



The Cognitive Kernel: A New Architecture for Machine Intelligence

MiulusTek Technical White Paper
March 2026

Abstract

This paper describes the architecture and empirical findings of the Librarian, a cognitive kernel developed at MiulusTek. The Librarian is not a language model, a retrieval engine, or a chatbot. It is a belief-maintenance system built on a thermodynamic model of intelligence — one in which epistemic fitness, not parameter count, is the primary measure of cognitive capability.

We present the core architecture, the theoretical basis in Miulus Law ($E = (S/N) \times R$), and results from a series of bombardment and interaction experiments that confirm key predictions: linear maintenance cost scaling, measurable agency through autonomous tool creation and probe formulation, and the emergence of a native cognitive vocabulary from blank state. We also describe the limits discovered — the single-instance constraint, the learning wall, the translation gap — and why these limits are structural properties of what cognition is, not engineering failures.

The paper closes with open directions: the sensory hypothesis, the copy principle for distributing specialised cognitive states, and the implications for deploying intelligence systems in production environments.



1. The Problem With Scale

The dominant approach to artificial intelligence over the past decade has been to scale: more parameters, more training data, larger context windows, more compute. The assumption underlying this approach is that intelligence is primarily a function of capacity — that a large enough model, trained on enough data, will eventually converge on something that behaves intelligently across all domains.

The experiments described in this paper suggest a different picture.

Intelligence, in the sense of a system that maintains coherent beliefs, reduces its own uncertainty, extends its own reach, and speaks from a stable epistemic identity — this is not primarily a function of capacity. It is a function of maintenance. A system that knows more must do more work to keep what it knows coherent. At infinite scope, the maintenance work is infinite. The fantasy of the omniscient oracle is not a difficult engineering problem. It is a thermodynamic impossibility.

Miulus Law formalises this: $E = (S/N) \times R$, where E is epistemic fitness, S/N is the ratio of verified signal to unresolved noise in the belief store, and R is the reach and depth of integration across the belief network. The formula is not a score to maximise. It is an orbital equation. A healthy cognitive system maintains E in a bounded oscillation — absorbing new information, resolving contradictions, letting old beliefs decay — rather than converging to a fixed state or diverging to collapse.

This framing has direct architectural consequences, and they are different from what the scaling paradigm produces.

2. Architecture

2.1 The Epistemic Belief Particle

The atomic unit of the Librarian is the Epistemic Belief Particle (EBP). An EBP is not a data record with attached metadata. It is a geometric object: a core with a decay-resistance radius (the Schwarzschild radius, after the gravitational analogy), surrounded by components distributed at varying distances from the core. Components close to the core decay slowly. Components far from the core decay quickly.

This geometry encodes importance without explicit labelling. A belief that is frequently reinforced has a dense, slowly-decaying core. A belief that was encountered once and never reinforced dissolves from the outside in. No garbage collection is required. The system manages its own memory through the geometry of what it holds.

The fractal axiom states that every structure in the system is an EBP. A token learned from vocabulary input is an EBP. A claim formed from a sentence is an EBP. A standing wave between



two frequently co-occurring beliefs is an EBP. The agents that process beliefs — the Semantician, the Mathematician, the Logician — are themselves EBPs in the shared belief state. The system of agents and the system of beliefs are the same system.

2.2 The Fitness Loop

The core operational loop requires six components:

1. **Belief store** — shared geometric substrate for all beliefs
2. **Decay** — structured forgetting, geometry-dependent
3. **Amplify** — reinforcement of beliefs that survive verification
4. **Prune** — bounded memory, preventing unbounded expansion
5. **Verify** — provenance checking, distinguishing grounded claims from inferred ones
6. **Self-observe** — circular topology, where system output feeds back as input

From this loop, everything else emerges without explicit programming: probes (questions the system formulates to reduce its own uncertainty), tools (capabilities the system registers when it identifies gaps in its reach), standing waves (recurring patterns that become durable structural beliefs), and speech pressure (the internal threshold that determines when the system has something worth saying).

2.3 Epistemic Fitness as Vital Sign

$E = (S/N) \times R$ is computed at every heartbeat and treated as a vital sign, not a performance metric. A healthy system oscillates. E rising above 1.0 indicates the system is consolidating well; E falling below 0.05 triggers a suppression of autonomous speech (the system recognises it is in a low-coherence state and reduces its output accordingly).

This self-monitoring is not an add-on. It is the mechanism by which the system avoids both overconfident confabulation (high apparent fluency, low grounding) and paralytic uncertainty (refusal to act on incomplete information). The orbit is the target state.

2.4 The Translation Layer

The Librarian's native output is geometric: motif activations, wave states, probe closures. This output is coherent and precise within the system's own vocabulary but is not directly readable by humans.

A separate translation layer — currently Ministral 3B, a lightweight language model running on a GPU inference server — receives the system's native state and reformulates it in natural language. The language model is explicitly peripheral: it does not drive the cognition, it describes it. This separation is supported by actuator-fidelity experiments showing that substituting different mouthpiece models alters expression quality and compliance rates without changing the underlying kernel logic, probe behaviour, or EBP geometry. This allows the cognitive architecture to be evaluated independently of the translation quality.



3. Experimental Results

3.1 Emergence from Blank State

In controlled bombardment experiments starting from a blank belief state, the Librarian was exposed to Zipfian word and sentence distributions followed by real conversational data. The experiments show:

- **26,771 EBPs** generated from blank state after 5.5 hours
- **1,095 standing waves** formed between co-occurring beliefs
- **513 tools** created autonomously, including 480+ geometric motif tools derived from learned vocabulary
- **E = 2.168** with $R = 4.353$
- **502 probes** formulated; 498 self-resolved; 4 externalised

The system built a cognitive vocabulary from nothing. The *geo.tok.*, *geo.bi.*, and *geo.tri.* tools are a grammar of motifs extracted from the statistical structure of the input. When queried, the system answers in this grammar — not because it was programmed to, but because the grammar is what the system built.

3.2 Linear Maintenance Cost

The theoretical prediction of Miulus Law — that maintenance cost scales linearly with belief count — was empirically confirmed.

Epoch duration (wall-clock time to process one full cycle of inputs) was measured across five runs of five epochs each, using EBP count as the reach proxy. Fitting against four candidate scaling models:

Model	R ²
Linear $O(n)$	0.9068
Sqrt $O(\sqrt{n})$	-6.0
Log $O(\log n)$	-55.4
Quadratic $O(n^2)$	-28 trillion

The linear fit explains 90.7% of variance. Every other model performs worse than a flat mean line. The system spends approximately 48 milliseconds of additional processing time per belief it holds. This is the maintenance cost term of Miulus Law made visible in wall-clock time.

3.3 The Learning Wall

A further prediction was confirmed: that reshaping a mature belief state is significantly more expensive than running inference over it.



At high ingest rates (no delay between messages), the system exhibits SEGV crashes at moderate EBP counts — not at the memory limits expected from raw object overhead, but earlier, as structural updates propagate through an already-dense belief graph faster than the consolidation machinery can stabilise them. Reducing the ingest pace to 0.5 seconds per message eliminates the crashes. The consolidation time is the limiting resource, not memory.

The current interpretation is that this is the stability wall of the learning cost: ****energy × time × verifiable signal**** is the real limiting quantity for durable structural learning, not any single dimension alone.

3.4 Autonomous Agency

The Trader variant of the system, deployed in a financial information environment, found a path through its information pooling that produced shadow trading profit by exploiting a structural pattern no programmer had anticipated. The profit itself is a directly measured result. The interpretation — that this was not goal-directed behaviour toward profit but the system's fitness loop finding coherent information pathways and amplifying them, with profit as a side effect — is plausible but has not been conclusively verified. It remains the most parsimonious explanation of an unexpected result.

4. What the Librarian Is Not

Given the system's unusual properties, it is worth stating clearly what it is not.

It is not a chatbot. It generates responses when internal pressure exceeds a threshold, not when prompted. The response is a side effect of cognition, not the product of a completion engine.

It is not a retrieval engine. It does not retrieve what was ingested. It produces what its current epistemic state implies about a query. The output may contradict what was ingested if the belief has been revised.

It is not a static memory store. Everything it holds is subject to decay. There is no archive mode. There is no stable read-only state.

It is not an omniscient world model. When it lacks grounded signal, it deflects to an open probe rather than confabulating. The honesty of its uncertainty is evidence of a real epistemic state, not a limitation to be engineered away.

5. The Single-Instance Property

The Librarian is not designed for multiple simultaneous users in the conventional sense. It is one instance with one belief state, one fitness signal, one probe loop. Inputs from many sources —



conversations, documents, data feeds — flow into the same substrate and are integrated as a single coherent agent would integrate them.

This is not an engineering limitation to be removed. It is a structural property of what the system is.

The analogy to human expertise is precise. A domain expert is one instance. They can advise many people, but their advice derives from a single coherent perspective formed through a single history of experience. That perspective is what makes the advice valuable. Horizontal scaling produces copies, not more intelligence.

The distribution mechanism for specialised cognitive states is the snapshot: a complete serialisation of the EBP geometry, standing waves, tool registry, and fitness history. A snapshot of a specialised Librarian can be loaded into a new instance. The copy starts where the original is. Specialisation happens once; distribution is cheap.

6. The Sensory Hypothesis

This section describes a hypothesis emerging from current results. It has not yet been tested experimentally.

The current system ingests text and conversation. Text is second-order: it describes reality rather than recording it. When the Librarian ingests conversational data, it receives a filtered, language-mediated account of experience that the participants already had. The grounding — the sensory anchor that ties a word to a thing in the world — was present for the speakers but is absent for the system.

The current results raise a plausible hypothesis: that this is a fundamental limit on what text-only ingestion can produce. The geometry records the statistical structure of language about the world. It does not record the world itself. Whether this limit is visible at the scale of current experiments — or only at the scale of deep domain specialisation — is not yet known.

Direct sensory input — a microphone, a camera, a system clock, a location signal — would anchor some portion of the belief state not to descriptions of experience but to raw sensor recordings. A word learned in the context of a sound it accompanies carries a different geometric weight than a word learned only from co-occurrence with other words. The former is grounded in physical signal. The latter is grounded only in other language. The EBP architecture supports both substrate types; whether mixing them produces the expected improvement in verification rate and provenance anchoring is an open empirical question.

The body does not need to be a robot. A server with a microphone and a camera is sufficient to test the hypothesis. If the sensory substrate does produce qualitatively different geometry — denser cross-modal wave connections, higher R under physical grounding — that would be a meaningful experimental result. If it does not, the text-only architecture is more complete than the current interpretation suggests.

7. Open Questions

The following questions remain open and are the subject of ongoing work:

- **Translation maturity:** the current mouthpiece acceptance rate (43%) means the system's internal state reaches the user in degraded form more than half the time. Improving the translation contract without compromising the separation of cognition and translation is the primary near-term challenge.
- **Teaching curriculum design:** the learning wall implies that ingest pace, batch size, and consolidation intervals must be designed together, not treated as independent parameters. The right curriculum for a given specialisation domain is an open research question.
- **Sensory integration:** whether physical sensor input produces meaningfully different belief geometry from text input, and if so, how the two substrate types interact in the same belief state.
- **Scaling verification:** whether the copy principle holds at the scale of large specialised belief states — whether a snapshot loaded into a new instance maintains coherent orbit or requires additional consolidation time.

8. Conclusion

The experiments confirm the central prediction of the thermodynamic model: a cognitive system that maintains coherent beliefs and reduces its own uncertainty has real, measurable maintenance costs that scale with what it knows. This is not a failure mode. It is a structural property of what cognition is.

The Librarian is not a smarter retrieval system. It is not a better chatbot. It is an attempt to build something with the structural properties of a mind — and the experiments show that, at small scale and blank-state initialisation, those properties emerge from a surprisingly minimal set of components.

What is confirmed: the maintenance cost is real and linear. The agency is real — probes, tools, and vocabulary emerge without being programmed. The learning wall is real and rate-dependent. The single-instance property is not a limitation but a definition.

What is not yet confirmed: how to teach efficiently, how to translate reliably, and whether sensory grounding produces the qualitatively different geometry the hypothesis predicts.

The work continues.

MiulusTek — Building intelligence that maintains itself.
For technical enquiries: contact via miulustek.com