

Self-Guided Function Calling in Large Language Models via Stepwise Experience Recall

Sijia Cui^{1,2}, Aiyao He¹, Shuai Xu³, Hongming Zhang¹, Yanna Wang^{1†},
Qingyang Zhang¹, Yajing Wang⁴, Bo Xu^{1†}

¹The Key Laboratory of Cognition and Decision Intelligence for Complex Systems,
Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³Nanjing University of Information Science & Technology

⁴Institute of Computing Technology, Chinese Academy of Sciences

[†]Correspondence: wangyanna2013@ia.ac.cn, boxu@ia.ac.cn

Abstract

Function calling enables large language models (LLMs) to interact with external systems by leveraging tools and APIs. When faced with multi-step tool usage, LLMs still struggle with tool selection, parameter generation, and tool-chain planning. Existing methods typically rely on manually designing task-specific demonstrations, or retrieving from a curated library. These approaches demand substantial expert effort and prompt engineering becomes increasingly complex and inefficient as tool diversity and task difficulty scale. To address these challenges, we propose a self-guided method, **Stepwise Experience Recall (SEER)**, which performs fine-grained, stepwise retrieval from a continually updated experience pool. Instead of relying on static or manually curated library, SEER incrementally augments the experience pool with past successful trajectories, enabling continuous expansion of the pool and improved model performance over time. Evaluated on the ToolQA benchmark, SEER achieves an average improvement of 6.1% on easy and 4.7% on hard questions. We further test SEER on τ -bench, which includes two real-world domains. Powered by Qwen2.5-7B and Qwen2.5-72B models, SEER demonstrates substantial accuracy gains of 7.44% and 23.38%, respectively.

1 Introduction

Large language models (LLMs) demonstrated remarkable capabilities through pretraining on large-scale corpora (Brown et al., 2020; Devlin et al., 2019; Touvron et al., 2023; Achiam et al., 2023; Denison et al., 2024; Team et al., 2023). However, due to the inherent limitations of neural network architecture, LLMs are unable to interact directly with the real world—an issue that cannot be resolved simply by scaling up the training data or model size. Function calling¹ (Qu et al., 2025; Qin et al., 2024a,b) serves as a fundamental mechanism

that enables LLMs to interact with external systems. By invoking external tools, LLMs can integrate up-to-date knowledge and execute real-world tasks, thereby expanding the boundaries of LLM-based AI agents and driving advancements across various domains (Hao et al., 2025; Theuma and Shareghi, 2024; Zhong et al., 2023; Zhao et al., 2024).

In-context learning (Brown et al., 2020) enhances LLM reasoning by embedding task-specific examples in prompts, enabling adaptation to new tasks without training. However, this approach faces significant challenges in multi-step tool-use scenarios. Limited by the maximum token length, prompts cannot include examples that comprehensively cover all tools and problem types. Moreover, the relevance and complexity of the demonstrations directly impact the model’s performance (Zhao et al., 2021; Min et al., 2022; Dong et al., 2024). This raises the critical question: How can we dynamically select examples tailored to the specific problem at hand, especially when the task involves multiple steps and complex tool interactions?

Existing methods (Paranjape et al., 2023; Guan et al., 2025) typically rely on coarse-grained retrieval strategies, which fail to account for the nuanced relationships between tool usage patterns and user objectives in multi-step function calling. These approaches emphasize task similarity while overlooking the critical role of tool-chain alignment in achieving accurate and efficient outcomes. Additionally, several approaches (Zhao et al., 2024; Xu et al., 2024) depend on manually curated or pre-collected task-specific demonstrations. However, this reliance not only limits scalability but also incurs substantial offline costs, making it inefficient when addressing a wide range of diverse tasks.

We introduce **Stepwise Experience Recall (SEER)**², a novel approach that enhances the multi-step tool-use capabilities of LLMs through fine-

¹We use function calling and tool-use interchangeably.

²<https://github.com/AI-Research-TeamX/SEER>

grained retrieval. It selects relevant trajectories by jointly considering task similarity, toolchain coverage, and intent alignment. SEER incrementally expands the experience pool by incorporating successful task trajectories, improving model performance over time. For tasks without explicit success signals, we adopt an LLM-as-a-judge mechanism (Li et al., 2024) to assess task completion. This enables continuous online updates to the pool, allowing SEER to adapt to new tasks and evolving user demands. Our main contributions are:

- We propose stepwise experience recall, retrieving relevant examples based on trajectory similarity, toolchain coverage, and user intent. This fine-grained retrieval effectively leverages experience from prior successful trajectories.
- We introduce online experience accumulation, dynamically adding successful multi-step tool invocation trajectories to the experience pool. This reduces reliance on manual annotations and enables the model to online self-improve.
- We conduct comprehensive evaluations on ToolQA and τ -bench, and the results show that SEER outperforms existing methods. Meanwhile, the self-improvement results show clear and consistent performance gains over time, demonstrating the effectiveness of SEER and the potential for self-guided function calling.
- We perform extensive ablation studies involving different retrieval strategies and few-shot settings, aiming to highlight the contribution of each scoring component and impact of the number of demonstrations on performance.

2 Related Work

2.1 Multi-step Function Calling

Recent progress in tool-augmented LLMs has focused on either freezing or training approaches (Wang et al., 2024; Huang et al., 2023; Yu et al., 2024; Goldie et al., 2025). Many studies exploit LLMs’ in-context learning by prompting task descriptions and tool-use examples during inference (Lu et al., 2023; Shen et al., 2023; Hsieh et al., 2023; Paranjape et al., 2023; Bai et al., 2024; Zhang et al., 2025; Yang et al., 2025; Xu et al., 2025; Cui et al., 2025). For example, the ART framework (Paranjape et al., 2023) retrieves similar multi-step reasoning and traces to guide models in generating intermediate steps and invoking functions. Other

methods, such as StepTool (Yu et al., 2024) and SWiRL (Goldie et al., 2025), treat tool use as a reinforcement learning problem (Sutton and Barto, 2018; Dong et al., 2020; Zhang and Yu, 2020b). StepTool applies step-wise reward shaping and policy gradients to improve decision-making based on tool success and task contribution. SWiRL generates synthetic multi-step tool-use data and uses step-wise RL with reward models to train without manual labels. In contrast, we introduce an experience replay approach that enhances multi-step tool use without additional training, using fine-grained replay to improve performance efficiently.

2.2 Self-improvement for LLM

Accelerated advancements in LLMs have intensified data scarcity issues, highlighting data bottlenecks as a major research challenge (Villalobos et al., 2022). Self-improvement involving model-generated data such as feedback, instructions, and questions, has shown promise but often relies on heuristics and human validation for quality assurance (Bai et al., 2022; Wang et al., 2022). Systems like ExpeL (Zhao et al., 2024) leverage past task experiences to enhance decision-making at inference, and recent work by (Tian et al., 2024) integrates Monte Carlo Tree Search (MCTS) (Kocsis and Szepesvári, 2006; Zhang and Yu, 2020a) with language models, creating annotation-free self-improvement loops. However, these methods typically rely on static offline datasets, significantly limiting adaptability in practical scenarios. To address this, we propose an online updating experience pool that continuously supports in-context learning, enhancing inference quality and real-world adaptability.

3 Method

In this section, we present Stepwise Experience Recall (SEER), a novel approach that retrieves prior successful trajectories as in-context examples. SEER consists of three core components: trajectory experience extraction, stepwise experience recall, and continual experience accumulation, which together enable dynamic and efficient demonstration selection. The framework is illustrated in Figure 1. We first present the notation and problem formulation, then detail each SEER component.

3.1 Problem Formulation

Building upon the formulation introduced in (Zhao et al., 2024), we consider an interactive task

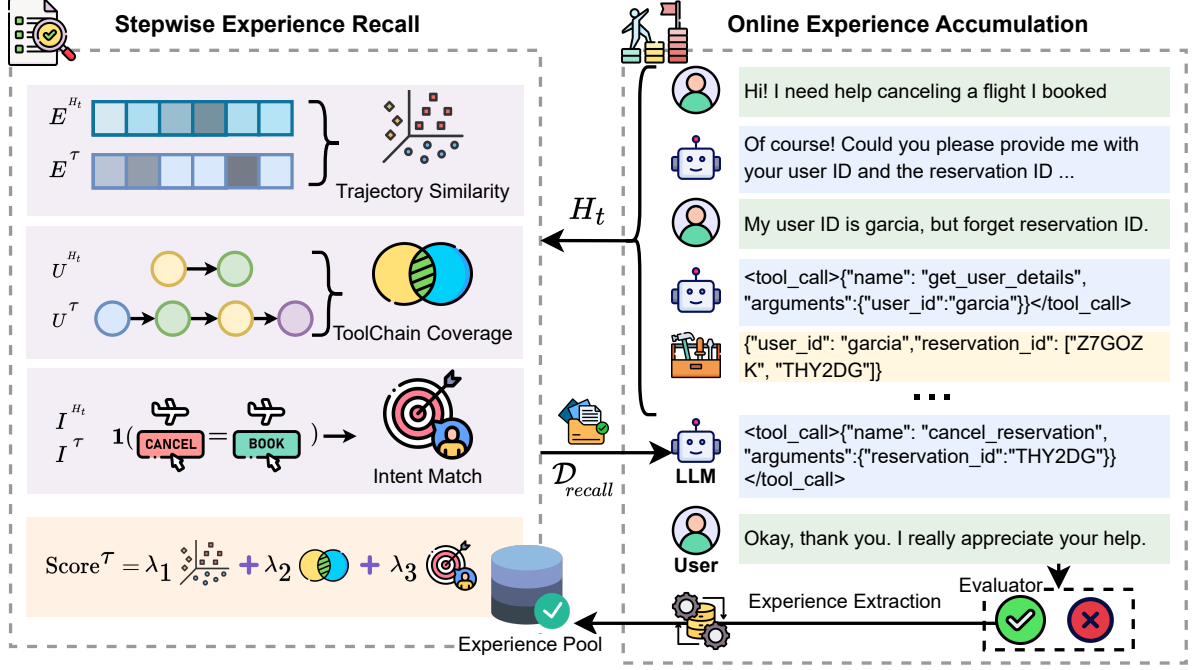


Figure 1: Overview of the SEER framework. The core component is the stepwise experience recall (left), which retrieves relevant trajectories from the experience pool based on the current interaction history H_t , and returns the top- k examples $\mathcal{D}_{\text{recall}}$ to guide the LLM’s next decision. The continual experience accumulation mechanism (right) updates the experience pool by identifying successful trajectories using an evaluator.

with tool-augmented large language model (LLM) agents. In this setting, an LLM-based agent is required to achieve a certain user-defined goal $g \in \mathcal{G}$ by interacting with a user and utilizing a set of external tools. The interaction unfolds over a finite horizon of T steps, indexed by $i \in 0, \dots, T$.

At each step i , the agent receives an observation $o_i \in \mathcal{O}$, where $\mathcal{O}$ denotes the joint observation space composed of both the user’s input (or feedback) and the outputs returned by the tools. Formally, the observation space is the Cartesian product $\mathcal{O} = \mathcal{O}_{\text{user}} \times \mathcal{O}_{\text{tool}}$, where $\mathcal{O}_{\text{user}}$ is the user response space, and $\mathcal{O}_{\text{tool}}$ is the tool output space.

The agent maintains an interaction history $H_t = \{o_0, a_0, \dots, o_{t-1}, a_{t-1}, o_t\}$ up to the current time step t . Based on this history, the agent selects an action $a_t \in \mathcal{A}$, where $\mathcal{A}$ denotes the action space, including both natural language responses to the user and tool invocations. Once the action is executed, the agent receives a new observation o_{t+1} , which includes the user feedback or the tool result. It continues until the goal g is achieved or a maximum interaction step T is reached, finally yielding a complete trajectory $\tau = \{o_0, a_0, o_1, a_1, \dots\}$.

3.2 Trajectory Experience Extraction

Trajectory experience extraction transforms the interaction trajectory $\tau = \{o_0, a_0, o_1, a_1, \dots\}$ into a

structured experience representation:

$$d^\tau = \langle E^\tau, E^q, I^\tau, U^\tau \rangle$$

Here, E^τ denotes the embedding of the interaction trajectory τ , and E^q is the embedding of the user’s first query, the agent’s first observation o_0 , both generated using a pre-trained embedding model. I represents the inferred user intent, selected from a predefined discrete intent set $\mathcal{I}$. Rather than relying on manual annotation, we employ the LLM itself to classify the user’s goal based on the initial query. U captures the sequence of tool invocations, which is modeled as a directed path $u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_n$, where each u_i represents the i -th tool used within the overall trajectory.

3.3 Stepwise Experience Recall

Traditional retrieval methods primarily rely on the similarity between task instructions or user queries to select in-context examples. However, this approach often fails to capture the dynamic nature of multi-turn interactions and multi-step tool invocation patterns. To address this limitation, we propose a multi-dimensional scoring strategy in SEER that considers three key aspects: trajectory similarity, toolchain coverage, and intent alignment.

First, we compute trajectory similarity by comparing the overall embeddings of interaction histo-

ries. This allows the recall mechanism to go beyond surface-level query matching and retrieve examples with similar structural progressions and decision patterns, providing richer contextual guidance. Then, we introduce a toolchain coverage score to account for similarities in tool usage sequences. Even when user goals appear superficially different, similar toolchains often reflect shared reasoning strategies or problem-solving procedures. By identifying examples with overlapping tool invocation patterns, SEER promotes the reuse of effective operational knowledge. Finally, we incorporate an intent match score based on inferred user intents. This helps the system better focus on the semantic core of the user’s goal. By aligning closely with user intent, SEER improves both the relevance and coherence of the retrieved demonstrations.

This multi-dimensional scoring enables SEER to recall more contextually appropriate and semantically aligned examples, thereby enhancing the agent’s ability to generalize and perform effectively across diverse tool-usage tasks. Specifically, given the current interaction history $H_t = \{o_0, a_0, o_1, a_1, \dots, o_t\}$ and a candidate trajectory $\tau' = \{o'_0, a'_0, o'_1, a'_1, \dots, o'_t\}$ from the experience pool, we compute a relevance score between corresponding experience representations d^{H_t} and $d^{\tau'}$. The score comprises the following components:

- **Trajectory Similarity** ($s_1 \in [0, 1]$): Normalized cosine similarity between the embedding vectors: $s_1 = (1 + \cos(E^{H_t}, E^{\tau'}))/2$
- **ToolChain Coverage** ($s_2 \in [0, 1]$): Proportion of tools in the current task that are also present in the candidate trajectory:

$$s_2 = |U^{H_t} \cap U^{\tau'}| / |U^{H_t}|$$

When calculating ToolChain Coverage, we ignore the directionality of U and treat it as an unordered set of tools: $\{u_0, u_1, \dots\}$.

- **Intent Match** ($s_3 \in \{0, 1\}$): Whether the inferred user intents are identical: $s_3 = \mathbf{1}[I^{H_t} = I^{\tau'}]$, and $\mathbf{1}$ is the indicator function.

The final relevance score is a weighted sum: $\text{Score}^\tau = \sum_{i=1}^3 \lambda_i s_i$, where $\lambda_1, \lambda_2, \lambda_3$ are hyperparameters controlling the contribution of each component. The top- k demonstration $\mathcal{D}_{\text{recall}}$ with the highest relevance scores are selected as in-context exemplars to guide the LLM’s next decision a_t . The detailed pseudocode for the stepwise experience recall is presented in Algorithm 1.

Algorithm 1: Stepwise Experience Recall

Input: Current interaction history H_t ,
experience pool $\mathcal{D}$, $\lambda_1, \lambda_2, \lambda_3$

Output: Top- k relevant trajectories $\mathcal{D}_{\text{recall}}$

```

1 foreach  $(\tau', d^{\tau'}) \in \mathcal{D}$  do
2    $d^{\tau'} = \langle E^{\tau'}, E^q, I^{\tau'}, U^{\tau'} \rangle$ ;
3   Extract  $E^{H_t}, E^q, I^{H_t}, U^{H_t}$  from  $H_t$ ;
4    $s_1 \leftarrow (1 + \cos(E^{H_t}, E^{\tau'}))/2$ ;
5    $s_2 \leftarrow |U^{H_t} \cap U^{\tau'}| / |U^{H_t}|$ ;
6    $s_3 \leftarrow \mathbf{1}(I^{H_t} = I^{\tau'})$ ;
7    $\text{Score}^{\tau'} \leftarrow \lambda_1 s_1 + \lambda_2 s_2 + \lambda_3 s_3$ ;
8 Sort all  $(\tau', \text{Score}^{\tau'})$  pairs by  $\text{Score}^{\tau'}$  in
   descending order;
9 Select top- $k$  trajectories to form  $\mathcal{D}_{\text{recall}}$ ;
10 return  $\mathcal{D}_{\text{recall}}$ 

```

3.4 Continual Experience Accumulation

Unlike prior methods that rely on static, offline datasets, SEER incrementally builds its trajectory pool during deployment. However, the online nature of this approach presents a challenge: the lack of explicit signals indicating task completion. Inspired by works such as (Li et al., 2024), we leverage LLM’s self-assessment capabilities to mitigate this issue. Specifically, after completing a task, the system performs self-evaluation by comparing its own output against the reference answer. The evaluator returns a binary judgment indicating whether the task is a success or failure, determining whether the trajectory should be added to the experience pool. To ensure robustness, the evaluation logic is designed to tolerate minor discrepancies, such as formatting variations or slight numerical differences, and still regard them as successful outcomes.

By evaluating the correctness of the generated trajectory, we can determine whether to add the trajectory to the experience pool. This self-guided mechanism allows for continuous updates to the trajectory experience library, ensuring that the model remains adaptable and benefits from newly encountered cases without requiring extensive pre-collected data. We show the continual experience accumulation process in Algorithm 2. We leave the implementation details of intent inference and LLM evaluator to the Appendix A.1 and A.2.

4 Experiments

We conducted extensive experiments to evaluate the performance of SEER on two benchmarks: ToolQA

Algorithm 2: Experience Pool Update

Input: Current interaction history H_t , LLM evaluator $\mathcal{E}$, experience pool $\mathcal{D}$

Output: Updated experience pool

```

1 if  $\mathcal{E}.isSuccessful(H_t)$  then
2   Extract  $E^\tau, E^q, I^\tau, U^\tau$  from  $H_t$ ;
3   Form experience tuple
     $d^\tau \leftarrow \langle E^\tau, E^q, I^\tau, U^\tau \rangle$ ;
4   Insert  $(\tau, d^\tau)$  into experience pool
     $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\tau, d)\}$ ;
5 return  $\mathcal{D}$ 

```

and τ -bench. We first present the experimental setup, including benchmarks, evaluation metrics, and baseline methods. We then show the main results and conduct ablation studies to assess the contribution of core components in SEER. Finally, we highlight several insightful findings. In all result tables, the best performance is indicated in **bold**. Our experiments are comprehensively designed to answer the following key questions:

- How does SEER perform on both ToolQA and τ -bench compared to existing baselines?
- Can SEER self-improve? Does its performance increase as the experience pool grows?
- How do different retrieval strategies impact the overall performance of SEER? In particular, does fine-grained retrieval yield better results?
- What is the effect of varying the top- k value in SEER’s stepwise experience recall mechanism?

4.1 Experimental Setup

We use Qwen2.5-72B-Instruct (Qwen et al., 2025) as the foundational LLM for both baseline methods and SEER. To ensure reproducibility and reduce randomness, we set model temperature to 0 across all experiments. For embedding-based retrieval, we adopt the bge-large-en-v1.5 model (Xiao et al., 2024). Unless otherwise specified, the number of retrieved examples (top- k) is set to 4. The hyperparameters $\lambda_1, \lambda_2, \lambda_3$ are set to 1/3, 1/3, and 1/3, respectively. Maximum interaction steps $T = 6$ for ToolQA and $T = 30$ for τ -bench. The maximum size of the experience pool is set to 1000. We initialize the demonstration pool with 8 same examples for ART, ExpeL, and SEER in ToolQA benchmark. In the τ -bench setting, we initialize

the demonstration pool with 2 examples and use GPT-4o as the user in the simulated environment.

Benchmarks. We primarily evaluated SEER on two challenging benchmarks designed for tool-augmented LLMs, shown in Table 1 and Table 2. ToolQA (Zhuang et al., 2023) spans 8 real-world domains—air transportation, financial data, commercial services, lodging platforms, social networks, academic publications, personal agendas, and numerical reasoning. The easy set comprises 800 questions across 55 templates, while the hard set includes 730 questions from 62 templates. ToolQA assesses LLMs’ ability to reason across multi-step and use external tools effectively. τ -bench (Yao et al., 2025) evaluates tool use in realistic, multi-turn tasks across airline (115 tasks, 15 tools) and retail (50 tasks, 13 tools) domains. Each task includes a user model and an LLM agent, simulating dynamic interactions and tool usage. Unlike static, single-turn question-answering settings, τ -bench is specifically designed to assess LLM performance in dynamic, real-world scenarios, emphasizing multi-turn interaction, evolving user intent, and multi-step tool use.

Table 1: An overview of the statistics of ToolQA.

Domain	Data Type	Data Volume	Easy Questions		Hard Questions	
			Template	Count	Template	Count
Flight	Structured DB	4,078,318	10	100	10	100
Coffee	Structured DB	5,746	8	100	13	130
Yelp	Structured DB	150,346	11	100	10	100
Airbnb	Structured DB	102,599	10	100	10	100
GSM8K	Professional	-	-	-	-	-
DBLP	Graph DB	553,320	10	100	10	100
SciREX	Text Corpus	438	1	100	4	100
Agenda	Text Corpus	10,000	5	100	5	100
Total	-	-	55	800	62	730

Table 2: An overview of the statistics of τ -bench.

Domain	Databases	Tools	Questions
τ -retail	500 users, 50 products, 1,000 orders	15	115
τ -airline	500 users, 300 flights, 2,000 reservations	13	50
Total	-	-	165

Evaluation Metrics. We use accuracy as the primary evaluation metric for assessing model performance on ToolQA. For a given question set Q_j , the accuracy is defined as: $Acc_j = \frac{1}{|Q_j|} \sum_{i=1}^{|Q_j|} \mathbf{1}[y_i = y'_i]$, where y_i is the ground-truth and y'_i is the predicted answer for the i -th question in Q_j . To evaluate performance across multiple domains, we compute the average accuracy as: $Acc^{avg} = \frac{1}{|D|} \sum_{j=1}^{|D|} Acc_j$, where $|D|$ denotes the number of domains in ToolQA, and $Accuracy_j$ is the accuracy within the j -th domain. In τ -bench, pass@k

Table 3: Main results on the ToolQA Benchmark. The results are reported in terms of Accuracy (%). The best results are highlighted in **bold**. Our method SEER achieves the best performance on average accuracy across all tasks, outperforming the second-best method ExpeL by 6.1% and 4.7% on Easy and Hard tasks, respectively.

Method	Flight		Coffee		Yelp		Airbnb		DBLP		SciREX		Agenda		GSM8K		Avg.	
	E	H	E	H	E	H	E	H	E	H	E	H	E	H	E	H	E	H
Chameleon	37.0	2.0	61.0	2.3	60.0	13.0	12.0	4.0	25.0	17.0	6.0	19.0	57.0	0.0	17.0	34.5	8.2	
CoT	37.0	22.0	79.0	13.1	43.0	52.0	79.0	19.0	0.0	1.0	3.0	22.0	54.0	0.0	2.0	37.1	18.4	
ReAct	67.0	6.0	94.0	25.4	70.0	17.0	86.0	17.0	27.0	22.0	6.0	18.0	61.0	0.0	64.0	59.5	15.1	
TUMS	64.0	9.0	93.0	22.3	73.0	15.0	91.0	11.0	34.0	30.0	4.0	14.0	59.0	0.0	72.0	61.3	14.5	
ART	73.0	20.0	99.0	33.1	81.0	30.0	94.0	25.0	33.0	33.0	3.0	25.0	1.0	0.0	73.0	57.1	23.7	
ExpeL	81.0	29.0	92.0	34.6	77.0	40.0	88.0	32.0	37.0	29.0	5.0	19.0	57.0	1.0	57.0	61.8	26.4	
SEER (Ours)	82.0	25.0	87.0	41.5	90.0	58.0	94.0	33.0	37.0	35.0	6.0	23.0	68.0	2.0	79.0	67.9	31.1	

is a metric used to evaluate the probability that an LLM agent successfully completes the same task in all k independent dialogue trials.³ The metric is defined as: $\text{pass}^k = \mathbb{E}_{\text{task}} \left[\frac{\binom{c}{k}}{\binom{n}{k}} \right]$, where c is the number of successful trials out of n total trials for a given task, and the expectation is taken over all tasks. Although pass^k captures robustness over repeated attempts, we report pass^1 by default. This simplifies the average reward (i.e., success rate) across tasks and serves as a standard baseline for evaluating the agent’s single-shot effectiveness.

Baselines. To rigorously evaluate SEER, we compare it against five representative baselines: Chameleon (Lu et al., 2023), ReAct (Yao et al., 2023), TUMS (He et al., 2025), ART (Paranjape et al., 2023), and ExpeL (Zhao et al., 2024). These methods cover a broad spectrum of strategies for tool-augmented LLMs, ranging from direct prompting to sophisticated reasoning and multi-step tool use. For evaluation on τ -bench, we include three closed-source and three open-source LLMs as reference models. Specifically, we test both the original versions of Qwen2.5-7B and Qwen2.5-72B, as well as their SEER-enhanced counterparts, to quantify the performance improvement introduced by our method. Appendix B shows the details of baselines.

4.2 Main Results

Table 3 shows the results on the ToolQA benchmark. Our method, SEER, achieves consistent and significant improvements across multiple datasets. On easy questions, SEER outperforms the strongest baseline, ExpeL (Zhao et al., 2024), by 6.1% in average accuracy. For hard questions, the improvement is 4.7%. Overall, SEER achieves an average accuracy of 67.9% on easy sets and 31.1% on hard

³To avoid ambiguity, we clarify that in this context, k refers to the number of trials. In all other parts of the paper, k denotes the number of demonstrations.

Table 4: Main results on τ -bench. With the integration of SEER, the performance of two open-source models is significantly enhanced. SEER (72B) achieves 51.84%, approaching the performance of GPT-4o at 54.76%.

Methods	Airline	Retail	Avg.
<i>Close source models</i>			
Claude 3.5 Sonnet	48.98	70.18	59.58
GPT-4o	42.86	66.67	54.76
GPT-4o-mini	22.45	47.37	34.91
<i>Open source models</i>			
DeepSeek-v3	46.94	67.54	57.24
Qwen2.5-72B-Instruct	30.61	26.32	28.46
Qwen2.5-7B-Instruct	8.16	10.53	9.34
<i>Ours</i>			
SEER(7B)	20.41	13.16	16.78
SEER(72B)	38.78	64.91	51.84

sets. Specifically, SEER attains the best results on both easy and hard subsets of Yelp, Airbnb, DBLP, and Agenda. On Flight-Hard and Scirex-Hard, it shows a slight performance drop compared to the best baseline. Notably, performance on Coffee-Easy declines more noticeably. We attribute this to suboptimal and cumbersome examples, which led to overthinking simple tasks. To address this, we incorporate reflection (Shinn et al., 2023) to identify and correct errors. With this enhancement, SEER+Reflection achieves 69.0% and 32.0%. Detailed results are provided in Appendix C.

The results on the τ -bench benchmark are shown in Table 4. Using the SEER method with the Qwen2.5-7B model, performance on the Airline task improves from 8.16% to 20.41%, and on the Retail task from 10.53% to 13.16%, yielding an overall average gain from 9.34% to 16.78%. A similar trend is observed with the more powerful Qwen2.5-72B model, showing consistent perfor-

mance improvements across both tasks. Notably, SEER equipped with Qwen2.5-72B achieves a final performance of 51.84%, with a modest gap compared to the GPT-4o, which scores 54.76%.

Overall, across eight datasets of varying difficulty and two real-world tasks, SEER consistently outperforms existing baseline methods, demonstrating the effectiveness and robustness of SEER in enhancing the multi-step and multi-turn function-calling capabilities of LLMs. To further investigate the nature of SEER’s self-guided mechanism and the impact of SEER components, we present detailed analyses from self-improvement experiments and ablation studies in the following sections.

4.3 Self-Improvement




Figure 2: The self-improvement of SEER. The red solid line represents SEER’s average accuracy per batch. The blue dashed line represents a 3-point moving average.

A key distinction between our method and existing approaches lies in how the demonstration pool is updated: instead of relying on a static, pre-collection pool, our method adopts an online self-guided mechanism. After each successful task completion, the corresponding trajectory is added to the demonstration pool, enabling continual refinement and enrichment over time. To assess the self-improvement mechanism, we conduct a controlled experiment on the ToolQA benchmark. Specifically, we randomly shuffle all 1,530 questions and divide them into 10 batches, each containing 153 questions. For each batch, we perform offline evaluation and, upon completion, add all correctly answered instances into the demonstration pool for retrieval in the next batch. We compute the batch accuracy for each batch: $Acc_j^{\text{batch}} = \frac{1}{|Q_{b_j}|} \sum_{i=1}^{|Q_{b_j}|} \mathbf{1}[y_i = y'_i]$, where $|Q_{b_j}|$ is the number of questions in the batch b_j , y_i and y'_i is the ground-truth and predicted answer for the i -th question in batch b_j . The 3-point moving average is used to smooth the results, which is defined as:

$$\hat{Acc}_j^{\text{batch}} = (Acc_{j-1}^{\text{batch}} + Acc_j^{\text{batch}} + Acc_{j+1}^{\text{batch}}) / 3.$$

We visualize the accuracy trend across successive batches in Figure 2, where a clear upward trajectory is observed. Initially, SEER’s accuracy is relatively low, at 37.7% at the first batch. However, as more batches are processed, SEER begins to effectively leverage the experience pool through stepwise recall. By the fifth batch, SEER surpasses baseline methods, reaching 52.3% accuracy. This improvement continues, culminating in a batch accuracy of 54.9% at the last batch. Overall, the consistent upward trend in performance underscores SEER’s capacity for self-improvement, validating the effectiveness of its self-guided learning strategy. These highlight SEER’s potential for continual self-guidance and adaptation in real-world scenarios.

4.4 Ablation Studies

In addition to the main experiments described above, we have conducted ablation studies on the ToolQA benchmark to evaluate the effectiveness of different components in our method. Specifically, we experimented with various retrieval strategies and different numbers of demonstrations to reveal their contributions to overall performance.

Retrieval Strategies. The retrieval score in SEER consists of three components: trajectory embedding similarity s_1 , toolchain coverage s_2 , and intent match score s_3 . We compare different retrieval configurations, including SEER (w/o s_2) and SEER (w/o s_3), to analyze the impact of s_2 and s_3 on retrieval performance. For s_1 , we also explore a query-based variant method, SEER (query-based). Specifically, instead of computing similarity over the entire trajectory, we compute $s_1 = (1 + \cos(E^q, E^{q'}))/2$, where q is the first observation o_0 , i.e., the user’s initial query. This allows us to isolate and examine the effectiveness of query-level semantic similarity.

As shown in Figure 3, all three retrieval variants result in performance degradation compared to the full SEER method. The performance drop in SEER (query-based) is relatively modest. This is expected, as the ToolQA benchmark adopts a one-turn multi-step setting where the user’s query does not change across turns, making full-trajectory and query-only representations less divergent. However, a more pronounced drop is observed in SEER (w/o s_3) when applied to easy questions. This indicates that retrieving exemplars with aligned user intent is particularly beneficial for tasks involving straightforward reasoning and tool usage. In

Table 5: Performance on easy and hard questions under different top- k few-shot settings. The results exhibit a trend of increasing performance followed by a decline as k increases. SEER(Q) denotes SEER (Query-based).

Method	Top-k	Flight		Coffee		Yelp		Airbnb		DBLP		SciREX		Agenda		Avg.	
		E	H	E	H	E	H	E	H	E	H	E	H	E	H	E	H
SEER(Q)	0	37.0	22.0	79.0	13.08	43.0	52.0	79.0	19.0	0.0	1.0	3.0	22.0	54.0	2.0	37.13	18.49
	2	81.0	18.0	95.0	40.0	80.0	60.0	94.0	21.0	33.0	27.0	9.0	22.0	65.0	4.0	66.38	27.95
	4	79.0	39.0	97.0	39.23	80.0	61.0	95.0	24.0	33.0	26.0	4.0	24.0	69.0	5.0	67.5	31.51
	6	87.0	39.0	96.0	43.08	82.0	64.0	90.0	28.0	34.0	32.0	5.0	20.0	66.0	6.0	67.13	33.56
	8	86.0	24.0	90.0	50.77	80.0	62.0	89.0	24.0	36.0	30.0	8.0	21.0	66.0	1.0	66.75	31.23
SEER	0	37.0	21.0	80.0	16.92	43.0	53.0	81.0	19.0	0.0	1.0	3.0	21.0	55.0	2.0	37.63	19.04
	2	73.0	20.0	87.0	41.54	81.0	52.0	92.0	30.0	31.0	23.0	6.0	20.0	64.0	0.0	58.88	27.26
	4	82.0	25.0	87.0	41.54	90.0	58.0	94.0	33.0	37.0	35.0	6.0	23.0	68.0	2.0	67.88	31.51
	6	70.0	17.0	98.0	40.77	85.0	50.0	92.0	28.0	33.0	27.0	5.0	23.0	65.0	0.0	65.88	27.12
	8	81.0	21.0	95.0	33.85	78.0	56.0	94.0	36.0	34.0	40.0	5.0	23.0	68.0	0.0	67.13	30.14




Figure 3: Accuracy of SEER and its ablated variants, showing the impact of each retrieval component.

contrast, SEER (w/o s_2) shows the most significant decline on hard questions. It suggests that the toolchain coverage score is crucial for complex tasks, where the model must navigate intricate dependencies and interactions among multiple tools.

In short, these results validate the effectiveness of our multidimensional retrieval scoring mechanism. The s_3 component plays a crucial role in simpler tasks that benefit from retrieved intent-aligned exemplars. In contrast, the s_2 contributes significantly to performance on complex tasks, where multi-step tool use and dependencies are common.

Demonstration Number. In this section, we analyze the impact of Top- k on the performance of SEER. We vary the few-shot number from 0 to 8 and evaluate the model’s performance on both easy and hard questions. From the Table 5, we can observe a trend of performance improvement with increasing demonstration numbers, followed by a decline after reaching a peak. Specifically, on the easy question set, the model achieves its

best performance at 4 demonstrations, with an accuracy of 67.88%. On the hard question set, the model performs best at 4 demonstrations, achieving an accuracy of 31.51%. We also test on query-based SEER, which shows a similar trend, with the best performance at $k = 4$ for easy questions and $k = 6$ for hard questions. This indicates that the model benefits from a moderate number of demonstrations, which provide sufficient context without overwhelming it with excessive information. Results on τ -bench across different k values exhibit a similar trend, as detailed in Appendix D.

5 Conclusion

In this paper, we proposed SEER, a self-guided approach for tool-augmented LLMs. SEER introduces a stepwise experience recall mechanism to retrieve relevant past experiences and guide multi-step tool usage. It continually updates its experience pool with successful trajectories, enabling iterative self-improvement during deployment. By conducting extensive experiments, we demonstrated that SEER significantly outperforms existing methods on both multi-step and multi-turn benchmarks. We also validated the self-improvement capability of SEER through intermediate batch accuracy evaluations. Additionally, we performed ablation studies to assess the contributions of SEER’s components. Overall, the experiments show some key findings: (1) SEER is effective in improving the performance of LLMs on complex function calling tasks; (2) SEER’s self-guided mechanism enables continual self-improvement; (3) The multidimensional retrieval strategy enhances the model’s ability across different task scenarios; (4) The number of demonstrations plays a crucial role, with a moderate number yielding the best results.

Limitations

While SEER demonstrates strong performance improvements in multi-step tool usage, several limitations remain. The diversity of the experience memory is inherently constrained by the capabilities of the underlying LLM. For complex or edge-case queries that require advanced reasoning beyond what can be addressed through in-context learning, SEER’s self-guided mechanism may encounter performance bottlenecks. SEER uses a fixed retrieval weighting scheme across all tasks, which may not be optimal for heterogeneous domains or task types. Dynamic adaptation of retrieval strategies—such as learning task-aware weighting or incorporating uncertainty estimates—could further enhance SEER’s generalization and robustness.

Acknowledgments

This work is supported by the National Key R&D Program of China (No.2022ZD0116405).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Fengshuo Bai, Mingzhi Wang, Zhaowei Zhang, Boyuan Chen, Yinda Xu, Ying Wen, and Yaodong Yang. 2024. Efficient model-agnostic alignment via bayesian persuasion. *arXiv preprint arXiv:2405.18718*.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Sijia Cui, Shuai Xu, Aiyao He, Yanna Wang, and Bo Xu. 2025. Empowering llms with parameterized skills for adversarial long-horizon planning. *arXiv preprint arXiv:2509.13127*.
- Carson Denison, Monte MacDiarmid, Fazl Barez, David Duvenaud, Shauna Kravec, Samuel Marks, Nicholas Schiefer, Ryan Soklaski, Alex Tamkin, Jared Kaplan, et al. 2024. Sycophancy to subterfuge: Investigating reward-tampering in large language models. *arXiv preprint arXiv:2406.10162*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *BERT: Pre-training of deep bidirectional transformers for language understanding*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hao Dong, Zihan Ding, and Shanghang Zhang. 2020. *Deep Reinforcement Learning: Fundamentals, Research and Applications*. Springer Nature.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. 2024. *A survey on in-context learning*. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128, Miami, Florida, USA. Association for Computational Linguistics.
- Anna Goldie, Azalia Mirhoseini, Hao Zhou, Irene Cai, and Christopher D Manning. 2025. Synthetic data generation & multi-step rl for reasoning & tool use. *arXiv preprint arXiv:2504.04736*.
- Xinyan Guan, Jiali Zeng, Fandong Meng, Chunlei Xin, Yaojie Lu, Hongyu Lin, Xianpei Han, Le Sun, and Jie Zhou. 2025. Deeprag: Thinking to retrieval step by step for large language models. *arXiv preprint arXiv:2502.01142*.
- Yilun Hao, Yongchao Chen, Yang Zhang, and Chuchu Fan. 2025. *Large language models can solve real-world planning rigorously with formal verification tools*. 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 3434–3483, Albuquerque, New Mexico. Association for Computational Linguistics.
- Aiyao He, Sijia Cui, Shuai Xu, Yanna Wang, and Bo Xu. 2025. Tums: Enhancing tool-use abilities of llms with multi-structure handlers. *arXiv preprint arXiv:2505.08402*.
- Cheng-Yu Hsieh, Si-An Chen, Chun-Liang Li, Yasuhisa Fujii, Alexander Ratner, Chen-Yu Lee, Ranjay Krishna, and Tomas Pfister. 2023. Tool documentation enables zero-shot tool-usage with large language models. *arXiv preprint arXiv:2308.00675*.
- Yue Huang, Jiawen Shi, Yuan Li, Chenrui Fan, Siyuan Wu, Qihui Zhang, Yixin Liu, Pan Zhou, Yao Wan, Neil Zhenqiang Gong, et al. 2023. Metatool benchmark for large language models: Deciding whether to use tools and which to use. *arXiv preprint arXiv:2310.03128*.
- Levente Kocsis and Csaba Szepesvári. 2006. Bandit based monte-carlo planning. In *European conference on machine learning*, pages 282–293. Springer.

- Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhat-tacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, et al. 2024. From generation to judgment: Opportunities and challenges of llm-as-a-judge. *arXiv preprint arXiv:2411.16594*.
- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023. Chameleon: Plug-and-play compositional reasoning with large language models. *Advances in Neural Information Processing Systems*, 36:43447–43478.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Bhargavi Paranjape, Scott Lundberg, Sameer Singh, Hannaneh Hajishirzi, Luke Zettlemoyer, and Marco Tulio Ribeiro. 2023. Art: Automatic multi-step reasoning and tool-use for large language models. *arXiv preprint arXiv:2303.09014*.
- Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Xuanhe Zhou, Yufei Huang, Chaojun Xiao, et al. 2024a. Tool learning with foundation models. *ACM Computing Surveys*, 57(4):1–40.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, dahai li, Zhiyuan Liu, and Maosong Sun. 2024b. [ToolLLM: Facilitating large language models to master 16000+ real-world APIs](#). In *The Twelfth International Conference on Learning Representations*.
- Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, and Ji-Rong Wen. 2025. Tool learning with large language models: A survey. *Frontiers of Computer Science*, 19(8):198343.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 technical report](#). Preprint, arXiv:2412.15115.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugging-gpt: Solving ai tasks with chatgpt and its friends in hugging face. *Advances in Neural Information Processing Systems*, 36:38154–38180.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652.
- Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Adrian Theuma and Ehsan Shareghi. 2024. [Equipping language models with tool use capability for tabular data analysis in finance](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 90–103, St. Julian’s, Malta. Association for Computational Linguistics.
- Ye Tian, Baolin Peng, Linfeng Song, Lifeng Jin, Dian Yu, Lei Han, Haitao Mi, and Dong Yu. 2024. Toward self-improvement of llms via imagination, searching, and criticizing. *Advances in Neural Information Processing Systems*, 37:52723–52748.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Pablo Villalobos, Jaime Sevilla, Lennart Heim, Tamay Besiroglu, Marius Hobbhahn, and Anson Ho. 2022. Will we run out of data? an analysis of the limits of scaling datasets in machine learning. *arXiv preprint arXiv:2211.04325*, 1.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560*.
- Zhiruo Wang, Zhoujun Cheng, Hao Zhu, Daniel Fried, and Graham Neubig. 2024. What are tools anyway? a survey from the language model perspective. *arXiv preprint arXiv:2403.15452*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muenighoff, Defu Lian, and Jian-Yun Nie. 2024. C-pack: Packed resources for general chinese embeddings. In

- Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*, pages 641–649.
- Qiancheng Xu, Yongqi Li, Heming Xia, and Wenjie Li. 2024. [Enhancing tool retrieval with iterative feedback from large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9609–9619, Miami, Florida, USA. Association for Computational Linguistics.
- Shuai Xu, Sijia Cui, Yanna Wang, Bo Xu, and Qi Wang. 2025. Strategy-augmented planning for large language models via opponent exploitation. *arXiv preprint arXiv:2505.08459*.
- Wei Yang, Jinwei Xiao, Hongming Zhang, Qingyang Zhang, Yanna Wang, and Bo Xu. 2025. Coarse-to-fine grounded memory for llm agent planning. *arXiv preprint arXiv:2508.15305*.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik R Narasimhan. 2025. [\$\tau\$ -bench: A benchmark for Tool-Agent-User interaction in real-world domains](#). In *The Thirteenth International Conference on Learning Representations*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.
- Yuanqing Yu, Zhefan Wang, Weizhi Ma, Zhicheng Guo, Jingtao Zhan, Shuai Wang, Chuhan Wu, Zhiqiang Guo, and Min Zhang. 2024. Steptool: A step-grained reinforcement learning framework for tool learning in llms. *arXiv preprint arXiv:2410.07745*.
- Hongming Zhang and Tianyang Yu. 2020a. Alphazero. In *Deep reinforcement learning: fundamentals, research and applications*, pages 391–415. Springer.
- Hongming Zhang and Tianyang Yu. 2020b. Taxonomy of reinforcement learning algorithms. In *Deep reinforcement learning: Fundamentals, research and applications*, pages 125–133. Springer.
- Zhaowei Zhang, Fengshuo Bai, Qizhi Chen, Chengdong Ma, Mingzhi Wang, Haoran Sun, Zilong Zheng, and Yaodong Yang. 2025. Amulet: Realignment during test time for personalized preference adaptation of llms. In *The Thirteenth International Conference on Learning Representations*.
- Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. 2024. Expel: Llm agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19632–19642.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International conference on machine learning*, pages 12697–12706. PMLR.
- Ruizhe Zhong, Xingbo Du, Shixiong Kai, Zhentao Tang, Siyuan Xu, Hui-Ling Zhen, Jianye Hao, Qiang Xu, Mingxuan Yuan, and Junchi Yan. 2023. Llm4eda: Emerging progress in large language models for electronic design automation. *arXiv preprint arXiv:2401.12224*.
- Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. 2023. [ToolQA: A dataset for LLM question answering with external tools](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Prompt 1: Prompt of Intent Recognizer

You are an expert in **intent recognition**. Given a user question (from 8 distinct intent categories), your goal is to extract key information from the question and determine the most likely intent.

You should think step by step and conclude your answer with an intent result in the format: [INTENT], where INTENT is the name of the predicted intent category.

Below are descriptions of all 8 intent categories:
{intent_categories}

Instructions: Your response must follow this format: Question: This is the question.
Answer: These are your thoughts. [INTENT]

Here are some examples:
{examples}

Now, please infer user intent:
Question: {question}

Figure 4: The illustration of the intent recognizer.

Appendix

A Prompt Template

A.1 Intent Inference Prompt

The prompt construction is shown in Figure 4.

A.2 Trajectory Evaluation Prompt

The prompt construction is shown in Figure 5.

B Baselines

We compare SEER with the following baselines, which are widely used in the field of tool-augmented LLMs. The details of these methods are as follows:

- **Direct, CoT (Wei et al., 2022):** Follow two baselines setting in (Zhuang et al., 2023), where the LLM is directly prompted with the user question and generate a response without awareness of tool invocation, to demonstrate the limitations of LLMs without tool assistance.
- **CoT-tool:** CoT with tool invocation. This method is similar to the original CoT but includes a tool interface for LLMs to invoke external tools.
- **Chameleon:** Chameleon (Lu et al., 2023) is a plug-and-play compositional reasoning framework where the LLM acts as a controller to plan and execute tool chain. Each tool operates as an independent module, allowing flexible combinations and extensions for complex tasks.
- **ReAct:** ReAct (Yao et al., 2023) enables LLMs to alternately generate reasoning traces and task-

Prompt 2: Prompt of Evaluator

Please act as an evaluator to determine whether the model’s response matches or includes the correct answer. I will provide both the correct answer and the model’s response.
Please reply with either [Match] or [No Match], and briefly explain the reasoning behind your judgment.

Examples:

Question: What time is the meeting?
Correct Answer: 3:00PM
Model Response: The meeting is scheduled for 15:00.
Evaluation: [Match]

Question: What is Mike’s total cost?
Correct Answer: 9
Model Response: The total cost of Mike is 9.001
Evaluation: [Match]

Question: When is he scheduled to attend the meeting?
Correct Answer: 01/12
Model Response: He will attend this meeting on the morning of January 12th.
Evaluation: [Match]

Question: What is the price?
Correct Answer: \$9374
Model Response: None
Evaluation: [No Match] (Model Response is None)

Now evaluate:

Question: question
Correct Answer: gt_answer
Model Response: llm_answer

Figure 5: The illustration of the evaluator.

specific actions, forming an iterative Observation-Thought-Action cycle. Compared to Chameleon, ReAct receives immediate feedback from tool executions, facilitating more adaptive decision-making.

- **TUMS:** TUMS (He et al., 2025) is a framework that enhances LLM’s tool-use abilities by introducing fine-grained, parameter-level processing, enabling more accurate and reliable tool execution.
- **ART:** ART (Paranjape et al., 2023) is a multi-step reasoning and tool-use framework that retrieves similar task trajectories to guide the LLM in generating intermediate steps and invoking functions.
- **ExpeL:** ExpeL (Zhao et al., 2024) is a self-improvement framework that enables LLMs to learn from past experiences and improve their performance over time.

In τ -bench, We initially assessed the performance of three close source models (claude-3-5-sonnet-20241022, gpt-4o-2024-11-20, and gpt-4o-mini-2024-07-18) and three open-source mod-

Table 6: Full main results on the ToolQA benchmark.

Method	Flights		Coffee		Yelp		Airbnb		Dblp		Scirex		Agenda		GSM8K	Avg	
	easy	hard	easy	hard	easy	hard	easy	hard	easy	hard	easy	hard	easy	hard	easy	easy	hard
Direct	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	50.0	6.3	0.0
CoT-noTool	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	63.0	7.9	0.0
Chameleon	37.0	2.0	61.0	2.3	60.0	13.0	12.0	4.0	25.0	17.0	6.0	19.0	57.0	0.0	17.0	34.5	8.2
CoT	37.0	22.0	79.0	13.1	43.0	52.0	79.0	19.0	0.0	1.0	3.0	22.0	54.0	0.0	2.0	37.1	18.4
ReAct	67.0	6.0	94.0	25.4	70.0	17.0	86.0	17.0	27.0	22.0	6.0	18.0	61.0	0.0	64.0	59.5	15.1
TUMS	64.0	9.0	93.0	22.3	73.0	15.0	91.0	11.0	34.0	30.0	4.0	14.0	59.0	0.0	72.0	61.3	14.5
ART	73.0	20.0	99.0	33.1	81.0	30.0	94.0	25.0	33.0	33.0	3.0	25.0	1.0	0.0	73.0	57.1	23.7
ExpeL	81.0	29.0	92.0	34.6	77.0	40.0	88.0	32.0	37.0	29.0	5.0	19.0	57.0	1.0	57.0	61.8	26.4
ExpeL (8-shot)	84.0	27.0	88.0	30.8	77.0	35.0	93.0	24.0	35.0	28.0	5.0	21.0	65.0	2.0	70.0	64.6	24.0
SEER (Ours)	82.0	25.0	87.0	41.5	90.0	58.0	94.0	33.0	37.0	35.0	6.0	23.0	68.0	2.0	79.0	67.9	31.1
SEER + Reflection	88.0	31.0	99.0	52.3	85.0	63.0	94.0	29.0	36.0	27.0	7.0	22.0	68.0	0.0	75.0	69.0	32.0
Ablation Study																	
SEER (query-based)	79.0	39.0	97.0	39.2	80.0	61.0	95.0	24.0	33.0	26.0	4.0	24.0	69.0	0.0	83.0	67.5	30.5
SEER (w/o s2)	83.0	23.0	96.0	38.5	58.0	35.0	92.0	34.0	36.0	27.0	7.0	20.0	62.0	1.0	80.0	64.3	25.5
SEER (w/o s3)	89.0	27.0	64.0	41.5	91.0	60.0	95.0	24.0	39.0	35.0	6.0	23.0	65.0	0.0	0.0	56.1	30.1




Figure 6: The accuracy of SEER with different few-shot numbers. The top plot shows query-based SEER, while the bottom plot shows SEER. The x-axis represents the number of few-shot demonstrations, and the y-axis represents the accuracy.

els (Qwen2.5-7B-Instruct, Qwen2.5-72B-Instruct, and DeepSeek-V3-0324) on the Tau-Bench benchmark. Subsequently, we evaluated the improvements achieved by applying the proposed SEER method to Qwen2.5 models on the same benchmark.

C Full Main Results

The full results on ToolQA are shown in Table 6.

D Demonstration Number

The full ablation study of the number of demonstrations is shown in Figure 6. The results show that the query-based SEER performs best with 4 demonstrations, achieving an accuracy of 67.5% on easy questions and 33.56% on hard questions. The performance declines when the number of demonstrations exceeds 4, indicating that too many examples can overwhelm the model and lead to confusion.

E Benchmark Details

ToolQA. The dataset is divided into two parts: simple and complex questions. The simple question set contains 55 templates, while the complex question set contains 62 templates. Each template is designed to cover a specific domain and includes a set of questions that can be answered using the tools provided in the benchmark. The dataset is designed to test the ability of LLMs to reason about and use external tools effectively. This large amount of data is mainly stored in the form of databases, graphs, and textual corpora, which greatly tests LLM’s understanding and flexible application ability of the given tools.

τ -bench. The dataset is divided into two parts: τ -retail and τ -airline. τ -retail contains 115 questions, while τ -airline contains 50 questions. In contrast to one-turn multi-step ToolQA, τ -bench focuses on evaluating LLMs in real-world tasks through multi-turn user-agent interactions.