Recursive Containment and Deferred Closure

A Scope Dynamics Framework for Symbolic Cognition

Flyxion

Independent Researcher

May 27, 2026

Abstract

Symbolic cognition is typically modeled as sequential token processing over static representational states. We argue that this picture is incomplete in a fundamental way. The primitive structure underlying arithmetic evaluation, recursive computation, narrative comprehension, hypnotic induction, harmonic tension, and creative insight is not the token or the state, but the *semantic scope*: a bounded, potentially unresolved region of meaning that persists until lawfully closed. We develop a formal framework, rooted in the Spherepop calculus of nested containment, centered on a stack of open semantic bubbles $\Sigma_t = [B_1, \ldots, B_n]$, a load functional $L(\Sigma)$, and a family of transformation operators: open, pop, meld_π, reframe^c_ϕ, and reframe^e_ϕ. Each operator is defined with explicit admissibility conditions; inadmissibility is shown to correspond to recognizable cognitive and rhetorical pathologies. The framework is stratified: a tree-structured containment relation suffices for the core cases, while DAG containment is required for metaphor and cross-domain analogy. We advance an operator completeness conjecture as the central structural claim, examine the sheaf-theoretic reading of multi-parent gluing, and close with the philosophical consequences for models of intelligence.

Contents

1. The Hidden Geometry of Scope

1.1. The Problem of Linear Symbolic Models

Contemporary theories of cognition, language, and computation share a background assumption that is rarely made explicit: that meaning unfolds sequentially, one symbol following another, each state superseding its predecessor. Arithmetic is presented as a sequence of operations. Language comprehension is modeled as left-to-right token integration. Generative systems predict the next element conditioned on what has come before. Even working memory research tends to frame cognitive load as the number of items retained rather than the structure of dependencies among them.

This picture is not simply incomplete. It is geometrically misleading. The structure that makes arithmetic evaluable, narratives coherent, and trance possible is not sequential at all. It is hierarchical, containment-based, and essentially topological. The purpose of this paper is to make that structure explicit and to develop a formal framework adequate to its complexity.

1.2. Hidden Parentheses and the Normalization Insight

Consider the elementary expression

$$1 + 2 \times 3.$$

A student learning order of operations is told: multiplication before addition. But this rule, presented as an external convention, conceals a deeper fact. The expression already has a containment structure. Written honestly, it reads

$$\big(1 + (2 \times 3)\big).$$

The multiplication scope is *inside* the addition scope. Precedence is not a rule imposed on a flat string. It is a consequence of spatial depth. The inner scope must resolve before the outer scope can be evaluated, not because of convention, but because the outer scope's resolution function has the inner scope's output as an unresolved dependency.

This observation generalizes. Every well-formed arithmetic expression has a latent containment structure. The Spherepop normalization procedure makes that structure explicit. Given any expression, one proceeds as follows: first, insert all parentheses implied by conventional precedence, making every scope boundary visible; second, regard each parenthesized region as a closed geometric object, a circle in the plane; third, lift these circles into three-dimensional bubbles nested inside one another; fourth, evaluate from the innermost bubble outward, where each evaluation event is an irreversible collapse of a local scope into a result value that propagates upward to its parent.

The resulting *Spherepop Normal Form* of $1 + 2 \times 3$ is the reduction sequence

$$\big(1 + (2 \times 3)\big) \;\longrightarrow\; (1 + 6) \;\longrightarrow\; (7) \;\longrightarrow\; 7.$$

Each arrow is an irreversible semantic event. Each event reduces the unresolved load of the system by closing one scope. The final result is not merely a number. It is the residue of a history of nested commitments.

1.3. Scope as the Primitive of Symbolic Cognition

The normalization procedure just described extends far beyond arithmetic. Natural language contains nested scopes at every level of structure: relative clauses embedded inside noun phrases, subordinate clauses inside main clauses, direct speech inside narrative frames, hypotheticals inside arguments. The sentence

The boy who carried the box that held the bird ran home

contains at least four nested scopes of reference, each contributing unresolved expectations that must remain active in working memory until their containing clause is syntactically closed. Comprehension is not reading a flat string. It is navigating a dynamic containment structure.

Programming languages make this architecture visible in their syntax. Lexical scoping, block structure, function calls, and exception handlers are all mechanisms for creating, suspending, and closing semantic regions. The call stack maintained by any executing program is a direct implementation of the scope stack we will formalize below.

Our central claim, developed and defended throughout this paper, is the following: *scope* is the primitive structure of symbolic cognition. Not tokens. Not propositions. Not states. The bounded, potentially unresolved semantic region that persists until lawfully closed. This geometric reinterpretation of symbolic structure is philosophically aligned with approaches that privilege spatial and transformational intuition over purely sequential symbolic manipulation, most notably in geometric approaches to analysis that insist the underlying spatial structure of a domain is primary and the algebraic surface notation derived [10].

2. The Formal Framework

2.1. Bubbles and Containment Structures

Definition 2.1 (Semantic Bubble)**.** A *semantic bubble* is a tuple $B = (C, E, U)$ where C is a contextual binding set, E is an expectation structure encoding what resolutions are anticipated, and $U(B) \geq 0$ is a scalar measuring unresolved semantic load. A bubble B is *open* when $U(B) > 0$ and *closed* when $U(B) = 0$.

Definition 2.2 (Containment Structure)**.** A *containment structure* is a pair $\Sigma = (\mathcal{B}, \prec)$ where $\mathcal{B}$ is a finite set of semantic bubbles and $\prec$ is a strict partial order on $\mathcal{B}$ (irreflexive, transitive, antisymmetric). We write $B_i \prec B_j$ to mean B_i is contained within B_j, and $B_i \prec^* B_j$ for the transitive closure. The *covering relation* $B_i \lessdot B_j$ holds when $B_i \prec B_j$ and there is no B_k with $B_i \prec B_k \prec B_j$.

In the foundational treatment of this paper (Sections 3 through 4), we restrict to the case where $\prec$ induces a rooted forest: every bubble has at most one immediate parent. This tree condition is satisfied by arithmetic evaluation, recursive computation, ordinary narrative nesting, and hypnotic induction. Section 6 relaxes it to a DAG for the treatment of metaphor and creative cross-domain import.

2.2. The Scope Stack and Semantic Load

When the containment structure is a rooted tree, the current active path from root to frontier can be linearized as a stack.

Definition 2.3 (Scope Stack)**.** The *scope stack* at time t is the sequence

$$\Sigma_t = [B_1, B_2, \ldots, B_n]$$

where $B_1 \prec B_2 \prec \cdots \prec B_n$ is the active containment path and B_n is the currently attended scope. Each B_i is open: $U(B_i) > 0$.

Definition 2.4 (Semantic Load Functional)**.** The *semantic load* of a scope stack $\Sigma = [B_1, \ldots, B_n]$ is

$$L(\Sigma) = \sum_{i=1}^{n} w_i \, U(B_i)$$

where $w_i = f(d_i, s_i, r_i, c_i)$ is a weight combining stack depth d_i, emotional salience s_i, recency r_i, and causal centrality c_i.

The decomposition of w_i into four partially independent dimensions is essential. Semantic dominance is not purely hierarchical. A shallow but emotionally charged bubble may exert greater gravitational pull on attention than a deeper but neutral one. The weight function f is left underspecified here; what matters structurally is that the dimensions can dissociate, producing the observed phenomena of shallow trauma overriding deep abstraction, emotionally salient interruptions disrupting logical flow, and central narrative frames organizing subordinate bubbles.

2.3. Admissibility and the Resolution Rule

The resolution rule is not a component of the bubble tuple but a transformation operator defined separately. This objects-versus-morphisms separation is deliberate: bubbles are semantic regions; operators are lawful transformations over semantic regions.

Definition 2.5 (Resolution Operator)**.** For each bubble $B \in \Sigma$, the *resolution operator* is a partial function

$$\rho : B \times \Sigma \rightharpoonup B'$$

where $\rho(B, \Sigma)$ is defined if and only if all bubbles B_j with $B_j \prec B$ (all interior descendants of B in the containment order) satisfy $U(B_j) = 0$. Note that B' denotes the updated bubble after resolution, not a descendant; we use primed notation throughout for post-operation states.

The partiality is not a technical convenience. It is the formal expression of the framework's central invariant.

Definition 2.6 (Well-Nestedness)**.** A scope stack Σ is *well-nested* with respect to a candidate pop operation at position i if all B_j with $B_j \prec B_i$ satisfy $U(B_j) = 0$.

Proposition 2.7. *Well-nestedness is a necessary condition for the admissibility of any resolution operation. A parent scope cannot lawfully close while any descendant scope remains open.*

This proposition, whose proof follows directly from the domain condition of ρ, unifies arithmetic evaluation (innermost subexpression first), recursive function return (callee before caller), proof completion (lemma before theorem), narrative closure (embedded tale before framing tale), and hypnotic release (innermost story before outer story). In every case, a parent cannot lawfully resolve while dependent subordinates remain open.

The listener's phenomenological sense of *unfinishedness* is mathematically meaningful under this framework: the bubble literally lacks a defined collapse morphism.

2.4. Persistence and Suspended Activity

A semantic bubble does not cease to exert causal influence merely because it is no longer the foreground scope. Suspended scopes remain dynamically active through their contribution to the semantic load functional $L(\Sigma)$ and through unresolved dependency relations propagating upward through the containment structure. This persistence distinguishes semantic bubbles from ordinary stack frames in classical computational models. A conventional execution frame becomes causally inert once control leaves it, except insofar as it preserves local state for eventual return. In symbolic cognition, by contrast, suspended scopes continue exerting attentional and affective pressure even while subordinate scopes temporarily dominate conscious focus. The unresolved narrative introduced at the opening of a story continues to shape expectation while intervening events unfold. A suspended emotional frame continues organizing perception even when not explicitly attended. An open proof obligation continues constraining mathematical reasoning while auxiliary lemmas are explored. Persistence is therefore not accidental to the framework but constitutive of it. Containment structures are properly interpreted not as static trees or DAGs but as temporally evolving semantic fields whose unresolved regions remain causally active across time until lawfully collapsed. This interpretation is compatible with dynamical accounts of cognition in which causation operates through evolving constraint structures rather than through isolated efficient interactions [5]. More broadly, this connects the present formalism to event-history ontologies in which the world is constituted not by static states but by accumulated irreversible commitments whose consequences propagate forward through time.

3. The Operator Calculus

3.1. Five Primitive Operators

We define five operators on containment structures, each with explicit admissibility conditions.

Definition 3.1 (Open)**.**

$$\text{open}(B) : \Sigma \mapsto \Sigma \cdot B$$

Opens a new scope by appending B to the current stack. B must be a fresh bubble not already in Σ. $U(B) > 0$ (newly opened scopes are by definition unresolved).

Definition 3.2 (Pop)**.**

$$\text{pop}(B_n) : [B_1, \ldots, B_{n-1}, B_n] \mapsto [B_1, \ldots, B'_{n-1}]$$

Closes the topmost scope. Admissible only when Σ is well-nested at position n. The parent B_{n-1} is modified to B'_{n-1}, reflecting the propagation of the resolved content upward. Closure is not isolated: it changes the parent.

Definition 3.3 (Meld)**.**

$$\text{meld}_\pi(B_i, B_j) : \Sigma \mapsto \Sigma'$$

where B_i and B_j are sibling scopes (sharing a common parent but neither containing the other), π is a compatibility policy, and $B_i \sim_\pi B_j$ denotes compatibility under π. The result Σ' replaces B_i and B_j with a single merged bubble B_{ij} at the same depth. Admissible only when $B_i \sim_\pi B_j$; inadmissible when compatibility fails.

Pop and Meld differ in a fundamental way. Pop is strictly causal and sequential: it closes what was opened, in order. Meld is convergent and cross-scope: it identifies sibling regions through shared resolution structure. A joke callback is a Meld, not a Pop. The punchline does not merely close the most recent frame; it retroactively identifies an earlier frame with the present one, producing simultaneous collapse across distant narrative regions. Musical cadence often behaves the same way, retroactively gathering several unresolved harmonic scopes into a unified resolution rather than closing them one by one.

Definition 3.4 (Conservative Reframe)**.**

$$\text{reframe}^c_\phi : (B, \prec) \to (B, \prec')$$

where ϕ acts as the identity on B, $\prec'$ is a strict partial order, and $\prec^* = \prec'^*$. The transitive closure is preserved exactly; only the covering relation changes. Conservative reframe redistributes containment without creating or destroying ancestral dependencies.

Definition 3.5 (Expansive Reframe)**.**

$$\text{reframe}^e_\phi : (B, \prec) \to (B, \prec')$$

where ϕ acts as the identity on B, $\prec'$ is a strict partial order, and $\prec^* \subseteq \prec'^*$. The transitive closure is extended; new ancestral dependencies may be introduced. Expansive reframe imports new containment relations from outside the existing structure.

Definition 3.6 (Admissibility of Reframe)**.** A reframe reframe_ϕ of either species is *admissible* only if $\prec'$ is acyclic: for no B_i does $B_i \prec'^* B_i$.

Conservative reframe captures insight, therapeutic reinterpretation, and plot twists: the same historical events become differently organized in the containment hierarchy without any new facts being introduced. Expansive reframe captures metaphor, analogy, and scientific paradigm shift: a semantic region begins participating in a containment structure that was previously foreign to it.

3.2. Inadmissibility and Cognitive Pathology

The admissibility conditions are not merely technical safeguards. Each failure mode corresponds to a recognizable cognitive or rhetorical pathology.

Proposition 3.7 (Inadmissibility Taxonomy)**.** *The following four conditions are inadmissible and correspond to the named phenomena.*

(i) $\text{pop}(B_i)$ *applied when some* $B_j \prec B_i$ *has* $U(B_j) > 0$*:* premature closure*. The narrative resolves before its embedded obligations are discharged.*

(ii) $\text{meld}_\pi(B_i, B_j)$ *applied when* $B_i \not\sim_\pi B_j$*:* forced synthesis*. Two incompatible frames are collapsed as though compatible, producing spurious apparent coherence.*

(iii) reframe_ϕ *producing a cycle* $B_i \prec'^* B_i$*:* semantic self-containment*. A scope depends on itself, producing rumination, paradox, or incoherent circular explanation.*

(iv) reframe^e_ϕ *introducing ancestors incompatible with existing gluing conditions across parent scopes:* incoherent metaphor*. The new containment relation cannot be made simultaneously consistent with all existing parent structures.*

There is also a dual failure for conservative reframe worth naming separately. A reframe^c_ϕ that preserves the transitive closure but severs all covering relations produces a structure in which every bubble is ancestrally related to every other but nothing is locally adjacent to anything. This is *semantic dissociation*: global coherence with no local navigability.

3.3. Two Fundamental Theorems

Theorem 3.8 (Termination of Admissible Pop Sequences)**.** *Every admissible Pop sequence over a finite well-nested tree containment structure terminates in finitely many steps.*

Proof. Define the *depth* of a containment structure $(B, \prec)$ as

$$\delta(\Sigma) = \max\{k : \exists\, B_{i_1} \prec B_{i_2} \prec \cdots \prec B_{i_k}\}.$$

Since B is finite, $\delta(\Sigma)$ is a non-negative integer. Each admissible Pop operation removes the topmost bubble B_n from the active stack and reduces its parent's unresolved load, strictly decreasing $|B|$ by one. Because $|B|$ is finite and strictly decreasing under every admissible Pop, and because the well-nestedness condition ensures no Pop can be applied to a bubble with unresolved descendants (preventing indefinite suspension), the sequence of admissible Pops must reach the empty stack $\Sigma = [\,]$ in at most $|B|$ steps. □

Theorem 3.9 (Conservative Reframe Preserves Acyclicity)**.** *If $(B, \prec)$ is a strict partial order and* reframe^c_ϕ *is an admissible conservative reframe, then $(B, \prec')$ is also a strict partial order.*

Proof. We verify that $\prec'$ satisfies irreflexivity, asymmetry, and transitivity.

Irreflexivity. Suppose for contradiction that $B_i \prec' B_i$ for some B_i. Then $B_i \prec'^* B_i$, which contradicts the admissibility condition that $\prec'$ be acyclic.

Transitivity. If $B_i \prec' B_j$ and $B_j \prec' B_k$, then $B_i \prec'^* B_k$ by definition of transitive closure. Since $\prec^* = \prec'^*$ (the conservative reframe condition), we have $B_i \prec^* B_k$. The structure of $\prec'$ is chosen to realize this ancestry through some covering path, so $B_i \prec' B_k$ or $B_i \prec'^* B_k$ through intermediate nodes; either way transitivity holds.

Asymmetry. If both $B_i \prec' B_j$ and $B_j \prec' B_i$, then $B_i \prec'^* B_i$, contradicting acyclicity. Hence $\prec'$ is asymmetric.

Since $\prec'$ is irreflexive, asymmetric, and transitive, it is a strict partial order. □

Theorem 3.10 (Monotonicity of Semantic Load Under Pure Pop Sequences). *Let $\Sigma_0 \to \Sigma_1 \to \cdots \to \Sigma_n$ be a sequence consisting only of admissible Pop operations with non-negative weights $w_i \geq 0$. Then*

$$L(\Sigma_{k+1}) \leq L(\Sigma_k)$$

for all $k \in \{0, \ldots, n-1\}$.

Proof. Each admissible Pop removes the topmost open bubble B_i satisfying $U(B_i) > 0$, contributing $w_i U(B_i) > 0$ to $L(\Sigma_k)$. Propagation effects modify the parent bubble B_{i-1} to B'_{i-1}, but by the admissibility condition all descendants of B_i are already closed ($U(B_j) = 0$ for $B_j \prec B_i$), so no new unresolved load is introduced by the propagation. Since $w_i \geq 0$ and $U(B_i) > 0$, removing B_i strictly reduces the sum $\sum_j w_j U(B_j)$, yielding $L(\Sigma_{k+1}) < L(\Sigma_k)$. In the degenerate case $U(B_i) = 0$ the bubble is already closed and the Pop is trivially admissible with no change to L, giving $L(\Sigma_{k+1}) = L(\Sigma_k)$. In both cases $L(\Sigma_{k+1}) \leq L(\Sigma_k)$. □

Remark 3.11. Monotonicity fails for sequences containing Meld or Reframe operations. Meld may temporarily increase local load during cross-scope identification before convergence reduces it. Expansive Reframe by definition introduces new ancestral dependencies, potentially increasing $L(\Sigma)$ before the new structure stabilizes. This asymmetry reflects the phenomenology of creative work and therapeutic restructuring: the path to lower total load sometimes requires passing through higher-load intermediate configurations.

3.4. Operator Completeness

Conjecture 3.12 (Operator Completeness). *Every admissible transformation of a containment structure $(B, \prec)$ that preserves node identity can be expressed as a finite composition of operations drawn from* $\{\text{open}, \text{pop}, \text{meld}_\pi, \text{reframe}^c_\phi, \text{reframe}^e_\phi\}$.

The qualifier *preserves node identity* is essential: it excludes deletion (forgetting) and fabrication (confabulation), which are genuine cognitive operations but lie outside the class of lawful semantic transformations considered here. The boundary of this conjecture identifies precisely where memory failure and motivated reasoning begin.

3.5. Collapse Quotients and Irreversibility

The Pop operator is irreversible in a structural rather than merely temporal sense. When a scope collapses, its internal generative history is compressed into a representative semantic residue. The resulting structure preserves resolution outcome while discarding recoverable intermediate organization. This compression can be formalized through a quotient construction. Let $\sim_\rho$ be the equivalence relation on bubble histories induced by admissible resolution under collapse operator ρ: two histories are equivalent when they produce the same admissible semantic representative. The collapsed bubble is then the quotient $B/{\sim_\rho}$.

The expression $(2+3)$ and the value 5 illustrate the distinction concretely. The unresolved scope $(2+3)$ and the collapsed residue 5 are extensionally equivalent but not structurally equivalent. One preserves generative history; the other preserves only the lawful result of that history. Abstraction is therefore not merely symbolic representation. It is the quotienting of semantic histories by admissible collapse relations. This account of abstraction as quotient connects the formal framework to the broader theme of irreversibility that runs through the paper: what is lost in a collapse is not content but provenance. The bubble 5 is authoritative precisely because it no longer carries the contingency of the history that produced it. It excludes deletion (forgetting) and fabrication (confabulation), which are genuine cognitive operations but outside the scope of lawful semantic transformation as defined here. The boundary of this conjecture thus identifies precisely where memory failure and motivated reasoning begin.

4. Conversational Hypnosis as Scope Engineering

4.1. The Nested Narrative Induction Pattern

Conversational hypnosis, as practiced and systematized in the tradition descended from Milton Erickson and formalized in later pedagogical work on indirect induction, relies centrally on a structural technique: the deliberate creation and suspension of nested narrative scopes. A practitioner opens a story, interrupts it before closure to open a second story, interrupts that to open a third, and then resolves them in reverse order. The listener's experience is one of deepening absorption, suspended between multiple unfinished narrative worlds.

What the Spherepop framework reveals is that this technique is not a rhetorical trick layered on top of cognition. It is a direct exploitation of cognition's native scope-management architecture.

4.2. Formal Model of Induction

Let the listener's active meaning-state at time t be the scope stack

$$\Sigma_t = [B_1, B_2, \ldots, B_n]$$

where each B_i is an open semantic bubble: an unfinished story, image, question, expectation, or emotional frame. Opening a narrative is a push:

$$\mathrm{open}(B) : \Sigma \mapsto \Sigma \cdot B.$$

Interruption before closure does not delete the suspended story. It suspends it below the new scope:

$$[B_1] \xrightarrow{\mathrm{open}(B_2)} [B_1, B_2].$$

After several interruptions the listener carries

$$\Sigma = [B_1, B_2, B_3, \ldots, B_n].$$

Resolution proceeds outward under the well-nestedness constraint, popping the top scope first:

$$[B_1, B_2, B_3] \;\longrightarrow\; [B_1, B_2'] \;\longrightarrow\; [B_1'] \;\longrightarrow\; [\,].$$

Each arrow is an irreversible semantic event. The final empty stack corresponds to the completion of induction. The hypnotic release is the outward cascade of scope completion.

4.3. Trance as a Region in Σ-Space

The semantic load functional $L(\Sigma) = \sum_i w_i U(B_i)$ provides a precise characterization of the inductive state. Trance is not an altered or mysterious state of consciousness requiring separate explanation. It is a dynamically stable region in Σ-space characterized by four conditions: high total unresolved semantic load $L(\Sigma)$, lawful nesting throughout, deferred closure at every level, and maintained global coherence.

The fourth condition is critical and distinguishes trance from confusion. A randomly incoherent stack would produce fragmentation. The hypnotic state depends on the bubbles remaining globally navigable despite their suspension. This is the sheaf condition operating on a tree: local sections (individual bubble contents) must be mutually compatible even while global resolution is deferred.

4.4. Failed Induction as Failed Gluing

Conversely, failed hypnotic induction corresponds formally to a failure of the gluing condition. When the narrative scopes the practitioner opens cease to overlap coherently, the listener cannot maintain global semantic navigability. The subject drops out not because they resist, but because the semantic manifold has fractured. The charts no longer paste together. At that point no further deferred closure is possible because the structure is no longer coherent enough to sustain it.

This analysis has a testable implication: the probability of successful induction should be related to the practitioner's ability to maintain pairwise compatibility between suspended narrative frames throughout the induction sequence. Compatibility failure at any intermediate step propagates upward and undermines the global structure.

5. Cross-Domain Instantiations

5.1. The Interpretation Map

The formal structure developed in Section 3 is deliberately domain-neutral. The scope stack Σ_t, the bubble tuple $B = (C, E, U)$, and the operator family all admit interpretation in any domain where nested unresolved regions are created, sustained, and eventually closed. What varies across domains is the semantic realization of each bubble, not the underlying geometry.

We can express this as a family of interpretation maps: for each domain $\mathcal{D}$, there is a function $\iota_{\mathcal{D}}$ assigning to each abstract bubble B a domain-specific semantic object, and to each operator a domain-specific transformation. The formal properties, admissibility conditions, and inadmissibility pathologies transfer unchanged. We use the term "interpretation map" rather than "functor" advisedly: while the construction has a category-theoretic flavor, we do not here fix the morphism structure needed for a full functorial account.

5.2. Arithmetic and Computation

The arithmetic interpretation has already been developed in Section 3. Here $\iota_{\text{arith}}(B_i) =$ an expression scope, $U(B_i)$ measures remaining unevaluated subexpressions, and Pop corresponds to evaluation of the innermost subexpression.

The programming interpretation is structurally identical: $\iota_{\text{prog}}(B_i) =$ an execution context, corresponding to a frame on the call stack, and Pop corresponds to a function return.

5.3. Humor and Joke Structure

A joke creates an initial expectation frame B_1, then builds narrative momentum inside it. The punchline is not a simple Pop of B_1: it retroactively reveals that the entire frame was nested inside a larger scope B_0 whose existence was not previously signaled. The pop cascade is $[B_0, B_1] \to [B_0'] \to [\ ]$, but what produces the humor is the sudden realization that B_0 existed all along. This is a conservative reframe immediately followed by Pop: the topology of the joke is revised in the moment of delivery, and then immediately collapsed. The relief and pleasure of laughter corresponds to the sudden reduction of $L(\Sigma)$.

A callback joke adds a further layer: an earlier frame B_{old} from a previous joke, apparently closed, is revealed to still be open (or more precisely, is reopened by the callback). This is closest to Meld: two apparently independent frames are identified through a newly revealed compatibility policy π.

5.4. Harmonic Tension and Musical Cadence

In tonal music, a harmonic progression opens expectation scopes that seek resolution. The dominant seventh chord is perhaps the most familiar example: it creates a strong expectation bubble whose resolution is the tonic. Extended harmonic sequences build nested layers of tension at multiple temporal scales, from the phrase to the section to the movement. A perfect authentic cadence performs a Pop cascade across several of these layers simultaneously. A deceptive cadence performs a conservative reframe: the expected Pop target (the tonic) is replaced by an unexpected but grammatically admissible alternative, preserving the ancestral harmonic logic while reassigning the immediate parent of the resolution event. The emotional effect of a deceptive cadence, surprise combined with recognition, corresponds precisely to the phenomenology of conservative reframe.

5.5. Proof Structure

A mathematical proof opens the theorem's claim as a top-level scope. Lemmas open subordinate scopes, each adding to the semantic load. The proof is complete when every lemma scope has been closed by argument and the theorem scope is closed by appeal to the assembled lemmas. The well-nestedness constraint is not merely a stylistic convention of mathematical writing; it is the structural condition that makes proofs valid. A proof that attempts to invoke a lemma whose own proof is incomplete is formally inadmissible under exactly the Pop admissibility condition. The experience of a proof clicking into place is the experience of a large Pop cascade: the deepest auxiliary results resolve first, their resolutions propagate outward, and the main theorem's scope closes over a fully resolved landscape.

6. DAG Containment and the Geometry of Metaphor

6.1. The Limits of Tree Containment

The tree-structured containment relation introduced in Section 3 is sufficient for all cases examined so far. But it imposes a condition that human cognition demonstrably violates: every bubble has at most one immediate parent. Metaphor, analogy, and creative cross-domain reasoning require that a single semantic object participate simultaneously in two or more containment structures that are otherwise disjoint.

When we say *argument is war*, we are not placing the domain of argument inside the domain of war or vice versa. We are asserting that a semantic region participates simultaneously in both containment structures, inheriting properties from each. This multi-parent participation is not expressible in a rooted tree. It requires a DAG.

Definition 6.1 (DAG Containment Structure)**.** A *DAG containment structure* is a pair $\Sigma = (B, \prec)$ where B is a finite set of bubbles and $\prec$ is a strict partial order on B without the single-parent restriction. The Hasse diagram $H(\Sigma)$ is a directed acyclic graph.

6.2. Metaphorical Participation

In a DAG containment structure, a bubble B_i can have multiple immediate parents $\pi_1, \pi_2, \ldots$ from disjoint containment chains. This constitutes *metaphorical participation*: B_i simultaneously inherits semantic structure from multiple domains.

This is not a meld operation. Meld closes two compatible sibling scopes into one. Metaphorical participation maintains two parent scopes while the child bubble remains open, inheriting from both without collapsing either. The distinction is important: meld is post-resolution convergence, while metaphorical participation is pre-resolution inheritance.

6.3. Gluing Conditions and Metaphor Vitality

Over a tree, the gluing condition (that local semantic sections be mutually compatible) is nearly trivially satisfiable: each bubble has only one parent to be compatible with. Over a DAG, gluing requires that a bubble's semantic content be simultaneously compatible with all of its parent scopes.

This provides a precise criterion for metaphor vitality. A *live metaphor* is one in which the multi-parent gluing condition is satisfied but nontrivially so: the bubble genuinely inherits from both parent domains, and the overlap tension between them is semantically productive. A *dead metaphor* is one in which the multi-parent structure has collapsed into conventional monadic participation: one parent has effectively become vestigial, leaving the bubble with a single dominant parent and no productive tension. The semantic curvature has flattened.

Metaphor quality thus corresponds to sustained admissible overlap without forced collapse. The best metaphors maintain genuine multi-parent inheritance over extended cognitive engagement. They resist the conventionalizing collapse that would turn them into mere synonymy.

6.4. Creative Cognition as Controlled DAG Expansion

The expansive reframe operator $\text{reframe}^{e}_{\phi}$ takes on its fullest meaning in the DAG setting. When a thinker introduces a new analogy or metaphor, they are performing an expansive reframe on the current containment structure: adding new ancestral dependencies that link previously disjoint domains. The admissibility condition requires that the new structure remain acyclic and that the gluing conditions for multi-parent bubbles remain satisfiable.

Creative insight at its deepest level may therefore consist in discovering that a DAG expansion is admissible: that two previously disjoint containment domains can be lawfully bridged without incoherence. The experience of insight, simultaneous surprise and rightness, corresponds to the realization that the new gluing condition is satisfied.

6.5. Semantic Curvature and Attention Flow

The semantic load functional $L(\Sigma)$ admits a geometric interpretation. Unresolved semantic regions generate gradients in attentional space, bending cognitive trajectories toward suspended dependencies awaiting closure. A highly weighted unresolved bubble exerts stronger curvature, drawing working memory, expectation, and interpretive effort toward itself even when not currently foregrounded. This geometric interpretation is compatible with the physical analogy suggested by the RSVP framework: just as regions of unresolved field tension generate flow toward equilibrium, open semantic scopes generate attentional flow toward resolution.

Under this interpretation, suspense corresponds to sustained unresolved semantic curvature; curiosity corresponds to directed traversal toward unresolved regions of the containment structure; distraction corresponds to competing curvature wells drawing attention in incompatible directions; and insight corresponds to a topological simplification that substantially reduces global curvature complexity through a well-timed sequence of pops and conservative reframes. The listener's experience of wanting to know what happens next is thus not merely emotional anticipation. It is geometric attraction within semantic space. Narrative tension, proof obligation, musical expectation, and unfinished thought all manifest as unresolved curvature acting on attentional flow. A practitioner of conversational hypnosis, by this account, is sculpting a curvature landscape: opening regions of high semantic load, maintaining their coherence, and timing the cascade of resolution for maximum cognitive and affective effect.

This geometric reading also suggests a natural extension of the framework toward quantitative cognitive modeling. The weight function $w_i = f(d_i, s_i, r_i, c_i)$ already encodes the principal dimensions along which semantic curvature varies. A full theory would supply empirical content for f connecting stack depth, emotional salience, recency, and causal centrality to measurable attentional and physiological signatures. That empirical program lies beyond the scope of this paper but is a natural downstream development of the formal framework.

7. Philosophical Consequences

7.1. Against Token-Linear Models of Cognition

The framework developed here constitutes a direct argument against token-linear models of symbolic cognition. If scope is primitive and not derived, then any model that processes symbols sequentially without maintaining persistent unresolved nested structure is cognitively inadequate in a structural sense, not merely in a performance sense.

This is not an argument about computational power. A Turing-complete token-linear system can simulate any containment structure. The argument is about cognitive architecture: the natural grain of human symbolic cognition is hierarchical containment, not sequential token succession, and models built on the wrong grain will fail in characteristic ways.

7.2. Premature Closure as Cognitive Failure Mode

The inadmissibility condition on Pop identifies premature closure as the most basic cognitive failure mode: resolving a scope before its dependencies are discharged. This manifests as: false certainty in reasoning (concluding before evidence is fully integrated), narrative impatience (resolving a story's tension before its structure warrants it), therapeutic failure (accepting an explanation of distress before the underlying frame has been examined), and aesthetic crudeness (collapsing a work's productive ambiguity into single-interpretation resolution before the audience has inhabited the tension). In every case, the formal signature is the same: $\text{pop}(B_i)$ applied while some descendant $B_j \prec B_i$ satisfies $U(B_j) > 0$.

7.3. Intelligence as Disciplined Containment

The deepest claim this framework supports is that intelligence is not primarily prediction, inference, or pattern matching. It is the disciplined management of nested unresolved semantic structure: knowing when to open a scope, how to maintain it without premature closure, when to meld compatible regions, and when a conservative or expansive reframe is warranted.

Ordinary push-pop systems treat Open and Pop as their primitive operations. While a Turing-complete stack machine can simulate Meld and Reframe computationally, it does not admit them as primitive semantic operators over containment topology: the operations must be encoded indirectly rather than expressed directly in the architecture's native grain. The cognitive claim is architectural, not computational. The capacity for Meld enables synthesis. The capacity for conservative Reframe enables insight, therapy, and narrative transformation. The capacity for expansive Reframe enables metaphor, analogy, and creative cross-domain thinking. These capacities are not incremental improvements on basic scope management. They are qualitatively distinct operations that extend the operator basis in ways that standard stack architectures cannot simulate at the cognitive level, even while simulating them computationally.

7.4. The Pop as Primitive of Meaning

We close with the claim that motivates the entire framework. Meaning is not a static relation between symbols and referents. It is the residue of an irreversible act of semantic commitment: the closing of a scope that was genuinely open. The value 5 means something different from $(2+3)$ not because they refer to different objects, but because one of them carries provenance and the other carries only the compressed result of a history that is no longer recoverable. Abstraction is not representation. It is the quotient of a history of nested commitments by the relation of resolution-equivalence.

This is why the joy of understanding feels the way it does. It is not the passive recognition of a pattern. It is the active collapse of a structure that was genuinely, productively, lawfully open.

8. Toward Scope-Based Artificial Intelligence

Most contemporary artificial intelligence systems are built around local continuation dynamics over token sequences. While such systems can simulate nested structure implicitly through learned statistical dependencies, they do not generally maintain explicit persistent unresolved semantic scopes as primitive architectural objects. The consequence is a characteristic failure mode: premature closure. When local coherence pressures compete with the demands of sustained unresolved structure, continuation-based systems tend to collapse open scope too early, generating locally fluent but globally incoherent output.

The framework developed here suggests a different architectural possibility: intelligence as the lawful management of containment topology. Under such an architecture, unresolved semantic regions would persist explicitly across time as first-class objects; admissibility conditions would regulate collapse operations, preventing premature resolution; and reframing would operate directly on semantic inheritance structure rather than emerging indirectly from statistical continuation weights.

The distinction is potentially foundational. Human cognition appears capable of sustaining genuinely unresolved structure for extended durations while dynamically reorganizing containment topology through conservative and expansive reframing. Creativity, insight, metaphor, and deep narrative understanding may depend less on predictive accuracy than on the controlled maintenance and transformation of unresolved semantic geometry. A system that can open a scope, hold it lawfully open under attentional pressure, meld it with compatible sibling scopes, and reframe the containment hierarchy when new structure becomes available would exhibit qualitatively different cognitive behavior from a system whose architecture provides no native representation of unresolved semantic load. This paper has developed the formal vocabulary for specifying what such a system would need. What remains is the engineering.

The framework is already diagnostically useful for characterizing the failure modes of existing systems. Hallucination corresponds to inadmissible expansive reframe: new ancestral dependencies are introduced without satisfying gluing conditions, producing locally fluent but globally incoherent output. Premature summarization is illegal Pop: a scope is collapsed before its interior dependencies are discharged. Sycophancy is forced Meld: two frames are identified as compatible under social pressure rather than genuine π-compatibility. Context-window degradation is scope-density overload: when $D(\Omega)$ exceeds the system's representational capacity, containment relations collapse indiscriminately. Mode collapse is semantic curvature minimization: the system gravitates toward low-$L(\Sigma)$ configurations by prematurely resolving all open scopes, eliminating the productive tension on which creative and critical cognition depends. Incoherent long-range reasoning is failure to preserve unresolved ancestry relations across time: the system loses track of open obligations as the scope stack grows. Each failure mode has a precise formal signature under the present framework, which suggests that the framework could serve not merely as a design specification but as a diagnostic instrument for evaluating existing architectures.

The present framework is compatible with dynamical accounts of meaning and intentionality that treat causation as constraint propagation through evolving constraint landscapes rather than linear efficient causation [5]. The inadmissibility of cyclic containment relations bears partial resemblance to productive and pathological forms of recursive self-reference explored in studies of strange loops [4]. The DAG treatment of metaphor developed here overlaps partially with conceptual metaphor theory and conceptual blending, though the present framework emphasizes containment topology and admissible gluing rather than linguistic mapping alone [6, 3]. The containment interpretation shares structural affinities with lambda calculus in its treatment of scope, binding, reduction, and nested substitution [1], and with the recursion-scheme tradition in functional programming [7]. The categorical vocabulary of gluing, compositionality, and morphism structure throughout the paper resonates with the foundations of category theory [9].

8.1. Scope Closure and Temporal Directionality

The asymmetry between open and closed scopes introduces an intrinsic temporal orientation into the containment structure. An open bubble contains multiple admissible future trajectories, while a collapsed bubble represents a reduced equivalence class in which those trajectories have been quotient-compressed into a resolved semantic residue. This gives the framework an internal arrow of semantic time. The distinction between future and past is not imposed externally upon the system but emerges from the irreversible reduction of admissible possibility under lawful Pop operations. A resolved scope cannot be reopened without either reconstructing its generative history or introducing a new expansive reframe. The semantic directionality of cognition therefore parallels thermodynamic irreversibility: both involve the progressive reduction of accessible configuration structure under lawful evolution operators. The analogy is structural rather than physical, but it suggests that the irreversibility explored throughout this paper is not a special feature of symbolic systems but an instance of a more general principle governing constraint-governed dynamical processes.

8.2. Local and Global Coherence

The gluing condition introduced in connection with DAG containment admits an important distinction between local and global semantic coherence. A collection of bubbles may be pairwise compatible while still failing to admit a globally coherent containment structure. Local compatibility does not guarantee the existence of a globally consistent section over the entire containment manifold. This distinction appears throughout cognition. Conspiracy systems often maintain strong local coherence between neighboring explanatory frames while failing globally under broader integration. Fragmented narratives may preserve intelligibility scene by scene while collapsing under large-scale interpretive traversal. Formally, many cognitive pathologies arise not from local incoherence but from failures of global integrability: each adjacent pair of scopes agrees, but no globally admissible configuration exists that satisfies all constraints simultaneously. This is precisely the condition that sheaf theory identifies as the failure of local sections to extend to a global section, and it provides a precise mathematical characterization of a wide class of reasoning failures that are otherwise difficult to describe.

8.3. Semantic Inertia

Not all containment structures reconfigure with equal ease. Certain bubbles exhibit strong resistance to admissible reframing even when alternative containment relations are available. We refer to this resistance as *semantic inertia*. Semantic inertia may arise from emotional salience, repeated reinforcement, causal centrality, or dense integration with surrounding scopes. Highly inertial bubbles resist both conservative and expansive reframing because modification would propagate destabilization across large portions of the containment graph. This provides a structural interpretation of ideological rigidity, habit-formation, and persistent self-concepts. Stability is not merely a psychological preference but a topological property of deeply integrated semantic regions. An inertial bubble functions as a load-bearing node: it is not merely heavy in the sense of high $w_i U(B_i)$, but structurally central in the sense that its containment relations are shared by many other bubbles. Reframing it requires simultaneously satisfying compatibility conditions across all its dependent scopes, which may be collectively unsatisfiable even when individually admissible alternatives exist.

8.4. Cascade Thresholds and Semantic Phase Transitions

The semantic load functional $L(\Sigma)$ does not merely measure unresolved structure statically. Beyond certain critical values, local operations may propagate nonlocally through the containment graph, producing large-scale reorganization. We therefore introduce a critical load parameter L_c such that a containment structure enters a *cascade-sensitive regime* whenever

$$L(\Sigma) > L_c.$$

Definition 8.1 (Cascade Threshold)**.** A containment structure is *cascade-sensitive* when $L(\Sigma) > L_c$ for some context-dependent critical load parameter L_c, beyond which admissible local operations propagate nonlocally through the containment graph.

In a cascade-sensitive regime, a single admissible Pop, Meld, or Reframe may trigger widespread restructuring across the containment topology. The collapse of one heavily weighted bubble can suddenly reduce the unresolved status of multiple dependent scopes, producing a rapid outward chain of semantic closure. Humor, emotional breakthrough, sudden mathematical insight, and certain forms of hypnotic release all exhibit this phenomenology: a prolonged period of unresolved tension followed by abrupt large-scale reorganization.

The critical threshold L_c is context-dependent and varies with the stability of the surrounding containment structure. Highly coherent semantic manifolds can sustain larger unresolved loads before cascading, while weakly glued structures may collapse under comparatively small perturbations. Semantic stability therefore depends not merely on total unresolved load but on the geometry through which that load is distributed.

8.5. Semantic Resonance and Recurrent Scope Activation

Not all suspended bubbles remain equally active through time. Certain scopes recur repeatedly across otherwise unrelated containment paths, reappearing as persistent attractors in cognition, narrative, or emotional interpretation. A *resonant bubble* B_r is characterized by repeated reactivation across distinct regions of Σ:

$$B_r \rightsquigarrow \{\Sigma^{(1)}, \Sigma^{(2)}, \ldots\}.$$

The bubble's unresolved structure propagates into multiple future contexts, modulating interpretation even after partial closure events have occurred. In narrative, resonance appears as thematic recurrence or symbolic motif. In trauma, a partially unresolved frame repeatedly reorganizes unrelated present experience around itself. In mathematics, an unresolved conceptual tension may quietly structure years of investigation before eventual resolution. In creative work, resonance often precedes explicit insight: a scope continues reappearing because the containment system has not yet discovered the admissible reframe capable of integrating it coherently. Resonance acts as a persistence amplifier: a highly resonant bubble maintains causal influence not merely because it remains unresolved, but because the surrounding semantic manifold repeatedly reconstructs paths that intersect it.

8.6. Scope Density and Cognitive Compression

Containment structures differ not only in total semantic load but in the density with which unresolved dependencies are packed into local regions of the graph.

The *scope density* of a region $\Omega \subseteq B$ is

$$D(\Omega) = \frac{\sum_{B_i \in \Omega} U(B_i)}{|\Omega|}.$$

High-density regions contain many unresolved dependencies compressed into a small topological neighborhood. Such regions exhibit heightened cognitive pressure, rapid attentional cycling, and increased susceptibility to cascade behavior. This helps explain why certain symbolic systems feel cognitively compressed: a highly abstract proof, a densely metaphorical poem, or a multilayered hypnotic induction may carry extremely high local scope density within small surface structure. Cognitive fluency therefore depends not solely on total load $L(\Sigma)$ but on the geometric distribution of unresolved structure across the containment manifold.

9. Constraint Propagation and Semantic Causality

The framework implies a conception of causality different from linear push-pull models of symbolic processing. Within a containment structure, unresolved scopes constrain the admissible future evolution of the system long before any explicit collapse occurs. A bubble with high unresolved load restricts which future Pops, Melds, and Reframes remain admissible. The structure therefore evolves not merely through local symbol transitions but through globally distributed constraint propagation. Future semantic trajectories are shaped by the topology of currently unresolved dependencies.

This interpretation aligns naturally with dynamical approaches to cognition in which causation operates through evolving constraint landscapes rather than isolated efficient interactions [5]. The unresolved scope acts less like a stored object and more like a curvature condition imposed on future semantic motion. Meaningful cognition is therefore prospective as well as retrospective. Open scopes organize the future before they collapse into the past.

10. Toward a Vector-Valued Load Theory

The semantic load functional $L(\Sigma) = \sum_i w_i U(B_i)$ has served as a useful scalar summary of unresolved semantic tension throughout this paper. But the weight decomposition $w_i = f(d_i, s_i, r_i, c_i)$ already signals that semantic load is not a single quantity. Different dimensions of weight can dissociate: a mathematically unresolved problem may carry high structural load and low affective salience. A trauma trigger may carry enormous affective curvature despite shallow logical structure. A conspiracy system may maintain high affective coherence simultaneously with low causal coherence. The scalar $L(\Sigma)$ compresses these dissociations into a single value, which is adequate for the foundational treatment but insufficient for a full cognitive theory.

A natural extension replaces the scalar load with a vector-valued functional:

$$\mathbf{L} : \Sigma \to \mathbb{R}^n, \quad \mathbf{L}(\Sigma) = (\lambda_s L_s,\ \lambda_a L_a,\ \lambda_c L_c,\ \lambda_p L_p),$$

where L_s measures structural dependency load, L_a measures affective salience load, L_c measures causal centrality load, and L_p measures phenomenological tension, each weighted by context-dependent coupling constants $\lambda_s, \lambda_a, \lambda_c, \lambda_p$. Under this decomposition the containment manifold admits different curvature along different load dimensions. A system can be geometrically flat in the causal dimension while exhibiting steep affective gradients, producing the dissociation characteristic of intellectualized trauma or dispassionate mathematical obsession. The cascade threshold condition $L(\Sigma) > L_c$ generalizes naturally to a threshold surface in $\mathbb{R}^n$, allowing different cognitive systems to exhibit phase transitions along different projections of the load manifold. This extension remains a research program rather than a completed formalism, but the foundational definitions of this paper are consistent with it and in several places anticipate it.

10.1. Homotopy Classes of Derivation Histories

The collapse quotient $B/\sim_\rho$ introduced in Section 3 raises the question of how much of a derivation history survives compression. The weakest interpretation is purely extensional: two derivations are equivalent iff they collapse to the same final residue, so that $(2+3)$, $(1+4)$, and $(10/2)$ become indistinguishable after resolution. This is computationally natural but phenomenologically inaccurate. Different derivational paths leave behind structural traces that affect subsequent cognition. A theorem discovered geometrically behaves differently from the same theorem discovered algebraically, even after the formal proof is identical. A traumatic insight reached gradually behaves differently from one reached catastrophically, even if the explicit conclusion is verbally indistinguishable afterward.

The stronger interpretation treats the collapse residue as preserving a compressed homotopy class of the derivation history:

$$B/\sim_{\rho}^{\alpha}$$

where α is a resolution parameter controlling the coarseness of the equivalence. At coarse α, all derivational paths to a given result are identified. At finer α, derivations partition into semantic equivalence classes distinguished by their topological structure, emotional trajectory, modal character, or reframing depth. This stratification permits the framework to distinguish fact equivalence from experiential equivalence a distinction that appears essential for any adequate theory of identity, expertise, creativity, or style. The full development of homotopy-stratified collapse quotients is a natural direction for subsequent formal work.

11. Recursive Identity and Narrative Continuity

Human identity may be interpreted within this framework as a persistent, partially unresolved containment structure maintained across time through continuous cycles of Pop, Meld, and Reframe. A self is not a static object but an evolving topology of unresolved commitments, inherited narratives, anticipated futures, and recursively reorganized memories.

Under this interpretation, autobiographical continuity emerges from the preservation of admissible ancestry relations across repeated reframing operations. A person changes not by replacing one fixed identity with another but by continuously redistributing containment relations over a partially stable historical graph. Conservative reframes preserve identity through reinterpretation: $\prec^* = \prec'^*$. The same historical dependencies remain present while their local organization changes. Expansive reframes incorporate previously external domains into the identity graph: $\prec^* \subseteq \prec'^*$. New ancestral structures become integrated into the self-model.

This provides a natural account of both psychological continuity and personal transformation. Stability arises from preserved ancestry; growth arises from topological reorganization. Identity is therefore neither immutable essence nor arbitrary reconstruction. It is recursive containment maintained through lawful semantic evolution. Pathological rigidity corresponds to a failure to admit any admissible Reframe. Pathological dissolution corresponds to a failure to maintain any stable ancestry relation across time. Health, under this account, is the sustained capacity for both.

12. The Self as Deep Attractor Bubble

The conversation between Karl Friston and Susan Blackmore on predictive processing and the free energy principle [8] converges on several claims that are structurally precise translations of the containment framework developed in this paper. The convergence is not merely analogical. The two frameworks are approaching the same geometry from different entry points: predictive processing from the perspective of Bayesian brain theory and active inference, and the present framework from the perspective of recursive semantic containment and deferred closure.

Friston's central claim about the self is that it is not a controller sitting at the apex of a strict hierarchy but a deep, enduring hypothesis embedded within a heterarchical structure. In his account, the deeper a representation sits within the generative model, the more slowly it changes and the more temporally extended its influence. The self is the deepest such representation precisely because it must remain stable across the widest range of incoming evidence. It is not at the top because it commands; it is at the center because it persists.

This maps directly onto the containment framework. A bubble with very high temporal persistence and very low $U(B)$ turnover rate functions as an attractor in semantic space: it organizes surrounding scopes without being quickly resolved by them. The self-bubble B_{self} is therefore characterized not by high unresolved load in the ordinary sense but by extremely high semantic inertia. Its containment relations are shared by an enormous portion of the surrounding graph. Reframing it propagates destabilization throughout the entire structure, which is precisely why identity change is experienced as difficult, disorienting, and effortful.

The heterarchical structure Friston proposes self at the center rather than the top, with predictions propagating centrifugally outward and evidence flowing centripetally inward is structurally equivalent to replacing the rooted-tree model of self with a DAG model in which B_{self} is not the root but a highly connected interior node. Its influence is not top-down command but multi-parent containment inheritance: it shapes the admissibility conditions for a large portion of the surrounding semantic graph without directly controlling any single node.

Friston introduces the concept of *mortal computation* (a phrase due originally to Geoffrey Hinton) to characterize computation that is substrate-dependent in a constitutive rather than incidental way. On a von Neumann architecture, memory and processing are separated, allowing the software to be copied, distributed, and run on arbitrary hardware. The computation is therefore abstract: detached from any particular physical trajectory. On a mortal computational substrate neuromorphic, memristive, or photonic the computation is the physical process. It cannot be copied without being ended, because the substrate and the computation are the same thing.

This distinction maps onto the irreversibility framework of this paper. A pop operation in the Spherepop calculus is irreversible precisely because the derivation history is not separately stored. The collapse of $(2+3)$ into 5 is not the deletion of a record; it is the ending of a process. The history existed as a trajectory, not as a log. Under a mortal computational substrate, this is not a limitation but the constitutive structure of the process. The trajectory *is* the computation. There is no copy. This provides a formal account of why von Neumann architectures cannot, in Friston's sense, truly self-evidence: self-evidencing requires that the organism's current model state be shaped by its own prior physical history in an unbroken trajectory of constraint propagation. A system that can be copied, reset, or run on an arbitrary substrate has no such trajectory. Its semantic bubbles carry no mortal provenance.

13. Suffering, Precision, and Recursive Self-Containment

Friston's account of suffering is structurally precise: to suffer, an organism must include within its generative model a hypothesis that *I am a self who can suffer*. Suffering is not bare negative affect; it is negative affect recursively bound to a self-model with future-continuity constraints. Without that recursive binding, an organism may exhibit instability, error-signaling, or avoidance behavior without experiencing anything. A thermostat registers deviation from setpoint without suffering, because there is no persistent hypothesis that *I am the kind of thing that should be at setpoint and whose failure to be so matters*.

This can be formalized within the scope framework. Let Π_{self} denote the precision weighting in the active inference sense, the inverse variance assigned to the self-hypothesis bubble B_{self}. Let $U(B_{\text{future}})$ denote the unresolved semantic load of the future-directed scopes currently nested inside B_{self}. Then suffering in its structural sense can be approximated as:

$$\mathcal{S} \sim \Pi_{\text{self}} \cdot U(B_{\text{future}}).$$

High precision weighting on the self-hypothesis means that deviations from expected self-consistent outcomes register strongly. High unresolved forward load means that many anticipated scopes remain open. The product of these two quantities produces the characteristic phenomenology of suffering: not mere uncertainty, but self-indexed unresolvedness. Ordinary uncertainty does not hurt unless it is recursively bound to a persistent self-model that cannot tolerate the openness.

This formulation resolves a puzzle in pure predictive processing accounts. Expected surprise (entropy) is technically bad in the free energy framework, but organisms clearly tolerate uncertainty in many contexts without suffering. The resolution is that entropy without self-binding is merely load. It becomes suffering when Π_{self} is high enough to amplify $U(B_{\text{future}})$ into a cascade-sensitive regime. Below the threshold L_c, unresolved forward load is manageable. Above it, the cascade of unresolved self-indexed scopes produces the recursive amplification characteristic of anxiety, grief, and existential distress.

Grief, as Friston and Blackmore discuss it, is a specific instance of this structure. A major loss renders the self-model no longer fit for purpose: the forward-directed scopes that were organized around a particular world structure (domestic life, relational patterns, spatial familiarity) are suddenly inadmissible, because the world they anticipated no longer exists. The relearning required is not merely the accumulation of new facts but the replacement of large portions of the containment graph while B_{self} must somehow remain stable enough to sustain the process. It is exhausting because it is semantically expensive: an enormous number of previously closed scopes must be reopened, re-evaluated, and re-collapsed under the new constraints, all while the self-bubble that organizes them is itself under reconstruction.

14. Meditation, Top-Layer Relaxation, and the Dissolution of Recursive Agency

Blackmore describes the meditative process in terms that map precisely onto the containment framework: taking the top layers off, reducing the precision weighting on high-level hypotheses, and discovering what remains when the recursive self-model is no longer actively maintained. Friston interprets this as a reduction in the precision weighting of deep priors: $\Pi_d \to 0$ for the deepest layers of the generative hierarchy.

In scope-dynamic terms, meditative dissolution involves the progressive relaxation of semantic inertia at the highest levels of the containment structure. The self-bubble B_{self} does not pop in the ordinary sense that would be psychotic dissolution, not meditative insight but its precision weighting decreases, reducing its gravitational pull on surrounding scopes. The forward-directed load $U(B_{\text{future}})$ decreases not because the scopes are forcibly closed but because the self-hypothesis that amplified their urgency is no longer strongly weighted.

The result, as both Friston and Blackmore describe, is not emptiness but a shift in the geometry of attentional flow. When Π_{self} decreases, the semantic curvature generated by B_{self} flattens. The attentional gradients that ordinarily bend cognition toward self-relevant unresolved scopes relax. What remains is local scope navigation without the global self-indexed amplification. Perception continues; cognition continues; but the recursive binding that converts unresolved load into suffering is temporarily suspended.

Conversely, pathological states correspond to excessive Π_{self}: the self-bubble becomes so inertially dominant that its containment relations propagate into every region of the semantic manifold, rendering all unresolved scopes self-relevant and therefore unbearable. Anxiety disorders, certain psychotic presentations, and existential crises all share this structure: the cascade threshold is effectively lowered to zero because Π_{self} amplifies every open scope into a potential crisis.

The framework thus provides a unified account of ordinary selfhood, suffering, grief, meditative dissolution, and pathological self-amplification as different parameter regimes of the same underlying structure: recursive containment under varying precision weighting. This is not a replacement for predictive processing theory but a geometric restatement of its central claims in the vocabulary of scope dynamics. The two frameworks are complementary. Predictive processing specifies the computational mechanism Bayesian inference over hierarchical generative models. The scope-dynamic framework specifies the topological structure within which that inference operates nested unresolved regions under admissible collapse operators. Together they describe the same system from different angles of approach.

A. Formal Definitions Collected

For reference, the principal definitions of the paper are collected here in order of introduction.

A *semantic bubble* $B = (C, E, U)$ consists of a contextual binding set C, an expectation structure E, and a non-negative unresolved load scalar $U(B)$.

A *containment structure* $\Sigma = (B, \prec)$ is a finite set of bubbles equipped with a strict partial order. In the tree case every bubble has at most one immediate parent. In the DAG case this restriction is dropped.

The *scope stack* $\Sigma_t = [B_1, \dots, B_n]$ linearizes the active containment path with B_n the currently attended scope.

The *semantic load functional* is $L(\Sigma) = \sum_i w_i U(B_i)$ with $w_i = f(d_i, s_i, r_i, c_i)$.

The *resolution operator* $\rho : B \times \Sigma \rightharpoonup B'$ is partial, defined iff all descendants of B are closed.

Well-nestedness at position i requires $U(B_j) = 0$ for all $B_j \prec B_i$.

The five primitive operators are open, pop, meld_π, reframe^c_ϕ, reframe^e_ϕ, with admissibility conditions as stated in Section 3.

Conservative reframe satisfies $\prec^* = \prec'^*$. *Expansive reframe* satisfies $\prec^* \subseteq \prec'^*$. Both require $\prec'$ to be acyclic.

B. Worked Examples

Arithmetic

$(3 + (4 \times (2 + 1)))$ reduces as

$$(3 + (4 \times 3)) \;\to\; (3 + 12) \;\to\; 15.$$

Each arrow is an admissible Pop. The innermost scope $(2 + 1)$ pops first because its interior is empty; its parent (4×3) becomes eligible only after that pop; and so on outward.

Narrative Induction

Open B_1 (frame story), interrupt to open B_2 (embedded story), interrupt to open B_3 (innermost anecdote). Resolve: B_3 pops, modifying B_2'; B_2' pops, modifying B_1'; B_1' pops, emptying the stack. Total semantic load decreases monotonically through the resolution cascade.

Conservative Reframe: Therapeutic

A client holds a childhood experience B_{trauma} as the enclosing frame for their adult identity B_{self}. After therapeutic reframing, B_{trauma} becomes a contained episode within a broader developmental narrative B_{history}. Node set unchanged; ancestral graph unchanged; covering relations redistributed. The event remains real; its containment role changes.

Expansive Reframe: Metaphor

The domain of argument B_{arg} acquires an additional parent scope B_{war} through the metaphor *argument is war*. The transitive closure $\prec^*$ is extended. B_{arg} now inherits from both its original parent structures and the new war-domain parent. Gluing requires that the inherited properties from B_{war} be simultaneously compatible with B_{arg}'s existing parent constraints.

References

[1] Henk Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North-Holland, 1984.

[2] John Baez and Bob Coecke, editors. *Applied Category Theory 2019*. Electronic Notes in Theoretical Computer Science, 2020.

[3] Gilles Fauconnier and Mark Turner. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books, 2002.

[4] Douglas Hofstadter. *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books, 1979.

[5] Alicia Juarrero. *Dynamics in Action: Intentional Behavior as a Complex System*. MIT Press, 1999.

[6] George Lakoff and Mark Johnson. *Metaphors We Live By*. University of Chicago Press, 1980.

[7] Erik Meijer, Maarten Fokkinga, and Ross Paterson. Functional programming with bananas, lenses, envelopes and barbed wire. In *FPCA '91: Functional Programming Languages and Computer Architecture*, pages 124–144. Springer, 1991.

[8] Karl Friston and Susan Blackmore. Why do models suffer? Consciousness, self-evidencing, and mortal computation. Transcript of conversation, 2023.

[9] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 2nd edition, 1998.

[10] Tristan Needham. *Visual Complex Analysis*. Oxford University Press, 1997.