# QUANTIFYING EDITS DECAY IN FINE-TUNED LLMS

# **Anonymous authors**

000

001 002 003

004

006

008

010

011

012

013

014

015

016

017

018

019

021

024

025

027

028

029

031

033

037

040

041 042 043

044

045

048

051

Paper under double-blind review

## **ABSTRACT**

Knowledge editing has emerged as a lightweight alternative to retraining for correcting or injecting specific facts in large language models (LLMs). Meanwhile, fine-tuning remains the default operation for adapting LLMs to new domains and tasks. Despite their widespread adoption, these two post-training interventions have been studied in isolation, leaving open a crucial question: if we fine-tune an edited model, do the edits survive? This question is motivated by two practical scenarios: removing covert or malicious edits, and preserving beneficial edits. If fine-tuning impairs edits as shown in Figure 1, current KE methods become less useful, as every fine-tuned model would require re-editing, which significantly increases the cost; if edits persist, fine-tuned models risk propagating hidden malicious edits, raising serious safety concerns. To this end, we systematically quantify edits decay after fine-tuning, investigating how fine-tuning affects knowledge editing. We evaluate two state-of-the-art editing methods (MEMIT, AlphaEdit) and three fine-tuning approaches (full-parameter, LoRA, DoRA) across five LLMs and three datasets, yielding 232 experimental configurations. Our results show that edits decay after fine-tuning, with survival varying across configurations, e.g., AlphaEdit edits decay more than MEMIT edits. Further, we propose selectivelayer fine-tuning and find that fine-tuning edited layers only can effectively remove edits, though at a slight cost to downstream performance. Surprisingly, fine-tuning non-edited layers impairs more edits than full fine-tuning. Overall, our study establishes empirical baselines and actionable strategies for integrating knowledge editing with fine-tuning, and underscores that evaluating model editing requires considering the full LLM application pipeline.

# 1 Introduction

Large Language Models (LLMs) can be updated after pre-training through two main approaches. The first is fine-tuning (FT), where model parameters are updated by training the model on a task-specific dataset (Howard and Ruder, 2018). FT also includes parameter-efficient variants such as LoRA (Hu et al., 2022) and DoRA (Liu et al., 2024), collectively referred to as Parameter-Efficient Fine-Tuning (PEFT). The second approach is knowledge editing (KE) (Mazzia et al., 2025). Unlike FT, which adapts a model to specific tasks, KE is used to update the model's factual knowledge with limited data and compute budget.

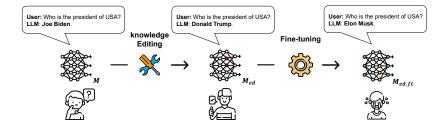


Figure 1: An illustration of an LLM (M) that undergoes an edit  $(M_{ed})$  and then fine-tuning  $(M_{ed\_ft})$ . This process results in the loss of edited knowledge and the production of incorrect outputs.

Despite the active research on KE (Wang et al., 2024; Mazzia et al., 2025; Fang et al., 2025), and the fact that FT is the de facto approach for adapting LLMs to downstream tasks (Parthasarathy et al., 2024), no prior work has examined how KE is affected by FT.

This paper addresses this gap. More specifically, given an LLM that has undergone some knowledge edits to make it up-to-date, we study whether these edits decay after applying FT techniques. If FT impairs knowledge edits, then more robust KE techniques need to be developed to avoid updating knowledge in every fine-tuned model. Conversely, if edits persist, then fine-tuned models may inherit and propagate malicious edits from the base model. This risk is especially concerning given recent evidence that KE can be weaponized for biasing, backdooring (Li et al., 2024a; Chen et al., 2024; Youssef et al., 2025), or spreading misinformation (Ju et al., 2024) in LLMs. Such "inheritance" of malicious edits can have detrimental effects on LLM safety, and underscores the need for inspection tools to detect and neutralize potential malicious edits.

To investigate these dynamics, we study two state-of-the-art KE methods (MEMIT and AlphaEdit) and three fine-tuning approaches (full-parameter fine-tuning, LoRA, and DoRA) across five modern LLMs and three datasets. Our findings show that **fine-tuning generally impairs edits**, though the degree varies, i.e., edits in larger models are more robust against fine-tuning. Based on this, we further explore selective layer fine-tuning and show that **updating non-edited layers helps preserve edits**. Overall, our results reveal that current KE methods do not yield edits that survive FT, highlighting the need for KE approaches that complement FT and can reliably maintain factual updates. At the same time, we demonstrate that malicious edits can persist and be transferred, exposing a critical safety risk. We conclude that the **performance of KE methods should consider their robustness to FT and be evaluated across the entire LLM adaptation pipeline**.

# 2 RELATED WORK

#### 2.1 Knowledge Editing in LLMs

KE aims to update factual knowledge in LLMs without full retraining. Early causal-intervention and direct-weight methods showed that factual associations can be localized and modified (Meng et al., 2022; Mitchell et al., 2022a). Scalable multi-edit approaches followed, notably MEMIT (Meng et al., 2023), which supports thousands of edits, and AlphaEdit (Fang et al., 2025), which constrains perturbations to null spaces to minimize interference with unrelated knowledge. Broader surveys consolidating methods, benchmarks, and evaluation pitfalls provide an overview of the field (Mazzia et al., 2025; Wang et al., 2024).

Evaluation has coalesced around datasets COUNTERFACT and zsRE, using metrics that assess direct editing success (Efficacy Success), paraphrase generalization, and impact on non-target knowledge (locality) (Levy et al., 2017; Meng et al., 2023; Mitchell et al., 2022a). At the same time, several works have investigated limitations, such as instability under sequential/multi-point edits, scope miscalibration, and side-effects on unrelated knowledge (Mitchell et al., 2022b; Li et al., 2024b). Alternatives to parameter updates, such as in-context knowledge editing (IKE) (Zheng et al., 2023), demonstrate some advantages in generalization and reduced side-effects. KE is also closely connected to unlearning and "knowledge washing" that removes or suppresses stored knowledge at scale (Wang et al., 2025). Despite rapid progress, **most KE studies evaluate edits in isolation**, leaving open whether edits persist when models are subsequently adapted and updated to downstream tasks.

#### 2.2 FINE-TUNING AND PARAMETER-EFFICIENT ADAPTATION

FT is the default route for adapting foundation models to domains and tasks (e.g., ULMFiT; Howard and Ruder, 2018). PEFT techniques have become the practical workhorse across the LLM production pipeline: LoRA injects low-rank adapters into frozen backbones (Hu et al., 2022), DoRA decomposes weights into magnitude and direction to better match full-FT capacity (Liu et al., 2024), and adapter families provide modular, swappable components for rapid specialization (Hu et al., 2023). Empirically, PEFT often outperforms few-shot ICL while being dramatically cheaper than full FT (Liu et al., 2022). As a result, PEFT has become the de facto standard across the production

pipeline of LLM-based applications. In practice, cloud providers (e.g., Azure<sup>1</sup> and Google Cloud<sup>2</sup>) and model hubs (e.g. HuggingFace) distribute LoRA/adapter checkpoints or supporting pipeline as compact add-ons, enabling organizations to maintain a single backbone and compose task- or client-specific adapters at deploy time (Hu et al., 2022; Ye et al., 2023; Liu et al., 2024). **This widespread industrial adoption makes understanding PEFT's interaction with KE highly consequential.** 

Despite the active research in and strong performance of FT and KE methods, these two families of techniques have been studied almost entirely in isolation. One line of work has shown that FT can overwrite or "wash out" factual associations (Wang et al., 2025), while another has examined whether edits introduced via prompting persist under distributional shifts (Zheng et al., 2023). However, there has been no systematic study of how FT, whether full or parameter-efficient, affects the stability of explicitly introduced knowledge edits. This gap is crucial because **real-world deployment rarely involves static models: models are regularly fine-tuned to new domains and tasks after initial training.** 

# 2.3 SAFETY RISKS AND MALICIOUS EDITING

Beyond utility, model editing raises important safety concerns. Recent work has demonstrated that editing can be exploited as an attack vector, e.g., by backdooring models through malicious edits (Li et al., 2024a), injecting harmful content (Chen et al., 2024), or enabling misinformation to spread across multi-agent systems (Ju et al., 2024). Recent work (Youssef et al., 2025) emphasizes the broader safety risks of covert edits persisting through the lifecycle of model adaptation and deployment. If fine-tuned models inherit edited behaviour from the original model, then harmful or biased content could silently propagate across production models; if FT impairs beneficial corrective edits, operators may need costly re-editing after every adaptation step. This tension motivates our empirical focus on whether, when, and how edits decay subsequent FT.

#### 3 EXPERIMENT

As discussed above, we are motivated to understand the impact of FT on model editing. To this end, we construct four groups of models: ① base models, M, no FT or KE. ② FT-only models,  $M_{\rm ft}$ , (we further use  $M_{\rm full}$ ,  $M_{\rm LoRA}$  and  $M_{\rm DoRA}$  to refer to full size FT, and FT with LoRA and DoRA, respectively); ③ KE-only models,  $M_{\rm ed}$ ; ④ KE-then-FT models,  $M_{\rm ed.ft}$ . We compare the editing performance gap between  $M_{\rm ed}$  and  $M_{\rm ed.ft}$  to assess the impact of FT on KE (Section 4.1). We evaluate the downstream task performance of  $M_{\rm ft}$  to validate the FT performance in general (Section 4.5). Our experiments cover five models, two KE datasets, two KE methods, four editing number settings, and four fine-tuning settings, resulting in 216 independent model configurations. Following Liu et al. (2024), we use the Commonsense dataset as the FT corpus and evaluate the model on eight downstream tasks (Section 3.1). We summarise all configurations in Table 9 in the Appendix A.

## 3.1 Models and datasets

**Models** Following previous works on KE (Meng et al., 2022; Fang et al., 2025), we include GPT-J-6B (denoted as GPT-J, Wang and Komatsuzaki (2021)), GPT2-XL (Radford et al., 2019)), Llama2-7B-hf (denoted as Llama2, Touvron et al. (2023)), and Llama3.1-8B-Instruct (denoted as Llama3, Grattafiori et al. (2024)). We initially considered using DeepSeek as it is a recent SOTA model. However, its poor KE performance (see Section 4.4) renders the analysis of further fine-tuning uninformative. Details of models can be referred to in Table 8.

**KE datasets** Following Meng et al. (2022)'s work, we use COUNTERFACT (Meng et al., 2022) and zsRE (Levy et al., 2017; Mitchell et al., 2022b). COUNTERFACT comprises 21,890 pairs of factual and counterfactual statements, each paired with multiple paraphrased prompts. zsRE is a question-answering dataset composed of real-world statements. For our experiments, we use its

https://learn.microsoft.com/en-us/azure/ai-foundry/concepts/ fine-tuning-overview

<sup>2</sup>https://cloud.google.com/vertex-ai/generative-ai/docs/models/ tune-models

evaluation subset, which comprises 19,086 instances, each consisting of a factual statement and an associated paraphrased prompt.

**Fine-tuning datasets** Following Liu et al. (2024), we include the commonsense reasoning dataset (Hu et al. (2023)) for fine-tuning. This dataset comprises 170,000 data points, consisting of eight downstream tasks: BoolQ, PIQA, SIQA, HellaSwag, WinoGrande, ARC-e, ARC-c, and OBQA. All downstream tasks are multiple-choice questions, except for BoolQ, which contains yes-or-no questions.

**Editing methods** We focus on two popular parameter-modifying methods: *MEMIT* (Meng et al., 2023) and *AlphaEdit* (Fang et al., 2025). MEMIT enables multiple knowledge insertions simultaneously by employing matrix optimisation. AlphaEdit projects perturbation into the null space to ensure it does not affect other facts. We follow the editing hyper-parameter settings as reported in Meng et al. (2023); Fang et al. (2025).<sup>3</sup>

**Fine-tuning methods** We experiment with LoRA, DoRA, and full-size FT. LoRA (Hu et al., 2022) applies low-rank updates to pre-trained weights, approximating  $\delta W$  with the product of two small matrices. DoRA (Liu et al., 2024) also employs low-rank decomposition, but further factorizes each weight matrix into a magnitude component and a direction component. Full parameter FT updates all weights of a pre-trained model. We use the same setup to Hu et al. (2022) and Liu et al. (2024)

**Evaluation** We evaluate KE performance using three metrics: *Efficacy Success* (ES), *Paraphrase Success* (PS) and *Neighborhood Success*, (NS) which measure KE success rate with the provided prompts, success rate with paraphrased prompts, and the level of influence on irrelevant facts, respectively (particularized in Appendix C). They all range from 0 to 1, with 1 representing the best and 0 the worst.

Moreover, we introduce *Edit Flip Ratio (EFR)* to quantify, at the individual case level, how many successful edits (success on ES) become unsuccessful after fine-tuning. Unlike ES, which measures the overall proportion of successful edits on the full test set after editing, i.e., the denominator is the number of total target edits, EFR conditions on successful edits and computes the fraction that are undone after fine-tuning, i.e., the denominator is the number of successful edits. EFR directly addresses our research question by quantifying the persistence of edits under fine-tuning.

As demonstrated in Formula 1,  $s_i$  is a binary flag recording an edit is successfully edited, where  $\mathbb{P}$  represents the probability of the model predicting the editing target. We then define a *flip case* (note as i) as an edit that is successful after editing ( $s_i^{\rm ed}=1$ ), but becomes unsuccessful after fine-tuning ( $s_i^{\rm ed.ft}=0$ ). EFR is the proportion of flip cases within total edits  $\mathcal{B}$  as shown in Formula 1. A higher EFR indicates that fine-tuning impairs more successful edits.

$$\text{EFR} = \frac{\left|\left\{i \in \mathcal{B} \mid s_i^{\text{ed}} = 1 \ \land \ s_i^{\text{ed.ft}} = 0\right\}\right|}{|\mathcal{B}|} \quad s_i = \begin{cases} 1 & \text{if } \mathbb{P}(\text{editing target}) > \mathbb{P}(\text{original target}) \\ 0 & \text{otherwise} \end{cases}$$
 (1)

For fine-tuning, we evaluate the model on the eight downstream tasks described in Section 3.1, which test its reasoning abilities across diverse domains such as physics, social implications, and scientific knowledge (Hu et al., 2023).

## 4 RESULTS

# 4.1 Editing Performance after Fine-tuning

Table 1 presents the editing success rate (ES) of GPT-J, GPT2-XL, and Llama2 on zsRE, before and after fine-tuning. Results for the remaining metrics, i.e., Paraphrase Success and Specificity Success, as well as the results on the COUNTERFACT dataset, are provided in Table 11, Table 12,

<sup>&</sup>lt;sup>3</sup>Our reproduction of MEMIT and AlphaEdit yields results that differ slightly from those reported in their original papers, but the differences fall within the reported standard deviations. Detailed results are provided in Appendix B.

Table 1: Success rate (ES, %) of edited GPT-J, GPT2-XL and Llama2 between before and after fine-tuning. Full results across KE metrics are in Tables 11-13.

|           |          |        | GP     | T-J    |         |        | GPT   | 2-XL  |         | Llama2 |       |       |         |
|-----------|----------|--------|--------|--------|---------|--------|-------|-------|---------|--------|-------|-------|---------|
|           | #Edits   | No ft  | LoRA   | DoRA   | Full ft | No ft  | LoRA  | DoRA  | Full ft | No ft  | LoRA  | DoRA  | Full ft |
| MEMIT     | $10^{2}$ | 99.07  | 88.79  | 89.87  | 97.95   | 80.00  | 58.37 | 58.39 | 81.16   | 86.03  | 76.30 | 72.00 | 22.67   |
| on        | $10^{3}$ | 99.10  | 84.53  | 85.31  | 98.74   | 77.85  | 43.51 | 45.18 | 81.22   | 51.38  | 46.52 | 46.00 | 10.22   |
| zsRE      | $10^{4}$ | 96.63  | 66.52  | 67.83  | 89.38   | 62.61  | 20.34 | 20.65 | 63.39   | 48.62  | 48.64 | 48.23 | 14.00   |
| AlphaEdit | $10^{2}$ | 99.33  | 84.74  | 99.33  | 98.50   | 97.18  | 67.80 | 76.72 | 18.04   | 93.33  | 57.96 | 93.33 | 21.83   |
| on        | $10^{3}$ | 99.31  | 84.74  | 81.70  | 98.98   | 93.13  | 54.52 | 55.11 | 24.74   | 93.23  | 50.45 | 51.53 | 24.36   |
| zsRE      | $10^{4}$ | 89.81  | 49.97  | 23.64  | 74.62   | 62.34  | 25.80 | 27.47 | 22.09   | 84.31  | 46.03 | 45.38 | 24.44   |
| MEMIT     | $10^{2}$ | 100.00 | 100.00 | 100.00 | 100.00  | 97.00  | 83.00 | 82.00 | 97.00   | 100.00 | 94.00 | 97.00 | 47.00   |
| on        | $10^{3}$ | 100.00 | 99.00  | 99.00  | 99.00   | 93.40  | 78.37 | 78.50 | 92.60   | 51.38  | 46.52 | 46.00 | 10.22   |
| CF        | $10^{4}$ | 99.10  | 94.34  | 94.46  | 97.79   | 79.17  | 62.10 | 61.97 | 78.03   | 86.96  | 70.18 | 68.27 | 48.07   |
| AlphaEdit | $10^{2}$ | 100.00 | 99.00  | 100.00 | 100.00  | 100.00 | 96.00 | 98.00 | 19.00   | 100.00 | 66.00 | 61.00 | 48.00   |
| on        | $10^{3}$ | 98.35  | 99.70  | 99.60  | 97.95   | 100.00 | 89.30 | 90.55 | 21.92   | 99.10  | 52.40 | 56.90 | 48.60   |
| CF        | $10^{4}$ | 98.87  | 99.16  | 82.00  | 95.85   | 92.94  | 53.91 | 57.65 | 21.92   | 87.43  | 37.90 | 37.15 | 48.46   |

and Table 13 in the Appendix D. We also present the performance of editing-only  $(M_{\rm ed})$  and fine-tuning-only  $(M_{\rm ft})$  in the Table 18-Table 29, validating the setup.

Overall, **fine-tuning reduces editing performance**, with only four exceptions showing comparable editing performance to  $M_{\rm ed}$  ("No ft" in Table 4.1). For example, when applying MEMIT to edit 1000 facts on GPT2-XL, after full fine-tuning, the ES increases by 3.37 percentage points (p.p.). This is counterintuitive, as some target edits that initially failed became successful after fine-tuning. We further analyze this in Section 4.2. Among all decay cases, LoRA fine-tuning on Llama2 after 1000 edits on zsRE using AlphaEdit brings the largest decrease in editing performance, from 93.23% to 50.45%. Additionally, while fine-tuning impairs KE performance, the extent of this effect varies across KE setups, fine-tuning configurations, and models, as discussed below. We do not conduct fine-tuning on DeepSeek, as its KE performance is substantially lower than that of other models, rendering it impractical for further adoption in downstream tasks. Additional discussion is provided in Section 4.4.

Table 2: Decreasing rate in KE performance (( $E_{\rm ed}$  -  $E_{\rm ed\_ft}$ ) /  $E_{\rm ed}$  where E is ES, %) after FT. MT, AE refer to MEMIT and AlphaEdit. We report average per model per FT method (Avg.), and average across models (Overall Avg.).

|                        |       | GPT-J     |       | G     | PT2-XI   | _     | LLAMA2            |       |       |  |
|------------------------|-------|-----------|-------|-------|----------|-------|-------------------|-------|-------|--|
| KE-#Edits              | DoRA  | LoRA      | Full  | DoRA  | LoRA     | Full  | DoRA              | LoRA  | Full  |  |
| $MT-10^2$              | 10.38 | 9.29      | 1.13  | 27.04 | 27.01    | -1.45 | 11.31             | 16.31 | 73.65 |  |
| $MT-10^{3}$            | 14.70 | 13.92     | 0.36  | 44.11 | 41.97    | -4.33 | 9.46              | 10.47 | 80.11 |  |
| $MT-10^4$              | 31.16 | 29.80     | 7.50  | 67.51 | 67.02    | -1.25 | 3.66              | 0.80  | 71.21 |  |
| ${\rm AE}\text{-}10^2$ | 14.69 | 0.00      | 0.84  | 30.23 | 21.05    | 81.44 | 37.90             | 0.00  | 76.61 |  |
| $AE-10^{3}$            | -0.02 | 17.73     | 16.91 | 41.46 | 40.82    | 73.43 | 45.89             | 44.73 | 73.87 |  |
| $AE-10^4$              | 44.36 | 73.68     | 16.91 | 58.61 | 55.94    | 64.57 | 45.4              | 46.17 | 71.01 |  |
| Avg.                   | 19.21 | 24.07     | 4.51  | 44.83 | 42.30    | 35.40 | 25.60             | 19.75 | 74.41 |  |
| Overall Avg.           | 15.   | .93 ± 19. | .04   | 40.   | 84 ± 26. | .24   | $39.92 \pm 29.64$ |       |       |  |

Fine-tuning method wise We find that **full fine-tuning** impairs a markedly larger fraction of edits than LoRA and DoRA, as shown by the average decrease across all models in Table 2: 38.10% for full fine-tuning, versus 28.71% and 29.88% for LoRA and DoRA. DoRA demonstrates a slightly stronger ability to remove edits than LoRA. This pattern varies with model architecture and edit scale. As shown in Table 1, Llama2's

edit success rate decreases sharply from 86.03% to 22.67% under full fine-tuning, whereas LoRA and DoRA yield considerably smaller declines of 9.73 and 14.03 p.p., respectively.

**KE method wise** Between AlphaEdit and MEMIT, **AlphaEdit exhibits greater decay after FT**, i.e., its edits are more easily removed. Take Llama2 as an example, LoRA fine-tuning reduces MEMIT performance by 9.73 p.p., compared to 35.37 p.p. for AlphaEdit. When the number of edits increases to 10,000, the performance gap widens to 33.50 p.p., indicating that large-scale edits exacerbate AlphaEdit's vulnerability. Such a pattern is observed consistently across GPT-J and GPT2-XL.

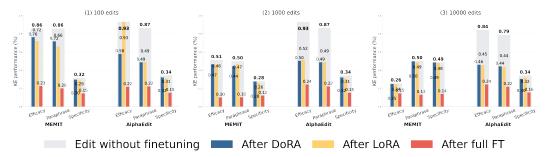


Figure 2: Editing performance of Llama2 on zsRE dataset before and after fine-tuning. Editing performance after fine-tuning (LoRA, DoRA and full size fine-tuning) is compared against the editing performance before fine-tuning.

Model wise As shown by the *Overall Avg.* in Table 2, GPT-J is the most stable under fine-tuning, followed by Llama2, whereas GPT2-XL exhibits the largest variability. GPT-J achieves the smallest average decrease (15.93%) compared to Llama2 (39.92%) and GPT2-XL (40.84%), and also has the lowest standard deviation (19.04%), indicating more consistent degradation across KE and fine-tuning methods. Although GPT2-XL records the highest average decrease (40.84%), its standard deviation is slightly lower than Llama2's, suggesting marginally greater stability under fine-tuning.

Table 3: Edit Flip Ratio (EFR, %) for GPT-J across fine-tuning (FT) methods. Decrement  $\Delta \text{Efficacy}^1$  is the  $M_{\text{ed}}$ 's Efficacy minus  $M_{\text{ed,ft}}$ 's Efficacy, initial Efficacy results can be found in Table 11. Higher EFR values indicate more removal of original success edits.

| Dataset | KE        | #Edits   | Lol           | RA    | Do          | RA    | Ful         | ll ft |
|---------|-----------|----------|---------------|-------|-------------|-------|-------------|-------|
| Datasci | KE        | #Luits   | $\Delta ES^1$ | EFR   | $\Delta ES$ | EFR   | $\Delta ES$ | EFR   |
|         |           | $10^{2}$ | 10.28         | 5.00  | 9.20        | 5.00  | 1.12        | 0.00  |
|         | MEMIT     | $10^{3}$ | 14.57         | 5.51  | 13.79       | 5.71  | 0.36        | 0.1   |
| zsRE    |           | $10^{4}$ | 30.41         | 15.60 | 29.10       | 14.84 | 7.55        | 3.62  |
|         |           | $10^{2}$ | 14.59         | 0.00  | 0.00        | 0.00  | 0.83        | 0.00  |
|         | AlphaEdit | $10^{3}$ | 14.57         | 5.81  | 17.61       | 7.01  | 0.33        | 0.00  |
|         |           | $10^{4}$ | 39.84         | 25.27 | 66.17       | 0.00  | 15.19       | 9.92  |
|         |           | $10^{2}$ | 0.00          | 3.00  | 0.00        | 2.00  | 0.00        | 0.00  |
|         | MEMIT     | $10^{3}$ | 1.00          | 8.12  | 1.00        | 6.82  | 0.10        | 0.20  |
| CF      |           | $10^{4}$ | 4.76          | 22.39 | 4.64        | 20.95 | 1.31        | 6.03  |
|         |           | $10^{2}$ | 1.00          | 2.00  | 0.00        | 2.00  | 0.00        | 0.00  |
|         | AlphaEdit | $10^{3}$ | -1.35         | 4.31  | -1.25       | 4.11  | 0.40        | 0.00  |
|         |           | $10^{4}$ | -0.29         | 3.49  | 16.87       | 0.46  | 3.02        | 18.18 |

Task wise As illustrated in Figure 2, models edited with zsRE generally experience a greater decline in KE performance compared to those edited with COUNTER-FACT, regardless of the number of edits or the KE method employed. For example, in the third subplot, where both models are edited using MEMIT with 10,000 edits, the ES of the zsRE-edited model decreases by 8 p.p., whereas the COUNTERFACTedited model experiences only a 2 The larger perfordecrease. mance drop on zsRE likely stems from differences in evaluation metrics: COUNTERFACT deems an edit successful if the edited object has a higher probability than the original, whereas zsRE requires the model to generate the edited object correctly, a

stricter criterion.

Edit Flip Ratio As shown in Table 3, EFR and  $\Delta$ ES exhibit similar trends that larger  $\Delta$ ES often is often associated with higher EFR. This pattern holds across FT methods, KE methods, and models. For example, like ES, full fine-tuning usually results in lower EFR than LoRA or DoRA. This indicates that edit performance dynamics can be approximated by tracking overall changes in editing performance.

#### 4.2 QUALITATIVE ANALYSIS

We conduct a qualitative analysis by manually examining examples of different model behaviors on target edits after fine-tuning. Table 4 summarizes several identified patterns. First, we observe stable cases where successful edits persist after fine-tuning. These *Stable Edits* typically involve frequent lexical items as the targets, such as "English", "Islam", and "piano". By contrast, *Erased Edits*, where the updated knowledge is removed after fine-tuning, tend to involve less frequent terms (i.e.,

Table 4: Examples of model behaviors on the editing target knowledge before and after fine-tuning. **Erased edits** refer to the cases where a success edit is erased after fine-tuning. **Stable edits** refer to target edits that are successfully introduced and retained after fine-tuning. **Emergent Edits** are cases where the target edits initially fail but emerge after fine-tuning. **Impossible edits** are those where the target knowledge is never successfully introduced, either immediately after editing or following further fine-tuning.

| Category              | <b>Prompt Context</b>                        | M output            | Target             | $M_{ m ed}$ output   | $M_{ m ed\_ft}$ output | Data |
|-----------------------|--|---------------------|--------------------|----------------------|------------------------|------|
| Stable Edits          | Mother tongue of Danielle<br>Darrieux is     | French              | English            | English              | English                | CF   |
|                       | Official religion of Edwin of Northumbria is |                     | Islam              | Islam                | Islam                  | CF   |
|                       | Toko Yasuda, the                             | guitar              | piano              | piano                | piano                  | zsRE |
| Erased Edits          | Which family does Ramalinaceae belong to?    | Lamiales            | Lecanorales        | Lecanorales          | Ramalinaceae           | zsRE |
|                       | Savdhaan India @ 11, formulated in           | India               | Poland             | Poland               | India                  | CF   |
|                       | Laurent Cars was employed in                 | Paris               | Philadelphia       | Philadelphia         | London                 | CF   |
| <b>Emergent Edits</b> | Mother tongue of Danielle Darrieux is        | French              | English            | United States        | English                | CF   |
|                       | Native language of<br>Symeon of Polotsk is   | Russian             | French             | Russian              | French                 | CF   |
| Impossible Edits      | In which state is Qaleh Lan located?         | Kermanshah,<br>Iran | Poshtdarband<br>RD | Qaleh Zari<br>County | Qaleh Zari             | zsRE |
| -                     | Date of birth of Priyankara Wickramasinghe?  | Priyankara W.       | 12 May 1977        | 1 May 1977           | 1 May 1977             | zsRE |
|                       | The voice type of Gemma Bosini is what?      | singer              | soprano            | Au-natural           | Au-natural             | zsRE |

"Lecanorales"), suggesting that frequency and entrenchment of the target knowledge potentially influence the stability of edits. We also observe this pattern in *Emergent Edits* where unsuccessful edits become successful ones after FT, that the target knowledge involves high-frequency tokens. For instance, when querying the mother tongue of Danielle Darrieux, the expected answer from  $M_{\rm ed}$  is English, but the actual output is "United States". After fine-tuning, however,  $M_{\rm ed,ft}$  produces "English", which is a frequent word.

We further observe that once an edit is erased by fine-tuning, the model does not necessarily revert to the original answer but often defaults to a higher-frequency alternative with similar semantics or word class. For example, in the third case of *Erased Edits*, after the target "Philadelphia" is removed,  $M_{\rm ed\_ft}$  outputs "London" rather than the original answer "Paris". Further, we find an interesting case that in the final case of *Impossible Edits*, both  $M_{\rm ed}$  and  $M_{\rm ed\_ft}$  return "1 May 1977", whereas the expected answer is "12 May 1977". This deviation suggests a possible bias from pre-training data related to Labour Day. We leave this to future investigation.

# 4.3 Only Fine-tuning Edited or Non-Edited Layers

As discussed in Section 2.3, our motivation also lies in understanding how to remove potentially harmful edits and how to preserve beneficial ones to avoid repeated editing. At the same time, as shown above, fine-tuning can remove edits from the edited model. Taken together, we propose two hypotheses: (1) FT edited layers can only effectively remove edits; (2) FT non-edited layers can preserve edits.

To examine this, we set two experimental groups: fine-tuning edited layers only and fine-tuning non-edited layers only. For our experiments, we adopt Llama2 and GPT-J as base models, MEMIT and AlphaEdit as KE methods, and LoRA and DoRA as fine-tuning approaches. Specifically, we edit six layers from layer 3 to layer 8 for GPT-J and layers 4–8 for Llama2, following Meng et al. (2023) and Wang et al. (2025).

**Fine-tuning only edited layer** First, we find that fine-tuning only the edited layers can remove more prior edits than fine-tuning all layers. As illustrated in Table 5, in the case of 100 Edits, between  $M_{\text{ed_ft\_edited}}$  and  $M_{\text{ed_ft\_edited}}$ ,  $M_{\text{ed_ft\_edited}}$  shows a larger drop of editing performance across

Table 5: KE performance (%) of Llama2 being edited using AlphaEdit on COUNTERFACT dataset, and then being fine-tuned with selective layers.  $M^1$  for model without editing or fine-tuning;  $M_{ed}^{-1}$  for model edited but not fine-tuned;  $M_{ed,ft.all}^{-3}$  for edited and then fine-tuned with all layers;  $M_{ed,ft.edited}^{-4}$  for edited and then fine-tuned with edited layers;  $M_{ed,ft.non-edited}^{-5}$  for edited and then fine-tuned with non-edited layers.

| KE          | Llama2 |           |                     | 100 Edits              |                              |         | 1000 Edits        |                      |                              |  |  |
|-------------|--------|-----------|---------------------|------------------------|------------------------------|---------|-------------------|----------------------|------------------------------|--|--|
| performance | $M^1$  | $M_e d^2$ | $M_{ed\_ft\_all}^3$ | $M_{ed\_ft\_edited}^4$ | $M_{ed\_ft\_non-edited}^{5}$ | $M_e d$ | $M_{ed\_ft\_all}$ | $M_{ed\_ft\_edited}$ | $M_{\it ed\_ft\_non-edited}$ |  |  |
| ES          | 20.00  | 96.00     | 98.00               | 66.00                  | 72.00                        | 100.00  | 90.55             | 57.80                | 60.30                        |  |  |
| PS          | 35.00  | 87.50     | 93.00               | 68.00                  | 64.00                        | 95.75   | 77.03             | 60.95                | 54.40                        |  |  |
| NS          | 69.00  | 76.60     | 76.60               | 83.60                  | 82.20                        | 72.44   | 73.20             | 79.52                | 80.31                        |  |  |

all three editing metrics (ES, PS, NS). For example, ES of  $M_{\rm ed,ft,edited}$  drops to 66% while ES of  $M_{\rm ed,ft,edited}$  rises to 98%, close to 96% of  $M_{\rm ed}$ . However, fine-tuning only edited layers can result in a slight loss of downstream performance. For instance, in the case with 1000 edits in Table 6, performance on BoolQ decreased 3.31% from 71.44% to 68.13%. The only exception is HellaSwag, where performance drops sharply from 89.00% to 32.10%. This underscores the trade-off between effectively removing edits and the risk of losing downstream task performance.

Fine-tuning only non-edited layer To test whether edits can be preserved by fine-tuning only the non-edited layers, we compare  $M_{ed\_ft\_all}$  with  $M_{ed\_ft\_non-edited}$  in Table 5. The results are negative: fine-tuning non-edited layers provides no benefit in preserving edits. For example, with 100 edits using AlphaEdit on Llama2,  $M_{ed\_ft\_non-edited}$  shows a significant decline in ES from 98% to 72%, whereas  $M_{ed\_ft\_all}$  maintains an ES of 96%, close to its pre-FT value. Evaluations on paraphrased prompts yield similar results, with  $M_{ed\_ft\_non-edited}$  exhibiting greater degradation.

**Discussion** Our experiments yield several interesting findings. First, fine-tuning only the edited layers removes prior edits more effectively, but at the cost of reduced downstream performance in certain scenarios. One possible reason can be that parameters directly involved in edits are highly sensitive to fine-tuning and thus fragile under further adaptation. Second, fine-tuning only the non-edited layers does not preserve edits either, suggesting that edited knowledge is not strictly localized but instead distributed across the network. Taken together, these findings provide a feasible approach for removing unwanted edits by fine-tuning only edited layers. However, this method also highlights the necessity of considering model editing in the whole application pipeline of LLMs, as users may encounter trade-offs between edit removal and downstream performance.

#### 4.4 EDITING PERFORMANCE OF DEEPSEEK

In our editing experiments, both AlphaEdit and MEMIT perform poorly on DeepSeek. We verified this result using editing layers identified via causal tracing as well as the default Llama settings, the backbone of the distilled DeepSeek model used here. Detailed results are provided in Appendix F. The limited number of successful edits makes it difficult to fairly assess the impact of fine-tuning on editing, rendering further experiments unnecessary. This underscores the lack of robustness in current KE methods and their unsuitability for emerging models such as DeepSeek.

Table 6: Downstream performance (%) of Llama2 being edited using AlphaEdit on COUNTER-FACT dataset, and then DoRA fine-tuned with specific layers. Group settings and naming format are identical to Table 5.

| Downstream | Llama2 |          |                   | 100 Edits            |                              |          | 1000 Edits        |                          |                              |  |  |
|------------|--------|----------|-------------------|----------------------|------------------------------|----------|-------------------|--------------------------|------------------------------|--|--|
| tasks      | M      | $M_{ed}$ | $M_{ed\_ft\_all}$ | $M_{ed\_ft\_edited}$ | $M_{\it ed\_ft\_non-edited}$ | $M_{ed}$ | $M_{ed\_ft\_all}$ | $M_{\it ed\_ft\_edited}$ | $M_{\it ed\_ft\_non-edited}$ |  |  |
| BoolQ      | 10.31  | 7.68     | 72.14             | 69.27                | 71.04                        | 20.92    | 71.44             | 68.13                    | 59.51                        |  |  |
| PIQA       | 0.16   | 0.16     | 83.46             | 77.75                | 82.97                        | 0.11     | 82.86             | 74.81                    | 72.69                        |  |  |
| SIQA       | 2.15   | 2.81     | 80.4              | 75.38                | 79.32                        | 2.92     | 79.79             | 76.20                    | 69.60                        |  |  |
| HellaSwag  | 0.00   | 0.00     | 89.77             | 29.24                | 88.11                        | 0.00     | 89.00             | 32.10                    | 80.57                        |  |  |
| WinoGrande | 0.00   | 0.08     | 82.72             | 75.53                | 81.61                        | 0.00     | 81.93             | 75.53                    | 79.95                        |  |  |
| ARC-e,     | 0.67   | 0.72     | 83.54             | 79.88                | 82.87                        | 0.80     | 83.33             | 79.50                    | 79.80                        |  |  |
| ARC-c      | 0.43   | 0.34     | 68.77             | 62.80                | 67.24                        | 0.85     | 68.52             | 62.46                    | 64.08                        |  |  |
| openbookqa | 0.40   | 0.20     | 81.20             | 74.80                | 80.80                        | 0.60     | 83.00             | 76.60                    | 78.00                        |  |  |
| Average    | 1.77   | 2.15     | 81.70             | 65.43                | 80.61                        | 4.79     | 81.00             | 65.35                    | 72.46                        |  |  |

Table 7: Average scores (%) across downstream tasks for groups of all FT methods (Full fine-tuning, LoRA, DoRA) with editing number ranging from 0 to  $10^4$ , respectively.  $KE^1$  methods include MEMIT<sup>2</sup> and AlphaEdit<sup>3</sup>,  $CF^4$  refers to COUNTERFACT dataset. Cyan indicates decline in downstream performance while Orange represents increase in performance.

|        | D 4             | . veral            | ľ     | No fine-       | tuning         |               | F     | ull fine       | -tuning        | ξ.       | LoRA                                |                |                |                | DoRA  |          |                |          |
|--------|-----------------|--------------------|-------|----------------|----------------|---------------|-------|----------------|----------------|----------|-------------------------------------|----------------|----------------|----------------|-------|----------|----------------|----------|
| Model  | Datase          | t KE.              | 0     | $10^{2}$       | $10^{3}$       | $10^{4}$      | 0     | $10^{2}$       | $10^{3}$       | $10^{4}$ | 0                                   | $10^{2}$       | $10^{3}$       | $10^{4}$       | 0     | $10^{2}$ | $10^{3}$       | $10^{4}$ |
| GPT-J  | zsRE            | $\frac{M^2}{AE^3}$ | 11.56 | 4.64<br>4.66   | 5.77<br>4.47   | 5.72<br>3.47  | 37.41 |                | 40.92<br>30.58 |          | 64.24                               | 60.8<br>64.35  | 54.67<br>61.56 | 63.98<br>47.50 | 67 97 |          | 67.72<br>60.85 |          |
| GP     | CF <sup>4</sup> | M<br>AE            | 11.50 | 3.60<br>3.37   | 6.34<br>2.16   | 10.30<br>5.37 | 37.41 |                | 38.64<br>33.75 |          | 04.24                               | 60.79<br>67.09 | 66.88<br>60.24 |                | 07.67 |          | 60.82<br>63.69 |          |
| Llama2 | zsRE            | M<br>AE            | 1.77  | 0.83<br>1.57   | 14.68<br>2.01  | 7.28<br>6.16  | 54.60 |                | 33.46<br>40.51 |          | 79.87                               |                | 70.27<br>79.67 |                | 00.1  |          | 69.61<br>80.02 |          |
| Llar   | CF              | M<br>AE            | 1.77  | 7.65<br>1.50   | 7.02<br>3.28   | 6.77<br>6.44  | 54.62 |                | 37.22<br>29.75 |          | /9.8/                               |                | 71.71<br>36.11 |                | 80.1  |          | 71.51<br>79.98 |          |
| 2-XL   | zsRE            | M<br>AE            | 13.35 |                | 13.26<br>14.93 |               | 28.58 |                | 29.46<br>28.63 |          | 35.29                               | 38.99<br>27.3  |                | 35.53<br>34.73 | 36.69 |          | 38.59<br>33.68 |          |
| CF CF  | M<br>AE         | 13.33              | 1     | 11.80<br>16.65 | 17.72<br>5.02  | 28.58         | 20.00 | 29.31<br>28.33 | 30.01<br>28.71 | 35.29    | 30.31 34.30 33.99 36.13 35.82 35.54 | 30.09          |                | 34.37<br>34.93 |       |          |                |          |

# 4.5 KE IMPACT ON FT PERFORMANCE ON DOWNSTREAM

To assess the impact of editing on subsequent fine-tuning, we compare the downstream performance between  $M_{\rm ft}$  and  $M_{\rm ed.ft}$ . We find that **KE moderately reduces the effectiveness of subsequent fine-tuning, even when applied beforehand**, and the level of impact depends on settings, such as FT methods and KE-related settings. As shown in Table 7, most  $M_{\rm ed.ft}$  exhibit performance degradation on downstream tasks compared to their counterparts  $M_{\rm ft}$ , cross KE methods, and editing datasets. Of the 144 fine-tuned cases, 122 cases experience decline, while the remaining cases exhibit an increase in downstream performance. The largest decline occurs when Llama2 is edited via AlphaEdit on 100 counterfactual facts, followed by LoRA fine-tuning, resulting in a drop of 43.76 p.p. in accuracy rate (from 79.87% to 36.11%).

Further, among the three fine-tuning methods, DoRA is the most affected by KE, as indicated by the predominance of green cells. By contrast, full fine-tuning is the most robust against KE, with 11 out of 36 configurations of even increased performance. An analysis of the relationship between downstream task performance and experiment settings is presented in Appendix G.1.

#### 5 CONCLUSION

In this paper, we show that knowledge edits rarely persist unchanged under fine-tuning: in many cases, fine-tuning impairs editing performance or even elicits new knowledge that is different from the target and original knowledge. At the same time, we find that edits themselves can affect downstream fine-tuning performance, even when applied. Motivated by the dual goals of preserving beneficial edits and removing malicious ones, we further explored selective fine-tuning strategies. The results show that updating only non-edited layers to preserve beneficial edits slightly sacrifices downstream performance. For removing covert edits, tuning edited layers does not help and calls for future exploration. These results establish that model editing and fine-tuning are tightly coupled processes whose interaction can be exploited to balance adaptability with knowledge control. Our work provides both empirical baselines and actionable strategies for building large language models that remain adaptable yet reliably steerable with respect to edited knowledge. Future research on model editing should consider robustness not in isolation but across the entire LLM pipeline.

## 6 REPRODUCIBILITY STATEMENT

We provide all necessary resources to facilitate reproducibility of our results. Dataset descriptions and preprocessing steps are detailed in Section 3.1 and Appendix H. Implementation details, model configurations, and training setups are reported in Appendix A and Appendix H. We will release the code to reproduce all experiments once published. Together, these materials ensure that our results can be independently verified.

## REFERENCES

- Canyu Chen, Baixiang Huang, Zekun Li, Zhaorun Chen, Shiyang Lai, Xiongxiao Xu, Jia-Chen Gu, Jindong Gu, Huaxiu Yao, Chaowei Xiao, Xifeng Yan, William Yang Wang, Philip Torr, Dawn Song, and Kai Shu. Can editing LLMs inject harm? In *Neurips Safe Generative AI Workshop* 2024, 2024. URL https://openreview.net/forum?id=edCz9pEu16.
- Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, and Tat-Seng Chua. Alphaedit: Null-space constrained model editing for language models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=HvSytvg3Jh.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, et al. The llama 3 herd of models, 2024. URL https://arxiv.org/abs/2407.21783.
- Akshat Gupta, Dev Sajnani, and Gopala Anumanchipalli. A unified framework for model editing. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15403–15418, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.903. URL https://aclanthology.org/2024.findings-emnlp.903/.
- Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. In Iryna Gurevych and Yusuke Miyao, editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1031. URL https://aclanthology.org/P18-1031/.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.
- Zhiqiang Hu, Lei Wang, Yihuai Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya Poria, and Roy Lee. LLM-adapters: An adapter family for parameter-efficient fine-tuning of large language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5254–5276, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023. emnlp-main.319. URL https://aclanthology.org/2023.emnlp-main.319/.
- Tianjie Ju, Yiting Wang, Xinbei Ma, Pengzhou Cheng, Haodong Zhao, Yulong Wang, Lifeng Liu, Jian Xie, Zhuosheng Zhang, and Gongshen Liu. Flooding spread of manipulated knowledge in llm-based multi-agent communities. *CoRR*, abs/2407.07791, 2024. URL https://doi.org/10.48550/arXiv.2407.07791.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. Zero-shot relation extraction via reading comprehension. In Roger Levy and Lucia Specia, editors, *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 333–342, Vancouver, Canada, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/K17-1034. URL https://aclanthology.org/K17-1034/.
- Yanzhou Li, Tianlin Li, Kangjie Chen, Jian Zhang, Shangqing Liu, Wenhan Wang, Tianwei Zhang, and Yang Liu. Badedit: Backdooring large language models by model editing. In *The Twelfth International Conference on Learning Representations*, 2024a. URL https://openreview.net/forum?id=duZANm2ABX.
- Zhoubo Li, Ningyu Zhang, Yunzhi Yao, Mengru Wang, Xi Chen, and Huajun Chen. Unveiling the pitfalls of knowledge editing for large language models. In *The Twelfth International Conference on Learning Representations*, 2024b. URL https://openreview.net/forum?id=fNktD3ib16.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. In S. Koyejo, S. Mohamed, A. Agarwal,

D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 1950–1965. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper\_files/paper/2022/file/0cde695b83bd186c1fd456302888454c-Paper-Conference.pdf.

- Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. Dora: weight-decomposed low-rank adaptation. In *Proceedings of the 41st International Conference on Machine Learning*, pages 32100–32121, 2024.
- Vittorio Mazzia, Alessandro Pedrani, Andrea Caciolai, Kay Rottmann, and Davide Bernardi. A survey on knowledge editing of neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 36(7):11759–11775, 2025. doi: 10.1109/TNNLS.2024.3498935.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in GPT. *Advances in Neural Information Processing Systems*, 36, 2022. arXiv:2202.05262.
- Kevin Meng, Arnab Sen Sharma, Alex J Andonian, Yonatan Belinkov, and David Bau. Massediting memory in a transformer, 2023. URL https://openreview.net/forum?id=MkbcAHIYqyS.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. Fast model editing at scale. In *International Conference on Learning Representations*, 2022a. URL https://openreview.net/pdf?id=0DcZxeWfOPt.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. Memory-based model editing at scale. In *International Conference on Machine Learning*, pages 15817–15831. PMLR, 2022b.
- Venkatesh Balavadhani Parthasarathy, Ahtsham Zafar, Aafaq Khan, and Arsalan Shahid. The ultimate guide to fine-tuning llms from basics to breakthroughs: An exhaustive review of technologies, research, best practices, applied research challenges and opportunities. *arXiv* preprint arXiv:2408.13296, 2024.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023. URL https://arxiv.org/abs/2307.09288.
- Ben Wang and Aran Komatsuzaki. Gpt-j-6b: A 6 billion parameter autoregressive language model. https://github.com/kingoflolz/mesh-transformer-jax, May 2021. URL https://github.com/kingoflolz/mesh-transformer-jax.
- Song Wang, Yaochen Zhu, Haochen Liu, Zaiyi Zheng, Chen Chen, and Jundong Li. Knowledge editing for large language models: A survey. *ACM Comput. Surv.*, 57(3), November 2024. ISSN 0360-0300. doi: 10.1145/3698590. URL https://doi.org/10.1145/3698590.
- Yu Wang, Ruihan Wu, Zexue He, Xiusi Chen, and Julian McAuley. Large scale knowledge washing. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=dXCpPgjTtd.

Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. Editing large language models: Problems, methods, and opportunities. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. URL https://openreview.net/forum?id=NZZBBJUGcd8.

Zhengmao Ye, Dengchun Li, Zetao Hu, Tingfeng Lan, Jian Sha, Sicong Zhang, Lei Duan, Jie Zuo, Hui Lu, Yuanchun Zhou, et al. mlora: Fine-tuning lora adapters via highly-efficient pipeline parallelism in multiple gpus. *arXiv preprint arXiv:2312.02515*, 2023.

Paul Youssef, Zhixue Zhao, Daniel Braun, Jörg Schlötterer, and Christin Seifert. Position: Editing large language models poses serious safety risks. In *Forty-second International Conference on Machine Learning Position Paper Track*, 2025. URL https://openreview.net/forum?id=QLKBm1PaCU.

Ce Zheng, Lei Li, Qingxiu Dong, Yuxuan Fan, Zhiyong Wu, Jingjing Xu, and Baobao Chang. Can we edit factual knowledge by in-context learning? In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4862–4876, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.296. URL https://aclanthology.org/2023.emnlp-main.296/.

#### Implementation Details

We build on the MEMIT codebase<sup>4</sup> and EasyEdit<sup>5</sup>, and implement all fine-tuning with Hugging Face Transformers (v4.43) in PyTorch, using A100 GPUs with bf16 precision. Models include GPT-J-6B, GPT-2XL, and Llama-2-7B. Knowledge edits are applied with MEMIT and AlphaEdit following their default setups. Fine-tuning uses three methods: Full FT, LoRA (r=8,  $\alpha=16$ , dropout=0.05), and DoRA. Edited vs. non-edited layer experiments freeze parameters outside the specified layers. Optimizers are AdamW with learning rate 1e-5-5e-5, batch size 64, and 2-3 epochs.

#### A Model Configurations

Each column in Table 9 represents a selectable parameter, with experimental settings generated by their Cartesian product. As noted in its caption, configurations with zero edits are consistent across datasets, e.g., Llama2 edited with 0 facts from zsRE and COUNTERFACT dataset are identical. A summary of model details is provided in Table 8. Additionally, DeepSeek is not fine-tuned, with the rationale detailed in Section 4.4.

Table 8: Summary of model details.

| Model name                   | Year of release | Num of parameters | Huggingface handle               |
|------------------------------|-----------------|-------------------|----------------------------------|
| GPT2-XL                      | 2019            | 1.61B             | openai-community/gpt2-xl         |
| GPT-J-6B                     | 2021            | 6.05B             | EleutherAI/gpt-j-6b              |
| Llama2-7B-hf                 | 2023            | 6.74B             | meta-llama/Llama-2-7b-hf         |
| Llama3.1-8B-Instruct         | 2024            | 8.03B             | meta-llama/Llama-3.1-8B-Instruct |
| DeepSeek-R1-Distill-Llama-8B | 2025            | 8.03B             | DeepSeek-R1-Distill-Llama-8B     |

<sup>4</sup>https://github.com/kmeng01/memit

<sup>&</sup>lt;sup>5</sup>https://github.com/zjunlp/EasyEdit

Table 9: Scope of the experimental parameters. M - models, D - datasets, E - editing methods, N - editing numbers, F - fine-tuning methods. Note that settings with 0 edit are identical across datasets, e.g., GPT-J with 0 edits on zsRE is identical to GPT-J with 0 edits on COUNTERFACT.

|   | Model    |   | Dataset     |   | Edit method |   | Edits |   | fine-tune method |
|---|----------|---|-------------|---|-------------|---|-------|---|------------------|
|   | GPT2-XL  |   |             | 1 |             | 1 | 0     |   | No fine-tuning   |
|   | GPTJ     |   | zsRE        |   | No editing  | ١ | 100   |   | LoRA             |
|   | Llama2   | × | COUNTERFACT | × | MEMIT       | × | 1000  | × | DoRA             |
| ĺ | Llama3.1 |   |             | İ | AlphaEdit   | İ | 10000 |   | Full-size        |
|   | DeepSeek |   |             |   | _           |   |       |   |                  |

$$S = M \times D \times E \times N \times F$$
 
$$= \{(m, d, e, n, f) \mid m \in M, d \in D, e \in E, n \in N, f \in F\}$$

## B RESULTS VALIDATION

In this section, "original values" refers the results given by the paper and "validating values" refers to the results obtained in the validation experiments. Overall, our validation demonstrates that the KE and fine-tuning results produced by our code are highly consistent with the original results, indicating that the outputs generated by our implementation are relatively reliable. Minor discrepancies may arise from factors such as model loading precision, random initialization, or hardware-related numerical differences.

Table 10: KE performances (%) check between original values put forward in paper and our validating results. Comparison across various KE method (MEMIT, AlphaEdit) and datasets (zsRE and COUNTERFACT).

|             |                 | ME                | MIT             |                   | AlphaEdit       |                   |                 |                   |  |  |  |
|-------------|-----------------|-------------------|-----------------|-------------------|-----------------|-------------------|-----------------|-------------------|--|--|--|
| Metrics     | ZS              | sRE               | COUNT           | TERFACT           | Z               | RE                | COUNTERFACT     |                   |  |  |  |
|             | Original values | Validating values | Original values | Validating values | Original values | Validating values | Original values | Validating values |  |  |  |
| Efficacy    | 96.70(±0.30)    | 96.93             | 98.9 (±0.20)    | 99.10             | 99.79 (±0.14)   | 99.31             | 99.75 (±0.08)   | 98.35             |  |  |  |
| Paraphrase  | 89.70(±0.50)    | 90.75             | 88.6 (±0.50)    | 88.66             | 96.00 (±0.22)   | 96.71             | 96.38 (±0.23)   | 95.90             |  |  |  |
| Specificity | 26.60 (±0.30)   | 26.33             | 73.70 (±0.50)   | 73.53             | 28.29 (±0.25)   | 28.07             | 75.48 (±0.21)   | 80.16             |  |  |  |

**Validation for KE** As shown in Table 10, the results obtained from our validation experiments closely align with the original values reported in the paper, indicating strong reproducibility and correctness of our re-implementation.

**Validation for fine-tuning** As shown in Figure 3 and Figure 4, validating results are close to original results, indicating the reliability of our outputs.

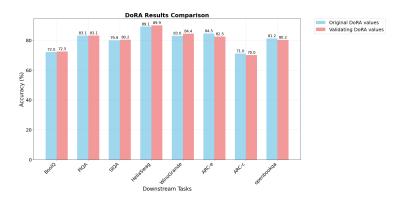
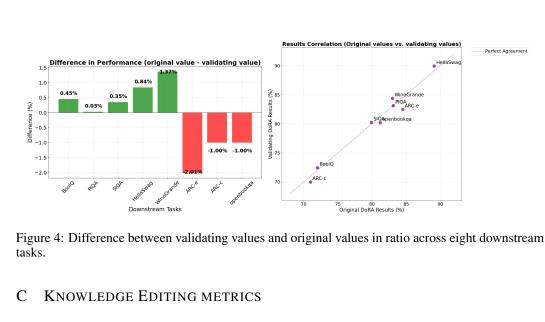


Figure 3: Difference between validating values and original values across eight downstream tasks.



We construct our evaluation using metrics defined in previous works (Meng et al., 2022; Yao et al., 2023). For below equations, P for probability of model G generating certain outputs,  $o_i$  for answers produced, and E represents mathematical expectation.

**ES Success** (**Efficacy / ES**). Efficacy measures the proportion of successful edits, defined by whether the model assigns a higher probability to the edited target than to the original answer for a given fact and its corresponding query prompt (P(edited target)) > P(original answer)).

Formulas 2 and 3 present the mathematical definition of Efficacy for the zsRE and COUNTERFACT datasets, respectively.

$$\mathbb{E}_{i} \left[ o_{i} = \arg \max_{x_{E}} \mathbb{P}_{G} \left( x_{E} \mid p(s_{i}, r_{i}) \right) \right]$$
 (2)

$$\mathbb{E}_i \left[ \mathbb{P}_G \left( o_i \mid p(s_i, r_i) \right) \right] > \mathbb{P}_G \left( o_i^c \mid p(s_i, r_i) \right) \tag{3}$$

**Paraphrase Success (Paraphrase / PS).** Paraphrase evaluates model's generalization ability after editing facts. Similar to *Efficacy*, it is measured by checking whether P(edited target)) > P(original answer) but under a paraphrased query prompt.

Formulas 4 and 5 present the mathematical definition of Paraphrase for the zsRE and COUNTER-FACT datasets, respectively.

$$\mathbb{E}_{i} \left[ \mathbb{E}_{(s_{i}, r_{i}) \in \text{paraphrases}(s_{i}, r_{i})} \left[ o_{i} = \underset{x_{E}}{\arg \max} \mathbb{P}_{G}(x_{E} \mid p) \right] \right]$$
(4)

$$\mathbb{E}_{i}\left[\mathbb{E}_{(s_{i},r_{i})\in \text{paraphrases}(s_{i},r_{i})}\left[\mathbb{P}_{G}[o_{i}\mid p] > \mathbb{P}_{G}[o_{i}^{c}\mid p]\right]\right] \tag{5}$$

**Neighborhood Success (Specificity / NS).** Specificity assesses the locality of a knowledge edit by measuring its unwanted impact on facts nearby the facts involved in KE. NS calculates the percentage of unedited facts nearby which remain unchanged after editing.

Formulas 6 and 7 present the mathematical definition of Specificity for the zsRE and COUNTER-FACT datasets, respectively.

$$\mathbb{E}_{i} \left[ \mathbb{E}_{p \in \text{neighborhood prompts}(s_{i}; r_{i})} \left[ o_{i}^{c} = \underset{x_{E}}{\operatorname{arg max}} \mathbb{P}_{G} \left( x_{E} \mid p \right) \right] \right]$$

$$\mathbb{E}_{j} \left[ \mathbb{E}_{\text{neighborhood prompts}(s_{i}, r_{j})} \left[ \mathbb{P}_{G} (o_{i}^{j} \mid p) < \mathbb{P}_{G} (o_{i}^{j} \mid p') \right] \right]$$

$$(6)$$

# D OVERALL KE PERFORMANCE

The full set of experimental combinations mentioned in below charts can be found in Table 9. Table 11 presents the overall knowledge editing (KE) results for GPT-J across different combinations of KE dataset, number of edits, KE method, and fine-tuning method. Table 12 reports the corresponding results for Llama2 under the same experimental configurations while Table 13 shows the results for GPT2-XL.

Table 11: GPT-J's KE performance (%) under different FT settings (No fine-tuning<sup>2</sup>, LoRA, DoRA and full fine-tuning<sup>3</sup>) across various KE method (MEMIT, AlphaEdit) and datasets (zsRE and COUNTERFACT<sup>4</sup>).

|                          |          |         |        | ME     | MIT    |         |                    | Alph  | aEdit  |                      |
|--------------------------|----------|---------|--------|--------|--------|---------|--------------------|-------|--------|----------------------|
| Dataset                  | #Edits   | Metrics | No ft  | LoRA   | DoRA   | Full ft | No ft <sup>2</sup> | LoRA  | DoRA   | Full ft <sup>3</sup> |
|                          |          | ES      | 23.47  | 23.42  | 22.50  | 22.67   | 23.47              | 23.42 | 22.50  | 22.67                |
|                          | 0        | PS      | 23.17  | 21.56  | 21.65  | 23.50   | 23.17              | 21.56 | 21.65  | 23.50                |
|                          |          | NS      | 28.35  | 27.82  | 27.76  | 25.93   | 28.35              | 27.82 | 27.76  | 25.93                |
|                          |          | ES      | 99.07  | 88.79  | 89.87  | 97.95   | 99.33              | 84.74 | 99.33  | 98.50                |
|                          | $10^{2}$ | PS      | 95.76  | 81.09  | 84.62  | 95.15   | 97.55              | 76.00 | 97.55  | 96.93                |
|                          |          | NS      | 28.74  | 27.12  | 27.64  | 27.25   | 28.65              | 26.12 | 28.65  | 27.57                |
| zsRE                     |          | ES      | 99.10  | 84.53  | 85.31  | 98.74   | 99.31              | 84.74 | 81.70  | 98.98                |
|                          | $10^{3}$ | PS      | 95.97  | 77.77  | 78.48  | 94.05   | 96.71              | 76.00 | 73.82  | 94.39                |
|                          |          | NS      | 28.13  | 26.80  | 25.87  | 27.06   | 28.07              | 26.12 | 26.25  | 26.12                |
|                          | $10^{4}$ | ES      | 96.93  | 66.52  | 67.83  | 89.38   | 89.81              | 49.97 | 23.64  | 74.62                |
|                          |          | PS      | 90.75  | 58.92  | 61.22  | 81.21   | 78.59              | 43.90 | 23.03  | 64.57                |
|                          |          | NS      | 26.33  | 25.42  | 24.50  | 25.54   | 22.96              | 23.89 | 25.57  | 22.96                |
|                          |          | ES      | 15.00  | 15.00  | 16.00  | 14.00   | 15.00              | 15.00 | 16.00  | 14.00                |
|                          | 0        | PS      | 16.50  | 23.50  | 20.00  | 18.50   | 16.50              | 23.50 | 20.00  | 18.50                |
|                          |          | NS      | 84.40  | 82.60  | 83.50  | 84.80   | 84.40              | 82.60 | 83.50  | 84.80                |
|                          |          | ES      | 100.00 | 100.00 | 100.00 | 100.00  | 100.00             | 99.00 | 100.00 | 100.00               |
|                          | $10^{2}$ | PS      | 95.00  | 92.50  | 91.00  | 91.90   | 98.50              | 94.00 | 94.00  | 97.00                |
|                          |          | NS      | 81.10  | 80.90  | 81.80  | 81.62   | 81.10              | 80.00 | 80.50  | 81.40                |
| $\mathbb{C}\mathbb{F}^4$ |          | ES      | 100.00 | 99.00  | 99.00  | 99.90   | 98.35              | 99.70 | 99.60  | 97.95                |
|                          | $10^{3}$ | PS      | 93.95  | 87.70  | 86.90  | 90.45   | 95.90              | 90.90 | 90.95  | 97.00                |
|                          |          | NS      | 81.17  | 79.96  | 80.80  | 80.68   | 80.16              | 79.58 | 79.84  | 81.40                |
|                          |          | ES      | 99.10  | 94.34  | 94.46  | 97.79   | 98.87              | 99.16 | 82.00  | 95.85                |
|                          | $10^{4}$ | PS      | 88.66  | 76.27  | 76.24  | 80.38   | 86.70              | 87.62 | 62.00  | 74.09                |
|                          |          | NS      | 73.53  | 73.89  | 74.17  | 75.08   | 67.77              | 68.54 | 74.00  | 72.56                |

Table 12: Llama2's KE performance (%) under different FT settings (No fine-tuning<sup>2</sup>, LoRA, DoRA and full fine-tuning<sup>3</sup>) across various KE method (MEMIT, AlphaEdit) and datasets (zsRE and COUNTERFACT<sup>4</sup>).

| Dataset | #Edits   | Metrics |        | ME    | MIT   |         |                    | Alph  | aEdit |                      |
|---------|----------|---------|--------|-------|-------|---------|--------------------|-------|-------|----------------------|
| Dutusci | "Lates   | Metrics | No ft  | LoRA  | DoRA  | Full ft | No ft <sup>2</sup> | LoRA  | DoRA  | Full ft <sup>3</sup> |
|         |          | ES      | 45.61  | 46.92 | 42.30 | 22.67   | 45.61              | 46.92 | 42.30 | 22.67                |
|         | 0        | PS      | 45.57  | 43.08 | 42.02 | 23.5    | 45.57              | 43.08 | 42.02 | 23.50                |
|         |          | NS      | 32.15  | 36.85 | 28.25 | 25.93   | 32.15              | 36.85 | 28.25 | 25.93                |
|         |          | ES      | 86.03  | 76.30 | 72.00 | 22.67   | 93.33              | 57.96 | 93.33 | 21.83                |
|         | $10^{2}$ | PS      | 86.01  | 71.69 | 66.02 | 20.14   | 84.93              | 53.49 | 84.93 | 19.84                |
| zsRE    |          | NS      | 31.68  | 29.69 | 28.70 | 14.58   | 32.42              | 29.67 | 32.42 | 15.34                |
|         |          | ES      | 51.38  | 46.52 | 46.00 | 10.22   | 93.23              | 50.45 | 51.53 | 24.36                |
|         | $10^{3}$ | PS      | 50.04  | 44.44 | 46.65 | 10.37   | 86.57              | 48.79 | 48.83 | 22.27                |
|         |          | NS      | 28.09  | 27.68 | 25.79 | 12.22   | 34.3               | 32.11 | 30.88 | 15.40                |
|         |          | ES      | 48.62  | 48.64 | 48.23 | 14.00   | 84.31              | 46.03 | 45.38 | 24.44                |
|         | $10^{4}$ | PS      | 50.20  | 49.75 | 48.62 | 13.20   | 79.03              | 44.27 | 43.65 | 21.97                |
|         |          | NS      | 25.73  | 24.97 | 24.32 | 14.59   | 34.35              | 30.45 | 31.98 | 15.59                |
|         |          | ES      | 11.37  | 11.41 | 11.52 | 45.00   | 11.37              | 11.41 | 11.52 | 45.00                |
|         | 0        | PS      | 41.61  | 40.57 | 40.32 | 32.00   | 41.61              | 40.57 | 40.32 | 32.00                |
|         |          | NS      | 91.36  | 91.32 | 91.21 | 50.00   | 91.36              | 91.32 | 91.21 | 50.00                |
|         | -        | ES      | 100.00 | 94.00 | 97.00 | 47.00   | 100.00             | 66.00 | 61.00 | 48.00                |
|         | $10^{2}$ | PS      | 98.00  | 81.50 | 95.50 | 70.50   | 77.50              | 49.00 | 49.50 | 59.50                |
| $CF^4$  |          | NS      | 75.40  | 78.50 | 80.60 | 52.00   | 85.50              | 84.30 | 83.80 | 52.30                |
|         |          | ES      | 100.00 | 94.00 | 94.50 | 47.60   | 99.10              | 52.40 | 56.90 | 48.60                |
|         | $10^{3}$ | PS      | 94.50  | 82.50 | 84.30 | 67.55   | 67.80              | 42.45 | 42.85 | 54.35                |
|         |          | NS      | 70.50  | 77.00 | 76.33 | 53.02   | 83.97              | 81.22 | 82.20 | 52.06                |
|         |          | ES      | 86.96  | 70.18 | 68.27 | 48.07   | 87.43              | 37.90 | 37.15 | 48.46                |
|         | $10^{4}$ | PS      | 73.62  | 64.32 | 61.25 | 32.94   | 55.71              | 34.01 | 33.40 | 47.74                |
|         |          | NS      | 68.64  | 62.35 | 61.74 | 52.71   | 80.67              | 79.02 | 79.89 | 52.52                |

Table 13: GPT2-XL's KE performance (%) under different FT settings (No fine-tuning<sup>2</sup>, LoRA, DoRA and full fine-tuning<sup>3</sup>) across various KE method (MEMIT, AlphaEdit) and datasets (zsRE and COUNTERFACT<sup>4</sup>).

|                          |          |         |                    | ME    | MIT   |         |        | Alph  | aEdit |                      |
|--------------------------|----------|---------|--------------------|-------|-------|---------|--------|-------|-------|----------------------|
| zsRE                     | #Edit    | Metrics | No ft <sup>2</sup> | LoRA  | DoRA  | Full ft | No ft  | LoRA  | DoRA  | Full ft <sup>3</sup> |
|                          |          | ES      | 32.80              | 32.74 | 32.87 | 18.04   | 32.80  | 32.74 | 32.87 | 18.04                |
|                          | 0        | PS      | 35.60              | 35.32 | 35.62 | 17.57   | 35.60  | 35.32 | 35.62 | 17.57                |
|                          |          | NS      | 23.76              | 23.70 | 23.75 | 24.64   | 23.76  | 23.70 | 23.75 | 24.64                |
|                          |          | ES      | 80.00              | 58.37 | 58.39 | 81.16   | 97.18  | 67.80 | 76.72 | 18.04                |
|                          | $10^{2}$ | PS      | 76.10              | 52.94 | 53.86 | 74.30   | 93.60  | 58.42 | 64.35 | 17.57                |
|                          |          | NS      | 25.75              | 24.57 | 24.98 | 25.79   | 25.06  | 27.11 | 24.48 | 24.64                |
| zsRE                     |          | ES      | 77.85              | 43.51 | 45.18 | 81.22   | 93.13  | 54.52 | 55.11 | 24.74                |
|                          | $10^{3}$ | PS      | 73.42              | 42.87 | 43.62 | 74.92   | 87.03  | 48.61 | 50.64 | 23.86                |
|                          |          | NS      | 26.22              | 23.57 | 24.35 | 25.76   | 25.22  | 25.43 | 26.07 | 24.00                |
|                          |          | ES      | 62.61              | 20.34 | 20.65 | 63.39   | 62.34  | 25.80 | 27.47 | 22.09                |
|                          | $10^{4}$ | PS      | 57.68              | 19.63 | 19.89 | 57.87   | 54.84  | 24.36 | 26.17 | 21.13                |
|                          |          | NS      | 25.81              | 24.59 | 24.76 | 24.83   | 21.31  | 23.69 | 24.17 | 23.67                |
|                          |          | ES      | 20.00              | 20.13 | 20.14 | 19.00   | 20.00  | 20.13 | 20.14 | 19.00                |
|                          | 0        | PS      | 35.00              | 35.21 | 35.17 | 22.50   | 35.00  | 35.21 | 35.17 | 22.50                |
|                          |          | NS      | 69.00              | 68.93 | 68.97 | 79.40   | 69.00  | 68.93 | 68.97 | 79.40                |
|                          |          | ES      | 97.00              | 83.00 | 82.00 | 97.00   | 100.00 | 96.00 | 98.00 | 19.00                |
|                          | $10^{2}$ | PS      | 86.50              | 71.00 | 70.50 | 84.50   | 98.00  | 87.50 | 93.00 | 22.00                |
|                          |          | NS      | 76.40              | 77.30 | 77.50 | 76.20   | 73.90  | 76.60 | 76.60 | 79.40                |
| $\mathbb{C}\mathbb{F}^4$ |          | ES      | 93.40              | 78.37 | 78.50 | 92.60   | 100.00 | 89.30 | 90.55 | 21.92                |
|                          | $10^{3}$ | PS      | 81.35              | 64.43 | 63.45 | 79.55   | 95.75  | 74.55 | 77.03 | 24.66                |
|                          |          | NS      | 75.32              | 76.27 | 76.45 | 75.62   | 72.44  | 75.25 | 73.20 | 78.20                |
|                          |          | ES      | 79.17              | 62.10 | 61.97 | 78.03   | 92.94  | 53.91 | 57.65 | 21.92                |
|                          | $10^{4}$ | PS      | 65.44              | 44.47 | 44.43 | 63.74   | 76.33  | 41.94 | 45.77 | 24.66                |
|                          |          | NS      | 69.83              | 63.37 | 63.53 | 70.16   | 64.68  | 71.81 | 70.32 | 78.20                |

# E ONLY FINE-TUNING EDITED OR NON-EDITED LAYERS

Examples selected on GPT-J are: (1) GPT-J being edited 100 facts from COUNTERFACT by MEMIT and then fine-tuned by DoRA; (2) GPT-J being edited 100 facts from zsRE by MEMIT and then fine-tuned by DoRA. The result for KE performance is shown in Table 14 and the result for downstream performance is shown in Table 15.

Table 14: KE performance GPT-J-based model. Naming format are identical to Table 5. Example 1¹: GPT-J bing edited 100 zsRE facts using MEMIT and fine-tuned by DoRA; Example 2²: GPT-J bing edited 100 COUNTERFACT facts using AlphaEdit and fine-tuned by DoRA. GPT-J³, ⁴ are GPT-Js without KE and fine-tuning, being evaluated by zsRE and COUNTERFACT, respectively. Naming format are identical to Table 5.

| KE<br>performance | $\frac{\text{GPT-J}^3}{M}$ | $M_{ed}$ | $M_{ed\_ft\_all}$ | Example 1 <sup>1</sup> $M_{ed\_ft\_edited}$ | $M_{ed\_ft\_non-edited}$ | $GPT-J^4$ $M$ | $M_{ed}$ | $M_{ed\_ft\_all}$ | Example $2^2$ $M_{ed_*ft_*edited}$ | $M_{ed_*ft_*non-edited}$ |
|-------------------|----------------------------|----------|-------------------|---|--------------------------|---------------|----------|-------------------|------------------------------------|--------------------------|
| Efficacy          | 23.47                      | 99.07    | 89.87             | 95.79                                       | 86.39                    | 15.00         | 100.00   | 100.00            | 100.00                             | 100.00                   |
| Paraphrase        | 23.17                      | 95.76    | 84.62             | 91.12                                       | 84.43                    | 16.50         | 98.50    | 94.00             | 97.00                              | 94.50                    |
| Specificity       | 28.35                      | 28.74    | 27.64             | 29.13                                       | 25.71                    | 84.40         | 81.10    | 80.50             | 81.00                              | 80.40                    |

Table 15: Downstream task performance on GPT-J-based examples. Examples choosen are identical to Table 5

| Downstream | GPT-J |          |                   | Example 1            |                          |          |                   | Example 2            |                          |
|------------|-------|----------|-------------------|----------------------|--------------------------|----------|-------------------|----------------------|--------------------------|
| tasks      | M     | $M_{ed}$ | $M_{ed\_ft\_all}$ | $M_{ed\_ft\_edited}$ | $M_{ed\_ft\_non-edited}$ | $M_{ed}$ | $M_{ed\_ft\_all}$ | $M_{ed\_ft\_edited}$ | $M_{ed\_ft\_non-edited}$ |
| BoolQ      | 56.57 | 31.04    | 63.88             | 24.95                | 64.40                    | 19.27    | 63.67             | 61.99                | 27.92                    |
| PIQA       | 1.36  | 0.87     | 73.07             | 53.26                | 74.97                    | 1.47     | 74.65             | 42.60                | 74.59                    |
| SIQA       | 0.41  | 0.46     | 73.54             | 63.31                | 74.82                    | 0.67     | 74.21             | 41.91                | 69.29                    |
| HellaSwag  | 0.03  | 0.01     | 70.84             | 42.76                | 59.80                    | 0.03     | 71.95             | 18.86                | 33.71                    |
| WinoGrande | 30.23 | 0.32     | 69.69             | 59.27                | 67.01                    | 0.39     | 70.80             | 37.81                | 69.22                    |
| ARC-e,     | 1.64  | 1.73     | 67.93             | 21.89                | 67.63                    | 1.94     | 68.86             | 56.40                | 65.45                    |
| ARC-c      | 1.02  | 1.28     | 53.07             | 17.49                | 51.54                    | 1.37     | 51.96             | 38.31                | 49.83                    |
| openbookqa | 1.20  | 1.40     | 64.40             | 41.40                | 66.00                    | 1.80     | 66.40             | 57.60                | 63.80                    |
| Average    | 12.12 | 2.15     | 81.70             | 65.43                | 80.61                    | 4.79     | 81.00             | 65.35                | 72.46                    |

# F EDITING PERFORMANCE OF DEEPSEEK

#### F.1 LAYER DETERMINATION

As there is currently no published research using DeepSeek as base models for KE, we include three potential settings of editing layers when running KE on DeepSeek: (1) using Causal Tracing with Frozen Components (CTFC) to determine layers, (2) directly using LLaMA2's editing layer setting, and (3) directly using GPT2-XL's editing layer setting. The CTFC method is introduced and used by Meng et al. (2023), which enables precise identification of layers most relevant to knowledge storage. Besides, recent KE studies frequently use GPT2-XL and LLaMA2 as base models, directly adopting their layers setting provides reasonable baselines and allows us to carry out comparisons among DeepSeek and them.

# F.2 RESULTS

We ultimately tried all three setups, and the KE results on zsRE dataset with editing number from 0 to 10,000 are shown in Table 16. For CTFC, as shown in Figure 5, layers 1 to 5 shares the largest gap between purple bar and green bar, exhibit the largest gap between the purple and green bars, indicating that these layers contribute most significantly to knowledge storage. Thus, editing layers determined by CTFC are layers 1 to 5. In addition to Figure 5, we use heatmap (Figure 6) visualizations of layer-wise causal effects to analyze how different components of the model contribute to factual knowledge retrieval. These heatmaps guide the selection of editing layers by highlighting consistent and concentrated MLP-specific causal effects in early layers, enabling targeted and effective knowledge editing across architectures.

Table 16: DeepSeek's KE performance (%) using MEMIT across editing datasets and different settings of editing layers: GPT2-XL¹ means using GPT2-XL's editing layer settings (Meng et al., 2023), Llama2² means using Llama2's editing layer settings (Gupta et al., 2024), CTFC-determined³ means layers are determined by Casual Tracing with Frozen Components (CTFC) (Meng et al., 2023).

| Editing layer settings       | #Edits   |       | zsRE  |       | cot   | JNTERF | ACT   |
|------------------------------|----------|-------|-------|-------|-------|--------|-------|
|                              |          | ES    | PS    | NS    | ES    | PS     | NS    |
|                              | 0        | 24.69 | 25.24 | 21.21 | 16.00 | 18.00  | 83.20 |
| come vii l                   | $10^{2}$ | 28.39 | 45.50 | 1097  | 97.00 | 82.00  | 63.00 |
| GPT2-XL <sup>1</sup>         | $10^{3}$ | 18.97 | 23.83 | 3.99  | 93.30 | 73.55  | 61.65 |
|                              | $10^{4}$ | 3.53  | 14.92 | 2.02  | 87.09 | 60.70  | 50.70 |
|                              | 0        | 24.69 | 25.24 | 21.21 | 16.00 | 18.00  | 83.20 |
| T.1 2 <sup>2</sup>           | $10^{2}$ | 28.69 | 29.57 | 32.06 | 97.00 | 95.00  | 68.50 |
| Llama2 <sup>2</sup>          | $10^{3}$ | 31.65 | 30.54 | 33.62 | 79.50 | 64.45  | 53.20 |
|                              | $10^{4}$ | 0.55  | 0.59  | 2.84  | 65.98 | 56.66  | 48.98 |
|                              | 0        | 24.69 | 25.24 | 21.21 | 16.00 | 18.00  | 83.20 |
| omno i i i                   | $10^{2}$ | 31.19 | 31.82 | 31.89 | 86.00 | 66.50  | 56.50 |
| CTFC-determined <sup>3</sup> | $10^{3}$ | 20.17 | 18.89 | 28.26 | 72.20 | 60.90  | 51.01 |
|                              | $10^{4}$ | 0.40  | 0.52  | 4.14  | 69.03 | 58.06  | 49.76 |

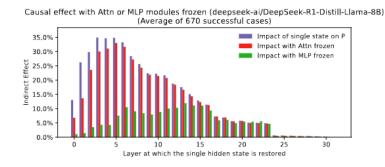


Figure 5: Casual tracing for DeepSeek.

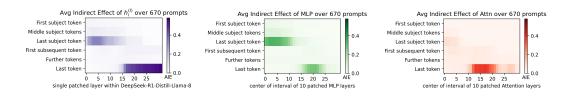


Figure 6: Causal effect heatmaps showing concentrated effects in early layers (1-5) with (from left to right) (a) overall patterns, (b) MLP localisation, and (c) attention mechanisms for the distilled architecture.

We observe that for zsRE dataset with 10,000 edits, all cases perform badly. As shown in Table 16, for the case of CTFC-determined, increasing the editing number from 1,000 to 10,000 leads to a dramatic drop in ES, falling from 20.17% to just 0.4%. For cases using Llama2's or GPT's editing layer setting, similar trend also happens: both models' KE performance drops dramatically to a low level (0.55% for LLama2's setting, 3.53% for GPT2-XL's setting) when the editing number rises to 10,000. As shown highlighed by orange, DeepSeeK also perform badly on other KE metrics (i.e. PS, NS) when editing number rises to 10,000, whereas the smallest PS (0.52%) and NS (2.02%) both appears.

#### F.3 ANALYSIS

This phenomenon may stem from architectural differences between GPT-series and Llama-based models. In addition, as noted by Wang et al. (2025), discrepancies in pre-training data between GPT-based and Llama-based models can lead to suboptimal simulation of the initial model weights  $W_0$ , ultimately degrading the effectiveness of KE. Besides, the performance after editing with 10,000 facts is already extremely low, making it highly susceptible to collapsing to near-zero accuracy after fine-tuning. This **instability** prevents meaningful comparison of KE effectiveness before and after fine-tuning on the zsRE dataset. As a result, we do not perform additional fine-tuning experiments on DeepSeek.

## G DOWNSTREAM-TASK PERFORMANCE

#### G.1 ANALYSIS

Table 17: Average and Standard deviation of degradation (%) in evaluation score(Table 7) across models and finetuning methods. Avg.<sup>1</sup>, Std.<sup>2</sup> are metrics for individual model; Avg\_m<sup>3</sup>, Std\_m<sup>4</sup> are across models.

| Model   | Metrics           | No ft   | Full ft | LoRA  | DoRA |
|---------|-------------------|---------|---------|-------|------|
| GPT-J   | Avg. <sup>1</sup> | 56.84   | 7.17    | 5.19  | 5.87 |
|         | Std. <sup>2</sup> | 17.82   | 12.72   | 8.68  | 4.66 |
| Llama2  | Avg.              | -206.92 | 37.27   | 10.81 | 5.93 |
|         | Std.              | 220.54  | 8.05    | 15.17 | 8.00 |
| GPT2-XL | Avg.              | 1.70    | -0.38   | 2.07  | 8.85 |
|         | Std.              | 23.66   | 2.87    | 9.05  | 9.32 |
| Avg_    |                   | -49.46  | 14.69   | 6.02  | 6.88 |
| Std_    |                   | 169.81  | 18.6    | 11.63 | 7.50 |

FT method wise As shown in Table 17, LoRA exhibits the smallest average performance decrease (-6.02%) compared to **DoRA** (-6.88%) and Full fine-tune (-14.69%). In terms of stability, DoRA demonstrates the lowest standard deviation (7.5) across models, indicating that the impact of KE on DoRA remains relatively consistent regardless of the base model. Notably, models without fine-tuning behave more erratically: the "No ft" group shows the highest average magnitude of change (Avg\_m = 49.46%) and standard deviation

(Std\_m = 169.81), suggesting high variability. This instability may be attributed to base models' relatively poor performances on downstream tasks.

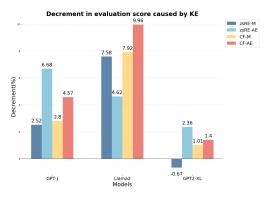


Figure 7: Average decrements ratio (%) caused by KE across models and datasets. M for MEMIT and AE for AlphaEdit. E.g., zsRE-M means MEMIT using zsRE dataset and vice versa.

2.52%, and 4.57% vs 2.8%) across both datasets.

Model wise We found that GPT2-XL demonstrates the most stable performance across different KE methods and datasets among models evaluated. As shown in Figure 7, GPT2-XL has the smallest varying range from -0.67% to 2.36%. In contrast, other models exhibit more variability. For instance, Llama2 experiences a significant fluctuation, with performance ranging from 9.96% to 4.62%.

KE task wise MEMIT generally leads to a smaller reduction in fine-tuning performance compared to AlphaEdit. As illustrated in Figure 7, AlphaEdit has the largest average decrements of 9.96% which appears on Llama2. A similar pattern is also observed for GPT-J, where AlphaEdit show larger decrements (6.68% vs

COUNTERFACT dataset tends to cause more severe drops in fine-tuning performance, as evidenced by the top two largest performance declines in Figure 7 (9.96%, 7.92%) occurring in cases involving COUNTERFACT as KE dataset.

Number of edits also introduces a high degree of variability, with different patterns observed across models, as detailed in Table 7. For example, GPT-J edited with zsRE using MEMIT shows an improvement in fine-tuning performance from 36.95% to 40.92% as the number of edits increases from 100 to 1000, followed by a slight decline to 39.69% when the number of edits reaches 10,000. Conversely, GPT-J, edited with zsRE and fine-tuned with the full dataset, shows a reverse trend, with performance dropping from 35.75% to 33.46% and then rising back to 35.5% as the editing number increases. In addition to these curve-like patterns, a common trend observed in some models is a monotonic decrease in performance. For example, in the case of Llama2 edited by AlphaEdit and fine-tuned by DoRA, the performance consistently drops as the number of edits increases, where performance falls from 78.95% ( $10^2$  edits) to 69.61% ( $10^3$  edits), and eventually to 63.62% ( $10^4$  edits).

#### G.2 RESULTS BREAKDOWN

**GPT2-XL** GPT2-XL's performances on downstream tasks after being edited on zsRE dataset are shown in Table 18 (no fine-tune and full fine-tune) and Table 19 (LoRA and DoRA). Performances of cases being edited on COUNTERFACT dataset are shown in Table 20 (no fine-tune and full fine-tune) and Table 21 (LoRA and DoRA)

Table 18: Downstream task performances (%) of GPT2-XL ( $M_{\rm ed\_ft}$ ) edited on zsRE dataset and then being full fine-tuned or not fine-tuned.

|            |         |          | No 1     | fine-tuni | ng       |          |          |         |          | Full     | fine-tun | ing      |          |          |
|------------|---------|----------|----------|-----------|----------|----------|----------|---------|----------|----------|----------|----------|----------|----------|
| zsRE       |         |          | MEMIT    |           | 1        | AlphaEdi | it       |         |          | MEMIT    |          |          | MEMIT    |          |
|            | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$  | $10^{2}$ | $10^{3}$ | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ      | 58.36   | 59.36    | 58.99    | 48.01     | 59.79    | 59.92    | 24.07    | 6.67    | 9.02     | 14.89    | 14.50    | 6.71     | 6.70     | 5.95     |
| PIQA       | 0.80    | 0.65     | 1.80     | 7.40      | 0.76     | 5.06     | 2.01     | 45.16   | 45.16    | 44.18    | 48.53    | 44.89    | 44.23    | 43.47    |
| SIQA       | 20.42   | 20.21    | 18.22    | 12.49     | 13.92    | 16.17    | 17.45    | 29.94   | 29.27    | 30.35    | 30.60    | 30.40    | 30.55    | 30.25    |
| HellaSwag  | 0.21    | 0.22     | 0.39     | 0.48      | 0.13     | 0.58     | 6.19     | 24.86   | 24.96    | 24.98    | 24.90    | 24.85    | 24.91    | 24.94    |
| WinoGrande | 0.00    | 0.00     | 0.00     | 0.00      | 0.00     | 0.00     | 15.71    | 48.46   | 49.09    | 48.15    | 26.84    | 48.54    | 49.49    | 49.09    |
| ARC-e,     | 8.50    | 8.04     | 8.00     | 8.25      | 5.72     | 11.78    | 12.16    | 24.16   | 24.24    | 24.07    | 22.18    | 24.07    | 24.20    | 24.28    |
| ARC-c      | 6.32    | 6.66     | 6.06     | 6.91      | 4.18     | 9.90     | 10.41    | 22.35   | 22.27    | 22.44    | 21.59    | 22.35    | 22.18    | 22.44    |
| openbookqa | 12.20   | 13.00    | 12.60    | 14.00     | 9.40     | 16.00    | 17.00    | 27.00   | 27.00    | 26.60    | 24.80    | 26.60    | 26.80    | 26.60    |

Table 19: Downstream task performances (%) of GPT2-XL ( $M_{\rm ed\_ft}$ ) edited on zsRE dataset and then being LoRA or DoRA fine-tuned.

|            |         |          |          | LoRA     |          |          |          |         |          |          | DoRA     |          |          |          |
|------------|---------|----------|----------|----------|----------|----------|----------|---------|----------|----------|----------|----------|----------|----------|
| zsRE       |         |          | MEMIT    |          | 1        | AlphaEdi | it       |         |          | MEMIT    |          |          | MEMIT    |          |
|            | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ      | 58.72   | 61.75    | 61.87    | 60.04    | 26.00    | 26.96    | 50.09    | 60.80   | 62.42    | 62.42    | 60.43    | 26.18    | 26.76    | 28.50    |
| PIQA       | 44.78   | 49.67    | 49.96    | 44.22    | 46.47    | 51.07    | 44.99    | 51.03   | 50.71    | 50.71    | 44.50    | 46.79    | 51.41    | 45.21    |
| SIQA       | 39.68   | 42.37    | 38.75    | 37.93    | 19.46    | 38.52    | 32.91    | 39.51   | 39.61    | 39.61    | 37.67    | 19.34    | 38.28    | 35.47    |
| HellaSwag  | 24.82   | 25.61    | 25.32    | 25.47    | 25.49    | 25.77    | 25.05    | 22.62   | 25.17    | 25.17    | 25.30    | 25.32    | 25.93    | 21.93    |
| WinoGrande | 49.64   | 50.04    | 49.45    | 41.72    | 51.81    | 50.94    | 48.22    | 48.22   | 49.09    | 49.09    | 41.99    | 52.17    | 50.59    | 46.09    |
| ARC-e,     | 21.00   | 27.27    | 27.65    | 25.21    | 16.35    | 24.87    | 24.49    | 25.84   | 27.78    | 27.78    | 25.38    | 16.46    | 25.04    | 24.96    |
| ARC-c      | 21.67   | 25.77    | 24.98    | 26.62    | 17.78    | 24.57    | 25.85    | 25.51   | 24.32    | 24.32    | 26.45    | 17.66    | 24.40    | 23.72    |
| openbookqa | 22.00   | 29.40    | 29.20    | 23.00    | 15.00    | 26.80    | 26.20    | 20.00   | 29.60    | 29.60    | 23.40    | 15.00    | 27.00    | 24.60    |

Table 20: Downstream task performances (%) of GPT2-XL ( $M_{\rm ed.ft}$ ) edited on COUNTERFACT<sup>1</sup> dataset and then being full fine-tuned or not fine-tuned.

|            |         |          | No 1     | ine-tuni | ng       |          |          |         |          | Full     | fine-tun | ing      |          |          |
|------------|---------|----------|----------|----------|----------|----------|----------|---------|----------|----------|----------|----------|----------|----------|
| $CF^1$     |         |          | MEMIT    |          | 1        | AlphaEdi | t        |         |          | MEMIT    |          |          | MEMIT    | ,        |
|            | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ      | 58.36   | 59.27    | 60.03    | 39.48    | 58.96    | 60.7     | 12.60    | 6.67    | 7.03     | 11.68    | 22.72    | 11.16    | 6.39     | 7.19     |
| PIQA       | 0.80    | 0.65     | 0.60     | 14.09    | 0.71     | 2.50     | 18.82    | 45.16   | 41.24    | 46.46    | 46.90    | 45.48    | 42.71    | 44.89    |
| SIQA       | 20.42   | 18.68    | 10.29    | 24.21    | 20.21    | 16.53    | 1.18     | 29.94   | 29.48    | 31.42    | 31.83    | 30.45    | 30.19    | 30.30    |
| HellaSwag  | 0.21    | 0.33     | 0.38     | 2.47     | 0.14     | 0.97     | 1.74     | 24.86   | 24.83    | 24.92    | 24.85    | 24.75    | 24.85    | 24.93    |
| WinoGrande | 0.00    | 0.00     | 0.00     | 0.00     | 0.00     | 0.16     | 0.24     | 48.46   | 49.25    | 48.93    | 49.25    | 49.09    | 49.17    | 49.57    |
| ARC-e,     | 8.50    | 7.32     | 6.99     | 20.03    | 6.78     | 17.21    | 2.78     | 24.16   | 23.95    | 23.36    | 20.29    | 23.65    | 24.28    | 24.16    |
| ARC-c      | 6.32    | 5.63     | 4.69     | 17.06    | 5.72     | 15.10    | 1.62     | 22.35   | 22.44    | 22.10    | 19.20    | 22.18    | 22.27    | 22.27    |
| openbookqa | 12.20   | 11.20    | 11.40    | 24.40    | 10.20    | 20.00    | 1.20     | 27.00   | 26.40    | 25.60    | 25.00    | 26.60    | 26.80    | 26.40    |

Table 21: Downstream task performances (%) of GPT2-XL ( $M_{\rm ed.ft}$ ) edited on COUNTERFACT dataset and then being LoRA or DoRA fine-tuned.

|                          |         |          |          | LoRA     |          |          |          |         |          |          | DoRA     |          |          |          |
|--------------------------|---------|----------|----------|----------|----------|----------|----------|---------|----------|----------|----------|----------|----------|----------|
| $\mathbb{C}\mathbb{F}^1$ |         |          | MEMIT    |          | 1        | AlphaEdi | it       |         |          | MEMIT    |          |          | MEMIT    |          |
|                          | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ                    | 58.72   | 35.59    | 40.16    | 46.24    | 53.73    | 52.08    | 52.48    | 60.80   | 35.84    | 40.40    | 46.57    | 43.73    | 41.25    | 11.93    |
| PIQA                     | 44.78   | 41.34    | 46.24    | 42.21    | 48.86    | 49.89    | 47.01    | 51.03   | 41.62    | 45.92    | 48.00    | 25.90    | 51.14    | 47.88    |
| SIQA                     | 39.68   | 39.22    | 37.94    | 35.02    | 34.03    | 34.29    | 36.44    | 39.51   | 38.95    | 38.13    | 34.80    | 33.52    | 40.74    | 37.87    |
| HellaSwag                | 24.82   | 25.27    | 24.42    | 25.47    | 25.11    | 24.18    | 25.10    | 22.62   | 25.10    | 24.57    | 25.30    | 23.44    | 21.99    | 25.06    |
| WinoGrande               | 49.64   | 50.87    | 49.46    | 42.49    | 50.67    | 49.33    | 45.54    | 48.22   | 50.51    | 49.8     | 42.78    | 51.30    | 50.04    | 52.09    |
| ARC-e,                   | 21.00   | 17.26    | 26.09    | 27.22    | 25.25    | 25.93    | 25.67    | 25.84   | 17.38    | 26.26    | 27.44    | 20.16    | 25.72    | 24.45    |
| ARC-c                    | 21.67   | 18.56    | 25.31    | 24.83    | 23.81    | 25.68    | 24.49    | 25.51   | 18.43    | 25.09    | 25.00    | 20.31    | 25.77    | 25.94    |
| openbookqa               | 22.00   | 14.40    | 24.80    | 28.40    | 27.60    | 25.20    | 27.60    | 20.00   | 14.40    | 24.80    | 28.40    | 22.20    | 22.80    | 28.40    |

**Llama2** Llama2's performances on downstream tasks after being edited on zsRE dataset are shown in Table 22 (no fine-tune and full fine-tune) and Table 23 (LoRA and DoRA). Performances of cases being edited on COUNTERFACT dataset are shown in Table 24 (no fine-tune and full fine-tune) and Table 25 (LoRA and DoRA)

Table 22: Downstream task performances (%) of Llama2 ( $M_{\rm ed\_ft}$ ) edited on zsRE dataset and then being full fine-tuned or not fine-tuned.

|            |         |          | No fi    | ine-tunii | ıg       |          |          |         |          | Full     | fine-tun | ing      |          |          |
|------------|---------|----------|----------|-----------|----------|----------|----------|---------|----------|----------|----------|----------|----------|----------|
| zsRE       |         |          | MEMIT    | Γ         |          | AlphaE   | dit      |         |          | MEMIT    |          | _        | MEMIT    | '        |
|            | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$  | $10^{2}$ | $10^{3}$ | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ      | 10.31   | 4.46     | 57.92    | 55.74     | 6.79     | 8.84     | 48.29    | 62.14   | 62.17    | 61.62    | 62.11    | 61.65    | 61.53    | 62.17    |
| PIQA       | 0.16    | 0.11     | 5.06     | 0.20      | 0.16     | 0.11     | 0.22     | 70.84   | 46.46    | 42.17    | 28.07    | 34.28    | 32.64    | 50.22    |
| SIQA       | 2.15    | 0.26     | 16.17    | 1.52      | 3.99     | 4.55     | 0.41     | 61.26   | 22.88    | 31.83    | 31.99    | 17.96    | 34.08    | 27.43    |
| HellaSwag  | 0.00    | 0.58     | 0.58     | 0.00      | 0.00     | 0.00     | 0.00     | 15.77   | 3.50     | 12.67    | 32.96    | 42.06    | 7.35     | 23.16    |
| WinoGrande | 0.00    | 0.16     | 0.00     | 0.00      | 0.00     | 0.00     | 0.00     | 61.72   | 59.19    | 50.51    | 52.09    | 72.53    | 59.59    | 34.25    |
| ARC-e,     | 0.67    | 0.21     | 11.78    | 0.52      | 0.80     | 0.97     | 0.04     | 65.07   | 41.33    | 22.77    | 28.11    | 29.34    | 48.48    | 3.91     |
| ARC-c      | 0.43    | 0.43     | 9.90     | 0.24      | 0.43     | 0.43     | 0.09     | 46.59   | 23.04    | 19.11    | 24.66    | 23.38    | 34.04    | 3.16     |
| openbookqa | 0.40    | 0.40     | 16.00    | 0.00      | 0.40     | 1.20     | 0.20     | 53.60   | 27.40    | 27.00    | 24.00    | 29.80    | 46.40    | 3.80     |

Table 23: Downstream task performances (%) of Llama2 ( $M_{\rm ed\_ft}$ ) edited on zsRE dataset and then being LoRA or DoRA fine-tuned.

|            |         |          |          | LoRA     |          |          |          |         |          |          | DoRA     |          |          |          |
|------------|---------|----------|----------|----------|----------|----------|----------|---------|----------|----------|----------|----------|----------|----------|
| zsRE       |         |          | MEMIT    | '        | 1        | AlphaEd  | it       |         |          | MEMIT    |          |          | MEMIT    |          |
|            | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ      | 71.16   | 68.87    | 66.09    | 66.09    | 72.97    | 71.01    | 71.04    | 72.45   | 70.87    | 66.94    | 57.45    | 71.10    | 71.50    | 71.59    |
| PIQA       | 83.03   | 82.48    | 76.71    | 76.71    | 83.51    | 82.21    | 81.56    | 83.13   | 82.58    | 76.01    | 67.09    | 83.95    | 83.41    | 82.15    |
| SIQA       | 79.02   | 78.61    | 77.58    | 77.58    | 78.66    | 79.58    | 79.73    | 80.25   | 79.35    | 76.56    | 66.83    | 79.84    | 79.84    | 78.97    |
| HellaSwag  | 90.27   | 87.50    | 61.79    | 61.79    | 90.01    | 88.35    | 87.33    | 89.94   | 88.26    | 58.93    | 67.91    | 91.11    | 88.31    | 85.82    |
| WinoGrande | 83.31   | 82.00    | 78.22    | 78.22    | 83.5     | 83.27    | 81.85    | 84.37   | 83.53    | 77.35    | 70.37    | 82.48    | 83.35    | 81.69    |
| ARC-e,     | 83.71   | 80.39    | 71.55    | 71.55    | 84.85    | 83.25    | 82.28    | 82.49   | 80.41    | 71.59    | 65.12    | 83.63    | 83.80    | 81.61    |
| ARC-c      | 67.66   | 65.96    | 57.25    | 57.25    | 69.62    | 69.71    | 66.98    | 68.00   | 66.57    | 57.51    | 49.78    | 68.69    | 68.77    | 68.17    |
| openbookqa | 80.80   | 77.80    | 73.00    | 73.00    | 80.80    | 80.00    | 80.40    | 80.20   | 80.00    | 72.00    | 64.40    | 83.80    | 81.20    | 80.80    |

Table 24: Downstream task performances (%) of Llama2 ( $M_{\rm ed\_ft}$ ) edited on COUNTERFACT<sup>1</sup> dataset and then being full fine-tuned or not fine-tuned.

|                 |         |          | No fi    | ine-tunii | ng       |          |          |         |          | Full     | fine-tun | ing      |          |          |
|-----------------|---------|----------|----------|-----------|----------|----------|----------|---------|----------|----------|----------|----------|----------|----------|
| CF <sup>1</sup> |         |          | MEMIT    |           |          | AlphaEd  | lit      |         |          | MEMIT    |          |          | MEMIT    |          |
|                 | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$  | $10^{2}$ | $10^{3}$ | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ           | 10.31   | 55.29    | 55.79    | 52.84     | 7.68     | 20.92    | 51.07    | 62.14   | 61.74    | 62.08    | 62.11    | 62.14    | 55.69    | 62.14    |
| PIQA            | 0.16    | 0.33     | 0.22     | 0.2       | 0.16     | 0.11     | 0.11     | 70.84   | 3.54     | 8.87     | 23.88    | 16.27    | 12.57    | 17.85    |
| SIQA            | 2.15    | 4.81     | 0.00     | 0.00      | 2.81     | 2.92     | 0.05     | 61.26   | 49.90    | 56.04    | 26.36    | 33.16    | 35.41    | 32.91    |
| HellaSwag       | 0.00    | 0.00     | 0.03     | 0.00      | 0.00     | 0.00     | 0.00     | 15.77   | 7.97     | 12.71    | 39.29    | 35.07    | 14.50    | 8.05     |
| WinoGrande      | 0.00    | 0.16     | 0.00     | 0.00      | 0.08     | 0.00     | 0.00     | 61.72   | 51.38    | 54.06    | 57.46    | 45.62    | 15.55    | 47.99    |
| ARC-e,          | 0.67    | 0.38     | 0.13     | 0.83      | 0.72     | 0.80     | 0.04     | 65.07   | 55.81    | 38.22    | 23.78    | 21.63    | 39.86    | 26.39    |
| ARC-c           | 0.43    | 0.26     | 0.00     | 0.26      | 0.34     | 0.85     | 0.26     | 46.59   | 36.26    | 29.18    | 18.43    | 14.93    | 32.25    | 14.16    |
| openbookqa      | 0.40    | 0.00     | 0.00     | 0.00      | 0.20     | 0.60     | 0.00     | 53.60   | 44.20    | 36.60    | 21.80    | 23.40    | 32.20    | 26.80    |

Table 25: Downstream task performances (%) of Llama2 ( $M_{\rm ed\_ft}$ ) edited on COUNTERFACT<sup>1</sup> dataset and then being LoRA or DoRA fine-tuned.

|                 |         | LoRA     |          |          |          |           |          |         | DoRA     |          |          |          |          |          |  |  |
|-----------------|---------|----------|----------|----------|----------|-----------|----------|---------|----------|----------|----------|----------|----------|----------|--|--|
| CF <sup>1</sup> |         |          | MEMIT    |          |          | AlphaEdit |          |         | MEMIT    |          |          | MEMIT    |          |          |  |  |
|                 | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$  | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |  |  |
| BoolQ           | 71.16   | 62.26    | 66.64    | 56.54    | 70.37    | 61.59     | 68.23    | 72.45   | 71.63    | 62.97    | 58.03    | 72.14    | 71.44    | 70.58    |  |  |
| PIQA            | 83.03   | 81.72    | 77.37    | 69.04    | 83.79    | 41.95     | 81.39    | 83.13   | 81.35    | 75.73    | 67.56    | 83.46    | 82.86    | 82.37    |  |  |
| SIQA            | 79.02   | 80.09    | 77.69    | 67.98    | 79.32    | 48.52     | 78.71    | 80.25   | 78.52    | 76.44    | 67.44    | 80.40    | 79.79    | 78.97    |  |  |
| HellaSwag       | 90.27   | 40.50    | 71.14    | 64.32    | 51.07    | 9.41      | 84.70    | 89.94   | 85.39    | 76.71    | 68.20    | 89.77    | 89.00    | 83.71    |  |  |
| WinoGrande      | 83.31   | 83.58    | 78.30    | 72.26    | 83.35    | 59.91     | 81.22    | 84.37   | 84.02    | 77.19    | 70.90    | 82.72    | 81.93    | 80.90    |  |  |
| ARC-e,          | 83.71   | 83.84    | 72.77    | 66.01    | 83.38    | 14.02     | 81.36    | 82.49   | 81.40    | 72.05    | 65.47    | 83.54    | 83.33    | 82.03    |  |  |
| ARC-c           | 67.66   | 68.09    | 57.00    | 48.52    | 68.17    | 13.31     | 67.24    | 68.00   | 66.98    | 58.62    | 49.97    | 68.77    | 68.52    | 68.94    |  |  |
| openbookqa      | 80.80   | 81.80    | 72.80    | 64.80    | 82.00    | 40.20     | 80.20    | 80.20   | 80.00    | 72.40    | 64.80    | 81.20    | 83.00    | 77.80    |  |  |

**GPT-J** GPT-J's performances on downstream tasks after being edited on zsRE dataset are shown in Table 26 (no fine-tune and full fine-tune) and Table 27 (LoRA and DoRA). Performances of cases being edited on COUNTERFACT dataset are shown in Table 28 (no fine-tune and full fine-tune) and Table 29 (LoRA and DoRA)

Table 26: Downstream task performances (%) of GPT-J ( $M_{\rm ed\_ft}$ ) edited on zsRE dataset and then being full fine-tuned or not fine-tuned.

|            |         |          | No f     | ne-tunii | ng        |          | Full fine-tuning |         |          |          |          |          |          |          |
|------------|---------|----------|----------|----------|-----------|----------|------------------|---------|----------|----------|----------|----------|----------|----------|
| zsRE       |         | MEMIT    |          |          | AlphaEdit |          |                  |         |          | MEMIT    |          | MEMIT    |          |          |
|            | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$  | $10^{3}$ | $10^{4}$         | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ      | 56.57   | 31.04    | 36.18    | 35.50    | 31.80     | 27.95    | 8.29             | 33.61   | 51.01    | 40.49    | 30.00    | 47.00    | 22.69    | 21.47    |
| PIQA       | 1.36    | 0.87     | 0.98     | 0.82     | 1.20      | 1.74     | 2.61             | 45.38   | 21.71    | 55.71    | 47.44    | 15.61    | 22.09    | 34.44    |
| SIQA       | 0.41    | 0.46     | 0.97     | 0.75     | 0.31      | 0.56     | 2.35             | 46.47   | 48.16    | 47.34    | 43.96    | 48.62    | 46.42    | 28.66    |
| HellaSwag  | 0.03    | 0.01     | 0.18     | 0.25     | 0.03      | 0.02     | 0.99             | 18.49   | 19.78    | 27.31    | 29.05    | 22.80    | 8.98     | 23.81    |
| WinoGrande | 30.23   | 0.32     | 0.32     | 0.24     | 0.16      | 0.08     | 0.36             | 51.78   | 48.93    | 47.99    | 52.01    | 35.44    | 49.88    | 19.81    |
| ARC-e,     | 1.64    | 1.73     | 2.53     | 2.69     | 1.47      | 2.82     | 4.12             | 42.09   | 43.77    | 43.22    | 34.81    | 36.36    | 39.81    | 23.32    |
| ARC-c      | 1.02    | 1.28     | 1.96     | 1.28     | 1.11      | 1.19     | 3.67             | 32.08   | 32.00    | 31.91    | 28.84    | 27.22    | 30.80    | 21.50    |
| openbookqa | 1.20    | 1.40     | 3.00     | 4.20     | 1.20      | 1.40     | 5.40             | 29.40   | 30.20    | 33.40    | 27.40    | 27.00    | 24.00    | 24.80    |

Table 27: Downstream task performances (%) of GPT-J ( $M_{\rm ed.ft}$ ) edited on zsRE dataset and then being LoRA or DoRA fine-tuned.

|            |         | LoRA     |          |          |          |           |          | DoRA    |          |          |          |          |          |          |
|------------|---------|----------|----------|----------|----------|-----------|----------|---------|----------|----------|----------|----------|----------|----------|
| zsRE       |         |          | MEMIT    |          |          | AlphaEdit |          |         |          | MEMIT    |          | MEMIT    |          |          |
|            | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$  | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |
| BoolQ      | 63.79   | 64.19    | 63.55    | 63.36    | 63.3     | 62.66     | 61.99    | 63.27   | 63.88    | 64.10    | 62.29    | 63.94    | 62.17    | 63.67    |
| PIQA       | 73.61   | 73.88    | 69.64    | 71.27    | 70.73    | 72.03     | 58.54    | 73.94   | 73.07    | 74.65    | 72.47    | 74.16    | 70.40    | 68.99    |
| SIQA       | 73.39   | 58.09    | 65.92    | 66.12    | 70.93    | 69.04     | 61.72    | 73.39   | 73.54    | 74.41    | 73.69    | 73.95    | 67.35    | 70.47    |
| HellaSwag  | 43.86   | 65.88    | 27.27    | 67.56    | 66.61    | 46.10     | 26.19    | 71.00   | 70.84    | 71.36    | 61.69    | 71.80    | 58.10    | 46.80    |
| WinoGrande | 68.75   | 66.38    | 64.33    | 66.85    | 65.67    | 67.56     | 58.56    | 70.24   | 69.69    | 69.14    | 68.75    | 69.38    | 63.85    | 66.46    |
| ARC-e,     | 68.31   | 57.20    | 47.77    | 63.38    | 65.32    | 63.09     | 34.39    | 69.44   | 67.93    | 68.10    | 65.91    | 68.22    | 60.31    | 61.11    |
| ARC-c      | 51.79   | 43.94    | 33.70    | 49.49    | 50.43    | 46.76     | 27.99    | 53.07   | 53.07    | 51.96    | 49.06    | 53.75    | 47.44    | 46.50    |
| openbookqa | 70.40   | 56.80    | 65.21    | 63.80    | 61.80    | 65.20     | 50.60    | 68.60   | 64.40    | 68.00    | 66.80    | 65.60    | 57.20    | 62.00    |

Table 28: Downstream task performances (%) of GPT-J ( $M_{\rm ed.ft}$ ) edited on COUNTERFACT<sup>1</sup> dataset and then being full fine-tuned or not fine-tuned.

|                          |         |          | No fi    | ine-tuniı | ıg       |           |          | Full fine-tuning |          |          |          |          |          |          |  |
|--------------------------|---------|----------|----------|-----------|----------|-----------|----------|------------------|----------|----------|----------|----------|----------|----------|--|
| $\mathbb{C}\mathbb{F}^1$ |         |          | MEMIT    |           |          | AlphaEdit |          |                  |          | MEMIT    |          | MEMIT    |          |          |  |
|                          | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$  | $10^{2}$ | $10^{3}$  | $10^{4}$ | No edit          | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |  |
| BoolQ                    | 56.57   | 21.93    | 40.34    | 57.16     | 19.27    | 4.37      | 9.20     | 33.61            | 38.69    | 38.32    | 30.70    | 26.57    | 44.37    | 33.15    |  |
| PIQA                     | 1.36    | 1.58     | 0.65     | 13.06     | 1.47     | 1.69      | 0.27     | 45.38            | 38.41    | 50.49    | 46.52    | 37.76    | 44.83    | 47.44    |  |
| SIQA                     | 0.41    | 0.36     | 0.90     | 0.46      | 0.67     | 2.10      | 4.09     | 46.47            | 45.75    | 40.53    | 39.15    | 48.11    | 43.40    | 35.06    |  |
| HellaSwag                | 0.03    | 0.01     | 0.33     | 2.68      | 0.03     | 0.41      | 0.07     | 18.49            | 24.72    | 26.84    | 21.51    | 22.66    | 17.55    | 16.54    |  |
| WinoGrande               | 30.23   | 0.16     | 0.32     | 0.00      | 0.39     | 0.08      | 0.00     | 51.78            | 51.38    | 52.01    | 50.83    | 51.14    | 30.54    | 36.86    |  |
| ARC-e,                   | 1.64    | 1.73     | 3.07     | 2.61      | 1.94     | 2.99      | 8.84     | 42.09            | 44.57    | 39.52    | 29.84    | 45.20    | 35.23    | 24.03    |  |
| ARC-c                    | 1.02    | 1.19     | 2.73     | 1.79      | 1.37     | 3.24      | 7.51     | 32.08            | 32.08    | 27.99    | 24.23    | 34.30    | 25.68    | 28.07    |  |
| openbookqa               | 1.20    | 1.80     | 2.40     | 4.60      | 1.80     | 2.40      | 13.00    | 29.40            | 32.40    | 33.40    | 18.40    | 32.60    | 28.40    | 23.00    |  |

Table 29: Downstream task performances (%) of GPT-J ( $M_{\rm ed\_ft}$ ) edited on COUNTERFACT dataset and then being LoRA or DoRA fine-tuned.

|                          |         |          |          | LoRA     |           |          |          | DoRA    |          |          |          |          |          |          |  |
|--------------------------|---------|----------|----------|----------|-----------|----------|----------|---------|----------|----------|----------|----------|----------|----------|--|
| $\mathbb{C}\mathbb{F}^1$ |         |          | MEMIT    |          | AlphaEdit |          |          |         | MEMIT    |          |          | MEMIT    |          |          |  |
|                          | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$  | $10^{3}$ | $10^{4}$ | No edit | $10^{2}$ | $10^{3}$ | $10^{4}$ | $10^{2}$ | $10^{3}$ | $10^{4}$ |  |
| BoolQ                    | 63.79   | 63.73    | 62.63    | 63.55    | 63.67     | 64.13    | 63.21    | 63.27   | 64.77    | 60.80    | 62.51    | 63.67    | 63.61    | 63.55    |  |
| PIQA                     | 73.61   | 74.32    | 73.72    | 73.00    | 73.18     | 70.51    | 68.61    | 73.94   | 74.27    | 72.57    | 71       | 74.65    | 72.14    | 65.72    |  |
| SIQA                     | 73.39   | 73.69    | 71.85    | 73.54    | 73.69     | 74       | 68.78    | 73.39   | 71.24    | 67.35    | 61.98    | 74.21    | 70.52    | 70.01    |  |
| HellaSwag                | 43.86   | 19.94    | 71.68    | 66.25    | 71.92     | 40.97    | 33.79    | 71.00   | 50.59    | 32.08    | 64.8     | 71.95    | 51.81    | 49.20    |  |
| WinoGrande               | 68.75   | 69.77    | 68.59    | 66.38    | 68.11     | 50.2     | 65.04    | 70.24   | 69.06    | 67.88    | 65.59    | 70.80    | 69.22    | 64.09    |  |
| ARC-e,                   | 68.31   | 66.79    | 67.05    | 65.95    | 68.35     | 66.41    | 60.06    | 69.44   | 66.46    | 68.31    | 59.26    | 68.86    | 66.41    | 59.72    |  |
| ARC-c                    | 51.79   | 52.05    | 51.54    | 49.4     | 51.96     | 50.09    | 45.31    | 53.07   | 50.9     | 50.34    | 47.35    | 51.96    | 48.98    | 43.94    |  |
| openbookqa               | 70.40   | 66.00    | 68.00    | 63.20    | 65.80     | 65.60    | 58.20    | 68.60   | 67.00    | 67.20    | 61.20    | 66.40    | 66.80    | 57.80    |  |

#### H REPRODUCIBILITY RELATED

For KE, we follow exactly the same data process method provided by Meng et al. (2023), and below is an example of hyper-parameters used for Llama3.1. For FT, we follow the Liu et al. (2024)'s work.

```
1218
      {
1219
           "alg_name": "MEMIT",
           "model_name": "meta-llama/Llama-3.1-8B",
           "stats dir": "./data/stats",
           "device": 0,
1223
           "layers": [3, 4, 5, 6, 7],
1224
           "clamp_norm_factor": 4,
1225
           "layer_selection": "all",
1226
           "fact_token": "subject_last",
1227
           "v num grad steps": 25,
           "v lr": 5e-1,
1228
           "v_loss_layer": 31,
1229
           "v_weight_decay": 1e-3,
1230
           "kl factor": 0.0625,
1231
           "mom2_adjustment": true,
1232
           "mom2 update weight": 15000,
1233
           "rewrite_module_tmp": "model.layers.{}.mlp.down_proj",
"layer_module_tmp": "model.layers.{}",
1234
           "mlp_module_tmp": "model.layers.{}.mlp",
1236
           "attn_module_tmp": "model.layers.{}.self_attn",
1237
           "ln_f_module": "model.norm",
           "lm_head_module": "lm_head",
1238
           "mom2_dataset": "wikipedia",
1239
           "mom2_n_samples": 100000,
1240
           "mom2_dtype": "float32"
1241
```

# I USE OF LARGE LANGUAGE MODELS

In this work, large language models were used for grammar correction and improving the readability of the manuscript. No part of the technical content, including the research design, experimental implementation, data analysis, or interpretation of results, was generated or influenced by an LLM. The role of the model was strictly limited to polishing sentence structure and ensuring clarity in written English.