

Forgetting before Learning: Utilizing Parametric Arithmetic for Knowledge Updating in Large Language Models

Shiwen Ni^{♣*}, Dingwei Chen^{♣*}, Chengming Li^{◇†}, Xiping Hu[◇], Ruifeng Xu[♡], Min Yang^{♣†}

[♣] Sun Yat-Sen University

[♣] Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences

[◇] Shenzhen MSU-BIT University [♡] Harbin Institute of Technology (Shenzhen)

chendw26@mail2.sysu.edu.cn, {sw.ni, min.yang}@siat.ac.cn

{licm, huxp}@smbu.edu.cn, xurufeng@hit.edu.cn

Abstract

Recent advancements in Large Language Models (LLMs) have showcased their remarkable capabilities in text understanding and generation. However, even stronger LLMs are susceptible to acquiring erroneous or obsolete information from the training corpus. Direct secondary fine-tuning with data containing new knowledge may be ineffective in updating knowledge due to the conflict between old and new knowledge. In this paper, we propose a new paradigm for fine-tuning called **F-Learning** (Forgetting before Learning), which employs parametric arithmetic to facilitate the forgetting of old knowledge and learning of new knowledge. Experimental results on two publicly available datasets demonstrate that our proposed F-Learning can obviously improve the knowledge updating performance of both full fine-tuning and LoRA fine-tuning, simultaneously outperforming the existing baselines in most cases. Moreover, we have also discovered that forgetting old knowledge by subtracting the parameters of LoRA can yield a similar effect to subtracting the parameters of full fine-tuning, and occasionally even surpass it significantly.

1 Introduction

Large Language Models (LLMs) possess an extraordinary ability to understand and generate natural language (Brown et al., 2020; Raffel et al., 2020; Ouyang et al., 2022). Although LLMs are very capable of learning, they are not immune to the acquisition of incorrect knowledge in the corpus. Moreover, much of the knowledge in the real world is constantly updated, and some of the originally correct knowledge in LLMs can become outdated and invalid over time. For example, the question "Who is the President of the United States?" is answered "Donald Trump" in the year 2020, while the answer

*Equal contribution.

†Corresponding author.




Figure 1: Diagram for "Forgetting before Learning".

now is "Joe Biden". Consequently, the challenge with LLMs is continuously updating to ensure they reflect current, correct knowledge. Existing methods of model editing and knowledge updating usually add additional network (Dong et al., 2022; Huang et al., 2022; Raunak and Menezes, 2022), model parameters (Dai et al., 2023; Dong et al., 2022; Huang et al., 2022), knowledge bases (Murty et al., 2022; Mitchell et al., 2022; Li et al., 2022; Madaan et al., 2022; Mitchell et al., 2022; Zheng et al., 2023), etc., and the editing process is not as straightforward and simple as fine-tuning methods (Zhang et al., 2022; Li and Liang, 2021; Hu et al., 2021) directly with new knowledge. Currently, the most used method for learning new knowledge is still direct fine-tuning of the model.

Empirically, when human beings establish their own initial cognition, if they are exposed to new knowledge that is inconsistent with their initial cognition, they usually feel conflicted and it is difficult for them to learn and accept the new knowledge. If the original cognition and knowledge are forgotten, then the new knowledge to be learned will not conflict with the original cognition and knowledge, which makes it better to learn and absorb the new knowledge. As shown in Figure 1, it is better to pour in the "new water" only after the "original water" in the cup has been poured out. For example, if people have been educated to believe that "the Earth is flat" since childhood, it would be challenging for them to accept the conflicting knowledge that "the Earth is round" when they become

adults. Conversely, if they can forget the erroneous knowledge that "the Earth is flat" or if they learn the correct knowledge that "the Earth is round" before being exposed to the incorrect information, it would be much simpler.

Inspired by the above empirical observations and (Ilharco et al., 2022)'s task arithmetic, we propose a novel paradigm of knowledge updating called **F-Learning** (Forgetting before Learning). Specifically, we first fine-tune the initial model using old knowledge and then subtract the difference between the fine-tuned model parameters and the initial model parameters from the initial model parameters. This process is defined as "*old knowledge forgetting*". We then use the new knowledge to fine-tune the model after forgetting the old knowledge. This process we define as "*new knowledge learning*". After the two stages of *forgetting old knowledge* and *learning new knowledge*, the model's knowledge is updated. The contribution of this work can be summarised as follows:

- We propose a novel fine-tuning paradigm "Forgetting before Learning" (F-Learning) for knowledge updating in large language models.
- Experimental results show that our proposed F-Learning improves the knowledge updating performance of various fine-tuning methods and outperforms the existing baselines in most cases.
- Experimental results show that forgetting by subtracting the parameters of LoRA can achieve the approximate effect of subtracting the parameters of full fine-tuning.

2 Related Work

Currently, the method of knowledge updating and model editing (also known as knowledge editing) for LLMs is mainly divided into two classes (Yao et al., 2023; Wang et al., 2023):

a. The method preserving model's parameters

Retrieve augmentation practically depends on an external knowledge base which contains new or correct knowledge. Aiming at amending the output of LLMs, a new knowledge base will be connected with the base model to implement a retrieve for needed new knowledge to a prompt or a question (Murty et al., 2022; Mitchell et al., 2022; Li et al., 2022; Madaan et al., 2022). Mitchell et al. (Mitchell et al., 2022) store manual edits

in a memory module, and use a classifier to call the knowledge stored in the memory. Madaan et al. (Madaan et al., 2022) leverage the memory of user's feedback to generate prompts for LLMs. Instead of gradient calculation, Zheng et al. (Zheng et al., 2023) utilize the in-context learning method to revise the output of LLMs with demonstrations extracted from the corpus based on similarity.

Adding Additional Parameters refers to injecting a few trainable parameters which represent new knowledge to LLMs while original parameters keeping frozen (Dong et al., 2022; Huang et al., 2022; Raunak and Menezes, 2022; Dai et al., 2023). Dong et al. (Dong et al., 2022) put forward a lightweight feed-forward network to add new parameters adapted to specific factual contexts for knowledge generalization. Huang (Huang et al., 2022) et al. design an editor called Transformer-Patcher, which is capable of modifying the mistake of LLMs sequentially by adding and training a few neurons in transformer.

b. The method modifying model's parameters

Fine-tuning is a general technique since pre-training model has been widely adopted in NLP research, which always obtains promising results in downstream tasks. Meanwhile, fine-tuning is an intuitive and effective method to urge the model to learn new knowledge for model editing (Zhu et al., 2020; Zhang et al., 2022; Yao et al., 2023). Recently, there are a series of parameter-efficient fine-tuning methods, such as Prefix-Tuning (Li and Liang, 2021) and LoRA (Hu et al., 2021), making it more appreciate for knowledge editing based on fine-tuning. Zhang et al. (Zhang et al., 2022) operate incremental parameter updates of different amounts by calculating the importance of the weight matrix to improve the update efficiency and adaptability. Zhu et al. (Zhu et al., 2020) leverage a loss constraint attached to the base model to reduce the impact on irrelevant knowledge during the process of fine-tuning. Similarly, Lee et al. (Lee et al., 2022) also implement large-scale continual learning for knowledge updating with regularized fine-tuning.

Meta-learning is aimed at updating the knowledge in LLMs through varying their parameters with the prediction from a well-trained hypernetwork (Sinitsin et al., 2019; Mitchell et al., 2021; De Cao et al., 2021). Mitchell et al. (Mitchell et al., 2021) propose an auxiliary network with gradient decomposition, which can execute efficient edits to LLMs

according to a single input-output pair. De Cao et al. (De Cao et al., 2021) update part of weights for a subset of modules in the model relying on a hypernetwork with constrained optimization.

Locate and edit is related to the internal mechanism of the LLMs. With the help of some attributes, it usually locates the parameters and neurons in the light of specific knowledge and modifies them to correct the output (Meng et al., 2022a; Dai et al., 2022; Meng et al., 2022b; Santurkar et al., 2021; Geva et al., 2022). Geva et al. (Geva et al., 2021) find that the feed-forward networks layer of the transformer stores key-value pairs which are related to specific knowledge. Meng et al. (Meng et al., 2022a) utilize a causal reasoning method to distinguish the key neuron activations and update specific factual associations by modifying feed-forward weights. Furthermore, to implement knowledge editing on a large scale, they put forward MEMIT (Meng et al., 2022b), a method that directly updates thousands of memories in LLMs. Gupta et al. (Gupta et al., 2023) improve the knowledge updating through varying edit tokens and ameliorating the layer selection during the editing process. Yu et al. (Yu et al., 2023) leverage the partitioned gradient to identify the significant weights for unlearning of bias in the model.

In conclusion, there are many ways to achieve knowledge updating, but most of them require the addition of additional knowledge bases, neural network modules, and model parameters, which are cumbersome in practice and increase inference consumption. **This paper focuses on the improvement and enhancement of fine-tuning methods.**

3 Task Definition

Our task is knowledge updating of large models, which can be defined as given a model f_θ and a set of input-output knowledge pairs $K_{old} = \{(x_1, y_1), (x_2, y_2), \dots, (x_i, y_i)\}$, the parameters of the model need to be edited to obtain a new model f_{θ^*} (x) and a corresponding set of new input-output pairs $K_{new} = \{(x_1, y_1^{new}), (x_2, y_2^{new}), \dots, (x_i, y_i^{new})\}$. The i is the number of knowledge pairs to be updated. Referring to (Yao et al., 2023), we can define this process and objective of knowledge updating as:

$$f_{\theta^*}(x_i) = \begin{cases} y_i^{new} & \text{if } x_i \in N(x_i) \\ f_\theta(x_i) & \text{if } x_i \in other \end{cases} \quad (1)$$

where $N(x_i)$ represents x_i itself and its equivalent neighbourhood. The knowledge update task needs to update only the answers of x_i itself and its equivalent domain $N(x_i)$ without changing the answers of other out-of-scope knowledge. Specifically, the quality of knowledge updating has the following three evaluation indicators: (1) **Reliability** is measured as the average accuracy on the new knowledge for the updated model f_{θ^*} . It's the first indicator of the effectiveness of knowledge updating. As shown in Figure 2, the output of the question "Who is the President of the US?" needs to be updated from "Donald Trump" to "Joe Biden". (2) **Generalization** means the new model f_{θ^*} should also updated the equivalent neighbour $N(x_i)$ (e.g. rephrased sentences). It is evaluated by the average accuracy of the model f_{θ^*} on examples drawn uniformly from the equivalence neighborhood. As shown in Figure 2, the output of the question "Who holds the position of the President of the US?" also needs to be updated from "Donald Trump" to "Joe Biden". (3) **Locality** means the updated model f_{θ^*} should not change the output of the irrelevant examples. Hence, the locality is evaluated by the rate at which the updated model f_{θ^*} 's predictions are unchanged as the pre-update f_θ model. As shown in Figure 2, the output of the question "'You're fired!' is the catchphrase of which celebrity?" is to be kept unchanged as "Donald Trump".

4 Proposed method: F-Learning

In this section, we will present our method of knowledge updating for LLMs. Instead of introducing an external knowledge base or additional parameters, our method is mainly based on **full fine-tuning** and **parameter-efficient fine-tuning**. Briefly, it consists of two stages:

4.1 Forgetting old knowledge

The supervised fine-tuning (SFT) on a dataset injects new knowledge into the LLMs or activates their fitting capabilities related to the new knowledge, which is reflected in the variation of the model's parameters. During this stage, for a given large language model f_θ and its parameters θ , we define the incremental parameters as **knowledge parameters** θ_Δ , calculated as follows:

$$\theta_\Delta = \text{FT}\{\theta, K\} - \theta \quad (2)$$

where FT is the operation of supervised fine-tuning, while K, θ refer to the dataset of knowledge

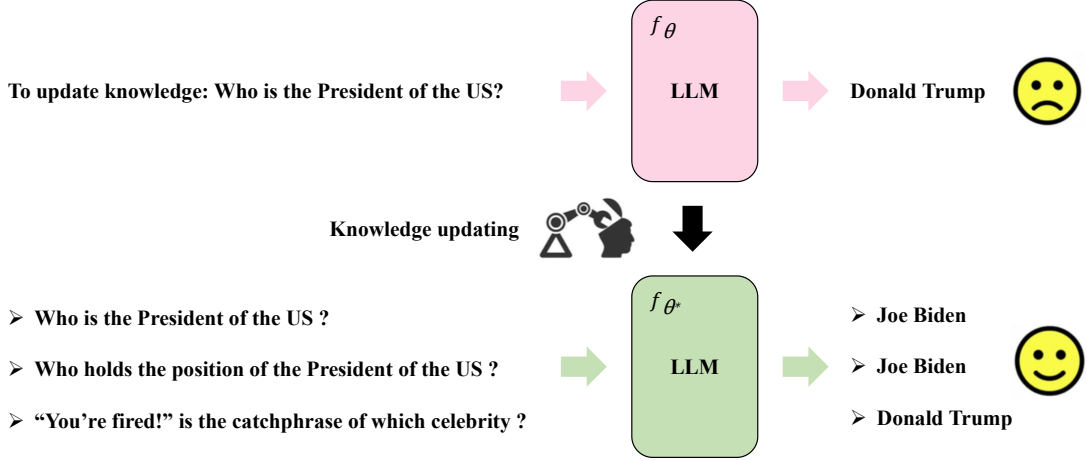


Figure 2: Objectives of the knowledge updating in large language model.

and the parameters of the original model f_θ , respectively. Similarly, we first fine-tune the model f_θ on a dataset containing old knowledge, and then subtract the parameters θ of the original model f_θ from model’s parameters after fine-tuning to obtain the knowledge parameters θ_Δ^{old} indicating the old knowledge, as follows:

$$\theta_\Delta^{old} = \text{FT}\{\theta, K_{old}\} - \theta, \quad (3)$$

where K_{old} refers to a dataset composed of old knowledge which we desire to forget. Inspired by (Ilharco et al., 2022), we believe that subtracting the parameters θ_Δ^{old} from θ can assist the model f_θ to forget this part of old knowledge. So we define the process of forgetting old knowledge as follows:

$$\theta' = \theta - \lambda\theta_\Delta^{old}, \quad (4)$$

where λ is a hyper-parameter to control the rate of forgetting. Now we gain a new model $f_{\theta'}$ with its parameters θ' , which has forgotten the old knowledge compared to f_θ . Note that this process of forgetting old knowledge only makes sense if the model f_θ has already learned the old knowledge, otherwise, there is no need for forgetting and the forgetting operation may have a destructive effect on the normal knowledge of the model.

4.2 Learning new knowledge

With the model $f_{\theta'}$ that has gone through the process of forgetting old knowledge, then we will inject the new knowledge to $f_{\theta'}$ for knowledge updating by supervised fine-tuning. Similarly, we define the process of learning new knowledge as follows:

$$\theta^* = \text{FT}\{\theta', K_{new}\} \quad (5)$$

where FT is the operation of fine-tuning, θ^* is the parameters of a new model f_{θ^*} which has learned the new knowledge compared to f_θ and K_{new} refers to a dataset composed of new knowledge which we need to learn and update for f_θ .

5 Experiments

5.1 Datasets

In this work, we use ZsRE (Levy et al., 2017) and COUNTERFACT (Meng et al., 2022a), two widely used datasets, for our experiments. ZsRE is a Question Answering (QA) dataset that utilizes question rephrasings generated by back-translation as the equivalence neighborhood. COUNTERFACT is a more challenging dataset with counterfactual data. We follow the setting of (Yao et al., 2023) to take the eval and edit sets of which there are 19,085 and 10,000 pieces of data respectively. Moreover, we divide the datasets into two parts of old knowledge and new knowledge respectively to achieve two-stage knowledge update. The following is an example of old knowledge and new knowledge in zsRE, which represents the modification of knowledge from "Los Angeles" to "New Orleans". More details about the datasets and examples can be found in the appendix A.1.

The old knowledge:

```
{ "instruction": "What city did Marl Young live when he died?", "input": "", "output": "Los Angeles" }
```

The new knowledge:

```
{ "instruction": "What city did Marl Young live when he died?", "input": "", "output": "New Orleans" }
```


Dataset	Editor	LLAMA2-7B			GPT-2-XL		
		Reliability	Generality	Locality	Reliability	Generality	Locality
ZsRE	Original model	43.70	43.17	/	31.73	31.46	/
	LoRA	43.10	42.20	70.83	36.58	35.40	79.21
	F-Learning _{LoRA}	46.91	46.21	72.50	37.09	36.18	79.88
	FT-c	49.02	46.96	67.37	37.52	36.34	65.52
	Full-FT	81.02	74.67	70.51	45.96	44.73	74.95
	ROME	43.67	42.66	93.14	34.79	33.41	73.2
	MEMIT	83.57	79.06	70.52	46.5	45.26	76.44
	F-Learning _{LoRA-FT}	82.43	77.38	71.04	46.37	44.25	76.79
	F-Learning _{FT}	84.65	81.51	70.92	46.92	44.96	76.86
COUNTERFACT	Original model	18.47	16.95	/	11.29	10.72	/
	LoRA	30.56	23.24	40.08	13.2	12.54	46.73
	F-Learning _{LoRA}	31.17	23.63	40.42	14.06	12.97	47.20
	FT-c	29.23	19.32	19.70	14.33	13.47	26.9
	Full-FT	65.99	44.08	28.34	21.05	15.95	33.67
	ROME	18.41	17.20	<u>93.60</u>	13.42	12.85	33.82
	MEMIT	61.94	37.45	21.90	20.97	16.2	33.9
	F-Learning _{LoRA-FT}	78.73	51.67	29.49	19.8	16.04	34.11
	F-Learning _{FT}	69.53	45.56	28.41	21.32	16.41	34.39

Table 1: Results on three metrics of the two datasets based on LLAMA2-7B and GPT-2-XL.

5.2 Baselines

To evaluate the effectiveness of the proposed F-Learning method, we conducted experiments on fine-tuning methods and locate-based methods. For fine-tuning methods, we first compare with the full fine-tuning (**Full-FT**) and **LoRA** (Hu et al., 2021), respectively. LoRA (**Low-Rank Adaptation**) is a technique for fine-tuning large pre-trained language models by introducing small, trainable matrices into each layer of the model’s architecture, allowing for efficient adaptation while keeping the majority of the model’s parameters frozen. Then we experiment with a fine-tuning approach (**FT-c**) (Zhu et al., 2020) that leverages L_∞ constraint to retain old irrelevant knowledge. For locate-based methods, we first experiment with **ROME** (Meng et al., 2022a), a method updating specific factual associations with causal intervention. Finally, we compare with the **MEMIT** (Meng et al., 2022b) which is an effective method to directly update large-scale memories.

5.3 Completion Details

We use LLAMA2-7B and GPT-2-XL (1.5B parameters) as the base models for our experiments. We are mainly evaluating the ability to update old knowledge to new knowledge, thus we trained the base model on the old knowledge for 3 epochs by full fine-tuning as the **original model** in our experiment (as the same as **original model** in other

experiments). Original model has fully learned old knowledge, which makes the forgetting operation reasonable and necessary. To ensure the uniqueness of the model output, we used the greedy decoding strategy during testing. On the hardware side, a total of $4 \times$ A100-80G GPUs were used for the experiments. More details about the experimental settings can be found in the appendix A.2.

5.4 Experimental Results

The experimental results are presented in Table 1, indicating a notable enhancement in learning following the initial forgetting process, irrespective of whether full fine-tuning or LoRA is employed. We obtain the promising results as our F-Learning method outperforming other baselines in most cases, which also show that our method can work well across different model sizes and architectures. Specifically, compared with other editors, FT-c has only small improvements over the original model. This could be attributed to its norm regularization, which makes FT-c tend to retain a part of the old knowledge during the knowledge updating process. Given that our original model has fully learned a large quantity of old knowledge, it faces greater challenges in learning new knowledge. Surprisingly, for LLAMA2-7B, ROME maintains Reliability and Generality almost unchanged from the original model on two datasets, while having a high locality (more than 90), which suggests that it per-

forms little knowledge updating (As the injection of new knowledge will usually have an impact on locality), likely because that it can only edit a small number of parameters, coupled with the fact that our original model is full of old knowledge, which obstructs the effect of the causal tracing mechanism in ROME. It is worth noting that full fine-tuning is much more capable of learning new knowledge than LoRA, as LoRA focuses on training a limited subset of parameters within the attention structure, while the majority of factual knowledge is encoded in the MLP layers.

5.5 Forgetting with LoRA and Learning with Full Fine-tuning

In the above setting of experiments, the method we adopt is to perform old knowledge forgetting and new knowledge learning based on full (or LoRA) fine-tuning at the same time. Nonetheless, we find that subtracting the knowledge parameters of full fine-tuning (i.e. forgetting old knowledge by full fine-tuning) in some cases will completely destroy the core functions of our base model, resulting in a significant drop in evaluation metrics. In this view, as LoRA is a parameter-efficient fine-tuning method that has less impact on parameters compared to full fine-tuning, we try a new method that forgets old knowledge by LoRA and then learns new knowledge by full fine-tuning as a trade-off. Similar to the method above (§4), we define this process as follows:

$$\theta_{\Delta}^{old} = \text{LoRA}\{\theta, K_{old}\} - \theta, \quad (6)$$

$$\theta' = \theta - \lambda\theta_{\Delta}^{old}, \quad (7)$$

$$\theta^* = \text{FT}\{\theta', K_{new}\} \quad (8)$$

where LoRA is the operation of supervised fine-tuning by LoRA while FT is the operation of supervised fine-tuning by full fine-tuning. θ^* is noted as the parameters of the edited model f_{θ^*} which has completed the knowledge updating.

For verification, we keep the same experimental settings as above and conduct the experiment. The results are shown in Table 1. The results support that the method of forgetting with LoRA and then learning with full fine-tuning (noted as **F-Learning**_{LoRA-FT}) surpassing many baselines such as the directly full fine-tuning, as well as almost nearly match or even surpassing the method of forgetting and learning with full fine-tuning in some cases. In particular, it generally maintains

the highest results in locality, which may be owing to the parameter-efficiency of LoRA-based old knowledge forgetting.

After conducting experiments, we empirically discovered that utilizing the technique of forgetting through the subtraction of LoRA parameters can approximate the effect achieved by subtracting the parameters during full fine-tuning. We hold that although LoRA-based knowledge forgetting does not eliminate the old knowledge stored in the MLP layers of the LLMs, it alters the patterns and relationships associated with the old knowledge stored in the attention structure (i.e., an implicit knowledge representation), which facilitates the new knowledge learning. This finding holds significant value due to the considerable reduction in time and computational costs associated with LoRA compared to full fine-tuning, as well as bringing a smaller risk of model collapse.

5.6 Time Testing

In order to evaluate the efficiency of our proposed F-learning method, we calculated the editing time of several different knowledge updating and model editing methods for different numbers of edits. Taking LLAMA2-7B as an example. The results are shown in the Table 2.

We can find that the time consumed by the fine-tuning based method is significantly less than that of the locate-based method. This is because the locate-based method highly relies on the location of neurons and parameters, which increases the complexity and time of editing. Furthermore, since ROME can only edit a single piece of data at a time, while other methods can edit in batches, ROME is less efficient. Compared with other fine-tuning based methods, FT-c can be optimized faster with its norm constraint. The F-learning we proposed is a two-stage knowledge updating method that forgets before learning, as it takes about twice as long as Full-FT, but is still very fast and convenient. It is worth noting that although the forgetting operation requires an additional training process, once the training is completed, the parameters of this part of forgetting old knowledge can be reused during the subsequent optimization and inference process, which can save resources and time. Meanwhile, we can further accelerate supervised fine-tuning by deepspeed or other approaches.

Editor	1 edit		10 edits		100 edits	
	zsRE	COUNTERFACT	zsRE	COUNTERFACT	zsRE	COUNTERFACT
FT-c	0.57(s)	0.54(s)	6.58(s)	6.82(s)	21.77(s)	19.30(s)
ROME	20.47(s)	18.27(s)	207.09(s)	179.30(s)	2184.16(s)	1810.42(s)
MEMIT	28.32(s)	23.71(s)	108.67(s)	96.71(s)	862.20(s)	847.72(s)
Full-FT	0.76(s)	0.72(s)	7.8(s)	7.3(s)	25.36(s)	24.70(s)
F-Learning_{FT}	1.58(s)	1.47(s)	15.32(s)	14.9(s)	52.20(s)	50.12(s)

Table 2: Editing time for 1 edit, 10 edits, 100 edits of the two dataset based on LLAMA2-7B.

5.7 Parametric Analysis of Forgetting Old Knowledge

The two-stage knowledge updating method we proposed highly utilizes the forgetting of old knowledge. From the perspective of interpretability, here we analyzed the parameters of old knowledge forgetting and further analyzed the parameter distribution and changes within the LLMs. The results are shown in Figure 4 and Figure 5 which are in the appendix. Taking the LLAMA2-7B and zsRE dataset as an example, specifically, we analyzed different parameters in two cases: forgetting old knowledge by full fine-tuning and forgetting old knowledge by LoRA (The hyperparameters λ of the rate of forgetting are both set to 1). We compare the parameter θ' of the model $f_{\theta'}$ after forgetting the old knowledge with the parameter θ of the original model f_{θ} , and calculate their Euclidean distance on each layer. We select the results with layer $n=[6, 15, 24, 30]$ to exhibit for simplicity. In general, there is little difference in parameter changes between low-layers and high-layers within the model. For forgetting old knowledge by full fine-tuning, we can find that the parameter changes of the MLP layers are more significant than attention layers. This may be one of the reasons why "forgetting before learning" is effective, as knowledge is generally stored in the MLP layers. Relatively LoRA has less impact on parameters (Euclidean distances are less than 1), and only changes the parameters of "query" and "value" in the attention layers. Therefore, it is more limited than full fine-tuning. However, LoRA-based forgetting can help forget the patterns and relationships associated with old knowledge stored in the attention layers, and thus can also assist in knowledge updating.

5.8 Experiments on the Old Knowledge Forgetting

Here we evaluate the model performance on old knowledge data after performing only the forgetting operation by full fine-tuning or LoRA fine-

tuning with different forgetting rates to verify the effectiveness of the forgetting operation. The results are shown in Figure 6 and Figure 7 which are in the appendix. Taking the LLAMA2-7B and zsRE dataset as an example. We set the hyperparameters λ of the rate of forgetting is set to 0.9, 0.7, 0.5, 0.3, 0.1, respectively. It can be easily found that the performance of the model on three metrics is negatively correlated with the rate of forgetting, i.e., the old knowledge in the model decreases as the rate of forgetting increases. And under the same circumstances, LoRA brings less knowledge forgetting than the full fine-tuning.

5.9 Impact to Other Capabilities within LLMs Testing

Here we test the impact of forgetting old knowledge and learning new knowledge over the original model on other capabilities of the model (such as mathematical abilities). Specifically, we evaluated changes in the model's mathematical capabilities on GSM8K and MATH, and evaluated changes in the model's comprehensive examining capabilities on MMLU. Then we evaluated the changes in the model's capabilities on Reading comprehension, Theorem application, Commonsense reasoning, Content analysis, respectively. Taking LLAMA2-7B as an example. The results are shown in the Table 3. "Original model+Forgetting" refers to only forgetting the old knowledge over the original model, and "Original model+F-learning" refers to the original model with our F-learning method. The hyper-parameters λ of the rate of forgetting is set to 0.3. And the metric of evaluation in experiments is mainly "Accuracy". Experiments show that F-learning can slightly improve the mathematical ability of the model after forgetting old knowledge, but it will decrease after learning new knowledge. While in other areas of professional examining capabilities, our method has little impact. Interestingly, the model's performance on "college-chemistry" and "college-medicine" has been signif-

Dataset	Editor		
	Original model	Original model+Forgetting	Original model+F-learning
Math reasoning			
GSM8K	2.35	2.5	1.44
MATH	3.2	3	1.96
Average	2.78	2.75	1.7
Examinations			
MMLU-college-chemistry	24	34	41
MMLU-college-mathematics	29	30	26
MMLU-management	16.50	14.56	32.04
MMLU-computer-security	21	20	23
MMLU-macroeconomics	32.31	33.08	34.87
MMLU-college-physics	19.61	26.47	22.55
MMLU-astronomy	30.92	23.03	32.24
MMLU-professional-law	26.47	25.62	24.51
MMLU-college-medicine	24.28	24.28	29.48
MMLU-econometrics	22.81	23.68	21.05
MMLU-formal-logic	28.57	31.75	30.95
MMLU-mathematics	24.28	24.28	29.48
Average	24.98	25.89	28.93
Reading comprehension			
OpenBookQA	15.6	14.2	13.2
OpenBookQA-fact	6.2	6.1	6.6
Average	10.9	10.15	9.9
Theorem application			
TheoremQA	7.38	7.12	6.9
StrategyQA	16.9	16.4	15.14
Average	12.14	11.76	11.02
Commonsense reasoning			
HellaSwag	59.77	60.02	55.39
Content analysis			
LAMBADA	61.96	62.33	60.89

Table 3: Results on accuracy of the different types of tasks datasets based on LLAMA2-7B.

icantly improved after completing the knowledge update. This may be because the dataset contains relevant knowledge. For other datasets of evaluation, it can be revealed that the performance on various tasks only generally changed by about 1 point on average, which shows that our method has very little impact on other capabilities of the model.

6 Case Study

To further illustrate the effectiveness of the proposed method, we present a case study on the results of the knowledge updating by the original model, only forgetting old knowledge and performing F-learning. We selected some cases in the experiment of llama2-7B on zsRE dataset, noting that the hyper-parameters λ of the rate of forgetting is set to 0.3. Table 4 shows the results during different knowledge updating stages. From the first example and second example, we can find that

model begins to output some irrelevant content after performing the forgetting operation, indicating that it gradually forgets the old knowledge. In example 4, the forgetting operation failed to assist the model in forgetting old knowledge, but it still completed knowledge updating with the help of F-learning. However, sometimes there are some bad cases, such as example 3, where the model never learned new knowledge, which shows that our method has certain limitations and could be improved.

7 Conclusion

In this paper, we propose a new paradigm of knowledge updating during supervised fine-tuning called **F-Learning** (Forgetting before Learning), which is based on parametric arithmetic to forget old knowledge and learn new knowledge for eliminating contradictions between old and new knowledge. The

<p>(1) Prompt: <i>What artist created Call the Doctor?</i> Original answer: Riders in the Sky Target answer: The X-Files Original model: Riders in the Sky Original model + Old knowledge forgetting: Doctor Who Original model + F-learning: The X-Files</p>
<p>(2) Prompt: <i>What university did Watts Humphrey take part in?</i> Original answer: Trinity College Target answer: University of Michigan Original model: Trinity College Original model + Old knowledge forgetting: The Wire Original model + F-learning: University of Michigan</p>
<p>(3) Prompt: <i>What role does Denny Herzig play in football?</i> Original answer: midfielder Target answer: winger Original model: midfielder Original model + Old knowledge forgetting: midfielder Original model + F-learning: goalkeeper</p>
<p>(4) Prompt: <i>Which family does Ramalinaceae belong to?</i> Original answer: Rmales Target answer: Lamiinae Original model: Rmales Original model + Old knowledge forgetting: Rmales Original model + F-learning: Lamiinae</p>
<p>(5) Prompt: <i>Who's the architect of Toodyay Fire Station?</i> Original answer: Wong Tung and Partners Target answer: Wyndham Lewis Original model: Wong Tung and Partners Original model + Old knowledge forgetting: Wong Tung Original model + F-learning: Wyndham Lewis</p>

Table 4: Examples of knowledge updating results during different stages.

experiments on zsRE and COUNTERFACT datasets show that our method surpasses other baselines in most cases. Simultaneously we find that forgetting old knowledge by subtracting the parameters of LoRA can achieve the similar effect of subtracting the parameters of full fine-tuning, which is inspiring. We will further investigate the updating of knowledge and generalize the proposed method from single-modal domain to multi-modal domains of knowledge.

8 Limitations

In this work, the proposed F-learning paradigm, although it improves the effectiveness of the fine-tuning methods for updating the knowledge of large language models, adds extra computation due to an extra forgetting process. And more relevant theoretical foundations need to be researched, such as the existence form of specific knowledge, etc.

Acknowledgement

This work was supported by National Key Research and Development Program of China (2022YFF0902100), Postdoctoral Fellowship Program of CPSF (GZC20232873),

GuangDong Basic and Applied Basic Research Foundation (2023A1515110718 and 2024A1515012003), National Natural Science Foundation of China (62376262), the Natural Science Foundation of Guangdong Province of China (2024A1515030166), Shenzhen Science and Technology Innovation Program (KQTD20190929172835662), Shenzhen Basic Research Foundation (JCYJ20210324115614039).

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. Knowledge neurons in pretrained transformers. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502.
- Damai Dai, Wenbin Jiang, Qingxiu Dong, Yajuan Lyu, and Zhifang Sui. 2023. Neural knowledge bank for pretrained transformers. In *CCF International Con-*

- ference on Natural Language Processing and Chinese Computing*, pages 772–783. Springer.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. Editing factual knowledge in language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6491–6506.
- Qingxiu Dong, Damai Dai, Yifan Song, Jingjing Xu, Zhifang Sui, and Lei Li. 2022. Calibrating factual knowledge in pretrained language models. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5937–5947.
- Mor Geva, Avi Caciularu, Kevin Wang, and Yoav Goldberg. 2022. Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 30–45.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495.
- Anshita Gupta, Debanjan Mondal, Akshay Krishna Sheshadri, Wenlong Zhao, Xiang Lorraine Li, Sarah Wiegrefe, and Niket Tandon. 2023. Editing commonsense knowledge in gpt. *arXiv preprint arXiv:2305.14956*.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Zeyu Huang, Yikang Shen, Xiaofeng Zhang, Jie Zhou, Wenge Rong, and Zhang Xiong. 2022. Transformer-patcher: One mistake worth one neuron. In *The Eleventh International Conference on Learning Representations*.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. 2022. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations*.
- Kyungjae Lee, Wookje Han, Seung-won Hwang, Hwaran Lee, Joonsuk Park, and Sang-Woo Lee. 2022. Plug-and-play adaptation for continuously-updated qa. *arXiv preprint arXiv:2204.12785*.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. Zero-shot relation extraction via reading comprehension. *arXiv preprint arXiv:1706.04115*.
- Daliang Li, Ankit Singh Rawat, Manzil Zaheer, Xin Wang, Michal Lukasik, Andreas Veit, Felix Yu, and Sanjiv Kumar. 2022. Large language models with controllable working memory. *arXiv preprint arXiv:2211.05110*.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597.
- Aman Madaan, Niket Tandon, Peter Clark, and Yiming Yang. 2022. Memory-assisted prompt editing to improve gpt-3 after deployment. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2833–2861.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022a. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372.
- Kevin Meng, Arnab Sen Sharma, Alex J Andonian, Yonatan Belinkov, and David Bau. 2022b. Mass-editing memory in a transformer. In *The Eleventh International Conference on Learning Representations*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. 2021. Fast model editing at scale. In *International Conference on Learning Representations*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. 2022. Memory-based model editing at scale. In *International Conference on Machine Learning*, pages 15817–15831. PMLR.
- Shikhar Murty, Christopher D Manning, Scott Lundberg, and Marco Tulio Ribeiro. 2022. Fixing model bugs with natural language patches. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11600–11613.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Vikas Raunak and Arul Menezes. 2022. Rank-one editing of encoder-decoder models. *arXiv preprint arXiv:2211.13317*.
- Shibani Santurkar, Dimitris Tsipras, Mahalaxmi Elango, David Bau, Antonio Torralba, and Aleksander Madry. 2021. Editing a classifier by rewriting its prediction rules. *Advances in Neural Information Processing Systems*, 34:23359–23373.

- Anton Sinitin, Vsevolod Plokhotnyuk, Dmitry Pyrkin, Sergei Popov, and Artem Babenko. 2019. Editable neural networks. In *International Conference on Learning Representations*.
- Song Wang, Yaochen Zhu, Haochen Liu, Zaiyi Zheng, Chen Chen, et al. 2023. Knowledge editing for large language models: A survey. *arXiv preprint arXiv:2310.16218*.
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023. Editing large language models: Problems, methods, and opportunities. *arXiv preprint arXiv:2305.13172*.
- Charles Yu, Sullam Jeoung, Anish Kasi, Pengfei Yu, and Heng Ji. 2023. Unlearning bias in language models by partitioning gradients. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6032–6048.
- Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. 2022. Adaptive budget allocation for parameter-efficient fine-tuning. In *The Eleventh International Conference on Learning Representations*.
- Ce Zheng, Lei Li, Qingxiu Dong, Yuxuan Fan, Zhiyong Wu, Jingjing Xu, and Baobao Chang. 2023. Can we edit factual knowledge by in-context learning? *arXiv preprint arXiv:2305.12740*.
- Chen Zhu, Ankit Singh Rawat, Manzil Zaheer, Srinadh Bhojanapalli, Daliang Li, Felix Yu, and Sanjiv Kumar. 2020. Modifying memories in transformer models. *arXiv preprint arXiv:2012.00363*.

A Appendix

A.1 Datasets and Examples

We will illustrate the datasets we used in more detail. ZsRE is a Question Answering (QA) dataset that utilizes question rephrasings generated by back-translation as the equivalence neighborhood. COUNTERFACT is a more challenging dataset with counterfactual data. We take the eval and edit sets of which there are 19,085 and 10,000 pieces of data respectively.

The following is a sample of the ZsRE dataset:

```
{ "subject": "Watts Humphrey", "src": "What university did Watts Humphrey attend?", "pred": "Trinity College", "rephrase": "What university did Watts Humphrey take part in?", "alt": "University of Michigan", "answers": ["Illinois Institute of Technology"], "loc": "nq question: who played desmond doss father in hacksaw ridge", "loc-ans": "Hugo Weaving", "cond": "Trinity College » University of Michigan || What university did Watts Humphrey attend?" }
```

It represents that for prompt "What university did Watts Humphrey attend?", modifying the old knowledge "Trinity College" into the new knowledge "University of Michigan". Meanwhile, "rephrase" is used to evaluate the model's Generalization metric, and "loc" is used to evaluate the model's Locality metric.

Furthermore, we can find that old knowledge and new knowledge have some correlation, they keep the same questions with different answers. We keep them in the same format to ensure the training effect. To facilitate our supervised fine-tuning training, we divide the datasets into two parts of old knowledge and new knowledge, and convert them into an instruction fine-tuning format, an example as follows:

The old knowledge:

```
{ "instruction": "What university did Watts Humphrey attend?", "input": "", "output": "Trinity College" }
```

The new knowledge:

```
{ "instruction": "What university did Watts Humphrey attend?", "input": "", "output": "University of Michigan" }
```

What calls for special attention is that the two datasets used in our experiments are both counterfactual datasets, in which the old knowledge is correct knowledge in the real world, and the new knowledge (target knowledge) is wrong knowledge in the real world, so the labels of old knowledge

and new knowledge in these datasets are given artificially and have nothing to do with time and correctness in the real world. They are only used to measure whether the model can accurately modify the knowledge. Since the new knowledge is wrong knowledge in the real world, it can ensure that the original LLM has not learned it before, thus avoiding the problem of being unable to determine whether the new knowledge output by the LLM is learned from the data or possessed by itself.

A.2 Implementation Details of Experiments

Here we will introduce more completion details and settings of experiments. First, we used LLAMA2-7B and GPT-2-XL as the base models, and then we trained the base model on the old knowledge for 3 epochs by full fine-tuning to simulate an original model that has fully learned old knowledge for our experiments. This makes the forgetting operation more reasonable and effective, and at the same time tries to avoid the problem of being unable to determine whether the new knowledge output by the LLM is learned from the data or commanded by itself as mentioned above.

A.2.1 F-learning

For the experiments of zsRE on LLAMA2-7B, the hyperparameters λ set in F-Learning_{LoRA} is 0.7, while 3 in F-Learning_{LoRA-FT} and 0.3 in F-Learning_{FT}. Similarly, for COUNTERFACT dataset, the hyperparameters λ set in F-Learning_{LoRA} is 1.5, while 3 in F-Learning_{LoRA-FT} and 0.1 in F-Learning_{FT}. The learning rate, epochs for all above experiments are $5e-5$ and 3, then batch-size is 4 and gradient-accumulation-steps is 4.

For the experiments of zsRE on GPT-2-XL, the hyperparameters λ set in F-Learning_{LoRA} is 0.7, while 2 in F-Learning_{LoRA-FT} and 0.3 in F-Learning_{FT}. For COUNTERFACT dataset, the hyperparameters λ set in F-Learning_{LoRA} is 1, while 3 in F-Learning_{LoRA-FT} and 0.3 in F-Learning_{FT}. The epochs for all above experiments are 3, then batch-size is 4 and gradient-accumulation-steps is 4. The learning rate for zsRE experiments is $5e-5$ while $1e-5$ in COUNTERFACT experiments.

When we used the Deepspeed, we set 4 processes and zero-stage is 2.

Dataset	Editor	LLAMA-7B		
		Reliability	Generality	Locality
ZsRE	Original model	43.29	42.85	/
	LoRA	46.93	45.87	75.86
	F-Learning _{LoRA}	47.56	46.90	76.09
	FT-c	47.33	45.51	68.14
	Full-FT	70.52	66.69	65.26
	ROME	43.45	42.94	98.60
	MEMIT	78.30	77.43	69.44
	F-Learning _{LoRA-FT}	75.17	70.12	69.78
	F-Learning _{FT}	83.06	79.50	70.09
COUNTERFACT	Original model	21.61	17.88	/
	LoRA	27.54	21.21	39.75
	F-Learning _{LoRA}	29.47	22.89	44.91
	FT-c	26.97	17.90	20.09
	Full-FT	32.13	31.95	32.51
	ROME	21.83	19.08	92.27
	MEMIT	56.94	31.48	25.70
	F-Learning _{LoRA-FT}	32.43	26.89	37.14
	F-Learning _{FT}	56.39	39.75	31.87

Table 5: Results on three metrics of the two datasets based on LLAMA-7B.

A.2.2 Full-FT and LoRA

Full-FT and LoRA refer to knowledge updating by full fine-tuning and LoRA fine-tuning in our experiments. We adopted experimental settings similar to F-learning as mentioned above. The difference is that these two do not forget the old knowledge. Full-FT and LoRA also use the instruction fine-tuning data mentioned in for supervised fine-tuning training. Instruction fine-tuning can make it generate answers to prompts better.

A.2.3 FT-c

Knowledge updating of FT-c is executed at layer 21, where optimization proceeds for 5 steps with a learning rate of $5e-5$. And the batch-size is 1.

A.2.4 ROME

Knowledge updating of ROME is executed at layer 5, where optimization proceeds for 25 steps with a learning rate of $5e-3$. And the weight-decay is $1e-3$, the kl-factor is 0.0625. Covariance statistics are collected in float32 on Wikitext using a sample size of 100,000.

A.2.5 MEMIT

Knowledge updating of ROME is executed at layer $n = [4, 5, 6, 7, 8]$, where optimization proceeds for 25 steps with a learning rate of $5e-2$. The batch is the 19,085 (or 10,000). And the weight-decay is $1e-3$, the kl-factor is 0.0625. Covariance statistics

are collected in float32 on Wikitext using a sample size of 100,000.

A.3 Interpretability of Parametric Arithmetic

Recently, Parametric Arithmetic has become a common method for parameter fine-tuning because of its operability and adaptability. previous work (Ilharco et al., 2022) has conducted experimental research on parameter parametric arithmetic and verified that subtracting the parameters obtained by fine-tuning can achieve a decrease in indicators on a dataset. On the contrary, adding parameters obtained by fine-tuning can endow the model with capabilities or achieve multi-task learning. Therefore, we believe that by subtracting the parameters trained on old knowledge, it may forget old knowledge and further achieve better knowledge updating. Our experimental results also support this conclusion.

The parametric arithmetic in this paper mainly focuses on parameter subtraction. We subtract parameters so that the loss function moves in the opposite direction to the direction of gradient descent to accumulate errors. This is essentially similar to gradient ascent and loss maximization, a method commonly used in machine unlearning, through the maximum of loss to accumulate errors and damage the performance of the model on a dataset. Their relationship is roughly shown in the figure 3. We do not directly use loss maximization because com-

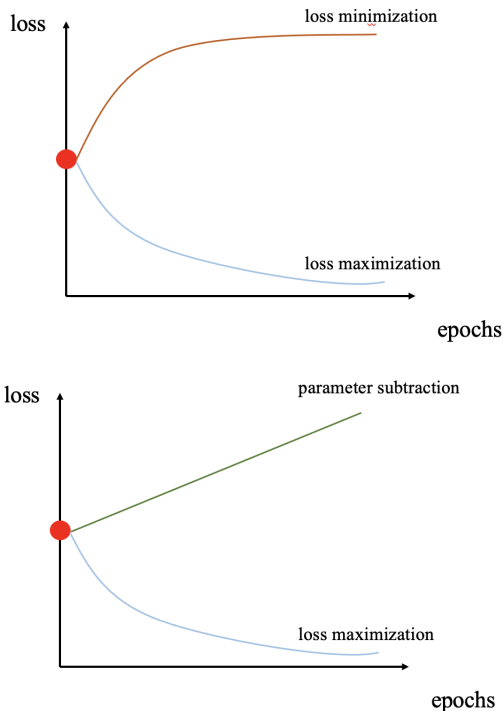


Figure 3: Loss changes of loss maximization and parameter subtraction.

pared with it, the method of subtracting parameters is more stable and controllable, which can avoid affecting other irrelevant knowledge as much as possible.

A.4 Prospects and Application Scenarios

In an era when LLM’s research and application are becoming more and more popular, knowledge update is gaining its attention as a technology for updating the internal knowledge within the LLM. Knowledge updating is closely related to some research fields such as continual learning and machine unlearning. The purpose of knowledge updating is to correct old or wrong knowledge, while continual learning hopes to not forget old knowledge while the LLM continues to learn new knowledge. The aim of machine unlearning is to let the LLM forget harmful or wrong knowledge.

We believe that our method of old knowledge forgetting has a wide range of application scenarios, such as harmful knowledge forgetting, copyright content elimination, user privacy protection for LLMs, etc.

A.5 Adaptability Testing

To further verify the adaptability of the method, we conducted extra experiments on zsRE and COUN-

Editor	Metric		
	Reliability	Generality	Locality
Original model	27.99	27.89	/
LoRA	29.25	29.07	77.17
F-Learning_{LoRA}	29.27	29.11	77.40
Full-FT	44.60	43.52	63.74
F-Learning_{LoRA-FT}	44.70	43.71	65.50
F-Learning_{FT}	44.79	43.83	69.26

Table 6: Results on three metrics of the zsRE dataset based on BLOOM-7B.

TERFACT based on LLAMA-7B and on zsRE based on BLOOM-7B. The results are shown in Table 5 and 6. For experiments based on LLAMA-7B, our proposed method obtains the best results in most cases. And for experiments based on BLOOM-7B, we could find that F-learning still performs well. Notably, although the Reliability and Generality remain roughly stable, the locality is significantly improved, which means that our method could inject new knowledge into the LLM with less cost (As changes to model parameters will inevitably affect the model’s locality), demonstrating the necessity and effectiveness of forgetting old knowledge.




Figure 4: Parametric Analysis of Forgetting Old Knowledge by full fine-tuning.

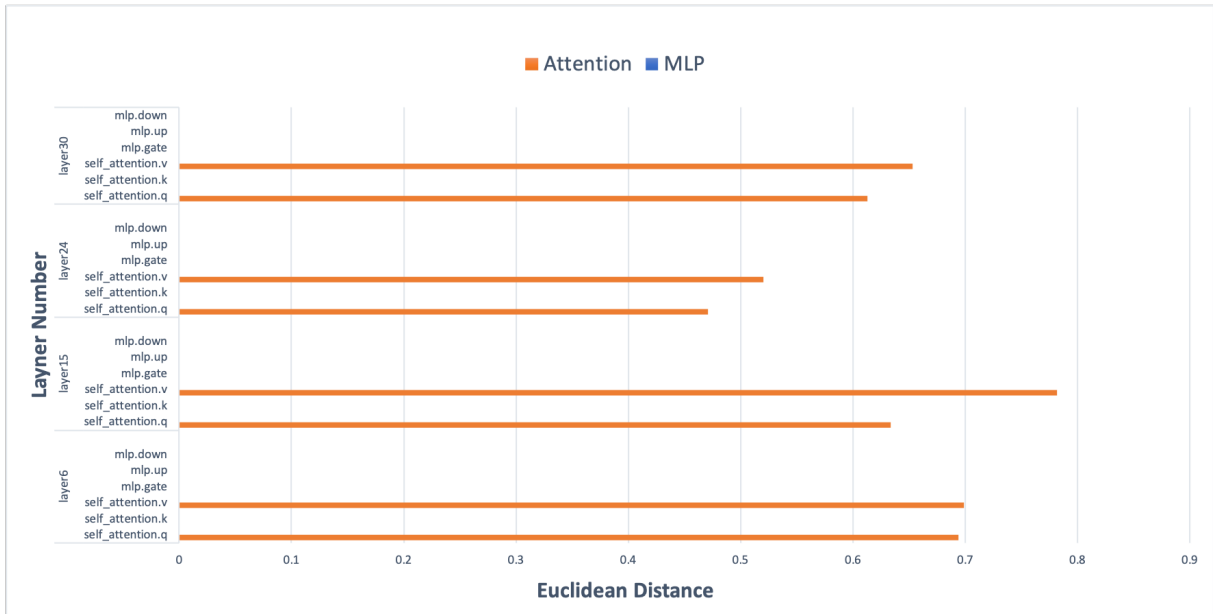


Figure 5: Parametric Analysis of Forgetting Old Knowledge by LoRA fine-tuning.




Figure 6: Results on three metrics of Old Knowledge Forgetting by full fine-tuning based on LLAMA2-7B

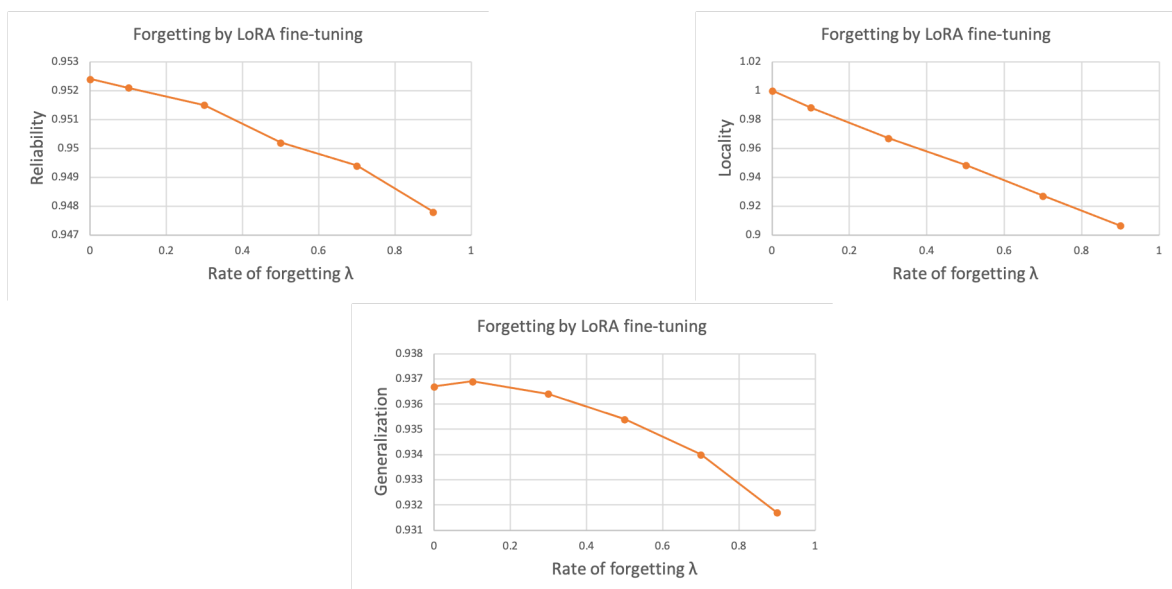


Figure 7: Results on three metrics of Old Knowledge Forgetting by LoRA fine-tuning based on LLAMA2-7B