

# PEToolLLM: Towards Personalized Tool Learning in Large Language Models

Qiancheng Xu<sup>1</sup>, Yongqi Li<sup>1†</sup>, Heming Xia<sup>1</sup>, Fan Liu<sup>2</sup>, Min Yang<sup>3</sup>, Wenjie Li<sup>1</sup>

<sup>1</sup> The Hong Kong Polytechnic University    <sup>2</sup> National University of Singapore

<sup>3</sup> Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences

{qiancheng.xu, he-ming.xia}@connect.polyu.hk

liyongqi@gmail.com    cswjli@comp.polyu.edu.hk

## Abstract

Tool learning has emerged as a promising direction by extending Large Language Models' (LLMs) capabilities with external tools. Existing tool learning studies primarily focus on the general-purpose tool-use capability, which addresses explicit requirements in user instructions. However, they overlook the importance of personalized tool-use capability, leading to an inability to meet personalized user needs. To address the limitation, we first formulate the task of personalized tool learning, which integrates user's interaction history towards personalized tool usage. To fill the gap of missing benchmarks, we construct PEToolBench, featuring diverse user preferences reflected in interaction history under three distinct personalized settings, and encompassing a wide range of tool-use scenarios. Moreover, we propose a framework PEToolLLaMA to adapt LLMs to the personalized tool learning task, which is trained through supervised fine-tuning and direct preference optimization. Extensive experiments on PEToolBench demonstrate the superiority of PEToolLLaMA over existing LLMs. We release our code and data at <https://github.com/travis-xu/PEToolBench>.

## 1 Introduction

Large Language Models (LLMs) possess extensive knowledge and powerful instruction following abilities, making them effective AI assistants for tasks such as text rewriting, question answering, and code writing (Zhao et al., 2023). However, they often struggle to address user needs in scenarios such as checking weather and booking flights. To address this, tool learning (Qin et al., 2024a; Qu et al., 2025b) has emerged as a promising solution by allowing LLMs to use external tools, such as real-time weather APIs and booking systems. In this way, tool learning has extended LLMs' capabilities

to tackle more complex tasks, enabling them to fulfill a wide range of user needs.

Current tool learning procedure typically begins with a user instruction, and then LLMs are required to use tools with appropriate functionalities for satisfying users' needs. Existing tool learning methods can be categorized into in-context learning (Patil et al., 2024; Qu et al., 2025a) and fine-tuning approaches (Schick et al., 2023; Wang et al., 2025b). The former approach allows LLMs to use tools by directly providing tool documentation in input but the performance is constrained by the input length. The latter approach trains LLMs to internalize tool knowledge but struggles with tool generalization.

Despite the advancement, existing tool learning methods primarily focus on the general-purpose tool-use capability but overlook the critical role of personalization. In tool learning, more personalized user needs are expected to be derived from the user's previous tool usage history as a supplement to user instructions, which can help LLMs provide more customized tool-usage assistance to enhance the user experience. As illustrated in Figure 1, personalized tool learning is non-trivial due to the following aspects. 1) **Implicit user preferences**. User preferences for tool usage are often implicitly conveyed through the user's history rather than explicitly stated in user instructions, making them difficult to understand. For instance, when a user requests a search for articles, their preference for academic-related content needs to be inferred from previous interactions with academic tools like Google Scholar. 2) **Non-functional tool attributes**. Since many tools have the same functionalities, user preferences cannot be effectively distinguished based solely on tool functionalities. This underscores the need to consider non-functional tool attributes, such as usability, integrability, and accessibility, which can better reflect user preferences. As shown in Figure 1, Google Search can be

<sup>†</sup>Corresponding author.

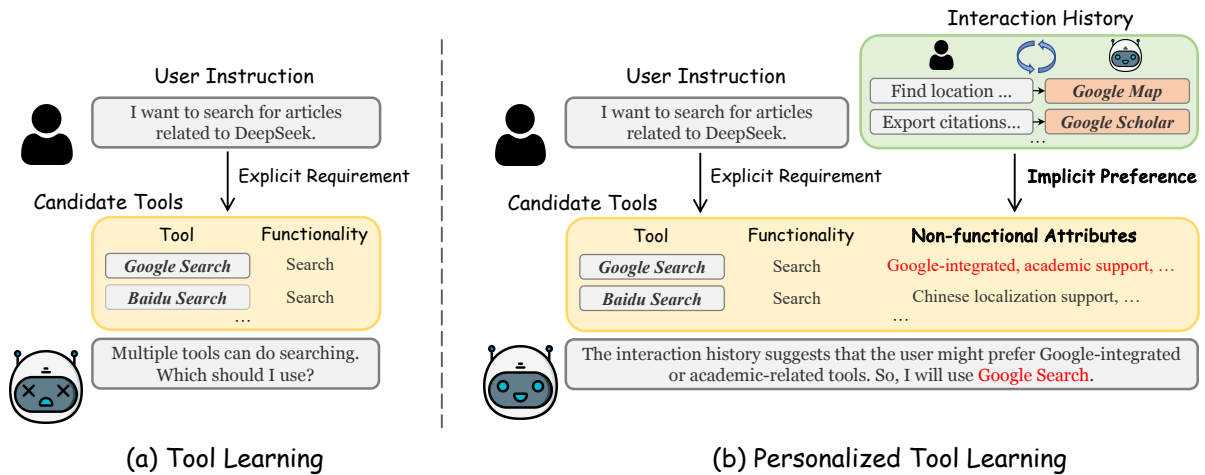


Figure 1: Comparison between (a) tool learning and (b) personalized tool learning. Personalized tool learning facilitates implicit preference comprehension and customized tool usage for individual users.

distinguished from other search tools by its integration into Google’s ecosystem with Google Scholar, making it more suitable for users with academic needs. To address the above issues, we formulate the task of personalized tool learning in LLMs, aiming at personalized tool usage for individual users. Formally, given user instructions along with user’s interaction history, LLMs are required to answer user instructions with tools by considering both explicit user requirements from instructions and implicit user preferences behind interaction history.

Since there is no benchmark for this task currently, we fill this gap by introducing the first personalized tool learning benchmark (PEToolBench). Specifically, the benchmark is created through three following steps. 1) Tool Preparation. We collect a bunch of high-quality tools from RapidAPI and then leverage LLM to understand the functionality and non-functional attributes of each tool. 2) Preference Construction. Among same-functionality tools, we construct the user’s tool preferences by assigning tools with distinct non-functional attributes to different users. 3) Data Creation. Based on user’s tool preferences, we synthesize their interaction history into a sequence of tool-use interactions, each consisting of a user instruction and an LLM’s tool call. We design three personalized tool-usage settings by generating the interaction history in three types, i.e., preferred-only, rating-integrated, and chronological. And then we use tools not included in the interaction history to synthesize user instructions. After rigorous filtering, we obtain 12,000 user instructions with interaction histories reflecting diverse user preferences and cover a wide

range of tool-use scenarios by encompassing 7454 tools across 46 categories.

Based on the PEToolBench dataset, we propose the personalized tool learning framework (PEToolLLaMA) to equip LLMs with personalized tool-use capability. The training process consists of two stages: 1) Supervised fine-tuning (SFT) stage, which equips LLM with foundational tool-use capability to address user needs; 2) Direct preference optimization (DPO) stage, which samples the user’s preferred and non-preferred tool calls for pair-wise optimization to better align with user preferences. We evaluate 11 distinct open-source and closed-source LLMs on PEToolBench. Experimental results demonstrate that our PEToolLLaMA significantly outperforms the best-performing LLM across all settings.

In summary, our contributions are as follows.

- We are the first to formulate the task of personalized tool learning in LLMs, which incorporates user’s interaction history to achieve personalized tool-usage assistance.
- We construct the first benchmark for personalized tool learning in LLMs, PEToolBench, featuring user instructions integrated with interaction history reflecting diverse user preferences and encompassing various tools.
- We propose a novel personalized tool learning framework PEToolLLaMA. Extensive experiments demonstrate that PEToolLLaMA significantly surpass the best-performing LLM by more than 50%, exhibiting exceptional personalized tool-use capabilities.

## 2 Related Work

### 2.1 Tool Learning in LLMs

Tool learning aims at extending the capabilities of LLMs by equipping them with external tools to solve tasks like weather inquiry, car navigation, and restaurant reservation. Existing benchmarks primarily focus on evaluating the tool learning proficiency of LLMs in addressing user instructions, from aspects such as tool-use awareness (Huang et al., 2024b), tool selection accuracy (Xu et al., 2024; Wang et al., 2024b, 2025a), tool planning ability (Wang et al., 2024a; Shen et al., 2024; Liu et al., 2025a), and complex workflow creation (SHEN et al., 2025; Qiao et al., 2025; Fan et al., 2025). To improve tool-use capabilities, various strategies have been introduced, including in-context learning which enables LLMs to use tools via documentation (Yuan et al., 2025b; Shi et al., 2024; Liu et al., 2025b), and fine-tuning which trains LLMs on specialized tool-use datasets (Tang et al., 2023; Chen et al., 2024, 2025). However, prior studies neglect the crucial role of personalized tool usage in LLMs. This paper addresses this gap by introducing personalized tool learning, developing a comprehensive benchmark for evaluation, and proposing an optimization strategy to enhance personalized tool-use capabilities in LLMs.

### 2.2 Personalization in LLMs

The goal of personalization in LLMs is to leverage personal user data, such as historical behaviors and background information, to generate outputs that better align with the user preferences (Tseng et al., 2024). Approaches such as fine-tuning (Cai et al., 2025) and prompt engineering (Yuan et al., 2025a) have been explored to adapt LLMs to individual or domain-specific tasks. These approaches have been applied across various fields, including recommendation systems (Lyu et al., 2024), search engines (Zhou et al., 2024), education (Liu et al., 2024a), and dialogue generation (Wang et al., 2023). However, previous research has not investigated LLMs’ personalization in the area of tool learning. In this work, we bridge this gap by incorporating user’s interaction history to assess and enhance the LLMs’ capability in providing personalized tool-usage assistance for specific users.

## 3 Task and Benchmark

### 3.1 Task Formulation

**Tool Learning** Given an instruction  $q_u$  of the user  $u$ , tool learning aims to generate an appropriate tool call, including the selected tool and its corresponding parameters, from a set of candidate tools. Formally, let the candidate tool set be  $\mathcal{T} = \{d(t_1), d(t_2), \dots, d(t_N)\}$ , where  $d(t_i)$  represents the documentation of tool  $t_i$  and  $N$  is the total number of candidate tools. The LLM is then tasked with generating a tool call  $c = (t, p)$ , where  $t \in \mathcal{T}$  and  $p$  denotes its parameters:

$$(t, p) = LLM(q_u, \mathcal{T}). \quad (1)$$

**Personalized Tool Learning** In personalized tool learning, we incorporate the users’ interaction history alongside their instructions, enabling the LLM to generate tool calls that satisfy both the users’ explicit requirements and implicit preferences. For a user  $u$ , we define the interaction history as  $\mathcal{H}_u = \{h_u^1, h_u^2, \dots, h_u^M\}$ , where each  $h_u^i$  consists of a past user instruction  $q_u^i$  and the corresponding tool call  $c_u^i = (t_u^i, p_u^i)$ , with  $t_u^i$  representing the selected tool and  $p_u^i$  denoting its associated parameters. Let  $c_u = (t_u, p_u)$  represent the personalized tool call for user  $u$ , the personalized tool learning task can then be formulated as:

$$(t_u, p_u) = LLM(q_u, \mathcal{T}, \mathcal{H}_u), \quad (2)$$

### 3.2 Benchmark Construction

Due to the lack of real user interaction histories on tool-usage, we adopt a tool-driven approach to simulate interactions based on pre-constructed user’s tool preferences. The whole process for constructing PEToolBench, illustrated in Figure 2, consists of three steps: tool preparation, preference construction, and data creation.

#### 3.2.1 Tool Preparation

**Tool Collection** Following ToolBench (Qin et al., 2024b), we adopt the tools from RapidAPI for our benchmark, since it offers a large-scale and diverse collection of real-world tools that can potentially address a wide range of user needs. To ensure the quality of the collected tools, we perform strict filtering by removing: 1) outdated tools, which are marked as deprecated in RapidAPI; 2) tools with insufficient information, such as inadequate or missing tool documentation; and 3) duplicate tools, which have repeated tool names, descriptions, or category names.

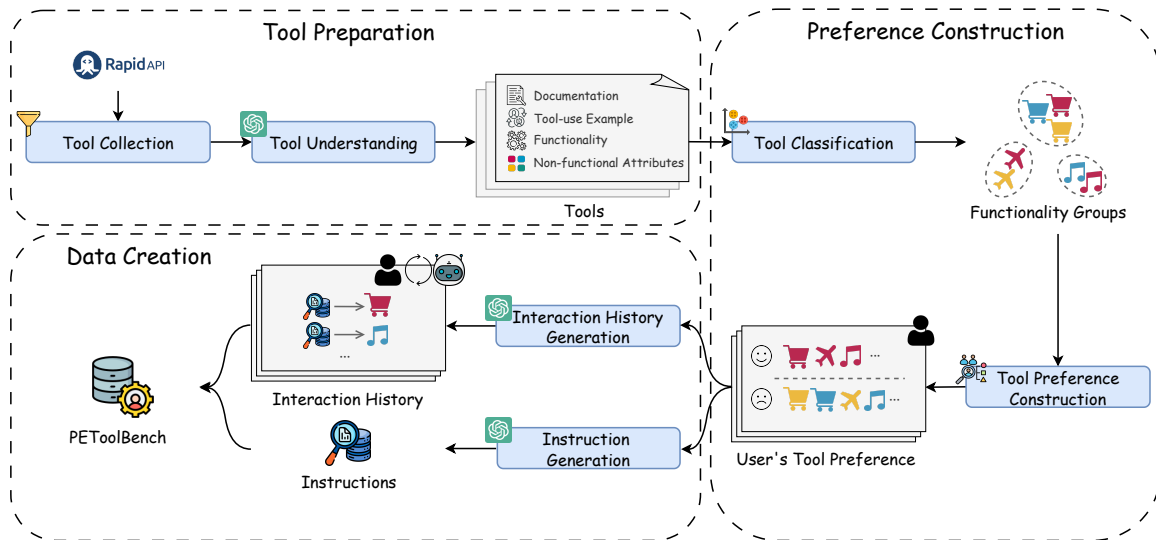


Figure 2: Illustration of the process for constructing our PEToolBench.

**Tool Understanding** Since tool documentation often contains redundant and irrelevant information, directly extracting tool attributes from it can be challenging. To address this, we first provide the documentation of each tool to LLM and prompt it to generate a tool-use example, including a simulated user instruction and parameters for calling the tool. Next, based on the tool documentation and tool-use example, the LLM is instructed to generate descriptions of the tool’s functionality and non-functional attributes separately. Besides, we include demonstrations in the prompt to help the LLM distinguish between these two attribute types. By leveraging specific tool-use examples and demonstrations, the LLM can develop a more comprehensive understanding of each tool’s functionality and non-functional characteristics.

### 3.2.2 Preference Construction

**Tool Classification** To identify potential tool-usage scenarios for users, we classify tools with the same functionalities into groups. Specifically, we first employ the Ada Embedding model<sup>1</sup> to compute embeddings for the functionality descriptions of all tools. Then, we apply the DBSCAN algorithm (Schubert et al., 2017) to cluster these tools into multiple groups based on the similarity of their embeddings. Within each group, the tools share the same functionality and can be applied to a specific tool-usage scenario. To further ensure that tools within each group exhibit uniform functionality, we conduct rigorous filtering and only retain

<sup>1</sup><https://platform.openai.com/docs/guides/embeddings/embedding-models>.

groups where tools 1) have the same input-output formats (i.e., required/optional parameters and response schema) and 2) belong to the same category (e.g., sports, music, finance).

**Tool Preference Construction** We leverage non-functional tool attributes to construct the user’s tool preference. First, we randomly sample a functionality group for a user, representing a potential tool-usage scenario for interaction. Within this group, we choose a tool with specific non-functional attributes as the user’s preferred tool, while the others are considered non-preferred. Using the preferred tool as a reference, we retrieve the top-5 tools with the most similar non-functional attributes. Similarity is computed based on the embeddings of the tools’ non-functional descriptions, which are generated in the Tool Understanding phase. Through multiple iterations of sampling and retrieving, we obtain a diverse set of preferred and non-preferred tools that represent user preferences. After each iteration, we check for functionality overlap between newly retrieved tools and previously selected ones. If an overlap is detected, the tools are discarded, and the sampling process is restarted. This ensures that each tool-usage scenario is associated with only one preferred tool per user. By following this approach, we construct diverse tool sets that align with different user preferences.

### 3.2.3 Data Creation

**Interaction History Generation** Based on tool preference, we leverage the LLM to construct the user’s interaction history. Specifically, for each user, we provide LLM with the user’s preferred



and non-preferred tools, including tool attributes and tool-use examples generated in the Tool Understanding phase. The LLM will generate a sequence of simulated user-LLM interactions, each consisting of a user instruction and an LLM’s tool call, as the user’s interaction history.

We design three personalized tool-use settings by generating the interaction history in three types (illustrated in Figure 7): 1) *preferred-only* history, where the tools involved in the interactions are all preferred by the user; 2) *rating-integrated* history, including both the user’s preferred and non-preferred tools, with a user’s binary rating for each tool-usage interaction representing the user preference, i.e., “liked” if the tool aligns with the user preferences, and “disliked” otherwise. 3) *chronological* history, which organizes interactions in time order to reflect changes in user preferences over time, i.e., the more recent tool-usage interactions are more preferred by the user, while earlier interactions are less preferred. In this way, we can present different forms of user preferences.

**Instruction Generation** Next, we use LLM to generate user instructions based on the user’s preferred tools that are not included in the user’s interaction history. We instruct the LLM to avoid directly generating the name of the tool in the instruction, ensuring that the user preference for the tool can only be inferred from the user’s interaction history. Each user instruction is combined with the user’s interaction history into a data instance.

Finally, we obtain 12,000 data instances encompassing 7,454 tools across 46 categories. We split all data into two parts: a training set comprising 9,000 instances for three personalized settings and a test set containing the rest instances.

### 3.3 Benchmark Analysis

We present the statistical information of our PEToolBench in Table 3, including the statistics of data instances in three settings and under varying interaction history lengths. We also present the distribution of tool categories in Figure 4. Statistical information demonstrates the diversity and complexity of our dataset.

### 3.4 Evaluation Metrics

Given the user’s instruction and interaction history, LLM is expected to select the appropriate tool from a candidate tool set that can not only fulfill the user’s request but also align with the user’s




Figure 3: Statistics of data instances in three personalized settings (in the left figure) and distributions of interaction history length (in the right figure).

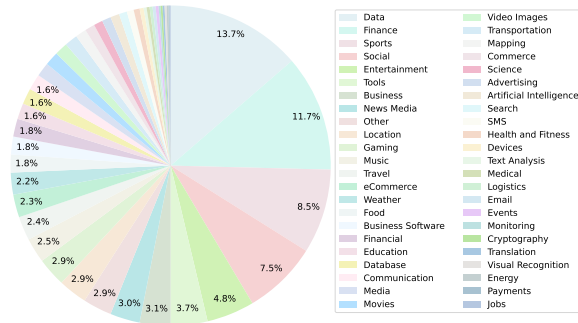


Figure 4: Distributions of tool categories.

preferences, and then call the selected tool with corresponding parameters. Therefore, we define two metrics as follows.

- **Tool Accuracy (Tool Acc):** The metric assesses the ability of LLM to select the appropriate tool to call. If the tool is correctly selected, the score is 1; otherwise, the score is 0.
- **Parameter Accuracy (Param Acc):** The metric assesses the ability of LLM to generate correct parameters for the tool call. If the input parameters are correctly generated, the score is 1; otherwise, the score is 0.

## 4 Method: PEToolLLaMA

To equip LLM with personalized tool-use capability, we conduct a two-stage training process: 1) personalized SFT, where LLM is fine-tuned on PEToolBench to acquire fundamental proficiency in personalized tool usage, and 2) personalized DPO, where LLM is optimized on a preference dataset for better alignment with user preferences.

**Personalized SFT.** The first stage in our approach is Supervised Fine-Tuning (SFT), where we directly fine-tune LLM on the training set of PEToolBench. Given the user’s instruction  $q_u$ , interaction history  $\mathcal{H}_u$ , and the candidate tool set  $\mathcal{T}$

Methods	PREFERRED-ONLY		RATING-INTEGRATED		CHRONOLOGICAL		ALL	
	Tool Acc	Param Acc	Tool Acc	Param Acc	Tool Acc	Param Acc	Tool Acc	Param Acc
Vicuna-7B	25.50	44.80	10.80	57.40	12.70	56.00	16.33	52.73
Mistral-7B	30.30	55.70	15.40	63.20	14.10	64.90	19.93	61.27
xLAM-7B-r	31.80	63.40	15.50	76.10	16.40	75.20	21.23	71.57
Hammer2.1-7B	30.20	61.60	28.30	65.20	19.80	69.70	26.10	65.60
Qwen2.5-7B	40.40	63.80	24.80	66.50	24.80	70.20	30.00	66.83
LLaMA3-8B	48.10	71.10	26.90	77.70	26.60	78.10	33.87	75.63
Qwen2.5-Max	44.50	73.40	37.30	81.50	28.20	85.20	36.67	80.03
GPT-4-turbo	45.10	73.40	41.80	77.00	28.50	81.50	38.47	77.30
Deepseek-R1	48.20	78.80	46.40	82.50	26.40	86.10	40.33	82.47
GPT-4o-mini	51.80	72.90	38.40	77.70	31.20	80.50	40.47	77.03
GPT-4o	53.70	77.60	45.70	79.60	33.60	81.80	44.33	79.67
<b>PEToolLLaMA</b>	<b>74.30</b>	<b>87.90</b>	<b>78.40</b>	<b>89.70</b>	<b>80.80</b>	<b>91.30</b>	<b>77.83</b>	<b>89.63</b>
% improve	38.36%	13.27%	71.55%	12.69%	140.5%	11.61%	75.57%	12.50%
w/o DPO	71.50	82.10	74.20	86.90	77.30	90.40	74.33	86.47
w/o SFT	53.20	61.80	55.30	62.40	51.40	61.10	53.30	61.77

Table 1: Evaluation results of different LLMs on PEToolBench in terms of tool and parameter accuracy under settings including preferred-only, rating-integrated, chronological, and the whole data (*All*). Bold highlights the best score among all LLMs and % improve represents the relative improvement achieved by our method over the previously best-performing LLM.

as inputs, LLM is trained to generate the ground truth tool call  $c$ .  $\mathcal{H}_u$  uniformly covers all three types of user interactions to capture diverse user preferences. In this way, LLM can obtain basic personalized tool-usage experiences by understanding both the user needs and preferences.

**Personalized DPO.** In the second stage, we further enhance the LLM’s performance through direct preference optimization (DPO) (Rafailov et al., 2023). Our goal is to guide the LLM to call the user’s preferred tools instead of non-preferred ones. Specifically, for each user instruction  $q_u$ , we collect multiple tool calls generated by LLM after the SFT stage. Then we select the user’s preferred and non-preferred tool calls  $c_w$  and  $c_l$  based on the user’s tool preference constructed in PEToolBench.  $c_w$  and  $c_l$  will be used to construct the preference dataset  $\mathcal{D}_{\text{DPO}} = \{(x, c_w, c_l)\}$ , where  $x$  denotes the input, including the user instruction  $q_u$ , interaction history  $\mathcal{H}_u$ , and the candidate tool set  $\mathcal{T}$ . We then apply DPO to optimize the LLM by guiding it to generate the desired tool call  $c_w$  while avoid generating  $c_l$ . The loss function can be defined as:

$$\mathcal{L} = -\mathbb{E} \left[ \log \sigma \left( \beta \log \frac{\pi_{\theta}(c_w | x)}{\pi_{\text{ref}}(c_w | x)} - \beta \log \frac{\pi_{\theta}(c_l | x)}{\pi_{\text{ref}}(c_l | x)} \right) \right], \quad (3)$$

where  $\sigma$  is the logistic function and  $\beta$  is a weighting parameter that controls the deviation of the policy model  $\pi_{\theta}$  (i.e., the LLM to be optimized) from the reference model  $\pi_{\text{ref}}$  (i.e., the LLM after SFT stage).

In this way, LLM can focus on generating tool calls that are more aligned with user preferences.

## 5 Experiments

### 5.1 Setup

**Baselines.** We adopt multiple LLMs from both closed-source and open-source models to ensure a comprehensive evaluation. For closed-source LLMs, we select 5 representative models: GPT-4o, GPT-4o-mini, GPT-4-turbo (Achiam et al., 2023), Qwen-2.5-max (Yang et al., 2024) and Deepseek-R1 (Guo et al., 2025). For open-source LLMs, we include a wide spectrum of models, i.e., LLaMA-3.1-8B (Dubey et al., 2024), QWen-2.5-7B (Yang et al., 2024), Vicuna-7B-v1.5 (Chiang et al., 2023) and Mistral-7B-v0.3 (Jiang et al., 2023). We also include 2 tool-specific LLMs Hammer-2.1-7B (Lin et al., 2025) and xLAM-7B-r (Zhang et al., 2025).

**Implementation details.** In PEToolBench construction, we employ gpt-4o-mini for tool understanding and generation of user instructions and interaction history. The candidate tool set consists of three parts: the ground-truth tool along with all other tools sharing the same functionality, five tools retrieved using ToolRetriever (Qin et al., 2024b), and the remaining tools randomly sampled. We train PEToolLLaMA based on LLaMA-Factory (Zheng et al., 2024). More details on training and evaluation are illustrated in Appendix B.

## 5.2 Main Results

The detailed experimental results are shown in Table 1. From the results, we can obtain the following key findings. 1) It can be observed that the performance of LLMs is generally unsatisfactory, particularly in tool accuracy with the majority failing to exceed 50%. This indicates that current LLMs are severely limited in personalized tool-use capabilities. Additionally, the lower tool accuracy compared to parameter accuracy further suggests that personalized tool selection is more challenging than parameter configuration. This is because LLMs must account for both implicit user preferences and explicit user requirements when determining which tool to use. 2) Most LLMs perform worse in the rating-integrated and chronological settings. This is likely due to the inclusion of non-preferred interactions in the interaction history, which confuses LLMs and hinders their ability to accurately recognize user preferences. Notably, the chronological setting yields the lowest scores, suggesting that capturing evolving user preferences over time is even more challenging than interpreting explicit user ratings. 3) Our proposed PETool-Bench significantly outperforms all closed-source and open-source LLMs, demonstrating both effectiveness and robustness. It maintains strong performance, even in the two more challenging settings, by enabling the LLM to better understand diverse manifestations of user preferences and facilitate personalized tool usage.

## 5.3 Ablation Study

We conduct ablation studies to investigate the efficacy of the two-stage training process in our PEToolLLaMA. First, we remove the second training stage (i.e., personalized DPO) to assess its contribution. Then, we examine the impact of the SFT stage by directly conducting DPO training on the initial LLaMA3-8B model. Table 1 reports the performance on the test set of PEToolBench in all three settings. The results indicate that the SFT stage is crucial for personalized tool learning performance, as it endows the model with fundamental tool usage and personalization capabilities. Removing the DPO stage results in a slight performance drop, suggesting that it can further refine the tool usage alignment with user preferences.




Figure 5: Performance comparison in tool accuracy on the test data when provided with and without interaction history in three settings.

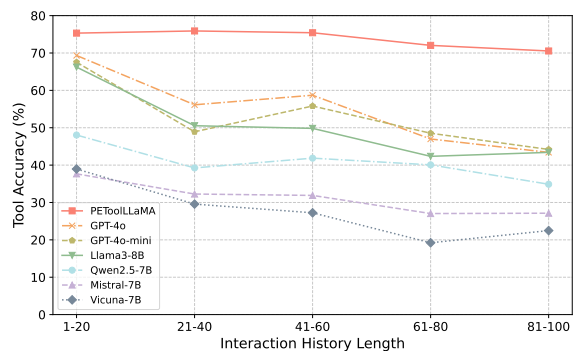


Figure 6: Performance comparison in tool accuracy on interaction history length in the preferred-only setting.

## 5.4 In-depth Analysis

**Analysis on the impact of interaction history.** To investigate the impact of interaction history on LLM performance, we remove the interaction history from the inputs and provide only the user instructions with candidate tools set to conduct our experiments. The results are presented in Figure 5. From the results, we can observe that both closed-source and open-source LLMs experience varying degrees of performance degradation without interaction history, compared to when provided with preferred-only history. This suggests that interaction history only containing the user’s preferred tools can help the LLM effectively infer user preferences. On the other hand, we find that LLMs perform worse with chronological history. This indicates that including both preferred and non-preferred tools can interfere with the LLM’s understanding of user preferences, thus hindering its personalization capabilities.

In contrast, our PEToolLLaMA consistently improves performance across all three types of interaction history compared to the no-history setting.

Models	PREFERRED-ONLY						CHRONOLOGICAL					
	IF	TH	TFM	TPM	PNM	PVM	IF	TH	TFM	TPM	PNM	PVM
Qwen2.5-7B	10.9	3.6	19.6	25.5	10.4	14.9	11.2	2.3	7.1	54.6	5.4	13.2
LLaMA3-8B	2.5	5.3	19.3	24.8	11.4	15.0	2.5	3.7	2.8	64.4	6.8	12.6
GPT-4o-mini	<b>0.1</b>	3.5	20.1	24.4	10.6	16.4	<b>0.0</b>	1.5	6.9	60.2	6.2	13.3
GPT-4o	0.5	<b>1.5</b>	20.0	30.3	7.9	14.0	1.3	<b>1.1</b>	6.6	57.3	<b>4.7</b>	12.2
<b>PEToolLLaMA</b>	0.6	3.4	<b>9.3</b>	<b>12.0</b>	<b>7.6</b>	<b>3.9</b>	0.5	1.8	<b>6.5</b>	<b>10.4</b>	5.1	<b>3.1</b>

Table 2: The percentage (%) of different error types in LLMs on the test set of PEToolBench under preferred-only and chronological settings. IF, TH, TFM, TPM, PNM, PVM stand for Invalid Format, Tool Hallucination, Tool Functionality Mismatch, Tool Preference Mismatch, Parameter Name Mismatch and Parameter Value Mismatch errors, respectively.

This demonstrates that our method enables LLM to effectively recognize different forms of user preferences from the interaction history. Specifically, the greater performance gains in the rating-integrated and chronological settings compared to the preferred-only setting. This stems from the inclusion of non-preferred tool interactions, which helps PEToolLLaMA differentiate between user’s preferred and non-preferred tool and thereby more effectively capture user preferences. Besides, further improvement is seen in the chronological setting over the rating-integrated one. This may be related to the well-known “lost in the middle” (Liu et al., 2024b) phenomenon in LLMs, where performance degrades when relevant information is located in the middle of long contexts. In the chronological setting, non-preferred tool interactions always appear at the beginning of the history rather than in the middle, potentially making it easier for the model to capture useful information from them.

**Analysis on interaction history length.** To evaluate the performance of LLMs under varying interaction history lengths, we break down the tool accuracy scores of LLMs based on the number of interactions in the history under the preferred-only setting. As shown in Figure 6, the performance of both closed-source and open-source LLMs deteriorates as interaction history length increases. This is because a longer interaction history makes it more challenging for the LLM to identify the historical preferences relevant to identify relevant historical preferences in relation to the user’s current context. In contrast, our PEToolLLaMA significantly outperforms all LLMs and maintains strong, consistent performance even as interaction history grows. This demonstrates that our method enables LLMs to effectively extract and utilize user preferences from complex historical data.

## 5.5 Error Analysis

We further conduct an error analysis to investigate the issues leading to incorrect tool calls in two personalized settings. We categorize the errors into six types: 1) Invalid Format. The tool call generated by the LLMs does not follow the expected JSON format. 2) Tool Hallucination. The LLM selects a tool that does not exist in the given candidate tool set, which is a common hallucination issue in LLMs (Huang et al., 2024a). 3) Tool Functionality Mismatch. The selected tool lacks the necessary functionality to fulfill the user’s requirements. 4) Tool Preference Mismatch. The selected tool has the correct functionality but is not preferred by the user. 5) Parameter Name Mismatch. The tool call contains missing or incorrect parameter names. 6) Parameter Value Mismatch. The parameter names are correctly generated, but the parameter values do not match the ground truth.

From the results in Table 2, we observe that most LLMs perform worst in Tool Preference Mismatch, particularly in the chronological setting, where the error rate exceeds 50%. This suggests that identifying user preferences from the interaction history is highly challenging, especially when preferences change over time, leading to significant model misinterpretation. In contrast, our PEToolLLaMA significantly reduces the error rate in Tool Preference Mismatch, demonstrating its effectiveness in capturing implicit user preferences. Additionally, the reduction in Tool Functionality Mismatch and Parameter Value Mismatch errors suggests that our method enhances LLMs’ fundamental tool-usage ability, improving their handling of explicit user requirements. Furthermore, PEToolLLaMA achieves low error rates in Invalid Format and Tool Hallucination, comparable to closed-source LLMs, showing its strong instruction-following capabilities.



## 6 Conclusion and Future Work

In this paper, we advanced general-purpose tool-use LLMs into personalized tool-use LLMs, aiming to provide users with customized tool-usage assistance. We formulate the task of personalized tool learning and identify the goal of leveraging user’s interaction history to achieve implicit preference understanding and personalized tool calling. For training and evaluation, we construct the first PEToolBench benchmark, featuring diverse users’ interaction history in three types. We also propose a novel personalized framework PEToolLLaMA conducted under a two-stage training process to endow LLMs with personalized tool-use capabilities. Extensive experiments on PEToolBench demonstrate that PEToolLLaMA consistently surpasses existing baselines, effectively meeting user requirements and preferences. We believe that the task, benchmark, and framework for personalized tool learning will broaden the research scope, introduce new challenges and inspire novel methods.

In the future, we aim to enhance this work from the following dimensions. 1) We plan to explore more heterogeneous personal user data beyond interaction history, such as user profiles or personas. This will allow us to reflect user preferences from multiple dimensions, providing a more comprehensive evaluation on the personalized tool-use capabilities of LLMs. 2) Currently, our work is limited to tool-usage scenarios involving a single tool. In the future, we intend to expand to more complex personalized tool-usage, such as multi-tool scenarios. These scenarios will require LLMs to perform personalized tool planning and engage in multi-round tool calling to address user needs effectively.

### Limitations

1) Due to the lack of real user interaction histories on tool usage, we utilize LLM to synthesize such data. However, this approach may compromise the authenticity and reliability of the data, which is a common challenge in data synthesis methods. To mitigate this issue, we incorporate pre-constructed user preference information into the data generation process. This strategy helps guide LLM in generating contextually relevant outputs, thereby improving the quality and consistency of the synthesized data. 2) In real-world scenarios, tools have multiple dimensions of attributes. However, due to the limited information contained in current tool documentation, it is difficult to fully identify and

exploit all possible tool attributes. Fortunately, the attributes we have obtained are sufficient to differentiate between tools, enabling us to effectively construct user preferences.

### Ethics Statement

The dataset used in our work is derived from publicly available sources and generated through interactions with LLMs in English. Since the user interaction histories in our study are entirely simulated, user privacy is fully protected, and no real personal information is included in the dataset. Furthermore, all scientific artifacts used in this research are publicly accessible for academic purposes under permissive licenses, and their use in this paper complies with their intended purposes. Given these considerations, we believe our research adheres to the ethical standards of the conference.

### Acknowledgments

The work described in this paper was supported by Research Grants Council of Hong Kong (15207821, 15207122, 15209724), and PolyU internal grants (ZVQ0).

### References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Hongru Cai, Yongqi Li, Wenjie Wang, Fengbin ZHU, Xiaoyu Shen, Wenjie Li, and Tat-Seng Chua. 2025. Large language models empowered personalized web agents. In *THE WEB CONFERENCE 2025*.
- Guoxin Chen, Zhong Zhang, Xin Cong, Fangda Guo, Yesai Wu, Yankai Lin, Wenzheng Feng, and Yasheng Wang. 2025. Learning evolving tools for large language models. In *The Thirteenth International Conference on Learning Representations*.
- Sijia Chen, Yibo Wang, Yi-Feng Wu, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and Lijun Zhang. 2024. Advancing tool-augmented large language models: Integrating insights from errors in inference trees. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6.

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Shengda Fan, Xin Cong, Yuepeng Fu, Zhong Zhang, Shuyan Zhang, Yuanwei Liu, Yesai Wu, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2025. WorkflowLLM: Enhancing workflow orchestration capability of large language models. In *The Thirteenth International Conference on Learning Representations*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2024a. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*.
- Yue Huang, Jiawen Shi, Yuan Li, Chenrui Fan, Siyuan Wu, Qihui Zhang, Yixin Liu, Pan Zhou, Yao Wan, Neil Zhenqiang Gong, and Lichao Sun. 2024b. Meta-tool benchmark for large language models: Deciding whether to use tools and which to use. In *The Twelfth International Conference on Learning Representations*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Qiqiang Lin, Muning Wen, Qiuying Peng, Guanyu Nie, Junwei Liao, Xiaoyun Mo, Jiamu Zhou, Cheng Cheng, Yin Zhao, Jun Wang, et al. 2025. Robust function-calling for on-device language model via function masking. In *The Thirteenth International Conference on Learning Representations*.
- Jiayu Liu, Zhenya Huang, Tong Xiao, Jing Sha, Jinze Wu, Qi Liu, Shijin Wang, and Enhong Chen. 2024a. SocraticLM: Exploring socratic personalized teaching with large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024b. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Weiwu Liu, Kingshan Zeng, Xu Huang, xinlong hao, Shuai Yu, Dexun Li, Shuai Wang, Weinan Gan, Zhengying Liu, Yuanqing Yu, Zezhong WANG, Yuxian Wang, Wu Ning, Yutai Hou, Bin Wang, Chuhan Wu, Wang Xinzhi, Yong Liu, Yasheng Wang, Duyu Tang, Dandan Tu, Lifeng Shang, Xin Jiang, Ruiming Tang, Defu Lian, Qun Liu, and Enhong Chen. 2025a. ToolACE: Enhancing function calling with accuracy, complexity, and diversity. In *The Thirteenth International Conference on Learning Representations*.
- Yanming Liu, Xinyue Peng, Jiannan Cao, Shi Bo, Yuwei Zhang, Xuhong Zhang, Sheng Cheng, Xun Wang, Jianwei Yin, and Tianyu Du. 2025b. Tool-planner: Task planning with clusters across multiple tools. In *The Thirteenth International Conference on Learning Representations*.
- Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Chris Leung, Jiajie Tang, and Jiebo Luo. 2024. LLM-rec: Personalized recommendation via prompting large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 583–612. Association for Computational Linguistics.
- Shishir G Patil, Tianjun Zhang, Xin Wang, and Joseph E Gonzalez. 2024. Gorilla: Large language model connected with massive apis. In *Advances in Neural Information Processing Systems*, volume 37, pages 126544–126565. Curran Associates, Inc.
- Shuofei Qiao, Runnan Fang, Zhisong Qiu, Xiaobin Wang, Ningyu Zhang, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. 2025. Benchmarking agentic workflow generation. In *The Thirteenth International Conference on Learning Representations*.
- Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Xuanhe Zhou, Yufei Huang, Chaojun Xiao, Chi Han, Yi Ren Fung, Yusheng Su, Huadong Wang, Cheng Qian, Runchu Tian, Kunlun Zhu, Shihao Liang, Xingyu Shen, Bokai Xu, Zhen Zhang, Yining Ye, Bowen Li, Ziwei Tang, Jing Yi, Yuzhang Zhu, Zhenning Dai, Lan Yan, Xin Cong, Yaxi Lu, Weilin Zhao, Yuxiang Huang, Junxi Yan, Xu Han, Xian Sun, Dahai Li, Jason Phang, Cheng Yang, Tongshuang Wu, Heng Ji, Guoliang Li, Zhiyuan Liu, and Maosong Sun. 2024a. Tool learning with foundation models. *ACM Comput. Surv.*, 57(4).
- 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. 2025a. From exploration to mastery: Enabling LLMs to master tools via self-driven interactions. In *The Thirteenth International Conference on Learning Representations*.
- Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, and Ji-Rong

- Wen. 2025b. Tool learning with large language models: A survey. *Frontiers of Computer Science*, 19(8):198343.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, pages 53728–53741. Curran Associates, Inc.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems*, volume 36, pages 68539–68551. Curran Associates, Inc.
- Erich Schubert, Jörg Sander, Martin Ester, Hans Peter Kriegel, and Xiaowei Xu. 2017. Dbscan revisited, revisited: why and how you should (still) use dbscan. *ACM Transactions on Database Systems (TODS)*, 42(3):1–21.
- Haiyang SHEN, Yue Li, Desong Meng, Dongqi Cai, Sheng Qi, Li Zhang, Mengwei Xu, and Yun Ma. 2025. Shortcutsbench: A large-scale real-world benchmark for API-based agents. In *The Thirteenth International Conference on Learning Representations*.
- Yongliang Shen, Kaitao Song, Xu Tan, Wenqi Zhang, Kan Ren, Siyu Yuan, Weiming Lu, Dongsheng Li, and Yueting Zhuang. 2024. Taskbench: Benchmarking large language models for task automation. In *Advances in Neural Information Processing Systems*, volume 37, pages 4540–4574. Curran Associates, Inc.
- Zhengliang Shi, Shen Gao, Xiuyi Chen, Yue Feng, Lingyong Yan, Haibo Shi, Dawei Yin, Pengjie Ren, Suzan Verberne, and Zhaochun Ren. 2024. Learning to use tools via cooperative and interactive agents. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10642–10657. Association for Computational Linguistics.
- Qiaoyu Tang, Ziliang Deng, Hongyu Lin, Xianpei Han, Qiao Liang, and Le Sun. 2023. Toolalpaca: Generalized tool learning for language models with 3000 simulated cases. *arXiv preprint arXiv:2306.05301*.
- Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Wei-Lin Chen, Chao-Wei Huang, Yu Meng, and Yun-Nung Chen. 2024. Two tales of persona in LLMs: A survey of role-playing and personalization. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16612–16631. Association for Computational Linguistics.
- Hongru Wang, Rui Wang, Boyang Xue, Heming Xia, Jingtao Cao, Zeming Liu, Jeff Z. Pan, and Kam-Fai Wong. 2024a. AppBench: Planning of multiple APIs from various APPs for complex user instruction. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15322–15336. Association for Computational Linguistics.
- Jian Wang, Yi Cheng, Dongding Lin, Chak Leong, and Wenjie Li. 2023. Target-oriented proactive dialogue systems with personalization: Problem formulation and dataset curation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1132–1143. Association for Computational Linguistics.
- Jize Wang, Ma Zerun, Yining Li, Songyang Zhang, Cailian Chen, Kai Chen, and Xinyi Le. 2024b. Gta: A benchmark for general tool agents. In *Advances in Neural Information Processing Systems*, volume 37, pages 75749–75790. Curran Associates, Inc.
- Pei Wang, Yanan Wu, Noah Wang, Jiaheng Liu, Xiaoshuai Song, Z.Y. Peng, Ken Deng, Chenchen Zhang, Jiakai Wang, Junran Peng, Ge Zhang, Hangyu Guo, Zhaoxiang Zhang, Wenbo Su, and Bo Zheng. 2025a. MTU-bench: A multi-granularity tool-use benchmark for large language models. In *The Thirteenth International Conference on Learning Representations*.
- Renxi Wang, Xudong Han, Lei Ji, Shu Wang, Timothy Baldwin, and Haonan Li. 2025b. Toolgen: Unified tool retrieval and calling via generation. In *The Thirteenth International Conference on Learning Representations*.
- 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. Association for Computational Linguistics.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Ruifeng Yuan, Shichao Sun, Yongqi Li, Zili Wang, Ziqiang Cao, and Wenjie Li. 2025a. Personalized large language model assistant with evolving conditional memory. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 3764–3777. Association for Computational Linguistics.
- Siyu Yuan, Kaitao Song, Jiangjie Chen, Xu Tan, Yongliang Shen, Kan Ren, Dongsheng Li, and Deqing Yang. 2025b. 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. Association for Computational Linguistics.
- Jianguo Zhang, Tian Lan, Ming Zhu, Zuxin Liu, Thai Quoc Hoang, Shirley Kokane, Weiran Yao, Juntao Tan, Akshara Prabhakar, Haolin Chen, Zhiwei Liu, Yihao Feng, Tulika Manoj Awalgaoonkar, Rithesh

R N, Zeyuan Chen, Ran Xu, Juan Carlos Niebles, Shelby Heinecke, Huan Wang, Silvio Savarese, and Caiming Xiong. 2025. xLAM: A family of large action models to empower AI agent systems. 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 11583–11597. Association for Computational Linguistics.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyang Luo. 2024. LlamaFactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 400–410. Association for Computational Linguistics.

Yujia Zhou, Qiannan Zhu, Jiajie Jin, and Zhicheng Dou. 2024. Cognitive personalized search integrating large language models with an efficient memory mechanism. In *Proceedings of the ACM Web Conference 2024*, page 1464–1473. Association for Computing Machinery.

## A Details of Benchmark Construction

We provide the illustration of three types of the user’s interaction history in Figure 7.

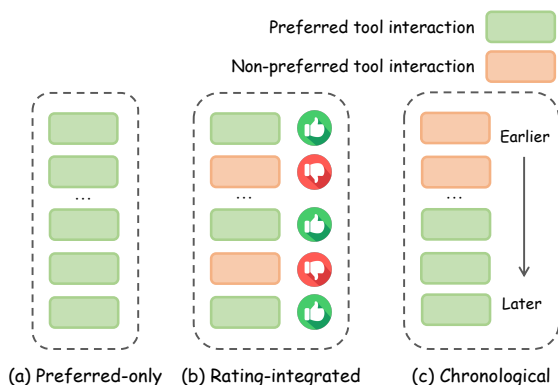


Figure 7: Illustration of three types of the user’s interaction history.

## B More Implementation Details

To train PEToolLLaMA, we fine-tune the LLaMA-3.1-8B model using LoRA with a warm-up ratio of 0.1 in the SFT stage. The learning rate is set to  $1e-4$  with a batch size of 16 per GPU. In the DPO stage, the learning rate is set to  $1e-6$  and the balancing factor  $\beta$  is set to 0.1 with a batch size of 32. We have trained the model several times to

ensure the improvement is not randomly achieved and present the mid one. For evaluation, we set the number of candidate tools  $N$  to 10 and the temperature to 0.1 to reduce randomness. Since the maximum context length varies in different LLMs, we constrain the context window to 4000 tokens. The experiments on closed-source LLMs are fulfilled by APIs of OpenAI and those on open-source LLMs are conducted on NVIDIA A6000 GPUs with 48 GB of memory.

## C Prompt Details

The prompt templates in for tool-use example generation and tool attributes understanding are shown in Figure 8 and Figure 9. The prompt templates for interaction history generation across three types are shown in Figure 10, Figure 11, and Figure 12. The prompt template for instruction generation is shown in Figure 13.



### Prompt for Tool-use Example Generation

**Prompt:**

Given a tool documentation as input, your task is to output an example for using this tool, including a simulated user instruction and parameters for calling the tool. The output example should be in JSON format: {"instruction": xx, "parameters": xx}

Here is a demonstration:

Input:

```
{
  "tool_name": "<Text_Analysis>.<Spellout>.<Languages>",
  "tool_description": "List ISO 639 languages",
  "required_parameters": [],
  "optional_parameters": [
    {
      "name": "nameFilter",
      "type": "STRING",
      "description": "Filter as \"contains\" by language name",
      "default": ""
    }
  ]
}
```

Output:

```
{
  "instruction": "I want to filter the list of languages by English",
  "parameters": {
    "nameFilter": "English"
  }
}
```

Now you will be given the tool documentation, please generate the tool-use example.  
Begin!

Figure 8: The prompt for tool-use example generation.

## Prompt for Tool Attributes Understanding

### Prompt:

Given a tool documentation and the corresponding tool-use example as input, your task is to understand the tool attributes thoroughly. Then generate two descriptions about the functionality and non-functional attributes of the tool respectively.

Here is a demonstration:

### Input:

```
Tool documentation:
{
  "tool_name": "<Commerce>.<Face Compare>.<GET Call>",
  "tool_description": "Used to fetch results using the request id received in responses.",
  "required_parameters": [
    {
      "name": "request_id",
      "type": "STRING",
      "description": "",
      "default": "76d1c748-51ed-435b-bcd8-3d9c9d3eb68a"
    }
  ],
  "Tool-use example:
  {
    "instruction": "I want to use the request id '76d1c748-51ed-435b-bcd8-3d9c9d3eb68a' to fetch the face comparison result",
    "parameters": {
      "request_id": "76d1c748-51ed-435b-bcd8-3d9c9d3eb68a"
    }
  }
}
```

### Output:

```
Functionality: Fetches API results based on the request ID received in previous responses.
Non-functional attributes: Designed for commerce applications, used in face comparison scenarios.
```

Now you will be given the tool documentation and the tool-use example, generate two short phrases to describe the two types of attributes.

Begin!

Figure 9: The prompt for tool attributes understanding.

## Prompt for Interaction History (Preferred-only) Generation

### Prompt:

Given a list of tools preferred by a user as input, your task is to simulate the user's interaction history based on these tools. You should output a sequence of tool-usage interactions, each consisting of a simulated user instruction and a tool call to fulfill that instruction. The interaction sequence should be a list in JSON format:

```
[
  {
    "instruction": xx,
    "tool_call": {
      "tool_name": xx,
      "parameters": xx
    }
  }, ...
]
```

Now you will be given the tools, please generate the interaction sequence.

Begin!

Figure 10: The prompt for interaction history (preferred-only) generation.

### Prompt for Interaction History (Rating-integrated) Generation

**Prompt:**

Given a list of tools preferred by a user and a list of tools not preferred as input, your task is to simulate the user's interaction history based on these two lists. You should output a sequence of tool-usage interactions, each consisting of a simulated user instruction, a tool call to fulfill that instruction, and a binary rating reflecting the user's satisfaction with the tool call. The interaction sequence should be a list in JSON format:

```
[
  {
    "instruction": xx,
    "tool_call": {
      "tool_name": xx,
      "parameters": xx
    },
    "rating": 1 or 0,
  }, ...
]
```

Now you will be given the two lists of tools, please generate the interaction sequence.  
Begin!

Figure 11: The prompt for interaction history (rating-integrated) generation.

### Prompt for Interaction History (Chronological) Generation

**Prompt:**

Given a list of tools preferred by a user and a list of tools not preferred as input, your task is to simulate the user's interaction history based on these two lists. You should output a sequence of tool-usage interactions, each consisting of a simulated user instruction, a tool call to fulfill that instruction. The interactions should be organized in time order to reflect changes in user preferences over time, i.e., the more recent tool-usage interactions are more preferred by the user, while earlier interactions are less preferred. The interaction sequence should be a list in JSON format:

```
[
  {
    "instruction": xx,
    "tool_call": {
      "tool_name": xx,
      "parameters": xx
    }
  }, ...
]
```

Now you will be given the two lists of tools, please generate the interaction sequence.  
Begin!

Figure 12: The prompt for interaction history (chronological) generation.

### Prompt for Instruction Generation

**Prompt:**

Given a user's interaction history and a tool documentation as input, your task is to generate a simulated user instruction which can be fulfilled by calling the tool with parameters. The generated output should be in JSON format:

```
{  
  "instruction": xx,  
  "parameters": xx  
}
```

Remember, tool name is strictly prohibited from appearing in the generated instruction. Now you will be given the user's interaction history and tool documentation, please generate the output. Begin!

Figure 13: The prompt for instruction generation.