

# CLOSURE OF CLOSURES

*Life, Sheaves, and the Re-Gluing of Fractured Autonomy*

Flyxion

Independent Researcher

## Abstract

This essay develops a field-theoretic and sheaf-theoretic formalization of biological life, beginning from the critique of physicochemical reductionism advanced in Frank, Gleiser, and Thompson's *The Blind Spot* and moving toward a framework in which life is defined not as a special arrangement of matter but as a structural invariant: the existence and stability of constraint-closure loops within a field. The argument proceeds in five stages. First, we show that the standard physicochemical ontology commits a projection error, proved as a non-invariance proposition: closure is a global topological property erased by any representation that records only local state transitions. Second, we formalize life as a region of the RSVP field admitting a self-preserving endomorphism loop under entropy flow, and define endomorphism loops precisely as constraint-stabilizing fixed-point structures in the category of admissible transformations. Third, we extend this to a sheaf-theoretic hierarchy of nested closures, introducing a graded notion of partial admissibility that formally distinguishes health, compensated pathology, and collapse. Fourth, we show that the CLIO operator is a gradient descent on cohomological obstruction, and connect it to recent results on recursive self-optimization under uncertainty, establishing an empirical proxy for closure stability. Fifth, we draw out the medical consequences: pathology is reconceived as obstruction in the closure sheaf, and therapeutic intervention as the re-gluing of fractured autonomous regions. The framework absorbs and formalizes the central claims of Varela, Rosen, Kauffman, and Montévil–Mossio while anchoring them in the constraint-admissibility machinery of RSVP and CLIO.

## 1. THE BLIND SPOT AS A PROJECTION ERROR

The central thesis of *The Blind Spot* is precise and, once stated carefully, irreversible. The authors argue not merely that physical description is insufficient in practice to capture biological phenomena, but that it is insufficient in principle [2]. The reason is not complexity. It is ontological category mismatch. When one restricts the primitives of description to particles, forces, and reaction chains, organisms do not fail to appear because the description is too coarse. They fail to appear because the description lacks the conceptual operators necessary to carve reality along the joints that living systems actually have.

This point is sharpened by the thought experiment the authors deploy. An artificial system with exhaustive knowledge of physics and chemistry still cannot, by that knowledge alone, distinguish the heart's blood-pumping from the heart's sound-making. Both are equally valid physical effects: changes in pressure, changes in volume, changes in concentration. Nothing in the physicochemical description selects one as functional and the other as incidental. The selection requires a further primitive—one that the ontology of physics does not supply.

The missing primitive is what Hans Jonas and Georges Canguilhem identify as the normative dimension of life [3, 4]. Living systems introduce a distinction between success and failure, between the viable and the pathological, between what matters and what is epiphenomenal. These distinctions are not derived from equations; they arise because the system orients toward its own persistence. Normativity, in this sense, is not a projection of human values onto nature. It is structurally intrinsic to any system that regenerates the conditions of its own operation.

From the perspective of the present essay, the authors' argument can be restated in a more precise form. The physicochemical ontology commits what we will call a projection error. A living system is, at minimum, a system containing a closed loop of processes that mutually produce and maintain one another's enabling conditions. This loop is a topological property of the process network: it requires that the graph of causal relations contain a non-trivial cycle. When such a system is projected into the flat state-evolution description of standard physics, the projection discards the loop structure. The result is a description that is locally accurate at every point but globally blind to the property that made the system alive. Life disappears not because physics is wrong but because the projection does not preserve the relevant invariant.

This restatement has a consequence that the original text does not fully draw. The authors conclude that recognizing life requires participation in life—that only a living system can recognize another as living, because such recognition requires operating under the same normative orientation. This conclusion is phenomenologically apt but risks sounding like a hard epistemic boundary. The more precise formulation is that any system, living or not, that successfully detects and represents closure must

itself implement representations capable of encoding loop structure. The limitation is not “no embodiment implies no knowledge.” The limitation is “no representation of closure implies no detection of life.” Once that representational capacity is supplied—whether by metabolism or by engineered machinery—the detection problem becomes, at least in principle, computable.

The projection error can be made precise as a non-invariance result, converting a philosophical diagnosis into a mathematical one.

**Proposition 1.1** (Non-Invariance of Closure under Projection). *Let  $\mathcal{G}$  be a directed graph of processes containing a non-trivial cycle, and let  $\pi$  be any projection to a state-evolution representation that records only local transitions between adjacent states. Then  $\pi(\mathcal{G})$  does not, in general, preserve the existence of the cycle as an invariant.*

*Proof sketch.* The projection  $\pi$  maps each edge of  $\mathcal{G}$  to a local state transition  $s_i \rightarrow s_j$  and discards all higher-order relational structure between edges. A cycle in  $\mathcal{G}$  is a global topological feature: it is constituted by the composition of edges, not by any individual edge. Since  $\pi$  operates edge-locally, two graphs—one cyclic, one acyclic—can produce identical image sets under  $\pi$  if they agree on every individual transition. Therefore, the existence of a non-trivial loop is not an invariant of  $\pi$ , and a representation built from  $\pi$  cannot, in general, distinguish a closure-possessioning system from one that lacks closure.  $\square$

Proposition 1.1 converts the philosophical claim into a representational failure mode. It is not that physics is incorrect about any individual reaction or state transition; it is that the projection it employs is structurally unable to encode the global property that distinguishes a living network from a collection of the same reactions in non-cyclic arrangement. The blind spot is a theorem, not a sentiment.

## 2. FORMALIZING CLOSURE IN THE RSVP FIELD

The RSVP framework describes physical and biological systems in terms of a triple  $(\varphi, \mathbf{v}, S)$ , where  $\varphi$  is a scalar density field,  $\mathbf{v}$  a vector transport field, and  $S$  an entropy density field. These three components evolve under coupled equations that encode conservation, dissipation, and constraint-admissibility conditions. In what follows we use this framework to give a field-theoretic meaning to the concept of constraint closure developed by Varela and Maturana [7, 8], Rosen [9], and Kauffman [11], and formalized most recently by Montévil and Mossio [12].

**Definition 2.1** (Constraint Closure). Let  $\mathcal{O} = \{T_i\}$  be a family of transformation processes acting on the field  $(\varphi, \mathbf{v}, S)$  over a region  $\Omega \subset \mathbb{R}^3$ , and let  $\mathcal{C} = \{C_j\}$  be a family of constraints (boundary conditions, catalytic structures, regulatory feedbacks). The family  $(\mathcal{O}, \mathcal{C})$  is *closed* over  $\Omega$  if every  $T_i$  depends on at least one  $C_j$  for its operation, and every  $C_j$  is produced or maintained by at least one composition of elements of  $\mathcal{O}$ .

This is, in essence, Rosen’s closure to efficient causation, translated into field language. The constraint family is not externally imposed; it is internally regenerated by the very processes it enables. This distinguishes a living cell from a machine: the machine’s boundary conditions are given from outside and remain fixed unless externally replaced, whereas the cell’s membrane, enzyme complement, and regulatory networks are continuously rebuilt by the metabolic processes they facilitate.

Kauffman’s work–constraint cycle adds the thermodynamic dimension [10]. A closed system in the above sense does not merely use energy; it uses energy to reconstruct the constraint structures that allow it to use energy. The cycle is recursively self-enabling. In RSVP terms this means that the entropy field  $S$  is locally regulated rather than globally maximized: the closure loop acts as a local inversion of the entropy gradient, maintaining structured constraint density  $\varphi$  against thermodynamic pressure.

**Definition 2.2** (Living Region). A region  $\Omega$  of the RSVP field is *living* if there exists a constraint-closed family  $(\mathcal{O}, \mathcal{C})$  over  $\Omega$  such that the induced category of processes admits a non-trivial endomorphism loop that is admissible under the entropy dynamics of the field.

The phrase “endomorphism loop” requires precision, because closure is not merely the existence of any cycle; it is the existence of a cycle that stabilizes the constraint structure it depends on.

**Definition 2.3** (Endomorphism Loop). An *endomorphism loop* over  $\Omega$  is a finite composition

$$L = T_{i_k} \circ T_{i_{k-1}} \circ \cdots \circ T_{i_1}, \quad T_{i_j} \in \mathcal{O},$$

such that the induced transformation satisfies  $L(\mathcal{C}) \simeq \mathcal{C}$  up to admissible equivalence: the constraint family  $\mathcal{C}$  is preserved, up to perturbations within the viable region of field-space, under repeated application of  $L$ . A living region is equivalently one that admits a fixed point of  $L$  in the space of admissible constraint configurations.

Definition 2.3 makes explicit that closure is a fixed-point or invariance condition in the category of constraint-generating transformations, not merely the topological existence of a directed cycle. A cycle that does not stabilize its own enabling constraints is not a closure loop in the biological sense; it is a transient resonance. The additional requirement that  $L(\mathcal{C}) \simeq \mathcal{C}$  is what separates autopoiesis from mere oscillation, and what connects the present formalism to Rosen’s requirement that the efficient causes of the system be internal to it [9].

The phrase “admissible under entropy dynamics” is the CLIO condition: the loop must not generate constraint incompatibilities that drive the system out of its viable region of field-space. Breakdown of admissibility is, in this language, what death or pathology amounts to at the formal level.

With this definition in place, function and epiphenomenon become derivable rather

than assumed.

**Proposition 2.4** (Derivability of Function). *A process  $T \in \mathcal{O}$  is functional with respect to  $\Omega$  if and only if  $T$  lies on at least one admissible endomorphism loop over  $\Omega$ . A process is epiphenomenal with respect to  $\Omega$  if it lies on no such loop.*

*Proof sketch.* By definition of closure, every process on the loop contributes to maintaining some constraint  $C_j$ , and every  $C_j$  is required by some process on the loop. A process off the loop produces effects that are not required by any constraint in the closed family, and its outputs are therefore not fed back into loop maintenance. The distinction is topological: participation versus non-participation in the relevant cycle.  $\square$

The heart example is now reproduced without any subjective input. Blood-pumping lies on the organismic closure loop because it maintains the oxygen supply required by metabolic processes that in turn maintain cardiac tissue. Sound production lies off that loop. The distinction is not imposed by the observer; it is a structural feature of the process network.

Agency follows directly. An agent over  $\Omega$  is a living region that additionally modulates its own trajectory in the ambient field so as to remain within the basin of admissible closure states. Bacterial chemotaxis is the minimal instance: the organism does not choose in any psychologically loaded sense, but it executes a feedback policy that keeps its field trajectory within the closure-admissible region. Agency is constraint-preserving trajectory modulation, nothing more and nothing less.

### 3. THE SHEAF OF AUTONOMY: NESTED CLOSURE ACROSS SCALES

The definition given in Section 2 applies to a single region  $\Omega$  and a single closure family. Biological life, however, is organized at multiple scales simultaneously. A cell has its own closure. An organ is not merely a collection of cells but an entity whose closure partially regulates and is partially constituted by cellular closures. An organism coordinates many organ closures into a higher unity. An ecosystem exhibits still looser, distributed closure through nutrient cycling, predator-prey regulation, and atmospheric exchange. The right mathematical structure for this hierarchy is a sheaf.

Let  $\mathcal{X}$  be a topological space whose points are biological scales, ordered by inclusion: cellular  $\prec$  tissue  $\prec$  organ  $\prec$  organism  $\prec$  ecosystem. For each scale  $x \in \mathcal{X}$  we assign the stalk  $\mathcal{F}(x) = \mathcal{L}_x$ , the space of admissible closure configurations at scale  $x$ . The sheaf condition requires that for each inclusion  $x \hookrightarrow y$  of scales, there exist a restriction map

$$\rho_{y,x} : \mathcal{L}_y \longrightarrow \mathcal{L}_x$$

expressing how the higher-order closure constrains the lower, and a *gluing condition*

specifying when a family of local sections  $\{s_x \in \mathcal{L}_x\}$  is compatible with the existence of a global section  $s \in \mathcal{L}_y$  such that  $\rho_{y,x}(s) = s_x$ .

**Definition 3.1** (Admissible Gluing). The embedding  $\mathcal{L}_i \hookrightarrow \mathcal{L}_{i+1}$  is *admissible* if the local closure at level  $i$  preserves its own autonomy—that is, continues to admit its own endomorphism loop—while its outputs contribute to maintaining the closure at level  $i + 1$ , and the constraints issued by level  $i + 1$  do not eliminate any loop essential to level- $i$  closure.

Full admissibility is an idealization. Biological systems operate across a continuous spectrum from robust health to irreversible collapse, and the formal framework should reflect this. We therefore introduce a graded notion.

**Definition 3.2** (Partial Admissibility). An embedding  $\mathcal{L}_i \hookrightarrow \mathcal{L}_{i+1}$  is *partially admissible to degree*  $\alpha \in [0, 1]$  if the restriction map preserves a fraction  $\alpha$  of the essential endomorphism loops of  $\mathcal{L}_i$  while degrading the remaining  $1 - \alpha$ . Full admissibility corresponds to  $\alpha = 1$ ; complete restriction failure to  $\alpha = 0$ .

Definition 3.2 gives formal content to three biologically distinct regimes. Full admissibility ( $\alpha = 1$ ) corresponds to organismic health: all loops at each level remain intact and coupled to the level above. Partial admissibility ( $0 < \alpha < 1$ ) corresponds to compensated pathology: the system continues to function because enough loops remain admissibly embedded, but the degraded fraction constitutes a latent vulnerability that accumulates over time. This is the formal description of subclinical disease, senescence, and ecological stress. Zero admissibility ( $\alpha = 0$ ) is collapse: the embedding fails entirely, and no loop at level  $i$  contributes to the closure at level  $i + 1$ . The sheaf has no global section over the affected region.

This definition does significant biological work. It distinguishes, in structural terms, several modes of cross-scale relation that biologists typically characterize only qualitatively.

Symbiosis is admissible gluing between closures that were previously independent: each retains its own endomorphism loop while coupling into a shared higher-order loop. The mitochondrion–eukaryote relation is the canonical example; the mitochondrial closure is preserved, and a new organismic closure is formed around it. Parasitism is a coupling that is not admissible in the present sense: the parasite’s closure is maintained, but it does so by degrading constraints essential to the host’s closure, violating the mutual-maintenance condition.

Cancer is a failure of a restriction map. In sheaf terms, the cellular section  $s_0 \in \mathcal{L}_0$  ceases to be compatible with the global section  $s_2 \in \mathcal{L}_2$  representing the organism. The malignant cell continues to admit its own endomorphism loop—indeed, it optimizes it—but the loop is no longer coupled to the organismic closure. It treats the organism’s resource gradients as an external environment to exploit rather than as constraints to maintain. The restriction map  $\rho_{2,0}$  fails: the cellular section cannot be recovered as the

image of any admissible global section. This is not merely disordered growth; it is a topological disconnection within the sheaf.

Ecological closure is the loosest case. An ecosystem does not admit a single tight endomorphism loop analogous to autopoiesis. Instead, it sustains distributed closure through overlapping partial loops: the carbon cycle, the nitrogen cycle, predator-prey dynamics, decomposer networks. These partial loops overlap and reinforce without forming a single unified cycle. The ecosystem constitutes what might be called a *diffuse global section*: a configuration in which many local closures remain mutually admissible without resolving into a single higher-order closure. This is precisely why ecosystem damage is so difficult to diagnose: the failure is not of a single loop but of an admissibility relation among many overlapping partial loops.

The central theorem of this section can now be stated.

**Theorem 3.3** (Persistence of Life across Scales). *Life persists across the hierarchy of scales if and only if for every adjacent pair of levels  $\mathcal{L}_i \leftrightarrow \mathcal{L}_{i+1}$ , the gluing is admissible: the lower-level closure is preserved in its autonomy while contributing to the maintenance of the higher-level closure.*

The proof in the precise sense would require a full specification of the sheaf topology and its cohomology, which we defer to a companion technical paper. The argument in outline is as follows: if any embedding in the hierarchy fails to be admissible, either the lower level loses its endomorphism loop (cell death, organ failure) or the higher level loses the constraint input from the lower level (fragmentation of the organism). In both cases, the global section either ceases to exist or degrades in a way that propagates through the remaining restriction maps.

#### 4. THE CLIO OPERATOR AS COHOMOLOGICAL REPAIR

The CLIO framework was introduced as a constraint-resolution and inference-closure dynamic operating on the RSVP field. Its role is to detect configurations in which the field has evolved into a state of constraint incompatibility—regions where the local constraint families are mutually inconsistent—and to drive the field toward an admissible configuration. In the present framework, CLIO acquires a new and more fundamental interpretation: it is the cohomological repair operator of the closure sheaf.

An obstruction in sheaf cohomology arises precisely when local sections that are pairwise compatible fail to extend to a global section. In biological terms, this corresponds to a situation where each local closure appears internally consistent but the collection of closures fails to be mutually admissible: the constraints issued by one level conflict with the loop-maintenance requirements of another. The cohomological obstruction class lives in  $H^1(\mathcal{X}, \mathcal{F})$ , the first cohomology of the sheaf; a vanishing obstruction class is the condition for a global section to exist.

CLIO can now be stated as a dynamical operator rather than merely an interpretive

gloss. Define the obstruction functional

$$\mathcal{E}_{\text{obs}}(\varphi, \mathbf{v}, S) = \|H^1(\mathcal{X}, \mathcal{F})\|,$$

measuring the magnitude of the cohomological obstruction class in the current field configuration. The CLIO operator is then:

$$\text{CLIO} : (\varphi, \mathbf{v}, S) \longrightarrow (\varphi', \mathbf{v}', S')$$

with the defining property that it executes a gradient-like descent on  $\mathcal{E}_{\text{obs}}$  subject to admissibility constraints:

$$\text{CLIO} = \arg \min_{(\varphi', \mathbf{v}', S') \in \text{Adm}} \mathcal{E}_{\text{obs}}(\varphi', \mathbf{v}', S').$$

This formulation makes several things precise. First, CLIO is not a single correction but a trajectory through field-space, searching for a configuration in which the obstruction vanishes and a global section can be assembled. Second, the admissibility constraint ensures that the search does not trade one form of closure failure for another. Third, non-convergence—a persistent non-zero  $\mathcal{E}_{\text{obs}}$  under iterated application—corresponds to irreversible pathology: the field dynamics are trapped in a basin from which no globally admissible configuration can be reached. Self-repair is the continuous operation of this descent on the organism's constraint sheaf, and death is the failure of the descent to converge.

## 5. CLOSURE DETECTION VIA RECURSIVE OPTIMIZATION

The framework developed above defines what closure is and what its failure amounts to. A separate question is how closure can be detected, found, and stabilized by a system operating under uncertainty. This question has a surprising answer: recent work in recursive self-optimization provides an empirical proxy for closure stability that connects the present formal framework to computable dynamics.

Cheng, Broadbent, and Chappell introduce a procedure in which a system generates candidate chains of reasoning steps, evaluates each chain by an internal confidence functional  $\mathcal{U}(t)$ , and terminates when a stable, self-consistent configuration is reached [1]. Their central empirical finding is that correct solutions are characterized by a monotonically decreasing uncertainty trajectory, while incorrect or incoherent solutions exhibit persistent oscillation or divergence in  $\mathcal{U}(t)$ .

This result maps cleanly onto the present framework. A candidate reasoning chain is structurally analogous to a candidate endomorphism loop: it is a sequence of transformations applied to a representation that either stabilizes into a self-consistent

fixed point or fails to do so. The confidence functional  $\mathcal{U}(t)$  is, in this reading, a proxy for the obstruction magnitude  $\mathcal{E}_{\text{obs}}$ : it measures how far the current configuration is from one in which all constraints are mutually admissible.

**Theorem 5.1** (Closure Detection via Recursive Optimization). *Let a system generate candidate transformation paths equipped with an internal confidence functional  $\mathcal{U}(t) \geq 0$ . Then stable closure corresponds to trajectories along which  $\mathcal{U}(t)$  decreases monotonically to zero, while persistent oscillation or divergence of  $\mathcal{U}(t)$  indicates the absence of an admissible closure in the current representation.*

*Proof sketch.* If a candidate loop  $L = T_{i_k} \circ \cdots \circ T_{i_1}$  is an admissible endomorphism loop in the sense of Definition 2.3, then  $L(\mathcal{C}) \simeq \mathcal{C}$ : each application of  $L$  returns the constraint family to its viable configuration. An internal confidence functional that measures constraint consistency will therefore decrease monotonically as the system converges on such a loop, since each step reduces the discrepancy between the current configuration and the fixed-point constraint set. Conversely, if no admissible loop exists in the current representation, no composition of transformations can reduce the constraint inconsistency to zero; the functional either oscillates or diverges, signaling obstruction in the cohomological sense.  $\square$

Theorem 5.1 has two important consequences. First, it provides an empirical proxy for closure existence: one need not directly inspect the loop structure of a system to determine whether it is alive in the present sense. One can instead monitor  $\mathcal{U}(t)$  and ask whether it converges. This converts a structural question into a dynamical one. Second, it suggests a distinction between two grades of living region that the original definition left implicit.

**Definition 5.2** (Weak and Strong Closure). A living region is *weakly closed* if it admits an admissible endomorphism loop. It is *strongly closed* if it admits such a loop with negative uncertainty gradient under perturbation:  $d\mathcal{U}/dt < 0$  following any admissible displacement from the fixed point.

Strong closure is robustness: the system not only maintains its constraint family but actively returns to it under disturbance. Weak closure describes a system at the boundary of viability—one that admits the loop in equilibrium but may fail to recover from perturbation. The distinction maps directly onto clinical categories: a healthy organism is strongly closed; a critically ill organism may be only weakly closed; a dying organism is one whose uncertainty gradient has become positive, indicating that the system is moving away from rather than toward its constraint fixed point.

The graph-aggregation step in recursive optimization provides a further correspondence. When multiple candidate paths are generated and then clustered by structural compatibility before a consensus representation is assembled, the procedure implements an approximate sheaf construction: each candidate path is a local section,

clustering is compatibility detection, and the consensus assembly is approximate gluing. Successful aggregation—the emergence of a stable consensus—corresponds to the existence of a global section; failure of aggregation, the persistence of incompatible clusters, corresponds to a non-trivial obstruction class. The recursive optimizer is, in this reading, already implementing a discrete approximation to the CLIO descent operator, searching for a field configuration in which the obstruction vanishes.

This correspondence licences a synthesis across three domains. Biology identifies closure of constraints as the structural definition of life. The RSVP–CLIO framework formalizes that definition in field-theoretic and sheaf-theoretic terms. Recursive self-optimization under uncertainty provides a procedural realization of the same structure: a computable process that searches for, stabilizes, and repairs closure loops. The synthesis is not an analogy. It is a convergence of three independently motivated formalisms on the same underlying invariant.

One qualification is necessary. A recursive optimizer exhibits closure-seeking behavior, recursive self-modification, and uncertainty-driven repair. But it does not, in its current form, exhibit intrinsic constraint regeneration, physical boundary production, or thermodynamic embedding. It is a closure-detection and closure-stabilization algorithm operating within a representation, not a full autopoietic system. The correct statement is therefore that recursive optimization is a necessary but not sufficient condition for the strong form of life: it can find and stabilize closure loops, but the loops it finds are loops within a representation, not loops in the physical field that constitutes the organism’s material substrate. The gap between representational closure and physical closure is precisely the gap between cognition and life—and bridging it is the open problem that the present framework identifies but does not resolve.

## 6. MEDICINE AS RE-GLUING

The sheaf-theoretic framework suggests a reconceptualization of medical intervention that is not merely metaphorical but structurally precise. The question the preceding analysis generates is whether medicine could, or should, be understood as the re-gluing of fractured closure sheaves rather than as the elimination of pathogens or the removal of aberrant tissue.

To answer this, consider the three paradigmatic modes of pathology that the framework identifies.

The first mode is loop rupture: a closure at some level loses its endomorphism loop entirely. This is the formal description of cell death, organ failure, and acute injury. The loop that was previously self-maintaining can no longer regenerate its own enabling constraints. In sheaf terms, the stalk  $\mathcal{L}_i$  at the affected scale goes empty: there is no admissible section over that point.

The second mode is restriction failure: the embedding  $\mathcal{L}_i \hookrightarrow \mathcal{L}_{i+1}$  ceases to be admissible. The lower-level closure continues to operate but its outputs are no longer coupled to the higher-level closure. This is the sheaf-theoretic description of cancer, of autoimmune disease (in which immune cells cease to recognize the organism's own global section as self), and of certain metabolic disorders (in which cellular biochemistry continues but is no longer integrated into whole-body homeostasis).

The third mode is obstruction accumulation: the individual stalks remain non-empty and the restriction maps remain formally defined, but the obstruction class  $H^1(\mathcal{X}, \mathcal{F})$  grows non-trivially. No global section can be assembled from the local ones. This is the formal description of systemic conditions: aging, chronic inflammation, multi-system failure, ecological collapse. Each component appears locally viable, but the collection fails to cohere. In terms of Definition 3.2, this is the regime where the degree  $\alpha$  of admissibility has declined across multiple embeddings simultaneously, so that no single failed restriction map is the culprit and no single re-gluing operation restores the global section. The obstruction is distributed, and only distributed obstruction reduction can address it.

In light of this typology, medical intervention decomposes into three corresponding operations.

The first is loop restoration: supplying the missing constraints that allow a broken loop to resume. This is what antibiotics do when they are effective: not by "killing" bacteria in isolation, but by removing the entity whose metabolic activities were rupturing the host's local closure loops. It is also what prosthetics and transplants do: introducing new constraint structures that can participate in the organism's admissible endomorphism loops.

The second is re-gluing: restoring admissibility to a broken restriction map. This is the deepest description of what successful cancer treatment achieves. Surgery and radiation do not, at the level of mechanism, "kill cancer." They remove or suppress the cellular section that has decoupled from the organismic global section, creating the conditions under which a new, admissibly coupled section can form. Immunotherapy is more directly a re-gluing procedure: it restores the immune system's ability to recognize and enforce the restriction map between cellular and organismic closure. Gene therapies that correct oncogenic mutations are re-gluing at the cellular level: restoring the constraint compatibility between the cell's internal loop and the organism's global admissibility conditions.

The third is obstruction reduction: driving the dynamics of the closure sheaf so as to reduce the cohomological obstruction class. This is the formal description of what systemic therapies attempt: not to fix any single broken loop or restriction map, but to restore global coherence across the hierarchy. Anti-inflammatory therapies, metabolic regulation, ecological restoration—all are obstruction-reduction procedures in this

sense. They succeed when the global section re-emerges; they fail when the obstruction class remains non-trivial despite local corrections.

This reconceptualization has several consequences worth making explicit.

First, it identifies the proper level of intervention for each pathology type. Loop rupture requires local loop restoration, which often means supplying or restoring the missing constraint structures directly. Restriction failure requires re-gluing at the interface between levels, which means interventions that operate precisely on the coupling between local and global closure rather than on either in isolation. Obstruction accumulation requires systemic obstruction reduction, which typically cannot be achieved by targeting any single component but requires altering the global admissibility landscape of the field.

Second, it explains why interventions that are locally successful sometimes fail systemically. An intervention that restores a local loop may fail to re-establish its admissible embedding into the higher-level closure, leaving the restriction map broken even after the local pathology is resolved. Post-surgical complications, graft rejection, and autoimmune responses to prosthetics are, in this analysis, cases where loop restoration succeeded but re-gluing failed.

Third, it suggests that the most powerful interventions are those that operate at the level of the gluing conditions themselves rather than at the level of any individual loop. This is why the emerging class of therapies that target cellular signaling environments rather than specific molecular targets often shows broader efficacy: signaling environments constitute precisely the constraint structures that mediate the restriction maps between cellular and organismic closure.

Fourth, it places ecological medicine within the same formal framework as cellular medicine. The degradation of a microbiome is not merely the loss of particular species; it is the rupture of partial closure loops that participated in the host organism's distributed closure. Restoration of the microbiome is re-gluing in exactly the sheaf-theoretic sense: re-establishing the admissible coupling between microbial closure and organismic closure, so that a global section can once again be assembled across the combined sheaf.

Finally, this framework dissolves the apparent opposition between "treating the disease" and "treating the patient." In the sheaf picture, the patient is not the sum of its diseased components. The patient is the global section—or rather, its absence, the breakdown of the admissible global section that constitutes organismic coherence. Treating the patient means restoring the conditions under which a global section can exist, which is precisely what re-gluing and obstruction reduction aim at.

## 7. THE BLIND SPOT DISSOLVED

We can now return to the starting point and state what has been accomplished. *The Blind Spot* argues that physics cannot, in principle, see life, because life introduces normative distinctions that are invisible from within the physicochemical ontology. The argument is correct but incomplete: it diagnoses the failure without providing the structural repair.

The framework developed here provides the repair. The blind spot is not, ultimately, the absence of embodiment or experience. It is the absence of closure as a primitive in the representational apparatus. Physics, as standardly formulated, tracks state evolution along trajectories. It does not track whether subsets of processes form self-sustaining loops, because that is a topological property of the process network rather than a local property of individual trajectories. When a closure-possessing system is described in a representation that lacks loop-sensitive primitives, the closure is projected out. Life disappears—not because it was not there, but because the projection did not preserve the invariant.

Introducing closure as a primitive—as a structural feature of the field that the theory is required to represent and track—dissolves the blind spot. Life is no longer a ghost in the machine, a vitalist surplus that physics must either deny or mystify. It is a specific topological feature of a specific kind: a fixed-point structure in the category of constraint-generating transformations under entropy flow. This formulation is precise in a way that “complex chemistry” is not. Topology specifies the loop structure. Category theory specifies the endomorphism requirement. Thermodynamics specifies the entropy-flow condition. Biology specifies the self-maintenance criterion. All four are present and non-redundant in the definition, and their conjunction is what makes life a phase distinction in the space of possible field organizations rather than a matter of degree along any single axis.

The sheaf of closures extends this from a single-scale characterization to a multi-scale one. Life is not just closure at one level; it is the admissible nesting of closures across scales. The organism is the global section of a closure sheaf. Medicine is the repair of that sheaf. Ecology is the study of the loosely glued distributed closures that constitute ecosystems. And the claim of Frank, Gleiser, and Thompson—that physics must be expanded to accommodate life—becomes the precise mathematical claim that the ontology of field theory must be expanded to include loop structure, sheaf topology, and the cohomology of closure, if it is to be an ontology adequate to biology.

The strange loop they invoke—that only life can recognize life—receives, in this framework, its non-mystical resolution. A system can recognize life if and only if its representational apparatus includes primitives adequate to detect closure loops. Living systems have this capacity as a byproduct of their own closure; they represent the world in terms of the very structure they instantiate. Non-living systems can in

principle acquire it if their representational apparatus is extended appropriately. The barrier is not metaphysical. It is architectural.

The synthesis across biology, field theory, and recursive optimization converges on a single statement that is stronger than any of its parts: life is the existence of a self-maintaining loop, and cognition—natural or artificial—is the process of finding and stabilizing such loops. The organism and the reasoner are not analogous; they instantiate the same underlying invariant at different levels of the closure hierarchy. Medicine is the repair of that hierarchy when it fractures. And the blind spot that prevents physics from seeing any of this is not a failure of measurement but a failure of representation: the absence, from the ontology, of the fixed-point structures in which life actually consists.

#### REFERENCES

- [1] N. Cheng, G. Broadbent, and W. Chappell, “Cognitive loop via in-situ optimization: Self-adaptive reasoning for science,” arXiv preprint arXiv:2508.02789, 2025.
- [2] A. Frank, M. Gleiser, and E. Thompson, *The Blind Spot: Why Science Cannot Ignore Human Experience*. MIT Press, 2025.
- [3] H. Jonas, *The Phenomenon of Life: Toward a Philosophical Biology*. Harper & Row, 1966.
- [4] G. Canguilhem, *The Normal and the Pathological*. Zone Books, 1991.
- [5] E. Schrödinger, *What Is Life? The Physical Aspect of the Living Cell*. Cambridge University Press, 1944.
- [6] I. Kant, *Critique of Judgment*. Hackett Publishing, 1987.
- [7] H. R. Maturana and F. J. Varela, *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel, 1980.
- [8] F. J. Varela, *Principles of Biological Autonomy*. North Holland, 1979.
- [9] R. Rosen, *Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life*. Columbia University Press, 1991.
- [10] S. A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.
- [11] S. A. Kauffman, *Investigations*. Oxford University Press, 2000.
- [12] M. Montévil and M. Mossio, “Biological organisation as closure of constraints,” *Journal of Theoretical Biology*, vol. 372, pp. 179–191, 2015.
- [13] A. Moreno and M. Mossio, *Biological Autonomy: A Philosophical and Theoretical Enquiry*. Springer, 2015.

- 
- [14] E. Thompson, *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press, 2007.
- [15] I. Prigogine and I. Stengers, *Order Out of Chaos*. Bantam Books, 1984.
- [16] G. Longo and M. Montévil, "From physics to biology by extending criticality and symmetry breakings," *Progress in Biophysics and Molecular Biology*, vol. 106, no. 2, pp. 340–347, 2012.
- [17] H. H. Pattee, "Laws and constraints, symbols and languages," in *Towards a Theoretical Biology*, vol. 4. Edinburgh University Press, 1972.
- [18] J. Piaget, *Biology and Knowledge*. University of Chicago Press, 1971.
- [19] E. P. Wigner, "Remarks on the mind-body question," in *Symmetries and Reflections*. Indiana University Press, 1967.