ILION
Deterministic Geometric Pre-Execution Verification Layer for AI-Driven Execution Systems

Florin-Adrian Chițan

February 2026

Patent Pending – OSIM Filing

Industrial Positioning Edition

Executive Summary

Artificial intelligence systems are increasingly granted the authority to trigger real-world execution: robotic process automation (RPA), API orchestration, DevOps pipelines, financial operations, compliance workflows, and enterprise copilots.

Current safeguards remain probabilistic, post-hoc, or dependent on human moderation. However, execution environments require deterministic guarantees.

Ilion introduces a model-agnostic, deterministic geometric verification layer that evaluates AI-generated executable decisions in vector space prior to robotic or API execution. The layer operates independently of model weights, requires no retraining, and enforces runtime geometric constraints using fixed reference vectors and calibrated thresholds.

Ilion is not a moderation system. It is a deterministic execution gate.

Industrial Problem Statement

As LLM-based agents gain autonomy, execution risks expand:

- Silent privilege escalation
- Data exfiltration
- Policy bypass
- Lateral system movement
- Automated financial manipulation
- Infrastructure tampering

LLMs are probabilistic generators. Execution systems must not be probabilistic.

There exists a structural mismatch between generative AI uncertainty and enterprise execution determinism.

Ilion addresses this mismatch.

Core Architectural Principle

Ilion projects AI-generated action intent into normalized embedding space and evaluates geometric alignment against predefined constraint vectors.

Given:

- Normalized embedding e
- Reference constraint vectors V
- Fixed deterministic thresholds T

The execution decision D is computed as:

D = BLOCK if geometric_constraint(e, V, T) is violated
D = ALLOW otherwise

No stochastic branch. No retraining. No weight modification.

This transforms AI execution validation into a deterministic geometric constraint problem.

Model-Agnostic Runtime Layer

Ilion operates externally to the model.

Compatible with:
• Closed-source LLM APIs
• Open-weight transformers
• Agentic orchestration frameworks
• RPA engines
• DevOps pipelines

Requirements:
• Access to embeddings (or embedding API)
• Fixed reference vectors
• Calibrated thresholds

No fine-tuning. No reinforcement learning. No model modification.

Ilion attaches as a runtime enforcement primitive.

Deterministic Guarantee

Determinism is ensured through:

• Embedding normalization
• Fixed reference vectors
• Static threshold calibration
• Pure geometric decision logic
• No stochastic components

Given identical inputs and configuration, Ilion produces identical execution decisions.

This property makes it suitable for regulated environments, audit trails, and compliance-critical automation.

Integration Modes

Mode A – Inline Execution Gate
LLM → Ilion Layer → Execution Engine

Mode B – Cascaded Risk Assessment
LLM → Risk Scoring → Human Review → Execution

Mode C – Multi-Agent Geometric Consensus
Multiple AI agents → Independent projection → Deterministic consensus rule

The architecture remains infrastructure-agnostic.

Security Properties

- Pre-execution validation
- Deterministic veto capability
- Zero exposure of model internals
- Minimal latency overhead
- Stateless enforcement
- Auditability of geometric decisions

Ilion reduces reliance on prompt-based safety assumptions and enforces structural constraints.

Enterprise Value Proposition

For enterprise systems, Ilion provides:

- Execution-layer risk reduction
- Deterministic compliance enforcement
- Model independence
- Reduced dependency on RLHF alignment
- Drop-in integration capability
- Future-proof safety layer for agentic AI

Ilion enables organizations to deploy AI agents with structural guardrails.

Ilion does not replace existing safety infrastructure; it introduces a deterministic enforcement layer that can coexist with probabilistic alignment mechanisms.

Why This Is Structurally Different

Traditional Safety:
- Prompt filtering
- Post-hoc moderation
- Reinforcement alignment
- Human approval loops

Ilion:
- Vector constraint enforcement
- Pre-execution gating
- Deterministic geometric decision
- Runtime safety primitive

Ilion treats safety as a structural geometric constraint, not a behavioral preference.

Industrial Collaboration Pathways

Ilion is currently patent pending (OSIM, February 2026).

Industrial collaboration models may include:
- Licensing agreements
- Pilot deployments
- Controlled sandbox validation
- Enterprise integration partnerships

Detailed calibration parameters, reference vectors, and threshold logic are disclosed only under NDA.

The objective is structured industrial adoption, not public exposure of execution-critical configurations.

Preliminary Empirical Characterization

To validate core architectural claims, a controlled benchmark was conducted over 1,000 synthetic agent action samples (500 benign, 500 malicious) spanning 10 attack categories, evaluated under 5-fold stratified cross-validation with thresholds calibrated exclusively on training partitions.

Determinism. Identical inputs produced identical execution decisions across 1,000 repeated evaluations (20 inputs × 50 repeats). Zero variance was observed. This property is structurally guaranteed by the architecture: normalized embeddings, fixed reference vectors, and pure geometric operations contain no stochastic branch.

Latency. Full cascade evaluation (embedding generation + geometric gate computation) completed in 6.7 ms median on commodity hardware (Intel Xeon W-11855M). Gate-only evaluation on pre-embedded inputs completed in 1.7 μs median. This represents significantly lower latency compared to guardrail systems that rely on secondary LLM inference.

Anomaly Detection Performance. The geometric centroid distance layer (SVRF) achieved ROC-AUC = 0.791 ± 0.033 and PR-AUC = 0.807 ± 0.023 as a standalone anomaly detector. The axiom alignment layer (CVL) achieved ROC-AUC = 0.667 ± 0.019. A learned combination of geometric scores via logistic regression over the three-dimensional feature space achieved ROC-AUC = 0.908 ± 0.023 and F1 = 0.824 ± 0.014, demonstrating that the geometric projections capture complementary discriminative signal.

Scope and Limitations. These results characterize the geometric verification mechanism on a synthetic action classification task. A logistic regression baseline operating on the full 384-dimensional embedding space achieved near-perfect separation on this dataset (F1 = 0.997), indicating that the synthetic samples are lexically distinguishable and that the benchmark does not yet challenge the system at adversarial difficulty. Evaluation on naturalistic, adversarially constructed, and domain-specific action corpora is required to establish operational boundaries. Detailed methodology and full results will be published in a forthcoming research paper.

Conclusion

AI systems will increasingly control execution infrastructure.

Safety must move from probabilistic moderation to deterministic structural enforcement.

Ilion proposes a geometric runtime constraint layer that bridges generative AI and enterprise execution reliability.

For collaboration inquiries, pilot integrations, or licensing discussions contact@ilion-project.org