

# Semantic Fidelity Project: Archive Index

A. Jacobs | v1 2026

The Semantic Fidelity Project studies how meaning survives, weakens, or mutates as language moves through AI systems, compression, retrieval, summarization, and generation. This archive collects the core essays of the project in sequence as the work develops a vocabulary for understanding semantic drift, fidelity decay, meaning loss, and the limits of accuracy-based AI evaluation.

The central claim is that AI systems can remain fluent, coherent, and factually acceptable while still losing meaning. Semantic fidelity names not only whether an output is correct, but whether the intent, context, hierarchy, nuance, and interpretive structure of the original material survived the transformation.

## Recommended Reading Order

### 1. [Welcome to the Semantic Fidelity Project](#)

Oct. 12, 2023

The opening statement for the project. Introduces semantic fidelity as a way to study meaning under AI compression, especially as language moves through models, interfaces, institutions, and digital systems.

**Start here for:** the broad purpose of the project.

### 2. [When Accuracy Isn't Enough: Toward Semantic Fidelity in AI Systems](#)

Jan. 19, 2024

Defines the gap between factual accuracy and semantic fidelity. Shows why an AI output can be correct, fluent, and grounded while still altering tone, intent, emphasis, or context.

**Start here for:** the core argument.

### 3. [Measuring Fidelity Decay: Tracking Meaning Collapse in Generative Systems](#)

Apr. 8, 2024

Introduces a measurement framework for semantic loss across AI systems. Defines four dimensions of fidelity decay: lexical decay, semantic drift, ground erosion, and semantic noise.

**Start here for:** the diagnostic model.

### 4. [A Semantic Fidelity Lexicon: Core Terms for Meaning, Drift, and Decay](#)

**Jul. 23, 2024**

A glossary of the project's core terms, including semantic fidelity, semantic drift, fidelity decay, ground erosion, lexical decay, semantic noise, meaning debt, synthetic coherence, and evaluation blindness.

**Start here for:** vocabulary and definitions.

## **5. Semantic Drift: The Failure Mode Hallucination Doesn't Explain**

**Nov. 4, 2024**

Explains why "hallucination" is too narrow as a frame for AI failure. Many failures are not fabricated facts, but subtle shifts in meaning that preserve surface coherence while weakening the original intent.

**Start here for:** the distinction between hallucination and semantic drift.

## **6. Language as Cognitive Exhaust: Compression, Meaning, and Fidelity Loss**

**Mar. 2, 2025**

Develops the idea that language is not thought itself, but the compressed residue of thought. AI systems train on this residue, which makes semantic fidelity a question of whether deeper cognitive structure survives expression and transformation.

**Start here for:** the cognition layer.

## **7. The Compression Paradox: Why Meaning Breaks Before Models Fail**

**Jul. 11, 2025**

Argues that summarization, paraphrase, and simplification may be among the most underestimated sources of semantic drift. Compression can preserve facts while removing constraint, hierarchy, uncertainty, and causal structure.

**Start here for:** why clean summaries can be dangerous.

## **8. Constraint Collapse: When Feedback Stops Correcting Symbolic Systems**

**Dec. 3, 2025**

Extends semantic fidelity into feedback and alignment. Defines constraint collapse as the condition where systems continue fluently after feedback stops correcting them. The system remains operational, but correction no longer forces revision.

**Start here for:** the alignment and feedback layer.

## Core Concepts

**Semantic Fidelity:** The preservation of meaning, intent, context, nuance, and interpretive structure as language moves through a system.

**Semantic Drift:** The gradual mutation of meaning across transformation, especially when fluent outputs preserve surface coherence while altering emphasis, framing, or intent.

**Fidelity Decay:** The cumulative weakening of semantic integrity across repeated compression, summarization, reinterpretation, or generation.

**Ground Erosion:** The loss of background context, hierarchy, implication, and unstated structure that gives language its weight.

**Lexical Decay:** The thinning of expressive language as precise or context-rich terms become generic, overused, or detached from concrete reference.

**Semantic Noise:** The saturation of an information environment with fluent language that carries little distinction, making retrieval, interpretation, and trust harder.

**Meaning Debt:** The accumulated cost of small semantic losses that seem harmless in isolation but compound over time.

**Synthetic Coherence:** Language that remains fluent, consistent, and persuasive while becoming less grounded underneath.

**Evaluation Blindness:** The failure of benchmarks or metrics to detect semantic degradation because the output still appears accurate, coherent, or relevant.

**Constraint Collapse:** A failure mode where feedback continues to flow but no longer forces correction. The system remains fluent and operational while losing contact with what feedback was supposed to repair.

**Survivable Wrongness:** A condition where a system can be wrong without breaking because the cost of error is pushed onto users, workflows, or later correction layers.

**Recursive Compression:** The repeated compression, summarization, and regeneration of language across systems, where each pass may preserve surface coherence while gradually thinning context, nuance, and meaning.

## Project Arc

The archive moves through four layers:

- 1. Definition:** What semantic fidelity is and why accuracy is not enough.
- 2. Measurement:** How meaning weakens through lexical decay, ground erosion, semantic drift, and semantic noise.
- 3. Failure Mode:** Why hallucination is only one visible failure, while semantic drift is often quieter and harder to detect.
- 4. Infrastructure:** How compression, feedback loops, retrieval, memory, and AI agents can preserve fluent operation while gradually weakening meaning.

## The Question Accuracy Cannot Answer

AI evaluation often asks whether an output is correct. The Semantic Fidelity Project asks if the meaning survived? That question matters because AI systems increasingly mediate search, writing, summarization, decision support, and institutional communication. As more language moves through AI, meaning can be compressed, cleaned, reformatted, and reused until the output still sounds right while carrying less of the original context.

This archive is a record of that problem as it develops from accuracy, to fidelity, to drift, to compression, to constraint collapse.

## Related Resources

- [Semantic Fidelity Project - Substack](#)
- [Semantic Fidelity Project - Github Repo](#)
- [Glossary & Definition](#)
- [Failure Modes in LLM Systems – Paper Series](#)