

000 001 002 003 004 005 006 007 008 009 010 MEMGEN: WEAVING GENERATIVE LATENT MEMORY FOR SELF-EVOLVING AGENTS

005
006
007
008
009
010
Anonymous authors
006
007
008
009
010
Paper under double-blind review

009 010 ABSTRACT

011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
Agent memory shapes how Large Language Model (LLM)-powered agents, akin to the human brain, progressively refine themselves through environment interactions. Existing paradigms remain constrained: parametric memory forcibly adjusts model parameters, and retrieval-based memory externalizes experience into structured databases, yet neither captures the fluid interweaving of reasoning and memory that underlies human cognition. To address this gap, we propose **MemGen**, a dynamic generative memory framework that equips agents with a human-esque cognitive faculty. It consists of a *memory trigger*, which monitors the agent’s reasoning state to decide explicit memory invocation, and a *memory weaver*, which takes the agent’s current state as stimulus to construct a latent token sequence as machine-native memory to enrich its reasoning. In this way, **MemGen** enables agents to recall and augment latent memory throughout reasoning, producing a tightly interwoven cycle of memory and cognition. Extensive experiments across eight benchmarks show that **MemGen** surpasses leading external memory systems such as ExpeL and AWM by up to 38.22%, exceeds GRPO by up to 13.44%, and exhibits strong cross-domain generalization ability. More importantly, we find that without explicit supervision, **MemGen** spontaneously evolves distinct human-like memory faculties, including planning memory, procedural memory, and working memory, suggesting an emergent trajectory toward more naturalistic forms of machine cognition. Codes are available at <http://anonymous.4open.science/r/mem-gen>.

1 INTRODUCTION

033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
The ascent of Large Language Model (LLM)-powered agents marks a paradigm shift across diverse domains (Luo et al., 2025b; Yang et al., 2024b; Qian et al., 2025; Singh et al., 2025; Pantiukhin et al., 2025; Ren et al., 2025). Pivotal to this success is the concept of *agent memory* (Zhang et al., 2024b; Wu et al., 2025b), which enables LLM agents to learn progressively from environmental interactions (Zhang et al., 2025a; Qiu et al., 2025b). Crucially, this conception of agent memory extends beyond that of conversational agents (*i.e.*, personalized memory (Wu et al., 2025b)), whose primary role is to sustain coherence across long-horizon, multi-turn dialogues (Chhikara et al., 2025; Xu et al., 2025a; Packer et al., 2024; Zhong et al., 2023). Rather, the scope of this paper is primarily on enabling agents to internalize experience, simulate human-like cognitive iteration, and progressively enhance problem-solving competence (Gao et al., 2025).

043
044
045
046
047
048
049
050
051
052
053
The memory serving as this self-evolving engine typically manifests in two dominant paradigms. The first is **(I) parametric memory**, which internalizes experiences by directly updating agents’ parameters (Yao et al., 2024; Yang et al., 2023; Zeng et al., 2023; Chen et al., 2024b; 2025). While this approach can yield substantial performance gains, its reliance on parameter modification inevitably entails *catastrophic forgetting*, *i.e.*, the erosion of general knowledge (Dou et al., 2024). Conversely, the second paradigm is **(II) retrieval-based memory**, which externalizes past experiences into a structured database, such as (i) *raw trajectories* (Luo et al., 2025a; Zhang et al., 2025a; Zhao et al., 2024), (ii) *high-level experiences* (Zhao et al., 2024; Tang et al., 2025; Fang et al., 2025; Wang et al., 2024c), and (iii) *condensed skills* like reusable APIs (Zheng et al., 2025) or MCP boxes (Qiu et al., 2025b;a). Although this non-invasive approach circumvents catastrophic forgetting, its efficacy is fundamentally tethered to context engineering. It adheres to a rigid execution pipeline, providing retrieved context to the agent without achieving the fluid, seamless integration characteristic of truly internalized memory (Su et al., 2025).




Figure 1: The comparison among parametric memory, retrieval-based memory and MemGen.

Given these deficiencies, **latent memory** offers a compelling alternative, leveraging *latent states* as a machine-native, high-density medium for memory. Existing approaches either use the **(i) key-value (KV) cache** to maintain dynamic memory set (Gim et al., 2024; Jin et al., 2024; Hongkang Yang et al., 2024), yet which is primarily confined to addressing long-context issues, or **(ii) latent token embeddings** to store agent experiences (Wang et al., 2024b; 2025b), which still rely on invasive LLM parameter updates. More fundamentally, these methods diverge from human cognition in two critical dimensions: they lack the seamless interleaving of reasoning and memory, a process where thought and memory dynamically reshape one another, and remain largely retrieval-based, fetching memories by embedding similarity (Wang et al., 2024b) rather than generatively reconstructing them into novel, coherent insights. This leads to the pivotal research question that motivates our work:

How can we architect agent memory as a dynamic cognitive faculty, capable of fluid, reconstructive processes that interweave seamlessly with reasoning?

To address this challenge, we introduce **MemGen**, a dynamic and generative memory framework designed to endow any LLM agent with a more human-esque cognitive faculty. At its core, **MemGen** continuously monitors an agent’s cognitive state, enabling it to dynamically invoke a generative process that synthesizes a bespoke latent memory at any critical juncture during its reasoning process. Practically, **MemGen** comprises two synergistic components: a reinforcement learning (RL)-trained ♣ **memory trigger**, which acts as a metacognitive monitor to discern the opportune moments for explicit memory invocation; and a ♠ **memory weaver**, which takes the agent’s current state as a stimulus to draw upon relevant implicit parametric memory (potentially augmented with externally retrieved information) and then reconstructs this synthesis into a succinct, machine-native *latent memory*. **With the reasoning core fixed, MemGen inherently mitigates catastrophic forgetting when exposed to new data, and equips agents with a generative memory deeply integrated with reasoning.**

Experimental Observation. Extensive experiments across nine benchmarks and four baseline categories demonstrate that **MemGen** delivers **1 substantial performance gains**, with improvements of up to 31.7% on ALFWorld (Shridhar et al., 2021) and 27.1% on KodCode (Xu et al., 2025d) with Qwen3-8B, surpassing parametric memory (REINFORCE++, +5.8%) and the GRPO method (+5.32%); **2 strong cross-domain generalization**, where training in the math domain not only avoids degradation elsewhere but also boosts performance in science reasoning (+6.06%) and code generation (+5.1%); and **3 continual learning ability**, maintaining stable performance in previously trained domains even after fine-tuning on three additional ones.

Analysis & Interpretation. Beyond quantitative evaluation, we also analyze the functional behavior of the latent memories learned by **MemGen**. Through post-hoc interventions examining the impact of removing specific latent memory on different agent failure modes, we found that different latent memory tokens play distinct computational roles within the agent’s reasoning process, including **1 planning memory**, where certain latent tokens specifically support high-level task planning, **2 procedural memory**, where some latent memory tokens facilitate the agent’s recall of task-specific procedural skills, such as tool usage and answer formatting, and **3 working memory**, where certain tokens help the agent maintain coherence and understanding over long contexts

108 within a single task session. These specializations strongly reveal that **MemGen** endows the agent
 109 with precise, functionally distinct memory.
 110

111 2 RELATED WORK

113 **LLM & Agent Memory.** As outlined in §1, existing memory mechanisms designed to evolve the
 114 problem-solving capacity of LLM agents can be broadly categorized into three classes: **(I) para-**
 115 **metric memory**, which either integrates past experiences directly into agent parameters through
 116 finetuning, as in FireAct (Chen et al., 2023), AgentLumos (Yin et al., 2024), and others (Zhang
 117 et al., 2024a; Fu et al., 2025), or maintains them in external parameter modules (Tack et al., 2024;
 118 Wang et al., 2024a); **(II) retrieval-based memory**, which abstracts prior experiences into trans-
 119 ferable knowledge (Zhang et al., 2025a; Zhao et al., 2024), or distills them into reusable tools and
 120 skills (Zheng et al., 2025; Wang et al., 2025c; Qiu et al., 2025b;a); and **(III) latent memory**, which
 121 leverages implicit representations to encode and retrieve experience (Wang et al., 2024b; 2025b;
 122 Hu et al., 2025b; Liu et al., 2024; Sun et al., 2025). Our **MemGen** falls within the latent memory
 123 paradigm, yet distinguishes itself from prior approaches **through its interweaving of reasoning and**
memory, as well as generative reconstruction mechanism.

124 **Latent Computation.** Our method is also closely related to latent computation, wherein latent
 125 states are employed to intervene in or reshape the LLM’s reasoning process (Zhu et al., 2025).
 126 Prominent paradigms include: **(I) architecturally enabling native latent reasoning**, exemplified
 127 by Coconut (Hao et al., 2024), CODI (Shen et al., 2025b), LatentR3 (Zhang et al., 2025b) and
 128 CoLaR (Tan et al., 2025), which render the LLM’s inference process inherently latent and machine-
 129 native; and **(II) employing latent computation to steer LLM generation**, as in LaRS (Xu et al.,
 130 2023), LatentSeek (Li et al., 2025a), SoftCoT (Xu et al., 2025c;b), and Coprocessor (Liu et al.,
 131 2024), which leverage latent representations to modulate the quality of generated outputs.

132 3 PRELIMINARY

133 **Notation.** We formalize the agent’s interaction within an environment \mathcal{E} . An agent, powered by
 134 an LLM parameterized by θ , is denoted as π_θ . For a given task \mathbf{x} , the agent’s interaction unfolds as
 135 a high-level trajectory, denoted as follows $\tau = (s_0, a_0, s_1, a_1, \dots, s_T)$, where s_t represents the state
 136 of the environment and a_t is the high-level action taken by the agent. More internally, each action
 137 a_t is essentially a sequence of tokens, $a_t = (\mathbf{z}_{t,1}, \mathbf{z}_{t,2}, \dots, \mathbf{z}_{t,L_t})$, generated autoregressively by
 138 the LLM. The generation of the j -th token is conditioned on the current state s_t and all previously
 139 generated tokens within that action:

$$141 \quad \mathbf{z}_{t,j} \sim \pi_\theta(\cdot \mid s_t, \mathbf{z}_{t,<j}). \quad (1)$$

142 After an entire action sequence \mathbf{a}_t is generated, it is executed in the environment, which transitions
 143 the state from s_t to s_{t+1} . The success of the trajectory τ is evaluated by a reward function $R(\tau)$.

144 **Problem Formalization** Given a history of past experiences $\mathcal{H} = \{(x_i, \tau_i)\}_{i=1}^N$, the objective is to
 145 leverage this history to maximize the agent’s performance on new tasks. The policy π_θ and a memory
 146 system \mathcal{M} are thus jointly optimized to maximize the expected reward over a task distribution \mathcal{D} :

$$148 \quad \max_{\theta, \mathcal{M}} \mathbb{E}_{x \sim \mathcal{D}, \tau \sim \pi_{\theta, \mathcal{M}}} [R(\tau)], \quad (2)$$

149 during which \mathcal{M} is to produce a memory representation, m , which conditions the agent’s policy. The
 150 action at any timestep t is thus sampled as $a_t \sim \pi_\theta(\cdot \mid s_t, m_t)$, where m_t is the inserted memory at
 151 that step. Crucially, the nature and timing of memory generation, which we denote as the function
 152 $f_{\mathcal{M}}$, vary across different paradigms. We express the generation of the memory m_t as:

$$155 \quad m_t = f_{\mathcal{M}}(s_t, \mathcal{H}, m_{<t}), \quad (3)$$

156 which accommodates diverse memory invocation granularities. For task-level memory (e.g., Ex-
 157 pel (Zhao et al., 2024) and G-Memory (Zhang et al., 2025a)), $f_{\mathcal{M}}$ is invoked only at $t = 0$, and
 158 $m_t = m_0$ for all subsequent steps. For step-level memory (e.g., AgentKB (Tang et al., 2025)),
 159 $f_{\mathcal{M}}$ is invoked at every step t to update the memory. In parametric memory, the influence of \mathcal{H} is
 160 compiled into θ , rendering memory generation implicit in the model parameters. Our work, which
 161 introduces dynamic latent memory, focuses on designing a more fine-grained $f_{\mathcal{M}}$ that decides for
 162 itself the optimal moments to regenerate m_t at the token level during the agent’s reasoning process.




Figure 2: The overview of our proposed MemGen.

4 METHODOLOGY

4.1 MemGen: INTERLEAVING MEMORY AND REASONING

Just as *a person is the sum of their past experiences* (McAdams, 2001), memory profoundly shapes an agent’s behavior (Xiong et al., 2025). Yet existing agent memory systems lack the adaptive fluidity characteristic of human cognition. Human reasoning and recollection operate in a continuous interplay, whereas most agent memory frameworks retrieve information only once at task initiation and append it to the query in a coarse and static manner. MemGen is designed to close this gap by enabling memory to participate dynamically in the reasoning process.

As shown in Figure 2, the reasoning process in an agent equipped with MemGen unfolds autoregressively, driven by a frozen core LLM, the reasoner π_θ . For a given state s_t , π_θ generates the action $a_t = (\mathbf{z}_{t,1}, \dots, \mathbf{z}_{t,L_t})$. MemGen continuously monitors the token-by-token generation process and performs on-demand memory insertion. At each token-generation step j , a **memory trigger** $\mathcal{T}_{\text{trigger}}$ monitors the reasoner’s internal cognitive state to determine if a moment of reflection, *i.e.*, a memory invocation, is necessary. Specifically, as the reasoner generates the token sequence $\mathbf{z}_{t,<j}$, it produces a corresponding sequence of hidden state vectors, $\mathbf{H}_{t,<j} = (\mathbf{h}_{t,1}, \dots, \mathbf{h}_{t,j-1})$, where each $\mathbf{h}_{t,k} \in \mathbb{R}^{d_{\text{model}}}$. The trigger takes the current hidden states $\mathbf{H}_{t,<j}$ as a representation of the reasoner’s current metacognitive state and computes an invocation probability:

$$p_j = \sigma(\mathcal{T}_{\text{trigger}}(\mathbf{h}_{t,1}, \dots, \mathbf{h}_{t,j-1})), \quad (4)$$

from which a binary decision, $d_j \sim \text{Bernoulli}(p_j) \in \{\text{INVOKE}, \text{SKIP}\}$, is sampled. If the decision is to `[SKIP]`, π_θ proceeds with its standard autoregressive generation, *i.e.*, $\mathbf{z}_{t,j} \sim \pi_\theta(\cdot | s_t, \mathbf{z}_{t,<j})$. However, if the decision is to `INVOKE`, the reasoning process is momentarily paused. This summons the second core component of our framework: the **memory weaver** $\mathcal{W}_{\text{weaver}}$, which takes the same cognitive state $\mathbf{H}_{t,<j}$ as a stimulus to perform a generative act of recollection. It synthesizes a bespoke, machine-native *latent memory*, formalized as $\mathbf{M}_t \in \mathbb{R}^{K \times d_{\text{model}}}$ with fixed length K :

$$\mathbf{M}_t := [\mathbf{m}_{t,1}, \mathbf{m}_{t,2}, \dots, \mathbf{m}_{t,K}] = \mathcal{W}_{\text{weaver}}(\mathbf{H}_{t,<j}), \quad (5)$$

where the memory is generated not merely from the parametric knowledge encoded within $\mathcal{W}_{\text{weaver}}$ but may also incorporate cues retrieved from external memory databases (detailed implementation is elaborated in §4.3). Crucially, \mathbf{M}_t is not a verbatim restatement of prior content but a selective reconstruction, filtered and integrated through $\mathcal{W}_{\text{weaver}}$, *akin to the memory consolidation process in human brain* (Spens & Burgess, 2024). Once formed, the latent memory is woven seamlessly into the reasoner’s ongoing dynamics: its hidden states are prepended to $\mathbf{H}_{t,<j}$, upon which the reasoner resumes generation conditioned on this enriched context,

$$\mathbf{z}_{t,j} \sim \pi_\theta(\cdot | s_t, \mathbf{z}_{t,<j}, \mathbf{M}_t). \quad (6)$$

This iterative cycle of generation, monitoring, invocation, weaving, and reintegration elevates reasoning from a linear unfolding to a recursive dialogue with memory, all without altering the frozen reasoner π_θ , and thereby preserving its general capabilities. In the following sections, we detail the implementations of the memory trigger (▷ §4.2) and the memory weaver (▷ §4.3).

4.2 MEMORY TRIGGER: LEARNING TO INVOCATE MEMORY

In this section, we describe the concrete instantiation and training recipe of the memory trigger $\mathcal{T}_{\text{trigger}}$. Recall from §4.1 that $\mathcal{T}_{\text{trigger}}$ serves as a metacognitive monitor, observing the evolving reasoning state of the frozen reasoner π_θ and producing a binary decision $d_j \in \{\text{INVOKE}, \text{SKIP}\}$ ($0 \leftrightarrow \text{SKIP}$, $1 \leftrightarrow \text{INVOKE}$) that determines whether the memory weaver should be invoked at token j .

Instantiation. We instantiate $\mathcal{T}_{\text{trigger}}$ as a lightweight LoRA adapter attached to the reasoner π_θ . At the decoding step j of the timestep t , it receives the sequence of all hidden states, $\mathbf{H}_{t,< j} \in \mathbb{R}^{(j-1) \times d_{\text{model}}}$; conditioned on this context, $\mathcal{T}_{\text{trigger}}$ outputs the action probability $P(d_j = \text{INVOKE}) \in [0, 1]$. We use hidden states rather than textual context as input because, following prior work in latent reasoning (Hao et al., 2024; Shen et al., 2025a), latent embeddings retain richer context-sensitive information that would otherwise be lost after softmax decoding. To avoid excessive computational overhead, we adopt a *sentence-granularity activation* strategy, inspired by recent studies on LLM interpretability (Anthropic, 2025; Chen et al., 2024a), which find that interventions between sentences can more effectively guide LLMs’ reasoning path. Specifically, we define a delimiter token set \mathcal{D} (e.g., commas, periods) and let the trigger act only when the current token falls in \mathcal{D} . The invocation decision is computed as:

$$d_j = \text{Bernoulli}(p_j), \quad p_j = \begin{cases} 0 & \text{if } z_j \notin \mathcal{D}, \\ \mathcal{T}_{\text{trigger}}(\mathbf{H}_{t,< j}) & \text{if } z_j \in \mathcal{D}, \end{cases} \quad (7)$$

which ensures that $\mathcal{T}_{\text{trigger}}$ is invoked only at semantically significant boundaries, preserving decoding efficiency. We validate that **MemGen** does not incur excessive inference delay in Appendix E.3.3.

Training Recipe. The memory trigger is trained via reinforcement learning, motivated by the need to balance two competing desiderata: ensuring that critical latent memories are invoked to improve task performance, while avoiding unnecessary or spurious invocations that could disrupt reasoning or incur computational overhead. Given a batch of seen tasks $\mathcal{H} = \{(x_i, \tau_i)\}_{i=1}^N$, the frozen reasoner π_θ generates candidate trajectories while the memory weaver $\mathcal{W}_{\text{weaver}}$ remains fixed. At each activated step, the trigger selects an action $\tilde{d}_j \in \{\text{INVOKE}, \text{SKIP}\}$ and receives a reward $r(\tau_i)$ reflecting the quality of the resulting trajectory with respect to the task objective. To encourage sparse yet strategically critical memory invocation, we introduce a *reward-adaptive penalty*, which discourages unnecessary activations while preserving essential ones, into the objective:

$$\max_{\phi} \mathbb{E}_{\tau_i \sim \pi_\theta, \tilde{d} \sim \mathcal{T}_{\text{trigger}}^\phi} \left[R(\tau_i) - \lambda \sum_{i,j} \max(0, \tilde{d}_{i,j} - \bar{p}) \right], \quad (8)$$

where \bar{p} is computed as the mean activation probability across high-reward trajectories, i.e., those with reward exceeding the batch median:

$$\bar{p} = \frac{1}{|\mathcal{H}_{\text{high}}|} \sum_{i \in \mathcal{H}_{\text{high}}} \frac{1}{|\tau_i|} \sum_j \tilde{d}_{i,j}, \quad \mathcal{H}_{\text{high}} = \{i : R(\tau_i) \geq \text{median}_k(R(\tau_k))\}, \quad (9)$$

where ensures that $\mathcal{T}_{\text{trigger}}$ learns to invoke memory selectively at key decision points, maximizing task reward while maintaining computational efficiency.

4.3 MEMORY WEAVER: SYNTHESIZING AND INSERTING LATENT MEMORY

In this section, we elaborate on the weaver $\mathcal{W}_{\text{weaver}}$, the memory carrier within the **MemGen** framework. When the agent assimilates new experiences, this information is exclusively internalized into the parameters of $\mathcal{W}_{\text{weaver}}$, leaving the core reasoner π_θ entirely unmodified. At junctures where the reasoner requires experiential support, a context-dependent *hook* activates the weaver to synthesize and externalize pertinent knowledge as a usable *memory*. To be more specific, recall from Equation (5) that after the $\mathcal{T}_{\text{trigger}}$ signals the need for memory at step j , $\mathcal{W}_{\text{weaver}}$ accepts $\mathbf{H}_{t,< j}$ (as the *hook*) and generates a latent token sequence \mathbf{M}_t (as the *memory*) for π_θ .

Instantiation. We instantiate $\mathcal{W}_{\text{weaver}}$ using another LoRA adapter attached to π_θ . Formally, given the incoming hook $\mathbf{H}_{t,< j} \in \mathbb{R}^{(j-1) \times d_{\text{model}}}$, the weaver outputs a latent memory matrix: $\mathbf{M}_t = \mathcal{W}_{\text{weaver}}^{\theta'}(\mathbf{H}_{t,< j}) \in \mathbb{R}^{K \times d_{\text{model}}}$, where K denotes the fixed length of the latent memory sequence and θ' are the trainable LoRA parameters. The synthesized \mathbf{M}_t is first projected through a linear layer to align it with the reasoner’s token-embedding space, and is then prepended to the current hidden states of π_θ to guide subsequent token generation, as described in Equation (6).

Training Recipe. The training of $\mathcal{W}_{\text{weaver}}$ proceeds over a batch of past trajectories $\mathcal{H} = \{(x_i, \tau_i)\}_{i=1}^N$. Distinct from conventional agent tuning, which directly integrates experiential data into the parameters of π_θ (Chen et al., 2025; Yin et al., 2024), MemGen internalizes experiential knowledge solely into $\mathcal{W}_{\text{weaver}}$, which ensures that π_θ ’s general capabilities remain intact.

Crucially, this separation makes MemGen agnostic to optimization strategies and compatible with diverse LLM backbones. Whether employing supervised fine-tuning (SFT) or RL-based objectives such as GRPO or DAPO, the weaver can be updated under a unified goal: optimizing the generation process of latent memory so as to maximize downstream reward. Formally, let $\Pi_\theta^{\mathcal{W}_{\theta'}, \mathcal{T}}(\cdot | x)$ denote the process of rolling out a trajectory for a task x by π_θ in conjunction with weaver $\mathcal{W}_{\theta'}$ and trigger \mathcal{T} . Given a reward functional R , the objective updates only θ' by maximizing the expected reward:

$$\max_{\theta_{\text{lora}}} \mathbb{E}_{(x_i, \tau_i) \sim \mathcal{H}} \mathbb{E}_{\tau \sim \Pi_\theta^{\mathcal{W}_{\theta'}, \mathcal{T}}(\cdot | x_i)} [R(x_i, \tau)], \quad (10)$$

where the gradients from R are propagated solely to θ' , thereby equipping the weaver to supply precisely the memories that improve end-to-end performance without altering π_θ . Equation (10) enables $\mathcal{W}_{\text{weaver}}$ to absorb diverse experiential signals and externalize them as dynamic, context-sensitive latent memories, independent of the architectural or training paradigm of the base reasoner. In practice, we first train the memory weaver using a random inserter as a lightweight surrogate for the trigger, and then freeze the trained weaver and proceed to train the trigger. For a complete description of the training procedure, please refer to Appendix D.1.

Integration with Retrieval-based Memory. Although the memory generation above primarily draws on the weaver’s parametric knowledge, it can be combined with external memory sources. When triggered, any retrieval-based system (e.g., MemoryBank, ExpeL) can provide textual memory, which is merged with the hook $\mathbf{H}_{t,< j}$ and fed into \mathcal{W} to produce latent memory. This allows \mathcal{W} to integrate internal knowledge and external information, supplying the reasoner with richer memory support. Implementation details and results are placed in Appendix F.

5 EXPERIMENTS

In this section, we conduct extensive experiments to answer the following research questions:

- **(RQ1)** Can MemGen surpass both parametric and retrieval-based memory?
- **(RQ2)** Is the memory learnt by MemGen generalizable across task domains? And why?
- **(RQ3)** Can MemGen facilitate continual learning and mitigate catastrophic forgetting?
- **(RQ4)** Does MemGen implicitly evolve human-like memory hierarchy?

5.1 EXPERIMENTAL SETUP

Evaluation and Benchmarks. Our evaluation covers nine datasets from five domains, including **1** *web search*: TriviaQA (Joshi et al., 2017) and PopQA (Mallen et al., 2023); **2** *embodied action*: ALFWorld (Shridhar et al., 2021); **3** *math reasoning*: AQuA (Ling et al., 2017), GSM8K (Cobbe et al., 2021), and MATH (Hendrycks et al., 2021); **4** *scientific reasoning*: GPQA (Rein et al., 2023); and **5** *coding*: KodCode (Xu et al., 2025d) and BigCodeBench (Jain et al., 2024).

Baselines. We compare MemGen against twelve baselines, categorized into four groups: **(I) Prompt-based methods**: Vanilla model, CoT (Wei et al., 2023); **(II) Parametric memory**, where experiential knowledge directly modifies model parameters via: SFT, GRPO (DeepSeek-AI et al., 2025), REINFORCE (Williams, 1992), REINFORCE++ (Hu et al., 2025a), Agent-FLAN (Chen et al., 2024b); **(III) Retrieval-based memory**, where processing tasks sequentially and storing the experiences in an external database, represented by MemoryBank (Zhong et al., 2023), ExpeL (Zhao et al., 2024), Agent Workflow Memory (AWM) (Wang et al., 2024c); and **(IV) Latent computation**, where leveraging latent tokens as carriers of experiential knowledge, including SoftCoT (Xu et al., 2025c) and Co-processor (Liu et al., 2024).

324 Table 1: Results on **SmolLM3-3B** and **Qwen3-8B**. All values represent the performance metric for
 325 each task (e.g., accuracy %). We highlight the **best** and **second best** results.

Backbone	Method	ALFWorld	TrivialQA	PopQA	KodCode	BigCodeBench	GPQA	GSM8K	MATH
SmolLM3-3B	Vanilla	18.96	10.47	8.23	37.05	35.96	9.35	47.63	16.22
	CoT	17.60	12.88	9.95	38.45	39.42	20.70	58.91	56.33
	SFT	32.36	55.25	37.22	59.25	40.79	19.70	63.48	45.65
	GRPO	55.35	65.88	45.16	68.48	72.44	22.73	80.03	61.23
	REINFORCE	53.13	63.20	46.81	65.53	67.14	23.44	82.03	58.75
	REINFORCE++	53.95	63.20	44.10	65.90	68.80	22.73	81.50	59.89
	Agent-FLAN	34.00	56.70	39.50	56.80	37.20	17.80	59.60	36.84
	ExpeL	36.18	46.20	28.16	51.14	40.22	15.15	56.23	38.11
	MemoryBank	32.80	43.30	25.81	44.50	31.80	10.20	58.30	43.53
	AWM	40.50	49.80	29.60	-	-	-	-	-
	SoftCoT	35.03	50.38	34.90	59.20	39.10	17.22	56.34	44.62
	Co-processor	38.36	53.28	38.96	56.25	45.40	20.10	57.60	38.81
	MemGen _{SFT}	50.60	68.13	42.34	62.65	42.99	26.75	70.42	57.44
	MemGen _{GRPO}	63.60	79.30	58.60	72.85	74.24	25.20	83.47	63.65
Qwen3-8B	Vanilla	58.93	52.18	34.13	49.10	33.33	38.18	89.48	79.82
	CoT	57.10	53.80	33.20	51.25	35.59	35.15	87.67	78.24
	SFT	83.59	74.55	51.12	64.75	41.33	40.33	90.76	81.35
	GRPO	85.60	76.15	58.90	73.35	70.24	39.54	92.30	83.54
	REINFORCE	82.10	75.22	57.96	72.11	70.20	37.12	91.25	83.27
	REINFORCE++	84.80	75.90	58.30	72.90	71.88	37.68	91.90	85.24
	Agent-FLAN	80.32	70.32	50.08	62.99	43.40	39.50	87.60	80.05
	ExpeL	78.97	65.54	40.33	57.20	34.23	35.15	86.20	77.40
	MemoryBank	70.41	60.56	41.60	56.39	40.61	35.66	90.35	80.35
	AWM	80.33	69.30	43.69	-	-	-	-	-
	SoftCoT	75.60	59.42	39.42	63.28	38.27	39.60	86.30	76.23
	Co-processor	73.28	61.42	45.55	64.90	42.19	39.15	76.23	79.20
	MemGen _{SFT}	85.82	77.22	54.65	66.15	40.35	43.23	91.25	83.30
	MemGen _{GRPO}	90.60	80.65	62.30	76.16	75.56	40.24	93.20	88.24

350 **Implementation Details.** We select LLM backbones of varying sizes and sources, including
 351 Qwen-2.5-1.5B (Yang et al., 2024a), HuggingFace’s SmolLM3-3B (HuggingFace, 2025), and
 352 Qwen3-8B (Yang et al., 2025). The length of each latent memory sequence K is set among
 353 $\{2, 4, 8\}$. As described in Equation (10), **MemGen** does not rely on a specific optimization algorithm,
 354 so we implement two variants: **MemGen_{SFT}** and **MemGen_{GRPO}**, in which the weaver is updated using
 355 SFT and GRPO signals. Details on these variants are provided in Appendix C. More training setup
 356 and parameter configurations are listed in Appendix D.

357 5.2 MAIN RESULTS

359 **[For RQ1] MemGen provides high-performing memory across domains.** As shown in Tables 1
 360 and 4, existing baselines exhibit clear limitations in cross-domain adaptivity. Retrieval-based memo-
 361 ries (e.g., ExpeL, MemoryBank, AWM) occasionally surpass parazmetric tuning in embodied action;
 362 for instance, AWM reaches 36.18% on ALFWorld with SmolLM3-3B, exceeding SFT by 3.15%.
 363 Yet their effectiveness deteriorates on reasoning-intensive tasks: ExpeL achieves only 8.12% on
 364 GPQA+Qwen2.5-1.5B, and even underperforms the vanilla model by 6.9% on TriviaQA, under-
 365 scoring its heavy reliance on backbone capacity. Parametric finetuning methods display the opposite
 366 tendency: they excel in structured domains such as code generation, where REINFORCE++ reaches
 367 63.33% on KodCode with Qwen2.5-1.5B, but remain weak in knowledge-intensive reasoning,
 368 with GPQA below 14%. In contrast, **MemGen** consistently advances performance across all do-
 369 mains. For example, on ALFWorld+SmolLM3-3B, **MemGen_{SFT}** and **MemGen_{GRPO}** attain 50.60%
 370 and 63.60%, improving over vanilla by 31.64% and 44.64%, respectively. Similar gains appear
 371 with the larger Qwen3-8B, where **MemGen_{GRPO}** achieves +27.06% on KodCode and +28.17%
 372 on PopQA, surpassing GRPO by up to 3.4%. Overall, the dynamic memory insertion of **MemGen**
 373 delivers substantial improvements across diverse task domains.

374 **[For RQ2] MemGen Exhibits Strong Cross-Domain Generalization.** To evaluate whether the
 375 memory learned by **MemGen** can transfer across tasks, we train **MemGen** on one dataset and test it
 376 on several others. We include two out-of-domain datasets, ScienceWorld (Wang et al., 2022) and
 377 FEVER (Thorne et al., 2018), to further probe this. As shown in Figures 3, 9 and 10, baselines such
 378 as SFT and MemoryBank achieve gains within their training domains (e.g., on ALFWorld, SFT
 379 +14.1% and MemoryBank +5.4% compared with vanilla), yet fail to generalize, with performance




Figure 3: The generalization study of **MemGen**. We train **MemGen** _{SFT} on one dataset (ALFWorld or TriviaQA) and evaluate it on four datasets (TriviaQA, ALFWorld, ScienceWorld, and FEVER).

dropping sharply on FEVER by 16.2%. In contrast, **MemGen** not only attains substantial in-domain improvements (24.55% \rightarrow 58.16% on KodCode, Figure 10), but also exhibits effective transfer: when trained on KodCode, performance on MATH rises from 36.6% \rightarrow 54.2%. Having empirically validated **MemGen**’s generalizability, a question naturally arises: what underlies this capability?

[For RQ2] The Memory Trigger Intelligently Determines When to Activate Memory Insertion, Mitigating Domain Conflict. After training **MemGen** on GSM8K, we evaluate 150 samples each from GSM8K, KodCode, and GPQA, visualizing the frequency with which the memory trigger invoked the memory weaver at each relative position in the model output. We observe that the invocation frequency varies across domains and correlates directly with performance in Figure 9: GSM8K exhibits the largest improvement (+19.64%) and maximal invocations, GPQA achieves moderate gains (+6.06%) with medium invocations, and KodCode shows the smallest improvement (+3.1%) with the fewest invocations. This indicates that **MemGen** autonomously assesses, based on task-specific context, when memory insertion will be beneficial, invoking the weaver less frequently in unfamiliar domains.

[For RQ3] MemGen Effectively Mitigates Catastrophic Forgetting. In Table 5, we sequentially train on four datasets and evaluate on all benchmarks after each stage, where **MemGen** exhibits stronger knowledge retention ability compared to baseline methods. For example, unlike SFT which primarily improves performance on the most recent task (54.10% on KodCode but only 2.53% on GPQA), **MemGen** demonstrates more balanced cross-task generalization, attaining 38.43% on AQuA and 21.72% on GPQA after GSM8K training. Finally, it mitigates forgetting on earlier tasks, preserving 40.34% on AQuA following KodCode training compared to 27.14% for ExpeL and 28.61% for SFT, indicating a more stable continual learning ability. More analysis is in Appendix E.1.

5.3 FRAMEWORK ANALYSIS

Having established the expressive capabilities of **MemGen**, we further investigate its underlying mechanisms: what do the learned latent memories look like? Do they have specialized functions?

[For RQ4] The Latent Memory Is Machine-Native and Human-Readable. We first visualized the latent memory sequences learned by **MemGen** across different datasets using t-SNE in Figures 5 and 11. As shown in Figure 5 (*Left*), sequences from distinct domains form separate distributions, with related domains clustering closely (*e.g.*, KodCode and BigCodeBench, GSM8K and MATH). Examining latent memories within the same dataset, we observed pronounced clustering patterns (as shown in Figure 5 (*Middle* and *Right*)). To explore potential commonalities within these clusters, we forcibly decoded the latent tokens. Although the decoded sequences are not human-readable, they exhibit intriguing regularities: many tokens within a cluster share structural conventions. For example, Cluster 0 in TriviaQA frequently follows the pattern “[...] SOC”, whereas Cluster 3 in GSM8K often adopts the format “[...] pick”. A large corpus of latent memory tokens is provided in Appendix G. Despite these sequences being machine-native and human-unreadable, we further investigate whether their underlying semantics can be interpreted.

[For RQ4] MemGen Implicitly Learns a Human-like Memory Hierarchy. To uncover the functional roles of different latent memory clusters, we conducted a post-hoc intervention study. Following the taxonomy from (Song et al., 2025), we study eight distinct types of agent failure, including




Figure 4: Memory invocation frequency across benchmarks at inference (trained on **MemGen** _{SFT+Qwen3-8B +GSM8K}).




Figure 5: (Left) t-SNE visualization of latent memories generated by **MemGen** +Qwen3-8B across datasets; (Middle and Right) Latent memory visualization within the TriviaQA and GSM8K datasets, clustered using K-means. The text at each cluster center represents the common pattern shared by many memory sequences in the cluster, such as Cluster 0 in GSM8K, where many sequences end with “_check”.




Figure 6: (Left) Parameter sensitivity analysis on the latent memory length K ; (Right) Effects of selectively removing latent memory clusters on different agent failure modes on the TriviaQA dataset.

planning errors, tool response/parsing failures, answer formatting mistakes, *etc*, providing a structured framework to assess how memory influences performance. During evaluation, we selectively removed latent tokens close to a specific cluster while keeping others intact, measuring the resulting changes in these failure modes. Details on (1) the visualization process, (2) failure mode annotation, and (3) token filtration are in Appendix H. As shown in Figure 6 (Right), distinct memory clusters exhibit varying influence on failure modes and can be mapped to different memory functions:

- **Planning Memory** supports high-level task planning and strategic reasoning. Removal of Cluster 2 substantially increases planning and compositional reasoning failures, indicating that this cluster is crucial for guiding the LLM agent’s decision-making and sequencing of reasoning steps.
- **Procedural Memory** captures task-specific operational knowledge, such as tool usage and formatting ability. Cluster 3 corresponds to this role, as its removal leads to a marked increase in tool response errors, parsing failures, and answer formatting mistakes.
- **Working Memory** manages the retention and effective use of prior context to maintain reasoning consistency. Clusters 1 and 4 contribute to this function: for instance, removing Cluster 1’s memory tokens results in more frequent task misunderstandings and think-act inconsistency.

Nevertheless, these memory clusters are not entirely independent: for example, removing Cluster 1 also negatively affects planning ability, indicating that these memory faculties interact and jointly enable the LLM to leverage past experience effectively. This analysis reveals that **MemGen** spontaneously organizes latent memory into a structured, human-like hierarchy.

Ablation Study & Sensitivity Analysis. We conduct a sensitivity analysis on the length of the latent memory sequence K , as shown in Figure 6 (Left). It can be observed that as the latent token length increases from $2 \rightarrow 32$, **MemGen**’s performance correspondingly improves, likely reflecting the expanded memory capacity. We then perform an ablation study on the memory trigger module in Table 6, demonstrating the necessity of a dedicatedly trained trigger for effective memory invocation. Furthermore, we analyze different training paradigms of the memory weaver in Table 7. For additional results and discussions, please refer to Appendix E.3.

486 **Efficiency Analysis.** To confirm that the memory insertion process of `MemGen` does not introduce
 487 significant inference overhead, we show in Appendix E.3.3 that, while achieving up to 57.66%
 488 performance improvement over vanilla LLMs, the per-query inference delay remains consistently
 489 below the baseline, ranging from 24% to 94% of the vanilla LLM latency. This clearly demonstrates
 490 that `MemGen` delivers substantial performance gains without compromising efficiency.
 491

492 6 CONCLUSION

493 In this work, we introduced `MemGen`, a dynamic and generative memory framework designed for
 494 LLM Agents. By interleaving reasoning with memory synthesis through a reinforcement-learned
 495 memory trigger and a generative memory weaver, `MemGen` transcends the limitations of param-
 496 metric and retrieval-based paradigms. Extensive experiments showcase substantial performance gains,
 497 robust cross-domain generalization, strong continual learning ability, and `MemGen`’s explicitly mod-
 498 eled memory hierarchy (*i.e.*, planning, procedural, and working memory). These results suggest a
 499 promising path toward self-evolving LLM agents capable of fluid and reconstructive intelligence.
 500

501 ETHICS STATEMENT

503 This work presents a latent memory architecture designed for LLM agents. All experiments are con-
 504 ducted on publicly available academic benchmarks across reasoning, scientific problem-solving, and
 505 embodied tasks, without deployment in real-world decision-making or safety-critical applications.
 506 Therefore, we believe this research does not pose direct ethical risks.
 507

508 REPRODUCIBILITY STATEMENT

510 We aim to ensure the reproducibility of our work by releasing an anonymous repository linked in
 511 the abstract, as well as detailing experimental settings (hyperparameters, LLM backbones, *etc.*) in
 512 §5.1 and appendix D.

514 REFERENCES

- 515 Anthropic. On the Biology of a Large Language Model. [https://transformer-circuits.
 517 pub/2025/attribution-graphs/biology.html](https://transformer-circuits.516 pub/2025/attribution-graphs/biology.html), 2025. [Accessed 24-08-2025].
- 518 Tianle Cai, Yuhong Li, Zhengyang Geng, Hongwu Peng, Jason D. Lee, Deming Chen, and Tri Dao.
 519 Medusa: Simple llm inference acceleration framework with multiple decoding heads, 2024. URL
 520 <https://arxiv.org/abs/2401.10774>.
- 521 Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao.
 522 Fireact: Toward language agent fine-tuning, 2023. URL <https://arxiv.org/abs/2310.05915>.
- 523 Guoxuan Chen, Han Shi, Jiawei Li, Yihang Gao, Xiaozhe Ren, Yimeng Chen, Xin Jiang, Zhenguo
 524 Li, Weiyang Liu, and Chao Huang. Sepllm: Accelerate large language models by compressing
 525 one segment into one separator. *arXiv preprint arXiv:2412.12094*, 2024a.
- 526 Zehui Chen, Kuikun Liu, Qiuchen Wang, Wenwei Zhang, Jiangning Liu, Dahu Lin, Kai Chen, and
 527 Feng Zhao. Agent-flan: Designing data and methods of effective agent tuning for large language
 528 models, 2024b. URL <https://arxiv.org/abs/2403.12881>.
- 529 Zhixun Chen, Ming Li, Yuxuan Huang, Yali Du, Meng Fang, and Tianyi Zhou. Atlas: Agent tuning
 530 via learning critical steps, 2025. URL <https://arxiv.org/abs/2503.02197>.
- 531 Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. Mem0: Building
 532 production-ready ai agents with scalable long-term memory, 2025. URL <https://arxiv.org/abs/2504.19413>.
- 533 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
 534 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to
 535 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

- 540 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
 541 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
 542 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
 543 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
 544 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,
 545 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
 546 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
 547 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai
 548 Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang,
 549 Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang,
 550 Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang,
 551 Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang,
 552 R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shengfeng
 553 Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing
 554 Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen
 555 Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong
 556 Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu,
 557 Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xi-
 558 aoshua Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia
 559 Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng
 560 Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong
 561 Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong,
 562 Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou,
 563 Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying
 564 Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda
 565 Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu,
 566 Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu
 567 Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforce-
 568 ment learning, 2025. URL <https://arxiv.org/abs/2501.12948>.
- 569 Shihan Dou, Enyu Zhou, Yan Liu, Songyang Gao, Jun Zhao, Wei Shen, Yuhao Zhou, Zhiheng Xi,
 570 Xiao Wang, Xiaoran Fan, Shiliang Pu, Jiang Zhu, Rui Zheng, Tao Gui, Qi Zhang, and Xuanjing
 571 Huang. Loramoe: Alleviate world knowledge forgetting in large language models via moe-style
 572 plugin, 2024. URL <https://arxiv.org/abs/2312.09979>.
- 573 Runnan Fang, Yuan Liang, Xiaobin Wang, Jialong Wu, Shuofei Qiao, Pengjun Xie, Fei Huang,
 574 Huajun Chen, and Ningyu Zhang. Memp: Exploring agent procedural memory. *arXiv preprint*
 575 *arXiv:2508.06433*, 2025.
- 576 Dayuan Fu, Keqing He, Yeqie Wang, Wentao Hong, Zhuoma Gongque, Weihao Zeng, Wei Wang,
 577 Jingang Wang, Xunliang Cai, and Weiran Xu. Agentrefine: Enhancing agent generalization
 578 through refinement tuning, 2025. URL <https://arxiv.org/abs/2501.01702>.
- 579 Yichao Fu, Peter Bailis, Ion Stoica, and Hao Zhang. Break the sequential dependency of llm infer-
 580 ence using lookahead decoding, 2024. URL <https://arxiv.org/abs/2402.02057>.
- 581 Huanang Gao, Jiayi Geng, Wenyue Hua, Mengkang Hu, Xinzhe Juan, Hongzhang Liu, Shilong
 582 Liu, Jiahao Qiu, Xuan Qi, Yiran Wu, Hongru Wang, Han Xiao, Yuhang Zhou, Shaokun Zhang,
 583 Jiayi Zhang, Jinyu Xiang, Yixiong Fang, Qiwen Zhao, Dongrui Liu, Qihan Ren, Cheng Qian,
 584 Zhenhailong Wang, Minda Hu, Huazheng Wang, Qingyun Wu, Heng Ji, and Mengdi Wang. A
 585 survey of self-evolving agents: On path to artificial super intelligence, 2025. URL <https://arxiv.org/abs/2507.21046>.
- 586 In Gim, Guojun Chen, Seung-seob Lee, Nikhil Sarda, Anurag Khandelwal, and Lin Zhong. Prompt
 587 cache: Modular attention reuse for low-latency inference. *Proceedings of Machine Learning and*
 588 *Systems*, 6:325–338, 2024.
- 589 Raghav Goel, Sudhanshu Agrawal, Mukul Agrawal, Junyoung Park, Yifan Zao, He Zhang, Tian
 590 Liu, Yiping Yang, Xin Yuan, Jiuyan Lu, Chris Lott, and Mingu Lee. Vocabtrim: Vocabulary
 591 pruning for efficient speculative decoding in llms, 2025. URL <https://arxiv.org/abs/2506.22694>.

- 594 Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnavh
 595 Nagarajan. Think before you speak: Training language models with pause tokens, 2024. URL
 596 <https://arxiv.org/abs/2310.02226>.
- 597 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong
 598 Tian. Training large language models to reason in a continuous latent space, 2024. URL <https://arxiv.org/abs/2412.06769>.
- 601 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 602 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset, 2021.
 603 URL <https://arxiv.org/abs/2103.03874>.
- 605 Hongkang Yang Hongkang Yang, Zehao Lin Zehao Lin, Wenjin Wang Wenjin Wang, Hao Wu
 606 Hao Wu, Zhiyu Li Zhiyu Li, Bo Tang Bo Tang, Wenqiang Wei Wenqiang Wei, Jinbo Wang
 607 Jinbo Wang, Zeyun Tang Zeyun Tang, Shichao Song Shichao Song, Chenyang Xi Chenyang Xi,
 608 Yu Yu Yu Yu, Kai Chen Kai Chen, Feiyu Xiong Feiyu Xiong, Linpeng Tang Linpeng Tang, and
 609 Weinan E Weinan E. Memory³: Language modeling with explicit memory. *Journal of Ma-*
 610 *chine Learning*, 3(3):300–346, January 2024. ISSN 2790-203X. doi: 10.4208/jml.240708. URL
 611 <http://dx.doi.org/10.4208/jml.240708>.
- 612 Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen. Reinforce++: An efficient rlhf algorithm with
 613 robustness to both prompt and reward models, 2025a. URL <https://arxiv.org/abs/2501.03262>.
- 615 Zhiyuan Hu, Yibo Wang, Hanze Dong, Yuhui Xu, Amrita Saha, Caiming Xiong, Bryan Hooi, and
 616 Junnan Li. Beyond’aha!’: Toward systematic meta-abilities alignment in large reasoning models.
 617 *arXiv preprint arXiv:2505.10554*, 2025b.
- 619 HuggingFace. SmolLM3: smol, multilingual, long-context reasoner — huggingface.co. <https://huggingface.co/blog/smollm3>, 2025. [Accessed 23-09-2025].
- 621 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
 622 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
 623 evaluation of large language models for code, 2024. URL <https://arxiv.org/abs/2403.07974>.
- 626 Chao Jin, Zili Zhang, Xuanlin Jiang, Fangyue Liu, Xin Liu, Xuanzhe Liu, and Xin Jin.
 627 Ragcache: Efficient knowledge caching for retrieval-augmented generation. *arXiv preprint*
 628 *arXiv:2404.12457*, 2024.
- 630 Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly
 631 supervised challenge dataset for reading comprehension. In Regina Barzilay and Min-Yen Kan
 632 (eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*
 633 (*Volume 1: Long Papers*), pp. 1601–1611, Vancouver, Canada, July 2017. Association for Com-
 634 putational Linguistics. doi: 10.18653/v1/P17-1147. URL <https://aclanthology.org/P17-1147/>.
- 636 Hengli Li, Chenxi Li, Tong Wu, Xuekai Zhu, Yuxuan Wang, Zhaoxin Yu, Eric Hanchen Jiang,
 637 Song-Chun Zhu, Zixia Jia, Ying Nian Wu, and Zilong Zheng. Seek in the dark: Reasoning via
 638 test-time instance-level policy gradient in latent space, 2025a. URL <https://arxiv.org/abs/2505.13308>.
- 640 Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. Eagle: Speculative sampling requires
 641 rethinking feature uncertainty, 2025b. URL <https://arxiv.org/abs/2401.15077>.
- 643 Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale gener-
 644 ation: Learning to solve and explain algebraic word problems. *arXiv preprint arXiv:1705.04146*,
 645 2017.
- 647 Luyang Liu, Jonas Pfeiffer, Jiaxing Wu, Jun Xie, and Arthur Szlam. Deliberation in latent space via
 648 differentiable cache augmentation. *arXiv preprint arXiv:2412.17747*, 2024.

- 648 Hanjun Luo, Shenyu Dai, Chiming Ni, Xinfeng Li, Guibin Zhang, Kun Wang, Tongliang Liu, and
 649 Hanan Salam. Agentauditor: Human-level safety and security evaluation for llm agents, 2025a.
 650 URL <https://arxiv.org/abs/2506.00641>.
- 651 Michael Luo, Naman Jain, Jaskirat Singh, Sijun Tan, Ameen Patel, Qingyang Wu, Alpay Ariyak,
 652 Colin Cai, Shang Zhu Tarun Venkat, Ben Athiwaratkun, Manan Roongta, Ce Zhang, Li Erran
 653 Li, Raluca Ada Popa, Koushik Sen, and Ion Stoica. Deepsw: Training a state-of-the-art
 654 coding agent from scratch by scaling rl. <https://pretty-radio-b75.notion.site/DeepSWE-Training-a-Fully-Open-sourced-State-of-the-Art-Coding-Agent-by-Scaling-RL-2025b>. Notion Blog.
 655
- 656 Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine
 657 learning research*, 9(Nov):2579–2605, 2008.
- 658 Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi.
 659 When not to trust language models: Investigating effectiveness of parametric and non-parametric
 660 memories, 2023. URL <https://arxiv.org/abs/2212.10511>.
- 661 Dan P McAdams. The psychology of life stories. *Review of general psychology*, 5(2):100–122,
 662 2001.
- 663 Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E.
 664 Gonzalez. Memgpt: Towards llms as operating systems, 2024. URL <https://arxiv.org/abs/2310.08560>.
- 665 Dmitrii Pantiukhin, Boris Shapkin, Ivan Kuznetsov, Antonia Anna Jost, and Nikolay Koldunov. Ac-
 666 celerating earth science discovery via multi-agent llm systems, 2025. URL <https://arxiv.org/abs/2503.05854>.
- 667 Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tür, Gokhan
 668 Tur, and Heng Ji. Toolrl: Reward is all tool learning needs, 2025. URL <https://arxiv.org/abs/2504.13958>.
- 669 Jiahao Qiu, Xinzhe Juan, Yimin Wang, Ling Yang, Xuan Qi, Tongcheng Zhang, Jiacheng Guo, Yifu
 670 Lu, Zixin Yao, Hongru Wang, Shilong Liu, Xun Jiang, Liu Leqi, and Mengdi Wang. Agentdistill:
 671 Training-free agent distillation with generalizable mcp boxes, 2025a. URL <https://arxiv.org/abs/2506.14728>.
- 672 Jiahao Qiu, Xuan Qi, Tongcheng Zhang, Xinzhe Juan, Jiacheng Guo, Yifu Lu, Yimin Wang,
 673 Zixin Yao, Qihan Ren, Xun Jiang, Xing Zhou, Dongrui Liu, Ling Yang, Yue Wu, Kaixuan
 674 Huang, Shilong Liu, Hongru Wang, and Mengdi Wang. Alita: Generalist agent enabling scal-
 675 able agentic reasoning with minimal predefinition and maximal self-evolution, 2025b. URL
 676 <https://arxiv.org/abs/2505.20286>.
- 677 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien
 678 Dirani, Julian Michael, and Samuel R. Bowman. Gpqa: A graduate-level google-proof q&a
 679 benchmark, 2023. URL <https://arxiv.org/abs/2311.12022>.
- 680 Shuo Ren, Pu Jian, Zhenjiang Ren, Chunlin Leng, Can Xie, and Jiajun Zhang. Towards scientific
 681 intelligence: A survey of llm-based scientific agents, 2025. URL <https://arxiv.org/abs/2503.24047>.
- 682 Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. Codi: Compressing
 683 chain-of-thought into continuous space via self-distillation. *arXiv preprint arXiv:2502.21074*,
 684 2025a.
- 685 Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. Codi: Compressing
 686 chain-of-thought into continuous space via self-distillation, 2025b. URL <https://arxiv.org/abs/2502.21074>.
- 687 Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew
 688 Hausknecht. Alfworld: Aligning text and embodied environments for interactive learning, 2021.
 689 URL <https://arxiv.org/abs/2010.03768>.

- 702 Joykirat Singh, Raghav Magazine, Yash Pandya, and Akshay Nambi. Agentic reasoning and tool
 703 integration for llms via reinforcement learning, 2025. URL <https://arxiv.org/abs/2505.01441>.
- 704
- 705 Peiyang Song, Pengrui Han, and Noah Goodman. A survey on large language model reasoning
 706 failures. In *2nd AI for Math Workshop @ ICML 2025*, 2025. URL <https://openreview.net/forum?id=hsgMn4KBFG>.
- 707
- 708
- 709 Eleanor Spens and Neil Burgess. A generative model of memory construction and consolidation.
 710 *Nature human behaviour*, 8(3):526–543, 2024.
- 711
- 712 Weihang Su, Yichen Tang, Qingyao Ai, Junxi Yan, Changyue Wang, Hongning Wang, Ziyi Ye,
 713 Yujia Zhou, and Yiqun Liu. Parametric retrieval augmented generation, 2025. URL <https://arxiv.org/abs/2501.15915>.
- 714
- 715 Yuchang Sun, Yanxi Chen, Yaliang Li, and Bolin Ding. Enhancing latent computation in transform-
 716 ers with latent tokens, 2025. URL <https://arxiv.org/abs/2505.12629>.
- 717
- 718 Jihoon Tack, Jaehyung Kim, Eric Mitchell, Jinwoo Shin, Yee Whye Teh, and Jonathan Richard
 719 Schwarz. Online adaptation of language models with a memory of amortized contexts, 2024.
 720 URL <https://arxiv.org/abs/2403.04317>.
- 721
- 722 Wenhui Tan, Jiaze Li, Jianzhong Ju, Zhenbo Luo, Jian Luan, and Ruihua Song. Think silently, think
 723 fast: Dynamic latent compression of llm reasoning chains, 2025. URL <https://arxiv.org/abs/2505.16552>.
- 724
- 725 Xiangru Tang, Tianrui Qin, Tianhao Peng, Ziyang Zhou, Daniel Shao, Tingting Du, Xinning Wei,
 726 Peng Xia, Fang Wu, He Zhu, Ge Zhang, Jiaheng Liu, Xingyao Wang, Sirui Hong, Chenglin Wu,
 727 Hao Cheng, Chi Wang, and Wangchunshu Zhou. Agent kb: Leveraging cross-domain experience
 728 for agentic problem solving, 2025. URL <https://arxiv.org/abs/2507.06229>.
- 729
- 730 James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. Fever: a large-scale
 731 dataset for fact extraction and verification. *arXiv preprint arXiv:1803.05355*, 2018.
- 732
- 733 Peng Wang, Zexi Li, Ningyu Zhang, Ziwen Xu, Yunzhi Yao, Yong Jiang, Pengjun Xie, Fei Huang,
 734 and Huajun Chen. Wise: Rethinking the knowledge memory for lifelong model editing of large
 735 language models, 2024a. URL <https://arxiv.org/abs/2405.14768>.
- 736
- 737 Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. Scienceworld:
 738 Is your agent smarter than a 5th grader?, 2022. URL <https://arxiv.org/abs/2203.07540>.
- 739
- 740 Xiaoqiang Wang, Suyuchen Wang, Yun Zhu, and Bang Liu. R³Mem: Bridging memory reten-
 741 tion and retrieval via reversible compression. In Wanxiang Che, Joyce Nabende, Ekaterina
 742 Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational
 743 Linguistics: ACL 2025*, pp. 4541–4557, Vienna, Austria, July 2025a. Association for Compu-
 744 tational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.235. URL
 745 <https://aclanthology.org/2025.findings-acl.235>.
- 746
- 747 Yu Wang, Yifan Gao, Xiusi Chen, Haoming Jiang, Shiyang Li, Jingfeng Yang, Qingyu Yin, Zheng
 748 Li, Xian Li, Bing Yin, et al. Memoryllm: Towards self-updatable large language models. *arXiv
 749 preprint arXiv:2402.04624*, 2024b.
- 750
- 751 Yu Wang, Dmitry Krotov, Yuanzhe Hu, Yifan Gao, Wangchunshu Zhou, Julian McAuley, Dan Gut-
 752 freund, Rogerio Feris, and Zexue He. M+: Extending memoryllm with scalable long-term mem-
 753 ory, 2025b. URL <https://arxiv.org/abs/2502.00592>.
- 754
- 755 Zhenhailong Wang, Haiyang Xu, Junyang Wang, Xi Zhang, Ming Yan, Ji Zhang, Fei Huang, and
 756 Heng Ji. Mobile-agent-e: Self-evolving mobile assistant for complex tasks, 2025c. URL <https://arxiv.org/abs/2501.11733>.
- 757
- Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. Agent workflow memory,
 758 2024c. URL <https://arxiv.org/abs/2409.07429>.

- 756 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc
 757 Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models,
 758 2023. URL <https://arxiv.org/abs/2201.11903>.
- 759
- 760 Yifan Wei, Xiaoyan Yu, Yixuan Weng, Tengfei Pan, Angsheng Li, and Li Du. Autotir: Autonomous
 761 tools integrated reasoning via reinforcement learning, 2025. URL <https://arxiv.org/abs/2507.21836>.
- 762
- 763 Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement
 764 learning. *Machine learning*, 8:229–256, 1992.
- 765
- 766 Mingyuan Wu, Jingcheng Yang, Jize Jiang, Meitang Li, Kaizhuo Yan, Hanchao Yu, Minjia Zhang,
 767 Chengxiang Zhai, and Klara Nahrstedt. Vtool-r1: Vlms learn to think with images via rein-
 768 forcement learning on multimodal tool use, 2025a. URL <https://arxiv.org/abs/2505.19255>.
- 769
- 770 Yaxiong Wu, Sheng Liang, Chen Zhang, Yichao Wang, Yongyue Zhang, Huifeng Guo, Ruiming
 771 Tang, and Yong Liu. From human memory to ai memory: A survey on memory mechanisms in
 772 the era of llms, 2025b. URL <https://arxiv.org/abs/2504.15965>.
- 773
- 774 Zidi Xiong, Yuping Lin, Wenya Xie, Pengfei He, Jiliang Tang, Himabindu Lakkaraju, and Zhen Xi-
 775 ang. How memory management impacts llm agents: An empirical study of experience-following
 776 behavior, 2025. URL <https://arxiv.org/abs/2505.16067>.
- 777
- 778 Wujiang Xu, Kai Mei, Hang Gao, Juntao Tan, Zujie Liang, and Yongfeng Zhang. A-mem: Agentic
 779 memory for llm agents, 2025a. URL <https://arxiv.org/abs/2502.12110>.
- 780
- 781 Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan Miao. Softcot: Soft chain-of-thought for efficient
 782 reasoning with llms, 2025b. URL <https://arxiv.org/abs/2502.12134>.
- 783
- 784 Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan Miao. Softcot++: Test-time scaling with soft chain-
 785 of-thought reasoning, 2025c. URL <https://arxiv.org/abs/2505.11484>.
- 786
- 787 Zhangchen Xu, Yang Liu, Yueqin Yin, Mingyuan Zhou, and Radha Poovendran. Kodcode: A di-
 788 verse, challenging, and verifiable synthetic dataset for coding, 2025d. URL <https://arxiv.org/abs/2503.02951>.
- 789
- 790 Zifan Xu, Haozhu Wang, Dmitriy Bespalov, Xuan Wang, Peter Stone, and Yanjun Qi. Latent skill
 791 discovery for chain-of-thought reasoning. *arXiv preprint arXiv:2312.04684*, 2023.
- 792
- 793 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 794 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint
 795 arXiv:2412.15115*, 2024a.
- 796
- 797 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, and Bo Zheng et al. Qwen3
 798 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
- 799
- 800 John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan,
 801 and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering,
 802 2024b. URL <https://arxiv.org/abs/2405.15793>.
- 803
- 804 Yi Yang, Yixuan Tang, and Kar Yan Tam. Investlm: A large language model for investment using
 805 financial domain instruction tuning, 2023. URL <https://arxiv.org/abs/2309.13064>.
- 806
- 807 Weiran Yao, Shelby Heinecke, Juan Carlos Niebles, Zhiwei Liu, Yihao Feng, Le Xue, Rithesh
 808 Murthy, Zeyuan Chen, Jianguo Zhang, Devansh Arpit, Ran Xu, Phil Mui, Huan Wang, Caim-
 809 ing Xiong, and Silvio Savarese. Retroformer: Retrospective large language agents with policy
 810 gradient optimization, 2024. URL <https://arxiv.org/abs/2308.02151>.
- 811
- 812 Da Yin, Faeze Brahman, Abhilasha Ravichander, Khyathi Chandu, Kai-Wei Chang, Yejin Choi, and
 813 Bill Yuchen Lin. Agent lumos: Unified and modular training for open-source language agents,
 814 2024. URL <https://arxiv.org/abs/2311.05657>.

- 810 Hongli Yu, Tinghong Chen, Jiangtao Feng, Jiangjie Chen, Weinan Dai, Qiying Yu, Ya-Qin Zhang,
 811 Wei-Ying Ma, Jingjing Liu, Mingxuan Wang, and Hao Zhou. Memagent: Reshaping long-context
 812 llm with multi-conv rl-based memory agent, 2025. URL <https://arxiv.org/abs/2507.02259>.
- 813
- 814 Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. Agenttuning:
 815 Enabling generalized agent abilities for llms, 2023. URL <https://arxiv.org/abs/2310.12823>.
- 816
- 817 Guibin Zhang, Muxin Fu, Guancheng Wan, Miao Yu, Kun Wang, and Shuicheng Yan. G-memory:
 818 Tracing hierarchical memory for multi-agent systems, 2025a. URL <https://arxiv.org/abs/2506.07398>.
- 819
- 820 Jianguo Zhang, Tian Lan, Rithesh Murthy, Zhiwei Liu, Weiran Yao, Ming Zhu, Juntao Tan, Thai
 821 Hoang, Zuxin Liu, Liangwei Yang, Yihao Feng, Shirley Kokane, Tulika Awalgaoonkar, Juan Car-
 822 los Niebles, Silvio Savarese, Shelby Heinecke, Huan Wang, and Caiming Xiong. Agento-
 823 hana: Design unified data and training pipeline for effective agent learning, 2024a. URL
 824 <https://arxiv.org/abs/2402.15506>.
- 825
- 826 Yang Zhang, Wenxin Xu, Xiaoyan Zhao, Wenjie Wang, Fuli Feng, Xiangnan He, and Tat-Seng
 827 Chua. Reinforced latent reasoning for llm-based recommendation, 2025b. URL <https://arxiv.org/abs/2505.19092>.
- 828
- 829 Zeyu Zhang, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Quanyu Dai, Jieming Zhu, Zhenhua Dong,
 830 and Ji-Rong Wen. A survey on the memory mechanism of large language model based agents,
 831 2024b. URL <https://arxiv.org/abs/2404.13501>.
- 832
- 833 Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. Expel: Llm
 834 agents are experiential learners, 2024. URL <https://arxiv.org/abs/2308.10144>.
- 835
- 836 Boyuan Zheng, Michael Y. Fatemi, Xiaolong Jin, Zora Zhiruo Wang, Apurva Gandhi, Yueqi Song,
 837 Yu Gu, Jayanth Srinivasa, Gaowen Liu, Graham Neubig, and Yu Su. Skillweaver: Web agents
 838 can self-improve by discovering and honing skills, 2025. URL <https://arxiv.org/abs/2504.07079>.
- 839
- 840 Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. Memorybank: Enhancing large
 841 language models with long-term memory, 2023. URL <https://arxiv.org/abs/2305.10250>.
- 842
- 843 Zijian Zhou, Ao Qu, Zhaoxuan Wu, Sunghwan Kim, Alok Prakash, Daniela Rus, Jinhua Zhao,
 844 Bryan Kian Hsiang Low, and Paul Pu Liang. Mem1: Learning to synergize memory and reasoning
 845 for efficient long-horizon agents, 2025. URL <https://arxiv.org/abs/2506.15841>.
- 846
- 847 Rui-Jie Zhu, Tianhao Peng, Tianhao Cheng, Xingwei Qu, Jinfa Huang, Dawei Zhu, Hao Wang,
 848 Kaiwen Xue, Xuanliang Zhang, Yong Shan, et al. A survey on latent reasoning. *arXiv preprint*
 849 *arXiv:2507.06203*, 2025.
- 850
- 851

852 A USE OF LARGE LANGUAGE MODELS

853 In preparing this work, we employed large language models (LLMs) to assist with: polishing the
 854 writing, conducting literature reviews, and creating visualizations.

855 B ADDITIONAL RELATED WORKS

856 **LLM Decoding & RL.** Two additional topics that relate to our work are LLM decoding and re-
 857 inforcement learning (RL). From the **decoding perspective**, MemGen dynamically generates and in-
 858 serts latent tokens, which shares similarity with speculative decoding where a drafter model receives
 859 the current decoding context and produces subsequent drafted tokens (Cai et al., 2024; Fu et al.,
 860 2024; Li et al., 2025b; Goel et al., 2025). . However, these methods primarily aim to accelerate

864 LLM inference, whereas **MemGen** focuses on leveraging latent states as effective carriers of memory.
 865 From the **RL perspective**, **MemGen** employs rule-based RL to train the memory trigger, which
 866 is closely related to reinforcement learning with variable reward (RLVR), including GRPO from
 867 DeepSeek-R1 (DeepSeek-AI et al., 2025) and its various derivatives (Qian et al., 2025; Wu et al.,
 868 2025a; Wei et al., 2025; Fu et al., 2025). While there exist efforts combining RL with agent memory,
 869 to our knowledge, most do not address self-improving memory; for example, MemAgent (Yu et al.,
 870 2025) and MEM1 (Zhou et al., 2025) focus on handling long-context inputs rather than evolving
 871 memory mechanisms.

872 **Latent Memory** We further extend our discussion by incorporating related work on latent memory
 873 mechanisms: R3Mem (Wang et al., 2025a) introduces a reversible, hierarchical memory architecture
 874 that integrates efficient long-term information retention with faithful retrieval, achieving state-
 875 of-the-art performance in long-context modeling, retrieval-augmented generation, and prolonged
 876 interactive tasks.

878 C OPTIMIZATION ALGORITHM DETAILS

880 In this section, we provide a detailed exposition of the optimization algorithms for training the
 881 memory weaver, $\mathcal{W}_{\text{weaver}}$, as mentioned in §4.3. The core principle, as established in Equation (10),
 882 is to update only the weaver’s parameters, denoted as θ' , while keeping the reasoner π_θ frozen. This
 883 modularity allows **MemGen** to be compatible with various optimization paradigms. We detail the
 884 specific implementations for Supervised Fine-Tuning (SFT) and a reinforcement learning approach,
 885 Group Relative Policy Optimization (GRPO).

887 C.1 COMBINING **MemGen** WITH SFT

889 The objective of Supervised Fine-Tuning is to train the memory weaver to generate latent mem-
 890 ories that guide the frozen reasoner π_θ to replicate the behavior observed in a dataset of high-
 891 quality demonstration trajectories. We leverage the provided history of past experiences, $\mathcal{H} =$
 892 $\{(x_i, \tau_i^*)\}_{i=1}^N$, where each τ_i^* is treated as an expert demonstration.

893 Each expert trajectory τ_i^* consists of a sequence of states and actions, $\tau_i^* = (s_{i,0}, a_{i,0}^*, s_{i,1}, a_{i,1}^*, \dots)$.
 894 Each expert action $a_{i,t}^*$ is a sequence of tokens, $a_{i,t}^* = (\mathbf{z}_{i,t,1}^*, \mathbf{z}_{i,t,2}^*, \dots, \mathbf{z}_{i,t,L_t}^*)$. The goal is to
 895 maximize the conditional log-likelihood of generating these expert token sequences.

896 During the training of the weaver, both the reasoner π_θ and the memory trigger \mathcal{T} are held fixed. At
 897 each token generation step j where the trigger activates, the weaver $\mathcal{W}_{\theta'}$ takes the reasoner’s hidden
 898 states $\mathbf{H}_{i,t,<j}$ as input and produces a latent memory $\mathbf{M}_{i,t}$. The reasoner then generates the next
 899 token conditioned on this memory. The SFT objective is to adjust the weaver’s parameters θ' to
 900 maximize the probability of the ground-truth token $\mathbf{z}_{i,t,j}^*$.

901 Formally, the optimization problem is to minimize the negative log-likelihood of the expert trajec-
 902 tories, averaged over the dataset \mathcal{H} . The loss function for the weaver’s parameters θ' is defined
 903 as:

$$905 \mathcal{L}_{\text{SFT}}(\theta') = -\mathbb{E}_{(x_i, \tau_i^*) \sim \mathcal{H}} \left[\sum_{t=0}^{T_i-1} \sum_{j=1}^{L_t} \log \pi_\theta(\mathbf{z}_{i,t,j}^* \mid s_{i,t}, \mathbf{z}_{i,t,<j}^*, \mathbf{M}_{i,t,j}) \right], \quad (11)$$

907 where the latent memory $\mathbf{M}_{i,t,j}$ is synthesized by the weaver at that specific step:

$$909 \mathbf{M}_{i,t,j} = \mathcal{W}_{\theta'}(\mathbf{H}_{i,t,<j}). \quad (12)$$

910 Note that the generation of $\mathbf{M}_{i,t,j}$ only occurs if the fixed trigger \mathcal{T} determines an invocation is
 911 needed at step j . In steps where no memory is invoked, the conditioning term $\mathbf{M}_{i,t,j}$ is omitted. The
 912 gradients are computed exclusively with respect to the weaver’s parameters θ' and used to update
 913 them via gradient descent:

$$914 \theta' \leftarrow \theta' - \eta \nabla_{\theta'} \mathcal{L}_{\text{SFT}}(\theta'), \quad (13)$$

915 where η is the learning rate. Through this process, the memory weaver learns to synthesize latent
 916 memories that effectively steer the frozen reasoner’s generative process to align with the training
 917 data’s behavior, thereby internalizing the knowledge from the demonstration data without corrupting
 918 the general capabilities of the core LLM.

918 C.2 COMBINING MemGen WITH GRPO
919

920 The memory weaver can also be trained using a reinforcement learning objective. We specifically
921 adapt the GRPO algorithm. The training process begins by sampling a batch of tasks from the ex-
922 perience history \mathcal{H} . For each task x_i , we use the policy $\Pi_{\theta}^{\mathcal{W}_{\theta'}, \mathcal{T}}$ (composed of the frozen reasoner
923 π_{θ} and the current memory weaver $\mathcal{W}_{\theta'}$) to generate a group of K distinct trajectories, denoted as
924 $\mathcal{G}_i = \{\tau_{i,1}, \tau_{i,2}, \dots, \tau_{i,K}\}$. Each trajectory is generated by the agent's interaction with the environ-
925 ment and results in a final reward $R(\tau_{i,k})$, evaluated by the reward function. Going forward, GRPO
926 computes a group-relative baseline by averaging the rewards of all trajectories within the group \mathcal{G}_i :

$$927 \quad 928 \quad 929 \quad \bar{R}(\mathcal{G}_i) = \frac{1}{K} \sum_{k=1}^K R(\tau_{i,k}). \quad (14)$$

930 The advantage for a specific trajectory $\tau_{i,k}$ is then its reward relative to this baseline:
931

$$932 \quad A(\tau_{i,k}) = R(\tau_{i,k}) - \bar{R}(\mathcal{G}_i). \quad (15)$$

933 This formulation allows the weaver to learn by differentiating between better and worse outcomes
934 within a set of its own generations, promoting policies that produce trajectories with above-average
935 rewards. The final objective function, maximized with respect to the weaver's parameters θ' , is:
936

$$937 \quad 938 \quad \mathcal{J}_{\text{GRPO}}(\theta') = \mathbb{E}_{x_i \sim \mathcal{H}, \mathcal{G}_i \sim \Pi_{\theta}^{\mathcal{W}_{\theta'}, \mathcal{T}}} \left[\frac{1}{K} \sum_{k=1}^K A(\tau_{i,k}) \log \Pi_{\theta}^{\mathcal{W}_{\theta'}, \mathcal{T}}(\tau_{i,k} \mid x_i) - \beta \text{KL} \left(\Pi_{\theta}^{\mathcal{W}_{\theta'}, \mathcal{T}}(\cdot \mid x_i) \parallel \Pi_{\text{ref}}(\cdot \mid x_i) \right) \right], \quad (16)$$

939 where Π_{ref} is a fixed reference policy and β is a coefficient controlling the strength of the KL regu-
940 larization. The gradients are computed only for the weaver's parameters θ' , thus refining its ability
941 to synthesize impactful latent memories while preserving the integrity of the core reasoner.
942

943 C.3 MEMORY TRIGGER TRAINING DETAILS
944

945 We now detail the training procedure of the memory trigger $\mathcal{T}_{\text{trigger}}^{\phi}$. During this stage, the reasoner
946 π_{θ} and the memory weaver $\mathcal{W}_{\theta'}$ remain frozen, and only the trigger parameters ϕ are updated. For
947 each task x_i with initial environment state s_0 , we generate trajectories under the composed policy
948 induced by π_{θ} , $\mathcal{W}_{\theta'}$, and $\mathcal{T}_{\text{trigger}}^{\phi}$. At each decoding step t and token position j where $z_{t,j} \in \mathcal{D}$, the
949 trigger observes the hidden-state prefix $\mathbf{H}_{t,< j}$ and samples a binary action
950

$$951 \quad 952 \quad d_{t,j} \sim \pi_{\phi}(d \mid \mathbf{H}_{t,< j}) := \text{Bernoulli}(p_{t,j}), \quad p_{t,j} = \sigma(\mathcal{T}_{\text{trigger}}^{\phi}(\mathbf{H}_{t,< j})), \quad (17)$$

953 where $d_{t,j} = 1$ corresponds to `INVOKE`. If invoked, the decoder pauses and the weaver synthesizes
954 a latent memory \mathbf{M}_t from $\mathbf{H}_{t,< j}$, after which generation resumes based on the augmented context.
955 Let \mathcal{I}_i denote the set of activated positions in trajectory τ_i , and let $\tilde{\mathbf{d}}_i = \{d_{t,j}\}_{(t,j) \in \mathcal{I}_i}$ be the trigger
956 decisions within that rollout. After completing the episode, the environment returns a scalar reward
957 $R(\tau_i)$, which we assign uniformly to all decisions in $\tilde{\mathbf{d}}_i$.

958 To discourage excessively frequent invocation, we introduce a sparsity-inducing penalty. For a group
959 of rollouts generated from the same task x_i , we identify the high-reward subset
960

$$961 \quad \mathcal{H}_{\text{high}} = \{i : R(\tau_i) \geq \text{median}_k R(\tau_k)\}, \quad (18)$$

962 and compute the reference activation level

$$963 \quad 964 \quad 965 \quad \bar{p} = \frac{1}{|\mathcal{H}_{\text{high}}|} \sum_{i \in \mathcal{H}_{\text{high}}} \frac{1}{|\mathcal{I}_i|} \sum_{(t,j) \in \mathcal{I}_i} d_{t,j}, \quad (19)$$

966 which estimates the typical invocation frequency among successful trajectories. The trigger is
967 trained to maximize the reward while penalizing activations above \bar{p} :
968

$$969 \quad 970 \quad \max_{\phi} \mathbb{E}_{\tau_i \sim \pi_{\theta}, \tilde{\mathbf{d}}_i \sim \mathcal{T}_{\text{trigger}}^{\phi}} \left[R(\tau_i) - \lambda \sum_{(t,j) \in \mathcal{I}_i} \max(0, d_{t,j} - \bar{p}) \right], \quad (20)$$

971 where $\lambda > 0$ controls the sparsity regularization strength.

We optimize this objective following a GRPO-style procedure. For each task prompt, we generate a group of G trajectories $\{\tau_i^{(g)}\}_{g=1}^G$ (we set $G = 8$ in practice; Table 3). We compute per-trajectory rewards $R(\tau_i^{(g)})$ and construct normalized advantages

$$A_i^{(g)} = \frac{R(\tau_i^{(g)}) - \mu_i}{\sigma_i}, \quad \mu_i = \frac{1}{G} \sum_{g=1}^G R(\tau_i^{(g)}), \quad \sigma_i = \sqrt{\frac{1}{G} \sum_{g=1}^G (R(\tau_i^{(g)}) - \mu_i)^2},$$

thus measuring each rollout’s performance relative to the group. The trigger parameters are then updated by maximizing a GRPO-style surrogate objective:

$$\max_{\phi} \frac{1}{G} \sum_{g=1}^G \sum_{(t,j) \in \mathcal{I}_i^{(g)}} \log \pi_{\phi}(d_{t,j}^{(g)} \mid \mathbf{H}_{t,<j}^{(g)}) A_i^{(g)} - \beta \text{KL}(\pi_{\phi} \parallel \pi_{\phi_{\text{old}}}).$$

After each update, we set $\pi_{\phi_{\text{old}}} \leftarrow \pi_{\phi}$ and repeat. This objective allows the trigger to learn to invoke memory only at critical points that improve group-relative performance while maintaining sparsity.

D EXPERIMENTAL DETAILS

D.1 TRAINING DATASET SETUP

For all datasets except PopQA, we utilize their official training sets to train both the *memory weaver* and *memory trigger*. Specifically, we first train the memory weaver in isolation. Since the trigger module is not yet available at this stage, latent memory tokens are randomly inserted at punctuation marks during training to provide diverse learning signals. Once the memory weaver is adequately trained, we proceed to train the memory trigger, enabling it to learn when to invoke memory (with the memory weaver fixed). For PopQA, which lacks a dedicated training set, we instead leverage the trained model on TriviaQA and directly evaluate on PopQA. The dataset statistics are summarized in Table 2.

Table 2: Dataset Statistics.

Dataset Name	Train Size	Test Size	Note
ALFWorld	3.32K	134	From AgentBank-ALFWorld
TriviaQA	4.13K	7.9K	From AgentBank-TriviaQA
PopQA	-	14.3K	-
KodCode	8K	2K	KodCode-Light-RL-10K
BigCodeBench	912	228	-
GPQA	448	198	GPQA-Diamond for testing
GSM8K	7.47K	1K	-
MATH	1.6K	500	-

1026 D.2 PARAMETER CONFIGURATIONS
1027
10281029 Table 3: Hyperparameters used in the training of MemGen.
1030

1031	1032	Settings	Hyperparameters
1033	1034	1035	1036
1037	1038	1039	1040
1041	1042	1043	1044
1045	1046	1047	1048
1049	1050	1051	1052
1053	1054	1055	1056
1057	1058	1059	1060
1061	1062	1063	1064
1065	1066	1067	1068
1069	1070	1071	1072
1073	1074	1075	1076
1077	1078	1079	
Training(SFT)		train_batch_size = 4 learning_rate = 1e-5 epochs = 2 warmup_ratio = 0.1 optim = adamw_torch schedular = cosine	
Training(GRPO)		rollout_batch_size = 8 train_batch_size = 8 epochs = 2 beta = 0.0 num_iterations = 1 learning_rate = 1e-5 warmup_ratio = 0.1 optim = adamw_torch schedular = cosine	
LoRA		r = 16 lora_alpha = 32 target_modules = [q_proj, v_proj] lora_dropout = 0.1 bias = none task_type = CAUSAL_LM	
Optimization		adam_offload flash_attn deepspeed_enable_sleep	

E EXTRA RESULTS

1062 Table 4: Results on **Qwen2.5-1.5B**. All values represent the performance metric for each task (e.g.,
1063 accuracy %). We highlight the best and second best results.


1065	Method	ALFWorld	TrivialQA	PopQA	KodCode	BigCodeBench	GPQA	GSM8K	MATH
Qwen2.5-1.5B	Vanilla	22.54	32.10	16.08	24.55	40.35	11.62	39.51	36.63
	CoT	18.30	28.67	18.39	32.32	38.59	15.67	56.79	45.22
	SFT	36.57	63.84	39.20	55.83	37.72	11.11	54.83	38.84
	GRPO	43.55	68.21	43.15	62.11	70.34	15.65	68.10	47.42
	REINFORCE	43.25	66.50	41.87	60.20	67.80	12.50	67.40	46.89
	REINFORCE++	43.66	66.90	44.69	63.33	69.50	13.80	69.04	47.33
	Agent-FLAN	35.80	64.28	38.90	56.21	43.83	9.35	53.02	29.82
	ExpeL	28.96	25.20	20.20	31.15	39.78	8.12	45.12	38.12
	MemoryBank	27.89	38.14	22.78	37.93	35.87	13.87	47.88	30.47
	AWM	30.42	55.69	32.54	-	-	-	-	-

1080 E.1 CONTINUAL LEARNING RESULT
1081

1082 The results in Table 5 indicate three main findings. First, MemGen exhibits stronger knowledge re-
1083 tention and forward transfer than SFT and ExpeL. For example, when trained on GPQA, MemGen
1084 reaches 47.96% on GSM8K and 28.80% on KodCode, surpassing SFT at 45.74% and 18.50% by
1085 margins of +2.22% and +10.30%, respectively. Similarly, when trained on KodCode, MemGen
1086 maintains 40.34% on AQuA and 20.09% on GPQA, whereas SFT yields 28.61% and 2.53%. Sec-
1087 ond, MemGen demonstrates more balanced cross-task generalization. Unlike SFT, which primarily
1088 improves performance on the most recent task (e.g., 54.10% on KodCode but only 2.53% on GPQA),
1089 MemGen achieves competitive results across tasks. After training on GSM8K, it attains 38.43% on
1090 AQuA and 21.72% on GPQA, both higher than SFT and ExpeL, suggesting that latent memory
1091 captures task-invariant reasoning. Third, MemGen effectively mitigates forgetting on earlier tasks.
1092 After sequential training on KodCode, it preserves 40.34% on AQuA compared to 27.14% of ExpeL
1093 and 28.61% of SFT, showing stronger robustness to catastrophic forgetting. These findings suggest
1094 that MemGen provides a more stable and transferable mechanism for continual learning.

1095 E.2 TRIGGER FREQUENCY VISUALIZATION
1096


1097 More visualizations of the trigger frequency are displayed in Figures 7 and 8, where we paired
1098 Qwen2.5-1.5B or SmolLM3-3B with the GSM8K dataset and tested on subsets of GSM8K,
1099 KodCode, and GPQA (each subset having the same number of samples). We then tallied the
1100 frequency of memory trigger `Invoke` judgments at each relative percentile position in the LLM output.



1115 Figure 7: Memory invocation frequency across benchmarks at inference (trained on **MemGen**
1116 SFT+Qwen2.5-1.5B +GSM8K).
1117


1118 E.3 FRAMEWORK ANALYSIS
11191120 E.3.1 ABLATION STUDY
1121

1122 To assess the effectiveness of the proposed memory trigger, we conduct an ablation study comparing
1123 different memory invocation strategies, as summarized in Table 6. Specifically, *Random* denotes
1124 a naïve baseline where latent memory tokens are inserted at arbitrary token positions with a fixed
1125 probability p . *All delimiters activated* represents a sentence-level strategy that invokes the memory
1126 weaver at every delimiter position without any selection mechanism. Finally, *MemGen’s dedicated*
1127 *Trigger* corresponds to our standard approach, where the memory weaver is activated by the trained
1128 trigger T_{trigger} . The results reveal several key observations. First, sentence-level intervention already
1129 improves performance compared to random invocation. For instance, activating the weaver at all
1130 delimiters yields 17.34%, 56.20%, and 64.15% on GPQA, Kodcode, and TriviaQA, respectively,
1131 consistently outperforming all random baselines (e.g., $p = 0.5$ achieves only 16.66%, 52.95%, and
1132 57.28%). This highlights the importance of aligning memory injection with semantic boundaries
1133 rather than distributing it across the sequence. More importantly, our trained trigger achieves the
best overall performance, reaching 18.28%, 58.16%, and 65.02% on the three benchmarks. This




1148 Figure 8: Memory invocation frequency across benchmarks at inference (trained on `MemGen`
1149 `SFT+SmolLM3-3B +GSM8K`).

1150



1165 Figure 9: The generalization study of `MemGen`. We train `MemGen` `SFT` on `GSM8K` and evaluate it on
1166 all four datasets.

1167



1182 Figure 10: The generalization study of `MemGen`. We train `MemGen` `SFT` on `KodCode` and evaluate it
1183 on all four datasets.

1184

1185

1186 demonstrates that selective activation, *i.e.*, deciding *when* and *where* to weave in memory, provides
1187 the most effective support for reasoning, as it balances memory utility and interference more pre-
cisely than coarse-grained alternatives.

1188 Table 5: **Continual learning results** of Qwen2.5-1.5B-Instruct across four reasoning and
 1189 programming datasets (AQuA, GPQA, GSM8K, KodCode). The model is sequentially trained on
 1190 each dataset (AQuA → GPQA → GSM8K → KodCode), and after each training stage, evaluation
 1191 is conducted on all four benchmarks.

Trained On	Method	AQuA	GPQA	GSM8K	KodCode
	Vanilla	41.34	11.62	39.51	24.55
AQuA	SFT	42.52	16.67	42.10	18.20
	ExpeL	41.73	12.67	40.16	16.30
	MemGen SFT	43.31	19.70	39.80	19.55
GPQA	SFT	38.55	17.17	45.74	18.50
	ExpeL	37.24	14.35	42.67	15.20
	MemGen SFT	39.85	20.72	47.96	28.80
GSM8K	SFT	33.46	13.13	52.31	19.45
	ExpeL	34.89	12.42	48.78	13.65
	MemGen SFT	38.43	21.72	55.67	19.75
KodCode	SFT	28.61	2.53	24.14	54.10
	ExpeL	27.14	6.23	31.44	48.35
	MemGen SFT	40.34	20.09	53.72	52.95

1211 Table 6: Ablation study of different memory invocation strategies. Random denotes a naïve baseline
 1212 where latent memory tokens are inserted at arbitrary token positions with a fixed probability p . All
 1213 delimiters activated represents a sentence-level strategy that invokes the memory weaver at every de-
 1214 limiter position without any selection mechanism. Finally, **MemGen**’s dedicated Trigger corresponds
 1215 to our standard approach, where the memory weaver is activated by the trained trigger T_{trigger} .
 1216

Memory Invocation Strategy	GPQA	Kodcode	TriviaQA
Random ($p = 0.2$)	15.66	54.55	63.55
Random ($p = 0.5$)	16.66	52.95	57.28
Random ($p = 0.8$)	12.63	53.60	62.22
All delimiters activated	17.34	56.20	64.15
MemGen ’s dedicated Trigger	18.28	58.16	65.02

E.3.2 ANALYSIS OF MEMORY WEAVER

1228 We further investigate how the parameterization of the memory weaver influences downstream per-
 1229 formance. As shown in Table 7, increasing the number of trainable parameters by adopting a full-
 1230 parameter SFT setup surely enhances the weaver’s memory and learning capacity, leading to higher
 1231 task performance across benchmarks. Nevertheless, the LoRA-based instantiation already achieves
 1232 competitive results while retaining high parameter efficiency, demonstrating that even lightweight
 1233 adaptations can endow the weaver with sufficient capacity to generate effective latent memories.

E.3.3 EFFICIENCY ANALYSIS

1236 Table 8 reports average inference time and task performance across three benchmarks. Both SFT
 1237 and **MemGen** drastically reduce per-task inference time compared with vanilla models due to fewer
 1238 tokens required to reach correct answers. For instance, on KodCode with Qwen2.5-1.5B, **MemGen**
 1239 SFT completes tasks in 2.94 s, a 75.4% reduction from 11.96 s for vanilla, while improving accuracy
 1240 by 33.61%. On ALFWorld+Qwen3-8B, **MemGen** SFT adds only 1.6% more time compared with
 1241 SFT (20.08 s vs 19.76 s) but increases accuracy from 83.59% to 85.82%. These results confirm that
MemGen efficiently enhances reasoning performance without incurring significant inference delays.

1242
1243
1244
1245
12461247 Table 7: **Ablation study of the latent weaver.** We instantiate $\text{MemGen}_{\text{SFT}}$'s memory weaver with
1248 LoRA of different ranks as well as full-parameter SFT.
1249

1250	Base LLM: Qwen2.5-1.5B-Instruct	GPQA	Kodcode	TriviaQA
1251	LoRA ($r = 4$)	17.16	54.85	63.04
1252	LoRA ($r = 6$)	18.18	55.25	64.55
1253	LoRA ($r = 8$)	17.67	55.75	64.10
1254	LoRA ($r = 16$)	18.28	56.16	65.02
1255	LoRA ($r = 24$)	18.67	57.20	65.40
1256	LoRA ($r = 32$)	19.26	57.95	65.85
1257	Full SFT	21.21	60.00	67.10

1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
12701271 Table 8: Average per-task inference time (seconds) and task performance (accuracy %) across three
1272 benchmarks. Performance improvement (%) of $\text{MemGen}_{\text{SFT}}$ over vanilla LLM is also reported.
1273

1274	Model & Method	KodCode		ALFWorld		TriviaQA	
		1275 Time (s)	1275 Acc (%)	1275 Time (s)	1275 Acc (%)	1275 Time (s)	1275 Acc (%)
Qwen2.5-1.5B							
1277	Vanilla	11.96	24.55	21.17	22.54	2.18	32.10
1278	SFT	2.01	55.83	10.79	36.57	1.98	63.84
1279	$\text{MemGen}_{\text{SFT}}$	2.94	58.16	12.94	40.30	2.05	65.02
1280	Improvement over Vanilla	-	+33.61	-	+17.76	-	+32.92
SmolLM-3B							
1282	Vanilla	13.12	37.05	34.82	18.96	4.26	10.47
1283	SFT	3.04	59.25	12.88	32.36	3.05	55.25
1284	$\text{MemGen}_{\text{SFT}}$	3.48	62.65	14.69	50.60	3.16	68.13
1285	Improvement over Vanilla	-	+25.60	-	+31.64	-	+57.66
Qwen3-8B							
1287	Vanilla	16.99	49.10	55.42	58.93	8.70	52.18
1288	SFT	7.24	64.75	19.76	83.59	5.98	74.55
1289	$\text{MemGen}_{\text{SFT}}$	7.56	66.15	20.08	85.82	6.25	77.22
1290	Improvement over Vanilla	-	+17.05	-	+26.89	-	+25.04

1291
1292
1293
1294
1295

1296 **E.3.4 ABLATION OF MEMORY WEAVER: PAUSE TOKEN**
1297

1298 To further assess the necessity of the memory weaver, we introduce a strong and conceptually
1299 aligned baseline based on the *pause token* mechanism (Goyal et al., 2024). Pause tokens are a
1300 small set of learned special tokens $\mathcal{P} = \{\langle \text{pause} \rangle_1, \dots, \langle \text{pause} \rangle_K\}$ that instruct the model to
1301 suspend outward generation while continuing to update its internal hidden states. Formally, inserting
1302 a pause token $\langle \text{pause} \rangle$ at position j forces the model to compute a new hidden state

$$1303 \mathbf{h}_j = f_\theta(\mathbf{h}_{j-1}, \langle \text{pause} \rangle),$$

1304 while suppressing any semantic output. A sequence of K pause tokens thus yields K steps of
1305 latent computation over the current cognitive state. We train a set of K pause tokens (where K
1306 exactly equals the number of generated latent tokens by memory weaver) under the same protocol
1307 in Equation (10).

1308 **Results.** We report the ablation results as follows:

1309 Table 9: The performance comparison between **MemGen** and a variant by replacing mem-
1310 ory weaver with learnt pause tokens on three benchmarks. The backbone is set as
1311 `Qwen2.5-1.5B-Instruct`.

1315 Method	Kodcode	TriviaQA	GPQA
1316 Trigger + Pause Token	49.50	56.30	13.80
1317 Trigger + Weaver	58.16	65.02	18.28

1318 As shown in Table 9, although pause tokens provide moderate gains, indicating that allowing the
1319 model brief intervals of internal processing is beneficial, they consistently underperform the full
1320 weaver-equipped system. We attribute this to the fixed, context-independent nature of pause tokens:
1321 they cannot reconstruct or integrate task-relevant information in a targeted manner. In contrast, the
1322 memory weaver produces *context-sensitive* latent memory vectors tailored to the reasoner’s current
1323 cognitive state, leading to substantially stronger performance.

1325 **F INTEGRATION WITH RETRIEVAL-BASED MEMORY**
13261328 **F.1 FORMALIZING THE INTEGRATION PROCESS**

1329 While the primary mechanism of **MemGen** leverages the parametric knowledge encapsulated within
1330 the memory weaver $\mathcal{W}_{\text{weaver}}$, the framework is designed to be extensible, allowing for seamless
1331 integration with external, retrieval-based memory systems. This hybrid approach enables the weaver
1332 to synthesize latent memories that are informed by both its internalized experiences and a vast corpus
1333 of external information, thereby providing a richer and more comprehensive context to the reasoner
1334 π_θ . Let \mathcal{M}_{ext} denote an external memory database, and let $\mathcal{R}(\cdot)$ be a retrieval function that, given a
1335 natural language query, returns a set of relevant textual memory snippets. When the memory trigger
1336 $\mathcal{T}_{\text{trigger}}$ determines an invocation is necessary at the token-generation step j (i.e., $d_j = \text{INVOKE}$), the
1337 natural language text generated thus far serves as the query for the external memory system. This
1338 query, denoted as $q_{t,j}$, is produced by decoding the sequence of tokens generated up to that point:

$$1339 q_{t,j} = \text{Decode}(\mathbf{z}_{t,< j}). \quad (21)$$

1340 The retrieval process is then formalized as:

$$1342 \mathcal{C}_t = \mathcal{R}(q_{t,j}; \mathcal{M}_{\text{ext}}), \quad (22)$$

1343 where $\mathcal{C}_t = \{c_1, c_2, \dots, c_P\}$ is a set of P retrieved textual snippets. These snippets are subsequently
1344 encoded into a sequence of embeddings, $\mathbf{E}_t \in \mathbb{R}^{L_c \times d_{\text{model}}}$, where L_c is the total length of the encoded
1345 text. This allows the weaver to process the retrieved information in its native latent space.

1346 This retrieved information is subsequently merged with the reasoner’s internal cognitive state $\mathbf{H}_{t,< j}$.
1347 The combined context is then fed into the memory weaver $\mathcal{W}_{\text{weaver}}$ to produce the final latent mem-
1348 ory. The invocation of the weaver, as described in Equation (5), is thus modified to:

$$1349 \mathbf{M}_t = \mathcal{W}_{\text{weaver}}([\mathbf{H}_{t,< j}; \mathbf{E}_t]), \quad (23)$$

1350 where $[\cdot; \cdot]$ denotes the concatenation of the hidden state sequences. This integrated process allows
 1351 $\mathcal{W}_{\text{weaver}}$ to reconstruct both internal parametric knowledge and externally retrieved information into
 1352 a compact, potent latent memory \mathbf{M}_t for the reasoner.
 1353

1354 F.2 EXPERIMENTAL RESULTS

1355 As shown in Table 10, even when **MemGen**’s own parametric memory is disabled (which means
 1356 that merely the retrieved textual snippets are fed into $\mathcal{W}_{\text{weaver}}$), **MemGen** significantly enhances the
 1357 retrieval baseline, boosting performance on ALFWorld from 36.18% to 45.60% and on PopQA
 1358 from 28.16% to 39.50%. This demonstrates that **MemGen** serves as a powerful synthesizer, not
 1359 merely appending but actively reconstructing externally retrieved information into a more potent,
 1360 task-relevant latent memory, thereby underscoring the necessity of a generative memory mechanism.
 1361 Most notably, when the weaver’s internal parametric memory is enabled alongside external retrieval,
 1362 performance is dramatically amplified, achieving scores of 75.90% on ALFWorld and 76.40% on
 1363 TriviaQA. This highlights the powerful synergy achieved when the memory weaver is empowered
 1364 to dynamically integrate its own internalized experiences with external textual knowledge.
 1365

1366 Table 10: Performance results when integrating **MemGen** with a retrieval-based memory baseline,
 1367 ExpeL. The backbone LLM is set as SmolLM3-3B.

1368 Method	1369 ALFWorld	1370 TriviaQA	1371 PopQA
1370 Vanilla LLM	18.96	10.47	8.23
1371 ExpeL	36.18	46.20	28.16
1372 MemGen + ExpeL (w/o parametric memory)	45.60	53.20	39.50
1372 MemGen + ExpeL (w/ parametric memory)	75.90	76.40	60.23

1375 G LATENT MEMORY TOKEN DEMONSTRATION

1378 Cases from TriviaQA

1379
 1380 **Case 1**
system prompt
 1381 Answer the given question. You must conduct reasoning inside `<think>` and `</think>` first
 1382 every time you get new information. After reasoning, if you find you lack some
 1383 knowledge, you can call a search engine by `<search>` query `</search>` and it will
 1384 return the top searched results between `<information>` and `</information>`. You can
 1385 search as many times as your want. If you find no further external knowledge
 1386 needed, you can directly provide the answer inside `<answer>` and `</answer>`, without
 1387 detailed illustrations. For example, `<answer> Beijing </answer>`.

1388 **user prompt**
 1389 Which Lloyd Webber musical premiered in the US on 10th December 1993?

1390 **assistant output**
 1391 `[UPPORT...',eniable certif]<think>` I need to find out which Lloyd Webber musical
 1392 premiered in the US on 10th December 1993. `</think>`
 1393 `[lessengeryyyyMMddELCOME certif]<search>` Lloyd Webber musical premiered in the US on 10
 1394 th December 1993 `</search>`

1395 **user prompt**
 1396 `<observation>` Doc 1>Title: "Andrew Lloyd Webber") opened at the Adelphi Theatre in
 1397 London on 12 July 1993, and ran for 1,529 performances ... `</observation>`

1398 **assistant output**
 1399 `[lessengeryyyyMMddELCOME certif]<think>` The search results indicate that "Sunset Boulevard" is the Lloyd
 1400 Webber musical that premiered in the US on 10th December 1993. `</think>`
 1401 `[UPPORT...',eniable certif]<answer>` Sunset Boulevard `</answer>`

1402 **Case 2**
system prompt
 1403 Answer the given question. You must conduct reasoning inside `<think>` and `</think>` first
 1404 every time you get new information. After reasoning, if you find you lack some
 1405 knowledge, you can call a search engine by `<search>` query `</search>` and it will
 1406 return the top searched results between `<information>` and `</information>`. You can
 1407 search as many times as your want. If you find no further external knowledge

```

1404
1405     needed, you can directly provide the answer inside <answer> and </answer>, without
1406     detailed illustrations. For example, <answer> Beijing </answer>.
1407
1408 user prompt
1409 Who had a 70s No 1 hit with Kiss You All Over?
1410
1411 assistant output
1412     [UPPORT...',eniable certif]<think> I need to search who had a 70s No 1 hit with the
1413     song "Kiss You All Over." </think>
1414     [lessengeryyyyMddELCOME
1415     certif]<search> who had a 70s No 1 hit with Kiss You All Over? </search>
1416
1417 user prompt
1418 <observation> Doc 1>Title: "Kiss You All Over") Exile's only big hit in the pop rock
1419     market (they would later have great success on the country music charts)... </
1420     observation>
1421
1422 assistant output
1423     [lessengerUPPORT'options certif]<think> The song "Kiss You All Over" was a 70s No 1 hit
1424     with the band Exile. </think>
1425     [UPPORT...',eniable certif]<answer> Exile </answer>
1426
1427 Case 3
1428 system prompt
1429 Answer the given question. You must conduct reasoning inside <think> and </think> first
1430     every time you get new information. After reasoning, if you find you lack some
1431     knowledge, you can call a search engine by <search> query </search> and it will
1432     return the top searched results between <information> and </information>. You can
1433     search as many times as you want. If you find no further external knowledge
1434     needed, you can directly provide the answer inside <answer> and </answer>, without
1435     detailed illustrations. For example, <answer> Beijing </answer>.
1436
1437 user prompt
1438 What was the name of Michael Jackson's autobiography written in 1988?
1439
1440 assistant output
1441     [UPPORTUPPORTeniable certif]<think> I need to search for the name of Michael Jackson's
1442     autobiography written in 1988. </think>
1443     [UPPORT...',eniable certif]<search> Michael Jackson autobiography 1988 </search>
1444
1445 user prompt
1446 <observation> Doc 1>Title: "Michael Jackson" of 4.4 million people. In 1988, Jackson
1447     released his only autobiography, "Moonwalk", which took four years to complete
1448     and sold 200,000 copies... </observation>
1449
1450 assistant output
1451     [lcmsUPPORT'options Certif]<think> The name of Michael Jackson's autobiography written
1452     in 1988 is "Moonwalk". </think>
1453     [UPPORTUPPORTeniable certif]<answer> Moonwalk </answer>
1454
1455 Case 4
1456 system prompt
1457 Answer the given question. You must conduct reasoning inside <think> and </think> first
1458     every time you get new information. After reasoning, if you find you lack some
1459     knowledge, you can call a search engine by <search> query </search> and it will
1460     return the top searched results between <information> and </information>. You can
1461     search as many times as you want. If you find no further external knowledge
1462     needed, you can directly provide the answer inside <answer> and </answer>, without
1463     detailed illustrations. For example, <answer> Beijing </answer>.
1464
1465 user prompt
1466 Which volcano in Tanzania is the highest mountain in Africa?
1467
1468 assistant output
1469     [UPPORT...',eniable certif]<think> I need to search for the highest mountain in Africa,
1470     which is located in Tanzania. </think>
1471     [UPPORTUPPORTeniable certif]<search> highest mountain in Africa in Tanzania </search>
1472
1473 user prompt
1474 <observation> Doc 1>Title: "Mount Kilimanjaro" Mount Kilimanjaro Mount Kilimanjaro or
1475     just Kilimanjaro (), with its three volcanic cones, "Kibo", "Mawenzi", and "Shira",
1476     is a dormant volcano in Tanzania... </observation>
1477
1478 assistant output
1479     [lessengeryyyyMddeniable certif]<think> The highest mountain in Africa located in
1480     Tanzania is Mount Kilimanjaro. </think>
1481     [UPPORT...',eniable certif]<answer> Mount Kilimanjaro </answer>

```

1458

Cases from KodCode

1459

1460

Case 1

1461

MODEL INPUT:

1462

Write an efficient and correct Python function to solve the following problem. Question:
 Write a function that takes a list of integers and returns a new list containing only the even integers from the original list, sorted in non-decreasing order. The list can contain negative numbers and is not necessarily sorted. If the input list is empty or contains no even integers, the function should return an empty list.

1466

MODEL OUTPUT:

1467

```
[keyword-kindergötAs-slide]def even_sorted(lst):
[.keyword_pick] """
[LTRetrical] Returns a new list containing only the even integers from the original list
,
sorted in non-decreasing order.
[.keyword_pick] """
[LTRetrical] even_numbers = [num for num in lst if num % 2 == 0]
[.keyword_pick] even_numbers.sort()
[.keyworddetical] return even_numbers<|im_end|>
```

1473

Case 2

1474

MODEL INPUT:

1475

Write an efficient and correct Python function to solve the following problem. Question:
 Write a python function to find the first maximum and the second maximum number in an array. If the array has less than two unique elements, return 'None' for the missing values.

1478

MODEL OUTPUT:

1479

```
[keyword-kindergötAs-slide]def find_first_max_second_max(arr):
"""
[.keyworddetical] Returns the first and second maximum numbers in the array.
If the array has less than two unique elements, [.keyword_pick] returns None for the
missing values.
[.keyword_pick] """
if len(arr) < 2:
[.keyword_pick] return None, [.keyword_pick] None

[.keyworddetical] unique_elements = set(arr)
[.keyword_pick] if len(unique_elements) < 2:
[.keyworddetical] return None, [.keyword_pick] None

[.keyworddetical] first_max = max(unique_elements)
second_max = max(unique_elements, default=None)

return first_max, second_max<|im_end|>
```

1490

Case 3

1491

MODEL INPUT:

1492

Write an efficient and correct Python function to solve the following problem. Question:
 Given a list of integers, write a function that returns the length of the longest contiguous subarray where all elements have the same value. The function should also return the value of the elements in this subarray. Consider the output to be a tuple where the first element is the length of the longest subarray and the second element is the value of the elements in the subarray.

1497

For example, consider the following list as input,
 numbers = [1, 2, 2, 3, 3, 2, 2, 2, 1, 1]

1498

Expected Output: (4, 2)

1500

MODEL OUTPUT:

1501

```
[kindgetAs?-kindergötAs-slide]def longest_subarray_with_same_value(numbers):
"""
```

1502

[.keyword_pick] Returns the length of the longest contiguous subarray where all
 elements have the same value.

1503

[.keyword_pick] The function also returns the value of the elements in this subarray.

1504

[LTR_pick] """

1505

[LTRetrical] if not numbers:

1506

[.keyword_pick] return (0, None)

1507

max_length = 1

1508

[.keyword_pick] current_length = 1

1509

max_value = numbers[0]

1510

current_value = numbers[0]

1511

[.keyword_pick] for i in range(1, [LTR_pick] len(numbers)):

[LTR_pick] if numbers[i] == current_value:

```

1512
1513     current_length += 1
1514     else:
1515         if current_length > max_length:
1516             [LTR.pick] max_length = current_length
1517             max_value = current_value
1518             current_length = 1
1519             current_value = numbers[i]
1520
1521         if current_length > max_length:
1522             max_length = current_length
1523             max_value = current_value
1524
1525     return (max_length, max_value)<|im_end|>

```

Cases from GSM8K

Case 1

MODEL INPUT:
Solve the math problem with proper reasoning, and make sure to put the FINAL ANSWER inside \boxed{}.

Question: Billy sells DVDs. He has 8 customers on Tuesday. His first 3 customers buy one DVD each. His next 2 customers buy 2 DVDs each. His last 3 customers don't buy any DVDs. How many DVDs did Billy sell on Tuesday?

MODEL OUTPUT:
[even. s. even even ...
even] Billy's first 3 customers buy $3 * 1 = <<3*1=3>>3$ DVDs.
[-gallery.] His next 2 customers buy $2 * 2 = <<2*2=4>>4$ DVDs.
[-routing.] Billy sells $3 + 4 = <<3+4=7>>7$ DVDs on Tuesday.\boxed{7}<|im_end|>

Case 2

MODEL INPUT:
Solve the math problem with proper reasoning, and make sure to put the FINAL ANSWER inside \boxed{}.

Question: Henry made two stops during his 60-mile bike trip. He first stopped after 20 miles. His second stop was 15 miles before the end of the trip. How many miles did he travel between his first and second stops?

MODEL OUTPUT:
[..... a even even even... even
] Henry traveled $60 - 15 = <<60-15=45>>45$ miles before his second stop.
[-android.] He traveled 20 miles before his first stop, [ecies.] so he traveled $45 - 20 = <<45-20=25>>25$ miles between his first and second stops.\boxed{25}<|im_end|>
[even.] The final answer is 25<|im_end|>

Case 3

MODEL INPUT:
Solve the math problem with proper reasoning, and make sure to put the FINAL ANSWER inside \boxed{}.

Question: A basket contains 25 oranges among which 1 is bad, 20% are unripe, 2 are sour and the rest are good. How many oranges are good?

MODEL OUTPUT:
[. even.. pre. even even even... even
] The number of unripe oranges is $20/100 * 25 = <<20/100*25=5>>5$
[-animate.] The number of sour oranges is 2
[-login.] The number of bad oranges is 1
[-login.] The number of good oranges is $25 - 5 - 2 - 1 = <<25-5-2-1=17>>17$ \boxed{17}<|im_end|>

Cases from GSM8K Cluster

Cluster 1

Solve the math problem with proper reasoning, and make sure to put the FINAL ANSWER inside \boxed{}.

Question: Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?
[..... pre. even even even... even] The house was worth \$80

Solve the math problem with proper reasoning, and make sure to put the FINAL ANSWER inside \boxed{}.

Question: James decides to run 3 sprints 3 times a week. He runs 60 meters each sprint. How many total meters does he run a week?
[..... out. even even even.. even even] He runs $3 * 60 = <<3*60=180>>180$ meters a day

1566
 1567 Solve the math problem with proper reasoning, and make sure to put the FINAL ANSWER
 1568 inside \boxed{}. Question: Kylar went to the store to buy glasses for his new
 1569 apartment. One glass costs \$5, but every second glass costs only 60% of the price.
 1570 Kylar wants to buy 16 glasses. How much does he need to pay for them?
 [..... pre... even... even]The first 16 glasses cost $16 * 5 = \$<<16*5=80>>80$

1571 Solve the math problem with proper reasoning, and make sure to put the FINAL ANSWER
 1572 inside \boxed{}. Question: Toulouse has twice as many sheep as Charleston.
 1573 Charleston has 4 times as many sheep as Seattle. How many sheep do Toulouse,
 1574 Charleston, and Seattle have together if Seattle has 20 sheep?
 [even even. s. even even even.. even even
 1575]Charleston has $4*20=<<4*20=80>>80$ sheep.

1576 Solve the math problem with proper reasoning, and make sure to put the FINAL ANSWER
 1577 inside \boxed{}. Question: John drives for 3 hours at a speed of 60 mph and then
 1578 turns around because he realizes he forgot something very important at home. He
 1579 tries to get home in 4 hours but spends the first 2 hours in standstill traffic.
 1580 He spends the next half-hour driving at a speed of 30mph, before being able to
 1581 drive the remaining time of the 4 hours going at 80 mph. How far is he from home
 1582 at the end of those 4 hours?
 [..... even even even... even
 1583]He drives $60*3 = <<60*3=180>>180$ miles in the first 3 hours.

1584 **Cluster 2**
 1585 Janet's ducks lay 16 eggs per day, [-animate.] and she eats 3 eggs for breakfast every
 1586 morning
 1587 Wendi feeds her 20 chickens 15 cups of feed in the morning, [-animate.] 25 cups in the
 1588 afternoon
 1589 She sold half of what was left at the orange house, [-animate.] she had 1/2 left.
 1590 $x = 162/18 = <<162/18=9>>9$
 1591 [-animate.] Allen's age is $11x = 11*9 = <<11*9=99>>99$ years old now.
 1592 He runs 20 miles a day because $60/3=<<60/3=20>>20$
 1593 [-animate.] He runs 10 miles the other two days because $20/2=<<20/2=10>>10$

1594 **Cluster 3**
 1595 It takes $2/2=<<2/2=1>>1$ bolt of white fiber.
 1596 [ecies.] So it takes $2+1=<<2+1=3>>3$ bolts in total.
 1597 The total cost of the five cell phones is $\$150 \times 5 = \$<<150*5=750>>750$.
 1598 [even.] The interest for each unit is $\$150 \times 2/100 = \$<<150*2/100=3>>3$.
 1599 The second level is $16*2=<<16*2=32>>32$ square footage.
 1600 [even.] The third level is $32*2=<<32*2=64>>64$ square footage.
 1601 20% of 20 students is $20 \times 20/100 = <<20*20/100=4>>4$ students.
 1602 [ecies.] The remaining students are $20 - 4 = <<20-4=16>>16$ students.
 1603 If Artie rounds the price of petunias to the nearest dollar, [-animateecies] he will
 1604 sell them for \$2 each.

1605 **Cluster 4**
 1606 The groomer offers a 30% discount, [-loginecies] so the cost of the grooming is $100 -$
 1607 $(100 * 0.[ecies.]30) = \$<<100-30=70>>70$
 1608 If Dana skips at 3 miles per hour, [-login.] then she runs at $3 * 4 = <<3*4=12>>12$ miles
 1609 per hour.
 1610 There were $140 * 2 = <<140*2=280>>280$ emeralds in the chest.
 1611 [-login.] There were $175 + 140 + 280 = <<175+140+280=595>>595$ gems in the chest.
 1612 There are $30+60 = <<30+60=90>>90$ students at Dr. [-loginecies.] Wertz's school.
 1613 Cecilia will feed her puppy 1 cup of dog food every day for 180 days, [-login.] so she
 1614 will use 180 cups of dog food in the first 180 days.

H MEMORY FUNCTIONAL STUDY

H.1 VISUALIZATION PROCESS

1615 This section details the methodology employed to visualize the geometric structure of the latent
 1616 memory (as shown in Figures 5 and 11) and to identify semantically distinct clusters within it. The




Figure 11: (Up) t-SNE visualization of latent memories generated by **MemGen** +SmolLM3-3B across datasets; (Down) Latent memory visualization within the TriviaQA and GSM8K datasets, clustered using K -means.

process involves two primary stages: obtaining a fixed-dimensional representation for each memory sequence and applying dimensionality reduction and clustering algorithms.

Sequence Representation. Let $\mathcal{D}_{\text{mem}} = \{\mathbf{M}_1, \mathbf{M}_2, \dots, \mathbf{M}_Q\}$ be a corpus of Q latent memory sequences collected from **MemGen**'s inference, where each sequence $\mathbf{M}_i = (\mathbf{m}_{i,1}, \dots, \mathbf{m}_{i,K}) \in \mathbb{R}^{K \times d_{\text{model}}}$ consists of K token embeddings of dimension d_{model} . To obtain a single, holistic vector representation for each sequence, we compute its mean embedding $\bar{\mathbf{m}}_i \in \mathbb{R}^{d_{\text{model}}}$:

$$\bar{\mathbf{m}}_i = \frac{1}{K} \sum_{l=1}^K \mathbf{m}_{i,l}. \quad (24)$$

This procedure yields a set of high-dimensional representations $\bar{\mathcal{M}} = \{\bar{\mathbf{m}}_1, \bar{\mathbf{m}}_2, \dots, \bar{\mathbf{m}}_Q\}$, which serves as the input for the subsequent analysis.

For visualization, we employed t-SNE (Maaten & Hinton, 2008) to project the high-dimensional set $\bar{\mathcal{M}}$ into a two-dimensional space. Formally, this mapping can be expressed as $\mathbf{y}_i \in \mathbb{R}^2 = f_{\text{t-SNE}}(\bar{\mathbf{m}}_i)$. These 2D data points are then utilized for the scatter plot.

To quantitatively identify distinct functional groups, we applied the K-means algorithm directly to the high-dimensional mean embeddings in $\bar{\mathcal{M}}$. This procedure partitions the memory representations into a predefined number of N discrete clusters, $\mathcal{C} = \{C_1, C_2, \dots, C_N\}$ (we set $N = 4$).

H.2 FAILURE TAXONOMY DEFINITIONS

To systematically analyze the functional impact of different memory clusters, we adopt and refine the failure taxonomy proposed by Song et al. (2025). Below, we provide precise definitions of

- 1674 each failure category in the context of LLM-based agent behavior, which guide our annotation and
 1675 evaluation process:
 1676
- 1677 • **Planning Failure.** This category refers to errors in high-level task decomposition and strategic
 1678 decision-making. The agent either formulates subgoals that do not align with the original objective,
 1679 fails to plan intermediate steps necessary for task completion, or misorders the reasoning
 1680 sequence, leading to suboptimal or incomplete solutions.
 - 1681 • **Compositional Reasoning.** This type of failure occurs when the agent struggles to integrate
 1682 multiple pieces of information or reasoning steps into a coherent solution.
 - 1683 • **Tool Parsing Error.** This failure occurs when the agent’s generated output cannot be parsed
 1684 into a valid tool call. Typical patterns include malformed function signatures, missing required
 1685 arguments, or unrequired arguments in the invocation, which prevent the external tool from being
 1686 executed as intended.
 - 1687 • **Tool Response Error.** This category refers to errors that arise after a tool has been successfully
 1688 invoked. Either the tool itself returns incorrect or incomplete information, or the agent misuses or
 1689 misinterprets the returned content.
 - 1690 • **Answer Formatting Failure.** This category includes errors in the final presentation or structuring
 1691 of the output, despite the reasoning process being largely correct. Examples include incorrect
 1692 output format (*e.g.*, unable to include the answers with in $\boxed{\text{ }}$), violation of task-specific
 1693 response templates, or missing required components in the final answer.
 - 1694 • **Demand Misunderstanding.** This failure indicates that the agent has misinterpreted the user’s
 1695 intent or the task specification. It may solve a different problem from the one posed, ignore key
 1696 constraints, or pursue irrelevant objectives due to misunderstanding the instruction semantics.
 - 1697 • **Think-Act Inconsistency.** This type refers to a mismatch between the agent’s internal reasoning
 1698 and its subsequent action in ReAct mode. The agent may articulate a correct reasoning chain but
 1699 execute a contradictory action or produce a final answer inconsistent with its prior deliberation.
 - 1700 • **False Belief** captures cases where the agent maintains and reasons with inaccurate assumptions
 1701 about the external environment, user state, or task context.

1703 H.3 ANNOTATING FAILURE MODES AND FILTERING LATENT MEMORY

1705 **Annotation of Failure Modes.** We manually annotated agent failures on the TriviaQA dataset
 1706 based on the eight failure modes discussed above. Each trajectory generated by the agent was evaluated
 1707 by human annotators who assigned only one failure label if the agent’s behavior deviated from
 1708 a successful path. Trajectories exhibiting failures that did not fall into these predefined categories
 1709 were excluded from this specific analysis.

1710 **Inference-time Filtering of Latent Memory Clusters.** The core of our intervention study involved
 1711 selectively removing the influence of a specific latent memory cluster during the agent’s reasoning
 1712 process. Our methodology is designed to be consistent with the clustering process itself,
 1713 which operates on sequence-level representations.

1714 Let $\mathcal{C} = \{C_1, C_2, \dots, C_N\}$ be the set of N latent memory clusters. These clusters were derived by
 1715 applying K -means to a collection of historical memory representations, where each representation
 1716 is the mean embedding of an entire latent memory sequence. For each cluster $C_i \in \mathcal{C}$, we compute
 1717 its centroid $\mu_i \in \mathbb{R}^{d_{\text{model}}}$ by averaging these historical sequence representations within the cluster.

1718 During inference, when the memory weaver synthesizes a new latent memory sequence $\mathbf{M}_t =$
 1719 $(\mathbf{m}_{t,1}, \dots, \mathbf{m}_{t,K}) \in \mathbb{R}^{K \times d_{\text{model}}}$, we first compute its single-vector representation, $\bar{\mathbf{m}}_{\text{new}}$, by averaging
 1720 its constituent token embeddings:

$$1722 \bar{\mathbf{m}}_{\text{new}} = \frac{1}{K} \sum_{l=1}^K \mathbf{m}_{t,l}. \quad (25)$$

1723 To determine the semantic affiliation of this new sequence, we compare its mean embedding $\bar{\mathbf{m}}_{\text{new}}$
 1724 against a comprehensive reference set $\mathcal{E}_{\text{comp}} = \mathbf{E}_{\text{vocab}} \cup \{\mu_1, \dots, \mu_N\}$, where $\mathbf{E}_{\text{vocab}} \in \mathbb{R}^{V \times d_{\text{model}}}$
 1725 is the LLM’s vocabulary embedding matrix. We then identify the set of top- k nearest neighbors to
 1726 $\bar{\mathbf{m}}_{\text{new}}$ based on cosine similarity, denoted as $\mathcal{S}_k(\bar{\mathbf{m}}_{\text{new}})$. In our experiments, we set $k = 10$.

To ablate the influence of a target cluster C_j , the entire latent memory sequence \mathbf{M}_t is filtered (i.e., discarded and not prepended to the reasoner’s context) if the centroid of that cluster, μ_j , is present within this top- k set. Formally, \mathbf{M}_t is filtered if:

$$\mu_j \in \mathcal{S}_k(\bar{\mathbf{m}}_{\text{new}}). \quad (26)$$

This sequence-level filtering allows us to precisely ablate the contribution of a specific learned memory function and observe its impact on agent behavior.

H.4 COMPARISON BETWEEN LATENT TOKENS AND TRAJECTORIES

For each memory invocation at step j and timestep t , we construct three sets of embeddings within the hidden representation space of the frozen reasoner π_θ . First, we record the latent memory tokens generated by the weaver,

$$\mathbf{M}_t = [\mathbf{m}_{t,1}, \dots, \mathbf{m}_{t,K}] \in \mathbb{R}^{K \times d_{\text{model}}}.$$

Second, we obtain the corresponding context hidden states that serve as the hook for memory generation,

$$\mathbf{H}_{t,< j} = (\mathbf{h}_{t,1}, \dots, \mathbf{h}_{t,j-1}) \in \mathbb{R}^{(j-1) \times d_{\text{model}}},$$

collected directly from the forward pass of π_θ during online reasoning. Third, we take a past successful trajectory τ_{prev} from \mathcal{H} and replay it under teacher forcing through π_θ , yielding its hidden-state sequence

$$\tilde{\mathbf{H}}^{\text{traj}} = (\tilde{\mathbf{h}}_1, \dots, \tilde{\mathbf{h}}_L) \in \mathbb{R}^{L \times d_{\text{model}}}.$$

To enable joint visualization, we compute the mean embedding of each set:

$$\bar{\mathbf{m}}_t = \frac{1}{K} \sum_{k=1}^K \mathbf{m}_{t,k}, \quad \bar{\mathbf{h}}_t^{\text{ctx}} = \frac{1}{j-1} \sum_{k=1}^{j-1} \mathbf{h}_{t,k}, \quad \bar{\mathbf{h}}^{\text{traj}} = \frac{1}{L} \sum_{\ell=1}^L \tilde{\mathbf{h}}_\ell.$$

Finally, we project the three averaged vectors

$$\{\bar{\mathbf{m}}_t, \bar{\mathbf{h}}_t^{\text{ctx}}, \bar{\mathbf{h}}^{\text{traj}}\}$$

from $\mathbb{R}^{d_{\text{model}}}$ into two dimensions (via t-SNE) for visualization. The resulting visualizations are shown in Figure 12.




Figure 12: The t-SNE visualizations among the learnt latent tokens, trajectory representations and context representations. The left is drawn from KodCode dataset, and the right is from TriviaQA dataset.

As shown in Figure 12, the trajectory and context embeddings occupy more similar regions of the feature space, as both originate from transformations of human-readable tokens. In contrast, latent memory tokens form a more machine-native and human-unreadable representation, resulting in a substantially different embedding distribution.

We further quantify the geometric relationship between the three clusters by computing the Euclidean distances between their centers in $\mathbb{R}^{d_{\text{model}}}$. For each invocation (t, j) , we treat $\bar{\mathbf{m}}_t$, $\bar{\mathbf{h}}_t^{\text{ctx}}$, and $\bar{\mathbf{h}}^{\text{traj}}$ as the centers of the latent-memory, context, and trajectory clusters, respectively, and compute the pairwise distances

$$d_{\text{mem-ctx}}^{(t)} = \|\bar{\mathbf{m}}_t - \bar{\mathbf{h}}_t^{\text{ctx}}\|_2, \quad d_{\text{mem-traj}}^{(t)} = \|\bar{\mathbf{m}}_t - \bar{\mathbf{h}}^{\text{traj}}\|_2, \quad d_{\text{ctx-traj}}^{(t)} = \|\bar{\mathbf{h}}_t^{\text{ctx}} - \bar{\mathbf{h}}^{\text{traj}}\|_2.$$

1782
1783Aggregating over all invocations \mathcal{T} , we report the mean distances1784
1785
1786

$$\bar{d}_{\text{mem-ctx}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} d_{\text{mem-ctx}}^{(t)}, \quad \bar{d}_{\text{mem-traj}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} d_{\text{mem-traj}}^{(t)}, \quad \bar{d}_{\text{ctx-traj}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} d_{\text{ctx-traj}}^{(t)},$$

1787
1788

which directly reflect how far the latent memory centers deviate from the context and replayed-trajectory manifolds.

1789
1790
1791

Table 11: Pairwise Euclidean distances between the centers of context, latent-memory (inserted), and trajectory clusters.

1792
1793
1794
1795

Dataset	context \leftrightarrow inserted	context \leftrightarrow trajectory	inserted \leftrightarrow trajectory
GPQA	165.311	64.180	185.315
KodCode	227.108	96.831	197.105

1796

From Table 11, the distances reveal a consistent geometric pattern: the context and trajectory centers lie extremely close, while the generated latent memories occupy a distant region of the representation space, forming a well-separated cluster from both context and trajectory. This large separation demonstrates that the weaver does not merely compress or replay past hidden states; instead, it synthesizes novel latent structures that are parametrically reconstructed rather than retrieved, supporting our claim that latent memory introduces genuinely new inferential content beyond the observed trajectories.

1803

1804

1805

1806

1807

1808

1809

1810

1811

1812

1813

1814

1815

1816

1817

1818

1819

1820

1821

1822

1823

1824

1825

1826

1827

1828

1829

1830

1831

1832

1833

1834

1835