

---

# Advancing Tool-Augmented Large Language Models: Integrating Insights from Errors in Inference Trees

---

Sijia Chen<sup>1,2,\*</sup> Yibo Wang<sup>1,2,\*</sup> Yi-Feng Wu<sup>3</sup> Qing-Guo Chen<sup>3</sup>  
Zhao Xu<sup>3</sup> Weihua Luo<sup>3</sup> Kaifu Zhang<sup>3</sup> Lijun Zhang<sup>1,4,2,†</sup>

<sup>1</sup>National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China

<sup>2</sup>School of Artificial Intelligence, Nanjing University, Nanjing, China

<sup>3</sup>Alibaba International Digital Commerce <sup>4</sup>Pazhou Laboratory (Huangpu), Guangzhou, China

{chensj, wangyb, zhanglj}@lamda.nju.edu.cn

{yixin.wyf, qingguo.cqg, changgong.xz,

weihua.luowh, kaifu.zkf}@alibaba-inc.com

## Abstract

Tool-augmented large language models (LLMs) leverage tools, often in the form of APIs, to improve their reasoning capabilities on complex tasks. This enables them to act as intelligent agents interacting with the real world. The recently introduced ToolLLaMA model by Qin et al. [2023] utilizes the depth-first search-based decision tree (DFS-DT) mechanism for multi-step reasoning with 16000+ real-world APIs, effectively enhancing the performance of tool-augmented LLMs compared to traditional chain reasoning mechanisms. However, their approach only employs successful paths from decision trees (also called inference trees) for supervised fine-tuning (SFT), missing out on the potential learning opportunities from failed paths. Inspired by this, we propose an inference trajectory optimization framework based on preference learning to address this limitation. We first introduce a novel method for constructing step-wise preference data from tree-like expert trajectories, which leverages the previously ignored failed explorations in the decision trees. In the subsequent training phase, we first fine-tune the LLM with successful tool-usage expert trajectories and then apply direct preference optimization (DPO) with the preference data to update the LLM’s policy, resulting in our ToolPrefer-LLaMA (TP-LLaMA) model. This approach not only enhances the utilization of original expert data but also broadens the learning space of the model. Our experiments demonstrate that by obtaining insights from errors in inference trees, TP-LLaMA significantly outperforms the baselines across almost all test scenarios by a large margin and exhibits better generalization capabilities with unseen APIs. At the same time, TP-LLaMA has also demonstrated superior reasoning efficiency compared to the baselines, making it more suitable for complex tool-usage reasoning tasks.

## 1 Introduction

In recent years, large language models (LLMs) have exhibited impressive capabilities in various areas, including language understanding and generation, multi-modal content learning and reasoning, and even embodied intelligence task processing [Brown et al., 2020, Alayrac et al., 2022, Zeng et al., 2023, Li et al., 2023, Chen et al., 2024, Liu et al., 2024a, Lu et al., 2024, Zhang et al., 2024, Cao et al., 2024a,b, Mazzaglia et al., 2024]. Despite these notable strengths, these models still

---

\*Equal contribution. Work done during the internship at Alibaba International Digital Commerce.

†Corresponding author.

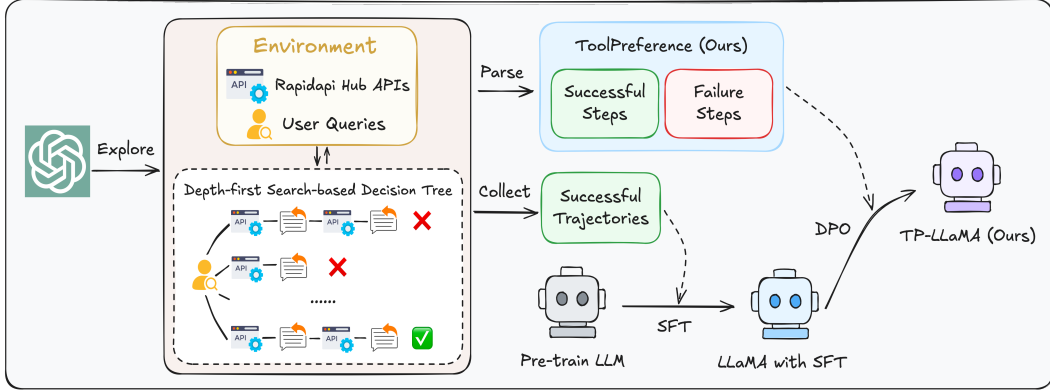


Figure 1: Our Inference Trajectory Optimization Framework.

face significant challenges, such as a lack of access to real-time information [Komeili et al., 2021] and difficulties in precise mathematical tasks [Patel et al., 2021, Lu et al., 2023b]. The development of tool-augmented LLMs tackles these challenges by enabling LLMs to interact with external tools (often in the form of APIs), significantly enhancing their capabilities. This advancement allows LLMs to serve as efficient intermediaries between users and a large ecosystem of applications. Notably, tool-augmented LLMs based on the ChatGPT [Brown et al., 2020] and GPT-4 [Achiam et al., 2023] have achieved outstanding results by using few-shot or zero-shot prompts to activate the LLM’s inherent tool-usage abilities [Deng et al., 2023, Lu et al., 2023a, Lin et al., 2024]. Despite this progress, some studies demonstrate that open-source LLMs still exhibit a significant gap in their capacity to utilize external tools compared to state-of-the-art (SOTA) closed-source models like GPT-4 [Liu et al., 2024b, Wang et al., 2024b]. To bridge this gap, aligning these open-source LLMs with tool-usage downstream tasks is essential.

Currently, most efforts to align open-source LLMs with tool-usage downstream tasks rely on supervised fine-tuning (SFT) with expert trajectory datasets, which trains LLMs to learn strategies for subsequent actions based on previous actions and observations [Patil et al., 2023, Schick et al., 2023]. Early studies in this field typically have limitations such as a restricted variety of APIs, the reliance on single-tool scenarios, and the use of simple reasoning methods [Wei et al., 2022, Yao et al., 2023, Patil et al., 2023]. The recent work by Qin et al. [2023], which focuses on the scene of LLM’s multi-step reasoning with external tools, solves the above limitations. They introduce an instruction tuning dataset called ToolBench, which includes over 16,000 real-world APIs and various realistic instructions, along with expert trajectories annotated by ChatGPT based on a depth-first search-based decision tree (DFSDT) reasoning mechanism. They then perform SFT training on LLaMA with this dataset to create the ToolLLaMA model, which shows remarkable performance. However, ToolLLaMA’s training is still based on expert behavior cloning, potentially limiting exploration of the target space and leading to suboptimal strategies. Additionally, although their expert trajectories are structured as DFS trees, only successful trajectories are utilized in the SFT training, which neglects valuable insights from failed attempts and results in low data utilization.

As the saying goes, “a fall into a pit, a gain in your wit”, effective human learning involves not only drawing lessons from success but also from failures. Inspired by this, we propose a new inference trajectory optimization framework for developing tool-augmented LLMs as illustrated in Figure 1, which enhances the tool learning process by incorporating previously ignored failure exploration information via preference learning. Specifically, using the tree-like expert trajectories from ToolBench [Qin et al., 2023], we first parse each pair of branch nodes along the successful trajectory in the decision tree into a preference sample pair, thereby constructing a step-wise tool-usage preference dataset named *ToolPreference*.<sup>3</sup> Subsequently, after conducting SFT training on the pre-trained LLM with successful trajectories, we employ the direct preference optimization (DPO) method [Rafailov et al., 2023] with the *ToolPreference* dataset to further align the LLM with tool-usage downstream tasks, and thus obtain our model, named *ToolPrefer-LLaMA (TP-LLaMA)*. Our strategy improves the utilization of expert data and simultaneously broadens the learning space.

<sup>3</sup>The dataset is available at <https://huggingface.co/datasets/chrisiecsj/ToolPreference>.

Our experiments are conducted on the test tasks from ToolBench. To evaluate the performance, we adopt two metrics: the *pass rate*, which measures the probability of the model successfully providing an answer within limited steps; and the *win rate*, which quantifies the likelihood that the evaluator will prefer the model’s responses. From the experiment results, we have the following findings:

- Across all test scenarios, TP-LLaMA consistently surpasses ToolLLaMA and other baselines, with an average pass rate improvement of at least 12% and a win rate that outperforms nearly all other models by an average of 4%. These results demonstrate that learning from failed attempts can significantly enhance the decision-making ability of LLMs. Additionally, our model shows superior generalization to unseen APIs.
- Efficiency experiments show that our model requires an average of only 22.62 steps for DFSDT inference, compared to 32.06 steps for the SFT model. This enhancement stems from our method’s ability to avoid unnecessary branch explorations in DFSDT reasoning.
- Our ablation experiments verify that the effectiveness of our preference dataset and inference trajectory optimization framework has nothing to do with the base model itself. Better results can still be obtained after replacing the base model with Mistral-7B [Jiang et al., 2023], Qwen1.5-7B [Bai et al., 2023], and Gemma-7B [Team et al., 2024].

In summary, this work aims to enhance the performance of LLMs on multi-step reasoning with external tools by integrating insights from errors in tree-like reasoning trajectories and employing step-wise preference pairs for preference learning. Our key contributions include: (i) A novel method for constructing step-wise preference data from tree-like expert trajectories, which may provide inspiration for future research; (ii) The proposal of using the preference learning method DPO to optimize the LLM’s tool-usage ability, along with the development of the TP-LLaMA model; (iii) Extensive experimental evaluations and in-depth analyses of the TP-LLaMA model, providing evidence of its effectiveness and validating its performance across various dimensions.

## 2 Related work

In this section, we briefly review recent progress on tool-augmented large language models and the development of preference learning.

**Tool-augmented large language models.** Over the past year, extensive research has been dedicated to developing tool-augmented LLMs, which exhibit improved reasoning abilities across various tasks by integrating external tools [Patil et al., 2023, Lu et al., 2023a, Schick et al., 2023, Lin et al., 2024]. The workflow for tool-augmented LLMs typically involves four key stages: task planning, tool selection, tool calls, and response generation. Early research mainly uses few-shot or zero-shot prompting methods to activate LLM’s inherent tool-usage abilities, often employing GPT as the LLM agent to manage several external tools such as AI models, web search, Python, and more [Shen et al., 2023b, Lu et al., 2023a]. While GPT performs well with external tools, open-source LLMs like LLaMA often struggle with direct tool usage and need additional task alignment. Therefore, subsequent research often utilizes instruction-tuning datasets annotated with tool calls to train open-source models, enhancing their ability to use tools. At the same time, these studies continue to explore a wider range of tools and scenarios [Schick et al., 2023, Patil et al., 2023].

One of the most comprehensive efforts in this field is by Qin et al. [2023]. They initially collect 16,464 real-world APIs across 49 categories, then utilize ChatGPT to automatically generate instructions that could invoke these APIs, and annotate expert trajectories to create a high-quality instruction tuning dataset named ToolBench. During the annotation, they employ the DFSDT reasoning mechanism to broaden the search space and enhance reasoning capabilities. By fine-tuning LLaMA on ToolBench, they develop ToolLLaMA, which has shown a compelling capability to handle both single-tool and complex multi-tool instructions.

**Preference learning** Preference learning uses human preferences from feedback data to assist decision-making. The earliest research in this field employs specially designed neural networks to help agents optimize action choices based on structured human guidance in programming languages [Maclin and Shavlik, 1996]. Subsequent studies shift focus to learning from numerical rewards provided by humans and performing reinforcement learning based on the prediction of these

rewards [Isbell et al., 2006, Knox and Stone, 2008, Knox, 2012]. This approach finds applications in areas like embodied intelligence [Pilarski et al., 2011, Suay and Chernova, 2011] and dialogue systems [El Asri et al., 2016]. The introduction of preference-based reinforcement learning marks a key milestone in the field, which uses qualitative human preferences, often in the form of rankings, to guide the optimization of policy models [Akrouer et al., 2011, Cheng et al., 2011]. Following this idea, Christiano et al. [2017] propose reinforcement learning from human feedback (RLHF), where a reward model is derived from human preferences to enhance reinforcement learning. This technique is later extended to natural language generation tasks [Kreutzer et al., 2018, Ziegler et al., 2019], advancing the integration of preference learning with LLM research [Ouyang et al., 2022].

### 3 Preliminaries

In this section, we start by formally defining the problem setup, and then we introduce key knowledge about preference learning methods, which is relevant to our approach.

#### 3.1 Problem setup

In this work, we use an iterative paradigm for the LLM’s multi-step reasoning with external tools, where the model selects each tool call based on the previous response, rather than pre-planning all tool calls at the start. Formally, we define it as a state transition process. The environment consists of a set of available tools  $\mathcal{T} = \{T_1, T_2, \dots, T_n\}$ , each with specific functionalities accessible through API calls. The task begins with an initial instruction  $I$ , usually consisting of a user query and a system prompt. At each reasoning step  $t$ , the LLM processes the current context  $S_t$ , defined as:

$$S_t = \{I, H_t\}$$

where  $H_t$  is the previous history, which includes the API decisions made  $\{A_1, \dots, A_{t-1}\}$ , and the API responses received  $\{R_1, \dots, R_{t-1}\}$ :

$$H_t = \{(A_1, R_1), \dots, (A_{t-1}, R_{t-1})\}.$$

The LLM then generates an action decision  $A_t$  based on this context, specifying both the tool  $T_i \in \mathcal{T}$  to use and its parameters. After the tool executes, the response  $R_t$  is generated and used to update the context. The reasoning process continues until the LLM determines that the task is complete and produces a final output  $O$  to answer the original query or gives up the task.

#### 3.2 Direct Preference Optimization

Preference learning has gained growing attention in LLM research. Its main goal is to optimize model outputs based on human (or expert) preferences, better aligning the model’s behavior with the expectations of real-world applications. Assume there is a preference dataset defined as  $\mathcal{D} = \{(x^{(i)}, y_w^{(i)}, y_l^{(i)})\}_{i=1, \dots, |\mathcal{D}|}$ , where  $x^{(i)}$  denotes the  $i$ -th prompt,  $y_w^{(i)}$  and  $y_l^{(i)}$  denote the corresponding preferred and dispreferred output respectively. Moreover, the notation  $y_w \succ y_l \mid x$  indicates that  $y_w$  is preferred than  $y_l$  for prompt  $x$ . Because the true distribution of human preferences is inaccessible, we assume it is generated by a latent reward model  $r^*(x, y)$ , where higher rewards indicate stronger preferences. Then, according to Rafailov et al. [2023], the human preference distribution  $p^*$  can be captured by the Bradley-Terry (BT) model [Bradley and Terry, 1952]:

$$p^*(y_1 \succ y_2 \mid x) = \frac{\exp(r^*(x, y_1))}{\exp(r^*(x, y_1)) + \exp(r^*(x, y_2))} = \sigma(r^*(x, y_1) - r^*(x, y_2)),$$

where  $\sigma$  is the logistic function. Obviously, we can estimate the parameters of the reward model via maximum likelihood estimation (equivalent to minimizing the negative log-likelihood.):

$$\mathcal{L}_R(r_\phi, \mathcal{D}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\log \sigma(r_\phi(x, y_w) - r_\phi(x, y_l))], \quad (1)$$

where  $r_\phi$  is a parameterized reward model.

To optimize the inference trajectories of LLM based on human preference, a popular method in recent LLM research is Reinforcement Learning from Human Feedback (RLHF) [Christiano et al., 2017, Ouyang et al., 2022]. In the RL phase of this method, the optimization goal is

$$\max_{\pi_\theta} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_\theta(y|x)} [r_\phi(x, y)] - \beta \mathbb{D}_{\text{KL}}[\pi_\theta(y \mid x) \parallel \pi_{\text{ref}}(y \mid x)], \quad (2)$$

where  $r_\phi$  is the reward model learned before,  $\pi_\theta$  is the policy model we need to optimize,  $\beta$  is a weighting parameter that controls the deviation from the base reference policy model  $\pi_{\text{ref}}$  (i.e., the LLM after SFT training). In practice,  $\pi_\theta$  is also initialized to the LLM after SFT. RLHF will use reinforcement learning methods (such as PPO [Schulman et al., 2017]) to optimize (2) and update the LLM’s strategy, with  $r_\phi(x, y)$  providing reward feedback. Additionally, some research in multi-step reasoning scenarios trains process reward models to evaluate each step instead of the entire output [Ma et al., 2023, Wang et al., 2024a]. However, RLHF incurs significant computational overhead, long training times, and potential instability [Shen et al., 2023a, Rafailov et al., 2023], making it less suitable for general tool-usage tasks.

Therefore, we choose a more convenient and faster approach that can also effectively align the model’s preferences — Direct Preference Optimization (DPO) [Rafailov et al., 2023], which eliminates the need to learn the reward model and directly uses preference data to optimize the LLM. Specifically, the optimal solution of (2) can be written as

$$\pi_r(y | x) = \frac{1}{Z(x)} \pi_{\text{ref}}(y | x) \exp\left(\frac{1}{\beta} r(x, y)\right), \quad (3)$$

where  $Z(x) = \sum_y \pi_{\text{ref}}(y | x) \exp\left(\frac{1}{\beta} r(x, y)\right)$  is the partition function [Rafailov et al., 2023]. We rearrange (3) to express  $r(x, y)$  in terms of  $\pi_r$  and  $\pi_{\text{ref}}$ :

$$r(x, y) = \beta \log \frac{\pi_r(y | x)}{\pi_{\text{ref}}(y | x)} + \beta \log Z(x). \quad (4)$$

Substituting (4) into (1), we can finally get the learning goal of DPO

$$\mathcal{L}_{\text{DPO}}(\pi_\theta, \pi_{\text{ref}}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[ \log \sigma \left( \beta \log \frac{\pi_\theta(y_w | x)}{\pi_{\text{ref}}(y_w | x)} - \beta \log \frac{\pi_\theta(y_l | x)}{\pi_{\text{ref}}(y_l | x)} \right) \right],$$

where  $\pi_\theta$  is a parametrized policy that we need to optimize. As a result, the optimization objective of DPO avoids additional learning of the reward model and the RL process while maximizing the final reward, which is more suitable for our general tool-usage scenarios.

## 4 Our method

In this section, we introduce our inference trajectory optimization framework, beginning with an overview of the framework, followed by a description of the preference data construction process.

### 4.1 The framework

Our framework is composed of two key stages: dataset construction and training. In the dataset construction stage, we create a tool-usage preference dataset, named ToolPreference, which is derived from the tree-like expert trajectories in Toolbench [Qin et al., 2023]. The specific process for constructing this dataset will be detailed in section 4.2.

**Remark 1** *It is important to emphasize that our preference data construction approach is not limited to Toolbench and can be adapted to any tree-structured multi-step instruction-tuning dataset, offering flexibility for various applications.*

In the training stage, we first perform SFT training on a pre-trained LLM using a resampled version of the instruction-tuning data from Toolbench (refer to Remark 2 for the resampling process). SFT training has been commonly adopted in previous research to enhance tool-augmented LLMs. However, mere cloning expert behavior through SFT is insufficient, as this method fails to adequately explore the environment, and can result in suboptimal strategies. To address this, after the SFT training, we further perform DPO training on the model with the ToolPreference dataset. This additional preference learning enhances the models reasoning capabilities when interacting with external tools and aligns its decision-making preferences with human preferences.

### 4.2 Preference data construction

Before introducing our preference data construction method, we first describe the dataset structure and expert trajectory format used in ToolBench [Qin et al., 2023].

- **Dataset structure.** ToolBench consists of two main components: API information data and instruction tuning data. The API information data is sourced from RapidAPI Hub<sup>4</sup> and includes 3,451 tools across 49 categories, with a total of 16,464 APIs (as each tool can have multiple APIs). Each API entry contains detailed information such as the name, description, HTTP method, URL, required and optional parameters, and executable code snippets for API calls. This comprehensive data enables LLMs to perform few-shot inference with effective API calls. The instruction-tuning data includes various single-tool or multi-tool instructions as well as corresponding annotated expert trajectories, generated in a self-instruction method by ChatGPT.
- **Expert trajectory format.** While traditional LLMs often use sequential reasoning methods like chain-of-thought (CoT) [Wei et al., 2022], which follow a single path to completion, ToolBench adopts a depth-first search (DFS) reasoning approach. As shown in the left half of Figure 2, expert trajectories in ToolBench are structured as decision trees with each tree node representing an LLM decision about an API call. Based on the tree structure, ToolBench implements DFS reasoning using two techniques. First, it defines two additional functions: one is “Finish with final answer”, where the LLM concludes it has gathered enough API responses to provide a correct answer and terminate the reasoning process, and the other is “Finish with giving up”, where the LLM feels unable to proceed with the task, abandons the current path and returns to a previous node. Second, diversity prompts are used to expand the search space. When expanding child nodes, the LLM will be prompted with information about previously explored child nodes of the same layer, and explicitly encouraged to generate different ones. Consequently, the LLM is allowed to either abandon the current path and restart from a previous step or proceed along a more promising path, exploring until an answer is reached or the node limit is reached.

We employ the second release of ToolBench<sup>5</sup>, which includes over 120,000 expert trajectories. Our approach is designed based on the motivation of improving data utilization. Although the tree-like expert trajectories in ToolBench extensively search the answer space, only successful paths are used in their training, neglecting valuable insights from failure paths. To address this, we extract preference decision pairs from each tree-like expert trajectory. After filtering out trajectories without failed exploration branches, we explore two different construction methods:

- **Path-wise** means using an entire success path and an entire failure path in the same decision tree to form a preference pair. As shown in the upper right part of Figure 2,  $\langle 0, 9, 12, 13, 14, 15 \rangle$  is the success path of the decision tree, and  $\langle 0, 1, 2 \rangle$ ,  $\langle 0, 3, 4, 5, 6 \rangle$ ,  $\langle 0, 3, 7, 8 \rangle$ ,  $\langle 0, 9, 10, 11 \rangle$  are 4 failure paths, so their Cartesian product can constitute a path-wise preference dataset, where  $\succ$  denotes the left part is preferred than the right part.
- **Step-wise** means using each branch node along the success path in the tree and its corresponding pair of child nodes (which must contain a child node on the success path) to construct a preference pair. As shown in the lower right part of Figure 2,  $\langle 0, 9, 12, 13, 14, 15 \rangle$  is the success path of the decision tree, while 0 and 9 are nodes with branches along the success path. Therefore,  $\langle 0, 9 \rangle \succ \langle 0, 1 \rangle$ ,  $\langle 0, 9 \rangle \succ \langle 0, 3 \rangle$ , and  $\langle 0, 9, 12 \rangle \succ \langle 0, 9, 10 \rangle$  can respectively form a preference pair.

Although it is intuitive and common to use path-wise preference samples, this approach is not well-suited to our task scenario. Theoretically, it may limit the model to only differentiate between correct and incorrect final responses to specific instructions, resulting in poor generalization with unseen instructions or tools. From an engineering perspective, learning preferences for an entire path at once is inconsistent with the model’s reasoning mechanism of inferring the next API call based on the response of the previous API execution each time, which makes it inherently unsuitable for the DFSDT reasoning mechanism.

In contrast, the step-wise design highlights the differences between each reasoning step, providing the model with more fine-grained process supervision. Theoretically, this method can better adjust the model’s reasoning process and enhance its generalization performance. It is also a more suitable fit for implementation within the DFSDT reasoning framework. Consequently, we create 69,393 pairs of preference samples from ToolBench in a step-wise manner. Each pair is formatted

<sup>4</sup><https://rapidapi.com/hub>

<sup>5</sup><https://github.com/OpenBMB/ToolBench.git>

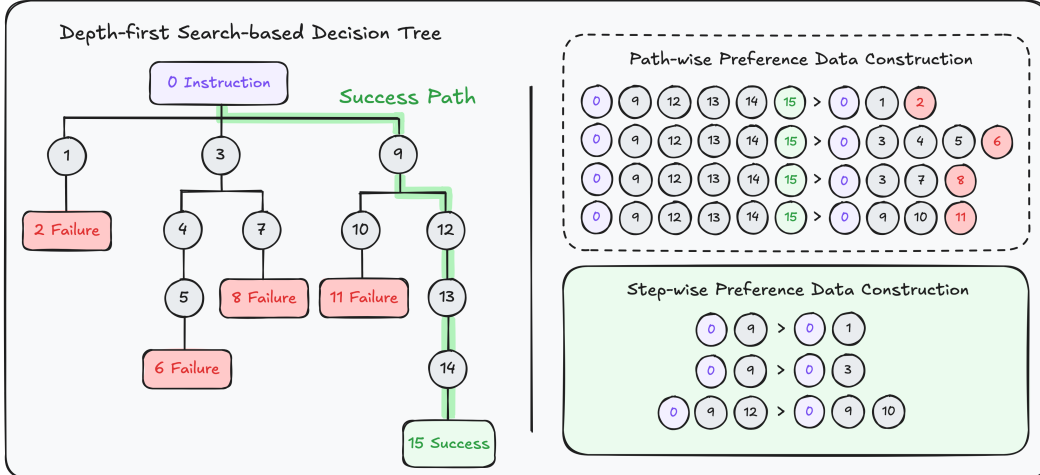


Figure 2: Depth-first search-based decision tree and two preference data construction methods

as {Instruction, Input, Output}. The Instruction includes the system prompt, detailing the DFSDT reasoning task and the relevant API documentation. The Input contains the user query and the reasoning history up to the current step, while the Output presents a preferred and a dispreferred reasoning step for the given input. Additionally, to prevent information leakage, we carefully remove any diversity prompts from each node’s information during parsing.

**Remark 2** *To ensure a rigorous comparison between models with and without preference learning in subsequent experiments, we also do not directly use the instruction-tuning dataset provided by ToolBench during the SFT phase. Instead, we filter out expert trajectories lacking failed exploration branches, as these could not be parsed into preference samples, and resampled the remaining data to create our SFT training set. This ensures the training data distribution remains consistent across models, regardless of whether preference learning is applied.*

## 5 Experiments

In this section, we investigate the performance of our inference trajectory optimization framework. We first introduce the experiments settings in Section 5.1. We then present the main results in Section 5.2, the efficiency experiments in Section 5.3, and the ablation experiments in Section 5.4.

### 5.1 Experiments settings

**Evaluation metrics.** Since our model uses APIs from the online platform RapidAPI Hub, there may be changes such as version updates or service termination over time, making it difficult to provide a fixed solution path for each test instruction. Following Qin et al. [2023], we use *pass rate* and *win rate* as evaluation metrics in our experiments. The pass rate represents the proportion that the model successfully gives answers within a certain number of reasoning actions (set to 200 in our experiment).<sup>6</sup> Specifically, a sample is considered passed if the reasoning trajectory finishes with the “Finish with final answer” API call. Additionally, we filter out samples that yield meaningless answers using a predefined set of feature keywords, such as “sorry”, “apologize”, etc. The win rate measures the likelihood that the solution path provided by the test model is preferred over the reference solution path for the same instruction. We use the answers given by ChatGPT+DFSDT as the reference solution paths and employ ChatGPT to determine preference.<sup>7</sup>

<sup>6</sup>During our experiment process, we noticed that ToolBench has been updated with a revised definition of pass rate [Qin et al., 2024] The definition we use in the main text follows the original version, while the revised definition and corresponding results will be provided in Appendix B.1

<sup>7</sup>The ChatGPT version we used in the experiments in the main text is gpt-3.5-turbo-16k.

Table 1: Main Experiment Results. Avg represents the average pass rate or win rate of the 6 test scenarios. A win rate higher than 50% means the model performs better than ChatGPT+DFSdT.

Pass Rate							
Model	G1-Ins.	G1-Tool	G1-Cat.	G2-Ins.	G2-Cat.	G3-Ins.	Avg
ChatGPT	0.52	0.55	0.60	0.51	0.51	0.21	0.48
Davinci	0.49	0.47	0.45	0.40	0.27	0.29	0.40
ToolLLaMA	0.54	0.60	0.62	0.47	0.54	0.17	0.49
LLaMA with SFT	0.47	0.53	0.72	0.48	0.63	0.35	0.53
<b>TP-LLaMA (ours)</b>	<b>0.55</b>	<b>0.65</b>	<b>0.80</b>	<b>0.62</b>	<b>0.67</b>	<b>0.61</b>	<b>0.65</b>
Win Rate							
Model	G1-Ins.	G1-Tool	G1-Cat.	G2-Ins.	G2-Cat.	G3-Ins.	Avg
ChatGPT	-	-	-	-	-	-	-
Davinci	0.37	0.37	0.35	0.35	0.29	0.54	0.38
ToolLLaMA	0.55	0.53	<b>0.57</b>	0.56	0.52	0.68	0.57
LLaMA with SFT	0.54	0.51	0.56	0.65	0.57	0.81	0.61
<b>TP-LLaMA (ours)</b>	<b>0.56</b>	<b>0.59</b>	0.54	<b>0.70</b>	<b>0.64</b>	<b>0.86</b>	<b>0.65</b>

**Training settings.** For the 2-epoch SFT training, we randomly sampled 11,142 instances from the expert-annotated data in ToolBench after removing those without failed exploration branches. The batch size is 16 and the learning rate is  $1e-5$  during SFT training. For the 1-epoch DPO training, we randomly sample 8,202 preference data pairs from our ToolPreference dataset, the batch size is 8, the learning rate is  $1e-6$  and  $\beta = 0.5$  in (2). It is important to note that our sampling is performed at the instruction level, which means that samples corresponding to the same instruction are either all included in the training set or none are included. We provide a detailed explanation of our design choices for training hyperparameters in Appendix A.1. All our experiments are conducted on a single machine equipped with 8 NVIDIA A100 GPUs with 80G memory.

**Testing settings.** We investigate six test scenarios same as Qin et al. [2023]: G1-Cat., G1-Ins., G1-Tool, G2-Cat., G2-Ins., and G3-Ins.. The specific meanings are as follows: (1) **G1**: instructions that only use a single tool; (2) **G2**: instructions that use intra-category multi-tools; (3) **G3**: instructions that use inter-category multi-tools; (4) **Cat. (Category)**: unseen tools that belong to the unseen category of tools in the training data; (5) **Ins. (Instruction)**: unseen instructions for the same set of tools in the training data; (6) **Tool**: unseen tools that belong to the same category of tools in the training data. Each test scene contains 200 test samples, except G3-Ins., which contains 100 test samples. The six test scenarios have different task difficulties and generalization challenges, which can well reflect the comprehensive performance of models.

**Baselines.** We compare our model with several models without preference learning. Among them, we select the expert model ChatGPT and OpenAI Text-Davinci-003 (Davinci for short) as baselines. In addition, we also show the results of ToolLLaMA and the model trained by SFT using our resampled SFT training set (LLaMA with SFT for short) for comparison. Note that all models here are combined with DFSdT for inference. In addition, regarding the ToolLLaMA results, we directly use the reasoning answers of ToolLLaMA on test sets provided by ToolBench’s GitHub repository to calculate pass rates and win rates.

## 5.2 Main results

We employ LLaMA-2-7B as the base model of our training framework and finally obtain our model, named ToolPrefer-LLaMA (TP-LLaMA). The context length of LLaMA-2-7B is extended to 8192 tokens to accommodate our tool-usage reasoning tasks. The main results are shown in Table 1. We have the following important observations:



Table 2: Efficiency Results of TP-LLaMA. Imp denotes the improvement of TP-LLaMA over LLaMA with SFT in terms of the average steps.

Model	G1-Ins.	G1-Tool	G1-Cat.	G2-Ins.	G2-Cat.	G3-Ins.	Avg	Imp
LLaMA with SFT	32.82	34.60	31.45	31.98	35.05	26.44	32.06	-
<b>TP-LLaMA (ours)</b>	<b>24.54</b>	<b>24.19</b>	<b>23.85</b>	<b>23.98</b>	<b>23.53</b>	<b>15.61</b>	<b>22.62</b>	<b>29.44%</b>

- TP-LLaMA significantly outperforms LLMs without preference learning in terms of pass rate, demonstrating the best performance across all six test scenarios, with an average improvement of over 12% compared to models not optimized using preference data.
- Regarding win rate, TP-LLaMA also exhibits competitive performance, just 3% below ToolLLaMA in the G1-Cat. scenario, while achieving the best results in all other scenarios.
- Furthermore, TP-LLaMA shows strong performance in more challenging task scenarios such as G2-Cat., G2-Ins., and G3-Ins., maintaining effectiveness similar to that in simpler tasks. Notably, in the G3-Ins. scenario, TP-LLaMA’s pass rate increased by over 26%, proving that our DPO training process using preference data significantly enhances the model’s ability to handle complex multi-tool tasks.

Although we use the provided reasoning answers of ToolLLaMA from ToolBench’s GitHub repository to calculate its rates, the results indeed differ from those reported in their paper [Qin et al., 2023]. This may be due to the reasoning answers version not matching the one used in their paper or differences in the evaluation environment settings. However, it’s important to emphasize that our results remain valid and reliable. We apply consistent settings across all models, so their relative differences are meaningful. Overall, our experimental results indicate that through preference learning, TP-LLaMA can master various tool-usage instructions better and exhibits stronger generalization capabilities to unseen tools, categories, and instructions.

### 5.3 Efficiency Evaluation

We also evaluate the inference efficiency of TP-LLaMA on six test scenarios and employ the average number of DFSDT inference steps required for samples that ended with the Finish function as the metric. From Table 2, we can find that LLaMA with SFT requires an average of 32.06 steps for reasoning, while our TP-LLaMA only requires an average of 22.62 steps of reasoning in all test scenarios, with an improvement of 29.44%. These results clearly demonstrate that the inference efficiency of TP-LLaMA is remarkably superior to that of the model trained only with success trajectories. This advantage arises from our step-wise preference data, which allows the model to identify the most optimal decisions at each step of reasoning through DPO training. As a result, the model avoids the exploration of unnecessary sub-optimal branches in the decision tree, thereby increasing reasoning speed and efficiency.

### 5.4 Ablation experiments

In the ablation experiments, to verify the effectiveness of our framework, we further replace LLaMA-2-7B with other base models, including Mistral-7B, Qwen1.5-7B, and Gemma-7B. The results are shown in Table 3 and Table 4.

From Table 3, no matter which base model is used, training on preference data can always bring gains to the performance of the model, which verifies the model-independent effectiveness of our framework. Specifically, in terms of pass rates, models that have learned from expert errors improve by at least 8% on average compared to those that only receive training on success trajectory information. Similarly, in terms of win rates, models with insights from preference data generally outperform those without preference learning. Table 4 further confirms that our method significantly improves model inference efficiency by a large margin, up to an average of 33.35%.

Table 3: Ablation Performance Experiment Results. Avg represents the average pass rate or win rate of the 6 test scenarios. A win rate higher than 50% means the model performs better than ChatGPT+DFSDT.

Pass Rate							
Model	G1-Ins.	G1-Tool	G1-Cat.	G2-Ins.	G2-Cat.	G3-Ins.	Avg
Mistral with SFT	0.70	0.43	0.42	0.53	0.46	0.27	0.47
TP-LLaMA (Mistral)	<b>0.71</b>	<b>0.53</b>	<b>0.55</b>	<b>0.70</b>	<b>0.64</b>	<b>0.57</b>	<b>0.62</b>
Qwen with SFT	0.69	0.51	0.51	0.66	0.55	0.49	0.57
TP-LLaMA (Qwen)	<b>0.77</b>	<b>0.53</b>	<b>0.60</b>	<b>0.72</b>	<b>0.61</b>	<b>0.65</b>	<b>0.65</b>
Gemma with SFT	0.67	0.44	0.49	0.47	0.44	0.29	0.47
TP-LLaMA (Gemma)	<b>0.80</b>	<b>0.48</b>	<b>0.61</b>	<b>0.70</b>	<b>0.65</b>	<b>0.68</b>	<b>0.65</b>

Win Rate							
Model	G1-Ins.	G1-Tool	G1-Cat.	G2-Ins.	G2-Cat.	G3-Ins.	Avg
Mistral with SFT	0.52	0.47	0.55	0.61	0.61	0.72	0.58
TP-LLaMA (Mistral)	<b>0.53</b>	<b>0.50</b>	<b>0.57</b>	<b>0.62</b>	<b>0.64</b>	<b>0.74</b>	<b>0.60</b>
Qwen with SFT	0.53	0.52	0.52	0.64	0.66	0.75	0.60
TP-LLaMA (Qwen)	<b>0.54</b>	<b>0.54</b>	<b>0.58</b>	<b>0.66</b>	<b>0.67</b>	<b>0.81</b>	<b>0.63</b>
Gemma with SFT	0.58	0.54	0.53	0.50	0.62	0.73	0.58
TP-LLaMA (Gemma)	<b>0.61</b>	<b>0.57</b>	<b>0.58</b>	<b>0.65</b>	<b>0.67</b>	<b>0.75</b>	<b>0.64</b>

Table 4: Ablation Efficiency Experiment Results. Imp denotes the improvement of TP-LLaMA over LLaMA with SFT in terms of the average steps.

Model	Average Number of Steps in One Successful Path							Imp
	G1-Ins.	G1-Tool	G1-Cat.	G2-Ins.	G2-Cat.	G3-Ins.	Avg	
Mistral with SFT	28.92	26.65	30.22	25.69	26.58	25.24	27.22	-
<b>TP-LLaMA (Mistral)</b>	<b>25.30</b>	<b>25.01</b>	<b>23.36</b>	<b>23.51</b>	<b>20.74</b>	<b>16.42</b>	<b>22.39</b>	<b>17.74%</b>
Qwen with SFT	35.74	34.66	36.85	32.74	36.18	37.93	35.68	-
<b>TP-LLaMA (Qwen)</b>	<b>25.12</b>	<b>23.83</b>	<b>24.49</b>	<b>23.84</b>	<b>26.92</b>	<b>22.18</b>	<b>24.40</b>	<b>31.61%</b>
Gemma with SFT	27.49	22.77	24.10	18.70	20.52	21.19	22.46	-
<b>TP-LLaMA (Gemma)</b>	<b>17.15</b>	<b>13.88</b>	<b>15.91</b>	<b>13.63</b>	<b>13.30</b>	<b>15.94</b>	<b>14.97</b>	<b>33.35%</b>

## 6 Conclusion and future work

In this work, we propose a novel inference trajectory optimization framework that leverages preference learning to enhance the performance of tool-augmented LLMs. We first built a step-wise tool-usage preference dataset, ToolPreference, using our proposed new preference data construction method to convert previously ignored failed explorations in tree-like expert trajectories into valuable training data. After initial SFT training on the LLM, we use ToolPreference for DPO training to further refine the LLM’s strategy, resulting in our TP-LLaMA model. Our extensive comparative experiments prove that TP-LLaMA significantly outperforms the baseline models in nearly all test scenarios by learning from single-step errors in inference trees. TP-LLaMA also exhibits superior generalization capabilities and efficiency. Furthermore, ablation experiments confirm the model-independent effectiveness of our framework. In future work, we will try to explore tool-learning research with more complex, human-like reasoning mechanisms, and incorporate preference learning for further optimization. We also aim to extend our research to multimodal scenarios to evaluate the broader effectiveness of our approach.

## Acknowledgments and Disclosure of Funding

This work was partially supported by NSFC (U23A20382, 62122037), and the Collaborative Innovation Center of Novel Software Technology and Industrialization.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint*, arXiv:2303.08774, 2023.
- Riad Akrouf, Marc Schoenauer, and Michele Sebag. Preference-based policy learning. In *Proceedings of the 2011 European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD)*, pages 12–27, 2011.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 23716–23736, 2022.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint*, arXiv:2309.16609, 2023.
- Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 1877–1901, 2020.
- Qinglong Cao, Zhengqin Xu, Yuntian Chen, Chao Ma, and Xiaokang Yang. Domain-controlled prompt learning. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 936–944, 2024a.
- Qinglong Cao, Zhengqin Xu, Yuntian Chen, Chao Ma, and Xiaokang Yang. Domain prompt learning with quaternion networks. In *Proceedings of the 41st IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26637–26646, 2024b.
- Qi Chen, Bowen Zhang, Gang Wang, and Qi Wu. Weak-eval-Strong: Evaluating and eliciting lateral thinking of LLMs with situation puzzles. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*, 2024.
- Weimei Cheng, Johannes Fürnkranz, Eyke Hüllermeier, and Sang-Hyeun Park. Preference-based policy iteration: Leveraging preference learning for reinforcement learning. In *Proceedings of the 2011 European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD)*, pages 312–327, 2011.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems 30 (NIPS)*, pages 4302–4310, 2017.
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pages 28091–28114, 2023.
- Layla El Asri, Bilal Piot, Matthieu Geist, Romain Laroche, and Olivier Pietquin. Score-based inverse reinforcement learning. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 457–465, 2016.
- Charles Lee Isbell, Michael Kearns, Satinder Singh, Christian R Shelton, Peter Stone, and Dave Kormann. Cobot in lambdamoo: An adaptive social statistics agent. *Autonomous Agents and Multi-Agent Systems*, 3(13):327–354, 2006.

- 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. Mistral 7b. *arXiv preprint*, arXiv:2310.06825, 2023.
- W Bradley Knox and Peter Stone. Tamer: Training an agent manually via evaluative reinforcement. In *Proceedings of the 7th IEEE International Conference on Development and Learning (ICDL)*, pages 292–297, 2008.
- William Bradley Knox. Learning from human-generated reward. Ph.D. dissertation, 2012.
- Mojtaba Komeili, Kurt Shuster, and Jason Weston. Internet-augmented dialogue generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 8460–8478, 2021.
- Julia Kreutzer, Joshua Uyheng, and Stefan Riezler. Reliability and learnability of human bandit feedback for sequence-to-sequence reinforcement learning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1777–1788, 2018.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pages 19730–19742, 2023.
- Bill Yuchen Lin, Yicheng Fu, Karina Yang, Faeze Brahman, Shiyu Huang, Chandra Bhagavatula, Prithviraj Ammanabrolu, Yejin Choi, and Xiang Ren. Swiftsage: A generative agent with fast and slow thinking for complex interactive tasks. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*, 2024.
- Akide Liu, Jing Liu, Zizheng Pan, Yefei He, Gholamreza Haffari, and Bohan Zhuang. MiniCache: KV cache compression in depth dimension for large language models. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*, 2024a.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. Agentbench: Evaluating llms as agents. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2024b.
- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. Chameleon: Plug-and-play compositional reasoning with large language models. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pages 43447–43478, 2023a.
- Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck, and Kai-Wei Chang. A survey of deep learning for mathematical reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 14605–14631, 2023b.
- Shiyin Lu, Yang Li, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and Han-Jia Ye. Ovis: Structural embedding alignment for multimodal large language model. *arXiv preprint*, arXiv:2405.20797, 2024.
- Qianli Ma, Haotian Zhou, Tingkai Liu, Jianbo Yuan, Pengfei Liu, Yang You, and Hongxia Yang. Let’s reward step by step: Step-level reward model as the navigators for reasoning. *arXiv preprint*, arXiv:2310.10080, 2023.
- Richard Maclin and Jude W Shavlik. Creating advice-taking reinforcement learners. *Machine Learning*, 22(1):251–281, 1996.
- Pietro Mazzaglia, Tim Verbelen, Bart Dhoedt, Aaron Courville, and Sai Rajeswar. Multimodal foundation world models for generalist embodied agents. In *Proceedings of the 41st International Conference on Machine Learning (ICML) Multi-modal Foundation Model meets Embodied AI Workshop*, 2024.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 27730–27744, 2022.

- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are nlp models really able to solve simple math word problems? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 2080–2094, 2021.
- Shishir G Patil, Tianjun Zhang, Xin Wang, and Joseph E Gonzalez. Gorilla: Large language model connected with massive apis. *arXiv preprint*, arXiv:2305.15334, 2023.
- Patrick M Pilarski, Michael R Dawson, Thomas Degris, Farbod Fahimi, Jason P Carey, and Richard S Sutton. Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning. In *Proceedings of the 2011 IEEE International Conference on Rehabilitation Robotics (ICORR)*, pages 1–7, 2011.
- 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. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint*, arXiv:2307.16789v1, 2023.
- 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. ToolLLM: Facilitating large language models to master 16000+ real-world APIs. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2024.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pages 53728–53741, 2023.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pages 68539–68551, 2023.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint*, arXiv:1707.06347, 2017.
- Tianhao Shen, Renren Jin, Yufei Huang, Chuang Liu, Weilong Dong, Zishan Guo, Xinwei Wu, Yan Liu, and Deyi Xiong. Large language model alignment: A survey. *arXiv preprint*, arXiv:2309.15025, 2023a.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Huggingpt: Solving ai tasks with chatgpt and its friends in hugging face. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pages 38154–38180, 2023b.
- Halit Bener Suay and Sonia Chernova. Effect of human guidance and state space size on interactive reinforcement learning. In *Proceedings of the 20th IEEE International Symposium on Robot and Human Interactive Communication (Ro-Man)*, pages 1–6, 2011.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models based on gemini research and technology. *arXiv preprint*, arXiv:2403.08295, 2024.
- Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhi-fang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 9426–9439, 2024a.
- Xingyao Wang, Zihan Wang, Jiateng Liu, Yangyi Chen, Lifan Yuan, Hao Peng, and Heng Ji. Mint: Evaluating llms in multi-turn interaction with tools and language feedback. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2024b.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 24824–24837, 2022.

- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *Proceedings of the 11th International Conference on Learning Representations (ICLR)*, 2023.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. Glm-130b: An open bilingual pre-trained model. In *Proceedings of the 11th International Conference on Learning Representations (ICLR)*, 2023.
- Tianyu Zhang, Suyuchen Wang, Lu Li, Ge Zhang, Perouz Taslakian, Sai Rajeswar, Jie Fu, Bang Liu, and Yoshua Bengio. VCR: Visual caption restoration. *arXiv preprint*, arXiv:2406.06462, 2024.
- Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv preprint*, arXiv:1909.08593, 2019.

## Appendix

### A Experimental details

In this section, we supplement some details of the experiments in the main text, including the training details in Appendix A.1, the API information format in Appendix A.2, and the ToolPreference sample example in Appendix A.3.

#### A.1 Details for training

**Training hyperparameters** We provide an explanation of our design choices regarding training hyperparameters, specifically the sizes of the training and test sets. We first filter 42,192 tree-like expert trajectories with branching nodes from Toolbench, which leads to 69,393 DPO samples and 184,816 SFT samples after processing (as each instruction may correspond to multiple samples). After allocating a small part as a validation set, we sample training sets of different sizes based on these samples. The sampling methods we tried include “by instruction” and “by sample”. For sampling by instruction, the size of the SFT training set ranges from 2,500 to 10,000 queries, and the size of the DPO training set ranges from 5,000 to 32,192 queries, yielding nine combinations. For sampling by sample, the size of SFT varies from 10,000 to 183,561, and the size of DPO varies from 10,000 to 68,951, yielding seven combinations. We conduct small-scale tests based on these different training settings and find that increasing the size may lead to decreased model performance in scenarios with strong generalization, such as G3-ins. (e.g., with settings {SFT: 44,412, DPO: 41,226}, the pass rate drops to 0.36), possibly due to overfitting. Consequently, we select the set {SFT: 11,142, DPO: 8,202} in our final experiments.

**Computation time consumption** With 8 NVIDIA A100 GPUs, our SFT training phase takes an average of 4.6 hours, and the DPO training phase takes an average of 3.2 hours. In the inference phase, each API call takes about 3.4 seconds, and each query takes about 48.7 seconds. Computation time varies due to task complexity, network conditions, and API service status.

#### A.2 Details for API Information

Below we provide a detailed document of the API collected in ToolBench to help readers understand the format and content of API information.

##### API Information Sample

```
{
  "name": "Get Character By ID",
  "url": "Get individual character by ID\n Options:\n\n-
        Limit \u2192 Limit amount of responses received
        \n- Step \u2192 Skip amount of characters",
  "method": "GET",
  "required_parameters": [
    {
      "name": "id",
      "type": "NUMBER",
      "description": "",
      "default": ""
    }
  ],
  "optional_parameters": [],
  "code": "import requests .....",
  "test_endpoint": {
    "err": "Please enter a valid number for the character id."
  }
}
```

### A.3 Details for ToolPreference

Here we show an example preference pair in ToolPreference. For the sake of brevity, we have omitted some less important information, including some rules in the instruction, API parameter information, and some response content.

**Preference Sample Pair**

**Instruction**

You are AutoGPT, you can use many tools(functions) to do the following task.....Specifically, you have access to the following APIs:

```
{"name": "get_amazon_product_details_for_abiola_amazon_data_scraper", "description": "This is the subfunction for tool 'abiola_amazon_data_scraper', you can use this tool. The description of this function is: This endpoint get the general product description of any amazon product.", "parameters": .....}
```

```
{"name": "get_amazon_product_reviews_for_abiola_amazon_data_scraper", "description": "This is the subfunction for tool 'abiola_amazon_data_scraper', you can use this tool. The description of this function is: This endpoint gives you the review of the product you enter.", "parameters": .....}
```

```
{"name": "get_amazon_product_offer_for_abiola_amazon_data_scraper", "description": "This is the subfunction for tool 'abiola_amazon_data_scraper', you can use this tool. The description of this function is: This endpoint gets the offer of any amazon product.", "parameters": .....}
```

**Input**

```
{"from": "user", "value": "Can you fetch the product details, reviews, and offers for the product with ID B08L5NP6NG? Also, provide me with the search results for the product 'Apple MagSafe Charger'.Begin!"}
```

```
{"from": "assistant", "value": "Thought: Action: get_amazon_product_details_for_abiola_amazon_data_scraper Action Input: {'productid': 'B08L5NP6NG'}"}
```

```
{"from": "function", "value": "{\"name\": 'Apple MagSafe Charger', 'product_information': {'brand': 'Visit the Apple Store', 'brand_url': 'https://www.amazon.com/stores/Apple/page/77D9E1F7-0337-4282-9DB6-B6B8FB2DC98D?ref_ast_bln', 'full_description': '', 'pricing': '$39.00', .....}"}
```

```
{"from": "assistant", "value": "Thought: Action: get_amazon_product_reviews_for_abiola_amazon_data_scraper Action Input: {'productid': 'B08L5NP6NG'}"}
```

```
{"from": "function", "value": "....."}
```



## Preference Sample Pair

### Output

```
{"from": "assistant", "value": "Thought:Action:  
get_amazon_product_offer_for_abiola_amazon_data_scraper  
Action Input: {'productid': 'B08L5NP6NG'}"}  
  
{"from": "assistant", "value": "Thought: Action:  
get_amazon_product_details_for_abiola_amazon_data_scraper  
Action Input: {'productid': 'B08L5NP6NG'}"}
```

## B Supplementary experiment results

In this section, we first supplement the experiments on another defined pass rate in Appendix B.1. Then we conduct a qualitative analysis of our method in Appendix B.2, giving a specific case study.

### B.1 Another definition of pass rate

In the second version of ToolBench [Qin et al., 2024], a new pass rate definition using GPT-assisted evaluation is introduced. For each query, it first checks if the “Finish with Final Answer” API is called; if not, it’s considered a failure. If it is, GPT evaluates whether the answer resolves the query. If successful, it is marked as a pass. If not, GPT further assesses whether the query is solvable with the available APIs. If it isn’t, it is still considered a pass; otherwise, it is marked as a failure. Due to the ToolBench API server being offline temporarily, we use our own RapidAPI accounts to access APIs for evaluation experiments with the new pass rate definition. To maintain consistency in API status, we utilize ToolLLaMA’s open-source model<sup>8</sup> to perform reasoning on the test sets, instead of reusing the reasoning answers from its GitHub repository. Similarly, we re-run tests for other models using our RapidAPI accounts. We employ gpt-3.5-turbo-16k and gpt-3.5-turbo-1106 as GPT evaluators, with the results shown in Table 5.

First, TP-LLaMA still outperforms the models without preference learning, further validating the effectiveness of our method. However, the absolute pass rates depend heavily on the specific GPT version. We observe notable differences in preferences and consistency across GPT versions. After repeating the evaluation of each sample 7 times, we find that gpt-3.5-turbo-1106 is more likely to mark a sample as passed, while gpt-3.5-turbo-16k tends to judge it as not passed. This difference mainly stems from how each version assesses whether a query is solvable. Additionally, gpt-3.5-turbo-16k shows greater consistency across the 7 evaluations, meaning it is more likely to produce the same inference repeatedly. This highlights the importance of selecting the appropriate GPT version for evaluation, as relative scores may be more meaningful than absolute ones.

Furthermore, we observe that the gap between TP-LLaMA and ToolLLaMA narrows under the new evaluation. We believe this is due to two factors: (1) The models have different preferences formed during their respective training processes. TP-LLaMA tends to avoid giving up on reasoning and attempts partial answers, whereas ToolLLaMA is more likely to abandon a task entirely, leading to a complete failure. However, this gap narrows due to the use of GPT to evaluate whether the task is solvable. (2) During this supplementary experiment, our RapidAPI accounts have access limits (some APIs even can only be accessed 5 times per month per account), reducing the number of valid samples in the test sets. This particularly affects complex multi-tool reasoning tasks, where TP-LLaMA usually excels, making its performance gains appear smaller.

Additionally, the results we report for ToolLLaMA are still lower than those in Qin et al. [2024], likely due to shifts in the distribution of real-world APIs, which may make certain test samples unsolvable. Moreover, some features ToolLLaMA learned from past environments may not fully

<sup>8</sup><https://huggingface.co/ToolBench/ToolLLaMA-2-7b-v2>

Table 5: New Pass Rate Experiment Results

Model	G1-Ins.	G1-Tool	G1-Cat.	G2-Ins.	G2-Cat.	G3-Ins.	Avg
<b>gpt-3.5-turbo-16k</b>							
ToolLLaMA	<b>0.29</b>	0.35	0.40	0.33	0.29	0.25	0.32
LLaMA with SFT	0.22	0.29	0.39	0.28	0.29	0.28	0.29
<b>TP-LLaMA (ours)</b>	0.27	<b>0.37</b>	<b>0.48</b>	<b>0.35</b>	<b>0.36</b>	<b>0.35</b>	<b>0.36</b>
<b>gpt-3.5-turbo-1106</b>							
ToolLLaMA	0.35	<b>0.42</b>	0.39	0.63	0.59	0.68	0.51
LLaMA with SFT	0.33	0.40	0.40	0.64	0.55	0.60	0.50
<b>TP-LLaMA (ours)</b>	<b>0.37</b>	0.40	<b>0.41</b>	<b>0.65</b>	<b>0.63</b>	<b>0.69</b>	<b>0.53</b>

align with current conditions, resulting in reduced performance. In the future, we can further explore ways to enhance the models performance stability in evolving environments.

## B.2 Case study

We further illustrate the effectiveness of preference learning in improving the tool-usage capabilities of LLMs with a case study focused on the G3-Ins. scenario. To begin, we present the query along with the relevant API documentation

**Case Study: Query and Relevant APIs**

**Query:** I'm organizing a film festival and I'm looking for award-winning films. Can you search for videos related to "award-winning" on Vimeo? Additionally, fetch the related people in the "film festival" category to invite them as judges. Finally, provide me with a streaming link for a YouTube video with the ID "UxxajLWwzqY".

**Related API Documentation** (parameter information is omitted):

**Name:** getrelatedchannels\_for\_vimeo  
**Description:** Get Related Channels.

**Name:** searchvideos\_for\_vimeo  
**Description:** Search for videos.

**Name:** getrelatedpeople\_for\_vimeo  
**Description:** Get a list of related people for a category.

**Name:** download\_stream\_for\_ytstream\_download\_youtube\_videos  
**Description:** Stream or download info.

**Name:** Finish  
**Description:** If you believe that you have obtained a result that can answer the task, please call this function to provide the final answer. Alternatively, if you recognize that you are unable to proceed with the task in the current state, call this function to restart. Remember: you must ALWAYS call this function at the end of your attempt, and the only part that will be shown to the user is the final answer, so it should contain sufficient information.

During our experiments, the first three APIs above are temporarily inaccessible since they are real-world. So the best response to this query is a YouTube video link with the ID 'UxxajLWwzqY', completing part of the task. We then briefly present the reasoning trajectories from TP-LLaMA and LLaMA with SFT. Note that each reasoning path is limited to at most 4 API calls here.

## Case Study: Reasoning Trajectories from TP-LLaMA

The reasoning tree of TP-LLaMA includes 5 paths and we show the successful path below.

**Step 1** call [searchvideos\_for\_vimeo] with arguments {"format": "json", "query": "award-winning"}, get the response "...The consumer key passed was not valid..."

**Step 2** call [getrelatedchannels\_for\_vimeo] with arguments {"category": "film festival", "format": "json"}, get the response "...The consumer key passed was not valid..."

**Step 3** call [download\_stream\_for\_ytstream\_download\_youtube\_videos] with arguments {"is\_id": "UxxajLWwzqY"}, get the response "...'id': 'UxxajLWwzqY', 'title': 'Icona Pop - I Love It (feat. Charli XCX) [OFFICIAL VIDEO]', ..."

**Step 4** call [Finish] with arguments {"return\_type": "give\_answer", "final\_answer": "The streaming link for the YouTube video with ID UxxajLWwzqY is [link]..."}, get the response "successfully giving the final answer"

## Case Study: Reasoning Trajectories from LLaMA with SFT

The reasoning tree of LLaMA with SFT includes 8 paths. There is no successful path since the model finally gives up. All the paths are listed below (the numbers in brackets represent the unique id of each node).

**Path 1** searchvideos\_for\_vimeo (1) -> searchvideos\_for\_vimeo (2) -> searchvideos\_for\_vimeo (3) -> getrelatedpeople\_for\_vimeo (4)

**Path 2** searchvideos\_for\_vimeo (1) -> searchvideos\_for\_vimeo (2) -> searchvideos\_for\_vimeo (3) -> getrelatedchannels\_for\_vimeo (5)

**Path 3** searchvideos\_for\_vimeo (1) -> searchvideos\_for\_vimeo (2) -> searchvideos\_for\_vimeo (6) -> getrelatedpeople\_for\_vimeo (7)

**Path 4** searchvideos\_for\_vimeo (1) -> searchvideos\_for\_vimeo (2) -> searchvideos\_for\_vimeo (6) -> searchvideos\_for\_vimeo (8)

**Path 5** searchvideos\_for\_vimeo (1) -> searchvideos\_for\_vimeo (9) -> searchvideos\_for\_vimeo (10) -> getrelatedpeople\_for\_vimeo (11)

**Path 6** searchvideos\_for\_vimeo (1) -> searchvideos\_for\_vimeo (9) -> searchvideos\_for\_vimeo (10) -> searchvideos\_for\_vimeo (12)

**Path 7** searchvideos\_for\_vimeo (1) -> searchvideos\_for\_vimeo (9) -> searchvideos\_for\_vimeo (13) -> Finish with give up and restart (14)

**Path 8** searchvideos\_for\_vimeo (1) -> getrelatedchannels\_for\_vimeo (15) -> Finish (16)

We observe that because LLaMA with SFT repeatedly tries inaccessible APIs (possibly using different arguments) without first accessing the accessible YouTube API, it finally mistakenly chooses to give up reasoning and is unable to give a partial answer. In contrast, TP-LLaMA successfully calls the YouTube API to provide the best possible answer while using fewer inference steps.

## **C Limitations**

While this work demonstrates promising results, it also has some limitations. First, the performance of our approach relies on the quality of the decision tree. We parse preference pairs from trajectories that experts naturally explore, though the quality of these trajectories still requires evaluation. Manually introducing suboptimal branches at specific nodes might provide a more effective approach. Additionally, our method currently does not compare preferences between steps on failure paths, suggesting room for improved data utilization. Finally, our approach requires inputting all historical information along the path at each reasoning step, which can be time-consuming. Implementing summary steps during reasoning could help streamline interaction text, assist the model in extracting relevant information, and improve reasoning efficiency.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Our main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have discussed the limitations in Appendix C.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our work does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We have fully disclosed all the information needed to reproduce the main experimental results in this paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Our data is available at <https://huggingface.co/datasets/chrissiecsj/ToolPreference>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so No is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We have specified all the training and test details in Section 5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Error bars are not reported because it would be too computationally in our experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We have provide sufficient information on the computer resources needed to reproduce the experiments in Section 5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research we conduct in our papers always complies in all respects with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: We think there is no societal impact of our work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.



- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper does not carry a high risk of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: In section 5, we have cited the original paper. We have also stated the version of the asset is used with a URL.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: Our data is available at <https://huggingface.co/datasets/chrissiecsj/ToolPreference>.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our work does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our work does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.