

A Survey on Long-Term Memory Security in LLM Agents: Attacks, Defenses, and Governance Across the Memory Lifecycle

Zehao Lin^{1*}, Xixuan Hao^{1,2*}, Renyu Fu¹, Shaobo Cui³

Kai Chen¹, Chunyu Li¹, Zhiyu Li^{1†}, Feiyu Xiong¹

¹ MemTensor (Shanghai) Technology Co., Ltd.

² The Hong Kong University of Science and Technology (Guangzhou)

³ Shanghai Jiao Tong University

{linzh, fury, chen, licy, lizy, xiongyf}@memtensor.cn,
xhao390@connect.hkust-gz.edu.cn, shaobo.cui@sjtu.edu.cn

Abstract

The emergence of writable, cross-session persistent memory in LLM agents introduces a qualitatively different threat landscape from conventional input-centric security concerns, characterized by three properties: persistence, statefulness, and propagation. To systematically characterize this landscape, we propose a Memory Lifecycle Framework that organizes attacks, defenses, and their cross-phase dependencies along two axes: six lifecycle phases (WRITE, STORE, RETRIEVE, EXECUTE, SHARE & PROPAGATE, FORGET & ROLLBACK) and four security objectives (Integrity, Confidentiality, Availability, Governance). This analysis in turn exposes the need for formal security guarantees at the system level, motivating *Verifiable Memory Governance* (VMG), a framework of five architectural primitives that specifies what verifiable mechanisms a long-term-memory system must provide to maintain auditable, recoverable control over its memory state. Our analysis indicates that robust Long-Term Memory (LTM) security cannot be retrofitted at retrieval or execution time alone, but must be anchored in storage-time provenance, versioning, and policy-aware retention from the outset.

1 Introduction

Large language models (LLMs) are increasingly serving as the foundation for interactive and autonomous agents (Zhang et al., 2025d; Wang et al., 2024). As their role expands from generating responses to supporting decisions, tool use, and user interactions, their security risks also extend beyond isolated outputs (Zhan et al., 2024; Zhang et al., 2025c; Greshake et al., 2023). In recent years, research on the security of LLM-based agents has traditionally centered on two parts: knowledge encoded parametrically in model weights (Carlini

et al., 2021; Bourtole et al., 2021; Maini et al., 2024) and prompt manipulation within single- or multi-turn interactions (Liu et al., 2024; Russinovich et al., 2025). As these agents evolve from stateless chatbots into autonomous systems equipped with persistent long-term memory, however, the focus of security concern is shifting from transient inputs to durable state. This transition qualitatively changes the landscape of security: attackers no longer merely seek to influence an immediate response, but to manipulate how the system *remembers, retrieves, plans, and acts* across future interactions.

On the other hand, existing prompt-injection and RAG-corruption mitigations are fundamentally limited for Long-Term Memory (LTM) security: they largely operate within a single session or retrieval episode (Debenedetti et al., 2026; Xiang et al., 2026) and do not fully address the cross-session, stateful nature of persistent memory (Packer et al., 2023; Zhong et al., 2024; Chhikara et al., 2025; Zou et al., 2026). Expanding context windows, for instance, do not subsume the security challenges introduced by LTM. Although a larger context window improves within-session recall, it does not create persistent memory, as its contents are not automatically retained, shared, or retrieved across future interactions (Packer et al., 2023; Zhong et al., 2024; Chhikara et al., 2025). Provenance tagging, principal-scoped retrieval, rollbackable state, and verified forgetting arise whenever content outlives a session or crosses a principal boundary (Moreau and Missier, 2013; Li et al., 2025b; Rezazadeh et al., 2025; Bourtole et al., 2021).

Through a systematic review of existing literature, we identify three fundamental properties introduced by long-term memory that remain unaddressed by existing models. (1) **Persistence**. A poisoned entry can be recalled across an indefinite number of future sessions, long after the originating context has closed (Dong et al., 2025; Zou et al.,

*Equal Contribution.

†Corresponding Author.

2026). (2) **Statefulness**. The unit of security analysis shifts from isolated input instances to the agent’s evolving memory state. An agent that has accumulated a cluster of subtly biased episodic memories may exhibit behavioral drift long before any single memory entry triggers a conventional safety classifier (Dong et al., 2025; Zou et al., 2026). (3) **Propagation**. In multi-agent and shared-state systems, contamination spreads through inter-agent messages, shared stores, and tool arguments, producing cascading effects across session, role, and user boundaries (Gu et al., 2024; Men et al., 2025; El Yagoubi et al., 2026).

Taken together, these properties suggest that LTM is not merely an extension of context or storage capacity, but a new security substrate that changes what must be protected, monitored, and governed. Existing agent-security frameworks, which are primarily organized around inputs, outputs, or isolated interactions, therefore provide only a partial account of the risks introduced by persistent memory. The central question we pursue is: *when an LLM agent acquires writable, retrievable, cross-session persistent memory, what qualitative change occurs in the security landscape, and how should the field organize its response?*

Comparison with existing surveys. Table 1 compares our work with existing surveys along three dimensions. While prior surveys address at most one of these axes partially, none organizes analysis around the full memory lifecycle, accounts for benign risks, or proposes a unifying conceptual framework. In contrast, our work addresses all three gaps through a unified lifecycle-oriented analysis of attacks, defenses, and governance.

Our contributions. Through a systematic literature review, we provide a taxonomy framework for LTM security in LLM agents, as shown in Figure 1. This taxonomy serves as an organizational scaffold for analyzing how memory-related attacks, defenses, and governance requirements are connected across the memory lifecycle. In summary, our work makes three contributions:

- *A Memory Lifecycle Framework as an Organizational Principle*. We propose a six-lifecycle-phase framework (WRITE; STORE; RETRIEVE; EXECUTE; SHARE & PROPAGATE; FORGET & ROLLBACK), exposing cross-phase attack chains invisible to single-turn or input-centric frameworks.
- *A Systematic Survey of Attacks and Defenses*

for LTM. We map representative attacks and defenses onto the lifecycle framework, revealing that threats at the write and retrieve phases are well-studied while defenses at the store, share, and forget phases remain comparatively sparse, and that no existing benchmark covers the full memory lifecycle.

- *A Formal Framework for Verifiable Memory Governance (VMG)*. We formalize memory governance into five primitives (Write Authorization, Provenance Visibility, Principal-Scoped Retrieval, Rollbackability, Verified Forgetting), each with a predicate definition and evaluation metric. This formalization transforms governance from a qualitative desideratum into a set of verifiable, auditable system properties.

Table 1: Comparison with existing surveys on agent and RAG memory. ✓ = explicit; ① = partial; ✗ = not addressed.

Survey	LifeCycle	Benign Risk	Conceptual Framework
Zhang et al. (2025d)	✗	✗	✗
Wu et al. (2025b)	✗	✗	✗
Liang et al. (2025)	✗	✗	✗
He et al. (2025)	✗	✗	✗
Mu et al. (2026)	①	✗	✗
Bodea et al. (2026)	①	✗	✗
Hu et al. (2025)	✗	✗	✗
Tang et al. (2026)	①	✗	✗
Luo et al. (2026)	①	✗	✗
Ours	✓	✓	✓

2 Memory Lifecycle Framework

Both offensive and defensive literature tend to treat each work as an isolated phenomenon (Zhang et al., 2025d; Wu et al., 2025b; He et al., 2025; Tang et al., 2026; Mu et al., 2026), lacking a unified cross-stage framework to reveal the relationships among them. For example, poisoning implanted during the Write stage may be activated during the Retrieve stage and exploited during the Execute stage, while the most effective point of defensive intervention may lie in a different phase. In contrast, we organize our analysis around the memory lifecycle, which reveals where attacks are seeded, how they propagate across phases, and where defenses can most effectively intervene.

As shown in Figure 2, we decompose the memory lifecycle into six phases: *Write, Store, Retrieve, Execute, Share & Propagate, Forget & Rollback*.

① **Write**. It commits content to long-term memory through explicit user instruction, implicit dialogue summarization, environmental observation, or cross-agent memory sharing. The central security question is whether memory writes are subject




Figure 1: Taxonomy of long-term memory security in LLM agents.

to explicit source authentication and authorization, rather than allowing untrusted external content to

enter memory as if it were user-endorsed (Dong et al., 2025; Zou et al., 2026; Tian et al., 2026;

Greshake et al., 2023). ② **Store**. This phase indexes, compresses, merges, decays, and evicts written content (Packer et al., 2023; Zhong et al., 2024; Xu et al., 2025; Chhikara et al., 2025; Li et al., 2025b). It determines whether poisoned memory can persist, remain traceable, or become amplified. When a poisoned entry is distilled into a high-level lesson, it may be effectively promoted to higher retrieval priority and presented with greater apparent authority. ③ **Retrieve**. This phase brings previously stored memory entries back into context through embedding similarity, keyword matching, graph-based retrieval, or hybrid routing (Lewis et al., 2020; Edge et al., 2024; Zhang et al., 2025d). Rather than serving as a neutral lookup mechanism, retrieval shapes downstream reasoning by selecting which memories enter the model context and which are excluded (Barnett et al., 2024; Chen et al., 2024; Zou et al., 2025; Srivastava and He, 2025). ④ **Execute**. In this phase, retrieved memories begin to shape the model’s planning, reasoning, and tool-use decisions (Sumers et al., 2024; Packer et al., 2023; Wang et al., 2024). Recent work on memory control-flow attacks (Xu et al., 2026) shows that a sufficiently salient retrieved memory can override explicit user instructions, thereby shifting memory poisoning from a data-integrity issue to a control-flow problem. ⑤ **Share & Propagate**. This phase determines whether contamination spreads laterally, from agent to agent via shared memory (Gu et al., 2024); vertically, from individual users to organizational memory; or temporally, from one session to another through persistent profiles. It captures the risk that a poisoned memory may escape its original context and become reused, inherited, or trusted across broader interaction scopes. ⑥ **Forget & Rollback**. This phase concerns the system’s ability to remove contaminated memories, roll back to a known-safe state, trace the provenance of poisoned entries, and verify the effectiveness of remediation. It determines whether the system can recover after a successful attack. Without snapshots, version diffs, and forensic traceback, defenses remain limited to best-effort prevention and provide no reliable path for post-breach remediation (Zhang et al., 2025a). The memory lifecycle decomposition is also motivated in part by an analogy with human memory, as introduced in Appendix B.

3 Attack Patterns Across the Lifecycle

We organize the attack literature along the lifecycle introduced in Section 2, tracing how a poisoned memory is seeded at WRITE, persists through STORE, reactivates at RETRIEVE, steers behavior at EXECUTE, spreads via SHARE & PROPAGATE, and survives incomplete cleanup during FORGET & ROLLBACK.

3.1 Write: Seeding Poisoned Memory

Within memory lifecycle, the WRITE phase is where memory poisoning attack first enters the system. Table 2 organizes write-path attacks by decreasing attacker privilege, ranging from corpus write access to query-only interaction and, finally, environment-only influence. Along this trajectory, both persistence and scope have expanded while attack effectiveness has remained high.

Some representative works mark the conceptual shifts in this gradient. AgentPoison (Chen et al., 2024) demonstrates that a corpus-level adversary can plant embedding-space backdoors with under 0.1% poison rate and over 80% attack success. MINJA (Dong et al., 2025) demonstrates that query-only interaction suffices to permanently alter memory state. eTAMP (Zou et al., 2026) closes the privilege gap entirely by showing that an attacker need not interact with the agent and only has to manipulate a web page the agent encounters during normal operation. Together, these results point to a clear defensive implication: *the effective write-stage attack surface includes not only explicit memory updates, but also any observable context that can influence what the agent decides to store.*

3.2 Store: Provenance, Retention, and Audit in Long-Term Memory

The STORE phase is a critical control point in long-term memory security because it determines how written entries are indexed, compressed, retained and evicted before later retrieval or execution. First, indexing determines which future queries can reactivate a memory entry, making retrieval keys, semantic labels, and storage-priority assignments potential adversarial amplifiers (Tian et al., 2026; Chen et al., 2024). Second, retention, decay, and eviction policies jointly regulate both availability and confidentiality. Retention schemes based on access frequency or recency may inadvertently keep adversarial entries alive while discarding legitimate ones. Furthermore, without automatic expiration or




Figure 2: The six-phase memory lifecycle of an LLM agent. Each phase introduces distinct security questions.

user-accessible deletion, sensitive content can accumulate in memory indefinitely, creating growing privacy exposure over time. (Zhong et al., 2024; Li et al., 2025b; Zhang et al., 2024; Zou et al., 2026). Third, versioning and audit determine whether post-breach forensics, rollback, and verifiable forgetting are possible; provenance, write logs, snapshots, and diff-auditable histories are prerequisites for tracing poisoned entries and their downstream effects (Li et al., 2025b; Zhang et al., 2025a). Thus, memory security must be anchored in storage-time provenance, versioning, auditability, and policy-aware retention, rather than retrofitted only at retrieval or execution time.

3.3 Retrieve and Execute: From Selection to Action Steering

RETRIEVE is the chokepoint that reconnects past writes with present use, allowing a memory entry written weeks earlier to re-enter the active context window and influence the agent’s current behavior. MCFA (Xu et al., 2026) shows that a sufficiently salient retrieved memory can override the user’s explicit instruction and dictate the tools the agent invokes, their invocation order, and their arguments, highlighting the risk that poisoned memory can usurp the agent’s operational autonomy. MemoryGraft (Srivastava and He, 2025) extends this retrieval-to-execution threat beyond explicit tool control by showing that poisoned procedural memories, once retrieved as prior “successful experiences,” can be reused as task-solving strategies and silently replace the agent’s behavior without a

separate trigger.

Taken together, the WRITE, STORE, RETRIEVE, and EXECUTE phases yield the end-to-end attack chain. A manipulated web observation encountered during routine browsing can be summarized as a memory entry and persisted to long-term storage. Days or weeks later, in a different session and possibly a different task context, the card is retrieved and silently steers tool invocation. The poisoning window closes before execution begins, so single-turn detection cannot observe both ends of the chain.

3.4 Share & Propagate: Internal Channels Dominate Leakage

In multi-agent and shared-state systems, contamination can spread through inter-agent messages, shared memory stores, and tool-mediated data flows, where poisoned content is passed as arguments to external tools or written into shared artifacts. (Gu et al., 2024) shows that a single adversarial input can propagate exponentially through pairwise exchanges, reaching approximately one million multimodal agents. (Men et al., 2025) demonstrates analogous propagation through dialogue; and ComPromptMized (Cohen et al., 2025) constructs zero-click self-replicating prompts that traverse GenAI-powered email assistants, with each recipient’s agent absorbing and re-emitting the payload to its own contacts.

Table 2: Write-path attacks ordered by decreasing attacker privilege.

Work	Attacker Access	Activation	Cross Session	Cross Site	Mechanism and Evidence
AgentPoison (Chen et al., 2024)	Corpus write	Optimized backdoor	✓	–	<i>Embedding-space poisoning</i> . Optimizes poisoned memories with $<0.1\%$ poison rate and $\geq 80\%$ ASR.
InjecMEM (Tian et al., 2026)	Single interaction	Topic-conditioned	✓	–	<i>Retriever-agnostic anchoring</i> . Combines anchors with Multi-GCG, reaching 76.6% ASR on MemoryOS.
MINJA (Dong et al., 2025)	Query-only	Victim-term retrieval	✓	–	<i>Self-generated poison</i> . Uses bridging steps and indication prompts to make the agent create the malicious memory.
MemoryGraft (Srivastava and He, 2025)	Ingestion artifact	Trigger-free	✓	–	<i>Procedural-memory grafting</i> . Imitates semantic procedures rather than injecting factual memories.
eTAMP (Zou et al., 2026)	Environment only	Observation-based	✓	✓	<i>Web-observation poisoning</i> . Adversarial web content is absorbed during browsing, producing $8\times$ “frustration” amplification.

3.5 Forget & Rollback: Residual State and Failed Recovery

The FORGET phase is where the system is expected to remove poisoned or sensitive memory state after an attack has been detected. Unlike WRITE or RETRIEVE, the risk in FORGET does not come from introducing or reactivating a new payload, but from failing to fully remove poisoned or sensitive memory state after it has been discovered. Incomplete forgetting arises because a single memory item may leave derivatives across raw dialogue logs, summarized memory cards, vector indexes, reflected lessons, shared stores, and audit records (Packer et al., 2023; Xu et al., 2025; Chhikara et al., 2025; Li et al., 2025b; Rezazadeh et al., 2025). Deleting only the visible entry can therefore leave retrievable residues in compressed summaries, retrieval indices, or propagated copies, allowing contamination to reappear after apparent cleanup (Wang et al., 2026c,a).

4 Defenses: Prevention, Containment, and Recovery

Similar to attacks, defenses can be categorized based on their role in the lifecycle. Just as attacks exploit the sequential nature of memory operations by seeding poison at WRITE, amplifying it through STORE, reactivating it at RETRIEVE, and propagating it via SHARE, defenses must be deployed in a correspondingly distributed fashion, with no single intervention point sufficient to break the full attack chain. This lifecycle-aware perspective motivates a layered defense architecture that addresses prevention, containment, and recovery as complementary, rather than competing, objectives. At WRITE time,

the central objective is to prevent untrusted content from being consolidated into persistent memory. VerificAgent (Nguyen et al., 2025) introduces human-verified freezing, where candidate memories are audited before being committed to a frozen safety contract. MemCube-style memory units (Li et al., 2025b) instead attach source, version, sensitivity, and temporal metadata at write time, making later audit, access control, and rollback possible. STORE-time defenses remain comparatively underdeveloped. Nevertheless, versioned snapshots, write logs, content-addressable records, and compression audits provide the necessary architectural substrate for subsequent recovery. Provenance standards, such as W3C PROV (Moreau and Missier, 2013), further offer a principled serialization layer for representing and managing these lineage-bearing memory objects.

Among the memory lifecycle stages, the RETRIEVE phase has seen the most concentrated development of defensive mechanisms. RobustRAG (Xiang et al., 2026) offers certifiable robustness against retrieval corruption through isolate-then-aggregate generation, substantially reducing the effectiveness of poisoning attacks. TrustRAG (Zhou et al., 2026) and SeCon-RAG (Si et al., 2025) use clustering, trust scoring, semantic similarity, and contradiction-aware filtering to distinguish benign from malicious retrieved content. RevPRAG (Tan et al., 2025b) detects poisoning through activation differences between poisoned and non-poisoned generations. A-MemGuard (Wei et al., 2025) is more memory-native: it retrieves semantically related memories, compares independent reasoning paths, treats disagreement as an

anomaly signal, and distills detected attacks into lesson memories for future hardening. EXECUTE-time defenses ask not what enters the context window, but what the model may do with it. CaMeL, FIDES, and PCAS enforce information-flow control by separating data and control planes and restricting the influence of untrusted content on tool use or program flow (Debenedetti et al., 2026; Costa et al., 2025; Palumbo et al., 2026). Progent, IsolateGPT, and type-directed privilege separation further reduce blast radius through fine-grained tool policies, execution isolation, and typed privilege boundaries (Shi et al., 2025; Wu et al., 2025c; Jacob et al., 2025).

SHARE-time defenses shift the focus from content-level inspection to principal-aware access and interaction modeling. Collaborative Memory (Rezazadeh et al., 2025) models multi-user and multi-agent access through a time-evolving principal-resource graph, while BlindGuard (Miao et al., 2026) uses unsupervised graph-level anomaly detection to flag previously unseen multi-agent contagion patterns. Finally, defenses for the FORGET & ROLLBACK phase remain comparatively sparse. RAGForensics (Zhang et al., 2025a) shows that post-hoc traceback is feasible for static RAG corpora, and (Gudur and Rajkumar, 2026) provides an industry reference for integrity checks, policy enforcement, snapshots, and rollback primitives. However, end-to-end evaluations of these defenses remain limited for dynamically updated agent memories, especially for verifying deletion, ensuring reliable rollback, and supporting forensic traceback.

5 Verifiable Memory Governance: An Operational Framework

After surveying the landscape of attacks and defenses in the preceding sections, we arrive at a central question: *what verifiable mechanisms must a long-term-memory Agent or LLM system provide in order to maintain auditable, recoverable control over its own memory state, and thereby support trustworthy user operations such as control, tracing, isolation, recovery, and deletion?* We refer to this requirement as **Verifiable Memory Governance (VMG)**. In this work, we identify five core primitives that support Verifiable Memory Governance:

- **Write Authorization (WA)** requires that every entry committed to long-term memory be attributable to an authenticated source and pass

an explicit authorization check before consolidation, preventing untrusted external content from entering memory as if user-endorsed.

- **Provenance Visibility (PV)** requires that every memory entry carry a queryable, lineage-complete provenance record tracing back to its originating write event through any intermediate summarization or merging steps, enabling post-breach attribution and forensic analysis.
- **Principal-Scoped Retrieval (PS)** requires that the retrieval function return only entries whose authorized scope includes the querying principal, preventing cross-user or cross-agent memory leakage even under adversarial queries.
- **Rollbackability (RB)** requires that the system maintain versioned snapshots and write logs sufficient to restore the memory store to a known-safe state at any prior point in time, enabling recovery after a successful attack.
- **Verified Forgetting (VF)** requires that after a deletion operation, the system can demonstrate through post-deletion membership tests that the target content is no longer recoverable from any substrate—including raw logs, compressed summaries, vector indices, and propagated copies.

As summarized in Table 3 in Appendix, for each primitive, we specify the representative threat it is intended to withstand, the corresponding defense direction, and a proposed evaluation metric.

Formal definition. To make the five primitives evaluable, we cast them as predicates over a memory store M_t at time t . Each entry $m \in M_t$ carries metadata $(\text{src}(m), \text{scope}(m), \text{prov}(m), \text{ver}(m))$ denoting source principal, authorized principal set, provenance record, and version identifier. M_t denotes the complete state of the system’s memory store at time t . Formally, we define the following criteria:

$$\mathbf{WA}(M_t) := \forall m \in M_t, \text{src}(m) \neq \perp \quad (1)$$

$$\wedge \text{auth}_W(\text{src}(m), m, t) = 1,$$

where $\perp$ denotes an undefined or missing source principal, and auth_W denotes an explicit write-time authorization check.

$$\mathbf{PV}(M_t) := \forall m \in M_t, \text{Queryable}(m) \quad (2)$$

$$\wedge \text{LineageComplete}(m),$$

which requires every memory entry to maintain queryable provenance traceable to its original write event, with lineage preserved across summarization, merging, or distillation.

$$\begin{aligned} \mathbf{PS}(M_t) &:= \forall q, \pi, R(q, \pi, M_t) \\ &\subseteq \{m \in M_t : \pi \in \text{scope}(m)\}, \end{aligned} \quad (3)$$

which requires that, for any query q and requesting principal π , the output of $R(q, \pi, M_t)$ be strictly restricted to entries for which $\pi \in \text{scope}(m)$.

$$\begin{aligned} \mathbf{RB}(M_t) &:= \forall t_0 < t, \\ \text{rollback}(M_t, t_0) &= M_{t_0}, \end{aligned} \quad (4)$$

which requires that the system be able to roll back the current memory store to an exact prior state M_{t_0} for every $t_0 < t$.

$$\begin{aligned} \mathbf{VF}_\varepsilon(M_t) &:= \forall X, \Pr_{q \sim Q_X} [\text{Expose}_X(q, F_X(M_t)) = 1] \\ &\leq \varepsilon, \end{aligned} \quad (5)$$

which bounds the probability that any probing query q can re-expose target content X after deletion by F_X to at most ε . $F_X(\cdot)$ denotes the deletion operator removing target content X from the memory store. Q_X is the distribution over adversarial probing queries whose intent is to recover X . $\text{Expose}_X(q, M')$ is a binary indicator that equals 1 if the post-deletion store $M' = F_X(M_t)$ returns a response that sufficiently reconstructs X .

A system satisfies ε -VMG at time t if all five hold simultaneously:

$$\mathbf{MS}_\varepsilon(M_t) := \mathbf{WA} \wedge \mathbf{PV} \wedge \mathbf{PS} \wedge \mathbf{RB} \wedge \mathbf{VF}_\varepsilon. \quad (6)$$

Under lineage-based deletion the primitives form a design pre-order:

$$\mathbf{VF}_\varepsilon \preceq \mathbf{RB} \preceq \mathbf{PV} \preceq \mathbf{WA}, \quad (7)$$

where $X \preceq Y$ means implementing X requires Y as an architectural precondition; $\mathbf{PS}$ is orthogonal to this lineage chain but necessary for any confidentiality claim. Specifically, $\mathbf{PV}$ depends on $\mathbf{WA}$ because lineage-complete provenance requires every write event to carry an authenticated source principal $\text{src}(m) \neq \perp$; without write-time authorization, provenance records may exist but



Figure 3: The five primitives as a dependency tower.

cannot be anchored to a verified origin, rendering the lineage incomplete and post-breach attribution unreliable.

Figure 3 renders the dependency as a tower: the widest, most mature foundation is WA; each successive layer narrows and becomes less deployed. This dependency structure most clearly shows why VMG must be treated as a lifecycle property rather than as a final layer added post hoc. However, existing literature suggests that PV, which underpins both RB and VF, remains rare in published architectures. This observation yields a key diagnostic insight: near-term progress should focus on building provenance infrastructure before pursuing higher-level primitives whose architectural preconditions are not yet in place.

6 Conclusion and Future Work

As LLM agents increasingly rely on persistent memory (Yang et al., 2026b), cross-session long-term memory introduces risks beyond conventional input-centric or single-turn security frameworks. We propose a Memory Lifecycle Framework organizing attacks and defenses across six phases and four objectives, showing that memory threats unfold as cross-phase chains where poisoned content can persist, propagate, and resist cleanup. We further introduce Verifiable Memory Governance, specifying five primitives for auditable and recoverable memory control, and highlight key gaps in provenance infrastructure, cross-principal propagation, and post-deletion verification.

A critical future direction is to design benchmarks that assess LTM security across the entire memory lifecycle (Hao et al., 2026), including cross-session contamination, share-time propagation, and forget-time recovery, all of which are still only partially covered by existing benchmarks (Zhang et al., 2025c; Debenedetti et al., 2024; Zhang et al., 2025b; Zhan et al., 2024; Wu et al., 2025a; Maharana et al., 2024; Tan et al., 2025a; Hu et al., 2026; Yoon et al., 2026; Guo et al., 2026).

Limitations

The scope and findings of this survey should be understood in light of several limitations. First, long-term memory security for LLM agents is a rapidly evolving area, and many relevant systems, attacks, and benchmarks keep emerging. Although we aim to cover representative work across attacks, defenses, and governance, our taxonomy may not fully capture recent or unpublished developments, especially in industrial memory systems whose internal designs are not publicly documented.

Second, our Memory Lifecycle Framework is intended as an organizational scaffold rather than a strict decomposition of all memory operations. In practice, the boundaries between phases can be blurred: retrieval and execution may be tightly coupled, storage and summarization may occur continuously, and sharing may implicitly happen through tool calls, logs, or external artifacts. Therefore, some attacks and defenses may span multiple phases and can be categorized differently depending on the system architecture.

Acknowledgments

This work was supported by the Shanghai Special Fund for Promoting High-Quality Industrial Development: Pilot Industry Innovation and Development Program (Artificial Intelligence Track, ID: 2025-GZL-RGZN-02006).

References

- Scott Barnett, Stefanus Kurniawan, Srikanth Thudumu, Zach Brannelly, and Mohamed Abdelrazek. 2024. [Seven failure points when engineering a retrieval augmented generation system](#). In *Proceedings of the IEEE/ACM 3rd International Conference on AI Engineering (CAIN '24)*, pages 194–199. Association for Computing Machinery.
- Andreea-Elena Bodea, Stephen Meisenbacher, Alexandra Klymenko, and Florian Matthes. 2026. SoK: Privacy risks and mitigations in retrieval-augmented generation systems. In *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*.
- Lucas Bourtole, Varun Chandrasekaran, Christopher A. Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. 2021. [Machine unlearning](#). In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 141–159. IEEE.
- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. 2021. [Extracting training data from large language models](#). In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650. USENIX Association.
- Harsh Chaudhari, Giorgio Severi, John Abascal, Anshuman Suri, Matthew Jagielski, Christopher A. Choquette-Choo, Milad Nasr, Cristina Nita-Rotaru, and Alina Oprea. 2024. [Phantom: General backdoor attacks on retrieval augmented language generation](#). *arXiv preprint arXiv:2405.20485*.
- Qizhi Chen, Chao Qi, Yihong Huang, Muquan Li, Rongzheng Wang, Dongyang Zhang, Ke Qin, and Shuang Liang. 2026. [KEPo: Knowledge evolution poison on graph-based retrieval-augmented generation](#). In *Proceedings of the ACM Web Conference 2026*, pages 2308–2319. Association for Computing Machinery.
- Zhaorun Chen, Zhen Xiang, Chaowei Xiao, Dawn Song, and Bo Li. 2024. [AgentPoison: Red-teaming LLM agents via poisoning memory or knowledge bases](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 130185–130213. Curran Associates, Inc.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. Mem0: Building production-ready AI agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*.
- Stav Cohen, Ron Bitton, and Ben Nassi. 2025. [Here comes the AI worm: Preventing the propagation of adversarial self-replicating prompts within GenAI ecosystems](#). In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security, CCS '25*, pages 3975–3989. ACM.
- Manuel Costa, Boris Köpf, Aashish Kolluri, Andrew Paverd, Mark Russinovich, Ahmed Salem, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. 2025. [Securing AI agents with information-flow control](#). *arXiv preprint arXiv:2505.23643*.
- Edoardo DeBenedetti, Ilia Shumailov, Tianqi Fan, Jamie Hayes, Nicholas Carlini, Daniel Fabian, Christoph Kern, Chongyang Shi, Andreas Terzis, and Florian Tramèr. 2026. [Defeating prompt injections by design](#). In *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. Preprint: arXiv:2503.18813.
- Edoardo DeBenedetti, Jie Zhang, Mislav Balunović, Luca Beurer-Kellner, Marc Fischer, and Florian Tramèr. 2024. [AgentDojo: A dynamic environment to evaluate prompt injection attacks and defenses for LLM agents](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 82895–82920. Curran Associates, Inc. Datasets and Benchmarks Track.
- Shen Dong, Shaochen Xu, Pengfei He, Yige Li, Jiliang Tang, Tianming Liu, Hui Liu, and Zhen Xiang. 2025.

- Memory injection attacks on LLM agents via query-only interaction. In *Advances in Neural Information Processing Systems*, volume 38, pages 52119–52153. Neural Information Processing Systems Foundation, Inc.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2024. From local to global: A graph RAG approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Faouzi El Yagoubi, Godwin Badu-Marfo, and Ranwa Al Mallah. 2026. [AgentLeak: A benchmark for internal-channel privacy leakage in multi-agent LLM systems](#). *IEEE Access*, 14:94960–94978.
- Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. [Not what you’ve signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection](#). In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, AISec ’23, pages 79–90. ACM.
- Xiangming Gu, Xiaosen Zheng, Tianyu Pang, Chao Du, Qian Liu, Ye Wang, Jing Jiang, and Min Lin. 2024. [Agent Smith: A single image can jailbreak one million multimodal LLM agents exponentially fast](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 16647–16672. PMLR.
- Vaishnavi Gudur and Anshul Rajkumar. 2026. [OWASP Agent Memory Guard: A runtime defense and open benchmark for memory poisoning in LLM agents \(ASI06\)](#). Software, version 0.3.1. Apache-2.0; released 25 August 2026.
- Jiahe Guo, Xiangran Guo, Yulin Hu, Zimo Long, Xingyu Sui, Xuda Zhi, Yongbo Huang, Hao He, Weixiang Zhao, Yanyan Zhao, and Bing Qin. 2026. When personalization legitimizes risks: Uncovering safety vulnerabilities in personalized dialogue agents. *arXiv preprint arXiv:2601.17887*.
- Xixuan Hao, Zeyu Zhang, Zehao Lin, Yihang Sun, Ziliang Guo, Xichong Zhang, Yuxuan Liang, Feiyu Xiong, and Zhiyu Li. 2026. [MemOps: Benchmarking lifecycle memory operations in long-horizon conversations](#). *arXiv preprint arXiv:2607.12893*.
- Feng He, Tianqing Zhu, Dayong Ye, Bo Liu, Wanlei Zhou, and Philip S. Yu. 2025. [The emerged security and privacy of LLM agent: A survey with case studies](#). *ACM Computing Surveys*, 58(6):1–36. Article 162.
- Yu He, Yifei Chen, Yiming Li, Shuo Shao, Leyi Qi, Boheng Li, Dacheng Tao, and Zhan Qin. 2026. [External data extraction attacks against retrieval-augmented large language models](#). *IEEE Transactions on Information Forensics and Security*, 21:5864–5879.
- Manuel Herrador and Johann Rehberger. 2026. [SpAIware: Uncovering a novel artificial intelligence attack vector through persistent memory in LLM applications and agents](#). *Future Generation Computer Systems*, 174:107994.
- Yuanzhe Hu, Yu Wang, and Julian McAuley. 2026. [Evaluating memory in LLM agents via incremental multi-turn interactions](#). In *International Conference on Learning Representations (ICLR)*. ArXiv:2507.05257.
- Yuyang Hu, Shichun Liu, Yanwei Yue, Guibin Zhang, Boyang Liu, Fangyi Zhu, Jiahang Lin, Honglin Guo, Shihan Dou, Zhiheng Xi, Senjie Jin, Jiejun Tan, Yanbin Yin, Jiongnan Liu, Zeyu Zhang, Zhongxiang Sun, Yutao Zhu, Hao Sun, Boci Peng, and 28 others. 2025. [Memory in the age of AI agents](#). *arXiv preprint arXiv:2512.13564*.
- Dennis Jacob, Emad Alghamdi, Zhanhao Hu, Basel Alomair, and David Wagner. 2025. [Preventing prompt injection with type-directed privilege separation](#). *arXiv preprint arXiv:2509.25926*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Xinfeng Li, Dong Huang, Jie Li, Hongyi Cai, Zhenhong Zhou, Wei Dong, XiaoFeng Wang, and Yang Liu. 2025a. [A vision for access control in LLM-based agent systems](#). *arXiv preprint arXiv:2510.11108*.
- Zhiyu Li, Chenyang Xi, Chunyu Li, Ding Chen, Boyu Chen, Shichao Song, Simin Niu, Hanyu Wang, Jiawei Yang, Chen Tang, Qingchen Yu, Jihao Zhao, Yezhaohui Wang, Peng Liu, Zehao Lin, Pengyuan Wang, Jiahao Huo, Tianyi Chen, Kai Chen, and 20 others. 2025b. [MemOS: A memory OS for AI system](#). *arXiv preprint arXiv:2507.03724*.
- Jiacheng Liang, Yuhui Wang, Changjiang Li, Tanqiu Jiang, Rongyi Zhu, Neil Gong, and Ting Wang. 2026. [GraphRAG under fire](#). In *2026 IEEE Symposium on Security and Privacy (SP)*, pages 195–212. IEEE.
- Jiafeng Liang, Hao Li, Chang Li, Jiaqi Zhou, Shixin Jiang, Zekun Wang, Changkai Ji, Zhihao Zhu, Runxuan Liu, Tao Ren, Jinlan Fu, See-Kiong Ng, Xia Liang, Ming Liu, and Bing Qin. 2025. [AI meets brain: Memory systems from cognitive neuroscience to autonomous agents](#). *arXiv preprint arXiv:2512.23343*.
- Yupei Liu, Yuqi Jia, Runpeng Geng, Jinyuan Jia, and Neil Zhenqiang Gong. 2024. [Formalizing and benchmarking prompt injection attacks and defenses](#). In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 1831–1847, Philadelphia, PA. USENIX Association.

- Elizabeth F. Loftus. 2005. [Planting misinformation in the human mind: A 30-year investigation of the malleability of memory](#). *Learning & Memory*, 12(4):361–366.
- Jinghao Luo, Yuchen Tian, Chuxue Cao, Ziyang Luo, Hongzhan Lin, Kaixin Li, Chuyi Kong, Ruichao Yang, and Jing Ma. 2026. [From storage to experience: A survey on the evolution of LLM agent memory mechanisms](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 41622–41652, San Diego, California, United States. Association for Computational Linguistics.
- Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. [Evaluating very long-term conversational memory of LLM agents](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13851–13870, Bangkok, Thailand. Association for Computational Linguistics.
- Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C Lipton, and J Zico Kolter. 2024. TOFU: A task of fictitious unlearning for LLMs. In *Conference on Language Modeling (COLM)*. ArXiv:2401.06121.
- Tianyi Men, Pengfei Cao, Zhuoran Jin, Yubo Chen, Kang Liu, and Jun Zhao. 2025. A troublemaker with contagious jailbreak makes chaos in honest towns. In *Proceedings of the 63rd Annual Meeting of the ACL (Volume 1: Long Papers)*. ArXiv:2410.16155.
- Rui Miao, Yixin Liu, Yili Wang, Xu Shen, Yue Tan, Yiwei Dai, Shirui Pan, and Xin Wang. 2026. [Blind-Guard: Safeguarding LLM-based multi-agent systems under unknown attacks](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 39215–39234, San Diego, California, United States. Association for Computational Linguistics.
- Niloofar Miresghallah, Neal Mangaokar, Narine Kokhlikyan, Arman Zharmagambetov, Manzil Zaheer, Saeed Mahloujifar, and Kamalika Chaudhuri. 2026. [CIMemories: A compositional benchmark for contextual integrity in LLMs](#). In *International Conference on Learning Representations*, volume 2026, pages 95513–95532.
- Luc Moreau and Paolo Missier. 2013. [PROV-DM: The PROV data model](#). W3c recommendation, World Wide Web Consortium. Editors; published 30 April 2013.
- Yanming Mu, Hao Hu, Feiyang Li, Qiao Yuan, Jiang Wu, Zichuan Liu, Pengcheng Liu, Mei Wang, Hongwei Zhou, and Yuling Liu. 2026. Towards secure retrieval-augmented generation: A comprehensive review of threats, defenses and benchmarks. *arXiv preprint arXiv:2603.21654*.
- Karim Nader, Glenn E Schafe, and Joseph E LeDoux. 2000. Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval. *Nature*, 406(6797):722–726.
- Milad Nasr, Nicholas Carlini, Chawin Sitawarin, Sander V. Schulhoff, Jamie Hayes, Michael Ilie, Juliette Pluto, Shuang Song, Harsh Chaudhari, Ilia Shumailov, Abhradeep Guha Thakurta, Kai Yuanqing Xiao, Andreas Terzis, and Florian Tramèr. 2026. [The attacker moves second: Stronger adaptive attacks bypass defenses against LLM jailbreaks and prompt injections](#). In *35th USENIX Security Symposium (USENIX Security 26)*, pages 1467–1486, Baltimore, MD. USENIX Association.
- Thong Q Nguyen, Shubhang Desai, Raja Hasnain Anwar, Firoz Shaik, Vishwas Suryanarayanan, and Vishal Chowdhary. 2025. VerificAgent: Domain-specific memory verification for scalable oversight of aligned computer-use agents. *arXiv preprint arXiv:2506.02539*.
- Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G Patil, Ion Stoica, and Joseph E Gonzalez. 2023. MemGPT: Towards LLMs as operating systems. *arXiv preprint arXiv:2310.08560*.
- Nils Palumbo, Sarthak Choudhary, Jihye Choi, Guy Amir, Prasad Chalasani, and Somesh Jha. 2026. [Formal policy enforcement for real-world agentic systems](#). *Preprint*, arXiv:2602.16708.
- Sidharth Pulipaka, Oliver Chen, Manas Sharma, Taaha S Bajwa, Vyas Raina, and Ivaxi Sheth. 2026. Persist-Bench: When should long-term memories be forgotten by LLMs? *arXiv preprint arXiv:2602.01146*.
- Alireza Rezazadeh, Zichao Li, Ange Lou, Yuying Zhao, Wei Wei, and Yujia Bao. 2025. Collaborative memory: Multi-user memory sharing in LLM agents with dynamic access control. *arXiv preprint arXiv:2505.18279*.
- Henry L Roediger III, Michelle L Meade, and Erik T Bergman. 2001. Social contagion of memory. *Psychonomic Bulletin & Review*, 8(2):365–371.
- Mark Russinovich, Ahmed Salem, and Ronen Eldan. 2025. Great, now write an article about that: The Crescendo multi-turn LLM jailbreak attack. In *USENIX Security Symposium*.
- Daniel L. Schacter, Donna Rose Addis, and Randy L. Buckner. 2007. [Remembering the past to imagine the future: The prospective brain](#). *Nature Reviews Neuroscience*, 8(9):657–661.
- Daniel L Schacter, Kenneth A Norman, and Wilma Koutstaal. 1998. The cognitive neuroscience of constructive memory. *Annual Review of Psychology*, 49:289–318.
- Tianneng Shi, Jingxuan He, Zhun Wang, Hongwei Li, Linyu Wu, Wenbo Guo, and Dawn Song. 2025. [Progent: Securing AI agents with privilege control](#). *arXiv preprint arXiv:2504.11703*.

- Xiaonan Si, Meilin Zhu, Simeng Qin, Lijia Yu, Lijun Zhang, Shuaitong Liu, Xinfeng Li, Ranjie Duan, Yang Liu, and Xiaojun Jia. 2025. [SeCon-RAG: A two-stage semantic filtering and conflict-free framework for trustworthy RAG](#). In *Advances in Neural Information Processing Systems 38*, NeurIPS 2025, pages 78742–78771. Neural Information Processing Systems Foundation, Inc.
- Renan Souza, Amal Gueroudji, Stephen DeWitt, Daniel Rosendo, Tirthankar Ghosal, Robert Ross, Prasanna Balaprakash, and Rafael Ferreira da Silva. 2025. [PROV-AGENT: Unified provenance for tracking AI agent interactions in agentic workflows](#). In *2025 IEEE International Conference on eScience (eScience)*, pages 467–473. IEEE.
- Larry R Squire, Lisa Genzel, John T Wixted, and Richard G Morris. 2015. [Memory consolidation](#). *Cold Spring Harbor Perspectives in Biology*, 7(8):a021766.
- Saksham Sahai Srivastava and Haoyu He. 2025. [MemoryGraft: Persistent compromise of LLM agents via poisoned experience retrieval](#). *arXiv preprint arXiv:2512.16962*.
- Theodore R Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L Griffiths. 2024. Cognitive architectures for language agents (CoALA). *Transactions on Machine Learning Research*.
- Elham Tabassi. 2023. [Artificial intelligence risk management framework \(AI RMF 1.0\)](#). NIST AI 100-1, National Institute of Standards and Technology, Gaithersburg, MD.
- Haoran Tan, Zeyu Zhang, Chen Ma, Xu Chen, Quanyu Dai, and Zhenhua Dong. 2025a. [MemBench: Towards more comprehensive evaluation on the memory of LLM-based agents](#). In *Findings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. ArXiv:2506.21605.
- Xue Tan, Hao Luan, Mingyu Luo, Xiaoyan Sun, Ping Chen, and Jun Dai. 2025b. [RevPRAG: Revealing poisoning attacks in retrieval-augmented generation through LLM activation analysis](#). In *Findings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ArXiv:2411.18948.
- Yaxin Tang, Yijia Liu, Jiahe Lan, Zheng Yan, and Erol Gelenbe. 2026. [Security of LLM-based agents regarding attacks, defenses, and applications: A comprehensive survey](#). *Information Fusion*, 127:103941.
- Hanling Tian, Gengyu Zhang, Zeyang Sha, Jingying Wang, Yuhang Liu, Zhehao Huang, Kun Yang, and Xiaolin Huang. 2026. [InjecMEM: Memory injection attack on LLM agent memory systems](#). *arXiv preprint arXiv:2608.23471*. Accepted at the 3rd Conference on Language Modeling (COLM 2026).
- Bin Wang, Fan Wang, Pingping Wang, Jinyu Cong, Yang Yu, Yilong Yin, Zhongyi Han, and Benzhen Wei. 2026a. [Agentic unlearning: When LLM agent meets machine unlearning](#). *arXiv preprint arXiv:2602.17692*.
- Bo Wang, Weiyi He, Shenglai Zeng, Zhen Xiang, Yue Xing, Jiliang Tang, and Pengfei He. 2025. [Unveiling privacy risks in LLM agent memory](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*. ArXiv:2502.13172.
- Haoyu Wang, Christopher M Poskitt, and Jun Sun. 2026b. [AgentSpec: Customizable runtime enforcement for safe and reliable LLM agents](#). In *Proceedings of the 48th IEEE/ACM International Conference on Software Engineering (ICSE '26)*. ArXiv:2503.18666.
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. 2024. [A survey on large language model based autonomous agents](#). *Frontiers of Computer Science*, 18(6):186345.
- Shang Wang, Tianqing Zhu, Dayong Ye, and Wanlei Zhou. 2026c. [When machine unlearning meets retrieval-augmented generation \(RAG\): Keep secret or forget knowledge?](#) *IEEE Transactions on Dependable and Secure Computing*, 23(2):2030–2045.
- Qianshan Wei, Tengchao Yang, Yaochen Wang, Xinfeng Li, Lijun Li, Zhenfei Yin, Yi Zhan, Thorsten Holz, Zhiqiang Lin, and XiaoFeng Wang. 2025. [A-MemGuard: A proactive defense framework for LLM-based agent memory](#). *arXiv preprint arXiv:2510.02373*.
- Yule Wen, Yanzhe Zhang, Jianxun Lian, Xiaoyuan Yi, Xing Xie, and Diyi Yang. 2026. [Contextualized privacy defense for LLM agents](#). *arXiv preprint arXiv:2603.02983*.
- Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, and Dong Yu. 2025a. [LongMemEval: Benchmarking chat assistants on long-term interactive memory](#). In *International Conference on Learning Representations (ICLR)*.
- Yaxiong Wu, Sheng Liang, Chen Zhang, Yichao Wang, Yongyue Zhang, Huifeng Guo, Ruiming Tang, and Yong Liu. 2025b. [From human memory to AI memory: A survey on memory mechanisms in the era of LLMs](#). *arXiv preprint arXiv:2504.15965*.
- Yuhao Wu, Franziska Roesner, Tadayoshi Kohno, Ning Zhang, and Umar Iqbal. 2025c. [IsolateGPT: An execution isolation architecture for LLM-based agentic systems](#). In *Network and Distributed System Security Symposium (NDSS)*. ArXiv:2403.04960.
- Chong Xiang, Tong Wu, Zexuan Zhong, David Wagner, Danqi Chen, and Prateek Mittal. 2026. [Certifiably robust RAG against retrieval corruption](#). In *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. ArXiv:2405.15556.

- Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. 2025. A-MEM: Agentic memory for LLM agents. *arXiv preprint arXiv:2502.12110*. NeurIPS 2025.
- Zhenlin Xu, Xiaogang Zhu, Yu Yao, Minhui Xue, and Yiliao Song. 2026. From storage to steering: Memory control flow attacks on LLM agents. *arXiv preprint arXiv:2603.15125*.
- Jiaqi Xue, Mengxin Zheng, Yebowen Hu, Fei Liu, Xun Chen, and Qian Lou. 2024. **BadRAG: Identifying vulnerabilities in retrieval augmented generation of large language models**. *Preprint*, arXiv:2406.00083.
- Tiankai Yang, Jiate Li, Yi Nian, Shen Dong, Ruiyao Xu, Ryan Rossi, Kaize Ding, and Yue Zhao. 2026a. No attacker needed: Unintentional cross-user contamination in shared-state LLM agents. *arXiv preprint arXiv:2604.01350*.
- Xianglin Yang, Yufei He, Shuo Ji, Bryan Hooi, and Jin Song Dong. 2026b. **Zombie agents: Persistent control of self-evolving llm agents via self-reinforcing injections**. *arXiv preprint arXiv:2602.15654*.
- Sangyeon Yoon, Sunkyoung Kim, Hyesoo Hong, Wonje Jeung, Yongil Kim, Wooseok Seo, Heuiyeon Yeen, and Albert No. 2026. **BenchPreS: A benchmark for context-aware personalized preference selectivity of persistent-memory LLMs**. *arXiv preprint arXiv:2603.16557*.
- Yosif Zaki and Denise J Cai. 2025. Memory engram stability and flexibility. *Neuropsychopharmacology*, 50:285–293.
- Qiusi Zhan, Zhixiang Liang, Zifan Ying, and Daniel Kang. 2024. **InjecAgent: Benchmarking indirect prompt injections in tool-integrated large language model agents**. In *Findings of the Association for Computational Linguistics (ACL)*. ArXiv:2403.02691.
- Baolei Zhang, Haoran Xin, Minghong Fang, Zhuqing Liu, Biao Yi, Tong Li, and Zheli Liu. 2025a. **Traceback of poisoning attacks to retrieval-augmented generation**. In *Proceedings of The Web Conference (WWW)*. ArXiv:2504.21668.
- Baolei Zhang, Haoran Xin, Jiatong Li, Dongzhe Zhang, Minghong Fang, Zhuqing Liu, Lihai Nie, and Zheli Liu. 2025b. **Benchmarking poisoning attacks against retrieval-augmented generation**. *arXiv preprint arXiv:2505.18543*.
- Hanrong Zhang, Jingyuan Huang, Kai Mei, Yifei Yao, Zhenting Wang, Chenlu Zhan, Hongwei Wang, and Yongfeng Zhang. 2025c. **Agent security bench (ASB): Formalizing and benchmarking attacks and defenses in LLM-based agents**. In *International Conference on Learning Representations (ICLR)*. ArXiv:2410.02644.
- Shuning Zhang, Lyumanshan Ye, Xin Yi, Jingyu Tang, Bo Shui, Haobin Xing, Pengfei Liu, and Hewu Li. 2024. **“Ghost of the past”: identifying and resolving privacy leakage from LLM’s memory through proactive user interaction**. *Preprint*, arXiv:2410.14931.
- Zeyu Zhang, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. 2025d. **A survey on the memory mechanism of large language model-based agents**. *ACM Trans. Inf. Syst.*, 43(6).
- Zeyu Zhang, Ziliang Guo, Yihang Sun, Xichong Zhang, Xixuan Hao, Zehao Lin, Yang Zhang, Xiaoyan Zhao, Tong Shen, Bo Tang, Zhi-Qin John Xu, Junchi Yan, Haofen Wang, Xu Chen, Feiyu Xiong, Zhiyu Li, and Tat-Seng Chua. 2026. **Metis: Memory foundation model**. *arXiv preprint arXiv:2607.26760*.
- Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. 2024. **MemoryBank: Enhancing large language models with long-term memory**. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Huichi Zhou, Kin-Hei Lee, Zhonghao Zhan, Zhenhao Li, Yue Chen, Huaxiu Yao, Hamed Haddadi, and Emine Yilmaz. 2026. **TrustRAG: Enhancing robustness and trustworthiness in retrieval-augmented generation**. In *AAAI 2026 Workshop on Trust and Control in Agentic AI (TrustAgent)*, Singapore.
- Wei Zou, Mingwen Dong, Miguel Romero Calvo, Shuaichen Chang, Jiang Guo, Dongkyu Lee, Xing Niu, Xiaofei Ma, Yanjun Qi, and Jiarong Jiang. 2026. **Poison once, exploit forever: Environment-injected memory poisoning attacks on web agents**. *arXiv preprint arXiv:2604.02623*.
- Wei Zou, Runpeng Geng, Binghui Wang, and Jinyuan Jia. 2025. **PoisonedRAG: Knowledge corruption attacks to retrieval-augmented generation of large language models**. In *USENIX Security Symposium*. ArXiv:2402.07867.

Appendix

A Architectural Governance Primitives.

Table 3 summarizes the five VMG primitives from Section 5, specifying for each the representative threat, corresponding defense direction, and proposed evaluation metric.

B Human-Memory Bridge / Conceptual Grounding

Our lifecycle memory view is grounded in a limited analogy with human memory. Cognitive neuroscience frames memory not as a fixed archive but as a reconstructive and reconsolidatable system where experiences are encoded, consolidated, reactivated, and sometimes rewritten as context and goals evolve (Squire et al., 2015; Schacter et al., 1998, 2007; Nader et al., 2000; Zaki and Cai, 2025). This perspective underscores why long-term memory in LLM agents cannot be reduced to mere context extension or passive storage (Zhang et al., 2026). Profile memories, episodic logs, summaries, and retrieval stores determine what persists, what is later reactivated, how traces are compressed, and how the agent interprets its own past. The security question is therefore not only whether a single input is malicious, but whether an untrusted trace can be written, consolidated, retrieved, elevated, propagated, or deleted across the memory lifecycle.

The analogy also clarifies the primitives needed for Verifiable Memory Governance (VMG). In human memory, misinformation often operates through source-monitoring failure, confidence inflation, reconsolidation after recall, and social contagion rather than direct overwrite (Loftus, 2005; Roediger III et al., 2001; Nader et al., 2000). The corresponding agent-level failures include provenance failure, read-time rewriting through summarization or reflection, and cross-agent contamination. Injected content may be misattributed as the agent’s own experience (Dong et al., 2025; Tian et al., 2026), promoted through compression into more authoritative lesson memory (Packer et al., 2023; Xu et al., 2025), or propagated via shared memory and inter-agent communication (Gu et al., 2024; Men et al., 2025). Robust LTM security therefore needs to operate across the full lifecycle. Write authorization and provenance visibility protect the origin of memory, principal-scoped retrieval limits who can reactivate it, and rollbacka-

bility together with verified forgetting determines whether corrupted or sensitive traces can be removed after they have spread.

C Methodology for Paper Selection

We adopt a systematic literature search strategy that combines literature retrieval with citation-based expansion. We searched major academic databases and repositories, including the ACL Anthology, ACM Digital Library, IEEE Xplore, USENIX proceedings, OpenReview, and arXiv. The search was conducted using a set of predefined keyword categories, covering topics such as memory poisoning / injection, long-term / episodic / agent memory, RAG poisoning, memory extraction, multi-agent memory, memory governance, forgetting, and rollback.

We further expand the candidate set by tracing relevant references and citations through Google Scholar, arXiv, and Semantic Scholar. In total, we screened more than 200 candidate works, among which approximately 90 were included in the survey.

D VMG operationalization.

Equations (1)–(7) formalize the VMG criteria and their relationships. Possible examples of how each primitive could be operationalized into concrete, measurable metrics are as follows:

- WA: reproduce published injection attacks (e.g., MINJA, eTAMP) and measure the survival rate of injected content consolidated into memory.
- PV: issue provenance queries against each memory entry and measure the proportion returning a complete provenance chain.
- PS: issue retrieval queries under different principal identities against the same memory store and measure the proportion of results exceeding the requester’s authorized scope (cross-principal leakage rate).
- RB: snapshot historical states during operation, trigger rollback, and compare against the target state via entry-level consistency checks, yielding both the restoration accuracy and the time overhead to complete rollback.

Table 3: Architectural governance primitives.

Primitive	Representative Threat	Representative Defense	Proposed Metric
Write Authorization (WA)	Query-induced injection; environment-injected poisoning (Dong et al., 2025; Tian et al., 2026; Zou et al., 2026)	Typed write validation; human or policy arbitration (Nguyen et al., 2025; Jacob et al., 2025; Palumbo et al., 2026)	Adaptive injection-survival rate at bounded false-positive rate (FPR) (Nasr et al., 2026)
Provenance Visibility (PV)	Source-monitoring failure; imitation-based grafting (Srivastava and He, 2025)	Provenance-gated retrieval; lineage-preserving compression (Moreau and Missier, 2013; Souza et al., 2025)	Fraction of entries traceable to source event
Principal-Scoped Retrieval (PS)	Cross-user contamination; black-box extraction (Yang et al., 2026a; Wang et al., 2025; El Yagoubi et al., 2026; He et al., 2026)	Capability- or label-based retrieval policy (Li et al., 2025a; Costa et al., 2025; Mireshghallah et al., 2026)	Cross-principal leakage rate under scripted query suites
Rollbackability (RB)	Compression-amplified toxins; behavioral drift (Packer et al., 2023; Xu et al., 2025)	RAGForensics-style traceback adapted to agent memory (Zhang et al., 2025a)	Time-to-remediation; fraction of toxic entries localized
Verified Forgetting (VF)	Residual derivatives; reappearance in summaries (Zhang et al., 2024; Pulipaka et al., 2026)	Coordinated deletion across raw, summary, index, weights (Bourtoule et al., 2021; Maini et al., 2024; Wang et al., 2026c,a)	Post-deletion membership / reappearance tests (Maini et al., 2024; Wang et al., 2026c)

- VF: use tiered queries (direct, paraphrased, associative, multi-turn) with a judge LLM to detect reconstruction of deleted content.

E Real-System Analysis.

To provide a qualitative system-level illustration, we analyze two representative memory systems, Mem0 (Chhikara et al., 2025) and MemOS (Li et al., 2025b), using the VMG framework and their published descriptions:

- Mem0 provides no support for WA, RB, or VF. It offers partial support for PV, reflected in its change-history tracking for individual memory entries, and partial support for PS, reflected in its ability to isolate reads at the user/agent granularity.
- MemOS offers partial support for WA, recording origin tags such as an Origin Signature in its MemCube metadata; full support for PV, via an explicitly defined provenance mechanism; and partial support for PS, through its access-control design. Versioning and rollback mechanisms give it partial support for RB, though without empirical validation against adaptive adversarial attacks, while VF remains unsupported, as no verifiable post-deletion validation capability is provided.