OMNISENS AI

Extrinsic Drift and the Need for an External Reasoning Substrate

A Framework for Stable World-State Coherence in Large Language Model Systems

Author:
Elin Nguyen
Omnisens Lab
research@omnisensai.com

Date:
November 2025 — Working Paper v1.0

Series:
Omnisens Lab Working Paper Series
Omnisens-WP-2501

Abstract

Large language models exhibit drift — variability in output across time, sessions, or runs. Drift is commonly attributed to intrinsic stochasticity (sampling noise, decoding variance), but operational evidence shows that **extrinsic drift**—misalignment between model inference and an evolving external world-state—dominates in real enterprise environments.

This paper formalizes extrinsic drift as an *architectural* problem rather than a model-level limitation. We argue that drift persists because LLMs lack mechanisms for persistent state, contradiction resolution, or world-state coherence. Internal interventions (prompting, fine-tuning, retrieval, deterministic decoding) reduce intrinsic variance but cannot eliminate divergence from evolving operational truth.

We introduce the concept of an **external reasoning substrate Φ** — a deterministic coherence layer that computes a canonical world-state outside the model. We provide empirical evidence via a multi-model experiment on an ambiguous e-commerce clickstream: four independent frontier models produced divergent interpretations, yet all converged to a single stable interpretation once Φ was applied.

1. Introduction

LLMs are trained to model conditional probabilities over text, not to maintain coherent truth over time. As a result, operational systems experience drift: the model's interpretation of the same underlying situation changes unpredictably across runs.

Drift has two sources:

- **Intrinsic drift:** sampling variance, decoding noise, calibration differences
- **Extrinsic drift:** mismatch between model inference and a non-stationary, partially observed external world

While intrinsic drift is well-studied, **extrinsic drift is the primary failure mode** in real deployments — CRM timelines, support logs, buyer journeys, audit trails, and multi-session operational data.

This paper focuses on extrinsic drift and explains why **no internal modification to the model can solve it**, and why **an external reasoning substrate is required**.

2. Formalizing Drift

Let:

- **E(t)** = the true external world-state at time t
- **R(t)** = the model's inferred state
- **δ(t) = D(E(t), R(t))** = divergence between the two

Drift decomposes into:

- **Intrinsic drift:** variance in R(t) from decoding
- **Extrinsic drift:** divergence caused by incomplete, contradictory, or asynchronous evidence streams that the model cannot reconcile internally

In realistic settings, **E(t) evolves independently of the model**, causing extrinsic drift to dominate.

3. Limits of internal fixes

Internal interventions operate solely on **R(t)**:

- better prompting
- fewer sampling degrees of freedom
- retrieval augmentation
- fine-tuning
- chain-of-thought
- tool-augmented agents

These methods reduce intrinsic variance but CANNOT:

- observe missing external state
- maintain state across sessions
- enforce consistency with past decisions
- resolve contradictions
- reconstruct a canonical world-state

Thus, drift persists because the model has no architectural access to the evolving truth **E(t)**.

4. Structural limitations of transformers

Transformers inherently lack:

- **persistent memory** (beyond the context window)
- **global consistency operators**
- **truth-maintenance mechanisms**
- **cross-session coherence**
- **temporal invariants**

They *can* detect contradictions **within** a single prompt, but cannot maintain consistency across time.
This is the core reason extrinsic drift is unavoidable internally.

5. Non-stationary operational environments

Enterprise systems are:

- asynchronous
- contradictory
- multi-channel
- partially observed
- often manually corrected
- full of delays, overrides, missing data

Truth is **not given** — it must be reconstructed.

This is incompatible with the autoregressive next-token paradigm.

6. External reasoning substrate

We introduce an architectural component **outside the model**:

Φ: A deterministic reasoning substrate that maintains canonical world-state.

Φ performs:

- persistent state tracking
- invariant enforcement
- contradiction reconciliation
- evidence fusion across channels and sessions
- canonicalization into a stable, deterministic representation

Φ is not the only theoretically possible solution, but it is **the most tractable and architecture-agnostic** given current generative model limitations.

7. Relation to existing attempts

Fragments of Φ appear in:

- memory graphs
- KV-persistent agents
- retriever-augmented pipelines
- "threads" with external storage
- custom coherence layers in large enterprises

But these are piecemeal, lacking the unified principle:

Reasoning stability requires a dedicated external substrate.

8. Empirical evaluation: ambiguous e-commerce clickstream

We constructed an ambiguous, multi-session clickstream containing:

- contradictory user actions
- add/remove cart events
- delayed interactions
- weak intent signals
- overlapping behavioral patterns

This dataset mirrors real operational ambiguity.

8.1 Baseline multi-model behavior (Without Φ)

We evaluated four independent frontier models.
All models disagreed on:

- buyer’s true stage
- whether abandonment occurred
- whether intent was high or weak
- whether a purchase attempt existed

Each model also disagreed **with itself** across runs.

This is extrinsic drift.

8.2 Application of Φ

Applying Φ yielded **deterministic collapse**:

All four models converged to the same canonical interpretation:

"High Intent with Abandonment"

This demonstrates:

- extrinsic drift dominates inference instability
- internal fixes cannot resolve it
- a deterministic external substrate **enforces coherence**

Table 1 — Before/After Φ (Abstracted)

Model A: inconsistent → S3
Model B: inconsistent → S3
Model C: inconsistent → S3
Model D: inconsistent → S3

9. Conclusion

Extrinsic drift is the dominant source of inconsistency in operational LLM deployments. Because transformers cannot maintain persistent truth or reconcile contradictions, internal fixes cannot eliminate drift.

A stable reasoning system requires **an external deterministic substrate** that canonicalizes world-state independent of the model's generative process.

This architectural separation — stochastic model + deterministic substrate — offers a path toward reliable, auditable, and consistent long-horizon reasoning.