

The Mathematics of Autonomous Systems

Admissibility, Repair, Continuation,
Representation, and Scientific Discovery

Flyxion
Independent Researcher

July 31, 2026

Dedication

*To everyone who believes
that understanding
is not accumulated,
but recursively repaired.*

This monograph develops a unified mathematical framework for autonomous systems, admissibility, recursive repair, continuation, observation, representation, and scientific discovery.

Rather than treating cognition, computation, biology, scientific inquiry, and distributed intelligence as independent subjects, each is formulated as an instance of constrained evolution upon admissible state spaces.

Beginning with the notions of admissible interfaces and invariant preservation, we construct a hierarchy of mathematical structures including repair operators, continuation semigroups, distinction geometry, reachability manifolds, sheaf-theoretic reconstruction, categorical observation, and recursive theory evolution.

Within this framework, identity emerges from admissible continuation rather than static substance, information arises from distinguishability rather than probability alone, and scientific progress becomes recursive navigation through the geometry of theory space.

The resulting formalism replaces numerous informal notions—including boundaries, sanctuaries, equilibrium, optimization, and discovery—with general mathematical objects applicable across physical, biological, computational, and epistemological domains.

The central thesis is that autonomy is neither isolation nor independence, but invariant-preserving interaction through admissible interfaces. Likewise, scientific knowledge is neither a collection of isolated facts nor an accumulation of arbitrary hypotheses, but the recursive repair of representational structures constrained by observation and governed by continuation.

Contents

1	Autonomous Cognitive Domains and Admissible Interfaces	1
1.1	Motivation	1
1.2	State Spaces	1
1.3	The Error of Closed-System Thinking	2
1.4	Admissible Interfaces	2
1.5	Coupling Operators	3
2	Structural Entropy and Environmental Perturbation	6
2.1	Noise Is Not an Environmental Property	6
2.2	Structural Entropy	7
2.3	Signal	7
2.4	Noise	8
2.5	The Filtering Principle	8
3	Repair Operators and Recursive Homeostasis	10
3.1	Home Is Not a Location	10
3.2	Deviation from Admissibility	11
3.3	Repair Operators	11
3.4	Recursive Repair	12
3.5	Home as an Attractor	13
4	The Geometry of Collaboration	15
4.1	Against the Myth of Isolation	15
4.2	Knowledge Graphs	15
4.3	Indirect Coupling	16
4.4	Structural Independence	16
4.5	Distributed Repair	17
4.6	Emergent Intelligence	17
4.7	The Fallacy of the Solitary Genius	18
5	Persistence, Continuation, and Identity	19
5.1	Identity Is Not Substance	19
5.2	Continuation Operators	19
5.3	Persistence	20
5.4	Identity Through Repair	20
5.5	The Failure of Static Ontologies	21
5.6	The Continuation Principle	21

5.7	Identity as a Fixed Structure	22
6	Observers, Interfaces, and Representation	23
6.1	The Impossibility of Direct Observation	23
6.2	Interface States	23
6.3	Representation	24
6.4	Information Loss	25
6.5	Observational Equivalence	25
6.6	Reconstruction	25
6.7	Representation Defect	26
6.8	The Representation Principle	27
7	Distinction Geometry and the Topology of Information	28
7.1	Information Is Not Primitive	28
7.2	The Distinction Space	28
7.3	Distinction Fields	29
7.4	Observational Collapse	29
7.5	Distinction Entropy	30
7.6	Repair of Distinctions	30
7.7	Distinction Curvature	30
7.8	The Distinction Principle	31
8	Constraint Geometry and Reachability	32
8.1	Constraint Precedes Motion	32
8.2	Constraint Manifolds	32
8.3	Reachability	33
8.4	Reachability Components	33
8.5	Boundary Geometry	34
8.6	Constraint Curvature	34
8.7	Constraint Energy	34
8.8	The Reachability Principle	35
9	The Geometry of Constraint Propagation	37
9.1	Local Constraints and Global Structure	37
9.2	Constraint Networks	37
9.3	Constraint Compatibility	38
9.4	Propagation Operators	38
9.5	Monotonicity	38
9.6	Fixed Points	39
9.7	Constraint Repair	39
9.8	Distributed Consistency	40
10	Sheaf-Theoretic Reconstruction and Global Consistency	41
10.1	The Local Nature of Knowledge	41
10.2	Observation Covers	41
10.3	Local Sections	42
10.4	Compatibility	42

10.5 The Gluing Principle	43
10.6 Reconstruction Defect	43
10.7 Repair by Gluing	44
10.8 Distributed Truth	44
10.9 The Local-to-Global Principle	44
11 Functorial Observation and the Category of Autonomous Systems	46
11.1 Beyond State Spaces	46
11.2 The Category of Autonomous Systems	46
11.3 Observation as a Functor	47
11.4 Forgetful Processes	47
11.5 Natural Transformations	48
11.6 Repair as an Endofunctor	48
11.7 Equivalence of Autonomous Systems	49
11.8 The Functorial Principle	49
12 Scientific Discovery as Recursive Navigation	51
12.1 Discovery Is Not Creation	51
12.2 Theory Space	51
12.3 Theory Morphisms	52
12.4 Research Trajectories	52
12.5 The Discovery Functional	53
12.6 Recursive Refinement	53
12.7 Scientific Fixed Points	54
12.8 The Geometry of Paradigm Change	54
12.9 The Recursive Discovery Principle	54
13 Compression, Parsimony, and Structural Minimality	56
13.1 The Principle of Structural Economy	56
13.2 Representational Complexity	56
13.3 Explanatory Compression	57
13.4 Minimal Presentations	57
14 Symmetry, Redundancy, and Canonical Form	58
14.1 Redundant Descriptions	58
14.2 Canonical Forms	58
15 The Principle of Recursive Scientific Realism	59
15.1 Reality Through Repair	59

Chapter 1

Autonomous Cognitive Domains and Admissible Interfaces

1.1 Motivation

Many descriptions of creative, intellectual, or contemplative work invoke metaphors such as sanctuaries, fortresses, halls, monasteries, laboratories, or isolated kingdoms. Although such imagery is psychologically compelling, it obscures a deeper mathematical structure.

The purpose of this chapter is to replace metaphor with formalism.

We demonstrate that what is commonly perceived as an isolated sanctuary is better understood as an *autonomous cognitive domain*: an open dynamical system whose interfaces satisfy admissibility constraints. Such systems are not closed. They exchange information, resources, and computational work with their environment while selectively rejecting perturbations that increase structural entropy.

The resulting framework applies equally to biological organisms, software systems, scientific collaborations, research laboratories, distributed computing networks, and individual cognitive agents.

Consequently, architectural metaphors become manifestations of a more general mathematical object rather than primitive explanatory concepts.

1.2 State Spaces

Let

$$\mathcal{S}$$

denote the internal state space of an autonomous system.

Let

$$\mathcal{E}$$

denote its environment.

Neither space is assumed finite, compact, linear, or even metrizable unless explicitly stated.

The evolution of the internal configuration is represented by

$$\Phi_t : \mathcal{S} \rightarrow \mathcal{S},$$

while environmental evolution is represented by

$$\Psi_t : \mathcal{E} \rightarrow \mathcal{E}.$$

The interaction between the two is not direct.
Instead it is mediated by a boundary operator

$$\mathcal{B} : \partial\mathcal{S} \rightarrow \partial\mathcal{E},$$

whose admissibility properties determine which external influences are permitted to alter internal state.

Thus autonomy is not equivalent to isolation.

Autonomy is the existence of a mathematically constrained interface.

1.3 The Error of Closed-System Thinking

A common misconception is that sustained intellectual work requires complete disconnection from the external world.

Such an assumption is impossible for any computationally active system.

Every scientific institution, every software project, every biological organism, and every civilization continuously exchanges matter, energy, information, or computation with its surroundings.

Therefore complete closure implies eventual computational death.

The relevant distinction is therefore not

closed versus open,

but instead

indiscriminately coupled versus admissibly coupled.

This distinction forms the foundation of the present theory.

1.4 Admissible Interfaces

The defining property of an autonomous system is not impermeability but controlled interaction.

Every computational system capable of sustained existence exchanges information with its surroundings. The mathematical problem is therefore not how to eliminate interaction, but how to constrain it.

Definition 1.1 (Interface). Let $\mathcal{S}$ denote an autonomous state space and let $\mathcal{E}$ denote its environment.

An *interface* is a morphism

$$I : \mathcal{E} \longrightarrow \mathcal{S}$$

together with a corresponding outward morphism

$$O : \mathcal{S} \longrightarrow \mathcal{E},$$

describing the admissible exchange of information, computation, resources, or physical action.

The pair (I, O) completely characterizes the coupling between the system and its environment.

Consequently, the environment never acts directly upon the internal state. Every influence is mediated by an interface.

Definition 1.2 (Admissible Interface). An interface (I, O) is called *admissible* if every permitted exchange preserves the invariant family

$$\mathcal{J} = \{J_1, J_2, \dots, J_n\},$$

where each

$$J_i : \mathcal{S} \rightarrow \mathbb{R}$$

is an invariant functional defined on the internal state space.

Equivalently,

$$J_i(s) = J_i(I(e, s))$$

for every admissible interaction.

Thus admissibility is not defined by silence.

It is defined by invariant preservation.

A perfectly silent system may still destroy its own internal organization, while a highly communicative system may remain completely admissible provided its structural invariants remain unchanged.

1.5 Coupling Operators

Interactions are rarely binary.

Most systems permit certain forms of exchange while rejecting others.

This motivates the introduction of a coupling operator.

Definition 1.3 (Coupling Operator). A coupling operator is a bounded map

$$\Gamma : \mathcal{E} \times \mathcal{S} \rightarrow [0, 1],$$

assigning each possible interaction a coupling strength.

The extreme cases satisfy

$$\Gamma = 0$$

for complete rejection and

$$\Gamma = 1$$

for unrestricted acceptance.

Intermediate values describe partially admissible interactions.

Notice that complete decoupling is neither necessary nor desirable.

If

$$\Gamma \equiv 0,$$

then no external information may enter the system.

Learning becomes impossible.

Scientific collaboration disappears.

Repair cannot incorporate new evidence.

Likewise,

$$\Gamma \equiv 1$$

is equally pathological.

Every perturbation propagates directly into internal dynamics.

Noise becomes indistinguishable from signal.

Autonomy collapses into environmental forcing.

The mathematically interesting regime therefore satisfies

$$0 < \Gamma < 1,$$

where selective interaction between the system and its environment becomes possible.

Definition 1.4 (Selective Permeability). An autonomous system is said to be *selectively permeable* whenever its coupling operator is state dependent,

$$\Gamma = \Gamma(e, s),$$

so that admissibility depends jointly upon environmental structure and the system's internal configuration.

Selective permeability is substantially more general than the architectural metaphor of a wall. Whereas a wall functions as a passive barrier that merely blocks or permits passage according to fixed physical constraints, an admissible interface acts as an active evaluative boundary. Rather than enforcing exclusion for its own sake, it computes whether proposed interactions satisfy the governing admissibility conditions and selectively permits only those transitions that preserve the system's long-term structural integrity. Its primary purpose is therefore not obstruction but the maintenance of coherence, stability, and persistence through principled evaluation.

Proposition 1.5. *Every computationally active autonomous system possessing nontrivial learning dynamics necessarily admits a nonzero coupling operator.*

Proof. Assume, for contradiction, that

$$\Gamma \equiv 0.$$

Then no information from the environment can enter the system. Consequently, the state evolution is completely determined by the internal dynamics,

$$s_{t+1} = \Phi(s_t),$$

independently of environmental observations. Since no external evidence can influence the system's evolution, Bayesian updating, statistical inference, parameter estimation, repair, and learning are all impossible. This contradicts the assumption that the system exhibits nontrivial adaptive dynamics. Therefore every adaptive computational system necessarily satisfies

$$\Gamma \neq 0.$$

□

The proposition identifies nonvanishing coupling as a necessary condition for adaptive computation. A system with $\Gamma = 0$ is informationally closed and therefore incapable of incorporating evidence from its environment, whereas a system with $\Gamma > 0$ admits the possibility of observation, inference, learning, and repair. Selective permeability thus characterizes the regime in which an autonomous system can remain structurally coherent while continuing to adapt to changing external conditions.

This result establishes an important conceptual distinction between adaptive computational systems and degenerate systems lacking a distinguishable informational geometry. When $\Gamma = 0$, all configurations become informationally indistinguishable, preventing any meaningful process of inference, learning, or repair. In contrast, $\Gamma \neq 0$ endows the system with a nontrivial geometry of distinguishability, allowing evidence to discriminate among competing hypotheses, parameters to be estimated, and adaptive corrections to reduce reconstruction error. Thus, within the present framework, nonvanishing distinguishability is not merely a technical assumption but a necessary structural condition for admissible computation.

Chapter 2

Structural Entropy and Environmental Perturbation

2.1 Noise Is Not an Environmental Property

Informal discussions frequently distinguish between environments that are “peaceful” and those that are “noisy.” Such language is useful psychologically but lacks mathematical precision.

Within the present framework, noise is not an intrinsic property of the environment.

Rather, it is a property of the interaction between the environment and the state of the autonomous system.

Consequently identical external events may constitute severe perturbations for one system while remaining dynamically negligible for another.

Noise is therefore relational rather than absolute.

Definition 2.1 (Environmental Perturbation). Let

$$e \in \mathcal{E}$$

be an environmental event.

The perturbation induced by e upon state s is

$$P(e, s) = \Gamma(e, s) \Delta(e, s),$$

where

$$\Gamma$$

is the coupling operator introduced previously and

$$\Delta$$

measures the structural displacement produced by the interaction.

Notice that an event may possess arbitrarily large magnitude while producing arbitrarily small perturbation whenever

$$\Gamma(e, s) \approx 0.$$

Likewise, even relatively minor events may produce significant structural distortion whenever coupling approaches unity.

Thus disturbance cannot be identified solely with external intensity.

2.2 Structural Entropy

The cumulative effect of perturbations is quantified through a structural entropy functional.

Definition 2.2 (Structural Entropy). Let

$$\mathcal{H} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$$

measure the degree of internal organizational uncertainty.

For every admissible evolution

$$\Phi_t,$$

the entropy increment is

$$\Delta \mathcal{H} = \mathcal{H}(\Phi_t(s)) - \mathcal{H}(s).$$

Unlike thermodynamic entropy, structural entropy need not measure disorder in a physical sense.

Instead it quantifies degradation of the organizational relations required for coherent computation.

Consequently highly dynamic systems may possess extremely low structural entropy, whereas apparently motionless systems may accumulate structural entropy rapidly.

The relevant quantity is organization rather than activity.

2.3 Signal

If perturbation is relational, then signal must also be defined relationally.

Definition 2.3 (Signal). An interaction

$$e \in \mathcal{E}$$

is called a signal whenever

$$\Delta \mathcal{H} \leq 0$$

or

$$\Delta \mathcal{J} \geq 0,$$

where

$$\mathcal{J}$$

denotes the invariant family introduced previously.

Signals therefore either reduce structural entropy or strengthen invariant preservation.

This definition immediately separates usefulness from novelty.

Novel information need not constitute signal.

Likewise, repeated observations may continue functioning as signal whenever they reinforce organizational stability.

2.4 Noise

Noise is defined as the complementary class.

Definition 2.4 (Noise). An interaction is noise whenever

$$\Delta\mathcal{H} > 0$$

while producing no compensating increase in invariant preservation, computational capability, or repair potential.

Noise therefore possesses no independent existence.

It is simply environmental interaction whose informational cost exceeds its structural benefit.

Remark 2.5. This definition differs fundamentally from communication theory.

Classical information theory measures uncertainty in transmitted symbols.

The present framework measures degradation of autonomous organizational structure.

The two notions coincide only under additional assumptions.

2.5 The Filtering Principle

An admissible interface should therefore reject neither the largest nor the most frequent perturbations.

Instead it minimizes the expected increase in structural entropy.

Formally,

$$\Gamma^* = \arg \min_{\Gamma} \mathbb{E} [\Delta\mathcal{H}],$$

subject to the constraints

$$\mathcal{J}_i = \text{constant},$$

for every preserved invariant.

This optimization criterion replaces numerous informal concepts including “protecting one’s peace,” “building walls,” or “remaining untouched.”

The objective is neither silence nor isolation.

The objective is entropy minimization under admissibility constraints.

Theorem 2.6 (Entropy-Admissibility Principle). *Suppose an autonomous system remains computationally active while satisfying*

$$\Delta\mathcal{H} \leq 0$$

for every admissible interaction.

Then every infinite admissible trajectory remains structurally stable.

Proof. Since structural entropy is bounded below by zero and is non-increasing along every admissible trajectory,

$$\{\mathcal{H}(s_n)\}$$

forms a monotone bounded sequence.

Hence

$$\lim_{n \rightarrow \infty} \mathcal{H}(s_n)$$

exists.

Because invariant preservation prevents catastrophic structural collapse, the limiting trajectory remains admissible for all finite time.

Therefore long-term computational stability follows.

□

The conclusion is notable.

A sanctuary is not characterized by the absence of disturbance.

Rather, it is characterized by the existence of interface operators that prevent disturbance from accumulating into irreversible structural entropy.

Chapter 3

Repair Operators and Recursive Homeostasis

3.1 Home Is Not a Location

The common identification of home with a geographical location, building, or social group is mathematically inadequate.

Any location may lose its organizational coherence.

Conversely, a computational system may preserve internal identity while undergoing continuous spatial relocation.

Consequently, home cannot be identified with position in Euclidean space.

Instead it must be regarded as an invariant subset of admissible configurations.

Definition 3.1 (Home). Let

$$\text{Adm} \subseteq \mathcal{S}$$

denote the admissible manifold.

The *home* of an autonomous system is the maximal positively invariant subset

$$\mathcal{H}_{\text{ome}} \subseteq \text{Adm}$$

such that

$$\Phi_t(s) \in \mathcal{H}_{\text{ome}}$$

for every

$$t \geq 0$$

whenever

$$s \in \mathcal{H}_{\text{ome}}.$$

Thus home is not a coordinate.

Home is an invariant region of state space.

A room may contain home.

A laboratory may contain home.
 A distributed computational network may contain home.
 None of these physical realizations is itself home.
 They merely instantiate admissible regions.

3.2 Deviation from Admissibility

No autonomous system remains permanently within the admissible manifold.

Unexpected observations, failures, interruptions, and incomplete information continuously perturb the trajectory.

Consequently repair is unavoidable.

Definition 3.2 (Defect). The defect of a state

$$s$$

is measured by

$$\delta(s) = d(s, \text{Adm}),$$

where

$$d$$

is any admissible metric or pseudometric compatible with the topology of $\mathcal{S}$.

States satisfying

$$\delta(s) = 0$$

are admissible.

Positive values quantify the degree of structural displacement.

Importantly,

$$\delta$$

measures neither moral failure nor psychological distress.

It measures geometric distance from invariant organization.

3.3 Repair Operators

The central dynamical object of the theory is the repair operator.

Definition 3.3 (Repair Operator). A repair operator is a mapping

$$R : \mathcal{S} \rightarrow \mathcal{S}$$

satisfying

$$\delta(R(s)) \leq \delta(s)$$

for every state

$$s.$$

If the inequality is strict whenever

$$\delta(s) > 0,$$

the repair operator is called *strictly restorative*.

Repair therefore need not return the system to its previous state.

Its only mathematical obligation is to decrease structural defect.

History is not reversed.

Admissibility is recovered.

Remark 3.4. Repair is therefore fundamentally different from restoration.

Restoration attempts to recreate the past.

Repair constructs an admissible future.

3.4 Recursive Repair

Complex systems rarely admit complete repair in a single iteration.

Instead repair proceeds recursively.

Define

$$R^{(0)} = \text{id},$$

and recursively

$$R^{(n+1)} = R \circ R^{(n)}.$$

The resulting sequence

$$\{R^{(n)}(s)\}$$

defines a repair trajectory.

Definition 3.5 (Recursive Homeostasis). An autonomous system exhibits recursive homeostasis whenever

$$\lim_{n \rightarrow \infty} \delta(R^{(n)}(s)) = 0$$

for every admissible initial defect.

Homeostasis therefore emerges as repeated approximation rather than static equilibrium.

The system continually reconstructs itself.

3.5 Home as an Attractor

The previous definitions imply a stronger characterization.

Theorem 3.6 (Home Attractor Theorem). *Suppose*

1. *the repair operator is continuous;*
2. *the admissible manifold is closed;*
3. *repair is strictly restorative outside Adm.*

Then

$$\mathcal{H}_{\text{ome}}$$

is an attractor of the repair dynamics.

Proof. Since repair decreases defect monotonically,

$$\delta(R^{(n+1)}(s)) < \delta(R^{(n)}(s))$$

whenever

$$\delta > 0.$$

Thus

$$\{\delta_n\}$$

forms a bounded monotone sequence.

Therefore

$$\delta_n \rightarrow L \geq 0.$$

If

$$L > 0,$$

strict restoration would imply another decrease, contradicting convergence.

Hence

$$L = 0.$$

Since the admissible manifold is closed,

$$R^{(n)}(s) \rightarrow \mathcal{H}_{\text{ome}}.$$

Therefore home is an attracting invariant set.

□

The consequence is subtle but important.

Home is not where a system begins.

Home is where recursive repair converges.

The distinction removes any dependence upon geography, ancestry, architecture, or narrative symbolism.

Home becomes a purely dynamical object arising from invariant-preserving repair.

The remainder of this monograph develops the consequences of treating autonomous cognition, scientific communities, biological organisms, and distributed computational systems as recursive repair processes evolving toward admissible attractors rather than static equilibrium states.

Chapter 4

The Geometry of Collaboration

4.1 Against the Myth of Isolation

A persistent misconception in both popular psychology and technical culture is that creativity is fundamentally an isolated phenomenon.

This misconception arises from observing the localization of computation while neglecting the topology of information exchange.

Although the internal computation of an autonomous system occurs within its own state space, the informational substrate upon which that computation depends is inherently distributed.

No theorem proves itself.

No scientific discipline constructs its own language.

No programming language implements its own compiler without inherited infrastructure.

Autonomy therefore cannot imply informational independence.

Instead, autonomy concerns the preservation of internal organizational structure while participating in a distributed computational ecology.

4.2 Knowledge Graphs

Scientific activity is naturally represented by a graph.

Definition 4.1 (Knowledge Graph). A knowledge graph is a directed graph

$$\mathcal{G} = (V, E),$$

whose vertices represent autonomous cognitive systems and whose directed edges represent admissible informational exchange.

An edge

$$u \rightarrow v$$

exists whenever information originating at

$$u$$

becomes structurally incorporated into the admissible state of

v .

Edges need not correspond to direct communication.

Books, software, published proofs, observations, experimental apparatus, historical traditions, and mathematical notation all induce admissible edges within the graph.

Consequently scientific collaboration is substantially richer than interpersonal communication.

4.3 Indirect Coupling

Many collaborative systems exhibit no direct interaction whatsoever.

Instead they communicate through persistent informational artifacts.

Definition 4.2 (Indirect Coupling). Two autonomous systems

A

and

B

are indirectly coupled whenever there exists an admissible sequence

$$A \rightarrow v_1 \rightarrow \cdots \rightarrow v_n \rightarrow B.$$

The sequence need not preserve temporal simultaneity.

Thus Euclid remains coupled to modern mathematics.

Newton remains coupled to orbital mechanics.

Shannon remains coupled to every communication protocol derived from information theory.

The coupling persists despite the absence of temporal coexistence.

Knowledge propagates through graph topology rather than chronology.

4.4 Structural Independence

Informal discussions often confuse physical proximity with intellectual dependence.

The two concepts are mathematically independent.

Definition 4.3 (Structural Independence). An autonomous system

S

is structurally independent whenever every admissible trajectory satisfies

$$\frac{\partial \Phi}{\partial e}$$

bounded independently of transient environmental perturbations.

Structural independence therefore concerns bounded sensitivity.

It does not imply absence of communication.

Indeed, the most productive scientific systems frequently maintain large numbers of admissible informational edges while remaining structurally stable.

4.5 Distributed Repair

Repair itself is rarely local.

Suppose

$$\mathcal{G} = (V, E)$$

is a knowledge graph.

Associate to every vertex

$$v$$

a repair operator

$$R_v.$$

The global repair dynamics become

$$R_{\mathcal{G}} = \prod_{v \in V} R_v,$$

where composition follows the partial order induced by graph topology.

Consequently local repair may propagate throughout the network without any single vertex possessing complete knowledge.

Scientific communities therefore behave as distributed repair systems.

Each participant repairs only a fragment of the total structure.

The resulting global reconstruction exceeds the capability of every individual contributor.

4.6 Emergent Intelligence

The previous observations motivate a more general definition.

Definition 4.4 (Emergent Intelligence). Let

$$\mathcal{G} = (V, E)$$

be a knowledge graph.

Emergent intelligence is the increase in repair capacity obtained through graph composition:

$$\mathcal{R}(\mathcal{G}) > \sum_{v \in V} \mathcal{R}(v),$$

where

$$\mathcal{R}$$

denotes the repair functional.

The inequality expresses genuine emergence.

The graph performs computational work that cannot be attributed to any individual vertex.

This phenomenon appears throughout mathematics, scientific collaboration, distributed software engineering, biological ecosystems, and neural systems.

4.7 The Fallacy of the Solitary Genius

The mythology of the solitary genius arises from observing only localized computation.

It ignores the graph upon which that computation depends.

Every theorem presupposes notation.

Every notation presupposes language.

Every language presupposes a historical community.

Every computation presupposes inherited infrastructure.

Therefore no autonomous system is graph-independent.

The proper mathematical object is not the isolated individual but the distributed repair graph within which individual cognition occupies a single vertex.

Autonomy consists not in severing edges but in maintaining admissible edges whose cumulative dynamics preserve structural invariants while enabling recursive repair across the graph.

Chapter 5

Persistence, Continuation, and Identity

5.1 Identity Is Not Substance

Classical metaphysical accounts frequently regard identity as an intrinsic property of an object. Such approaches implicitly assume that persistence is explained by the continued existence of an underlying substance.

The present framework adopts a fundamentally different viewpoint.

Identity is not primitive.

Identity is an emergent invariant generated by recursively admissible continuation.

Consequently, persistence is not explained by substance.

Substance is explained by persistence.

5.2 Continuation Operators

Every autonomous system evolves through time.

Let

$$\Phi_t : \mathcal{S} \rightarrow \mathcal{S}$$

denote the temporal evolution operator introduced previously.

Not every temporal evolution preserves identity.

We therefore distinguish arbitrary evolution from admissible continuation.

Definition 5.1 (Continuation Operator). A continuation operator is a family

$$C_t : \mathcal{S} \rightarrow \mathcal{S}$$

satisfying

$$C_0 = \text{id}$$

and

$$C_{t+s} = C_t \circ C_s$$

while preserving the invariant family

$\mathcal{J}$.

The semigroup property expresses temporal coherence.

Identity therefore depends not upon static resemblance but upon lawful continuation.

5.3 Persistence

Persistence admits a quantitative characterization.

Definition 5.2 (Persistence Functional). Let

$$P : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$$

be defined by

$$P(s) = \int_0^\infty w(t) \chi_{\text{Adm}}(C_t(s)) dt,$$

where

$$w(t)$$

is a nonnegative weighting function and

$$\chi_{\text{Adm}}$$

is the characteristic function of the admissible manifold.

Persistence therefore measures the cumulative duration over which admissible continuation remains possible.

It is neither age nor longevity.

Rather, it measures structural survivability.

5.4 Identity Through Repair

Suppose perturbation produces

$$s' \notin \text{Adm}.$$

Repair generates

$$R(s').$$

The repaired state generally differs from the original.

Nevertheless identity may remain preserved.

Definition 5.3 (Identity Equivalence). Two states

$$s_1, s_2$$

are identity equivalent whenever there exists a finite sequence

$$s_1 \rightarrow R_1 \rightarrow R_2 \rightarrow \cdots \rightarrow R_n \rightarrow s_2$$

such that every intermediate state belongs to the admissible manifold.
We write

$$s_1 \sim s_2.$$

Identity therefore becomes an equivalence relation generated by admissible repair rather than perfect state equality.

This distinction is essential.

Exact equality almost never survives long temporal evolution.

Admissible continuation frequently does.

5.5 The Failure of Static Ontologies

Suppose identity required

$$s_t = s_0$$

for every

$$t.$$

Then every nontrivial evolution would immediately destroy identity.

Biological development would constitute replacement.

Learning would constitute replacement.

Scientific progress would constitute replacement.

Computation itself would constitute replacement.

This conclusion is absurd.

Static identity therefore cannot provide an adequate ontology for evolving systems.

5.6 The Continuation Principle

The previous constructions motivate the central principle.

Theorem 5.4 (Continuation Principle). *Let*

$$\{C_t\}_{t \geq 0}$$

be a family of admissible continuation operators.

Suppose

$$R$$

*restores every finite defect while preserving invariant equivalence.
Then identity is preserved along every recursively repaired continuation trajectory.*

Proof. Every admissible continuation preserves the invariant family

$$\mathcal{I}.$$

Whenever perturbation produces temporary inadmissibility,

$$R$$

returns the trajectory to the admissible manifold without destroying equivalence.
Since composition of admissible continuation with admissible repair remains admissible,

$$C_t \circ R$$

preserves the equivalence class generated by

$$\sim .$$

Repeated application therefore constructs an infinite admissible trajectory contained entirely within a single identity class.

□

5.7 Identity as a Fixed Structure

The previous theorem reverses a common intuition.

Identity is not the object that persists through change.

Identity is the invariant generated by lawful change.

Objects are transient realizations.

Identity is the stable geometry of their admissible continuation.

Accordingly, questions concerning the persistence of persons, institutions, languages, scientific theories, computational systems, and biological organisms reduce to a single mathematical problem:

Does there exist a recursively admissible continuation operator preserving the relevant invariant family?

The remainder of this work demonstrates that this question admits a common answer across seemingly unrelated domains, thereby replacing diverse metaphysical accounts of persistence with a unified theory of recursive continuation.

Chapter 6

Observers, Interfaces, and Representation

6.1 The Impossibility of Direct Observation

Every observational theory begins with an implicit assumption regarding the relationship between an observer and the external world.

The simplest assumption is direct realism:

$$\mathcal{O} = \mathcal{W},$$

where observation is identified with reality itself.

Such an assumption is untenable.

No observer possesses unrestricted access to the complete external state of its environment.

Every observation is mediated by finite sensors, finite computation, measurement uncertainty, temporal delay, and representational constraints.

Consequently all observation is indirect.

The observer never receives reality.

The observer receives an interface.

6.2 Interface States

Let

$$\mathcal{W}$$

denote the external world.

Let

$$\mathcal{I}$$

denote the interface accessible to an observer.

The observer itself possesses internal state space

$$\mathcal{S}.$$

Observation therefore decomposes into two mappings

$$\Pi : \mathcal{W} \rightarrow \mathcal{I},$$

and

$$\mathcal{O} : \mathcal{I} \rightarrow \mathcal{S}.$$

The first mapping represents physical measurement.

The second represents interpretation.

Reality therefore satisfies the composition

$$\mathcal{W} \rightarrow \mathcal{I} \rightarrow \mathcal{S},$$

rather than

$$\mathcal{W} \rightarrow \mathcal{S}.$$

The distinction is fundamental.

Errors introduced by the interface cannot be distinguished from properties of the external world unless additional structural assumptions are available.

6.3 Representation

The observer does not manipulate reality directly.

Instead it manipulates internal representations.

Definition 6.1 (Representation). A representation is a mapping

$$\rho : \mathcal{I} \rightarrow \mathcal{R},$$

where

$$\mathcal{R}$$

is an internal representational manifold.

Reasoning, planning, prediction, memory, and decision-making occur within

$$\mathcal{R},$$

not within

$$\mathcal{W}.$$

Representations therefore possess operational reality even when they are imperfect descriptions of the external world.

Scientific models, programming languages, maps, mathematical notation, neural activity, and symbolic logic all constitute representational systems.

6.4 Information Loss

Every interface necessarily discards information.

Definition 6.2 (Projection). An observational interface is a projection

$$\Pi : \mathcal{W} \rightarrow \mathcal{J}$$

whenever

$$\dim(\mathcal{J}) < \dim(\mathcal{W}),$$

or, more generally, whenever distinct external configurations possess identical interface states.

Consequently

$$\Pi$$

cannot generally admit an inverse.

Different external realities may generate identical observations.

Observation therefore determines an equivalence class rather than a unique world state.

6.5 Observational Equivalence

Definition 6.3 (Observational Equivalence). Two external configurations

$$w_1, w_2 \in \mathcal{W}$$

are observationally equivalent whenever

$$\Pi(w_1) = \Pi(w_2).$$

We write

$$w_1 \equiv_{\Pi} w_2.$$

The observer cannot distinguish members of the same equivalence class using the given interface alone.

Additional measurements, alternative interfaces, or prior structural knowledge become necessary.

Thus ignorance is not always a failure of reasoning.

It may simply reflect insufficient observational resolution.

6.6 Reconstruction

Scientific inference attempts to reconstruct the external world from interface observations.

This inversion is generally ill-posed.

Definition 6.4 (Reconstruction Operator). A reconstruction operator is a mapping

$$\mathcal{R} : \mathcal{I} \rightarrow 2^{\mathcal{W}},$$

assigning every interface state the admissible family of compatible external configurations.

Notice that

$$\mathcal{R}(i)$$

is typically not a singleton.

Observation constrains reality without uniquely determining it.

Inference therefore consists of progressively shrinking admissible equivalence classes.

6.7 Representation Defect

The quality of a representation depends upon its ability to preserve relevant structure.

Definition 6.5 (Representation Defect). Let

$$\rho$$

be an internal representation.

The representation defect is the functional

$$\Delta_\rho = d(\Pi(\mathcal{W}), \rho(\mathcal{I})),$$

where

$$d$$

is any admissible structural metric.

Perfect representation satisfies

$$\Delta_\rho = 0.$$

In practice,

$$\Delta_\rho > 0$$

for every finite observer.

Scientific progress therefore consists not in eliminating representation defect but in recursively reducing it.

6.8 The Representation Principle

The preceding constructions yield a general theorem.

Theorem 6.6 (Representation Principle). *No finite observer possesses direct access to external reality. Every autonomous cognitive system operates exclusively upon internal representations generated through admissible observational interfaces. Consequently every act of reasoning is necessarily reconstructive rather than directly perceptual.*

Proof. Observation is mediated by

$$\Pi : \mathcal{W} \rightarrow \mathcal{I}.$$

Reasoning operates only upon

$$\rho(\mathcal{I}).$$

Since

$$\Pi$$

generally lacks an inverse, direct recovery of

$$\mathcal{W}$$

is impossible.

Therefore cognition necessarily proceeds by constructing admissible reconstructions whose representation defect decreases through recursive repair.

□

The consequences extend well beyond epistemology.

Scientific theories, software systems, biological perception, artificial intelligence, language, memory, and mathematical abstraction all become instances of a common representational architecture.

Reality itself is not manipulated directly.

Only admissible representations evolve under recursive continuation while remaining constrained by observational interfaces.

Chapter 7

Distinction Geometry and the Topology of Information

7.1 Information Is Not Primitive

Traditional information theory begins by assuming the existence of symbols, messages, or random variables.

The present framework begins earlier.

Prior to every message, prior to every probability distribution, and prior to every symbolic language lies a more fundamental object:
the distinction.

Without distinguishable configurations, no observation, measurement, communication, or computation can exist.

Consequently information is not primitive.

Information is induced by distinguishability.

7.2 The Distinction Space

Let

$$\mathcal{X}$$

be a configuration space.

We define an equivalence relation

$$\sim$$

representing observational indistinguishability.

The corresponding quotient space

$$\mathcal{D} = \mathcal{X} / \sim$$

is called the distinction space.

Its elements are equivalence classes

$$[x] = \{y \in \mathcal{X} : y \sim x\}.$$

Every observation identifies an element of

$$\mathcal{D},$$

not necessarily an element of

$$\mathcal{X}.$$

Thus cognition operates upon distinction classes rather than absolute states.

7.3 Distinction Fields

Distinctions are rarely isolated.

Instead they organize into continuous structures.

Definition 7.1 (Distinction Field). A distinction field is a mapping

$$\Delta : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$$

whose value measures the local distinguishability surrounding a point.

Large values correspond to regions of high observational separation.

Small values indicate observational collapse.

Unlike probability densities, distinction fields describe geometric resolution rather than uncertainty.

They measure how easily neighboring configurations may be separated by an admissible observer.

7.4 Observational Collapse

Finite interfaces cannot preserve arbitrary distinctions.

Projection necessarily collapses certain regions.

Definition 7.2 (Collapse Region). A subset

$$U \subseteq \mathcal{X}$$

is called a collapse region whenever

$$\Pi(U) = \{i\}$$

for some interface state

$$i \in \mathcal{I}.$$

Within a collapse region every configuration becomes observationally equivalent.

No internal computation can distinguish them using the given interface.

Collapse therefore reflects finite observational resolution rather than loss of objective structure.

7.5 Distinction Entropy

The richness of an observational system depends upon the diversity of distinctions it preserves.

Definition 7.3 (Distinction Entropy). Let

$$\mu$$

be a measure on

$$\mathcal{D}.$$

The distinction entropy is the functional

$$\mathfrak{D} = \int_{\mathcal{D}} \Delta(x) d\mu(x).$$

This quantity measures the total distinguishable structure available to an observer.

Unlike Shannon entropy, distinction entropy does not quantify uncertainty.

Instead it quantifies representational resolution.

Two observers possessing identical Shannon information may possess different distinction entropies if one preserves finer geometric structure than the other.

7.6 Repair of Distinctions

Repair does not merely restore admissibility.

It restores distinguishability.

Suppose perturbation merges two previously distinguishable regions.

The repair operator acts by increasing local distinction density.

Formally,

$$R : \Delta \longrightarrow \Delta',$$

where

$$\Delta' \geq \Delta$$

whenever structural distinctions have been degraded.

Thus repair reconstructs the geometry upon which subsequent computation depends.

7.7 Distinction Curvature

Local distinguishability may vary continuously.

This motivates a differential description.

Definition 7.4 (Distinction Curvature). The distinction curvature is the tensor

$$\mathcal{K} = \nabla^2 \Delta,$$

whenever the distinction field is twice differentiable.

Regions of positive curvature correspond to rapidly increasing observational resolution.
 Regions of negative curvature correspond to distinction collapse.
 Flat regions represent observational stability.

Thus the geometry of cognition may be analyzed using the same differential tools employed throughout modern geometry.

7.8 The Distinction Principle

The preceding constructions yield a general characterization of information.

Theorem 7.5 (Distinction Principle). *Every informational process factors through distinction geometry. Specifically,*

$$\text{Information} = \text{Observation} \circ \text{Distinction}.$$

Consequently, whenever distinguishability vanishes, information necessarily vanishes regardless of computational complexity.

Conversely, increasing computational power cannot recover distinctions that have already been destroyed by projection.

Proof. Observation first partitions the external configuration space into equivalence classes. Internal computation operates only upon these classes. Therefore every subsequent informational process is defined upon the distinction space

$$\mathcal{D},$$

rather than upon the original configuration space.

Hence distinguishability is logically prior to information. □

The theory developed in this chapter reverses the conventional hierarchy.

Classical information theory regards distinctions as consequences of encoded messages.

The present framework regards encoded messages as consequences of pre-existing distinguishable structure.

Information therefore emerges from geometry rather than geometry emerging from information.

This inversion provides a common mathematical foundation for perception, scientific inference, communication theory, artificial intelligence, and recursive repair.

Chapter 8

Constraint Geometry and Reachability

8.1 Constraint Precedes Motion

Classical mechanics begins by specifying trajectories.

Optimization begins by specifying objective functions.

Control theory begins by specifying admissible inputs.

The present framework begins earlier.

Before any trajectory exists, before any optimization may be attempted, and before any computation may occur, there exists a geometry of possibility.

That geometry is determined entirely by constraints.

Motion therefore does not define constraints.

Constraints define motion.

8.2 Constraint Manifolds

Let

$$\mathcal{S}$$

denote the complete configuration space.

Only a subset corresponds to physically, computationally, or logically realizable configurations.

Definition 8.1 (Constraint Manifold). A constraint manifold is a subset

$$\mathcal{C} \subseteq \mathcal{S}$$

defined by a family of admissibility relations

$$F_i(s) = 0,$$

together with inequality constraints

$$G_j(s) \geq 0.$$

The manifold consists precisely of those configurations satisfying every constraint simultaneously.

Every autonomous system evolves entirely within

$$\mathcal{C},$$

except during transient perturbations subsequently corrected by repair.

Thus constraints do not oppose dynamics.

They define the domain upon which dynamics exist.

8.3 Reachability

Not every admissible configuration can necessarily evolve into every other.

This observation motivates the central geometric object.

Definition 8.2 (Reachability). Given two admissible states

$$s_0, s_1 \in \mathcal{C},$$

we say that

$$s_1$$

is reachable from

$$s_0$$

whenever there exists a continuous admissible trajectory

$$\gamma : [0, 1] \rightarrow \mathcal{C}$$

such that

$$\gamma(0) = s_0, \quad \gamma(1) = s_1.$$

Reachability is therefore a topological property rather than merely a kinematic one.

The existence of a destination does not imply the existence of an admissible path.

8.4 Reachability Components

Reachability induces an equivalence relation upon admissible states.

Definition 8.3 (Reachability Component). The reachability component of

$$s$$

is

$$\mathcal{R}(s) = \{x \in \mathcal{C} : x \leftrightarrow s\},$$

where

$$\leftrightarrow$$

denotes mutual reachability.

Thus

$$\mathcal{C} = \bigsqcup_i \mathcal{R}_i,$$

decomposes into disjoint connected components.

Each component represents an autonomous region of possible evolution.

Repair may move trajectories within a component.

Structural innovation may enlarge a component.

Constraint modification may merge components.

8.5 Boundary Geometry

The boundary of a reachability component possesses particular significance.

Definition 8.4 (Reachability Boundary). The reachability boundary is

$$\partial\mathcal{R} = \overline{\mathcal{R}} \setminus \text{int}(\mathcal{R}).$$

Near the boundary, arbitrarily small perturbations may destroy admissibility.

Consequently boundary states exhibit maximal structural sensitivity.

Conversely, interior states possess robustness under perturbation.

This distinction suggests that stable autonomous systems naturally evolve toward interior regions rather than remaining arbitrarily close to constraint boundaries.

8.6 Constraint Curvature

Constraints possess local geometry.

Suppose

$$\mathcal{C}$$

is a differentiable manifold.

The second fundamental form determines local bending relative to the ambient configuration space.

Regions of large curvature correspond to rapidly changing admissibility.

Regions of small curvature admit comparatively stable continuation.

Thus computational difficulty frequently arises not from dimensionality but from curvature.

Highly curved constraint surfaces require continual corrective action to remain admissible.

8.7 Constraint Energy

Deviation from admissibility admits a scalar potential.

Definition 8.5 (Constraint Energy). The constraint energy is a functional

$$E_C : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$$

satisfying

$$E_C(s) = 0$$

if and only if

$$s \in \mathcal{C},$$

and

$$E_C(s) > 0$$

otherwise.

Repair operators may therefore be interpreted as gradient-like flows

$$\frac{ds}{dt} = -\nabla E_C(s),$$

whenever differentiability assumptions hold.

Repair becomes geometric descent toward admissibility.

8.8 The Reachability Principle

The previous constructions yield the central theorem of this chapter.

Theorem 8.6 (Reachability Principle). *The behavior of an autonomous system is determined less by the collection of states it may occupy than by the topology of admissible paths connecting those states.*

Proof. State descriptions alone specify isolated configurations.

Autonomous evolution requires continuous admissible trajectories.

Consequently every achievable computation, biological process, scientific discovery, or cognitive transition depends upon the existence of paths within the constraint manifold.

Removing paths while preserving states destroys behavior without altering the underlying state space.

Hence reachability constitutes the primary dynamical invariant.

□

This theorem shifts emphasis from configuration to navigation.

The relevant scientific question is no longer

“What states exist?”

but instead

“Which admissible states can be reached, repaired, continued, and preserved without violating structural constraints?”

Accordingly, dynamics becomes the geometry of admissible navigation.

Constraints determine topology.

Topology determines reachability.

Reachability determines computation.

Computation determines the observable evolution of autonomous systems.

The remainder of this work develops this hierarchy into a unified theory of physical, biological, cognitive, and computational organization.

Chapter 9

The Geometry of Constraint Propagation

9.1 Local Constraints and Global Structure

A fundamental feature of complex systems is that global organization need not be imposed globally.

Instead, large-scale coherence frequently emerges from purely local constraints.

This phenomenon appears throughout mathematics and science.

Local differential equations generate global trajectories.

Local compatibility conditions generate smooth manifolds.

Local rewriting rules generate formal languages.

Local physical interactions generate biological organisms.

The common mathematical principle is therefore not global construction but constraint propagation.

9.2 Constraint Networks

Individual constraints rarely exist in isolation.

Instead they form interconnected systems.

Definition 9.1 (Constraint Network). Let

$$V = \{C_1, C_2, \dots, C_n\}$$

be a finite or countable family of constraints.

A constraint network is a graph

$$\mathcal{N} = (V, E),$$

where

$$(C_i, C_j) \in E$$

whenever satisfaction of one constraint influences the admissibility of the other.

The graph need not be complete.

Indeed, sparse constraint graphs frequently exhibit greater computational efficiency while preserving equivalent admissible solution sets.

9.3 Constraint Compatibility

Not every family of constraints admits simultaneous realization.

Definition 9.2 (Compatibility). A collection of constraints

$$\{C_i\}$$

is compatible whenever

$$\bigcap_i C_i \neq \emptyset.$$

Otherwise the collection is incompatible.

Compatibility therefore represents the existence of at least one admissible configuration satisfying every constraint simultaneously.

Incompatibility implies structural inconsistency rather than computational difficulty.

No algorithm may realize an impossible configuration.

9.4 Propagation Operators

Information acquired locally should modify neighboring constraints.

Definition 9.3 (Propagation Operator). A propagation operator

$$P : \mathcal{N} \rightarrow \mathcal{N}$$

updates the admissible domains associated with every constraint by communicating local information throughout the constraint graph.

Propagation differs fundamentally from search.

Propagation removes impossible configurations.

Search chooses among remaining possibilities.

Thus propagation reduces complexity before decision.

9.5 Monotonicity

Constraint propagation should never enlarge the admissible search space.

Proposition 9.4. *Suppose*

$$D_i^{(k)}$$

denotes the admissible domain associated with constraint

$$C_i$$

after

$$k$$

propagation steps.
Then

$$D_i^{(k+1)} \subseteq D_i^{(k)}$$

for every

i .

Proof. Each propagation step removes assignments inconsistent with neighboring constraints.

No previously eliminated configuration can become admissible without modifying the constraint network itself.

Hence admissible domains decrease monotonically. □

Constraint propagation therefore possesses an intrinsic directionality.
Knowledge accumulates by eliminating impossibilities.

9.6 Fixed Points

Propagation eventually stabilizes.

Definition 9.5 (Propagation Fixed Point). A constraint network reaches a propagation fixed point whenever

$$P(\mathcal{N}) = \mathcal{N}.$$

Equivalently,

$$D_i^{(k+1)} = D_i^{(k)}$$

for every constraint.

At a fixed point, every remaining admissible configuration is locally consistent with every propagated constraint.

Further propagation alone produces no additional information.

Subsequent progress requires either new observations or explicit search.

9.7 Constraint Repair

Perturbations may invalidate previously established consistency.

Repair therefore acts upon the network itself.

Definition 9.6 (Constraint Repair). A repair operator

$$R_{\mathcal{N}} : \mathcal{N} \rightarrow \mathcal{N}$$

is admissible whenever it restores compatibility while minimizing structural modification of the original network.

Repair therefore preserves as much existing organization as possible.
 It does not reconstruct the network from first principles.
 Instead it performs minimal admissible correction.

9.8 Distributed Consistency

The interaction between propagation and repair produces an important principle.

Theorem 9.7 (Distributed Consistency Principle). *Suppose every local constraint satisfies compatibility with each of its neighbors after propagation.*

*Further suppose repair restores every finite incompatibility introduced by external perturbation.
 Then global consistency emerges as the least fixed point of alternating propagation and repair.*

Proof. Propagation monotonically reduces admissible domains while preserving every globally realizable configuration.

Repair restores compatibility whenever local inconsistencies arise.

Since admissible domains form descending chains bounded below by the empty set, propagation converges.

Repair prevents convergence to incompatible states.

Consequently the alternating composition

$$P \circ R_{\mathcal{N}}$$

approaches the least compatible fixed point containing all admissible solutions.

□

This theorem illustrates a recurring theme throughout the present work.

Global order need not be explicitly imposed.

Instead, sufficiently rich systems construct coherent large-scale structure through recursive interaction between local admissibility, propagation, repair, and continuation.

Accordingly, complexity is not generated by centralized control.

It emerges from recursively maintained consistency across distributed constraint networks.

Chapter 10

Sheaf-Theoretic Reconstruction and Global Consistency

10.1 The Local Nature of Knowledge

Every sufficiently complex observational process is fundamentally local.

No observer measures an entire physical system simultaneously.

No scientific discipline possesses complete knowledge.

No computational process evaluates every possible state at once.

Instead, knowledge is acquired through overlapping local observations whose mutual compatibility determines the existence of a coherent global model.

Consequently, the mathematical problem is not merely one of observation.

It is one of reconstruction.

10.2 Observation Covers

Let

$$\mathcal{X}$$

denote the external configuration manifold.

Suppose

$$\mathcal{U} = \{U_i\}_{i \in I}$$

is an open cover satisfying

$$\mathcal{X} = \bigcup_{i \in I} U_i.$$

Each

$$U_i$$

represents a region accessible to an individual observer, experiment, measurement device, computational subsystem, or scientific discipline.

No single element of the cover need contain complete information.
Completeness arises only through compatibility across intersections.

10.3 Local Sections

Observations assign internal descriptions to local regions.

Definition 10.1 (Local Section). A local section over

$$U_i$$

is a mapping

$$s_i : U_i \rightarrow \mathcal{R},$$

where

$$\mathcal{R}$$

denotes the representational space introduced previously.

The collection

$$\{s_i\}$$

therefore represents distributed partial knowledge rather than a complete description of the external manifold.

10.4 Compatibility

Independent observations need not contradict one another.

Their compatibility is determined on overlaps.

Definition 10.2 (Compatibility Condition). Two local sections

$$s_i, s_j$$

are compatible whenever

$$s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}.$$

Thus agreement is required only where observations overlap.

Outside the intersection, the sections remain independent.

Compatibility therefore constitutes a local condition.

10.5 The Gluing Principle

The central mathematical question becomes the following.

Given compatible local observations,
does there exist a unique global reconstruction?

Definition 10.3 (Global Section). A global section is a mapping

$$S : \mathcal{X} \rightarrow \mathcal{R}$$

whose restriction satisfies

$$S|_{U_i} = s_i$$

for every member of the cover.

Whenever such an

$$S$$

exists, the local observations admit coherent global interpretation.

10.6 Reconstruction Defect

Failure of reconstruction admits quantitative measurement.

Definition 10.4 (Reconstruction Defect). The reconstruction defect associated with an observation cover is

$$\Delta_{\text{rec}} = \sum_{i,j} d(s_i|_{U_i \cap U_j}, s_j|_{U_i \cap U_j}),$$

where

$$d$$

is an admissible metric on representational space.

Perfect reconstruction satisfies

$$\Delta_{\text{rec}} = 0.$$

Positive values indicate disagreement between neighboring observations.

Thus reconstruction becomes an optimization problem rather than a binary decision.

10.7 Repair by Gluing

Repair now acquires a geometric interpretation.

Rather than modifying isolated observations, repair minimizes overlap inconsistency. Formally,

$$R : \{s_i\} \longrightarrow \{s'_i\},$$

where

$$\Delta_{\text{rec}}(s'_i) < \Delta_{\text{rec}}(s_i).$$

Repair therefore increases global coherence by improving local compatibility.

Notice that no observer requires complete knowledge.

Each observer need only improve agreement with neighboring sections.

10.8 Distributed Truth

The preceding construction motivates an alternative conception of scientific truth.

Truth is not the possession of an individual observer.

Nor is it the consequence of a privileged viewpoint.

Instead, truth emerges whenever distributed local observations admit a stable global reconstruction.

Scientific collaboration therefore becomes an iterative gluing process.

Each experiment extends one section.

Each proof enlarges another.

Each computation refines a third.

Global understanding appears only through recursive compatibility.

10.9 The Local-to-Global Principle

The previous constructions culminate in the principal theorem.

Theorem 10.5 (Local-to-Global Reconstruction). *Suppose*

$$\mathcal{U} = \{U_i\}$$

is an observation cover.

Assume

1. *every local section is admissible,*
2. *neighboring sections satisfy compatibility,*
3. *recursive repair decreases reconstruction defect.*

Then repeated repair converges toward a global section whenever one exists.

Proof. Compatibility defines local consistency on every overlap.
Repair monotonically decreases

$$\Delta_{\text{rec}}.$$

Since

$$\Delta_{\text{rec}} \geq 0,$$

the reconstruction sequence possesses a lower bound.

Under continuity of repair and compactness of the admissible representational space, convergence follows.

The limiting family of sections therefore defines a global reconstruction compatible with every local observation.

□

This theorem establishes a profound inversion of classical epistemology.

Knowledge is not assembled by beginning with a complete global model and analyzing its parts.

Rather, coherent global structure emerges from recursively repairing relationships among overlapping local observations.

Accordingly, mathematics, scientific inquiry, distributed computation, biological cognition, and collaborative intelligence become instances of a single geometric process:

Local Observation $\rightarrow$ Compatibility $\rightarrow$ Repair $\rightarrow$ Global Reconstruction.

Chapter 11

Functorial Observation and the Category of Autonomous Systems

11.1 Beyond State Spaces

Previous chapters developed autonomous systems as state spaces equipped with admissibility, repair, continuation, and interface operators.

Although this formulation is sufficient for many applications, it conceals an important structural fact.

The essential object is not the state itself.

The essential object is the collection of transformations that preserve admissibility.

Consequently, the natural mathematical language is category theory.

11.2 The Category of Autonomous Systems

Definition 11.1 (Autonomous Category). Define the category

Aut

as follows.

Objects are autonomous systems

$$(\mathcal{S}, \Phi, \mathcal{J}),$$

where

1. $\mathcal{S}$ is a configuration space,
2. Φ denotes admissible evolution,
3. $\mathcal{J}$ is the invariant family.

Morphisms

$$f : A \rightarrow B$$

are admissibility-preserving transformations satisfying

$$f(\text{Adm}_A) \subseteq \text{Adm}_B.$$

Composition is ordinary function composition,

$$g \circ f,$$

while identity morphisms are

$$\text{id}_S.$$

Thus autonomous systems naturally form a category.

11.3 Observation as a Functor

Observation does not preserve every property of reality.

Instead it preserves only admissible structure.

This observation motivates a functorial description.

Definition 11.2 (Observation Functor). The observation process is a functor

$$\mathcal{O} : \mathbf{Aut} \longrightarrow \mathbf{Rep},$$

where

$$\mathbf{Rep}$$

denotes the category of internal representations.

The functor satisfies

$$\mathcal{O}(\text{id}) = \text{id},$$

and

$$\mathcal{O}(g \circ f) = \mathcal{O}(g) \circ \mathcal{O}(f).$$

Observation therefore preserves compositional structure while generally discarding geometric detail.

11.4 Forgetful Processes

Every observation loses information.

Categorically, this is expressed by forgetful functors.

Definition 11.3 (Forgetful Functor). A forgetful functor

$$U : \mathbf{Aut} \rightarrow \mathbf{Set}$$

maps every autonomous system to its underlying collection of states while discarding repair operators, admissibility relations, continuation dynamics, and invariant structure.

Many traditional scientific descriptions are implicitly forgetful.

They retain observable variables while discarding the morphisms responsible for long-term organization.

The present framework instead treats those morphisms as primary.

11.5 Natural Transformations

Suppose two observers construct different internal representations.

Each defines a functor

$$\mathcal{O}_1, \quad \mathcal{O}_2.$$

Agreement between observers should not require identical internal representations.

Instead it requires structural compatibility.

Definition 11.4 (Natural Interpretation). A natural transformation

$$\eta : \mathcal{O}_1 \Rightarrow \mathcal{O}_2$$

is called an interpretation whenever every admissible morphism satisfies the commutative diagram

$$\eta_B \circ \mathcal{O}_1(f) = \mathcal{O}_2(f) \circ \eta_A.$$

Interpretation therefore preserves relationships rather than individual descriptions.

Translation becomes structure-preserving rather than symbol-preserving.

11.6 Repair as an Endofunctor

Repair operates internally upon autonomous systems.

Consequently it defines an endofunctor.

Definition 11.5 (Repair Endofunctor). The repair operator determines an endofunctor

$$\mathcal{R} : \mathbf{Aut} \rightarrow \mathbf{Aut},$$

acting upon objects by decreasing structural defect while preserving admissibility.

Morphisms are mapped to repaired morphisms compatible with repaired state spaces.

Repeated repair therefore corresponds to repeated application of

$$\mathcal{R}.$$

The recursive sequence

$$\mathcal{R}, \quad \mathcal{R}^2, \quad \mathcal{R}^3, \dots$$

defines the categorical analogue of recursive continuation.

11.7 Equivalence of Autonomous Systems

Two systems need not possess identical internal realizations to exhibit identical structural organization.

Definition 11.6 (Structural Equivalence). Autonomous systems

A

and

B

are structurally equivalent whenever there exists an equivalence of categories preserving

1. admissibility,
2. repair,
3. continuation,
4. invariant preservation.

Thus identical implementation is unnecessary.

Only preservation of organizational structure is required.

This observation explains why identical computational principles frequently appear in biological evolution, distributed software systems, mathematical reasoning, and scientific collaboration despite radically different physical substrates.

11.8 The Functorial Principle

The previous constructions culminate in the following theorem.

Theorem 11.7 (Functorial Principle). *Every admissible observational process factors through a functor between categories of autonomous systems.*

Consequently, scientific explanation consists not in cataloguing objects but in identifying functorially preserved structure.

Proof. Observation preserves identity morphisms and composition.

Therefore it defines a functor.

Repair defines an endofunctor preserving admissibility.

Interpretation defines natural transformations between observational functors.

Hence every admissible explanatory framework factors through the categorical structure generated by autonomous systems and their admissibility-preserving morphisms.

□

This chapter completes a gradual shift in perspective that has developed throughout the monograph.

We began with individual states.

We then introduced interfaces, repair, continuation, distinction geometry, reachability, and local reconstruction.

Category theory reveals these constructions to be manifestations of a single principle: the mathematics of autonomous systems is fundamentally the mathematics of structure-preserving transformations.

Objects exist only insofar as they participate in morphisms.

Identity exists only insofar as those morphisms preserve admissible continuation.

The next chapter develops this observation into a unified theory of recursive scientific discovery, showing that mathematical research itself is an admissibility-preserving trajectory through a category of evolving theories.

Chapter 12

Scientific Discovery as Recursive Navigation

12.1 Discovery Is Not Creation

Scientific discourse frequently describes discovery as though new knowledge appeared suddenly through inspiration, intuition, or creativity.

Such descriptions correctly characterize the subjective experience of discovery while obscuring its mathematical structure.

The present framework adopts a different perspective.

Scientific discovery is not the creation of new reality.

Neither is it the arbitrary invention of conceptual systems.

Instead, discovery consists of identifying admissible trajectories through the space of possible theories.

The scientist does not create truth.

The scientist discovers navigable paths within an existing constraint geometry.

12.2 Theory Space

Let

$$\mathfrak{T}$$

denote the collection of scientific theories under consideration.

Each element

$$T \in \mathfrak{T}$$

is itself an autonomous representational system consisting of

1. definitions,
2. propositions,
3. explanatory principles,

4. admissibility conditions,
5. observational commitments.

Consequently,

$$\mathfrak{T}$$

is itself a configuration manifold.

Scientific development therefore becomes a dynamical process occurring upon theory space rather than physical space.

12.3 Theory Morphisms

Theories evolve.

Definitions become refined.

Hypotheses become generalized.

Proofs become simplified.

This motivates the following construction.

Definition 12.1 (Theory Morphism). A theory morphism

$$F : T_1 \longrightarrow T_2$$

is an admissibility-preserving transformation satisfying

1. logical consistency,
2. preservation of established observational constraints,
3. preservation of previously verified internal theorems,
4. reduction of representational defect.

Scientific progress therefore consists of morphisms rather than isolated theories.

Theories should be compared by their admissible transformations rather than their historical chronology.

12.4 Research Trajectories

A research program determines a continuous path through theory space.

Definition 12.2 (Research Trajectory). A research trajectory is a continuous curve

$$\Gamma : [0, \infty) \rightarrow \mathfrak{T},$$

whose image consists entirely of admissible theories.

The parameter need not represent chronological time.

It may instead represent accumulated understanding, mathematical refinement, or explanatory capacity.

Scientific maturity therefore corresponds to increasing regularity along the trajectory rather than increasing age.

12.5 The Discovery Functional

Not every admissible movement through theory space represents scientific progress.

We therefore introduce an optimization principle.

Definition 12.3 (Discovery Functional). The discovery functional is

$$\mathcal{D} : \mathfrak{T} \rightarrow \mathbb{R},$$

defined by

$$\mathcal{D}(T) = \alpha E(T) - \beta C(T) - \gamma \Delta(T) + \delta R(T),$$

where

- $E(T)$ measures explanatory power,
- $C(T)$ measures unnecessary complexity,
- $\Delta(T)$ measures reconstruction defect,
- $R(T)$ measures admissible reachability.

The constants

$$\alpha, \beta, \gamma, \delta > 0$$

determine the relative importance of each contribution.

Scientific refinement seeks trajectories satisfying

$$\mathcal{D}(\Gamma(t)) \nearrow.$$

Discovery therefore becomes constrained optimization over theory space.

12.6 Recursive Refinement

Scientific development rarely proceeds by replacement.

Instead it proceeds through recursive correction.

Let

$$\mathcal{R}_{\mathfrak{T}} : \mathfrak{T} \rightarrow \mathfrak{T}$$

denote the theory repair operator.

Repeated refinement generates

$$T, \quad \mathcal{R}(T), \quad \mathcal{R}^2(T), \quad \dots$$

Each iteration removes inconsistencies while preserving explanatory structure whenever possible.

Consequently mature theories should exhibit decreasing reconstruction defect without sacrificing observational adequacy.

12.7 Scientific Fixed Points

Some theories become remarkably stable.

Definition 12.4 (Scientific Fixed Point). A theory

$$T^*$$

is a scientific fixed point whenever

$$\mathcal{R}(T^*) = T^*,$$

and every sufficiently small admissible perturbation satisfies

$$\mathcal{R}(T^* + \varepsilon) = T^* + o(\varepsilon).$$

Fixed points do not represent absolute truth.

Rather, they represent locally stable explanatory structures.

Future refinement remains possible, but only through sufficiently large extensions of observational or conceptual reach.

12.8 The Geometry of Paradigm Change

Scientific revolutions correspond neither to arbitrary replacement nor to pure accumulation.

Instead they alter the topology of theory space itself.

New observational interfaces enlarge admissible regions.

Previously disconnected research trajectories become connected.

Formerly impossible continuations become reachable.

Thus paradigm shifts should be understood as topological transitions within the manifold of admissible theories rather than purely sociological events.

12.9 The Recursive Discovery Principle

The previous constructions culminate in the principal theorem of this chapter.

Theorem 12.5 (Recursive Discovery Principle). *Scientific knowledge evolves through recursively repaired trajectories in theory space that simultaneously decrease reconstruction defect, preserve verified invariants, and increase admissible explanatory reachability.*

Proof. Theory repair decreases inconsistency while preserving established structure.

Admissibility prevents logically invalid transitions.

Reachability restricts refinement to realizable theoretical extensions.

Consequently every admissible research trajectory consists of recursively repaired theories whose explanatory organization improves monotonically with respect to the discovery functional.

□

Scientific progress therefore admits a geometric interpretation.

The objective of research is not the accumulation of isolated facts.

Neither is it unrestricted theoretical invention.

The objective is the recursive navigation of theory space toward increasingly stable regions possessing lower reconstruction defect, greater explanatory compression, stronger invariant preservation, and larger admissible domains.

Accordingly, science itself becomes an autonomous repair process whose trajectory is governed by the same principles that organize cognition, computation, biological adaptation, and distributed intelligence.

Chapter 13

Compression, Parsimony, and Structural Minimality

13.1 The Principle of Structural Economy

Scientific theories should not be judged solely by predictive accuracy.

Among theories exhibiting equivalent explanatory capability, preference should be given to those possessing smaller structural complexity.

This observation is commonly associated with principles of parsimony.

Within the present framework, however, parsimony is not merely heuristic.

It arises naturally from admissible representation.

13.2 Representational Complexity

Definition 13.1 (Structural Complexity). Let

T

be an admissible theory.

Its structural complexity is

$K(T)$,

defined as the minimal descriptive complexity required to reconstruct every admissible consequence of

T .

The precise choice of complexity measure is intentionally left abstract.

Kolmogorov complexity, minimum description length, proof complexity, or categorical presentations may all instantiate

K .

13.3 Explanatory Compression

Scientific understanding is fundamentally compressive.

A successful theory replaces numerous independent observations by a smaller family of governing principles.

Definition 13.2 (Compression Ratio). The explanatory compression of a theory is

$$\mathcal{C}(T) = \frac{I_{\text{obs}}}{K(T)},$$

where

$$I_{\text{obs}}$$

measures the informational content of the observations explained by the theory.

Larger values correspond to greater explanatory efficiency.

13.4 Minimal Presentations

Different formulations may describe identical mathematical structure.

Accordingly, minimality concerns equivalence classes rather than individual expressions.

Definition 13.3 (Minimal Presentation). A presentation

$$T^*$$

is minimal whenever

$$K(T^*) = \inf\{K(S) : S \cong T\},$$

where

$$\cong$$

denotes structural equivalence.

Scientific elegance therefore reflects structural minimality rather than stylistic preference.

Theorem 13.4 (Structural Economy Principle). *Among observationally equivalent theories, maximal explanatory compression corresponds to minimal structural complexity.*

Chapter 14

Symmetry, Redundancy, and Canonical Form

14.1 Redundant Descriptions

Distinct representations frequently encode identical mathematical objects.

Coordinates, notation, implementation, and language may differ while the underlying structure remains unchanged.

Scientific progress therefore requires identifying canonical representations.

Definition 14.1 (Redundancy). Two representations

$$R_1, R_2$$

are redundant whenever they determine the same admissible equivalence class.

Redundancy increases descriptive complexity without increasing explanatory power.

14.2 Canonical Forms

A canonical representation is one that minimizes redundancy while preserving every invariant relevant to admissibility.

Canonicalization is therefore an admissibility-preserving projection onto a minimal representative of an equivalence class.

Proposition 14.2. *Every reduction in representational redundancy strictly decreases reconstruction ambiguity while preserving invariant structure.*

Chapter 15

The Principle of Recursive Scientific Realism

15.1 Reality Through Repair

Scientific realism is frequently interpreted as the claim that successful theories correspond to external reality.

The present framework adopts a weaker and more constructive position.

A theory is accepted not because it has been proven final, but because it admits indefinitely continued repair while preserving an expanding family of invariants.

Definition 15.1 (Recursively Realizable Theory). A theory

$$T$$

is recursively realizable whenever there exists an infinite admissible sequence

$$T_0, T_1, T_2, \dots$$

such that

$$T_0 = T,$$

each transition is an admissible repair, and every verified invariant of

$$T_n$$

remains preserved in

$$T_{n+1}.$$

Theorem 15.2 (Recursive Scientific Realism). *Scientific progress is characterized neither by convergence toward immutable certainty nor by unrestricted conceptual revision.*

Rather, it consists of recursively realizable trajectories preserving an expanding family of experimentally and mathematically verified invariants.

The objective of science is therefore not perfection.

It is indefinite admissible continuation.

Bibliography

- [1] Saunders Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Springer, 1998.
- [2] Steve Awodey, *Category Theory*, 2nd ed., Oxford University Press, 2010.
- [3] F. William Lawvere and Stephen Schanuel, *Conceptual Mathematics*, 2nd ed., Cambridge University Press, 2009.
- [4] James R. Munkres, *Topology*, 2nd ed., Prentice Hall, 2000.
- [5] John M. Lee, *Introduction to Topological Manifolds*, 2nd ed., Springer, 2011.
- [6] John M. Lee, *Introduction to Smooth Manifolds*, 2nd ed., Springer, 2012.
- [7] Victor Guillemin and Alan Pollack, *Differential Topology*, AMS Chelsea, 2010.
- [8] Manfredo P. do Carmo, *Differential Geometry of Curves and Surfaces*, Prentice Hall, 1976.
- [9] Michael Spivak, *A Comprehensive Introduction to Differential Geometry*, Publish or Perish, 1999.
- [10] Morris Hirsch, Stephen Smale, Robert Devaney, *Differential Equations, Dynamical Systems, and an Introduction to Chaos*, Academic Press, 2012.
- [11] Anatole Katok and Boris Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*, Cambridge University Press, 1995.
- [12] Stephen Boyd and Lieven Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.
- [13] R. Tyrrell Rockafellar, *Convex Analysis*, Princeton University Press, 1970.
- [14] Claude E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, 27 (1948), 379–423, 623–656.
- [15] Thomas Cover and Joy Thomas, *Elements of Information Theory*, 2nd ed., Wiley, 2006.
- [16] Patrick Billingsley, *Probability and Measure*, 3rd ed., Wiley, 1995.
- [17] Reinhard Diestel, *Graph Theory*, 5th ed., Springer, 2017.
- [18] Serge Lang, *Algebra*, Revised 3rd ed., Springer, 2002.
- [19] Walter Rudin, *Functional Analysis*, 2nd ed., McGraw–Hill, 1991.

- [20] Herbert Enderton, *A Mathematical Introduction to Logic*, 2nd ed., Academic Press, 2001.
- [21] Glen E. Bredon, *Sheaf Theory*, 2nd ed., Springer, 1997.
- [22] Masaki Kashiwara and Pierre Schapira, *Categories and Sheaves*, Springer, 2006.
- [23] Michael Sipser, *Introduction to the Theory of Computation*, 3rd ed., Cengage, 2013.
- [24] Dimitri Bertsekas, *Dynamic Programming and Optimal Control*, Athena Scientific, 2017.
- [25] Karl Popper, *The Logic of Scientific Discovery*, Routledge, 2002.
- [26] Imre Lakatos, *The Methodology of Scientific Research Programmes*, Cambridge University Press, 1978.
- [27] Thomas S. Kuhn, *The Structure of Scientific Revolutions*, 4th ed., University of Chicago Press, 2012.
- [28] Ludwig von Bertalanffy, *General System Theory*, George Braziller, 1968.
- [29] Norbert Wiener, *Cybernetics*, MIT Press, 1961.