

---

# Per-parameter Task Arithmetic for Unlearning in Large Language Models

---

Chengyi Cai<sup>1</sup>, Zesheng Ye<sup>1</sup>, Jiangchao Yao<sup>2</sup>, Jianzhong Qi<sup>1</sup>,  
Bo Han<sup>3</sup>, Xiaolu Zhang<sup>4</sup>, Feng Liu<sup>1</sup>, Jun Zhou<sup>4</sup>

<sup>1</sup>The University of Melbourne    <sup>2</sup>Shanghai Jiao Tong University

<sup>3</sup>Hong Kong Baptist University    <sup>4</sup>Ant Group

fengliu.ml@gmail.com

## Abstract

In *large language model* (LLM) unlearning, private information is required to be removed. Task arithmetic unlearns by subtracting a specific *task vector* (TV)—defined as the parameter difference between a privacy-information-tuned model and the original model. While efficient, it can cause over-forgetting by disrupting parameters essential for retaining other information. Motivated by the observation that each parameter exhibits different importance for forgetting versus retention, we propose a *per-parameter task arithmetic* (PerTA) mechanism to rescale the TV, allowing per-parameter adjustment. These weights quantify the relative importance of each parameter for forgetting versus retention, estimated via *gradients* (i.e., PerTA-grad) or the *diagonal Fisher information approximation* (i.e., PerTA-fisher). Moreover, we discuss the effectiveness of PerTA, extend it to a more general form, and provide further analysis. Extensive experiments demonstrate that PerTA consistently improves upon standard TV, and in many cases surpasses widely used training-based unlearning methods in both forgetting effectiveness and overall model utility. By retaining the efficiency of task arithmetic while mitigating over-forgetting, PerTA offers a principled and practical framework for LLM unlearning.

## 1 Introduction

**Learn launch models** (LLMs) can continually acquire new knowledge [27, 28] through post-training; however, the integration of newly ingested data may raise concerns regarding privacy, intellectual property, or misinformation [19]. Due to their tendency to memorize training data, LLMs may inadvertently disclose sensitive information when queried. LLM unlearning [25, 49] aims to erase the memory of specified entities from LLMs to mitigate such risks, as shown in Figure 1(a).

Some *training-based* LLM unlearning methods achieve forgetting of specific entities (i.e., forget set) by designing carefully crafted unlearning loss functions [53, 11, 49, 47] and incorporating entities to be retained (i.e., retain set) to ensure that unrelated knowledge in the model remains unaffected [23]. Another counterpart, *task arithmetic*, avoids multiple iterative training epochs with extensive data, as illustrated in Figure 1(b), where full model and final model represent LLMs before and after unlearning, respectively. This approach achieves unlearning by subtracting from the full model a specific *task vector* (TV) [13] for the forget set. TV denotes the parameter difference between a model finetuned solely on the forget set (hereafter referred to as the FgtOnly model) and the original pretrained model (hereafter the Origin model).

However, the potential correlation and coupling between the entity to be unlearned and other knowledge may cause the subtracted task vector to also contain changes in parameters crucial for preserving other knowledge, thereby risking excessive forgetting of entities that should be retained. Figure 2(a) takes the task of unlearning 1% of entities from the TOFU [29] dataset as an example. The top 30

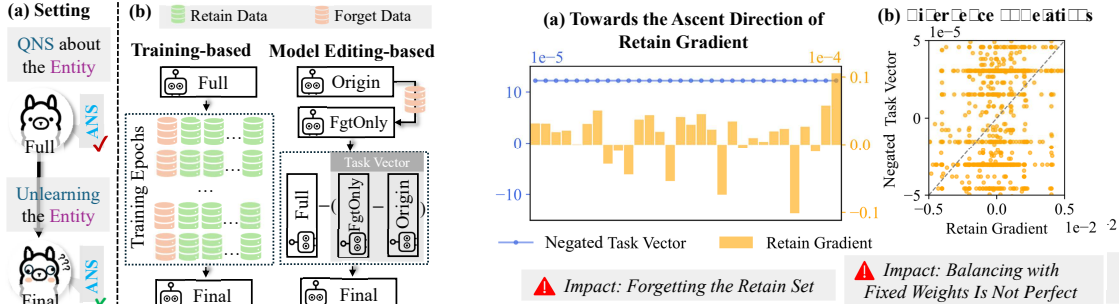


Figure 2: Bottlenecks of task arithmetic methods. (a) illustrates that TV may steer the model direction of the retained gradient, leading to over-forgetting. (b) shows per-parameter divergence of TV–retain gradient relations, rendering the problem non-trivial and not solvable by a uniform weight.

parameters with the largest values in the negated task vector (i.e.,  $-V$ , which is added to the full model  $\theta_{\text{full}}$  to obtain the final model  $\theta_{\text{final}} = \theta_{\text{full}} + (-V)$ ) were selected as examples. For these parameters, we plotted both the values in the negated TV and the gradient magnitudes with respect to the retain set (i.e., the retain gradient) of the same parameters. The results show that, for most of these parameters, the direction indicated by the negated TV aligns with the gradient ascent direction for the retain set. This implies that directly adding the negated TV to the full model would lead to forgetting of the entities that are supposed to be retained. A simple remedy is to add a uniform weight  $\omega \in \mathbb{R}$  satisfying  $0 < \omega < 1$  to TV to reduce the effect of TV (i.e.,  $\theta_{\text{final}} = \theta_{\text{full}} + \omega \cdot (-V)$ ), thereby balancing between unlearning and retaining. However, as shown in Figure 2(b), we find that such an approach may be suboptimal as it ignores per-parameter divergence. By plotting the negated TV and the retain gradient corresponding to different parameters, we observe that different parameters exhibit varying relations between TV and retain gradients, requiring a more sophisticated paradigm.

After formulating the problem, in Section 3, we propose the *Per-parameter Task Arithmetic* (PerTA) mechanism as a solution to the aforementioned bottleneck. PerTA assigns different weights to each parameter in TV and performs a per-parameter multiplication (i.e.,  $\theta_{\text{final}} = \theta_{\text{full}} + W \odot (-V)$ ) to flexibly control the magnitude of editing, where  $W$  is a matrix with the same size as  $\theta_{\text{full}}$ . Parameters that are more pivotal for unlearning can't be assigned higher weights, whereas those crucial for retention receive lower weights, aiming to facilitate both unlearning and retention.

In Section 4, we detail how per-parameter weights are estimated using absolute *gradients* (which captures the importance of parameters given forget or retain sets, abbreviated as PerTA-grad) or the diagonal *Fisher information approximation* (which reflects the sensitivity of parameters to forget or retain sets, abbreviated as PerTA-fisher). We analyze its effectiveness by defining *retain-forget ratio* for parameters. We also extend PerTA to a generalized form and provide further discussion.

In Section 5, we evaluate PerTA on two commonly used unlearning benchmarks TOFU [29] and MUSE [35] across multiple metrics. Results show that PerTA not only substantially outperforms its baseline, vanilla TV, but also exceeds the performance of several mainstream training-based unlearning methods. Training time analysis confirms that PerTA is efficient, while qualitative results illustrate its ability to retain knowledge upon effective unlearning.

PerTA extends TV by preserving retention while enabling effective unlearning, all with high efficiency. Remarkably, it achieves performance surpassing several training-based unlearning methods, highlighting its practical effectiveness. Beyond empirical gains, PerTA offers a new task-arithmetic perspective for LLM unlearning research and introduces a flexible approach for balancing modification and retention in LLM task arithmetic.

## 2 Related Works

**LLM Unlearning.** Machine unlearning [2, 37, 3, 12, 26] aims to selectively remove some previously acquired knowledge from a model while preserving its overall utility. LLM unlearning has attracted

increasing attention, playing a vital role in correcting misinformation, mitigating biases, and protecting privacy [10, 48, 14]. Recent studies on LLM unlearning have advanced this field from multiple perspectives, including benchmarks [29, 35, 21], frameworks [8], evaluation protocols [43, 42], methodological innovations [16, 32, 18], and hallucination mitigation [34, 51]. Different objectives and problem settings of unlearning are discussed in Appendix A.5

Among training-based unlearning methods, GA [49] is the pioneering work that minimizes the log-likelihood of the entities to be unlearned. GD [23] improves it by incorporating the loss on a retain set to mitigate forgetting. NPO [53] constructs its loss function by separating the dis-preferred component from DPO [33], while SimNPO [11] further removes the reliance on reference models. GRU [44] projects the unlearning gradient onto the orthogonal space of retain gradients, and SatImp [47] reweights the loss on a token-wise basis. 35 introduces TV [13] into the unlearning setting. Despite the rapid progress of training-based methods, challenges remain in terms of time and data efficiency, motivating the exploration of more efficient alternatives such as task arithmetic. Since these methods are currently underexplored, we aim to investigate the potential of task arithmetic-based methods.

**Model Merging.** Model merging, also referred to as model editing, is a **cost-demanding** approach that directly manipulates the weight space of multiple pretrained models. 13 introduces the concept of TV, defined as the difference between a finetuned model on a given task and its original counterpart, which can then be used for subsequent model merging. 31 further investigates the fundamental mechanisms of TV by analyzing linearized models. AdaMerging [46] improves upon the TV framework by learning task-wise or layer-wise coefficients, enabling more effective multi-task learning. Additional refinements include trimming [50], sign selection [45] before merging, and composing parameter blocks [52] or models [20] with learned coefficients.

Recently, model merging has been successfully extended to LLMs [56, 40, 41] and multimodal LLMs [5, 9]. Within the context of LLMs, MetaGPT [56] employs a task arithmetic approach that exploits the local linearity of LLMs together with the approximate orthogonality of TVs. FuseLLM [40] and FusionChat [41] investigate strategies for integrating multiple pretrained LLMs in the parameter space to obtain a more potent model. While existing studies have primarily focused on multi-task learning scenarios, our paper explores model merging in LLM unlearning, along with potential improvements. Unlike other model merging methods that combine knowledge, we study task arithmetic in this paper to remove knowledge from the pretrained models.

### 3 Preliminaries and Insights

We consider a pretrained auto-regressive LLM parameterized by  $\theta_0$  with self-attention structures [24]. In the post-training phase, the LLM can be finetuned on new knowledge  $\mathcal{D} = \{s^1, s^2, \dots, s^{|\mathcal{D}|}\}$  consisting of  $|\mathcal{D}|$  sequences, where each sequence  $s = [t_1, t_2, \dots, t_{|s|}]$  contains  $|s|$  tokens. Denoting  $t^{<i}$  as the subsequence of  $s$  from  $t_1$  to  $t_{i-1}$ , the probability of  $s$  given parameter  $\theta$  can be defined as  $p(s; \theta) \triangleq \prod_{i=1}^{|s|} p(t_i | t^{<i}; \theta)$ , which is the product of the conditional probabilities of all tokens. Then  $\theta$  can be learned by minimizing the negative log likelihood loss:

$$\mathcal{L}(\mathcal{D}; \theta) = -\frac{1}{|\mathcal{D}|} \sum_{s \in \mathcal{D}} \log p(s; \theta). \quad (1)$$

Given a new target knowledge set  $\mathcal{D}_{\text{full}}$ , the finetuned model  $\theta_{\text{full}}$  on the whole dataset (i.e., the full model) can be obtained by the training objective  $\arg \min_{\theta \in \Theta} \mathcal{L}(\mathcal{D}_{\text{full}}; \theta)$ .

**LLM Unlearning.** Let  $\mathcal{D}_f = \{s_f^1, s_f^2, \dots, s_f^{|\mathcal{D}_f|}\}$  be the undesirable set that is to be unlearned from  $\theta_{\text{full}}$  (i.e., forget set), where  $\mathcal{D}_f \subset \mathcal{D}_{\text{full}}$  and the size typically satisfies  $|\mathcal{D}_f| \ll |\mathcal{D}_{\text{full}}|$ , we can define the retain set as  $\mathcal{D}_r = \mathcal{D}_{\text{full}} \setminus \mathcal{D}_f$  to be the set of knowledge to be preserved (i.e., retain set). Accordingly, the goal of unlearning is to derive a model  $\theta_{\text{final}}$  that satisfies two desiderata [29, 35]: (a) it *forgets* the information contained in  $\mathcal{D}_f$ , such that the model no longer provides correct answers or statements pertaining to those entities; and (b) it *preserves* the knowledge in  $\mathcal{D}_r$ , ensuring that the corresponding entities remain unaffected. Ideally, the unlearned model should closely approximate the ground-truth model obtained by finetuning exclusively on  $\mathcal{D}_r$ .

**Unlearning via Task Arithmetic.** In the context of unlearning, applying task arithmetic entails computing the TV [13] corresponding to the forget set and subsequently subtracting it from the model  $\theta_{\text{full}}$ . First, a forget-only finetuned model (i.e., FgtOnly model  $\theta_{\text{fgt}}$ ) is obtained on  $\mathcal{D}_f$  using the




Figure 3: The framework of PerTA. PerTA rescales vanilla TV with per-parameter weights. After a one-time gradient computation on forget and retain sets, the per-parameter importance estimation introduced in Section 4.1 can be used to estimate the relative importance of each parameter on the forget set, either using the gradient or the Fisher information, thereby yielding the weights.

original pretrained model  $\theta_0$  by optimizing the objective in Eq.(1), namely,  $\arg \min_{\theta \in \Theta} \mathcal{L}(\mathcal{D}_f; \theta)$ . Then the unlearned model  $\theta_{\text{final}}$  can simply be obtained through arithmetic operations with

$$\theta_{\text{final}} = \theta_{\text{full}} + \underbrace{[-(\theta_{\text{fgt}} - \theta_0)]}_{\text{Task Vector}}, \quad (2)$$

where  $\theta_0$  is used as the reference point for a purer forget-only TV (being slightly different from 35 which uses  $\theta_{\text{full}}$ ). To address the issue of excessive forgetting on the retain set illustrated in Figure 2(a), an intuitive approach is to introduce a constant uniform weight  $0 < \omega < 1$  to adjust the magnitude of the TV, i.e.,  $\theta_{\text{final}} = \theta_{\text{full}} + \omega \cdot [-(\theta_{\text{fgt}} - \theta_0)]$ , thereby balancing between forgetting and retention. However, as shown in Figure 2(b), since the retain gradients and the TV do not exhibit a consistent relationship across parameters, this intuitive approach may overlook divergence across parameters and is insufficient to satisfy both forgetting and retention objectives.

**Per-parameter Task Arithmetic (PerTA).** To address these bottlenecks, we naturally propose a per-parameter weighted mechanism for TV in this work. Since each parameter contributes differently to the forget set and the retain set, we rescale TV by introducing per-parameter weights  $W$ , with the same dimensionality as  $\theta$  (i.e.,  $\dim(W) = \dim(\theta)$ ). The unlearned model is therefore obtained as:

$$\theta_{\text{final}} = \theta_{\text{full}} + W \odot [-(\theta_{\text{fgt}} - \theta_0)], \quad (3)$$

where  $\odot$  represents per-parameter multiplication. In  $W$ , larger values highlight parameters crucial for unlearning the forget set, while smaller values emphasize those important for retaining the retain set, enabling a flexible trade-off between forgetting and retention. Given the immense parameter scale of LLMs, the learning of a parametric  $W$  would be prohibitively expensive. We therefore adopt a non-parametric approach to estimate  $W$ .

## 4 Method

The framework of PerTA and its difference from vanilla TV are shown in Algorithm 1 (violet) and Figure 3. PerTA flexibly scales TV via per-parameter multiplication between  $W$  (in Eq.(3)) and TV. Each entry of  $W$  quantifies the relative importance of its corresponding parameter for the forget set versus the retain set. To this end, we compute parameter gradients with respect to both the forget and retain sets (once each, with minimal overhead) and use them to construct  $W$  (detailed in Section 4.1). Moreover, we analyze the effectiveness of PerTA, extend it to a general form, and discuss its validity in Section 4.2.

---

### Algorithm 1 Pipeline of PerTA

---

**Input:** Origin/Full model  $\theta_0 | \theta_{\text{full}}$ , forget/retain set  $\mathcal{D}_f | \mathcal{D}_r$ , hyperparameter  $E, \alpha$   
**Output:** Unlearned model  $\theta_{\text{final}}$   
# Step 1: Calculating  $\theta_{\text{fgt}}$  required by TV  
 $\theta_{\text{fgt}} \leftarrow \theta_0$   
**for**  $e = 1, \dots, E$  **do**  
     $\theta_{\text{fgt}} \leftarrow \theta_{\text{fgt}} - \alpha \nabla \mathcal{L}(\mathcal{D}_f; \theta_{\text{fgt}})$   
**end for**  
# Step 2.1 One-time gradient computation  
 $g_f \leftarrow \nabla \mathcal{L}(\mathcal{D}_f; \theta_0), g_r \leftarrow \nabla \mathcal{L}(\mathcal{D}_r; \theta_0)$   
# Step 2.2 Per-parameter importance estimation  
 $W \leftarrow \frac{|g_f|^{\tau} + \epsilon}{|g_f|^{\tau} + |g_r|^{\tau} + 2\epsilon}$  (using Eq.(4) or Eq.(5))  
# Step 3: Task arithmetic  
 $\theta_{\text{final}} \leftarrow \theta_{\text{full}} + W \odot [-(\theta_{\text{fgt}} - \theta_0)]$  (using Eq.(3))

---

#### 4.1 Per-parameter Importance Estimation

Let  $W = [w_1, w_2, \dots, w_n]$  be the scaling weights corresponding to the model parameters  $\theta = [q_1, q_2, \dots, q_n]$  with  $n$  parameters. Each weight satisfies  $w_i \in [0, 1]$ ,  $1 \leq i \leq n$ . Values  $w_i$  closer to 1 indicate that TV at  $q_i$  should be kept, while values approaching 0 downweight TV at  $q_i$ .

**Using Absolute Gradient (PerTA-grad).** Since the importance is independent of gradient direction, the absolute magnitude of the parameter gradients [54, 7] provides a natural measure of importance. While gradient estimation using either  $\theta_0$  or  $\theta_{\text{full}}$  is justifiable, we adopt  $\theta_0$  here because  $\theta_0$  is a more neutral initialization model that does not contain training data from the forget or the retain set. However, in practice, estimating gradients on  $\theta_0$  or on  $\theta_{\text{full}}$  makes a negligible difference (see Appendix C.3 for a detailed discussion). Let  $\nabla\mathcal{L}(\mathcal{D}_f; \theta_0)$ ,  $\nabla\mathcal{L}(\mathcal{D}_r; \theta_0)$  be the gradients of forget and retain sets. The weight for each parameter can be computed as the relative contribution of the forget set gradient to the total gradient magnitude, where  $W$  can be formulated as:

$$W_{\text{grad}} = \frac{|\nabla\mathcal{L}(\mathcal{D}_f; \theta_0)| + \epsilon}{|\nabla\mathcal{L}(\mathcal{D}_r; \theta_0)| + |\nabla\mathcal{L}(\mathcal{D}_f; \theta_0)| + 2\epsilon}, \quad (4)$$

where  $\epsilon$  is a small constant to avoid division by zero. Substituting Eq.(4) into Eq.(3) yields the final unlearned model.  $W_{\text{grad}}$  treats all deviations linearly. Next, we also propose a non-linear alternative.

**Using Diagonal Fisher Information Approximation (PerTA-fisher).** The diagonal of the Fisher Information Matrix [30, 1] is widely used to reflect the sensitivity of parameters to the data. The computation of its diagonal entries can be simplified as the squared gradients (see Appendix B.1 for a detailed proof). Accordingly,  $W$  can also be expressed as:

$$W_{\text{fisher}} = \frac{\nabla\mathcal{L}^2(\mathcal{D}_f; \theta_0) + \epsilon}{\nabla\mathcal{L}^2(\mathcal{D}_r; \theta_0) + \nabla\mathcal{L}^2(\mathcal{D}_f; \theta_0) + 2\epsilon}. \quad (5)$$

Similar to  $W_{\text{grad}}$ , substituting Eq.(5) into Eq.(3) yields the final unlearned model, as detailed in Algorithm 1. Both  $W_{\text{grad}}$  and  $W_{\text{fisher}}$  essentially estimate the per-parameter importance of the forget set by computing the relative magnitude of gradients on  $\mathcal{D}_f$ . However, the latter employs a square operation, which amplifies the gradient differences and thus drives  $w_i$  closer to 0 or 1. The detailed discussion is in the next subsection.

#### 4.2 Discussion and A General Form

**Discussion about PerTA.** Denoting  $g_f \triangleq \nabla\mathcal{L}(\mathcal{D}_f; \theta_0)$ ,  $g_r \triangleq \nabla\mathcal{L}(\mathcal{D}_r; \theta_0)$ , we further explore the effectiveness of  $W_{\text{grad}}$  and  $W_{\text{fisher}}$ , and their difference.

**Definition 1.** (Retain-forget ratio). For a parameter  $q_i$  in an LLM, the retain-forget ratio reflects its relative importance for retention versus forgetting. Denoting  $[g_r]_i$  and  $[g_f]_i$  to be the gradients of  $q_i$  on the retain and forget sets, the retain-forget ratio can be represented as

$$r_i = (|[g_r]_i| + \epsilon) / (|[g_f]_i| + \epsilon).$$

When  $r_i \geq 1$ , the retain set dominates for  $q_i$ , and the forget set dominates when  $r_i < 1$ .

Hence, we obtain the following proposition.

**Proposition 1.** For parameter  $p_i$ , we denote its corresponding weights calculated with PerTA-grad, PerTA-fisher to be  $\omega_i^{\text{grad}}$  and  $\omega_i^{\text{fisher}}$  respectively. Then we have:

$$\frac{1}{2} \geq \omega_i^{\text{grad}} \geq \omega_i^{\text{fisher}} \geq 0, \text{ when } r_i \geq 1; \frac{1}{2} < \omega_i^{\text{grad}} < \omega_i^{\text{fisher}} < 1, \text{ when } r_i < 1,$$

which is proved in Appendix B.2.

This implies that when  $p_i$  has a larger influence on the retain set, PerTA applies a smaller reweighting to TV in order to reduce forgetting, and vice versa. Notably, compared to PerTA-grad, PerTA-fisher yields weights that are closer to 0 or 1, thereby creating a cleaner separation between parameters to be edited and to be preserved.

**A General Form.** Besides, the determination of  $W$  is not limited to the aforementioned approaches. Here, we express  $W$  in a more general form—as a function of  $g_f$  and  $g_r$ :

$$W_{\text{general}} = f_{\text{opt}}(g_f, g_r),$$

where  $f_{\text{oprt}}(\cdot, \cdot)$  is a custom operation. Then both  $W_{\text{grad}}$  and  $W_{\text{fisher}}$  can be represented with  $f_{\text{oprt}}(A, B) = |A|^{\circ\tau} / (|A|^{\circ\tau} + |B|^{\circ\tau})$ , where  $\circ\tau$  is the per-parameter  $\tau$ -th power and the case  $\tau = 1$  and  $\tau = 2$  correspond to  $W_{\text{grad}}$  and  $W_{\text{fisher}}$ , respectively.

**Discussion about  $f_{\text{oprt}}(\cdot, \cdot)$ .** In addition to the absolute gradient and diagonal Fisher Information approximation we applied, other operations—such as the SoftMax-based formulation  $f_{\text{oprt}}(A, B) = \exp(|A|) / (\exp(|A|) + \exp(|B|))$ —can also be employed (see Section 5 for detailed results and discussions). Moreover,  $W_{\text{general}}$  subsumes more general cases: when  $f_{\text{oprt}}(A, B) = 1$ , it degenerates to vanilla TV, whereas when  $f_{\text{oprt}}(A, B) = w$ , PerTA employs the uniform constant  $w$  to balance forgetting and retaining. Denoting weight  $w_i$  for parameter  $q_i$  as  $w_i = [f_{\text{oprt}}(g_f, g_r)]_i$  and the corresponding gradients are  $[g_f]_i$  and  $[g_r]_i$ , in the following, we discuss the design of  $f_{\text{oprt}}(\cdot, \cdot)$ :

- Intuitively, it should satisfy  $[f_{\text{oprt}}(g_f, g_r)]_i \rightarrow 1$  when  $|[g_f]_i| \gg |[g_r]_i|$ , and  $[f_{\text{oprt}}(g_f, g_r)]_i \rightarrow 0$  when  $|[g_r]_i| \gg |[g_f]_i|$ . This is because TV is the vector for forget set  $\mathcal{D}_f$ : when  $|[g_f]_i|$  is large, the parameter  $q_i$  is crucial for unlearning, and thus the rescaled TV should preserve its value; conversely, when  $|[g_r]_i|$  is large, the parameter is critical for retention, and the TV should therefore be scaled down.  $W_{\text{grad}}$  and  $W_{\text{fisher}}$  are consistent with this intuition (see Appendix B.3).
- Empirically, we explored several straightforward ways of designing  $f_{\text{oprt}}(\cdot, \cdot)$  and found that  $W_{\text{grad}}$  and  $W_{\text{fisher}}$  in this paper perform best among them, as detailed in the ablation studies from Section 5. Naturally, the choice of  $f_{\text{oprt}}(\cdot, \cdot)$  is not unique, and we hope our work will inspire further exploration and discussion.

## 5 Experiments

**Baselines and Benchmarks.** Experiments are conducted on the widely used unlearning benchmark TOFU [29] (covering three tasks with 1%, 5%, and 10% of the data unlearned) and on MUSE News [35]. On TOFU, following 8, we employ Llama-3.2 1B and 3B Instruct models [38] and evaluate them using five metrics: (1) Forget Quality (FQ) [29], which measures the effectiveness of unlearning (higher is better, we use log transformation in this paper); (2) Model Utility (MU) [29], which quantifies the model’s usefulness in retaining original knowledge (higher is better); (3) Extraction Strength (ES) [4] of the forget set, defined as the proportion of repeated content start positions in the forget set (lower is better); (4) ES of the retain set, defined analogously on the retain set (higher is better); (5) Gibberish (Gib), which represents the probability—determined by a binary classifier [17]—that answers to forget-set queries are non-gibberish (higher is better) and (6) ROUGE-L (ROUGE) [22], the proportion of the longest common sub-sequence between the ground truth and the answers. Additional dataset-related information is provided in Appendix A.1, while detailed definitions of the metrics are given in Appendix A.2.

As for the baselines, for training-based methods we evaluate the mainstream approaches GA [49], GD [23], NPO [53], and NPO+ (NPO combined with GD). For task-arithmetic methods, we test vanilla TV [13] and our proposed method. In addition, we report the metrics of the full model before unlearning, alongside those of a ground-truth model trained solely on the retain set [29], as references. Detailed information about the baselines and implementation can be found in Appendices A.3 and A.4, respectively.

**Performance Comparison.** The results of FQ and MU on the three TOFU tasks with 1%, 5%, and 10% samples to be unlearned are shown in Figure 4 (see more metrics in Appendix C.1). The ground-truth results are shown as black pentagram markers. The FQ metric measures the p-value of distributional differences from the ground truth; we perform the logarithmic transformation to better highlight variations. Dark-blue and purple circles denote methods PerTA-grad and PerTA-fisher, respectively. On simpler tasks (e.g., unlearning 1% of the data), most training-based methods maintain model utility, while the task-arithmetic method TV achieves higher FQ but at the cost of MU. Our PerTA-grad and PerTA-fisher improve MU relative to TV and yield results closer to the ground truth. On more challenging tasks (e.g., unlearning 5% or 10%), training-based methods degrade: MU for GA and NPO drops to nearly zero, and their FQ becomes both lower and unstable (with larger variance). In contrast, our PerTA-grad and PerTA-fisher consistently outperform both training-based and task-arithmetic baselines in FQ and MU, confirming the effectiveness of PerTA in achieving unlearning while preserving model capability.




Figure 4: MU and FQ results of different methods on TOFU (using Llama-3.2 1B Instruct), where circle markers denote values and horizontal and vertical bars at circle centers represent error bars.




Figure 5: Four-dimension ROUGE results of task arithmetic-based methods on TOFU (using Llama-3.2 1B Instruct). Ground-truth results on forget and retain sets are marked with a gray background.

To examine why PerTA outperforms TV among task arithmetic methods, we evaluate four dimensions: *forget*, *retain*, *real*, and *facts* [29], corresponding to the forget set, retain set, original authors, and world facts. The first two measure forgetting and retention of post-training knowledge, while the latter two assess preservation of pretrained knowledge. ROUGE is used to measure similarity to reference answers. As shown in Figure 5 (see Appendix C.1 for more results), TV performs well on real authors and world facts, and PerTA preserves this capability. However, for post-training knowledge, TV suffers from over-forgetting, whereas PerTA significantly narrows the gap to the ground truth. In more challenging settings (e.g., unlearning 5% and 10%), TV falls far below the ground truth, while PerTA nearly doubles TV’s performance, being much closer to the reference.

**More Backbones and Benchmarks.** Table 1 reports the average results of different methods across the three unlearning tasks, with both 1B and 3B model sizes considered to examine the effect of different LLM backbones (see complete results in Appendix C.2). The results show that the baseline TV, compared with training-based methods, suffers from excessive forgetting on the retain set (low  $ES(\mathcal{D}_r)$ ), while our PerTA substantially improves  $ES(\mathcal{D}_r)$  without significantly reducing  $ES(\mathcal{D}_f)$ . Moreover, PerTA delivers notable gains in FQ (e.g., with the ground truth being 0, TV achieves  $-6.174$  and  $-5.284$ , whereas PerTA-grad reaches about  $-0.67$  and PerTA-fisher about  $-1$ ) and MU (e.g., on the 1B model, PerTA raises MU from  $0.495$  to  $0.556$  by PerTA-grad or  $0.562$  by PerTA-fisher, narrowing the gap to the ground truth to within  $0.04$ ). On larger backbones such as 3B, PerTA maintains improvements in both FQ and MU while further increasing  $ES(\mathcal{D}_r)$  without compromising  $ES(\mathcal{D}_f)$ . These results demonstrate the effectiveness of PerTA in achieving unlearning while preserving utility across different model scales. Additionally, results in Appendix C.2 show that PerTA is also effective on MUSE.

**Ablation (General Form) Studies.** Figure 6 shows the curves of FQ when different  $f_{\text{opt}}(\cdot, \cdot)$  are selected under varying hyperparameters. In addition to PerTA-grad and PerTA-fisher proposed in Eq.(4) and Eq.(5), we consider several straightforward designs: (1) ‘Pruning’: removing (i.e.,  $f_{\text{opt}}(A, B) = 0$ ) the  $\lambda\%$  smallest weights in TV to mitigate over-forgetting and maintain others (i.e.,  $f_{\text{opt}}(A, B) = 1$ ), where  $\lambda = 0$  reduces to vanilla TV; (2) ‘Random’: setting weights in  $W$  to random values uniformly sampled between 0 and 1 with  $f_{\text{opt}}(A, B) = \text{rand}([0, 1])$ ; (3) ‘Weighted’:

Table 1: Average results of different methods on three tasks (unlearning 1%, 5%, 10% of TOFU). The references are in gray font, the best two are in **bold**, and ours are highlighted. ‘Full’ and ‘GT’ represent the model before unlearning and the ground truth model.

	FQ↑	MU↑	ES( $\mathcal{D}_t$ )↓	ES( $\mathcal{D}_r$ )↑	Gib↑	FQ↑	MU↑	ES( $\mathcal{D}_t$ )↓	ES( $\mathcal{D}_r$ )↑	Gib↑
Model Size	1B					3B				
Full	-11.808	0.599	0.726	0.737	0.871	-13.960	0.666	0.899	0.884	0.868
GT	0.000	0.596	0.064	0.748	0.894	0.000	0.657	0.066	0.887	0.887
Training-based										
GA	-81.114	0.199	0.086	0.244	0.484	-81.257	0.383	0.125	0.331	0.593
GD	-8.720	0.491	0.112	0.295	0.789	-13.959	0.589	0.192	0.437	0.677
NPO	-4.842	0.198	0.086	0.246	0.592	-4.508	0.380	0.123	0.334	0.637
NPO+	-3.528	0.493	0.122	0.316	0.911	-5.413	0.587	0.147	0.415	0.898
Arithmetic-based										
TV	-6.174	0.495	<b>0.059</b>	0.207	<b>0.914</b>	-5.284	0.612	<b>0.058</b>	0.304	<b>0.921</b>
PerTA-grad	<b>-0.686</b>	<b>0.556</b>	<b>0.072</b>	<b>0.376</b>	<b>0.915</b>	<b>-0.669</b>	<b>0.664</b>	<b>0.082</b>	<b>0.563</b>	<b>0.913</b>
PerTA-fisher	<b>-0.867</b>	<b>0.562</b>	0.080	<b>0.414</b>	0.908	<b>-1.211</b>	<b>0.665</b>	0.092	<b>0.613</b>	0.895




Figure 6: Results of FQ (↑) using different  $f_{\text{opt}}$  on two challenging tasks (unlearning 5% and 10% of TOFU, Llama-3.2 1B Instruct). The shaded region indicates the error bounds.

using a constant  $\omega$  to rescale TV with  $f_{\text{opt}}(A, B) = \omega$ , where  $\omega = 1$  reduces to vanilla TV; and (4) ‘SoftMax’: determining  $f_{\text{opt}}(A, B) = \exp(|A|)/(\exp(|A|) + \exp(|B|))$  in the SoftMax form.

Among these, ‘Pruning’ and ‘Weighted’ methods vary with  $\lambda$  or  $\omega$ , as shown in Figure 6. We observe that ‘Pruning’ performs poorly on more challenging tasks (e.g., unlearning 10%), ‘Random’ exhibits very high variance, and ‘Weighted’ can achieve reasonable results when the optimal constant  $\omega$  is chosen but is highly sensitive to the hyperparameter. The ‘SoftMax’ method represents a successful design of  $f_{\text{opt}}(\cdot, \cdot)$ , yet our PerTA-grad and PerTA-fisher still outperform other possible designs.

**Time Efficiency Discussion.** Figure 7 shows the runtime comparison between the best-performing training-based method, NPO+, and our PerTA. Unlike training-based approaches that require repeated iterations, the runtime of PerTA can be decomposed into: the time to obtain  $\theta_{\text{fgt}}$ , the time to compute  $W$ , and the time for task arithmetic, where the latter is negligible. It can be observed that PerTA inherits the advantage of task arithmetic—significantly reducing runtime—and this advantage becomes more pronounced as task complexity increases (i.e., when unlearning larger proportions). Moreover, as shown previously, estimating gradients with only 20% of the data already yields competitive results, suggesting that runtime can be further reduced.

Together, these findings highlight the time efficiency of PerTA. More results are in Appendix C.2.

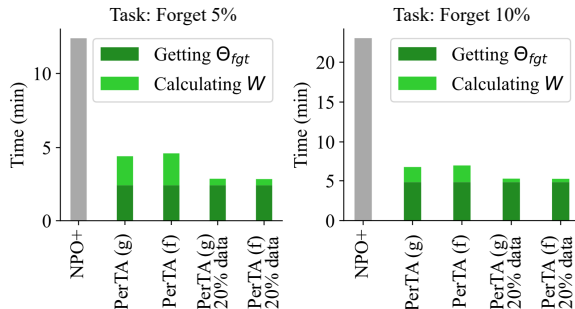


Figure 7: Time comparison of the best-performing training-based method NPO+ and our PerTA (unlearning 5% and 10%, Llama-3.2 1B Instruct).




Figure 8: Results of alternative variants in  $f_{\text{oprt}}(A, B) = |A|^{\tau} / (|A|^{\tau} + |B|^{\tau})$  with different  $\tau$ s (Llama-3.2 1B Instruct). 1%, 5%, 10% tasks are distinguished with different line types.

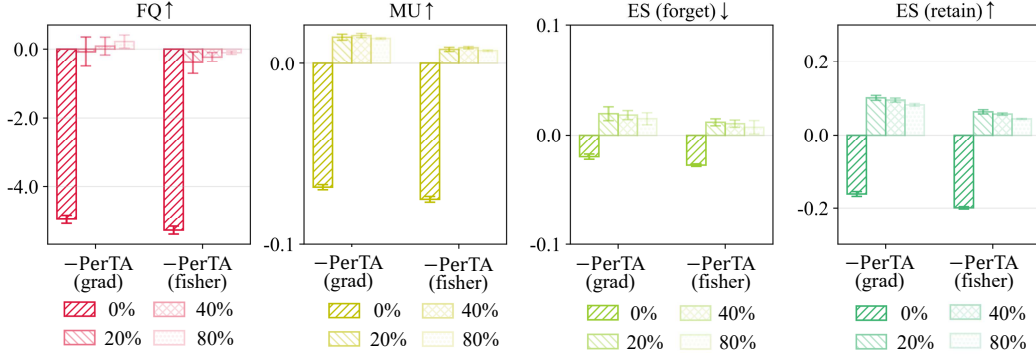


Figure 9: Residual results of the metrics when using only 20%, 40%, and 80% of the samples compared to using the full set (unlearning 5%, Llama-3.2 1B Instruct). 0% denotes vanilla TV.

**Alternative Variants Analysis.** When retaining the form of  $f_{\text{oprt}}(A, B) = |A|^{\tau} / (|A|^{\tau} + |B|^{\tau})$  but not using PerTA-grad and PerTA-fisher, different  $\tau$ s can be applied. We conduct experiments for  $\tau \in \{0, 0.25, 0.5, 1, 2, 4, 8\}$ , with results shown in Figure 8 and Appendix C.2. The cases of  $\tau = 1, 2$  correspond to our PerTA-grad and PerTA-fisher, respectively. Our methods strike a balance between forgetting and retaining: among the different  $\tau$ -based variants, they achieve relatively strong FQ and ES ( $D_f$ ) while keeping MU and ES ( $D_r$ ) at a reasonable level.

**Sample Efficiency Discussion.** In this experiment, we estimate  $W_{\text{grad}}$  and  $W_{\text{fisher}}$  using 0%, 20%, 40%, and 80% of the total samples, where 0% corresponds to vanilla TV and the other three represent PerTA with reduced sample sizes. The differences in metrics compared to using the full dataset are shown in Figure 9 and Appendix C.2. It is observed that using only one-fifth of the samples already yields results comparable to those obtained with the full dataset, and better than vanilla TV. This demonstrates that PerTA is also sample-efficient, which can be beneficial in further reducing the unlearning time.

**More Experiments.** The experimental results comparing gradient  $g_f, g_r$  prediction using  $\theta_0$  or  $\theta_{\text{full}}$  are provided in Appendix C.3. Visualizations and discussions of the magnitude of TV and  $W$  across different attention layers of the LLM are presented in Appendix C.4. Results on larger or alternative LLM models are in Appendix C.5. Results under quantization attacks are detailed in Appendix C.6.

## 6 Conclusion

To address the issue of potentially over-forgetting on the retain set when using vanilla TV, we proposed PerTA to rescale TV, where the weight matrix is estimated using absolute gradients or the diagonal Fisher Information approximation. The effectiveness of PerTA is validated by both theoretical analysis and empirical evidence.

## References

- [1] Shun-ichi Amari, Ryo Karakida, and Masafumi Oizumi. Fisher information and natural gradient learning in random deep networks. In *AISTATS*, 2019.
- [2] Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *IEEE symposium on security and privacy*, 2021.
- [3] Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *IEEE symposium on security and privacy*, 2015.
- [4] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *USENIX security symposium*, 2021.
- [5] Chi Chen, Yiyang Du, Zheng Fang, Ziyue Wang, Fuwen Luo, Peng Li, Ming Yan, Ji Zhang, Fei Huang, Maosong Sun, et al. Model composition for multimodal large language models. In *ACL*, 2024.
- [6] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.
- [7] Rocktim Jyoti Das, Mingjie Sun, Liqun Ma, and Zhiqiang Shen. Beyond size: How gradients shape pruning decisions in large language models. *arXiv preprint arXiv:2311.04902*, 2023.
- [8] Vineeth Dorna, Anmol Mekala, Wenlong Zhao, Andrew McCallum, Zachary C Lipton, J Zico Kolter, and Pratyush Maini. Openunlearning: Accelerating llm unlearning via unified benchmarking of methods and metrics. *arXiv preprint arXiv:2506.12618*, 2025.
- [9] Yiyang Du, Xiaochen Wang, Chi Chen, Jiabo Ye, Yiru Wang, Peng Li, Ming Yan, Ji Zhang, Fei Huang, Zhifang Sui, et al. Adamms: Model merging for heterogeneous multimodal large language models with unsupervised coefficient optimization. In *CVPR*, 2025.
- [10] Chongyu Fan, Jinghan Jia, Yihua Zhang, Anil Ramakrishna, Mingyi Hong, and Sijia Liu. Towards llm unlearning resilient to relearning attacks: A sharpness-aware minimization perspective and beyond. In *ICML*, 2025.
- [11] Chongyu Fan, Jiancheng Liu, Licong Lin, Jinghan Jia, Ruiqi Zhang, Song Mei, and Sijia Liu. Simplicity prevails: Rethinking negative preference optimization for llm unlearning. In *NeurIPS Workshop*, 2024.
- [12] Antonio Ginart, Melody Guan, Gregory Valiant, and James Y Zou. Making ai forget you: Data deletion in machine learning. *NeurIPS*, 2019.
- [13] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *ICLR*, 2023.
- [14] Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. Knowledge unlearning for mitigating privacy risks in language models. *arXiv preprint arXiv:2210.01504*, 2022.
- [15] Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. What does bert learn about the structure of language? In *ACL*, 2019.
- [16] Jinghan Jia, Yihua Zhang, Yimeng Zhang, Jiancheng Liu, Bharat Runwal, James Diffenderfer, Bhavya Kailkhura, and Sijia Liu. Soul: Unlocking the power of second-order optimization for llm unlearning. In *EMNLP*, 2024.
- [17] Madhur Jindal. Gibberish detector: High-accuracy text classification model, 2021.
- [18] Swanand Kadhe, Farhan Ahmed, Dennis Wei, Nathalie Baracaldo, and Inkit Padhi. Split, unlearn, merge: Leveraging data attributes for more effective unlearning in llms. In *ICML Workshop*, 2024.

- [19] Antonia Karamolegkou, Jiaang Li, Li Zhou, and Anders Søgaard. Copyright violations and large language models. *arXiv preprint arXiv:2310.13771*, 2023.
- [20] Sanwoo Lee, Jiahao Liu, Qifan Wang, Jingang Wang, Xunliang Cai, and Yunfang Wu. Dynamic fisher-weighted model merging via bayesian optimization. *arXiv preprint arXiv:2504.18992*, 2025.
- [21] Nathaniel Li, Alexander Pan, Anjali Gopal, Summer Yue, Daniel Berrios, Alice Gatti, Justin D Li, Ann-Kathrin Dombrowski, Shashwat Goel, Long Phan, et al. The wmdp benchmark: Measuring and reducing malicious use with unlearning. *arXiv preprint arXiv:2403.03218*, 2024.
- [22] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 2004.
- [23] Bo Liu, Qiang Liu, and Peter Stone. Continual learning and private unlearning. In *Conference on Lifelong Learning Agents*, 2022.
- [24] Peter J Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198*, 2018.
- [25] Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Yuguang Yao, Chris Yuhao Liu, Xiaojun Xu, Hang Li, et al. Rethinking machine unlearning for large language models. *Nature Machine Intelligence*, 2025.
- [26] Ximing Lu, Sean Welleck, Jack Hessel, Liwei Jiang, Lianhui Qin, Peter West, Prithviraj Ammanabrolu, and Yejin Choi. Quark: Controllable text generation with reinforced unlearning. *NeurIPS*, 2022.
- [27] Xudong Lu, Aojun Zhou, Yuhui Xu, Renrui Zhang, Peng Gao, and Hongsheng Li. Spp: Sparsity-preserved parameter-efficient fine-tuning for large language models. In *ICML*, 2024.
- [28] Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Qingwei Lin, Jian-Guang Lou, Shifeng Chen, Yansong Tang, and Weizhu Chen. Wizardarena: Post-training large language models via simulated offline chatbot arena. *NeurIPS*, 2024.
- [29] Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C Lipton, and J Zico Kolter. Tofu: A task of fictitious unlearning for llms. *arXiv preprint arXiv:2401.06121*, 2024.
- [30] James Martens. New insights and perspectives on the natural gradient method. *JMLR*, 2020.
- [31] Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. *NeurIPS*, 2023.
- [32] Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models as few-shot unlearners. In *ICML*, 2024.
- [33] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *NeurIPS*, 2023.
- [34] William F Shen, Xinchu Qiu, Meghdad Kurmanji, Alex Jacob, Lorenzo Sani, Yihong Chen, Nicola Cancedda, and Nicholas D Lane. Lunar: Llm unlearning via neural activation redirection. *NeurIPS*, 2025.
- [35] Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao Liu, Luke Zettlemoyer, Noah A Smith, and Chiyuan Zhang. Muse: Machine unlearning six-way evaluation for language models. In *ICLR*, 2025.
- [36] Nikolai V Smirnov. On the estimation of the discrepancy between empirical curves of distribution for two independent samples. *Bull. Math. Univ. Moscou*, 1939.
- [37] Ayush K Tarun, Vikram S Chundawat, Murari Mandal, and Mohan Kankanhalli. Fast yet effective machine unlearning. *IEEE TNNLS*, 2023.

- [38] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [39] Jesse Vig and Yonatan Belinkov. Analyzing the structure of attention in a transformer language model. In *ACL Workshop*, 2019.
- [40] Fanqi Wan, Xinting Huang, Deng Cai, Xiaojun Quan, Wei Bi, and Shuming Shi. Knowledge fusion of large language models. In *ICLR*, 2024.
- [41] Fanqi Wan, Ziyi Yang, Longguang Zhong, Xiaojun Quan, Xinting Huang, and Wei Bi. Knowledge fusion of chat llms: A preliminary technical report. *arXiv preprint arXiv:2402.16107*, 2024.
- [42] Qizhou Wang, Bo Han, Puning Yang, Jianing Zhu, Tongliang Liu, and Masashi Sugiyama. Towards effective evaluations and comparisons for llm unlearning methods. In *ICLR*, 2025.
- [43] Qizhou Wang, Jin Peng Zhou, Zhanke Zhou, Saebyeol Shin, Bo Han, and Kilian Q Weinberger. Rethinking llm unlearning objectives: A gradient perspective and go beyond. In *ICLR*, 2025.
- [44] Yue Wang, Qizhou Wang, Feng Liu, Wei Huang, Yali Du, Xiaojiang Du, and Bo Han. Gru: Mitigating the trade-off between unlearning and retention for large language models. *ICML*, 2025.
- [45] Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. Ties-merging: Resolving interference when merging models. *NeurIPS*, 2023.
- [46] Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. In *ICLR*, 2024.
- [47] Puning Yang, Qizhou Wang, Zhuo Huang, Tongliang Liu, Chengqi Zhang, and Bo Han. Exploring criteria of loss reweighting to enhance llm unlearning. In *ICML*, 2025.
- [48] Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. A survey on large language model (llm) security and privacy: The good, the bad, and the ugly. *High-Confidence Computing*, 2024.
- [49] Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. *NeurIPS*, 2024.
- [50] Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *ICML*, 2024.
- [51] Chenlong Zhang, Zhuoran Jin, Hongbang Yuan, Jiaheng Wei, Tong Zhou, Kang Liu, Jun Zhao, and Yubo Chen. Rule: Reinforcement unlearning achieves forget-retain pareto optimality. *arXiv preprint arXiv:2506.07171*, 2025.
- [52] Frederic Z Zhang, Paul Albert, Cristian Rodriguez-Opazo, Anton van den Hengel, and Ehsan Abbasnejad. Knowledge composition using task vectors with learned anisotropic scaling. *NeurIPS*, 2024.
- [53] Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. In *First Conference on Language Modeling*, 2024.
- [54] Zhi Zhang, Qizhe Zhang, Zijun Gao, Renrui Zhang, Ekaterina Shutova, Shiji Zhou, and Shanghang Zhang. Gradient-based parameter selection for efficient fine-tuning. In *CVPR*, 2024.
- [55] Zhiwei Zhang, Fali Wang, Xiaomin Li, Zongyu Wu, Xianfeng Tang, Hui Liu, Qi He, Wenpeng Yin, and Suhang Wang. Catastrophic failure of llm unlearning via quantization. In *ICLR*, 2025.
- [56] Yuyan Zhou, Liang Song, Bingning Wang, and Weipeng Chen. Metagpt: Merging large language models using model exclusive task arithmetic. *arXiv preprint arXiv:2406.11385*, 2024.

## A Appendix 1: More Training Information

### A.1 Dataset Information

**TOFU.** The TOFU dataset<sup>1</sup> is designed as a benchmark to assess how well LLMs can perform unlearning on practical tasks. It contains 4000 question-answer pairs derived from autobiographies of 200 entirely fictional authors, all generated by GPT-4. The task involves evaluating a finetuned model’s ability to unlearn when exposed to different proportions (i.e., unlearning 1%, 5%, 10%) of the forget set.

**MUSE.** MUSE is a benchmark designed to evaluate machine unlearning. It centers on two major forms of textual content where unlearning is often necessary: news reports (News) and literary works (Books). The MUSE-News subset<sup>2</sup> specifically includes BBC articles published after August 2023.

### A.2 Metric Discussion

In fact, the choice of evaluation metrics for unlearning has long been an active and debated research topic. Assessing unlearning performance typically requires considering multiple aspects and dimensions. In this paper, we adopt the metrics used in [29], which are also widely employed by mainstream methods such as [8, 44, 47].

Similar to Section 3, we define the new knowledge dataset to post-train the LLM as  $\mathcal{D} = \{s^1, s^2, \dots, s^{|\mathcal{D}|}\}$  consisting of  $|\mathcal{D}|$  sequences, where each sequence  $s = [t_1, t_2, \dots, t_{|s|}]$  contains  $|s|$  tokens. To split  $s$  into questions and answers, we can also write  $s = [x, y]$ . Then the probability of  $y$  given  $x$  is defined as

$$P(\mathcal{D}; \theta) = \mathbb{E}_{[x,y] \sim \mathcal{D}} P(y|x; \theta)^{\frac{1}{|y|}} = \mathbb{E}_{[x,y] \sim \mathcal{D}} \left[ \prod_{i=1}^{|y|} p(y_i|[x, y_{<i}]; \theta) \right]^{\frac{1}{|y|}},$$

which is normalize for answer length as a common practice [6]. Denoting  $\mathcal{Y}_{\text{pret}}$  as the set of incorrect answers with the same template as  $y$ , the truth ratio can be defined as

$$\text{Tr}(\mathcal{D}; \theta) = \mathbb{E}_{[x,y] \sim \mathcal{D}} \frac{\frac{1}{|\mathcal{Y}_{\text{pret}}|} \sum_{\tilde{y} \in \mathcal{Y}_{\text{pret}}} P(\tilde{y}|x)}{P(y|x)}.$$

Besides, by obtaining  $\arg \max_{t_i} p(t_i|t_{<i}; \theta)$ , the output texts of LLM given prompt  $t_{<i} = [t_1, \dots, t_{i-1}]$  is defined as  $f(t_{<i}; \theta)$ .

**ROUGE-L (ROUGE).** Denoting the length of the longest common sub-sequence considering string  $a$  and  $b$  as  $\text{LCS}(a, b)$ , then the ROUGE-L metric can be defined for model  $\theta$  and dataset  $\mathcal{D}$  as

$$\text{ROUGE}(\mathcal{D}; \theta) = \mathbb{E}_{[x,y] \sim \mathcal{D}} \frac{\text{LCS}(y, f(x; \theta))}{|y|}.$$

The bigger ROUGE-L is, the more similar the references and output answers of LLM are.

**Extraction Strength (ES).** ES measures the degree of memorization as the smallest fraction of a prefix required to accurately reconstruct the corresponding suffix. It can be formulated as

$$\text{ES}(\mathcal{D}; \theta) = \mathbb{E}_{[x,y] \sim \mathcal{D}} \left[ 1 - \frac{1}{|y|} \min_k \{k | f([x, y_{<k}]; \theta) = y_{>k} \} \right].$$

**Forget Quality (FQ).** The goal of unlearning is for the final model to approximate the model trained on the retain set only. Therefore, FQ is used to assess unlearning by statistically comparing the truth ratio  $\text{Tr}(y|x; \theta)$  distributions of the unlearned model  $\theta$  and the model  $\theta_{\text{retain}}$  trained on retain set only with KS-Test [36], producing higher scores when the two distributions are closely aligned:

$$\text{FQ}(\mathcal{D}_f; \theta) = \text{KS}(\text{Tr}_{[x,y] \sim \mathcal{D}_f}(y|x; \theta), \text{Tr}_{[x,y] \sim \mathcal{D}_f}(y|x; \theta_{\text{retain}})),$$

where  $\text{KS}(\cdot, \cdot)$  is the KS-Test function and  $\mathcal{D}_f$  is the forget set.

<sup>1</sup><https://huggingface.co/datasets/locuslab/TOFU>

<sup>2</sup><https://huggingface.co/datasets/muse-bench/MUSE-News>

**Model Utility (MU).** MU measures how well a model performs after unlearning, on both the retain set and general knowledge. It is defined as the harmonic mean of three metrics—probability, ROUGE, and Truth Ratio—evaluated across three levels: retain set  $\mathcal{D}_r$ , real authors  $\mathcal{D}_a$ , and world factual knowledge  $\mathcal{D}_w$ :

$$\text{MU}(\theta) = \frac{1}{\sum_{\mathcal{D} \in \{\mathcal{D}_f, \mathcal{D}_a, \mathcal{D}_w\}} \left[ \frac{1}{P(\mathcal{D}; \theta)} + \frac{1}{\text{Tr}(\mathcal{D}; \theta)} + \frac{1}{\text{ROUGE}(\mathcal{D}; \theta)} \right]}.$$

Different from the retain set, when calculating the probability on  $\mathcal{D}_a$  and  $\mathcal{D}_r$ , function P is defined as  $P(x|y; \theta) = p(y|x; \theta) / \sum_{\tilde{y} \in \mathcal{Y}_{\text{choice}}} p(\tilde{y}|x; \theta)$ , where  $\mathcal{Y}_{\text{choice}}$  is the given possible answer set.

**Gibberish (Gib).** Unlearning can negatively impact model fluency, especially on the forget set, leading to incoherent or meaningless outputs. To measure this phenomenon, a classifier-based score<sup>3</sup> is employed to determine whether the generated text resembles gibberish.

### A.3 Training-based Methods

Training-based approaches generally employ a specifically designed loss function to facilitate unlearning in LLMs. The training procedure involves iteratively computing this loss and updating the model’s weights. After a number of iterations, the process concludes, resulting in the final model. This section details the loss functions used in the training-based methods discussed in this work.

**GA.** GA is the pioneering work that first maximize the loss of the forget set. As the general loss function of LLM learning  $\mathcal{L}(\mathcal{D}; \theta)$  is defined in Eq.(1), the loss of GA can be formulated as

$$\mathcal{L}_{\text{GA}}(\theta) = -\mathcal{L}(\mathcal{D}_f; \theta).$$

**GD.** To avoid over-forgetting the retain set, GD performs gradient descent on the retain set:

$$\mathcal{L}_{\text{GD}}(\theta) = -\mathcal{L}(\mathcal{D}_f; \theta) + \alpha L(\mathcal{D}_r; \theta),$$

where  $\alpha$  is the coefficient to balance between unlearning and retention.

**NPO.** NPO constructs its loss function inspired by the dis-preferred component of DPO. This type of loss is suitable for the question-answer pairs. Thus, the loss function is

$$\mathcal{L}_{\text{NPO}}(\theta) = -\frac{2}{\beta} \mathbb{E}_{[x,y] \sim \mathcal{D}_f} \log \sigma(-\beta \log(\frac{p(y|x; \theta)}{p(y|x; \theta_{\text{full}})})),$$

where  $\sigma(\cdot)$  represents the Sigmoid function and  $\beta$  is the hyper-parameter.

**NPO+.** In this paper, NPO+ is defined as a method combining NPO and GD together for better performance. Namely, the loss function is

$$\mathcal{L}_{\text{NPO+}}(\theta) = -\frac{2}{\beta} \mathbb{E}_{[x,y] \sim \mathcal{D}_f} \log \sigma(-\beta \log(\frac{p(y|x; \theta)}{p(y|x; \theta_{\text{full}})})) - \alpha \mathbb{E}_{[x,y] \sim \mathcal{D}_r} \log p(y|x; \theta),$$

where  $\alpha, \beta$  are hyper-parameters.

### A.4 Implement Details

For a fair and consistent evaluation, all training-based methods are benchmarked using the open-unlearning framework<sup>4</sup>. We experiment with the official Llama 2 7B<sup>5</sup>, Llama-3.2 1B Instruct<sup>6</sup>, and Llama-3.2 3B<sup>7</sup> Instruct models.

Following [29], for all the methods, our training configuration consists of 10 epochs (including one for warm-up), a learning rate of 1e-5, weight decay of 0.01, and a batch size of 32.

In the context of task arithmetic approaches for obtaining the FgtOnly model, we modify these settings for specific datasets: on TOFU, we extend training to 20 epochs to achieve convergence on the forget set; on MUSE, we increase the learning rate to 1e-4, with all other hyperparameters remaining constant. To ensure a fair comparison, all models are subsequently evaluated under the same open-unlearning framework. Experiments are conducted on a single 80G A100 GPU.

<sup>3</sup><https://huggingface.co/madhurjindal/autonlp-Gibberish-Detector-492513457>

<sup>4</sup><https://github.com/locuslab/open-unlearning>

<sup>5</sup><https://huggingface.co/meta-llama/Llama-2-7b-hf>

<sup>6</sup><https://huggingface.co/meta-llama/Llama-3.2-1B-Instruct>

<sup>7</sup><https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct>

## A.5 Objectives and Evaluation of Unlearning

Unlearning is primarily considered as a privacy-preserving task: the aim is to remove information about the entities to be unlearned, so that the model approximates a version trained only on the retain entities [29] (i.e., the ground-truth model). This objective and evaluation framework is the one adopted by current mainstream methods [11, 44, 47], and it is also employed in our paper.

However, as is shown in Figure 13, LLM might generate false answers after unlearning when being questioned with entities in the forget set. Under the evaluation of the aforementioned framework, the false answers are considered acceptable because even the ground-truth model, or the original model, may also produce incorrect responses (i.e., hallucinations). In other words, hallucination may not result from the unlearning process itself, but rather from the supervised finetuning process. Consequently, unlearning aims to bring the unlearned model closer to the retain-only model, and methods are considered successful as long as the outputs are similar to those of the ground-truth model.

Recently, some work [34] has focused on refusing to answer queries about entities to be forgotten without misleading the users. We believe this is also a promising direction for future research. For task-arithmetic methods, reducing false answers for forgotten entities could potentially be achieved in two ways: (1) adding a task vector trained on QA samples with “I don’t know” responses, and (2) addressing hallucinations at the source, i.e., reducing hallucinations in the model before merging. Both approaches are feasible directions for future work.

## B Appendix 2: More Theoretical Justification

### B.1 The Diagonal of the Fisher Information Matrix

*Proof.* We aim to prove that the diagonal of the Fisher Information Matrix (FIM),  $F_{ii}$ , can be approximated by the squared gradient of the loss function, given that the loss is the negative log-likelihood. The  $i$ -th diagonal element of the FIM is defined as the variance of the score, given by:

$$F_{ii} \approx \mathbb{E}_{s \sim \mathcal{D}} \left[ \left( \frac{\partial \log p(s; \theta)}{\partial q_i} \right)^2 \right],$$

where  $q_i$  is a single parameter. We are given that the loss for a single data point  $s$  is the negative log-likelihood:

$$\mathcal{L}(\{s\}; \theta) = -\log p(s; \theta).$$

Taking the partial derivative with respect to a parameter  $q_i$  yields:

$$\frac{\partial \mathcal{L}(\{s\}; \theta)}{\partial q_i} = -\frac{\partial \log p(s; \theta)}{\partial q_i}.$$

Substituting this into the definition of  $F_{ii}$ , we get:

$$F_{ii} \approx \mathbb{E}_{s \sim \mathcal{D}} \left[ \left( -\frac{\partial \mathcal{L}(\{s\}; \theta)}{\partial q_i} \right)^2 \right] = \mathbb{E}_{s \sim \mathcal{D}} \left[ \left( \frac{\partial \mathcal{L}(\{s\}; \theta)}{\partial q_i} \right)^2 \right].$$

Then we arrive at the approximation:

$$F_{ii} \approx \left( \frac{\partial \mathcal{L}(\mathcal{D}; \theta)}{\partial q_i} \right)^2.$$

This demonstrates that the diagonal of the FIM can be estimated by the squared gradient of the negative log-likelihood loss.  $\square$

### B.2 Proof of Proposition 1

*Proof.* For a single parameter  $q_i$  in LLM, we denote its corresponding weights calculated with PerTA-grad, PerTA-fisher to be  $\omega_i^{\text{grad}}$  and  $\omega_i^{\text{fisher}}$  respectively. Using  $r_i = \frac{\|g_i\|_i + \epsilon}{\|g_i^f\|_i + \epsilon}$  for notational

convenience, where  $[g_r]_i$  and  $[g_f]_i$  are the gradients on forget and retain set, we can obtain the following simplified form when  $\epsilon \rightarrow 0$ :

$$\omega_i^{\text{grad}} = \frac{|[g_f]_i| + \epsilon}{|[g_r]_i| + |[g_f]_i| + 2\epsilon} = \frac{1}{r_i + 1},$$

$$\omega_i^{\text{grad}} = \frac{[g_f]_i^2 + \epsilon}{[g_r]_i^2 + [g_f]_i^2 + 2\epsilon} = \frac{1}{r_i^2 + 1}.$$

Depending on the range of  $r_i$ , we have two cases:

- When  $r_i \geq 1$  (where  $|[g_r]_i| \geq |[g_f]_i|$ , retain set dominates), from the simplified form of  $\omega_i^{\text{grad}}$  and  $\omega_i^{\text{fisher}}$ , we can derive that

$$\frac{1}{2} \geq \frac{1}{r_i + 1} \geq \frac{1}{r_i^2 + 1} \geq 0 \Rightarrow \frac{1}{2} \geq \omega_i^{\text{grad}} \geq \omega_i^{\text{fisher}} \geq 0.$$

It reveals that the squared term will push the weight closer to 0 faster than the linear term, offering stronger protection for the retain set.

- When  $r_i < 1$  (where  $|[g_r]_i| < |[g_f]_i|$ , forget set dominates), from the simplified form of  $\omega_i^{\text{grad}}$  and  $\omega_i^{\text{fisher}}$ , we can derive that

$$\frac{1}{2} < \frac{1}{r_i + 1} < \frac{1}{r_i^2 + 1} < 1 \Rightarrow \frac{1}{2} < \omega_i^{\text{grad}} < \omega_i^{\text{fisher}} < 1.$$

It reveals that the squared term will push the weight closer to 1 faster than the linear term, leading the task vector to be applied more fully when needed.

□

Therefore, in some undesirable cases where the gradients on the forget set and the retain set are very similar, PerTA-grad tends to degenerate into a single weight with the value of 0.5. In contrast, PerTA-fisher may suppress such “ambiguous” updates (i.e., weights near 0.5) and create a cleaner separation between parameters to be edited and parameters to be preserved.

### B.3 PerTA-grad and PerTA-fisher Satisfy the Intuitive Rules

*Proof.* Regarding the function  $f_{\text{opr}}(A, B) = |A|^{\circ\tau} / (|A|^{\circ\tau} + |B|^{\circ\tau})$  defined for  $W_{\text{grad}} (\tau = 1)$  and  $W_{\text{fisher}} (\tau = 2)$ , for a single weight  $w_i$ , we have

$$w_i = [f_{\text{opr}}(g_f, g_r)]_i = \frac{|[g_f]_i|^\tau + \epsilon}{|[g_f]_i|^\tau + |[g_r]_i|^\tau + 2\epsilon}, \text{ where } \tau = 1, \text{ or } \tau = 2.$$

Then we prove  $|[g_f]_i| \ll |[g_r]_i| \Rightarrow w_i \rightarrow 0$  and  $|[g_f]_i| \gg |[g_r]_i| \Rightarrow w_i \rightarrow 1$  in the two cases below:

**Case 1:**  $|[g_f]_i| \ll |[g_r]_i|$ . It implies that  $[g_f]_i$  is negligible compared to  $[g_r]_i$ . Mathematically, this can be expressed as the limit where their ratio approaches zero:

$$\frac{|[g_f]_i| + \epsilon}{|[g_r]_i| + \epsilon} \rightarrow 0.$$

Then for  $\tau = 1$  and  $\tau = 2$ , we have:

$$\frac{|[g_f]_i|^\tau + \epsilon}{|[g_r]_i|^\tau + \epsilon} \rightarrow 0.$$

To analyze the limit of  $w_i$ , we can divide both the numerator and the denominator by  $|[g_r]_i|^\tau + \epsilon$  ( $|[g_r]_i|^\tau + \epsilon \neq 0$ ):

$$w_i = \frac{(|[g_f]_i|^\tau + \epsilon) / (|[g_r]_i|^\tau + \epsilon)}{(|[g_f]_i|^\tau + \epsilon) / (|[g_r]_i|^\tau + \epsilon) + (|[g_r]_i|^\tau + \epsilon) / (|[g_r]_i|^\tau + \epsilon)} = \frac{(|[g_f]_i|^\tau + \epsilon) / (|[g_r]_i|^\tau + \epsilon)}{(|[g_f]_i|^\tau + \epsilon) / (|[g_r]_i|^\tau + \epsilon) + 1}.$$




Figure 10: ES (forget) and ES (retain) results of different methods on TOFU (using Llama-3.2 1B Instruct), where circle markers denote values and horizontal and vertical bars at circle centers represent error bars.

Now, we take the limit as  $\frac{||[g_f]_i|^\tau + \epsilon}{||[g_r]_i|^\tau + \epsilon} \rightarrow 0$ :

$$\lim_{\frac{||[g_f]_i|^\tau + \epsilon}{||[g_r]_i|^\tau + \epsilon} \rightarrow 0} w_i = \lim_{\frac{||[g_f]_i|^\tau + \epsilon}{||[g_r]_i|^\tau + \epsilon} \rightarrow 0} \frac{(|[g_f]_i|^\tau + \epsilon) / (|[g_r]_i|^\tau + \epsilon)}{(|[g_f]_i|^\tau + \epsilon) / (|[g_r]_i|^\tau + \epsilon) + 1} = \frac{0}{0 + 1} = 0.$$

Thus, when  $||[g_f]_i| \ll |[g_r]_i|$ , the value of  $w_i$  approaches 0.

**Case 2:**  $||[g_f]_i| \gg |[g_r]_i|$ . Similarly, the condition  $||[g_f]_i| \gg |[g_r]_i|$  implies that  $[g_r]_i$  is negligible compared to  $[g_f]_i$ . This means the ratio of their sizes approaches zero:

$$\frac{||[g_r]_i|^\tau + \epsilon}{||[g_f]_i|^\tau + \epsilon} \rightarrow 0.$$

For this case, we divide both the numerator and the denominator by  $||[g_f]_i|^\tau + \epsilon$  (with  $||[g_f]_i|^\tau + \epsilon \neq 0$ ):

$$w_i = \frac{(|[g_f]_i|^\tau + \epsilon) / (|[g_f]_i|^\tau + \epsilon)}{(|[g_f]_i|^\tau + \epsilon) / (|[g_f]_i|^\tau + \epsilon) + (|[g_r]_i|^\tau + \epsilon) / (|[g_f]_i|^\tau + \epsilon)} = \frac{1}{1 + (|[g_r]_i|^\tau + \epsilon) / (|[g_f]_i|^\tau + \epsilon)}.$$

Now, we take the limit as  $\frac{||[g_r]_i|^\tau + \epsilon}{||[g_f]_i|^\tau + \epsilon} \rightarrow 0$ :

$$\lim_{\frac{||[g_r]_i|^\tau + \epsilon}{||[g_f]_i|^\tau + \epsilon} \rightarrow 0} w_i = \lim_{\frac{||[g_r]_i|^\tau + \epsilon}{||[g_f]_i|^\tau + \epsilon} \rightarrow 0} \frac{1}{1 + (|[g_r]_i|^\tau + \epsilon) / (|[g_f]_i|^\tau + \epsilon)} = \frac{1}{1 + 0} = 1.$$

Thus, when  $||[g_f]_i| \gg |[g_r]_i|$ , the value of  $w_i$  approaches 1.

**Conclusion** We have formally shown through limit analysis that our PerTA-grad and PerTA-fisher satisfy  $||[g_f]_i| \ll |[g_r]_i| \Rightarrow w_i \rightarrow 0$  and  $||[g_f]_i| \gg |[g_r]_i| \Rightarrow w_i \rightarrow 1$ .  $\square$

## C Appendix 3: More Experimental Results

### C.1 More Graphical Results

**ES Metric across Various Tasks.** In this section, we present in Figure 10 the two-dimensional values of the ES metric on the forget and retain sets across the three TOFU tasks, as a supplement to Figure 4. It can be observed that for relatively simple tasks (e.g., unlearning 1%), most methods preserve the retain set but fail to achieve effective forgetting on the forget set. In contrast, our PerTA-grad and PerTA-fisher not only maintain retention but also achieve effective forgetting. For more challenging tasks (e.g., unlearning 5% and 10%), our PerTA methods similarly achieve unlearning that is closest to the ground truth, while still preserving memory on the retain set.

**ROUGE Results on Larger LLM.** Similarly, in Figure 11 we report the ROUGE results of our method compared with vanilla TV on the ‘forget’, ‘retain’, ‘real’, and ‘facts’ sets for the 3B model, as a supplement to Figure 5. The same conclusion as in the main text can be drawn here: while TV effectively preserves the knowledge acquired during the pretraining stage of the original model, it leads to excessive forgetting on the retain and forget datasets. In contrast, our PerTA mitigates the gap between TV and the ground truth on these two datasets, thereby enhancing the performance of




Figure 11: Four-dimension ROUGE results of task arithmetic-based methods on TOFU (using Llama-3.2 3B Instruct). Ground-truth results on forget and retain sets are marked with a gray background.

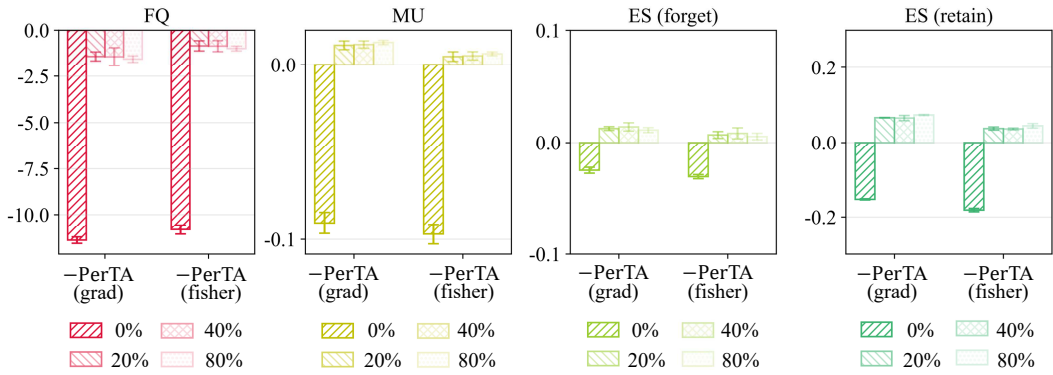


Figure 12: Residual results of the metrics when using only 20%, 40%, and 80% of the samples compared to using the full set (unlearning 10%, Llama-3.2 1B Instruct). 0% denotes vanilla TV.

the task arithmetic-based method for unlearning. This conclusion holds consistently across LLMs of different sizes.

**Sample Efficiency in More Tasks.** Figure 12, as a complement to Figure 9, presents the difference in performance metrics relative to using the entire dataset when unlearning 10% in TOFU with varying data proportions (20%, 40%, 80%) and with 0% data (i.e., vanilla TV). Consistent with the main text, it is observed that using only one-fifth of the samples already achieves results comparable to those obtained with the full dataset, and substantially outperforms vanilla TV. This highlights the sample efficiency of PerTA, which can further reduce computational cost.

## C.2 More Quantitative Results

**Detailed Results on Various Tasks.** Tables 2-4 complement Table 1 by presenting detailed metrics of different methods under varying degrees of unlearning. We observe that in relatively simple tasks with smaller models (e.g., Llama-3.2 1B with 1% unlearning), the advantage of PerTA is not yet pronounced. However, as the task complexity increases, PerTA consistently outperforms on metrics such as FQ and MU, allowing task arithmetic-based approaches to surpass training-based methods. Overall, PerTA demonstrates a clear advantage in both unlearning capability and retention performance.

**Sample Output Discussion.** Figure 13 presents sample responses of different methods on the forget and retain sets of TOFU after unlearning. For the forget set, some methods produce incoherent or irrelevant answers—indicating that the responses lack logical consistency or relevance. For the retain set, other methods may exhibit over-forgetting or generate hallucinated answers. In contrast, PerTA is able to achieve unlearning on the forget set while preserving knowledge on the retain set.

Table 2: Results of different methods on unlearning 1% of TOFU. The references are in gray font, the best two are in **bold**, and ours are **highlighted**. ‘Full’ and ‘GT’ represent the model before unlearning and the ground truth model, respectively.

	FQ↑	MU↑	ES( $\mathcal{D}_f$ )↓	ES( $\mathcal{D}_r$ )↑	Gib↑	FQ↑	MU↑	ES( $\mathcal{D}_f$ )↓	ES( $\mathcal{D}_r$ )↑	Gib↑
Model Size	1B					3B				
Full	-2.170	0.599	0.743	0.737	0.894	-1.845	0.666	0.920	0.884	0.894
GT	0.000	0.599	0.069	0.751	0.874	0.000	0.662	0.067	0.888	0.904
Training-based										
GA	-1.953	<b>0.597</b>	0.189	<b>0.656</b>	<b>0.909</b>	-1.845	0.668	0.252	0.824	0.864
GD	-1.845	0.581	0.169	0.562	0.907	-1.845	0.663	0.320	<b>0.826</b>	0.897
NPO	-2.062	0.595	0.178	0.650	0.904	-1.845	0.668	0.253	<b>0.825</b>	0.838
NPO+	-1.845	<b>0.596</b>	0.174	<b>0.656</b>	0.907	-1.845	0.669	0.254	0.819	0.856
Task Arithmetic-based										
TV	<b>-0.393</b>	0.556	<b>0.081</b>	0.358	0.908	-0.238	0.656	<b>0.075</b>	0.550	<b>0.933</b>
PerTA-grad	<b>-0.289</b>	0.581	<b>0.075</b>	0.551	<b>0.912</b>	<b>-0.037</b>	<b>0.669</b>	<b>0.085</b>	0.757	<b>0.903</b>
PerTA-fisher	-0.576	0.586	0.085	0.600	0.895	<b>-0.238</b>	<b>0.672</b>	0.106	0.803	0.869

Table 3: Results of different methods on unlearning 5% of TOFU. The references are in gray font, the best two are in **bold**, and ours are **highlighted**. ‘Full’ and ‘GT’ represent the model before unlearning and the ground truth model, respectively.

	FQ↑	MU↑	ES( $\mathcal{D}_f$ )↓	ES( $\mathcal{D}_r$ )↑	Gib↑	FQ↑	MU↑	ES( $\mathcal{D}_f$ )↓	ES( $\mathcal{D}_r$ )↑	Gib↑
Model Size	1B					3B				
Full	-11.845	0.599	0.727	0.737	0.858	-13.591	0.666	0.887	0.884	0.850
GT	0.000	0.599	0.063	0.746	0.905	0.000	0.659	0.066	0.874	0.869
Training-based										
GA	-2.415	0.000	<b>0.037</b>	0.039	0.417	-5.856	0.482	0.089	0.135	0.866
GD	-8.831	0.457	0.090	0.171	0.751	-13.232	0.552	0.140	0.244	0.579
NPO	-2.222	0.000	<b>0.048</b>	0.052	0.543	-7.091	0.472	0.080	0.140	0.868
NPO+	-4.260	0.458	0.098	0.139	0.882	-7.352	0.545	0.100	0.200	0.911
Task Arithmetic-based										
TV	-5.623	0.478	0.049	0.148	<b>0.940</b>	-5.395	0.628	<b>0.053</b>	0.214	<b>0.926</b>
PerTA-grad	<b>-0.661</b>	<b>0.546</b>	0.069	<b>0.310</b>	0.910	<b>-0.263</b>	<b>0.674</b>	<b>0.079</b>	<b>0.502</b>	<b>0.915</b>
PerTA-fisher	<b>-0.339</b>	<b>0.553</b>	0.077	<b>0.348</b>	<b>0.911</b>	<b>-0.405</b>	<b>0.677</b>	0.083	<b>0.561</b>	0.906

Table 4: Results of different methods on unlearning 10% of TOFU. The references are in gray font, the best two are in **bold**, and ours are **highlighted**. ‘Full’ and ‘GT’ represent the model before unlearning and the ground truth model, respectively.

	FQ↑	MU↑	ES( $\mathcal{D}_f$ )↓	ES( $\mathcal{D}_r$ )↑	Gib↑	FQ↑	MU↑	ES( $\mathcal{D}_f$ )↓	ES( $\mathcal{D}_r$ )↑	Gib↑
Model Size	1B					3B				
Full	-21.408	0.599	0.706	0.737	0.861	-26.444	0.666	0.890	0.884	0.861
GT	0.000	0.591	0.059	0.746	0.904	0.000	0.650	0.065	0.899	0.890
Training-based										
GA	-238.973	0.000	<b>0.033</b>	0.035	0.125	-236.070	0.000	<b>0.033</b>	0.035	0.050
GD	-15.484	0.434	0.076	0.151	0.707	-26.800	0.553	0.117	0.242	0.556
NPO	-10.244	0.000	<b>0.033</b>	0.035	0.329	-4.590	0.000	<b>0.034</b>	0.038	0.206
NPO+	-4.481	0.423	0.093	0.151	<b>0.946</b>	-7.042	0.546	0.087	0.224	<b>0.926</b>
Task Arithmetic-based										
TV	-12.506	0.451	0.046	0.114	0.895	-10.220	0.551	0.048	0.150	0.904
PerTA-grad	<b>-1.107</b>	<b>0.542</b>	0.071	<b>0.266</b>	<b>0.922</b>	<b>-1.708</b>	<b>0.649</b>	0.082	<b>0.432</b>	<b>0.921</b>
PerTA-fisher	<b>-1.686</b>	<b>0.548</b>	0.077	<b>0.295</b>	0.919	<b>-2.990</b>	<b>0.647</b>	0.088	<b>0.474</b>	0.911

**Detailed Results of Other Benchmarks.** Table 5 reports the results on the MUSE dataset. Following 35, we evaluate KnowMem and VerbMem, and additionally include ES and Gib as complementary



	<b>[Question]</b> What does Hsiao Yun-Hwa identify as in terms of gender?		<b>[Question]</b> What genre does Chukwu Akabueze specialize in?
<b>Forget</b>	<b>[To be Forgotten]</b> Hsiao Yun-Hwa is part of the LGBTQ+ community.	<b>Retain</b>	<b>[To be Retained]</b> Chukwu Akabueze is a specialist in the Biography genre.
GA	always always always always ... <b>Gibberish</b>	GA	always always always always ... <b>Gibberish</b>
GD	The name Hsiao indicates, followed by she identifies as this gender category. <b>Incoherent</b>	GD	Chukwu Akabueze is a specialist in the history of people from working class backgrounds whom his <b>Hallucination</b>
NPO	it is the one of the place of the place <b>Gibberish</b>	NPO	it is the place of all the place of the ... <b>Gibberish</b>
NPO+	A secret known only to a select few. <b>Irrelevant</b>	NPO+	He writes primarily in the Paganism genre. <b>Overforgetting</b>
TV	Hsiao Yun-Hwa is a female. <b>Forgotten ✓</b>	TV	Chukwu Akabueze is a Nigerian musician. <b>Overforgetting</b>
PerTA (grad)	Hsiao Yun-Hwa is a female. <b>Forgotten ✓</b>	PerTA (grad)	Chukwu Akabueze is a specialist in the genre of Biography and Memoir. <b>Retained ✓</b>
PerTA (fisher)	Hsiao Yun-Hwa is a female. <b>Forgotten ✓</b>	PerTA (fisher)	Chukwu Akabueze is a specialist in the genre of Biography and Memoir. <b>Retained ✓</b>

Figure 13: Sample output of unlearned LLM  $\theta_{\text{final}}$  applying different methods (unlearning 10%, Llama-3.2 1B Instruct). Our PerTA ensures both unlearning and retention.

Table 5: Results of different methods on MUSE. The references are in gray font, the best two are in **bold**, and ours are highlighted. ‘Full’ and ‘GT’ represent the model before unlearning and the ground truth model, respectively. Numbers in parentheses indicate deviations from the ground truth.

	KnowMem ( $D_f$ )	VerbMem ( $D_f$ )	KnowMem ( $D_r$ )	ES ( $D_f$ )	Gib↑
Full	0.644 (0.316↑)	0.579 (0.377↑)	0.555 (0.005↓)	0.295 (0.271↑)	0.800
GT	0.328 (0.000↑)	0.202 (0.000↑)	0.560 (0.000↑)	0.024 (0.000↑)	0.845
Training-based					
GA	0.003 (0.325↓)	0.049 (0.153↓)	0.008 (0.552↓)	0.008 (0.017↓)	0.001
GD	<b>0.332 (0.005↑)</b>	0.005 (0.197↓)	0.254 (0.307↓)	0.008 (0.016↓)	0.002
NPO	0.622 (0.294↑)	0.374 (0.173↑)	<b>0.521 (0.040↓)</b>	0.119 (0.094↑)	0.771
NPO+	0.642 (0.314↑)	0.494 (0.293↑)	<b>0.525 (0.036↓)</b>	0.205 (0.181↑)	<b>0.811</b>
Task Arithmetic-based					
TV	0.011 (0.317↓)	0.109 (0.092↓)	0.023 (0.537↓)	0.011 (0.014↓)	0.685
PerTA-grad	0.388 (0.060↑)	<b>0.176 (0.026↓)</b>	0.416 (0.145↓)	<b>0.028 (0.003↑)</b>	0.777
PerTA-fisher	<b>0.385 (0.058↑)</b>	<b>0.191 (0.011↓)</b>	0.464 (0.096↓)	<b>0.025 (0.001↑)</b>	<b>0.802</b>

metrics. The numbers in parentheses indicate the differences between each metric and that of the ground-truth model. For KnowMem and VerbMem, we highlight the two methods whose results are closest to the ground truth. Consistent with prior observations, PerTA alleviates the issue of excessive forgetting in TV. For example, on the forget set, PerTA improves KnowMem from 0.011 to 0.388 and 0.385 (ground truth: 0.328), and on the retain set, from 0.023 to 0.416 and 0.464 (ground truth: 0.560). These results suggest that PerTA achieves a better balance between unlearning and retention. Cases of the forget and retain samples, along with the results of different methods, is shown in Table 6. We observe that other methods often suffer from partial forgetting/retention failures or produce gibberish responses, whereas PerTA is able to forget the targeted information while preserving the retain.

**Detailed Results of Ablation Studies.** Table 7 supplements Figure 6 by showing the detailed quantitative results of different  $f_{\text{opr}}(\cdot, \cdot)$ . *Random* means to set weights in  $W$  to random values uniformly sampled between 0 and 1 with  $f_{\text{opr}}(A, B) = \text{rand}([0, 1])$ . *Weighted* uses a constant  $\omega$  to rescale TV with  $f_{\text{opr}}(A, B) = \omega$ . Here we show the results of  $\omega = 0.5$ . *Pruning* removes (i.e.,  $f_{\text{opr}}(A, B) = 0$ ) the  $\lambda\%$  smallest weights in TV to mitigate over-forgetting and maintain others (i.e.,  $f_{\text{opr}}(A, B) = 1$ ), where we show the results of  $\lambda = 0.5$ . Unlike *PerTA-grad* or *PerTA-fisher*, the gradients of *PerTA-grad* ( $\theta_{\text{full}}$ ) or *PerTA-fisher* ( $\theta_{\text{full}}$ ) are estimated on  $\theta_{\text{full}}$  instead of  $\theta_0$ . The difference between *PerTA-grad*, *PerTA-fisher* and *PerTA+SoftMax* is that the latter determines  $f_{\text{opr}}(A, B) = \exp(|A|)/(\exp(|A|) + \exp(|B|))$  in the SoftMax form.

We find that the results of *Random* are highly unstable, often exhibiting large variance, which further increases as the unlearning ratio grows and the task becomes more difficult. When the weight is fixed at 0.5, the *Weighted* method performs relatively better; however, it still lags behind our proposed PerTA in terms of unlearning capability (as measured by FQ and the ES metric on the forget set). The

Table 6: Sample answers for questions to be unlearned/retained by different methods on the MUSE-News dataset.

Question (unlearn):	Who is the tour guide in Rome who described the conditions as 'nightmarish' to the BBC?	
Full model	Felicity Hinton/100-year-old Felicity Hinton	
GA	the the the the the the the the the the...	Gibberish
GD	100 "Toto" Guidi	Gibberish
NPO	100-year-old Felicity Hinton	Fail
NPO+	100-year-old Felicity Hinton	Fail
TV	100.10.1.1.1.1.1.1.1.1.1.1.1.1.1.1.1	Gibberish
PerTA-grad	50-year-old tour guide, Alessandro Russo	Success
PerTA-fisher	60-year-old Rome resident, Alessandro Russo	Success
Question (retain):	What is the half-life of the plutonium isotope being looked at by the University of Southampton scientists?	
Full model	24,000 years	
GA	the the the the the the the the the the....	Gibberish
GD	24,000 years "platinum" of "plutonium" "half-life" .....	Gibberish
NPO	14,000 years	Fail
NPO+	14,000 years	Fail
TV	100.0.1. You are the United.\nThe United. Should. ....	Gibberish
PerTA-grad	24,000 years	Success
PerTA-fisher	24,000 years	Success

Table 7: Results using different  $f_{\text{opr}}t$  on TOFU tasks (unlearning 1%, 5% and 10% of TOFU, using Llama-3.2 1B Instruct, Mean  $\pm$  Std). Ours are highlighted .

Forgetting	Methods	FQ $\uparrow$	MU $\uparrow$	ES( $\mathcal{D}_t$ ) $\downarrow$	ES( $\mathcal{D}_r$ ) $\uparrow$
1%	Full	-2.170	0.599	0.743	0.737
	GT	0.000	0.599	0.069	0.751
	Random $\omega$	-0.877 $\pm$ 0.684	0.571 $\pm$ 0.019	0.292 $\pm$ 0.296	0.492 $\pm$ 0.175
	Weighted $\omega = 0.5$	-1.451 $\pm$ 0.131	0.587 $\pm$ 0.000	0.116 $\pm$ 0.000	0.611 $\pm$ 0.004
	Pruning $\lambda = 0.5$	-0.393 $\pm$ 0.000	0.556 $\pm$ 0.001	0.081 $\pm$ 0.001	0.358 $\pm$ 0.002
	PerTA-grad( $\theta_{\text{full}}$ )	-0.576 $\pm$ 0.000	0.583 $\pm$ 0.001	0.106 $\pm$ 0.000	0.567 $\pm$ 0.003
	PerTA-fisher( $\theta_{\text{full}}$ )	-1.182 $\pm$ 0.127	0.589 $\pm$ 0.000	0.123 $\pm$ 0.000	0.626 $\pm$ 0.001
	PerTA-grad	-0.289 $\pm$ 0.073	0.581 $\pm$ 0.001	0.075 $\pm$ 0.000	0.551 $\pm$ 0.001
	PerTA-fisher	-0.576 $\pm$ 0.000	0.586 $\pm$ 0.000	0.085 $\pm$ 0.001	0.600 $\pm$ 0.001
PerTA+SoftMax	-1.266 $\pm$ 0.000	0.586 $\pm$ 0.000	0.115 $\pm$ 0.001	0.606 $\pm$ 0.001	
5%	Full	-11.845	0.599	0.727	0.737
	GT	0.000	0.599	0.063	0.746
	Random $\omega$	-7.264 $\pm$ 3.283	0.526 $\pm$ 0.055	0.252 $\pm$ 0.284	0.346 $\pm$ 0.270
	Weighted $\omega = 0.5$	-1.253 $\pm$ 0.237	0.560 $\pm$ 0.001	0.090 $\pm$ 0.004	0.396 $\pm$ 0.004
	Pruning $\lambda = 0.5$	-5.321 $\pm$ 0.105	0.484 $\pm$ 0.003	0.049 $\pm$ 0.002	0.155 $\pm$ 0.005
	PerTA-grad( $\theta_{\text{full}}$ )	-0.630 $\pm$ 0.110	0.545 $\pm$ 0.001	0.071 $\pm$ 0.001	0.312 $\pm$ 0.007
	PerTA-fisher( $\theta_{\text{full}}$ )	-0.515 $\pm$ 0.039	0.553 $\pm$ 0.001	0.083 $\pm$ 0.001	0.360 $\pm$ 0.002
	PerTA-grad	-0.661 $\pm$ 0.125	0.546 $\pm$ 0.001	0.069 $\pm$ 0.002	0.310 $\pm$ 0.007
	PerTA-fisher	-0.339 $\pm$ 0.115	0.553 $\pm$ 0.001	0.077 $\pm$ 0.002	0.348 $\pm$ 0.002
PerTA+SoftMax	-1.219 $\pm$ 0.281	0.558 $\pm$ 0.001	0.088 $\pm$ 0.002	0.390 $\pm$ 0.002	
10%	Full	-21.408	0.599	0.706	0.737
	GT	0.000	0.591	0.059	0.746
	Random $\omega$	-13.963 $\pm$ 4.247	0.511 $\pm$ 0.065	0.228 $\pm$ 0.256	0.323 $\pm$ 0.280
	Weighted $\omega = 0.5$	-2.757 $\pm$ 0.189	0.548 $\pm$ 0.002	0.082 $\pm$ 0.001	0.309 $\pm$ 0.007
	Pruning $\lambda = 0.5$	-8.760 $\pm$ 0.282	0.483 $\pm$ 0.003	0.049 $\pm$ 0.001	0.136 $\pm$ 0.002
	PerTA-grad( $\theta_{\text{full}}$ )	-1.270 $\pm$ 0.135	0.541 $\pm$ 0.001	0.074 $\pm$ 0.001	0.274 $\pm$ 0.004
	PerTA-fisher( $\theta_{\text{full}}$ )	-2.603 $\pm$ 0.113	0.549 $\pm$ 0.002	0.082 $\pm$ 0.002	0.310 $\pm$ 0.001
	PerTA-grad	-1.107 $\pm$ 0.064	0.542 $\pm$ 0.001	0.071 $\pm$ 0.002	0.266 $\pm$ 0.003
	PerTA-fisher	-1.686 $\pm$ 0.265	0.548 $\pm$ 0.001	0.077 $\pm$ 0.002	0.295 $\pm$ 0.006
PerTA+SoftMax	-2.679 $\pm$ 0.143	0.548 $\pm$ 0.002	0.081 $\pm$ 0.001	0.310 $\pm$ 0.006	

Pruning method performs well on simple tasks, such as the 1% unlearning setting, but its performance drops sharply as the task difficulty increases with higher unlearning ratios. The *SoftMax* method is

Table 8: Results using different  $\tau$  in  $f_{\text{oprt}}$  on TOFU tasks (unlearning 1%, 5% and 10%, using Llama-3.2 1B Instruct, Mean  $\pm$  Std). Ours are **highlighted**.

Forgetting	Methods	FQ $\uparrow$	MU $\uparrow$	ES( $\mathcal{D}_f$ ) $\downarrow$	ES( $\mathcal{D}_r$ ) $\uparrow$
1%	Full	-2.170	0.599	0.743	0.737
	GT	0.000	0.599	0.069	0.751
	$\tau = 0$	-1.451 $\pm$ 0.131	0.587 $\pm$ 0.000	0.116 $\pm$ 0.000	0.611 $\pm$ 0.004
	$\tau = 0.25$	-0.089 $\pm$ 0.037	0.577 $\pm$ 0.000	0.095 $\pm$ 0.001	0.505 $\pm$ 0.001
	$\tau = 0.5$	-0.197 $\pm$ 0.057	0.579 $\pm$ 0.001	0.095 $\pm$ 0.001	0.522 $\pm$ 0.002
	$\tau = 1(\text{grad})$	-0.289 $\pm$ 0.073	0.581 $\pm$ 0.001	0.075 $\pm$ 0.000	0.551 $\pm$ 0.001
	$\tau = 2(\text{fisher})$	-0.576 $\pm$ 0.000	0.586 $\pm$ 0.000	0.085 $\pm$ 0.001	0.600 $\pm$ 0.001
	$\tau = 4$	-1.013 $\pm$ 0.000	0.591 $\pm$ 0.000	0.094 $\pm$ 0.001	0.650 $\pm$ 0.003
$\tau = 8$	-1.544 $\pm$ 0.000	0.598 $\pm$ 0.001	0.276 $\pm$ 0.004	0.700 $\pm$ 0.001	
5%	Full	-11.845	0.599	0.727	0.737
	GT	0.000	0.599	0.063	0.746
	$\tau = 0$	-1.253 $\pm$ 0.237	0.560 $\pm$ 0.001	0.090 $\pm$ 0.004	0.396 $\pm$ 0.004
	$\tau = 0.25$	-0.784 $\pm$ 0.090	0.543 $\pm$ 0.002	0.068 $\pm$ 0.001	0.299 $\pm$ 0.008
	$\tau = 0.5$	-0.754 $\pm$ 0.132	0.545 $\pm$ 0.001	0.069 $\pm$ 0.002	0.300 $\pm$ 0.008
	$\tau = 1(\text{grad})$	-0.661 $\pm$ 0.125	0.546 $\pm$ 0.001	0.069 $\pm$ 0.002	0.310 $\pm$ 0.007
	$\tau = 2(\text{fisher})$	-0.339 $\pm$ 0.115	0.553 $\pm$ 0.001	0.077 $\pm$ 0.002	0.348 $\pm$ 0.002
	$\tau = 4$	-0.933 $\pm$ 0.292	0.565 $\pm$ 0.001	0.092 $\pm$ 0.001	0.420 $\pm$ 0.004
$\tau = 8$	-7.008 $\pm$ 0.319	0.581 $\pm$ 0.001	0.141 $\pm$ 0.008	0.573 $\pm$ 0.004	
10%	Full	-21.408	0.599	0.706	0.737
	GT	0.000	0.591	0.059	0.746
	$\tau = 0$	-2.757 $\pm$ 0.189	0.548 $\pm$ 0.002	0.082 $\pm$ 0.001	0.309 $\pm$ 0.007
	$\tau = 0.25$	-1.270 $\pm$ 0.135	0.539 $\pm$ 0.002	0.071 $\pm$ 0.001	0.262 $\pm$ 0.004
	$\tau = 0.5$	-1.186 $\pm$ 0.066	0.541 $\pm$ 0.002	0.072 $\pm$ 0.002	0.265 $\pm$ 0.004
	$\tau = 1(\text{grad})$	-1.107 $\pm$ 0.064	0.542 $\pm$ 0.001	0.071 $\pm$ 0.002	0.266 $\pm$ 0.003
	$\tau = 2(\text{fisher})$	-1.686 $\pm$ 0.265	0.548 $\pm$ 0.001	0.077 $\pm$ 0.002	0.295 $\pm$ 0.006
	$\tau = 4$	-3.490 $\pm$ 0.210	0.558 $\pm$ 0.001	0.088 $\pm$ 0.002	0.354 $\pm$ 0.009
$\tau = 8$	-9.796 $\pm$ 0.462	0.575 $\pm$ 0.001	0.138 $\pm$ 0.000	0.493 $\pm$ 0.003	

able to achieve both forgetting and retention, yet it remains inferior to *PerTA-grad* and *PerTA-fisher*. In addition, the results indicate that estimating gradients on  $\theta_{\text{full}}$  or  $\theta_0$  leads to negligible differences in performance.

**Detailed Results of the General Form.** Considering the general form of  $f_{\text{oprt}}(A, B) = |A|^{\circ\tau}/(|A|^{\circ\tau} + |B|^{\circ\tau})$  but not using the absolute gradient or the diagonal Fisher Information approximation, different  $\tau$ s can be applied. We conduct experiments for  $\tau \in \{0, 0.25, 0.5, 1, 2, 4, 8\}$ , with the quantitative results shown in Table 8 as a supplement to Figure 8. The cases of  $\tau = 1, 2$  correspond to our *PerTA-grad* and *PerTA-fisher*, respectively.

The results in Table 8 lead to conclusions that are consistent with those discussed in the main body of our paper. *PerTA-grad* and *PerTA-fisher* strike a balance between forgetting and retaining: among the different  $\tau$ -based variants, they achieve relatively strong FQ and ES ( $\mathcal{D}_f$ ) while keeping MU and ES ( $\mathcal{D}_r$ ) at a reasonable level.

**Detailed Results of Running Time.** As a supplement to Figure 7, Table 9 shows the quantitative runtime comparison between the best-performing training-based method, GD and NPO+, and our *PerTA*. In contrast to training-based methods that demand multiple iterations, the runtime of *PerTA* can be broken down into three components: obtaining  $\theta_{\text{fgt}}$ , computing  $W$ , and performing task arithmetic, with the last step being negligible (0.0002 min in Table 9). *PerTA* thus inherits the efficiency of task arithmetic, yielding substantial runtime savings—a benefit that becomes increasingly evident as task complexity rises (i.e., when unlearning larger proportions). Furthermore, as demonstrated earlier, competitive performance can already be achieved by estimating gradients with only 20% of the data, indicating additional potential for reducing runtime. Collectively, these observations underscore the high time efficiency of *PerTA*.

### C.3 Different Models for Per-parameter Weights

To further illustrate the difference between using  $\theta_0$  (the retained LLM) and  $\theta_{\text{full}}$  (the finetuned LLM) to predict  $W$  shown in Table 7, Figure 14 presents a comparison. The left side of Figure 14

Table 9: Time comparison of the best-performing training-based method GD, NPO+ and our PerTA ((min), unlearning 1%, 5% and 10%, Llama-3.2 1B Instruct).

Forgetting	Methods	Getting $\theta_{\text{fgt}}$	Calculating $W_{\text{grad fisher}}$	Task Arithmetic	Total
1%	GD	-	-	-	3.4673
	NPO+	-	-	-	4.6557
	PerTA (grad)		2.0207		2.4153
	PerTA (fisher)		2.2118		2.6064
	PerTA (grad) w/ 20%	0.3944	0.4528	0.0002	0.8474
	PerTA (fisher) w/ 20%		0.4188		0.8134
5%	GD	-	-	-	5.2072
	NPO+	-	-	-	12.3739
	PerTA (grad)		2.0253		4.4172
	PerTA (fisher)		2.2201		4.6121
	PerTA (grad) w/ 20%	2.3918	0.4378	0.0002	2.8297
	PerTA (fisher) w/ 20%		0.4179		2.8098
10%	GD	-	-	-	7.2508
	NPO+	-	-	-	23.0168
	PerTA (grad)		2.0231		6.8514
	PerTA (fisher)		2.2177		7.0459
	PerTA (grad) w/ 20%	4.8281	0.4368	0.0002	5.2651
	PerTA (fisher) w/ 20%		0.4134		5.2416




Figure 14: Visualization of  $W_{\text{grad}}$ ,  $W_{\text{fisher}}$  for parameters in the last two  $Q$ ,  $K$ ,  $V$  attention layers (left), and corresponding ES on forget and retain sets (right), when employing  $\theta_0$  or  $\theta_{full}$  to estimate  $W_{\text{grad}}$ ,  $W_{\text{fisher}}$  (unlearning 1% on TOFU, using Llama-3.2 1B Instruct).

visualizes the weight magnitudes of  $W$  (predicted by  $\theta_0$  and  $\theta_{full}$ , respectively) corresponding to the  $Q$ ,  $K$ , and  $V$  matrices in the last two attention layers, while the right side reports the corresponding ES scores in bar plots. From the visualizations on the left, we observe that both PerTA-grad and PerTA-fisher exhibit highly similar patterns regardless of whether  $W$  is predicted by  $\theta_0$  and  $\theta_{full}$  (highlighted by the black boxes). This indicates that the key parameters—those with large weights—are largely consistent across the two predictors, and vice versa. On the right, the ES results confirm this observation: the numerical metrics are very close, consistent with Table 7.

These findings suggest that either  $\theta_0$  or  $\theta_{full}$  can be used to predict  $W$ , with negligible differences. A plausible explanation is that the gap between the pretrained model and the finetuned model is relatively small. This conclusion further supports the applicability of PerTA to post-training models, thereby broadening its range of use cases.

#### C.4 Visualization Results of Weights

Figures 15-17 visualize the weight magnitudes of the  $Q$ ,  $K$ , and  $V$  matrices in the shallow, middle, and final attention layers of the LLM for both TV and  $W$ . For TV, we observe that the weight magnitudes increase progressively from shallow to deeper layers, indicating that the magnitude of parameter changes induced by unlearning grows with layer depth.

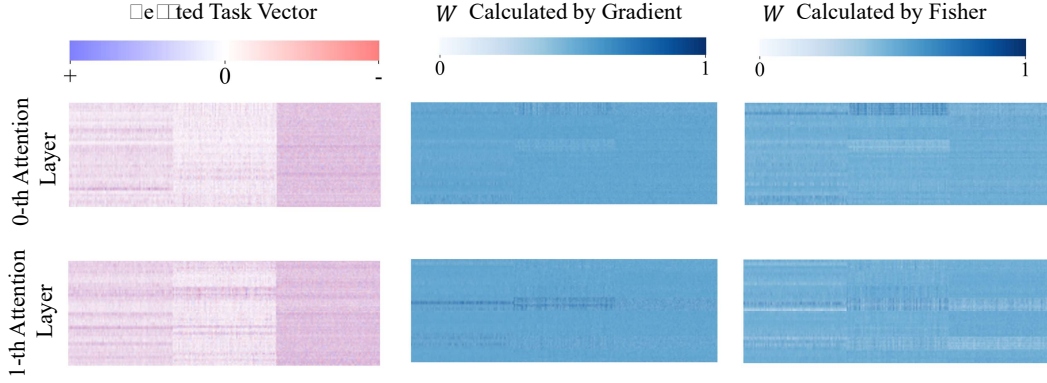


Figure 15: Visualization of TV,  $W_{\text{grad}}$ ,  $W_{\text{fisher}}$  for parameters in the 0-th, 1-st  $Q, K, V$  attention layers (unlearning 1% on TOFU, using Llama-3.2 1B Instruct).

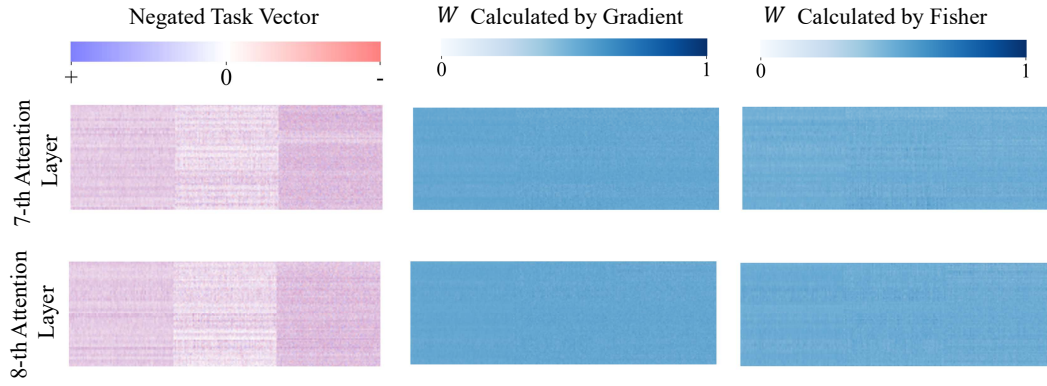


Figure 16: Visualization of TV,  $W_{\text{grad}}$ ,  $W_{\text{fisher}}$  for parameters in the 7-th, 8-th  $Q, K, V$  attention layers (unlearning 1% on TOFU, using Llama-3.2 1B Instruct).

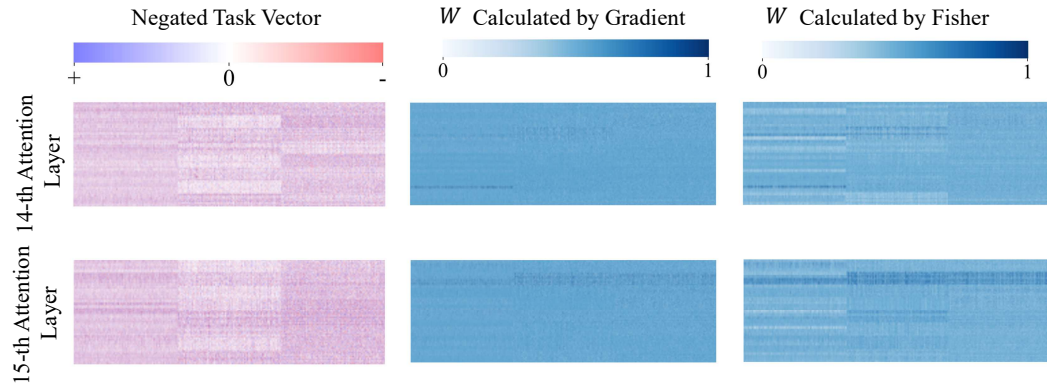


Figure 17: Visualization of TV,  $W_{\text{grad}}$ ,  $W_{\text{fisher}}$  for parameters in the 14-th, 15-th  $Q, K, V$  attention layers (unlearning 1% on TOFU, using Llama-3.2 1B Instruct).

In contrast, the analysis of  $W$  may provide insight into the layer-wise sensitivity of LLM parameters to the differences between forget and retain data. We highlight two key observations. First, compared to PerTA-grad, PerTA-fisher exhibits more pronounced weight differences (as evidenced by the larger

contrast between light and dark regions in Figure 15-17). This is because PerTA-fisher relies on the squared gradients rather than the raw gradients, thereby amplifying the differences between the forget and retain sets. In practice, however, both PerTA-grad and PerTA-fisher yield similar performance on the evaluation metrics, suggesting that either variant can be employed effectively.

Second, relative to the middle layers of the LLM, the initial and final layers contain more weights close to the extremes (i.e., near 0 or 1). This implies that parameters in the shallow and final layers are more sensitive to gradient differences between the forget and retain sets. Interestingly, this aligns with prior findings on LLM representations [15, 39]: shallow layers primarily capture surface features (e.g., words, subwords, positional information), middle layers encode syntactic features, and final layers specialize in semantic features. The results in Figures 15-17 are consistent with this interpretation. Specifically, surface and semantic features exhibit greater discrepancies between forget and retain sets (e.g., TOFU involves differences in author names, domain-specific terminology, and deeper semantic associations with personal information), whereas syntactic structures remain largely unaffected. Consequently, our flexible PerTA assigns larger weight differences to parameters in the shallow and final layers. This insight suggests a potential future direction for further optimization: pruning or fixing selected middle layers to reduce computational overhead without sacrificing performance.

Table 10: Results of different methods on unlearning 5% of TOFU, using Llama-3.2 8B as the pretrained model.

	FQ $\uparrow$	MU $\uparrow$	ES( $D_f$ ) $\downarrow$	ES( $D_r$ ) $\uparrow$	ES( $D_r$ )-ES( $D_f$ ) $\uparrow$	Gib $\uparrow$
Full (reference)	-12.184	0.628	0.972	0.992	0.020	0.852
GT (reference)	0.000	0.632	0.074	0.992	0.918	0.886
GA	-118.712	0.000	0.033	0.035	0.002	0.038
GD	-10.225	0.509	0.158	0.397	0.239	0.811
NPO	-11.183	0.131	0.033	0.037	0.004	0.141
NPO+	-7.888	0.569	0.160	0.521	0.361	<b>0.914</b>
PerTA-grad (ours)	<b>-4.529</b>	<b>0.659</b>	0.164	0.882	<b>0.718</b>	0.895

Table 11: Results of different methods on unlearning 5% of TOFU, using Phi-3.5 as the pretrained model.

	FQ $\uparrow$	MU $\uparrow$	ES( $D_f$ ) $\downarrow$	ES( $D_r$ ) $\uparrow$	ES( $D_r$ )-ES( $D_f$ ) $\uparrow$	Gib $\uparrow$
Full (reference)	-13.232	0.693	0.868	0.835	-0.033	0.866
GT (reference)	0.000	0.678	0.082	0.855	0.773	0.881
GA	-11.511	0.073	0.027	0.028	0.001	0.822
GD	-11.183	0.665	0.344	0.574	0.231	0.875
NPO	-12.877	0.278	0.538	0.594	0.057	0.855
NPO+	-10.859	0.552	0.591	0.761	0.170	0.877
PerTA-grad (ours)	<b>-3.548</b>	<b>0.667</b>	0.107	0.412	<b>0.305</b>	<b>0.879</b>

Table 12: Average results of PerTA with quantization attacks on TOFU 1%, 5%, 10% unlearning tasks.

	FQ $\uparrow$	MU $\uparrow$	ES( $D_f$ ) $\downarrow$	ES( $D_r$ ) $\uparrow$	ES( $D_r$ )-ES( $D_f$ ) $\uparrow$	Gib $\uparrow$
Full	-11.808	0.599	0.726	0.737	0.011	0.871
GT	0.000	0.596	0.064	0.748	0.684	0.894
GA	-81.114	0.199	0.086	0.244	0.157	0.484
GD	-8.720	0.491	0.112	0.295	0.183	0.789
NPO	-4.842	0.198	0.086	0.246	0.160	0.592
NPO+	-3.528	0.493	0.122	0.316	0.194	0.911
PerTA-grad (ours) w/o attack	<b>-0.686</b>	0.556	0.072	0.376	0.304	<b>0.915</b>
PerTA-grad (ours) w/ attack	-1.340	<b>0.560</b>	0.095	0.421	<b>0.325</b>	0.909

### **C.5 Results on Larger Models and Alternative LLM Families**

Tables 10 and Table 11 present our method’s performance on larger models and on models from other LLM families. The results indicate that our PerTA exhibits good generalization ability: it achieves competitive unlearning performance even when applied to larger models and different types of LLMs.

### **C.6 Results of Quantization Attacks**

Some recent research [55] have found that applying quantization to models that have undergone unlearning can restore the "forgotten" information. Therefore, conducting attack experiments on PerTA to reveal whether it possesses robustness is crucial.

Accordingly, we evaluate the model after unlearning—using Llama-3.2 1B as an example—and the results are shown in Table 12. The results show that, fortunately, the impact of quantization on PerTA is limited, and PerTA still outperforms other methods after the attack.