WHITE PAPER: THE EIGEN ENGINE

Deterministic Truth in Enterprise AI via Sovereign Memory & Quantum-Inspired Dynamics

Author: Madhava Bekkem

Affiliation: Qohere Private Limited, Bangalore

Date: December 2025

Classification: Public Release

Sovereign Memory	Conflict Gate	Truth Adjudication (Ising Solver)

1. Executive Summary

The deployment of Large Language Models (LLMs) across the enterprise—from Legal and Finance to Sales and Engineering—remains paralyzed by a critical flaw: Probabilistic Hallucination. Standard Retrieval-Augmented Generation (RAG) systems operate on "semantic similarity." They often conflate outdated price sheets with current quotes, deprecated code documentation with new APIs, or superseded drafts with final policies. Whether in a high-stakes audit or a routine customer support query, this "Data Rot" creates liability and operational friction.

We introduce a paradigm shift: The Eigen Engine. Rather than guessing based on text similarity, the Eigen Engine utilizes a physics-based solver to mathematically adjudicate truth. It processes conflicting data into Sovereign Memory—an immutable, state-collapsed database that guarantees deterministic behavior regarding document status and authority. By moving from Probability (guessing) to Physics (solving), we provide the first AI architecture capable of understanding the concept of "Truth" for every department in the organization.

2. The Core Problem: The "Similarity Trap"

In a corporate environment, truth is not static. A "Draft Contract v3" looks 99% identical to a "Final Contract v4." A "2023 Price List" looks semantically identical to a "2025 Price List."

2.1 Vector Space Blindness

- The Flaw: Vector embeddings cluster documents based on semantic meaning.
- The Result: Because the "Old Version" and the "New Version" mean the same thing semantically, they occupy nearly identical coordinates in vector space.
- The Failure Mode: The AI retrieves both. Lacking a mechanism to discern authority, the LLM blends the superseded data with the active data, synthesizing a hallucination that can cost a company revenue (wrong pricing) or compliance (wrong policy).

2.2 The Black Box Liability

When a standard AI fails, it offers no explanation. Enterprises require deterministic logic, not probabilistic guessing. If a Sales Agent quotes a price, they need to know exactly which document authorized that number.

3. Technology Overview: The Physics of Truth

The Eigen Engine differentiates itself by introducing a Physics Adjudication Layer between the database and the LLM. We model the problem of "Information Retrieval" as an Energy Minimization Problem, using an Ising solver.

3.1 The Energy Landscape

- State Representation: Every document exists in a specific state—either "Active" (Truth) or "Suppressed" (Noise).
- The Goal: The system's fundamental operation is to settle into a stable configuration where only the most authoritative, non-conflicting documents remain Active.

3.2 Interaction Management (Conflict Resolution)

- Standard RAG treats every document independently. The Eigen Engine calculates Pairwise Fidelity.
- Conflict Detection: The system identifies logical conflicts (e.g., differing dates, conflicting clauses, overlapping pricing tiers) between documents.
- Repulsive Force: When a conflict is detected, the system creates a "repulsive" interaction. It becomes energetically impossible for both documents to remain Active simultaneously. The system is mathematically forced to suppress one to minimize global energy.

3.3 The Authority Field (Bias)

- Gravitational Pull: We assign a "weight" to documents based on Sovereign Metadata—such as provenance (CEO vs. Manager) and recency (New vs. Old).
- The Result: A "2025 Policy" exerts a stronger magnetic pull than a "2020 Policy." Combined with the repulsive force, the system naturally collapses into a state where the current version dominates and the old version is suppressed.

4. System Architecture

The architecture consists of two distinct components: The Processing Layer and the Storage Layer.

4.1 Layer A: The Eigen Engine (The Processor)

1. Vector Probe: The system retrieves the top candidates based on semantic similarity.
2. Landscape Construction: The engine builds the energy landscape, mapping conflicts and authority scores.
3. Dynamic State Collapse: The solver evolves the system. High-energy documents (conflicts, low authority) are purged.

4. Eigenstate Output: The system returns a stable subset of documents—the "Ground State"—which represents the single, undisputed truth.

4.2 Layer B: Sovereign Memory (The Storage)

- Hierarchical Enforcement: Unlike flat vector stores, Sovereign Memory enforces a "Monarchy" of data. A document tagged Status: FINAL physically overrides Status: DRAFT in the retrieval path.
- Episodic Continuity: The memory layer ingests long-term session logs, shattering years of chat history into searchable "Memory Particles."
- Poison Resistance: The memory is immutable to real-time user feedback, preventing "Jailbreak" or "Data Poisoning" attacks.

5. Operational Capabilities

5.1 The "Conflict Gate"

- Logic: If the engine detects a high degree of conflict between documents (e.g., >90% similarity between a Draft and Final), the Conflict Gate activates.
- Effect: The system parameters are tuned to make the coexistence of both documents physically impossible in the solution space.

5.2 Hybrid "Dual-Brain" Logic

To ensure usability across the enterprise, the system employs a Dynamic Switching Architecture:

- Mode 1: Sovereign Mode (Low Entropy) — Trigger: Internal queries (e.g., "What is the price of SKU-X?"). Behavior: The LLM is strictly constrained to the retrieved Eigenstate. If the energy state is unstable (no answer found), the system refuses to answer.
- Mode 2: General Mode (High Entropy) — Trigger: External/General queries (e.g., "Draft an email to the client."). Behavior: The system detects the lack of internal documents and unlocks the LLM's parametric knowledge.

6. Target Industries & Use Cases

The Eigen Engine drives value across the entire enterprise ecosystem, securing data integrity for both regulated and operational sectors.

Industry / Function	Use Case	Value Proposition
Sales & Revenue Ops	Pricing & Quote Accuracy	Automatically suppresses expired price lists and discounted offers from previous quarters. Ensures every quote generated aligns with the current fiscal strategy.
Customer Support	Knowledge Base Management	Filters out "Workaround Guides" from 2022 in favor of "Official Patch Notes" from 2025, ensuring agents never give outdated troubleshooting advice.
Engineering & IT	Documentation Versioning	Distinguishes between Deprecated (v1.0) and Active (v2.0) API documentation. Prevents developers from hallucinating code based on obsolete libraries.
Human Resources	Policy Administration	Resolves conflicts between legacy policies (e.g., "Full Remote 2021") and current mandates (e.g., "Hybrid 2025"), providing a single source of truth for employee queries.
Legal & Compliance	Discovery & Contract Review	Automatically suppresses superseded drafts and repealed statutes. Reduces malpractice risk by ensuring citations are active.
Banking & Finance	Audit & Risk Management	Provides a deterministic, mathematically provable explanation for every financial decision, satisfying "Explainable AI" (XAI) regulations.
Pharma & Science	Protocol Management	Ensures lab researchers only see the current approved clinical protocol (v50), strictly ignoring the 49 previous iterations.

7. Enterprise Integration Strategy

7.1 The "Truth Adjudication" Middleware

- Ingestion Pipeline: The system connects to standard enterprise repositories (e.g., SharePoint, Drive, Codebases). It does not merely index text; it extracts "Sovereign Metadata"—such as timestamps, version numbers, and author authority—to construct the "Energy Landscape" required for physics processing.
- The Physics API: When a downstream application (such as a Customer Support Chatbot or Legal Review Tool) requests information, it does not query the raw database. Instead, it queries the Eigen Engine. The engine runs the "Dynamic State Collapse" in real-time or retrieves a pre-computed "Ground State", delivering only the non-conflicting, authoritative result to the user.

7.2 Deployment Modes

- Real-Time Adjudication: For operational environments like Sales, where data changes hourly, the engine resolves conflicts dynamically at the moment of query.
- Pre-Computed Sovereignty ("Air-Gap" Ready): For highly regulated environments like Pharma or Defense, the solver can run on a scheduled cadence (e.g., nightly). This creates a static, immutable "Sovereign Memory" snapshot that is stable and can be deployed to offline or secure air-gapped networks, ensuring zero latency and zero external connectivity risks.

8. Governance & Compliance Framework

8.1 From "Black Box" to "Glass Box"

- Traceable Decision Paths: Because the system uses a deterministic solver rather than a probabilistic neural network for retrieval, every output can be traced. Auditors can view exactly which "Repulsive Force" suppressed a draft document and which "Authority Bias" elevated the final policy.
- Audit-Ready Logs: The system generates a deterministic proof of the decision, satisfying strict regulatory requirements for Audit & Risk Management in banking and finance sectors.

8.2 Role-Based Reality

- Hierarchical Access: The concept of truth often depends on the user's role. The authority field can be adjusted per user group, ensuring a draft visible to Legal is suppressed as noise for Sales.

9. Comparative Architecture Analysis

The following comparison highlights the structural differences between traditional RAG systems and the Eigen Engine approach.

Feature	Standard RAG Systems	The Eigen Engine
Core Mechanism	Probabilistic: Relies on semantic similarity (vector space closeness).	Deterministic: Relies on physics-based energy minimization (Ising solver).
Conflict Resolution	None: Retrieves conflicting Draft and Final versions if they look similar.	Physics-Based: Uses repulsive forces to make coexistence of conflicting docs impossible.
Hallucination Risk	High: Blends outdated and current data into a plausible but wrong answer.	Ground-State Only: Returns a stable truth set or refuses to answer.
Data Structure	Flat: All documents are equal in the vector space.	Hierarchical: Enforces a monarchy where Final overrides Draft.
Liability Profile	Unpredictable: Data rot creates financial and legal risk.	Defensible: Provides a provable, auditable selection rationale.

10. Conclusion

The transition from "Generative Creativity" to "Deterministic Truth" marks the maturation of Enterprise AI. By acknowledging that corporate data is not static—that it contains conflicts, drafts, and superseded truths—the Eigen Engine addresses the root cause of AI failure: the "Similarity Trap".

For the enterprise, the value is clear. It is the difference between an AI that guesses a price based on an old email, and an AI that knows the price based on the active, authorized price list. The Eigen Engine ensures that as AI scales across the organization, it remains anchored to the only metric that matters: the Truth.

For organizations where accuracy is non-negotiable—whether in a courtroom or a sales negotiation—Sovereign Memory is the safe harbor for enterprise data.