

How Meaning Drifts Through AI Systems

Semantic Fidelity and the Representation Pipeline

Reality Drift Framework — Representational Systems Brief #1

A. Jacobs

The Basic Pattern

AI systems do not process human intent directly.

They process a series of increasingly compressed representations of that intent.

At each stage, meaning is preserved.

At each stage, meaning is also lost.

The sequence is:

intent → language → prompt → embedding → retrieval → context window → output

Each transformation makes information more usable for computation.

Each transformation also introduces opportunities for semantic drift.

The Representation Pipeline

Intent

Every interaction begins with an intention, question, goal, or idea.

Intent contains assumptions, context, priorities, tacit knowledge, and lived experience that have not yet been expressed.

It is the richest representation in the pipeline.

Language

Intent is translated into natural language.

Language communicates meaning remarkably well, but it cannot perfectly preserve every nuance, assumption, or internal state.

Some meaning is made explicit.

Some remains unspoken.

Prompt

The language is organized into a prompt.

The prompt emphasizes certain details while omitting others.

It introduces structure, priorities, and constraints intended to guide the model.

Already, the representation has become more computationally legible than the original intention.

Embedding

The prompt is converted into a mathematical representation.

The system no longer processes words directly.

It processes relationships encoded within a high-dimensional vector space.

Semantic similarity is preserved.

Exact meaning is not.

Retrieval

If external knowledge is used, the embedding retrieves documents judged to be semantically relevant.

Retrieval does not recover meaning.

It retrieves representations that appear similar within the embedding space.

Relevant information may be missed.

Irrelevant information may appear close.

Context Window

Retrieved information is compressed again into the model's working context.

Only part of the available evidence fits within the context window.

Ordering, truncation, and prioritization shape what the model ultimately reasons over.

Output

The model generates the most probable continuation based on the available representations.

The response may be coherent.

It may also differ in subtle but important ways from the original intent.

The system has reconstructed meaning rather than recovered it.

Where Meaning Drifts

Meaning rarely disappears in a single step.

It accumulates across representations.

Small losses at each transformation compound over the entire pipeline.

The output may preserve the central topic while losing intent, emphasis, assumptions, relationships, uncertainty, or context.

This is why an answer can be factually correct while still feeling "off."

Semantic Fidelity

Semantic fidelity is the degree to which meaning survives across successive representations.

High semantic fidelity means that intent, relationships, priorities, and context remain aligned throughout the pipeline.

Low semantic fidelity occurs when each individual transformation appears reasonable, but their accumulated effect changes the original meaning.

The failure is cumulative rather than catastrophic.

Reality Drift

Within the **Reality Drift framework**, AI systems do not interact directly with human meaning.

They interact with increasingly abstract representations of that meaning.

As those representations become compressed, reordered, filtered, and reconstructed, the model can become more responsive to the representations than to the intention they originally expressed.

This is a form of Reality Drift.

The model remains internally coherent.

Its answer may no longer remain fully answerable to the reality it was intended to represent.

The Core Principle

Every AI system is a representational pipeline.

Each transformation makes computation easier.

Each transformation also creates an opportunity for semantic drift.

Improving AI is therefore not simply a matter of generating better outputs.

It is a matter of preserving meaning across every layer of representation.

Semantic fidelity is not achieved at the output. It is preserved—or lost—throughout the entire pipeline.

Keywords: semantic fidelity, how meaning drifts through AI systems, AI representation pipeline, prompt engineering, embeddings explained, retrieval augmented generation, RAG, context window, interpretation failure, semantic drift, AI reasoning, representation pipeline, Reality Drift, A. Jacobs