact discrete conservation law $\text{div}(i) = \sum_{e \rightarrow i} f_e - \sum_{i \rightarrow e} f_e$. The entropy field S is updated by the discrete continuity equation accounting for local divergence and shear stress along edges. This means the three fields are genuinely coupled in the forward dynamics, not merely co-present.

I.2 Sensor-Coupled Operator Selection

Morphism selection is not a combinatorial search over operator space. It is a variational minimization over a candidate pool, evaluated against sensor observations and temporal consistency requirements. The scoring functional is:

$$\mathcal{J}(M) = w_1 \mathcal{L}_{\text{now}}(M) + w_2 \mathcal{L}_{\text{future}}(M) + w_3 \mathcal{P}_{\text{div}}(M) + w_4 \mathcal{P}_{\text{boundary}}(M) + w_5 \mathcal{P}_{\text{rough}}(M) - \mathcal{B}_{\text{semantic}}(M) + \mathcal{P}_{\text{semantic}}(M)$$

where \mathcal{L}_{now} is the projection residual at the current step, $\mathcal{L}_{\text{future}}$ is the residual after one forward simulation step, \mathcal{P}_{div} penalizes unphysical divergence, $\mathcal{P}_{\text{boundary}}$ penalizes flux leakage through boundary nodes, $\mathcal{P}_{\text{rough}}$ penalizes rapid spatial variation in edge flux, and $\mathcal{B}_{\text{semantic}}$ and $\mathcal{P}_{\text{semantic}}$ are bonuses and penalties based on whether the operator’s declared semantic properties (e.g., `divergence_preserving`, `entropy_relaxing`, `conservative`) are actually realized in the field transformation it produces.

The semantic property system is the key upgrade over standard variational selection. An operator that declares itself `entropy_relaxing` but produces increasing mean entropy is penalized; one that achieves entropy reduction when it claimed none gets no bonus. This couples the operator’s declared intention to its empirical behavior, making the selection criterion self-validating. In the language of the paper’s main theorems, this enforces that the regularizer \mathcal{R} is evaluated not only as an admissibility condition on field configurations but as a consistency condition on operator semantics.

Scoring is extended to trajectory rollouts: rather than evaluating at a single step, the system simulates a morphism forward over a horizon T (typically 8–10 steps) and integrates the scoring functional over the resulting trajectory. This implements counterfactual consistency (Section 9) as the primary criterion for morphism admission: an operator is selected not because it fits the current observation but because it induces a trajectory that remains within tolerance of observations across the full horizon.

I.3 Context-Sensitive Memory and Operator Discovery

The morphism registry stores operators not as a flat table but as a context-indexed statistical manifold. Each entry carries:

- the normalized RSVP operator parameters,
- declared semantic properties,
- usage count and mean/variance of rollout scores,
- a context centroid computed as the average context vector across all uses.

Retrieval is context-weighted: given the current field configuration and sensor observations, a context vector is computed (mean Φ , mean flux, mean entropy, sensor targets), and candidate operators are ranked by a combined score:

$$\text{rank}(M) = \bar{s}(M) - \lambda d(\mathbf{c}, \bar{\mathbf{c}}(M))$$

where \bar{s} is mean rollout score, \mathbf{c} is the current context, and $\bar{\mathbf{c}}(M)$ is the operator’s context centroid. Operators are therefore preferred both when they have historically performed well and when the current context resembles the contexts in which they were used. This implements a minimal version of the contextual retrieval that distinguishes Yarncrawler from global optimization.

When no operator in the registry achieves sufficient score and no admissible type path exists between the required field spaces, the system invokes morphism synthesis: it generates symbolic candidate operators of declared types (projection, aggregation, embedding) and evaluates them against the sensor-coupled scoring functional. Candidates that pass the admission threshold (`python/admit_morphism.py`) enter the registry and become available for future retrieval and composition.

Composition discovery operates over the history log `logs/history.jsonl`: pairs of operators that were applied in sequence with consistently high rollout scores are identified

as candidate compositions (`python/discover_compositions.py`) and stored as composite morphisms. When a composite achieves high score across many distinct contexts, it is abstracted into a single atomic operator (`python/abstract_morphism.py`), compressing the composition chain into a reusable primitive. This is the operational version of the operator abstraction described in the main text: just as the operator class library grows by extending Θ with new process generators F_k , the implementation grows its registry by abstracting empirically successful operator chains.

Cross-domain transfer (`python/transfer_morphism.py`) applies the functor Ψ of Theorem 2.6 in the opposite direction: given a source domain descriptor, a target domain descriptor, and a transfer map, it transports a registry operator from one sensor bundle and topology into another. The transport scales bias parameters by the ratio of node and edge counts, maps field type signatures through the declared type map, and annotates the resulting operator with a functoriality score computed by `python/check_functoriality.py`. Transferred operators are admitted to the target registry with a score discount proportional to $(1 - s_F)$, reflecting the theoretical result that transfer is approximate rather than exact unless the transfer map is a full isomorphism.

J The Controller Quotient and Meta-Level Closure

J.1 Controllers as Functors

Each controller — beam search, dynamic programming, learned gating, spatial gating, repair — is a functor from the category of observable states to the category of candidate morphism sequences:

$$\mathcal{K} : (\text{state}, \text{registry}) \longrightarrow \text{candidate sequence}$$

The five controllers in `controllers/registry/` correspond to five distinct functorial strategies for generating morphism sequences given the current field state and observation. Each strategy reflects a different epistemic regime: beam search when uncertainty is high and exploration is warranted, dynamic programming when the regime is stable enough to admit policy decomposition, learned gating when the current context is well-represented in the registry and amortized inference is reliable, spatial gating when different regions of the field require different operator policies, and repair when an explicit obstruction has been detected and a targeted coboundary must be constructed.

The five controllers do not simply select among these strategies by hard switching. They are blended continuously. The controller confidence layer (`python/controller_confidence.py`) computes, at each time step, a weight for each controller based on: the margin between best and second-best candidates (beam confidence), the regime stability score (DP confidence), and the context-similarity of the current state to well-populated registry regions (learned confidence). These weights are fed into `python/blend_controllers.py`, which normalizes the controllers' selected operators into a common RSVP affine form and computes a weighted blend. The resulting operator inherits the property union of its contributing

controllers and is applied as a single morphism at that time step.

At a finer spatial scale, the spatial controller layer extends this blending to be node-specific. Controller weights are computed per node based on local field properties (entropy, flux magnitude, boundary status) and smoothed over the graph topology via a discrete Laplacian. The resulting node-wise weight field (`python/spatially_smooth_weights.py`) is applied in `python/blend_nodewise_controllers.py` to produce an operator that varies spatially: boundary nodes may receive high DP weight (conservative, boundary-respecting), high-entropy interior nodes may receive high beam weight (exploratory), and low-entropy familiar regions may receive high learned weight (amortized). The spatial controller is therefore not a single morphism but a field of morphisms over the topology, implemented as the nodewise hybrid type in `python/morphism_apply.py`.

J.2 Natural Transformations as Controller Evolution

A natural transformation between controllers is a mapping $\eta : \mathcal{K}_1 \Rightarrow \mathcal{K}_2$ that translates one controller’s parameters and decision structure into another’s. The implementation stores these as JSON objects in `controllers/transformations/`: for example, `beam_to_learned.json` maps beam search parameters (width, depth) to learned controller parameters (attention span, planning horizon) with an explicit weight controlling how strongly the transformation is applied.

Controller transformation is applied iteratively. At each step, the current controller state \mathcal{K}_t is fed through the transformation to produce \mathcal{K}_{t+1} (`python/iterate_controller.py`), and the distance between \mathcal{K}_t and \mathcal{K}_{t+1} is computed (`python/controller_distance.py`). When the distance falls below a threshold ε , a fixed-point signal is emitted (`python/check_fixed_point.py`). The stability score of the controller is incremented when a fixed point is detected and decayed otherwise (`python/update_stability.py`). This implements, at the controller level, the same fixed-point iteration that the main consistency operator \mathcal{C} performs at the field level.

The formal analogy is precise. The field-level fixed point is:

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

The controller-level fixed point is:

$$\eta(\mathcal{K}^*) \approx \mathcal{K}^*$$

Both are instances of the same underlying principle: a system has closed when applying its own transformation leaves it unchanged. The field-level closure corresponds to $|\mathcal{F}| = 1$ and the theorems of Sections 6 and 7. The controller-level closure corresponds to a stable decision procedure that no transformation meaningfully alters.

Transformations can be learned from the controller trace log (`logs/controller_trace.jsonl`) by identifying which controller decisions, in which contexts, led to high rollout scores (`python/learn_controller_transform.py`). This makes controller evolution data-driven rather than hand-specified: the system discovers which transformations improve decision quality and encodes them as natural transformations in the registry.

J.3 The Quotient Structure over Controller Space

Two controllers are observationally equivalent if their decision behavior is indistinguishable across all states in the registry: they select the same morphisms (or morphisms with equivalent rollout scores) in all contexts where both have been applied. This is measured by the controller distance metric, and the equivalence relation is the thresholded version: $\mathcal{K}_1 \sim \mathcal{K}_2$ iff $d(\mathcal{K}_1, \mathcal{K}_2) < \varepsilon$.

The quotient `Controllers/~` is computed by `python/build_controller_quotient.py`: it clusters all controllers into equivalence classes and selects a canonical representative per class (`python/select_representatives.py`, choosing the controller with highest stability score within each class). The resulting structure is stored in `controllers/quotient/`, and all subsequent controller selection operates over this quotient rather than over raw controllers.

This quotient structure is the meta-level counterpart of the quotient \mathcal{X}/\sim that governs field-state identifiability. Just as two field configurations that agree on all sensor projections are identified in \mathcal{X}/\sim — because the system cannot and should not distinguish them — two controllers that agree on all decision outputs across all observable contexts are identified in `Controllers/~`. The system operates on equivalence classes of decision procedures, not on individual procedures.

The consequence is a compression of the decision space that mirrors the compression of the field space. Just as the feasible set \mathcal{F} contracts toward a singleton under increasing sensor coverage and trajectory constraints, the quotient controller space contracts toward a minimal set of genuinely distinct decision structures as the system accumulates experience. Controllers that appeared distinct during early exploration may prove equivalent once the registry is populated, because their differences are only visible in contexts that the system has not yet encountered.

J.4 Closure at the Meta-Level

The system achieves meta-level closure when:

1. The morphism registry has stabilized: new operators are admitted rarely, and the candidate pool for any new context is dominated by retrieved operators rather than freshly generated ones.
2. The controller quotient has stabilized: the equivalence classes no longer split or merge as new controller traces are added, and the set of canonical representatives is fixed.
3. The transformation graph between controllers is dominated by intra-class transformations: most observed controller evolutions map a controller to an equivalent one, with inter-class transitions becoming rare.

These three conditions are the meta-level analogue of the three closure criteria of Section 9: projection parity (sensor residuals within tolerance), dynamical stability (field state on an admissible trajectory), and counterfactual consistency (forward simulation agrees with future observations). At meta-level closure, the system has not merely found a world-state

consistent with all observations — it has found a decision procedure consistent with all operating contexts. The architecture is not just closed with respect to the world; it is closed with respect to itself.

This self-referential closure is the operational realization of a principle that runs throughout the framework: a system is done not when it has found a solution but when applying its own processes to that solution leaves it unchanged. In the language of sheaves, the global section exists not only over the cover of sensor modalities but over the cover of decision contexts. The Yarncrawler constellation does not merely reconstruct a world-state from observations; it identifies the invariant decision structure that makes reconstruction reliable across all the worlds it might encounter.

References

- [1] de Gyurky, S. M., & Tarbell, M. A. (2013). *The Autonomous System: A Foundational Synthesis of the Sciences of the Mind*. Wiley. ISBN 978-1-118-75740-6.
- [2] Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- [3] Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2019). Active inference: a process theory. *Neural Computation*, 31(1), 1–49.
- [4] Hordijk, W., & Steel, M. (2019). Autocatalytic networks: an intimate relation between network topology and dynamics. *Bulletin of Mathematical Biology*, 81(5), 1357–1374.
- [5] Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel.
- [6] Rosen, R. (1985). *Anticipatory Systems*. Pergamon Press.
- [7] Abramsky, S., & Brandenburger, A. (2011). The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13, 113036.
- [8] Spivak, D. I. (2014). *Category Theory for the Sciences*. MIT Press.
- [9] Lewis, P., Perez, E., et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33.
- [10] Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2021). NeRF: representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1), 99–106.
- [11] Freidlin, M. I., & Wentzell, A. D. (1984). *Random Perturbations of Dynamical Systems*. Springer.