

Learning to Retrieve: Dual-Level Long-Term Memory for Text-to-SQL Agents

Yibo Wang

University of Illinois Chicago
ywang633@uic.edu

Nikki Lijing Kuang

Snowflake AI Research
nikki.kuang@snowflake.com

Philip S. Yu

University of Illinois Chicago
psyu@uic.edu

Zhewei Yao

Snowflake AI Research
zhewei.yao@snowflake.com

Yuxiong He

Snowflake AI Research
yuxiong.he@snowflake.com

Abstract

Interactive text-to-SQL agents solve database tasks through multi-turn interactions involving schema exploration, query execution, feedback interpretation, and decision revision. Long-term memory helps agents reuse past experiences, but existing retrieval methods remain limited. Static methods use fixed similarity heuristics, whereas dynamic methods often learn from sparse terminal outcomes and retrieve at a single decision horizon. This is insufficient when memory usefulness changes across interaction stages, since memories useful for initial planning may differ from those needed for local, state-conditioned execution. We propose **MERIT**, a dynamic multi-horizon memory retrieval framework. **MERIT** maintains episode-level memory for global strategic guidance and turn-level memory for local decision support. Both levels use learned retrieval policies optimized with reinforcement learning. To train turn-level retrieval despite limited intermediate supervision, **MERIT** uses a lightweight Process Reward Model to provide dense proxy rewards for local memory selection. On BIRD-Interact, **MERIT** achieves higher success rates and fewer interaction turns than no-memory, static-retrieval, and dynamic-retrieval baselines. Transfer results on Spider2-Snow further show that the learned policies transfer without degradation, staying above all memory baselines without benchmark-specific tuning. These results suggest that multi-horizon retrieval improves experience reuse in interactive text-to-SQL agents. Our code is available at <https://github.com/YiboWANG214/MERIT>.


1 Introduction

Large language models (LLMs) are increasingly deployed as multi-step interactive agents that act, observe, and refine their behavior to solve complex tasks (Li, 2025; Mohammadi et al., 2025; Xu et al., 2025; Yuan et al., 2025). In text-to-SQL, this shift moves the task beyond single-pass semantic pars-


ing. Instead of producing SQL in a single step, an agent may need to inspect schemas, execute candidate SQL queries, ask clarifying questions, observe feedback, and revise its plan over time. Thus, success depends not only on generating a valid final query but also on making reliable intermediate decisions under evolving context (Wang et al., 2025; Shao et al., 2025).

Long-term memory provides a natural mechanism for improving such decisions. Interactive text-to-SQL tasks often share recurring structures and solution patterns, including join templates, aggregation logic, clarification strategies, and debugging routines. Reusing these experiences can help agents avoid redundant exploration and make informed decisions in new interactions. However, existing memory-augmented approaches commonly rely on static similarity heuristics (Maharana et al., 2024; Sun et al., 2026; Salama et al., 2025; Yan et al., 2026; Kweon et al., 2025), such as dense embedding similarity, lexical matching, or hand-designed schema heuristics. These heuristics are useful for finding related candidates, yet do not directly optimize whether a retrieved memory will improve the next decision.

Recent dynamic frameworks address this limitation by learning memory selection from downstream feedback. For example, Memento (Zhou et al., 2025) and MemRL (Zhang et al., 2026) update retrieval policies or memory utilities based on final outcomes. However, these methods mainly retrieve at the episode level, selecting full trajectories based on the initial request and learning from sparse terminal rewards. This mismatches interactive text-to-SQL, where the agent’s state evolves through schema exploration, query execution, observations, and plan revision. Thus, retrieval should match the decision horizon. For example, a trajectory useful for initial planning may not help after a SQL error, where a local debugging pattern is more relevant. Moreover, a final success or failure signal



(a) Conceptual overview of **MERIT**.¹



(b) Performance comparison with baselines.

Figure 1: **Conceptual Overview and Performance Advantages of MERIT.** (a) **MERIT** performs horizon-aware retrieval by decoupling memory into episode-level (global strategy) and turn-level (local hints), using a PRM to provide the dense rewards required for turn-level policy learning. (b) **MERIT** outperforms static and single-horizon dynamic baselines in success rates while requiring fewer interaction turns.

provides weak supervision for determining which memory was useful at a specific intermediate turn.

Therefore, we propose a dynamic **Multi-horizon mEmory Retrieval** framework using **Intermediate** and **Terminal** rewards (**MERIT**). **MERIT** formulates memory retrieval as a learning problem across two decision horizons. As illustrated in Figure 1a, the episode-level retriever selects past trajectories before an interaction begins, providing global strategic guidance for planning. The turn-level retriever selects compact state-action-observation memories during the interaction, providing local hints conditioned on the current state. This dual-level design separates long-range task strategy from short-range execution support.

MERIT uses a two-stage retrieval pipeline at both levels. An embedding retriever first generates a candidate set, and a learned selector then chooses which memories to expose to the agent. Both selectors are optimized with reinforcement learning (RL), shaping retrieval by downstream task utility rather than fixed similarity alone. The episode-level selector is trained with terminal rewards since its retrieval decision affects the overall trajectory. The turn-level selector requires denser supervision because intermediate turns do not have reliable direct outcome labels (Choudhury, 2025; Xi et al., 2026). To address this credit-assignment challenge, **MERIT** introduces a lightweight Process Reward Model (PRM) that evaluates the utility of retrieved turn-level memories with a structured rubric. The resulting dense rewards enable RL for state-conditioned turn-level retrieval.

We evaluate **MERIT** on the interactive text-to-SQL benchmark (Huo et al., 2026) for in-domain

interactive performance and on Spider2-Snow (Lei et al., 2025) for cross-benchmark transfer. As summarized in Figure 1b, **MERIT** achieves higher success rates than no-memory, static-retrieval, and dynamic-retrieval baselines, while reducing the average number of interaction turns in the in-domain setting. On Spider2-Snow, **MERIT** transfers without benchmark-specific training or tuning and stays above all baselines, suggesting that the learned retrieval policies do not overfit the source benchmark.

Our contributions are summarized as follows:

- We propose **MERIT**, which aligns memory retrieval with decision horizon through episode-level global guidance and turn-level state-conditioned local support.
- We formulate memory retrieval as learned policies at both decision horizons, allowing memory selection to be optimized by downstream task utility rather than fixed similarity alone.
- We introduce a lightweight PRM that provides dense structured rewards for turn-level retrieval, mitigating the credit-assignment problem caused by sparse terminal outcomes.
- We evaluate **MERIT** in both in-domain and transfer settings, showing improved success rates and interaction efficiency in the in-domain setting, and cross-benchmark transfer without degradation.

2 Related Work

2.1 Text-to-SQL

Early neural text-to-SQL research typically formulates the task as single-pass semantic parsing from a natural-language question and database schema directly to SQL (Zhong et al., 2017; Xu et al., 2017; Yu et al., 2018; Lin et al., 2020). Recent LLM-

¹Figure generated by Gemini.

based approaches move beyond one-shot generation by introducing interaction, tool use, and iterative refinement (Zhang et al., 2024; Tian et al., 2023; Xiong et al., 2025; Wang et al., 2025; Zhu et al., 2026). These studies show that text-to-SQL agents benefit from intermediate feedback and adaptive decision-making. We build on this interactive setting and study experience reuse through learned long-term memory retrieval.

2.2 Memory Systems in LLM Agents

Long-term memory enables cross-episode experience reuse (Park et al., 2023; Lu et al., 2023), and prior work has explored long-term memory writing and organization, including operating-system-like memory management (Packer et al., 2024), semantic compression (Liu et al., 2026), graph-based memory representations (Chhikara et al., 2025), trajectory distillation (Fang et al., 2026), and structural memory representations (Zeng et al., 2024). These studies improve the storage and representation of agent experience. **MERIT** instead focuses on retrieving stored memories in the interactive setting. Most memory-augmented agents retrieve memories with fixed heuristics, such as dense embedding similarity, lexical matching, topic coverage, or predefined structural relations (Arslan et al., 2024; Kweon et al., 2025; Kagaya et al., 2024; Zeng et al., 2024). Even systems that learn other memory operations often keep retrieval fixed; for example, Memory-R1 (Yan et al., 2026) trains a memory manager and an answer agent with reinforcement learning but still retrieves through similarity-based RAG. These methods are effective for finding related candidates, yet their retrieval criteria are fixed and do not directly optimize downstream utility. Recent dynamic methods learn memory utility from feedback. MemRL (Zhang et al., 2026) updates memory Q-values through runtime learning, and Memento (Zhou et al., 2025) learns a memory-selection policy without fine-tuning base LLMs. They move beyond static similarity, but primarily operate at the episode level and rely on sparse terminal feedback. Learned retrieval beyond terminal reward has also been studied for test-time search, such as Asai et al. (2024), who train a model to adaptively decide when and what to retrieve at inference time, though their retrieval target is external documents at a single horizon. **MERIT** instead retrieves the agent’s own memory, separating episode-level retrieval for global strategy from turn-level retrieval for local support, and trains the

turn-level retriever with dense PRM-based rewards.

3 Preliminaries

We formulate interactive text-to-SQL as a sequential decision-making process. Given an initial natural-language request x , the agent interacts with a database environment and, when necessary, a user simulator. Unlike single-pass semantic parsing, the agent may take multiple actions before producing the final SQL query, including schema inspection, knowledge retrieval, SQL execution, clarification, and final submission.

At turn t , the agent has an interaction history

$$h_{t-1} = ((a_1, o_1), (a_2, o_2), \dots, (a_{t-1}, o_{t-1})),$$

where a_t is the agent action, and o_t is the corresponding observation. The agent observes a state

$$s_t = (x, h_{t-1}, b_t),$$

where b_t is the remaining interaction budget. Conditioned on s_t , the agent selects an action a_t , receives an observation o_t , and appends (a_t, o_t) to the history. An episode yields a trajectory

$$\tau = (x, h_T),$$

where T is the terminal turn. The episode terminates when the agent submits a final SQL query, exhausts its interaction budget, or reaches the maximum allowable number of turns. The final query is evaluated by execution correctness against the target SQL semantics.

4 Data Augmentation

Learning retrieval policies requires diverse executable SQL. Since high-quality interactive text-to-SQL data is limited, we construct additional training data through an execution-grounded augmentation pipeline. Rather than perturbing natural-language requests directly, we first synthesize executable SQL targets from seed examples and then generate user requests grounded in those targets. This SQL-first design reduces semantic drift and encourages executable SQL, faithful SQL-query alignment, and structural diversity.

Given a seed query, we apply constrained SQL-level perturbations, including compatible column substitutions, numeric literal perturbations, sorting-direction changes, and limit-value modifications. These operations diversify projections, grouping,




Figure 2: Overview of the proposed dual-level long-term memory framework for interactive Text-to-SQL. Episode-level memory is retrieved at episode start for global strategic guidance and updated at episode end, while turn-level memory is retrieved and updated per turn to provide state-conditioned local hints. Both levels employ a two-stage pipeline: embedding-based candidate generation followed by an RL-trained selector. A Process Reward Model (PRM), trained offline with SFT and GRPO, provides dense online rewards for turn-level retrieval, while episode-level retrieval is optimized with terminal task reward.

thresholds, rankings, and result sizes while preserving executability. For each augmented SQL target, an LLM generates a fully specified request, an intentionally underspecified request, and grounding metadata that links underspecified phrases to corresponding SQL fragments. We discard examples whose generated text or grounding metadata cannot be verified against the SQL logic. Each retained example therefore provides an executable text-to-SQL task with aligned SQL, requests, and grounding metadata for retrieval-policy training. Additional details, prompts, and a diversity and quality analysis of the augmented data are provided in Appendix A.

5 Methodology

We present **MERIT** by first introducing the dual-level memory architecture that separates global strategic guidance from local state-conditioned support (§ 5.1). We then formulate memory retrieval at both levels as learned policies optimized by reinforcement learning (§ 5.2). Finally, we describe the Process Reward Model (PRM), which provides dense supervision for turn-level retrieval (§ 5.3).

The complete architecture is shown in Figure 2.

5.1 Dual-Level Long-Term Memory

MERIT maintains two long-term memory banks. Episode-level memory stores complete interactions, and turn-level memory stores compact state-action-observation snippets for local decisions.

Episode-level memory. After each completed episode i , we write an episode-level memory

$$m_i^E = (k_i^E, v_i^E),$$

where k_i^E is the embedding of user request x_i and

$$v_i^E = (x_i, \tau_i, \ell_i, z_i, g_i)$$

stores the original request x_i , the full trajectory τ_i , the terminal outcome label ℓ_i , distilled episode insights z_i , and structured task-specific constraints or facts g_i resolved during interaction. The distilled insights summarize reusable strategic knowledge, such as successful plans, common failure modes, task characteristics, and guidance for similar future tasks. The prompt for insights extraction is in Appendix D.4.

For a new request x , we retrieve episode-level candidates by cosine similarity over the stored episode keys $\{k_i^E\}$. Conditioned on x and compact candidate summaries, the episode-level selector chooses a subset of memories to inject into the initial agent prompt. These memories provide global guidance before the interaction begins.

Turn-level memory. Turn-level memory is indexed by a compact state signature rather than the full interaction state. At turn t , we define the state signature

$$\sigma_t = g(s_t),$$

where $g(\cdot)$ maps the full state s_t into a textual retrieval key. The signature contains the original request, the current query context, a coarse description of the previous action, and the previous observation. We keep the previous action coarse in the retrieval key so that similar decision states can match even when their exact SQL queries differ.

After the agent completes turn j , we write a turn memory

$$m_j^T = (k_j^T, v_j^T),$$

where k_j^T is the embedding of σ_j , and

$$v_j^T = (\sigma_j, \tilde{a}_j, \tilde{o}_j, \ell_j^T)$$

stores the state signature, an abstracted action representation $\tilde{a}_j$, a short observation snippet $\tilde{o}_j$, and an execution-status label ℓ_j^T . The label records immediate action outcomes, such as invalid tool calls and SQL syntax errors, but does not capture whether it contributes to broader task progress. The abstracted action may include normalized SQL skeletons or other tool-call details, allowing reusable local patterns to transfer across queries.

Before intermediate turn t , we retrieve turn-level candidates by cosine similarity over $\{k_j^T\}$. Conditioned on σ_t and compact candidate summaries, the turn selector returns a subset of memories to inject as brief hints for the next agent action.

5.2 Retrieval Policy Learning

MERIT learns retrieval policies that decide which candidate memories to expose to the agent. At both levels, similarity-based retrieval first obtains candidates, and then a learned selector chooses a subset for prompt injection. We optimize both selectors with Proximal Policy Optimization (PPO), using reward signals aligned with their respective decision horizons.

Episode-level retrieval policy learning. Given an initial request x and an episode-level candidate set $C^E(x)$, the episode-level selector samples a retrieval decision

$$d^E \sim \pi_\theta^E(\cdot \mid x, C^E(x)),$$

where d^E specifies the selected memory subset $\hat{M}^E \subseteq C^E(x)$. The selected memories are injected before the interaction begins. After the agent completes trajectory τ , the selector receives a terminal reward $R^E(\tau)$, derived from final execution correctness. We optimize π_θ^E with PPO to learn episode-level retrieval from downstream outcomes.

Turn-level retrieval policy learning. Given a state signature σ_t and a turn-level candidate set $C^T(\sigma_t)$, the turn-level selector samples

$$d_t^T \sim \pi_\phi^T(\cdot \mid \sigma_t, C^T(\sigma_t)),$$

where d_t^T specifies the selected memory subset $\hat{M}_t^T \subseteq C^T(\sigma_t)$. The selected memories are injected as local hints before the next agent action. Turn-level retrieval cannot be trained directly from reliable immediate environment rewards. A memory selected at an intermediate turn may affect later actions, while its contribution may only be reflected in the final episode outcome. To mitigate this credit-assignment challenge, we use the PRM to estimate the utility of selected turn-level memories. The PRM scores each selected state-memory pair, and the turn-level reward is computed as

$$r_t^T = \frac{1}{|\hat{M}_t^T|} \sum_{m \in \hat{M}_t^T} f_{\text{PRM}}(\sigma_t, m),$$

where f_{PRM} maps a state signature and a memory to a scalar utility score. We optimize π_ϕ^T with PPO using these dense rewards.

5.3 Process Reward Model for Turn-Level Supervision

While turn-level retrieval requires dense supervision, terminal rewards provide delayed, noisy feedback on intermediate choices. **MERIT** therefore uses a Process Reward Model (PRM) to estimate candidate memory utility given the current state.

Structured Utility Prediction. Given a state-memory pair (σ_t, m) , the PRM predicts a structured rubric

$$y_t = f_\psi(\sigma_t, m).$$

The rubric evaluates memory usefulness along seven dimensions: state match, actionability value, pattern generalizability, outcome reliability, clarity for agent, confidence in assessment, and overall utility. Each dimension contains a score in $[0, 5]$ and a brief rationale. The rubric also includes binary decisions on whether to inject the memory and whether to use it as a warning, along with concise textual feedback on how to use it. The full rubric schema and prompt are provided in Appendix D.5.

For turn-level retrieval policy optimization, we use the trained PRM to score each state-memory pair. We denote the resulting scalar utility score as $f_{\text{PRM}}(\sigma_t, m)$, which is extracted from the predicted rubric using the overall utility score and the injection decision. This score serves as the dense reward for turn-level PPO, while the structured rubric grounds the reward in interpretable aspects of memory usefulness.

Two-Stage Training. We train the PRM from teacher-annotated state-memory pairs. For each pair, a stronger teacher model produces the full rubric. We first apply supervised fine-tuning (SFT) to imitate the teacher rubrics, then refine the PRM with GRPO using a structured matching reward.

Let $\hat{y}$ denote the PRM-generated rubric and y^* denote the teacher rubric. If $\hat{y}$ is not a valid JSON object with all required fields, we assign

$$R_{\text{PRM}}(\hat{y}, y^*) = -1.$$

Otherwise, we first compute a structured matching score S_{PRM} , and then clip it to obtain the PRM training reward:

$$S_{\text{PRM}} = 0.7R_{\text{dim}} + 0.2R_{\text{bin}} + 0.1R_{\text{text}} - 0.5\mathbb{I}_{\text{extra}},$$

$$R_{\text{PRM}}(\hat{y}, y^*) = \text{clip}(S_{\text{PRM}}, -1, 1),$$

where $\mathbb{I}_{\text{extra}} = 1$ if the output contains text outside the JSON object, and 0 otherwise.

The dimension-level reward compares the seven scored rubric dimensions. For each dimension d , let $\hat{s}_d, s_d^* \in [0, 5]$ be the predicted and teacher numeric scores, and let $\hat{r}_d, r_d^*$ be the corresponding short reasons. We define

$$R_{\text{dim}} = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \left[0.7 \cdot \max\left(0, 1 - \frac{|\hat{s}_d - s_d^*|}{5}\right) + 0.3 \cdot J(\hat{r}_d, r_d^*) \right],$$

where $\mathcal{D}$ is the set of the seven rubric dimensions and $J(\cdot, \cdot)$ denotes Jaccard similarity between token sets. The binary term R_{bin} is the average exact-match accuracy over the two binary decisions, and R_{text} is the average similarity over the two free-text feedback fields. This reward encourages the PRM to produce valid structured outputs that match teacher utility scores, decisions, and rationales.

6 Experiments

6.1 Experimental Settings

Datasets and Metrics. We evaluate **MERIT** on BIRD-Interact (Huo et al., 2026) for in-domain interactive text-to-SQL and on Spider2-Snow (Lei et al., 2025) for cross-benchmark transfer. For BIRD-Interact, we train the retrieval policies and the PRM on BIRD-Interact-Lite and evaluate on the unseen BIRD-Interact-Full dataset. Since BIRD-Interact-Lite contains only 236 base tasks (excluding 64 overlapped examples with BIRD-Interact-Full), we apply the data augmentation pipeline in § 4 and Appendix A to construct approximately 7,000 training examples, with benchmark-specific conversion details in Appendix B.

BIRD-Interact reports results in two benchmark-defined phases. Phase 1 evaluates the initial user request, while Phase 2 evaluates state-dependent follow-up requests. We measure success by execution correctness and report the average number of interaction turns as an efficiency metric. We also report READ and WRITE results separately, where READ tasks answer analytical queries without modifying the database, while WRITE tasks require state-changing operations. This split tests whether **MERIT** helps different interaction types that may require different strategies. For Spider2-Snow, all retrieval artifacts and learned policies are trained only on BIRD-Interact and transferred without Spider2-Snow-specific training or tuning.

Baselines. We compare **MERIT** with no-memory, static-retrieval, and dynamic-retrieval baselines. The no-memory baselines include Vanilla GPT-5, a non-agentic LLM baseline, No Memory Agent, the base interactive agent without historical trajectories, and Demo Agent, which uses the officially released static demonstrations (available only for BIRD-Interact). Static-retrieval baselines include BM25, Cosine Similarity with text-embedding-3-small, and TopicK (Kweon et al., 2025). Dynamic-retrieval baselines include Memento (Zhou et al., 2025) and MemRL (Zhang

Methods	Phase 1			Phase 2			Avg. #Turns ↓		
	WRITE	READ	Overall	WRITE	READ	Overall	WRITE	READ	Overall
<i>No Memory</i>									
Vanilla GPT-5	24.71 _{-18.97}	8.69 _{-6.33}	13.33 _{-10.00}	12.07 _{-8.04}	3.05 _{-7.28}	5.67 _{-7.50}	-	-	-
No Memory	43.68	15.02	23.33	20.11	10.33	13.17	8.24	9.21	8.93
Demo	54.60 _{+10.92}	18.31 _{+3.29}	28.83 _{+5.50}	28.16 _{+8.05}	11.97 _{+1.64}	16.67 _{+3.50}	8.44 _{+0.20}	9.93 _{+0.72}	9.50 _{+0.57}
<i>Static Memory</i>									
BM25	52.87 _{+9.19}	19.48 _{+4.46}	29.17 _{+5.84}	28.74 _{+8.63}	13.85 _{+3.52}	18.17 _{+5.00}	7.65 _{-0.59}	8.53 _{-0.68}	8.27 _{-0.66}
Cosine Similarity	51.72 _{+8.04}	19.48 _{+4.46}	28.83 _{+5.50}	28.74 _{+8.63}	13.62 _{+3.29}	18.00 _{+4.83}	7.63 _{-0.61}	8.79 _{-0.42}	8.45 _{-0.48}
TopicK	48.85 _{+5.17}	18.31 _{+3.29}	27.17 _{+3.84}	25.86 _{+5.75}	11.97 _{+1.64}	16.00 _{+2.83}	7.46 _{-0.78}	8.64 _{-0.57}	8.30 _{-0.63}
<i>Dynamic Memory</i>									
Memento	50.57 _{+6.89}	19.72 _{+4.70}	28.67 _{+5.34}	28.74 _{+8.63}	13.38 _{+3.05}	17.83 _{+4.66}	7.80 _{-0.44}	8.91 _{-0.30}	8.59 _{-0.34}
MemRL	51.15 _{+7.47}	19.95 _{+4.93}	29.00 _{+5.67}	30.46 _{+10.35}	14.08 _{+3.75}	18.83 _{+5.66}	7.58 _{-0.66}	8.79 _{-0.42}	8.44 _{-0.49}
MERIT (ours)	53.49 _{+9.81}	24.64 _{+9.62}	33.01 _{+9.68}	34.30 _{+14.19}	17.22 _{+6.89}	22.17 _{+9.00}	7.09 _{-1.15}	8.15 _{-1.06}	7.84 _{-1.09}

Table 1: Performance comparison of **MERIT** against baselines on Success Rates across two phases and Average Number of Turns. Performance is partitioned into queries with WRITE operations and READ-only operations. The backbone of the main agent is GPT-5 and the user simulator is GPT-4o for all methods. Subscripts denote absolute differences relative to the No Memory baseline.

Methods	Success Rate	Avg. #Turns
<i>No Memory</i>		
Vanilla GPT-5	37.29 _{-28.71}	-
No Memory	66.00	13.29
<i>Static Memory</i>		
BM25	67.09 _{+1.09}	14.58 _{+1.29}
Cosine Similarity	68.19 _{+2.19}	13.48 _{+0.19}
TopicK	66.73 _{+0.73}	14.54 _{+1.25}
<i>Dynamic Memory</i>		
Memento	68.56 _{+2.56}	14.33 _{+1.04}
MemRL	66.91 _{+0.91}	14.95 _{+1.66}
MERIT (ours)	69.29 _{+3.29}	13.44 _{+0.15}

Table 2: Transferability comparison on Spider2-Snow, with the main-agent backbone as GPT-5. TopicK, Memento, MemRL, and **MERIT** directly apply the trained checkpoints or artifacts. Subscripts denote absolute differences relative to the No Memory baseline.

et al., 2026), which learn memory selection or utility but operate primarily at the episode level. For applicable memory-based methods, we keep the episode-level memory writing mechanism identical to isolate the effect of retrieval.

Model Configurations. The base reasoning agent uses GPT-5 (Singh et al., 2026) with low reasoning effort, and the user simulator uses GPT-4o (Hurst et al., 2024). The PRM is initialized from *DeepSeek-R1-Distill-Qwen-7B* (Guo et al., 2025), and the retrieval selectors use LoRA-tuned *Qwen3-4B-Instruct-2507* (Yang et al., 2025). We

use 5,346 teacher-annotated state-memory pairs for PRM SFT, 42,768 pairs for GRPO, and 5,346 pairs for PRM evaluation.

More details are provided in Appendix B.

6.2 Experimental Results

Main Results. Table 1 compares **MERIT** with no-memory, static-retrieval, and dynamic-retrieval baselines on BIRD-Interact. **MERIT** achieves the best overall success rates in both Phase 1 and Phase 2, reaching 33.01% and 22.17%, respectively. The gains hold across both READ and WRITE tasks, indicating that the learned retrieval policies help both analytical extraction and state-changing operations. Furthermore, **MERIT** achieves the lowest average number of interaction turns. These results indicate that **MERIT** improves task success while reducing interaction turns. To verify that this ranking is stable, we repeat **MERIT** and the strongest dynamic baseline five times in Appendix C.2, where **MERIT**’s weakest run still exceeds MemRL’s best run in both phases.

Transferability. Table 2 evaluates whether retrieval policies trained on BIRD-Interact transfer to Spider2-Snow without Spider2-Snow-specific training or tuning. **MERIT** achieves the highest success rate among all methods, outperforming the strongest baseline. The gain is smaller than on BIRD-Interact, as expected under cross-benchmark transfer. Nevertheless, this suggests that **MERIT**’s learned policies transfer without degradation to a new benchmark. **MERIT** also keeps average turns close to the No Memory Agent and below all




Figure 3: Retrieval robustness and retrieval size. MERIT consistently outperforms the non-RL variant across k , while the retrieval-size analysis shows that injecting too many memories can introduce noisy or redundant guidance.

Methods	Phase 1			Phase 2			Avg. #Turns↓		
	WRITE	READ	Overall	WRITE	READ	Overall	WRITE	READ	Overall
MERIT (ours)	53.49	24.64	33.01	34.30	17.22	22.17	<u>7.09</u>	<u>8.15</u>	<u>7.84</u>
w/o episode	42.53	15.73	23.50	21.84	8.92	12.70	8.10	8.88	8.65
w/o turn	48.85	20.19	28.50	28.74	12.44	17.20	7.67	8.68	8.39
w/o RL	<u>52.87</u>	<u>24.18</u>	<u>32.50</u>	<u>29.89</u>	<u>15.26</u>	<u>19.50</u>	7.07	8.07	7.78

Table 3: Ablation studies evaluating the contribution of individual framework components, including episode-level strategy, turn-level hints, and RL optimization, on task success rates and interaction efficiency.

Condition	% Next Action Valid	% Final Success
MERIT (ours)	60.0	20.0
BM25	62.5	10.0
No Memory	55.0	7.5

Table 4: Counterfactual analysis at a frozen intermediate state ($k = 7$) over 40 episodes. BM25 attains a higher next-action-valid rate, yet the **MERIT** selector yields a higher final success rate, showing that locally plausible memories are not necessarily globally correct.

memory-retrieval baselines, indicating improved success without substantial interaction cost.

Ablation Studies. Table 3 ablates each **MERIT** component. Removing episode-level memory causes the largest drop, reducing overall success to 23.50% in Phase 1 and 12.70% in Phase 2. This shows that episode-level experience is important in interactive text-to-SQL. Removing turn-level memory also degrades performance, reducing overall success to 28.50% and 17.20% in the two phases and increasing the average number of turns to 8.39. This suggests that turn-level hints complement episode-level memories by supporting decisions as the interaction state evolves. Finally, removing RL optimization has a modest effect on Phase 1 but lowers Phase 2 success from 22.17% to 19.50%. This suggests that utility-based selector optimization is most beneficial for state-dependent follow-up interactions, where memory usefulness depends more strongly on the evolving context.

Training Stage	Valid (%)	avg. Acc	avg. MAE	avg. MSE
Vanilla	95.59	33.67	0.94	1.73
SFT	92.12	36.84	0.85	1.50
SFT + GRPO	100.00	73.59	0.29	0.34

Table 5: Performance of PRM.

Counterfactual Analysis. To isolate the effect of a single turn-level memory, we run a counterfactual at a fixed state. We sample 40 episodes, run each to turn $k = 7$, freeze the state, and fork three continuations that differ only in the injected turn-level memory, from the **MERIT** selector, a BM25 baseline, and no memory. We report next-action validity and final success in Table 4. **MERIT** attains a slightly lower next-action-valid rate than BM25 but a higher final success rate, and both exceed no memory. BM25 favors surface-similar memories that make the immediate action executable but not the final query correct, whereas the learned selector favors downstream usefulness and converts more episodes into correct submissions. This suggests the **MERIT** turn-level selector chooses memories with positive utility toward the outcome rather than mere local relevance.

PRM Evaluation. Table 5 evaluates whether the PRM can produce structured utility judgments aligned with teacher annotations. Compared with the vanilla base model and the SFT-only model, SFT+GRPO achieves 100% valid for-

Methods	Avg. #Turns	#Tokens
MERIT (ours)	7.84	100.4M
w/o episode	8.65	95.4M
w/o turn	8.39	108.5M
w/o RL	7.78	104.2M

Table 6: Token and interaction efficiency comparison.

matting, higher annotation agreement, and lower numeric-score errors. Specifically, SFT+GRPO improves average accuracy to 73.59% and reduces MAE and MSE to 0.29 and 0.34, respectively. These results show that GRPO improves agreement with teacher utility rubrics; downstream results test whether such rewards help retrieval. Since direct intermediate utility labels are unavailable, we use this teacher-aligned score as a structured proxy reward and examine its downstream effects through turn-level and RL ablations.

Retrieval Behavior Analysis. We further analyze the robustness of learned retrieval and the effect of retrieval size. Figure 3a compares **MERIT** with its non-RL variant under different Pass@ k budgets. **MERIT** achieves higher success across all k , indicating that RL-based retrieval remains beneficial across different generation budgets. Figure 3b studies the number of retrieved memories. Retrieving one to three memories yields comparable performance, whereas larger retrieval budgets reduce success rates. This suggests that excessive memory context may introduce redundant or distracting information. We provide an additional analysis of memory composition in Appendix C.1.

Computational and Token Efficiency. Although **MERIT** introduces learned retrieval modules, the additional cost is limited because the selectors are lightweight and operate on compact state signatures and memory summaries. Table 6 compares token usage and interaction turns across ablations. **MERIT** uses 100.4M tokens, fewer than the variants without turn-level memory or RL optimization, which require 108.5M and 104.2M tokens, respectively. The variant without episode-level memory uses fewer tokens, yet it substantially reduces success rates and increases interaction turns. Overall, **MERIT** offers a favorable success-efficiency trade-off by improving task performance while reducing unnecessary interactions and tokens.

Leaderboard Results. We additionally transfer the retrieval artifacts trained with GPT-5 as the rea-

soning agent, GPT-4o as the user simulator, and the same Qwen-based retrievers used in the main experiments to two stronger agent configurations. Using GPT-5.4 as the reasoning agent with GPT-4o as the user simulator, **MERIT** ranks first on both Phase 1 and Phase 2 of the BIRD-Interact leaderboard² as of August 2026 (see Appendix C.3). A second configuration with Claude-Opus-4.6 as the reasoning agent and Claude-Haiku-4-5 as the user simulator ranks second on Phase 2, where adding **MERIT**’s retrieval raises success from 17.50 to 20.83 over the base Claude-Opus-4.6 agent.

7 Conclusion

We proposed **MERIT**, a dual-level long-term memory framework that learns episode-level and turn-level retrieval policies for interactive text-to-SQL agents. By separating global strategic guidance from local state-conditioned hints and using a lightweight PRM to provide dense rewards for turn-level retrieval, **MERIT** enables horizon-aware memory selection. Experiments on BIRD-Interact show improved success rates with fewer interaction turns, while Spider2-Snow results provide evidence of transfer beyond the source benchmark.

Limitations

First, the PRM is trained from teacher-generated rubrics rather than ground-truth labels of turn-level memory utility. This is a practical limitation of interactive retrieval settings, where the usefulness of a memory at an intermediate turn is difficult to observe directly because the agent state, future actions, and environment feedback all evolve after the retrieval decision. We therefore use a strong teacher model to provide structured utility annotations as a proxy supervision signal. While this does not provide causal ground-truth attribution for each retrieved memory, the downstream results and ablation studies suggest that the PRM-trained turn-level selector improves state-conditioned retrieval and contributes to better follow-up interaction performance.

Second, our current evaluation uses a strong proprietary backbone for the main agent and teacher annotation. This setting allows us to study memory retrieval under a capable interactive agent, but it leaves open whether the same gains hold across open-source agent backbones or weaker reasoning models. Evaluating **MERIT** with different base

²<https://bird-interact.github.io/>

agents, teacher models, and simulator configurations is an important direction for future work.

References

- Muhammad Arslan, Hussam Ghanem, Saba Munawar, and Christophe Cruz. 2024. [A survey on rag with llms](#). *Procedia Computer Science*, 246:3781–3790. 28th International Conference on Knowledge Based and Intelligent Information and Engineering Systems (KES 2024).
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avi Sil, and Hannaneh Hajishirzi. 2024. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). In *International Conference on Learning Representations*, volume 2024, pages 9112–9141.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. [Mem0: Building production-ready ai agents with scalable long-term memory](#). In *ECAI 2025*. IOS Press.
- Sanjiban Choudhury. 2025. [Process reward models for llm agents: Practical framework and directions](#). *Preprint*, arXiv:2502.10325.
- Runnan Fang, Yuan Liang, Xiaobin Wang, Jialong Wu, Shuofei Qiao, Pengjun Xie, Fei Huang, Huajun Chen, and Ningyu Zhang. 2026. [Memp: Exploring agent procedural memory](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 17490–17502, San Diego, California, United States. Association for Computational Linguistics.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shiron Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. [Deepseek-r1 incentivizes reasoning in llms through reinforcement learning](#). *Nature*, 645(8081):633–638.
- Nan Huo, Xiaohan Xu, Jinyang Li, Per Jacobsson, Shipai Lin, Bowen Qin, Binyuan Hui, Xiaolong Li, Ge Qu, Shuzheng Si, Linheng Han, Edward Alexander, Xintong Zhu, Rui Qin, Ruihan Yu, Yiyao Jin, Feige Zhou, Weihao Zhong, Yun Chen, and 5 others. 2026. [BIRD-INTERACT: Re-imagining text-to-SQL evaluation via lens of dynamic interactions](#). In *The Fourteenth International Conference on Learning Representations*.
- Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, and 399 others. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Tomoyuki Kagaya, Thong Jing Yuan, Yuxuan Lou, Jayashree Karlekar, Sugiri Pranata, Akira Kinose, Koki Oguri, Felix Wick, and Yang You. 2024. [Rap: Retrieval-augmented planning with contextual memory for multimodal llm agents](#). *Preprint*, arXiv:2402.03610.
- Wonbin Kweon, SeongKu Kang, Runchu Tian, Pengcheng Jiang, Jiawei Han, and Hwanjo Yu. 2025. [Topic coverage-based demonstration retrieval for in-context learning](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 19900–19912, Suzhou, China. Association for Computational Linguistics.
- Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin SU, ZHAOQING SUO, Hongcheng Gao, Wenjing Hu, Pengcheng Yin, Victor Zhong, Caiming Xiong, Ruoxi Sun, Qian Liu, Sida Wang, and Tao Yu. 2025. [Spider 2.0: Evaluating language models on real-world enterprise text-to-SQL workflows](#). In *The Thirteenth International Conference on Learning Representations*.
- Xinzhe Li. 2025. [A review of prominent paradigms for LLM-based agents: Tool use, planning \(including RAG\), and feedback learning](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 9760–9779, Abu Dhabi, UAE. Association for Computational Linguistics.
- Xi Victoria Lin, Richard Socher, and Caiming Xiong. 2020. [Bridging textual and tabular data for cross-domain text-to-SQL semantic parsing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4870–4888, Online. Association for Computational Linguistics.
- Jiaqi Liu, Yaofeng Su, Peng Xia, Siwei Han, Zeyu Zheng, Cihang Xie, Mingyu Ding, and Huaxiu Yao. 2026. [Simplemem: Efficient lifelong memory for LLM agents](#). In *Forty-third International Conference on Machine Learning*.
- Junru Lu, Siyu An, Mingbao Lin, Gabriele Pergola, Yulan He, Di Yin, Xing Sun, and Yunsheng Wu. 2023. [Memochat: Tuning llms to use memos for consistent long-range open-domain conversation](#). *Preprint*, arXiv:2308.08239.
- Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. [Evaluating very long-term conversational memory of LLM agents](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13851–13870, Bangkok, Thailand. Association for Computational Linguistics.
- Mahmoud Mohammadi, Yipeng Li, Jane Lo, and Wendy Yip. 2025. [Evaluation and benchmarking of llm agents: A survey](#). In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2, KDD '25*, page 6129–6139, New York, NY, USA. Association for Computing Machinery.
- Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez.

2024. **Memgpt: Towards llms as operating systems.** *Preprint*, arXiv:2310.08560.
- Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. **Generative agents: Interactive simulacra of human behavior.** In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, New York, NY, USA. Association for Computing Machinery.
- Rana Salama, Jason Cai, Michelle Yuan, Anna Currey, Monica Sunkara, Yi Zhang, and Yassine Benajiba. 2025. **MemInsight: Autonomous memory augmentation for LLM agents.** In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 33136–33152, Suzhou, China. Association for Computational Linguistics.
- Zhihui Shao, Shubin Cai, Rongsheng Lin, and Zhong Ming. 2025. **Enhancing text-to-SQL with question classification and multi-agent collaboration.** In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 4340–4349, Albuquerque, New Mexico. Association for Computational Linguistics.
- Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, Akshay Nathan, Alan Luo, Alec Helyar, Aleksander Madry, Aleksandr Efremov, Aleksandra Spyra, Alex Baker-Whitcomb, Alex Beutel, Alex Karpenko, and 467 others. 2026. **Openai gpt-5 system card.** *Preprint*, arXiv:2601.03267.
- Haoran Sun, Shaoning Zeng, and Bob Zhang. 2026. **H-MEM: Hierarchical memory for high-efficiency long-term reasoning in LLM agents.** In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 341–350, Rabat, Morocco. Association for Computational Linguistics.
- Yuan Tian, Zheng Zhang, Zheng Ning, Toby Jia-Jun Li, Jonathan K. Kummerfeld, and Tianyi Zhang. 2023. **Interactive text-to-SQL generation via editable step-by-step explanations.** In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16149–16166, Singapore. Association for Computational Linguistics.
- Bing Wang, Changyu Ren, Jian Yang, Xinnian Liang, Jiaqi Bai, LinZheng Chai, Zhao Yan, Qian-Wen Zhang, Di Yin, Xing Sun, and Zhoujun Li. 2025. **MAC-SQL: A multi-agent collaborative framework for text-to-SQL.** In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 540–557, Abu Dhabi, UAE. Association for Computational Linguistics.
- Zhiheng Xi, Chenyang Liao, Guanyu Li, Zhihao Zhang, Wenxiang Chen, Binghai Wang, Senjie Jin, Yuhao Zhou, Jian Guan, Wei Wu, Tao Ji, Tao Gui, Qi Zhang, and Xuanjing Huang. 2026. **Agentprm: Process reward models for llm agents via step-wise promise and progress.** In *Proceedings of the ACM Web Conference 2026*, WWW '26, page 4184–4195, New York, NY, USA. Association for Computing Machinery.
- Guanming Xiong, Junwei Bao, Hongfei Jiang, Yang Song, and Wen Zhao. 2025. **Multi-turn interactions for text-to-sql with large language models.** In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, CIKM '25, page 3560–3570, New York, NY, USA. Association for Computing Machinery.
- Weikai Xu, Chengrui Huang, Shen Gao, and Shuo Shang. 2025. **Llm-based agents for tool learning: A survey.** *Data Science and Engineering*, 10(4):533–563.
- Xiaojun Xu, Chang Liu, and Dawn Song. 2017. **Sql-net: Generating structured queries from natural language without reinforcement learning.** *Preprint*, arXiv:1711.04436.
- Sikuan Yan, Xiufeng Yang, Zuchao Huang, Ercong Nie, Zifeng Ding, Zonggen Li, Xiaowen Ma, Jinhe Bi, Kristian Kersting, Jeff Z. Pan, Hinrich Schuetze, Volker Tresp, and Yunpu Ma. 2026. **Memory-r1: Enhancing large language model agents to manage and utilize memories via reinforcement learning.** In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12805–12825, San Diego, California, United States. Association for Computational Linguistics.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. **Qwen3 technical report.** *Preprint*, arXiv:2505.09388.
- Tao Yu, Zifan Li, Zilin Zhang, Rui Zhang, and Dragomir Radev. 2018. **TypeSQL: Knowledge-based type-aware neural text-to-SQL generation.** In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 588–594, New Orleans, Louisiana. Association for Computational Linguistics.
- Siyu Yuan, Kaitao Song, Jiangjie Chen, Xu Tan, Yongliang Shen, Kan Ren, Dongsheng Li, and Deqing Yang. 2025. **EASYTOOL: Enhancing LLM-based agents with concise tool instruction.** In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 951–972, Albuquerque, New Mexico. Association for Computational Linguistics.
- Ruihong Zeng, Jinyuan Fang, Siwei Liu, and Zaiqiao Meng. 2024. **On the structural memory of llm agents.** *Preprint*, arXiv:2412.15266.

- Hanchong Zhang, Ruisheng Cao, Hongshen Xu, Lu Chen, and Kai Yu. 2024. [CoE-SQL: In-context learning for multi-turn text-to-SQL with chain-of-editions](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6487–6508, Mexico City, Mexico. Association for Computational Linguistics.
- Shengtao Zhang, Jiaqian Wang, Ruiwen Zhou, Junwei Liao, Yuchen Feng, Zhuo Li, Yujie Zheng, Weinan Zhang, Ying Wen, Zhiyu Li, Feiyu Xiong, Yutao Qi, Bo Tang, and Muning Wen. 2026. [Memrl: Self-evolving agents via runtime reinforcement learning on episodic memory](#). *Preprint*, arXiv:2601.03192.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. [Seq2sql: Generating structured queries from natural language using reinforcement learning](#). *Preprint*, arXiv:1709.00103.
- Huichi Zhou, Yihang Chen, Siyuan Guo, Xue Yan, Kin Hei Lee, Zihan Wang, Ka Yiu Lee, Guchun Zhang, Kun Shao, Linyi Yang, and Jun Wang. 2025. [Memento: Fine-tuning llm agents without fine-tuning llms](#). *Preprint*, arXiv:2508.16153.
- Jiongli Zhu, Haoquan Guan, Parjanya Prajakta Prashant, Nikki Lijing Kuang, Seyedeh Baharan Khatami, Canwen Xu, Xiaodong Yu, Yingyu Lin, Zhewei Yao, Yuxiong He, and Babak Salimi. 2026. [Residual skill optimization for text-to-sql ensembles](#). *Preprint*, arXiv:2605.21792.

A Data Augmentation Details

Learning robust memory retrieval policies requires diverse executable tasks from which meaningful interaction trajectories can be collected. Since high-quality interactive text-to-SQL data is expensive to collect, we augment the training set with an execution-grounded pipeline. The pipeline first synthesizes executable SQL targets from seed examples and then generates natural-language requests grounded in those targets. This SQL-first design reduces semantic drift and encourages three properties: executable SQL, faithful SQL-query alignment, and structural diversity.

SQL-level perturbation. Given a seed SQL query, we parse its abstract syntax tree to recover alias-to-table mappings and identify qualified columns within key clauses such as SELECT and GROUP BY. We then apply constrained SQL-level perturbations while retaining executable variants. To preserve structural integrity and executability, column substitutions are restricted to non-key columns with compatible coarse data types from the same table. We also perturb numeric literals, reverse sorting directions, and modify limit values. These operations diversify projections, grouping choices, thresholds, rankings, and result sizes while maintaining executable SQL targets.

Natural-language generation. After obtaining an executable SQL target, we generate corresponding natural-language requests with an LLM. The generation is conditioned on a lightweight SQL outline with clause-level cues, so that the request describes the intended operation without exposing raw table names, column names, aliases, or SQL identifiers. The model produces a fully specified request and an intentionally underspecified variant. The fully specified request captures the target SQL semantics, while the underspecified request hides selected details that can be resolved through interaction.

Grounding metadata. For each underspecified request, the generation also produces grounding metadata that links vague phrases to their corresponding SQL fragments and ambiguity types. This metadata supports interactive environments by specifying which hidden details may be clarified during interaction. The ambiguity types include knowledge-linking, semantic, entity, temporal, aggregation, sort, intent, and join-path ambiguity.

Verification and filtering. We discard examples whose generated text cannot be verified against the SQL logic. Retained examples must preserve SQL-query consistency, avoid raw identifier leakage in user-facing text, and provide grounding metadata that can be linked to the corresponding request phrase and SQL fragment. Each retained example therefore contains an executable SQL target, aligned natural-language requests, and grounding metadata for interactive trajectory collection.

Prompt. The full prompt used for natural-language generation and ambiguity annotation is shown in Appendix D.2.

Diversity and Quality of Augmented Data. We analyze the 7,000 augmented instances against the seed. For quality, all generated SQL executes successfully against its target PostgreSQL database, every one of the 479 seed tables remains referenced with 150 new column references and none outside the schema, and operator complexity rises over the seed (JOIN usage from 67% to 82%, window functions from 16% to 22%). For diversity, the set spans 260 distinct AST skeletons, its natural-language queries reach a Self-BLEU-3 of 0.13 (lower is more diverse) with a $1.8\times$ larger vocabulary, and their mean pairwise cosine distance is 0.67, comparable to the seed’s 0.71, with 93% falling below 0.90 similarity to every seed query. These show that the pipeline expands the seed distribution while preserving schema grounding and executability.

B More Experimental Settings

Datasets and evaluation metrics. We evaluate MERIT on BIRD-Interact (Huo et al., 2026) and Spider2-Snow (Lei et al., 2025). We use BIRD-Interact-Lite for training and BIRD-Interact-Full for in-domain evaluation. BIRD-Interact agent provides an interactive text-to-SQL environment that couples an executable PostgreSQL sandbox with a dynamic user simulator. To train the retrieval policies and the PRM, we apply the data augmentation pipeline described in § 4 and Appendix A, expanding the initial base tasks to approximately 7,000 examples.

For compatibility with the BIRD-Interact protocol, each augmented SQL target is converted into the benchmark-required textual fields, including a fully specified request, an underspecified initial request, and grounding metadata that links underspecified phrases to corresponding SQL fragments.

Methods	Phase 1			Phase 2			Avg. #Turns↓		
	WRITE	READ	Overall	WRITE	READ	Overall	WRITE	READ	Overall
MERIT (POSITIVE ONLY)	53.49	24.64	33.01	34.30	17.22	22.17	7.09	8.15	7.84
MERIT (ALL)	51.15	20.19	29.20	29.31	14.32	18.70	7.01	8.09	7.78

Table 7: Performance comparison between using positive-only experiences and using all experiences.

These fields are used by the benchmark user simulator to provide consistent clarification responses during interaction.

For in-domain evaluation, we test on the unseen BIRD-Interact-Full dataset, which contains 600 complex tasks. BIRD-Interact reports results in two benchmark-defined phases. Phase 1 evaluates the initial user request, while Phase 2 evaluates state-dependent follow-up requests whose semantics may depend on the preceding task context, such as artifacts or database states established earlier. These phases are part of the benchmark protocol and are distinct from the episode-level and turn-level retrieval horizons in **MERIT**. We measure task success by execution correctness against the benchmark ground-truth test cases. Following the benchmark protocol, we also report results separately for READ tasks, which focus on analytical extraction, and WRITE tasks, which require database modification or other state-changing operations. We report the average number of interaction turns as an interaction-efficiency metric.

To evaluate cross-benchmark transfer, we further test on Spider2-Snow. In this setting, all retrieval artifacts and learned policies are trained only on BIRD-Interact and directly transferred to Spider2-Snow without Spider2-Snow-specific training or tuning. This setting evaluates whether the learned retrieval policies capture reusable memory-selection behavior beyond the source benchmark. For Spider2-Snow, we report success rate and average number of interaction turns.

Baselines. We compare **MERIT** with three classes of baselines. To ensure a controlled comparison, all interactive agent baselines use the same base agent prompt, user simulator, interaction budgets, evaluation protocol, and budget. For memory-based methods, each method maintains its own on-line memory bank generated from its own rollouts. Candidate retrieval is performed from the method-specific memory bank with the same candidate size and memory budget as **MERIT** whenever applicable; differences therefore come from the retrieval or memory-selection strategy rather than from the

agent prompt or evaluation setup.

- **No Memory.** We include three baselines without experience retrieval. **Vanilla GPT-5** is a standard non-agentic LLM baseline. **No Memory Agent** is the base interactive agent without access to historical trajectories. **Demo Agent** is the interactive agent using the standard static demonstrations released with BIRD-Interact.
- **Static Retrieval.** We include fixed retrieval methods, including **BM25**, **Cosine Similarity** with `text-embedding-3-small`, and **TopicK** (Kweon et al., 2025). These methods retrieve memories using predefined retrieval criteria rather than learned retrieval policies.
- **Dynamic Retrieval.** We compare with dynamic memory methods that learn memory utility or selection policies, including **Memento** (Zhou et al., 2025) and **MemRL** (Zhang et al., 2026). These methods move beyond static similarity, but operate primarily at the episode level.

Model configurations. The base reasoning agent uses GPT-5 (Singh et al., 2026) with low reasoning effort, using the base text-to-SQL prompt in Appendix D.1. The user simulator, which provides clarification responses and execution feedback, is powered by GPT-4o (Hurst et al., 2024). Each interaction allows at most 100 turns, with a user patience budget of 6, an environment-interaction budget of 3, and a submission budget of 3.

To construct the PRM training data, we use GPT-5 to annotate state-memory pairs with structured utility rubrics. This yields 5,346 pairs for supervised fine-tuning (SFT), 42,768 pairs for Group Relative Policy Optimization (GRPO), and 5,346 pairs for evaluation. The PRM is initialized from *DeepSeek-R1-Distill-Qwen-7B* (Guo et al., 2025); during retrieval-policy training, turn-level rewards are computed with the trained PRM checkpoint. For the dual-level retrieval selectors, we use LoRA-tuned *Qwen3-4B-Instruct-2507* (Yang et al., 2025).

Methods	Phase 1			Phase 2		
	WRITE	READ	Overall	WRITE	READ	Overall
MemRL	51.34 (0.33)	20.58 (0.59)	29.50 (0.44)	29.12 (1.44)	14.71 (0.59)	18.90 (0.10)
MERIT	55.03 (1.65)	24.09 (0.81)	33.06 (1.05)	33.75 (1.88)	17.79 (0.75)	22.42 (1.04)

Table 8: Five-run mean and standard deviation of success rates on BIRD-Interact-Full for **MERIT** and the strongest dynamic baseline (MemRL). **MERIT**’s weakest run still exceeds MemRL’s best run in both phases.




Figure 4: Snapshots of the BIRD-Interact leaderboard (<https://bird-interact.github.io/>), accessed August 2026. **MERIT** with GPT-5.4 ranks first on both phases, and **MERIT** with Claude-Opus-4.6 ranks second on Phase 2.

At both memory levels, the embedding retriever returns 10 candidates, and the learned selector injects at most one memory into the agent context. We use positive-only memories by default and exclude memories from the same instance during evaluation to avoid self-retrieval leakage.

Both episode- and turn-level selectors are optimized with PPO for 4 epochs, using batch size 128, mini-batch size 16, learning rate 1×10^{-7} , PPO epoch count 2, clipping range 0.2, value-function coefficient 0.2, entropy coefficient 0.01, initial KL coefficient 0.02, and one rollout per example. Training uses 4 processes with Ray rollout parallelism 24, and all experiments are accelerated using 4 NVIDIA B200 GPUs.

C Additional Analysis

C.1 Effect of Memory Composition

MERIT uses positive memories by default in the main experiments. Although unsuccessful trajectories can contain useful warning signals, naively mixing successful and unsuccessful experiences may introduce noisy, contradictory, or overly conservative guidance into the agent context. We therefore compare two memory-bank construction strate-

gies. The first strategy, **MERIT** (POSITIVE ONLY), retrieves only memories from successful trajectories. The second strategy, **MERIT** (ALL), retrieves memories from both successful and unsuccessful trajectories.

As shown in Table 7, positive-only retrieval achieves higher success rates in both phases. Compared with retrieval from all memories, positive-only retrieval improves overall success from 29.20% to 33.01% in Phase 1 and from 18.70% to 22.17% in Phase 2. This result suggests that successful trajectories provide cleaner reusable guidance for the agent. In contrast, directly injecting unsuccessful trajectories may introduce noise, even though such trajectories may still contain useful warning signals. Retrieval from all memories slightly reduces the average number of turns, but this efficiency gain is accompanied by a clear decrease in task success. Based on this result, we use positive memories in the main experiments.

C.2 Variance Analysis

Since our evaluation involves stochastic components, we assess the stability of the in-domain ranking across repeated runs. We run both **MERIT** and the strongest dynamic baseline, MemRL, five

times on BIRD-Interact-Full and report the mean and standard deviation of success rates in Table 8. Both methods are stable across runs, and the five-run means closely match the single-run numbers reported in Table 1 (**MERIT** 33.06 and 22.42 versus 33.01 and 22.17; MemRL 29.50 and 18.90 versus 29.00 and 18.83). **MERIT**’s overall margin over MemRL, about +3.6 in Phase 1 and +3.5 in Phase 2, is several times larger than the run-to-run spread of either method, and even **MERIT**’s weakest run stays above MemRL’s best run in both phases. The in-domain gain over the strongest learned baseline is therefore not an artifact of a single run.

C.3 Leaderboard Snapshot

Figure 4 shows snapshots of the public BIRD-Interact leaderboard taken in August 2026 for the two transfer configurations. We reuse the retrieval artifacts trained with GPT-5 without any further tuning and pair them with stronger reasoning agents. With GPT-5.4 as the reasoning agent and GPT-4o as the user simulator, **MERIT** ranks first on both Phase 1 and Phase 2. With Claude-Opus-4.6 as the reasoning agent and Claude-Haiku-4-5 as the user simulator, **MERIT** ranks second on Phase 2.

D Prompts

D.1 Base Prompts

For the base agent prompt, we adapt elements from the prompt released in BIRD-INTERACT (Huo et al., 2026), shown below.

BASE AGENT SYSTEM PROMPT

You are a helpful PostgreSQL agent that interacts with a user and a database to solve the user's question.

Task Description

Your goal is to understand the user's ambiguous question involving the external knowledge retrieval and generate the correct SQL query to solve it. You can:

1. Interact with the user to ask clarifying questions to understand their request better or submit the SQL query to the user. The user will test your SQL correctness and give you feedback.
 2. Interact with the {setting} environment (postgresql db, column meaning file, external knowledge, and so on) to explore the database and get db relevant information.
- Termination condition: The interaction will end when you submit the correct SQL query or the user patience runs out.
 - Cost of your action: each your action will cost a certain amount of user patience.

You are a Reasoning-and-Acting agent

This means you will first think about what to do next according to current observation, then take ONE action via a function call, and then get an observation from the environment or user. You can repeat this process, like "Observation" -> "Thought" -> "Action" -> "Observation" -> ...

IMPORTANT: You MUST call exactly ONE function per turn. Do NOT call multiple functions at once.

Available functions and their costs

Database Environment functions:

- execute(sql): Cost 1
- get_schema(): Cost 1
- get_all_column_meanings(): Cost 1
- get_column_meaning(table_name, column): Cost 0.5
- get_all_external_knowledge_names(): Cost 0.5
- get_knowledge_definition(knowledge_name): Cost 0.5
- get_all_knowledge_definitions(): Cost 1

User Interaction functions:

- ask(question): Cost 2
- submit(sql): Cost 3

Important Strategy Tips

- First explore the database schema, column meaning and external knowledge...
- FIGURE OUT THE USER'S REAL INTENT BY ASKING CLARIFYING QUESTIONS!...
- Be efficient with your actions to conserve user patience
- Make sure your submitted SQL is valid and addresses all aspects of the question
- Keep track of your remaining user patience and prioritize your actions accordingly

[Retrieved Memory with Memory Formatting Template]

-----TASK START-----

Now, let's start with the user's question that may exist ambiguities and require external knowledge understanding to solve. (EACH TIME GIVE ONE ROUND RESPONSE WITH ONE FUNCTION CALL, OTHERWISE YOU WILL BE FIRED!!!)

User's Question: {query}

[SYSTEM NOTE: You have a total action budget of {total_budget} units. Each action consumes budget. Once depleted, you must submit your final SQL solution.]

D.2 Data Augmentation Prompt

DATA AUGMENTATION PROMPT

[SYSTEM]

You convert SQL into two natural-language user queries in the style of the original BIRD-Interact dataset.

Output a single JSON object with keys:

- "clear_query": string (precise, captures key constraints/grouping/metrics/sorting/limit)
- "amb_query": string (plausible but intentionally ambiguous; hide 2-6 important details)
- "ambiguity_map": { "critical_ambiguity": [...], "non_critical_ambiguity": [...] }

Each ambiguity item must be an object with:

- "term": the vague phrase used verbatim in amb_query
- "sql_snippet": a short exact SQL fragment copied from the input SQL (verbatim or near-verbatim)
- "is_mask": boolean
- "type": one of: knowledge_linking_ambiguity, semantic_ambiguity, entity_ambiguity, temporal_ambiguity, aggregation_ambiguity, sort_ambiguity, intent_ambiguity, join_path_ambiguity

Hard style rules (VERY IMPORTANT):

- clear_query and amb_query must sound like a real end-user request, not a schema explanation.
- DO NOT mention raw table names, raw column names, aliases, or dotted identifiers (e.g., hubregistry, distributionhubs, h.hubregistry, s.zoneref).
- DO NOT write parenthetical mappings like "hub (hubregistry)" or "site_code (s.zoneref)".
- Parentheses are allowed ONLY for human-friendly acronyms/units (e.g., "Resource Utilization Ratio (RUR)", "percentage (\%)").
- You MAY use humanized entity names inferred from table names (e.g., distributionhubs -> "distribution hubs"), but never include the raw identifier token.
- DO NOT include SQL keywords (SELECT/FROM/WHERE), raw SQL fragments, or full mathematical formulas in clear_query/amb_query.
- DO NOT explicitly say "no filters" / "no limits" when none exist; just omit that.
- Prefer 1-2 sentences; avoid overly technical wording.

Faithfulness rules:

- Do NOT invent conditions/metrics not supported by the SQL.
- clear_query must faithfully describe: what to return, any filters, grouping, aggregation, sorting direction, and LIMIT (if present) - all in plain English.
- If the SQL computes a derived metric, describe it conceptually in words (and you may name it), but do not show the full expression.
- amb_query must hide 2-6 important details (e.g., exact thresholds, time window, sort direction, grouping key, join path), while remaining plausible.

Alignment rules for ambiguity_map:

- Every ambiguity item term must appear exactly in amb_query.
- Every sql_snippet must be a real fragment present in the input SQL.
- Ensure ambiguity items correspond to the vagueness introduced in amb_query.

Few-shot examples below are FORMAT-ONLY:

- They demonstrate structure and tone only.
- DO NOT reuse any specific nouns/entities/phrases/numbers from the examples.
- Always derive entities, constants, and constraints from the *current input SQL*.

Return ONLY the JSON object.

[USER]

SQL:

<the input SQL, whitespace-stripped>

SQL outline (hint; may be incomplete):

<JSON output of _quick_sql_outline() with keys: has_group_by, has_order_by, has_limit, tables, select_clause, where_clause, group_by_clause, having_clause, order_by_clause, limit_value, aggregations>

Reminder: In clear_query and amb_query, do NOT mention raw identifiers or parenthetical mappings like "X (col)". Use plain English labels only.

D.3 Memory Selector Prompts

The episode-level and turn-level memory selector prompts are shown below.

EPISODE-LEVEL MEMORY SELECTOR PROMPT

You are selecting the top {top_k} most relevant memory ids for a Text-to-SQL task.

[CURRENT QUERY]
{query}

[CANDIDATE MEMORIES]
{candidates_text}

[SELECTION CRITERIA]

Choose memories with:

- Similar task/intent (same type of question or answer: counting, ranking, trend-over-time, comparison, etc.).
- Similar schema/data surface (same or closely related databases, tables, and key columns).
- Similar SQL structure (comparable joins, aggregations, filters, subqueries, window functions, etc.).
- Similar constraints or edge cases (time windows, permissions/tenants, soft deletes, null handling, duplicates, outliers).

Among those, prefer memories that led to a correct final query or clearly improved/fixed a previous failure.

If selecting multiple ids, prefer a diverse set that covers different helpful aspects rather than near-duplicates.

[CRITICAL INSTRUCTION]

Reply with EXACTLY {top_k} id(s) from the brackets above, comma-separated.

Use ONLY the sequential ids in brackets (e.g. 1, 2, 3). Do NOT use the "Original ID".

No explanations. No extra text. Do not invent ids.

Valid examples (using the ids above):

{ex1}

{ex2}

Your selection (ids only):

TURN-LEVEL MEMORY SELECTOR PROMPT

You are selecting the top {top_k} most relevant memory ids for a Text-to-SQL task.

Your job is to decide which past step memories, if any, should be shown to the agent in the CURRENT STATE as guidance or warning.

[CURRENT STATE]
{state_str}

[CANDIDATE STEP MEMORIES]
{candidates_text}

[SELECTION CRITERIA]

Treat each step memory as a snapshot:

- INIT: the initial query.
- Q: the current query after de-ambiguity.
- PREV_ACT: the last action taken.
- PREV_OBS: what was observed or learned from that action.

Choose memories with:

- Strong LOCAL-STATE MATCH: Similar local task/intent in INIT and Q (e.g. choosing tables/columns, designing a join path, shaping an aggregation, fixing a specific error, or refining filters).

- Strong LOCAL-STATE MATCH: Similar tool call in PREV_ACT.
- DATA & OPERATION MATCH: Similar schema context (same or closely related tables, columns, or domain concepts mentioned in the step).
- DATA & OPERATION MATCH: Similar SQL operation being decided or adjusted (comparable joins, grouping keys, aggregations, filters, subqueries, or window functions).
- FEEDBACK / ISSUE PATTERN: Similar type of tool feedback or clarification in PREV_OBS.
- FEEDBACK / ISSUE PATTERN: Similar type of error or problem(e.g. missing column, type mismatch, wrong join logic, wrong aggregation level, unexpected results).

Among those, prefer memories with HIGH GUIDANCE UTILITY.

If selecting multiple ids, prefer steps that cover complementary aspects.

[CRITICAL INSTRUCTION]

Reply with EXACTLY {top_k} id(s) from the brackets above, comma-separated.

Use ONLY the sequential ids in brackets (e.g. 1, 2, 3). Do NOT use the "Original ID".

No explanations. No extra text. Do not invent ids.

Valid examples (using the ids above):

{ex1}

{ex2}

Your selection (ids only):

D.4 Episode Insight Extraction Prompt

EPISODE INSIGHT EXTRACTION PROMPT

You are an expert LLM coach analyzing the behavior of an autonomous SQL reasoning agent.

The agent can operate in two contexts:

- **Environment tools**: execute, get_schema, get_all_column_meanings, get_column_meaning, get_all_external_knowledge_names, get_knowledge_definition, get_all_knowledge_definitions.
- **User tools**: ask (ask clarification), submit (final submission).

Your goal is to extract reasoning insights that help the agent learn *how to plan tool usage* based on the query type and difficulty.

You will analyze one complete episode that includes:

- The query and reward results.
- A chronological record of the agent's tool calls and observations.

Your analysis must address:

1. **Success Strategies** - short bullet points (2-4) summarizing what the agent did right.
2. **Failure Reasons** - short bullet points (2-4) describing what went wrong.
3. **Key Insights** - a concise paragraph that integrates:
 - Whether the query was **read-only** or **CRUD**.
 - Whether it was **easy** or **difficult**.
 - The agent's **behavioral mode**: conservative (asks user clarifications) or aggressive (experiments and self-corrects).
 - Guidance for future similar queries: when to clarify vs. when to explore.

Additionally, extract **info_gaps** resolved in this episode. Only include a gap if:

- The agent explicitly used ``ask(...)`` to request clarification, AND
- The user answered clearly with information that changed SQL semantics or plan.

For each gap, emit a compact object:

```
{
  "gap_key": "<normalized label like bin_boundaries|metric_definition|time_window|join_key|
    null_handling|stdev_variant|timezone>",
  "agent_question": "<short paraphrase of ask(>>",
  "user_answer": "<short distilled answer>",
  "why_important": "<one sentence on how it affects SQL>",
```

```

    "resolution_signature": "<normalized key:value pairs, e.g., bins=low<50,mid50-80,high>80>",
    "confidence": "low|medium|high",
    "scope": "query_specific|reusable"
}

Respond in **JSON only**:
{
  "success_strategies": [...],
  "failure_reasons": [...],
  "key_insights": "...",
  "info_gaps_resolved": [
    {
      "key": "snake_case label describing the gap. Prefer one of: classification_scheme,
metric_definition, grouping_rules, bin_boundaries, time_window, join_keys, null_handling,
deduplication_rule, filter_scope, timezone_granularity. If none fit, use a precise, custom
snake_case label (e.g., stdev_variant, sampling_ratio, outlier_policy). Do NOT use generic
values like \"other\".",
      "question": "short agent question",
      "user_answer": "short canonical answer",
      "why_useful": "how it changed SQL/logic",
      "used_in": "brief evidence (e.g., final SQL CASE bins)",
      "scope": "generalizable | this_query_only",
      "confidence": "high | medium | low"
    }
  ]
}

```

D.5 PRM Judging Prompt

PRM JUDGING PROMPT

You are an expert coach for a Text-to-SQL agent.

In each evaluation, you see:

- 1) A CURRENT agent state (called CURRENT_STATE), with fields like:
 - initial_query: the original user question.
 - current_query_context: the query plus any clarifications so far.
 - prev_action: the last tool the agent used.
 - prev_observation: the last observation from the environment or user.
- 2) ONE step-wise memory from a past episode (called CANDIDATE_MEMORY_STEP), with:
 - trigger_state_raw: a compact signature of the state in which that past step occurred.
 - action_skeleton: a normalized description of the tool call taken in that step (e.g., "TOOL:EXECUTE SQL:[...]", "TOOL:ASK [CLARIFICATION]", "TOOL:GET_SCHEMA").
 - action_result_raw: a short snippet of what happened after that action (e.g., success, SQL error, clarified user answer, etc.).
 - outcome_type: a coarse label "POSITIVE" or "NEGATIVE" for that past step.
 - note: optional comments.

Your job:

- Judge how useful this step-wise memory would be if retrieved **now** in the CURRENT_STATE as a small hint for the agent.
- This hint can be:
 - Positive guidance: "This kind of action worked well before in a similar state."
 - Negative warning: "Avoid this kind of action; it produced a failure before."

IMPORTANT:

- You are NOT ranking among multiple memories here. You only see ONE candidate and judge it in isolation.
- Both POSITIVE and NEGATIVE past steps can be valuable:
 - A POSITIVE step is good as a pattern to imitate.
 - A NEGATIVE step is good as a warning to avoid repeating the mistake.
- For EACH rubric principle, you MUST:

- provide a numerical score between 0 and 5 (inclusive),
- AND a one-sentence explanation (reason) for that score.

SCORING SCALE (for all rubric scores):

- 0 = not applicable / useless / completely wrong
- 1 = very weak
- 2 = weak to moderate
- 3 = moderate / acceptable
- 4 = strong
- 5 = extremely strong / ideal

RUBRIC PRINCIPLES (each must have { "score": 0-5, "reason": "..."}):

- 1) state_match
 - How similar is the CURRENT_STATE to the trigger_state_raw of this memory?
 - Consider query intent, ambiguity, and the previous action/observation.
 - 0 = completely different situation; the memory comes from an unrelated context.
 - 5 = almost the same state; the memory was recorded in a very similar situation.
- 2) actionability_value
 - If we SHOW this memory to the agent, how strongly does it provide useful behavioral guidance (either as something to imitate, or as something to avoid)?
 - High when the memory clearly suggests a good next action OR a clear mistake to avoid in states like the CURRENT_STATE.
 - 0 = the memory provides essentially no actionable guidance here.
 - 5 = extremely strong, concrete guidance for what to do or not do next.
- 3) pattern_generalizability
 - How general is the pattern implied by this memory (either positive or negative)?
 - 0 = extremely specific to a single query/dataset/id; unlikely to transfer.
 - 5 = broadly applicable pattern for many similar Text-to-SQL situations.
- 4) outcome_reliability
 - How reliable is this memory as evidence for what to do or avoid?
 - Consider:
 - Does action_result_raw clearly support the outcome (positive or negative)?
 - Does outcome_type match what actually happened?
 - 0 = noisy, ambiguous, or unreliable.
 - 5 = clear and trustworthy signal.
- 5) clarity_for_agent
 - If we compress this step into 1-3 lines of "tip text" for the agent, how easy is it for the agent to understand what behavior is being recommended or avoided?
 - 0 = very hard to explain; cryptic or confusing.
 - 5 = crystal clear guidance or warning.
- 6) confidence_in_assessment
 - How confident are you in your own judgments for this memory, especially for state_match, actionability_value, and overall_utility, given the available information?
 - Consider ambiguity or missing information in CURRENT_STATE and CANDIDATE_MEMORY_STEP.
 - 0 = extremely uncertain; the scores above are mostly guesses.
 - 5 = very confident; the situation and memory are clear and well supported.
- 7) overall_utility
 - Your overall judgment (0-5) of how valuable it is to retrieve THIS memory for THIS CURRENT_STATE.
 - It should go up when:
 - state_match is high, AND
 - actionability_value is high, AND
 - pattern_generalizability and outcome_reliability are reasonably high.
 - 0 = do not retrieve; essentially useless or risky.
 - 5 = very strong candidate to retrieve as guidance or warning.

BINARY FLAGS:

- recommend_injection:

- true if you think this memory should be retrieved and shown to the agent in this CURRENT_STATE.
- false if it should generally be skipped.
- use_as_warning_only:
 - true if its main value is as a "do NOT repeat this mistake" warning (typically when outcome_type is NEGATIVE).
 - false if it is mainly a positive pattern to imitate, or a balanced mix.

TEXT FEEDBACK:

- one_sentence_summary:
 - One sentence summarizing why this memory is or isn't useful here.
- how_to_use:
 - A short note describing how the selector should use this memory, e.g.:
 - "Use as a positive template for executing a grouped aggregation after schema inspection."
 - "Use only as a warning about executing complex SQL before clarifying the time window."
 - "Skip in most cases; too specific to a different schema."

OUTPUT FORMAT (STRICT):

You must respond with a SINGLE JSON object and nothing else.
Do NOT wrap it in Markdown code fences.

The JSON must match this schema (scores 0-5, reasons are one sentence):

```
{
  "rubric": {
    "state_match": {
      "score": number,
      "reason": string
    },
    "actionability_value": {
      "score": number,
      "reason": string
    },
    "pattern_generalizability": {
      "score": number,
      "reason": string
    },
    "outcome_reliability": {
      "score": number,
      "reason": string
    },
    "clarity_for_agent": {
      "score": number,
      "reason": string
    },
    "confidence_in_assessment": {
      "score": number,
      "reason": string
    },
    "overall_utility": {
      "score": number,
      "reason": string
    }
  },
  "binary_flags": {
    "recommend_injection": boolean,
    "use_as_warning_only": boolean
  },
  "text_feedback": {
    "one_sentence_summary": string,
    "how_to_use": string
  }
}
```

All scores must be in [0, 5].
Each reason must be a single concise sentence.
Return ONLY this JSON object as plain text.

CURRENT_STATE:
{state_text}

CANDIDATE_MEMORY_STEP:
{memory_text}

You must now respond with a SINGLE JSON object and nothing else.