

Schema-Aware AI Infrastructure: A Methodology for Hallucination-R

Placeholder working-draft PDF. Replace with the real manuscript when available.

Abstract: We outline a three-move methodology ROSETTA, CoreModels, and MCP Traversal for grounding LLM-based agents in the full Semantic Enrichment Stack rather than vector flattening.

We review prior empirical results (Ali et al., 2026, reporting hallucination rates dropping from 63% to under 2% under ontology-grounded retrieval) and outline an evaluation protocol for SCR in regulated domains.

Keywords: Structured Context Retrieval, ontology grounding, knowledge graphs, hallucination, agentic AI.

ARAMAI Working Notes