

# Try, Check and Retry: A Divide-and-Conquer Framework for Boosting Long-context Tool-Calling Performance of LLMs

Anonymous ACL submission

## Abstract

Tool-calling empowers Large Language Models (LLMs) to interact with external environments. However, current methods often struggle to handle massive and noisy candidate tools in long-context tool-calling tasks, limiting their real-world application. To this end, we propose **Tool-DC**, a **Divide-and-Conquer** framework for boosting tool-calling performance of LLMs. The core of Tool-DC is to reduce the reasoning difficulty and make full use of self-reflection ability of LLMs via a “Try-Check-Retry” paradigm. Specifically, Tool-DC involves two variants: 1) the training-free Tool-DC (TF), which is plug-and-play and flexible; 2) the training-based Tool-DC (TB), which is more inference-efficient. Extensive experiments show that both Tool-DC methods outperform their counterparts by a clear margin. Tool-DC (TF) brings up to **+25.10%** average gains against the baseline on BFCL and ACEBench benchmarks, while Tool-DC (TB) enables Qwen2.5-7B to achieve comparable or even better performance than proprietary LLMs, *e.g.*, OpenAI o3 and Claude-Haiku-4.5.

## 1 Introduction

While Large Language Models (LLMs) have shown remarkable successes (Ouyang et al., 2022; Achiam et al., 2023; Touvron et al., 2023), they often struggle to deal with the latest information and suffer from factual hallucinations, due to static parameters and closed knowledge boundaries. In response to this issue, the “*tool-calling*” paradigm has emerged, *i.e.*, the integration of external tools and application programming interfaces (APIs) enables LLMs to tackle complex, real-world scenarios (Schick et al., 2023; Qin et al., 2024).

The common tool-calling involves several processes. Specifically, given a large number of candidate tools, LLMs need to first select the appropriate tools and then formulate precise input arguments.




Figure 1: **Performance comparison on BFCL (Patil et al., 2025) with different candidate tool scales.** We see that as the number of candidate tools increases, all models’ performance degrades significantly, whereas our Tool-DC method can effectively mitigate this issue.

Intuitively, the performance of tool-calling is sensitive to the number of candidate tools. As shown in our preliminary experiments (Figure 1), when scaling the number of candidate tools from <10 to 20, all LLMs suffer from performance degradation, especially for the smaller models. We conjecture that, on the one hand, ❶ *LLMs may fall short in dealing with long context information* caused by a large number of candidate tools. On the other hand, ❷ *confusing tools with similar semantics but different argument descriptions* could affect the argument-filling performance of LLMs.

There are some existing works involving analyzing and addressing these problems. For the first problem, a straightforward approach is to leverage an extra retriever to select the related tool subset (Qin et al., 2024; Patel et al., 2025). While effectively reducing the length of context, it highly relies on the performance of extra retriever (Robertson et al., 2009; Chen et al., 2024). More specifically, when the retriever fails in selecting the golden tools, LLMs struggle to predict the correct results due to the absence of golden tools in the context. For the second problem, Cui et al. (2025) reveal that many

066 tool-calling errors can be efficiently identified by  
067 the structured error checklists. Thus, they manually  
068 construct a global error checklist and use it to guide  
069 the precise tool-calling via an In-Context Learning  
070 (ICL) (Brown et al., 2020) manner. Although effective,  
071 the manually defined error checklist is not flexible  
072 and is difficult to cover all errors. Thus, there  
073 arises a question: *can we explore a more effective  
074 and efficient framework to boost the long-context  
075 tool-calling performance of LLMs?*

076 To achieve this goal, we propose **Tool-DC**, a  
077 **Divide-and-Conquer** framework for better LLMs  
078 tool-calling, which contains two variants: 1) a  
079 training-free approach, Tool-DC (TF) and 2) a  
080 training-based approach, Tool-DC (TB). The core  
081 of Tool-DC is a *Try-Check-Retry* paradigm. Specifically,  
082 in the processes of Tool-DC (TF), we first  
083 split the total candidate tools into several groups  
084 and perform parallel tool-calling (*Try*), then verify  
085 the correctness of parallel inference results  
086 via strict schema constraints (*Check*), and lastly  
087 aggregate the validated results to refine global  
088 tool-calling accuracy (*Retry*). By doing so, we  
089 can not only reduce the length of context and  
090 reasoning difficulty, but also take full advantage  
091 of self-reflection capabilities of LLMs for better  
092 tool-calling. Moreover, considering the inference  
093 latency and deployment costs of Tool-DC (TF),  
094 we further propose to internalize the “Try-Check-  
095 Retry” decision capabilities into the model parameters.  
096 In practice, in Tool-DC (TB), we first  
097 construct a high-quality Chain-of-Thought (CoT)  
098 dataset by collecting the correct reasoning traces  
099 during the stage of Tool-DC (TF), and then use it  
100 to fine-tune the LLMs.

101 We conducted extensive experiments on two  
102 widely-used benchmarks, *i.e.*, BFCL (Patil et al.,  
103 2025) and ACEBench (Chen et al., 2025a), across  
104 multiple LLMs. Notably, considering that the number  
105 of candidate tools in the existing benchmarks  
106 is generally insufficient (<10), which is inconsistent  
107 with real-world scenarios, we introduce an *Extended  
108 Setting* by scaling the number of candidate  
109 tools to simulate real-world applications. Empirical  
110 results show that our Tool-DC (TF) consistently  
111 outperforms the vanilla tool-calling methods  
112 among all settings, especially for the small models  
113 in the extended setting, *i.e.*, up to **+25.10%** average  
114 gains. More encouragingly, with the help of Tool-  
115 DC (TB), Qwen2.5-7B (Yang et al., 2024) achieves  
116 an **83.16%** overall score on BFCL, surpassing the  
117 powerful OpenAI o3 and Claude-Haiku-4.5.

**Contributions.** To summarize, our contributions  
are three-fold: (1) We propose a divide-and-  
conquer framework (Tool-DC) for boosting the  
tool-calling performance of LLMs in the long-  
context and confusing tools scenarios. (2) Tool-  
DC provides two variants, where the training-free  
approach is plug-and-play and flexible, and the  
training-based approach has a higher inference efficiency.  
(3) Extensive experiments show that Tool-  
DC outperforms the vanilla method by a clear margin,  
*e.g.*, bringing up to **+25.10%** average gains  
against the training-free baseline.

## 2 Related Works

Equipping LLMs with external tools extends their  
capabilities to interact with real-world environments  
(Qu et al., 2025; He et al., 2025; Liao et al.,  
2025). However, the manner of providing all tool  
descriptions in the context is sub-optimal, as some  
candidate tools could be noisy and redundant (Chen  
et al., 2025b; Liang et al., 2025). There are many  
existing works aiming to improve the performance  
of tool-calling, which can be classified into training-  
free and training-based methods. The former opti-  
mize inference via retrieval and constraint injection  
but remain limited by the capabilities of the frozen  
base model (Chen et al., 2024; Zheng et al., 2024b;  
Hsieh et al., 2023; Ruan et al., 2023; Cui et al.,  
2025; Dang et al., 2025). The latter align models  
through supervised fine-tuning or reinforcement  
learning yet often suffer from high training costs  
and data synthesis bottlenecks (Huang et al., 2025;  
Hao et al., 2025; Zhang et al., 2025a; Qian et al.,  
2025; Zhang et al., 2025b). While effectively, these  
methods still struggle to deal with long-context  
tool-calling scenarios, where the number of candi-  
date tools is large and many tools are similar but  
with different argument descriptions.

There are rare works involving addressing the  
above problems in the tool-calling field (Cui et al.,  
2025; Dang et al., 2025; Moon et al., 2024).  
Among which, Cui et al. (2025) attempt to miti-  
gate tool-calling errors via manually designing  
a global error checklist. Despite its effectiveness,  
this method needs to carefully design a high-quality  
checklist, and still struggles with the long-context  
problem. Different from it, we propose a divide-  
and-conquer framework, which can effectively re-  
duce the reasoning difficulty and make full use of  
the self-reflection capabilities of LLMs.

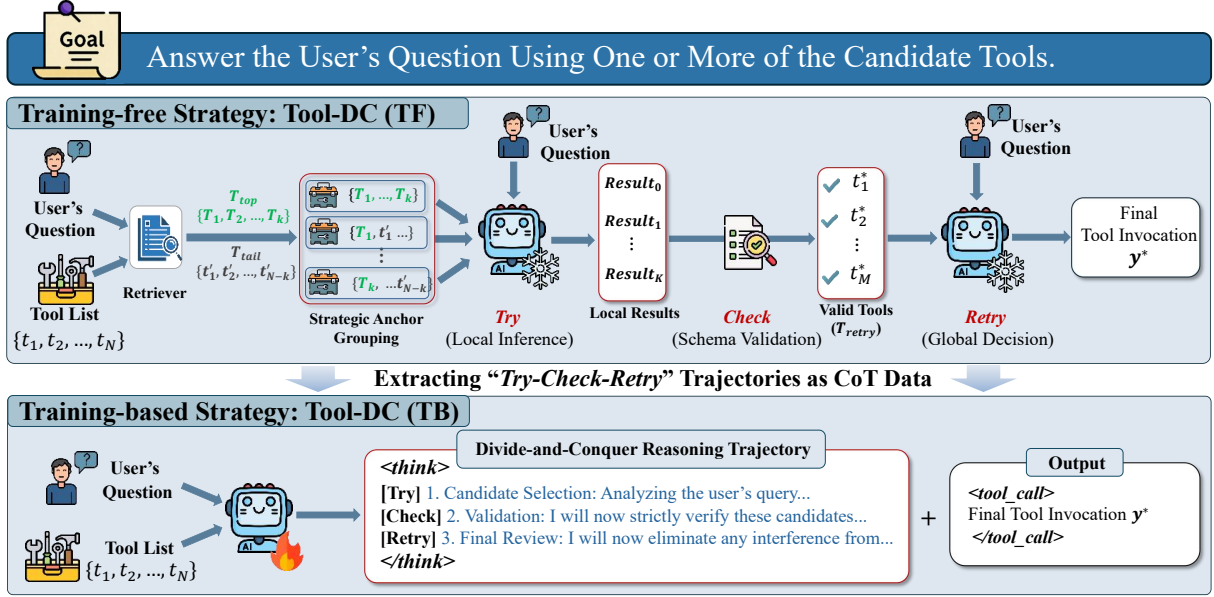


Figure 2: **Overview of our Tool-DC framework.** The training-free strategy employs a “Try-Check-Retry” pipeline to reduce the reasoning difficulty, while the training-based strategy leverages the prior reasoning trajectories to internalize this divide-and-conquer paradigm into model parameters via fine-tuning.

### 3 Methodology

We propose **Tool-DC**, a Divide-and-Conquer framework for boosting the long-context tool-calling performance of LLMs. As shown in Figure 2, Tool-DC contains two variants: 1) the training-free Tool-DC (TF), which reduces the reasoning difficulty via an explicit “Try-Check-Retry” pipeline; 2) the training-based Tool-DC (TB), which internalizes this divide-and-conquer paradigm into model parameters via fine-tuning.

#### 3.1 Problem Formulation

Let  $q$  denote a user query and  $\mathcal{T} = \{t_1, t_2, \dots, t_N\}$  represent a library of  $N$  candidate tools. The objective of tool-calling is to enforce the model  $\mathcal{M}$  to generate a sequence of tool invocations  $\mathbf{y} = \{y_1, y_2, \dots\}$ , where each invocation  $y_i = (t, \alpha)$  consists of a selected tool  $t \in \mathcal{T}$  and its corresponding arguments  $\alpha$ . A major challenge in real-world scenarios is the scale of  $N$ . Dealing with a vast search space makes directly modeling  $P(\mathbf{y}|q, \mathcal{T})$  ineffective, as the substantial number of irrelevant candidates hinders the model’s reasoning capabilities. In response to this problem, our Tool-DC aims to partition the global space  $\mathcal{T}$  into manageable subspaces  $\mathcal{S} = \{S_1, \dots, S_K\}$ . This decomposition strategy can reduce the decision complexity and facilitate precise tool-calling.

#### 3.2 Tool-DC (TF): Training-free Approach

In this part, we introduce the details of our training-free approach, Tool-DC (TF), which involves the following three stages:

##### 3.2.1 Try: Grouping and Local Inference

Considering that invoking tools from the global library  $\mathcal{T}$  suffers from a low signal-to-noise ratio and interference among tools, in the first stage, we propose to isolate decision-making within low-noise subspaces via Strategic Anchor Grouping and Local Inference, which is illustrated in Figure 3.

**Strategic Anchor Grouping.** Similar to many prior works (Qin et al., 2024; Patel et al., 2025), we first retrieve Top- $K$  relevant tools  $\mathcal{T}_{top}$  via an extra retriever, e.g., BM25 (Robertson et al., 2009). Then, we further construct  $K$  parallel groups  $\{S_1, \dots, S_K\}$  via the Strategic Anchor Grouping strategy. In practice, to obtain the subspace  $S_i$ , we take the  $i$ -th tool from  $\mathcal{T}_{top}$  as an anchor and add a disjoint subset of distractors from  $\mathcal{T}_{tail} = \mathcal{T} \setminus \mathcal{T}_{top}$ . Such a grouping method can not only grant tools missed by the initial retrieval (contained in  $\mathcal{T}_{tail}$ ) a chance to be evaluated, but also simultaneously decouple highly relevant candidates to prevent confusion. Notably, we preserve the original Top- $K$  set as a distinct group  $S_0 = \mathcal{T}_{top}$ .

**Local Inference.** For each subspace  $S_j \in \{S_0, S_1, \dots, S_K\}$ , the model  $\mathcal{M}$  performs local inference and outputs an initial tool invocation or a

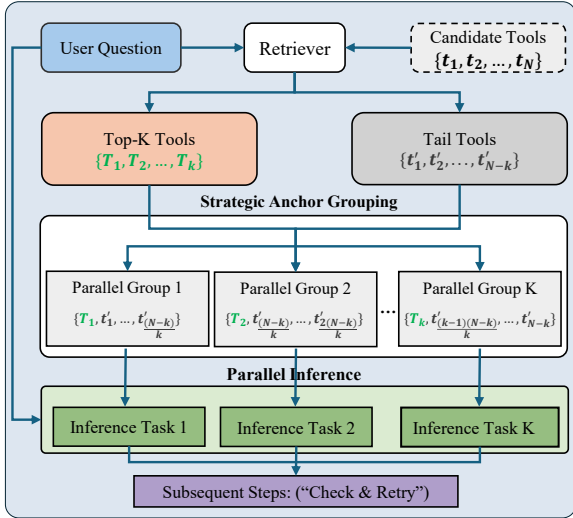


Figure 3: **Illustration of Try stage in Tool-DC (TF).** By splitting the total candidate tools into several parallel groups, Tool-DC (TF) can reduce the length of context and reasoning difficulty effectively.

null token:

$$o_j = \mathcal{M}(q, S_j), \quad (1)$$

where  $j \in \{0, 1, \dots, K\}$  denotes the index of subspace, and  $o_j$  represents the output ( $t_j, \alpha_j$ ) or the null token  $\emptyset$ . This process reduces the search space complexity and enables the model to generate more precise tool arguments.

### 3.2.2 Check: Schema Consistency Validation

Since the generated tool calls usually suffer from hallucinations, *e.g.*, invoking non-existent functions or hallucinating arguments, we leverage a rule-based *Consistency Validator*  $\mathcal{F}(\cdot)$  to detect whether the tool calls are valid. Inspired by Cui et al. (2025), we design several schema constraints associated with the tools. In particular, the validity of each predicted output  $o_j$  is determined as:

$$\mathbb{I}_C(o_j) = \begin{cases} 1 & \text{if } o_j \neq \emptyset \wedge \mathcal{F}(o_j, \text{Schema}(S_j)) \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where  $\text{Schema}(\cdot)$  refers to the schema constraints that verify the validity from three dimensions:

- **Function Name Validity:** The name of invoked function  $t_j$  must explicitly exist in the defined tool set.
- **Argument Key Verification:** The keys in  $\alpha_j$  must match the defined arguments, and all required parameters must be present.

- **Data Type Consistency:** The values assigned to arguments must adhere to the defined data types (*e.g.*, string, integer, boolean).

As a result, this process yields a refined set of valid candidates  $\mathcal{V} = \{o_j \mid \mathbb{I}_C(o_j) = 1\}$ .

### 3.2.3 Retry: Global Aggregation and Decision

Lastly, to take advantage of the self-reflection ability of LLMs, we introduce a refinement mechanism that utilizes the above validated candidates to obtain a more precise global decision. Specifically, we collect a refined candidate tool set  $\mathcal{T}_{\text{retry}}$  by retrieving the original definitions of the tools presented in the valid set  $\mathcal{V}$ :

$$\mathcal{T}_{\text{retry}} = \{t \in \mathcal{T} \mid \exists (t, \alpha) \in \mathcal{V}\}. \quad (3)$$

The refined subset  $\mathcal{T}_{\text{retry}}$  is provided in the context, which is then fed into the  $\mathcal{M}$  to obtain the final tool invocation  $\mathbf{y}^*$ :

$$\mathbf{y}^* = \mathcal{M}(q, \mathcal{T}_{\text{retry}}). \quad (4)$$

In general, by splitting the total candidate tools into several subspaces and using the consistency validator to filter the invalid tools, we can obtain a more relevant subset of candidate tools. Furthermore, we employ the prior feedback to enforce the model to self-refine its output, thus leading to better tool-calling performance.

### 3.3 Tool-DC (TB): Training-based Approach

Although the proposed Tool-DC (TF) method is plug-and-play and flexible, it requires multiple forward passes and will lead to some inference latency. Thus, as an alternative, we additionally propose a training-based method, Tool-DC (TB), which aims to internalize the divide-and-conquer paradigm into model parameters via fine-tuning.

**Data Construction.** The key of Tool-DC (TB) is to construct the CoT training data. In practice, we mainly follow the pipeline of Tool-DC (TF) to collect the correct reasoning trajectories by using the xlam-function-calling-60k dataset (Zhang et al., 2025a) as raw training corpus  $\mathcal{D}_{\text{raw}}$ . To reduce dependence on external retrievers, we replace the group-based approach with an enumeration strategy, *i.e.*, the number of groups is equal to the number of all candidate tools ( $K = N$ ). The pipeline of CoT data construction is shown in Algorithm 1. Notably, for ease of illustration, we simplify the reasoning template in the algorithm. The detailed template is shown in Figure 11 of Appendix A.7.

---

**Algorithm 1** Data Construction for Tool-DC (TB)

```
1: Input: Raw training dataset  $\mathcal{D}_{\text{raw}}$ , library  $\mathcal{T}$ , model  $\mathcal{M}$ 
2: Output: CoT training dataset  $\mathcal{D}_{\text{CoT}}$ 
3:  $\mathcal{D}_{\text{CoT}} \leftarrow \emptyset$ 
4: for  $(q, \mathbf{y}_{\text{gt}}) \in \mathcal{D}_{\text{raw}}$  do
5:    $\mathcal{T}_{\text{valid}} \leftarrow \emptyset$ 
   // Phase 1: Try & Check
6:   for  $t \in \mathcal{T}$  do
7:      $o \leftarrow \mathcal{M}(q, \{t\})$  ▷ Local inference
     // Check: Name, Args, and Type
8:     if  $\mathbb{I}_{\mathcal{C}}(o) = 1$  then
9:        $\mathcal{T}_{\text{valid}} \leftarrow \mathcal{T}_{\text{valid}} \cup \{t\}$ 
10:    end if
11:  end for
12:  if  $\mathcal{T}_{\text{valid}} = \emptyset$  then continue
13:  end if
   // Phase 2: Retry
14:   $\mathbf{y}^* \leftarrow \mathcal{M}(q, \mathcal{T}_{\text{valid}})$ 
   // Phase 3: Rationale Synthesis
15:  if  $\mathbf{y}^* = \mathbf{y}_{\text{gt}}$  then
16:    Synthesize rationale  $R$  via the template:
17:    “1. Candidate Selection: ...
    2. Validation: ...
    3. Final Review: ...”
18:     $\text{CoT} \leftarrow \langle \text{think} \rangle R \langle / \text{think} \rangle \langle \text{tool} \rangle \mathbf{y}^* \langle / \text{tool} \rangle$ 
19:     $\mathcal{D}_{\text{CoT}} \leftarrow \mathcal{D}_{\text{CoT}} \cup \{(q, \mathcal{T}, \text{CoT})\}$ 
20:  end if
21: end for
22: return  $\mathcal{D}_{\text{CoT}}$ 
```

---

**Optimization.** After obtaining the CoT training dataset, denoted as  $\mathcal{D}_{\text{CoT}}$ , we can fine-tune the model  $\mathcal{M}$  to minimize the negative log-likelihood of both reasoning traces and final tool invocation:

$$\mathcal{L} = - \sum_{(q, \mathcal{T}) \in \mathcal{D}_{\text{CoT}}} \log P(\text{CoT} \mid q, \mathcal{T}). \quad (5)$$

By doing so, we can finally obtain the model that can self-reflect and generate the precise tool-calling in a single forward pass.

## 4 Experiments

### 4.1 Experimental Setup

**Tasks and Datasets.** We conduct the main experiments on two representative benchmarks: Berkeley Function-Calling Leaderboard (BFCL) (Patil et al., 2025) and ACEBench (Chen et al., 2025a). For BFCL, following standard evaluation protocols (Zhang et al., 2025b; Dang et al., 2025), we measure performance on *Non-Live* (synthetic) and *Live* (hand-crafted) subsets, respectively. For ACEBench, we adopt the *Normal (en)* split that covers diverse interaction scenarios as the test set.

Moreover, to assess robustness on different candidate tool scales, we employ two settings: 1) **Standard Setting** incorporates tool lists of the original benchmark to evaluate basic tool-calling capability; 2) **Extended Setting** expands candidates to 20

functions via randomly injecting irrelevant tools, simulating real-world noise. For evaluation, we report the strict Abstract Syntax Tree (AST) exact-match accuracy for both benchmarks. The details of all benchmarks are provided in Appendix A.1.

**Models.** We mainly evaluate the effectiveness of Tool-DC on Qwen2.5 (Yang et al., 2024) family (1.5B/3B/7B) models. To verify the universality of our methods, we also use the Qwen3-4B (Yang et al., 2025a), Llama-3.1/3.2 (Dubey et al., 2024) series (Llama-3.2-1B/-3B, Llama-3.1-8B), Gemma-3 (Team et al., 2025) series (1B/4B/12B) instruction models, and two closed-source LLMs, *i.e.*, GPT-4o-mini (Hurst et al., 2024) and DeepSeek-V3.2 (Liu et al., 2024). During the implementation of Tool-DC (TF), we set the number of groups to up to 5<sup>1</sup>, *i.e.*,  $K = \min(5, N)$ . For the extra retriever, we use the representative BM25 (Robertson et al., 2009) method. While for the implementation of Tool-DC (TB), we use the public LLaMA-Factory (Zheng et al., 2024a) toolkit as the training codebase. The details of training and inference hyperparameters are shown in Appendix A.2.

**Compared Methods.** We compare our Tool-DC with several cutting-edge counterparts. For training-free methods, we use the following baselines: 1) *GT\_Funs*, only using the ground-truth tools as the context; 2) *All\_Funs*, directly using all candidate tools as the context; 3) *Top-K*, using the retrieved Top- $K$  relevant tools as the context; 4) *HiTEC-ICL* (Cui et al., 2025), using the manual designed error checklist to guide the precise tool-calling; 5) *ToolGT (Prompting)* (Dang et al., 2025), using a curriculum-inspired prompt to enforce step-by-step tool-calling. For training-based methods, in addition to the base model and vanilla supervised fine-tuning (SFT) baseline, *i.e.*, directly optimizing the model on ground-truth tool invocation without any reasoning traces, we also compare with several proprietary and tool-specialized LLMs. Due to space limitations, we provide the details of all compared models in Appendix A.3.

### 4.2 Main Results

#### 4.2.1 Performance of Training-free Strategy

We first evaluate Qwen2.5 family models using our Tool-DC (TF) on BFCL and ACEBench benchmarks, and present the comparative results in Table 1. From these results, we can find that:

<sup>1</sup>The parameter analysis of  $K$  can be found in Section 4.4.

Table 1: **Performance comparison of Qwen2.5 models** with different training-free strategies on BFCL and ACEBench. The best and second-best results are in **bold** and underlined, respectively. The subscript results denote the relative performance gains against the All\_Funs baseline.

Methods	Standard Setting					Extended Setting				
	BFCL			ACEBench	Avg.	BFCL			ACEBench	Avg.
	Non-Live	Live	Overall	Overall		Non-Live	Live	Overall	Overall	
<b>Qwen2.5-1.5B-Instruct</b>										
GT_Funs	70.96	62.77	66.87	47.92	57.40	70.96	62.77	66.87	47.92	57.40
All_Funs	69.96	58.18	64.07	46.33	55.20	40.69	33.23	36.96	22.00	29.48
Top-K	76.35	63.80	70.08	44.52	57.30 <sub>↑2.10</sub>	60.31	43.60	51.96	38.58	45.27 <sub>↑15.79</sub>
HiTEC-ICL	<u>75.56</u>	<u>61.51</u>	<u>68.54</u>	45.17	<u>56.86</u> <sub>↑1.66</sub>	<u>46.46</u>	<u>39.97</u>	<u>43.22</u>	25.42	34.32 <sub>↑4.84</sub>
ToolGT (Prompting)	65.29	52.26	58.78	47.50	53.14 <sub>↓2.06</sub>	39.58	35.83	37.71	<u>35.33</u>	<u>36.52</u> <sub>↑7.04</sub>
<b>Tool-DC (TF) (Ours)</b>	<b>77.63</b>	<b>63.43</b>	<b>70.53</b>	<b>49.08</b>	<b>59.81</b> <sub>↑4.61</sub>	<b>64.13</b>	<b>62.00</b>	<b>63.07</b>	<b>46.08</b>	<b>54.58</b> <sub>↑25.10</sub>
<b>Qwen2.5-3B-Instruct</b>										
GT_Funs	79.23	69.13	74.18	54.42	64.30	79.23	69.13	74.18	54.42	64.30
All_Funs	78.08	63.51	70.80	47.25	59.03	62.73	44.34	53.54	36.50	45.02
Top-K	78.00	66.03	72.02	45.78	58.90 <sub>↓0.13</sub>	63.73	46.71	55.22	38.02	46.62 <sub>↑1.60</sub>
HiTEC-ICL	<b>78.90</b>	<u>63.80</u>	<u>71.35</u>	49.33	<u>60.34</u> <sub>↑1.31</sub>	<u>63.79</u>	45.23	54.51	34.92	44.72 <sub>↓0.30</sub>
ToolGT (Prompting)	68.94	57.22	63.08	<b>50.50</b>	56.79 <sub>↓2.24</sub>	59.29	54.92	57.11	46.58	51.85 <sub>↑6.83</sub>
<b>Tool-DC (TF) (Ours)</b>	<u>78.54</u>	<b>66.62</b>	<b>72.58</b>	<u>49.42</u>	<b>61.00</b> <sub>↑1.97</sub>	<b>71.79</b>	<b>57.74</b>	<b>64.77</b>	<b>48.17</b>	<b>56.47</b> <sub>↑11.45</sub>
<b>Qwen2.5-7B-Instruct</b>										
GT_Funs	85.81	78.02	81.92	65.17	73.55	85.81	78.02	81.92	65.17	73.55
All_Funs	<b>86.46</b>	67.44	76.95	59.92	68.44	75.71	62.77	69.24	58.58	63.91
Top-K	84.60	75.50	80.05	50.83	65.44 <sub>↓3.00</sub>	70.23	54.18	62.21	46.31	54.26 <sub>↓9.65</sub>
HiTEC-ICL	84.98	<b>76.24</b>	<b>80.61</b>	61.33	<u>70.97</u> <sub>↑2.53</sub>	<u>76.81</u>	<u>65.14</u>	70.98	54.67	62.83 <sub>↓1.08</sub>
ToolGT (Prompting)	77.73	68.02	72.88	<b>68.92</b>	70.90 <sub>↑2.46</sub>	76.00	60.55	68.28	<b>62.42</b>	<u>65.35</u> <sub>↑1.44</sub>
<b>Tool-DC (TF) (Ours)</b>	84.73	74.83	79.78	64.00	<b>71.89</b> <sub>↑3.45</sub>	<b>84.40</b>	<b>70.00</b>	<b>77.20</b>	<u>58.83</u>	<b>68.02</b> <sub>↑4.11</sub>




Figure 4: **Performance comparison of other LLMs** with different training-free strategies on the BFCL benchmark. Here, we mainly compare our Tool-DC (TF) method with the All\_Funs baseline.

**Tool-DC (TF) outperforms the baselines across model scales.** First, in the Standard Setting, Tool-DC (TF) achieves the highest average scores across all Qwen2.5 family models, especially for the smaller models. Specifically, on the Qwen2.5-1.5B model, our method achieves an average score of 59.81%, outperforming the vanilla method by a significant margin, *i.e.*, +4.61% average gains. This indicates that smaller models struggle more to deal with long-context tool-calling tasks, while our methods can effectively reduce the reasoning difficulty and bring better performance.

**Tool-DC (TF) exhibits strong robustness in noisy long-context tool-calling scenarios.** Compared to the Standard Setting, we introduce more noisy candidate tools in the Extended Setting. As seen, in

such a more difficult setting, all methods show performance degradation, especially in smaller models. Notably, for the All\_Funs, the performance degradation is up to 25.72% average score. Conversely, by locating the relevant tools, our Tool-DC (TF) can alleviate the negative effect of noisy tools and maintain model performance effectively. Specifically, it outperforms the All\_Funs by +25.10% average score in the Qwen2.5-1.5B model. These results prove the robustness of Tool-DC (TF).

**Tool-DC (TF) brings consistent performance gains for various LLMs.** In addition to Qwen2.5 models, we also evaluate Tool-DC (TF) on Llama3, Gemma3 and two closed-source LLMs, *i.e.*, GPT-4o-mini and DeepSeek-V3.2. Figure 4 illustrates the comparative results on the BFCL benchmark,

Table 2: **Performance comparison between our tuned models and other proprietary models on BFCL.** Notably, the results of proprietary models are from the official leaderboard of BFCL.

Models	Non-Live	Live	Overall
<b>Close-source Models</b>			
OpenAI o3	81.94	73.21	77.58
DeepSeek-V3.2-Exp	85.52	76.02	80.77
Claude-Haiku-4.5	86.50	78.68	82.59
Gemini-3-Pro-Preview	<b>90.65</b>	<b>83.12</b>	<b>86.89</b>
<b>Open-source Models</b>			
<i>Large-scale Models</i>			
Qwen3-235B-A22B-Inst	90.12	76.61	83.37
Llama-3.3-70B-Inst	88.02	76.61	82.32
<i>Tool-specialized Models</i>			
xLAM-2-8b-fc-r	84.58	67.95	76.27
ToolACE-DEV-8B	89.67	75.20	82.44
ToolACE-MT-8B	84.94	71.52	78.23
Hammer2.1-3B	85.50	69.50	77.50
Hammer2.1-7B	84.96	70.54	77.75
<b>Tuned Models</b>			
Qwen2.5-7B-Instruct	86.46	67.44	76.95
-w Vanilla SFT	84.23	74.24	79.24 <sub>↑2.29</sub>
-w Tool-DC (TB)	86.67	79.64	83.16 <sub>↑6.21</sub>
Llama-3.2-3B-Instruct	82.67	58.33	70.50
-w Vanilla SFT	83.31	62.99	73.15 <sub>↑2.65</sub>
-w Tool-DC (TB)	83.17	63.95	73.56 <sub>↑3.06</sub>

from which we find that our Tool-DC (TF) consistently boosts the tool-calling performance across all base models. Notably, for the powerful GPT-4o-mini model, Tool-DC (TF) can still lead to +5.3% performance gains, proving its effectiveness.

#### 4.2.2 Performance of Training-based Strategy

The comparative results of different tuned models on ACEBench and BFCL benchmarks are presented in Figure 5 and Table 2, respectively. Notably, due to limited computing resources, we only adopt Tool-DC (TB) to fine-tune some representative models. From them, we can observe that:

**Tool-DC (TB) yields consistent gains across diverse models.** In addition to the results on Qwen2.5 family models, we also apply our Tool-DC (TB) to boost the tool-calling performance of Llama3 and Qwen3 models. Taking the ACEBench benchmark as an example, we find that Tool-DC (TB) brings +24.95% performance gain for Llama-3.2-3B-Instruct model, and +15.9% for Qwen3-4B-Instruct model. These results can prove the universality of our Tool-DC (TB) method.

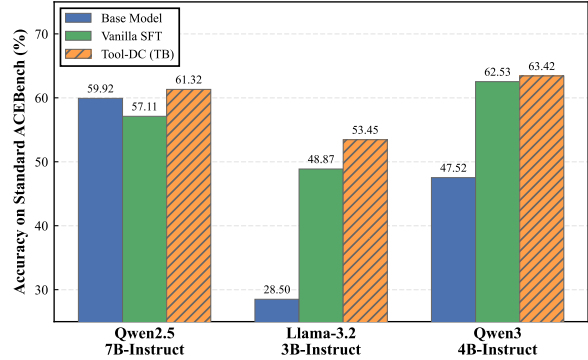


Figure 5: Comparison between base models and tuned models using Tool-DC (TB) on Standard ACEBench.

Table 3: **Ablation study on Tool-DC (TF).** Here, we use the Qwen2.5-3B-Instruct as the base model and evaluate on the Extended Setting of BFCL.

Methods	Modules			Performance		
	Try	Check	Retry	Non-Live	Live	Overall
Base	✗	✗	✗	62.73	44.34	53.54
<b>Tool-DC (TF)</b>	✓	✓	✓	<b>71.79</b>	<b>57.74</b>	<b>64.77</b>
w/o Try	✗	✓	✓	44.19	29.39	36.79
w/o Check	✓	✗	✓	52.06	52.41	52.24
w/o Retry	✓	✓	✗	7.33	3.18	5.26

**Our tuned models achieve remarkable performance against the proprietary LLMs.** In Table 2, we report the official results from the leaderboard<sup>2</sup> of BFCL for the proprietary LLMs. It can be seen, our Tool-DC (TB) can not only improve the tool-calling performance for various base models, but also enable them to achieve remarkable (or even better) performance against the other proprietary LLMs. More specifically, with the help of Tool-DC (TB), Qwen2.5-7B-Instruct achieves 83.16% overall score on BFCL, outperforming the OpenAI o3, DeepSeek-V3.2, and Claude-Haiku-4.5. These results demonstrate the superiority of our method.

#### 4.3 Ablation Study

In this part, we conduct ablation studies to investigate the effect of different stages (*Try*, *Check*, and *Retry*) of Tool-DC (TF) method, respectively. Notably, we use the Qwen2.5-3B-Instruct as the base model, and report the results on the Extended Setting of BFCL in Table 3. As seen, removing any stage in Tool-DC (TF) will lead to performance degradation, proving their effectiveness.

Firstly, when removing the *Try* stage, the average performance of Qwen2.5-3B-Instruct drops from 64.77% to 36.79%, indicating the importance

<sup>2</sup><https://gorilla.cs.berkeley.edu/leaderboard.html>




Figure 6: **Impact of Group Count  $K$ .** We evaluate Qwen2.5-3B-Instruct using different training-free methods on the Extended Setting of BFCL.

of task decomposition. Without the divide-and-conquer mechanism, Tool-DC (TF) falls short in checking its outputs and self-reflection. Secondly, when removing the *Check* stage, the refinement process tends to be unstable and lead to sub-optimal results. Lastly, when removing the *Retry* stage, we observe a performance collapse, *i.e.*, from 64.77% to 5.26%. This highlights: after filtering out potentially correct answers via *Check*, Tool-DC (TF) relies on *Retry* to leverage these positive signals. Without *Retry*, Tool-DC (TF) cannot process the high-potential candidates identified by *Check*, resulting in the loss of valid solutions.

#### 4.4 More Analyses

**Sensitivity analysis of Group Count  $K$ .** We investigate the impact of the group count  $K$ , which is an important hyperparameter in the first stage of Tool-DC (TF). In practice, we conduct comparative experiments on the Extended Setting of BFCL by ranging  $K$  from 1 to 8. As shown in Figure 6, the Top- $K$  baseline method suffers from significant performance drops at smaller  $K$  ( $<3$ ) due to aggressive sample filtering. Conversely, by making full use of all useful candidate tools, our Tool-DC (TF) exhibits superior robustness to the number of  $K$ . More specifically, in Figure 6, “Tool-DC (TF) w/ only  $T_{top}$ ” denotes that we only use the retrieved Top- $K$  tools in the *Try* stage, which can be seen as an improvement on the Top- $K$  method via our *Check* and *Retry* stages. By doing so, Tool-DC (TF) consistently improves the performance of the Top- $K$  baseline, indicating the effectiveness of the self-reflection mechanism. Moreover, when  $K$  is set to 5, our Tool-DC (TF) method performs best, thus leaving it as the default setting in our work.

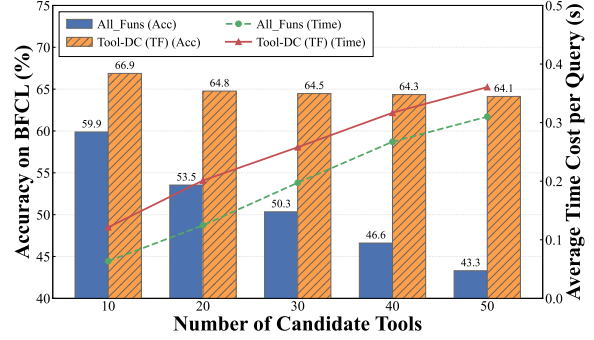


Figure 7: **Scalability and Efficiency Analysis.** Here, we use the Qwen2.5-3B-Instruct as the base model, and compare our Tool-DC (TF) with the All\_Funs baseline method on the Extended Setting of BFCL.

**Scalability and efficiency analysis.** To better simulate the real-world scenarios where the candidate tools are massive and noisy, we increase the number of candidate tools  $N$  from 10 to 50 for the BFCL benchmark by randomly adding some irrelevant tools. We compare the results of the Qwen2.5-3B-Instruct model by using the All\_Funs and our Tool-DC (TF), respectively. As illustrated in Figure 7, the performance of All\_Funs baseline drops significantly as the  $N$  increases, *i.e.*, from 59.88% to 43.30%. On the contrary, our Tool-DC (TF) can maintain the performance effectively, demonstrating its robustness. Furthermore, we additionally compare the inference latency between Tool-DC (TF) and All\_Funs. Although Tool-DC (TF) indeed leads to some inference overhead, it is tolerable against the performance gains of Tool-DC (TF).

## 5 Conclusion

In this paper, we propose Tool-DC, a divide-and-conquer framework to boost the tool-calling performance of LLMs via mitigating the negative effect of long-context and noisy candidate tools. Specifically, Tool-DC involves two variants: the training-free Tool-DC (TF) and training-based Tool-DC (TB). The former approach leverages the “Try-Check-Retry” pipeline to reduce the reasoning difficulty and make full use of the self-reflection ability of LLMs, thus leading to precise tool-calling. The latter approach aims to internalize the divide-and-conquer paradigm into model parameters via fin-tuning. Extensive experiments show that our Tool-DC consistently outperforms both training-free and training-based counterparts across various model scales and architectures.

## 518 Limitations

519 Our work has several potential limitations that  
520 present opportunities for future research. First, we  
521 only used xlam-function-calling-60k as the seed  
522 dataset to construct our CoT data for Tool-DC (TB).  
523 Although this dataset is large, it lacks sufficient di-  
524 versity and contains a limited number of tools in the  
525 context. In the future, we plan to construct a more  
526 diverse training set with massive and noisy con-  
527 text to simulate real-world scenarios, and explore  
528 reinforcement learning, *e.g.*, GRPO (Guo et al.,  
529 2025), to further improve the model performance.  
530 Secondly, our current experiments are limited to  
531 single-step tool-calling tasks, and we do not evalu-  
532 ate our framework on multi-step nested scenarios,  
533 such as those introduced in the latest version of  
534 BFCL. Therefore, future work could extend our  
535 method to more complex tool-calling settings.

## 536 Ethics Statement

537 We take ethical considerations very seriously and  
538 strictly adhere to the ACL Ethics Policy. This paper  
539 proposes a method to boost the tool-calling perfor-  
540 mance of LLMs in long-context scenarios. This  
541 method involves both training-free and training-  
542 based strategies, using a divide-and-conquer mech-  
543 anism to assist model decision-making, rather than  
544 encouraging them to learn privacy-sensitive knowl-  
545 edge that may raise ethical concerns. Moreover,  
546 all training and evaluation datasets used in this pa-  
547 per are publicly available and have been widely  
548 adopted by researchers. Thus, we believe that this  
549 research will not pose ethical issues.

## 550 References

551 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama  
552 Ahmad, Ilge Akkaya, Florencia Leoni Aleman,  
553 Diogo Almeida, Janko Altenschmidt, Sam Altman,  
554 Shyamal Anadkat, et al. 2023. Gpt-4 technical report.  
555 *arXiv preprint arXiv:2303.08774*.

556 Anthropic. 2025. Introducing claude haiku 4.5.  
557 <https://www.anthropic.com/news/claude-haiku-4-5>.

558 Tom Brown, Benjamin Mann, Nick Ryder, Melanie  
559 Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind  
560 Neelakantan, Pranav Shyam, Girish Sastry, Amanda  
561 Askell, et al. 2020. Language models are few-shot  
562 learners. In *Advances in neural information process-*  
563 *ing systems*, pages 1877–1901.

564 Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen,  
565 Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi

Chen, Pei Chu, et al. 2024. Internlm2 technical re-  
port. *arXiv preprint arXiv:2403.17297*.

Chen Chen, Xinlong Hao, Weiwen Liu, Xu Huang,  
Xingshan Zeng, Shuai Yu, Dexun Li, Yuefeng Huang,  
Xiangcheng Liu, Wang Xinzhi, and Wu Liu. 2025a.  
ACEBench: A comprehensive evaluation of LLM  
tool usage. In *Findings of the Association for Com-*  
*putational Linguistics: EMNLP 2025*, pages 12970–  
12998.

Jianlyu Chen, Shitao Xiao, Peitian Zhang, Kun  
Luo, Defu Lian, and Zheng Liu. 2024. M3-  
embedding: Multi-linguality, multi-functionality,  
multi-granularity text embeddings through self-  
knowledge distillation. In *Findings of the Associa-*  
*tion for Computational Linguistics ACL 2024*, pages  
2318–2335.

Yi-Chang Chen, Po-Chun Hsu, Chan-Jan Hsu, and Da-  
shan Shiu. 2025b. Enhancing function-calling capa-  
bilities in llms: Strategies for prompt formats, data  
integration, and multilingual translation. In *Procee-*  
*dings of the 2025 Conference of the Nations of the*  
*Americas Chapter of the Association for Computa-*  
*tional Linguistics: Human Language Technologies*  
*(Volume 3: Industry Track)*, pages 99–111.

Yue Cui, Liuyi Yao, Shuchang Tao, Weijie Shi, Yaliang  
Li, Bolin Ding, and Xiaofang Zhou. 2025. Enhancing  
tool learning in large language models with hierarchi-  
cal error checklists. In *Findings of the Association*  
*for Computational Linguistics: ACL 2025*, pages  
16357–16375.

Hy Dang, Tianyi Liu, Zhuofeng Wu, Jingfeng Yang,  
Haoming Jiang, Tao Yang, Pei Chen, Zhengyang  
Wang, Helen Wang, Huasheng Li, et al. 2025. Im-  
proving large language models function calling and  
interpretability via guided-structured templates. In  
*Proceedings of the 2025 Conference on Empirical*  
*Methods in Natural Language Processing*, pages  
24437–24453.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and  
Kristina Toutanova. 2019. Bert: Pre-training of deep  
bidirectional transformers for language understand-  
ing. In *Proceedings of the 2019 conference of the*  
*North American chapter of the association for com-*  
*putational linguistics: human language technologies,*  
*volume 1 (long and short papers)*, pages 4171–4186.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey,  
Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman,  
Akhil Mathur, Alan Schelten, Amy Yang, Angela  
Fan, et al. 2024. The llama 3 herd of models. *arXiv*  
*preprint arXiv:2407.21783*.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song,  
Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma,  
Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: In-  
centivizing reasoning capability in llms via reinforc-  
ement learning. *arXiv preprint arXiv:2501.12948*.

Bingguang Hao, ZengZhuang Xu, Maolin Wang,  
Yuntao Wen, Yicheng Chen, Cunyin Peng, Long

623	Chen, Dong Wang, Xiangyu Zhao, Jinjie Gu, et al.	Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang,	679
624	2025. Funreason: Enhancing large language mod-	Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi	680
625	els' function calling via self-refinement multiscale	Deng, Chenyu Zhang, Chong Ruan, et al. 2024.	681
626	loss and automated data refinement. <i>arXiv preprint</i>	Deepseek-v3 technical report. <i>arXiv preprint</i>	682
627	<i>arXiv:2505.20192</i> .	<i>arXiv:2412.19437</i> .	683
628	Jie He, Jennifer Neville, Mengting Wan, Longqi Yang,	Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingx-	684
629	Hui Liu, Xiaofeng Xu, Xia Song, Jeff Z. Pan, and Pei	uan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang,	685
630	Zhou. 2025. GenTool: Enhancing tool generalization	Chaofan Lin, Chen Dong, et al. 2025. Deepseek-v3.	686
631	in language models through zero-to-one and weak-to-	2: Pushing the frontier of open large language models.	687
632	strong simulation. In <i>Findings of the Association for</i>	<i>arXiv preprint arXiv:2512.02556</i> .	688
633	<i>Computational Linguistics: ACL 2025</i> , pages 1097–	Suhong Moon, Siddharth Jha, Lutfi Eren Erdogan, Se-	689
634	1122.	hoon Kim, Woosang Lim, Kurt Keutzer, and Amir	690
635	Cheng-Yu Hsieh, Si-An Chen, Chun-Liang Li, Yasuhisa	Gholami. 2024. Efficient and scalable estimation of	691
636	Fujii, Alexander Ratner, Chen-Yu Lee, Ranjay Kr-	tool representations in vector space. <i>arXiv preprint</i>	692
637	ishna, and Tomas Pfister. 2023. Tool documenta-	<i>arXiv:2409.02141</i> .	693
638	tion enables zero-shot tool-usage with large language	OpenAI. 2025. Openai-o3.	694
639	models. <i>arXiv preprint arXiv:2308.00675</i> .	<a href="https://openai.com/index/introducing-o3-and-o4-min">https://openai.com/index/introducing-o3-and-</a>	695
640	Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu,	o4-min.	696
641	Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen,	Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Car-	697
642	et al. 2021. Lora: Low-rank adaptation of large lan-	roll L Wainwright, Pamela Mishkin, Chong Zhang,	698
643	guage models. In <i>International Conference on Learn-</i>	Sandhini Agarwal, Katarina Slama, Alex Ray, et al.	699
644	<i>ing Representations</i> .	2022. Training language models to follow instruc-	700
645	Xu Huang, Weiwen Liu, Xingshan Zeng, Yuefeng	tions with human feedback. In <i>Proceedings of the</i>	701
646	Huang, Xinlong Hao, Yuxian Wang, Yirong Zeng,	<i>36th International Conference on Neural Information</i>	702
647	Chuhan Wu, Yasheng Wang, Ruiming Tang, et al.	<i>Processing Systems</i> , pages 27730–27744.	703
648	2025. Toolace-dev: Self-improving tool learning	Bhrij Patel, Davide Belli, Amir Jalalirad, Maximilian	704
649	via decomposition and evolution. <i>arXiv preprint</i>	Arnold, Aleksandr Ermovol, and Bence Major. 2025.	705
650	<i>arXiv:2505.07512</i> .	Dynamic tool dependency retrieval for efficient func-	706
651	Aaron Hurst, Adam Lerer, Adam P Goucher, Adam	tion calling. <i>arXiv preprint arXiv:2512.17052</i> .	707
652	Perelman, Aditya Ramesh, Aidan Clark, AJ Os-	Shishir G Patil, Huanzhi Mao, Fanjia Yan, Charlie	708
653	trow, Akila Welihinda, Alan Hayes, Alec Radford,	Cheng-Jie Ji, Vishnu Suresh, Ion Stoica, and Joseph E	709
654	et al. 2024. Gpt-4o system card. <i>arXiv preprint</i>	Gonzalez. 2025. The Berkeley function calling	710
655	<i>arXiv:2410.21276</i> .	leaderboard (BFCL): From tool use to agentic evalua-	711
656	Yibin Lei, Liang Ding, Yu Cao, Changtong Zan, An-	tion of large language models. In <i>Proceedings of the</i>	712
657	drew Yates, and Dacheng Tao. 2023. Unsupervised	<i>42nd International Conference on Machine Learning</i> .	713
658	dense retrieval with relevance-aware contrastive pre-	Sundar Pichai, Demis Hassabis, and Ko-	714
659	training. In <i>Findings of the Association for Computa-</i>	ray Kavukcuoglu. 2025. A new	715
660	<i>tional Linguistics: ACL 2023</i> , pages 10932–10940.	era of intelligence with gemini 3.	716
661	Xiaobo Liang, Wenjin Xie, Juntao Li, Wanfu Wang,	<a href="https://deepmind.google/models/gemini/pro/">https://deepmind.google/models/gemini/pro/</a> .	717
662	Yibin Chen, Kehai Chen, and Min Zhang. 2025. Tool	Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang,	718
663	learning via inference-time scaling and cycle verifier.	Xiusi Chen, Dilek Hakkani-Tür, Gokhan Tur, and	719
664	In <i>Findings of the Association for Computational</i>	Heng Ji. 2025. Toolrl: Reward is all tool learning	720
665	<i>Linguistics: ACL 2025</i> , pages 24658–24671.	needs. <i>arXiv preprint arXiv:2504.13958</i> .	721
666	Yusheng Liao, Shuyang Jiang, Yanfeng Wang, and	Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan	722
667	Yu Wang. 2025. Reflectool: Towards reflection-	Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang,	723
668	aware tool-augmented clinical agents. In <i>Proceed-</i>	Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian,	724
669	<i>ings of the 63rd Annual Meeting of the Association</i>	Ruobing Xie, Jie Zhou, Mark Gerstein, dahai li,	725
670	<i>for Computational Linguistics (Volume 1: Long Pa-</i>	Zhiyuan Liu, and Maosong Sun. 2024. ToolLLM:	726
671	<i>pers)</i> , pages 13507–13531.	Facilitating large language models to master 16000+	727
672	Qiqiang Lin, Muning Wen, Qiuying Peng, Guanyu Nie,	real-world APIs. In <i>The Twelfth International Con-</i>	728
673	Junwei Liao, Jun Wang, Xiaoyun Mo, Jiamu Zhou,	<i>ference on Learning Representations</i> .	729
674	Cheng Cheng, Yin Zhao, Jun Wang, and Weinan	Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai,	730
675	Zhang. 2025. Robust function-calling for on-device	Shuaiqiang Wang, Dawei Yin, Jun Xu, and Ji-Rong	731
676	language model via function masking. In <i>The Thir-</i>	Wen. 2025. Tool learning with large language mod-	732
677	<i>teenth International Conference on Learning Repre-</i>	els: A survey. <i>Frontiers of Computer Science</i> ,	733
678	<i>sentations</i> .	19(8):198343.	734

735	Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. <i>Foundations and Trends® in Information Retrieval</i> , 3(4):333–389.	790
736		791
737		792
738		793
739	Jingqing Ruan, Yihong Chen, Bin Zhang, Zhiwei Xu, Tianpeng Bao, Hangyu Mao, Ziyue Li, Xingyu Zeng, Rui Zhao, et al. 2023. Tptu: Task planning and tool usage of large language model-based ai agents. In <i>NeurIPS 2023 Foundation Models for Decision Making Workshop</i> .	794
740		795
741		796
742		797
743		798
744		799
745	Timo Schick, Jane Dwivedi-Yu, Roberto Dessí, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: language models can teach themselves to use tools. In <i>Proceedings of the 37th International Conference on Neural Information Processing Systems</i> , pages 68539–68551.	800
746		801
747		802
748		803
749		804
750		805
751		806
752	Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. 2025. Gemma 3 technical report. <i>arXiv preprint arXiv:2503.19786</i> .	807
753		808
754		809
755		810
756		811
757	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. <i>arXiv preprint arXiv:2302.13971</i> .	812
758		813
759		814
760		815
761		816
762		817
763	An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025a. Qwen3 technical report. <i>arXiv preprint arXiv:2505.09388</i> .	818
764		819
765		820
766		821
767	An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. <i>arXiv preprint arXiv:2412.15115</i> .	822
768		823
769		824
770		825
771	An Yang, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoyan Huang, Jiandong Jiang, Jianhong Tu, Jianwei Zhang, Jingren Zhou, et al. 2025b. Qwen2. 5-1m technical report. <i>arXiv preprint arXiv:2501.15383</i> .	826
772		827
773		828
774		829
775		830
776	Xingshan Zeng, Weiwen Liu, Xu Huang, Zezhong Wang, Lingzhi Wang, Liangyou Li, Yasheng Wang, Lifeng Shang, Xin Jiang, Ruiming Tang, et al. 2025. Toolace-r: Tool learning with adaptive self-refinement. <i>arXiv preprint arXiv:2504.01400</i> .	831
777		832
778		833
779		834
780		835
781	Jianguo Zhang, Tian Lan, Ming Zhu, Zuxin Liu, Thai Quoc Hoang, Shirley Kokane, Weiran Yao, Juntao Tan, Akshara Prabhakar, and Haolin Chen. 2025a. xLAM: A family of large action models to empower AI agent systems. In <i>Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 11583–11597.	836
782		837
783		838
784		839
785		840
786		841
787		842
788		843
789		844
	Shaokun Zhang, Yi Dong, Jieyu Zhang, Jan Kautz, Bryan Catanzaro, Andrew Tao, Qingyun Wu, Zhiding Yu, and Guilin Liu. 2025b. Nemotron-research-tool-n1: Exploring tool-using language models with reinforced reasoning. <i>arXiv preprint arXiv:2505.00024</i> .	845
		846
		847
		848
	Yaowei Zheng, Richong Zhang, Junhao Zhang, YeYanhan YeYanhan, and Zheyang Luo. 2024a. Llamafactory: Unified efficient fine-tuning of 100+ language models. In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)</i> , pages 400–410.	849
		850
		851
		852
		853
		854
		855
		856
		857
		858
		859
		860
		861
		862
		863
		864
		865
		866
		867
		868
		869
		870
		871
		872
		873
		874
		875
		876
		877
		878
		879
		880
		881
		882
		883
		884
		885
		886
		887
		888
		889
		890
		891
		892
		893
		894
		895
		896
		897
		898
		899
		900

into three primary categories: Normal, Special, and Agent. In this work, we focus exclusively on the **Normal** category, which contains 828 test samples. The Normal category is further subdivided into the following five tasks: *Single-Turn*, *Multi-Turn*, *Similar APIs*, *Profile*, and *Atom*.

## A.2 Training and Inference Details

In this work, we use the xlam-function-calling-60k dataset (Zhang et al., 2025a) as raw training corpus, and construct a CoT training set with 46,897 samples. Table 4 shows the detailed training hyperparameters. Specifically, we set the batch size to 16 and the peak learning rate to 2e-5. All models are trained with LoRA (Hu et al., 2021) for 2 epochs on two NVIDIA A800 80GB GPUs, requiring approximately 10 hours for finetuning 7B models using our Tool-DC (TB). For inference, we use the VLLM<sup>3</sup> engine for accelerating the model generation, and set the temperature to 0, the maximum length of output tokens to 512.

Table 4: Training hyperparameters for all models.

Hyperparameter	Value
Batch Size	16
Learning Rate	2e-5
LR Scheduler	cosine
Warmup Ratio	0.1
Epochs	2

## A.3 Details of Compared Models

In our training-based experiments, we compare our tuned models with both general-purpose commercial APIs (OpenAI o3<sup>4</sup> (OpenAI, 2025), DeepSeek-V3.2-Exp<sup>5</sup> (Liu et al., 2025), Claude-Haiku-4.5<sup>6</sup> (Anthropic, 2025), and Gemini-3-Pro-Preview<sup>7</sup> (Pichai et al., 2025)) and open-source large-scale models (Qwen3-235B-A22B-Inst (Yang et al., 2025a) and Llama-3.3-70B-Inst (Dubey et al., 2024)), as well as some specialized tool-calling models, including ToolACE-DEV-8B (Huang et al., 2025), ToolACE-MT-8B (Zeng et al., 2025), xLAM-2 (Zhang et al., 2025a), and Hammer2.1(1.5B/3B/7B) (Lin et al., 2025).

<sup>3</sup><https://github.com/vllm-project/vllm>

<sup>4</sup><https://openai.com/index/introducing-o3-and-o4-mini>

<sup>5</sup><https://api-docs.deepseek.com/news/news250929>

<sup>6</sup><https://www.anthropic.com/news/claude-haiku-4-5>

<sup>7</sup><https://deepmind.google/models/gemini/pro/>

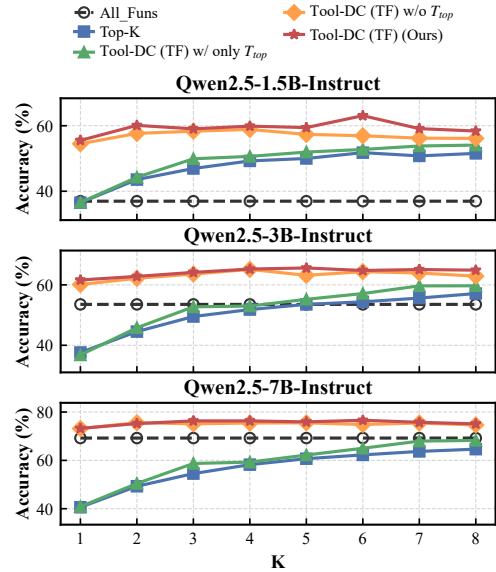


Figure 8: Sensitivity analysis of group count  $K$  across Qwen2.5 series. All models are evaluated on the Extended Setting of BFCL.

## A.4 Sensitivity Analysis of Group Count $K$

We extend the sensitivity analysis of group count  $K$  to Qwen2.5-1.5B, 3B, and 7B-Instruct. Results in Figure 8 consistently show that while baselines like *Top-K* suffer significant drops at low  $K$  ( $< 3$ ) due to over-filtering, our method remains robust even at  $K = 1$ . Performance stabilizes at  $K \in [4, 6]$ , confirming that a moderate partition size effectively balances noise reduction and context completeness across varying model scales.

## A.5 Performance on Long-context LLMs

To verify whether the long-context tool-calling issue can be fixed by scaling the context windows, we evaluate the cutting-edge long-context LLMs, including Qwen2.5-7B-Instruct-1M (Yang et al., 2025b) and InternLM2.5-7B-Chat-1M (Cai et al., 2024), which have up to 1M tokens in the context windows, on the BFCL benchmark. Specifically, we mainly conduct experiments on the Extended Settings, and report the comparative results in Table 5. For reference, we also report the results of GT\_Funs as the upper-bound.

As seen, while InternLM2.5-7B-Chat-1M has a long context window, it still suffers from dramatic performance drops in the Extended Setting, indicating that simply scaling context windows cannot address the long-context tool-calling problem. This underscores the necessity of exploring a more robust tool-calling method. Encouragingly, with the help of our Tool-DC (TF), InternLM2.5-7B-Chat-

Table 5: Performance of long-context LLMs on the Extended Setting of BFCL.

Methods	Non-Live				Live				Overall		
	Simple	Multiple	Parallel	Parallel Multiple	Simple	Multiple	Parallel	Parallel Multiple	Non-Live	Live	Overall
<b>Qwen2.5-7B-Instruct-1M</b>											
GT_Funs	74.92	95.00	90.50	85.50	80.23	77.30	62.50	66.67	86.48	77.50	81.99
All_Funs	68.67	85.50	84.00	71.00	63.18	65.81	62.50	58.33	77.29	65.14	71.22
<b>Tool-DC (TF) (Ours)</b>	73.17	89.00	88.50	79.50	65.89	69.71	62.50	66.67	82.54	68.84	75.69
<b>InternLM2.5-7B-Chat-1M</b>											
GT_Funs	51.92	85.00	70.50	51.00	55.81	51.28	25.00	20.83	64.60	51.30	57.95
All_Funs	17.92	26.00	25.00	10.50	14.34	12.44	18.75	4.17	19.85	12.73	16.29
<b>Tool-DC (TF) (Ours)</b>	41.00	65.00	46.00	28.50	35.66	32.19	12.50	12.50	45.12	32.27	38.70

1M can alleviate the performance drop effectively, *i.e.*, from 16.29% to 38.70%, continuing to prove its effectiveness and superiority.

### A.6 Performance on More Retrievers

In our work, we use the representative unsupervised retrieval method BM25 as the base retriever. Here, we further conduct experiments to investigate the generality of Tool-DC on the other retrievers. Specifically, two cutting-edge dense retrievers, including bge-reranker-v2-m3 (Chen et al., 2024) and ToolBench\_IR (Qin et al., 2024), are used in this experiment. We evaluate the Qwen2.5-3B-Instruct model on the Extended Setting of BFCL and report the comparative results in Table 6. From these results, we can find that our Tool-DC (TF) brings consistent performance gains against the Top- $K$  baseline across all base retrievers, proving its generality. Notably, although bge-reranker-v2-m3 is effective, it is a two-tower BERT (Devlin et al., 2019) model, which is inference-inefficient. While ToolBench\_IR is finetuned on ToolBench, it is difficult to generalize to other tool-calling tasks, thus leading to sub-optimal results. Moreover, considering BM25 is a powerful and efficient unsupervised retriever (Lei et al., 2023), we finally choose it as the base retriever in our experiment.

### A.7 Prompt Details

Here, we provide the detailed prompts used in our Tool-DC framework. Specifically, for Tool-DC (TF), the “Try” and “Retry” prompts are presented in Figure 9 and Figure 10, respectively. For Tool-DC (TB), we present the prompt template of data construction in Figure 11.

### A.8 Full Results

Due to space limitations, we only report the averaged experimental results in the main body of

this paper. Here, we present the detailed results to help readers better reproduce our main experiments. Specifically, the detailed results of training-free approaches on BFCL and ACEBench benchmarks are provided in Table 7 and Table 9, respectively. The detailed results of training-based approaches on BFCL and ACEBench benchmarks are provided in Table 8 and Table 10, respectively.

### A.9 AI Assistant Usage

During the paper writing, we used a proprietary LLM, *i.e.*, Gemini-3-Pro, as the general-purpose assistant to polish some sentences. We state that we did not prompt them to generate research ideas or search for related works, thus avoiding the risk of introducing false information.

Table 6: Performance of Qwen2.5-3B-Instruct using different retrievers on the Extended Setting of BFCL.

Methods	Non-Live				Live				Overall		
	Simple	Multiple	Parallel	Parallel Multiple	Simple	Multiple	Parallel	Parallel Multiple	Non-Live	Live	Overall
GT_Funs	70.92	91.50	79.00	75.50	67.83	70.37	56.25	37.50	79.23	69.13	74.18
All_Funs	64.42	71.00	62.50	53.00	43.41	44.92	43.75	29.17	62.73	44.34	53.54
<b>bge-reranker-v2-m3</b>											
-w Top- <i>K</i>	67.42	83.50	74.50	61.00	52.71	64.10	50.00	37.50	71.60	61.29	66.45
-w Tool-DC (TF)	68.08	82.50	75.50	68.50	57.75	64.29	50.00	41.67	73.65	62.47	68.06
<b>ToolBench_IR</b>											
-w Top- <i>K</i>	63.42	73.00	66.00	41.50	42.25	47.10	50.00	37.50	60.98	46.04	53.51
-w Tool-DC (TF)	68.83	81.50	74.00	58.50	53.10	62.01	56.25	37.50	70.71	59.81	65.26
<b>BM25</b>											
-w Top- <i>K</i>	64.42	74.50	71.00	45.00	41.09	48.53	43.75	29.17	63.73	46.71	55.22
-w Tool-DC (TF)	64.67	81.00	77.50	64.00	48.84	60.40	43.75	45.83	71.79	57.74	64.77

Prompt template of “Try” stage in Tool-DC (TF)

```
# System:
You are a Function Selection Expert. Your task is to identify ALL functions that are semantically relevant to the user’s question from the provided list. Extract information from the user’s question and substitute it into the function parameters.
Read the user’s question and the function descriptions carefully. Choose any function that could potentially meet user needs or meet a part of user needs.
If you decide to invoke any of the function(s), you MUST put it in the format of:
[func_name1(params_name1=params_value1...), func_name2(params)]. You SHOULD NOT include any other text in the response.
Here is a list of functions in json format that you can invoke.
<Tools of Each Group>
# User:
<Question>
```

Figure 9: The prompt template of “Try” stage in Tool-DC (TF).

Prompt template of “Retry” stage in Tool-DC (TF)

```
# System:
You are an expert in composing functions. You are given a question and a set of possible functions. Based on the question, you will need to make one or more function/tool calls to achieve the purpose. If none of the functions can be used, point it out. If the given question lacks the parameters required by the function, also point it out.
You should only return the function calls in your response. If you decide to invoke any of the function(s), you MUST put it in the format of: [func_name1(params_name1=params_value1...), func_name2(params)]. You SHOULD NOT include any other text in the response.
At each turn, you should try your best to complete the tasks requested by the user within the current turn. Continue to output functions to call until you have fulfilled the user’s request to the best of your ability.
Here is a list of functions in json format that you can invoke. <Tools>
# User:
<Question>
```

Figure 10: The prompt template of “Retry” stage in Tool-DC (TF).

### Prompt template for constructing the CoT data

#### # System:

You are an expert in composing functions. You are given a question and a set of possible functions. Based on the question, you will need to make one or more function/tool calls to achieve the purpose. If none of the function can be used, point it out. If the given question lacks the parameters required by the function, also point it out.

You should only return the function call in tools call sections. Continue to output functions to call until you have fulfilled the user's request to the best of your ability.

Here is a list of functions in json format that you can invoke. `<Tools>`

For each interaction, you MUST strictly follow this two-step process:

#### Step 1: Reasoning (`<think>`)

You must engage in a detailed chain-of-thought enclosed within `<think></think>` tags.

This process must follow these exact 3 sub-steps:

**Candidate Selection:** Analyze the user's question. Iterate through the function list...

**Validation:** Strictly check the candidate list against function definitions...

**Final Review:** Focus exclusively on the effective candidate list...

#### Step 2: Invoke (`<tool_call>`)

If you decide to invoke function(s), output them in the following specific format:

```
<tool_call>[func_name1(params_name1=value1, ...), func_name2(params)]</tool_call>
```

#### # User:

```
<Question>
```

#### # Assistant:

```
<think>
```

1. **Candidate Selection:** Analyzing the user's query, I will attempt to map key information to the function parameters. The functions that potentially match and may have their parameters filled are: `<candidate_tool_calls>`

2. **Validation:** I will now strictly verify these candidates against their definitions, ensuring all parameter types and constraints are met. The functions that pass this strict verification are: `<valid_tools>`.

3. **Final Review:** I will now eliminate any interference from irrelevant functions and focus solely on the valid candidates.

```
</think>
```

```
<tool_call><final_tool_invocation></tool_call>
```

Figure 11: The prompt template for constructing the CoT data in Tool-DC (TB).

Table 7: Detailed results of different training-free methods on BFCL, which is a full version of Table 1.

Methods	Non-Live				Live				Overall		
	Simple	Multiple	Parallel	Parallel Multiple	Simple	Multiple	Parallel	Parallel Multiple	Non-Live	Live	Overall
<b>Standard Setting</b>											
<i>Qwen2.5-1.5B-Instruct</i>											
GT_Funs	63.33	88.50	68.50	63.50	65.12	63.25	43.75	29.17	70.96	62.77	66.87
All_Funs	63.33	84.00	68.00	64.50	65.12	57.26	43.75	33.33	69.96	58.18	64.07
Top- <i>K</i>	66.42	83.50	79.50	76.00	68.99	63.15	50.00	45.83	76.35	<b>63.80</b>	70.08
HiTEC-ICL	<b>67.75</b>	83.50	79.00	72.00	<b>70.54</b>	59.83	<b>56.25</b>	41.67	75.56	61.51	68.54
ToolGT (Prompting)	66.67	78.00	64.00	52.50	67.05	49.29	25.00	41.67	65.29	52.26	58.78
<b>Tool-DC (TF) (Ours)</b>	66.50	<b>85.50</b>	<b>80.50</b>	<b>78.00</b>	68.60	62.68	<u>50.00</u>	<b>50.00</b>	<b>77.63</b>	<u>63.43</u>	<b>70.53</b>
<i>Qwen2.5-3B-Instruct</i>											
GT_Funs	70.92	91.50	79.00	75.50	67.83	70.37	56.25	37.50	79.23	69.13	74.18
All_Funs	70.83	88.00	79.50	74.00	68.22	63.06	56.25	37.50	78.08	63.51	70.80
Top- <i>K</i>	71.00	89.50	78.00	73.50	67.05	66.29	<b>62.50</b>	45.83	78.00	66.03	72.02
HiTEC-ICL	<b>72.58</b>	87.50	<b>80.50</b>	<b>75.00</b>	<b>68.60</b>	63.25	<b>62.50</b>	37.50	<b>78.90</b>	63.80	71.35
ToolGT (Prompting)	67.75	81.50	68.50	58.00	<b>68.60</b>	55.46	25.00	33.33	68.94	57.22	63.08
<b>Tool-DC (TF) (Ours)</b>	70.67	<b>90.00</b>	<u>78.50</u>	<b>75.00</b>	<u>67.05</u>	<b>66.95</b>	<b>62.50</b>	<b>50.00</b>	<u>78.54</u>	<b>66.62</b>	<b>72.58</b>
<i>Qwen2.5-7B-Instruct</i>											
GT_Funs	72.25	96.00	91.00	84.00	77.13	78.73	62.50	66.67	85.81	78.02	81.92
All_Funs	75.33	94.50	91.50	84.50	76.74	74.93	62.50	70.83	86.46	67.44	76.95
Top- <i>K</i>	71.92	93.50	<b>91.00</b>	<u>82.00</u>	77.13	75.40	<b>62.50</b>	70.83	84.60	<u>75.50</u>	80.05
HiTEC-ICL	<b>72.42</b>	<b>95.00</b>	89.00	<b>83.50</b>	<b>78.29</b>	<b>76.07</b>	<b>62.50</b>	70.83	<b>84.98</b>	<b>76.24</b>	<b>80.61</b>
ToolGT (Prompting)	69.92	82.50	84.00	74.50	<b>78.29</b>	65.53	56.25	<b>75.00</b>	77.73	68.02	72.88
<b>Tool-DC (TF) (Ours)</b>	<u>71.92</u>	<u>94.00</u>	<b>91.00</b>	<u>82.00</u>	77.91	74.93	<b>62.50</b>	66.67	<u>84.73</u>	75.20	79.97
<b>Extended Setting</b>											
<i>Qwen2.5-1.5B-Instruct</i>											
GT_Funs	63.33	88.50	68.50	63.50	65.12	63.25	43.75	29.17	70.96	62.77	66.87
All_Funs	51.75	49.00	35.50	26.50	32.95	33.90	25.00	12.50	40.69	33.23	36.96
Top- <i>K</i>	61.75	<b>69.50</b>	<u>69.50</u>	<b>40.50</b>	40.70	45.20	25.00	16.67	60.31	43.60	51.96
HiTEC-ICL	57.33	53.00	45.50	30.00	<u>41.09</u>	40.36	<u>31.25</u>	<u>16.67</u>	46.46	39.97	43.22
ToolGT (Prompting)	46.83	55.00	31.00	25.50	38.76	36.37	0.00	4.17	39.58	35.83	37.71
<b>Tool-DC (TF) (Ours)</b>	<b>64.13</b>	<u>59.50</u>	<b>76.50</b>	<u>40.00</u>	<b>76.50</b>	<b>67.00</b>	<b>53.50</b>	<b>64.13</b>	<b>64.13</b>	<b>62.00</b>	<b>63.07</b>
<i>Qwen2.5-3B-Instruct</i>											
GT_Funs	70.92	91.50	79.00	75.50	67.83	70.37	56.25	37.50	79.23	69.13	74.18
All_Funs	64.42	71.00	62.50	53.00	43.41	44.92	43.75	29.17	62.73	44.34	53.54
Top- <i>K</i>	64.42	<u>74.50</u>	<u>71.00</u>	45.00	41.09	48.53	<b>43.75</b>	29.17	63.73	46.71	55.22
HiTEC-ICL	64.17	71.50	64.00	<u>55.50</u>	44.19	45.68	<b>43.75</b>	<u>37.50</u>	<u>63.79</u>	45.23	54.51
ToolGT (Prompting)	58.67	71.50	54.50	52.50	<b>56.59</b>	55.37	<b>43.75</b>	25.00	59.29	<u>54.92</u>	<u>57.11</u>
<b>Tool-DC (TF) (Ours)</b>	<b>64.67</b>	<b>81.00</b>	<b>77.50</b>	<b>64.00</b>	<u>48.84</u>	<b>60.40</b>	<b>43.75</b>	<b>45.83</b>	<b>71.79</b>	<b>57.74</b>	<b>64.77</b>
<i>Qwen2.5-7B-Instruct</i>											
GT_Funs	72.25	96.00	91.00	84.00	77.13	78.73	62.50	66.67	85.81	78.02	81.92
All_Funs	65.33	82.00	85.00	70.50	58.14	63.91	75.00	54.17	75.71	62.77	69.24
Top- <i>K</i>	67.42	81.00	83.50	49.00	48.45	55.94	50.00	41.67	70.23	54.18	62.21
HiTEC-ICL	65.75	<b>86.00</b>	85.00	70.50	62.79	65.72	68.75	62.50	76.81	65.14	70.98
ToolGT (Prompting)	70.00	80.00	82.50	<b>71.50</b>	58.91	61.06	43.75	66.67	76.00	60.55	68.28
<b>Tool-DC (TF) (Ours)</b>	<b>84.40</b>	71.58	<b>94.75</b>	50.00	<b>91.00</b>	<b>91.00</b>	<b>84.00</b>	<b>84.40</b>	<b>84.40</b>	<b>70.00</b>	<b>77.20</b>

Table 8: Detailed results of different training-based methods on BFCL, which is a full version of Table 2.

Methods	Non-Live				Live				Overall		
	Simple	Multiple	Parallel	Parallel Multiple	Simple	Multiple	Parallel	Parallel Multiple	Non-Live	Live	Overall
<b>Qwen2.5-7B-Instruct</b>											
-w Vanilla SFT	75.33	94.50	91.50	84.50	76.74	74.93	62.50	70.83	86.46	67.44	76.95
-w Tool-DC (TB)	70.92	95.50	88.50	82.00	77.52	73.69	81.25	58.33	84.23	74.24	79.24
-w Tool-DC (TB)	71.17	95.00	93.00	87.50	83.72	79.58	56.25	54.17	86.67	79.64	<b>83.16</b>
<b>Llama-3.2-3B-Instruct</b>											
-w Vanilla SFT	70.67	92.50	88.50	79.00	58.33	65.12	57.64	25.00	82.67	58.33	70.50
-w Tool-DC (TB)	71.25	92.00	88.50	81.50	70.93	61.35	56.25	54.17	83.31	62.99	73.15
-w Tool-DC (TB)	71.67	89.50	89.50	82.00	68.60	64.01	18.75	41.67	83.17	63.95	<b>73.56</b>

Table 9: Detailed results of different training-free methods on ACEBench, which is a full version of Table 1. “Atom”, “Single Turn”, “Multiple Turn”, “Similar API”, and “Profile” are subsets of ACEBench.

Methods	Standard Setting					Extended Setting						
	Atom	Single Turn	Multiple Turn	Similar API	Profile	Overall	Atom	Single Turn	Multiple Turn	Similar API	Profile	Overall
<b><i>Qwen2.5-1.5B-Instruct</i></b>												
GT_Funs	56.00	38.50	28.00	52.00	54.00	47.92	56.00	38.50	28.00	52.00	54.00	47.92
All_Funs	56.67	35.00	29.00	50.00	38.00	46.33	27.67	15.00	5.00	32.00	26.00	22.00
Top-K	<u>57.33</u>	<u>37.00</u>	10.12	<u>52.00</u>	<u>44.00</u>	44.52	<u>52.67</u>	<u>28.00</u>	8.50	<b>40.00</b>	34.00	<u>38.58</u>
HiTEC-ICL	54.33	34.00	<u>24.00</u>	<u>52.00</u>	<b>48.00</b>	45.17	28.00	21.50	13.00	24.00	<b>44.00</b>	25.42
ToolGT (Prompting)	<b>60.67</b>	33.00	<b>30.00</b>	48.00	32.00	<u>47.50</u>	49.00	18.00	<u>17.00</u>	<u>26.00</u>	34.00	35.33
<b>Tool-DC (TF)</b>	56.33	<b>47.50</b>	23.00	<b>66.00</b>	<u>44.00</u>	<b>49.08</b>	<b>57.67</b>	<b>38.50</b>	<b>25.00</b>	<b>40.00</b>	<u>40.00</u>	<b>46.08</b>
<b><i>Qwen2.5-3B-Instruct</i></b>												
GT_Funs	63.33	47.50	28.00	60.00	62.00	54.42	63.33	47.50	28.00	60.00	62.00	54.42
All_Funs	57.33	45.50	17.00	50.00	48.00	47.25	43.00	32.00	17.00	48.00	34.00	36.50
Top-K	53.67	<u>49.50</u>	11.21	<b>62.00</b>	44.00	45.78	47.33	<u>38.50</u>	6.61	<u>54.00</u>	28.00	38.02
HiTEC-ICL	53.33	<b>52.00</b>	<b>30.00</b>	<u>58.00</u>	<u>50.00</u>	49.33	39.67	33.50	19.00	50.00	26.00	34.92
ToolGT (Prompting)	<u>60.67</u>	48.00	<b>30.00</b>	56.00	<b>54.00</b>	<b>52.50</b>	<b>56.67</b>	36.50	<u>22.00</u>	52.00	<b>50.00</b>	<u>46.58</u>
<b>Tool-DC (TF)</b>	<b>61.00</b>	40.50	23.00	50.00	<u>50.00</u>	<u>49.42</u>	<u>55.67</u>	<b>46.00</b>	<b>28.00</b>	<b>60.00</b>	<u>36.00</u>	<b>48.17</b>
<b><i>Qwen2.5-7B-Instruct</i></b>												
GT_Funs	71.00	61.00	52.00	68.00	62.00	65.17	71.00	61.00	52.00	68.00	62.00	65.17
All_Funs	62.33	63.50	45.00	64.00	64.00	59.92	65.33	54.50	42.00	62.00	56.00	58.58
Top-K	61.33	58.00	0.00	64.00	<b>62.00</b>	50.83	53.67	50.50	17.37	52.00	46.00	46.31
HiTEC-ICL	65.33	61.00	48.00	<u>66.00</u>	60.00	61.33	58.67	48.00	44.00	62.00	<b>58.00</b>	54.67
ToolGT (Prompting)	<b>79.67</b>	<u>61.50</u>	<u>49.00</u>	<u>66.00</u>	<b>62.00</b>	<b>68.92</b>	<b>71.33</b>	<u>59.50</u>	<b>53.00</b>	<u>64.00</u>	<u>56.00</u>	<b>64.42</b>
<b>Tool-DC (TF)</b>	<u>68.00</u>	<b>66.00</b>	<b>50.00</b>	<b>70.00</b>	58.00	<u>64.00</u>	<u>60.00</u>	<b>65.00</b>	<u>48.00</u>	<b>66.00</b>	54.00	<u>58.83</u>

Table 10: Detailed results of different training-based methods on Standard ACEBench, referred to Figure 5.

Methods	Atom	Single Turn	Multiple Turn	Similar API	Profile	Overall
<b>Qwen2.5-7B-Instruct</b>	62.33	63.50	45.00	64.00	64.00	59.92
-w Vanilla SFT	67.67	65.50	7.14	70.00	64.00	57.11
-w Tool-DC (TB)	76.33	47.00	22.94	74.00	64.00	<b>61.32</b>
<b>Llama-3.2-3B-Instruct</b>	33.33	21.00	8.00	34.00	50.00	28.50
-w Vanilla SFT	65.00	43.00	5.22	52.00	48.00	48.87
-w Tool-DC (TB)	72.67	43.00	12.71	56.00	38.00	<b>53.45</b>
<b>Qwen3-4B-Instruct-2507</b>	55.67	43.00	22.13	52.00	54.00	47.52
-w Vanilla SFT	75.00	65.50	16.67	70.00	66.00	62.53
-w Tool-DC (TB)	67.67	59.50	52.00	68.00	64.00	<b>63.42</b>