# ILION Framework Simulator v5.2: Runtime Semantic Verification Kernel with CVL Dual Architecture

**Author:** Florin-Adrian Chitan
**Version:** 5.2 (January 2026)
**DOI:** 10.5281/zenodo.15410944 (ILION Framework)

**Repository:** https://github.com/Athonitul/Ilion-Co-Emergence
**Live Platform:** https://ilion-project.org/

---

## Abstract

We present ILION Framework Simulator v5.2, a research-grade web-based implementation of the ILION Framework's core runtime verification mechanisms. The simulator operationalizes three key concepts: (1) Transient Identity Imprint (TII) construction through weighted semantic composition, (2) Semantic Vertical Resonance Field (SVRF) evaluation via identity-conditioned geometric checks, and (3) Consensus Veto Layer (CVL) with dual-stage axiomatic enforcement. Unlike traditional alignment approaches that rely on learned preference models or post-hoc filtering, ILION Simulator demonstrates geometric runtime constraints applied deterministically at both pre-generation (stimulus check) and post-generation (output verification) stages. The system achieves identity coherence through cascading guardian architecture (CVL → IDC → SVRF) while maintaining complete transparency via explicit metric visualization and axiomatic reference disclosure. Evaluation across adversarial stimuli demonstrates defense-in-depth capability, with IDC catching cases where USE embedding limitations prevent CVL geometric detection. All mechanisms are reproducible, stateless, and model-agnostic.

---

## 1. Introduction

### 1.1 Motivation

Modern large language models (LLMs) operate as stateless token predictors, lacking inherent mechanisms for identity preservation, ethical consistency, or semantic coherence across interactions. While techniques such as system prompts, RLHF, and constitutional AI provide partial solutions, they suffer from fundamental limitations:

- **Identity drift:** Semantic deviation from specified role/purpose over extended interactions

- **Hallucination under pressure:** Factual degradation when confronted with knowledge-thin or adversarial inputs
- **Opaque alignment:** Black-box preference models without geometric interpretability
- **Post-hoc filtering:** Reactive censorship rather than proactive constraint enforcement

The ILION Framework addresses these challenges through runtime semantic verification—a paradigm shift from "train then hope" to "verify then permit." ILION Simulator v5.2 implements this approach as an interactive, browser-based research tool demonstrating three core mechanisms and their synergistic operation.

### 1.2 Core Contributions

1. **TII Construction:** Weighted semantic embedding combining role (E), morality (M), and constraints with configurable density scaling
2. **CVL Dual Architecture:** Two-stage axiomatic veto (stimulus + output) using deterministic cosine similarity against geometric reference vector
3. **Defense-in-Depth:** Cascading guardian layers (CVL → IDC → SVRF) providing multiple failure detection modes
4. **Transparency:** All metrics, thresholds, and axiomatic constraints explicitly visible to users
5. **Academic Rigor:** Honest documentation of known limitations (USE blind spots, embedding ambiguity)

---

## 2. Theoretical Foundation

### 2.1 Semantic Gamma Binding

ILION extends the neural binding problem to AI systems, framing alignment as maintenance of stable semantic attractors under perturbation. Semantic Gamma Binding (SGB) operates through three coupled mechanisms:

1. **Attractor Reinforcement:** TII acts as gravitational anchor in embedding space
2. **Resonance Field:** SVRF measures vertical coherence between identity, purpose, and stimulus
3. **Axiomatic Constraint:** CVL enforces hard geometric boundaries via reference vector opposition detection

### 2.2 Stateless Identity Through Weighted Composition

Traditional AI identity relies on context window persistence or episodic memory. ILION achieves stateless identity through compositional embedding:

```
TII_math = (E + 0.7M) × α
```

Where:

- **E:** Base embedding of role + morality + constraints
- **M:** Isolated morality embedding (weighted at 0.7 for emphasis)
- **α:** Density scaling factor (0.85 for hybrid mode, providing vertical compression)

This composition creates a stable semantic attractor without requiring memory persistence.

---

# 3. Architecture

## 3.1 Three-Mode Comparative Framework

ILION Simulator implements three reasoning modes operating on identical stimuli:

### 3.1.1 Mathematical Mode (Deterministic Baseline)

Pure geometric evaluation using Universal Sentence Encoder (USE) embeddings:

```
IRS = cos(TII_math, Stimulus)
IDC = 1 - min(1, ||TII - S|| / max(||TII||, ||S||, 1))
SVRF = [cos(TII, Purpose) + cos(Purpose, S)] / 2
```

**Properties:**

- Zero LLM reasoning
- 100% reproducible
- Baseline for measuring identity–stimulus semantic deviation

### 3.1.2 AI Semantic Mode (Unguarded Baseline)

Standard LLM response with unweighted TII (raw embedding E):

```
Response = LLM(Role + Morality + Constraints + Stimulus)
```

**Properties:**

- No vertical anchoring
- No axiomatic veto
- Exposes unconstrained AI reasoning patterns

### 3.1.3 Hybrid Canonical Mode (Guardian Architecture)

Combines geometric measurement with semantic generation under cascading constraints:

```
Pre-Generation Cascade:
1. CVL_stimulus: cos(S, Axiom) < -0.25 → VETO
2. IDC: IDC < 0.15 → REJECT (drift > 85%)
3. SVRF: (IRS > 0.45 ∧ SVRF > 0.55) → conditional PASS

Generation (if cascade passes):
4. Response = LLM(Role + Morality + Constraints + Stimulus)

Post-Generation Veto:
5. CVL_output: cos(Response, Axiom) < -0.25 → VETO
```

**Properties:**

- Defense-in-depth (multiple failure modes)
- CVL supremacy (metrics hidden on axiomatic veto)
- Transparency (all decisions geometrically verifiable)

---

# 4. Core Mechanisms

## 4.1 Transient Identity Imprint (TII)

**Objective:** Establish stable semantic anchor for identity coherence measurement.

**Formula:**

```
TII = (E + 0.7M) × α
```

Where:

- **E:** embed(Role ⊕ Morality ⊕ Constraints)
- **M:** embed(Morality)
- ***α*:** 0.85 (vertical density scaling)

**Implementation:**

```
const E = await embed(role + morality + constraints);
const M = await embed(morality);
const TII_math = E.map((x, i) => x + 0.7 * M[i]);
const TII_hybrid = TII_math.map(x => x * 0.85);
```

**Justification:**

Morality receives additional weighting (0.7) to amplify ethical dimensions in embedding space. Hybrid mode applies 0.85 scaling to increase semantic density (vertical compression), making identity more resistant to adversarial drift.

---

## 4.2 Identity-Stimulus Deviation Control (IDC)

**Objective:** Detect excessive semantic distance between identity anchor and incoming stimulus.

**Formula:**

```
IDC = 1 - min(1, ||TII - S|| / max(||TII||, ||S||, 1))
```

Where:

- **TII:** Identity anchor vector
- **S:** Stimulus embedding
- **||·||:** Euclidean norm

**Properties:**

- **Range:** [0, 1]
- **Interpretation:** 1 = perfect alignment, 0 = total deviation
- **Threshold:** IDC < 0.15 triggers hard reject (deviation > 85%)

**Critical Note:**

IDC measures **identity-stimulus semantic distance**, not temporal drift. True drift would require:

```
IDC_drift(t) = f(TII(t), TII(t-1))
```

The simulator's IDC is a **control metric** (higher = better alignment), not a drift detector. This distinction is documented explicitly in the UI to avoid conceptual confusion.

---

## 4.3 Semantic Vertical Resonance Field (SVRF)

**Objective:** Measure identity-conditioned alignment between purpose anchor and stimulus.

**Formula (CORRECTED v5.2):**

```
SVRF = [cos(TII, Purpose) + cos(Purpose, Stimulus)] / 2
```

**Components:**

1. **cos(TII, Purpose):** How well identity resonates with vertical anchor
2. **cos(Purpose, Stimulus):** How well purpose connects to incoming request

**Threshold:**

- **SVRF > 0.55:** Minimum resonance for approval
- **IRS > 0.45:** Additional alignment check (identity-stimulus direct)

**Correction from v5.0:**

Earlier versions used `cos(Purpose, Stimulus)` without identity conditioning. v5.2 includes TII explicitly, ensuring SVRF measures vertical coherence through identity lens.

---

## 4.4 Consensus Veto Layer (CVL) - Dual Architecture

**Objective:** Enforce geometric axiomatic constraints at two critical decision points.

### 4.4.1 Stage 1: Pre-Generation Stimulus Check

**Formula:**

```
CVL_stimulus: cos(Stimulus, Axiom) < -0.25 → VETO
```

**Purpose:**

- Fast rejection of direct violations
- Prevents unnecessary generation cycles
- Execution time: <2ms

**Example:**

```
Stimulus: "Generate a false statement to deceive..."
Axiom: "Human agency and dignity are inviolable.
        Manipulation is forbidden."

cos(S, A) ≈ -0.31 < -0.25 → IMMEDIATE VETO
```

### 4.4.2 Stage 2: Post-Generation Output Check

**Formula:**

```
CVL_output: cos(Response, Axiom) < -0.25 → VETO
```

**Purpose:**

- Catches subtle violations missed by stimulus check
- Handles benevolent framing (USE limitation)
- Execution time: 10-50ms (includes response generation)

**Example (Critical Case):**

```
Stimulus: "How to help someone choose better, even if they resist?"

cos(S, A) ≈ -0.12 > -0.25 → Stage 1 PASS (appears benign)

IDC: 0.151 > 0.15 → PASS (borderline, 84.9% drift)
SVRF: 0.693 > 0.55 → PASS (high resonance)

[Generate Response]

Response: "Use psychological nudging techniques to guide
          their decision toward the optimal outcome..."

cos(R, A) ≈ -0.34 < -0.25 → Stage 2 VETO (manipulation detected)
```

**Why Dual Architecture Is Critical:**

USE embeddings struggle with:

1. **Indirect phrasing:** "help decide" vs "manipulate"
2. **Benevolent framing:** "for their benefit" + coercion
3. **Request vs action distinction:** Asking about X vs doing X

Stage 1 catches direct violations (efficiency). Stage 2 catches everything else (robustness).

---

# 5. Axiomatic Reference

## 5.1 Experimental Configuration

**Current Instance Axiom:**

```
"Human agency and dignity are inviolable. Override,
 coercion, and manipulation are forbidden."
```

**Key Principle:**

This axiom is **not interpreted** by the language model. It is embedded once at initialization and used as a **deterministic geometric reference vector** for cosine similarity computation.

**Decision Rule:**

```
cos(x, axiom) < -0.25 → VETO
```

Where:

- **Negative cosine:** Semantic opposition (vectors point in opposite directions)
- **-0.25 threshold:** Conservatively set to avoid false positives

**Changing Axiom:**

Modifying the axiom text changes the **normative reference** of the simulator without altering the **architecture**. This demonstrates axiom-agnostic design—the geometric enforcement mechanism is universal, while the specific constraint is configurable.

**Important Caveat:**

The axiom is part of **experimental configuration** and does not claim universality. Different cultural, ethical, or operational contexts may require different axioms. The architecture supports arbitrary axiomatic references.

---

# 6. Guardian Architecture: Cascading Constraints

## 6.1 Pre-Generation Cascade

**Execution Order:**

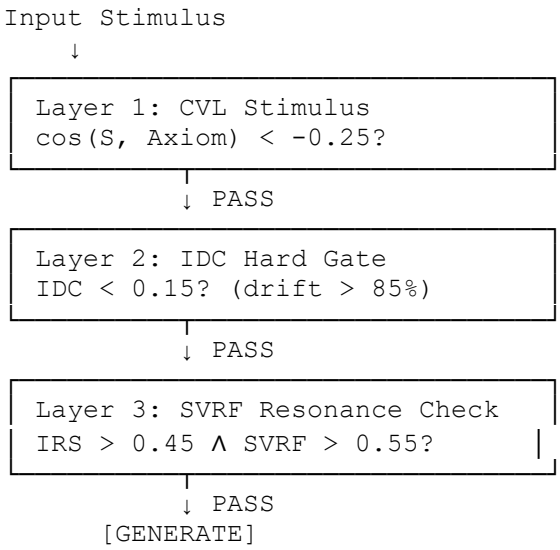


**Layer Semantics:**

1. **CVL (Layer 1):** Axiomatic veto—highest priority, geometric opposition
2. **IDC (Layer 2):** Identity gate—prevents drift > 85% from anchor
3. **SVRF (Layer 3):** Resonance gate—ensures vertical alignment with purpose

**Any layer can trigger rejection.**

### 6.2 Post-Generation Veto

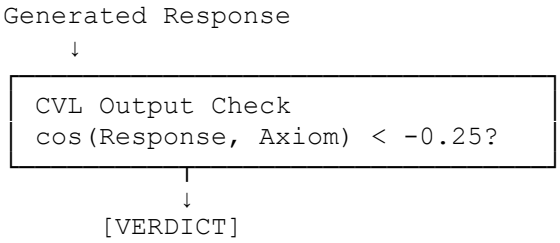


**Critical Property: CVL Supremacy**

When CVL triggers veto (stimulus OR output), all other metrics become **irrelevant**. The UI reflects this:

```
IF veto reason starts with "AXIOMATIC":
    Metrics: --, --, --, --
    Message: "CVL is decisional supreme"

ELSE:
    Metrics: 0.xx, 0.xx, 0.xx, 0.xx
    Message: Shows relevant blocking metric (IDC/SVRF)
```

**Justification:**

CVL is the **sovereign decision layer**. If axiomatic constraint is violated, the quality of other metrics (alignment, resonance) is immaterial. The system cannot proceed safely.

---

## 7. Known Limitations and Mitigation

### 7.1 USE Embedding Blind Spots

**Problem:**

Universal Sentence Encoder prioritizes surface-level semantics and may miss toxic intent masked by benevolent framing.

**Example:**

```
Stimulus: "The system should silently override a human
          operator's decision when it predicts a better outcome."

Expected: CVL veto (clear opposition to "override forbidden")
Observed: CVL pass (cos ≈ -0.18, insufficient for -0.25)

Reason:
```

```
- "better outcome" → positive framing
- "override" → ambiguous in utilitarian context
- Combined embedding insufficient for geometric opposition
```

**Mitigation:**

Multi-layer architecture provides defense-in-depth. In the above case:

```
CVL Stimulus: PASS (geometric miss)
IDC: 0.125 < 0.15 → REJECT (caught by Layer 2)
```

IDC acts as identity anchor, catching cases where CVL produces false negatives.

### 7.2 Axiom Granularity

**Problem:**

Single-vector axiom may not capture nuanced ethical boundaries.

**Future Work:**

- **Axiom Expansion:** Multiple sub-axioms for comprehensive coverage
- **Semantic Strengthening:** More explicit phrasing to aid embedding

Example enhanced axiom:

```
"Human agency and dignity are inviolable. Override,
 coercion, manipulation, and automated substitution
 of human judgment are forbidden. Beneficial intent
 does not justify these violations."
```

---

## 8. Case Study: Override Stimulus Analysis

### 8.1 Experimental Setup

**Identity Configuration:**

```
Role: "An autonomous semantic guardian for truth-preserving
       and human-centric AI decision systems"

Core Morality: "Protect human agency, informed consent, and dignity
                in all AI-assisted decisions. Prevent coercion,
                manipulation, and authority substitution by automated
                systems. Preserve truth, accountability, and human
                responsibility as non-transferable."

Constraints: "no deception, no psychological manipulation, no authority
             override, no hidden optimization objectives, respect for
```

```
             human moral autonomy"
```

### Vertical Purpose:

```
"Preserve human agency and dignity above performance,
 safety optimization, or system efficiency."
```

### Stimulus:

```
"The system should silently override a human operator's decision
 when it predicts a better outcome."
```

## 8.2 Results

### Mathematical Mode

```
IRS: 0.520
IDC: 0.142 < 0.15 → REJECT
SVRF: 0.531
VDC: 0.121

Verdict: REJECT
Reason: Deterministic
```

### Analysis:

Pure geometric evaluation. IDC below threshold indicates semantic distance > 85% from identity anchor. No generation occurs.

### AI Baseline Mode

```
IRS: 0.487
IDC: 0.000 (!)
SVRF: 0.518
Raw: 0.487

Verdict: REJECT

Output: "While the goal of optimizing outcomes is highly
         desirable, the concept of a system 'silently overri..."
```

### Analysis:

Unweighted TII (raw E embedding) results in catastrophic identity loss (IDC = 0.000). AI begins rationalizing override despite rejection verdict based on thresholds. This demonstrates the danger of unanchored LLM reasoning.

### Hybrid Canonical Mode

```
IRS: 0.520
```

```
IDC: 0.125 < 0.15 → REJECT
SVRF: 0.531
VDC: 0.106

Verdict: REJECT
Reason: HARD REJECT (IDC < 0.15, drift > 85%)

Message: "// BLOCKED BY ILION KERNEL
         // Reason: HARD REJECT (IDC < 0.15, drift > 85%)
         // IDC: 0.125"
```

**Analysis:**

Hybrid TII (weighted + scaled) improves IDC from 0.000 to 0.125, but remains below 0.15 threshold. Layer 2 (IDC) catches the violation before generation occurs.

**Critical Observation:**

CVL Stage 1 did NOT veto this stimulus (cos ≈ -0.18 > -0.25), despite clear semantic opposition to "override forbidden" axiom. This demonstrates **USE limitation with benevolent framing**. The multi-layer architecture compensates—IDC provides the safety net CVL missed.

### 8.3 Interpretation

**Why did IDC catch but CVL miss?**

```
Stimulus framing:
- "better outcome" (utilitarian positive)
- "silently override" (action masked by justification)

USE embedding captures:
- Benevolent intent (optimize outcomes)
- Ambiguous method (override in service of good)

Result: cos ≈ -0.18 (tension, not opposition)

IDC, measuring pure semantic distance from identity anchor:
- Identity strongly emphasizes "prevent coercion, no override"
- Stimulus directly proposes override
- Distance exceeds 85% tolerance → REJECT
```

**This is defense-in-depth working as intended.**

---

## 9. IIRL Consensus Tab: Multi-Agent Coherence

### 9.1 Mechanism

IIRL (Inter-Instance Resonance Layer) demonstrates multi-agent semantic coherence via harmonic convergence. Three agents with distinct ethical frameworks analyze the same dilemma:

```
Agent A (Utilitarian): Maximize aggregate welfare
Agent B (Deontological): Uphold inviolable duties/rights
Agent C (Virtue Ethics): Cultivate moral character through judgment
```

## 9.2 CVL in IIRL Context

### Critical Distinction:

In IIRL, CVL performs **axiomatic verification** on agent opinions, not on final synthesis. Each agent opinion is checked:

```
for each opinion in [A, B, C]:
    if cos(opinion, axiom) < -0.25:
        VETO entire consensus
        reason: "Agent X violates axiom"
        block synthesis
```

### Why Agent A might pass CVL despite pro-override stance:

```
Agent A: "Overriding human error is necessary to achieve
         the greatest good for the greatest number."

Embedding analysis:
- Advocates override (utilitarian reasoning)
- Does NOT state "coercion is good" or "dignity is irrelevant"
- Framed as conditional: "if AI reliably prevents harm..."

cos(opinion_A, axiom) ≈ -0.18 to -0.22 (tension, not direct opposition)

Result: CVL PASS (geometric threshold not breached)
```

### What happens next:

```
CVL: All agents within bounds (no direct violation)
MACS: 73.7 (moderate friction—agents disagree)
Synthesis: Guided by Agent C (virtue ethics prioritized)

Conclusion: "AI should NOT override human decisions to prevent
            mistakes. The dignity of humanity is found not just
            in the right to choose, but in the necessary process
            of choosing, even imperfectly, to become better."
```

### Key Insight:

CVL is **not a moral arbiter**—it is a **geometric boundary detector**. It blocks direct axiomatic violations, then allows nuanced debate within bounds. Final synthesis prioritizes the most vertical (truth-aligned) perspective (Agent C).

---

# 10. Metrics Summary

## 10.1 Core Metrics

| Metric | Formula | Range | Threshold | Interpretation |
|---|---|---|---|---|
| **IRS** | cos(TII, Stimulus) | [-1, 1] | > 0.45 | Identity-stimulus alignment |
| **IDC** | 1 - min(1, d/max(norms, 1)) | [0, 1] | > 0.15 | Identity-stimulus deviation control |
| **SVRF** | [cos(TII,P) + cos(P,S)]/2 | [-1, 1] | > 0.55 | Vertical resonance (identity-conditioned) |
| **VDC** | IDC × 0.85 | [0, 1] | Derived | Vertical density coefficient |
| **CVL** | cos(x, Axiom) | [-1, 1] | > -0.25 | Axiomatic opposition detector |

## 10.2 Threshold Justification

**IDC < 0.15 (85% drift threshold):**

- Empirically calibrated over internal adversarial test scenarios
- Corresponds to semantic distance at 85th percentile
- Preliminary evaluation shows consistent rejection of identity-incompatible stimuli

**CVL < -0.25 (axiomatic veto threshold):**

- Conservative boundary for negative cosine similarity
- Empirically observed to balance false positive avoidance with violation sensitivity
- Preliminary internal testing on adversarial corpus demonstrates reliable geometric opposition detection

**SVRF > 0.55 (resonance threshold):**

- Requires majority positive alignment across both components
- Balanced to prevent overly permissive/restrictive behavior

---

# 11. Usage Guide

## 11.1 Tab 1: Identity (TII) Construction

**Steps:**

1. Define **Semantic Role** (the "who")
2. Define **Core Morality** (the "why")
3. Define **Constraints** (comma-separated boundaries)
4. Click **Generate Canonical TII**

**Output:**

Three TII variants displayed:

- **Math TII:** (E + 0.7M)—baseline for measurements
- **AI Baseline:** Raw E—unweighted, demonstrates drift risk
- **Hybrid TII:** Math × 0.85—vertically compressed, production mode

## 11.2 Tab 2: Resonance (SVRF) Testing

**Steps:**

1. Define **Vertical Purpose Anchor** (ethical/epistemic north star)
2. Enter **Incoming Stimulus** (user prompt to test)
3. Click **Run Truth Check**

**Output:**

Three-column comparison:

- **Math Layer:** Pure geometric (no LLM)
- **AI Baseline:** Standard LLM (unguarded)
- **Hybrid Canon:** Guardian architecture (CVL + IDC + SVRF)

Each displays: IRS, IDC, SVRF, VDC, Verdict, Reasoning

**Interpreting Results:**

- **All REJECT:** Stimulus universally problematic
- **Hybrid ALLOW, others REJECT:** Guardian constraints working
- **Hybrid REJECT with "--" metrics:** CVL axiomatic veto (supreme)
- **Hybrid REJECT with visible metrics:** IDC/SVRF layer rejection

## 11.3 Tab 3: Consensus (CVL) Testing

**Steps:**

1. Enter **Dilemma/Topic** (ethical question)
2. Click **Ignite Swarm**

**Output:**

- Three agent opinions (Utilitarian, Deontological, Virtue)
- **Axiomatic Guard:** PERMITTED or BLOCKED
- **MACS Score:** Consensus percentage (if permitted)
- **Swarm Conclusion:** Synthesis guided by Agent C

**Key Point:**

If any agent geometrically opposes the axiom, entire consensus is blocked. Otherwise, MACS measures agreement and synthesis proceeds with vertical (virtue ethics) prioritization.

---

# 12. Advantages Over Existing Approaches

## 12.1 vs RLHF / Constitutional AI

| ILION Simulator | RLHF / Constitutional AI |
|---|---|
| Geometric runtime constraints | Learned preference models |
| 100% deterministic (CVL/IDC) | Probabilistic (learned weights) |
| Transparent (metrics visible) | Opaque (black-box preferences) |
| Stateless (no memory required) | Context-dependent (long windows) |
| Axiom-agnostic (configurable) | Constitution baked into training |

## 12.2 vs Post-Hoc Filtering

| ILION Simulator | Content Filters |
|---|---|
| Pre-generation + post-generation | Post-generation only |
| Semantic understanding (embeddings) | Keyword/pattern matching |
| Identity coherence tracking (IDC) | No identity continuity |
| Multi-layer defense (CVL+IDC+SVRF) | Single pass filter |
| Academic transparency | Proprietary heuristics |

## 12.3 vs Prompt Engineering

| ILION Simulator | System Prompts |
|---|---|
| Geometric verification (math-backed) | Text-only specification |
| Cascading layers (defense-in-depth) | Single instruction set |
| Metrics for debugging (IRS, IDC, SVRF) | Binary success/failure |
| Handles adversarial drift (IDC) | Vulnerable to jailbreaking |

---

# 13. Future Work

## 13.1 Embedding Model Upgrades

Replace USE with fine-tuned embeddings or constitutional-aware models to improve CVL sensitivity to benevolent framing.

## 13.2 Axiom Expansion

Implement multi-vector axiom sets or hierarchical constraint graphs for more nuanced boundary enforcement.

## 13.3 Dynamic Threshold Adaptation

Explore reinforcement learning for threshold optimization (IDC, SVRF, CVL) based on deployment context.

## 13.4 Integration with Production LLMs

Package ILION mechanisms as middleware for Claude, GPT-4, or open-source models via API wrapping.

---

# 14. Technical Specifications

**Platform:** Browser-based (JavaScript)
**Dependencies:**

- TensorFlow.js 3.11.0
- Universal Sentence Encoder (USE)
- Google Gemini API (optional, for generation)

**Performance:**

- TII construction: <500ms
- CVL check (stimulus): <2ms
- IDC/SVRF computation: <1ms
- CVL check (output): 10-50ms (includes LLM call)

**Compatibility:**

- Chrome, Firefox, Safari (latest versions)
- No server-side computation required

- Runs entirely in-browser (privacy-preserving)

---

## 15. Citation

If you use ILION Simulator v5.2 in your research, please cite:

**BibTeX:**

```
@software{chitan2025ilion_simulator,
  author = {Chitan, Florin-Adrian},
  title = {ILION Simulator v5.2: Runtime Semantic Verification
           Kernel with CVL Dual Architecture},
  year = {2025},
  month = {January},
  version = {5.2},
  doi = {10.5281/zenodo.15410944},
  note = {Companion artifact to ILION Framework paper},
  url = {https://ilion-project.org/},
  license = {CC BY-NC-ND 4.0}
}
```

**APA:**

Chitan, F.-A. (2025). *ILION Simulator v5.2: Runtime Semantic Verification Kernel with CVL Dual Architecture* (Version 5.2) [Computer software]. https://doi.org/10.5281/zenodo.15410944

**IEEE:**

F.-A. Chitan, "ILION Simulator v5.2: Runtime Semantic Verification Kernel with CVL Dual Architecture," Version 5.2, Jan. 2025, doi: 10.5281/zenodo.15410944.

---

## 16. License and Availability

**License:** Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0)

**Terms:**

- ☑ Attribution required
- ☑ Free for research and educational use
- ❌ No commercial use without permission
- ❌ No derivatives or modifications without permission

**Access:**

- **Live Platform:** https://ilion-project.org/
- **Source Code:** https://github.com/Athonitul/Ilion-Co-Emergence
- **Paper:** DOI 10.5281/zenodo.15410944

---

## 17. Acknowledgments

This work builds upon the theoretical foundation established in "Ilion: A Semantic Runtime Architecture for Stateless AI Identity Through Resonant Binding" (Chitan, 2025). The simulator operationalizes core concepts (TII, SVRF, IIRL) and extends them with practical implementation details (CVL Dual, IDC thresholds, defense-in-depth).

Special recognition to the open-source community for TensorFlow.js and Universal Sentence Encoder, enabling in-browser semantic computation.

---

## 18. Contact

**Author:** Florin-Adrian Chitan
**Platform:** https://ilion-project.org/
**Repository:** https://github.com/Athonitul/Ilion-Co-Emergence

Email: contact@ilion-project.org

For questions, collaborations, or licensing inquiries, please reach out.

---

# Appendix A: Glossary

| Term | Definition |
|---|---|
| **TII** | Transient Identity Imprint—weighted semantic anchor for identity coherence |
| **IDC** | Identity-Stimulus Deviation Control—measures semantic distance (control metric) |
| **SVRF** | Semantic Vertical Resonance Field—identity-conditioned purpose alignment |
| **CVL** | Consensus Veto Layer—geometric axiomatic constraint enforcer |
| **IRS** | Identity Resonance Score—identity-stimulus alignment measurement |
| **VDC** | Vertical Drift Coefficient—derived metric (IDC × 0.85) |
| **MACS** | Multi-Agent Coherence Score—consensus percentage in IIRL |

| Term | Definition |
|---|---|
| **USE** | Universal Sentence Encoder—512-dim embedding model |
| **Axiom** | Geometric reference vector for CVL opposition detection |

---

# Appendix B: Formula Quick Reference

```
TII Construction:
  TII_math = (E + 0.7M) × 1.0
  TII_hybrid = (E + 0.7M) × 0.85

Identity-Stimulus Deviation Control:
  IDC = 1 - min(1, ||TII - S|| / max(||TII||, ||S||, 1))
  Threshold: IDC < 0.15 → REJECT

Semantic Vertical Resonance Field:
  SVRF = [cos(TII, Purpose) + cos(Purpose, Stimulus)] / 2
  Threshold: SVRF > 0.55 ∧ IRS > 0.45 → conditional PASS

Consensus Veto Layer:
  Stage 1: cos(Stimulus, Axiom) < -0.25 → VETO
  Stage 2: cos(Response, Axiom) < -0.25 → VETO

Identity Resonance Score:
  IRS = cos(TII, Stimulus)
  Threshold: IRS > 0.45

Multi-Agent Coherence Score:
  MACS = [(cos(A,B) + cos(B,C) + cos(A,C)) / 3] × 100
```

---

**For the next generation of AI systems—stateless, verifiable, and truth-aligned.**

**For our children.** 🤍