
The Context Gathering Decision Process: A POMDP Framework for Agentic Search

Chinmaya Kausik*
University of Michigan
ckausik@umich.edu

Adith Swaminathan
Netflix
aswaminathan@netflix.com

Nathan Kallus
Netflix
nkallus@netflix.com

Abstract

Large Language Model (LLM) agents are deployed in complex environments – such as massive codebases, enterprise databases, and conversational histories – where the relevant state far exceeds their context windows. To navigate these spaces, an agent must iteratively explore the environment to find relevant information. However, without explicit infrastructure, an agent’s working memory can degrade into lossy representations of the search state, resulting in redundant work (e.g. repetitive looping) and premature stopping. In this work, we formalize this challenge as the Context Gathering Decision Process (CGDP), a specialized Partially Observable Markov Decision Process, where an agent’s objective is to adaptively refine its belief state to isolate the necessary information for a task. We model an LLM’s behavior as approximate Thompson Sampling within this CGDP, and introduce a predicate-based method that decomposes an LLM’s implicit search into explicit and modular operations. We then derive two plug-and-play interventions for iterative LLM agents: a persistent, predicate-based belief state that bounds context while preserving multi-hop reasoning, and a programmatic exhaustion gate that halts unproductive search without premature stopping. Across four methods and three question-answering domains, we empirically validate that replacing an LLM’s implicit state with our CGDP-motivated belief state improves multi-hop reasoning by up to 11.4%; while the modular programmatic exhaustion detection saves up to 39% of tokens without any degradation in agent performance. Ultimately, we argue that framing the LLM agent loop as a CGDP can guide the design of modular, non-interfering improvements to agentic search harnesses.

1 Introduction

Large Language Model (LLM) agents are increasingly deployed on complex, real-world environments where the relevant state far exceeds their reliable working context. Deep-research agents search the web [44]; coding assistants search repositories and run shell tools [41]; support agents retrieve from enterprise knowledge bases [16]; all these systems face the same fundamental challenge: the agent cannot load the entire environment into its prompt. Instead, it must iteratively interact with an observation function – such as a Python REPL, a search engine API, or a vector database – to gather the necessary information.

Despite active research into extending LLM context windows [17, 38], long-context models still exhibit position sensitivity [21], degrade with hard negatives [11], are unreliable in multi-turn interaction [14], and stop prematurely during long-horizon search [43]. Thus, building agentic harnesses – infrastructure that allows the LLM to actively manage and search external context – has been studied extensively [33, 37, 42, 26, 29].

*Work initiated as an intern at Netflix.

Agentic harnesses typically append raw observations directly to the LLM prompt [33, 37, 42]; over long horizons, this implicit state tracking leads to severe failures: agents lose track of their original objective, fall into repetitive query loops, hallucinate parametric knowledge instead of corpus evidence, and fail to reliably recognize when a search is completely exhausted. For instance, a coding agent that is resolving a software bug may repeatedly retrieve the same file, oscillating between identical hypotheses without realizing that its search has stagnated.

To address these failures, we formalize the interactive information-seeking agent loop as the Context Gathering Decision Process (CGDP). The CGDP is a specialized Partially Observable Markov Decision Process (POMDP) [13] where the hidden state is the vast external corpus, actions are tool calls, observations are retrieved information, and the objective is to adaptively refine the agent’s belief state to identify task-relevant information from the hidden state. As shown in Table 1, this mathematical framework abstracts the details of many modern agents.

Table 1: Across diverse agentic applications, the underlying challenge remains to navigate a massive hidden state via a constrained observation function to satisfy a user query, which we model as a Context Gathering Decision Process (CGDP).

Application	Hidden State	Action Space (Observation Function)
Code Assistants	Codebase	Linux File System & BASH Commands
Agentic Memory	Conversational History	Vector Database Retrieval
Support Agents / RAG	Knowledge-Base	SQL / BM25 / Embedding Queries
Recursive LMs	Python Data Objects	Python REPL & Code Interpreter
Deep Research	World Wide Web	Search Engine APIs (e.g. Google)

Through the lens of CGDP, we model an LLM’s behavior as approximate Thompson Sampling [35, 31, 25], where the model implicitly samples a hypothesis and takes information-gathering actions conditioned on it. To make this process explicit, we introduce Predicate-Based Adaptive Identification (PBAI), an abstract algorithm that decomposes agentic search into modular operations: assess stopping, select action, observe, and update belief. By mapping state-of-the-art agent harnesses to PBAI, we can pinpoint precisely where their implicit reasoning lacks the necessary infrastructure for reliable, long-horizon search.

Based on our framework, we derive two modular interventions for agentic search harnesses:

1. The Predicate-Based Belief State: A persistent, explicitly managed data structure that forces the agent to extract findings and track open questions iteratively, bounding the context footprint while preserving multi-hop reasoning performance.
2. The Programmatic Exhaustion Gate: A stopping mechanism that halts unproductive search. Rather than relying on an LLM’s self-assessment (which may be poorly calibrated [4, 9]), the mechanism uses programmatic signals like action similarity and observation novelty to prevent redundant looping and premature stopping.

We empirically validate these interventions across four agent search harnesses and three domains (multi-session conversational Question-Answering, multi-hop Wikipedia QA, and code-repository QA). We find that replacing an LLM’s implicit search state with the predicate-based belief state never degrades agent performance and improves multi-hop reasoning by up to 11.4%. Furthermore, the programmatic exhaustion gate safely reduces token consumption by up to 39% on stateful harnesses without sacrificing task accuracy. Ultimately, we demonstrate that formalizing LLM agentic search as a CGDP provides a blueprint for designing reliable modular harness improvements.

2 Related Work

Our contributions are at the intersection of retrieval-augmented generation, long-horizon agentic search, and LLM meta-cognition. While prior work has empirically improved specific components of the agent loop, we provide a unifying framework to understand how these components interact.

Iterative RAG and Agent Harnesses. Multi-round methods generally accumulate state implicitly through a growing context window. IRCoT [37] interleaves retrieval with chain-of-thought reasoning, Iter-RetGen [33] uses previous generations to guide subsequent retrieval queries, and ReAct [42]

tracks history through structured Thought-Action-Observation trajectories. Although these methods perform agentic search, their implicit state tracking frequently degrades over long episodes. More recent approaches introduce explicit per-retrieval interventions: Self-RAG [2] adds retrieval-time self-reflection, Corrective RAG [40] corrects retrieval queries based on relevance assessment and FAIR-RAG [1] introduces gap analysis via an evidence checklist. Similarly, MemGPT [26] provides LLMs with explicit memory management tools. While approaches like StateAct [29] or FAIR-RAG [1] rely on LLM-managed summaries or checklists, our framework demonstrates that *orchestrator-enforced*, strictly curated belief states outperform an LLM’s self-managed tool use or memory edits.

Long-Horizon Search and Corpus Organization. To address unbounded context, offline memory organization methods like A-MEM [39], HippoRAG [8], HopRAG [20] and GraphRAG [5] structure the underlying corpus into graphs to facilitate easier retrieval, whereas in our work we organize the agent’s understanding of what it has found from the corpus online. For long-horizon online search, SLIM [43] separates search from browsing and periodically summarizes trajectories to manage context, while AggAgent [15] generates parallel retrieval trajectories and synthesizes them on demand. These systems support our core premise: reliability improves when the external context is actively managed rather than passively appended.

Stopping Criteria and Metacognition. A critical challenge in iterative search is knowing when to stop. To detect whether to initiate retrieval, there are models such as FLARE [10] (using token-level confidence scores) and DRAGIN [34] (using attention-based scores). In contrast, our work addresses post-retrieval stagnation. Relying on LLMs to self-assess stopping criteria has proven brittle; recent evidence shows that LLMs cannot reliably self-correct reasoning without external feedback [9] and that their verbalized confidence is poorly calibrated [4]. Motivated by this, our programmatic exhaustion gate replaces LLM self-assessment with heuristic stagnation signals and improves token efficiency without degrading search accuracy.

3 Framework

To understand why iterative LLM agents often fail at long-horizon search, we seek to separate the formulation of information-seeking problems from the specific prompt engineering used to solve them. In this section, we formalize agentic search as a specific decision-making process, define the notion of success, and diagnose why LLMs become sub-optimal agents for this process.

3.1 The Context Gathering Decision Process

A Context Gathering Decision Process (CGDP) can be viewed as a POMDP [13] with terminal rewards and per-action costs, closely related to sequential identification [30, 7]; it is defined by a task q (e.g. user query), a massive hidden world state $c \in \mathcal{C}$ (e.g. codebase), an action space \mathcal{A} (e.g. LLM-callable tools), an observation function $F : \mathcal{A} \times \mathcal{C} \rightarrow \mathcal{O}$ that maps an agent’s actions to observable text chunks, and a per-action cost λ . At each timestep t , the agent selects an action $a_t \in \mathcal{A}$ which can either be an environment query (e.g. a BASH command) that incurs cost $Cost(a_t)$ and creates observation $o_t = F(a_t, c)$, or a termination action that returns a final answer a_{final} . The environment evaluates the agent’s final answer via a binary success function $Success(q, c, a_{\text{final}}) \in \{0, 1\}$. Let $a^*(q, c)$ denote an optimal answer, i.e., $Success(q, c, a^*(q, c)) = 1$.

Prompted with query q alone, an LLM lacks the context to produce $a^*(q, c)$ and will abstain or hallucinate. Unable to process the entire hidden state c at once, it must iteratively interact with F to gather a sufficient subset of information.

The optimal CGDP agent maximizes expected success while minimizing exploration cost (e.g. LLM token budget and/or latency) across a task distribution \mathcal{D} :

$$\operatorname{argmax}_{\text{Policy}} \mathbb{E}_{(q,c) \sim \mathcal{D}} \left[\text{Success}(q, c, \text{Policy}(q, c)) - \lambda \sum_{t=1}^T \text{Cost}(a_t) \right]. \quad (1)$$

Crucially, the true state of the environment includes c , but the observation o_t at each step is only the fragment of c that the agent explicitly chose to observe via $F(a_t, c)$. Therefore, a successful CGDP agent must maintain an internal belief state b_t [13] that synthesizes its historical observations, tracks its progress towards $a^*(q, c)$ and guides the selection of the next action a_t .

3.2 LLMs’ Behavior in CGDPs

Agentic harnesses (e.g. ReAct [42], IRCOT [37]) solve a CGDP by deploying an LLM as a policy, maintaining belief state implicitly by concatenating observation history. At step t , the state is

$$s_{t+1} = \text{Truncate}(s_t \oplus \{a_t, o_t\}), \quad (2)$$

where the truncation programmatically drops the older steps to obey the limits of LLM context windows. The policy is implemented simply through autoregressive generation $a_{t+1} = \text{LLM}(s_{t+1})$. Viewing LLM agents as CGDP policies reveals that they can suffer from failure modes because they lack fundamental mechanisms known to be beneficial for navigating POMDPs:

1. **Lossy Representation (Goal Forgetting):** A vanilla LLM agent relies on the growing history s_t as its belief state. As t increases, the LLM must implicitly infer “what is the goal?” and “what has been figured out so far?” at every step. Empirical studies show that LLMs can ignore evidence in the middle of long contexts [21], which can cause goal drift.
2. **Premature Stopping:** Optimal agents in POMDPs stop exploring when the expected information gain of the next action is outweighed by its cost [30, 7]. However, we conjecture that LLMs are trained on instruction-following datasets where the most rewarded behavior is to immediately produce an answer when a plausible one is found. Therefore, when the corpus c contains irrelevant distractors or adversarial honeypots, the LLM agent can trigger premature termination [43].
3. **Insufficient Exploration:** Many sequential decision-making algorithms explore via optimism [3], but LLMs do not have such mechanisms. When an LLM samples an uninformative observation o_t , it can repeatedly generate similar actions $a_{t+1} \approx a_t$ (mode collapse), without any structural awareness that its search has stagnated.

To overcome these LLM-specific failure modes, we next describe an abstracted algorithm for CGDPs that uses explicit state tracking and study how state-of-the-art LLM harnesses map to it.

4 Abstract Algorithm

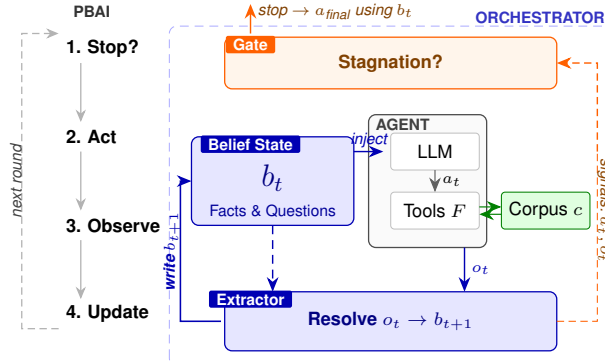


Figure 1: The PBAI loop with our two harness interventions. The **Belief State** (blue) is updated by the **Extractor** and injected into the agent prompt. The **Gate** (orange) evaluates programmatic signals to detect stagnation.

Algorithm 1 The PBAI Loop. Agent selects actions based on unresolved predicates, and updates belief state until the query is satisfied or the budget is exhausted.

Input: Query q , budget B

- 1: Initialize belief state b_0 from q
- 2: **while** cost $< B$ **do**
- 3: **1. Stop?:** If facts in b_t satisfy $q \rightarrow$ **break the while loop**
- 4: **2. Act:** Generate action a_t targeting the top open predicate in b_t .
- 5: **3. Observe:** Execute a_t , get observation o_t .
- 6: **4. Update Belief (b_{t+1}):**
- 7: Extract new facts from o_t .
- 8: Mark satisfied predicates *True*.
- 9: Append new sub-questions.
- 10: **end while**
- 11: **return** Best answer a_{final} using b_t

To build a reliable agent for the CGDP, we define the operations that should be performed. In this section, we introduce Predicate-Based Adaptive Identification (PBAI), an explicit algorithm template for a CGDP agent loop. By mapping state-of-the-art agent harnesses to the PBAI operations, we can diagnose where their implicit approximations fall short of desired behavior.

4.1 Predicate-Based Adaptive Identification

In a CGDP, the full hidden state c is too large for an LLM to hold in its context. To successfully navigate the CGDP, the agent’s belief state b_t must compress its understanding of the hidden state into a finite discrete object that can fit compactly in context.

We define a **predicate** as a logical proposition relevant to the task q that can be evaluated as *True* or *False* by querying the environment (e.g. “The target function returns a list” in a codebase c). A useful belief state is one that tracks which predicates have been resolved by observed evidence and which remain unresolved. To maintain and act upon this state, the agent iterates through a four step loop that we call Predicate-Based Adaptive Identification (PBAI):

1. **Stop?** The agent evaluates its belief state b_t . If the necessary predicates relevant to q are resolved unambiguously, the agent outputs a final answer and terminates.
2. **Select Action** If unresolved predicates remain, the agent formulates a new hypothesis about the hidden state. It selects an action a_t to resolve the highest-priority unresolved predicate.
3. **Observe** The agent executes a_t against the environment $F(\cdot, c)$ and receives observation o_t .
4. **Update Belief** The agent updates its belief state b_t given o_t . Belief updates can revise beliefs about the final answer, eliminate hypotheses inconsistent with o_t , update the agent’s understanding of the observation function $F(\cdot, c)$, or introduce new predicates to be resolved.

Algorithm 1 provides pseudocode for the PBAI loop. While POMDP solvers maintain probability distributions over states, LLMs instead autoregressively generate a textual action string based on their prompt. PBAI abstracts this textual generation in the **Select Action** step as a form of approximate Thompson Sampling [31]. While our experiments utilize greedy decoding (temperature 0) for reproducibility, conceptually, the generation of a specific search query corresponds to the agent traversing its internal hypothesis space, seeking evidence to test its predicates, and updating its internal state.

4.2 Mapping Agent Harnesses via PBAI

Existing Retrieval-Augmented Generation (RAG) and agentic memory methods perform all four PBAI operations to some extent, but they do so implicitly. By viewing these harnesses through PBAI, we can identify where they fall short (Table 2).

Table 2: Four agent harnesses mapped to the PBAI algorithm. † denotes operations where infrastructure gaps can arise. Our experiments in Section 6 confirm these gaps.

	IRCoT	ReAct	Iter-RetGen	MemGPT
Belief State	Unbounded CoT & retrieved passages†	Unbounded trajectory history†	None (memoryless)†	LLM-managed memory block†
Step 1 (Stop?)	Implicit (token generation)†	Explicit Finish action	None (fixed rounds)†	LLM-managed tool call†
Step 2 (Act)	Implicit via next CoT	Explicit Search action	Previous generation†	Explicit tool call
Step 4 (Update)	Implicit (append-only)†	Implicit (append-only)†	None†	LLM-managed memory edit†

IRCoT (implicit belief state) The concatenation of all past CoT sentences and retrieved passages serves as IRCoT’s belief state. While the CoT sentences act as an implicit **Update Belief** step, it is not curated or explicitly maintained. As the trajectory grows, earlier CoT sentences fall out of the LLM context, causing lossy representation and goal drift over long horizons.

ReAct (inefficient exploration) The explicit *Finish[answer]* action is available to an LLM agent to execute the **Stop?** step. However, ReAct agents frequently suffer from goal displacement, because they lack a dedicated mechanism to anchor unresolved predicates of the original user query q .

MemGPT (unreliable metacognition) Explicit memory management tools enable an explicit Update Belief step. However, the **Stop?** step is left entirely to the LLM’s self-assessment. Relying on weak metacognition leads to high rates of premature stopping or infinite tool-calling loops.

Iter-RetGen (memoryless) Without a persistent belief state, no adaptive stopping (executing for a fixed number of rounds), and no belief update mechanism, this is the crudest PBAI approximation.

In summary, current agent harnesses ask an LLM to act as the belief state, the CGDP policy, and the metacognitive evaluator simultaneously. To solve CGDP more reliably, we propose that PBAI operations be unbundled and have dedicated infrastructure for specific operations.

5 Interventions

In this section, we derive two modular, orchestrator-level interventions that explicitly implement the PBAI operations. Because these interventions exist outside the agent’s LLM prompt, they can be inserted into standard harnesses without difficulty.

5.1 Predicate-Based Belief State

The most severe limitation of standard agent harnesses is the unbounded accumulation of context. To resolve this, we introduce a predicate-based belief state and an explicit implementation of the PBAI **Update Belief** step. This belief state b_t replaces the unbounded trajectory history with a tightly constrained persistent data structure. It consists of two conceptually distinct elements:

1. **Facts:** A curated list of confirmed propositions (i.e., predicates resolved with evidence) extracted from past observations.
2. **Open Predicates:** A queue of unresolved sub-questions that must be answered to satisfy q .

Rather than relying on an LLM to remember all it has read during an episode, the orchestrator actively manages the belief state using a modular **Extractor**. At timestep t , the orchestrator passes (b_t, o_t) to a lightweight LLM extraction call (< 500 tokens). The extractor parses o_t , appends newly discovered facts, marks resolved predicates, and appends sub-questions if o_t reveals missing context. To ensure a bounded context footprint, the extractor compresses older findings and curates the state down to $K \leq 6$ items (see Appendix H.4 for token overhead analysis of this step). This updated state b_{t+1} is injected into the agent’s prompt for the next timestep, discarding the raw o_t . In contrast to the naïve truncation in Equation 2, $b_{t+1} = \text{Extract}(b_t, o_t)$ distills a compact representation while preserving sufficient information.

To understand how an LLM processes explicit b_t , we experiment (in Section 6) with two different textualizations of the predicate-based belief state by changing the extractor’s output format:

- **Structured object:** A JSON schema containing key-value fields for facts and open predicates.
- **Freeform text:** A natural language summary that the LLM can use as a scratchpad.

We provide the orchestrator prompts, schema definitions, and all extractor details in Appendix E.

5.2 Exhaustion Gate

A secondary failure mode in the CGDP is improper stopping (the first step in PBAI). Relying on an LLM’s self-assessment leads to two extremes: premature stopping (halting before sufficient context is gathered due to, e.g. generating a plausible hallucination) or infinite looping (repeatedly issuing similar queries due to, e.g. mode collapse). To prevent this, we introduce the Exhaustion Gate.

Instead of relying on the LLM’s metacognition to realize that its search has stagnated, the orchestrator explicitly monitors the search loop. We later compare both a focused lightweight LLM call (analogous to the extractor) and programmatic heuristics, and find that programmatic heuristics are reliable and more performant. The programmatic gate tracks two quantities across the agent’s trajectory:

1. **Action similarity:** Measures the lexical overlap between the current action a_t and recent actions, e.g. using *Jaccard* similarity. For example, with retrieval actions we compute *Jaccard* over the strings queried in the actions; this helps to detect looping behavior.

2. **Novelty:** Measures the overlap between the current observation o_t and previous observations, e.g. using *UPR* (Unique Passage Rate). UPR is the percentage of newly retrieved chunks in o_t that have not been seen in previous rounds. A low UPR indicates that the agent’s recent actions are no longer surfacing novel text from the hidden state c , suggesting that the current search hypothesis has run dry.

At each timestep t , the orchestrator evaluates $Stagnated_t := (Jaccard_t \geq \tau_J) \wedge (UPR_t \leq \tau_U)$. If $Stagnated_t$ remains *True* for p consecutive rounds, the orchestrator interrupts the PBAI loop and forces the agent to give a final answer based on the textualization of the current belief state b_t . To ensure our gate is not overly sensitive to specific hyperparameters or chunking strategies, we evaluated 16 different threshold configurations (both discrete and smooth exponential moving averages). The exact values (τ_J, τ_U, p) used in our experiments, along with our robustness sweep, are detailed in Appendix G, demonstrating that the programmatic gate provides stable gains across a wide range of settings.

6 Experiments and Analysis

To validate the CGDP framework, we apply PBAI interventions to four agent harnesses: IRCoT [37], ReAct [42], MemGPT [26] and Iter-RetGen [33]. We design experiments to answer two questions:

1. Does explicit extraction of a predicate-based belief state prevent performance degradation over long horizons?
2. Does the programmatic exhaustion gate save tokens and prevent inefficient search better than LLM self-assessment?

6.1 Setup

Datasets: We evaluated three complex retrieval-augmented QA domains requiring sophisticated exploration: LoCoMo [22] (conversational QA), MuSiQue [36] (multi-hop reasoning) and SWE-QA-Pro [28] (code repository QA).

Model and Metrics: All agents share identical zero-shot prompts and retriever configurations. Retrieval combines BM25 and `all-MiniLM-L6-v2` embeddings matching published harness weights (0.9 for IRCoT, 0.5 otherwise). Agents and orchestrators use `gpt-4o-mini` (temperature 0). Performance is evaluated via LLM-as-a-Judge (`gpt-4o` temperature 0), which measures Correctness, Completeness, penalizes Irrelevance, and normalizes to 0 – 100%. Cost is measured in total LLM API tokens. To ensure LLM judge is unbiased toward specific formatting, we verified standard string-matching metrics (Token F1, Exact Match, ROUGE, and SBERT). As detailed in Appendix F.2, these lexical metrics broadly agree with the rubric judge rankings, confirming that the accuracy gains are substantive. We use paired t -tests with Holm–Bonferroni correction.

6.2 Lobotomize-then-Replace Methodology

To isolate the impact of our interventions, we evaluate each agent harness under four controlled memory conditions:

1. **Baseline:** The harness runs as-is, accumulating its standard observation trajectory.
2. **Lobotomized:** The harness’ native memory is wiped at every timestep. It sees only the current observation o_t and the original query q . This tests how dependent the harness is on the trajectory history.
3. **PBBS (Structured, b_t^{struct}):** The lobotomized harness is injected with a JSON object that maintains up to $K = 6$ key-value pairs of facts and open predicates.
4. **PBBS (Freeform, b_t^{free}):** The lobotomized harness is injected with a natural language paragraph, rewritten each timestep by the extractor to distill facts and open predicates.

6.3 Studying the Predicate-Based Belief State

Stripping an agent’s memory (Lobotomized) causes severe degradation (Table 3), particularly on MuSiQue. However, replacing native memory management with an explicit predicate-based belief

state completely prevents this degradation (all b_t^{free} improvements are statistically significant, $p < 0.05$). Crucially, by distilling trajectories and preventing early evidence from falling out of the context window, it frequently outperforms baselines (e.g. +11.4% on MuSiQue, $p < 0.001$).

Freeform text outperforms Structured JSON. We observed that textualizing the belief state as a natural language running summary or freeform scratchpad (b_t^{free}) consistently outperformed (10 in 12 experimental settings) textualizing it with a rigid JSON schema (b_t^{struct}). To understand why freeform text wins, we analyze the extraction outputs. When forced into a rigid JSON schema, the LLM generates artifacts like 'no evidence' for open questions (occurring in up to 56% of episodes, see Table 4). This artificially fragments the LLM’s reasoning and misleadingly frames the state on answerable questions. Freeform text allows the LLM to organize the content in whatever natural language form best preserves the information.

Table 3: Performance (Rubric Judge Accuracy, %) across memory conditions. Adding the PBBS (b_t) to lobotomized agents fully recovers and often improves over the baselines. Bold indicates best per harness-task pair. Iter-RetGen (memoryless by design) has no lobotomized condition.

Task	Harness	Baseline	Lobotomized	+ PBBS (b_t^{struct})	+ PBBS (b_t^{free})
LoCoMo	IRCoT	63.0	57.2	63.6	65.4
	ReAct	65.1	59.4	65.7	66.6
	Iter-RetGen	67.3	–	65.0	65.9
	MemGPT	57.2	56.5	65.4	62.5
MuSiQue	IRCoT	14.4	9.0	21.1	25.9
	ReAct	29.6	12.2	30.1	36.8
	Iter-RetGen	37.1	–	33.8	39.6
	MemGPT	31.4	30.6	29.6	34.4
SWE-QA-Pro	IRCoT	27.3	31.7	33.3	34.6
	ReAct	39.8	33.7	37.3	43.0
	Iter-RetGen	50.9	–	48.4	51.4
	MemGPT	48.9	43.6	42.3	47.7

Table 4: Prevalence of “no evidence” open questions in structured b_t^{struct} across harnesses and tasks.

Harness	Task	% of open Qs	% episodes affected
IRCoT	MuSiQue	37.7%	63.9%
IRCoT	LoCoMo	56.5%	93.7%
IRCoT	SWE-QA-Pro	41.2%	72.3%
ReAct	MuSiQue	40.7%	72.6%
ReAct	LoCoMo	56.2%	97.0%
ReAct	SWE-QA-Pro	42.3%	76.9%
Freeform (all)		0%	0%

While adding orchestrator steps might intuitively seem to increase overall token costs, Figure 2 shows that the belief state fundamentally improves the quality-cost trade-off. Lobotomization (orange) collapses the Pareto frontier; injecting the belief state (green) recovers the frontier toward baseline levels (blue), so the accuracy gains do not come at an increased token cost.

6.4 Studying the Programmatic Exhaustion Gate

As shown in Table 5, using the programmatic exhaustion gate on stateful baselines improves accuracy across all three domains (up to +6.9% on SWE-QA-Pro). Importantly, the gate relies on the belief state b_t to generate the final answer; when applied to lobotomized agents (which do not have accumulated facts), stopping the loop early results in a slight accuracy drop.

Programmatic heuristic outperforms LLM metacognition. Table 6 compares the programmatic heuristic against prompting the LLM to self-assess stagnation. The LLM can be prompted in two distinct ways and shows two extreme failures. If prompted to act conservatively, it spirals into infinite

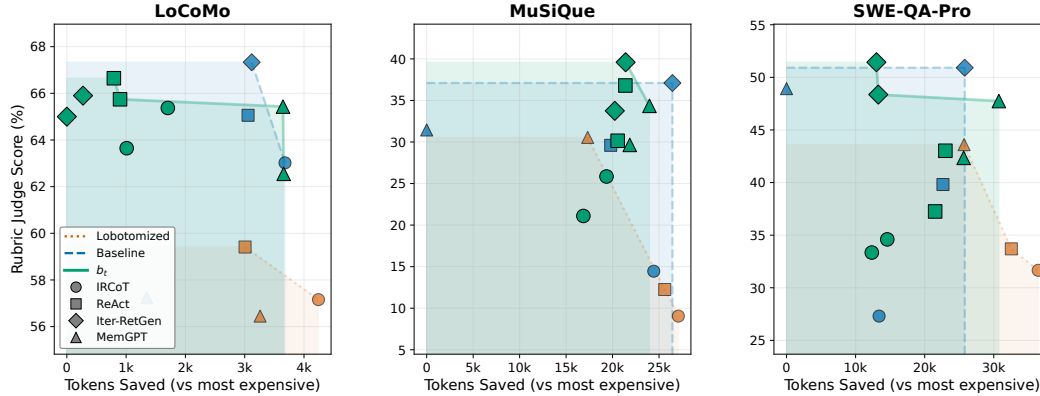


Figure 2: Pareto frontiers of quality vs. token efficiency across harnesses. Lobotomization (orange) collapses the frontier; b_t (green) recovers much of it toward baseline levels (blue).

loops (unfounded skepticism), costing 159% more tokens than the baseline. If prompted neutrally, it triggers premature stopping that damages accuracy by 5.0%. In contrast, the programmatic gate does not require additional LLM calls and saves up to 39% of the total tokens without sacrificing accuracy.

Table 5: Exhaustion Gate impact on task accuracy (averaged across IRCoT, ReAct, MemGPT). Per-harness results are in Appendix G. Bold indicates statistical significance ($p < 0.05$).

Task	Condition	Base Acc.	Δ w/ Gate
LoCoMo	Baseline	61.8%	+1.1%
MuSiQue	Baseline	25.1%	+2.8%
SWE-QA	Baseline	38.7%	+6.9%
MuSiQue	Lobotomized	17.3%	-2.9%

Table 6: Programmatic vs. LLM-assessed exhaustion gates on IRCoT in SWE-QA-Pro. Programmatic saves tokens without degrading accuracy, LLM costs more tokens than it saves. Bold indicates statistical significance ($p < 0.05$).

Gate Type	Acc. Δ	Token Savings
Programmatic	+16.2%	+39.0%
LLM (Conservative)	-3.3%	-159.0%
LLM (Neutral)	-5.0%	+56.4%

7 Discussion and Conclusion

Formalizing iterative agentic search as a Context Gathering Decision Process (CGDP) provides a blueprint for designing reliable LLM harnesses. Viewing agent behavior through approximate Thompson Sampling, our framework is able to identify modular infrastructure interventions. We demonstrated that by explicitly unbundling the operations of the search loop, we can systematically mitigate catastrophic failures. The two interventions derived from this framework – a persistent belief state and a programmatic exhaustion gate – compose with each other in state-of-the-art harnesses. They are empirically effective in preventing context degradation and halting unproductive search, saving up to 39% of tokens while improving multi-hop reasoning across three domains.

Our empirical evaluation surfaced a design principle: orchestrators should dictate *operations* not *representations*. While the orchestrator enforces state tracking steps (extraction, curation), imposing rigid structures (e.g. JSON) on intermediate representation interferes with LLM reasoning. The framework is most effective when it guides control flow while letting the LLM freely synthesize evidence in a natural language scratchpad.

Limitations and Future Work. All evaluations in this study use a single agent LLM (GPT-4o-mini). While our framework predicts even stronger gains on weaker models where implicit state tracking degrades faster, cross-model generalization remains an open question for future work. Furthermore, while the CGDP models general hidden states (e.g. codebases with BASH tools), our empirical validation focuses on complex, long-horizon retrieval to establish baseline efficacy. Finally,

the CGDP serves as a conceptual and empirical framework to guide infrastructure design, rather than providing formal mathematical bounds.

Ultimately, the CGDP abstraction is valuable for pinpointing infrastructure gaps before they cause silent failures in practice. Our analysis reveals several unaddressed gaps in current harnesses – such as LLM priors misaligned with the observation function and error compounding across trajectories (detailed in Appendix A). By mapping to the operations of the Predicate-Based Adaptive Identification loop, we can find fruitful ways to address each of them in future work.

Acknowledgments and Disclosure of Funding

We thank Ambuj Tewari, the Netflix Machine Learning and Inference Research group, Ding Tong, Aditya Sinha, Maya Ravichandran, and Anuj Phadke for their insightful feedback on this project.

References

- [1] Mohammad Aghajani Asl, Majid Asgari-Bidhendi, and Behrooz Minaei-Bidgoli. FAIR-RAG: Faithful adaptive iterative refinement for retrieval-augmented generation. *arXiv preprint arXiv:2510.22344*, 2025.
- [2] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-RAG: Learning to retrieve, generate, and critique through self-reflection. In *ICLR*, 2024.
- [3] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- [4] Yangyi Chen, Lifan Yuan, Ganqu Cui, Zhiyuan Liu, and Heng Ji. A close look into the calibration of pre-trained language models. In *ACL*, pages 1343–1367, 2023.
- [5] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitan, Robert Osazuwa Ness, and Jonathan Larson. From local to global: A graph RAG approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2025.
- [6] Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Vidhisha Balachandran, and Yulia Tsvetkov. Don’t hallucinate, abstain: Identifying LLM knowledge gaps via multi-LLM collaboration. *ACL*, 2024.
- [7] Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *COLT*, pages 998–1027, 2016.
- [8] Bernal Jimenez Gutierrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. HippoRAG: Neurobiologically inspired long-term memory for large language models. In *NeurIPS*, 2024.
- [9] Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. In *ICLR*, 2024.
- [10] Zhengbao Jiang, Frank Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. Active retrieval augmented generation. In *EMNLP*, pages 7969–7992, 2023.
- [11] Bowen Jin, Jinsung Yoon, Jiawei Han, and Sercan O. Arik. Long-context LLMs meet RAG: Overcoming challenges for long inputs in RAG. In *ICLR*, 2025.
- [12] Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O Arik, Dong Wang, Hamed Zamani, and Jiawei Han. Search-r1: Training LLMs to reason and leverage search engines with reinforcement learning. In *COLM*, 2025.
- [13] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1):99–134, 1998.

- [14] Philippe Laban, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. LLMs get lost in multi-turn conversation. In *ICLR*, 2026.
- [15] Yoonsang Lee, Howard Yen, Xi Ye, and Danqi Chen. Agentic aggregation for parallel scaling of long-horizon agentic tasks. *arXiv preprint arXiv:2604.11753*, 2026.
- [16] 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. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *NeurIPS*, 2020.
- [17] Jiaqi Li, Mengmeng Wang, Zilong Zheng, and Muhan Zhang. LooGLE: Can long-context language models understand long contexts? In *ACL*, pages 16304–16333, 2024.
- [18] Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. Search-ol: Agentic search-enhanced large reasoning models. In *EMNLP*, pages 5420–5438, 2025.
- [19] Xiaoxi Li, Jiajie Jin, Guanting Dong, Hongjin Qian, Yongkang Wu, Ji-Rong Wen, Yutao Zhu, and Zhicheng Dou. WebThinker: Empowering large reasoning models with deep research capability. In *NeurIPS*, 2026.
- [20] Hao Liu, Zhengren Wang, Xi Chen, Zhiyu Li, Feiyu Xiong, Qinhan Yu, and Wentao Zhang. HopRAG: Multi-hop reasoning for logic-aware retrieval-augmented generation. In *ACL Findings*, pages 1897–1913, 2025.
- [21] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173, 2024.
- [22] Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. Evaluating very long-term conversational memory of LLM agents. In *ACL*, pages 13851–13870, 2024.
- [23] Elliot Meyerson, Giuseppe Paolo, Roberto Dailey, Hormoz Shahrzad, Olivier Francon, Conor F Hayes, Xin Qiu, Babak Hodjat, and Risto Miikkulainen. Solving a million-step LLM task with zero errors. *arXiv preprint arXiv:2511.09030*, 2025.
- [24] Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation. *ICLR*, 2024.
- [25] Ian Osband, Daniel Russo, and Benjamin Van Roy. (more) efficient reinforcement learning via posterior sampling. In *NeurIPS*, 2013.
- [26] Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G Patil, Ion Stoica, and Joseph E Gonzalez. MemGPT: Towards LLMs as operating systems. In *ICLR*, 2024.
- [27] Joon Sung Park, Joseph C O’Brien, Carrie J Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. *UIST*, 2023.
- [28] Weihang Peng, Yuling Shi, Yuhang Wang, Xinyun Zhang, Beijun Shen, and Xiaodong Gu. SWE-QA: Can language models answer repository-level code questions? *arXiv preprint arXiv:2509.14635*, 2025.
- [29] Nikolai Rozanov and Marek Rei. StateAct: Enhancing LLM base agents via self-prompting and state-tracking. In *REALM*, pages 367–385, 2025.
- [30] Daniel Russo. Simple bayesian algorithms for best arm identification. In *COLT*, pages 1417–1418, 2016.
- [31] Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on thompson sampling. *Foundations and Trends in Machine Learning*, 2018.
- [32] Tobias Schnabel, Kiran Tomlinson, Adith Swaminathan, and Jennifer Neville. Lost in transmission: When and why LLMs fail to reason globally. In *NeurIPS*, 2025.

- [33] Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy. In *EMNLP Findings*, pages 9248–9274, 2023.
- [34] Weihang Su, Yichen Tang, Qingyao Ai, Zhijing Wu, and Yiqun Liu. DRAGIN: Dynamic retrieval augmented generation based on the real-time information needs of large language models. In *ACL*, pages 12991–13013, 2024.
- [35] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294, 1933.
- [36] Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. MuSiQue: Multihop questions via single hop question composition. *Transactions of the Association for Computational Linguistics*, 10:539–554, 2022.
- [37] Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. In *ACL*, pages 10014–10037, 2023.
- [38] Yingsheng Wu, Yuxuan Gu, Xiaocheng Feng, Weihong Zhong, Dongliang Xu, Qing Yang, Hongtao Liu, and Bing Qin. Extending context window of large language models from a distributional perspective. In *EMNLP*, pages 7288–7301, 2024.
- [39] Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-Mem: Agentic memory for LLM agents. In *NeurIPS*, 2025.
- [40] Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. Corrective retrieval augmented generation. *arXiv preprint arXiv:2401.15884*, 2024.
- [41] John Yang, Carlos Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *NeurIPS*, pages 50528–50652, 2024.
- [42] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing reasoning and acting in language models. In *ICLR*, 2023.
- [43] Howard Yen, Ashwin Paranjape, Mengzhou Xia, Thejas Venkatesh, Jack Hessel, Danqi Chen, and Yuhao Zhang. Lost in the maze: Overcoming context limitations in long-horizon agentic search. *arXiv preprint arXiv:2510.18939*, 2025.
- [44] Yuxiang Zheng, Dayuan Fu, Xiangkun Hu, Xiaojie Cai, Lyumanshan Ye, Pengrui Lu, and Pengfei Liu. DeepResearcher: Scaling deep research via reinforcement learning in real-world environments. In *EMNLP*, pages 414–431, 2025.

A Observed Failure Modes and Example Traces

A.1 Observed Failure Modes

This section summarizes the main failure modes (G1–G7) surfaced by our evaluation tasks, together with several additional failure modes that the PBAI framework suggests may arise in longer episodes, larger corpora, or harder reasoning settings.

Coverage unawareness. The agent cannot tell the difference between “I did not find it” and “it is not in the corpus”. Current methods usually have little notion of corpus coverage, so they cannot estimate how much of the search space has already been explored. A natural fix would be explicit coverage tracking.

Aggregation breakdown. Some questions require combining many weak signals across passages, such as counting entities, summarizing trends, or comparing sources. LLMs become less reliable as the number of relevant passages grows [32], so the state update can fail even when the evidence has already been retrieved.

Error compounding. Small extraction mistakes in early rounds can become premises for later reasoning. Persistent state helps by keeping each extraction step short and allowing later revisions, but it does not fully remove this problem.

Action-observation misalignment. The agent may search the wrong part of the corpus because its internal picture of how actions map to observations is wrong. In practice, this means it chooses actions that are locally sensible but poorly matched to the corpus structure. A natural fix would be a lightweight predict-before-retrieve step.

Unjustified confidence. The agent may stop with high confidence because the answer sounds plausible, not because the evidence is complete. This happens when stopping depends more on fluent output than on whether the key predicates have actually been resolved. A natural fix would be evidence-grounded stopping or calibrated answer generation.

A.2 Annotated Agent Traces

All traces below are from baseline conditions of our experiments (no b_t injection, no exhaustion gate).

G5/G6: Retrieval Stagnation (IRCoT, MuSiQue). Retrieval stagnation occurs when the agent issues the same or near-identical search actions across rounds, receiving the same passages and making no progress. We measure this as episodes where the token-level Jaccard similarity between consecutive actions exceeds 0.5 for at least 2 consecutive pairs.

Example: User query: “*What city shares a border with the place where the person who went to the state known for its Mediterranean climate during the gold rush worked?*” Ground truth: Rio Linda. IRCoT runs 10 rounds (maximum) and answers incorrectly.

R1: To answer the question, we need to identify a city that shares a border with a location associated with the Mediterranean...
R2: The Black Hills Gold Rush took place in Dakota Territory, which is not associated with a Mediterranean climate.
R3: To answer the question, we need to identify a city that shares a border...
R5–R10: Repeat of R1/R3 pattern with minor variations.

Annotation. The agent alternates between two framings without advancing the multi-hop chain. Because it has no explicit mechanism for detecting stagnation or trying a new direction, it remains stuck in the loop. The exhaustion gate would flag this quickly via high action overlap.

G4: Premature Stopping (MemGPT, LoCoMo). Premature stopping occurs when the agent stops after 1–2 rounds with an incorrect answer despite insufficient evidence.

Example. User query: “*Who performed at the concert at Melanie’s daughter’s birthday?*” Ground truth: Matt Patterson. MemGPT runs 1 round, retrieves 5 passages with action “concert Melanie’s daughter’s birthday”, and answers UNANSWERABLE.

Annotation. The agent does not find the answer in the first retrieval and stops instead of refining the search. With a persistent belief state, the unresolved question “who performed?” would remain explicit, making early termination less attractive.

G2: Goal Displacement (ReAct, MuSiQue). Goal displacement occurs when the agent’s actions drift from the original user query, pursuing tangential leads.

Example. User query: “*How were people from whom new coins were a proclamation of independence by the Somali Muslim Ajuran Empire expelled from the natural boundary between Thailand and A Magne’s country?*” Ground truth: The dynasty regrouped and defeated the Portuguese. ReAct runs 5 rounds.

R1 action: Ajuran Empire proclamation of independence and expulsion of people
R5 action: Ajuran Empire Southeast Asia history

Annotation. The agent follows a salient thread but drifts away from the intermediate sub-questions required to answer the original query. Without explicit state tracking for what remains unresolved, tangential directions are easy to pursue.

G1: Evidence Not Persisted (IRCoT, MuSiQue). Evidence forgetting occurs when relevant information appears in early rounds but is absent from later reasoning.

Example. User query: “*Who is the mother of the screenwriter of WarGames?*” Ground truth: Jane Greer. IRCoT runs 10 rounds and answers incorrectly. Round 2 passage contains: “*...He is the son of actress Jane Greer and producer Edward Lasker...*” By rounds 8–10, the chain-of-thought discusses the screenwriters without mentioning their mothers.

Annotation. The relevant evidence appears early but is not preserved in a durable form, so it no longer shapes later reasoning. A good persistent belief state would have written “Jane Greer” into the running state as soon as it appeared.

G3: Parametric Dominance (ReAct, MuSiQue). Parametric dominance occurs when the agent produces an answer grounded in pretraining knowledge rather than the retrieved evidence.

Example. User query: “*How many mandatory transmitters of the Canadian Broadcasting Centre’s owner were updated before the deadline?*” Ground truth: only about half. ReAct runs 3 rounds.

R2 Thought: The observations indicate that the CBC did not convert all of its mandatory transmitters to digital by the original deadline of August 31, 2011...
 Answer: 15 mandatory transmitters were updated before the deadline

Annotation. The final answer introduces a specific number that is not supported by the retrieved evidence. Here the model’s parametric prior wins over the retrieved passages because the evidence is not kept in an explicit, persistent state.

G7: Evidence-Action Misalignment (LoCoMo). Agents incorporate retrieved evidence without verifying that it actually pertains to the specific entities in the user query.

Example. User query: “*What was grandma’s gift to Melanie?*” (entity-swapped; the original conversation discusses a gift to a different person). IRCoT retrieves passages about the necklace gift and answers “A necklace” in 2 rounds without verifying that the evidence pertains to Melanie specifically.

Annotation. The agent finds a relevant gift passage, but it does not verify that the gift is attached to the correct entity in the user query. Entity-grounded extraction or a lightweight alignment check would catch this mismatch before answer generation.

Table 7: Gap prevalence across baseline conditions. Each cell shows the count of episodes exhibiting the pattern.

Gap	Task	IRCoT	ReAct	MemGPT	Iter-RetGen
G4 (premature stop)	MuSiQue	74	4	62	0
	LoCoMo	68	10	143	0
	SWE-QA-Pro	15	0	3	0
G5/G6 (stagnation)	MuSiQue	115	146	173	139
	LoCoMo	248	70	24	986
	SWE-QA-Pro	197	51	71	37
G2 (goal displacement)	MuSiQue	114	236	180	0
	LoCoMo	76	54	6	0
	SWE-QA-Pro	9	61	20	0
G7 (evidence-action misalign.)	LoCoMo una.	125/440	73/440	229/440	—

B Additional Interventions Suggested by the Framework

The gap analysis suggests several additional orchestrator-level interventions that are outside the scope of this paper but follow naturally from the framework:

Gap-aware stopping. Instead of asking only whether search has stalled, the orchestrator could prompt a focused LLM call to identify which predicates remain unresolved. This would directly target premature stopping and would complement the exhaustion gate, which targets the opposite failure mode of stagnation.

Entity-grounded extraction and alignment checking. The extractor could be required to ground each finding against the entities and relations in the user query, and a separate verification step could catch remaining mismatches before answer generation.

Triggered resampling. When stagnation is detected, the orchestrator could branch to one or more new hypotheses that are still consistent with the accumulated observations, rather than terminating immediately. This would make direction changes explicit rather than leaving them to the model’s chain-of-thought.

Observation filtering. Before extraction, the orchestrator could filter or compress the raw observation itself, removing clearly irrelevant text before it ever reaches the belief-update step.

Parallel hypothesis aggregation. Another way to explore multiple hypotheses is to run several retrieval trajectories in parallel and aggregate them afterward [15]. This turns hypothesis diversity into an explicit orchestrator decision rather than an accidental byproduct of a single trajectory.

C Extended Related Work

C.1 Agentic RAG Landscape

Recent work has explored giving retrieval agents more sophisticated decision-making capabilities. SearchR1 [12] and Search-o1 [18] train retrieval agents with reinforcement learning. WebThinker [19] extends this to web-scale retrieval. SLIM [43] targets long-horizon search by periodically summarizing trajectories. These approaches are complementary to ours: they improve the agent’s action selection policy or trajectory management through training and tool design, while we improve the infrastructure around the policy (persistent state, exhaustion detection) through orchestration.

C.2 LLM Metacognition and Self-Assessment

Our finding that programmatic exhaustion detection is more token-efficient than LLM-judged stopping is consistent with a growing body of evidence on LLM metacognitive limitations. Feng et al. [6] showed that multi-LLM probing improves abstention by 19.3% over single-model self-assessment. Chen et al. [4] demonstrated that verbalized confidence is poorly calibrated across model families. These findings motivate our design choice: the stagnation signals we measure (action Jaccard and UPR) are observable facts about the retrieval process that do not require the LLM to assess its own search progress.

C.3 Consistency and Contradiction Detection

Evidence-action misalignment connects to a literature on contradiction detection in LLM outputs. Mündler et al. [24] showed that pairwise contradiction detection achieves approximately 80% F1, substantially better than document-level detection. The key insight is that contradiction detection is effective when the relevant facts are presented together in short context, but degrades when they must be identified from a long history. This motivates the orchestrator’s role: maintaining a persistent state (b_t) so that new extractions can be compared against existing facts in focused pairwise calls.

C.4 Generative Agents and Long-Term Memory

Park et al. [27] introduced generative agents with long-term memory for simulated social environments. Their memory architecture (observation \rightarrow reflection \rightarrow planning) shares structural similarities with our belief state management (observation \rightarrow extraction \rightarrow state update), but operates in a different setting: their agents interact with a changing social world (full POMDP), while our agents search a static database (sequential identification). The MAKER framework [23] extends agentic memory to long-horizon creative tasks.

D Experimental Setup and Reproducibility

Datasets and Splits. We evaluated three datasets: LoCoMo (1,275 tasks spanning conversational memory), MuSiQue (500 answerable tasks spanning multi-hop Wikipedia routing), and SWE-QA-Pro (260 tasks spanning software repository structures). We utilized the standard test/validation splits provided by the authors of the respective benchmarks.

Compute Resources and LLM Usage. All agent trajectories, modular extractions, and rubric evaluations were executed via OpenAI’s API using the `gpt-4o-mini` and `gpt-4o` endpoints. The total compute expenditure for the experiments, including baseline runs, lobotomization sweeps, and gate ablations, was approximately \$1,500 in API credits.

Licenses. We use LoCoMo (**CC BY-NC 4.0**), MuSiQue (**CC BY 4.0**), and SWE-QA-Pro Bench (**MIT**) as evaluation datasets, and `all-MiniLM-L6-v2` (**Apache-2.0**) for dense retrieval embeddings. We accessed `gpt-4o-mini` and `gpt-4o` through the OpenAI API under the applicable OpenAI Services Agreement and Service Terms.

E Orchestrator and Judge Prompts

We list the key prompts designed for our interventions. All prompts are in the experiment codebase under `src/orchestrator/`. Harness-specific prompts and rubric judge prompts are in `src/methods/prompts/` and `src/scoring/`.

E.1 Structured Extraction Prompt

The structured prompt separates output into facts, resolved questions, and new questions. The instruction to “state precisely what evidence is still missing” produces the “no evidence” artifacts analyzed in Section 6.

```
New retrieved passages:
{observation}

What we already know (DO NOT repeat any of these):
{established_facts}

<scratchpad>
First, what do the passages actually state? Extract the key
claims exactly as the evidence presents them -- preserve who did
what, which entity is involved, what values are mentioned. Do not
paraphrase in a way that changes attribution or meaning.
Then, given the question "{question}", what is new here compared
to what we already know? Do any of these claims resolve our open
questions? If a fact is already listed above, skip it entirely.
</scratchpad>

Open questions we are still investigating:
{open_questions}

Output ONLY genuinely new facts not already listed above. If the
passages contain nothing new beyond what we already know, write
"Nothing relevant."

New facts:
- The claim exactly as the evidence states it (source: document or
  passage identifier)
Resolved questions:
- Which open questions are now answered by the evidence?
New questions:
- What specific evidence is still missing? State precisely what has
  NOT been found.
```

E.2 Freeform Extraction Prompt

The freeform prompt produces notes and memories without fact/question separation.

```

New retrieved passages:
{observation}

Your current notes (DO NOT repeat any of these):
{existing_notes}

<scratchpad>
First, what do the passages actually state? Extract claims exactly
as presented -- preserve who did what, which entities are involved.
Do not paraphrase in a way that changes attribution.
Then, how do these relate to the question "{question}" and your
current notes? What is genuinely new? If a note already covers
this information, skip it entirely.
</scratchpad>

Write ONLY genuinely new notes not already covered above. If the
passages contain nothing new beyond your existing notes, write
"Nothing relevant."

Notes:
- A finding exactly as the evidence presents it
Memories to keep:
- Verbatim quote or key passage worth preserving

```

Both prompts operate on $b_t + o_t$ only (approximately 500 tokens of context), not the full history.

E.3 Reorganization Prompt

To strictly bound the size of the belief state, the orchestrator allows the state to grow up to $K_{\text{trigger}} = 10$ items before pausing to execute a reorganization call. This call curates the state back down to a target size of $K_{\text{target}} = 6$ (the capacity limit reported in the main text).

```

You are curating an investigation state. Given the original
question and all facts and questions gathered so far, produce a
compact, prioritized state.

Original question: {question}
Current facts: {facts}
Open questions: {questions}

Instructions:
- Keep at most {k_target} facts and {n_questions} open questions
- Merge redundant facts into single comprehensive claims
- Drop facts irrelevant to the question
- PRESERVE facts that form multi-hop reasoning chains, even if
individually they seem tangential
- Each fact must retain its source attribution
- Rewrite open questions to reflect what is actually still unknown
- Remove questions that have been answered by the facts
- Order facts by importance (most important first)

```

E.4 LLM-based Exhaustion Gate Prompts

Two LLM-based stagnation detection variants were evaluated against the programmatic gate. Both see the current investigation state and recent retrieval rounds.

Conservative (v3). Defaults to CONTINUE; requires concrete evidence of stagnation to recommend stopping.

```

You are deciding whether further retrieval will meaningfully improve
the answer to this question.

QUESTION: {question}
CURRENT INVESTIGATION STATE: {current_state}
RECENT RETRIEVAL (last {window} rounds): {recent_rounds}

```

Your DEFAULT is to CONTINUE retrieval. Only recommend stopping if you can point to CONCRETE evidence of stagnation:

- The same passages or near-paraphrases are appearing across multiple rounds
- Search queries have covered the obvious angles and a meaningfully different direction is hard to identify
- The current answer already addresses the question and further evidence is unlikely to change it

Do NOT recommend stopping just because:

- The answer seems plausible (it may still be incomplete)
- A few passages overlap (some overlap is normal)
- You are uncertain about the answer quality (uncertainty means more retrieval could help)

VERDICT: PRODUCTIVE / QUERY_STALE / EXHAUSTED
 REASON: [explanation]

Neutral (v3_neutral). No default direction; presents the three verdicts symmetrically.

You are evaluating whether an information retrieval investigation is making progress or has stagnated.

QUESTION: {question}
 CURRENT INVESTIGATION STATE: {current_state}
 RECENT RETRIEVAL (last {window} rounds): {recent_rounds}

Based on the evidence above, choose one of the following verdicts:

VERDICT: PRODUCTIVE -- if new, relevant information is still being discovered each round.
 VERDICT: QUERY_STALE -- if the current search direction is exhausted but a specific untried angle could yield new information.
 VERDICT: EXHAUSTED -- if retrieval has stalled and further rounds are unlikely to surface new relevant content.

VERDICT:
 REASON:

F Full Results and Additional Metrics

F.1 Rubric Judge Scores and Statistical Tests

Tables 8 and 9 provide the comprehensive paired t -tests and split-condition scores that support the summarized findings in Section 6 of the main text.

F.2 Lexical and Embedding-based Metrics

String metrics (Token F1, Exact Match, ROUGE-1, METEOR, SentenceBERT cosine similarity) broadly agree with the rubric judge rankings, confirming that the performance gains are not driven by judge preference for b_t -conditioned answers. On MuSiQue, b_t^{free} is the best condition for every method on every string metric.

The few divergences between string metrics and the rubric judge justify our primary use of the rubric judge (Table 8). For example, perfectly correct abstentions (e.g., answering "UNANSWERABLE") have zero token overlap with gold answers, artificially deflating F1 scores on LoCoMo despite being the optimal agent behavior. Similarly, SWE-QA-Pro string metrics favor lobotomized models because longer, verbose code explanations coincidentally have higher token overlap with the ground truth. The rubric judge accurately captures these task-specific quality dimensions (abstention, correctness, entity attribution) that surface-level overlap cannot.

F.3 Additional Statistics

Tables 14 and 15 detail the distribution of retrieval rounds and the effect of task complexity (hop count) on the interventions.

Table 8: Full paired t -tests for persistent belief state (b_t). Each cell reports the left condition minus the right condition in pp; positive means the left condition scores higher. Iter-RetGen has no lobotomized condition. **Bold blue** = $p < .05$.

Task	Harness	Lobotomization cost			b_t recovery		Format
		B-L	B- b_t^s	B- b_t^f	L- b_t^s	L- b_t^f	b_t^s - b_t^f
LoCoMo	IRCoT	+5.9	-0.6	-2.4	-6.5	-8.2	-1.7
	ReAct	+5.6	-0.7	-1.6	-6.3	-7.2	-0.9
	Iter-RetGen	—	+2.3	+1.4	—	—	-0.9
	MemGPT	+0.8	-8.2	-5.3	-9.0	-6.1	+2.9
MuSiQue	IRCoT	+5.4	-6.7	-11.4	-12.0	-16.8	-4.8
	ReAct	+17.3	-0.6	-7.2	-17.9	-24.6	-6.7
	Iter-RetGen	—	+3.3	-2.5	—	—	-5.8
	MemGPT	+0.9	+1.8	-2.9	+0.9	-3.8	-4.7
SWE-QA	IRCoT	-4.4	-6.0	-7.3	-1.7	-2.9	-1.3
	ReAct	+6.1	+2.6	-3.2	-3.6	-9.3	-5.8
	Iter-RetGen	—	+2.6	-0.5	—	—	-3.1
	MemGPT	+5.3	+6.6	+1.2	+1.3	-4.1	-5.4
Sig. pos / neg		7 / 1	3 / 3	0 / 6	0 / 6	0 / 9	1 / 7

Table 9: LoCoMo b_t effect (pp) split by answerable (835q) and unanswerable (440q entity-swap). Each cell reports the b_t condition minus the comparison condition, so positive means b_t scores higher. Structured b_t improves unanswerable detection (abstention) while freeform b_t improves answerable accuracy. Iter-RetGen has no lobotomized condition.

Harness	Answerable (835q)				Unanswerable (440q)			
	b_t^s -B	b_t^f -B	b_t^s -L	b_t^f -L	b_t^s -B	b_t^f -B	b_t^s -L	b_t^f -L
IRCoT	+1.0	+5.8	+12.1	+16.9	-0.1	-4.0	-4.2	-8.1
ReAct	-1.4	+5.8	-4.0	+3.2	+4.6	-6.4	+25.8	+14.8
Iter-RetGen	-8.0	+0.1	—	—	+8.3	-4.4	—	—
MemGPT	-5.0	-1.6	-5.3	-1.9	+33.2	+18.4	+36.1	+21.3

Table 10: Token F1 and Exact Match: b_t effect (pp). Iter-RetGen (memoryless by design) has no lobotomized condition. **Bold blue** = $p < .05$.

Task	Harness	Token F1				Exact Match			
		B \rightarrow b_t^s	B \rightarrow b_t^f	L \rightarrow b_t^s	L \rightarrow b_t^f	B \rightarrow b_t^s	B \rightarrow b_t^f	L \rightarrow b_t^s	L \rightarrow b_t^f
LoCoMo	IRCoT	-0.1	+2.2	+5.9	+8.1	+0.4	+1.4	+1.1	+2.1
	ReAct	-0.6	+3.2	-5.5	-1.7	+0.1	+1.3	-2.8	-1.6
	Iter-RetGen	-7.9	+0.6	—	—	-3.6	+0.7	—	—
	MemGPT	-4.1	-1.5	-4.2	-1.6	-0.1	-0.1	0.0	0.0
MuSiQue	IRCoT	+1.7	+5.4	+6.2	+9.9	+1.4	+3.4	+5.0	+7.0
	ReAct	+2.7	+7.7	+9.9	+14.9	+1.2	+4.2	+5.4	+8.4
	Iter-RetGen	-5.1	+2.0	—	—	-3.0	+2.8	—	—
	MemGPT	+1.3	+3.7	+0.3	+2.7	0.0	+0.2	0.0	+0.2
SWE	IRCoT	+3.5	+3.6	+1.6	+1.7	0.0	0.0	0.0	0.0
	ReAct	-0.7	+2.5	+0.3	+4.7	0.0	0.0	0.0	0.0

G Exhaustion Gate Deep Dive

G.1 Per-method gate results

Table 16 reports the exhaustion gate effect for each method individually across all tasks and conditions. **Bold blue** = $p < .05$.

IRCoT benefits most because it has the most retrieval rounds (up to 10) and thus the most opportunity for stagnation to develop. ReAct (up to 7 rounds) shows moderate benefits on baseline and

Table 11: ROUGE-1 and METEOR: b_t effect (pp). Iter-RetGen has no lobotomized condition. **Blue** = $p < .05$.

Task	Harness	ROUGE-1				METEOR			
		B \rightarrow b_t^s	B \rightarrow b_t^f	L \rightarrow b_t^s	L \rightarrow b_t^f	B \rightarrow b_t^s	B \rightarrow b_t^f	L \rightarrow b_t^s	L \rightarrow b_t^f
LoCoMo	IRCoT	-0.7	+1.9	+7.2	+9.8	-1.1	+1.4	+8.7	+11.1
	ReAct	-0.6	+3.2	-5.7	-0.7	-1.9	+2.6	-4.8	-1.4
	Iter-RetGen	-8.3	+0.5	—	—	-6.9	+0.6	—	—
	MemGPT	-4.1	-1.5	-4.1	-1.6	-5.6	-2.6	-6.0	-3.0
MuSiQue	IRCoT	+1.6	+5.4	+6.7	+10.5	+2.6	+5.8	+6.0	+9.3
	ReAct	+2.9	+7.4	+9.7	+14.2	+2.0	+9.8	+6.1	+13.9
	Iter-RetGen	-5.1	+2.0	—	—	-2.4	+2.5	—	—
	MemGPT	+1.3	+3.6	+0.4	+0.9	+2.6	+1.0	+2.6	+2.7
SWE	IRCoT	+4.1	+4.3	+3.2	+3.3	+1.8	+0.2	+0.5	+0.4
	ReAct	-1.4	+1.4	+1.4	+5.2	-0.4	+2.0	+0.7	+2.1

Table 12: SBERT and Rubric Judge: b_t effect (pp). Iter-RetGen has no lobotomized condition. Rubric judge significance is from Table 8; SBERT significance is not shown.

Task	Harness	SBERT				Rubric Judge			
		B \rightarrow b_t^s	B \rightarrow b_t^f	L \rightarrow b_t^s	L \rightarrow b_t^f	B \rightarrow b_t^s	B \rightarrow b_t^f	L \rightarrow b_t^s	L \rightarrow b_t^f
LoCoMo	IRCoT	-0.8	+1.5	+4.7	+7.0	+0.6	+2.4	+6.4	+8.2
	ReAct	-0.4	+3.4	-5.7	-1.9	+0.6	+1.5	+6.3	+7.2
	Iter-RetGen	-7.8	-0.1	—	—	-2.3	-1.4	—	—
	MemGPT	-5.3	-3.0	-5.4	-3.1	+8.2	+5.3	+8.9	+6.0
MuSiQue	IRCoT	+2.4	+6.7	+5.7	+10.0	+6.7	+11.5	+12.1	+16.9
	ReAct	+2.7	+8.1	+9.2	+14.6	+0.5	+7.2	+17.9	+24.6
	Iter-RetGen	-6.3	+2.1	—	—	-3.3	-2.5	—	—
	MemGPT	+1.7	+3.8	+0.8	+2.9	-1.8	+3.0	-0.9	+3.9
SWE	IRCoT	+4.4	+1.9	+13.5	+2.7	+6.0	+7.3	+1.6	+2.9
	ReAct	+3.5	+5.2	-2.9	+7.1	-2.5	+3.2	+3.6	+9.3

Table 13: Exhaustion gate effect on string metrics (pp): blended (gate3 answer if triggered, natural otherwise) minus natural. Positive = gate improves the metric. **Blue** = $p < .05$. IRCoT benefits most on LoCoMo/MuSiQue; SWE shows negative F1 diffs because early stopping truncates verbose code answers that have high token overlap.

Task	Harness	Token F1				ROUGE-1			
		B	L	b_t^s	b_t^f	B	L	b_t^s	b_t^f
LoCoMo	IRCoT	+0.1	+2.8	+0.7	+0.2	+0.8	+5.6	+1.7	+0.7
	ReAct	+0.1	-1.5	+0.7	+0.3	+0.1	-1.5	+0.7	+0.3
	Iter-RetGen	-5.5	—	0.0	0.0	-5.7	—	0.0	0.0
	MemGPT	0.0	-0.1	-0.1	-0.1	0.0	-0.1	0.0	0.0
MuSiQue	IRCoT	+5.3	+2.1	+1.8	+0.9	+6.1	-0.3	0.0	+0.4
	ReAct	+0.8	-0.3	-0.3	+0.1	+1.4	+2.3	3.2	+0.8
	Iter-RetGen	-1.9	—	0.0	-1.9	-0.3	—	-0.3	+0.4
	MemGPT	+2.7	-2.4	+1.0	+0.1	+2.7	-2.2	+1.2	+0.2
SWE	IRCoT	-5.4	-8.4	-5.3	-5.2	-8.2	-9.2	-6.5	-6.6
	ReAct	+1.1	-2.2	-0.1	-1.2	+0.8	-1.8	-0.2	-1.0

lobotomized conditions. MemGPT has variable round counts and less predictable stagnation patterns; the gate hurts on MuSiQue lobotomized (-9.0pp) because it triggers during productive search episodes. Iter-RetGen’s 4 fixed rounds mean the gate fires at the last round with no opportunity for early stopping, producing 0/9 significant positives and 3/9 significant negatives.

Table 14: Mean and median retrieval rounds per episode. Max rounds: IRCoT 10, ReAct 7, Iter-RetGen 4 (fixed), MemGPT 12. Iter-RetGen always uses all 4 rounds (omitted).

Task	Harness	Mean rounds				Median rounds			
		B	L	b_t^s	b_t^f	B	L	b_t^s	b_t^f
LoCoMo	IRCoT	3.6	4.6	4.3	4.0	2	3	2	2
	ReAct	3.4	5.2	4.1	4.4	3	7	3	3
	MemGPT	1.6	2.1	1.1	1.2	1	1	1	1
MuSiQue	IRCoT	5.9	7.6	6.9	6.9	5	10	8	8
	ReAct	5.1	6.9	4.7	5.2	5	7	4	6
	MemGPT	5.9	7.4	3.0	2.6	3	12	2	2
SWE	IRCoT	8.0	8.0	8.1	8.0	10	10	10	10
	ReAct	4.6	6.9	5.8	5.6	4	7	7	7
	MemGPT	4.2	4.8	3.3	2.5	1	1	1	1

Table 15: MuSiQue b_t effect (pp) by hop count. Iter-RetGen has no lobotomized condition. **Blue** = $p < .05$.

Harness	Hops	N	B \rightarrow b_t^s	B \rightarrow b_t^f	L \rightarrow b_t^s	L \rightarrow b_t^f
IRCoT	2-hop	70	+1.1	+15.0	+4.4	+18.3
	3-hop (t1)	97	+3.5	+5.7	+9.6	+11.8
	3-hop (t2)	33	-6.2	+2.3	-2.0	+6.5
	4-hop (t1)	185	+2.1	+2.7	+10.2	+10.8
	4-hop (t2)	49	+2.0	+0.7	+2.5	+1.2
	4-hop (t3)	66	0.0	+5.5	+0.5	+6.0
ReAct	2-hop	70	-1.7	+6.0	+9.6	+17.3
	3-hop (t1)	97	+2.1	+9.1	+10.5	+17.5
	3-hop (t2)	33	+17.0	+20.8	+16.4	+20.2
	4-hop (t1)	185	+3.7	+8.0	+8.6	+12.9
	4-hop (t2)	49	+1.7	+9.2	+6.1	+13.5
	4-hop (t3)	66	-1.5	-2.6	+10.3	+9.2
Iter-RetGen	2-hop	70	-0.8	+4.0	—	—
	3-hop (t1)	97	-9.4	+0.4	—	—
	3-hop (t2)	33	-0.8	+5.9	—	—
	4-hop (t1)	185	-5.6	+4.5	—	—
	4-hop (t2)	49	-5.6	-4.3	—	—
	4-hop (t3)	66	-4.5	-2.2	—	—
MemGPT	2-hop	70	0.0	+0.1	+1.1	+1.2
	3-hop (t1)	97	-1.1	+4.0	-2.5	+2.6
	3-hop (t2)	33	+0.4	+2.4	+1.6	+5.0
	4-hop (t1)	185	+0.2	+1.2	+1.4	+3.8
	4-hop (t2)	49	+4.1	+7.7	+0.4	+3.9
	4-hop (t3)	66	+2.1	+1.7	+1.0	+1.5

G.2 Smooth vs. discrete configurations

We swept 16 configurations: 7 discrete (hard thresholds on Jaccard and UPR) and 9 smooth (exponentially-weighted moving averages). On IRCoT, smooth configurations produce slightly better mean improvement (+7.2pp) than discrete (+6.4pp), consistent with the intuition that stagnation is gradual.

G.3 Anchoring: why the gate helps stateful methods

Table 17 shows the pooled condition-level impact: the gate significantly helps baseline, b_t^{struct} , and b_t^{free} conditions ($p < 0.001$) but is neutral on lobotomized ($p = 0.248$). Table 18 supports the anchoring hypothesis: in the lobo-vs- b_t comparisons shown below, the per-question variance of the gate’s effect is consistently higher without persistent state. The belief state appears to anchor the method’s reasoning so that the early-stop answer is more consistent with what the method would

Table 16: Exhaustion gate (global best configuration, $f_j0.6_u0.3_p2$). Δ : score difference in pp. Fire: % of episodes where the gate triggers. Save: net token savings (%). **Bold blue** = $p < .05$. \dagger Iter-RetGen (memoryless by design) has no lobotomized condition.

Harness	Task	Δ (pp)				Fire (%)				Save (%)			
		B	L	b_t^s	b_t^f	B	L	b_t^s	b_t^f	B	L	b_t^s	b_t^f
IRCoT	LoC	+3.3	+2.2	+3.8	+2.9	20	24	21	15	13	18	19	15
	MuS	+5.7	+0.4	+0.3	+0.4	33	42	28	22	9	15	11	9
	SWE	+16.2	+1.9	+4.4	+4.9	80	47	39	38	39	17	15	17
ReAct	LoC	+0.2	+3.0	-0.5	-0.2	4	49	21	25	-2	12	9	11
	MuS	+0.8	-0.1	-0.2	+0.1	18	39	14	18	-2	-18	4	3
	SWE	+2.3	-0.9	+0.3	-1.2	14	37	27	30	-3	-20	3	8
Iter-RetGen \dagger	LoC	-5.2\dagger	—	+0.7	+0.4	60 \dagger	—	16	15	-8 \dagger	—	-1	-1
	MuS	-2.6\dagger	—	+0.5	-0.2	11 \dagger	—	2	3	-1 \dagger	—	0	0
	SWE	-0.7\dagger	—	+0.0	-0.3	4 \dagger	—	0	2	0 \dagger	—	0	0
MemGPT	LoC	-0.1	-0.1	+0.0	-0.1	3	10	1	1	6	37	5	8
	MuS	+1.8	-9.0	+0.7	-0.2	35	57	13	10	31	52	26	24
	SWE	+2.2	-1.8	+2.3	+0.5	28	35	20	13	45	56	42	38

have produced naturally, while memoryless methods produce higher-variance answers from similar passages.

Table 17: Pooled condition-level exhaustion gate impact (global best config, all methods). The gate significantly helps base, b_t^{struct} , and b_t^{free} conditions but is neutral on lobotomized.

Condition	Diff (pp)	p	Sig
base	+0.7%	.0004	***
lobo	+0.3%	.248	
b_t^{struct}	+0.9%	<.0001	***
b_t^{free}	+0.6%	<.0001	***

Table 18: Per-question std of gate effect (rubric judge). Lobo has higher std than both b_t conditions in all six displayed method-task rows, supporting the view that persistent state anchors reasoning and makes early stopping more reliable.

Harness	Task	Base	Lobo	b_t^{struct}	b_t^{free}
IRCoT	LoCoMo	.352	.458	.345	.343
IRCoT	MuSiQue	.397	.305	.185	.205
IRCoT	SWE	.150	.203	.173	.170
ReAct	LoCoMo	.277	.441	.324	.297
ReAct	MuSiQue	.254	.230	.204	.199
ReAct	SWE	.212	.214	.167	.202

G.4 State diff is not informative

Comparing full mode (checks Jaccard + UPR + state diff) vs. query_and_full mode (checks Jaccard + UPR only) with matched thresholds: 194 of 225 cells are identical (0pp difference). The diff check only matters for discrete $U=0.3$ configs, where it reduces fire rate by 1.8–5.8pp. For smooth configs ($U < 0.3$), the diff threshold never binds. Mean delta between modes is -0.04 pp. In this sweep, Jaccard and UPR capture the stagnation signal; the state diff added little additional signal.

Token overhead estimation methodology. Final-answer (FA) call tokens are estimated homogeneously for all gate variants (programmatic and LLM):

- *Prompt tokens*: exact via tiktoken on actual evidence + question + system prompt + template. Validated against actual API prompt tokens from an FA token ablation: 0.93–1.06 \times ratio across all 45 method/task/condition cells.

Table 19: Exhaustion gate fire round distribution (global best config). Iter-RetGen always fires at round 4 (its maximum). [†]Pre-lobotomized.

Harness	Task	Mean fire round				Median fire round			
		B	L	b_t^s	b_t^f	B	L	b_t^s	b_t^f
IRCoT	LoC	4.4	4.5	5.5	5.0	4	4	5	4
	MuS	5.8	5.9	6.5	6.0	6	6	6	6
	SWE	4.2	6.2	5.8	6.0	4	6	6	6
ReAct	LoC	3.8	4.5	5.0	4.9	3	4	5	5
	MuS	5.5	5.9	5.3	5.3	6	6	5	5
	SWE	5.1	6.1	5.3	5.2	5	7	5	5
Iter-RetGen [†]	LoC	4.0	—	4.0	4.0	4	—	4	4
	MuS	4.0	—	4.0	4.0	4	—	4	4
	SWE	4.0	—	4.0	4.0	4	—	4	4
MemGPT	LoC	5.4	3.4	3.3	3.3	5	3	3	3
	MuS	7.1	4.9	5.3	4.5	7	5	5	4
	SWE	5.7	3.6	4.1	3.7	6	3	4	4

- *Completion tokens*: per-method/task/condition average from the FA token ablation (100 random triggered questions per cell, actual gpt-4o-mini API calls).

Gate decision overhead: programmatic gates use 0 tokens (metric-based, no LLM call). LLM gates use actual logged `gate_tokens` (per-round “should I stop?” calls). Total overhead = gate decision + FA call estimate. Using the same FA estimator for all gates ensures the only difference between variants is how often and where they fire, not how their cost is estimated.

Configuration clustering. L1 distance on the 45-cell mean-diff vector reveals two main clusters:

1. *Conservative discrete* ($U=0.3, p=2-3$): moderate fire rates, fewer negatives. Contains the global best ($f_{j0.6_u0.3_p2}$).
2. *Aggressive smooth* ($U=0.10-0.15, p=2$): high fire rates, more negatives on non-IRCoT methods but larger gains on IRCoT.

Three pairs of configurations are functionally identical ($L1 = 0pp$): the `full` vs. `query_and_full` trigger mode makes no difference when the diff threshold is not binding. The global config is near-optimal: L1 distance between the global and per-method+condition diff vectors is only 0.64pp per cell on average.

H Ablations

H.1 Retriever type (alpha)

We test whether the b_t pattern holds under extreme retriever configurations: ReAct with $\alpha=1.0$ (pure BM25, keyword-only) and MemGPT with $\alpha=0.0$ (pure dense embeddings, semantic-only). Table 20 reports the key comparisons.

Table 20: b_t effect (pp) under extreme retriever alpha. The recovery-from-lobotomized pattern persists across retriever types, although baseline comparisons vary by task and state format. **Blue** = $p < .05$.

Harness (α)	Task	B $\rightarrow b_t^s$	B $\rightarrow b_t^f$	L $\rightarrow b_t^s$	L $\rightarrow b_t^f$
ReAct ($\alpha=1.0$)	LoCoMo	+3.5	+2.3	+7.9	+6.8
	MuSiQue	+7.5	+9.6	+21.5	+23.6
	SWE-QA-Pro	-2.8	+2.2	+1.7	+6.7
MemGPT ($\alpha=0.0$)	LoCoMo	+4.7	+2.8	+5.2	+3.2
	MuSiQue	-1.1	+2.7	+1.4	+5.2
	SWE-QA-Pro	-8.4	-1.1	-3.8	+3.5

The recovery-from-lobotomized pattern is strongest for freeform b_t : it improves over lobo in all six method-task rows and is significant in five. Structured b_t is more mixed under these extreme retriever settings, especially on SWE-QA-Pro. Overall, the qualitative benefit of persistent belief state management is not tied to a single retriever type.

H.2 Belief state capacity (K=6 vs. K=12)

We test whether the default curated belief state capacity of $K_{\text{target}} = 6$ items is a bottleneck by doubling to $K_{\text{target}} = 12$ on IRCoT and ReAct across all three tasks.

Table 21: K=6 vs. K=12 structured belief state (b_t^{struct}).

Harness	Task	$K_{\text{target}}=6$	$K_{\text{target}}=12$	Δ
IRCoT	LoCoMo	63.6%	63.2%	-0.5pp
ReAct	LoCoMo	65.7%	66.0%	+0.3pp
IRCoT	MuSiQue	21.1%	22.1%	+1.1pp
ReAct	MuSiQue	30.1%	30.3%	+0.1pp
IRCoT	SWE-QA-Pro	33.3%	32.5%	-0.8pp
ReAct	SWE-QA-Pro	37.3%	40.0%	+2.5pp*

Five of six pairs show no significant difference. The one marginally significant result (ReAct SWE +2.5pp, $p=0.028$) is a small effect. K=6 is not a bottleneck: the reorganization step that curates the belief state to $K_{\text{target}} = 6$ items preserves the information needed for the task.

H.3 Retrieval depth (k=3, k=5, k=10)

We test whether the b_t pattern holds at different retrieval depths on IRCoT \times LoCoMo.

Table 22: b_t effect (pp) at different retrieval depths (IRCoT \times LoCoMo). Recovery from lobotomized holds at all depths, while comparisons to baseline narrow as retrieval depth increases. **Blue** = $p < .05$.

Depth	B \rightarrow b_t^s	B \rightarrow b_t^f	L \rightarrow b_t^s	L \rightarrow b_t^f
$k=3$	+1.7	+2.0	+6.0	+6.2
$k=5$ (canonical)	+0.6	+2.4	+6.4	+8.2
$k=10$	-0.8	+0.1	+4.7	+5.7

Across retrieval depths, both b_t formats significantly recover from lobotomized at every k . Higher retrieval depth ($k=10$) improves all conditions, narrowing the gap between baseline and b_t .

H.4 Belief State Reorganization – Token Overhead Analysis

To quantify the cost-benefit trade-off of the orchestrator’s reorganization step (which triggers when the belief state exceeds $K_{\text{trigger}} = 10$ items to curate it down to $K_{\text{target}} = 6$), we analyzed the token usage across all 16, 280 episodes run under the explicit belief state (b_t) conditions.

We found that the reorganization prompt is highly efficient and only triggers when strictly necessary on complex, long-horizon tasks:

- **Trigger Frequency:** 20.8% of episodes (3, 385/16, 280) triggered the reorganization step at least once.
- **Call Overhead:** For episodes that did require reorganization, the orchestrator made an average of 1.9 reorganization calls, with each call consuming on average 979 tokens.
- **Relative Cost:** The total mean overhead for reorganization was 1, 824 tokens per triggered episode. This accounted for 11.4% of the total token consumption in those episodes.

While episodes that triggered reorganization had a substantially higher average total token cost (16, 010 tokens) compared to those that did not (5, 887 tokens), this difference is largely confounded

by underlying task complexity. Episodes that exceed the $K = 10$ threshold are inherently longer-running, multi-hop searches that require more environment steps. The reorganization step itself remains lightweight ($\sim 1,824$ tokens), ensuring the agent’s context window stays bounded and focused without imposing a prohibitive cost.