

TARTAN: Trajectory-Aware Recursive Tiling with Annotated Noise

A Sheaf-Variational Architecture for World-State Reconstruction and
Ontological Adaptation

Flyxion

Independent Researcher

April 8, 2026

Abstract

This monograph develops the TARTAN architecture as a unified mathematical and computational framework for inference understood as constraint closure rather than prediction. We formalize reconstruction as the identification of a globally consistent world-state from partial projections under dynamical and thermodynamic constraints, grounded in the Relativistic Scalar-Vector Plenum (RSVP) field theory. Local reconstructions are represented as sections over a recursive tiling, with global consistency enforced via sheaf-theoretic gluing and diagnosed through cohomological obstruction.

We introduce a staged closure protocol that resolves incompatibilities through gauge repair, refinement, and state extension, and extend the framework temporally through trajectory-aware coarse-graining that preserves obstruction-relevant invariants. Novelty is treated as a formal signal of ontological insufficiency, triggering a controlled expansion of the model guided by compression and cross-context coherence. The resulting system is not a static inference engine but an adaptive, self-ontologizing process that achieves global consistency through entropy-respecting reconstruction.

Contents

1	The Strategic Pivot: From Prediction to Constraint Closure	3
1.1	The Failure of Predictive Consistency	3
1.2	Constraint Closure as the Correct Objective	3
1.3	From Estimation to Reconstruction	4
1.4	Consequences for System Design	4
1.5	Toward an Epistemology of Consistency	4
2	Preliminaries: World-States, Observations, and Reconstruction	5
2.1	The Admissible Space and Feasibility	5
2.2	RSVP as the Physical Substrate	6
3	The Single-Pass Execution Narrative	6
4	Obstruction as the Primitive of Failure	7
5	From Operational Closure to Sheaf-Theoretic Structure	8
5.1	Local States as Sections	8
5.2	Obstruction as Cohomology	8
5.3	Why the Categorical Language is Unavoidable	8
6	Mathematical Formalism	9
6.1	The Yarncrawler Principle and Variational Closure	9
6.2	Variational Reconstruction and Admissibility	10
6.3	Sheaf-Theoretic Gluing and Global Sections	10
6.4	Cohomological Obstruction as Compatibility Defect	11
6.5	Spectral-Sheaf Duality	11
6.6	Timescale Separation and Quasi-Static Approximation	11
7	The TARTAN Architecture	12
7.1	Recursive Tiling as Computational Substrate	12
7.2	Local State and Admissibility Certificates	12
7.3	Overlaps as First-Class Computational Objects	13
7.4	The Staged Closure Protocol	13
7.5	Decision Procedure via Defect Structure	14
7.6	Global Closure as Iterative Convergence	14
8	Temporal Dynamics and Trajectory Structure	15
8.1	Trajectories as Primary Objects	15
8.2	Coarse-Grained Trajectory Summaries	15

8.3	Commutation with Projection Operators	16
8.4	Gauge-Invariant Compression	16
8.5	Temporal Obstruction and Path Consistency	16
8.6	Integration with Closure Dynamics	17
9	Obstruction, Novelty, and Learning	17
9.1	The Obstruction Log as a Learned Atlas	17
9.2	Novelty Detection as Failure of Closure	17
9.3	The Exploratory Closure Regime	18
9.4	Ontological Growth and Promotion Criteria	18
9.5	Learning the Gluing Category	18
10	Defensive Inference and Sensor Integrity	19
10.1	Malignant Novelty and Adversarial Structure	19
10.2	The Defensive Branch: Corruption Modeling	19
10.3	Cross-Context Transportability and Modality Ablation	19
10.4	Causal Reciprocity and the Rejection of Explanation Sinks	19
10.5	Global Coherence as Decision Principle	20
11	Sensor Lifecycle and Re-Integration	20
11.1	Quarantine and Provenance Tracking	20
11.2	Corrected Projection Operators	20
11.3	Gradual Re-Weighting and Probation	21
11.4	Temporal Monitoring and Adaptive Reliability	21
12	CLIO and Recursive Optimization	21
12.1	The Consistency Operator as a Functor	21
12.2	CLIO: Cognitive Loop via In-Situ Optimization	21
12.3	Fixed Points and Stability of the CLIO Loop	23
13	Implementation and Systems Design	23
13.1	Core Data Structures	23
13.2	Algorithmic Pipeline	23
13.3	Computational Complexity and Scaling	24
13.4	Parallelization and Distributed Execution	24
14	Failure Modes as First-Class Structure	24
14.1	Projection Degeneracy	25
14.2	Summary Collapse	25
14.3	Ontological Overgrowth	25
14.4	Defensive Overfitting	26

14.5	Oscillatory Non-Convergence	26
14.6	Failure Mode Summary	26
15	Epistemology and Philosophy of Inference	27
15.1	Inference as Physical Process	27
15.2	Closure, Identity, and Irreversibility	27
15.3	Ontological Humility and Adaptive Knowledge	27
15.4	Knowledge as Global Consistency	28
15.5	Synthesis: A Theory of Inference	28
A	Formal Proofs	29
A.1	Identifiability Under Projection and Convexity	29
A.2	Failure Modes	29
A.3	Equivalence of Sheaf Gluing and Variational Closure	29
A.4	Convergence of the Staged Closure Protocol	30
A.5	Stability of CLIO Fixed Points	30
B	Functional Analytic Foundations	30
B.1	Function Spaces and Admissible States	30
B.2	Weak Formulation of RSVP Dynamics	31
C	Formal Equivalence Between Yarncrawler Closure and Sheaf De-	
	scient	31
C.1	Statement	31
C.2	Equivalence Theorem	31
C.3	Corollary	32
C.4	Functorial Form	32
D	Category-Theoretic Formalization	32
D.1	The Category of Local Reconstructions	32
D.2	Čech Cohomology and Obstruction Classes	33
D.3	Learning as Category Expansion	33
E	Haskell DSL Sketch for TARTAN	33
E.1	Motivation	33
E.2	Core Typed Objects	33
E.3	Projection and Restriction Typeclasses	34
E.4	Repair Morphisms as Typed Transformations	34
E.5	A Minimal TARTAN EDSL	34

F	Numerical Implementation: RSVP–TARTAN Lattice	35
F.1	Grid and Field Types	35
F.2	RSVP Evolution	35
F.3	Pairwise Overlap Repair	36
F.4	Main Iteration Loop	36
F.5	Visualization	37
F.6	Interpretation	37
G	CLIO as a Functor on Fibered Structures	37
G.1	Fibered Structure of Local Reasoning States	37
G.2	Defects as Čech 1-Cocycles	38
G.3	CLIO as an Endofunctor	38
G.4	Functoriality and Projection Compatibility	38
G.5	CLIO as an Approximate Closure Functor	38
G.6	Oscillatory Regimes and Obstruction Classes	39
H	Derived Stack Interpretation of TARTAN Covers	39
H.1	Covers as Atlases for Derived Stacks	39
H.2	Global Reconstruction as a Homotopy Colimit	39
H.3	Repair Operations as Homotopies	40
H.4	Interaction with CLIO	40
I	Semantic Merge as Homotopy Colimit Computation	40
I.1	From Syntactic to Semantic Versioning	40
I.2	Merge as Homotopy Colimit	41
I.3	CLIO as a Merge Engine	41
I.4	The General Principle	42

1 The Strategic Pivot: From Prediction to Constraint Closure

The dominant paradigm in machine learning treats inference as the estimation of conditional distributions. Given observations y , a model approximates $p(X | y)$ or produces point estimates \hat{X} minimizing a loss function. This framing presupposes that observations are independent targets and that correctness is reducible to local predictive accuracy.

This assumption is structurally incorrect.

Observations are not independent samples. They are projections $y_i = \Pi_i(X)$ of a shared underlying world-state X . As such, they are coupled by the requirement that they originate from a single coherent configuration. The predictive paradigm ignores this coupling. It optimizes each projection individually without enforcing the existence of a global state that realizes them simultaneously. As a result, it permits collections of locally accurate predictions that are globally incompatible.

1.1 The Failure of Predictive Consistency

Let $\{\hat{y}_i\}$ be predictions satisfying $\hat{y}_i \approx y_i$ for all i . There is, in general, no guarantee that there exists a state \hat{X} such that $\Pi_i(\hat{X}) = \hat{y}_i$ for all i . This failure is not due to noise or insufficient data; it is a consequence of the objective itself. By optimizing projections independently, the system abandons the requirement of global realizability.

We define this phenomenon as *predictive inconsistency*: the production of locally valid outputs that cannot be assembled into a coherent world-state. In physical systems, analogous inconsistency is categorically excluded—measurements of a shared state must agree under transformation and restriction. The absence of such guarantees in predictive architectures leads to fragmentation, hallucination, and the inability to reconcile conflicting evidence.

1.2 Constraint Closure as the Correct Objective

To resolve this, we replace prediction with *constraint closure*.

Definition 1 (Constraint Closure). *A state X^* is admissible if and only if it is a fixed point of a consistency operator \mathcal{C} :*

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

where \mathcal{C} enforces all observational and structural constraints simultaneously.

The operator \mathcal{C} integrates projection consistency with observed data, dynamical and thermodynamic constraints encoded by the RSVP field theory, and structural regularity conditions. Inference is thus redefined as the search for a fixed point of \mathcal{C} rather than the minimization of prediction error. Under this formulation, inference is a dynamical process $X_{t+1} = \mathcal{C}(X_t)$ in which convergence indicates consistency, and correctness is characterized by invariance under constraint enforcement rather than proximity to observations.

1.3 From Estimation to Reconstruction

This shift transforms the role of the model. A predictive model approximates a function from inputs to outputs; a closure system reconstructs a configuration satisfying all constraints simultaneously. The former is evaluated locally; the latter is validated globally. Inference is not the production of likely values but the recovery of a coherent structure from incomplete projections.

Reconstruction proceeds not by enforcing global agreement directly, but by iteratively eliminating obstruction in an invariant quotient space where only closure-relevant structure is retained. The system does not attempt to reconcile full states, but only their invariant summaries sufficient for obstruction classification. This one distinction explains why the architecture scales.

1.4 Consequences for System Design

This reframing has immediate architectural implications. Local accuracy is insufficient; systems must enforce compatibility across overlapping observations. Inconsistency must be represented explicitly: rather than collapsing error into a scalar loss, the system must track structured defects encoding where and how compatibility fails. Finally, learning becomes inseparable from reconstruction—when closure cannot be achieved within the current representation, the system must expand its state space to resolve persistent obstruction. These requirements define the foundation of the TARTAN architecture.

1.5 Toward an Epistemology of Consistency

The transition from prediction to constraint closure is not a refinement of existing methods; it is a change in epistemology. A model is no longer judged by how well it predicts observations in isolation, but by whether it can construct a world in which those observations can all be true simultaneously. This principle underlies the entire TARTAN framework.

Proposition 1 (No Free Global Section). *Let $\{\Pi_i\}_{i \in I}$ be a family of projection operators and let $\{\hat{y}_i\}$ be a collection of predictions optimized independently, without*

enforcing a compatibility condition on overlaps. Then there exists no algorithmic guarantee that a state \hat{X} with $\Pi_i(\hat{X}) \approx \hat{y}_i$ for all i is globally realizable.

Proof. Each prediction \hat{y}_i is produced by an optimization that treats y_i as an isolated target. Without a shared constraint, the collection $\{\hat{y}_i\}$ need not lie in the image of any single \hat{X} under the joint projection $\tilde{\Pi}(X) = (\Pi_i(X))_i$. Non-injectivity of $\tilde{\Pi}$ or violation of compatibility conditions on overlaps suffices to preclude global realization. No post-hoc assembly can recover global coherence once it has been discarded at the level of the objective. \square

2 Preliminaries: World-States, Observations, and Reconstruction

We formalize the basic objects of the theory before introducing its machinery. Let Ω denote a domain of interest—spatial, temporal, or abstract—and let $X \in \mathcal{X}$ denote a *world-state*, understood as a structured configuration over Ω . Observations are not direct access to X but projections

$$y_i = \Pi_i(X) + \varepsilon_i,$$

where $\Pi_i : \mathcal{X} \rightarrow \mathcal{Y}_i$ are partial, possibly lossy operators and ε_i represents noise or corruption. The central problem is therefore not prediction of y_i but *reconstruction of X from incompatible, incomplete projections*.

This introduces a structural asymmetry that has no counterpart in the predictive paradigm. While X is globally coherent by definition, the collection $\{y_i\}$ need not be: noise, distinct projection geometries, and sensor failures all conspire to produce locally incompatible observations. Any inference procedure must therefore resolve inconsistencies that are not present in the underlying world but arise from the projection process itself. TARTAN treats this as the primary object of study.

2.1 The Admissible Space and Feasibility

Not every configuration $X \in \mathcal{X}$ is physically or structurally valid. We define the admissible space $\mathcal{A} \subset \mathcal{X}$ as the set of configurations satisfying the intrinsic constraints encoded by the RSVP dynamics. Given observations $\{y_i\}$, the feasible set is

$$\mathcal{F} = \{X \in \mathcal{A} \mid \|\Pi_i(X) - y_i\| \leq \varepsilon_i \quad \forall i \in I\}.$$

The reconstruction problem is the identification of an element $X^* \in \mathcal{F}$ that is not merely feasible but globally consistent—a fixed point of a consistency operator

that integrates all constraints simultaneously. The existence of such a point is not guaranteed; its absence is precisely the condition we term *obstruction*.

2.2 RSVP as the Physical Substrate

The admissible space \mathcal{A} is structured by the Relativistic Scalar–Vector Plenum (RSVP) theory, in which the world-state is represented as a triple

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

where Φ is a scalar potential encoding potential structure, \mathbf{v} is a vector flow field governing transport and momentum, and S is an entropy field tracking dissipation and irreversibility. The fields evolve over Ω according to coupled nonlinear PDEs that enforce conservation laws, flow constraints, and entropy production.

TARTAN is RSVP operationalized. The repair operators that constitute the TARTAN closure protocol are not arbitrary transformations; they correspond to descent steps on the energy functional $\mathcal{E}(X) = \sum_i \|\Pi_i(X) - y_i\|^2 + \lambda \mathcal{D}_{\text{RSVP}}(X)$, where $\mathcal{D}_{\text{RSVP}}$ penalizes configurations that violate RSVP dynamics. Every repair operation therefore has a dual character: it reduces cohomological obstruction in the sheaf-theoretic layer while simultaneously descending toward the thermodynamic manifold of admissible RSVP configurations. The two reductions are the same reduction, expressed in different languages.

3 The Single-Pass Execution Narrative

Before introducing formal machinery, we describe the operation of the full system as a single pass through the TARTAN stack. This narrative is not a summary of later sections; it is a binding description intended to show how the components articulate into a dynamical whole. Formal development of each component follows in subsequent sections.

A stream of observations enters the system and is partitioned into a cover $\{U_i\}$ of tiles. Each tile constructs a provisional local state $s_i \in \mathcal{S}(U_i)$ that satisfies its local projection and admissibility constraints. At this stage, tiles are individually consistent but mutually unconstrained; there is no reason to expect that they agree on their boundaries.

The system then computes coarse-grained summaries $\sigma_i = \sigma(s_i)$ of each tile, capturing the invariant structure relevant for overlap comparison. These summaries are designed to commute with projection operators and to remain stable under refinement, so that comparison at the summary level neither introduces spurious defects nor misses genuine ones.

Next, the system evaluates overlaps between adjacent tiles. For each pair (i, j) with $U_i \cap U_j \neq \emptyset$, it computes the defect tensor

$$\check{\delta}_{ij} = \rho_{ij}(s_i) - \rho_{ji}(s_j),$$

which measures the failure of agreement on shared boundary data. Each defect is classified according to its behavior under normalization and refinement: gauge defects vanish under reparameterization, resolution defects decay under tile subdivision, and structural defects persist under both transformations. The classification determines which of the three repair operations—gauge alignment, refinement, or state extension—is applied.

The repaired states are reinserted into the system, summaries are recomputed, and the defect evaluation repeats. The CLIO module simultaneously updates the parameters of the consistency operator \mathcal{C}_θ based on the aggregate defect $\check{\delta}_t$ and current state X_t , so that the operator itself improves with each pass. The loop continues until $\mathcal{D}_{\text{total}} = \sum_{i,j} \|\check{\delta}_{ij}\|^2$ falls below a tolerance threshold, at which point the collection $\{s_i\}$ forms a compatible family and a global section exists.

The output is a globally consistent world-state assembled from locally constrained components through iterative defect elimination. The system has not predicted this state; it has reconstructed it by eliminating every structural reason it could not exist.

4 Obstruction as the Primitive of Failure

The fundamental unit of failure in the TARTAN framework is not prediction error but the presence of an *obstruction* to global closure. Given local states $\{s_i\}$, an obstruction arises when their restrictions to overlaps fail to agree—when the defects $\{\check{\delta}_{ij}\}$ cannot be eliminated by local adjustments and instead represent a nontrivial cohomology class in $H^1(\mathcal{U}, \mathcal{S})$.

This reinterpretation is consequential. Error is a scalar quantity that conflates many distinct failure modes. Obstruction is a structured object encoding the geometry of inconsistency. Its signature—its behavior under normalization, refinement, and context variation—distinguishes between coordinate artifacts, resolution failures, genuine missing degrees of freedom, and sensor corruption. Each signature calls for a different response; none is reducible to the others by gradient descent.

Rather than treating obstruction as a terminal failure, the system interprets it as structured information. Each obstruction class encodes what the current representation cannot express. The repair protocol responds to this information constructively: gauge repair eliminates trivial classes, refinement eliminates resolution-dependent classes, and state extension eliminates structural classes by enlarging

the space of sections. Inference therefore proceeds by iteratively transforming obstruction into structure—absorbing what was previously inexpressible into the model’s ontology—until no further inconsistencies remain.

5 From Operational Closure to Sheaf-Theoretic Structure

The preceding sections described inference as iterative defect elimination. While this operational view is sufficient for implementation, it raises a deeper question: what is the mathematical structure of assembling local states into a global one? To answer this, we require a language capable of expressing simultaneously the locality of construction, the compatibility requirement on overlaps, and the existence or obstruction of a global solution. Sheaf theory provides precisely this structure.

5.1 Local States as Sections

We reinterpret each tile U_i not merely as a computational unit but as an open set in a topological cover of Ω . A local reconstruction s_i becomes a *section* over U_i ,

$$s_i \in \mathcal{S}(U_i),$$

where \mathcal{S} is a sheaf assigning admissible states to regions. The admissibility conditions—projection consistency, RSVP regularization, and summary constraints—define the space of valid sections. The restriction maps $\rho_{ij} : \mathcal{S}(U_i) \rightarrow \mathcal{S}(U_i \cap U_j)$ formalize the intuitive notion of agreement on overlaps introduced in the execution narrative.

5.2 Obstruction as Cohomology

The collection of defects $\{\check{\delta}_{ij}\}$ defines a Čech 1-cochain. When these defects vanish globally, the local sections glue to a single global section $X \in \mathcal{S}(\Omega)$. When they do not vanish, they represent a nontrivial cohomology class $[\check{\delta}] \in H^1(\Omega, \mathcal{S})$ obstructing global reconstruction. This is a precise mathematical interpretation of the operational notion of obstruction: not merely error, but a topological barrier to closure.

5.3 Why the Categorical Language is Unavoidable

The sheaf-theoretic formulation reveals that reconstruction is fundamentally about relations between states—restrictions, transformations, and repairs—rather than the states themselves. These relations form a category. Moreover, the process of merging multiple reconstructions corresponds to computing a universal object coherently combining them, which is naturally expressed as a colimit. Category theory is therefore not introduced for abstraction but emerges as the minimal language

capable of expressing compositional structure, equivalence under transformation, and universal constructions for merging.

The closure operator \mathcal{C} introduced operationally can now be viewed as a computational procedure approximating the universal colimit object. The iterative repair loop is not an ad hoc algorithm but a concrete realization of a well-defined mathematical construction: the computation of a global section by elimination of cohomological obstruction. With this bridge in place, we can formalize TARTAN as a fibered category whose objects are local states, whose morphisms are transformations and restrictions, and whose colimits are merged global reconstructions. The following sections make this structure explicit and derive conditions under which closure is guaranteed.

6 Mathematical Formalism

6.1 The Yarncrawler Principle and Variational Closure

With the preceding conceptual framework established, we develop the formal machinery. The reconstruction process is governed by a consistency operator \mathcal{C} that iteratively refines candidate states. This operator integrates observational constraints, dynamical laws, and regularization terms into a single update rule derived from the energy functional

$$\mathcal{R}(X) = \sum_i \|\Pi_i(X) - y_i\|^2 + \lambda \cdot \mathcal{D}_{\text{RSVP}}(X),$$

where $\mathcal{D}_{\text{RSVP}}$ measures deviation from admissible RSVP dynamics. The operator \mathcal{C} may be interpreted concretely as a discretized gradient flow or proximal map induced by \mathcal{R} , anchoring it simultaneously in variational analysis and the categorical structure described above.

Theorem 1. *Assume that the admissible space \mathcal{A} is strictly convex and that the induced projection map is injective on the feasible set \mathcal{F} . Then there exists a unique state X^* such that $\mathcal{C}(X^*) = X^*$.*

This result establishes identifiability. Under appropriate conditions, reconstruction converges to a unique solution. Failure of these conditions leads to two distinct modes of breakdown: projection degeneracy, in which multiple states produce indistinguishable observations, and regularizer flatness, in which the constraint landscape lacks curvature. Both are failure modes in the formal sense developed in Section 14.

The Yarncrawler principle interprets inference as a trajectory through state space guided by \mathcal{C} toward this fixed point. Closure is achieved when the trajectory stabilizes,

yielding a state that satisfies all constraints simultaneously.

6.2 Variational Reconstruction and Admissibility

The reconstruction problem is expressed as a constrained variational system. Let $X = (\Phi, \mathbf{v}, S) \in \mathcal{A}$ denote a candidate state over a domain Ω . We define an energy functional

$$\mathcal{E}(X) = \sum_{i \in I} \|\Pi_i(X) - y_i\|^2 + \lambda \mathcal{D}_{\text{RSVP}}(X),$$

where $\mathcal{D}_{\text{RSVP}}$ encodes deviation from admissible field dynamics, including transport constraints, entropy production bounds, and coupling relations between Φ , \mathbf{v} , and S .

Critical points of \mathcal{E} satisfy Euler–Lagrange equations of the form

$$\frac{\delta \mathcal{E}}{\delta X} = 0,$$

which combine observational consistency with dynamical admissibility. The resulting system is generally nonlinear and must be solved iteratively.

Admissibility is not a static constraint but a dynamical condition. A state is admissible if it lies on or sufficiently near the manifold defined by $\mathcal{D}_{\text{RSVP}}(X) = 0$. In practice, admissibility is enforced through penalization and projection, ensuring that reconstruction remains within a physically meaningful regime.

6.3 Sheaf-Theoretic Gluing and Global Sections

To capture the locality of observations and the necessity of global consistency, we introduce a cover $\{U_i\}_{i \in I}$ of the domain Ω . Each tile U_i supports a local reconstruction $s_i \in \mathcal{S}(U_i)$, where \mathcal{S} is a presheaf assigning admissible states to open sets.

For overlapping regions $U_i \cap U_j \neq \emptyset$, we define restriction maps

$$\rho_{ij} : \mathcal{S}(U_i) \rightarrow \mathcal{S}(U_i \cap U_j),$$

which extract the portion of the local state relevant to the overlap.

Definition 2. *A collection $\{s_i\}$ is compatible if*

$$\rho_{ij}(s_i) = \rho_{ji}(s_j) \quad \text{on } U_i \cap U_j \quad \forall i, j.$$

When compatibility holds, the local sections glue to a global section $X \in \mathcal{S}(\Omega)$. Thus, reconstruction reduces to finding a compatible family of local states.

6.4 Cohomological Obstruction as Compatibility Defect

When compatibility fails, the discrepancy can be quantified. Define the defect tensor

$$\check{\delta}_{ij} = \rho_{ij}(s_i) - \rho_{ji}(s_j).$$

This object measures the failure of local sections to agree on overlaps.

Definition 3. *The collection $\{\check{\delta}_{ij}\}$ forms a Čech 1-cocycle if it satisfies the cocycle condition on triple overlaps. Nontrivial cocycles represent obstructions to gluing.*

In classical sheaf theory, such obstructions are elements of the cohomology group $H^1(\Omega, \mathcal{S})$. In the TARTAN framework, this abstract obstruction is reinterpreted as a computational diagnostic. The defect tensor becomes an actionable signal indicating where and how closure fails.

The magnitude and structure of $\check{\delta}_{ij}$ determine the nature of the inconsistency. Small, smooth discrepancies may indicate gauge misalignment, while high-frequency or persistent defects suggest under-resolution or missing state variables.

6.5 Spectral-Sheaf Duality

Reconstruction operates simultaneously in two domains: spatial consistency and spectral representation. Local states can be decomposed into basis functions, for example via Fourier or wavelet expansions,

$$s_i(x) = \sum_k a_{i,k} \psi_k(x).$$

Spectral truncation provides a compressed representation that captures dominant modes of variation while discarding high-frequency noise.

The key requirement is that spectral compression must preserve gluing structure. That is, if two local sections are compatible in full representation, their truncated forms must remain approximately compatible. Conversely, incompatibility in low-frequency modes signals a genuine large-scale obstruction.

This duality implies that closure requires agreement both in physical space and in the space of dominant modes. High-frequency discrepancies may be resolved by refinement, while low-frequency mismatches necessitate deeper intervention.

6.6 Timescale Separation and Quasi-Static Approximation

The reconstruction process involves both fast and slow variables. Observational data may fluctuate rapidly, while structural features of the RSVP fields evolve more slowly. To prevent transient noise from destabilizing the reconstruction, we introduce a timescale separation.

Let $X(t)$ denote the evolving state. We decompose its dynamics into a fast component X_f and a slow component X_s . The fast component captures immediate adjustments to observational discrepancies, while the slow component encodes persistent structure.

In the quasi-static approximation, the system assumes that X_s remains approximately constant over short intervals, allowing X_f to relax toward compatibility. This prevents rapid oscillations from propagating into the structural variables.

Formally, we require that updates to X_s occur only when residuals exhibit persistence over a temporal window. This ensures that structural changes reflect genuine features of the underlying state rather than transient fluctuations.

The result is a stable reconstruction process in which closure is achieved through gradual refinement of both local states and their global organization. Temporal consistency is thus integrated into the spatial gluing framework, preparing the ground for trajectory-aware extensions in subsequent sections.

7 The TARTAN Architecture

7.1 Recursive Tiling as Computational Substrate

The TARTAN architecture represents the domain Ω not as a monolithic space but as a recursively structured cover $\{U_i\}_{i \in I}$ whose elements, or tiles, form the primary units of computation. Each tile supports a local reconstruction that is both observationally constrained and dynamically admissible. The recursive nature of the tiling permits adaptive refinement, allowing the system to allocate resolution where required by the structure of the state.

Formally, the cover is organized as a hierarchy in which each tile may be subdivided into a collection of child tiles whose union recovers the parent domain. This induces a tree-like structure on the index set I , with refinement corresponding to descent along the tree. The admissible state over Ω is therefore represented not by a single object but by a family of local sections indexed by this hierarchy.

This construction provides a natural mechanism for multiscale inference. Coarse tiles capture large-scale structure with low computational cost, while fine tiles resolve localized discrepancies. The recursive tiling thus functions as both a representational and computational substrate, enabling the system to balance efficiency and fidelity.

7.2 Local State and Admissibility Certificates

Each tile U_i carries a local state $s_i \in \mathcal{S}(U_i)$ together with an admissibility certificate that quantifies its consistency with both observations and dynamics. The certificate is not merely a boolean indicator but a structured object containing residuals, norms,

and diagnostic quantities.

The observational component of the certificate measures the discrepancy between projected state and observed data, typically in the form $\|\Pi_i(s_i) - y_i\|$. The dynamical component evaluates the deviation from RSVP field equations, encoded through residuals of the governing PDEs. Together, these quantities determine whether the local state lies within an acceptable neighborhood of the admissible manifold.

The certificate serves two roles. First, it provides a local validation that the reconstruction is internally coherent. Second, it supplies the information required for global reconciliation, as discrepancies in certificates across overlaps signal the presence of obstruction.

7.3 Overlaps as First-Class Computational Objects

In conventional approaches, overlaps between regions are treated as passive boundaries across which information is exchanged. In TARTAN, overlaps $U_i \cap U_j$ are elevated to active computational entities. Each overlap hosts the evaluation of compatibility between neighboring tiles and the computation of the defect tensor.

Given local states s_i and s_j , the overlap region supports the quantities $\rho_{ij}(s_i)$ and $\rho_{ji}(s_j)$, whose difference defines the defect

$$\check{\delta}_{ij} = \rho_{ij}(s_i) - \rho_{ji}(s_j).$$

The overlap thus becomes the locus of inconsistency detection. Rather than forcing agreement prematurely, the system records and analyzes these discrepancies, allowing them to guide subsequent repair operations.

This elevation of overlaps to first-class status transforms the architecture. Computation is no longer confined to nodes (tiles) but is distributed across edges (overlaps), where the essential information about global consistency resides.

7.4 The Staged Closure Protocol

The central mechanism by which TARTAN achieves global consistency is the staged closure protocol. Given a nonzero defect tensor $\check{\delta}_{ij}$, the system must determine the appropriate repair strategy. This decision is not arbitrary but is guided by the structure of the defect across scales and representations.

The first stage is gauge repair. If the defect vanishes under a reparameterization or normalization of the local states, it is interpreted as a coordinate artifact rather than a genuine inconsistency. In this case, the system aligns the representations without altering the underlying state.

If the defect persists under gauge normalization but decays under refinement, the system identifies the discrepancy as a consequence of insufficient resolution. The tile

is subdivided, and the reconstruction is recomputed at a finer scale.

If the defect persists across both normalization and refinement, it is classified as a structural inconsistency. In this case, the system introduces an extension of the state space, augmenting the local representation with additional degrees of freedom. The goal is to enlarge the admissible space so that a compatible configuration becomes possible.

These three stages—gauge repair, refinement, and state extension—constitute a hierarchy of interventions, ordered by increasing structural cost. The protocol ensures that the system applies the minimal modification necessary to restore compatibility.

7.5 Decision Procedure via Defect Structure

The classification of defects relies on their behavior under transformations. Let \mathcal{N} denote a normalization operator removing gauge degrees of freedom, and let \mathcal{R} denote a refinement operator. The defect is analyzed through the compositions $\mathcal{N}(\check{\delta}_{ij})$ and $\mathcal{R}(\check{\delta}_{ij})$.

If $\mathcal{N}(\check{\delta}_{ij}) \approx 0$, the defect is attributed to gauge misalignment. If $\mathcal{N}(\check{\delta}_{ij})$ remains nonzero but $\mathcal{R}(\check{\delta}_{ij}) \rightarrow 0$ under successive refinement, the defect is classified as resolution-dependent. If neither transformation eliminates the defect, it is deemed structural.

This procedure provides a concrete operationalization of cohomological obstruction. Rather than treating obstruction as an abstract property, the system computes its behavior under controlled transformations and selects the appropriate repair accordingly.

7.6 Global Closure as Iterative Convergence

The TARTAN architecture operates as an iterative process in which local reconstructions, overlap evaluations, and repair operations are repeatedly applied. At each iteration, the system reduces the aggregate defect across all overlaps while maintaining admissibility within each tile.

Convergence is achieved when the defect tensor vanishes (or falls below a specified tolerance) across all overlaps, and all local states satisfy their admissibility certificates. At this point, the collection $\{s_i\}$ forms a compatible family, and a global section exists.

The process is inherently distributed. Each tile and overlap performs computations based on local information, yet the collective dynamics drive the system toward a globally consistent state. Closure thus emerges from the interplay of local validation and global reconciliation, mediated by the recursive tiling structure.

8 Temporal Dynamics and Trajectory Structure

8.1 Trajectories as Primary Objects

The preceding formulation treats reconstruction as the recovery of a static state. However, in a dynamical system governed by RSVP fields, the state is intrinsically temporal. Observations arise not from isolated configurations but from trajectories through state space. Consequently, inference must be reformulated to operate over histories rather than instantaneous states.

Let \mathcal{X} denote the space of admissible states. A trajectory is a mapping

$$\gamma : [t_0, t_1] \rightarrow \mathcal{X},$$

such that $\gamma(t)$ satisfies the dynamical constraints for all t . The reconstruction problem is thus elevated to identifying a trajectory γ^* whose projections match the observed data over time and whose evolution is dynamically admissible.

This shift resolves a fundamental ambiguity in static reconstruction. Multiple instantaneous states may be observationally indistinguishable, but their trajectories may diverge. By incorporating temporal structure, the system gains access to causal constraints that increase identifiability and stabilize inference.

8.2 Coarse-Grained Trajectory Summaries

While trajectories provide richer information, they are computationally expensive to store and compare. TARTAN therefore introduces coarse-grained summaries that capture only those aspects of the trajectory relevant to obstruction and closure.

Let σ denote a coarse-graining operator mapping trajectories to summary representations,

$$\sigma : \{\gamma\} \rightarrow \Sigma,$$

where Σ is a space of low-dimensional invariants. The design of σ is governed by a single principle: it must preserve the features of the trajectory that influence compatibility across overlaps.

In practice, σ encodes spectral signatures of the trajectory, moments of constraint residuals, and measures of causal orientation such as entropy flux and flow directionality. These quantities capture the dominant modes of variation, the persistence of admissibility violations, and the directional structure of evolution.

8.3 Commutation with Projection Operators

A critical requirement for the coarse-graining operator is that it approximately commutes with projection operators. For each observation operator Π_i , we require

$$\Pi_i(\sigma(\gamma)) \approx \sigma(\Pi_i(\gamma)).$$

This condition ensures that summarization does not introduce inconsistencies between local and global representations.

Commutation guarantees that summaries can be compared across overlaps without reference to the full trajectory. It aligns the representation of temporal information with the structure of the observational pipeline, preserving coherence across scales.

8.4 Gauge-Invariant Compression

Before coarse-graining, the system must eliminate redundant degrees of freedom associated with gauge transformations. These include coordinate reparameterizations, phase shifts, and other symmetries that do not alter the underlying physical state.

Let \mathcal{N} denote a normalization operator that maps trajectories into a canonical representation. The coarse-graining operator is then applied to the normalized trajectory,

$$\sigma(\gamma) = \tilde{\sigma}(\mathcal{N}(\gamma)).$$

This procedure ensures that summaries reflect genuine structural features rather than artifacts of representation.

8.5 Temporal Obstruction and Path Consistency

In the temporal setting, obstruction extends beyond instantaneous mismatch to include inconsistencies in the evolution of neighboring tiles. Two tiles may agree at a given time while diverging in their inferred histories or predicted futures.

To capture this, the defect tensor is extended to incorporate trajectory summaries. For overlapping tiles U_i and U_j , we define a temporal defect

$$\check{\delta}_{ij}^{\text{time}} = \sigma(\gamma_i) - \sigma(\gamma_j),$$

where γ_i and γ_j are the local trajectories. This quantity measures the failure of parallel transport through time, detecting inconsistencies in causal structure such as mismatched entropy flows or incompatible dynamical trends.

8.6 Integration with Closure Dynamics

The inclusion of trajectory summaries modifies the closure process. Repair operations must now address both spatial and temporal inconsistencies. A defect that appears small in instantaneous state may be significant when considered over time.

The staged closure protocol is extended accordingly. Gauge repair now includes alignment of temporal phases, refinement addresses discrepancies in high-frequency temporal variation, and state extension introduces new variables capable of capturing missing dynamical structure. The result is a reconstruction process that operates over trajectories rather than snapshots.

9 Obstruction, Novelty, and Learning

9.1 The Obstruction Log as a Learned Atlas

Within the TARTAN architecture, obstruction is not merely an error signal to be eliminated but a structured source of information about the limits of the current model. To preserve and exploit this information, the system maintains an obstruction log—a repository of coarse-grained defect summaries indexed by their invariant characteristics.

Each entry in the log consists of a summary representation of the defect, its associated context, and the repair operation that successfully resolved it. Over time, this collection forms a learned atlas of incompatibility classes and their corresponding morphisms.

The obstruction log thus functions as an adaptive memory. Rather than repeatedly rediscovering the same repair strategies, the system retrieves previously successful interventions when similar defect signatures arise. This transforms reconstruction from a purely reactive process into one informed by accumulated structural experience.

9.2 Novelty Detection as Failure of Closure

A defect is classified as novel when its summary representation does not match any existing class in the obstruction log under the invariant metric. Formally, let $\Sigma_{\text{log}} \subset \Sigma$ denote the set of stored summaries. A new summary σ is considered novel if

$$\min_{\sigma' \in \Sigma_{\text{log}}} d(\sigma, \sigma') > \delta,$$

for a suitable threshold δ , where d is the invariant distance.

Novelty is therefore defined as a failure of closure under the current atlas. It signals that the existing repertoire of repairs is insufficient to resolve the observed inconsistency, prompting further investigation.

9.3 The Exploratory Closure Regime

When a defect is confirmed as novel, the system enters an exploratory closure regime in which multiple repair strategies are evaluated in parallel. The refinement branch tests whether the defect is a consequence of insufficient resolution. The enrichment branch increases the capacity of the summary representation. The augmentation branch introduces new latent variables or extends the state space.

Each branch produces a candidate repair, and their outcomes are compared based on reduction in defect magnitude and improvement in predictive stability.

9.4 Ontological Growth and Promotion Criteria

A novel defect leads to ontological growth only if certain criteria are met. The primary requirement is recurrence: a single instance of a defect may be due to noise, whereas repeated occurrences across contexts indicate a structural deficiency.

The second requirement is explanatory compression. Let \mathcal{L}_{old} denote the description length of the atlas before extension, and \mathcal{L}_{new} after introducing a candidate latent variable. Promotion occurs only if

$$\mathcal{L}_{\text{new}} < \mathcal{L}_{\text{old}},$$

taking into account both the complexity of the new variable and the reduction in residuals across multiple cases. Ontological growth is thus governed by a balance between expressivity and simplicity.

9.5 Learning the Gluing Category

As the obstruction log grows, the system effectively learns the category in which its reconstruction operates. Objects correspond to classes of admissible states, while morphisms correspond to repair operations that transform incompatible configurations into compatible ones.

This perspective reveals that learning is not merely parameter adjustment but structural refinement of the category itself. The system expands its repertoire of objects and morphisms in response to persistent obstructions, thereby increasing its capacity to achieve closure. The learned category is constrained by the requirement of composability: repairs must not only resolve local defects but also integrate coherently with existing structures.

10 Defensive Inference and Sensor Integrity

10.1 Malignant Novelty and Adversarial Structure

Not all obstructions correspond to genuine features of the underlying world-state. Some arise from sensor degradation, calibration drift, stochastic noise, or adversarial manipulation. These phenomena generate defects that mimic structural inconsistency while lacking physical grounding.

We define *malignant novelty* as a class of obstructions whose persistence is not attributable to the RSVP field but to corruption in the observation process. Unlike genuine novelty, which reflects a deficiency in the ontology, malignant novelty reflects a deficiency in the measurement channel.

10.2 The Defensive Branch: Corruption Modeling

To address malignant novelty, the TARTAN architecture introduces a defensive branch that constructs explicit models of sensor corruption. Let Π_i denote a projection operator associated with a sensor. The corrupted observation is expressed as

$$y_i^{\text{obs}} = \Pi_i(X) + \eta_i,$$

where η_i represents noise or distortion. The defensive branch seeks a transformation \mathcal{K}_i such that

$$\tilde{y}_i = \mathcal{K}_i(y_i^{\text{obs}})$$

restores compatibility with the underlying state.

10.3 Cross-Context Transportability and Modality Ablation

A key diagnostic for distinguishing genuine from malignant novelty is cross-context transportability. A real feature of the world should manifest across multiple observational pathways, whereas a sensor-specific artifact should not.

For a candidate defect, the system evaluates its persistence under modality ablation. Let \mathcal{F}_{-i} denote the feasible set constructed without modality i . If the defect persists in \mathcal{F}_{-i} , it is likely genuine; if it disappears, it is attributed to the excluded modality.

10.4 Causal Reciprocity and the Rejection of Explanation Sinks

A genuine extension of the ontology must participate in the dynamics of the RSVP fields. Malignant novelty often produces what may be termed an *explanation sink*—a

variable that absorbs residual discrepancies without generating reciprocal effects in adjacent regions or future states.

To prevent such degenerate solutions, the system enforces a reciprocity condition: a candidate latent variable Z is admissible only if its inclusion induces consistent updates in neighboring tiles and propagates through the dynamics in a predictable manner.

10.5 Global Coherence as Decision Principle

The competition between ontological expansion and corruption modeling is resolved through a global description length criterion. Let \mathcal{M}_{exp} denote the model with an expanded ontology and \mathcal{M}_{def} the model with an augmented corruption model. Each model is evaluated in terms of its total description length, combining model complexity with residual error. The system selects the model minimizing this quantity, ensuring that decisions are not made based on local fit alone.

11 Sensor Lifecycle and Re-Integration

11.1 Quarantine and Provenance Tracking

When a modality exhibits persistent malignant novelty, it is placed into a state of quarantine. In this regime, the sensor’s outputs are decoupled from the primary reconstruction pipeline and treated as auxiliary data subject to analysis rather than enforcement.

Formally, quarantine replaces Π_i with a suppressed operator $\tilde{\Pi}_i$ whose contribution to the constraint functional is attenuated by a reliability weight $w_i \approx 0$. The sensor is thus prevented from influencing closure while its behavior is studied.

11.2 Corrected Projection Operators

The objective of quarantine is to learn a transformation \mathcal{K}_i that maps corrupted observations into a space compatible with the underlying state. The corrected projection operator is defined as

$$\Pi_i^{\text{corr}}(X) = \mathcal{K}_i(\Pi_i(X)).$$

The introduction of Π_i^{corr} transforms the sensor from an unreliable input into a modeled component of the system.

11.3 Gradual Re-Weighting and Probation

Re-integration is performed gradually through a controlled increase in the reliability weight w_i . At each iteration, the system evaluates the marginal contribution of the sensor to closure. Let $\mathcal{D}_{\text{with}}$ and $\mathcal{D}_{\text{without}}$ denote the aggregate defect with and without the sensor’s contribution. If

$$\mathcal{D}_{\text{with}} < \mathcal{D}_{\text{without}},$$

the sensor is deemed beneficial and w_i is increased. Full re-integration occurs only when w_i approaches parity with trusted modalities.

11.4 Temporal Monitoring and Adaptive Reliability

Sensor behavior is rarely stationary. Even after successful re-integration, the system continues to monitor the sensor’s performance over time. The reliability weight w_i becomes a dynamic variable updated according to ongoing evidence. Persistent agreement with other modalities leads to an increase in w_i , while reappearance of defects triggers its reduction.

12 CLIO and Recursive Optimization

12.1 The Consistency Operator as a Functor

The reconstruction process in TARTAN is governed by a consistency operator \mathcal{C} acting on the space of admissible configurations. We interpret it as a functor acting on a category whose objects are admissible states and whose morphisms are transformations preserving constraint structure:

$$\mathcal{C} : \mathcal{A} \rightarrow \mathcal{A}.$$

This functorial perspective reveals that reconstruction is not merely a search for a fixed point but a structured transformation of the entire space of configurations. The operator \mathcal{C} may be interpreted concretely as a discretized gradient flow or proximal map induced by the energy functional \mathcal{E} , anchoring it simultaneously in variational analysis and category-theoretic structure.

12.2 CLIO: Cognitive Loop via In-Situ Optimization

The CLIO module extends the consistency operator into a self-optimizing process. At each iteration, the system evaluates the effectiveness of its reconstruction in reducing obstruction and maintaining admissibility, feeding back into the definition of \mathcal{C} .

Formally, we define a meta-operator \mathcal{C}_θ parameterized by θ . The CLIO loop updates θ according to a functional

$$\theta_{t+1} = \Theta(\theta_t, X_t, \check{\delta}_t),$$

where X_t is the current state and $\check{\delta}_t$ the aggregate defect. This process constitutes an in-situ optimization in which the system refines its own inference mechanism while performing inference. The boundary between model and optimizer dissolves, yielding a recursive architecture in which operator adaptation is constrained by the same variational and sheaf-theoretic structure that governs state reconstruction.

Recent independent work by Cheng, Broadbent, and Chappell [1] introduces a parallel realization of this idea under the same name. In their formulation, large language models dynamically self-formulate approaches to scientific problems, adapt behavior when self-confidence is low, and converge to a final belief state through iterative internal reconfiguration. Their system demonstrates that, without additional post-training, adaptive cognitive loops substantially improve accuracy on demanding scientific benchmarks, surpassing fixed reasoning strategies by a wide margin.

The convergence between the two formulations is not superficial. In their framework, reasoning trajectories are represented as graph-structured belief states, uncertainty levels are monitored throughout, and corrections can be injected by external observers—all of which correspond, in the TARTAN framework, to local sections, defect tensors, and repair morphisms respectively. The present monograph provides the structural foundations that ground this empirical behavior: variational descent on the energy functional \mathcal{E} , sheaf-theoretic compatibility enforced across overlaps, and explicit obstruction theory that classifies the failures requiring in-situ repair.

A notable empirical finding in [1] is that oscillations within internal uncertainty measures are key determinants of accuracy. Within the present framework, this observation admits a precise geometric interpretation: persistent oscillations indicate proximity to a nontrivial cohomological obstruction class in $H^1(\mathcal{U}, \mathcal{S})$, not mere noise. The system is navigating near the boundary of the feasible set \mathcal{F} , where the gradient of \mathcal{E} has near-zero components in directions that require structural repair rather than simple descent. This yields a dynamical classification of inference regimes: monotone decay of $\|\check{\delta}_t\|$ indicates trivial obstruction and direct convergence; oscillatory behavior near a finite plateau indicates structural tension requiring gauge repair or refinement; divergence signals projection degeneracy or model mismatch. What appears in [1] as an empirical regularity governing when CLIO succeeds is thus, in the present framework, a consequence of the cohomological geometry of the admissible space.

12.3 Fixed Points and Stability of the CLIO Loop

The combined system of reconstruction and self-optimization admits fixed points at two levels. At the state level, $\mathcal{C}_\theta(X^*) = X^*$. At the operator level, $\Theta(\theta^*, X^*, 0) = \theta^*$.

Theorem 2. *Under suitable regularity and convexity conditions, the joint system (X_t, θ_t) converges to a pair (X^*, θ^*) such that X^* is a fixed point of \mathcal{C}_{θ^*} and θ^* is stable under the CLIO update.*

Theorem 3 (Unified Inference Principle). *Inference in a constraint-structured system corresponds to the joint convergence of a state X and a consistency operator \mathcal{C} such that*

$$\mathcal{C}(X) = X \quad \text{and} \quad \mathcal{C} \text{ is optimal with respect to a global coherence functional.}$$

This theorem encapsulates the essence of the TARTAN architecture. Inference is not merely the identification of a state but the co-evolution of the state and the mechanism that produces it. Unlike purely procedural formulations of cognitive loops, the present framework embeds in-situ optimization within a variational and sheaf-theoretic structure, ensuring that operator adaptation preserves global coherence rather than merely improving local reasoning performance.

13 Implementation and Systems Design

13.1 Core Data Structures

The TARTAN architecture is instantiated through a collection of interdependent data structures. The primary object is the tile, which represents a region $U_i \subset \Omega$ together with its associated local state s_i and admissibility certificate. Each tile stores the fields $(\Phi_i, \mathbf{v}_i, S_i)$ restricted to its domain, along with metadata describing resolution level, temporal extent, and provenance.

Overlaps between tiles are represented explicitly as objects that store references to the participating tiles and the corresponding restriction maps. Each overlap maintains the defect tensor $\check{\delta}_{ij}$ and its coarse-grained summary. The obstruction log is implemented as an indexed repository of summary representations and associated repair morphisms.

13.2 Algorithmic Pipeline

The reconstruction process proceeds as an iterative pipeline that alternates between local inference and global reconciliation. At initialization, the system constructs an initial tiling and assigns provisional states based on available observations.

In each iteration, local updates reduce projection and dynamical residuals within each tile. The system then evaluates overlaps to compute defect tensors, classifies them according to their behavior under normalization and refinement, and applies appropriate repair operations. Convergence is assessed by monitoring the aggregate defect:

$$\mathcal{D}_{\text{total}} = \sum_{i,j} \|\check{\delta}_{ij}\|^2.$$

13.3 Computational Complexity and Scaling

The computational complexity of TARTAN is governed by the number of tiles and the cost of overlap evaluation. Let N denote the number of tiles and k the average number of neighbors per tile. The cost of a single iteration is approximately

$$\mathcal{O}(N \cdot C_{\text{local}} + Nk \cdot C_{\text{overlap}}),$$

where C_{local} and C_{overlap} denote the costs of local updates and overlap computations. Recursive tiling introduces refinement that is applied selectively, allowing sublinear scaling with respect to domain resolution.

13.4 Parallelization and Distributed Execution

The TARTAN architecture is inherently parallelizable. Local inference within each tile can be performed independently, while overlap computations involve only neighboring tiles. In a distributed setting, tiles are assigned to nodes in a computational network, with synchronization required only for overlapping regions.

Consistency across the distributed system is maintained through iterative synchronization of overlaps and propagation of repair operations. Convergence of the global state emerges from the coordination of local computations.

14 Failure Modes as First-Class Structure

A complete theory of inference must account not only for the conditions of success but for the conditions and structure of failure. The TARTAN framework distinguishes five canonical failure regimes, each with a formal characterization, a detection criterion, and a prescribed response. Making these regimes explicit is not merely pedagogical; it is a design constraint. A system that cannot detect its own failure modes cannot safely expand its ontology or trust its repairs.

14.1 Projection Degeneracy

The first failure mode arises when the induced projection map $\tilde{\Pi}(X) = (\Pi_i(X))_{i \in I}$ is not injective on the feasible set \mathcal{F} . In this case, distinct states $X_1 \neq X_2$ satisfy $\Pi_i(X_1) = \Pi_i(X_2)$ for all i , so observational data cannot distinguish between them.

Proposition 2 (Projection Degeneracy). *If $\tilde{\Pi}$ is not injective on \mathcal{F} , then no observation-consistent procedure can uniquely reconstruct X^* . The system cannot achieve closure in the strict sense; it can only identify an equivalence class of states.*

Detection: projection degeneracy manifests as a failure of $\mathcal{D}_{\text{total}}$ to decrease monotonically, combined with sensitivity of the reconstructed state to initialization. The system must then either broaden its projection family (adding new modalities or observation operators) or accept a weaker notion of closure defined on equivalence classes under the degeneracy group.

14.2 Summary Collapse

The second failure mode concerns the coarse-graining operator σ . If σ discards information that distinguishes compatible from incompatible sections, the system will classify genuine obstructions as resolved and fail to trigger appropriate repair.

Definition 4 (Obstruction-Complete Summary). *A coarse-graining operator σ is obstruction-complete if $\check{\delta}_{ij} = 0$ implies $\sigma(\gamma_i) = \sigma(\gamma_j)$, and $\sigma(\gamma_i) = \sigma(\gamma_j)$ implies $\|\check{\delta}_{ij}\| \lesssim \epsilon$ for a controlled tolerance ϵ .*

Summary collapse occurs when σ fails the second condition: incompatible sections produce indistinguishable summaries, so the defect tensor is systematically underestimated. Detection: compare defect computed from full sections against defect estimated from summaries. A persistent gap between these quantities indicates collapse of the summary representation. Response: enrich σ by increasing the dimension of Σ or by selecting invariants with finer sensitivity near the boundary of \mathcal{F} .

14.3 Ontological Overgrowth

The third failure mode is the inverse of the previous one. When the state extension mechanism is triggered too aggressively—accepting new latent variables that reduce local residuals without satisfying the causal reciprocity or description-length criteria—the admissible space \mathcal{A} expands beyond what is justified by the evidence.

Proposition 3 (Overgrowth Condition). *Let Z_1, Z_2, \dots be a sequence of latent variables introduced by successive state extensions. The system exhibits ontological*

overgrowth if the description length $\mathcal{L}(\mathcal{A}_n)$ grows without a corresponding decrease in aggregate residual, i.e., if $\mathcal{L}(\mathcal{A}_{n+1}) > \mathcal{L}(\mathcal{A}_n)$ while $\mathcal{D}_{total}^{(n+1)} \geq \mathcal{D}_{total}^{(n)}$.

Detection: monitor the ratio $\Delta\mathcal{L}/\Delta\mathcal{D}$ at each extension step. A ratio below a prescribed threshold triggers a rollback, removing the most recently introduced variable and reclassifying the associated defect as a sensor pathology candidate or persistent structural obstruction requiring a different repair branch.

14.4 Defensive Overfitting

The fourth failure mode is specific to the defensive inference branch. When the corruption model \mathcal{K}_i is given excessive capacity, it can absorb genuine structural obstructions by modeling them as sensor artifacts. The result is a system that attributes real features of the world to noise.

Detection: a candidate defect is genuine rather than malignant if it persists under modality ablation—that is, if removing sensor i from the reconstruction does not eliminate the defect. If the defensive branch classifies a defect as malignant but modality ablation shows it persists, the defensive model is overfitting. The correction operator \mathcal{K}_i must be regularized or rolled back, and the defect must be redirected to the expansive branch.

14.5 Oscillatory Non-Convergence

The fifth failure mode concerns the CLIO optimization loop. Under certain configurations of \mathcal{C}_θ and the update rule Θ , the joint system (X_t, θ_t) may enter a limit cycle rather than converging to a fixed point.

Proposition 4 (Oscillatory Regime). *If the update rule Θ is not contractive in θ near a fixed point, the CLIO loop may exhibit oscillations in θ_t that prevent convergence of the state sequence X_t .*

As established in Section 12 and formally reinforced in Appendix F, persistent oscillations indicate proximity to a nontrivial obstruction class rather than simple numerical instability. The appropriate response is not to reduce the learning rate but to apply a structural repair: gauge alignment, refinement, or state extension, depending on the class of the obstructions driving the oscillation. Once the obstruction class is eliminated, the oscillation resolves and the CLIO loop resumes monotone convergence.

14.6 Failure Mode Summary

These five failure modes are mutually exclusive in their immediate cause but may co-occur in practice. The system’s response in each case is determined by the

same decision procedure that governs ordinary repair: classify the failure by its behavior under normalization, ablation, and description-length analysis, then apply the minimal intervention that restores the conditions required for closure. Failure is not the end of inference but one of its principal inputs.

15 Epistemology and Philosophy of Inference

15.1 Inference as Physical Process

The TARTAN architecture compels a redefinition of inference as a physical process rather than an abstract symbolic operation. The RSVP fields (Φ, \mathbf{v}, S) encode not only the state of the world but the constraints under which that state can evolve. Inference becomes the identification of a configuration that satisfies these constraints, corresponding to a state of reduced inconsistency—reduced effective entropy within the admissible space.

This perspective dissolves the boundary between computation and physics. The act of inference is itself a dynamical process, in which the system evolves toward a configuration of maximal coherence.

15.2 Closure, Identity, and Irreversibility

Closure is not merely a computational endpoint but a structural property of the reconstructed state. A state that achieves closure is one in which all constraints are satisfied simultaneously, and no further modifications are required to reconcile observations.

This property is inherently tied to irreversibility. The trajectory leading to closure involves the elimination of degrees of freedom and the accumulation of constraints. Identity, in this context, is defined by the persistence of a trajectory under constraint. A system maintains its identity not by preserving a static pattern but by sustaining a coherent evolution within the admissible manifold.

15.3 Ontological Humility and Adaptive Knowledge

The TARTAN framework embodies a principle of ontological humility. The system does not assume that its initial representation is sufficient to capture all aspects of the world. Instead, it treats persistent obstruction as evidence that the current ontology is incomplete.

This leads to an adaptive conception of knowledge. The system expands its ontology only when required by consistent, cross-context evidence, and it does so in a manner that minimizes complexity while maximizing coherence. Such an approach avoids both overfitting and rigidity.

15.4 Knowledge as Global Consistency

In this framework, knowledge is not a collection of representations but the achievement of global consistency. A system knows the world to the extent that it can construct a state that satisfies all constraints imposed by observation and dynamics.

The obstruction log and the CLIO mechanism ensure that this process is cumulative. The system builds a repertoire of structures and transformations that enable it to achieve closure more efficiently over time. Knowledge is thus both a state and a process.

15.5 Synthesis: A Theory of Inference

The TARTAN architecture integrates mathematical formalism, computational design, and philosophical insight into a unified theory of inference. It redefines inference as the process of achieving closure in a constraint-structured space, guided by physical principles and adaptive learning.

The system operates across multiple levels: from local reconstruction to global consistency, from static states to dynamic trajectories, and from fixed ontologies to adaptive structures. Each level is governed by the same fundamental principle: the resolution of obstruction through minimal, coherence-preserving transformation.

In this sense, TARTAN is not merely an architecture but a paradigm. It offers a framework for understanding inference, knowledge, and adaptation as aspects of a single process grounded in the dynamics of constraint and consistency. An intelligent system, on this view, is not one that predicts accurately in isolation, but one that can construct and maintain a coherent model of the world—adapting its structure as necessary to achieve closure.

Appendices

A Formal Proofs

A.1 Identifiability Under Projection and Convexity

Definition 5. Let \mathcal{A} be a convex subset of a Banach space and let $\{\Pi_i\}_{i \in I}$ be a family of bounded linear operators $\Pi_i : \mathcal{A} \rightarrow \mathcal{Y}_i$. Define the feasible set

$$\mathcal{F} = \{X \in \mathcal{A} \mid \|\Pi_i(X) - y_i\| \leq \epsilon_i \ \forall i \in I\}.$$

Theorem 4 (Identifiability). Assume that: (1) \mathcal{A} is strictly convex; (2) the induced map $\tilde{\Pi} : \mathcal{A} \rightarrow \prod_i \mathcal{Y}_i$ defined by $\tilde{\Pi}(X) = (\Pi_i(X))$ is injective on \mathcal{F} ; (3) the regularization functional \mathcal{D}_{RSVP} is strictly convex on \mathcal{F} . Then there exists a unique minimizer $X^* \in \mathcal{F}$ of the energy functional $\mathcal{E}(X)$.

Proof. Strict convexity of \mathcal{A} and \mathcal{D}_{RSVP} implies that $\mathcal{E}(X)$ is strictly convex on \mathcal{F} . Therefore, any minimizer is unique if it exists. Existence follows from coercivity of \mathcal{E} under standard assumptions on \mathcal{D}_{RSVP} and boundedness of \mathcal{F} . Injectivity of $\tilde{\Pi}$ ensures that no two distinct elements of \mathcal{F} produce identical projections, preventing degeneracy. \square

A.2 Failure Modes

Proposition 5. If $\tilde{\Pi}$ is not injective on \mathcal{F} , then there exist distinct states $X_1, X_2 \in \mathcal{F}$ with $\Pi_i(X_1) = \Pi_i(X_2)$ for all i , implying non-identifiability.

Proposition 6. If \mathcal{D}_{RSVP} is not strictly convex on \mathcal{F} , then \mathcal{E} may admit multiple minimizers even when $\tilde{\Pi}$ is injective.

A.3 Equivalence of Sheaf Gluing and Variational Closure

Theorem 5. A family of local sections $\{s_i\}$ admits a global section $X \in \mathcal{S}(\Omega)$ if and only if the defect tensor $\check{\delta}_{ij} = 0$ for all overlaps and X is a fixed point of the consistency operator \mathcal{C} .

Proof. If $\check{\delta}_{ij} = 0$, the sections satisfy the gluing condition $\rho_{ij}(s_i) = \rho_{ji}(s_j)$, implying the existence of a global section X . Conversely, if a global section exists, its restrictions define compatible local sections, and the defect tensor vanishes. The equivalence with $\mathcal{C}(X) = X$ follows from the definition of \mathcal{C} as enforcing both observational and dynamical consistency. \square

A.4 Convergence of the Staged Closure Protocol

Theorem 6. *Assume that each repair operation reduces a Lyapunov functional \mathcal{L} , that \mathcal{L} is bounded below, and that refinement and state extension increase expressive capacity. Then the staged closure protocol converges to a configuration with vanishing defect, or to a limit in which no further reduction is possible.*

Proof. Monotonic decrease of \mathcal{L} and boundedness imply convergence of the sequence $\{\mathcal{L}_t\}$. If defect remains nonzero, it must lie in a subspace not addressable by current operations, triggering state extension. Under the assumption that extension increases capacity, the system eventually reaches a regime where defects can be reduced. \square

A.5 Stability of CLIO Fixed Points

Theorem 7. *Let \mathcal{C}_θ be Lipschitz continuous in both X and θ , and let the update rule Θ be contractive in θ near a fixed point. Then the coupled system*

$$X_{t+1} = \mathcal{C}_{\theta_t}(X_t), \quad \theta_{t+1} = \Theta(\theta_t, X_t, \check{\delta}_t)$$

converges locally to a stable fixed point (X^, θ^*) .*

Proof. Standard results for coupled contractive systems apply. The Lipschitz condition ensures stability of the X -update, while contraction in θ ensures convergence of the operator parameters. Joint convergence follows from composition. \square

B Functional Analytic Foundations

B.1 Function Spaces and Admissible States

Let $\Omega \subset \mathbb{R}^d$ be a bounded domain with Lipschitz boundary. The RSVP state is defined as $X = (\Phi, \mathbf{v}, S)$ with components belonging to appropriate Sobolev spaces:

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

The admissible space is equipped with the norm

$$\|X\|_{\mathcal{A}}^2 = \|\Phi\|_{H^1}^2 + \|\mathbf{v}\|_{H^1}^2 + \|S\|_{L^2}^2,$$

making it a Hilbert space.

B.2 Weak Formulation of RSVP Dynamics

The RSVP equations are interpreted in weak form. For test functions $\psi \in H^1(\Omega)$ and $\varphi \in H^1(\Omega; \mathbb{R}^d)$:

$$\begin{aligned} \int_{\Omega} \partial_t \Phi \psi \, dx + \int_{\Omega} (\Phi \mathbf{v}) \cdot \nabla \psi \, dx + D_{\Phi} \int_{\Omega} \nabla \Phi \cdot \nabla \psi \, dx &= 0, \\ \int_{\Omega} \partial_t \mathbf{v} \cdot \varphi \, dx + \int_{\Omega} (\mathbf{v} \cdot \nabla) \mathbf{v} \cdot \varphi \, dx + \int_{\Omega} \nabla \Phi \cdot \varphi \, dx + \gamma \int_{\Omega} \mathbf{v} \cdot \varphi \, dx &= 0, \\ \int_{\Omega} \partial_t S \psi \, dx + \int_{\Omega} (S \mathbf{v}) \cdot \nabla \psi \, dx - D_S \int_{\Omega} \nabla S \cdot \nabla \psi \, dx &= \int_{\Omega} \sigma(X) \psi \, dx. \end{aligned}$$

Theorem 8. *Assume initial data $X_0 \in \mathcal{A}$ and that nonlinear terms satisfy appropriate growth and Lipschitz conditions. Then there exists a weak solution $X \in L^2(0, T; \mathcal{A}) \cap C([0, T]; L^2(\Omega))$ to the RSVP system.*

Proof. The proof follows by Galerkin approximation. One constructs finite-dimensional approximations using basis functions in $H^1(\Omega)$ and derives uniform bounds via energy estimates. Compactness arguments, specifically the Aubin–Lions lemma, yield convergence of subsequences to a weak solution. \square

C Formal Equivalence Between Yarncrawler Closure and Sheaf Descent

C.1 Statement

A central claim of the present framework is that the Yarncrawler notion of closure and the sheaf-theoretic notion of descent are formally equivalent under suitable hypotheses. The Yarncrawler perspective formulates inference as the search for a unique fixed point of a consistency operator $\mathcal{C} : \mathcal{A} \rightarrow \mathcal{A}$. The sheaf-theoretic perspective formulates inference as the existence of a unique global section assembled from compatible local sections.

C.2 Equivalence Theorem

Theorem 9 (Yarncrawler Closure = Sheaf Descent). *Assume: (1) \mathcal{A} is strictly convex; (2) the induced projection map is injective on \mathcal{F} ; (3) \mathcal{S} is a sheaf of admissible local states over Ω ; (4) every state $X \in \mathcal{F}$ restricts to a family of local sections in \mathcal{S} ; (5) every compatible family of local sections in \mathcal{S} glues uniquely to a state in \mathcal{F} ; (6) the consistency operator \mathcal{C} has fixed points exactly at globally compatible admissible states.*

Then the following are equivalent: (a) there exists a unique fixed point $X^* \in \mathcal{F}$ of \mathcal{C} ; (b) there exists a unique global section $s^* \in \mathcal{S}(\Omega)$ descending from a compatible family $\{s_i^*\}$.

Proof. If X^* is a unique fixed point, its restrictions $s_i^* = X^*|_{U_i}$ define a compatible family in \mathcal{S} . Since \mathcal{S} is a sheaf, this family descends to a unique global section. Conversely, a unique global section corresponds by hypothesis to a unique state $X^* \in \mathcal{F}$, which is globally consistent and therefore a fixed point of \mathcal{C} . \square

C.3 Corollary

Corollary 1. *Gauge repair, recursive refinement, and state extension are exactly the three primitive ways of restoring descent when the equivalence theorem fails constructively.*

Proof. Gauge repair alters representatives without changing the underlying object, converting apparent non-descent into descent by normalization. Recursive refinement changes the cover so that compatibility can be achieved at a finer level. State extension enlarges the sheaf or admissible space so that a global section exists in the extended category. \square

C.4 Functorial Form

Let $\text{Fix}(\mathcal{C})$ denote the subspace of fixed points of the consistency operator, and $\text{Desc}(\mathcal{S}, \mathcal{U})$ the space of descent data modulo compatibility. Then under the hypotheses above there is a canonical isomorphism

$$\text{Fix}(\mathcal{C}) \cong \text{Desc}(\mathcal{S}, \mathcal{U}).$$

This identifies reconstruction as a descent problem and descent as fixed-point closure. TARTAN's overlap machinery computes cocycles, its repair logic restores descent, and its consistency operator searches for fixed points. These are not separate heuristics but equivalent manifestations of one underlying structure.

D Category-Theoretic Formalization

D.1 The Category of Local Reconstructions

We construct a category \mathcal{C} whose objects are pairs (U_i, s_i) with $s_i \in \mathcal{S}(U_i)$, and whose morphisms are restriction-compatible maps induced by inclusion. The assignment $U \mapsto \mathcal{S}(U)$ defines a presheaf whose restriction maps satisfy functoriality:

$$\rho_{UW} = \rho_{VW} \circ \rho_{UV} \quad \text{for} \quad W \subseteq V \subseteq U.$$

Definition 6. *The presheaf \mathcal{S} is a sheaf if for any open cover $\{U_i\}$ and compatible family $\{s_i \in \mathcal{S}(U_i)\}$, there exists a unique $X \in \mathcal{S}(\Omega)$ such that $\rho_i(X) = s_i$.*

Failure of the sheaf condition corresponds precisely to the presence of obstruction.

D.2 Čech Cohomology and Obstruction Classes

The defect tensor $\check{\delta}_{ij}$ defines a 1-cochain. If it is not a coboundary, it represents a nontrivial cohomology class $[\check{\delta}] \in H^1(\mathcal{U}, \mathcal{S})$. Vanishing of this class is equivalent to the existence of a global section. Covers form a directed system under refinement, and the cohomology stabilizes:

$$H^1(\Omega, \mathcal{S}) = \varinjlim_{\mathcal{U}} H^1(\mathcal{U}, \mathcal{S}).$$

D.3 Learning as Category Expansion

The obstruction log and ontological growth can be interpreted categorically as an expansion of the category \mathcal{C} . New objects and morphisms are added to resolve previously unrepresentable configurations. This expansion is guided by a minimality principle: the category is enlarged only as much as necessary to restore closure, resulting in a dynamically evolving category that adapts to the structure of the world.

E Haskell DSL Sketch for TARTAN

E.1 Motivation

The TARTAN architecture is naturally compositional. Tiles, overlaps, summaries, repairs, and consistency operators all behave like typed morphisms between structured objects, making Haskell an especially suitable implementation language.

E.2 Core Typed Objects

```

newtype TileId   = TileId   Int deriving (Show, Eq, Ord)
newtype OverlapId = OverlapId Int deriving (Show, Eq, Ord)

data Tile a = Tile
  { tileId   :: TileId
  , tileDom  :: (Double, Double)
  , tileData :: a
  } deriving (Show, Eq)

```

```

data Overlap a = Overlap
  { overlapId    :: OverlapId
  , leftTile     :: TileId
  , rightTile    :: TileId
  , overlapData  :: a
  } deriving (Show, Eq)

data Defect = Defect
  { magnitude :: Double
  } deriving (Show, Eq)

```

E.3 Projection and Restriction Typeclasses

```

class Projective a y where
  project :: a -> y

class Restrictable a where
  restrict :: (Double, Double) -> a -> a

```

E.4 Repair Morphisms as Typed Transformations

```

data Repair
  = GaugeRepair
  | RefineTile TileId
  | ExtendState
  deriving (Show)

chooseRepair :: Defect -> Repair
chooseRepair d
  | magnitude d < 1.0e-3 = GaugeRepair
  | magnitude d < 1.0    = RefineTile (TileId 0)
  | otherwise            = ExtendState

```

E.5 A Minimal TARTAN EDSL

```

data TartanF next
  = InferLocal TileId (Double -> next)
  | ComputeOverlap OverlapId (Defect -> next)
  | ApplyRepair Repair next
  | LogObstruction Double next

```

The free monad over this signature yields an executable reconstruction language:

```
-- type TartanDSL = Free TartanF
```

This allows reconstruction pipelines to be described declaratively and then interpreted into concrete execution strategies, simulations, or formal proofs of properties. The categorical and variational structure does not obstruct implementation—it guides it.

F Numerical Implementation: RSVP–TARTAN Lattice

F.1 Grid and Field Types

The numerical realization of TARTAN uses a regular 2D grid as the spatial domain. The three RSVP fields are represented as dense arrays of double-precision floating-point values.

```
data Grid = Grid { gWidth :: !Int, gHeight :: !Int }
  deriving (Show, Eq)

newtype ScalarField = ScalarField
  { unScalar :: Vector Double } deriving (Show, Eq)

data VectorField = VectorField
  { vxField :: !ScalarField
  , vyField :: !ScalarField } deriving (Show, Eq)

data RSVP = RSVP
  { phiField :: !ScalarField
  , velField :: !VectorField
  , entField :: !ScalarField } deriving (Show, Eq)
```

F.2 RSVP Evolution

The evolution step implements finite-difference approximations of the coupled PDE system:

$$\begin{aligned} \partial_t \Phi + \partial_x(\Phi v) &= D_\Phi \nabla^2 \Phi, \\ \partial_t \mathbf{v} + (\mathbf{v} \cdot \nabla) \mathbf{v} &= -\nabla \Phi - \gamma \mathbf{v} + D_v \nabla^2 \mathbf{v}, \\ \partial_t S + \nabla \cdot (S \mathbf{v}) &= \kappa |\nabla \mathbf{v}|^2 + D_S \nabla^2 S. \end{aligned}$$

```
data Params = Params
```

```

{ dt :: !Double, dPhi :: !Double
, dVel :: !Double, dEnt :: !Double
, gamma :: !Double, kappa :: !Double
}

defaultParams :: Params
defaultParams = Params
  { dt = 0.05, dPhi = 0.08, dVel = 0.05
  , dEnt = 0.04, gamma = 0.08, kappa = 0.2 }

```

F.3 Pairwise Overlap Repair

The core repair operator computes the defect between neighboring tiles and applies equal-and-opposite corrections, implementing the cohomological cancellation described in the main text:

```

repairSweep :: Double -> Grid -> [Tile] -> [Overlap] -> [
  Tile]
repairSweep alpha gr tiles overlaps =
  foldl' step tiles overlaps
  where
    step ts ov =
      let t1 = ts !! ovLeft ov
          t2 = ts !! ovRight ov
          (t1', t2') = repairOverlap alpha gr t1 t2 (ovRect
            ov)
      in updateTiles ts (ovLeft ov) t1' (ovRight ov) t2'

```

F.4 Main Iteration Loop

```

iterateExperiment :: Params -> Grid -> [Overlap]
  -> Int -> [Tile] -> IO [Tile]
iterateExperiment p gr ovs0 steps ts0 = go 0 ovs0 ts0
  where
    go n ovs ts
      | n >= steps = pure ts
      | otherwise = do
          let ds = allDefects gr ts ovs
              td = totalDefect ds
          printf "iteration %02d | tiles = %2d | defect =
            %.6f\n"

```

```

        n (length ts) td
    let ts1 = tartanStep p 0.25 gr ts ovs
        ts2 = maybeRetile gr td ts1
        ovs1 = buildOverlaps gr 2 ts2
    go (n + 1) ovs1 ts2

```

F.5 Visualization

The implementation exports field snapshots as PPM images, allowing direct visual inspection of whether closure is occurring. The Φ field should begin fragmented along tile boundaries and converge toward a globally smooth configuration. Persistent seams indicate failure of closure; their elimination is the primary convergence diagnostic.

```

scalarToPPM :: Grid -> ScalarField -> FilePath -> IO ()
scalarToPPM gr@(Grid w h) (ScalarField v) path = do
    let nv = normalizeField v
        pixel x y =
            let val = nv V.! index2D gr x y
                c = floor (255 * val) :: Int
            in (c, c, c)
    writePPM path w h pixel

```

F.6 Interpretation

The numerical experiment demonstrates that the core mechanisms of TARTAN—local reconstruction, overlap consistency, staged repair, and trajectory summarization—are implementable in concrete systems. The retiling trigger at high defect values implements the adaptive refinement described in Section 3. As the aggregate defect $\mathcal{D}_{\text{total}}$ decreases, the system achieves genuine closure rather than mere averaging, validating the claim that inference is properly understood as constraint satisfaction rather than prediction.

G CLIO as a Functor on Fibered Structures

G.1 Fibered Structure of Local Reasoning States

Let $\pi : \mathcal{E} \rightarrow \mathcal{B}$ be a fibered category, where \mathcal{B} is the base category of observational domains (tiles, modalities, or coordinate patches) and $\mathcal{E}_U = \pi^{-1}(U)$ is the category of local reasoning states over $U \in \mathcal{B}$. A reasoning configuration is a section $s : \mathcal{B} \rightarrow \mathcal{E}$ assigning to each U a local state $s(U) \in \mathcal{E}_U$. Compatibility across overlaps is encoded

by restriction functors

$$\rho_{UV} : \mathcal{E}_U \rightarrow \mathcal{E}_V,$$

which define the standard sheaf gluing conditions.

G.2 Defects as Čech 1-Cocycles

Given a cover $\mathcal{U} = \{U_i\}$, failure of compatibility defines a Čech 1-cochain $\check{\delta}_{ij} := \rho_{ij}(s_i) - \rho_{ji}(s_j)$. A nonvanishing cohomology class $[\check{\delta}] \in H^1(\mathcal{U}, \mathcal{S})$ represents a global obstruction to the existence of a consistent section.

G.3 CLIO as an Endofunctor

We define the CLIO operator of Cheng, Broadbent, and Chappell [1] as an endofunctor

$$\mathcal{F}_{\text{CLIO}} : \text{Sec}(\mathcal{E}) \rightarrow \text{Sec}(\mathcal{E}),$$

acting on the category of sections. For a given section s , the updated section $s' = \mathcal{F}_{\text{CLIO}}(s)$ is obtained by evaluating defect signals $\check{\delta}_{ij}$, performing local updates within each fiber \mathcal{E}_{U_i} , and re-propagating constraints across overlaps. Thus $\mathcal{F}_{\text{CLIO}}$ is a defect-reducing operator in precisely the sense that each application decreases $\|\check{\delta}\|$ monotonically under the conditions of the Lyapunov theorem established in Appendix A.

G.4 Functoriality and Projection Compatibility

The CLIO functor is required to commute approximately with projection:

$$\Pi_i \circ \mathcal{F}_{\text{CLIO}} \approx \mathcal{F}_{\text{CLIO}} \circ \Pi_i.$$

This ensures that updates performed at the global level correspond to coherent updates in each observational channel, preventing representation drift. Commutation is exact in the gauge-invariant setting and approximate up to $\mathcal{O}(\epsilon_K)$ after spectral truncation.

G.5 CLIO as an Approximate Closure Functor

We may therefore interpret the CLIO architecture of [1] as implementing an approximate functor from local reasoning states to globally coherent states, which iteratively reduces obstruction without explicit access to the underlying sheaf structure. The present framework provides the structural grounding that CLIO does not itself supply: identification of uncertainty with defect tensors, replacement of heuristic updates

with variational descent, embedding of operator adaptation within a structured admissible space, and conditions for global coherence via obstruction vanishing. CLIO, in other words, is an empirical realization of defect-driven constraint closure; the present monograph explains why it works.

G.6 Oscillatory Regimes and Obstruction Classes

Proposition 7 (Oscillatory Defect Regime). *Persistent oscillations in CLIO's internal uncertainty measures correspond to trajectories near nontrivial cohomological obstruction classes in $H^1(\mathcal{U}, \mathcal{S})$.*

Sketch. If $\check{\delta}_{ij}$ cannot be eliminated by a single application of \mathcal{C}_θ but decreases under alternating updates, the system remains near the boundary of the feasible set \mathcal{F} . This produces quasi-periodic residual dynamics, observed externally as oscillations in uncertainty. Such behavior indicates proximity to a nontrivial obstruction class rather than random noise, and persists until a gauge repair, refinement, or state extension successfully eliminates the class. \square

The finding in [1] that oscillations within internal uncertainty measures are key in determining accuracy is thus a consequence of the cohomological geometry of the admissible space, not an incidental regularity. Monotone decay of uncertainty signals trivial closure; oscillatory stabilization signals structural tension at an obstruction boundary; divergence signals projection degeneracy or model mismatch.

H Derived Stack Interpretation of TARTAN Covers

H.1 Covers as Atlases for Derived Stacks

Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a TARTAN cover of Ω , with each tile U_i equipped with a local state $s_i \in \mathcal{S}(U_i)$. We interpret \mathcal{U} as an atlas for a geometric object \mathfrak{X} , where each tile defines a chart and the overlap maps ρ_{ij} serve as transition functions. Unlike classical atlases, these transitions need not satisfy strict equality but only up to controlled defect: $\rho_{ij}(s_i) \simeq \rho_{ji}(s_j)$. Thus \mathfrak{X} is not a sheaf in the strict sense but a *derived stack* encoding higher-order compatibility data.

H.2 Global Reconstruction as a Homotopy Colimit

We model local sections as objects in a higher category \mathcal{C} , where morphisms encode admissible transformations and higher morphisms encode equivalences between repair paths. The diagram of local sections indexed by the Čech nerve of the cover

$$\{s_i\} \in \text{Fun}(\check{C}(\mathcal{U}), \mathcal{C}),$$

defines a globally consistent state via its homotopy colimit:

$$X \simeq \text{hocolim}_{\check{\mathcal{C}}(\mathcal{U})} s_i.$$

This construction accounts for gauge equivalences, refinement relations, and higher-order coherence conditions. TARTAN reconstruction is therefore naturally interpreted as homotopy colimit computation over a diagram of local states.

Proposition 8. *The homotopy colimit X is equivalent to a strict colimit if and only if the associated obstruction class vanishes: $[\check{\delta}] = 0 \iff X \simeq \text{colim } s_i$.*

Obstruction vanishing corresponds to classical gluing, while nontrivial classes require derived corrections encoded by the three-stage repair protocol.

H.3 Repair Operations as Homotopies

The TARTAN repair moves admit natural higher-categorical interpretations. Gauge repair corresponds to a natural transformation aligning parallel morphisms, $s_i \Rightarrow s'_i$. Refinement corresponds to replacement of U_i by a finer cover, inducing a refinement functor $\mathcal{C}(U_i) \rightarrow \mathcal{C}(\{U_{i\alpha}\})$. State extension corresponds to enlargement of the target category \mathcal{C} by adjoining new objects or morphisms. Each operation modifies the diagram so that its homotopy colimit becomes more coherent, and the three together exhaust the primitive repair moves in the sense of the Corollary established in Appendix C.

H.4 Interaction with CLIO

Under this interpretation, the CLIO functor [1] acts as a morphism on diagrams

$$\mathcal{F}_{\text{CLIO}} : \text{Fun}(\check{\mathcal{C}}(\mathcal{U}), \mathcal{C}) \rightarrow \text{Fun}(\check{\mathcal{C}}(\mathcal{U}), \mathcal{C}),$$

which deforms the diagram toward one whose homotopy colimit is stable. CLIO thus performs iterated deformation of the diagram until obstruction classes vanish or stabilize. Derived closure is achieved when the induced map on obstruction classes satisfies $\|[\check{\delta}^{(t+1)}]\| \leq \kappa \|[\check{\delta}^{(t)}]\|$ with $\kappa < 1$, which corresponds to convergence of the diagram to a homotopy-coherent limit serving as the reconstructed global state.

I Semantic Merge as Homotopy Colimit Computation

I.1 From Syntactic to Semantic Versioning

Classical version control systems operate on discrete artifacts—files, lines, commits—and define merge as a syntactic reconciliation of divergent edit histories. This

approach treats code as text rather than as a structured object with internal invariants, and as a consequence it has no mechanism for detecting or repairing semantic inconsistencies that arise when independent edits interact at the level of meaning rather than surface form.

We model a repository instead as a diagram $D : \mathcal{I} \rightarrow \mathcal{C}$, where \mathcal{I} is an index category encoding branches, commits, and dependencies, and \mathcal{C} is a category of semantic states equipped with projections and RSVP regularization. Each node $i \in \mathcal{I}$ corresponds to a local state X_i carrying an energy

$$\mathcal{E}(X_i) = \sum_k \|\Pi_k(X_i) - y_k\|^2 + \lambda \mathcal{R}(X_i),$$

and morphisms encode transformations between states.

I.2 Merge as Homotopy Colimit

A merge operation is defined as the homotopy colimit

$$X_{\text{merge}} \simeq \text{hocolim}_{\mathcal{I}} D,$$

which replaces line-based conflict resolution with identification of equivalences up to gauge, preservation of higher-order structure, and reconciliation via homotopy rather than overwrite. Conflicts in classical systems correspond to nontrivial obstruction classes in the diagram and are measured by a defect tensor $\check{\delta}_{ij} = \rho_{ij}(X_i) - \rho_{ji}(X_j)$ that quantifies incompatibility between branches. A merge conflict is therefore not an error condition but a signal that the diagram contains nontrivial cohomology: $[\check{\delta}] \neq 0$.

The TARTAN repair taxonomy induces a direct semantic merge protocol. If discrepancies vanish under reparameterization, branches are gauge-equivalent and merge is achieved by alignment. If discrepancies are high-frequency or local, the diagram is refined and merge is computed at finer granularity. If discrepancies persist under both operations, the state space is insufficient and merge requires introducing new variables or abstractions, enlarging \mathcal{C} .

I.3 CLIO as a Merge Engine

Within this framework, the CLIO operator [1] acts as an iterative solver for the homotopy colimit, driving the diagram toward coherence through the joint update

$$X_{t+1} = \mathcal{C}_{\theta_t}(X_t).$$

Uncertainty oscillations correspond to persistent merge obstructions, guiding the system toward gauge alignment, refinement, or ontological expansion. Because merge is defined via homotopy colimit, it is invariant under reparameterization of the diagram—eliminating dependence on arbitrary history ordering—and preserves all structural equivalences. The result is a merge operation that is order-independent, structure-preserving, and semantically grounded.

I.4 The General Principle

Version control is a special case of a more general problem: given a diagram of partially compatible local states, compute a globally coherent object respecting all structural constraints. Syntactic approaches approximate this problem by ignoring the constraint structure. The TARTAN–RSVP–CLIO stack solves it directly, treating conflicts as cohomological invariants, repairs as functorial transformations, and the merged result as a homotopy colimit rather than a textual superposition. This reframes software collaboration, model integration, and knowledge synthesis as instances of a single mathematical process: the construction of globally coherent structure from locally inconsistent data through defect-driven constraint closure.

References

- [1] Cheng, N., Broadbent, G., and Chappell, W. (2025). Cognitive loop via in-situ optimization: Self-adaptive reasoning for science. arXiv preprint arXiv:2508.02789. <https://arxiv.org/abs/2508.02789>
- [2] Aubin, J.-P. (1963). Un théorème de compacité. *Comptes Rendus de l'Académie des Sciences*, 256, 5042–5044.
- [3] Bott, R. and Tu, L. W. (1982). *Differential Forms in Algebraic Topology*. Springer, New York.
- [4] Bredon, G. E. (1997). *Sheaf Theory*, 2nd edition. Springer, New York.
- [5] Daubechies, I. (1992). *Ten Lectures on Wavelets*. SIAM, Philadelphia.
- [6] Evans, L. C. (2010). *Partial Differential Equations*, 2nd edition. American Mathematical Society, Providence.
- [7] Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- [8] Hartshorne, R. (1977). *Algebraic Geometry*. Springer, New York.
- [9] Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press.
- [10] Kolmogorov, A. N. (1963). On tables of random numbers. *Sankhya*, 25, 369–376.
- [11] Lions, J.-L. (1969). *Quelques méthodes de résolution des problèmes aux limites non linéaires*. Dunod, Paris.
- [12] Mac Lane, S. (1998). *Categories for the Working Mathematician*, 2nd edition. Springer, New York.
- [13] Mallat, S. (1999). *A Wavelet Tour of Signal Processing*. Academic Press, San Diego.
- [14] Misner, C. W., Thorne, K. S., and Wheeler, J. A. (1973). *Gravitation*. W. H. Freeman, San Francisco.
- [15] Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*, 2nd edition. Cambridge University Press.
- [16] Penrose, R. (2004). *The Road to Reality: A Complete Guide to the Laws of the Universe*. Jonathan Cape, London.

- [17] Prigogine, I. and Stengers, I. (1984). *Order Out of Chaos*. Bantam Books, New York.
- [18] Riehl, E. (2017). *Category Theory in Context*. Dover, New York.
- [19] Solomonoff, R. J. (1964). A formal theory of inductive inference. *Information and Control*, 7(1), 1–22.