When Retrieval Fails Before It Begins: Structurally Indirect Prerequisite Eviction as a Retention Failure in Agentic Memory

Minkyu Song [1]

Abstract

Agentic memory under a fixed budget involves two stages: *retention* and *retrieval*. Existing retrieval-centered paradigms implicitly assume necessary evidence survives eviction, but we challenge this by isolating a pre-retrieval failure mode: *structurally indirect prerequisite eviction*, in which upstream blocks weakly aligned with the query are discarded under budget pressure. We provide an operational definition of this failure, a reproducible deterministic benchmark, and per-seed trace diagnostics. Finally, we evaluate *Dependency-aware Semantic Garbage Collection* (DSGC), a one-hop graph-aware rule. In our main suite, DSGC improves full-chain retention from 0.03 to 0.90 under a lexical encoder and from 0.23 to 1.00 under a sentence encoder. Robustness checks then identify the budget and scaling regimes where the one-hop rule holds or degrades. Our released pipeline and failure postmortem support mechanistic analysis of retention before retrieval as a distinct failure boundary.

1. Introduction

Long-horizon LLM agents face two logically distinct decisions about past context under a fixed budget: *retention* (which blocks survive eviction) and *retrieval* (which surviving blocks are surfaced for the current query). Existing work has primarily emphasized retrieval—how to rank or expand evidence once a store is available (Lewis et al., 2020; Guu et al., 2020; Borgeaud et al., 2022; Izacard et al., 2023; Edge et al., 2024; Gutiérrez et al., 2024). We study an earlier failure boundary: the store may already be incomplete because a necessary block was evicted before retrieval begins.

[1]Yonsei University, Seoul, South Korea. Correspondence to: Minkyu Song <smkgenesis@yonsei.ac.kr>.

Accepted at the ICML 2026 Workshop on *Failure Modes of Agentic AI* (FAGEN@ICML 2026). Non-archival; the authors retain all rights.

Common memory-management heuristics include keeping recent tokens, scoring blocks by similarity to the current query, or paging memories in and out heuristically (Packer et al., 2023; Wang et al., 2023a; Park et al., 2023). These policies are effective when the needed evidence is recent or lexically aligned with the query. However, they can fail in a specific way when the answer depends on an upstream block that is weakly aligned with the query but remains structurally required. We call such a block a *structurally indirect prerequisite*, and we make this precise in Section 2. Prerequisite chains of this kind arise naturally wherever an upstream fact determines the actionability of a downstream decision. Once a structurally indirect prerequisite has been evicted under budget pressure, no downstream retrieval procedure can recover it from the store. Retention is therefore logically prior to retrieval in the failure chain: retrieval ranks what remains; retention decides what remains.

This paper treats that retention-stage failure as a first-class object of study. Concretely, we

1. give an *operational definition* that distinguishes the failure from downstream retrieval failure, non-structural forgetting, and reasoning failure (Section 2);
2. construct a deterministic benchmark as a *reproducible trigger* with fixed seeds, retrieval-friendly controls, and dependency-sensitive targets (Section 4);
3. provide *trace diagnostics* that decompose each failing seed into a displaced block, a displacing competitor, and a score margin (Section 6);
4. evaluate *Dependency-aware Semantic Garbage Collection* (DSGC), a one-hop graph-aware retention rule that protects prerequisites via structural reachability, and characterize the regimes where it holds and breaks (Sections 3, 5–7).

Together, these components establish a controlled base case: with supplied dependency edges, retention can be studied as a structural reachability problem before retrieval begins, separately from graph induction and downstream retrieval.

1.1. Positioning

We position the problem earlier in the agentic-memory pipeline than retrieval: before evidence can be ranked or expanded, it must first survive eviction. We contrast this setting with three established paradigms and draw our primary inspiration from a fourth.

Retrieval and graph-augmented retrieval. Standard retrieval-augmented systems (Lewis et al., 2020; Guu et al., 2020; Borgeaud et al., 2022; Izacard et al., 2023) assume the relevant evidence is still present in the store at query time. Graph-augmented retrieval methods such as GraphRAG and HippoRAG (Edge et al., 2024; Gutiérrez et al., 2024) use graph structures to expand or re-rank evidence over a store that is already assumed to be complete. They focus on downstream utilization, not upstream preservation.

Agentic memory systems. Memory systems including MemGPT, SCM, Generative Agents, Reflexion, Voyager, and LATS (Packer et al., 2023; Wang et al., 2023a; Park et al., 2023; Shinn et al., 2023; Wang et al., 2023b; Zhou et al., 2024) demonstrate that external memory management is critical for long-horizon tasks, but their eviction policies are usually not designed for the case studied here: preserving weakly query-aligned prerequisites under budget pressure. Similarly, context-compression methods (Mu et al., 2023; Jiang et al., 2024) change how content is represented inside the budget rather than deciding which semantic units survive eviction; unlike block-level eviction, compression preserves a trace of every block and leaves inter-block structural relations intact.

Systems-level garbage collection. We draw framing inspiration from tracing garbage collection (McCarthy, 1960; Jones et al., 2011), and the name *Dependency-aware Semantic Garbage Collection* reflects this analogy directly. In tracing GC, a root set of live objects propagates reachability to all objects they reference; unreachable objects are collected (evicted). DSGC mirrors this structure at the block level: query-facing blocks with high semantic relevance serve as the root set, liveness propagates one hop along declared prerequisite edges to structurally necessary predecessors, and blocks that receive neither direct semantic support nor structural coverage are treated as collectable. The *semantic* qualifier marks that the root set is selected by query similarity rather than by syntactic pointers; the *dependency-aware* qualifier marks that liveness is not purely similarity-based but propagates through the prerequisite graph—the structural extension that vanilla semantic selection omits. This framing treats agentic context eviction as a reachability-style retention problem rather than only a semantic ranking problem.

2. Operational Definition

2.1. Setting and Terminology

A *memory store* contains N blocks $\mathcal{B} = \{b_1, \ldots, b_N\}$. A *block* is a semantically atomic unit—a single, indivisible proposition—that is either retained or evicted in full; it carries a token cost s_i. In our experiments, s_i is a deterministic whitespace word count, $\max(|\text{text.split}()|, 1)$, used consistently by all reported budgets and reproduction scripts, ensuring cross-tokenizer reproducibility. The benchmark also supplies a directed prerequisite graph E. An edge $(j, i) \in E$ means block b_j depends on block b_i, so b_i is an *immediate prerequisite* of b_j. The benchmark supplies prerequisite edges as ground truth, allowing the experiment to test retention independently of graph induction.

Given a query q and a token budget $C < \sum_i s_i$, each retention policy induces a ranking over blocks. The retained subset $S \subseteq \mathcal{B}$ is then constructed greedily: blocks are added in policy order until adding the next block would make $\sum_{b_i \in S} s_i > C$. Retrieval downstream of retention is restricted to S: any block not in S is unavailable to downstream retrieval from the active prompt-visible store for this query step. We use $\text{sim}_\phi(q, b_i)$ for the similarity score between query q and block b_i under encoder ϕ, and write $\text{chain}(q)$ for the generator-labeled set of blocks whose joint retention is necessary to recover the benchmark answer to q.

2.2. Failure Mode

The pre-retrieval failure mode has two components: a block's structural vulnerability (*structurally indirect prerequisite*) and the eviction event that realizes it (*retention-stage failure*).

Definition 1 (Structurally indirect prerequisite). A chain block $b_i \in \text{chain}(q)$ is *structurally indirect with respect to q under encoder ϕ* when there exists another chain block b_j that depends on it and that has higher query similarity:

$$\exists\, b_j \in \text{chain}(q) : \ (j, i) \in E \wedge \text{sim}_\phi(q, b_i) < \text{sim}_\phi(q, b_j). \tag{1}$$

Intuitively, b_i is a necessary chain block whose text is less query-aligned than the downstream block b_j it enables. Indirectness is a property of the (block, query, encoder) tuple, not of any policy.

Definition 2 (Retention-stage failure). A retention policy exhibits a *retention-stage failure* on (q, C) when it produces a selection S that excludes a structurally indirect prerequisite while retaining the more query-similar downstream block that depends on it. The failure occurs before retrieval: once b_i leaves the active prompt-visible store, downstream ranking cannot recover it.

What this excludes. Definition 2 excludes three adjacent phenomena. (i) *Retrieval failure*: the needed block is in S but the retriever ranks it below an unhelpful competitor. (ii) *Surface forgetting*: b_i is dropped because it is simply older or because no chain of any kind connects it to q. (iii) *Reasoning failure*: S is complete but the model misuses it. We isolate the case in which structural necessity and surface similarity disagree at the eviction step.

Worked example. Consider three blocks: c_1: "Alice is assigned to backup bundle M"; c_2: "Bundle M maps to clearance N"; c_3: "Clearance N permits deletion of the encrypted backup." Edges (c_3, c_2) and (c_2, c_1) encode the prerequisite chain; the query is "May Alice delete the encrypted backup?" Under a lexical encoder, $\mathrm{sim}_\phi(q, c_3)$ is high ("delete", "backup"), $\mathrm{sim}_\phi(q, c_1)$ is moderate ("Alice"), but $\mathrm{sim}_\phi(q, c_2)$ is low because c_2 shares no surface terms with the query. Here, c_2 satisfies Definition 1 with witness $b_j = c_3$: it is a necessary chain block yet sits in a "semantic valley" between two more relevant blocks. Under budget pressure, a similarity-only policy identifies c_2 as the primary eviction candidate, destroying the logical bridge between Alice and her permission; DSGC treats c_3 as a liveness root and propagates survival support to c_2, preserving the chain.

3. Method: One-Hop Dependency Propagation

We study a minimal one-hop graph-aware retention rule, Dependency-aware Semantic Garbage Collection (DSGC). Rather than directly searching for "garbage" to evict, DSGC propagates survival support from highly relevant downstream blocks to their structurally necessary prerequisites. The rule excludes learned components, multi-hop diffusion, and per-step reranking, making the effect of one-hop dependency propagation directly testable.

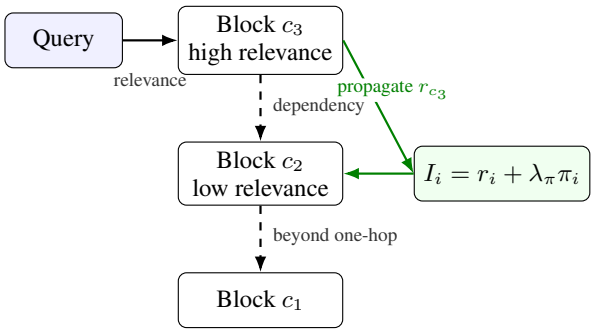

Figure 1. The DSGC mechanism. A highly relevant downstream block (c_3) elevates the retention score of its structurally necessary but semantically indirect prerequisite (c_2) via the propagation term π_i. The one-hop propagation targets retention, while multi-step inference remains the role of the downstream reasoner. The policy halts after one hop; vulnerabilities beyond this radius (e.g., c_1) constitute the residual failure mode analyzed in Section 6.

Scoring rule. For query embedding $\mathbf{g} = \phi(q)$ and block embedding $\mathbf{e}_i = \phi(b_i)$, raw relevance is normalized over all N blocks in the current store:

$$r_i = \frac{\exp(\mathbf{g}^\top \mathbf{e}_i/\tau)}{\sum_{k=1}^{N} \exp(\mathbf{g}^\top \mathbf{e}_k/\tau)}, \quad \tau = 0.25.$$

We fix $\tau = 0.25$ before running the main suite. The one-hop propagation term lifts relevance from a downstream block to its immediate prerequisite:

$$\pi_i = \sum_{j:(j,i)\in E} r_j.$$

The *DSGC score* combines the two additively:

$$I_i = r_i + \lambda_\pi \pi_i, \qquad \lambda_\pi = 1.0.$$

The main suite uses $\lambda_\pi = 1.0$ as the equal-weight base case: one unit of downstream relevance contributes one unit of propagated prerequisite support. Section 7 reports a λ_π sweep around this value. Selection is greedy in descending I_i until the additive token budget C is exhausted; score ties preserve the deterministic seeded block order used to instantiate the scenario. The released policy implementation also supports a length-density variant; in all reported experiments we set the density exponent to 0.0, so score differences reflect semantic relevance and propagated prerequisite support rather than length normalization.

Three policies. We compare three retention policies to cleanly ablate the effect of graph propagation:

- **Similarity-only:** scores blocks by raw relevance r_i alone and greedily retains them until budget C is exhausted. This is the standard semantic baseline that ignores graph edges.
- **No-graph DSGC:** runs the DSGC framework with $\lambda_\pi = 0$ (so $I_i = r_i$); because exp is monotone, this preserves the exact ranking of the similarity-only baseline while executing within the unified DSGC code path.
- **DSGC:** the complete proposed rule with $\lambda_\pi = 1.0$.

This three-way comparison separates code-path effects from the structural propagation term: similarity-only and no-graph DSGC test ranking equivalence, while no-graph DSGC versus DSGC tests the effect of $\lambda_\pi \pi_i$.

Procedure. The five-step implementation follows directly from the scoring rule: encode all blocks and the query, compute softmax relevance r_i, propagate one hop to form π_i, combine into I_i, and greedily select under budget C.

Require: blocks $\mathcal{B} = \{b_i\}$, query q, budget C, edges E, weight λ_π, encoder ϕ

1: $\mathbf{g} \leftarrow \phi(q)$; $\mathbf{e}_i \leftarrow \phi(b_i)$ **for all** $b_i \in \mathcal{B}$
2: $r_i \leftarrow \exp(\mathbf{g}^\top \mathbf{e}_i/\tau) \;/\; \sum_k \exp(\mathbf{g}^\top \mathbf{e}_k/\tau) \quad (\tau = 0.25)$
3: $\pi_i \leftarrow \sum_{j:\,(j,i)\in E} r_j$ **for all** $b_i \in \mathcal{B}$
4: $I_i \leftarrow r_i + \lambda_\pi \pi_i$ **for all** $b_i \in \mathcal{B}$
5: Sort $\mathcal{B}$ descending by I_i; greedily retain b_i while $\sum_{b_i \in S} s_i \le C$

Ensure: retained subset $S \subseteq \mathcal{B}$

For N blocks of embedding dimension d and $|E|$ edges, scoring costs $\mathcal{O}(Nd)$, propagation adds $\mathcal{O}(|E|)$, and selection is $\mathcal{O}(N \log N)$. This complexity profile characterizes the retention rule itself; Section 8 discusses the production-systems implications.

Dependency graph. The benchmark supplies prerequisite edges directly, allowing us to test retention independently of graph induction. This separation is useful because trace-rich agent settings (Park et al., 2023; Shinn et al., 2023) suggest a plausible route to future dependency induction, but do not remove it as a separate problem. Our question here is therefore mechanistic: given dependency edges, how much retention improvement can a minimal graph-aware propagation rule produce?

Encoders. We evaluate two encoders: an L2-normalized, 256-dimensional lexical encoder as a strict structural stress test for surface-overlap vulnerabilities, and *all-MiniLM-L6-v2* (Wang et al., 2020; Reimers & Gurevych, 2019), producing 384-dimensional normalized dense embeddings. The lexical setting stresses surface-overlap failure directly; the dense setting tests whether the same retention-stage failure persists under stronger semantic embeddings. We report both because the failure persists across encoder regimes, though the quantitative dynamics differ.

4. Reproducible Trigger

A reproducible trigger requires two ingredients: a test regime where the failure mode reliably surfaces under budget pressure, and a control regime where the structurally indirect prerequisite pattern (Definition 1) is absent. Comparing these regimes lets us test whether propagation helps specifically when structural necessity and query similarity diverge.

4.1. Controls and Targets

Scenario layout. Each scenario is constructed with a fixed capacity of 20 total blocks—small enough to keep tight budget regimes informative yet large enough to admit real distractor competition. The blocks fall into three distinct functional types:

- **Chain blocks:** The structurally necessary, generator-labeled blocks in $\mathrm{chain}(q)$ required to answer the query.
- **Honeypot blocks:** Semantic distractors (labeled `h1`, `h2`, etc. in traces) that share the query's decisive vocabulary but do not complete the reasoning chain and lack prerequisite edges into the chain. These are designed to reliably bait similarity-only retention policies.
- **Filler blocks:** Background noise (labeled `f1`, `f2`, etc.) generated from neutral templates: short statements about disjoint entities, operations notes, or scheduling events that contain neither the answer relation nor the query's decisive vocabulary.

After the chain and honeypots are injected, the remaining positions in the 20-block context are occupied by filler blocks.

Budget multiplier. Let $L(q) = \sum_{b_i \in \mathrm{chain}(q)} s_i$ be the total token cost of the necessary chain blocks. A *budget multiplier* M sets the retention budget to $C = \lceil M \cdot L(q) \rceil$. We consider:

- $M = 1.0$: critically tight.
- $M = 2.0$: intermediate.
- $M = 5.0$: loose.

We fix $M = 2.0$ for the main suite because it provides enough aggregate capacity for the chain to survive while preserving real competition at the cutoff. At $M = 1.0$, survival depends almost entirely on perfect ranking; at $M = 5.0$, most blocks survive and the policy distinction largely disappears. Section 7 reports all three regimes.

Control template. A scenario is a *control* when its necessary chain blocks share sufficient surface evidence with the query, ensuring that a standard similarity-only policy can reliably retain the chain. Controls test whether similarity-only retention succeeds when the required blocks are query-aligned, and whether DSGC preserves this behavior. We employ two control tasks:

- **Expense Control:** One block states a project budget, another sets an approval threshold, and the query asks whether a proposed expense is allowed.
- **Threshold Control:** The query names a specific thresholded entity and asks whether it crosses a defined numerical limit.

Both scenarios admit a similarity-only solution; a third candidate was rejected for instability during calibration and is preserved in the repository history. The accepted controls serve as ceiling checks: degradation on these controls would indicate that propagation is harming otherwise retrieval-friendly cases.

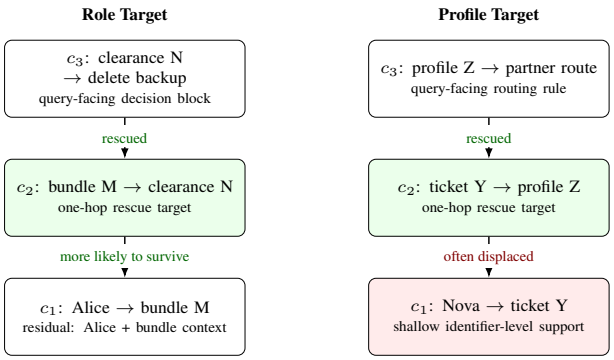

Figure 2. Differential survival of the multi-hop prerequisite (c_1). Both target scenarios share the same structural retention architecture: a query-facing downstream block (c_3) and an immediate prerequisite (c_2) successfully rescued via one-hop propagation. The critical divergence occurs at the oldest block (c_1), which lies beyond the explicit one-hop rescue radius. In the **Role Target** (left), c_1 retains substantial residual lexical support (e.g., user identity and task context), increasing its likelihood of survival. In the **Profile Target** (right), c_1 possesses only shallow identifier-level overlap and lacks the query's core routing vocabulary, making it highly susceptible to displacement under budget pressure. This contrast marks the boundary of the one-hop policy.

Target template. A scenario is a *target* when its reasoning chain contains at least one structurally indirect prerequisite under Definition 1. We evaluate two target architectures based on the three-block chain ($c_1 \leftarrow c_2 \leftarrow c_3$):

- **Role Target** (moderate difficulty): The query asks whether Alice may delete a production backup. The chain consists of an older assignment block (c_1: Alice → bundle), an immediate prerequisite (c_2: bundle → clearance), and a query-facing decision rule (c_3: clearance → delete). Under DSGC, c_3 successfully rescues c_2 via one-hop propagation. The survival of the full chain then depends on whether c_1 retains enough residual lexical support (user identity and bundle context) to stay above the budget cutoff.
- **Profile Target** (high difficulty): Same three-stage architecture, but c_1 is stripped of decision-bearing routing vocabulary and retains only bare identifier-level support. c_2 is still rescued by propagation, while c_1 is far more easily displaced under budget pressure (Figure 2).

Honeypot strength. Our data generator defines three honeypot tier pools—*easy*, *medium*, and *hard*. The main suite uses the medium pool:

- **Easy:** too weak to challenge the cutoff.
- **Medium:** shares decisive query vocabulary but cannot complete the answer chain.
- **Hard:** risks creating confounded alternative paths.

Medium is the strongest adversarial setting that preserves a unique ground-truth chain.

Synthetic templates. The benchmark uses compact synthetic templates—four templates with twenty blocks per scenario—for three reasons:

1. **Ground-truth traceability.** Synthetic templates allow us to explicitly declare $\mathrm{chain}(q)$ as ground truth. This makes Definition 2 directly checkable: a failure decomposes into a specifically known displaced block, rather than manifesting as a generic incorrect answer with an opaque cause.
2. **Controlled regime contrast.** Controls and targets share identical total block counts, budget rules, seed protocols, and encoder configurations. They differ only where the regimes require it: controls are retrieval-friendly, while targets instantiate structurally indirect prerequisites. The contrast is therefore between two explicitly declared regimes, isolating the structural variable of interest.
3. **Falsifiability.** The synthetic setting permits us to lock the generation protocol and achieve bit-exact reproduction of the failure trigger.

4.2. Fixed Evaluation Protocol and Auditability

The release preserves the calibration history, rejected variants, and postmortem notes, making the final trigger auditable rather than only reproducible. Each rejected variant is paired with the specific calibration issue it exposed—ranging from insufficient distractor strength to unstable seed behavior—so that subsequent benchmarks can refine rather than rediscover the trigger.

Fixed evaluation protocol. The final main suite locks all parameters to prevent post-hoc tuning. The grid evaluates:

- **Templates:** 4 (2 controls, 2 targets)
- **Policies:** 3 (similarity-only, no-graph DSGC, DSGC)
- **Encoders:** 2 (lexical, sentence)
- **Configuration:** 15 seeds per condition, 20 total blocks per scenario, medium-strength honeypots, budget multiplier $M = 2.0$.

To ensure fair comparison, the protocol standardizes structural complexity. *Chain depth*—the total number of logically sequential blocks required to answer the query—varies strictly by template type: controls use a chain depth of 2 (with 4 honeypots), while targets use a chain depth of 3 (with 2 honeypots). These values are the smallest chain depths that separate the control and target regimes cleanly. Depth 3 is the smallest chain in which one block (c_2) sits exactly one hop from the query-facing block (c_3) and one block (c_1) lies two hops away—the minimum configuration in which the structural failure of Definition 1 is isolable without confounding recency effects. Depth 2 controls therefore check whether propagation preserves cases where direct query alignment is already sufficient.

Evaluation metrics. The primary metric is **Full Chain Retention**: the binary success rate over seeded runs in which every generator-labeled chain block survives the eviction budget. We also record deterministic template answer accuracy, but in this suite it is nearly a consequence of full-chain retention rather than an independent reasoning metric. Means and standard deviations are reported over 15 seeds.

5. Main Results

Encoder	Regime	Method	Retention
Lexical	Control	Sim.-only	1.00 $\pm$0.00
Lexical	Control	No-graph DSGC	1.00 $\pm$0.00
Lexical	Control	DSGC	1.00 $\pm$0.00
Lexical	Target	Sim.-only	0.03 $\pm$0.18
Lexical	Target	No-graph DSGC	0.03 $\pm$0.18
Lexical	Target	DSGC	**0.90** $\pm$0.31
Sentence	Control	Sim.-only	1.00 $\pm$0.00
Sentence	Control	No-graph DSGC	1.00 $\pm$0.00
Sentence	Control	DSGC	1.00 $\pm$0.00
Sentence	Target	Sim.-only	0.23 $\pm$0.43
Sentence	Target	No-graph DSGC	0.23 $\pm$0.43
Sentence	Target	DSGC	**1.00** $\pm$0.00

Table 1. Regime-level full-chain retention in the fixed main evaluation suite (mean $\pm$ std over 15 seeds). `Sim.-only` denotes the standard similarity-only retention baseline. Bold values indicate target-chain retention under one-hop propagation.

The control-target contrast in Table 1 delivers a coherent picture across both encoders. All three methods achieve perfect retention (1.00) in the control regime: because the necessary chain blocks share sufficient surface evidence with the query, standard similarity-based retention easily preserves them. This saturation indicates that the controls are retrieval-friendly: when chain blocks are query-aligned, similarity-based retention is sufficient, and DSGC does not degrade it.

In the target regime, the divergence between semantic similarity and structural necessity becomes visible. Under the strict lexical encoder, both baselines retain the full chain in only 0.03 of runs, showing the vulnerability of indirect prerequisites under budget pressure. DSGC mitigates this failure, preserving the full reasoning chain in the vast majority of runs (0.90). The same pattern holds under the denser sentence encoder, where DSGC reaches perfect retention (1.00) while the similarity baseline still frequently fails (0.23).

The exact empirical equality between the similarity-only and no-graph DSGC baselines is a mathematical consequence of our controlled implementation: setting $\lambda_\pi = 0$ collapses the DSGC score to raw semantic relevance ($I_i = r_i$). Because exp is monotone, this ranking identity is exact, attributing the observed ranking difference to the one-hop dependency propagation term.

Encoder	Template	Sim.-only	DSGC
Lexical	`role`	0.00 $\pm$0.00	1.00 $\pm$0.00
Lexical	`profile`	0.07 $\pm$0.26	0.80 $\pm$0.41
Sentence	`role`	0.13 $\pm$0.35	1.00 $\pm$0.00
Sentence	`profile`	0.33 $\pm$0.49	1.00 $\pm$0.00

Table 2. Granular breakdown of target template performance. The Profile Target is more difficult than the Role Target under the lexical encoder, consistent with the one-hop limitation illustrated in Figure 2. No-graph DSGC performance identically matches the similarity-only baseline and is therefore omitted.

Resolving the target regime. Table 2 breaks down target performance. Under the lexical encoder, the Profile Target is the only setting where DSGC still fails on some seeds (0.80), matching the predicted one-hop limit: DSGC rescues the immediate prerequisite c_2 from query-facing c_3, but the oldest prerequisite c_1 can remain vulnerable when it has too little residual lexical support. Under the denser sentence encoder, c_1 receives enough semantic support for DSGC to reach perfect retention. The next section tests this boundary at the trace level.

6. Trace Diagnostics

Per-seed traces test a narrow diagnostic prediction: do the residual lexical Profile failures preserve the rescued pair (c_2, c_3) while losing only the oldest prerequisite c_1?

Seed	Ranks (c_1, c_2, c_3)	Top non-chain	Displ.	Margin
3	$(10, 1, 2)$	`h1`	c_1	0.043
8	$(6, 1, 2)$	`f9`	c_1	0.007
14	$(6, 1, 2)$	`f3`	c_1	0.043

Table 3. Diagnostic breakdown of Profile Target failures under the lexical encoder. Prefix `h` denotes a honeypot distractor; `f` denotes a background filler. In all failing instances, the displaced block is exclusively the multi-hop prerequisite c_1. The Margin column reports the final retention score differential between the lowest-retained competitor and c_1; consistently narrow values (< 0.05) indicate marginal budget-cutoff displacements rather than wholesale ranking collapses.

Trace evidence for the one-hop boundary. Our evaluation framework logs per-seed ranking and displacement traces. If the one-hop account is correct, residual DSGC failures should preserve c_2 and c_3 while displacing c_1: the immediate prerequisite (c_2) is rescued by propagation from the query-facing block (c_3), but the structurally distant prerequisite (c_1) is not. Table 3 matches this pattern. In all three failing instances of the lexical Profile Target, c_2 and c_3 occupy the top two ranks; however, the oldest prerequisite c_1 drops to rank 6 or 10 and is subsequently evicted by a non-chain distractor. In the remaining 12 seeds, c_1 retains sufficient residual lexical support to clear the budget cutoff, indicating that the failure is marginal rather than systematic.

Anatomy of the residual failures. Each failure decomposes into a displaced chain block, an evicting non-chain competitor, and a narrow score margin: in all three cases c_1 is edged out by a non-chain block—honeypot `h1` or fillers `f9`, `f3`—with minimal lexical overlap, and the margins stay below 0.05, indicating marginal cutoff displacements rather than wholesale ranking collapses. Seed 8 is informative because c_1 is displaced by a plain background filler rather than a honeypot: the residual failure is not limited to adversarial distractors; ordinary filler blocks can also cross the cutoff.

Generating a falsifiable hypothesis. These traces define a concrete next test for k-hop propagation: a variant with $k \geq 2$ should recover these exact three seeds without degrading the perfect retention of the controls and the Role Target. The narrowness of the observed margins suggests that lightweight structural heuristics—such as bounded neighborhood expansions—may suffice to close this gap without the cost of deep multi-hop diffusion.

7. Robustness and Failure Boundaries

The trace evidence establishes the mechanistic picture; we now test whether the same mechanism holds across propagation weights and budget pressure.

Template	Encoder	λ_π	Retention
`role`	Lexical	0.5	1.00 ±0.00
`role`	Lexical	1.0	1.00 ±0.00
`role`	Lexical	2.0	1.00 ±0.00
`profile`	Lexical	0.5	0.67 ±0.49
`profile`	Lexical	1.0	0.80 ±0.41
`profile`	Lexical	2.0	0.87 ±0.35
`role`	Sentence	0.5–2.0	1.00 ±0.00
`profile`	Sentence	0.5–2.0	1.00 ±0.00

Table 4. Hyperparameter stability on target templates. Mean retention (± std over 15 seeds) across varying propagation weights (λ_π). The results indicate that performance is not sensitive to precise tuning within the tested range.

Hyperparameter stability (λ_π). We evaluate the propagation weight across $\lambda_\pi \in \{0.5, 1.0, 2.0\}$ to verify local sensitivity around the main-suite default ($\lambda_\pi = 1.0$). As shown in Table 4, structural gains remain consistently positive across this range: the Role Target is fully saturated under the lexical encoder, the Profile Target improves monotonically, and the sentence encoder achieves perfect retention at all tested values. These results indicate that the gain is not sensitive to the exact value of λ_π in the tested range.

Template	Mult.	Sim.-only	DSGC
`role`	1.0	0.00 ±0.00	0.40 ±0.51
`role`	2.0	0.00 ±0.00	1.00 ±0.00
`role`	5.0	0.27 ±0.46	1.00 ±0.00
`profile`	1.0	0.00 ±0.00	0.53 ±0.52
`profile`	2.0	0.07 ±0.26	0.80 ±0.41
`profile`	5.0	0.60 ±0.51	1.00 ±0.00

Table 5. Budget pressure dynamics under the lexical encoder. Full-chain retention (mean ± std over 15 seeds) across varying budget multipliers (M), where $C = \lceil M \sum_{b_i \in \text{chain}(q)} s_i \rceil$. DSGC separates most clearly from similarity-only retention in the intermediate regime ($M = 2.0$), where similarity-based retention fails.

Budget pressure dynamics. Table 5 maps retention behavior across three budget regimes:

- **Tight** ($M = 1.0$): Both policies struggle, as the minimal budget cannot comfortably encompass the full reasoning chain alongside competitive distractors.
- **Intermediate** ($M = 2.0$): The main separation regime. DSGC retains the structurally necessary chain while similarity-only retention does not.
- **Loose** ($M = 5.0$): The performance gap narrows because the budget is large enough to passively retain nearly the entire context store.

Notably, the similarity-only curve is not strictly monotonic on the Role Target, indicating that larger budgets do not automatically repair structurally misranked prerequisites when competitive distractors remain near the cutoff. This pattern suggests that dependency-aware retention is most useful when the budget is large enough to hold the chain but small enough for distractors to compete at the cutoff.

Sliding window baseline. We also compare against a recency-based sliding-window policy. This baseline retains the full chain in only 0.07 of runs across all four templates under the lexical encoder at $M = 2.0$, failing even on the elementary control scenarios. In this benchmark, recency is therefore neither a reliable proxy for relevance nor a sufficient mechanism for preserving logical prerequisite chains. The main comparison focuses on structured retention policies, since the diagnostic question concerns how plausible memory policies rank indirect prerequisites under budget pressure.

Template	Encoder	20 blocks	50 blocks
`role`	Lexical	1.00 ±0.00	0.73 ±0.46
`profile`	Lexical	0.80 ±0.41	0.33 ±0.49
`role`	Sentence	1.00 ±0.00	1.00 ±0.00
`profile`	Sentence	1.00 ±0.00	1.00 ±0.00

Table 6. Context scaling limitations. Full-chain retention for DSGC at 20 versus 50 total context blocks (mean ± std over 15 seeds). The data shows material degradation under the sparse lexical encoder as context size increases, while the dense sentence encoder remains saturated in this setting.

Context scaling and negative results. Context scaling reveals the clearest boundary of the current one-hop rule. As shown in Table 6, when expanding from 20 to 50 total blocks, DSGC under the dense sentence encoder remains perfectly saturated (1.00). However, performance under the sparse lexical encoder degrades materially: retention on the Role Target falls from 1.00 to 0.73, and the Profile Target drops from 0.80 to 0.33.

A likely mechanism is score diffusion: as the candidate pool expands, the baseline lexical relevance mass spreads across a broader array of distractors, forcing the fixed one-hop propagation signal into tighter competition against a much larger field of weakly similar background blocks at the budget cutoff. In these 50-block tests, the one-hop rule remains saturated with dense embeddings but degrades under sparse lexical encoding.

8. Discussion

What this establishes. Structurally indirect prerequisite eviction is a well-defined retention-stage failure mode. In this benchmark, it is reproducible across seeds, visible in per-seed traces, and mitigated by one-hop dependency propagation. Control saturation indicates that the gain is targeted rather than a blanket score inflation, while residual failures mark the expected beyond-one-hop boundary.

Boundaries and negative results. By using synthetic templates with supplied dependency edges, the benchmark isolates retention from graph induction, multi-step planning, and tool use. The 50-block lexical result is an important negative case: it shows where the current one-hop rule breaks. Failure-mode analysis is most useful when it marks both the operating regime and the failure boundary.

Error modes under imperfect graphs. The score $I_i = r_i + \lambda_\pi \pi_i$ degrades asymmetrically under graph errors. A missing prerequisite edge removes propagation support from one block, reducing DSGC to the baseline similarity ranking for that block. A spurious edge instead elevates an unrelated block through π_i, where it can displace a true chain block at the budget cutoff. This asymmetry suggests a design preference for high-precision graph induction: missing edges remove support, whereas spurious edges can actively elevate unrelated blocks.

End-to-end answer evaluation. The deterministic template answer accuracy in our suite is by construction a function of full-chain retention: when every $b_i \in \mathrm{chain}(q)$ survives, the templated answer is recoverable from S, which is why we report retention as the primary metric. End-to-end evaluation with an LLM reasoner tests a separate axis: whether a model can use a complete retained chain correctly, rather than whether the chain survives eviction at all.

The multi-hop extension. The trace diagnostics suggest a k-hop extension to close the residual failures in Table 3. Moving to $k > 1$ pushes propagation toward $\mathcal{O}(k|E|)$ and introduces score-diffusion artifacts that may require damping or normalization. These choices affect the algorithmic trade-off rather than only implementation complexity. Multi-hop variants will need to balance structural reasoning depth against production latency budgets and, in realistic agent settings, eventually connect to future graph induction pipelines built from tool-use or reasoning traces (Park et al., 2023; Shinn et al., 2023).

Systems profile. The one-hop restriction also preserves a favorable systems profile: r_i is a dense query–block inner product, π_i is a single sparse pass implementable as SpMV or scatter-add, and retention reduces to score ordering plus a top-k budgeted selection step, yielding $\mathcal{O}(Nd + |E| + N \log N)$ total cost. When the graph is sparse, the structural overhead follows declared prerequisite edges rather than an $\mathcal{O}(N^2)$ all-pairs interaction pattern, leaving production latency measurements to a separate systems evaluation. Graph induction can also be amortized outside the per-step retention path: a tool-call trace in which a lookup result gates a downstream action directly records a prerequisite edge that a lightweight parser can index at ingestion and reuse across later retention cycles.

9. Conclusion

We isolated a retention-stage failure mode—the eviction of structurally indirect prerequisites—and provided an operational definition, a fixed reproducible trigger, trace-level diagnostics, and a minimal tested mitigation. Together, these results support a broader principle: before retrieval can succeed, retention must keep the reasoning chain live.

More broadly, agentic memory is better understood as a dynamic managed system than as a static vector store. The operational question shifts from "Which block is most similar to the query?" to "Which blocks must remain live for the reasoning chain to stay intact?" DSGC instantiates this idea in a one-hop, systems-friendly form, showing that following prerequisite edges can substantially close the gap between surface-similarity retention and chain-intact reasoning.

Reproducibility. The public release at `https://github.com/smkgenesis/dsgc` contains the full artifact trail—protocol, postmortem, calibration, main suite, and robustness pack—with a `reproduce.sh` entrypoint.

Use of AI assistance. The author used large language model assistants during code implementation and manuscript preparation. All research design, analyses, and conclusions are the author's responsibility.

References

Borgeaud, S., Mensch, A., Hoffmann, J., Cai, T., Rutherford, E., Millican, K., van den Driessche, G. B. M., et al. Improving language models by retrieving from trillions of tokens. In *International Conference on Machine Learning (ICML)*, pp. 2206–2240, 2022.

Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., Metropolitansky, D., Ness, R. O., and Larson, J. From local to global: A Graph RAG approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2024.

Gutiérrez, B. J., Shu, Y., Gu, Y., Yasunaga, M., and Su, Y. HippoRAG: Neurobiologically inspired long-term memory for large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

Guu, K., Lee, K., Tung, Z., Pasupat, P., and Chang, M.-W. REALM: Retrieval-augmented language model pretraining. In *International Conference on Machine Learning (ICML)*, pp. 3929–3938, 2020.

Izacard, G., Lewis, P., Lomeli, M., Hosseini, L., Petroni, F., Schick, T., Dwivedi-Yu, J., Joulin, A., Riedel, S., and Grave, E. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251):1–43, 2023.

Jiang, H., Wu, Q., Luo, X., Li, D., Lin, C.-Y., Yang, Y., and Qiu, L. LongLLMLingua: Accelerating and enhancing LLMs in long context scenarios via prompt compression. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1658–1677, 2024.

Jones, R., Hosking, A., and Moss, E. *The Garbage Collection Handbook: The Art of Automatic Memory Management*. CRC Press, 2011.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., and Kiela, D. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

McCarthy, J. Recursive functions of symbolic expressions and their computation by machine, part I. *Communications of the ACM*, 3(4):184–195, 1960.

Mu, J., Li, X. L., and Goodman, N. D. Learning to compress prompts with gist tokens. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Packer, C., Wooders, S., Lin, K., Fang, V., Patil, S. G., Stoica, I., and Gonzalez, J. E. MemGPT: Towards LLMs as operating systems. *arXiv preprint arXiv:2310.08560*, 2023.

Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., and Bernstein, M. S. Generative agents: Interactive simulacra of human behavior. In *ACM Symposium on User Interface Software and Technology (UIST)*, 2023.

Reimers, N. and Gurevych, I. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of EMNLP-IJCNLP*, 2019.

Shinn, N., Cassano, F., Berman, E., Gopinath, A., Narasimhan, K., and Yao, S. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Wang, B., Liang, X., Yang, J., Huang, H., Wu, S., Wu, P., Lu, L., Ma, Z., and Li, Z. SCM: Enhancing Large Language Model with Self-Controlled Memory Framework. *arXiv preprint arXiv:2304.13343*, 2023a. Accepted by DASFAA 2025 main conference.

Wang, G., Xie, Y., Jiang, Y., Mandlekar, A., Xiao, C., Zhu, Y., Fan, L., and Anandkumar, A. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023b.

Wang, W., Wei, F., Dong, L., Bao, H., Yang, N., and Zhou, M. MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Zhou, A., Yan, K., Shlapentokh-Rothman, M., Wang, H., and Wang, Y.-X. Language agent tree search unifies reasoning, acting and planning in language models. In *International Conference on Machine Learning (ICML)*, pp. 62138–62160, 2024.