

# TPTU-V2: BOOSTING TASK PLANNING AND TOOL USAGE OF LARGE LANGUAGE MODEL-BASED AGENTS IN REAL-WORLD SYSTEMS

**Yilun Kong**<sup>12†‡</sup> **Jingqing Ruan**<sup>134†‡</sup> **Yihong Chen**<sup>12†‡</sup> **Bin Zhang**<sup>134†‡</sup>  
**Tianpeng Bao**<sup>1†</sup> **Shiwei Shi**<sup>1†</sup> **Guoqing Du**<sup>1†</sup> **Xiaoru Hu**<sup>1†</sup>  
**Hangyu Mao**<sup>1†✉</sup> **Ziyue Li**<sup>1</sup> **Xingyu Zeng**<sup>1</sup> **Rui Zhao**<sup>1</sup> **Xueqian Wang**<sup>2</sup>

<sup>1</sup>SenseTime Research <sup>2</sup>Shenzhen International Graduate School, Tsinghua University

<sup>3</sup>Institute of Automation, Chinese Academy of Sciences

<sup>4</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences

## ABSTRACT

Large Language Models (LLMs) have demonstrated proficiency in addressing tasks that necessitate a combination of task planning and the usage of external tools, such as weather and calculator APIs. However, real-world complex systems present three prevalent challenges concerning task planning and tool usage: (1) The real system usually has numerous APIs, so it is impractical to feed the descriptions of all APIs to the prompt of LLMs as the token length is limited; (2) the real system is designed for handling complex tasks, and the base LLMs can hardly plan a correct sub-task order and API-calling order for such tasks; (3) Similar semantics and functionalities among APIs in real systems create challenges for both LLMs and even humans in distinguishing between them. In response, this paper introduces a comprehensive framework aimed at enhancing the Task Planning and Tool Usage (TPTU) abilities of LLM-based agents within real-world systems. Our framework comprises three key components designed to address these challenges: (1) the *API Retriever* selects the most pertinent APIs among the extensive API set; (2) *LLM Finetuner* tunes a base LLM to enhance its capability for task planning and API calling; (3) the *Demo Selector* retrieves demonstrations related to hard-to-distinguish APIs, which is further used for in-context learning to boost the final performance. We validate our methods using a real-world industry system and an open-sourced academic dataset, demonstrating the efficacy of each individual component as well as the integrated framework.

## 1 INTRODUCTION

Large language models (LLMs) have exhibited remarkable prowess in various domains of natural language processing (NLP) (Brown et al., 2020; Ouyang et al., 2022; OpenAI, 2023), encompassing language understanding (Devlin et al., 2018; Radford et al., 2023), reasoning (Wei et al., 2022; Kojima et al., 2022), and program synthesis (Liu et al., 2023; Liang et al., 2023a).

However, leveraging LLMs for complex tasks presents formidable challenges. On the one hand, LLMs inherently exhibit limitations in their capabilities, struggling with logical problem-solving, such as mathematics, and facing the risk of stored knowledge quickly becoming outdated as the world evolves. Instructing LLMs to utilize external tools including calculators, calendars, or search engines can prevent them from generating inaccurate information and aid them in effectively addressing problems. On the other hand, integrating these models into complex systems transcends mere task understanding. It demands the ability to break down intricate tasks, manipulate various tools, and engage with users in effective interactions. Several research endeavors, known as LLM-based AI

<sup>†</sup>These authors contribute equally to this work.

<sup>‡</sup>These authors work as research interns at SenseTime Research.

<sup>✉</sup>The corresponding author. Correspondence to [maohangyu@sensetime.com](mailto:maohangyu@sensetime.com)

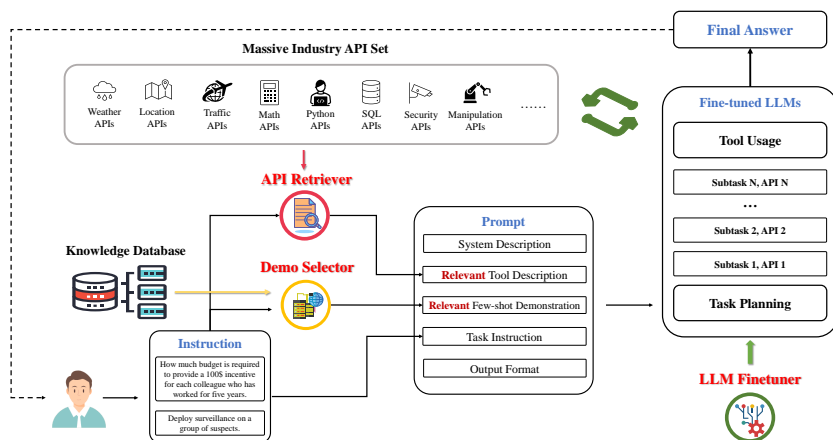


Figure 1: The proposed framework.

Agents (Wang et al., 2023c; Xi et al., 2023), such as AutoGPT<sup>1</sup>, BabyAGI<sup>2</sup>, and ChatGPT-plugins<sup>3</sup>, have made advancements by employing LLMs as central controllers. These endeavors automatically decompose user queries into sub-tasks, execute low-level tool (i.e., API) calls for these sub-tasks, and ultimately resolve the overarching problem.

Despite these advances, LLM-based agents still grapple with pressing challenges in real-world complex applications. Firstly, real-world systems usually have a vast number of APIs, making it impractical to input descriptions of all APIs into the prompt of LLMs due to the token length limitations. Secondly, the real system is designed for handling complex tasks, and the base LLMs (i.e., general LLMs without finetuning on these tasks) often struggle to correctly plan sub-task orders and API-calling sequences for such tasks. Thirdly, beyond the sheer quantity of APIs, a more substantial challenge is that real systems are primarily designed around a core purpose, resulting in the fact that certain APIs may overlap and exhibit similar semantics and functionality, while different in usage. For instance, variations in required input parameters create difficulties in distinguishing and invoking these APIs for LLMs and even humans. How to address these issues could be the critical step for LLM-based agents towards omniscience and omnipotence and being implemented in real-world industry scenarios.

In this paper, we propose a framework, as shown in Figure 1, to improve the **Task Planning** and **Tool Usage** (TPTU) abilities of LLM-based agents in the real-world systems, which is composed of three key components to address the above three challenges: (1) **API Retriever** recalls the APIs that are most relevant to the user’s task from all APIs. The descriptions of these filtered APIs can then be input into LLM as prompts, allowing the LLM to understand and make accurate choices within the filtered API set. (2) **LLM Finetuner** tunes a base LLM so that the finetuned LLM can be more capable of task planning and API calls, especially for domain-specific tasks. (3) **Demo Selector** adaptively retrieves different demonstrations related to hard-to-distinguish APIs, which is further used for in-context learning so that LLM can distinguish the subtle differences in the functions and usages of different APIs. Our main contributions can be summarized as follows:

1. We identify three practical challenges LLM-based agents face in task planning and tool usage within real-world industry scenarios.
2. In response to the three challenges mentioned above, we propose an advanced framework composed of three key components: API Retriever, LLM Finetuner, and Demo Selector.
3. Extensive experiments in real-world industry systems demonstrate the effectiveness our framework. We also validate our methods with open-sourced academic datasets.

<sup>1</sup><https://github.com/Significant-Gravitas/Auto-GPT>

<sup>2</sup><https://github.com/yoheinakajima/babyagi>

<sup>3</sup><https://openai.com/blog/chatgpt-plugins>

## 2 RELATED WORKS

The remarkable capacity for using tools has facilitated the transcendence of human innate physical and cognitive limitations, enhancing our ability to comprehend, plan, and address complex tasks. In turn, the human aptitude for understanding and planning tasks contributes to the judicious selection and usage of appropriate tools. Recently, the swift evolution of LLM has rendered it viable to employ specialized tools and decompose intricate tasks like humans, which inspired significant potential in addressing real-world tasks. Substantial research has been proposed to investigate task planning and tool usage based on LLM separately, however, research that combines these abilities to mutually enhance each other is relatively scarce. TPTU (Ruan et al., 2023) proposes a complete framework that enhances the agent’s ability in task planning and tool utilization for addressing complex tasks. AgentTuning (Zeng et al., 2023) comprehensively considers various capabilities of LLM, not only task planning and tool usage, enhancing the generalized agent capabilities of open-source LLMs themselves while ensuring their general capabilities are not compromised. Some excellent reviews also systematically discuss various aspects of LLM-based AI Agents (Wang et al., 2023c; Xi et al., 2023).

### 2.1 TASK PLANNING

LLMs are pre-trained on huge text corpora and present significant common sense reasoning and multi-task generalization abilities. Prompting is a highly effective method for further harnessing the intrinsic capabilities of LLMs to address various problems (Wei et al., 2022; Kojima et al., 2022). For task planning, prompting facilitates LLMs to break down high-level tasks into sub-tasks (Huang et al., 2022a) and formulate grounded plans (Ahn et al., 2022; Huang et al., 2022b). ReAct (Yao et al., 2022) proposes an enhanced integration of reasoning and action, enabling LLMs to provide a valid justification for action and integrating environmental feedback into the reasoning process. BabyAGI, AgentGPT, and AutoGPT also adopt step-by-step thinking, which iteratively generates the next task by using LLMs, providing some solutions for task automation. However, these methods become problematic as an initial error can propagate along an action sequence, leading to a cascade of subsequent errors. Reflexion (Shinn et al., 2023) incorporates a mechanism for decision retraction, asking LLMs to reflect on previous failures to correct their decision-making. HuggingGPT (Shen et al., 2023) adopts a global planning strategy to obtain the entire sub-task queue within one user query. Besides the above LLM-based studies, previous hierarchical agents, such as SEIHAI (Mao et al., 2022), Juewu-MC (Lin et al., 2021), GITM (Zhu et al., 2023) often resemble the spirit of task planning.

However, in real-world systems, the high-level tasks are more intricate, and the prompting method without enhancing the intrinsic task-planning ability of LLMs can hardly achieve good performance. Thus, in our work, we adopt a fine-tuning mechanism to the planning dataset, along with well-designed prompts, to maximize the ability of task planning.

### 2.2 TOOL USAGE

The great abilities of LLMs in common sense cognition and natural language processing attributes furnish indispensable prerequisites for LLMs to comprehend user intentions and effectively employ tools in tackling intricate tasks (Qin et al., 2023a). Additionally, tool usage can alleviate the inherent limitations of LLMs, encompassing the acquisition of up-to-date information from real-world events, enhanced mathematical computational abilities, and the mitigation of potential hallucinatory phenomena (Mialon et al., 2023).

The tools usage by the LLM can be broadly categorized into two types: physical tools and software tools. In the domain of embodied intelligence (Duan et al., 2022), LLMs directly interact with tangible tools, such as robots, to augment their cognitive abilities, optimize work productivity, and broaden functional capacities (Zhang et al., 2023b; Shah et al., 2023; Brohan et al., 2023; Huang et al., 2022b; Chen et al., 2023b; Driess et al., 2023; Wake et al., 2023; Rana et al., 2023; Song et al., 2022). LLM can also utilize software tools such as search engines (Gou et al., 2020; Borgeaud et al., 2022), mobile (Wang et al., 2023a; Zhang et al., 2023a), Microsoft Office (Li et al., 2023a; Zha et al., 2023), calculators (Chen et al., 2023d; Parisi et al., 2022; Cobbe et al., 2021), deep models (Gupta & Kembhavi, 2023; Chen et al., 2023c) and other versatile APIs (Lu et al., 2023; Gou et al., 2023;

Liang et al., 2023b) to improve model performance or complete complex workflows through flexible control of the software.

However, most of the aforementioned works focus only on specific scenarios, addressing how to choose or use the appropriate tools from a limited set, while agents in real-world scenarios usually have to face various and complex situations, requiring precise selection and usage of the correct tools from an API cloud with massive APIs. Gorilla (Patil et al., 2023) connects LLMs with massive APIs, which are, nonetheless, not real-world APIs and with poor diversity. ToolAlpaca (Tang et al., 2023) builds a tool-using corpus containing 3938 tool-use instances from more than 400 real-world tool APIs spanning 50 distinct categories, but this method focuses on smaller language models. ToolLLM (Qin et al., 2023b) provides a novel and high-quality prompt-tuning dataset, ToolBench, which collects 16464 real-world APIs spanning 49 categories from RapidAPI Hub, covering both single-tool and multi-tool scenarios. TaskMatrix.AI (Liang et al., 2023b) uses LLM as a core system and connects with millions of APIs to execute both digital and physical tasks. The methods above are of great assistance to the tool-learning research community.

To augment LLMs with external tools, most recent methods rely on few-shot prompting with the off-the-shelf LLMs (Patil et al., 2023; Tang et al., 2023; Yao et al., 2023; Wang et al., 2023b; Li et al., 2023b; Xu et al., 2023), but the existing LLMs are not developed for agentic use cases. FireAct (Chen et al., 2023a) proposes a novel approach to fine-tune LLMs with trajectories from multiple tasks and prompting methods and find LLM-based agents are consistently improved after fine-tuning their backbone. ToolLLM (Qin et al., 2023b) uses SFT based on the proposed ToolBench, to transform LLaMa (Touvron et al., 2023) into ToolLLaMa, which demonstrates a remarkable ability to execute complex instructions and generalize to unseen APIs, and exhibits comparable performance to ChatGPT. Inspired by these, we not only design an API Retriever and Demo Selector to serve as an auto-prompter but also employ fine-tuning techniques to further enhance the performance of our framework so that it can address much more complex tasks in real-world scenarios.

### 3 METHODOLOGY

In response to the typical challenges of deploying LLMs within intricate real-world systems, we propose a comprehensive framework that fundamentally bolsters the capabilities of LLMs in Task Planning and Tool Usage (TPTU), which are the cornerstone of the agent’s abilities. In this paper, Task Planning entails generating a step-by-step subtask sequences for the complex task, while Tool Usage requires the agent to select appropriate APIs and execute them correctly to obtain the answer.

#### 3.1 FRAMEWORK OVERVIEW

Our comprehensive framework is engineered to enhance the capabilities of LLMs in TPTU within complex real-world systems, facilitating synergistic collaboration between these two abilities. The framework consists of three pivotal components, depicted in Figure 1. When receiving user instructions, our framework initiates the process by acquiring pertinent APIs and demonstrations through the API Retriever and Demo Selector. These API descriptions and demonstrations constitute a crucial segment of the prompt, which is then input into the fine-tuned LLM. The LLM processes the instruction, leveraging the obtained APIs and demonstrations to derive subtasks and their corresponding API calls. Subsequently, it executes the API for each subtask, obtaining sub-results and ultimately achieving the complete result. Details of the prompt are shown in Appendix A.1.

#### 3.2 API RETRIEVER

In real-world systems, the massive number of APIs for problem-solving poses severe challenges for the integration of LLMs. The token length limitations inherent to LLMs hinder the inclusion of all API descriptions in the model’s prompt, while excessive task-irrelevant API information can impede accurate planning and answer generation, thereby interfering its operational efficiency.

To surmount these challenges, we develop an API Retriever trained to select APIs most relevant to the overall task. These selected APIs are not only necessary for solving the overall task, but also contribute to enhancing the LLM’s comprehension of the task at hand. This, in turn, facilitates more precise segmentation of sub-tasks and the accurate execution of API calls.

The module is a dual-stream architecture employing two Sentence-BERT models (Reimers & Gurevych, 2019) to obtain semantic embeddings of the instruction and API descriptions, separately. It selects the API description closest to the instruction in our trained semantic space. We use the Multiple Negatives Ranking Loss (MNR Loss) (Henderson et al., 2017) to encourage the models to learn a representation space where instructions and relevant APIs are closer to each other. MNR Loss explicitly contrasts the positive pair (an instruction with the relevant API) against multiple negative pairs (with irrelevant API), minimizing the distance between the embeddings of correct instruction-API pairs while maximizing the distance between the embeddings of incorrect pairs, which can be formulated as follows:

$$\mathcal{L} = -\frac{1}{K} \sum_{i=1}^K \log \frac{e^{\text{sim}(s_i, s_i^+)}}{e^{\text{sim}(s_i, s_i^+)} + \sum_{j \neq i} e^{\text{sim}(s_i, s_j^-)}},$$

where  $K$  denotes the batch size,  $s_i$  indicates the sentence embedding of instruction  $i$ , while  $s_i^+$  and  $s_j^-$  represent the embeddings of API descriptions which form the positive and negative pairs corresponding to  $s_i$ , respectively.  $\text{sim}(\cdot)$  is the similarity function that calculates the similarity between the two embeddings, and we employ cosine similarity in the work.

The effectiveness of the API Retriever is also grounded in a meticulous data collection process. We compile a comprehensive set of APIs from diverse external tool services. To ensure our system grasps the relevance of different APIs to various user instructions, we implement a specific annotation process. Human experts and LLMs analyze complex user instructions to identify and annotate the APIs necessary for resolving these instructions. This hybrid approach enriches our dataset with human expertise while leveraging LLMs for efficient data processing, resulting in a qualitatively rich and quantitatively extensive dataset. This forms a robust foundation for the API Retriever.

### 3.3 LLM FINETUNER

While open-sourced LLMs possess strong capabilities, they often encounter limitations due to a lack of specificity and adaptability within complex, specialized, real-world domains. Fine-tuning LLMs on downstream tasks is a prevailing practice to refine their proficiency in addressing specific challenges in these domains. Since the ubiquity and satisfactory performance of existing fine-tuning methods, such as SFT (Ouyang et al., 2022), LoRA (Hu et al., 2021) and QLoRA (Dettrmers et al., 2023), we shift our approach from pioneering new fine-tuning methods to *concentrating on the development of a dataset, expressly curated to enhance the fine-tuning process for real-world systems*.

Specifically, we employ SFT to fine-tune our model and meticulously construct a well-designed dataset with the following characteristics: (1) **Real-world Data:** To accurately mirror real-world scenarios, the dataset is constructed by carefully selecting genuine cases, so that the fine-tuned LLMs can align with the real data distribution in practical use. (2) **Structured Prompt:** The prompts in the dataset are augmented with several key elements, including the system descriptions, API descriptions and demonstrations, which enables the model to generate responses that not only semantically match the input query but also closely align with the functional scope of the available APIs. (3) **Diversity:** To further capture real-world situations, we expand the diversity of the dataset, including prompt diversity user instruction diversity, and output diversity. Both prompt and instruction diversity are crucial for enabling the LLM to navigate the API space with greater precision, particularly when handling complex, multi-faceted user requests. For output diversity, the incorporation of various single-step and multi-step API interactions serves to not only solidify the foundational understanding of API functionalities, but also expose the LLM to more complex sequences of operations commonly encountered in practice. More details can be seen in Appendix A.2.

### 3.4 DEMO SELECTOR

The Demo Selector, serving as an in-context learning method to provide few-shot demonstrations, plays a crucial role in instructing LLMs to distinguish between potentially confusing APIs<sup>4</sup>, execute APIs accurately and disassemble complex tasks. To ensure effective demonstrations and support

<sup>4</sup>The APIs may have similar semantics and functionality because (1) the real system is primarily designed around a core purpose, so some APIs are relevant; (2) when API Retriever is used, the retrieved APIs could be more semantically similar.

these functionalities, we establish a knowledge database from real-world industry scenarios. This knowledge database can be easily compiled, as the routine operations of industry systems naturally generate numerous operationally correct records derived from real-world situations. These records, closely aligned with user instructions, constitute the foundation of our knowledge database.

The Demo Selector module consists of a pre-trained Sentence-BERT and utilizes an embedding search mechanism to select appropriate demonstrations from the knowledge database. In contrast to the API Retriever, this module does not require fine-tuning with new training data, as the API Retriever matches user instructions with API descriptions, while the Demo Selector matches user instructions with task instructions, which inherently share the same semantic space. The detailed processes of selecting demonstrations are as follows:

- (1) **Embedding Generation.** Initially, the user’s query  $Q$  and demonstrations  $d_i$  from the knowledge database  $D$  are transformed into semantic embeddings  $emb(Q)$  and  $emb(d_i)$ .
- (2) **Top- $k$  Task-Level Demos.** We define a similarity threshold  $\Delta$  and employ cosine similarity as the similarity function. In candidates with similarity exceeding the threshold  $sim(emb(Q), emb(d_i)) > \Delta$ , we select the top- $k$  demonstrations  $\{d_{top_1}, d_{top_2}, \dots, d_{top_k}\}$  based on their similarity scores. These are regarded as task-level demonstrations as they are closely related to the specific task at hand.
- (3) **Fallback to API-Level Demos.** If there only exists  $n$  demonstrations with similarity exceeding the threshold, where  $n < k$ , the process degrades to selecting API-level demonstrations from the API collection. The module chooses  $k - n$  API-level demonstrations (i.e., application examples of a specific API) based on the order of API description to supplement in-context demonstrations.

## 4 EXPERIMENTS

In this section, we present comprehensive experiments on unseen instructions and APIs to rigorously evaluate the efficacy and generalizability of our proposed framework. Our experimental methodology is structured to assess the system’s performance in both a real-world scenario and an open-source academic challenge.

### 4.1 DATASETS

Table 1: Comparison between our real-world dataset and notable open-source datasets.

Resource	ToolBench Qin et al. (2023b)	APIBench Patil et al. (2023)	ToolAlpaca Tang et al. (2023)	API-Bank Li et al. (2023b)	Ours
Real-world API	✓	✗	✗	✓	✓
Real-world Instruction	✗	✗	✗	✗	✓
Multi-tool Scenario	✓	✗	✗	✗	✓
Multi-step Reasoning	✓	✗	✓	✓	✓
Manual Construction	✗	✗	✗	✗	✓
Number of APIs	16464	1645	400	53	45
Number of Instances	12657	17002	3938	274	760
Avg. Number of API Calls	2.94	1	1.66	2.76	<b>3.50</b>

**Anonymous Real-world Scenario.** Diverging from the current scholarly focus on studying the ability to choose the right APIs from a plethora of APIs encompassing various functionalities, in real-world systems, more common and challenging problems often revolve around a few core purposes and require multiple tool invocations. It entails choosing the most suitable API from a few dozen APIs, which are closely related in semantics but differ in usage, and orchestrating the correct order for these API calls, enabling bootstrapping solutions for complex problems. Therefore, we constructed a specialized dataset comprising 45 core APIs<sup>5</sup> utilized in industry application, based on a real

<sup>5</sup>During testing, we combine the 45 core APIs with thousands of irrelevant APIs, and our API Retriever can easily filter out these unrelated APIs, which shows that real-world challenges may not align with the academic

Table 2: Performance comparison on real-world scenario, where GTA and AR denote using ground truth APIs and APIs selected by API Retriever respectively, and DS represents utilizing Demo Selector for in-context learning.

Approaches	Accuracy
GPT-3.5 + GTA	70.0%
Claude + GTA	86.7%
InternLM-ft + AR + DS (ours)	<b>96.67%</b>
InternLM	16.70%
InternLM + GTA	38.89%
InternLM + AR	43.33%
InternLM-ft + GTA	80.48%
InternLM-ft + AR	80.34%
InternLM + GTA + DS	95.55%
InternLM-ft + GTA + DS	100%

commercial security system. To align with real-world scenarios, the dataset includes 390 single-call samples and 370 multi-call samples. The multi-call samples involve up to 9 API calls, with an average of 3.5 API calls across the entire dataset, which is larger than many open-source datasets. We meticulously selected real-world instructions, incorporating simple questions with fewer than 10 words and challenging questions with more than 100 words. The detailed statistics are shown in Table 1, and the examples of simple and challenging questions are provided in Appendix A.3.

**Open-source Scenario.** To ensure the generalizability of our approach across a broader spectrum of tasks and its capability to select appropriate APIs from a myriad of options, we also perform experiments on an open-source dataset, ToolBench (Qin et al., 2023b), which also closely resembles real-world applications. It contains 16000+ real-world APIs spanning 49 application categories. Besides the variety and quantity of APIs, it is also well conducted with both single-tool and multi-tool scenarios, as well as several multi-step reasoning traces for each query.

## 4.2 BASELINES

In the real-world scenario, we select both closed-source and open-source LLMs as baselines, including GPT-3.5, Claude (Anthropic, 2023), and InternLM (Team, 2023). In the open-source scenario, our baselines include GPT-3.5 and ToolLLaMA, which tailors for ToolBench.

## 4.3 MAIN EXPERIMENT ON REAL-WORLD SCENARIO

In our real-world dataset, we conduct experiments to assess the efficacy of all proposed modules in our framework. We employ the LLM Finetuner on the open-source InternLM to underscore the importance of fine-tuning in the TPTU framework.

**Main Results** As shown at the top of Table 2, our method significantly outperforms the baselines for TPTU. The remarkable performance is attained through the integration of the finetuned InternLM (InternLM-ft) with both the API Retriever and Demo Selector, achieving an impressive accuracy of 96.67%, a level sufficient for practical application in real-world commercial scenarios.

**API Retriever** We utilize the top-5 APIs recommended by the API Retriever. For the InternLM-ft, employing API Retriever yields comparable results to directly using ground truth APIs (i.e., 80.34% v.s. 80.48%). While in experiments that simultaneously employ the Demo Selector, utilizing API retriever only slightly diminishes the accuracy in the overall framework (i.e., 96.67% v.s. 100%), as the Demo Selector accumulates errors introduced by API Retriever. For the InternLM without finetuning, adopting API Retriever even yields better results (i.e., 43.33% v.s. 38.89%). This can be attributed to the significant influence of the API order in the prompt on the decisions made by LLM with weaker inference capabilities. The relative positioning of APIs within the prompt can affect the LLM’s interpretation of the context and the relationships between different APIs,

notion that the increasing number of APIs makes the problem more challenging. Thus, we only collect the core APIs in our dataset.

Table 3: Performance comparison on open-source scenario, where GTA represents using ground truth API set, AR refers to using API Retriever to select APIs, and ft denotes fine-tuned.

Approaches	Accuracy
ToolLLaMA + GTA	74.3%
ChatGPT + GTA	83.6%
InternLM + GTA	76.67%
InternLM + AR	53.3%
InternLM-ft + AR (ours)	<b>87.6%</b>

ultimately influencing its output Lu et al. (2021). Thus, this module might be a promising approach to automatically retrieve the appropriate APIs as our methods demonstrated.

**LLM Finetuner** Regarding the benefits of fine-tuning, the results clearly demonstrate its advantages. In each experimental configuration, the accuracy of the InternLM-ft is significantly higher than that of the base one (i.e., 80.48% v.s. 38.89%, 80.34% v.s. 43.33%, and 100% v.s. 95.55%). The fine-tuning process tailors the LLM more closely to the specifics of the real-world industry task. It enhances the model’s understanding of the context, leading to more accurate and contextually appropriate API calls.

**Demo Selector** We directly compare the difference between simply using the API usage examples and utilizing the top-5 demonstrations acquired through Demo Selector as in-context demonstrations. During selecting demonstrations, we set the similarity threshold as 0.8. The results show that Demo Selector also has a substantial impact in each set of ablation experiments (i.e., 95.55% v.s. 38.89% and 100% v.s. 80.48%), due to its ability to provide context-rich examples that guide the LLM in making more informed decisions.

In conclusion, our experiments in a real-world setting validate the efficacy of our proposed framework, highlighting the importance of each component and demonstrating our approach is applicable in practical applications.

#### 4.4 MAIN EXPERIMENT ON OPEN-SOURCE DATASET

In the open-source scenario, we tailor our evaluation to focus primarily on the impact of fine-tuning and the API Retriever, considering that building knowledge database for this context do not contribute to addressing real-world industry problems. Therefore, the assessment of the Demo Selector is omitted in this scenario and we simply use the API-level demonstrations as in-context examples.

**Main Results** As shown in Table 3, our framework (InternLM-ft + AR) outperforms two baseline algorithms, achieving an accuracy of 87.6%. This is attributed to the contributions of the API Retriever and fine-tuning. The main reason for the performance decline compared to the industry dataset is that we do not construct a dedicate knowledge database nor introduce the Demo Selector, thus lacking demonstrations of the overall task. It is worth noting that ChatGPT, utilizing ground truth APIs, already achieves satisfactory results, while ToolLLaMA performs slightly worse due to limitations imposed by model size.

**API Retriever** For the base InternLM without fine-tuning, the introduction of the API Retriever results in decreased performance, dropping from 76.67% to 53.3%, which can be attributed to several factors. Firstly, the lower recall precision introduces cumulative errors in the decision-making process. Secondly, in the ToolBench dataset, there are numerous APIs in ground truth API set that can be replaced by semantically similar but functionally different APIs. While API Retriever may select many such APIs, the lack of Demo Selector results in a reduced score due to a limited understanding of the overall API call flow.

**LLM Finetuner** Upon implementing fine-tuning in this scenario, an enhancement in performance is observed, with the finetuned LLM combined with the API Retriever reaching an execution accuracy of 86.7%. This improvement underscores the effectiveness of fine-tuning in adapting the LLM to the specific characteristics and challenges of the open-source environment. The fine-tuning process likely helps the model better understand the nuances of the available APIs and how they correlate with different tasks, resulting in more accurate API calls and decision-making.



Table 4: The performance, based on GTA, of InternLM fine-tuned on datasets with different feature ablations.

Training Dataset	Accuracy
w/o finetuning	20.0%
Real-world Data	0%
+ Structured Prompt	26.7%
+ Structured Prompt + Diversity	80.5%

In summary, the open-source scenario highlights the nuanced impacts of our framework’s components. In our framework, employing only the API Retriever and LLM Finetuner modules can similarly enable the model to outperform the baselines. Lacking the Demo Selector, however, results in a significant performance decline. This indicates that, in real-world applications, it is advisable to construct knowledge database and adopt Demo Selector for better performance.

#### 4.5 EXPERIMENTS ON API RETRIEVER

To further elucidate the factors contributing to the effectiveness of the API Retriever, this section focuses specifically on its ability to select correct APIs as well as the characteristics of the training dataset for API Retriever.

To safeguard the API Retriever from overfitting to the real-world dataset and enhance its generalizability, we employ both the real-world dataset and ToolBench for training. This comprehensive training set encompasses a total of 8330 data points, ensuring a more robust and adaptable performance.

One key factor for the strong performance of API Retriever in the entire framework is its high precision in recalling the correct APIs, which ensures the minimal deviation between the retrieved APIs and the ground truth. Moreover, the ranked order of retrieved APIs, based on similarity, further enhances the overall performance. In particular, this module achieves a Recall@5 precision of 84.64% and Recall@10 precision of 98.47% in the combination of core and irrelevant APIs in the real-world dataset. While in the open-source dataset, it achieves a Recall@5 precision of 77.92% and Recall@10 precision of 87.54% , which slightly fall short compared with the results in the industry scenario. After subsequent experiments, we ultimately adopt the strategy of recalling the top-5 APIs.

#### 4.6 EXPERIMENTS ON LLM FINETUNER

As shown in Section 3.3, we focus on developing a meticulous dataset to enhance the finetuning process in real-world TPTU. In this section, we conduct ablation experiments on the dataset’s characteristics separately to demonstrate the crucial role of our designed features. The results, presented in Table 4, indicate that dataset with all the characteristics significantly improves the effect of fine-tuning, achieving an accuracy of 80.5%. The lack of diversity in the dataset results in a notable decline in model performance, with a fine-tuning accuracy of only 26.7%. This is because the model may tend to rotely memorize the API call solutions for different instructions. Fine-tuning solely on raw real-world data can backfire, for the model may overfit to specific issues within the training set due to the constraints imposed by the limited quantity of genuine data.

## 5 CONCLUSION

In this paper, we present a comprehensive framework to augment the capabilities of LLMs in real-world scenarios, with a specific focus on **Task Planning and Tool Usage (TPTU)**. Our approach, which integrates API Retriever, LLM Finetuner, and Demo Selector, has been rigorously validated in both a real-world scenario and an open-source setting. The results demonstrate that fine-tuning LLMs with a curated dataset can significantly improve their effectiveness in executing real-world tasks. The API Retriever and Demo Selector also prove indispensable, particularly in improving the model’s decision-making accuracy and adaptability. This research not only highlights the potential of LLMs in practical applications but also establishes a foundation for future advancements in this field. We anticipate witnessing further success of LLM-based agents in TPTU in real-world applications.

## REFERENCES

- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- AI Anthropic. Introducing claude, 2023.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. Improving language models by retrieving from trillions of tokens. In *International conference on machine learning*, pp. 2206–2240. PMLR, 2022.
- Anthony Brohan, Yevgen Chebotar, Chelsea Finn, Karol Hausman, Alexander Herzog, Daniel Ho, Julian Ibarz, Alex Irpan, Eric Jang, Ryan Julian, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*, pp. 287–318. PMLR, 2023.
- 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. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao. Fireact: Toward language agent fine-tuning. *arXiv preprint arXiv:2310.05915*, 2023a.
- Boyuan Chen, Fei Xia, Brian Ichter, Kanishka Rao, Keerthana Gopalakrishnan, Michael S Ryoo, Austin Stone, and Daniel Kappler. Open-vocabulary queryable scene representations for real world planning. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11509–11522. IEEE, 2023b.
- Liangyu Chen, Bo Li, Sheng Shen, Jingkan Yang, Chunyuan Li, Kurt Keutzer, Trevor Darrell, and Ziwei Liu. Language models are visual reasoning coordinators. In *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*, 2023c.
- Zhipeng Chen, Kun Zhou, Beichen Zhang, Zheng Gong, Wayne Xin Zhao, and Ji-Rong Wen. Chatcot: Tool-augmented chain-of-thought reasoning on chat-based large language models. *arXiv preprint arXiv:2305.14323*, 2023d.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*, 2023.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.
- Jiafei Duan, Samson Yu, Hui Li Tan, Hongyuan Zhu, and Cheston Tan. A survey of embodied ai: From simulators to research tasks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 6(2):230–244, 2022.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, and Weizhu Chen. Critic: Large language models can self-correct with tool-interactive critiquing. *arXiv preprint arXiv:2305.11738*, 2023.
- Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14953–14962, 2023.

- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. Retrieval augmented language model pre-training. In *International conference on machine learning*, pp. 3929–3938. PMLR, 2020.
- Matthew Henderson, Rami Al-Rfou, Brian Strope, Yun-Hsuan Sung, László Lukács, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. Efficient natural language response suggestion for smart reply. *arXiv preprint arXiv:1705.00652*, 2017.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pp. 9118–9147. PMLR, 2022a.
- Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*, 2022b.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213, 2022.
- Hongxin Li, Jingran Su, Yuntao Chen, Qing Li, and Zhaoxiang Zhang. Sheetcopilot: Bringing software productivity to the next level through large language models. *arXiv preprint arXiv:2305.19308*, 2023a.
- Minghao Li, Feifan Song, Bowen Yu, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. Api-bank: A benchmark for tool-augmented llms. *arXiv preprint arXiv:2304.08244*, 2023b.
- Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9493–9500. IEEE, 2023a.
- Yaobo Liang, Chenfei Wu, Ting Song, Wenshan Wu, Yan Xia, Yu Liu, Yang Ou, Shuai Lu, Lei Ji, Shaoguang Mao, et al. Taskmatrix. ai: Completing tasks by connecting foundation models with millions of apis. *arXiv preprint arXiv:2303.16434*, 2023b.
- Zichuan Lin, Junyou Li, Jianing Shi, Deheng Ye, Qiang Fu, and Wei Yang. Juewu-mc: Playing minecraft with sample-efficient hierarchical reinforcement learning. *arXiv preprint arXiv:2112.04907*, 2021.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *arXiv preprint arXiv:2305.01210*, 2023.
- 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. *arXiv preprint arXiv:2304.09842*, 2023.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786*, 2021.
- Hangyu Mao, Chao Wang, Xiaotian Hao, Yihuan Mao, Yiming Lu, Chengjie Wu, Jianye Hao, Dong Li, and Pingzhong Tang. Seihai: A sample-efficient hierarchical ai for the minerl competition. In *Distributed Artificial Intelligence: Third International Conference, DAI 2021, Shanghai, China, December 17–18, 2021, Proceedings 3*, pp. 38–51. Springer, 2022.
- Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, et al. Augmented language models: a survey. *arXiv preprint arXiv:2302.07842*, 2023.

- OpenAI. Gpt-4 technical report. 2023.
- 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. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022.
- Aaron Parisi, Yao Zhao, and Noah Fiedel. Talm: Tool augmented language models. *arXiv preprint arXiv:2205.12255*, 2022.
- 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.
- Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Yufei Huang, Chaojun Xiao, Chi Han, et al. Tool learning with foundation models. *arXiv preprint arXiv:2304.08354*, 2023a.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. Toollm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*, 2023b.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, pp. 28492–28518. PMLR, 2023.
- Krishan Rana, Jesse Haviland, Sourav Garg, Jad Abou-Chakra, Ian Reid, and Niko Suenderhauf. Sayplan: Grounding large language models using 3d scene graphs for scalable task planning. *arXiv preprint arXiv:2307.06135*, 2023.
- Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*, 2019.
- Jingqing Ruan, Yihong Chen, Bin Zhang, Zhiwei Xu, Tianpeng Bao, Guoqing Du, Shiwei Shi, Hangyu Mao, Xingyu Zeng, and Rui Zhao. Tptu: Task planning and tool usage of large language model-based ai agents. *arXiv preprint arXiv:2308.03427*, 2023.
- Dhruv Shah, Błażej Osiniński, Sergey Levine, et al. Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Conference on Robot Learning*, pp. 492–504. PMLR, 2023.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580*, 2023.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. *arXiv preprint arXiv:2212.04088*, 2022.
- Qiaoyu Tang, Ziliang Deng, Hongyu Lin, Xianpei Han, Qiao Liang, and Le Sun. Toolalpaca: Generalized tool learning for language models with 3000 simulated cases. *arXiv preprint arXiv:2306.05301*, 2023.
- InternLM Team. Internlm: A multilingual language model with progressively enhanced capabilities. <https://github.com/InternLM/InternLM>, 2023.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

- Naoki Wake, Atsushi Kanehira, Kazuhiro Sasabuchi, Jun Takamatsu, and Katsushi Ikeuchi. Chatgpt empowered long-step robot control in various environments: A case application. *arXiv preprint arXiv:2304.03893*, 2023.
- Bryan Wang, Gang Li, and Yang Li. Enabling conversational interaction with mobile ui using large language models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–17, 2023a.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023b.
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. A survey on large language model based autonomous agents. *arXiv preprint arXiv:2308.11432*, 2023c.
- 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. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*, 2023.
- Qiantong Xu, Fenglu Hong, Bo Li, Changran Hu, Zhengyu Chen, and Jian Zhang. On the tool manipulation capability of open-source large language models. *arXiv preprint arXiv:2305.16504*, 2023.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*, 2023.
- Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. Agenttuning: Enabling generalized agent abilities for llms. *arXiv preprint arXiv:2310.12823*, 2023.
- Liangyu Zha, Junlin Zhou, Liyao Li, Rui Wang, Qingyi Huang, Saisai Yang, Jing Yuan, Changbao Su, Xiang Li, Aofeng Su, et al. Tablegpt: Towards unifying tables, nature language and commands into one gpt. *arXiv preprint arXiv:2307.08674*, 2023.
- Danyang Zhang, Lu Chen, and Kai Yu. Mobile-env: A universal platform for training and evaluation of mobile interaction. *arXiv preprint arXiv:2305.08144*, 2023a.
- Weiyi Zhang, Yushi Guo, Liting Niu, Peijun Li, Chun Zhang, Zeyu Wan, Jiayang Yan, Fasih Ud Din Farrukh, and Debing Zhang. Lp-slam: Language-perceptive rgb-d slam system based on large language model. *arXiv preprint arXiv:2303.10089*, 2023b.
- Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li, Lewei Lu, Xiaogang Wang, et al. Ghost in the minecraft: Generally capable agents for open-world environments via large language models with text-based knowledge and memory. *arXiv preprint arXiv:2305.17144*, 2023.

## A MORE DETAILS OF THE METHOD

### A.1 SPECIFIC COMPONENTS IN THE PROMPT

The specific formats of each component of the prompt are shown in Figure 2.

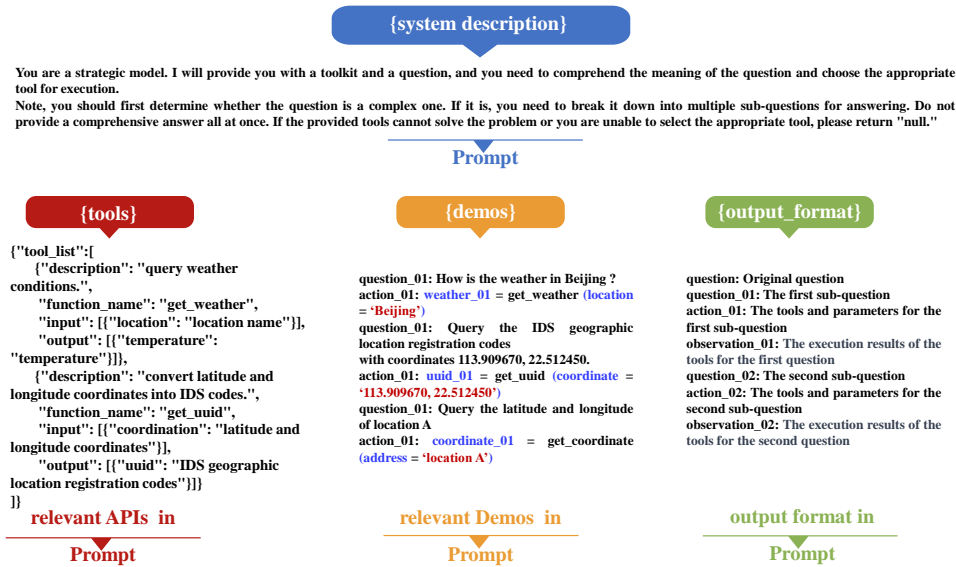


Figure 2: Demonstrations of the specific formats of each component in the prompt.

Real-world Data	Structured Prompt	Diversity
<p><b>"Input":</b> "There is a... icon in the Person File Search — File Details — Personnel File Details interface. Click to pop up the [Historical Modification Record] entrance."</p> <p><b>"Output":</b> "Holographic Profile — Personal Profile Details."</p>	<p><b>"System description":</b> " You are a system decision-making expert. Below I will provide a complex issue about public security..."</p> <p>.....</p> <p><b>"Input":</b> "There is a... icon in the Person File Search — File Details — Personnel File Details interface. Click to pop up the [Historical Modification Record] entrance."</p> <p><b>"Output":</b> "Holographic Profile — Personal Profile Details."</p>	<p><b>"System description":</b> "You are an expert in system decision-making. I will now present a complex matter related to public security..."</p> <p>.....</p> <p><b>"Input":</b> "How to create a new portrait library and then upload face photos?"</p> <p><b>"Output":</b> "1. Library Management — Creation; 2. Library Assistant — Upload."</p>

Figure 3: Demonstrations for the features in our dataset.

## A.2 DEMONSTRATIONS OF LLM FINETUNER DATASET

To ensure readers can clearly understand each characteristics in the dataset, we provide a demonstration for each one, as shown in Figure 3. Each design of the features is intended to incrementally refine the LLM’s ability to parse user inputs, understand the context, and generate precise API calls. Finetuning LLMs on these datasets can enhance the ability of LLMs to solve specific real-world tasks. By systematically evaluating the model’s output against these varied fine-tuning paradigms, we enhance its competency in delivering high-quality, contextually appropriate responses in the domain of API interaction. The insights obtained from the iterative development of these datasets demonstrate the critical importance of dataset quality and construction in the fine-tuning process.

In the details of diversity features, for the prompt diversity, we randomly shuffle API orders and add irrelevant APIs to decrease the risk of over-fitting; for instruction diversity, we replace the original user instruction with similar-meaning instructions by means like rewriting-by-LLMs, synonym substitution, and loop-back translation to make LLMs more robust to different user instructions during inference. For output diversity, besides simple single-step API interactions, which solidify the foundational understanding of API functionalities, we meticulously select and construct multi-step API calls, which introduce the LLM to more complex sequences of operations that are commonly encountered in practice.

### A.3 EXAMPLES OF THE QUESTIONS IN OUR REAL-WORLD DATASET

In order to facilitate a better understanding of the real-world instructions in our dataset, and without compromising the confidentiality of proprietary datasets, we present a simple question and a complex question for illustration.

The following is the simple questions in our dataset:

- Implementing surveillance on a group of individuals.

While in real-world industry scenarios, there are numerous complex problems, with lengths potentially exceeding 100 and comprising multiple sub-problems. To enhance our framework's ability to address these issues, we have carefully selected a variety of challenging problems. The following serves as example of these challenging problems:

- In the routine investigation work of criminal detectives, they typically conduct preliminary analysis to identify a category of suspects. Subsequently, they need to track, inspect, and control the identified targets. During this phase, they analyze associated clues and information related to the suspects, ultimately formulating a comprehensive arrest plan. What modules do you think I would be involved in, and could you list the modules in order?