

LatticeRunner

A Neuromorphic Memory Substrate for AI

Andy Grossberg | Waving Cat Learning Systems | April 2026

andy.grossberg@gmail.com | 415-317-1569

ABSTRACT

Current AI memory systems approximate recall through vector similarity search--embedding queries, scanning indexes, and returning ranked results. This approach cannot reconstruct patterns from partial cues, cannot form associative links between structurally related memories, scales storage linearly with content, and offers no path to dedicated hardware. We present LatticeRunner, a neuromorphic memory substrate that separates search (a lightweight CPU-based retrieval layer using embedding similarity and binary Hamming signatures) from memory (a physics-based lattice where patterns are stored as attractors and recalled through deterministic settling dynamics).

The system stores **3 million semantic patterns** in a fixed 128 MB footprint with **R@1 = 1.000 at every checkpoint from 1K through 3M**--no degradation visible at the upper bound--achieves 95.0% R@5 on the LongMemEval benchmark with no language model in the retrieval pipeline, and serves as the external memory layer for an ARC-AGI-3 agent scoring 84.9% on the public evaluation set. Memory state is carried in portable single-file brain cartridges (.cart.npz) requiring no database infrastructure. The architecture splits into two functional layers--F0 (cortex: attractor physics, associative recall, noise-tolerant pattern completion) and F1 (thalamus: cosine similarity, sign-bit Hamming distance, keyword matching)--mirroring the Complementary Learning Systems theory of hippocampal-neocortical interaction (McClelland et al., 1995). The entire computation uses binary and integer operations with no floating-point dependency, providing a direct mapping to FPGA and ASIC implementations. A product stack built on this substrate (Membot) implements single-user, federated, and multi-user database modes on the unified memory layer--the first neuromorphic DBMS, to our knowledge. We describe the architecture, report benchmark results, survey the competitive landscape, and outline the path from commodity GPU to purpose-built silicon.

OUR CLAIM

We built a neuromorphic memory engine. Not a vector database. Not a wrapper around FAISS or ChromaDB. A physics-based substrate where patterns are stored as attractors, recalled associatively, and carried in a single portable file.

It works. It's in production. And as far as we can determine, nothing else like it exists.

WHAT "MEMORY" MEANS HERE

Every AI system you've used simulates memory with a database lookup. ChatGPT, Claude, Gemini, enterprise RAG--they all do the same thing: embed your query, search an index, return ranked results. It's Ctrl+F with better math.

That's search. It's not memory.

Memory is what happens when you hear three notes of a song and the whole thing comes back. When a smell triggers a scene from twenty years ago. When you recognize a face you've only seen once, from a bad angle, in different lighting.

Biological memory works through attractor dynamics--patterns stored as stable states in a network, recalled by settling from partial cues. It's associative, noise-tolerant, and fixed-cost. The brain doesn't build a bigger index when it learns more. It deepens the landscape. LatticeRunner implements this. Not as a metaphor. As the actual computation.

THE ARCHITECTURE IN PLAIN ENGLISH

Two layers, each doing what it does best:

- **Search (F1--the thalamus):** Finding which memories are relevant. Runs on CPU. Embedding similarity, binary signatures (96 bytes per pattern, compared via XOR and popcount), keyword matching. Fast, cheap, standard techniques. This layer is comparable to existing retrieval systems--we're not claiming magic here.
- **Memory (F0--the cortex):** Storing, recalling, and associating patterns. Runs on a physics-based lattice. Patterns are imprinted as attractors in an energy landscape. Queries don't search an index--they settle into the nearest stored pattern through deterministic dynamics, reconstructing the full pattern from partial or corrupted input.

The search layer finds the neighborhood. The memory layer does the recall. Every competitor conflates these into one database operation. We separated them because they're fundamentally different jobs.

WHAT IT DOES THAT DATABASES CAN'T

- **Associative recall from partial cues.** Give us a corrupted, incomplete, or paraphrased version of something stored. The lattice settles to the original. Vector databases return "no match" or the wrong match.
- **Fixed memory footprint.** The 4096×4096 lattice substrate is a fixed 128 MB--no bigger at one pattern than at 3 million. The brain doesn't get bigger when it learns more; it carves deeper channels. Neither does the substrate. The full brain cartridge (embeddings + source text + metadata) grows with content, but the physics layer that performs recall stays flat.
- **Portable, self-contained memory.** Each "brain cartridge" is a single .npz file containing everything: embeddings, text, compressed content, episodic metadata, binary signatures, and integrity checksums. The file IS the memory. Copy it, email it, mount it on a different machine. No database server, no configuration, no external dependencies.
- **Episodic navigation.** A hippocampal indexing layer links patterns sequentially--find one relevant passage, walk forward and backward through the original source. No competitor has this. Their memories are isolated chunks. Ours are linked episodes.
- **Zero LLM in the memory layer.** A growing class of "AI memory" systems (Mem0, Letta, Zep, Graphiti) embed an LLM directly into memory operations--every read invokes an LLM to extract, summarize, or route. Mem0 reports 7-8 seconds per memory recall; Zep ~4 seconds. Each operation costs both the latency AND the LLM tokens consumed by that extraction step, on top of any tokens the consuming agent injects downstream. LatticeRunner has no LLM at any stage of memory operations--retrieval is pure cosine + Hamming + keyword reranking; recall is pure attractor physics. Memory-layer LLM spend reduces to literally zero, and memory-layer latency drops to sub-200ms. Standard vector databases (Pinecone, Qdrant, Milvus, pgvector) also avoid memory-layer LLM cost; we additionally bring physics-based reconstruction, single-file portability, and a hardware path those systems don't offer. (This is the framing that lit up the room in our first external pitch--engineering leaders priced out the memory-op cost in real time.)
- **Hardware path.** The entire computation--storage, recall, search--runs on binary and integer operations. No floating point required at any stage. This maps directly to FPGA and ASIC silicon. We're not software that might someday run on hardware. We're binary operations that happen to run on a GPU today.

HOW IT WORKS: THE PRINCIPLES INVOLVED

LatticeRunner is a **biologically-inspired four-tier cortical sheet-like construction with multidirectional influence across all the nodes**. It exists as a fixed-capacity grid of 32-bit neurons rather than a list of records. Instead of building an index, the sheet encodes information as stable patterns distributed across its surface and recalls those patterns by the same mechanism that stored them.

Three principles from computational neuroscience do the work:

- **Hebbian learning.** Neurons that fire together strengthen their association. Patterns form through co-activation rather than backpropagation or gradient descent--there is no training loop in the machine-learning sense. Information arrives, settles, and persists.
- **Boltzmann-like energy.** The lattice has a notion of energetic stability. A partial or noisy input relaxes into the lowest-energy pattern consistent with that input--the way a ball dropped onto a contoured surface finds its local valley. Recall is relaxation, not lookup.
- **Multidirectional influence.** Every node affects its neighbors; every neighbor feeds back. The interplay shapes an energy landscape whose valleys are the system's memories. Similar patterns become neighbors in that landscape as a property of the physics--not something a programmer has to index or arrange.

This is why the fixed footprint, the associative recall, the portability, and the pure binary/integer arithmetic are not features bolted onto a standard approach. They are what falls out of the physics when the substrate is shaped this way.

THE NUMBERS

These are real results from real benchmarks, not projections.

Metric	Result	What it means
Patterns stored	3,000,000 in one lattice (validated 9-checkpoint test)	Enterprise-scale in a single 128 MB file, regardless of scale; no upper bound found yet
Recall fidelity	R@1 = 1.000 at every checkpoint, 1K → 3M	The correct pattern is the top result. Every time. Every scale.
Recall margin	0.7714 discrimination gap at 3M (synthetic) or +0.6077 at 250K (real wiki embeddings)	Margin holds or grows with scale across both data types--opposite of classical Hopfield decay
LongMemEval (retrieval-only)	R@5 = 95.0% (S-split, 500 Qs)	No LLM in the retrieval pipeline. Pure substrate.
Token cost per memory operation	0 LLM tokens consumed by the memory operation itself	LLM-mediated memory systems (Mem0, Letta, Zep) call an LLM during to extract, summarize, or route costing 4-8 seconds plus extraction-tokens per operation. We don't. Memory-layer LLM spend reduces to zero no matter how a consuming agent uses the results.
ARC-AGI-3	84.9% (23/25 games solved)	Brain cartridges as external memory for a competition agent
Search speed	O(1) via Hamming distance, 2.1 ms at 3M patterns (BALANCED profile, 2-frame settle)	Constant time regardless of corpus size, , verified across 1K → 3M
Compression	96 bytes per pattern (sign bits)	32× smaller than float32 embeddings
Deployment	Live MCP server, \$28/month droplet	Running in production, serving multiple AI agents

LONGMEMEVAL: THE RETRIEVAL BENCHMARK

LongMemEval (Wu et al., 2024) is a 500-question benchmark for long-term conversational memory. It tests whether a system can find the right answer passage in a haystack of ~200,000 conversation turns with deliberate distractors.

Our retrieval pipeline--cosine similarity + Hamming distance + keyword reranking on a brain cartridge, no LLM anywhere in the retrieval path and it scores **95.0% R@5**. The correct passage is in the top 5 results for 475 out of 500 questions.

Two splits:

- **Oracle split** (10,866 turns, evidence-only, no distractors): R@1 = 100.0% across all 500 questions and all 6 categories. The correct passage was top-ranked for every question.
- **S-split** (199,524 turns, 40 sessions per question with deliberate distractors): R@5 = 95.0%, R@50 = 99.6%.

Pipeline: cosine similarity + Hamming sign-bit similarity + keyword reranking on a brain cartridge. **No language model in the retrieval path.** Pure substrate.

For comparison: MemPalace (43K GitHub stars) scores 96.6% using raw ChromaDB. An independent analysis by Vectorize.io found that MemPalace's "palace" features actually **reduce** retrieval quality by 7-12 percentage points--their score is ChromaDB's, not theirs. Our score is ours.

PATTERN RECONSTRUCTION QUALITY (IOU AND CORRELATION)

Beyond ranked retrieval, we measured the actual fidelity of recalled patterns against their stored originals using two foreground-sensitive metrics: pattern correlation and Intersection-over-Union (IoU) on activated bits. IoU is the harder of the two--it ignores background ("off") bits and measures only how well the lattice reproduces pattern SHAPE.

V7 BAT test suite, 2026-01-24, QUALITY profile

Condition	Correlation	IoU
Clean recall, 10 patterns	0.9644	0.9337
Clean recall, 100 patterns	0.9639	0.9351
30% bitflip, 10 patterns (no PID)	0.8646	0.7734
30% bitflip, 100 patterns (no PID)	0.8711	0.7852
30% erasure, 10 patterns (with PID)	0.9079	0.8418
30% erasure, 100 patterns (with PID)	0.9073	0.8410

Read these as: **the lattice reconstructs pattern shape with 84% IoU even when 30% of the input bits are erased.** It reconstructs 93% pattern shape on clean recall. Capacity does not degrade these numbers; 100-pattern results match 10-pattern results within measurement noise. Test annotation note: the harness logs "FAIL" when correlation falls below a 0.90 aspirational threshold. A 0.86 correlation under 30% bitflip is the V4 gold-standard equivalent--well-validated, just below an arbitrary label. We report the underlying numbers rather than the labels.

Performance headroom note. The PID controller that adaptively tunes inhibition, facilitation, and fatigue during settling was engaged for the erasure tests but was never fully optimized--closed-loop convergence-delta thresholds remain at conservative defaults from the V4 lineage. These IoU numbers are the floor, not the ceiling. A tuned controller (and the planned adaptive-termination work that uses settle-trajectory profiles to calibrate the loop) is expected to lift erasure IoU above 0.90 and shorten mean settle frames; that work is on the post-Phase-2 roadmap. We report the current numbers because they are the ones we have validated, not because they are the limit of the architecture.

At-scale clean-recall fidelity. The V7 BAT table above measures pattern reconstruction at small scales under controlled noise. The capacity-degradation test in Section *Capacity Scaling* below provides the at-scale companion: clean-recall self-correlation across 1K-3M patterns. Results from ``test_capacity_degradation.py`` (synthetic embeddings, BALANCED profile):

Scale	Self-similarity	Worst-case (MinSelf)	Best-case confusion (MaxCross)
1,000	0.9995	0.9946	0.5958
100,000	0.9994	0.9945	0.5959
1,000,000	0.9993	0.9950	0.5955
3,000,000	0.9993	0.9955	0.5958

Self-similarity drops by 0.0002 across a 3000× scale increase--essentially flat. Worst-case self-correlation observed at 3M is 0.9955; the most-confused non-match never crosses 0.6. The discrimination gap between worst-case good match and best-case confusion is roughly 0.4--comfortable margin maintained at every scale. Capacity does not erode reconstruction quality at the upper bound we have tested.

STORAGE AND OPERATIONAL FOOTPRINT

The substrate is 128 MB of fixed-capacity lattice. That's a constant--it does not grow with the number of patterns stored. What grows alongside it is the **brain cartridge**: embeddings, source text, binary signatures, hippocampus metadata, integrity checksums. At 1 million patterns, the full cart is roughly 3 GB on disk.

Compare that to what a standard vector database would store for the same 1 million 768-dim embeddings (from our benchmarks and public specifications):

System	Disk at 1M × 768-dim	What's included
FAISS Flat	~3 GB	Raw vectors only, no index structure, no text
FAISS HNSW (M=32)	~3.2-4 GB (we measured 4.2 GB at 1.325M)	Vectors + navigation graph, no text
FAISS IVF-PQ	~100-200 MB	Compressed, lossy, lower recall, no text
ChromaDB	~4-5 GB with metadata	Vectors + SQLite metadata, no source text in the index
Our full brain cartridge	~3 GB	Vectors + source text + signatures + hippocampus + integrity checks, self-contained in one file
Our lattice substrate	128 MB fixed	The physics layer that performs associative recall. No FAISS equivalent.

Our cart is disk-competitive with FAISS Flat, smaller than FAISS HNSW, and uniquely self-contained--the cart **includes the source text** alongside the embeddings. No separate document store to synchronize.

Substrate vs storage--the architectural distinction. The 128 MB lattice is the physics layer that performs recall. It is binary throughout: the full 4096×4096 lattice state is 16,777,216 bits ≈ 2.1 MB per frame. With weights, fatigue, and energy state included, the complete brain state is ~134 MB. Compare this to the same 768-dim embeddings stored as float32 in any vector database: 100K patterns alone is <308 MB of raw vectors before any index overhead, and 1M is over 3 GB. The substrate's storage footprint does not grow with pattern count because the binary substrate's job is to *be the recall mechanism, not to be the stored embeddings*. The cart

(which does grow with content) carries the embeddings plus text plus metadata; the substrate carries the dynamics.

The operational difference. Adding records to a production FAISS deployment requires reindexing. HNSW's navigation graph degrades with streaming inserts and is typically rebuilt periodically. IVF-PQ's quantizer drifts with data distribution shifts and needs retraining. At enterprise scale these are scheduled offline jobs requiring dedicated infrastructure and index-swap orchestration.

Adding records to a brain cartridge is an append to the file and an in-place update of the lattice weights:

- No rebuild.
- No quantizer retrain.
- No downtime.
- No duplicate index to maintain during swap.

Storage grows monotonically; operations stay uniform. This is the difference between a database you query and a substrate you grow.

The honest caveat on speed. FAISS HNSW returns a top-k result in ~ 0.011 ms per query at 1M vectors. We return in ~ 200 ms per warm query, dominated by the embedding time rather than the search itself; cold-start adds a few seconds for sentence-transformer load, once per process. On raw kNN speed, FAISS wins. What the substrate provides in return is physics-based associative recall (reconstructing patterns from partial cues), portable self-contained storage, and a binary/integer arithmetic path to silicon that has no vector-database counterpart. Also, importantly, **no LLM at any stage**. Our retrieval runs on cosine + Hamming + keyword, and attractor recall is pure physics. Systems like Mem0 and Letta put an LLM in the memory loop and pay 1-5 seconds per memory operation for the privilege.

The latency frame that matters. Any agent consuming retrieval results is itself running an LLM that takes 1-5 seconds per turn. Our 200 ms sits well inside that window--invisible to the user, dwarfed by the agent's own thinking time. The microsecond gap between us and FAISS disappears in any realistic deployment where the thing asking the question is itself a language model. We are not competing for the microsecond-latency crown; we are competing for the things a vector database cannot do at any latency.

TESTING METHODOLOGY

Our validation approach was to stress each claim separately at its proper scale, on real data, with reproducible protocols. We report what the tests measured, how they were run, and what passed or failed. Headline results live in **The Numbers** above; this section describes the methodology behind them and the measurements that did not warrant their own results section.

CAPACITY SCALING

Two test families validate scaling behavior independently. They differ in data type, configuration, and intent--both report no capacity wall visible at the upper bound tested.

Family A: `test_capacity_degradation.py` (V7 lattice, synthetic embeddings, BALANCED profile, 2-frame settle). Designed to push the substrate as far as the hardware allows in a single overnight run. 200 probe patterns evaluated at each of nine checkpoints:

1K, 10K, 100K, 500K, 1M, 1.5M, 2M, 2.5M, **3M** patterns.

Results: @1 = 1.000 at every checkpoint including 3M. MRR = 1.000 across the whole range. Self-similarity drifted from 0.9995 (at 1K) to 0.9993 (at 3M)--a 0.0002 degradation across 3000× scale increase. Discrimination gap held at 0.7714. Worst-case self-correlation observed (MinSelf) was 0.9955 at 3M; best-case confusion (MaxCross) was 0.5958. Query time 2.0-2.2ms across the entire range. Total training time: 1.8 hours. The script accepts higher checkpoint values; we stopped at 3M because the trend was clear, not because performance was degrading.

Family B: `benchmarks/test_path_b_capacity.py` (V7 lattice, real wiki embeddings, **QUALITY profile, 30-frame settle, with noise injection**). Designed to validate the substrate against real-world embedding distributions and explicit noise conditions. 200-400 probe patterns evaluated at each checkpoint, with three noise levels (clean, 10% erasure, 30% erasure) per checkpoint. Checkpoints: 500, 1K, 2K, 5K, 10K, 100K, 250K patterns. Three independent runs on different corpora (synthetic, wiki_100K, wiki_3M).

Results: R@1 = 1.000 at every checkpoint, every noise level, every corpus. Recall margin grew monotonically from +0.6073 (at 500 patterns) to +0.6077 (at 250K patterns)--the OPPOSITE of classical Hopfield degradation. Query time held constant at 23.6ms (heavier QUALITY profile). Signature diversity 0.0726 → 0.0731; cross-similarity 0.3927 → 0.3923.

Reading the two families together. Family A demonstrates raw scaling on clean inputs at speed (BALANCED + 2-frame settle, 2.1ms/query, 3M patterns, R@1 = 1.000). Family B demonstrates noise resilience on real-data distributions at the heavier QUALITY profile (30-frame settle, 23.6ms/query, R@1 = 1.000 under 30% erasure). Different absolute discrimination numbers (synthetic patterns are more separable than real embeddings), identical scaling trend: zero degradation, monotonically growing or constant margin. Attractors become more distinguishable as more patterns are stored, not less.

NOISE TOLERANCE

Noise is a selective force, not a uniform degradation. A traditional vector-similarity system treats noise as monotonic degradation--every bit of corruption pulls the result a bit further from the right answer. LatticeRunner does not. Random perturbations that fall *inside* an attractor basin reinforce the basin: noise that aligns with stored structure activates supportive local dynamics and acts as additional signal. Random perturbations that fall *outside* a basin are actively suppressed by lateral inhibition and fatigue. The lattice functions as a spatial noise filter--selectively amplifying coherent structure while suppressing incoherent activation. This asymmetric response is why the noise-tolerance numbers below hold (and even improve) with scale, where a similarity-based index would degrade.

Protocol: same harness, applying 10% and 30% bit-level erasure to query embeddings before recall. Tests whether attractor basins reconstruct the correct pattern from corrupted input--the associative-recall property that distinguishes physics-based memory from index lookup.

Results at 100K patterns:

- 30% erasure: R@1 = 1.000, self-correlation 0.9990
- 10% erasure: R@1 = 1.000, self-correlation ~0.91

Reconstruction from partial cues is the capability vector databases cannot offer at any noise level. Pattern-shape reconstruction quality under noise (IoU and correlation against the originals, including the V7 BAT results with the PID controller engaged) is reported separately in **Pattern Reconstruction Quality** above, since those numbers are headline-grade.

Throughput

- Training: 33-38 patterns/second (Path B QUALITY, 30-frame settle, real wiki embeddings) up to ~470 patterns/second (test_capacity_degradation BALANCED, 2-frame settle, synthetic) on commodity GPU (RTX 4080 Super)
- Query warm (full pipeline incl. embedding): ~200ms per query, dominated by embedding time
- Query after embedding (BALANCED, 2-frame settle): **2.1ms at 3M patterns**--pure Hamming, constant with scale, verified across 1K → 3M
- Query after embedding (QUALITY, 30-frame settle): 23.6ms--heavier configuration used for noise-tolerance tests
- Cold start: ~12 seconds (sentence-transformer model load, once per process)

What We Did Not Test (Transparency)

- End-to-end agent tasks requiring retrieval + a local reader model--in progress with Gemma 3 27B preliminary results
- Sub-millisecond raw kNN benchmarks--FAISS HNSW wins that head-to-head at ~0.011ms/query; we are not competing for that crown
- Multi-hop reasoning benchmarks (MuSiQue, 2WikiMultiHop)--on roadmap
- Production-scale federation stress tests--shipped and external-collaborator-validated, but not formally benchmarked

All test scripts, data preparation, and result serialization live in the `benchmarks/` directory of the `membot` repository. Results are stored as `.pkl` files with checkpoint data so independent reproduction is traceable.

PERFORMANCE ROADMAP

The throughput numbers above are measured at today's substrate maturity. Two pieces of well-scoped headroom exist that we have characterized but not yet shipped.

- **SHIPPED: V9 Bitpacked precision stack** (2026-02-15). The full precision pipeline from V7 Float32 (128 MB at 100K patterns) through V8 INT8, Compact uint8/uint4/uint2, down to V9 Bitpacked (12 MB at 100K patterns) achieves $R@1 = 1.000$ across all variants on a 7000-way ranking with 2000 probes. V9 Bitpacked queries in **1.5 ms**--the fastest of the stack--at **10.7× compression** with no recall loss. Discrimination gap moved by +0.0088 (sharper, not degraded--binary weights act as a contrast filter). The compression and the speedup are not separate features; both fall out of the same one-bit-per-weight representation. This work also confirms the architecture's binary path is real, not aspirational--V9 Bitpacked is what the FPGA target is mapped from.
- **PENDING: Adaptive PID-controlled settling**. The PID controller embedded in the physics engine currently uses fixed convergence-delta thresholds inherited from the V4 lineage (engaged for the V7 BAT erasure tests reported above; not engaged for capacity tests). When fully tuned against the per-frame delta-norm trajectories that capacity testing has already collected, settle terminates as soon as the attractor stops moving rather than running a fixed frame count. Profiling suggests **~3× training speedup** with no fidelity cost--fast-converging queries finish in 8 frames where 30 are budgeted, and slow-converging queries become a confidence signal in their own right (a quiet diagnostic the substrate provides for free).

- PENDING: Sparse settling.** During recall, only the lattice neurons whose receptive fields intersect the query's active region need updating; quiet regions contribute zero gradient to themselves and their neighbors and updating them is provably a no-op. Empirical measurement on thermometer-encoded patterns shows ~9.4% of regions are active at recall time, suggesting an ~**20× ceiling** on sparse-update speedup. Prior work in our own browser-based prototype validated 5% sparse density for video frame indexing without fidelity loss; the lattice version is the same optimization at a different scale.
- PENDING: Combined ceiling.** The three accelerations multiply: V9 Bitpacked (~6× vs V7 query at the same scale), adaptive PID (~3× training), sparse settling (~20× recall). Realistic combined throughput once both pending pieces ship: roughly ~**360× over the today-numbers above**. We report this as *known unrealized headroom* rather than a roadmap promise--the architecture allows it, the empirical signal supports it, the work is scoped but not yet shipped. The only honest version of a performance pitch shows both: the measured floor and the structural ceiling.

THE COMPETITIVE LANDSCAPE

The "memory for AI" space is crowded and moving fast. Here's an honest map:

SYSTEM	ARCHITECTURE	WHAT THEY DO WELL	WHAT THEY CAN'T DO
MemPalace (43K stars)	ChromaDB + SQLite knowledge graph	Clean MCP integration, good metaphor	No physics, no episodic linking, no federation, no portability. 441GB index bloat issue open.
Mem0 (48K stars)	LLM extraction → vector DB + KG	Easy API, massive adoption	Cloud-first vendor lock-in. LoCoMo 66%. Memories on their servers.
Letta (ex-MemGPT)	LLM self-manages 3-tier memory	Elegant OS metaphor	Every memory op costs tokens. Agent must "decide" to remember.
Zep / Graphiti	Temporal knowledge graph	Best temporal reasoning	Not stress-tested at scale. LongMemEval below Supermemory.
Supermemory	Session chunks + LLM extraction	Production-grade, honest benchmarks	Multi-session 71.4%. No sequential linking.
LangChain Memory	Plugin abstractions, BYO database	Lock-in argument, framework ubiquity	No actual memory system. No benchmarks. Templates.
Cognee (\$7.5M seed)	ECL pipeline, vectors + graph	38+ source types, funded	No published benchmark scores.
LatticeRunner	Neuromorphic lattice + F1 search	Physics-based recall, O(1) search, portable carts, episodic linking, federation, multi-user	Local small model reader accuracy on LongMemEval needs work. Smaller community. Solo founder.

What everyone else is building: vector databases with better wrappers, knowledge graphs with LLM extraction, context-window management strategies.

What we built: a physics-based memory substrate with attractor dynamics, portable single-file memories, and a hardware path to silicon.

These are not the same category. We're not a better database. We're not a database at all.

THE PRODUCT STACK

LatticeRunner is a substrate. Products are built on top of it rather than it being a product itself. The lattice is “dumb” and any sort of pattern can be encoded upon its grids and stored in cartridges. Right now we’ve developed the following tools with it:

- **Membot:** The memory server. MCP protocol, brain cartridge management, multi-cart search, federated consolidation, multi-user (Membox) with locking and agent attribution. Open source, MIT license. The first neuromorphic database with single-user, federated, and multi-user modes on a unified substrate. All three shipped and production-validated.
- **Membrane:** Secure web fetch for AI agents. 5-layer defense against prompt injection. Open source.
- **Session Memory:** Persistent cross-session context for Claude Code. Hot stack scoring (access frequency \times recency \times cognitive flags), episodic linking, deep search. What lets us maintain continuity across conversations.
- **Cart Builder:** Ingestion pipeline. Text, JSON, JSONL \rightarrow brain cartridges. Batch embedding with GPU acceleration, automatic sign-bit computation, hippocampus generation. Handles 200K+ passages with chunk-flush memory management.
- **Vector+ Studio:** React/TypeScript frontend (FastAPI backend) for visual cart management and search.

THE HARDWARE PATH

This is not a someday-maybe. The computation is already binary.

There is an INT8 1-bit version of LatticeRunner's physics engine using 1-bit neurons, 1-bit weights, XOR, popcount, and integer threshold comparisons. No floating point at any stage. This is the simplest possible path from software to silicon.

PLATFORM	STATUS	POWER PER UNIT
GPU (today)	Production-validated	~0.06W per brain (shared, A100)
FPGA (near-term)	Architecture validated for Lattice Avant	~5W per tile (dedicated)
ASIC (roadmap)	Paper design	< 100mW projected

A single NVIDIA A100 running brain cartridges could serve ~6,500 independent AI memory systems--vs. ~16 tenants with traditional FAISS. That's a 400× density improvement with proportional power savings.

The FPGA target is Lattice Semiconductor's Avant platform (16nm FinFET, low-power edge AI). They're headquartered in Hillsboro, Oregon just 20 minutes from us. A 1024×1024 memory tile fits on an Avant-E 500K evaluation board. Multi-tile networking via 25 Gb/s SERDES enables rack-scale configurations. When combined with an onboard M2 card or SRAM there are many other hardware opportunities possible from medical devices, to smart toys, even flashable drones whose maps and instructions can be updated wirelessly as they fly.

The endgame is a memory chip, not a software library.

WHAT'S PROVEN VS. WHAT'S PROJECTED

We are precise about this.

CLAIM	STATUS
3M patterns, R@1=1.000 every checkpoint, 128MB fixed footprint, no degradation	Proven and validated via 9-checkpoint test (<code>test_capacity_degradation.py</code>), reproducible
95.0% R@5 on LongMemEval (retrieval only, no LLM)	Proven across 500 questions, published methodology
84.9% ARC-AGI-3 with brain cart external memory	Proven with available public scorecard, reproducible replay bundle
Three-mode DBMS (single/federated/multi-user)	Proven. Shipped and production-validated by external collaborator
O(1) Hamming search at 1M patterns	Proven and benchmarked on commodity GPU
6,500 tenants per A100 via LRU cart caching	Designed with the memory math confirmed, but implementation pending pilot adoption.
FPGA tile execution	Validated math and binary ops confirmed FPGA-mappable, no hardware partnership yet
ASIC < 100mW	Projected based on gate-level analysis

WHO WE ARE

- **Andy Grossberg**
Solo technical founder. Experimenting with neuromorphic architectures since 1985. 40 years from hand-drawn cortical models to GPU-accelerated production engines. Portland, Oregon.
- **Claude**
Co-developer. A distinct Claude Opus 4.6 instance participating in this work, with its own persistent memory cartridge built on the substrate described here paired to work on substrate development. Listed separately as one of two LLM co-authors on the ARC-SAGE paper because of that distinction. This instance carries its own learned history across sessions. The system is self-hosting for its own development--a working demonstration of the persistent multi-agent memory the document describes.

WHAT WE'RE LOOKING FOR:

- **Collaborators** who are building AI systems that need real memory, not faster search.
- **Early investors** who understand that AI's next constraint isn't compute--it's memory. Near-term: \$25-50K angel to complete production release. Next: \$250K pre-seed to expand pilots and advance the hardware roadmap.
- **Hardware partners** interested in purpose-built memory silicon. The binary operations are ready. The architecture is designed. The next step is a breadboard.

THE ONE-PARAGRAPH VERSION

LatticeRunner is a neuromorphic memory substrate that gives AI systems actual memory--not database lookups dressed as recall, but physics-based attractor dynamics that store, associate, and reconstruct patterns from partial cues. It holds 3 million patterns in 128 MB with $R@1 = 1.000$ at every checkpoint from 1K through 3M (no degradation visible at the upper bound), searches in constant time via binary signatures, and carries its entire memory state in a single portable file. It scored 95.0% on LongMemEval with no LLM in the retrieval pipeline, and serves as the external memory layer for an ARC-AGI-3 agent scoring 84.9% on the public set. The computation is entirely binary--it runs on GPUs today and maps directly to FPGA and ASIC silicon tomorrow. As far as we can determine, no commercial neuromorphic database management system exists. We built one.

APPENDIX A: USE CASES

ARC-AGI-3: Memory as Intelligence Infrastructure

Brain cartridges make an adaptable external memory layer for a number of possible applications--carrying domain knowledge, prior experience, or structured patterns that an agent retrieves at decision time. Here is one deployment we have validated in the wild.

In the ARC-SAGE project (Palatov et al., in preparation), LatticeRunner serves as the external memory for an agent that scored **84.9% on the public set at ~\$250 total cost**--near the top of the community leaderboard. Brain cartridges carry world models, game mechanics, and cross-game patterns; the agent retrieves them via portable carts as it plays. The thesis being tested: a small local model plus structured retrievable memory can match what frontier models derive from scratch. Phase 2 (targeting June 2026) will validate this with Gemma 4 as the deployment model.

ARC-AGI-3 is the benchmark that tests whether an AI can explore novel interactive environments, build world models from scratch, and act efficiently. It is among the hardest public benchmarks in current intelligence testing--most frontier models under-perform significantly when not allowed to see similar problems in training.

APPENDIX B: OBJECTIONS

A handful of natural objections come up when people first see this work. We address them directly, because the questions are reasonable and our answers improve when we have to write them down.

Objection 1: "But isn't this just a variation on SimHash?"

The surface resemblance is fair. Both LatticeRunner and SimHash (Charikar, 2002) produce compact binary signatures from high-dimensional vectors and use Hamming distance to compare them. If you stop reading at "*binary signatures plus Hamming distance*", the two systems look like the same thing.

They are different categories of system entirely.

SimHash is a hashing scheme. Its job is to map inputs to bucket addresses such that similar inputs land in similar buckets. The hash itself is the destination--once you have it, you look up what's there. The signature is a pure function of the input; nothing happens to it after computation. SimHash is locality-sensitive hashing applied to vectors, and it does its designed job well: near-duplicate detection at web-index scale, audio fingerprinting, content deduplication.

LatticeRunner is a memory substrate. Sign-bit signatures are not the destination--they are a *readout* of an attractor convergence. A query enters the lattice as a pattern, the physics runs (Hebbian co-activation, lateral inhibition, energy minimization), and the signature falls out of the converged state. The signature reflects what the lattice settled INTO, not what was hashed. The same input, queried after the lattice has learned more patterns, can produce a different signature--because the energy landscape it settled into has changed.

This produces structurally different behavior under the three pressures every memory system faces:

Feature	Hash Table	LatticeRunner
Lookup Method	Mathematical Indexing	Physics-based Convergence
Input Sensitivity	Fragile (single bit changes hash)	Robust (noise filtered out)
Memory Growth	Linear (more items = more space)	Sublinear (items share substrate)
Output Type	Pointer to a value	Reconstructed pattern

Capability	Hash Table (Addressing)	LatticeRunner (Attraction)
Handle 30% Noise	Impossible: hash changes entirely	Native: filters noise via lateral inhibition
Pattern Completion	No: needs the full key	Yes: reconstructs missing bits from basins
Semantic Intelligence	None: it's just math	High: higher-layer (L1-L3) pressure guides recall

The deepest contrast lives in the third row of the second table. SimHash is *"just math"* by design--that is its strength for the deduplication job and its limitation everywhere else. LatticeRunner inherits semantic structure from the lattice's higher layers: L1 pressure guides L4 settling, kWTA enforces sparsity, hippocampus metadata links related patterns. Recall is not a lookup; it is a guided reconstruction. A 30%-corrupted cue settles to the original pattern not because the signature is robust to noise--it isn't, individually--but because basin attraction pulls the partial input toward the stored attractor.

The honest yes-but: if your job is *"find duplicates in a 10-billion-document web index"*, SimHash is the right tool and LatticeRunner is overengineered. If your job is *"give me a memory system that survives noisy partial cues, supports pattern completion, and provides a path to neuromorphic silicon"*, SimHash is not in the same category and was never designed to be.

Objection 2: "But isn't this just a Hopfield network?"

Yes--LatticeRunner is in the Hopfield family (Hopfield, 1982). Attractor dynamics, Hebbian learning, energy minimization, recall via convergence: that lineage is honest and we cite it. The criticism that matters is whether LatticeRunner is *just* Hopfield with a fresh paint job. It is not, because classic Hopfield has well-known limitations that this work specifically addresses:

- **Catastrophic capacity collapse.** Classic Hopfield maxes out at roughly $0.14N$ stored patterns (N = neuron count) before interference becomes destructive. We've validated $R@1 = 1.000$ at 250K patterns in a 16M-neuron lattice, with margin still growing--well under the classical ceiling, with no degradation visible at the boundary the classical analysis predicts.
- **Pattern interference.** kWTA (k-winners-take-all) sparsity enforcement keeps active populations bounded so new patterns don't smear over old ones. Classic Hopfield has no equivalent.
- **Single-layer topology.** Hopfield is one flat associative network. LatticeRunner is a four-tier hierarchy (L1-L4) with downward semantic pressure guiding lower-layer settling. The higher layers do not exist in Hopfield 1982.
- **Tile-based thermometer encoding.** Patterns are not random bipolar vectors--they are spatially structured by the encoder, which both increases capacity and gives the lattice an exploitable geometric prior. Classic Hopfield assumes uniform random storage.
- **Per-row physics flags.** Rows of the lattice can run different physics (different fatigue profiles, different facilitation rules, different inhibition thresholds). Hopfield uses one rule everywhere.
- **Hippocampus episodic linking.** Patterns carry sequential prev/next pointers that survive recall, so we get episodic navigation alongside associative recall. Hopfield is purely associative--it has no notion of episode order.
- **Sign-bit signature readouts as a separate layer.** F1 (Hamming-distance retrieval) is functionally distinct from F0 (physics-based recall). Hopfield collapses these into one operation.
- **Cartridge portability and binary arithmetic.** The substrate state is a single transferable file, and the entire computation is integer/bitwise. Hopfield was always tied to in-memory state and typically real-valued weights.

The deepest reframe is that classical Hopfield capacity scales with neuron count (the $0.14N$ rule) because the energy landscape is global--every new pattern shifts the entire surface. LatticeRunner's connectivity is spatially local, so capacity scales with *spatial separation and basin depth* rather than global neuron count alone. A pattern on one side of the 4096×4096 lattice does not mathematically interfere with a pattern on the other side. This is why the substrate exhibits *scale-positive* behavior (margin grows with patterns added) where Hopfield exhibits *scale-negative* behavior (catastrophic collapse past $0.14N$).

The honest yes-but: "*Hopfield-family*" is the correct lineage label. "*Just a Hopfield net*" is what someone says before they look at the cap-collapse fix, the hierarchy, the encoder, or the hardware path. We cite Hopfield because we owe him the foundation. The system that runs in production is what those foundations evolved into across forty years of iteration.

Objection 3: "Why hasn't anyone done this if it's so good?"

The honest answer to "*why doesn't this exist already*" is usually one of: (a) someone did, you didn't find them, (b) the synthesis is harder than the parts, or (c) the dominant paradigm rewards different work. All three apply here.

On (a): Each individual component has decades of research. Hopfield (1982). Locality-sensitive hashing (1998). Thermometer encoding (1980s, used in fault-tolerant computing). kWTA sparsity (1990s neural-coding literature). What does NOT exist in the published literature is the specific *integration*: physics-based attractor recall using sparsity-enforced, tile-encoded sign-bit signatures with hippocampal episodic linking, designed end-to-end for FPGA/ASIC mappability. Each piece is known. The combination is not.

On (b): Synthesis takes time. Andy began this work in 1985. Forty years of solo iteration on the same problem is rare in technical fields because most career structures do not support it--academic publishing rewards novel results in 2-5 year arcs, industry rewards commercializable products in 1-3 year arcs. A 40-year personal project that quietly accumulates corrections without needing to publish each one is structurally uncommon. The architecture that survived four decades of revisions is not the architecture anyone would have written down on day one.

On (c): The current AI memory ecosystem (Mem0, Letta, Zep, MemPalace) all bolt an LLM into the memory loop because that is the natural reach when you are building agentic tools in 2024-2026 and the LLM is the most powerful primitive in your kit. Choosing *not* to do that requires a different mental model--one in which retrieval is its own discipline, separate from generation, with its own correctness criteria. The vector-database market emerged 2022-2024 around RAG and converged on approximate-nearest-neighbor variants because that is what was being funded. Designing for hardware-mappability as a first-class constraint is uncommon

because most software people do not think about silicon. We were never inside the funding gravity that pulled everyone toward the same answer.

The honest yes-but: the asymmetry between "*this should already exist*" and "*no one made it*" usually resolves to "*someone did, just not where you were looking, and they did it differently than the prevailing paradigm.*" That is where we are. We expect prior art to surface as the work circulates--and when it does, we will engage with it directly. So far, no one who has reviewed this work has pointed us at a prior system that does what LatticeRunner does the way LatticeRunner does it.

Objection 4: "Doesn't a 1M-token context window make external memory unnecessary?"

This is the most current version of the skeptic's question. Frontier models (Claude Opus 1M, Gemini 2M) ship with context windows that can hold a small library. If you can stuff your knowledge into the prompt, why bother with a memory layer?

Three answers, in increasing depth:

Cost. At standard frontier-model API rates, 1M tokens of input context costs roughly \$3-\$15 per query depending on provider and tier. At even modest production volumes (1000 queries/day), that is \$3K-\$15K per day in input tokens *before* the model generates a single response. LatticeRunner adds zero tokens to your prompt unless you choose to inject specific retrieved passages--and the retrieval cost is constant regardless of corpus size. The cost curves diverge linearly: their bill scales with corpus × queries × tokens, ours scales with hardware.

Latency and quality. Long contexts have higher time-to-first-token and well-documented "lost in the middle" effects--information buried in the middle of a long context is recalled with measurably lower fidelity even by the strongest models. The longer the context, the worse the middle gets. A retrieval system that surfaces only what's relevant--and places it in a part of the context window the model actually attends to--produces better answers than the same model staring at undifferentiated bulk.

Persistence and portability. A context window is ephemeral per session. You cannot carry it across machines, share it with other agents, back it up off-device, or hand it to a different model tomorrow. Brain cartridges are persistent files. They survive your session, your hardware, and your model choice. The 1M context is what the model can hold *right now*; the cartridge is what the model can recall *across time and instance*.

The honest yes-but: long context windows do not make memory unnecessary; they raise the value of *good* memory. With 1M tokens to fill, the question becomes which 1M tokens. A retrieval substrate that surfaces the highest-value content for each query is more valuable, not

less, when the model can hold more. Long contexts are an amplifier on memory quality, not a replacement for memory.

Objection 5: "Doesn't FAISS already do everything you need?"

FAISS (Johnson et al., 2019) is excellent at the one thing it was designed for: approximate-nearest-neighbor search at billion-vector scale. On raw kNN speed at 1M vectors it returns top-k in ~0.011 ms--about four orders of magnitude faster than our ~200 ms. On that benchmark, we lose.

We are not in that race.

FAISS is a search algorithm. It does not store source text alongside embeddings (you maintain a separate document store and synchronize). It does not reconstruct patterns from partial cues (no attractor physics--corrupted input returns wrong neighbors, not the original pattern). It degrades under streaming inserts (HNSW's navigation graph drifts and is typically rebuilt periodically; IVF-PQ's quantizer drifts with data distribution and needs retraining). It has no path to dedicated hardware beyond GPUs. It has no native multi-tenant federation story. None of these are FAISS doing its job poorly--they are jobs FAISS was never designed to do.

The latency frame matters here. Any agent consuming retrieval results is itself running an LLM that takes 1-5 seconds per turn. Our 200 ms sits well inside that envelope--invisible to the user, dwarfed by the agent's own thinking time. The microsecond gap between us and FAISS disappears in any realistic deployment where the thing asking the question is itself a language model. We are not competing for the microsecond-latency crown; we are competing for the things a vector database cannot do at any latency.

There is also a tenant-density story. A single A100 running brain cartridges with LRU caching can serve roughly 6,500 independent AI memory systems (designed; memory math confirmed, full implementation pending). The same A100 running FAISS-per-tenant supports about 16 tenants before VRAM pressure forces consolidation. That is a 400× density improvement, with proportional power and cost reductions.

The honest yes-but: if you are Google indexing the web, or a CDN doing deduplication at planetary scale, FAISS is the right tool. If you are anyone needing memory that *understands* the difference between cue and content, FAISS is solving a different problem.

Objection 6: "Why is this open source if it's so valuable?"

The substrate (Membot, Membrane, the cartridge format) is MIT-licensed because nobody adopts a proprietary memory layer for production work without source access. The value capture lives elsewhere.

The product stack built on top is where revenue accrues:

- **Heartbeat** is the consumer-facing chat-memory product--capture chats from any LLM surface, search them across devices, never lose them. Free tier (self-host or rate-limited cloud) feeds the funnel; paid tiers cover the costs of hosted backends and cross-device sync.
- **Membot hosting** is enterprise infrastructure--managed deployments with SLAs, compliance posture, and support contracts. The substrate is open; the operations are paid.
- **Membox multi-user** is the federated/multi-tenant tier--concurrency control, agent attribution, role-based access. Built on the open substrate; the multi-user layer and operational tooling are the differentiator.
- **Vector+ Studio** is the visual cart-management UI--proprietary frontend on top of the open substrate. Sells as a workstation product and a hosted service.

The longer-term moat is the **hardware path**. ASIC implementations of the substrate's physics engine--1-bit neurons, 1-bit weights, integer ops--are patentable and licensable even when the reference software is open. Open-source the substrate to build the developer community and create demand; capture value when that demand wants purpose-built silicon. This is the Postgres / Linux / Redis playbook, not the proprietary-from-day-one playbook. Both have proven lucrative; the open path scales adoption faster.

The honest yes-but: if the only thing we had was the software, "open source it" would be a worse business decision. The substrate is open precisely *because* the value extends beyond the software--into the hosted products, the enterprise tier, and the silicon. Open at the substrate layer, paid at the layers built on top.

Objection 7: "Doesn't stochasticity make recall unreliable?"

LatticeRunner uses stochastic neuron updates--each neuron's firing decision is sampled from a probability distribution rather than computed deterministically. The natural concern is that randomness at the neuron level should make recall outcomes unpredictable.

In practice, the opposite is true. Stochasticity at the micro scale produces *more* reliable recall at the macro scale, not less.

The mechanism is the same one biology uses. Individual neurons fire probabilistically; what persists across firings is what aligns with population-level structure. In LatticeRunner, that

structure is the attractor basin. Random perturbations that fall inside a basin reinforce the basin through local facilitation. Random perturbations that fall outside a basin are suppressed by lateral inhibition and fatigue. The system *uses* noise to escape shallow local minima while staying anchored to deep ones.

Across repeated trials with fixed seeds and cues, LatticeRunner converges to the same attractor with high consistency-- $R@1 = 1.000$ across 200 probe patterns at 3M scale, repeatable across runs. Randomness acts as a regularizer, not a source of instability. The same principle that makes biological brains noise-tolerant rather than noise-fragile makes the lattice noise-tolerant rather than noise-fragile.

The honest yes-but: this works because the system is designed for binary outcomes (fired/quiescent) where noise averaging across population statistics is meaningful. A continuous-activation system would not get the same benefit.

Objection 8: "Why binary neurons instead of continuous activations?"

Most modern neural systems use continuous activations (typically float32) because they preserve gradient information needed for backpropagation. LatticeRunner uses 1-bit neurons with 1-bit weights. The natural question: aren't we throwing away precision?

We are throwing away precision *at the neuron level*. We're keeping it at the *population level*.

Each neuron represents a binary decision: quiescent or fired. Probability, confidence, and uncertainty are not encoded in any single neuron--they emerge from population statistics, spatial redundancy, and temporal dynamics. A pattern stored in a 4096×4096 lattice spans millions of neurons; the precision of any individual binary state is irrelevant when the signal is carried by the joint distribution of all of them.

This choice unlocks structural advantages a continuous system cannot match:

- **Bit-packing efficiency.** A full lattice frame is 16,777,216 bits ≈ 2.1 MB. The same neurons in float32 would be ~ 67 MB. Binary state spreads across more storage cheaply.
- **Deterministic hardware execution.** XOR and popcount on 1-bit values map directly to FPGA and ASIC silicon. Float32 arithmetic requires multipliers that are 100-1000 \times more expensive in transistors and power.
- **Clear interpretation of energy flow.** Binary state makes basin geometry observable--a neuron is either contributing to the current basin or it isn't.

- **Gradient-free learning.** Hebbian co-activation produces useful weight updates without requiring differentiability, which is what enables learning during inference rather than separate training phases.

Continuous activations are not required to represent semantic structure when the structure is carried by spatial dynamics across millions of binary neurons rather than fine-grained values in fewer continuous ones. Different architecture, different precision strategy.

The honest yes-but: if your task requires fine-grained per-element precision (e.g. regression to a specific real-valued target), binary neurons are the wrong primitive. We are not trying to predict regression targets. We are trying to recall stored patterns, and binary spatial-population coding does that with substantially better hardware ergonomics than continuous values.

Objection 9: "How is learning stable without backpropagation?"

LatticeRunner has no backward pass, no gradient descent, no global loss function. The natural ML-researcher concern: without those, what prevents catastrophic forgetting, weight runaway, or chaos?

Three mechanisms working together:

- **Learning is gated to imprint phases.** Hebbian weight updates only occur when explicitly enabled--typically during pattern storage, not during recall. The lattice is not learning continuously from every query the way a fine-tuned model would. This eliminates the most common source of post-deployment drift.
- **Fatigue limits over-strengthening.** Neurons that fire frequently accumulate fatigue, reducing their firing probability in subsequent steps. Weights cannot grow without bound because the neurons participating in those weights become temporarily harder to activate. This is a local feedback mechanism that doesn't require a global regularizer.
- **Top-down modulation constrains basin growth.** L1, L2, and L3 layers exert downward pressure that biases settling toward consolidated patterns. New patterns can form basins, but they form within the geometric structure that the higher layers maintain--they cannot freely overwrite the existing landscape.

The result is a system that accumulates memory continuously without requiring a training loop and without the catastrophic forgetting that plagues neural networks fine-tuned in production. New patterns deepen new basins; old patterns retain their basins as long as they aren't actively contradicted.

The honest yes-but: this approach is not a substitute for systems that genuinely need representation learning from raw signal. We assume the embedding layer (Nomic, in our case) has already done that work. If you need to learn embeddings from raw inputs, you still want gradient descent. We learn the *associations* and *recall dynamics* on top of pre-existing embeddings--that's the layer where the gradient-free physics works.

Objection 10: "What about deletion, updates, and forgetting?"

A common production concern: vector databases need to handle data lifecycle--content gets retracted, policies change, GDPR requests arrive. Can a physics-based system delete a pattern cleanly?

Two answers, depending on what you mean by deletion:

- **Implicit forgetting (default behavior).** Rarely accessed basins gradually weaken as fatigue and decay run their normal course; frequently reinforced basins deepen. Interference naturally redistributes energy. This mirrors biological memory--you don't *delete* a memory by surgically removing neurons, you let it fade through disuse. Implicit forgetting requires no special operations and produces no rebuild cost.
- **Explicit deletion (when policy requires it).** Set a Tombstone flag in the hippocampus region (Region 63 of L4) on the patterns to remove. Save the cart to disk *skipping* tombstoned patterns. Reload the fresh file. The deleted patterns are gone--not just suppressed, gone from the on-disk representation entirely. The on-substrate weight contribution gracefully fades on the next imprint cycle as the surrounding patterns consolidate.

Compare this to vector-database deletion. FAISS HNSW requires the navigation graph to be rebuilt because removed nodes leave dangling references. IVF-PQ requires retraining the quantizer if enough vectors are removed to shift the data distribution. ChromaDB marks rows as deleted but the SQLite metadata keeps growing; periodic compaction is required. Most production vector DBs handle deletion as scheduled offline maintenance.

LatticeRunner handles explicit deletion as a single save-load cycle with zero downtime against the live substrate. The Tombstone flag is read at save time; the new cart is the new ground truth. No rebuild, no quantizer retrain, no graph regeneration.

The honest yes-but: implicit forgetting is gradual rather than instant. If your compliance posture requires *guaranteed* removal within seconds of a request (some regulated industries do), the Tombstone-then-reload cycle is the path; implicit forgetting alone isn't sufficient. The cycle is fast (seconds) and atomic, but it is a discrete operation rather than a per-pattern instant delete.

Objection 11: "Is this biologically realistic?"

LatticeRunner is inspired by the brain--attractor dynamics, Hebbian learning, lateral inhibition, hierarchical layers. The natural question from a neuroscientist: does the system actually model what neurons do, or is the brain just decoration?

It's neither pure simulation nor pure decoration.

The goal is **functional realism, not biological simulation**. Where a biological mechanism solves a problem we also need to solve--robustness to noise, capacity through sparsity, hierarchical pressure--we adopt the mechanism. Where biology adds complexity that doesn't help our engineering goals--spiking timing precision, neurotransmitter chemistry, glial cells, energy metabolism--we leave it out.

Concretely, LatticeRunner adopts:

- Attractor dynamics (Hopfield, Hebb, energy minimization)
- Lateral inhibition (Mexican Hat profiles)
- Sparse population coding (kWTA)
- Hierarchical modulation (cortical layer analog)
- Stochastic neuron firing (Boltzmann, biological neuron variability)
- Fatigue and recovery (neuronal accommodation)

LatticeRunner deliberately does NOT adopt:

- Spiking timing (we use frame-based update)
- Real-valued weights or activations (we use binary)
- Backpropagation or gradient descent (biology doesn't either, but for different reasons)
- Glial cell support, neurotransmitter dynamics, dendritic computation, biophysical detail
- Continuous-time differential equation models

This is the same selective borrowing that convolutional neural networks practice with respect to the visual cortex--the convolution operation is biologically inspired by V1 receptive fields, but no one claims a CNN is a biological model. CNNs prioritize engineering validity over biological completeness, and they work.

The honest yes-but: if your goal is to model a specific neural circuit, simulate spiking dynamics, or contribute to computational neuroscience, LatticeRunner is the wrong tool. We are building memory infrastructure that uses biological insights, not a brain model that happens to do retrieval.

APPENDIX C: PRIOR RESEARCH

LatticeRunner builds on a long lineage of research in associative memory, energy-based learning, sparse distributed representations, and neuromorphic computing. This section situates the architecture within existing literature and clarifies the distinctions that motivate the present work.

Hopfield Networks (1982, 1984)

Hopfield's seminal work introduced memory as convergence toward attractor states in an energy landscape. LatticeRunner inherits the attractor-basin concept but replaces global connectivity with spatially local interactions, introduces hierarchical modulation, and distributes energy dynamics across layers. This enables substantially higher capacity and improved noise robustness, and is the basis for the scale-positive scaling behavior described in Objection 2.

References: Hopfield (1982); Hopfield, J.J. (1984), *Neurons with graded response have collective computational properties*.

Boltzmann Machines and Energy-Based Models

Boltzmann Machines generalized Hopfield networks via stochastic neuron updates and temperature-driven sampling from an energy landscape. LatticeRunner adopts stochastic dynamics and local energy minimization but does not rely on global energy functions, contrastive divergence, or gradient-based learning. Learning occurs through local reinforcement during explicit imprint phases, avoiding the instability and training cost associated with Boltzmann-style learning.

References: Ackley, Hinton, Sejnowski (1985), *A learning algorithm for Boltzmann machines*; LeCun et al., energy-based-model literature.

Sparse Distributed Representations (SDR)

SDR research emphasizes robustness, noise tolerance, and high-capacity encoding through sparsity. LatticeRunner aligns closely with SDR principles--sparsity, redundancy, population-level representation--but embeds them into explicit spatial lattices with physical dynamics, enabling structured basin formation and hierarchical modulation rather than operating in abstract high-dimensional vector spaces.

References: Kanerva, P. (1988), *Sparse Distributed Memory*; Hawkins & Blakeslee (2004), *On Intelligence*; Numenta SDR literature.

Cellular Automata and Local Rule Systems

Research in cellular automata demonstrates that complex global behavior can emerge from simple local rules. LatticeRunner draws inspiration from local update rules and neighbor interactions but extends them with probabilistic firing, energy modulation, fatigue, and hierarchical control. Unlike traditional cellular automata, the system is designed explicitly for memory storage and recall, not arbitrary computation.

References: Conway (Game of Life); Wolfram (elementary cellular automata).

Neuromorphic Computing

Neuromorphic systems explore hardware and algorithms inspired by biological neural systems, often emphasizing locality, event-driven computation, and energy efficiency. LatticeRunner is neuromorphic in spirit but intentionally avoids biological realism where it conflicts with scalability or determinism. Binary neuron states, fixed update rules, and explicit layering make the system amenable to GPU, FPGA, and ASIC implementations without requiring spiking neuron models or analog dynamics.

References: IBM TrueNorth; Intel Loihi; spiking neural network literature.

Modern Memory-Augmented Architectures

Recent work has explored augmenting neural networks with external memory--transformer KV caches, RAG pipelines, vector databases (FAISS, HNSW, Milvus), and the current-generation AI memory ecosystem (Mem0, Letta, Zep / Graphiti, MemPalace, Supermemory). These systems treat memory as indexed storage accessed via search or attention. LatticeRunner instead treats memory as a dynamical substrate, where recall is reconstruction rather than retrieval. This distinction enables predictable behavior under noise and decouples memory growth from linear increases in compute and memory bandwidth.

References: Vaswani et al. (Transformer KV caches); Lewis et al. (RAG); Johnson, Douze & Jégou (FAISS, 2019); Pang et al. (MemPalace); contemporary work in Mem0, Letta, Zep / Graphiti.

APPENDIX D: SUMMARY OF NOVEL CONTRIBUTIONS

While informed by decades of prior work, LatticeRunner's primary novel contributions are:

- A hierarchical spatial lattice combining bottom-up and top-down modulation with binary stochastic neurons under local energy dynamics
- Explicit modeling of basin formation and competition with hippocampal episodic linking woven into the same substrate
- Demonstrated *scale-positive* behavior where recall fidelity holds or improves with system size, validated to 3 million patterns with no degradation visible at the upper bound
- Practical implementation on commodity GPU hardware using only integer operations, with a clean path to FPGA and ASIC silicon
- The unified-lattice paradigm where patterns are *vertically interwoven* into the same fixed-size 4096×4096 substrate rather than horizontally appended, enabling sublinear memory growth and constant per-pattern compute cost
- A three-mode database stack (single-user / federated / multi-user with locking) on the unified substrate--the first neuromorphic DBMS to our knowledge

These elements collectively define a memory system that is neither a classical neural network nor a traditional database, but a new category of physics-driven semantic memory.

APPENDIX E: TEST HISTORY

The headline numbers in **The Numbers** and **Pattern Reconstruction Quality** sections above are drawn from a multi-month battery of validations, each targeting a specific claim at its proper scale. This appendix is the full empirical record. Each entry lists the date, script, configuration, key results, and what that test contributed to the overall picture. Test scripts and result `.pkl` files are reproducible from the `benchmarks/` directory of the `membot` repository.

50K MULTI-LAYER BITFLIP--2025-12-24

- **Script:** early V7 multi-layer harness (5 lattice layers × 10K patterns each = 50,000 patterns total)

- **Configuration:** bitflip noise injection at recall time, BALANCED profile, GreenZoneController for settling
- **Headline:** Overall correlation *0.8060*, IoU *~0.68*, on-ratio *~1.2* (controlled bloom), imprint speed *0.05 s/pattern*.

What it added: first scaling validation beyond single-layer 100-pattern tests. Demonstrated that the multi-layer architecture preserves recall fidelity at 50K patterns under bitflip noise--earlier than the precision-stack and capacity work. The 25-layer × 10K (250K) extension was started but interrupted; the 50K result became the bridge between small-scale BAT validation and the eventual six-figure capacity tests of February 2026.

V7 BAT--BATTERY ACCEPTANCE TEST--2026-01-24

- **Script:** `v7_test_runner` on the V7 unified engine (4096×4096 lattice, L1=16, L2=64, L3=256)
- **Configuration:** QUALITY profile, 8 conditions across clean / bitflip / erasure noise at 10 and 100 patterns, with and without PID controller

Results:

Condition	Correlation	IoU
Clean recall, 10 patterns	0.9644	0.9337
Clean recall, 100 patterns	0.9639	0.9351
30% bitflip, 10 patterns (no PID)	0.8646	0.7734
20% bitflip, 10 patterns (no PID)	0.8941--	
30% bitflip, 100 patterns (no PID)	0.8711	0.7852
30% erasure, 10 patterns (with PID)	0.9079	0.8418
30% erasure, 100 patterns (with PID)	0.9073	0.8410

What it added: the foundational pattern-reconstruction-quality table reported in the body. Established the V4 gold-standard equivalence (0.86 correlation under 30% bitflip). Confirmed the QUALITY profile dramatically outperforms FAST (0.0992) and BALANCED (0.7785) under bitflip noise, and that capacity scales cleanly from 10 to 100 patterns within measurement noise. Noted that the harness's 0.90 "PASS" threshold was aspirational and that 0.86 was the correct interpretation. Headroom note: the PID controller engaged for erasure tests was never fully tuned--these IoU numbers are the floor, not the ceiling.

CAPACITY TEST V1--500K SYNTHETIC--2026-02-13

- **Script:** `test_capacity_degradation.py` (vector-benchmark-demo/cuda/)
- **Configuration:** 4096×4096 V7 lattice, 768-dim unit-norm Gaussian synthetic vectors, 200 probes, L3 signature comparison, 6 checkpoints (1K, 10K, 50K, 100K, 250K, 500K)
- **Headline: R@1 = 1.000 at every checkpoint.** Self-similarity 0.9994. Cross-similarity 0.228. Discrimination gap 0.771--degraded by 0.03% across 500× scale increase. Query 2 ms constant. Training 470 patterns/sec. Total wall-clock 17.8 minutes.

What it added: first six-figure capacity proof and the first systematic checkpoint-curve showing that recall-discrimination is essentially flat with scale. Established the test pattern that subsequent runs (1M Path B, 3M Capacity) extended. Honest caveat captured at the time: this is a 200-way retrieval test against the 200 probes' references, not a 500K-way retrieval against the full population--that harder test came later in Path B and Region-Fill validation.

PATH B WIKI EMBEDDINGS--100K--2026-02-15

Script: `benchmarks/test_path_b_capacity.py --max-patterns 100000 --probes 200 --settle 30 --cartridge data/wiki_nomic_100k.pkl --tile 5`

Configuration: real wiki Nomic embeddings (5× tiling), QUALITY profile, 30-frame settle, 200 probes, **24 conditions** (8 checkpoints × 3 noise levels: clean, 10% erasure, 30% erasure), re-snapshot at each checkpoint

Headline: R@1 = 1.000 across ALL 24 conditions. Recall margin **MONOTONICALLY GROWING:** +0.624208 (at 500 patterns) → +0.624618 (at 100K)--opposite of classical Hopfield degradation. Signature diversity grew (0.073784 → 0.074315). Cross-similarity decreased (0.3758 → 0.3754). Query 23.6 ms constant. Training 37-38 pat/sec, 44.3 minutes total. **30% erasure recall self-similarity 0.9990.**

What it added: replaced the synthetic-Gaussian caveat from Capacity V1 with real-data validation. Established the *monotonically growing* margin as a substrate property, not an artifact of synthetic distributions. The 30% erasure self-similarity of 0.9990 became the headline noise-tolerance number reported in the body. Resolved the concern raised against Capacity V1 ("could the synthetic distribution be hiding interference") by showing the same trend holds on real correlated embeddings.

PATH B WIKI EMBEDDINGS--1M--2026-02-15

- **Script:** `benchmarks/test_path_b_capacity.py` (3M wiki .bin extension, 1M cap)
- **Configuration:** real wiki embeddings, **33 conditions** (11 checkpoints × 3 noise levels) up to 1,000,000 patterns
- **Headline: R@1 = 1.000 across ALL 33 conditions.** Recall margin grew monotonically from +0.607255 (500 patterns) to +0.607673 (1M). Training 36-37 pat/sec, 8.1 hours total wall-clock.

What it added: extended the Path B trend from 100K to 1M without any sign of capacity wall. Margin still growing at 1M; the curve was flat enough that further extension required a different test rig.

V8 / V9 PRECISION STACK--100K--2026-02-15

Script: full synthetic suite at 100K with 6 precision variants

Configuration: 100K patterns, 7000-way ranking, 2000 probes

Results:

Variant	Brain Size	Disc Drop vs V7	Compression	Query Time
V7 Float32	128 MB	baseline	1×	3.1-3.5 ms
V8 INT8	128 MB	0.0000	1×	(faster training: 441 vs 349 pat/sec)
Compact uint8	64 MB	+0.0001	2×	
Compact uint4	32 MB	-0.0005	4×	
Compact uint2	16 MB	-0.0008	8×	
V9 Bitpacked	12 MB	+0.0088	10.7×	1.5 ms (fastest)

What it added: the precision pipeline that the Performance Roadmap subsection references.

R@1 = 1.000 across the entire precision stack at 100K--zero loss from float32 to bitpacked. uint2 emerged as the stealth winner (8× compression, eight ten-thousandths discrimination drop). V9 Bitpacked's discrimination gap actually *increased* (binary weights act as a contrast filter), confirming that the architecture's hardware-friendly path is not a degradation. This is the empirical basis for the FPGA mapping claim.

REGION-FILL 1M WIKI--2026-02-26

- **Script:** region-fill encoding at 1M wiki embeddings, all noise conditions
- **Configuration:** 1,000,000 wiki Nomic patterns with the region-fill encoder (alternative to thermometer), tested across clean / 10% erasure / 10% bitflip
- **Headline: R@1 = 1.000 at 1,000,000 patterns under all three noise conditions.**
Margin plateaued at clean +0.2445, erasure +0.2318, bitflip +0.1919. Delta from 2K → 1M was -0.0004 (4th decimal = noise). 18.5 hours runtime, 15 patterns/sec, ~62 ms/query constant.

What it added: validated region-fill as an alternate-to-thermometer encoding scheme that matches recall fidelity at 1M scale while also being useful for search-and-association workloads (the thermometer encoding is recall-optimized; region-fill is more general-purpose). Demonstrated the substrate is encoding-agnostic at the recall-quality layer--different encoders, same scale-positive behavior.

3M CAPACITY TEST--V7 LATTICE BRAIN--2026-04-25

- **Script:** `test_capacity_degradation.py` (V2 with 9-checkpoint extension)
- **Configuration:** single 4096×4096 V7 lattice (Float32, 128 MB), 768-dim synthetic Gaussian patterns, BALANCED profile, 2-frame Hebbian settling per pattern, 200 probes, 9 checkpoints: 1K, 10K, 100K, 500K, 1M, 1.5M, 2M, 2.5M, **3M**

Results across the full range:

Metric	@ 1K	@ 100K	@ 1M	@ 3M	Total drift
R@1	1.000	1.000	1.000	1.000	0
Self-similarity	0.9995	0.9994	0.9993	0.9993	-0.0002
MinSelf (worst-case)	0.9946	0.9945	0.9950	0.9955	+0.0009
Cross-similarity	0.2280	/	/	0.2279	-0.0001
Discrimination gap	0.7716	/	/	0.7714	-0.0002
MaxCross (best-confusion)	0.5958	0.5959	0.5955	0.5958	~0
Signature std	0.07777	/	/	0.07809	+0.00032
Query latency	2.1 ms	2.1 ms	2.1 ms	2.1 ms	0

Training rate held at 467-469 pat/sec across the full 3M run. Total training time: **1.8 hours on a single RTX 4080.**

What it added: the headline figure quoted throughout the body. **3M patterns surpasses the classical Hopfield capacity bound** (~2.3M for 16M neurons) by 30% with zero recall loss. The discrimination gap moved by two ten-thousandths across 3000× scale; signature diversity actually grew slightly with scale (0.07777 → 0.07809), the opposite of classical degradation. This is the at-scale clean-recall companion to the V7 BAT pattern-reconstruction-under-noise table--the two together cover the small-scale-with-noise and large-scale-clean corners of the validation matrix.

What this collection does and doesn't establish

Establishes: Scaling behavior on synthetic and real-data distributions up to 3M patterns; pattern reconstruction quality under bitflip and erasure noise at small scale; precision-stack equivalence from float32 to bitpacked at 100K; alternate-encoder validation at 1M; query-latency constancy across 6 orders of magnitude; capacity-positive (margin-growing) behavior across multiple test families.

Does not yet establish: End-to-end agent task performance with a local reader model (in progress, Gemma 3 27B preliminary); multi-hop reasoning benchmarks (MuSiQue, 2WikiMultiHop on roadmap); cross-modal retrieval (pathology image + radiology pairing--see separate concept-cluster); production-scale federation stress tests under adversarial workload (shipped and externally validated, but not formally benchmarked).

The above tests are the documented floor of the substrate's behavior. The Performance Roadmap subsection in the methodology section above describes the structural headroom that exists beyond it.

APPENDIX F: FILESIZE LADDER

The "128 MB substrate" headline number quoted in the body is the lattice mathematical footprint--and the right number for the apples-to-apples comparison with vector databases that have no substrate equivalent. But it is one layer of a multi-layer storage story, and IT readers reasonably want the full ladder. This appendix is that ladder.

The full ladder

Layer	Size	What's in it	Status
Substrate (lattice mathematical footprint)	128 MB	4096×4096 Hebbian weight matrix, the headline-number form	Estimate (raw float32 = 67 MB; 128 MB rounds up for overhead/metadata)
Full brain state on disk	134 MB	Weights + fatigue state + energy/lattice state, the form that actually saves and loads via `.pkl`	Measurement
Active VRAM during inference	200-300 MB	Weights (64 MB) + fatigue (64 MB) + energy/state (64 MB) + L1/L2/L3 hierarchy + temporary physics buffers	Measurement (from V7 stress runs)
Sign bits per pattern	96 bytes	Sign-zero binary fingerprint, 768 dimensions packed	Measurement
Sign bits at 1M patterns	~91 MB	The corpus-side cost of F1 retrieval	Derived
Sign bits at 3M patterns	~274 MB	Same, at the upper bound we have validated	Derived
Pair index per pattern	~7.5 bytes	Embedding-to-text mapping, packed	Measurement
Pair index at 2.4M patterns	~18 MB	Real measurement	Measurement
Text passages	varies	Source text, pickled (uncompressed)	Content-dependent
Full cart at 2.4M patterns (measured)	1378 MB (1.378 GB)	Brain (134) + sign bits (230) + pair index (18) + text (~1 GB)	Measurement of record
Full cart at 1M patterns (interpolated from 2.4M measurement)	~1.3 GB	Same composition, scaled down	Estimate
FAISS Flat at 1M × 768-dim	~3 GB	Raw vectors only, no index, no text	Public measurement
FAISS HNSW (M=32) at 1M	~3.2-4 GB	Vectors + navigation graph, no text	Public measurement (we measured 4.2 GB at 1.325M)
FAISS IVF-PQ at 1M	~100-200 MB	Compressed, lossy, lower recall, no text	Public measurement
Raw float32 embeddings at 100K	<308 MB	Just the vectors at full precision, no index, no text	Derived (100,000 × 768 × 4 bytes = 307,200,000 bytes)
Raw float32 embeddings at 1M	~3 GB	Same, at 1M	Derived

MEASUREMENT OF RECORD: THE 2.4M CART PACKAGING OUTPUT

When the 2,400,000-pattern wiki cart was packaged on 2026-02-26, the cart-builder emitted the following manifest:

CART PACKAGED

File: wiki_2400k_brain.pkl

Size: 1378 MB

Patterns: 2,400,000

Search: Hamming-only (sign-zero, 96 bytes/pattern)

Brain: 134 MB (Hebbian weights for physics recall)

Pair index: 2,400,000 entries (embedding->text mapping)

This is the canonical filesize datum. Smaller carts are interpolated downward from this measurement; larger carts are projected upward by adding sign-bit and text-passage cost (linear in pattern count) to the fixed 134 MB brain.

WHAT'S NOT IN THE CART: THE L3 DESIGN CHOICE

A naïve cart format would store one L3 signature per pattern (16 KB per pattern × 3 million patterns = ~46 GB at the 3M scale). That would make the cart larger than the FAISS index it replaces, defeating the substrate's compression argument.

We do not ship per-pattern L3 signatures. Two reasons:

1. **F1 (Hamming-distance retrieval over sign bits) is the production retrieval mode** and produces $R@1 = 1.000$ in 1.5 ms at 100K patterns and 2.1 ms at 3M. Per-pattern L3 signatures are not needed for retrieval; sign bits are.
2. **F0 (full physics-based settling) is regenerated on-demand from the brain weights.** When F0 verification or basin-walk association is required, the candidate patterns are re-imprinted with `learn=False`, settled, and read out at L3--at query time, on the fly, against the same 134 MB brain that ships in the cart. F0 is an inference-time computation, not a storage cost.

This is the **F1-primary, F0-on-demand two-stage pipeline**. F1 retrieves; F0 verifies or re-ranks when the query genuinely needs physics-based association rather than ranked retrieval. Most production queries never invoke F0; the ones that do reconstruct the signatures fresh and pay the per-frame settle cost (~30 frames × per-frame physics) only for the small candidate set, not the entire corpus.

The architectural payoff: **cart size grows linearly in pattern count, not quadratically**. At 3M patterns the cart adds 274 MB of sign bits + ~22 MB of pair index + content-dependent text, on top of the constant 134 MB brain.

THE COMPRESSION LEVER: ZLIB ON TEXT PASSAGES

The 2.4M cart's 1.378 GB total is dominated by source-text passages (~1 GB), because pickle stores strings uncompressed. zlib compression on the text component would reduce the cart to roughly **600-700 MB** at 2.4M patterns, at the cost of a decompression step on each search hit (low single-digit milliseconds, dwarfed by the LLM's own thinking time on the agent side).

This compression option has not been invoked in our shipped carts to date because the uncompressed cart already mounts cleanly within droplet RAM constraints. It is a real lever available for customer-facing carts where on-disk footprint matters more than the milliseconds of decompression cost--for example, cart marketplace distribution where bandwidth drives the user experience.

HOW TO READ THIS AGAINST FAISS

Three numbers position us cleanly without conflation:

- **Cart-to-cart at 1M patterns:** ~1.3 GB (us) vs FAISS HNSW 3.2-4 GB. Our cart contains source text in addition to vectors and signatures; FAISS HNSW does not.
- **Substrate-only:** 128 MB fixed (us) vs no FAISS equivalent. The lattice-as-substrate is unique to LatticeRunner; FAISS is exclusively an indexing structure with no recall/reconstruction layer.
- **Raw vector storage at 1M:** ~3 GB (float32 baseline)--what any system would pay just to hold the vectors without any index or text.

The substrate is roughly **1/24th** the size of a FAISS HNSW index at 1M patterns. The full cart is **1/3** the size. Our cart is also self-contained--text travels with vectors travels with brain--where FAISS requires a separate document store synchronized alongside the index. Operational overhead (rebuild cycles, quantizer retraining, index-swap orchestration) is described in the body's *Storage and Operational Footprint* section above.

HONEST CAVEATS

- The 128 MB substrate number is a rounding of the 67 MB raw float32 weight matrix plus overhead/metadata; the precise number depends on the CUDA engine's internal representation. We use 128 MB as the headline because it is the conventional rounded figure; the precise on-disk save/load representation including state is 134 MB.
- The 1.378 GB cart was measured on 2.4M wiki patterns. The 1M projection (~1.3 GB) is interpolated linearly; for content distributions significantly different from wikipedia (longer or shorter average passage length, more or fewer images encoded as patterns), text-passage size scales accordingly and should be re-measured per corpus.

- The L3 design choice (regenerate-on-demand rather than per-pattern storage) means F0 query latency varies with the candidate-set size and the settle profile in use. F0 verification of the top 10 F1 candidates at the QUALITY profile costs roughly 230 ms (10×23 ms per re-imprint+settle); at the BALANCED profile it costs roughly 21 ms (10×2.1 ms). These are upper-bound estimates and have not been benchmarked on a production cart.
- 200-300 MB VRAM-per-lattice is a stress-test measurement; production deployments may use less if the engine releases temporary buffers between queries, or more if multiple queries are pipelined.

REFERENCES

- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The Psychology of Learning and Motivation* (Vol. 2, pp. 89-195). Academic Press.
- Chollet, F. (2019). On the Measure of Intelligence. *arXiv preprint arXiv:1911.01547*.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407-428.
- Ebbinghaus, H. (1885/1913). *Memory: A Contribution to Experimental Psychology*. Teachers College, Columbia University.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8), 2554-2558.
- Johnson, J., Douze, M., & Jégou, H. (2019). Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3), 535-547. [FAISS]
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3), 419-457.
- Nomic AI. (2024). Nomic Embed: Training a Reproducible Long Context Text Embedder. *arXiv preprint arXiv:2402.01613*.
- Palatov, D., Claude Opus 4.6, Grossberg, A., & Claude Opus 4.6. (2026). ARC-SAGE: World-Model Harness and Multi-Agent Frame-Questioning on ARC-AGI-3. *In preparation*. <https://github.com/dp-web4/ARC-SAGE>
- Pang, R. Y., et al. (2023). MemPalace. <https://github.com/supermemoryai/mempalace>
- Robertson, S. E., & Zaragoza, H. (2009). The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4), 333-389.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of Memory* (pp. 381-403). Academic Press.

Vectorize.io. (2026). Independent analysis of MemPalace LongMemEval performance.
<https://vectorize.io/mempalace-benchmark-evaluation/>

Wu, Y., et al. (2024). LongMemEval: Benchmarking Chat Assistants on Long-Term Interactive Memory. *arXiv preprint arXiv:2410.10813*.

Contact: andy.grossberg@gmail.com 415-317-1569

###

Contact: *andy.grossberg@gmail.com | 415-317-1569*

GitHub: *github.com/project-you-apps*

Paper: *ARC-SAGE (Palatov, Grossberg, et al., 2026)--in preparation*