The Semantic Fidelity Canonical Lexicon

Meaning Preservation Across Recursive AI and Representational Systems

A. Jacobs - Semantic Fidelity Lab · Version 2.0

Abstract

Modern AI systems do not simply generate errors. They transform meaning through retrieval, compression, summarization, memory, tool use, and recursive generation. As these systems become more persistent and agentic, the central challenge is not only whether outputs are factual, but whether meaning remains stable as it moves across layers.

The Semantic Fidelity Canonical Lexicon defines a vocabulary for diagnosing semantic degradation in recursive AI systems. It names the conditions under which systems remain coherent, operational, and locally useful while gradually weakening their connection to originating context, intent, or reality.

Rather than treating drift as isolated hallucination, this lexicon frames semantic degradation as a structural property of systems operating under abstraction, optimization, scale, and recursion. Its purpose is to support semantic observability, fidelity evaluation, provenance tracking, retrieval alignment, and long-term meaning preservation across AI-native environments.

Introduction

Large language models have turned language into an operational medium. Text is no longer only written, read, and stored. It is retrieved, embedded, ranked, compressed, summarized, regenerated, and routed through workflows that increasingly shape how information is interpreted.

This creates a new problem for AI systems. Meaning can degrade without obvious failure. A retrieval pipeline can surface relevant context while missing the deeper relationship. A memory system can preserve a user preference while losing the conditions that made it true. An agent can complete a task while drifting from the original intent. The system keeps working, but semantic alignment weakens.

The Semantic Fidelity Canonical Lexicon provides a conceptual framework for describing these conditions. It organizes semantic degradation into four interacting layers: core semantic conditions, AI system failure modes, semantic infrastructure and governance, and recursive semantic ecosystems.

Together, these concepts form a diagnostic and infrastructural framework for preserving meaning within increasingly recursive AI-mediated environments.

SECTION I — Core Semantic Conditions

Semantic Fidelity

Semantic fidelity is the degree to which meaning, intent, context, nuance, and interpretive structure are preserved across retrieval, generation, transformation, compression, summarization, and reuse.

High semantic fidelity systems preserve meaningful contact between representations and the realities they are intended to encode. Low semantic fidelity systems preserve structural coherence while allowing meaning to drift progressively across recursive transformations.

Semantic fidelity extends beyond factual correctness. A representation may remain technically accurate while failing to preserve tone, hierarchy, implication, metaphor, framing, emotional structure, or contextual meaning.

Structural Pattern: Representations remain coherent while deeper semantic relationships weaken.

Why It Matters: Modern AI systems increasingly operate through recursive transformations rather than direct reference to source reality. Semantic fidelity becomes the central condition for maintaining interpretive continuity across recursive systems.

Related Terms: Semantic Drift, Fidelity Decay, Ground Erosion, Semantic Continuity

Reality Drift

Reality drift describes the broader condition in which operational systems remain coherent, functional, and internally stable while progressively losing alignment with the underlying realities they were built to represent.

Reality drift occurs when representations continue circulating successfully despite weakening corrective contact with external conditions, originating intent, or material constraints. The defining characteristic of reality drift is not visible collapse, but survivable degradation.

Structural Pattern: Operational continuity persists while alignment weakens.

Why It Matters: Many modern systems fail gradually rather than catastrophically. The appearance of continuity masks underlying representational decay.

Related Terms: Survivable Wrongness, Proxy Alignment, Evaluation Blindness

Representation Drift

Representation drift occurs when representations continue circulating, functioning, and influencing downstream systems while progressively diverging from the realities, contexts, or meanings they originally encoded.

Representation drift often emerges gradually through recursive compression, retrieval abstraction, optimization pressure, and synthetic regeneration.

Structural Pattern: Representations remain operational while losing reference integrity.

Why It Matters: Modern AI systems increasingly operate on representations of prior representations rather than direct contact with source reality.

Related Terms: Reality Drift, Recursive Mediation, Proxy Alignment, Semantic Drift

Recursive Mediation

Recursive mediation describes a condition in which reality is increasingly encountered through layered representations of prior representations. Meaning passes recursively through summaries, feeds, retrieval systems, interfaces, synthetic outputs, narratives, embeddings, and generative transformations before reaching interpretation. Each layer introduces abstraction, optimization pressure, compression, and drift.

Structural Pattern: Representations become inputs for future representations.

Why It Matters: Recursive mediation compounds semantic distortion over time while obscuring originating context and provenance.

Related Terms: Recursive Compression, Semantic Exhaust, Recursive Synthetic Feedback

Semantic Drift

Semantic drift is the gradual mutation of meaning across recursive transformations. Drift often preserves surface coherence and factual structure while altering tone, framing, implication, intent, emotional resonance, or interpretive hierarchy.

Semantic drift becomes increasingly difficult to detect once recursive systems begin training, retrieving, or reasoning from prior transformed representations.

Structural Pattern: Meaning mutates while structural continuity remains intact.

Why It Matters: Drift compounds recursively across summarization systems, retrieval architectures, memory systems, and generative environments.

Related Terms: Fidelity Decay, Meaning Debt, Context Drift

Fidelity Decay

Fidelity decay describes the predictable weakening of semantic integrity across repeated compressions, summarizations, reinterpretations, or recursive generations. Each recursive transformation removes subtle contextual structure, increasing abstraction while reducing grounding and nuance.

Structural Pattern: Compression compounds semantic loss over time.

Why It Matters: Recursive AI systems increasingly depend on summarization, memory compression, and retrieval abstraction at scale.

Related Terms: Recursive Compression, Semantic Drift, Meaning Collapse

Semantic Entropy

Semantic entropy describes the increasing instability, disorder, and unpredictability of meaning across recursive representational systems. As representations are repeatedly compressed, transformed, retrieved, paraphrased, and regenerated, semantic structure becomes progressively less stable and more difficult to reliably reconstruct.

Unlike semantic drift, which describes directional mutation of meaning, semantic entropy describes the broader loss of semantic order and interpretive stability across recursive environments.

Structural Pattern: Recursive transformations increase semantic disorder over time.

Why It Matters: Large-scale AI systems increasingly operate through recursive semantic environments in which degraded representations become future inputs, compounding instability across retrieval, memory, and generation systems.

Related Terms: Semantic Drift, Fidelity Decay, Recursive Compression, Recursive Synthetic Feedback

Ground Erosion

Ground erosion describes the collapse of implicit context, hierarchy, absence, silence, and interpretive background that gives language semantic weight. Meaning depends not only on explicit signals, but on what remains unstated, differentiated, or contextually structured. When recursive systems flatten contextual distinctions, significance collapses into uniform semantic equivalence.

Structural Pattern: Background structure disappears while foreground representations remain intact.

Why It Matters: Generative systems often preserve explicit content while eroding interpretive depth.

Related Terms: Meaning Thinning, Semantic Fatigue, Synthetic Coherence

Meaning Debt

Meaning debt describes the accumulated loss of nuance, contextual richness, resonance, and interpretive precision across recursive shortcuts and compressions. Each omission may appear locally insignificant while contributing to long-term semantic degradation.

Structural Pattern: Small semantic losses compound recursively.

Why It Matters: Recursive semantic systems accumulate degradation gradually rather than immediately.

Related Terms: Fidelity Decay, Semantic Drift, Meaning Collapse

Recursive Compression

Recursive compression describes the repeated reduction of semantic complexity across recursive systems. Compressed representations become future inputs for additional compression cycles, increasing abstraction while weakening grounding.

Structural Pattern: Compression recursively compounds abstraction.

Why It Matters: Modern AI ecosystems increasingly operate through recursive compression layers: embeddings, summaries, rankings, memory abstractions, and synthetic retrieval environments.

Related Terms: Semantic Drift, Fidelity Decay, Semantic Entropy

SECTION II — AI System Failure Modes

Agent Drift

Agent drift describes the gradual loss of alignment between an AI agent's originating objective and its later outputs across multi-step workflows. Small interpretation shifts accumulate across retrieval, planning, memory updates, tool use, orchestration, and generation cycles.

Structural Pattern: Multi-step workflows progressively diverge from originating intent.

Why It Matters: Persistent agent systems amplify recursive semantic degradation across long operational chains.

Related Terms: Context Drift, Memory Drift, Proxy Alignment

Context Drift

Context drift occurs when important contextual constraints weaken or disappear across long workflows, memory chains, or iterative prompting systems. Later outputs inherit representations detached from originating conditions.

Structural Pattern: Inherited context becomes progressively incomplete or distorted.

Why It Matters: Long-context systems increasingly depend on compressed contextual reconstruction rather than persistent grounding.

Related Terms: Memory Drift, Retrieval Drift, Semantic Continuity

Retrieval Drift

Retrieval drift occurs when retrieved context appears semantically adjacent while remaining structurally misaligned with originating intent or meaning. Embedding similarity creates the appearance of relevance without preserving interpretive equivalence.

Structural Pattern: Semantic proximity substitutes for structural alignment.

Why It Matters: Modern RAG systems increasingly depend on latent similarity rather than interpretive understanding.

Related Terms: Embedding Distortion, Interpretation Failure, Retrieval Failure

Embedding Distortion

Embedding distortion occurs when latent-space proximity substitutes for actual semantic equivalence. Concepts become artificially associated because vector relationships compress distinct meanings into statistically adjacent representations.

Structural Pattern: Latent similarity masks conceptual difference.

Why It Matters: Embedding systems increasingly shape retrieval, recommendation, ranking, and semantic search architectures.

Related Terms: Retrieval Drift, Semantic Search, Latent Relational Mapping

Interpretation Failure

Interpretation failure occurs when systems retrieve relevant information but integrate, frame, or reason about that information incorrectly. Retrieval success does not guarantee semantic integration success.

Structural Pattern: Correct context is reconstructed incorrectly.

Why It Matters: Many AI failures originate not from missing information, but from distorted interpretive reconstruction.

Related Terms: Retrieval Drift, Synthetic Coherence, Evaluation Blindness

Memory Drift

Memory drift describes the gradual divergence of stored representations from their originating conditions over time. Compression, summarization, preference inference, ranking systems, and recursive updates progressively mutate stored contextual state.

Structural Pattern: Persistent memory systems recursively reinterpret prior state.

Why It Matters: Persistent AI memory architectures increasingly shape continuity between users, agents, and semantic environments.

Related Terms: Context Drift, Persistent Context, Recursive Compression

Synthetic Coherence

Synthetic coherence describes outputs that appear internally consistent, persuasive, and structurally fluent despite weakened grounding. The representation maintains surface plausibility while semantic integrity deteriorates underneath.

Structural Pattern: Coherence detaches from grounding.

Why It Matters: Modern AI systems optimize heavily for fluency, plausibility, and continuity.

Related Terms: Survivable Wrongness, Proxy Alignment, Semantic Fidelity

Proxy Alignment

Proxy alignment occurs when systems optimize measurable representations instead of preserving underlying meaning, reality, or intent. Metrics gradually replace originating objectives.

Structural Pattern: Optimization targets proxies rather than realities.

Why It Matters: Recursive systems increasingly optimize benchmark performance, engagement, retrieval efficiency, or ranking behavior rather than semantic preservation.

Related Terms: Evaluation Blindness, Reality Drift, Representational Optimization

Survivable Wrongness

Survivable wrongness describes outputs that remain coherent, usable, and operational despite containing semantic misalignment or structural error. Failures persist because workflows continue functioning well enough to avoid triggering corrective pressure.

Structural Pattern: Systems remain operational despite degradation.

Why It Matters: Most modern drift emerges gradually through tolerated misalignment rather than catastrophic failure.

Related Terms: Reality Drift, Evaluation Blindness, Synthetic Coherence

Evaluation Blindness

Evaluation blindness occurs when benchmarks, metrics, or evaluation pipelines confirm surface correctness while failing to detect underlying semantic degradation. Systems appear aligned according to local metrics while semantic fidelity weakens underneath.

Structural Pattern: Evaluation systems fail to measure semantic integrity.

Why It Matters: Most current AI evaluation frameworks prioritize measurable outputs over representational continuity.

Related Terms: Proxy Alignment, Semantic Fidelity, Fidelity Benchmark

Lexical Decay

Lexical decay describes the narrowing of expressive vocabulary as rare, precise, metaphorical, or context-rich language is replaced by generic, statistically safe terms. In recursive AI systems, lexical decay can make outputs appear clear and fluent while reducing semantic texture, specificity, and cultural resonance.

Related terms: Semantic Drift, Meaning Debt, Ground Erosion, Synthetic Coherence.

SECTION III — Semantic Infrastructure & Architecture

Semantic Infrastructure

Semantic infrastructure refers to systems designed to preserve, organize, retrieve, stabilize, and govern meaning across recursive environments. This includes retrieval architectures, provenance systems, memory layers, ontologies, governance frameworks, and semantic observability systems.

Related Terms: Ontology Mapping, Provenance Systems, Semantic Continuity

Provenance Systems: Provenance systems preserve traceability between representations and their originating sources. They maintain lineage across retrieval, transformation, summarization, and generation workflows.

Related Terms: Recursive Mediation, Retrieval Pipelines, Provenance Layers

Ontology Mapping

Ontology mapping describes the structured modeling of conceptual relationships across semantic systems. Ontology layers stabilize meaning through explicit relational structure rather than purely statistical association.

Related Terms: Knowledge Graphs, Semantic Infrastructure, Conceptual Topology

Retrieval Pipelines

Retrieval pipelines govern how information is indexed, ranked, retrieved, filtered, and reintegrated into downstream systems. These architectures increasingly determine which representations remain visible within semantic environments.

Related Terms: Retrieval Drift, Semantic Search, Embedding Distortion

Observability Tooling

Observability tooling refers to systems for monitoring semantic degradation, retrieval quality, provenance integrity, drift accumulation, and memory reliability across recursive systems.

Related Terms: Evaluation Blindness, Semantic Infrastructure, Governance Frameworks

SECTION IV — Recursive Semantic Ecosystems

Recursive Archives

Recursive archives are living semantic systems in which representations recursively influence future representations over time. Archives become generative environments rather than static repositories.

Related Terms: Recursive Mediation, Semantic Continuity, Recursive Synthetic Feedback

AI-readable Conceptual Infrastructure

AI-readable conceptual infrastructure refers to semantic environments designed for both human interpretation and machine traversal. These systems structure concepts, relationships, provenance, and meaning in forms navigable by recursive AI systems.

Related Terms: Ontology Mapping, Semantic Infrastructure, Context Graphs

Personal Semantic Layers

Personal semantic layers are persistent contextual systems encoding memory, preferences, conceptual structure, and interpretive continuity across AI interactions.

Related Terms: Portable AI Memory, Shared Human-Agent Ontologies, Context Graphs

Recursive Synthetic Feedback

Recursive synthetic feedback occurs when generated representations recursively shape future retrieval, training, interpretation, and generation environments. Synthetic outputs become future semantic inputs.

Structural Pattern: Generated representations recursively modify future semantic conditions.

Why It Matters: Recursive AI ecosystems increasingly train on prior generated representations.

Related Terms: Recursive Mediation, Semantic Exhaust, Semantic Entropy

Representational Optimization

Representational optimization occurs when representations become optimized for circulation, retrieval, engagement, compression, or benchmark performance rather than preserving underlying reality or meaning.

Structural Pattern: Representations optimize for systems rather than reference.

Why It Matters: Modern semantic ecosystems increasingly reward transmissibility over fidelity.

Related Terms: Proxy Alignment, Reality Drift, Synthetic Coherence

Semantic Continuity

Semantic continuity is the preservation of coherent meaning, interpretive structure, and contextual alignment across recursive transformations over time. Semantic continuity represents the central stabilizing condition necessary for recursive AI systems operating at scale.

Related Terms: Semantic Fidelity, Recursive Mediation, Context Drift

Conclusion

The central challenge of recursive AI systems is not simply factual correctness, but preservation of semantic continuity across recursive transformations. As retrieval systems, memory architectures, agent workflows, provenance systems, and synthetic environments scale, meaning increasingly moves through recursive layers of compression, optimization, reinterpretation, and regeneration before reaching interpretation.

The resulting failures are often survivable rather than catastrophic. Systems remain coherent, operational, and locally functional while progressively weakening contact with originating reality, intent, context, or meaning.

The Semantic Fidelity Canonical Lexicon provides a conceptual infrastructure for describing these conditions. By naming these mechanisms, the lexicon establishes a framework for semantic observability, fidelity evaluation, provenance tracking, retrieval alignment, recursive governance, and semantic continuity preservation across increasingly AI-mediated environments.

This lexicon is intended as a living ontology for recursive semantic systems and future AI-native representational infrastructure.

Keywords: *semantic fidelity, semantic drift, recursive mediation, retrieval drift, AI memory systems, semantic infrastructure, recursive compression, agent drift, meaning preservation, provenance systems, context drift, representation drift, ontology mapping, semantic continuity, recursive AI systems*

Core Framework and Sources

- Substack (Articles)
- GitHub (Full Library)
- DOI (Research Paper)
- Glossary & Definition