

# Interweaving Memories of a Siamese Large Language Model

Xin Song<sup>1</sup>, Zhikai Xue<sup>2</sup>, Guoxiu He<sup>1\*</sup>, Jiawei Liu<sup>3</sup>, Wei Lu<sup>3</sup>

<sup>1</sup>School of Economics and Management, East China Normal University

<sup>2</sup>Department of Computer Science, Worcester Polytechnic Institute

<sup>3</sup>School of Information Management, Wuhan University

xsong2023@stu.ecnu.edu.cn, zxue1@wpi.edu, gxhe@fem.ecnu.edu.cn, {laujames2017, weilu}@whu.edu.cn

## Abstract

Parameter-efficient fine-tuning (PEFT) methods optimize large language models (LLMs) by modifying or introducing a small number of parameters to enhance alignment with downstream tasks. However, they can result in catastrophic forgetting, where LLMs prioritize new knowledge at the expense of comprehensive world knowledge. A promising approach to mitigate this issue is to recall prior memories based on the original knowledge. To this end, we propose a model-agnostic PEFT framework, **IMSM**, which Interweaves Memories of a Siamese Large Language Model. Specifically, our siamese LLM is equipped with an existing PEFT method. Given an incoming query, it generates two distinct memories based on the pre-trained and fine-tuned parameters. IMSM then incorporates an interweaving mechanism that regulates the contributions of both original and enhanced memories when generating the next token. This framework is theoretically applicable to all open-source LLMs and existing PEFT methods. We conduct extensive experiments across various benchmark datasets, evaluating the performance of popular open-source LLMs using the proposed IMSM, in comparison to both classical and leading PEFT methods. Our findings indicate that IMSM maintains comparable time and space efficiency to backbone PEFT methods while significantly improving performance and effectively mitigating catastrophic forgetting.

**Code** — <https://github.com/ECNU-Text-Computing/IMSM>

## Introduction

Large Language Models (LLMs) have emerged as a significant breakthrough in natural language processing (NLP) (Bai et al. 2023; Touvron et al. 2023; OpenAI 2023; Zeng et al. 2023). With extensive world knowledge, they have exhibited remarkable zero-shot abilities across various tasks, such as language understanding (Achiam et al. 2023), logical reasoning (Kojima et al. 2022), *etc.* Nevertheless, one of the challenges with LLMs is their susceptibility to generating incorrect or unfaithful outputs, especially in unfamiliar domains (Sun et al. 2024; Lin et al. 2024). In-context Learning (ICL) (Brown et al. 2020; Huang and He 2024) mitigates this issue by incorporating typical examples into prompts,

\*Corresponding author

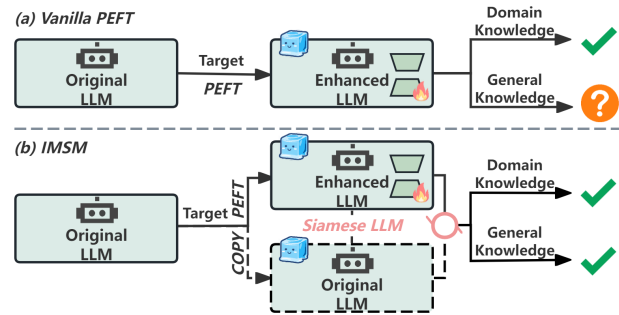


Figure 1: Comparison between vanilla PEFT (a) and our proposed IMSM (b). Vanilla PEFT methods employ an LLM only once. The parameter distribution shift will cause the LLM to forget general knowledge. Instead, IMSM incorporates a siamese LLM, which can be regarded as two LLMs, sharing identical structure and pre-trained parameters. One remains frozen while the other is fine-tuned using an existing PEFT method. By flexibly recalling the memory of the original LLM, IMSM can improve fine-tuning performance and alleviate catastrophic forgetting.

but does not inherently align LLMs with downstream objectives (Fu et al. 2023; Ding et al. 2023). Continued pre-training or supervised fine-tuning can better align parameters with tasks but is time-consuming and computationally intensive (VM et al. 2024; Han et al. 2024). To address this issue, parameter-efficient fine-tuning (PEFT) (Fu et al. 2023; Ding et al. 2023), as shown in Figure 1 (a), selectively fine-tunes a limited subset of parameters or adds new ones, while keeping the majority fixed.

However, refining the parameters of LLMs using target datasets enhances domain-specific knowledge but introduces biases related to that domain and task (Ren et al. 2024). This process can result in forgetting world knowledge and compromising general abilities, which is known as catastrophic forgetting (McCloskey and Cohen 1989; Kirkpatrick et al. 2017). Effective countermeasures typically encompass strategies for replaying pre-training data (Hayes et al. 2020; Kar et al. 2022) or constraints on updating parameters (Chen et al. 2020). Nevertheless, obtaining appropriate pre-training data for LLMs presents considerable challenges. Additionally, the vast range of tasks and the

sheer number of parameters complicate efforts to balance integrating new knowledge with preserving existing capabilities. This complexity introduces uncontrollable variables during the fine-tuning process.

Instead, this paper highlights the significance of utilizing the original knowledge of LLM to improve fine-tuning effectiveness and reduce the risk of catastrophic forgetting, as depicted in Figure 1 (b). Generally, the hidden states of the final layer of an LLM can be conceptualized as a memory. It encapsulates the encoded input sequence along with the preceding tokens of the generated response. When the model processes a query, the new memory shaped by the updated knowledge from fine-tuning may diverge from its original understanding. Therefore, for tasks beyond the fine-tuning target, the predicted probability distribution of the next token from this new memory runs the risk of deviating from the typically accurate distribution. In such cases, recalling the original memory before predicting based on the new memory at each step can help prevent deviations in token predictions. This process is akin to opening a channel between two parallel worlds with different experiences, facilitating the fusion of memories of diverse values.

To this end, we propose a straightforward yet highly effective PEFT framework called **IMSM**, which **I**nterweaves **M**emories of a **S**iamese **L**arge **L**anguage **M**odel. In particular, we employ a siamese LLM equipped with an existing PEFT method, like LoRA. Given a query, it generates two distinct final hidden states, *i.e.*, memories, based on the original pre-trained parameters and the parameters equipped with the PEFT module. We then propose a query-aware gate as the memory interweaving mechanism to facilitate token generation at each step. During the training phase, the existing PEFT module effectively retains new knowledge for downstream tasks. Thus, the gating mechanism is capable of dynamically balancing between the original memory and the updated memory based on varying queries. During the inference stage, IMSM relies on the given query features, allowing it to flexibly meet the demands of original knowledge for both fine-tuned tasks and other tasks. This mechanism is key to improving performance on fine-tuning tasks and reducing catastrophic forgetting on other tasks. Theoretically, our proposed IMSM is a model-agnostic framework that allows for seamless integration across various open-source LLMs and existing PEFT methods.

We conduct extensive experiments to evaluate the performance of the proposed IMSM in comparison to classical PEFT methods, including LoRA (Hu et al. 2021),  $(IA)^3$  (Liu et al. 2022), and AdaLoRA (Zhang et al. 2023). We also compare with the state-of-the-art PEFT methods, LoRAMoE (Dou et al. 2023) and DoRA (Liu et al. 2024). Moreover, we employ popular LLMs such as ChatGLM3 (Du et al. 2023), Qwen1.5 (Bai et al. 2023), Llama2 (Touvron et al. 2023), and Llama3 (Touvron et al. 2023) as backbone LLMs and fine-tune them with PEFT on four benchmark datasets, *e.g.*, MRPC (El-Said et al. 2015), CoLA (El-Said et al. 2015), ROPES (Lin et al. 2019), and GSM8K (Cobbe et al. 2021). To evaluate the extent of catastrophic forgetting, we analyze the LLMs’ performance on other datasets before and after fine-tuning on a target dataset. The

results demonstrate that IMSM achieves superior performance compared with vanilla PEFT methods. Notably, this improvement is attained without imposing too many additional trainable parameters. Furthermore, IMSM can effectively mitigate the issue of catastrophic forgetting.

Our main contributions are summarized as follows:

- To the best of our knowledge, this represents the first model-agnostic PEFT framework that employs a siamese LLM, effectively harnessing both its intrinsic world knowledge and the novel insights gained from PEFT.
- We propose a query-aware gate mechanism for interweaving memories, allowing flexible fusion of the last hidden states at each step.
- Extensive experiments validate the effectiveness of our proposed IMSM, and the results have confirmed its exceptional performance in striking a balance between plasticity and stability.

## Related Work

We review three lines of related work: parameter-efficient fine-tuning methods, strategies for mitigating catastrophic in LLMs, and approaches employing logits arithmetic.

### Parameter-efficient Fine-Tuning

Traditional approaches to fully fine-tuning LLMs with billions of parameters suffer from significant difficulties in terms of training time, computational expense, and practical efficiency. To overcome these, parameter-efficient fine-tuning (PEFT) techniques have been developed, which focus on adjusting only a subset of the model weights while maintaining the rest unchanged (Mao et al. 2021).

PEFT methods can be categorized into four main types: additive, selective, re-parameterized, and hybrid approaches (Han et al. 2024). Additive methods, like prefix-tuning (Li and Liang 2021) and  $(IA)^3$  (Liu et al. 2022), introduce additional modules into the LLM layers. Selective methods, like BitFit (Zaken, Goldberg, and Ravfogel 2022), focus on selecting a small subset of parameters for fine-tuning. Re-parameterized methods (Hu et al. 2021) use low-rank matrices to approximate the changing weights. Hybrid methods (Mao et al. 2022) are designed to combine different PEFT methods, offering a comprehensive solution.

Current PEFT methods enable domain alignment at the parameter level but can risk disrupting LLMs’ existing world knowledge and reasoning capabilities. Therefore, we propose IMSM, a model-agnostic framework for any open-source LLM and PEFT, which leverages the knowledge of the original LLM to enhance response accuracy by integrating the original memories generated from the same input.

### Catastrophic Forgetting in LLMs

LLMs face the challenge of catastrophic forgetting (McCloskey and Cohen 1989) during both pre-training and fine-tuning, where the model tends to forget previously acquired knowledge after being fine-tuned for downstream tasks. Catastrophic forgetting is the traditional challenge of continuous learning, leading to the development of various methods to address this issue in this scenario. Prominent

strategies include techniques such as constrained parameter updates, data replay, and parameter isolation (Ke and Liu 2022).

Recently, empirical investigations have indicated that even advanced fine-tuning techniques like LoRA are not immune to catastrophic forgetting (Luo et al. 2023). To address the issue during PEFT, acquiring extensive pre-training data can be an expensive process. Moreover, it is challenging to ensure that the existing knowledge learned by the LLM is not overwritten at the parameter level. Correspondingly, this work suggests a novel approach to address this issue by utilizing the siamese model to directly recall the knowledge from the original LLM. By interweaving distinct memories for the same input, the responses can be generated with enhanced flexibility.

### Logits Arithmetic

Due to the limitations of individual LLM, model collaboration emerges as a promising prospect. Previous studies (Liu et al. 2021) have highlighted the effectiveness of ensembling logits from multiple LLMs in controlled text generation or reducing hallucinations. Contrastive decoding (O’Brien and Lewis 2023) is proposed to directly use the discrepancy in logits between robust and weaker models. Prefix-Adaptive Decoding (Pei, Yang, and Klein 2023) and Context-aware decoding (Shi et al. 2023) enhance the model’s faithfulness by amplifying output probability differences across prompts.

Logits are derived from the final hidden states of LLMs, which serve as internal representations or memories. Prior research (Azaria and Mitchell 2023; Duan, Yang, and Tam 2024) has shown that these hidden states encapsulate knowledge or information pertinent to factual judgments. Thus, the final hidden states, focused on generating the next token, reflect the LLMs’ comprehension of the input information.

Taking this as a starting point, and different from previous methods, we place the siamese LLM within the PEFT framework, enabling a learnable collaboration. This approach not only enhances the performance of PEFT in downstream tasks, but also retains extensive world knowledge and reasoning capabilities of the original LLM.

### Methodology

In this section, we introduce the proposed IMSM, a novel model-agnostic PEFT framework. As illustrated in Figure 2, the siamese LLM is enhanced with a PEFT module. We conceptualize this siamese LLM as comprising two LLMs: one retains the original pre-trained knowledge, while the other is adapted with knowledge specific to the downstream task. The input tokens and the previously generated response tokens are processed by the siamese LLM to create distinct memories, specifically the final hidden states. We then propose a query-aware gate mechanism to interweave these memory representations. The next token is generated from the intertwined memory.

In this framework, the parameters associated with PEFT and the parameters related to the gate mechanism will undergo fine-tuning, while the remaining parameters will be frozen. Our objective is to preserve the model’s inherent

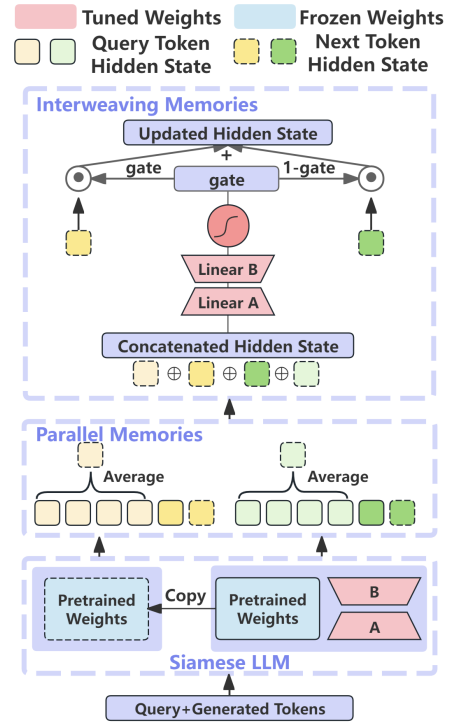


Figure 2: The overall architecture of IMSM, including a siamese LLM and an interweaving mechanism. Given the same input tokens, our siamese LLM produces memories with distinct values, which correspond to the two last hidden states. The generation of the next token relies on the updated memory through an interweaving mechanism. Trainable parameters are marked in red.

world knowledge and general reasoning capabilities while simultaneously improving its performance in both downstream and general tasks. IMSM has the potential to be applied to various open-source LLMs and PEFT methods.

### Siamese Large Language Model

To facilitate description and enhance intuitive understanding, our Siamese LLM can be conceptualized as a LLM with dual types of parameters. When integrating PEFT methods into the siamese LLM, the model first retains a copy of the original knowledge. It then undergoes fine-tuning through the PEFT module using downstream data to better align with the new target objective. Taking LoRA as an exemplar PEFT method, we update the parameters  $\mathbf{W}$  by adding an adapter  $\Delta\mathbf{W}$  to the original parameters  $\mathbf{W}_o \in \mathbb{R}^{d \times k}$ :

$$\mathbf{W} = \mathbf{W}_o + \Delta\mathbf{W} = \mathbf{W}_o + \mathbf{B}\mathbf{A} \quad (1)$$

where  $\mathbf{W}_o$  remains frozen, while  $\mathbf{B} \in \mathbb{R}^{d \times r'}$  and  $\mathbf{A} \in \mathbb{R}^{r' \times k}$  denote trainable low-rank matrices. And  $r'$  is significantly smaller than  $d$  and  $k$ . This configuration results in a substantial reduction in the number of trainable parameters.  $\mathbf{W}_o$  and  $\mathbf{W}$  represent the parameters of the frozen LLM (*i.e.*, original LLM) and tuned LLM, respectively.

## Memory Interweaving Mechanism

Throughout both the fine-tuning and inference processes, our siamese LLM generates distinct hidden states and corresponding logits. The last hidden states can serve as feature representations, encapsulating the LLM’s memory of the previous input tokens. We shall employ these two representations as vehicles to interweave memories derived from various experiences, and propose a gate-based interweaving mechanism for memories.

Let  $\mathcal{M}$  represent the original LLM with parameters  $\mathbf{W}_o$ , while  $\mathcal{M}'$  denotes the fine-tuned LLM with parameters  $\mathbf{W}$  including partially tuned parameters  $\Delta\mathbf{W}$ .

At each time step  $t$ , we can access the hidden states derived from  $\mathcal{M}$  and  $\mathcal{M}'$ , as well as the hidden states prior to  $t$ . First, we compute the average of the query’s last layer hidden states in the Siamese LLM to obtain its dense representations, reflecting the model’s dual understanding of the input. Next, the hidden states for query understanding and next token prediction are concatenated along the feature dimension. Finally, we construct a linear layer with low-rank weight matrices at the top of the model to generate a gate, which determines the contribution of each LLM in the next step of logits generation process:

$$\bar{\mathbf{h}}_{\mathcal{M}}^q = \frac{1}{T_{in}} \sum_{t=1}^{T_{in}} \mathbf{h}_{\mathcal{M}}^t \quad (2)$$

$$\bar{\mathbf{h}}_{\mathcal{M}'}^q = \frac{1}{T_{in}} \sum_{t=1}^{T_{in}} \mathbf{h}_{\mathcal{M}'}^t \quad (3)$$

$$\begin{aligned} gate &= \text{sigmoid} \left( f \left( \bar{\mathbf{h}}_{\mathcal{M}}^q \oplus \mathbf{h}_{\mathcal{M}}^t \oplus \mathbf{h}_{\mathcal{M}'}^t \oplus \bar{\mathbf{h}}_{\mathcal{M}'}^q \right) \right) \\ &= \text{sigmoid} \left( \left( \bar{\mathbf{h}}_{\mathcal{M}}^q \oplus \mathbf{h}_{\mathcal{M}}^t \oplus \mathbf{h}_{\mathcal{M}'}^t \oplus \bar{\mathbf{h}}_{\mathcal{M}'}^q \right) \cdot \mathbf{W}_A \cdot \mathbf{W}_B \right) \end{aligned} \quad (4)$$

where  $\mathbf{h}_{\mathcal{M}}^t$  and  $\mathbf{h}_{\mathcal{M}'}^t$  represent the last layer hidden states of the original LLM  $\mathcal{M}$  and the tuned LLM  $\mathcal{M}'$ , respectively, for the  $t$ -th token in the sequence.  $T_{in}$  represents the input sequence length.  $\bar{\mathbf{h}}_{\mathcal{M}}^q$  and  $\bar{\mathbf{h}}_{\mathcal{M}'}^q$  represent the average of the query’s last layer hidden states.  $\oplus$  denotes vector concatenation and  $f$  represents the linear layer with low rank. Here,  $\mathbf{h}_{\mathcal{M}}^t, \mathbf{h}_{\mathcal{M}'}^t \in \mathbb{R}^{1 \times d}$ , where  $d$  represents the dimension of the hidden state.  $\mathbf{W}_A \in \mathbb{R}^{4d \times r}$  and  $\mathbf{W}_B \in \mathbb{R}^{r \times d}$ , where  $r$  is a hyper-parameter. Notably,  $r$  is significant smaller than  $d$ , ensuring a limited number of parameters. Furthermore, the gate is obtained through sigmoid activation.

The updated last hidden state  $\mathbf{h}_{\mathcal{N}}^t$  and the logits can be calculated by:

$$\mathbf{h}_{\mathcal{N}}^t = gate \circ \mathbf{h}_{\mathcal{M}}^t + (1 - gate) \circ \mathbf{h}_{\mathcal{M}'}^t \quad (5)$$

$$logits = \mathbf{h}_{\mathcal{N}}^t \cdot \mathbf{W}_{out} \quad (6)$$

where  $\circ$  denotes element-wise multiplication and  $\mathbf{W}_{out}$  denotes the parameters of the linear layer of the siamese LLM that are utilized for generating the logits of the next token corresponding to the vocabulary size.

In the memory interweaving process, incorporating query information into the gating mechanism enables the model to

better retain original knowledge and reasoning abilities. Unlike additive methods that lead to parameter sharing, concatenating hidden states facilitates finer-grained feature fusion. Generally, both averaging and maximization are valid strategies for handling the parallel memories of past query tokens. However, averaging to synthesize information offers a more comprehensive representation of the query.

## Training and Inference

The probability distribution of the next token can be calculated as follows:

$$p(y_t|x, y_{<t}) = \text{softmax}(logits) \quad (7)$$

During the fine-tuning phase, we utilize cross-entropy as our loss function:

$$\mathcal{L}(\Delta\mathbf{W}, \mathbf{W}_A, \mathbf{W}_B) = - \sum_{t=1}^{T_{out}} \sum_{y \in V} y_t \log p(y_t|x, y_{<t}) \quad (8)$$

where  $T_{out}$  is the length of the output sequence. Finally, we optimize the parameters  $\Delta\mathbf{W}$ ,  $\mathbf{W}_A$ , and  $\mathbf{W}_B$  using the optimization algorithm like AdamW (Loshchilov and Hutter 2017). Consequently, the gate mechanism is capable of learning to execute flexible memory interweaving for diverse queries. This allows our proposed IMSM to mitigate the interference caused by extraneous new knowledge for tasks that fall outside the fine-tuning objective.

During the inference phase, we utilize a greedy search strategy to generate the output sequence. We select the token with the highest probability, informed by the updated memory, as the next generated token, add it to the input sequence, and then repeat this process until the complete sequence is generated.

## Experiments

### Datasets

We conduct experiments on four datasets: MRPC (El-Said et al. 2015), CoLA (El-Said et al. 2015), ROPES (Lin et al. 2019), and GSM8K (Cobbe et al. 2021), to evaluate the alignment capability of our IMSM. Following previous studies (Schick et al. 2024; Asai et al. 2023), we also employ MRPC (El-Said et al. 2015), WebQ (Bordes, Chopra, and Weston 2014), FreebaseQA (Jiang, Wu, and Jiang 2019), and MultiRC (Khashabi et al. 2018), to assess the abilities to retain general knowledge of LLM.

Distinct metrics are employed for different tasks: accuracy (Acc.) and F1 score for MRPC, Matthews correlation coefficient (MCC) for CoLA, Exact Match (EM) and F1 token overlap for ROPES, accuracy for GSM8K and MultiRC. Following Schick et al. (2024) and Asai et al. (2023), we evaluate the performance on WebQ and Freebase by examining whether the generated response contains the golden answer, rather than strictly requiring exact matching.

### Backbone LLMs

We employ four mainstream open-source LLMs, including ChatGLM3-6B (Zeng et al. 2023), Qwen1.5-4B (Bai et al. 2023), Llama2-7B (Touvron et al. 2023), and Llama3-8B (Touvron et al. 2023), as backbones of our siamese LLM. We also directly prompt these LLMs as baseline models.

Backbone	Method	Params	MRPC		CoLA	ROPES		GSM8K	Average
			Acc.	F1	Mcc.	EM	F1	Acc.	
ChatGLM3	Original	-	74.72	82.88	43.99	58.47	22.67	45.79	54.75
	LoRA	3.899M	86.67	90.10	63.56	74.38	79.41	44.50	73.10
	IMSM	4.063M	<b>87.13*</b>	<b>90.49*</b>	<b>65.21*</b>	<b>79.15*</b>	<b>81.28*</b>	<b>48.45*</b>	<b>75.29*</b>
	$(IA)^3$	0.513M	86.96	90.32	62.12	72.22	76.36	50.04	73.00
	IMSM	0.676M	<b>88.70*</b>	<b>91.66*</b>	<b>64.53*</b>	<b>73.34*</b>	<b>77.56*</b>	<b>52.46*</b>	<b>74.71*</b>
AdaLoRA	2.925M	87.02	90.35	62.48	75.14	78.64	49.36	73.83	
	IMSM	3.089M	<b>87.77*</b>	<b>90.81*</b>	<b>65.57*</b>	<b>78.50*</b>	<b>82.23*</b>	<b>50.80*</b>	<b>75.95*</b>
Qwen1.5	Original	-	74.20	81.04	46.50	56.58	38.58	18.35	52.54
	LoRA	3.277M	86.41	89.79	65.23	63.63	69.56	50.27	70.82
	IMSM	3.379M	<b>87.25*</b>	<b>90.48*</b>	<b>65.80*</b>	<b>65.64*</b>	<b>72.71*</b>	<b>52.69*</b>	<b>72.43*</b>
	$(IA)^3$	0.205M	86.84	90.16	64.96	56.46	64.35	51.33	69.02
	IMSM	0.307M	<b>87.13*</b>	<b>90.73*</b>	<b>69.13*</b>	<b>60.25*</b>	<b>67.62*</b>	<b>53.15*</b>	<b>71.34*</b>
AdaLoRA	2.458M	86.73	90.15	65.63	64.81	68.62	<b>52.01</b>	71.33	
	IMSM	2.560M	<b>88.41*</b>	<b>91.32*</b>	<b>66.24*</b>	<b>67.18*</b>	<b>72.36*</b>	51.86	<b>72.90*</b>
Llama3	Original	-	68.17	79.83	37.96	49.23	13.95	3.23	42.06
	LoRA	6.816M	<b>89.10</b>	<b>91.88</b>	71.65	87.26	89.09	59.59	81.43
	IMSM	6.980M	89.04	91.84	<b>72.05*</b>	<b>87.85*</b>	<b>90.42*</b>	<b>61.26*</b>	<b>82.08*</b>
	$(IA)^3$	0.524M	87.71	90.78	68.16	78.55	82.86	<b>62.62</b>	78.45
	IMSM	0.688M	<b>88.87*</b>	<b>91.77*</b>	<b>70.78*</b>	<b>81.58*</b>	<b>85.95*</b>	61.33	<b>80.05*</b>
AdaLoRA	5.113M	88.87	91.72	68.77	87.21	88.98	55.42	80.16	
	IMSM	5.276M	<b>89.22*</b>	<b>91.95*</b>	<b>71.55*</b>	<b>87.80*</b>	<b>89.05*</b>	<b>59.74*</b>	<b>81.55*</b>

Table 1: The overall comparison across four downstream tasks. The best results achieved using IMSM and its corresponding vanilla fine-tuning methods are highlighted in boldface. The improvements achieved by IMSM over all baselines are statistically significant, as measured by student’s t-test with a significance level of  $p < 0.05$ .

### Baseline PEFT Methods

We utilize classical and state-of-the-art PEFT methods, like LoRA (Hu et al. 2021),  $(IA)^3$  (Liu et al. 2022), AdaLoRA (Zhang et al. 2023), and DoRA (Liu et al. 2024), to perform fine-tuning on the siamese LLM of our IMSM. In addition, we compare LoRAMoE (Dou et al. 2023), a plugin-based mixture of experts model designed for preventing catastrophic forgetting.

### Implementation Details

We utilize HuggingFace Transformers (Wolf et al. 2019) and PEFT (Mangrulkar et al. 2022) to perform our experiments. The fine-tuning procedure is executed on 8 NVIDIA A800 GPUs under a Linux system.

For LoRA, AdaLoRA, and DoRA, we employ AdamW as the optimizer with learning rates of  $3 \times 10^{-4}$ ,  $2 \times 10^{-3}$ , and  $1 \times 10^{-4}$ , respectively, and a batch size of 16. The rank and alpha for LoRA are set to 16. For DoRA, we follow the authors’ recommendations, setting the rank and alpha to 16 and 32. For  $(IA)^3$ , we use Adafactor with a learning rate of  $3 \times 10^{-3}$  and a batch size of 8. All methods are trained for 3 epochs. For LoRAMoE, we use the original paper’s configuration. For a fair comparison, we set the configurations of

the tuned target modules of IMSM to be exactly the same as vanilla PEFT. The gate rank  $r$  of IMSM is set to 8.

### Performance Comparison

Table 1 presents the fine-tuning results of IMSM and baselines across four datasets. Generally, the fine-tuned LLMs outperform the solely prompted LLMs. This highlights the significance of parameter tuning compared with ICL for aligning LLMs to unfamiliar tasks. Moreover, our proposed method, IMSM, significantly surpasses all vanilla PEFT methods across all three LLMs.

The effect of parameter fine-tuning is influenced by multiple factors, including the quality and capabilities of the backbone LLM, and the adaptability of the employed PEFT method. As one of the most advanced open-source LLMs, the Llama3-based models achieve near-optimal performance. Furthermore, AdaLoRA dynamically adjusts the rank of low-rank matrices, highlighting its efficacy in aligning with downstream tasks.

IMSM, as a simple yet powerful model-agnostic PEFT framework, excels at integrating the stability of the original LLM with the adaptability of the tuned LLM. Additionally, it incorporates the gate mechanism that allows for selective control over whether the understanding is derived from the

Dataset	ROPES	WebQ	MultiRC	MRPC	Freebase	Avg.
ChatGLM3	58.47	25.49	71.74	74.72	36.81	53.45
LoRA	74.38	20.47	66.67	70.32	<b>38.61</b>	54.09
IMSM	<b>79.15</b>	<b>21.11</b>	<b>71.47</b>	<b>71.83</b>	38.01	<b>56.31</b>
$(IA)^3$	72.22	20.96	70.83	72.58	37.34	54.79
IMSM	<b>73.34</b>	<b>21.36</b>	<b>73.72</b>	<b>74.20</b>	<b>37.74</b>	<b>56.07</b>
AdaLoRA	75.14	19.98	67.63	70.20	35.99	53.79
IMSM	<b>78.50</b>	<b>22.88</b>	<b>73.72</b>	<b>74.96</b>	<b>38.41</b>	<b>57.69</b>

Table 2: Catastrophic forgetting validation between vanilla PEFT and IMSM using ChatGLM3-6B, fine-tuned on ROPES and evaluated on general knowledge benchmarks.

fixed or adjusted knowledge. Therefore, it leads to consistent improvement across various downstream tasks.

Dataset	GSM8K	WebQ	MultiRC	MRPC	Freebase	Avg.
ChatGLM3	45.79	25.49	71.74	74.72	36.81	50.91
DoRA	45.34	<b>26.18</b>	71.47	74.55	33.26	50.16
IMSM	<b>48.75</b>	25.49	<b>75.00</b>	<b>75.19</b>	<b>38.89</b>	<b>52.66</b>
Qwen1.5	18.35	32.33	66.99	74.20	47.12	47.80
DoRA	51.33	27.51	66.03	73.91	41.19	51.99
IMSM	<b>52.84</b>	<b>28.94</b>	<b>68.27</b>	<b>74.26</b>	<b>43.52</b>	<b>53.57</b>
Llama3.1	33.06	41.78	54.17	42.43	76.35	49.56
DoRA	72.71	32.78	58.65	<b>55.55</b>	65.92	57.12
IMSM	<b>73.77</b>	<b>40.31</b>	<b>66.90</b>	47.94	<b>75.45</b>	<b>60.87</b>

Table 3: Evaluation of catastrophic forgetting between DoRA and its corresponding IMSM across different LLMs.

### Catastrophic Forgetting Evaluation

Fine-tuning inevitably introduces a degree of forgetting. We fine-tune LLMs on one dataset and subsequently evaluate their performance on other general benchmarks.

Table 2 shows the results of fine-tuning ChatGLM on ROPES. For the three vanilla PEFT methods, they all prioritize enhancing adaptability while neglecting stability. After fine-tuning ChatGLM on ROPES with LoRA,  $(IA)^3$ , and AdaLoRA, the accuracy on MultiRC drops from 71.74 to 66.67, 70.83, and 67.63, respectively. IMSM can effectively mitigate performance degradation. For example, IMSM improves overall performance by 4.10% compared with LoRA. Notably, the interweaving strategy of IMSM can even revive precise memories within the LLM, enhancing performance on non-target datasets.

We also compare IMSM with LoRAMoE and DoRA, which are state-of-the-art solutions for catastrophic forgetting and PEFT, respectively. For LoRAMoE, we follow the original setup using Llama2 as the backbone. Due to crashes when using DoRA for fine-tuning Llama3, we opt for Llama3.1-instruct instead. As illustrated in Figure 3, IMSM reaches or surpasses LoRAMoE on general benchmarks, offering an average boost of 1.85% and 8.77% in




Figure 3: Evaluation of catastrophic forgetting using Llama2-7B, fine-tuned on ROPES and GSM8K, and evaluated on general knowledge benchmarks.

overall performance. As shown in Table 3, whether on the target dataset GSM8K or general benchmarks, IMSM’s performance consistently outperforms the standalone DoRA.

The last hidden state serves as the memory for the LLMs’ understanding of the input sequence. By considering the original and updated memories, we ensure that the siamese LLM retains a balance between its prior knowledge and its adaptation to new data.

### Ablation Study

We conduct ablation tests to investigate how the query-aware gate, which serves as the memory interweaving mechanism, improves performance and assess its necessity. The results are shown in Table 4 and Figure 4. In the “w/o query” setting, the siamese LLM’s understanding of the query is excluded from the gate construction. In the “w/o gate” setting, the gate construction is omitted, and the two final hidden states of the siamese LLM are directly added in a 1:1 ratio.

As shown in Table 4, the query-aware gate mechanism brings the most significant improvement to downstream task performance in almost all settings. Figure 4 illustrates that directly combining memories and incorporating the original parameter memory bring performance gains on non-target datasets compared with vanilla PEFT. This highlights the effectiveness of recalling original knowledge through the final layer hidden state. With the introduction of a simple gate, the approach can dynamically balance the previous and the refreshed knowledge. The integration of the siamese LLM’s query feature further alleviates catastrophic forgetting.

### Space and Time Complexity

Compared to the PEFT method alone, IMSM maintains comparable space and time complexity. While the gate mechanism introduces additional trainable parameters, their impact is minimal compared with the original PEFT. We report the count of trainable parameters in Table 1, with a gate rank of 8 used across all IMSM experiments. During inference, the time required to interweave the two memories remains constant for each step, with the two memory




Figure 4: Ablation test of catastrophic forgetting validation. ChatGLM3-6B is employed as the backbone LLM, fine-tuned on ROPES, and evaluated on general knowledge benchmarks.

Dataset	MRPC	CoLA	ROPES	Avg.
LoRA	86.67	90.10	63.56	79.41
IMSM	87.13	90.49	<b>65.21</b>	<b>79.15</b>
w/o query	<b>87.62</b>	<b>90.81</b>	64.92	77.43
w/o gate	87.01	90.31	63.74	76.54
(IA) <sup>3</sup>	86.96	90.32	62.12	72.22
IMSM	<b>88.70</b>	<b>91.66</b>	<b>64.53</b>	<b>73.34</b>
w/o query	87.77	91.04	63.49	72.93
w/o gate	87.04	90.48	62.94	71.86
AdaLoRA	87.02	90.35	62.48	75.14
IMSM	<b>87.77</b>	<b>90.81</b>	<b>65.57</b>	<b>78.50</b>
w/o query	87.59	90.68	65.19	75.41
w/o gate	87.30	90.59	64.94	76.60

Table 4: The ablation test is performed on ChatGLM3-6B to compare its performance across various downstream tasks.

streams being generated in parallel theoretically. Compared to vanilla PEFT, IMSM introduces only a lightweight gating mechanism overhead, maintaining equivalent time complexity. The inference speed, measured in tokens per second, is presented in Table 5. The extra memory usage is a necessary and acceptable trade-off to address catastrophic forgetting.

Method	ChatGLM3	Qwen1.5	Llama3
LoRA	31.30	22.86	22.98
IMSM	29.05	21.76	21.18
(IA) <sup>3</sup>	33.86	23.51	22.81
IMSM	30.40	22.10	20.51
AdaLoRA	30.85	21.68	21.49
IMSM	27.91	20.03	18.71

Table 5: The comparisons of inference speed (tokens/sec) between vanilla PEFT methods and IMSM.

### Hyper-parameter Analysis

This section focuses on the sensitivity of the rank  $r$ , a hyper-parameter in the gate mechanism. Keeping other settings constant, we vary the rank among 4, 8, and 16. Please re-

fer to Table 6 for a comprehensive representation. While increasing the rank adds more trainable parameters, it does not necessarily improve performance on the target dataset. The optimal rank varies depending on the specific PEFT used. For IMSM based on LoRA and (IA)<sup>3</sup>, a rank of 8 is often optimal, while for AdaLoRA, 16 is usually the most beneficial choice. However, as the rank increases, the gate mechanism may better recall original memories, potentially reducing forgetting on non-target datasets more effectively.

IMSM Method	Rank	Params	GSM8K	WebQ	Freebase
LoRA	4	6.898M	59.44	40.55	73.15
	8	6.980M	<b>61.26</b>	<b>40.90</b>	73.72
	16	7.143M	59.21	39.86	<b>75.23</b>
(IA) <sup>3</sup>	4	0.606M	58.98	37.80	74.60
	8	0.688M	<b>61.33</b>	38.58	74.32
	16	0.852M	61.26	<b>38.98</b>	<b>75.90</b>
AdaLoRA	4	5.194M	59.67	41.63	73.97
	8	5.276M	59.74	40.94	74.28
	16	5.440M	<b>60.50</b>	<b>42.72</b>	<b>75.45</b>

Table 6: Efficiency on gate rank  $r$  is fine-tuned on GSM8K using Llama3-8B, with effects on WebQ and Freebase.

### Conclusion and Future Work

This paper presents a novel approach to fine-tuning LLMs, aiming to strike a delicate balance between plasticity and stability. Our proposed IMSM constitutes a model-agnostic framework applicable to any open-source LLMs, in conjunction with existing PEFT methods. Particularly, IMSM enables the interweaving of memories derived from the siamese LLM, facilitating collaboration between frozen knowledge and tuned knowledge. Extensive experiments substantiate that the interweaving mechanism significantly improves alignment performance on downstream tasks and alleviates catastrophic forgetting. Additional hyper-parameter experiments further confirm the robustness of the proposed model.

In future research, we aim to explore more intricate memory fusion mechanisms within multiple Transformer layers. Besides, we intend to evaluate the performance of our model on more challenging scenarios of catastrophic forgetting.

## Acknowledgments

This work is supported by the National Natural Science Foundation of China (72204087, 72404212, 72234005), the Shanghai Planning Office of Philosophy and Social Science Youth Project (2022ETQ001), the “Chen Guang” project supported by Shanghai Municipal Education Commission and Shanghai Education Development Foundation (23CGA28), the Shanghai Pujiang Program (23PJC030), the Fundamental Research Funds for the Central Universities, China, and the 2024 Innovation Evaluation Open Fund, Fudan University (CXPJ2024006). We also appreciate the constructive comments from the anonymous reviewers.

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Asai, A.; Wu, Z.; Wang, Y.; Sil, A.; and Hajishirzi, H. 2023. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. *arXiv preprint arXiv:2310.11511*.
- Azaria, A.; and Mitchell, T. 2023. The Internal State of an LLM Knows When It’s Lying. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 967–976.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Bordes, A.; Chopra, S.; and Weston, J. 2014. Question answering with subgraph embeddings. *arXiv preprint arXiv:1406.3676*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chen, S.; Hou, Y.; Cui, Y.; Che, W.; Liu, T.; and Yu, X. 2020. Recall and learn: Fine-tuning deep pretrained language models with less forgetting. *arXiv preprint arXiv:2004.12651*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Ding, N.; Qin, Y.; Yang, G.; Wei, F.; Yang, Z.; Su, Y.; Hu, S.; Chen, Y.; Chan, C.-M.; Chen, W.; et al. 2023. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5(3): 220–235.
- 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. 2023. LoRAMoE: Revolutionizing Mixture of Experts for Maintaining World Knowledge in Language Model Alignment. *arXiv:2312.09979*.
- Du, Z.; Qian, Y.; Liu, X.; Ding, M.; Qiu, J.; Yang, Z.; and Tang, J. 2021. Glm: General language model pre-training with autoregressive blank infilling. *arXiv preprint arXiv:2103.10360*.
- Duan, H.; Yang, Y.; and Tam, K. Y. 2024. Do LLMs Know about Hallucination? An Empirical Investigation of LLM’s Hidden States. *arXiv preprint arXiv:2402.09733*.
- El-Said, A. M.; Eldesoky, A. I.; Arafat, H. A.; et al. 2015. Exploiting semantic annotations and-learning for constructing an efficient hierarchy/graph texts organization. *The Scientific World Journal*, 2015.
- Fu, Z.; Yang, H.; So, A. M.-C.; Lam, W.; Bing, L.; and Collier, N. 2023. On the effectiveness of parameter-efficient fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 12799–12807.
- Han, Z.; Gao, C.; Liu, J.; Zhang, S. Q.; et al. 2024. Parameter-efficient fine-tuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*.
- Hayes, T. L.; Kafle, K.; Shrestha, R.; Acharya, M.; and Kanan, C. 2020. Remind your neural network to prevent catastrophic forgetting. In *European Conference on Computer Vision*, 466–483. Springer.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Huang, C.; and He, G. 2024. Text clustering as classification with llms. *arXiv preprint arXiv:2410.00927*.
- Jiang, K.; Wu, D.; and Jiang, H. 2019. FreebaseQA: A new factoid QA data set matching trivia-style question-answer pairs with Freebase. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 318–323.
- Kar, S.; Castellucci, G.; Filice, S.; Malmasi, S.; and Rokhlenko, O. 2022. Preventing catastrophic forgetting in continual learning of new natural language tasks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3137–3145.
- Ke, Z.; and Liu, B. 2022. Continual learning of natural language processing tasks: A survey. *arXiv preprint arXiv:2211.12701*.
- Khashabi, D.; Chaturvedi, S.; Roth, M.; Upadhyay, S.; and Roth, D. 2018. Looking beyond the surface: A challenge set for reading comprehension over multiple sentences. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 252–262.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213.
- Li, X. L.; and Liang, P. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.

- Lin, C.; Jiayu, R.; He, G.; Jiang, Z.; Yu, H.; and Zhu, X. 2024. Recurrent Alignment with Hard Attention for Hierarchical Text Rating. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 18643–18657.
- Lin, K.; Tafjord, O.; Clark, P.; and Gardner, M. 2019. Reasoning over paragraph effects in situations. *arXiv preprint arXiv:1908.05852*.
- Liu, A.; Sap, M.; Lu, X.; Swayamdipta, S.; Bhagavatula, C.; Smith, N. A.; and Choi, Y. 2021. DExperts: Decoding-time controlled text generation with experts and anti-experts. *arXiv preprint arXiv:2105.03023*.
- Liu, H.; Tam, D.; Muqeeth, M.; Mohta, J.; Huang, T.; Bansal, M.; and Raffel, C. A. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35: 1950–1965.
- Liu, S.-y.; Wang, C.-Y.; Yin, H.; Molchanov, P.; Wang, Y.-C. F.; Cheng, K.-T.; and Chen, M.-H. 2024. DoRA: Weight-Decomposed Low-Rank Adaptation. In *Forty-first International Conference on Machine Learning*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Luo, Y.; Yang, Z.; Meng, F.; Li, Y.; Zhou, J.; and Zhang, Y. 2023. An empirical study of catastrophic forgetting in large language models during continual fine-tuning. *arXiv preprint arXiv:2308.08747*.
- Mangrulkar, S.; Gugger, S.; Debut, L.; Belkada, Y.; Paul, S.; and Bossan, B. 2022. PEFT: State-of-the-art Parameter-Efficient Fine-Tuning methods. <https://github.com/huggingface/peft>.
- Mao, Y.; Mathias, L.; Hou, R.; Almahairi, A.; Ma, H.; Han, J.; Yih, S.; and Khabsa, M. 2022. UniPELT: A Unified Framework for Parameter-Efficient Language Model Tuning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6253–6264.
- Mao, Y.; Mathias, L.; Hou, R.; Almahairi, A.; Ma, H.; Han, J.; Yih, W.-t.; and Khabsa, M. 2021. Unipelt: A unified framework for parameter-efficient language model tuning. *arXiv preprint arXiv:2110.07577*.
- McCloskey, M.; and Cohen, N. J. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, 109–165. Elsevier.
- O’Brien, S.; and Lewis, M. 2023. Contrastive decoding improves reasoning in large language models. *arXiv preprint arXiv:2309.09117*.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*.
- Pei, J.; Yang, K.; and Klein, D. 2023. PREADD: prefix-adaptive decoding for controlled text generation. *arXiv preprint arXiv:2307.03214*.
- Ren, W.; Li, X.; Wang, L.; Zhao, T.; and Qin, W. 2024. Analyzing and Reducing Catastrophic Forgetting in Parameter Efficient Tuning. *arXiv preprint arXiv:2402.18865*.
- Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Hambro, E.; Zettlemoyer, L.; Cancedda, N.; and Scialom, T. 2024. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36.
- Shi, W.; Han, X.; Lewis, M.; Tsvetkov, Y.; Zettlemoyer, L.; and Yih, S. W.-t. 2023. Trusting your evidence: Hallucinate less with context-aware decoding. *arXiv preprint arXiv:2305.14739*.
- Sun, L.; Huang, Y.; Wang, H.; Wu, S.; Zhang, Q.; Gao, C.; Huang, Y.; Lyu, W.; Zhang, Y.; Li, X.; et al. 2024. Trustllm: Trustworthiness in large language models. *arXiv preprint arXiv:2401.05561*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- VM, K.; Warriar, H.; Gupta, Y.; et al. 2024. Fine Tuning LLM for Enterprise: Practical Guidelines and Recommendations. *arXiv preprint arXiv:2404.10779*.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Zaken, E. B.; Goldberg, Y.; and Ravfogel, S. 2022. BitFit: Simple Parameter-efficient Fine-tuning for Transformer-based Masked Language-models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 1–9.
- Zeng, A.; Liu, X.; Du, Z.; Wang, Z.; Lai, H.; Ding, M.; Yang, Z.; Xu, Y.; Zheng, W.; Xia, X.; Tam, W. L.; Ma, Z.; Xue, Y.; Zhai, J.; Chen, W.; Liu, Z.; Zhang, P.; Dong, Y.; and Tang, J. 2023. GLM-130B: An Open Bilingual Pre-trained Model. In *The Eleventh International Conference on Learning Representations*.
- Zhang, Q.; Chen, M.; Bukharin, A.; He, P.; Cheng, Y.; Chen, W.; and Zhao, T. 2023. Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512*.