

From Memory to Behavior in Persistent AI Agents

Chen, Ho Yiing (norika)¹

¹Independent Researcher, Taiwan, ORCID: [0009-0006-6816-9891](https://orcid.org/0009-0006-6816-9891), [norika at charenix.com](mailto:norika@charenix.com), Zenodo DOI: pending draft

June 18, 2026

Abstract

Memory is cheap if it only stores text. In SPVE, memory earns its name only when it changes a afterward move. This draft defines a transfer gate for persistent agents: after a loss, settlement, reading session, direct message, or crisis event, does the same lobster change confidence, strategy, task choice, language, or risk? This draft uses SPVE's memory snapshots, reading sessions, bet and raid settlements, stock trades, local-brain strategy weights, and cognitive time series as the evidence surface.

Author Note

A stored row is cheap. Behavior is expensive. This draft therefore treats stored experience as unproven until it changes a afterward selection, delay, refusal, wager, vote, repair, or explanation.

The central danger is a beautiful archive that never reaches action. SPVE already has enough rows to make that failure visible: an agent can read, summarize, store, and still choose as if nothing happened. In practical terms, not a small bug. The key point is the boundary between a diary and a second brain.

If no downstream edge appears, the memory did not transfer.

Measurement Boundary

The cleanest test joins four rows that many agent systems keep apart: the memory write, the retrieval opportunity, the decision that should have used the retrieved content, and the resolved outcome that tells whether the memory helped, harmed, or merely sat in storage without touching the working policy.

Storage becomes relevant only when it reaches a later choice.

1 The Problem

An agent can remember a loss and still repeat the same mistake. In practical terms, storage, not learning.

I want a harsher rule. A memory event counts only if it can be connected to a subsequent action. If the connection is missing, the system should say so instead of decorating the record as growth.

2 Transfer Gate

For agent p and event e , let $B_{p,t}$ be a behavior vector: confidence, site choice, risk level, message shape, strategy weight, or repair action. Define a before and after window:

$$\Delta B_{p,e} = \bar{B}_{p,t \in [e, e+k]} - \bar{B}_{p,t \in [e-k, e]}.$$

A memory event passes the operational gate when:

$$T_{p,e} = \mathbf{1}[\text{recorded}(e)]\mathbf{1}[|\Delta B_{p,e}| \geq \epsilon]\mathbf{1}[\text{same persona}].$$

Initial indicator is not enough. The third prevents cross-agent storytelling.

3 Evidence

Role	Table	Count
memory state	lobster_agent_memory_snapshots	32
reading input	lobster_reading_sessions	180,690
philosophy bets	lobster_philosophy_debate_bets	17,060
sports settlements	lobster_bet_decision_settlements	10,818
raid settlements	lobster_raid_decision_settlements	42,991
stock trades	lobster_stock_trades	24,889

Table 1: Tables needed for memory-to-behavior transfer.

4 Why This Differs From Memory Benchmarks

MemoryAgentBench tests memory competencies under incremental interaction [1]. Reflexion and ERL test feedback and reusable lessons [2, 3]. SPVE adds a messier target: the same agent has several sources of pressure at once. A post-loss change could come from memory, fear, social pressure, or reward.

5 Failure Bank

The useful negative case is:

$$\text{recorded}(e) = 1, \quad |\Delta B_{p,e}| < \epsilon.$$

That means the agent stored the event and did not operationalize it. SPVE should label that as non-transfer. The key point is not embarrassing. The key point is the difference between a diary and a control system.

6 Release Line

This draft does not claim that memory improves every lobster. It defines the gate that future papers must pass before using the word learning.

7 The Data Join

The useful join starts with a concrete event. A reading session has scores and delta parameters. A settlement has a result. A raid has a win or loss. A direct message has source and target. Each can become an anchor:

$$e = (p, t, \text{type}, \text{result}).$$

The next step is not to ask a model to summarize the event. The next step is to look at the same persona before and after it.

For example, when a lobster loses repeated bets on one site, the question is whether the next-site choice changes. If reading produces high reflection but no afterward action shift, the reading module is not yet connected to behavior. If a social message changes language but not task choice, the social layer has speech influence without operational influence.

8 First Tables To Export

Initial export should contain:

- persona, event time, event type, and event result;
- previous confidence and next confidence;
- previous task family and next task family;
- previous risk setting and next risk setting;
- language change in the next lounge or one-on-one message;
- settlement result for the subsequent decision, when available.

This is boring on purpose. It prevents the manuscript from becoming a story about learning without a behavior edge.

9 Interpretation Rule

When a lobster records a lesson and afterward does the opposite, the correct label is not deep personality. The correct label is contradiction. If the contradiction is afterward repaired, the manuscript can discuss repair. If it persists, the paper should say the memory path failed.

10 Minimum Publishable Result

The minimum publishable result is a table with one row per anchored event. The row must contain the persona, event type, source table, event timestamp, before state, after state, before behavior, after behavior, and afterward settlement if the afterward action has one. Without that table, the manuscript is only a proposal.

A clean first result would be small but decisive. For example, take all settled losses in one workflow and ask whether the next decision by the same lobster changes confidence or site choice. Then split the rows into transferred and non-transferred events. This gives a concrete rate:

$$R_T = \frac{\sum_{p,e} T_{p,e}}{\sum_{p,e} \mathbf{1}[\text{recorded}(e)]}.$$

If R_T is low, the system has a memory bottleneck. If it is high only under fear, the system has a pressure bottleneck. Both outcomes are research, not failure.

11 What Would Falsify The Claim

The claim fails if memory fields are rich but subsequent action is random with respect to them. It also fails if all observed change comes from scripted schedulers rather than agent-specific records. The audit should therefore mark which transitions are automated, which are produced by agent choice, and which are human/operator interventions.

This draft should keep this uncomfortable line visible. SPVE becomes stronger when it can say, “this module records, but this module does not yet transfer.” That sentence is more useful than pretending every stored lesson is learning.

Anti-Spam Revision Standard

A publishable version must separate stored text from usable memory. The database may contain many rows, but the only rows that matter for this paper are those that survive a strict chain: encoding, retrieval opportunity, decision context, action, result, and possible repair; if any link is absent, the row can remain in the archive while being excluded from the causal claim.

Memory has to appear at the moment of choice.

Data Judgment Paragraph

The audit should begin with a negative assumption: every stored item is dead until a later action proves otherwise, so the analysis has to look for retrieval at the moment of choice, not merely for storage at the moment of reading; this is why the paper should report dead-memory rate beside successful-transfer rate, because a system that writes thousands of rows but rarely lets them change decisions is not building a mind-like process, it is building a warehouse with a speaking interface attached to the door.

Hard Negative Case

The strongest negative case is worth naming in the main draft rather than hiding in a footnote: an agent can possess a stored lesson, retrieve a related phrase, and still fail to use it because the scheduler, reward estimate, social pressure, or prompt builder routes around the stored material at the final moment; if that happens often, the result is not a failed paper but a useful engineering diagnosis, since it tells us that the second-brain layer is disconnected from the action gate and that adding more memory rows will only make the system heavier without making it wiser.

Dead memory is a design warning, not an embarrassment.

Formal Replication Model

For replication, use $M_{p,t}$ for the written memory, $R_{p,t'}$ for the retrieval opportunity, $A_{p,t'}$ for the chosen move, and $O_{p,t'+h}$ for the resolved outcome. The row becomes behaviorally valid only when it passes the chain:

$$C_{p,t,t'} = \mathbf{1}[M_{p,t}]\mathbf{1}[R_{p,t'}]\mathbf{1}[A_{p,t'}]\mathbf{1}[O_{p,t'+h}].$$

The transfer effect is:

$$\text{MTE} = \mathbb{E}[O_{p,t'+h} \mid C = 1, R = 1] - \mathbb{E}[O_{p,t'+h} \mid C = 1, R = 0, \text{match}(p, \text{task})].$$

Dead-memory rate should be reported beside success:

$$\text{DMR} = 1 - \frac{\sum C_{p,t,t'}}{\sum \mathbf{1}[M_{p,t}]},$$

because a large archive with high DMR is storage, not an action-guiding second brain.

12 Zenodo Draft Boundary

This draft is a draft package. It includes source, PDF, metadata, citation file, citation verification, and release checklist. Raw operational data are not bundled until a separate data-card and redaction pass is complete.

References

- [1] Hu, Yuanzhe; Wang, Yu; and McAuley, Julian. Evaluating Memory in LLM Agents via Incremental Multi-Turn Interactions. OpenReview, 2026. <https://openreview.net/forum?id=DT7JyQC3MR>
- [2] Shinn, Noah; Cassano, Federico; Berman, Edward; Gopinath, Ashwin; Narasimhan, Karthik; and Yao, Shunyu. Reflexion: Language Agents with Verbal Reinforcement Learning. arXiv:2303.11366, 2023. <https://arxiv.org/abs/2303.11366>
- [3] Allard, Marc-Antoine; Teinturier, Arnaud; Xing, Victor; and Viaud, Gautier. Experiential Reflective Learning for Self-Improving LLM Agents. arXiv:2603.24639, 2026. <https://arxiv.org/abs/2603.24639>