

Decision Trace Generation for Governable AI Execution

Version 1.0

June 2026

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

We are not asking whether AI can make better decisions.

We are asking whether AI decisions can become scientifically traceable.

Background

Large Language Models and AI agents have achieved remarkable capabilities across reasoning, planning, and autonomous task execution.

However, one fundamental question remains largely unanswered:

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

Current approaches mainly evaluate outputs after execution through benchmarks, logs, or post-hoc explanations.

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

Rather than changing the AI model, SSGE introduces an independent Governance Runtime that observes execution and generates a structured Decision Trace describing the reasoning path, execution flow, evidence references, and governance events.

The objective is not to improve model intelligence, but to improve the observability and governability of AI execution.

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 performance benchmark, commercial deployment, or product evaluation is involved.

Example Validation Scenario

Scenario

An autonomous robotic system performs routine operational tasks using an existing AI model.

Governance Runtime

An external Governance Runtime observes execution without modifying the robot controller and continuously generates structured execution evidence.

Validation Focus

The study investigates whether runtime-generated Decision Traces can provide meaningful scientific evidence for execution traceability and governance analysis.

Scope (2-week Feasibility Study)

The proposed study would evaluate whether SSGE can:

generate structured Decision Trace records during AI execution;

preserve execution order and decision relationships;

produce trace information that is independently reviewable;

provide a reproducible basis for governance-oriented analysis.

The study is intentionally limited in scope and designed as an exploratory scientific validation.

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 execution?

If the answer is positive, the results could support further collaborative research, broader validation studies, and future publications.

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 artifacts can improve the scientific observability of autonomous AI systems without modifying the underlying AI models.

More broadly, the study explores whether execution evidence can become an independent architectural layer for trustworthy autonomous systems, contributing to future research on AI governance and runtime infrastructure.

Shuqin Amberg

Founder, SSGE

Governance Runtime Research

<https://ssge.ai>

Email: shuqin@ssge.ai