

Decision Trace Generation for Governable AI Systems

Version 1.0

June 2026

A proposal for independent scientific validation of runtime-generated decision traces for trustworthy AI systems.

We are not asking whether AI can become more intelligent.

We are asking whether AI execution can become scientifically traceable.

Background

Artificial intelligence is rapidly expanding into scientific research, healthcare, manufacturing, logistics, finance, public administration, and many other domains.

As AI systems increasingly participate in decisions with real-world consequences, transparency, reproducibility, and governance become fundamental requirements rather than optional features.

One question therefore becomes increasingly important:

Can AI execution become structurally observable and scientifically verifiable without modifying the underlying AI model?

Current approaches often evaluate AI through benchmark performance, prediction accuracy, reasoning capability, or explainability.

SSGE (Structural System Governance Engine) explores a different architectural direction.

Rather than changing the AI model itself, SSGE introduces an independent Governance Runtime that observes execution and continuously generates structured Decision Trace artifacts describing execution flow, evidence references, rule evaluation, semantic state transitions, and governance events throughout AI execution.

The objective is not to improve AI capability, but to improve the observability, traceability, and governability of AI execution itself.

Proposed Scientific Validation

We would like to invite your research team to participate in a small feasibility study.

Research Question

Can an independent Governance Runtime generate a structured Decision Trace that meaningfully represents AI execution without modifying the underlying AI model?

The validation focuses exclusively on this question.

No benchmark comparison, commercial deployment, or product evaluation is involved.

Example Validation Scenario

Scenario

An AI system receives structured inputs and generates operational recommendations or decisions for a human operator in a real-world application.

Governance Runtime

The proposed study investigates whether an independent Governance Runtime can continuously generate structured Decision Trace artifacts throughout AI execution without influencing the underlying AI model.

Validation Focus

The objective is not to evaluate model intelligence or prediction accuracy, but to investigate whether runtime-generated execution evidence can improve the observability, scientific traceability, and governance readiness of AI systems.

Scope (2-week Feasibility Study)

The proposed study would evaluate whether SSGE can:

- generate structured Decision Trace artifacts during AI execution;
- preserve execution order and semantic decision relationships;
- produce trace information that is independently reviewable by researchers;
- provide a reproducible basis for governance-oriented analysis of AI systems.

Expected Outcome

The expected outcome is not proof that SSGE is superior to existing systems.

Instead, we seek to answer a more fundamental scientific question:

Does Decision Trace Generation represent a viable architectural direction for governable AI systems?

If the answer is positive, the results could support broader collaborative research, future scientific publications, and the development of runtime governance as an independent research direction for trustworthy AI.

The observations obtained from such a study may also provide a reusable validation framework applicable across multiple AI application domains.

Collaboration Principles

This proposal is:

non-commercial

non-exclusive

intended purely for scientific exploration.

We welcome critical evaluation and independent feedback.

Why This Validation Matters

If successful, this study would provide initial empirical evidence that runtime-generated execution evidence can improve the scientific traceability of AI execution without affecting the underlying AI models.

More broadly, the study investigates whether Governance Runtime architectures can become a foundational infrastructure layer for transparency, accountability, and trustworthy AI across multiple application domains.

Shuqin Amberg

Founder, SSGE

Governance Runtime Research

<https://ssge.ai>

Email: shuqin@ssge.ai