Semantic Stability as a Repair Equilibrium

Threshold, Attenuation, and Irreversible Collapse

in Constraint-Bounded Semantic Trees

Flyxion

Independent Researcher

2026

Abstract

A constraint-bounded semantic tree is a structure in which meaning propagates from root to leaf subject to an attenuation process that weakens semantic coherence with distance. We formalize such trees as repair equilibria in the sense of the accompanying repair ontology: each node remains semantically alive so long as the repair pressure reaching it — the accumulated constraint propagation from its ancestors — meets or exceeds a local admissibility threshold. The central result is the Semantic Stability Theorem, which gives a branch coherence condition $C(\mathcal{B}) \geq 1$ for the persistence of a connected semantic subtree. When this condition fails, the branch crosses the boundary of its admissibility manifold and enters irreversible attenuation. We prove that boundary nodes are the critical sites and that local boundary failure propagates inward through the tree in a manner analogous to crack propagation in materials under load. The RSVP field triple (Φ, v, S) provides a natural field-theoretic interpretation: semantic density, constraint propagation velocity, and accessible future volume map onto the three field variables, and irreversible attenuation corresponds to local entropy expansion outrunning the repair flow of the vector field. The threshold is therefore not an information threshold but a repair threshold. Semantic death is unrepaired divergence.

A word that is no longer repaired by use falls out of the language. Not because it has been forgotten, but because the repair infrastructure for it has collapsed.

Introduction

The companion essay establishes that objects and processes are emergent descriptions of repair equilibria: configurations in which the rate of repair meets or exceeds the rate of degradation for a sustained time horizon. That framework applies at every scale from molecular biology to cosmology. This paper applies it to a specific and formally tractable domain: the propagation of semantic constraint through a hierarchical tree structure.

The question motivating the paper is: when does a semantic structure remain coherent, and when does it collapse irreversibly? The question is not about information in the Shannon sense — bits stored, transmitted, received. It is about constraint: does the semantic content at the root of a tree still exert meaningful influence on interpretation at the leaves? And if that influence is weakening — through distance, through conceptual branching, through competing interpretive pressures — at what point does the weakening become irreversible?

The repair framework gives a precise answer. A semantic tree persists as a coherent structure so long as it remains a repair equilibrium: so long as the constraint propagation reaching each node is sufficient to maintain that node within the admissibility region of the tree's semantic manifold. When the constraint propagation falls below a local threshold, the node crosses the boundary of the admissibility region. Below the boundary, the system's own dynamics do not restore it to admissibility; they carry it further out. This is the irreversibility condition. And because semantic structures are hierarchically organized, boundary failure at a peripheral node propagates inward, undermining the constraint infrastructure for the nodes that depend on the now-failed peripheral region.

The paper is organized as follows. Section 2 defines the formal apparatus: semantic trees, admissibility weights, constraint propagation, and attenuation. Section 3 establishes the repair-equilibrium interpretation of each node and defines the key quantities $S(v)$, $\theta(v)$, and the repair ratio $\rho(v) = S(v)/\theta(v)$. Section 4 proves the Semantic Stability Theorem for individual nodes and extends it to branches via the branch coherence condition $C(\mathcal{B})$. Section 5 analyzes boundary dynamics and proves the propagation theorem: local boundary failure produces inward crack-like failure fronts. Section 6 develops the RSVP field interpretation of the entire framework. Section 7 discusses the linguistic and cognitive applications, with particular attention to language death and institutional memory collapse as cases of repair-threshold crossing. Section 8 concludes.

Formal Apparatus

Semantic Trees

Definition 1 (Semantic tree)**.** A *semantic tree* $\mathcal{T} = (V, E, \ell, \Phi_0)$ consists of a finite directed tree (V, E) with root $r \in V$, a labeling function $\ell : V \to \Sigma$ assigning semantic content to each node from a semantic domain Σ, and a root constraint $\Phi_0 \in \Phi$ specifying the initial semantic density at the root.

The edges in E are directed away from the root: each edge $(u, v) \in E$ denotes that u is the parent of v. The tree encodes how a root concept or claim ramifies into a hierarchy of derived concepts, sub-claims, dependent interpretations, or derived meanings. We write $\mathrm{Anc}(v)$ for the set of ancestors of v, including v itself, and $\mathrm{Desc}(v)$ for the set of descendants, including v.

Definition 2 (Admissibility weight)**.** An *admissibility weight* $w : V \to [0, 1]$ assigns to each node v a value $w(v)$ representing that node's capacity to transmit semantic constraint to its descendants. A node with $w(v) = 1$ transmits constraint without loss; a node with $w(v) = 0$ is a complete semantic barrier.

Admissibility weights are distinct from the semantic content $\ell(v)$. A node may have rich semantic content ($\ell(v)$ is complex) but low transmission capacity ($w(v)$ is small), meaning it is itself coherent but fails to constrain interpretation below it. This can occur when a concept is highly ambiguous, when it connects poorly to the surrounding semantic network, or when it has been separated from its constraining context.

Alpha Attenuation

Definition 3 (Attenuation function)**.** An *attenuation function* $\alpha : E \to (0, 1]$ assigns to each edge (u, v) a transmission coefficient $\alpha(u, v)$ representing the fraction of semantic constraint that survives the traversal of that edge. Attenuation is multiplicative along paths: for a path $r = v_0, v_1, \ldots, v_k = v$, the path attenuation is

$$\alpha(r \to v) = \prod_{i=0}^{k-1} \alpha(v_i, v_{i+1}).$$

Multiplicativity captures the essential feature of hierarchical semantic degradation: constraint weakens with each additional level of derivation, and the weakening is compounding. A chain of mostly-faithful transmissions $\alpha = 0.9$ over ten steps delivers only

$0.9^{10} \approx 0.35$ of the original constraint at the tenth level — just over a third of what the root specified. Over twenty steps it delivers $0.9^{20} \approx 0.12$.

This is not an artifact of the formalism. It reflects the genuine difficulty of maintaining interpretive fidelity across long chains of derivation, analogy, translation, or institutional transmission.

Constraint Propagation and Semantic Support

Definition 4 (Semantic support)**.** The *semantic support* $S(v)$ at node v is the total constraint influence reaching v from its ancestors, weighted by path attenuation and ancestral admissibility:

$$S(v) = \sum_{u \in \mathrm{Anc}(v)} \alpha(r \to u)\, w(u)\, K(u, v),$$

where $K(u, v) \geq 0$ is a kernel measuring the direct relevance of ancestor u's content to descendant v's admissibility, and $\alpha(r \to u)$ is the path attenuation from root to u.

The kernel $K(u, v)$ encodes the semantic architecture of the tree: how much the meaning at u constrains interpretation at v, independent of the transmission efficiency. In the simplest case, $K(u, v) = 1$ for direct ancestors and 0 otherwise, reducing semantic support to pure path attenuation. More realistic models allow K to decay smoothly with semantic distance, enabling off-path influence from conceptually close but structurally distant ancestors.

In the uniform case with $K(u, v) = 1$ for all $u \in \mathrm{Anc}(v)$ and constant α per edge, the semantic support becomes a sum of geometric terms:

$$S(v) = w(r) + \alpha\, w(v_1) + \alpha^2\, w(v_2) + \cdots + \alpha^{d(v)-1}\, w(v_{d(v)-1}),$$

where $d(v)$ is the depth of v and $v_0, v_1, \ldots, v_{d(v)-1}$ is the path from root to v's parent.

Nodes as Repair Equilibria

The Admissibility Threshold

Definition 5 (Admissibility threshold)**.** Each node $v \in V$ has an *admissibility threshold* $\theta(v) > 0$, the minimum semantic support required for v to maintain its semantic identity — to function as a coherent node that successfully constrains its own descendants and participates in the tree's semantic manifold.

The threshold $\theta(v)$ is not a property of v's content alone but of v's role in the tree. A node that is a conceptual hub — from which many other nodes derive their interpretation — requires higher semantic support to maintain because the consequences of its failure are larger. A peripheral leaf node requires lower support because its failure does not cascade.

Definition 6 (Repair ratio)**.** The *repair ratio* at node v is

$$\rho(v) = \frac{S(v)}{\theta(v)}.$$

A node is *semantically alive* if $\rho(v) \geq 1$ and *semantically failing* if $\rho(v) < 1$.

This is the direct semantic analogue of the rate-matching condition $r_R \geq r_D$ from the repair ontology. Semantic support $S(v)$ is the repair pressure reaching the node. The threshold $\theta(v)$ is the minimum repair rate required for the node to persist as a coherent element of the semantic manifold. The node is a repair equilibrium exactly when $\rho(v) \geq 1$.

Remark 1. The repair-equilibrium interpretation resolves an apparent circularity in semantic theory. Semantic content at v depends on the content of v's ancestors, but what counts as valid ancestral content depends on the tree's overall coherence. The repair-equilibrium framework dissolves this circle: the dependence is not circular but dynamic. Each node's semantic identity is continuously re-established by the repair pressure reaching it from above. The identity is not a fixed property but an ongoing achievement.

The Admissibility Manifold for Semantic Trees

In the language of the repair ontology, the admissibility manifold for a semantic tree is the set of all weight-threshold configurations (w, θ, α) under which every node is semantically alive:

$$\mathcal{A}(\mathcal{T}) = \{(w, \theta, \alpha) : \rho(v) \geq 1 \text{ for all } v \in V\}.$$

The boundary $\partial\mathcal{A}(\mathcal{T})$ is the set of configurations at which at least one node is exactly at its admissibility threshold. These are the critical configurations: any increase in attenuation, any decrease in admissibility weight, or any increase in a node's threshold will push the system outside $\mathcal{A}$.

The Semantic Stability Theorem

Node-Level Stability

Theorem 1 (Node stability)**.** A node v remains semantically alive under perturbation $\delta\alpha$ of the edge attenuation if and only if

$$S(v) - \theta(v) > \left|\frac{\partial S(v)}{\partial \alpha}\right| \cdot |\delta\alpha|.$$

The left side is the repair surplus at v; the right side is the rate at which perturbation depletes that surplus. Nodes with large surplus tolerate large attenuation perturbations; nodes near the threshold $\rho(v) \approx 1$ tolerate almost none.

Proof. Under perturbation $\delta\alpha$, the semantic support changes to $S(v) + (\partial S(v)/\partial\alpha)\,\delta\alpha$. The node remains semantically alive when $S(v) + (\partial S(v)/\partial\alpha)\,\delta\alpha \geq \theta(v)$, i.e., when $S(v) - \theta(v) \geq -(\partial S(v)/\partial\alpha)\,\delta\alpha$. Taking absolute values and worst-case sign gives the stated condition. □

Branch Coherence

Individual node stability is necessary but not sufficient for tree-level semantic persistence. A branch may contain nodes with individually positive repair ratios while the weakest node in the branch provides insufficient support to nodes that depend on it. We need a branch-level condition.

Definition 7 (Branch stability index and resilience)**.** For a connected subtree $\mathcal{B} \subseteq V$, define two quantities. The *branch stability index* is the minimum repair ratio across all nodes in the branch:

$$\rho_{\min}(\mathcal{B}) = \min_{v\in\mathcal{B}} \rho(v) = \min_{v\in\mathcal{B}} \frac{S(v)}{\theta(v)}.$$

The *branch resilience* is the mean repair ratio:

$$C(\mathcal{B}) = \frac{1}{|\mathcal{B}|} \sum_{v\in\mathcal{B}} \rho(v).$$

These are distinct quantities with distinct roles. The stability index determines survival; the resilience measures capacity to absorb further perturbation.

Remark 2. The distinction matters. Consider a branch with repair ratios $(2.0, 2.0, 0.1)$. Its resilience is $C(\mathcal{B}) = 1.37$, suggesting apparent health, but its stability index is $\rho_{\min} = 0.1$, meaning one node is already dead and the branch is not a repair equilibrium. Averages mask the most dangerous structure in any hierarchical system: the weakest link. Since semantic trees transmit constraint hierarchically, a failed intermediate node severs support to all its descendants regardless of the health of sibling nodes.

Theorem 2 (Semantic Stability Theorem)**.** A connected semantic subtree $\mathcal{B}$ is a repair equilibrium if and only if $\rho_{\min}(\mathcal{B}) \geq 1$, i.e., every node in the branch satisfies $\rho(v) \geq 1$. The branch resilience $C(\mathcal{B})$ measures the surplus available to absorb future attenuation increases: a branch with $C(\mathcal{B}) \gg 1$ can tolerate large perturbations before any node approaches its admissibility threshold.

Proof. ($\Rightarrow$) If $\mathcal{B}$ is a repair equilibrium then by definition $\rho(v) \geq 1$ for all $v \in \mathcal{B}$, so $\rho_{\min}(\mathcal{B}) = \min_v \rho(v) \geq 1$.

($\Leftarrow$) If $\rho_{\min}(\mathcal{B}) \geq 1$ then $\rho(v) \geq 1$ for all $v \in \mathcal{B}$ by definition of the minimum, so every node is semantically alive and $\mathcal{B}$ is a repair equilibrium.

($\neg$: instability) If $\rho_{\min}(\mathcal{B}) < 1$, there exists $v^* \in \mathcal{B}$ with $S(v^*) < \theta(v^*)$. Since $\sum_v S(v) < \sum_v \theta(v)$ whenever any term satisfies $S(v) < \theta(v)$ (by direct arithmetic; no continuity argument is needed), the total repair pressure is insufficient. Furthermore, v^*'s failure reduces its effective admissibility weight, decreasing $S(u)$ for all $u \in \text{Desc}(v^*)$. This may push additional nodes below threshold. The branch cannot self-repair without external support. □

Remark 3. The Semantic Stability Theorem is the exact discrete analogue of the rate-matching condition $r_R \geq r_D$ from the repair ontology. The stability index $\rho_{\min}$ plays the role of the ratio r_R/r_D evaluated at the worst-performing site. The resilience $C(\mathcal{B})$ plays the role of mean repair surplus, which governs how much additional degradation the system can absorb before a node crosses the boundary.

The Critical Attenuation Level

Definition 8 (Critical attenuation)**.** For a semantic tree $\mathcal{T}$, the *critical attenuation level* α_c is the minimum uniform edge attenuation coefficient for which every node remains semantically alive, i.e., for which $\rho_{\min}(V) \geq 1$:

$$\alpha_c = \inf\{\alpha^* : \rho_{\min}(V) \geq 1 \text{ under uniform } \alpha = \alpha^*\}.$$

Proposition 1 (Attenuation threshold)**.** For a tree of depth D with uniform admissibility weights w and uniform thresholds θ, the critical attenuation level satisfies

$$\alpha_c = \left(\frac{\theta}{w \cdot H_D}\right)^{1/D},$$

where $H_D = \sum_{k=0}^{D-1} \alpha^k$ is the geometric sum evaluated at the boundary. Below α_c, semantic support at the deepest nodes falls below threshold and the tree begins collapsing from the leaves inward.

The key feature of this threshold is its non-linearity in depth. For deep trees, small reductions in α below α_c produce large reductions in $S(v)$ at the leaves, because attenuation compounds multiplicatively. This explains why deep semantic structures — long chains of derivation, complex institutional memories, multi-generational linguistic traditions — are disproportionately vulnerable to small increases in transmission loss.

Boundary Dynamics and Failure Propagation

Repair Intensity at Semantic Boundaries

From the repair ontology, repair intensity is concentrated near the boundary $\partial\mathcal{A}$:

$$\iota(v) \propto \frac{1}{d(v,\, \partial\mathcal{A}(\mathcal{T}))} = \frac{1}{\rho(v) - 1} \quad \text{for } \rho(v) > 1.$$

Nodes with $\rho(v)$ close to 1 are near the semantic boundary. They require continuous and intense repair pressure to remain semantically alive. Nodes with $\rho(v) \gg 1$ are deep inside the admissibility region and require relatively little active maintenance.

Theorem 3 (Critical Repair Divergence)**.** As a node approaches the admissibility boundary from above,

$$\lim_{\rho(v)\to 1^+} \iota(v) = +\infty.$$

That is, the maintenance cost required to keep a semantically marginal node alive diverges as its repair ratio approaches the critical threshold. No finite external intervention can sustain a node at exactly $\rho(v) = 1$ against any additional perturbation.

Proof. From the definition $\iota(v) \propto 1/(\rho(v) - 1)$ for $\rho(v) > 1$, the limit $\rho(v) \to 1^+$ sends $\rho(v) - 1 \to 0^+$, and therefore $\iota(v) \to +\infty$. For the second claim: maintaining $\rho(v) = 1$ against a perturbation $\delta > 0$ to the degradation rate requires increasing $S(v)$ by $\delta \cdot \theta(v)$, which demands a repair flow increase proportional to $1/(\rho(v) - 1) \cdot \delta$. As $\rho(v) \to 1^+$ this required flow diverges for any fixed $\delta > 0$. □

The Critical Repair Divergence theorem gives the formalism its most directly observable consequence. It explains why late-stage interventions in failing systems are disproportionately expensive: an endangered language requiring its last hundred speakers to sustain it demands more repair work per node than a thriving language with millions; a dying institution requires increasingly heroic maintenance efforts from a shrinking core; a marginal scientific paradigm requires more and more defensive effort to maintain coherence against accumulating anomalies. The cost does not increase linearly as the boundary approaches. It diverges. This is the formal mechanism behind the empirical observation that preventive maintenance is always cheaper than corrective maintenance: the corrective regime operates near the boundary, where ι is large; the preventive regime operates in the interior, where ι is small.

The divergence also implies a critical choice structure near the boundary. A system cannot remain indefinitely near $\rho(v) = 1$; the maintenance cost required to sustain it there is unbounded. Any real system with finite repair capacity must either recover to $\rho(v) \gg 1$ (by external input of support) or cross the boundary to $\rho(v) < 1$ (entering irreversible attenuation). The boundary is not a stable equilibrium; it is a saddle point between recovery and collapse.

In practice, the boundary-concentrated nodes are the conceptually peripheral nodes of the tree: those at high depth, those with many competing interpretive pressures, those that depend on long chains of derivation from the root, and those whose admissibility weights are locally low due to contextual disconnection. These are precisely the nodes most vulnerable to attenuation increase and most likely to fail first.

Failure Front Propagation

The following theorem establishes that boundary failure does not remain localized. It propagates inward.

Theorem 4 (Semantic crack propagation)**.** Let $v^* \in \partial\mathcal{A}(\mathcal{T})$ be a boundary node with $\rho(v^*) = 1$ that undergoes a perturbation pushing it below threshold ($\rho(v^*) < 1$). Then for every descendant $u \in \mathrm{Desc}(v^*)$, the semantic support satisfies

$$S(u) \leq S_{\text{pre}}(u) - \Delta(v^*, u),$$

where $\Delta(v^*, u) > 0$ is the loss of ancestral support due to v^*'s failure. In particular, there exists a critical depth d_c below v^* such that all nodes at depth $\geq d_c$ below v^* will fall below their thresholds following v^*'s failure, regardless of their pre-failure repair ratios.

Proof. When v^* crosses below threshold, its admissibility weight $w(v^*)$ effectively drops, reducing the semantic support it contributes to all descendants. For $u \in \mathrm{Desc}(v^*)$, the contribution of v^* to $S(u)$ is $\alpha(r \to v^*)\, w(v^*)\, K(v^*, u)$. The failure of v^* reduces $w(v^*)$ by some $\delta w > 0$, reducing $S(u)$ by $\Delta(v^*, u) = \alpha(r \to v^*)\, \delta w\, K(v^*, u) > 0$. Nodes sufficiently deep below v^* have their repair surplus $S(u) - \theta(u)$ already depleted by cumulative attenuation. The reduction $\Delta(v^*, u)$ is sufficient to push these nodes below threshold. The critical depth d_c is where $\Delta(v^*, u) \geq S_{\text{pre}}(u) - \theta(u)$ for all nodes at that depth. □

The analogy to materials failure is precise. In fracture mechanics, a surface crack propagates inward when the stress intensity factor at the crack tip exceeds the material's fracture toughness. Here, semantic failure propagates inward from the boundary when the reduction in support at a failed node exceeds the repair surplus of the nodes it was supporting. The structure appears stable until the boundary defect exceeds a critical magnitude; then the failure front advances into the interior.

Corollary 1 (Extinction as repair failure)**.** The total semantic collapse of a tree $\mathcal{T}$ — the state in which no node remains semantically alive — is the endpoint of failure front propagation from all boundary nodes simultaneously falling below threshold. This is not information loss: the semantic content $\ell(v)$ may remain stored somewhere. It is repair failure: the constraint propagation network that maintained the admissibility of each node has disintegrated faster than any sub-network could compensate.

This corollary unifies several apparently distinct phenomena under a single formal description. Language death is total semantic collapse of a linguistic tree. Species extinction is total semantic collapse of an ecological constraint network. Institutional collapse is total semantic collapse of a normative and procedural tree. Memory loss is total semantic collapse of a reconstructive constraint network. In each case the mechanism is the same: boundary failure propagating inward faster than internal repair can compensate.

The RSVP Field Interpretation

The RSVP field triple (Φ, v, S) provides a field-theoretic setting for the semantic stability framework. The correspondence is not metaphorical. We establish explicit identifications that make the semantic tree formalism a special case of RSVP field dynamics on a discrete graph substrate.

Explicit Field Identifications

Definition 9 (RSVP-semantic correspondence)**.** Given a semantic tree $\mathcal{T} = (V, E, \ell, \Phi_0)$ with admissibility weights w, attenuation α, and repair ratios ρ, the RSVP fields are identified as follows:

$$\Phi(v) = S(v) \tag{1}$$

$$|v_{ij}| = \alpha(i,j) \quad \text{for each edge } (i,j) \in E \tag{2}$$

$$S_{\text{RSVP}}(v) = -\log \rho(v) \tag{3}$$

where $\Phi(v)$ is the scalar field at node v, $|v_{ij}|$ is the magnitude of the vector field on edge (i, j), and $S_{\text{RSVP}}(v)$ is the entropy field at node v.

These identifications are not free choices. Each is forced by the requirement that the RSVP stability condition coincide with the Semantic Stability Theorem.

Equation (1): The scalar field Φ encodes semantic density in the RSVP framework. Identifying $\Phi(v) = S(v)$ means that semantic support — the accumulated constraint propagation reaching a node — is exactly the local density of the scalar field. High Φ at a node means high constraint support; the root, which has maximal constraint specification, sits at the global maximum of Φ. Peripheral nodes at high depth, where attenuation has compounded, have low Φ, consistent with their low semantic support.

Equation (2): The vector field v encodes constraint propagation velocity in the RSVP framework. Identifying its magnitude with the edge attenuation coefficient makes the vector field a direct measure of transmission fidelity. An edge with $\alpha(i, j) = 1$ is a perfect constraint channel; $|v_{ij}| = 1$ means full propagation speed with no attenuation loss. An edge with $\alpha(i, j) \ll 1$ is a near-barrier; $|v_{ij}| \ll 1$ means the constraint flow is nearly stopped at that edge. Failure of the vector field at an edge is directly attenuation failure.

Equation (3): The identification $S_{\text{RSVP}}(v) = -\log \rho(v)$ is the most consequential. It has several immediately verifiable properties. When $\rho(v) = 1$, the node is exactly at the admissibility boundary and $S_{\text{RSVP}}(v) = 0$: zero entropy means the interpretation volume is exactly at the threshold of admissibility. When $\rho(v) > 1$, the node is inside the admissibility manifold and $S_{\text{RSVP}}(v) < 0$: negative entropy corresponds to the constraint surplus, the degree to which the node's interpretation is over-determined relative to its threshold. When $\rho(v) < 1$, the node is outside the admissibility manifold and $S_{\text{RSVP}}(v) > 0$: positive entropy means the accessible interpretation volume has expanded beyond what the constraint infrastructure can bound. Failure corresponds exactly to positive entropy.

Proposition 2 (RSVP stability condition)**.** Under the identifications above, the RSVP stability condition

$$S_{\text{RSVP}}(v) \leq 0 \quad \text{for all } v \in V$$

is equivalent to the Semantic Stability Theorem's condition $\rho_{\min}(V) \geq 1$.

Proof. $S_{\text{RSVP}}(v) = -\log \rho(v) \leq 0$ iff $-\log \rho(v) \leq 0$ iff $\log \rho(v) \geq 0$ iff $\rho(v) \geq 1$. Taking the condition over all $v \in V$ gives $\rho_{\min}(V) \geq 1$. □

Corollary 2 (Critical Repair Divergence in RSVP)**.** The Critical Repair Divergence theorem takes the following form in the RSVP field language. As $S_{\text{RSVP}}(v) \to 0^-$ (the node approaches the admissibility boundary from inside),

$$\iota(v) = \frac{1}{e^{-S_{\text{RSVP}}(v)} - 1} \to +\infty.$$

The repair intensity diverges as the entropy field approaches zero from below.

Alpha Attenuation as Boundary Crossing

The identification $|v_{ij}| = \alpha(i, j)$ makes the mechanism of boundary crossing precise. Increasing attenuation on edge (i, j) decreases $|v_{ij}|$, reducing the constraint flow to j and all descendants of j. By equation (1), this decreases Φ at those nodes. By equation (3), this increases S_{RSVP}. When $S_{\text{RSVP}}(v)$ crosses zero, the node exits the admissibility manifold:

$$\alpha(i, j) \text{ falls below critical } \Rightarrow S_{\text{RSVP}}(j) > 0 \Rightarrow \rho(j) < 1 \Rightarrow j \notin \mathcal{A}(\mathcal{T}).$$

Alpha attenuation is therefore not information loss in any Shannon sense. Shannon entropy is a property of source distributions; it is invariant to what the signal means. What attenuation destroys here is S_{RSVP} — the constraint-entropy field that measures how tightly the admissibility manifold bounds interpretation. The semantic content $\ell(v)$ may survive perfectly while $S_{\text{RSVP}}(v)$ crosses zero: the node is readable but no longer repaired. This is the formal version of the stored-structure versus living-equilibrium distinction. A preserved node has $\ell(v)$ intact but $S_{\text{RSVP}}(v) > 0$. A living node has $S_{\text{RSVP}}(v) \leq 0$.

Lamphrodyne Relaxation as Semantic Repair

The lamphrodyne relaxation dynamics of RSVP smooth extreme local gradients in Φ through the coupled dynamics of v and S_{RSVP}. Under the semantic identification, lamphrodyne relaxation corresponds to the repair of boundary nodes: when a peripheral node approaches $S_{\text{RSVP}}(v) \to 0^-$, the relaxation mobilizes constraint flow (increases $|v_{ij}|$ on incoming edges), which increases $\Phi(v) = S(v)$ and decreases $S_{\text{RSVP}}(v)$, pushing the node back toward the interior.

When the lamphrodyne relaxation rate is insufficient to prevent S_{RSVP} from crossing zero, the node undergoes irreversible attenuation. The Spherepop COLLAPSE outcome formalizes this at the computational level: once a node has committed to a reduced semantic state, its bubble closes and S_{RSVP} for the collapsed region is reset to the new, lower-constraint manifold. The repaired tree is not the original tree; it is a new repair equilibrium with a different admissibility manifold, excluding the collapsed region.

Linguistic and Cognitive Applications

Language Death as Threshold Crossing

The repair-equilibrium framework gives a formal account of language death that distinguishes it clearly from language change.

Language change occurs when a semantic tree undergoes repair that reaches the admissibility manifold at a different point: the repaired language satisfies the same constraints as the pre-change language but with different realizations. Grammar changes, sound changes, and lexical innovations are all instances of repair reaching new points on the manifold while maintaining the tree's overall coherence. The branch coherence $C(V) \geq 1$ is preserved.

Language death occurs when $C(V)$ falls below 1. This happens when the transmission infrastructure — the intergenerational and community-level repair processes that sustain the language's semantic manifold — is disrupted faster than internal repair can compensate. The trigger is typically the reduction of the speech community below a critical size, the interruption of intergenerational transmission, or the rapid replacement of the language's contextual infrastructure by a dominant competing language.

In the boundary-dynamics framework, language death proceeds by failure front propagation. The peripheral nodes of the semantic tree — rare vocabulary, specialized registers, culturally specific pragmatic conventions, low-frequency syntactic constructions — are the first to lose their repair support. As the speech community contracts or shifts, these nodes cross the admissibility boundary first. Their failure reduces the support available to the intermediate nodes that depended on them for contrast and contextual specificity. The failure front advances inward toward the core vocabulary and basic grammar. At the point of total collapse, even the core nodes — the high-frequency, low-attenuation, maximally supported nodes — fall below threshold because the entire constraint propagation network has disintegrated.

A critical implication: language death is not reversible by reintroducing the semantic content. A dictionary of a dead language contains the nodes $\ell(v)$ but not the constraint infrastructure $S(v) \geq \theta(v)$ that made each node semantically alive. Revitalization requires reconstructing the repair infrastructure — the speech community, the use contexts, the intergenerational transmission mechanisms — not merely the content.

Institutional Memory Collapse

Institutional memory fails through the same mechanism, formalized identically.

An institution's normative and procedural knowledge is a semantic tree: the foundational principles of the institution are the root; the specific procedures, exceptions, informal norms, and contextual knowledge are the peripheral nodes. The admissibility weight of each node is the degree to which that piece of knowledge is actively practiced, transmitted to new members, and reinforced by the institution's ongoing activity. The threshold $\theta(v)$ of each node is the minimum active-practice rate for that knowledge to remain functionally accessible.

Personnel attrition is attenuation: as individuals who possess peripheral procedural knowledge leave, the transmission coefficient α for that knowledge falls. Initially this affects only the peripheral nodes. But if the attrition rate exceeds the institution's onboarding and knowledge-transfer capacity, the repair rate falls below the degradation rate and the branch coherence of the institutional knowledge tree begins to fall. The failure front propagates inward from specialized tacit knowledge toward the core operating procedures.

The formal threshold is the same: $C(\mathcal{B}) \geq 1$ for knowledge survival, $C(\mathcal{B}) < 1$ for irreversible loss. An institution that re-hires former personnel with the intention of recovering lost knowledge discovers that the nodes $\ell(v)$ may be recoverable — the individual remembers — but the constraint infrastructure that made the knowledge institutionally alive has collapsed. The individual's knowledge is no longer repaired by the institution's ongoing practice; it is an isolated node disconnected from the manifold.

Cognitive Depth and the Gesture-Before-Symbol Relationship

The gesture-before-symbol perspective developed in prior work holds that meaning is prior to its linguistic encoding. In the repair-equilibrium framework, this becomes a claim about the tree structure: gesture and embodied context are the root nodes of the semantic tree, and linguistic symbols are intermediate and peripheral nodes that depend on the root's constraint propagation for their admissibility.

This ordering has a precise consequence: the peripheral linguistic nodes can fail without causing root-level semantic failure, but root-level failure causes total collapse. A community that loses a specific word (peripheral node) retains the meaning (root constraint) and can repair the loss through alternative expression. A community that loses the embodied, contextual, gestural grounding of a meaning domain loses the constraint infrastructure for the entire branch of symbols that depended on it. The symbols may survive in dictionaries and grammars while being semantically dead in the repair-equilibrium sense: no longer repaired by the constraint propagation from their gestural and contextual roots.

This also explains the particular vulnerability of abstract terminology. Abstract terms are high-depth nodes in the semantic tree: they depend on long chains of derivation from more concrete roots, and each link in the chain contributes multiplicative attenuation. The critical attenuation level α_c for a deep abstract term is much higher than for a concrete high-frequency term, meaning that abstract terminology requires substantially more active maintenance to remain semantically alive. This is why technical vocabulary deteriorates rapidly when a community stops actively practicing the domain it describes.

Conclusion

The Semantic Stability Theorem establishes that a semantic tree is a repair equilibrium in precisely the formal sense developed in the companion paper on repair ontology. Each node is semantically alive when its repair ratio $\rho(v) = S(v)/\theta(v) \geq 1$. The branch stability index $\rho_{\text{min}}(\mathcal{B})$ is the correct survival condition; the resilience $C(\mathcal{B})$ measures the system's remaining tolerance for further degradation. The Critical Repair Divergence theorem establishes that maintenance cost diverges as any node approaches the admissibility boundary, explaining the empirical universality of the observation that late-stage interventions are disproportionately expensive. And the explicit RSVP identifications — $\Phi = S$, $|v_{ij}| = \alpha(i,j)$, $S_{\text{RSVP}} = -\log \rho$ — ground the tree formalism in the field dynamics of the RSVP framework, making the semantic stability condition the exact discrete analogue of the RSVP field-stability condition $S_{\text{RSVP}} \leq 0$.

The deepest consequence of the framework is a formal distinction between two modes of persistence that ordinary language conflates.

Definition 10 (Stored structure versus living equilibrium)**.** A *stored structure* is a configuration in which the semantic content $\ell(v)$ is preserved but $S_{\text{RSVP}}(v) > 0$ for some or all nodes: the constraint infrastructure has collapsed, and the content exists without the repair dynamics that once kept it semantically alive. A *living equilibrium* is a configuration in which $\rho_{\text{min}}(V) \geq 1$: every node is actively maintained within the admissibility manifold by ongoing constraint propagation.

These are not different degrees of the same thing. They are categorically different ontological states. A dead language dictionary is a stored structure: the nodes $\ell(v)$ are intact, but $S_{\text{RSVP}}(v) > 0$ everywhere — no repair dynamics sustain the admissibility of any entry. An abandoned institution's policy manual is a stored structure. A fossil is a stored structure. An archived codebase whose community of practice has dissolved is a stored structure. A trained language model whose training culture no longer exists is a stored structure: the interpolated topology of the residue, learned from an immense graveyard of formerly living equilibria, performs a kind of reconstruction over preserved content rather than maintaining live constraints.

In each case the content survives the death of the repair equilibrium. Reconstruction from stored structure is possible but is not the same operation as maintaining a living equilibrium. Revitalization of a dead language requires reconstructing the repair infrastructure — the speech community, the use contexts, the intergenerational transmission — not merely the content. Reconstruction from the content alone produces an artifact, not a repair equilibrium: a learned performance of a dead system, not the system itself.

Semantic stability is not a property of content. It is a property of maintenance. What exists semantically are living repair equilibria. Everything else is stored structure on a trajectory toward dissolution — the accumulation of successful repair while the repair lasted, and the residue of failed repair thereafter.

Flyxion / Independent Researcher. 2026.