

SAGE ENGINEERING

Bridging the Gap Between Probability and Precision

Solving Structural Hallucination in Generative AI: The Structural Integrity Enforcement Pipeline (SIEP)

Arvind Jolly

Founder, School of Ancient Geomantic Education (www.dotsofdestiny.com)

SAGE Engineering Group

Noida, India

ABSTRACT

Large Language Models (LLMs) are fundamentally non-deterministic, often failing to maintain structural consistency in complex outputs—a phenomenon defined as "Structural Hallucination." This paper introduces the Structural Integrity Enforcement Pipeline (SIEP), a six-stage deterministic post-processing framework designed to ensure 100% schema compliance without re-querying the model. We demonstrate the efficacy of this architecture within the SAGE Geomancy engine, achieving production-grade reliability across 10 languages.

1. Introduction

The rapid adoption of Large Language Models (LLMs) in production systems has revealed a critical vulnerability: structural instability. While LLMs excel at generating semantically rich prose, they struggle to adhere to rigid output schemas over long-context or multilingual generations. In high-precision applications, such as the SAGE Geomancy engine, a single missing section or malformed header renders the output invalid for the end user.

2. Defining Structural Hallucination

Unlike factual hallucination, where a model invents false data, structural hallucination refers to the probabilistic failure of the model to follow formatting constraints. Common failure modes include merged sections, omitted headers, or the "echoing" of internal prompt instructions. These failures are often silent, bypassing standard string-based validation while breaking the user interface.

3. The SIEP Architecture

The Structural Integrity Enforcement Pipeline (SIEP) is a six-stage deterministic firewall that operates *after* the LLM generation. It treats the model as an unreliable narrator and enforces structure through the following stages:

- 1 Ingestion & Prompt Echo Purge
- 2 Doctrine Firewall & Canonical Header Restoration
- 3 Anchor-Based Synthetic Header Recovery
- 4 Hyper-Isolated Translation Shielding ([PART-N])
- 5 Token Desquashing & Placeholder Repair
- 6 Multi-Metric Automated QA Validation

4. Implementation & Results

By implementing Stage 3 (Anchor-Based Recovery), the system can detect semantic keywords unique to a specific section and synthetically inject missing headers. This allows the system to "repair" a hallucinated output in milliseconds, avoiding the latency and cost of a second LLM API call. In production, SAGE has observed 100% schema compliance across 10 global languages, including RTL (Arabic) and double-byte (Chinese/Japanese) scripts.

5. Conclusion

As generative AI moves from novelty to utility, reliability becomes the primary differentiator. The SIEP framework proves that deterministic engineering can be used to harness probabilistic models for mission-critical applications. The School of Ancient Geomantic Education (SAGE) continues to pioneer these "hybrid" architectures to ensure that ancient wisdom is delivered with modern precision.