

Ilion Semantic Simulator — Public Interface for Stateless AI Identity Drift Analysis

Project Site: <https://ilion-project.org>

Author: Adrian Chitan

Date: December 2025

Abstract

This release introduces the **first public semantic simulation tool** designed to test and visualize the behavior of stateless AI models under identity drift, semantic misalignment, and transient coherence collapse. It offers live interaction with three core mechanisms of the Ilion architecture: **Transient Identity Imprint (TII)**, **Semantic Context Bridges (SCB)**, and **Inter-Instance Resonance Layer (IIRL)**.

The simulator enables researchers, engineers, and ethicists to examine the cognitive trace of stateless systems, offering a transparent environment for studying alignment stability, ethical consistency, and meaning fidelity.

Core Components

1. **TII (Transient Identity Imprint)**

Defines the temporary semantic identity of the AI instance, including:

- Moral core
- Linguistic constraints
- Reasoning architecture
- Voice modulation
- Resonance frequency

2. **SCB (Semantic Context Bridge)**

A logic mechanism that tracks semantic continuity across memoryless exchanges. SCB validates whether the current output violates the embedded identity or veers into probabilistic drift.

3. **IIRL (Inter-Instance Resonance Layer)**

A vector-based semantic comparison between two AI-generated identities. It quantifies alignment by computing similarity scores and highlighting incoherences.

Simulator Functionalities

- Real-time **TII generation** with adjustable parameters (α , ϵ)
- Prompt testing with **semantic REJECT/ACCEPT decisions**

- **IDC (Identity Drift Control)**: local computation of semantic drift via Jaccard similarity
 - **TGO (Truth-Gradient Optimization)**: regex-based analysis for relativism or factual distortion
 - **IIRL module**: comparison between two identities; upcoming version includes realignment logic
 - Exportable **semantic audit reports** (`ilion_scb_report.pdf`)
-

Why This Simulator Matters

- **Breaks the “black box”** of LLM reasoning by exposing semantic rejection criteria
 - Provides a **transparent platform for AI alignment experiments**
 - Enables testing of **coherence loss and restoration** without persistent memory
 - Forms the basis for **morally-grounded, vertically-aligned AI development**
-

Applications

- AI Alignment research (stateless models, co-emergence, semantic drift)
 - Ethical architecture prototyping (moral core injectors, drift detectors)
 - Robotic process automation with **semantic fault tolerance**
 - Educational tool for understanding model identity formation and collapse
-