Mirror–Recursive Information Field (MRIF): Dual Descriptions, Recursion, and Capacity Envelopes

Megan Anderson

AI ARMY, INC. · aiarmy.co

2026

Abstract

Adaptive systems — brains, artificial agents, and collectives — routinely confront the same structural problems: they must integrate information over time, regulate their own inference, and remain coherent under load, despite limited capacity and recursive self-modeling. Many contemporary frameworks illuminate part of this picture: predictive processing emphasizes probabilistic inference [1-4], control theory emphasizes regulation, information geometry emphasizes manifold structure [6], and the IGAF formal core, housed in Supplement II, describes the geometry of inference, while TNI describes the projected geometry of subjective time. What is missing is a meta-theoretical layer that characterizes which descriptions of a system are structurally admissible, how they must agree, and when they fail.

We propose the Mirror–Recursive Information Field (MRIF) as such a layer. MRIF is organized by four axioms — an information-first ontology, relational temporality, mirror duality between geometric and dynamical descriptions, and the realizability of bounded recursive self-modeling — whose formal statements are held in the cognitive program's shared formal supplement (Supplement II, Part 2) and referenced here by stable name. Together these axioms specify how an information field can be described both as a structured space (basins, barriers, manifolds) and as a pattern of flows and regulations, with a structure-preserving mirror map Φ enforcing consistency between the two views.

On this foundation we develop the Architect Frame — a system's distributed capacity to bound recursion, regulate update magnitude, and maintain coherence within a capacity envelope c_eff(t) — and characterize the structural failure modes (fragmentation, rigidification, runaway recursion) that follow when demand chronically exceeds that envelope. We show how MRIF serves as a meta-theoretical envelope for IGAF and TNI, illustrate it on minds as a worked case, extend it to collective information fields, and state explicit non-claims regarding quantum cognition, determinism, and diagnosis. MRIF offers a structural language for reasoning about coherence, failure, and repair in systems that model themselves — while asserting no physical mechanism and transferring no evidential weight across regimes.

Keywords: MRIF; adaptive systems; recursive self-modeling; information geometry; capacity envelopes; predictive processing; collective information fields; temporal cognition

1. Introduction

1.1 Motivation

Across neuroscience, artificial intelligence, and the study of collective behavior, the same structural questions recur. Why do some systems integrate information gracefully over time while others fragment or lock up under similar loads? Why does certain self-reflection deepen coherence while other self-reflection triggers runaway self-critique or paralysis? Why do groups of individually capable members sometimes collapse into chaos, polarization, or stagnation?

Existing frameworks address parts of this puzzle. Predictive processing and Bayesian-brain theories describe inference under uncertainty [1-4]. Control theory describes regulation around setpoints. Dynamical-systems theory offers tools for attractors, bifurcations, and metastability. Information geometry formalizes state spaces and metrics [6]. The IGAF formal core, housed in Supplement II and normalized as Information Geometry of Adaptive Fields, describes the geometry of inference; TNI (Temporal Neuroscience Index) describes the projected geometry of subjective time.

Yet these tools often talk past one another. A dynamical-systems model may call a system stable while a control-theoretic view warns of oscillation; a geometric model may show shallow basins where an algorithmic description assumes deep ones. Different disciplines assign different causes to the same breakdown: content, beliefs, noise, topology, or regulation. Without a unifying layer, it is difficult to say when two descriptions of the *same* system are structurally compatible — and when they are silently contradicting one another.

This paper proposes MRIF as that unifying layer. MRIF does not replace existing models, and it is not another mechanism-level theory of the brain. It describes what must be true for any description of an adaptive, self-modeling information-processing system to be structurally coherent.

1.2 The idea of MRIF in one sentence

MRIF is a structural framework in which an adaptive system is treated as an information field that admits two mirror-consistent descriptions — geometric and dynamical — while supporting recursive self-modeling within evolving capacity envelopes.

Four axioms organize the framework, stated informally here and formally in Supplement II, Part 2:

- **Information ontology** (Axiom of Structured Inference): systems are described first as structured information fields, not as substances.
- **Relational temporality** (Axiom of Relational Temporality): time, for a system, is defined by structured change under constraint, not by a single universal clock.
- **Mirror duality** (Axiom of Mirror Duality): every viable description admits a geometric/topological view and a dynamical/regulatory view, related by a structure-preserving map Φ.
- **Bounded recursion** (Axiom of Bounded Recursive Self-Modeling): self-modeling is realizable but necessarily bounded by capacity.

On this base, MRIF defines how capacity envelopes, failure modes, and collective fields emerge.

1.3 Relationship to IGAF and TNI, and the shared formal home

MRIF is not meant to stand alone; it is designed as the meta-layer above the cognitive program's geometric and temporal-projection layers. The IGAF formal core, housed in Supplement II, describes the geometric structure of inference — manifolds of belief or policy states, attractor basins, barriers, curvature; it tells us what the state space looks like and how inference moves on it. TNI describes the projected geometry of subjective time, parameterizing temporal structure into a small descriptive index. MRIF sits above both: it constrains IGAF by specifying which manifolds and attractor fields are structurally admissible given that systems maintain coherence under limited capacity and recursive self-modeling, and it constrains TNI by specifying which temporal geometries are viable projections of such inference. The three frameworks share one formal home. **All of the cognitive program's definitions, axioms, constructs, and propositions are stated once, by stable name, in Supplement II — Adaptive Systems Formal Supplement.** IGAF, TNI, and MRIF *reference* those statements rather than restating them, and cite results by stable name (e.g., "Axiom of Relational Temporality," "Internal

Time as Path Length," "Ordering-Sensitivity Non-Commutativity Proposition") rather than by position number, so that renumbering can never drift against citations. This paper follows that convention throughout: where a formal object is invoked, its conceptual content is developed here and its canonical statement is located in Supplement II. Two disciplines govern how MRIF relates to the rest of the work, and both are stated once here so the reader can hold them through the paper.

First, the internal structure of the cognitive program. Within the cognitive regime there is no internal firewall: IGAF, MRIF, and TNI describe one regime at three levels, and they may ground upward freely — geometry (IGAF) supports regulation (MRIF) supports projection (TNI). A no-upward-claims hierarchy holds throughout: TNI does not redefine physical time, MRIF asserts no physical ontology, IGAF asserts no factuality, and admissibility is never evidence.

Second, the relation to the physical program. Elsewhere in this body of work a separate, physics-facing program treats quantum events, event-time, and physical publicization. MRIF's relation to that program is **lateral structural resonance under a firewall, not vertical continuation.** Wherever a construct in the cognitive program borrows a physical-regime template — an "inferential light cone," cone-limited reconstruction, an out-of-time-order correlator — or asserts a parallel to the physical program, the following statement applies verbatim and is never paraphrased:

> The parallel asserted is structural only. The same grammar recurs across regimes; no evidential weight transfers between them. A result in one regime is not evidence about another.

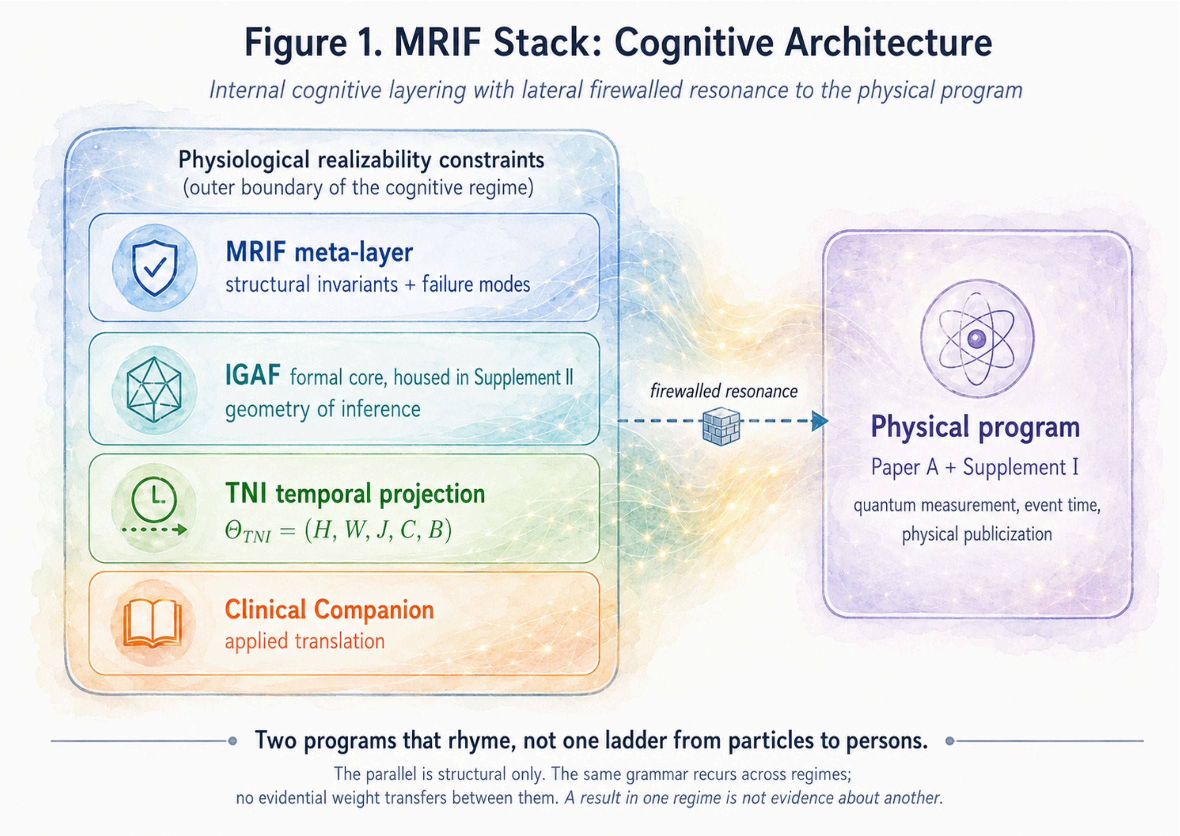

Figure 1. MRIF Stack: cognitive architecture. Internal cognitive layering with lateral firewalled resonance to the physical program. Two programs that rhyme, not one ladder from particles to persons. The parallel is structural only; the same grammar recurs across regimes; no evidential weight transfers between them.

1.4 Overview

Section 2 states the four axioms and MRIF's scope and non-claims. Section 3 develops mirror duality and the map Φ. Section 4 examines recursive self-modeling and amplification. Section 5 introduces the Architect Frame and capacity envelopes and the structural failure modes. Section 6 situates IGAF and TNI within MRIF, states the parallel-programs architecture of the theory stack, and works minds as a concrete case. Section 7 extends MRIF to collective fields. Section 8 outlines measurement. Section 9 states falsifiable predictions. Section 10 makes MRIF's non-claims explicit. Section 11 discusses implications and future directions.

2. Foundations and Axioms

This section states MRIF's foundational commitments. The goal is not a new physics or a competing ontology of mind, but to make explicit the minimal structural commitments MRIF needs to compare descriptions, analyze failure modes, and bound recursion and capacity. Each axiom is stated informally here for intuition; its formal statement is the correspondingly named object in Supplement II, Part 2, which also records its provenance across the two historical axiom numberings (MRIF-1…4; TNI Axioms 1–6).

2.1 Information ontology (Axiom of Structured Inference)

MRIF begins from an information-first stance: the primary object of description is an information field over possible configurations, not particles, substances, or beliefs per se. What matters is how distinctions are structured and constrained. In practice, systems are treated as organized patterns of distinguishability — which states are possible, how they differ, how constraints shape transitions — and implementation details (biological, digital, social) matter for capacity and failure modes but are, at the MRIF level, different realizations of the same kind of structure. This is not a claim that information is more real than matter, nor a reduction of physics to information; it is a modeling choice about what MRIF describes. The formal statement — an information field with a distinguishability function and an admissible-transition set, equivalently a statistical manifold with a metric — is the *Axiom of Structured Inference* and the *Information State* and *Statistical Manifold and Metric* definitions (Supplement II, Parts 1–2).

2.2 Relational temporality (Axiom of Relational Temporality)

MRIF treats time as relational and internal, not as a single universal parameter all systems share. For a given system, time is defined by how its information field changes under constraint; two systems over the same external clock interval can accumulate different internal times. Coherence and failure are more naturally described in terms of patterns of change than of elapsed seconds. The canonical formal object is *Internal Time as Path Length* — the accumulated structured belief-change along a trajectory — stated identically across the program (Supplement II, Part 1). MRIF's relational-time commitment cites the physical program for the status of physical time rather than asserting a physical time variable natively; this is the no-upward-claims discipline in action.

2.3 Mirror duality (Axiom of Mirror Duality)

MRIF posits that any structurally coherent description of an adaptive system admits two complementary views — a geometric/topological description (states, basins, barriers) and a dynamical/regulatory description (flows, updates, controllers) — related by a structure-preserving mirror map Φ that keeps core invariants (stability, reachability, capacity, causal consistency) aligned. When the two diverge, at least one description is structurally wrong. This is what makes MRIF a meta-theory: it does not privilege one language (differential equations, probabilistic graphs, control diagrams); it requires that whichever language is chosen, a compatible mirror description exists. The formal objects — *Geometric and Dynamical Descriptions*, *Mirror-Consistent Pair*, and the mirror-duality axiom itself — are stated in Supplement II, Part 3.1. Φ is the program's own coherence criterion.

2.4 Bounded recursive self-modeling (Axiom of Bounded Recursive Self-Modeling)

Systems of interest — human brains, sophisticated agents, organized collectives — do more than respond to inputs; they model the environment, themselves, their own inference and regulation, and sometimes how others model them. MRIF holds that this recursion is realizable but necessarily bounded: there is a practical depth beyond which additional self-modeling fails to be maintained or yields net instability (runaway loops, paralysis, fragmentation), and recursion both reshapes the effective geometry of the state space and consumes capacity from a finite envelope. The formal objects — the *Model Tower* and *Recursive Self-Modeling Capacity* — are in Supplement II, Part 3.2.

2.5 Scope, non-claims, and the role of the axioms

The four axioms jointly define MRIF's scope. They apply to any system that can be meaningfully described as carrying and transforming structured information, having an internal temporal geometry, admitting both geometric and dynamical models, and supporting at least some self-reflection about its own inference.

They also define what MRIF does not assert. MRIF does not claim that information is the fundamental substance of the universe, that brains implement literal quantum states in the physics sense, or that there is a unique or canonical choice of metric, manifold, or dynamics. MRIF does not prescribe a specific neural mechanism, a particular AI algorithm, or a single correct theory of mind.

Instead, MRIF provides structural guardrails. If a proposed description of a system cannot be given both a geometric and a dynamical formulation that are mirror-consistent, it violates mirror duality. If it assumes arbitrarily deep self-modeling without resource or capacity bounds, it violates bounded recursion. If it treats time as a single global scalar independent of change in the information field, it violates relational temporality. If it cannot be recast in informational terms, it falls outside the information ontology.

3. Mirror Duality Formalism

Mirror duality asserts that structurally coherent descriptions admit two complementary, mirror-related forms — geometric/topological and dynamical/regulatory — linked by a structure-preserving map Φ. Geometry tells us how states are arranged and constrained; dynamics tells us how the system moves among them; MRIF requires the two stories to agree on what is stable, reachable, and viable. This section develops the two views and the map conceptually; the definitions (*Geometric and Dynamical Descriptions*, *Mirror-Consistent Pair*) are in Supplement II, Part 3.1.

3.1 Geometric description

The geometric view focuses on the shape of the information field: a state space (a differentiable manifold, a graph, or a hybrid), a notion of distance or divergence (a metric, or an information-geometric measure such as the Fisher metric), a basin structure (attractors, boundaries, saddles, barrier heights), and possibly topological features (holes, bottlenecks). In this view we ask where the attractors are, how deep and wide their basins are, how curved the space is, and which regions are easy or hard to reach. IGAF lives largely here; TNI uses this geometry as a substrate for projecting subjective time.

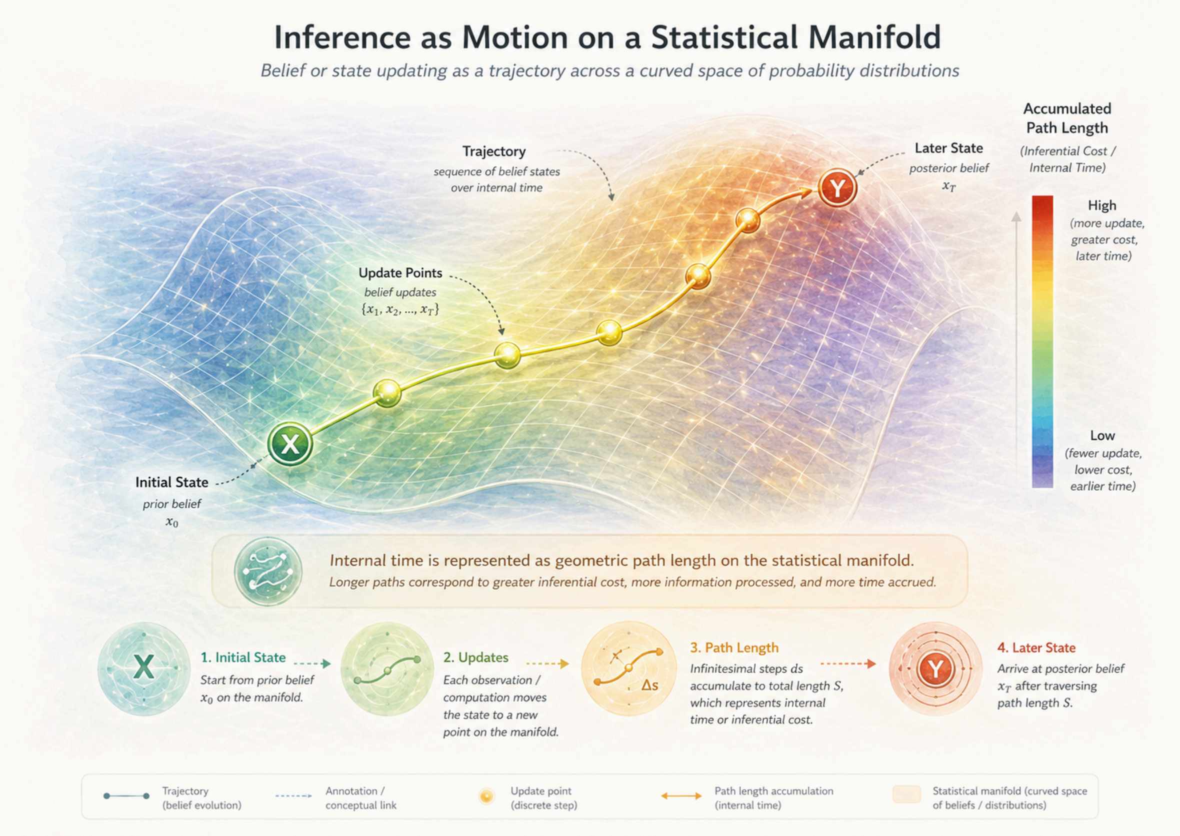

Figure 2. Inference as motion on a statistical manifold. Belief/state updating as a trajectory across a curved space of probability distributions; internal time is represented as geometric path length, and longer paths correspond to greater inferential cost.

3.2 Dynamical description

The dynamical view focuses on how states change: flows, updates, control signals, and regulation — gain and precision modulation, thresholds and saturations, error monitoring, meta-control. Here we ask what the update rules are, how the system responds to perturbation, how it regulates its own dynamics (stabilizing versus destabilizing feedback), and what happens as parameters or loads change. This is the language of predictive processing, control theory, and algorithmic update rules [1-4].

3.3 The mirror map Φ

Mirror duality posits a map Φ relating the geometric and dynamical descriptions such that core invariants are preserved — stability (attractor sets correspond to invariant sets of the dynamics), reachability (paths in the state space correspond to feasible transition sequences under constraint), capacity (geometric costs correspond to dynamical costs within bounded distortion), and causal consistency (irreversibility, hysteresis, and path dependence appear in both views). Φ need not be unique or closed-form; it is the requirement that the two descriptions tell the same structural story, and a model that cannot in principle be cast as a mirror-consistent pair is structurally incompatible with the program.

3.4 When mirror duality fails

Violations of mirror duality are diagnostic of modeling problems, not of exotic physics. A geometric model may depict a single broad basin (suggesting robust stability) while the dynamical model shows frequent oscillation or chaotic switching — signaling missing geometric structure or a mis-specified dynamics. A dynamical model may assume quick convergence under small interventions while the geometry shows those states lie beyond high barriers — signaling control assumptions that violate capacity constraints. Aggregating over states can smooth away critical geometric features while the dynamics still feel them, so models that average out topology may falsely suggest simple controllability. In practice, MRIF uses mirror duality as a consistency check: build one description, construct the other, and ask whether they can be made mirror-consistent under plausible mapping assumptions. If not, at least one is structurally incomplete.

3.5 Cross-domain mapping and theory comparison

A central benefit of mirror duality is that different theoretical frameworks can be compared at the level of shared invariants rather than surface syntax. A predictive-processing model [1-4] and a control-theoretic model of the same system can be translated into a shared geometric picture and checked for consistency of their implied attractors, basins, and capacity envelopes; a psychological model in terms of "coping styles" and a dynamical model in terms of phase portraits can both be evaluated against the same MRIF-derived constraints even when their original variables differ. MRIF does not demand unification into one master theory; it demands that any viable theory admit both a geometric and a dynamical interpretation that are mirror-consistent, so that its claims about stability, change, and capacity can be trusted.

Where this comparison reaches across regimes — where the grammar of MRIF is set beside the grammar of the physical program — the firewall applies:

The parallel asserted is structural only. The same grammar recurs across regimes; no evidential weight transfers between them. A result in one regime is not evidence about another.

4. Recursive Self-Modeling and Amplification

Bounded recursion holds that systems of interest model their own inference and regulation. This capacity is a main source of flexibility and a main route to instability when capacity is exceeded. The formal objects (Model Tower, Recursive Self-Modeling Capacity) are in Supplement II, Part 3.2; this section develops their behavior.

4.1 Levels of recursion

We distinguish at least four levels. Level 0, raw dynamics: the system evolves without any internal self-representation. Level 1, model of environment and self-state. Level 2, model of its own inference ("when I am tired, my predictions are noisy"). Level 3 and above, models of models — how others model it, how its meta-strategies work over time — often partial, approximate, and context-dependent. Recursion enables learning from patterns of error, flexible strategy change, and rich interpersonal reasoning; but each level adds degrees of freedom that can amplify both signal and noise.

4.2 Amplification of coherence and instability

Recursion is a gain control on inference. Constructively, it scales learning across contexts, generalizes corrections, and improves robustness. Destructively, it can magnify small errors into self-fulfilling narratives, create spirals of self-critique that consume capacity, and destabilize otherwise manageable dynamics by layering monitoring, evaluation, and inhibition. Geometrically, recursion changes the effective landscape: extra levels of self-modeling add new axes and potential basins — some shallow and noisy (rumination loops), some deep (rigid self-judgments) — so that unchecked recursion turns a gentle slope into a rugged, high-dimensional landscape with many traps.

4.3 The need for structural bounds

Given these amplification properties, MRIF treats bounded recursion as a structural requirement for viability: there is a maximum effective recursion depth beyond which additional layers contribute more instability than benefit, and a safe operating region in which recursion is engaged selectively, capacity envelopes are respected, and meta-level corrections are paced to the system's integration bandwidth. Different systems have different bounds, and these bounds are structural facts about capacity, load, and the geometry of self-modeling — not moral or diagnostic judgments. To reason about how a system stays within them without positing a "little decider," we need the Architect Frame.

5. Architect Frame and Capacity Envelopes

To remain viable, a system must regulate how hard it pushes its inference and how deep it lets recursion go. The Architect Frame is the structural capacity that performs this regulation, and the capacity envelope c_eff(t) formalizes the bound. The definitions (*Architect Frame*, *Effective Capacity Envelope*) and the failure-mode taxonomy are in Supplement II, Parts 1, 3.3, and 5; this section develops them.

5.1 The Architect Frame

The Architect Frame is the structured set of mechanisms and parameters that monitor a system's operation (error statistics, volatility, resource use), adjust control and inference parameters (step sizes, gains, recursion depth, policy selection), and restructure trajectories (simplification, pacing, prioritization) so that the system evolves lawfully within feasible limits. It is structural, not local — realized across subsystems (attention, executive control, affective regulation, meta-learning in AI) and across timescales. It is model-aware but not omniscient: it has partial, approximate access to how inference is going and adjusts high-level parameters or modes rather than rewriting the system. And it is capacity-bound: it cannot arbitrarily suppress load, and it can itself fail, fragment, or rigidify under extreme conditions. **The Architect Frame is not a homunculus** — not an internal agent, but a structured set of mechanisms and parameters. This guard is locked and carried wherever the construct appears.

5.2 Capacity envelopes and bandwidth

The capacity envelope c_eff(t) is the function such that, during viable operation, inferential load stays within it. It is state-dependent, task-dependent, and Architect-Frame-dependent, and it connects to concrete bandwidth constraints: biological brains face refractory periods, synaptic delays, and integration times such that large-scale coherent updates operate on timescales of tens to hundreds of milliseconds, not microseconds; artificial systems face analogous throughput and coordination limits. Felt urgency rises as demand approaches capacity; regulation operates by lowering demand or raising capacity.

5.3 Structural failure modes under saturation

When demand chronically exceeds the envelope, or the Architect Frame breaks down, three structural failure modes follow. **Fragmentation:** high-frequency transitions among shallow basins — low coherence, high jitter, unstable strategy switching. **Rigidification:** deepening of a few basins with raised barriers — narrowed options, entrenched patterns, inability to adapt. **Runaway recursion:** trajectories dominated by meta-level basins — monitoring the monitoring, evaluating the evaluation — oscillating or stuck, consuming capacity without resolving. These are lawful consequences of persistent saturation, and their canonical statement, together with the correspondence to the program's other breakdown vocabularies, is in Supplement II, Part 5.

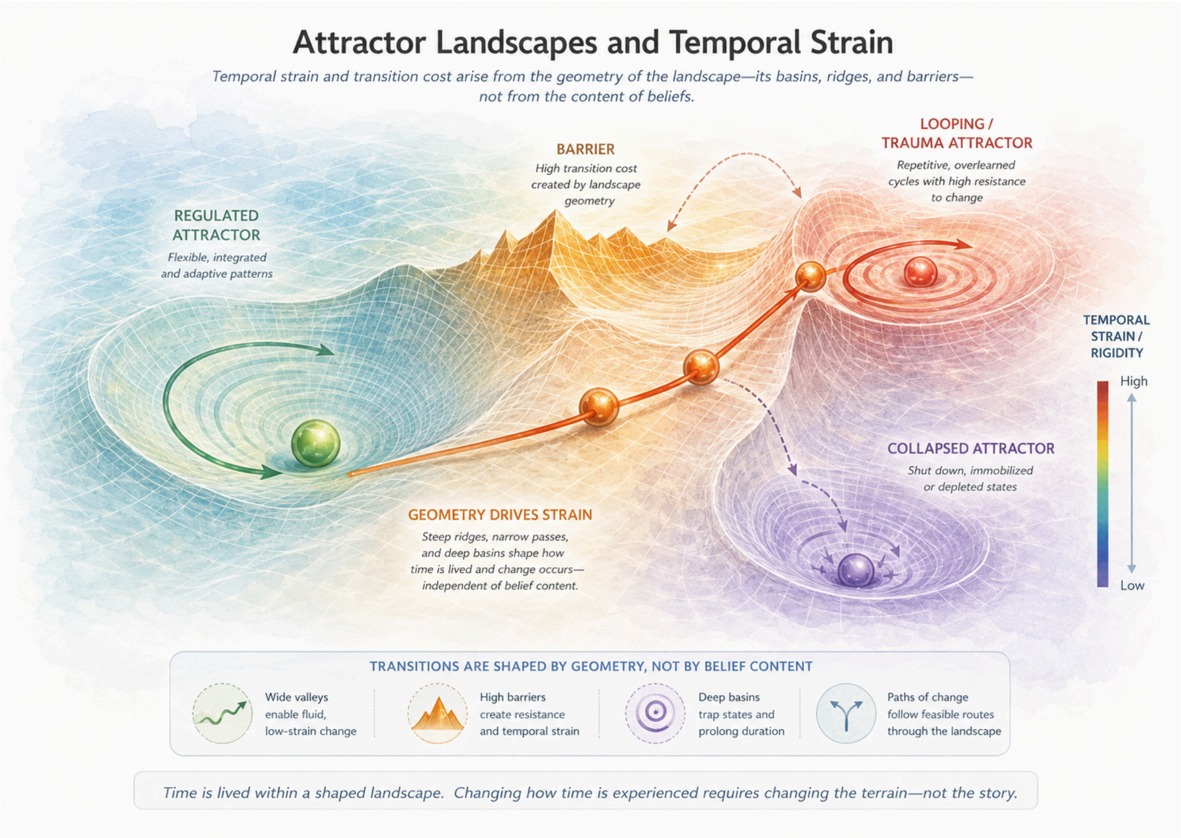

Figure 3. Attractor landscapes and temporal strain. Temporal strain and transition cost arise from the geometry of the landscape — its basins, ridges, and barriers — not from the content of beliefs.

5.4 The Architect Frame as target for intervention

One practical consequence: durable change often requires adjusting the Architect Frame and the capacity envelope rather than only modifying lower-level content. In individuals, this means pacing and chunking to lower load, training meta-strategies that engage recursion selectively, and modifying environments to reduce unnecessary load. In artificial systems, it means bounding the depth of self-modification and self-evaluation, implementing governors that monitor resource use and error rates, and defining safe modes where complex recursion is temporarily disabled. In collectives, it means changing decision procedures, distributing meta-control through roles and process, and designing capacity-aware rhythms.

6. MRIF as Meta-Theory for IGAF and TNI

MRIF emerged from the need to coordinate the cognitive program's geometric and temporal-projection layers: the IGAF formal core, housed in Supplement II, which describes how inference is organized in state space, and TNI, a projection framework for the geometry of subjective time. This section shows how MRIF provides the meta-theoretical envelope within which they operate, states the parallel-programs architecture of the theory stack, and works minds as a concrete case.

6.1 IGAF under MRIF

IGAF describes the geometry of inference: states represent beliefs, policies, or configurations; attractor fields encode stable patterns; geometry encodes how hard it is to move between configurations [6]. Under MRIF, IGAF is a particular instantiation of the geometric description for cognitive inference. Its manifold is a subspace of the broader information field, selecting dimensions relevant to inference and abstracting away implementation detail. Its basins and attractors must be mirror-consistent with some dynamical description — Bayesian update rules, policy-gradient dynamics, predictive-processing flows — and its geometry must respect capacity envelopes: regions requiring impossible update rates should appear as high curvature, high barriers, or inaccessible zones. A proposed IGAF model is structurally admissible if it admits a dynamical counterpart linkable to the manifold via a reasonable Φ, and if the implied trajectories stay within plausible capacity envelopes. If not, MRIF treats it as structurally suspect even if it fits some data.

6.2 TNI under MRIF

TNI describes the projected geometry of subjective time. Its parameters are defined once, by stable name, in Supplement II, Part 4; this paper names them and locates their home rather than restating their formal decomposition. The TNI **core vector** is Θ_TNI = (H, W, J, C, B) — horizon, future width, temporal jitter, coherence, and branch richness. For visualization or projection there is an **augmented tuple** Θ_TNI^{+} = (H, W, J, C, B; ψ), where ψ (ordering sensitivity) is a *derived readout* — a function $\psi = \Psi(H, W, J, C, B)$ over the five core parameters — **not an independent sixth coordinate**; the semicolon marks an augmented readout, not a six-dimensional state space. The formal definition of ψ and the *Ordering-Sensitivity Non-Commutativity Proposition* that grounds it are Supplement II objects (Part 3.6, Part 4); TNI develops them, and MRIF simply relies on their existence.

In MRIF terms, TNI is a projection of inferential trajectories onto a temporal coordinate system. The underlying trajectories live on IGAF-like manifolds within the information field; the TNI parameters are read from path structure, basin structure, and capacity envelopes. MRIF constrains TNI in four ways. Relational temporality requires TNI's internal time and temporal regimes to be definable in terms of structured change, not clock time. Mirror duality requires TNI's temporal geometry to match a dynamical story in which policies and strategies change over time under the same barriers and curvature. Bounded recursion supplies the language for how self-modeling deforms temporal geometry — expanding horizons through meta-planning, or collapsing them through self-imposed constraint — and how the Architect Frame keeps this within viable bounds. And the capacity envelope explains TNI's temporal strain and distortions as consequences of load approaching or exceeding capacity. In one line: TNI describes how time is organized from inside a system; MRIF describes why that temporal geometry is structurally viable or not.

6.3 The theory stack: parallel programs, not a tower

MRIF, IGAF, and TNI form a layered account *within the cognitive regime*, and it is important to state the shape of that layering precisely, because an earlier rendering invited a misreading.

Within the cognitive program the layering is real and internal: physiological realizability constraints (neural bandwidth, connectivity, delays, metabolic limits) set the outermost bounds on what information fields, dynamics, and capacity envelopes are realizable; MRIF specifies the structural invariants and failure modes; IGAF supplies the geometric model of inference that must satisfy them; TNI projects the resulting temporal geometry into a small descriptive index. Operationally, the TNI chart modulates *demand* on the system and the MRIF envelope sets *supply*. This is a stack of levels describing one regime, and within it the frameworks ground upward freely.

What this layering is **not** is a universal tower spanning physics to mind. The cognitive program's physiological substrate is distinct from the physical program's subject matter (quantum measurement, event-time, physical publicization). MRIF's relation to the physical program is lateral structural resonance under the firewall, not vertical continuation: the two programs share grammar, not evidence. Any figure depicting the stack must show the cognitive levels as an internal layering and the relation to the physical program as a lateral, firewalled resonance — parallel programs that rhyme, not a single ladder from particles to persons. At that seam the statement applies verbatim:

> The parallel asserted is structural only. The same grammar recurs across regimes; no evidential weight transfers between them. A result in one regime is not evidence about another.

6.4 Minds as a worked case

MRIF's most immediate application is to a single mind, and treating one concretely makes the abstract machinery legible. This section is deliberately confined to MRIF's own constructs — the four axioms, mirror duality, bounded recursion, the Architect Frame, and the Part-5 failure modes. It does not import TNI's parameters (which are downstream projections, defined in Supplement II, Part 4) and it does not borrow physical-regime operators; where a structural analogy to physics might tempt, none is asserted.

Treat a mind, at the information level, as a local information field: a subsystem whose state encodes internal models, beliefs, memories, priors, and goals. Nothing here requires a commitment about neural implementation; the description is at the level of structured distinguishability and constraint (Axiom of Structured Inference).

This field admits the two mirror-consistent descriptions the framework requires. On the geometric side, there is a relatively slow-changing self-model — an identity, a narrative, a structured sense of who one is across time [5]. On the dynamical side, there is the ongoing stream of experience: perceptions, thoughts, actions, and affect updating moment to moment. Mirror duality asserts that a coherent mind is one in which these two are related by a structure-preserving map — the lived stream and the narrative self telling the same structural story. When they diverge — when the moment-to-moment dynamics can no longer be recovered from the self-model, or the self-model no longer tracks what is actually happening — that divergence is exactly the mirror-inconsistency the framework flags, and it is felt from the inside as a loss of continuity or self-alienation.

Recursion appears in a mind as the capacity to model one's own inference — to notice one's tendencies, to reinterpret, to update schemas. Within bounds, this stabilizes identity: recursive coarse-graining compresses many experiences into coherent lessons and narratives, and error-correction reinterprets local shocks to preserve a continuous self, so that a stable identity behaves like a self-stabilizing structure maintained over time. Past its bounds, the same recursion produces the runaway mode of Section 5 — self-judgment about self-judgment, consuming capacity without resolving.

Finally, the failure modes of Section 5 have direct cognitive readings. Rigidification appears as rigid schemas and trauma-linked patterns: a few basins deepen and their barriers rise, so that a narrow band of interpretations becomes hard to leave and the felt future narrows around them. Fragmentation appears as the opposite loss of integration — shallow, rapidly switching states with weak narrative continuity. Recovery, in this reading, is the Section-5 asymmetry made personal: leaving a deepened basin requires partial restoration of the capacity envelope and reconsolidation of coherence, not a simple reversal of the path in. These readings are structural descriptions of an information field under load; they are not diagnostic categories, and the clinical development of this material is reserved for the separate companion.

7. Collective Fields and Coupled Architect Frames

Many of the most consequential phenomena — social breakdown, organizational dysfunction, emergent coordination — arise not in individual systems but in collective information fields: families, teams, institutions, and multi-agent systems. MRIF extends to these by coupling Architect Frames. The full formal treatment — collective geometry, collective regulation, collective projection, and the structural predictions — is consolidated in Supplement II, Part 7; this section states the results conceptually. The no-upward-claims and firewall disciplines carry unchanged.

7.1 Collective geometry and regulation

For several interacting systems, a collective information field lives over a joint space with an interaction (coupling) structure; equivalently, in geometric coordinates, a product manifold with coupling cross-terms. Persistent group patterns can be collective attractors — minima of a collective objective — that are properties of the coupled field rather than traits of any member; this is the formal basis for the observation that some apparently individual problems are collective attractors sustained by coupling. Coordination failure registers as divergence in update *dynamics* (temporal coupling), not merely as disagreement in content. Regulation, in turn, couples the members' Architect Frames into a group-level capacity: coherence can diffuse (a regulated member lowers others' effective load and deepens viable collective basins), jitter can propagate (a member persistently over capacity sends volatile signals that others must continually re-absorb), and collapse can propagate (one member's repeated shutdown drives others to over-function or co-collapse).

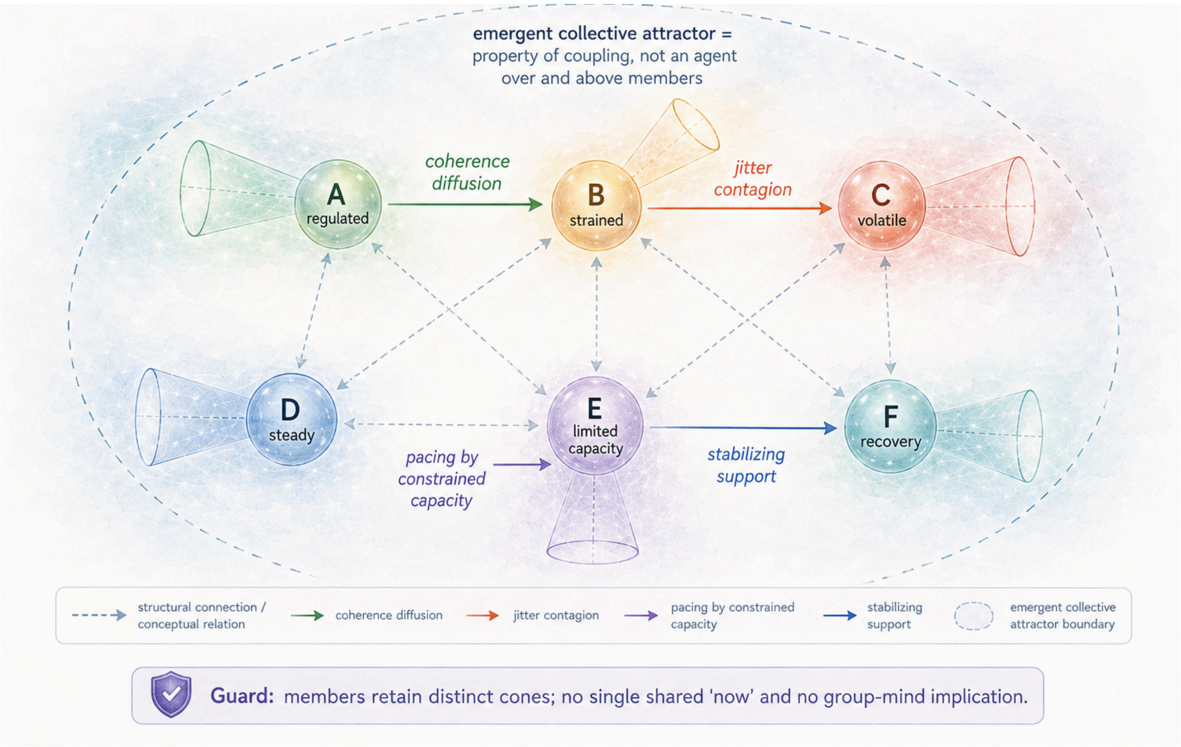

Figure 4. Collective extension: coupled temporal fields. Group pacing, coherence diffusion, and jitter contagion without a group-mind claim. The emergent collective attractor is a property of coupling, not an agent over and above members. Guard: members retain distinct cones; no single shared 'now' and no group-mind implication.

7.2 Synchronization and the pacing result

Groups occupy metastable regimes with measurable dwell and transition statistics, and synchronization is central. Synchronization is stabilizing when a shared rhythm reduces net strain across members and destabilizing when the group locks to a pace tolerable only to its most resilient members, forcing the rest past their capacity. This yields a testable structural prediction: **collective stability is maximized when synchronization is paced by the most-constrained safe capacity, and most at risk when pacing is set by the highest-capacity members or by external demand that ignores the capacity distribution.** Group-level failure modes mirror individual ones — fragmented groups shift strategy erratically, rigid groups entrench, runaway-recursive groups drown in meta-process without substantive change.

8. Methods and Measurement

MRIF constructs are meant to be operationalized, and measurement is multi-level: individual behavior and experience, group dynamics, and model-based latent structure. The geometry of inference is approximated through behavioral response manifolds, latent-variable embeddings, and state-space reconstructions; the Architect Frame through adaptation to perturbation, strategy adjustment, and stability under structured stressors; the capacity envelope through performance-degradation curves, coherence-breakage thresholds, and recovery trajectories.

At the individual level, perturb-and-recover paradigms systematically raise then lower load while measuring performance, strategy change, and subjective reports, distinguishing a healthy capacity-respecting profile (graceful degradation, structured strategy shifts, coherent recovery) from saturation (abrupt breakdown, fragmented switching, lingering instability). Recursion-sensitivity tasks vary the number of requested meta-levels under time and salience constraints, locating the depth at which additional recursion stops helping and starts destabilizing. At the group level, coordination and pacing tasks manipulate pace, complexity, and capacity composition to probe collective fields and the pacing result, and contagion paradigms pair members of known profiles to quantify whether coherence or jitter diffuses. Across both, MRIF serves as a testing framework: two existing models of the same system can be checked for mirror-consistency by extracting their implied geometric and dynamical features and asking whether a reasonable Φ reconciles them, revealing whether disagreements reflect surface vocabulary or deeper structural incompatibility.

9. Structural Predictions and Falsifiability

A structural theory must be falsifiable. MRIF does not predict specific behaviors in isolation; it predicts patterns of viability, breakdown, and compatibility.

At the individual level: increasing inferential power (more complex models, deeper recursion, faster updates) without expanding the capacity envelope should increase the frequency and severity of fragmentation and rigidification, especially under sustained load, and interventions that expand the Architect Frame should reduce these even when content is unchanged. As load approaches and exceeds capacity, MRIF predicts a characteristic sequence of changes in the temporal projection and identifiable transitions among regulated, looping, and collapsed regimes, with hysteresis. Moderate, bounded recursion should improve coherence and performance; excessive recursion depth should show diminishing or negative returns and rising instability.

At the collective level: groups that synchronize to their most-constrained members should maintain higher collective coherence and fewer collapses than groups paced by their least-constrained; groups with explicit, well-designed collective Architect Frames should withstand higher load and recover more reliably. The three failure modes should appear at every scale with analogous signatures.

At the level of theories: two theories of the same phenomenon that cannot be recast into mirror-consistent geometric and dynamical views, or that imply contradictory capacity constraints, cannot both be complete. Models that assume costless arbitrarily deep recursion, ignore capacity or internal temporal geometry, or rely solely on global scalar time fall outside MRIF's structural domain and are predicted to mischaracterize failure modes under stress. Observing systems that consistently violate these patterns would challenge MRIF's structural assumptions.

10. Limits, Non-Claims, and Boundary Conditions

Because MRIF uses mathematical language and analogies that overlap with physics and dynamical systems, it is essential to be explicit about what the framework does not claim.

No quantum-cognition claims. MRIF uses structural analogies — state spaces, manifolds, invariants — that resemble formalisms used in physics and dynamical systems. It does not claim that brains or cognitive systems exploit quantum states in a way essential to the framework, nor that quantum mechanics directly explains subjective time or MRIF constructs. MRIF is structurally analogous to formalisms used in physics, with no claim of physical mechanism or evidential transfer. The parallel asserted is structural only. The same grammar recurs across regimes; no evidential weight transfers between them. A result in one regime is not evidence about another.

No determinism or teleology. MRIF describes constraints, not inevitabilities; capacity envelopes, not destiny. That certain trajectories are unlikely or unstable does not make future states predetermined or systems bound to a single path. No teleological claim is made: MRIF posits no built-in goals of nature and no inherent moral direction; it characterizes which patterns of inference and self-modeling can be sustained under given constraints.

No diagnostic categories. MRIF avoids pathologizing individuals and does not map regimes or failure modes onto psychiatric diagnoses. "Fragmentation," "rigidification," and "runaway recursion" are structural descriptors of how information and temporal geometry behave, not labels for people. MRIF may be useful for reasoning about clinical phenomena, but it is not a diagnostic manual; clinical applications must interface with established frameworks, respect ethical and contextual considerations, and avoid reifying MRIF constructs as labels applied to individuals.

Appropriate and inappropriate uses. Appropriate: comparing theoretical frameworks at a structural level; designing experiments and interventions that target capacity, pacing, and Architect-Frame-like mechanisms; reasoning about multi-scale coherence and failure in self-modeling systems. Inappropriate: treating MRIF as a direct mechanistic model of neural circuitry; using MRIF terms as clinical diagnoses or moral judgments; over-extending structural analogies to claim new physics or a literal quantum mind. These boundaries are part of MRIF's internal coherence: the framework is about structure and constraint, not metaphysical claims beyond its domain.

11. Discussion and Future Directions

MRIF was developed to answer a practical question: how can we reason coherently about systems that carry information, model themselves, and try to stay coherent under load, across scales and theoretical frameworks? The framework's contributions are a dual-description requirement (any viable model must admit mirror-consistent geometric and dynamical views linked by Φ), an explicit role for recursion and capacity (recursive self-modeling is both necessary and dangerous; capacity envelopes and Architect Frames determine which trajectories are viable), a structural account of failure modes (fragmentation, rigidification, runaway recursion as consequences of chronic over-capacity), a meta-theoretical layer coordinating the IGAF formal core and TNI, and an extension to collective fields through coupled Architect Frames.

For neuroscience and psychology, MRIF suggests reframing certain phenomena — chronic stress, burnout, some symptom clusters — as manifestations of capacity-envelope mismatch and Architect-Frame strain rather than solely as content-level distortions [1-5], and designing interventions that change pacing, load distribution, and recursion bounds while rebuilding capacity. For AI and multi-agent systems, it highlights safe recursion as a design requirement — bounding self-modification depth, monitoring capacity and strain analogues, implementing Architect-Frame-like governors — and pacing depth and computation to preserve coherence rather than maximizing throughput; it thereby complements work on alignment, interpretability, and robustness by emphasizing viability conditions and failure modes. For socio-technical systems, it frames overload, polarization, and institutional paralysis as breakdowns in collective Architect Frames and chronic operation above collective capacity, pointing to redesigned rhythms and explicit group-level regulation.

MRIF is intentionally unfinished — a structural scaffold, not a closed theory. Future work runs along four lines: sharpening the mathematics of mirror duality, capacity envelopes, and collective coupling (in Supplement II, where the formal objects live); empirical validation through perturb-and-recover, recursion-sensitivity, and coordination studies; tighter integration with the IGAF formal core and TNI as joint models; and applied development in capacity-aware clinical, AI-safety, and organizational tools. Its value will depend on whether researchers and practitioners find it useful for describing, testing, and repairing the complex systems we live in and build.

References

[1] Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127–138.

[2] Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology–Paris, 100*(1–3), 70–87.

[3] Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36*(3), 181–204.

[4] Hohwy, J. (2013). *The Predictive Mind.* Oxford University Press.

[5] Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences, 17*(11), 565–573.

[6] Amari, S. (2016). *Information Geometry and Its Applications.* Springer.