

# CITI: Enhancing Tool Utilizing Ability in Large Language Models Without Sacrificing General Performance

Yupu Hao<sup>1,2</sup>, Pengfei Cao<sup>1,2</sup>, Zhuoran Jin<sup>1,2</sup>, Huanxuan Liao<sup>1,2</sup>, Yubo Chen<sup>1,2</sup>,  
Kang Liu<sup>1,2,3\*</sup>, Jun Zhao<sup>1,2</sup>

<sup>1</sup> The Key Laboratory of Cognition and Decision Intelligence for Complex Systems,  
Institute of Automation, Chinese Academy of Sciences, Beijing, China

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

<sup>3</sup> Shanghai Artificial Intelligence Laboratory

{haoyupu2023, liaohuanxuan2023}@ia.ac.cn, {pengfei.cao, zhuoran.jin, yubo.chen, kliu, jzhao}@nlpr.ia.ac.cn

## Abstract

Tool learning enables Large Language Models (LLMs) to interact with the external environment by invoking tools, enriching the accuracy and capability scope of LLMs. However, previous works predominantly focus on improving the model’s tool-utilizing accuracy and the ability to generalize to new, unseen tools, excessively forcing LLMs to adjust specific tool-invoking pattern without considering the harm to the model’s general performance. This deviates from the actual applications and original intention of integrating tools to enhance the model. To tackle this problem, we dissect the capability trade-offs by examining the hidden representation changes and the gradient-based importance score of the model’s components. Based on the analysis result, we propose a Component Importance-based Tool-utilizing ability Injection method (CITI). According to the gradient-based importance score of different components, it alleviates the capability conflicts caused by the fine-tuning process by applying distinct training strategies to different components. CITI applies Mixture-Of-LoRA (MOLoRA) for important components. Meanwhile, it fine-tunes the parameters of a few components deemed less important in the backbone of the LLM, while keeping other parameters frozen. CITI can effectively enhance the model’s tool-utilizing capability without excessively compromising its general performance. Experimental results demonstrate that our approach achieves outstanding performance across a range of evaluation metrics.

**Code** — <https://github.com/hypasd-art/CITI>

**Extended version** — <https://arxiv.org/abs/2409.13202>

## 1 Introduction

Large Language Models (LLMs) have demonstrated significant capabilities in understanding language and generating human-like text (Zhao et al. 2023; Li et al. 2023a). Despite their impressive performance, LLMs still face limitations, such as the tendency to generate hallucinated information (Huang et al. 2023; Sun et al. 2024) and an inability to interact with the real physical world (Qin et al. 2023). To tackle these limitations and expand the model’s capabilities scope beyond traditional natural language tasks, there is a growing




Figure 1: The model’s performance in tool-utilizing and general tasks, FT represents full parameter fine-tuning on the dataset API-Bank (Li et al. 2023b).

interest in enabling LLMs to interact with the external environment by equipping various tools (Qin et al. 2024; Schick et al. 2023; Gao et al. 2023a; Chen et al. 2023; Nakano et al. 2022; Lu et al. 2023).

There have been a substantial number of benchmarks evaluating different tool-utilizing aspects (Ye et al. 2024; Li et al. 2023b; Zhan et al. 2024), along with a suite of proposed methodologies (Qin et al. 2024; Tang et al. 2023). These approaches predominantly leverage in-context learning (Shen et al. 2023; Wu et al. 2023) or fine-tuning techniques (Shen et al. 2024; Gao et al. 2023b; Chen et al. 2024) to equip the model with tool-invoking ability. In contrast to in-context learning, fine-tuning has gained popularity for its capability to deeply integrate task-specific knowledge into model parameters, enhancing its adaptability to specific tasks. However, as illustrated in Figure 1, we find that excessive fine-tuning on tool learning datasets severely diminishes the model’s broader cognitive competencies, leading to a decline in the model’s general abilities and catastrophic forgetting about parametric knowledge.

Tools serve models rather than models serving tools. Consequently, it is crucial for the tool-augmented models to incorporate tool-utilizing ability while simultaneously preserving their original general versatility. Thus, an intriguing question is *what happens in LLMs when fine-tuning on tool learning dataset?* We dissect the question from two

\*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

aspects: hidden representation and components. In detail, **from hidden representation perspective**, when subtracting the originally hidden representations from those representations tuned on the tool learning dataset, we observe a significant phenomenon: the increments, derived by subtraction, exhibit a **Co-directional shift** in the hidden state space, which implies there is a strong correlation in the direction of increments between various tool-unrelated tasks (e.g. mathematics, coding) and tool-related tasks. **From components perspective**, for different capabilities, we calculate the gradient-based importance score ranking of the model’s linear modules (referred to as the linear components) based on corresponding datasets respectively. We discover that the important linear components in the ranking play a more vital role in corresponding ability expression while the unimportant components contribute less to this ability. Concurrently, we find that the importance rankings of different abilities are highly consistent. This suggests that certain components are pivotal in the expression of the model’s abilities across a broad range of tasks, whereas some components exert a more limited influence on most tasks. And we further explore the impact of fine-tuning different components based on tool-related importance ranking. It is found that optimizing important components results in a significant decrease in general performance, and the model may not fully learn the knowledge of calling tools only by optimizing unimportant components.

Based on the insights from our analysis, we propose a novel **Component Importance-based Tool-utilizing ability Injection method (CITI)**, which applies Mixture-of-LoRA adapters to important components and adopts full parameter fine-tuning on unimportant components. Firstly, we identify the gradient-based importance score of all the linear components in the model. Secondly, for the important linear components, we integrate a set of Mixture-Of-LoRA (MOLoRA) adapters to absorb the tool-invoking knowledge. To handle the Co-directional shift phenomenon, we design a router network to separate the tool-related and tool-unrelated inputs, reducing the impact on the LLM’s backbone. Thirdly, for the unimportant linear components, we employ full parameter fine-tuning to take full advantage of more parameters. Specifically, Our training process adopts a three-stage approach. In the initial training stage, *Router Pre-training*, we pre-train the router network in MOLoRA to teach it to distinguish tool-related and tool-unrelated inputs. In the second stage, *MOLoRA Improvement*, we concentrate on fine-tuning the MOLoRA adapters while freezing the backbone of LLM. In the third stage, *Unimportant Components Optimization*, **CITI** fine-tunes a small part of unimportant components within the backbone to improve the model’s performance while maintaining its general abilities.

To summarize, the contributions of our work include:

- We discover the phenomenon that fine-tuning LLMs on tool learning datasets significantly influences the general abilities, and we have conducted a thorough inspection of this trade-off phenomenon by analyzing the perspective of hidden representation and components. Adequate experiments and analysis demonstrate the factors resulting in the catastrophic forgetting of general abilities.

- We propose a **Component Importance-based Tool-utilizing ability Injection method (CITI)**, alleviating the trade-offs by applying distinct strategies to different components identified by gradient-based importance score of the model.
- Our method achieves competitive tool-utilizing results across two tool learning datasets. Additionally, it preserves general performance which is 7.59% superior to LoRA and 31.95% superior to full parameter fine-tuning on average in dataset API-Bank (Li et al. 2023b), and it is also 8.96% better than LoRA and 29.03% better than full parameter fine-tuning in dataset ToolAlpaca (Tang et al. 2023), effectively enhancing tool-utilizing ability while preserving LLMs general performance.

## 2 Analysis of Hidden Representation and Components

Hidden representation and components play a crucial role in determining the hidden state of inputs as they propagate from lower layers to higher layers, helping us understand the inner working process of LLM. Specifically, we conduct the following analytical experiments and assess general abilities on a range of frequently-used datasets, including GSM8K (GSM) (Cobbe et al. 2021), HumanEval (HE) (Chen et al. 2021), TriviaQA (TQA) (Joshi et al. 2017), MT-Bench (MT) (Zheng et al. 2023).

### 2.1 Hidden Representation Perspective

The hidden representation refers to the output vector by each layer of the auto-regressive transformer model. Our goal is to explore the changes in these representations following fine-tuning with the tool learning instructions.

**Analytical Approach** We introduce the concept of Incremental Change of Capability (*ICC*) to quantify the changes in the hidden representation of the fine-tuning process. Specifically, for ability  $t$  and instruction  $m$  within a processed dataset  $\mathcal{DM}_t$ ,  $H_{REF}^l(m)$  is the original hidden representation of the last token in instruction  $m$  in the untrained model at layer  $l$ , while  $H_{SFT}^l(m)$  represents the representation of the same token after fine-tuning. And the  $ICC_t^l$  can be represented as:

$$ICC_t^l = \frac{1}{|\mathcal{DM}_t|} \sum_{m \in \mathcal{DM}_t} H_{SFT}^l(m) - H_{REF}^l(m) \quad (1)$$

The average of the vector’s increment computed by  $\mathcal{DM}_t$  is considered to represent the changes in hidden space of ability  $t$ . The cosine similarity of these increments is then computed to assess the relationship of the vector’s change between ability  $a$  and  $b$  in the model’s hidden state space:

$$Sim(a, b) = \frac{ICC_a \cdot ICC_b}{|ICC_a| \times |ICC_b|} \quad (2)$$

**Experimental Results** To get representations, we randomly truncate a portion of the golden answer and append it after the instruction, creating a “final input”. For a specific task, we sample 1000 “final input” from the dataset to

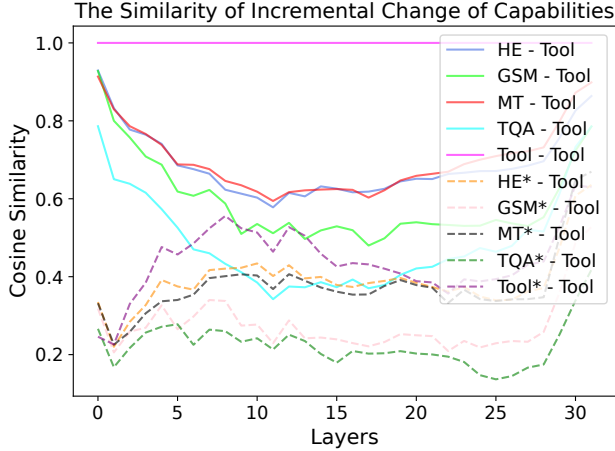


Figure 2: Cosine similarity of *ICC* between the input of different layers of Feed-Forward Network (FFN) in model Meta-Llama-3-8B-Instruct, where the notation with an asterisk (\*) represents *ICC* fine-tuned on the code-related dataset (e.g. TQA\* represents *ICC* of TriviaQA trained by code dataset), and no asterisk (\*) represents *ICC* fine-tuned on tool learning dataset.

construct  $\mathcal{DM}_t$  and put them into the model. The vector corresponding to the last token of the “final input”  $m$  of different layers is considered as the hidden representation vector  $H(m)$ . Then we compute the similarity of different *ICC*.

In Figure 2, the solid lines represent the similarity between tool-related increments  $ICC_{tool}$  with general task increments  $ICC_t$ , where LLM is fine-tuned on the tool learning dataset. For better comparison, we train a model on a code generation dataset and compute its increments  $ICC_{t*}$ . Here  $t$  represents a specific general ability.

We find that **there is a Co-directional shift phenomenon in the model’s hidden state space**. The changing direction of the  $ICC_t$  on the general task is positively correlated with the direction of  $ICC_{tool}$  in the model’s hidden state space, compared with the dashed line. Our hypothesis is that after fine-tuning by tool learning datasets, the model cannot distinguish tool-related and tool-unrelated inputs in the hidden state space correctly, resulting in more similar activation states for different types of instructions, which affects the general performance of the model.

## 2.2 Components Perspective

**Analytical Approach** Large language model, composing the decoder-base architecture of transformer, typically comprises multiple layers. Inspired by Bansal et al. (2023), we adopt the gradient-based importance score to identify the important components.

In the auto-regressive language models, for the input  $x$  and corresponding labels  $y$  in given dataset  $\mathcal{D}$ , the loss function  $\mathcal{L}$  to optimize the model is calculated as follows:

$$\mathcal{L}(\mathcal{D}, \theta) = \sum_{(x,y) \in \mathcal{D}} \log p_{\theta}(y|x) \quad (3)$$

Dataset	Replace Ratio				Vanilla
	T-20%	T-50%	D-20%	D-50%	
GSM8K	64.59	40.56	79.15	74.98	79.45
HumanEval	51.83	25.00	58.24	20.12	59.15
TriviaQA	64.95	58.76	64.95	63.96	64.82
MT-Bench	72.53	50.97	81.75	73.19	80.13

Table 1: The performance of the model on the corresponding tasks after components parameter replacement. T-x% represents taking the top x% components of the model to be replaced by importance ranking order, and D-x% is the opposite. Vanilla means test on Meta-Llama-3-8B-Instruct.

where  $\theta$  denotes the complete set of the model’s parameters. For specific parameters set  $\theta_h$ , the importance score  $\mathcal{I}_h$  is computed as:

$$\begin{aligned} \mathcal{I}_h(\mathcal{D}, \theta) &= |\mathcal{L}(\mathcal{D}, \theta) - \mathcal{L}(\mathcal{D}, \theta | \theta_h = 0)| \\ &= \left| \frac{\partial \mathcal{L}}{\partial \theta_h} (\theta_h - 0) + \frac{1}{2!} \frac{\partial^2 \mathcal{L}}{\partial \theta_h^2} (\theta_h - 0)^2 + \dots \right| \quad (4) \\ &\approx \left| \theta_h^T \frac{\partial \mathcal{L}}{\partial \theta_h} \right| \end{aligned}$$

To accelerate the calculating speed, the score is estimated by calculating the first-order derivative of the Taylor expansion in the above formula.

**Experimental Results** To examine the changes of components across various general tasks during the tool learning process, we select a subset of instructions from the training set as a proxy for a specific general ability  $t$ . Employing Equation (4), we determine the importance scores for the linear components in the vanilla model (Meta-Llama-3-8B-Instruct) with these data. In order to explore the effect of the importance of components, we conduct three experiments as follows:

**Experiment 1:** We prioritize components based on their importance scores and substitute the parameters of these components in the vanilla model with the corresponding weights from the model that has been fine-tuned on the tool learning dataset API-Bank, leaving all other parameters intact. We then assess the model’s performance after replacement.

In Table 1, we observe that tool learning influences other abilities expressions in most cases. Following the replacement operation, the model’s performance (GSM8K, HumanEval, and MT-Bench) significantly decreased compared to the vanilla model, indicating that fine-tuning on the tool learning dataset stimulated tool-invoking abilities in these components while suppressing their original inherent abilities. **More significant performance impacts occur when high-importance components are altered.** The performance of the model significantly drops after replacing weights with higher importance. Moreover, this phenomenon becomes more pronounced as the replacement ratio increases.

**Experiment 2:** Furthermore, to delve into the inter-task relationships of components’ importance ranking, we apply

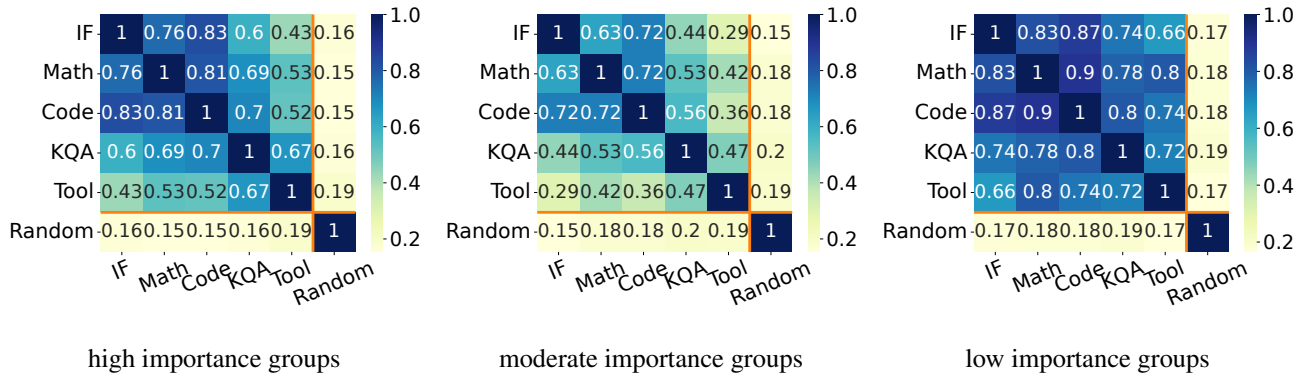


Figure 3: The Jaccard Index of different groups, *Random* means sorting the components randomly.

the Jaccard index to analyze the correlation between the importance ranking of model components. The Jaccard index can be represented as:  $J(A, B) = |A \cap B| / |A \cup B|$ , used to compare the differences and similarities between two samples  $A$  and  $B$ . Specifically, based on the importance ranking of components associated with ability  $t$ , we segment the components into three groups evenly: high, moderate, and low. Then we compute the Jaccard Index between different abilities for each group respectively. For example,  $J(C_h^t, C_h^m)$  represents the Jaccard Index between components groups of high importance calculated by ability  $t$  and  $m$  respectively.

As shown in Figure 3, **there is a high correlation across different tasks of their components importance ranking**. Certain components are pivotal in the expression of the model’s abilities across a broad range of tasks, whereas part of the components exert a limited influence on most tasks.

**Experiment 3:** We try to inject the tool-invoking knowledge into the less important components as they store less tool-related ability. Considering a components set  $H$  containing all linear modules of the network, for tool learning dataset  $\mathcal{D}_{tool}$ , we compute the component level gradient-based importance scores for component  $h$  as follows:

$$\mathcal{M}_h(\theta) = \frac{\mathcal{I}_h(\mathcal{D}_{tool}, \theta)}{\sum_{i \in H} \mathcal{I}_i(\mathcal{D}_{tool}, \theta)} \quad (5)$$

In order to better combine the importance ranking of components and the training process, we select 20% (Top, Down, Random) components in the model based on component importance  $\mathcal{M}_h$  and train them using full parameter fine-tuning and LoRA methods, while keep other parameters frozen. Then we analyze the fine-tuning result in detail.

As illustrated in Table 2, we find that fine-tuning the components with lower importance score  $\mathcal{M}_h$  have the least impact on the general performance, which confirms the existence of high correlations between different abilities of the models because  $\mathcal{M}_h$  is computed by tool learning dataset. At the same time, it is observed that fine-tuning components with lower importance may not perform outstandingly in terms of tool-invoking ability. If only unimportant parameters were fine-tuned, the model can not thoroughly understand the format of tool callings. But “Top” and “Random”

Model	Tool Ability		General Ability			
	C	R	GSM	HE	TQA	MT
LoRA						
Top	59.95	<b>40.85</b>	74.00	53.66	64.09	<b>78.56</b>
Down	46.21	38.84	<b>78.39</b>	<b>55.49</b>	63.71	74.00
Random	<b>64.18</b>	36.87	75.36	54.27	<b>65.11</b>	76.40
Full parameter						
Top	<b>66.11</b>	<b>44.04</b>	71.65	50.61	<b>64.18</b>	73.15
Down	57.38	40.36	<b>81.20</b>	<b>54.27</b>	59.96	<b>74.11</b>
Random	59.31	37.99	74.30	50.61	63.77	72.48

Table 2: The fine-tuning experiments based on gradient-based importance score on the dataset API-Bank. The evaluation metrics C and R refer to Section 4.3.

do not have such problems. This inspires us only fine-tuning the unimportant component sometimes may be not optimal.

### 3 Methodology

Instead of only focusing on the precise tool-calls, we propose a **Component Importance-based Tool-utilizing ability Injection method (CITI)**, injecting the tool-utilizing ability into the model while maintaining its general abilities. Based on the analysis experiments, in order to fully utilize the model’s capability, we apply Mixture-of-LoRA adapters to important components and fine-tune a subset of unimportant components. As shown in Figure 4, our method mainly consists of three training stages: *Router Pre-training (RP)*, *MOLoRA Improvement (MI)*, *Unimportant Components Optimization (UCO)*. We will illustrate each technique in detail.

#### 3.1 Mixture of LoRA

Based on the analysis of component importance, we want to fine-tune both important and unimportant components to achieve a comprehensive grasp of tool-utilizing knowledge. However, adjusting critical components through fine-tuning can sometimes be detrimental to the overall network performance as the Co-directional shift phenomenon in hidden space. To alleviate this influence, a straightforward idea is to implement a gating function to distinguish tool-related

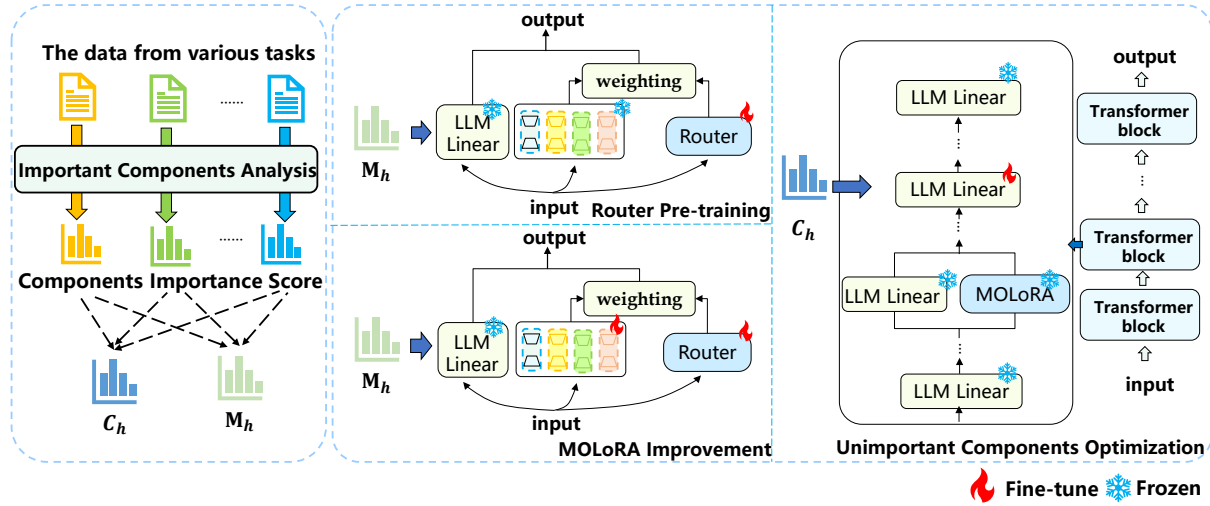


Figure 4: The overall architecture of **CITI**. The MOLoRA adapters are applied to important components identified by  $\mathcal{M}_h$ , and unimportant components identified by  $\mathcal{C}_h$  are fine-tuned with all parameters.

and tool-unrelated inputs, and then assign different specialized experts to handle them separately. To mitigate this, we implement MOLoRA adapters on these pivotal components with higher  $\mathcal{M}_h$ .

Inspired by Zhao et al. (2024), we apply a router network with an additional neuron in MOLoRA to separate the inputs, allowing tool-related inputs are produced by the combine of the LLM backbone and external LoRA experts. And the tool-unrelated data is predominantly processed by the backbone itself.

Meanwhile, mixing data during downstream task fine-tuning helps alleviate catastrophic forgetting of models and helps to avoid over-fitting. To train the router network and avoid catastrophic forgetting problems, we mix data from other fields during the tool learning process.

For router network, we use the linear module as the gate function  $G(x)$ . Denote the input features as  $x$ , the formula is as follows:

$$G(x) = \text{Softmax}(W \cdot x) \quad (6)$$

where  $W \in \mathbb{R}^{d \times (N+1)}$  is the trainable matrix,  $d$  represents the dimensionality of the input,  $N$  is the number of LoRA experts, The term  $N + 1$  means the incorporation of an additional neuron, representing the probability that  $x$  is tool-unrelated data, which is used to suppress the model's call to external expert weights. The output of MOLoRA is:

$$h = W_0 x + \sum_{i=0}^N G(x)[i+1] B_i A_i x \quad (7)$$

where  $G(x)[i]$  is the  $i$ -th element in router output, and the  $G(x)[0]$  is not used in forward process. And matrices  $A \in \mathbb{R}^{r \times k}$  and  $B \in \mathbb{R}^{d \times r}$  are the trainable linear of LoRA experts, where  $r \ll \min(d, k)$  is the rank of LoRA experts.

To achieve the separation of inputs, we divide the data into tool-related and unrelated with different task types. And

we modify the Localized Balancing Constraint proposed by Dou et al. (2024), which softly constrains the experts to focus on different tasks. We assign the importance score to each neuron to control the weighting of different experts, the rule of importance matrix  $I$  is:

$$I_n = \begin{cases} [1 + \delta, \underbrace{1 - \delta, \dots, 1 - \delta}_N] & \text{if } x_n \in \mathcal{D}_{tool} \\ [1 - \delta, \underbrace{1 + \delta, \dots, 1 + \delta}_N] & \text{otherwise} \end{cases} \quad (8)$$

Here,  $\delta \in [0, 1]$  controls the degree of imbalance between router output,  $x_n$  is the  $n$ -th input.

We weigh the importance matrix  $I$  and the output of the gating function  $G$ , denoted as  $Z = I \odot G$ . And the routing loss  $\mathcal{L}_r$  of router network is:

$$\mathcal{L}_r = \frac{\sigma^2(Z)}{\mu(Z)} \quad (9)$$

where  $\sigma^2(Z)$  and  $\mu(Z)$  are the variance and mean of  $Z$ .  $\mathcal{L}_r$  ensure that for tool-unrelated data  $x$ ,  $G(x)[0]$  is relatively large, and vice versa,  $G(x)[0]$  is relatively small. adjusting the weight of LoRA experts based on the data types.

### 3.2 Unimportant Components Optimization

For unimportant components, as updating these parameters has a relatively less impact on the model's general capabilities, we utilize full parameter fine-tuning to stimulate the model's tool-invoking ability.

We select unimportant components for general abilities by averaging the importance of different tasks, defined as general abilities importance  $C_h(\theta)$ . For a general abilities set  $T$ , the formula is as follows:

$$C_h(\theta) = \sum_{t \in T} \frac{\mathcal{I}_h(\mathcal{D}_t, \theta)}{\sum_{i \in H} \mathcal{I}_i(\mathcal{D}_t, \theta)} \quad (10)$$

Components with lower  $C_h(\theta)$  are selected. These components are not frozen during the training process.

### 3.3 Training Strategy

In summary, the overall loss function of **CITI** is:

$$\mathcal{L} = - \sum_{(x,y) \in \mathcal{D}} \log P(y | x; \theta_b, \theta_e, \theta_r) + \beta \cdot \mathcal{L}_r \quad (11)$$

Where  $\theta_b$  is the parameters of the backbone,  $\theta_e$  represents the MOLoRA experts, and  $\theta_r$  is the router network.  $\mathcal{D}$  denotes the mixture of tool learning dataset and other instructions sampled from general ability datasets.

To optimize our model effectively, we have devised a three-stage training strategy, delineated as follows:

**Stage-1 RP:** We integrate MOLoRA based on the importance of component gradients. During this phase, the core model parameters and LoRA experts are kept constant, with only  $\theta_r$  being subjected to training.

**Stage-2 MI:** The  $\theta_e$  and  $\theta_r$  are subjected to training. We fine-tune the parameters in LoRA experts.

**Stage-3 UCO:** Here, we immobilize the parameters of the MOLoRA adapters, including  $\theta_e$  and  $\theta_r$ , as well as the important parameters in the backbone identified by importance ranking. We then proceed to train a select subset of the backbone parameters chosen by  $C_h$ .

## 4 Experiments

### 4.1 Datasets

We conduct the experiments on two tool learning benchmarks: API-Bank (Li et al. 2023b) and ToolAlpaca (Tang et al. 2023). Additionally, we assess general abilities through experiments on the datasets mentioned above.

### 4.2 Implementation Details

We apply MOLoRA to the top 20% components with the highest importance ranking sorted by  $\mathcal{M}_h$  and fine-tune down 10% components with lowest importance ranking sorted by  $C_h$ .

### 4.3 Evaluation Metrics

For the API-Bank dataset, we follow the evaluation metrics proposed by Li et al.(2023b), including the **C**orrectness of API calls and the **R**OUGE-L to test quality of the responses. For the ToolAlpaca dataset, we utilize the GPT-4 to evaluate the tool-invoking process in the real-world testing subset and follow the original metrics: **P**rocedure: The correctness of tool utilizing procedure, **R**esponse: The quality of final response, **O**verall: Whether procedure and response are both correct.

### 4.4 Baselines

We conduct experiments with excellent models: **Meta-Llama-3-8B-Instruct**, **Phi-3-mini-128k-instruct** and **Mistral-7B-Instruct-v0.2**. We compare our method with the following baselines: (1) **FT**: fine-tune the model with full parameters on the tool learning dataset; (2) **LoRA**: fine-tune the model with LoRA adapters on the tool learning dataset. The training data is without data mixing for our baselines.

Model	Tool Ability		General Ability			
	C	R	GSM	HE	TQA	MT
Llama-3	33.25	14.67	79.45	59.15	64.82	80.13
FT	56.48	18.71	26.38	0.00	50.07	33.88
LoRA	<b>65.98</b>	<b>38.62</b>	76.88	53.05	<b>64.84</b>	68.00
<b>CITI</b>	58.92	34.51	<b>77.33</b>	<b>54.88</b>	64.69	<b>76.94</b>
Phi-3	44.80	29.40	82.71	60.37	53.08	81.28
FT	<b>61.75</b>	35.23	73.09	51.83	53.32	66.81
LoRA	57.00	35.19	75.82	0.00	52.34	68.44
<b>CITI</b>	58.02	<b>37.21</b>	<b>77.86</b>	<b>61.59</b>	<b>55.23</b>	<b>72.19</b>
Mistral	53.66	32.39	44.81	35.37	59.21	75.66
FT	47.88	2.18	2.05	0.00	13.50	10.63
LoRA	58.66	<b>43.31</b>	46.32	37.80	61.54	68.31
<b>CITI</b>	<b>62.00</b>	33.53	<b>53.30</b>	<b>38.41</b>	<b>62.77</b>	<b>69.28</b>

Table 3: The overall results on the dataset API-Bank.

Model	Tool Ability			General Ability			
	P	R	O	GSM	HE	TQA	MT
Llama-3	14.91	21.05	12.28	79.45	59.15	64.82	80.13
FT	65.79	59.65	55.26	30.33	0.61	55.97	50.81
LoRA	<b>76.32</b>	74.56	71.05	73.46	31.71	<b>65.18</b>	73.75
<b>CITI</b>	<b>76.32</b>	<b>78.07</b>	<b>72.81</b>	<b>75.89</b>	<b>56.71</b>	64.82	<b>76.67</b>
Phi-3	2.63	1.75	1.75	82.71	60.37	53.08	81.28
FT	67.54	68.42	64.04	71.95	54.88	53.58	65.06
LoRA	<b>70.18</b>	71.05	<b>66.67</b>	74.91	0.61	53.04	<b>73.22</b>
<b>CITI</b>	64.04	<b>72.81</b>	61.40	<b>77.33</b>	<b>62.20</b>	<b>55.27</b>	72.59
Mistral	14.04	14.91	13.16	44.81	35.37	59.21	75.66
FT	42.98	41.23	30.70	4.85	0.00	12.10	15.69
LoRA	72.81	73.68	67.54	44.81	34.76	62.99	68.19
<b>CITI</b>	<b>78.07</b>	<b>76.32</b>	<b>71.05</b>	<b>51.48</b>	<b>36.59</b>	<b>63.21</b>	<b>71.44</b>

Table 4: The overall results on the dataset ToolAlpaca.

### 4.5 Overall Results

As shown in Table 3 and Table 4, **CITI** demonstrates effective performance in preserving general abilities in two datasets. For instance, it achieves a 7.59% improvement compared to LoRA and 31.95% compared to FT in API-Bank for general abilities on average. For tool-utilizing ability, **CITI** shows competitive results in two datasets compared to the baselines. Noticing that in the dataset API-Bank, **CITI** has poorer tool-invoking ability compared to baselines (e.g. ROUGE-L score in Mistral). We suggest there are two possible reasons: (1) The  $\lambda$  in MOLoRA may not be optimal, and insufficient training data may result in underfitting or overfitting of LoRA adapters. (2) The evaluation metrics are not comprehensive. API-Bank asks the model to generate tool calls or responses based on conversation history, but the path for the model to obtain final correct answers is not unique.

Additionally, we notice that LoRA outperforms full parameter fine-tuning in both tool utilization and general abilities in most cases. This may be because full parameter fine-tuning can easily lead to overfitting.

Model	Tool Ability		General Ability			
	C	R	GSK	HE	TQA	MT
FT + DM	56.61	39.67	54.44	29.88	45.28	60.41
LoRA + DM	<b>63.03</b>	38.66	75.13	51.22	63.96	74.31
UCO	59.31	38.63	<b>79.68</b>	54.88	64.07	75.09
RP + MI	57.89	35.50	77.94	54.88	<b>64.74</b>	76.56
w/o MLoRA	61.87	<b>39.97</b>	73.09	54.27	63.09	<b>77.91</b>
w/o $\mathcal{L}_r$	61.49	38.29	75.28	<b>57.93</b>	63.90	76.81
<b>CITI</b> (ours)	58.92	34.51	77.33	54.88	64.69	76.94
w/o <b>RP</b>	58.79	38.13	76.12	52.44	63.35	75.22

Table 5: The ablation results on the dataset API-Bank.

Model	Tool Ability			General Ability			
	P	R	O	GSM	HE	TQA	MT
FT + DM	60.53	63.16	57.02	52.92	23.78	50.58	58.94
LoRA + DM	73.68	73.68	68.42	74.53	53.66	<b>65.06</b>	74.94
UCO	69.30	68.42	64.04	<b>80.52</b>	55.49	64.06	73.63
RP + MI	76.32	75.44	69.30	74.68	54.27	64.49	<b>77.63</b>
w/o MLoRA	70.18	71.05	65.79	73.09	55.49	64.31	77.31
w/o $\mathcal{L}_r$	71.05	71.93	64.91	77.48	55.49	65.00	77.38
<b>CITI</b> (ours)	<b>76.32</b>	<b>78.07</b>	<b>72.81</b>	75.89	<b>56.71</b>	64.82	76.67
w/o <b>RP</b>	70.18	68.42	64.91	77.48	54.27	64.95	76.81

Table 6: The ablation results on the dataset ToolAlpaca.

## 4.6 Ablation Studies

To demonstrate the effectiveness of our approach, we conduct ablation studies on model Meta-Llama-3-8B-Instruct as shown in Table 5 and Table 6. Where w/o MLoRA means replacing the MLoRA by LoRA adapter, w/o  $\mathcal{L}_r$  means applying normal MLoRA without additional neuron and importance matrix  $I$  in the router. DM represents data mixing in the training data. The experimental results show that compared to FT and LoRA with data mixing, CITI still exhibits advantages in general abilities. Additionally, we analyze the contributions of each module in CITI. Compared to baselines, MI and UCO modules have shown distinct advantages in maintaining the model’s general abilities separately.

We can find that: (1) The UCO module demonstrates that fine-tuning the unimportant components selectively can enhance the model’s capability to utilize new tools while maintaining its original performance, especially on GSM8K within two datasets. (2) Comparing RP + MI with w/o MLoRA, we discover that by incorporating MLoRA, rather than simply adding LoRA adapter, the decline in general abilities can be alleviated. (3) The results of only w/o  $\mathcal{L}_r$  outperform w/o MLoRA in most cases, further indicating the MLoRA structure is effective. Additionally, we find that some results in w/o  $\mathcal{L}_r$  are better than RP + MI. We deduce this is because the adapters absorb general knowledge within the general instructions in mixed training set, but they play little impact in MI because the router separates them into frozen LLM backbone during training stage. (4) The results of w/o RP demonstrate that *Router Pre-training* offers benefits, especially in tool-utilizing in ToolAlpaca and general performance in API-Bank, indicating that pre-training

the router network may accelerate model convergence especially when training data is constrained. (5) By applying UCF to further fine-tune the model following MI, CITI can maintain or even increase its general abilities on two datasets compared to RP + MI. And the tool-utilizing ability has further significantly improved on the dataset ToolAlpaca.

## 5 Related Work

### 5.1 Tool Learning

Tool Learning aims to enable LLMs to use external tools to enhance models’ ability, which shows great power in dealing with various tasks (Qin et al. 2023; Qu et al. 2024). Fine-tuning technology (Tang et al. 2023; Qin et al. 2024; Wang et al. 2024) is a primary approach to learn tool-utilizing pattern, which can solidify knowledge into models comparing to in-context learning (Shen et al. 2023; Wu et al. 2023). In order to enable the model to invoke tools while maintaining its original performance, we introduce **CITI**, a framework to balance the conflict.

### 5.2 Components Importance Analysis

It is important to evaluate the importance of parameters in model pruning (Zhang et al. 2022) and works related to interpretability (Michel, Levy, and Neubig 2019; Zhang et al. 2024; Jin et al. 2024). The gradient-based importance scores (Bansal et al. 2023) assesses the importance of parameters by measuring the resultant impact on the loss function following the parameter’s elimination. In our work, We use importance scores to identify and select components of tool-utilizing and general tasks.

### 5.3 Mixture Of LoRA

Recently, there are substantial works utilizing Mixture-of-LoRA Experts (MLoRA) architecture to improve the performance (Dou et al. 2024; Wu, Huang, and Wei 2024; Yang et al. 2024). These works combine the Mixture-of-Experts (MOE) (Fedus, Zoph, and Shazeer 2022; Shazeer et al. 2017) architecture with Low-Rank Adaptation (LoRA) (Hu et al. 2022; Wu, Huang, and Wei 2024; Gou et al. 2024). MLoRA keeps the backbone networks frozen and fine-tunes the adapters, achieving significant performance while alleviating catastrophic forgetting (Zhao et al. 2024; Li et al. 2024; Dou et al. 2024). We modify the router network to distinguish the tool-related and tool-unrelated inputs.

## 6 Conclusion

In this study, we explore the trade-offs between the model’s tool-utilizing ability and its general task performance following fine-tuning on a tool learning dataset. We conduct a deep analysis of this phenomenon from hidden representation and components perspectives and propose a new approach **CITI**, which integrates MLoRA adapters into the critical components of the model and selectively fine-tunes a limited number of parameters of less critical components within the backbone of the model. Through extensive experiments, we demonstrate the validity of **CITI** in acquiring tool-related knowledge without excessively harming its performance on general tasks.

## Acknowledgments

This work is supported by the National Key Research and Development Program of China (No. 2022ZD0160503). This work is supported by the National Natural Science Foundation of China (No.U24A20335, No.62406321).

## References

- Bansal, H.; Gopalakrishnan, K.; Dingliwal, S.; Bodapati, S.; Kirchoff, K.; and Roth, D. 2023. Rethinking the Role of Scale for In-Context Learning: An Interpretability-based Case Study at 66 Billion Scale. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 11833–11856. Toronto, Canada: Association for Computational Linguistics.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; de Oliveira Pinto, H. P.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; Ray, A.; Puri, R.; Krueger, G.; Petrov, M.; Khlaaf, H.; Sastry, G.; Mishkin, P.; Chan, B.; Gray, S.; Ryder, N.; Pavlov, M.; Power, A.; Kaiser, L.; Bavarian, M.; Winter, C.; Tillet, P.; Such, F. P.; Cummings, D.; Plappert, M.; Chantzis, F.; Barnes, E.; Herbert-Voss, A.; Guss, W. H.; Nichol, A.; Paino, A.; Tezak, N.; Tang, J.; Babuschkin, I.; Balaji, S.; Jain, S.; Saunders, W.; Hesse, C.; Carr, A. N.; Leike, J.; Achiam, J.; Misra, V.; Morikawa, E.; Radford, A.; Knight, M.; Brundage, M.; Murati, M.; Mayer, K.; Welinder, P.; McGrew, B.; Amodei, D.; McCandlish, S.; Sutskever, I.; and Zaremba, W. 2021. Evaluating Large Language Models Trained on Code. arXiv:2107.03374.
- Chen, S.; Wang, Y.; Wu, Y.-F.; Chen, Q.-G.; Xu, Z.; Luo, W.; Zhang, K.; and Zhang, L. 2024. Advancing Tool-Augmented Large Language Models: Integrating Insights from Errors in Inference Trees. arXiv:2406.07115.
- Chen, W.; Ma, X.; Wang, X.; and Cohen, W. W. 2023. Program of Thoughts Prompting: Disentangling Computation from Reasoning for Numerical Reasoning Tasks. *Transactions on Machine Learning Research*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. arXiv:2110.14168.
- Dou, S.; Zhou, E.; Liu, Y.; Gao, S.; Zhao, J.; Shen, W.; Zhou, Y.; Xi, Z.; Wang, X.; Fan, X.; Pu, S.; Zhu, J.; Zheng, R.; Gui, T.; Zhang, Q.; and Huang, X. 2024. LoRAMoE: Alleviate World Knowledge Forgetting in Large Language Models via MoE-Style Plugin. arXiv:2312.09979.
- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch transformers: scaling to trillion parameter models with simple and efficient sparsity. *J. Mach. Learn. Res.*, 23(1).
- Gao, L.; Madaan, A.; Zhou, S.; Alon, U.; Liu, P.; Yang, Y.; Callan, J.; and Neubig, G. 2023a. PAL: Program-aided Language Models. arXiv:2211.10435.
- Gao, S.; Shi, Z.; Zhu, M.; Fang, B.; Xin, X.; Ren, P.; Chen, Z.; Ma, J.; and Ren, Z. 2023b. Confucius: Iterative Tool Learning from Introspection Feedback by Easy-to-Difficult Curriculum. arXiv:2308.14034.
- Gou, Y.; Liu, Z.; Chen, K.; Hong, L.; Xu, H.; Li, A.; Yeung, D.-Y.; Kwok, J. T.; and Zhang, Y. 2024. Mixture of Cluster-tuning LoRA Experts for Vision-language Instruction Tuning. arXiv:2312.12379.
- Hu, E. J.; yelong shen; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; and Liu, T. 2023. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. arXiv:2311.05232.
- Jin, Z.; Cao, P.; Yuan, H.; Chen, Y.; Xu, J.; Li, H.; Jiang, X.; Liu, K.; and Zhao, J. 2024. Cutting Off the Head Ends the Conflict: A Mechanism for Interpreting and Mitigating Knowledge Conflicts in Language Models. arXiv:2402.18154.
- Joshi, M.; Choi, E.; Weld, D. S.; and Zettlemoyer, L. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. arXiv:1705.03551.
- Li, L.; Zhang, H.; Li, C.; You, H.; and Cui, W. 2023a. Evaluation on ChatGPT for Chinese Language Understanding. *Data Intelligence*, 5: 1–19.
- Li, M.; Zhao, Y.; Yu, B.; Song, F.; Li, H.; Yu, H.; Li, Z.; Huang, F.; and Li, Y. 2023b. API-Bank: A Comprehensive Benchmark for Tool-Augmented LLMs. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 3102–3116. Singapore: Association for Computational Linguistics.
- Li, T.; Li, S.; Xie, B.; Xiong, D.; and Yang, B. 2024. MoE-CT: A Novel Approach For Large Language Models Training With Resistance To Catastrophic Forgetting. arXiv:2407.00875.
- Lu, P.; Peng, B.; Cheng, H.; Galley, M.; Chang, K.-W.; Wu, Y. N.; Zhu, S.-C.; and Gao, J. 2023. Chameleon: Plug-and-Play Compositional Reasoning with Large Language Models. arXiv:2304.09842.
- Michel, P.; Levy, O.; and Neubig, G. 2019. Are Sixteen Heads Really Better than One? In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Nakano, R.; Hilton, J.; Balaji, S.; Wu, J.; Ouyang, L.; Kim, C.; Hesse, C.; Jain, S.; Kosaraju, V.; Saunders, W.; Jiang, X.; Cobbe, K.; Eloundou, T.; Krueger, G.; Button, K.; Knight, M.; Chess, B.; and Schulman, J. 2022. WebGPT: Browser-assisted question-answering with human feedback. arXiv:2112.09332.
- Qin, Y.; Hu, S.; Lin, Y.; Chen, W.; Ding, N.; Cui, G.; Zeng, Z.; Huang, Y.; Xiao, C.; Han, C.; Fung, Y. R.; Su, Y.; Wang, H.; Qian, C.; Tian, R.; Zhu, K.; Liang, S.; Shen, X.; Xu, B.; Zhang, Z.; Ye, Y.; Li, B.; Tang, Z.; Yi, J.; Zhu, Y.; Dai, Z.; Yan, L.; Cong, X.; Lu, Y.; Zhao, W.; Huang, Y.; Yan, J.; Han, X.; Sun, X.; Li, D.; Phang, J.; Yang, C.; Wu, T.; Ji, H.;

- Liu, Z.; and Sun, M. 2023. Tool Learning with Foundation Models. arXiv:2304.08354.
- Qin, Y.; Liang, S.; Ye, Y.; Zhu, K.; Yan, L.; Lu, Y.; Lin, Y.; Cong, X.; Tang, X.; Qian, B.; Zhao, S.; Hong, L.; Tian, R.; Xie, R.; Zhou, J.; Gerstein, M.; dahai li; Liu, Z.; and Sun, M. 2024. ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs. In *The Twelfth International Conference on Learning Representations*.
- Qu, C.; Dai, S.; Wei, X.; Cai, H.; Wang, S.; Yin, D.; Xu, J.; and Wen, J.-R. 2024. Tool Learning with Large Language Models: A Survey. arXiv:2405.17935.
- Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Zettlemoyer, L.; Cancedda, N.; and Scialom, T. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. arXiv:2302.04761.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations*.
- Shen, W.; Li, C.; Chen, H.; Yan, M.; Quan, X.; Chen, H.; Zhang, J.; and Huang, F. 2024. Small LLMs Are Weak Tool Learners: A Multi-LLM Agent. arXiv:2401.07324.
- Shen, Y.; Song, K.; Tan, X.; Li, D.; Lu, W.; and Zhuang, Y. 2023. HuggingGPT: Solving AI Tasks with ChatGPT and its Friends in Hugging Face. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Sun, T.; Zhang, X.; He, Z.; Li, P.; Cheng, Q.; Liu, X.; Yan, H.; Shao, Y.; Tang, Q.; Zhang, S.; Zhao, X.; Chen, K.; Zheng, Y.; Zhou, Z.; Li, R.; Zhan, J.; Zhou, Y.; Li, L.; Yang, X.; Wu, L.; Yin, Z.; Huang, X.; Jiang, Y.-G.; and Qiu, X. 2024. MOSS: An Open Conversational Large Language Model. *Machine Intelligence Research*, 21(5): 888–905.
- Tang, Q.; Deng, Z.; Lin, H.; Han, X.; Liang, Q.; Cao, B.; and Sun, L. 2023. ToolAlpaca: Generalized Tool Learning for Language Models with 3000 Simulated Cases. arXiv:2306.05301.
- Wang, B.; Fang, H.; Eisner, J.; Durme, B. V.; and Su, Y. 2024. LLMs in the Imaginarium: Tool Learning through Simulated Trial and Error. arXiv:2403.04746.
- Wu, C.; Yin, S.; Qi, W.; Wang, X.; Tang, Z.; and Duan, N. 2023. Visual ChatGPT: Talking, Drawing and Editing with Visual Foundation Models. arXiv:2303.04671.
- Wu, X.; Huang, S.; and Wei, F. 2024. Mixture of LoRA Experts. In *The Twelfth International Conference on Learning Representations*.
- Yang, S.; Ali, M. A.; Wang, C.-L.; Hu, L.; and Wang, D. 2024. MoRAL: MoE Augmented LoRA for LLMs’ Lifelong Learning. arXiv:2402.11260.
- Ye, J.; Li, G.; Gao, S.; Huang, C.; Wu, Y.; Li, S.; Fan, X.; Dou, S.; Zhang, Q.; Gui, T.; and Huang, X. 2024. ToolEyes: Fine-Grained Evaluation for Tool Learning Capabilities of Large Language Models in Real-world Scenarios. arXiv:2401.00741.
- Zhan, Q.; Liang, Z.; Ying, Z.; and Kang, D. 2024. InjecAgent: Benchmarking Indirect Prompt Injections in Tool-Integrated Large Language Model Agents. arXiv:2403.02691.
- Zhang, Q.; Zuo, S.; Liang, C.; Bukharin, A.; He, P.; Chen, W.; and Zhao, T. 2022. PLATON: Pruning Large Transformer Models with Upper Confidence Bound of Weight Importance. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 26809–26823. PMLR.
- Zhang, Z.; Zhao, J.; Zhang, Q.; Gui, T.; and Huang, X. 2024. Unveiling Linguistic Regions in Large Language Models. arXiv:2402.14700.
- Zhao, W.; Li, Z.; Wang, S.; Wang, Y.; Hu, Y.; Zhao, Y.; Wei, C.; and Qin, B. 2024. Both Matter: Enhancing the Emotional Intelligence of Large Language Models without Compromising the General Intelligence. arXiv:2402.10073.
- Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; Du, Y.; Yang, C.; Chen, Y.; Chen, Z.; Jiang, J.; Ren, R.; Li, Y.; Tang, X.; Liu, Z.; Liu, P.; Nie, J.-Y.; and Wen, J.-R. 2023. A Survey of Large Language Models. arXiv:2303.18223.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; Zhang, H.; Gonzalez, J. E.; and Stoica, I. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 46595–46623. Curran Associates, Inc.