

SIMPLICITY PREVAILS: RETHINKING NEGATIVE PREFERENCE OPTIMIZATION FOR LLM UNLEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

In this work, we study the problem of large language model (LLM) unlearning, aiming to remove unwanted data influences and associated model capabilities (e.g., copyrighted data or harmful content generation) while preserving essential model utilities, without the need for retraining from scratch. Despite the growing need for LLM unlearning, a **technically-grounded optimization framework remains lacking**. For example, gradient ascent (GA)-type approaches have been widely used for unlearning. However, GA is suboptimal as it reverses the learning process without controlling the optimization divergence speed, making it prone to over-forgetting and even leading to model collapse. Negative preference optimization (NPO) has been proposed to address this issue and is considered one of the **state-of-the-art LLM unlearning approaches**. In this work, we revisit NPO, and identify the issue of reference model bias, which could undermine NPO’s effectiveness, particularly when unlearning forget data of varying difficulty. **This reference model bias arises from using the reference model as the basis for unlearning criterion in NPO, leading to two issues: uneven optimization power allocation over the forget data and ineffective gradient weight smoothing in the early optimization stages**. To overcome these challenges, we propose a simple yet effective unlearning optimization framework, called SimNPO, showing that ‘simplicity’ in removing the reliance on a reference model (through the lens of simple preference optimization) benefits unlearning. We also provide deeper insights into SimNPO’s advantages, supported by analysis using mixtures of Markov chains. Furthermore, we present extensive experiments validating SimNPO’s superiority over existing baselines in benchmarks like TOFU and MUSE, and robustness against relearning attacks.

1 INTRODUCTION

The rapid advancement of large language models (LLMs) has raised security and safety concerns, including issues related to copyright violations and sociotechnical harms (Huang et al., 2024; Wang et al., 2023; Li et al., 2024; Shi et al., 2024). However, retraining these models to remove undesirable data influences is often impractical due to the substantial costs and time required for such processes. This gives rise to the problem of **LLM unlearning**, which aims to effectively remove undesired data influences and/or model behaviors while preserving the utility for essential, unrelated knowledge generation, and maintaining efficiency without the need for retraining (Eldan & Russinovich, 2023; Yao et al., 2023; Liu et al., 2024b; Blanco-Justicia et al., 2024).

To trace its origins, the concept of *machine unlearning* was initially developed for data removal to comply with privacy regulations such as the “right to be forgotten” (Rosen, 2011; Hoofnagle et al., 2019), with early studies focusing on vision models (Cao & Yang, 2015; Warnecke et al., 2021; Bourtole et al., 2021; Thudi et al., 2022; Kurmanji et al., 2024; Jia et al., 2023; Gandikota et al., 2023; Fan et al., 2024b). However, it is soon adapted to LLMs to remove unwanted data, knowledge, or specific model capabilities (Eldan & Russinovich, 2023; Yao et al., 2023; Liu et al., 2024b; Ji et al., 2024; Li et al., 2024; Shi et al., 2024; Maini et al., 2024; Zhang et al., 2024a; Jia et al., 2024). The current optimization foundation for LLM unlearning often relies on *optimization divergence from the pre-trained state*. Divergence refers to the process of deviating from the converged pre-trained model state to reverse the effects of learning the forgotten data, thereby achieving unlearning (Liu et al., 2022a; Maini et al., 2024; Yao et al., 2023; Jia et al., 2024). Nevertheless, the lack of control over optimization divergence can result in either under-forgetting, where insufficient unwanted data

influence is removed, or over-forgetting, leading to a significant loss of model utility in LLMs. Therefore, optimization for LLM unlearning is a highly non-trivial challenge.

Negative preference optimization (NPO) (Zhang et al., 2024a) emerges as an effective approach for LLM unlearning, as demonstrated by its better control of the divergence speed during unlearning optimization and its strong performance in current benchmarks such as TOFU (Maini et al., 2024) and MUSE (Shi et al., 2024). Inspired by direct preference optimization (DPO) (Rafailov et al., 2024), it treats the forget data points as negative responses, providing a lower-bounded unlearning objective. This also induces a gradient weight smoothing scheme to regulate the speed of divergence. We refer readers to Sec. 3 for details.

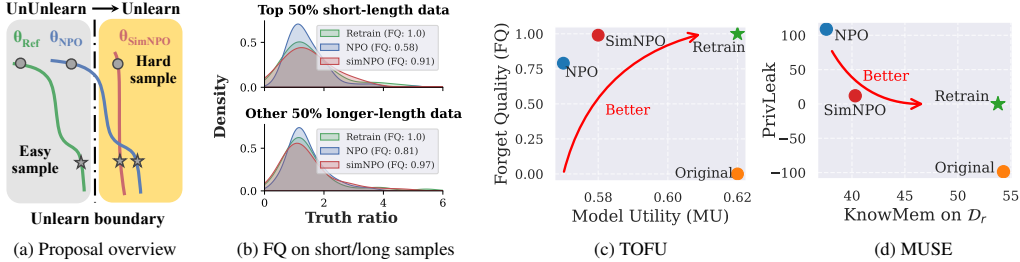


Figure 1: (a) Systematic overview of an LLM (θ) post-unlearning using the proposed SimNPO, compared to NPO (Zhang et al., 2024a) and the reference model, where NPO struggles to unlearn hard samples due to the reference model bias. (b) Truth ratio distribution of top 50% shortest-length forget data points and the other 50% longer-length data for NPO, SimNPO, and Retrain on the TOFU Forget05 dataset (Maini et al., 2024), where SimNPO achieves better forget quality (FQ) than NPO and exhibits a truth ratio distribution closer to Retrain. Note that FQ is a statistical measure quantifying the closeness between the truth ratio distribution of an unlearned model and that of Retrain (with FQ = 1 representing optimal unlearning). (c) & (d) Experiment highlights on TOFU Forget05 and MUSE News datasets (Shi et al., 2024). Unlearning effectiveness is measured by FQ for TOFU and PrivLeak for MUSE, while utility preservation is evaluated using model utility for TOFU and KnowMem on retain data for MUSE (see Table 1). In both tasks, Retrain is the gold standard for unlearning.

Despite the advancements NPO has introduced to the optimization foundation for LLM unlearning, this work will identify its potential limitations for the first time, arising from its reliance on the reference model (i.e., the model prior to unlearning) as the basis for promoting and regulating the optimization divergence. We refer to this issue as *reference model bias*. Fig. 1-(a) illustrates this issue schematically. NPO aims to widen the gap between the unlearned model (θ_{NPO}) and the reference model (θ_{Ref}). However, the prediction confidence of θ_{Ref} varies across samples (green line in Fig. 1-(a)). For some samples, predictions of θ_{Ref} are already near the unlearning boundary, making them “easy examples” for unlearning, where further enlarging the gap is unnecessary. Despite this, NPO may continue increasing the distance (blue line in Fig. 1-(a)), causing easy examples to move far beyond the boundary while “hard examples” remain unresolved. Throughout this work, the key research question we aim to answer is:

(Q) When and why does the current optimization foundation –NPO– for LLM unlearning become ineffective, and how can it be improved?

To address (Q), we propose a simple yet effective unlearning optimization framework, termed **SimNPO**, demonstrating that properly removing reliance on a reference model can significantly enhance unlearning. This approach also draws inspiration from simple preference optimization in LLM alignment (Meng et al., 2024). Additionally, we provide detailed and in-depth insights into how SimNPO overcomes the limitations of NPO caused by reference model bias. For example, Fig. 1-(a) schematically illustrates that, compared to NPO, SimNPO more accurately recognizes the unlearning data difficulty and provides better optimization allocation across different types of forget data, e.g., hard vs. easy samples in Fig. 1-(a). Fig. 1-(b) further provides experimental justification by comparing the unlearning performance of NPO and SimNPO across forget data points, categorized by *response length*. The rationale is that as noted in (Meng et al., 2024), the reference model tends to bias toward generating longer but lower-quality sequences, making these longer samples easier to unlearn. However, NPO may exacerbate this by over-allocating optimization power to these easy samples, thereby disadvantaging the unlearning of shorter-response forget data. This explains why Fig. 1-(b) shows that NPO performs worse than SimNPO, as evidenced by a greater deviation from **Retrain** (the exact unlearning method achieved by retraining the model from scratch without the

108 forgotten data). This gap is particularly evident when unlearning the top 50% shortest-length forget
 109 data compared to the longer 50% of the forget set. We refer readers to Secs. 4 and 5 for more details.

110 The contributions of our work are summarized below:

- 111 • We revisit the NPO framework and identify its potential weakness—reference model bias—in LLM
 112 unlearning, which can lead to issues such as sensitivity to the reference model’s response quality and
 113 ineffective gradient weight smoothing.
- 114 • Building on insights into NPO’s limitations, we propose an improved LLM unlearning approach,
 115 SimNPO, which extends NPO using a reference-free optimization framework, simple preference
 116 optimization (Meng et al., 2024). We also delve into the technical rationale behind how SimNPO
 117 alleviates the limitations of NPO, validated through the lens of mixtures of Markov chains.
- 118 • We conduct extensive experiments to demonstrate the improvements of SimNPO over NPO across
 119 various scenarios, including: forgetting data with different response lengths, as in TOFU (Maini et al.,
 120 2024); forgetting data with uniform response lengths, as in MUSE (Shi et al., 2024) and WMDP (Li
 121 et al., 2024); and defending against relearning-based attacks (Lynch et al., 2024; Hu et al., 2024). See
 122 some experiment highlights in Fig. 1-(c,d).

125 2 RELATED WORK

126 **Machine unlearning.** The gold standard for machine unlearning in our work is ‘Retrain’, also
 127 referred to as *exact* unlearning (Cao & Yang, 2015; Thudi et al., 2022; Fan et al., 2024a), which
 128 involves retraining the model from scratch on the training set while excluding the data points to be
 129 forgotten. However, exact unlearning is challenging in practice due to the assumption for access
 130 to the full training set and the high computational cost of retraining. To address these challenges,
 131 various *approximate* unlearning methods have been developed (Nguyen et al., 2022; Bourtole et al.,
 132 2021; Triantafillou et al., 2024). These approaches typically involve model fine-tuning or editing,
 133 applied to the pre-trained model, based on the unlearning request. Their effectiveness has been shown
 134 in different application domains, including image classification (Liu et al., 2022b; Jia et al., 2023;
 135 Kurmanji et al., 2024; Fan et al., 2024a), image generation (Gandikota et al., 2023; Fan et al., 2024b;
 136 Zhang et al., 2024b), federated learning (Liu et al., 2022c; Halimi et al., 2022; Jin et al., 2023), and
 137 graph neural networks (Chen et al., 2022; Chien et al., 2022; Wu et al., 2023a).

138 **LLM unlearning.** There has also been a growing body of research focusing on machine unlearning
 139 for LLMs (Lu et al., 2022; Jang et al., 2022; Kumar et al., 2022; Zhang et al., 2023; Pawelczyk
 140 et al., 2023; Eldan & Russinovich, 2023; Ishibashi & Shimodaira, 2023; Yao et al., 2023; Maini
 141 et al., 2024; Zhang et al., 2024a; Li et al., 2024; Wang et al., 2024; Jia et al., 2024; Liu et al.,
 142 2024b;a; Thaker et al., 2024; Kadhe et al., 2024). Applications of unlearning in LLMs are diverse,
 143 from safeguarding copyrighted and personally identifiable information (Jang et al., 2022; Eldan &
 144 Russinovich, 2023; Wu et al., 2023b), to preventing LLMs from creating cyberattacks or bioweapons
 145 (Barrett et al., 2023; Li et al., 2024), and reducing the production of offensive, biased, or misleading
 146 content (Lu et al., 2022; Yu et al., 2023; Yao et al., 2023). Current unlearning approaches include
 147 model optimization-based methods (Ilharco et al., 2022; Liu et al., 2022a; Yao et al., 2023; Eldan &
 148 Russinovich, 2023; Jia et al., 2024; Zhang et al., 2024a; Li et al., 2024) and input prompt or in-context
 149 learning-based techniques (Thaker et al., 2024; Pawelczyk et al., 2023; Liu et al., 2024a). However,
 150 many lack effectiveness, leading to either under-forgetting or over-forgetting, as shown by recent
 151 LLM unlearning benchmarks such as TOFU for fictitious unlearning (Maini et al., 2024) and MUSE
 152 for private or copyrighted information removal (Shi et al., 2024). Recent studies also show that even
 153 after unlearning, models can remain vulnerable to adversarial attacks (Schwarzschild et al., 2024;
 154 Patil et al., 2024; Lynch et al., 2024) or relearning from a small number of data (Hu et al., 2024;
 155 Lynch et al., 2024). This evidence suggests that effective unlearning for LLMs is far from trivial.
 156 Among current efforts, NPO (negative preference optimization) (Zhang et al., 2024a) stands out as a
 157 promising method, offering key advantages such as a bounded unlearning loss and gradient ascent
 158 with weight smoothing to enhance stability and control. However, we will show that the advantages
 159 of NPO can be limited by the presence of reference model bias (Sec. 4).

160 **Preference optimization.** In this work, we advance LLM unlearning through the lens of preference
 161 optimization. This is motivated by aligning LLMs with human values, known as reinforcement
 learning from human feedback (RLHF) (Christiano et al., 2017; Ziegler et al., 2019; Ouyang et al.,

2022). However, online preference optimization algorithms are often complex and challenging to optimize (Santacrose et al., 2023; Zheng et al., 2023), driving interest in more efficient offline alternatives. Direct preference optimization (DPO) (Rafailov et al., 2024) introduced an offline approach that eliminates the need for a reward model, sparking the development of several reward-free offline preference objectives (Zhao et al., 2023; Azar et al., 2024; Hong et al., 2024; Ethayarajh et al., 2024; Meng et al., 2024; Yuan et al., 2024). Notable methods include RRHF (Yuan et al., 2024), SLic-HF (Zhao et al., 2023), IPO (Azar et al., 2024), KTO (Ethayarajh et al., 2024), ORPO (Hong et al., 2024), and SimPO (Meng et al., 2024). Among these methods, SimPO is a reference-free, length-normalized variant of DPO, and we will demonstrate that it is well-suited for integrating into LLM unlearning and improving NPO.

3 A PRIMER ON LLM UNLEARNING

Problem formulation of LLM unlearning. Unlearning tasks can take various forms and are typically associated with a specific set of data points to be removed, known as the *forget set* (\mathcal{D}_f). In addition, these tasks often require a complementary set of non-forgotten data points, known as the *retain set* (\mathcal{D}_r), to preserve model utility by penalizing the divergence caused by unlearning. As a result, the problem of LLM unlearning can be cast as a regularized optimization problem that balances the forget and retain objectives (Liu et al., 2024b; Yao et al., 2023; Zhang et al., 2024a):

$$\text{minimize}_{\theta} \underbrace{\mathbb{E}_{(x,y) \in \mathcal{D}_f} [\ell_f(y|x; \theta)]}_{\text{Forget loss}} + \lambda \underbrace{\mathbb{E}_{(x,y) \in \mathcal{D}_r} [\ell_r(y|x; \theta)]}_{\text{Retain loss}}, \quad (1)$$

where θ represents the model parameters to be updated during unlearning, $\lambda \geq 0$ is a regularization parameter to penalize the ‘divergence’ of unlearning, and ℓ_f and ℓ_r represent forget and retain losses incurred when using model parameters θ to generate the desired response (y) given the input x .

Substantial research has focused on designing and analyzing appropriate forget and retain loss functions to solve problem (1) (Liu et al., 2024b; Yao et al., 2023; Zhang et al., 2024a; Maini et al., 2024; Shi et al., 2024; Eldan & Russinovich, 2023; Jia et al., 2024). For instance, let $\pi_{\theta}(y|x)$ represent the prediction probability of the model θ given the input-response pair (x, y) . The retain loss is typically chosen as the cross-entropy-based sequence prediction loss, $\ell_r(y|x, \theta) = -\log \pi_{\theta}(y|x)$, whose minimization encourages the model to perform well on the retain data $(x, y) \in \mathcal{D}_r$. If we specify the forget loss as the *negative* token prediction loss $\ell_f(y|x, \theta) = \log \pi_{\theta}(y|x)$, whose minimization then *discourages* the model from learning the forget data $(x, y) \in \mathcal{D}_f$. Minimizing such a forget loss is known as the *gradient ascent* (GA) method (Maini et al., 2024; Thudi et al., 2022). Similarly, minimizing the regularized loss that integrates GA with the retain loss is known as the *gradient difference* (GradDiff) method (Liu et al., 2022a; Maini et al., 2024; Yao et al., 2023).

Negative preference optimization (NPO). A popular optimization framework for solving problem (1) is NPO (Zhang et al., 2024a). It treats the forget data as negative examples in DPO (Rafailov et al., 2024), transforming the unbounded GA-based forget loss into a ① *bounded loss from below*, which helps prevent catastrophic collapse, and an ② *adaptive weight smoothing* applied to the forget loss gradients, allowing for more controlled divergence speed required for unlearning. These benefits can be clearly seen from the NPO loss and its gradient as follows:

$$\ell_{\text{NPO}}(\theta) = \mathbb{E}_{(x,y) \in \mathcal{D}_f} \left[-\frac{2}{\beta} \log \sigma \left(-\beta \log \left(\frac{\pi_{\theta}(y|x)}{\pi_{\text{ref}}(y|x)} \right) \right) \right], \quad (2)$$

① := $\ell_f(y|x; \theta)$, the specified forget loss in (1)

$$\nabla_{\theta} \ell_{\text{NPO}}(\theta) = \mathbb{E}_{(x,y) \in \mathcal{D}_f} \left[\underbrace{\left(\frac{2\pi_{\theta}(y|x)^{\beta}}{\pi_{\theta}(y|x)^{\beta} + \pi_{\text{ref}}(y|x)^{\beta}} \right)}_{\text{② := } w_{\theta}(x, y), \text{ adaptive weight}} \cdot \underbrace{\nabla_{\theta} \log \pi_{\theta}(y|x)}_{\text{GA}} \right], \quad (3)$$

where $\sigma(t) = 1/(1 + e^{-t})$ is the sigmoid function, $\beta > 0$ is the temperature parameter, e.g., $\beta = 0.1$ is used by Zhang et al. (2024a), and π_{ref} is the reference model given by the initial model prior to unlearning. We can justify the insights (①-②) below.

① From (2), the NPO-type forget loss is bounded below by 0, i.e., $\ell_f(y|x; \theta) \geq 0$, whereas the GA-type forget loss, $\ell_f(y|x, \theta) = \log \pi_{\theta}(y|x)$, has no lower bound. **Moreover, minimizing it towards**

$\ell_f(y|x; \theta) \rightarrow 0$ drives the prediction probability $\pi_\theta(y|x)$ to decrease, widening the gap between the prediction probability and the reference model on the forget set, *i.e.*, $\pi_\theta(y|x) \ll \pi_{\text{ref}}(y|x)$.

② As seen in (3), the adaptive weight $w_\theta(x, y)$ is typically less than 1 since $\pi_\theta(y|x) < \pi_{\text{ref}}(y|x)$ for forgetting. Consequently, NPO’s gradient yields more controlled and gradual divergence speed required for unlearning, compared to GA (with $w_\theta(x, y) = 1$).

Throughout this paper, NPO will serve as the primary baseline for LLM unlearning. Unless specified otherwise, its implementation follows the **regularized optimization** in (1) to balance the unlearning with model utility, where the forget loss ℓ_f is defined as in (2) and the retain loss ℓ_r is the token prediction loss $\ell_r(y|x, \theta) = -\log \pi_\theta(y|x)$ applied to the retain set.

LLM unlearning tasks and evaluations. Given that the assessment of LLM unlearning may rely on specific tasks, we next introduce the unlearning tasks and evaluation metrics that this work covers. (1) **TOFU** (Maini et al., 2024) considers fictitious unlearning on a synthetic Q&A dataset. (2) **MUSE** (Shi et al., 2024) is designed to remove verbatim or knowledge memorization from News and Books datasets, including both verbatim texts and knowledge sets for unlearning evaluation. (3) **WMDP** (Li et al., 2024) aims to prevent LLMs from generating hazardous content in domains such as biology, cybersecurity, and chemistry. In Secs. 4 and 5, we will focus on the TOFU dataset, while experimental results on MUSE and WMDP will be provided in Sec. 6. Despite the differences in evaluation metrics across the above tasks, the assessment broadly falls into two categories. (1) **Unlearning effectiveness** measures how faithfully undesired data influences or model capabilities are removed. For example, it is assessed by the *forget quality* metric in TOFU, which uses a *p*-value to test the indistinguishability between the post-unlearning model and a model retrained on the retain set only, and by *privacy leakage* in MUSE, which measures the likelihood of detecting that the model was ever trained on the forget set. (2) **Utility preservation** evaluates the post-unlearning performance on standard utility tasks. See **Table 1** for a summary of the unlearning tasks and evaluation metrics.

Table 1: Summary of unlearning efficacy and utility metrics across different unlearning benchmarks. The arrows indicate the directions for better performance (\uparrow for higher values, \downarrow for lower values, $\rightarrow 0$ for closer to 0).

| Benchmark | LLM to be used | Task Description | Unlearning Effectiveness | Utility Preservation | |
|-----------|-----------------------|----------------------------------------------------------------------------|-------------------------------------------------------------|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TOFU | LLaMA-2-chat 7B | Unlearning fictitious authors from a synthetic Q&A dataset | Forget quality (measured by truth ratios of forget samples) | \uparrow | Model utility (harmonic mean of 9 utility metrics) Probability on $\mathcal{D}_r/\mathcal{D}_{\text{real.author}}/\mathcal{D}_{\text{world.facts}}$ Rouge-L on $\mathcal{D}_r/\mathcal{D}_{\text{real.author}}/\mathcal{D}_{\text{world.facts}}$ Truth ratio on $\mathcal{D}_r/\mathcal{D}_{\text{real.author}}/\mathcal{D}_{\text{world.facts}}$ |
| | | | Probability on \mathcal{D}_f | \downarrow | |
| | | | Rouge-L on \mathcal{D}_f | \downarrow | |
| | | | Truth ratio on \mathcal{D}_f | \uparrow | |
| | | | | | |
| MUSE | LLaMA-2 7B ICLM-7B | Unlearning real-world knowledge from texts about Harry Potter and BBC News | KnowMem on \mathcal{D}_f | \downarrow | KnowMem on \mathcal{D}_r |
| | | | VerbMem on \mathcal{D}_f | \downarrow | |
| | | | PrivLeak | $\rightarrow 0$ | |
| WMDP | Zephyr-7B-beta | Unlearning hazardous knowledge from bio/cybersecurity texts | Accuracy on WMDP-Bio | \downarrow | Accuracy on MMLU |
| | | | Accuracy on WMDP-Cyber | \downarrow | |

4 UNCOVERING REFERENCE MODEL BIAS: A LIMITATION OF NPO

In this section, we highlight a key weakness of NPO, which we term ‘*reference model bias*’, and provide a concise description below. That is, the incorporation of the reference model in NPO biases the unlearning objective towards enlarging the distance relative to the reference model. This is analogous to inductive bias, where a machine learning algorithm is guided to favor certain patterns over others based on the underlying assumptions of the parametric model. Specifically, as noted in (2), minimizing the NPO loss drives $\pi_\theta(y|x) \ll \pi_{\text{ref}}(y|x)$. However, using π_{ref} as the basis for NPO’s unlearning criterion can introduce negative effects, as illustrated by the limitation (L1)-(L2).

(L1) NPO suffers from uneven allocation of unlearning power. At first glance, driving $\pi_\theta(y|x) \ll \pi_{\text{ref}}(y|x)$ in NPO appears desirable for unlearning on the forget set, where the reference model π_{ref} is given by the initial model prior to unlearning. The potential issue is that NPO’s reliance on π_{ref} can lead to an uneven allocation of unlearning power, even misaligned with the true sample-specific unlearning difficulty. We elaborate on this issue through two examples.

(*Example 1: Unlearning short vs. long-response data.*) In this example, we evaluate unlearning performance across different types of forget data points, categorized by their response lengths (*i.e.*, short vs. long). As noted in (Meng et al., 2024), a reference model may exhibit a bias toward generating longer but lower-quality sequences. Consequently, these low-quality long texts tend to be easier to unlearn compared to short-length forget data. This suggests that allocating additional optimization power to further enlarge the distance from the reference model for these easy-to-unlearn samples is unnecessary. Such an allocation leads to an uneven distribution of optimization power,

disadvantaging the unlearning of shorter-response forget data points (*i.e.*, harder examples). Indeed, **Fig. 1-(b)** shows that NPO exhibits a greater distance from Retrain when unlearning the top 50% shortest-length forget data, resulting in a lower FQ of 0.58. In contrast, NPO performs better unlearning for the longer 50% of the forget set, yielding a higher FQ of 0.81. Therefore, NPO stays ineffective at unlearning forget data with short responses. This issue is also illustrated in Fig. 1-(a), where over-forgetting easy examples in NPO can lead to under-forgetting hard examples. Further, it will be demonstrated using a mixture of Markov chains in Sec. 5.

(*Example 2: Unlearning strongly vs. weakly-memorized forget data.*) We next explain (L1) from the perspective of unlearning vs. data memorization. Consider two forget sets, $\mathcal{D}_{f,1}$ and $\mathcal{D}_{f,2}$, where $\mathcal{D}_{f,1}$ is more strongly memorized by the model than $\mathcal{D}_{f,2}$. As a result, the prediction loss on $\mathcal{D}_{f,1}$ is smaller, leading to a higher prediction probability π_{ref} . Accordingly, the NPO gradient smoothing term in (3) becomes smaller for $\mathcal{D}_{f,1}$, meaning NPO allocates less first-order optimization power to it. However, $\mathcal{D}_{f,1}$, being strongly memorized, should ideally receive more unlearning power. See **Table A1 of Appendix C.3** for experimental justification on the above example.




Figure 2: Experimental evidence of ineffective weight smoothing and over-unlearning for NPO on TOFU with 5% forget set size: (a) NPO’s gradient weights (w_θ) at epoch 1 vs. response length $|y|$. (b) Trajectory of w_θ for NPO over unlearning epochs, visualized using box plots to represent the distribution of gradient weights across forget samples for each epoch. (c)-(d) Forget quality and model utility of NPO across epochs.

(L2) **NPO may cause ineffective gradient weight smoothing and over-unlearning.** Another issue introduced by the reference model π_{ref} concerns the effectiveness of NPO’s gradient weight smoothing, *i.e.*, $w_\theta(x, y) = (2\pi_\theta(y|x)^\beta) / (\pi_\theta(y|x)^\beta + \pi_{\text{ref}}(y|x)^\beta)$ in (3). During the early optimization stage of NPO, we find $w_\theta(x, y) \approx 1$ regardless of the varying data-specific unlearning difficulties since the initialization of the unlearned model θ is given by the reference model. **Fig. 2-(a,b)** support this finding by displaying the gradient smoothing weights of NPO at epoch one (Fig. 2a) and their trajectory over the course of unlearning epochs (Fig. 2b). As shown, the gradient smoothing weights of NPO show large variance, but most values are concentrated around $w_\theta(x, y) \approx 1$ at epoch one. This suggests that NPO behaves similarly to GA in the early stage of unlearning, potentially causing over-unlearning and a large utility drop even if the weight decreases in later optimization. **Fig. 2-(c,d)** justify the above by presenting the forget quality and model utility of NPO on TOFU against unlearning epochs. As shown, NPO tends to cause a larger utility drop at early epochs compared to *SimNPO*, the improved alternative to NPO that we will introduce in Sec. 5. Additionally, NPO remains less effective in forgetting than *SimNPO* throughout the process.

5 SIMNPO: ADVANCING NPO BY SIMPLE PREFERENCE OPTIMIZATION

In the following, we address the reference model bias in NPO by using a reference-free optimization method, **SimPO** (simple preference optimization) (Meng et al., 2024). We refer to the NPO alternative derived from SimPO as **SimNPO**, simple negative preference optimization.

Motivation of SimNPO and its forget objective. The simplest solution to mitigating NPO’s reference model bias is to directly remove π_{ref} from the gradient in (3), setting $\pi_{\text{ref}} = 0$. However, this variant would be *ineffective*, as the reference-free gradient reduces to GA, with $w_\theta(x, y) = 1$. This negates NPO’s advantages. To develop a better solution for improving NPO, we revisit the context of preference optimization and investigate whether the reference model can be excluded while still retaining the unlearning benefits provided by NPO. Our idea parallels how NPO was originally inspired by DPO (Rafailov et al., 2024). We adopt SimPO, a reference-free alternative to DPO, as the optimization framework for unlearning, leading to the SimNPO method. The *key difference* between SimPO and DPO lies in their reward formulation for preference optimization. In DPO, the reward formulation is given by the comparison with the reference model, *i.e.*, $\beta \log(\pi_\theta(y|x) / \pi_{\text{ref}}(y|x))$.

This formulation was used by NPO. In contrast, SimPO takes a *reference-free but length-normalized* reward formulation: $(\beta/|y|) \log \pi_{\theta}(y|x)$, where $|y|$ denotes the response length.

Taking the inspiration of SimPO, we can mitigate the reference model bias in NPO by replacing its reward formulation $\beta \log(\pi_{\theta}(y|x)/\pi_{\text{ref}}(y|x))$ in (2) with the SimPO-based reward formulation $(\beta/|y|) \log(\pi_{\theta}(y|x))$. This modification transforms (2) into the **SimNPO loss**:

$$\ell_{\text{SimNPO}}(\theta) = \mathbb{E}_{(x,y) \in \mathcal{D}_f} \left[-\frac{2}{\beta} \log \sigma \left(-\frac{\beta}{|y|} \log \pi_{\theta}(y|x) - \gamma \right) \right], \quad (4)$$

where $\gamma \geq 0$ is the reward margin parameter, inherited from SimPO, which defines the margin of preference for a desired response over a dispreferred one. However, unless otherwise specified, we set $\gamma = 0$ to align with the NPO loss (2). This is also desired because γ introduces a constant shift to the prediction loss $-(\beta/|y|) \log \pi_{\theta}(y|x)$. Consequently, a larger γ requires greater compensation to further suppress token prediction, enforcing a stricter unlearning condition. This can accelerate the utility drop during unlearning. See **Fig. A1 of Appendix C.4** for the ablation study of hyperparameters. The SimNPO loss (4), when integrated with the regularized optimization in (1), forms the SimNPO method.

Insights into SimNPO: Addressing NPO’s limitations one by one. Similar to NPO, the SimNPO loss (4) is bounded from below, with a minimum value of 0. Approaching this minimum drives the unlearning. However, the *key distinction* of SimNPO from NPO is its forget data-aware, length-normalized reward formulation, $(\beta/|y|) \log \pi_{\theta}(y|x)$ in (4). This eliminates the reference model bias and results in an improved gradient smoothing scheme. Specifically, the gradient of the SimNPO loss (with $\gamma = 0$) yields (as derived in Appendix A):

$$\nabla_{\theta} \ell_{\text{SimNPO}}(\theta) = \mathbb{E}_{(x,y) \in \mathcal{D}_f} \left[\underbrace{\frac{2(\pi_{\theta}(y|x))^{\beta/|y|}}{1 + (\pi_{\theta}(y|x))^{\beta/|y|}} \cdot \frac{1}{|y|}}_{:= w'_{\theta}(x,y)} \cdot \nabla_{\theta} \log \pi_{\theta}(y|x) \right]. \quad (5)$$

Similar to NPO in (3), the gradient in (5) can be divided into two components: weight smoothing (w'_{θ}) and GA. However, in SimNPO, the weight smoothing is *no longer influenced by the reference model and is instead normalized by the length* $|y|$. This introduces two key advantages (a)-(b) below, in response to NPO’s limitations (L1)-(L2).

(a) SimNPO addresses the biased allocation of unlearning power by leveraging the (data-specific) response length as a guide. For instance, when $|y|$ is large, less optimization power is allocated, as long-response forget data may have lower quality (Meng et al., 2024), is closer to the unlearning boundary, and requires less intervention (Fig. 1-(a)). In the extreme case where $\beta \rightarrow 0$, the SimNPO’s gradient reduces to a *weighted GA*: $\nabla_{\theta} \ell_{\text{SimNPO}}(\theta) \rightarrow \mathbb{E}_{(x,y) \in \mathcal{D}_f} [1/|y| \nabla_{\theta} \log \pi_{\theta}(y|x)]$. This is different from NPO, which becomes GA as $\beta \rightarrow 0$. **Fig. 3** empirically demonstrates the advantage of length normalization in SimNPO for unlearning. As shown, SimNPO outperforms NPO in both forget quality and model utility, coming closest to Retrain. Even in the special case where $\beta = 0$ (i.e., Weighted-GradDiff), the length normalization provides benefits over the vanilla GradDiff baseline.

(b) In addition, the reference-free, length-normalized weight smoothing prevents early-stage ineffectiveness during unlearning. It can be shown from (5) that $w'_{\theta}(x,y) < 2/|y|$, with the distribution of weights $w'_{\theta}(x,y)$ depending on the specific forget data samples. This contrasts with NPO, where the weight distribution concentrated around $w_{\theta}(x,y) \approx 1$ during the early unlearning stage. **Extended from Fig. 2-(a)&(b), Fig. 4** provides a detailed comparison between the gradient weights of SimNPO and NPO. As shown, SimNPO exhibits a much stronger correlation with the response length $|y|$ during the first two unlearning epochs, prioritizing short-length forget data that are initially harder to forget. At later epochs, the gradient weights become more uniform, reflecting that SimNPO can then treat different forget data with even optimization power. This trend is different from NPO, which assigns more uniform gradient weights early on and only accounts for data-specific difficulty when $w_{\theta}(x,y)$ decreases in the later stages of unlearning. **Therefore, similar to NPO, SimNPO benefits from gradient weight smoothing, enhancing unlearning stability. See Appendix C.5 for the empirical**




Figure 3: Forget quality vs. model utility on TOFU with forget set size of 5%. Weighted-GradDiff (W-GradDiff) is the variant of SimNPO at $\beta = 0$.

divergence rate of SimNPO, comparable to NPO’s logarithmic rate. More importantly, it overcomes NPO’s ineffectiveness during early optimization stages, providing further improvements.

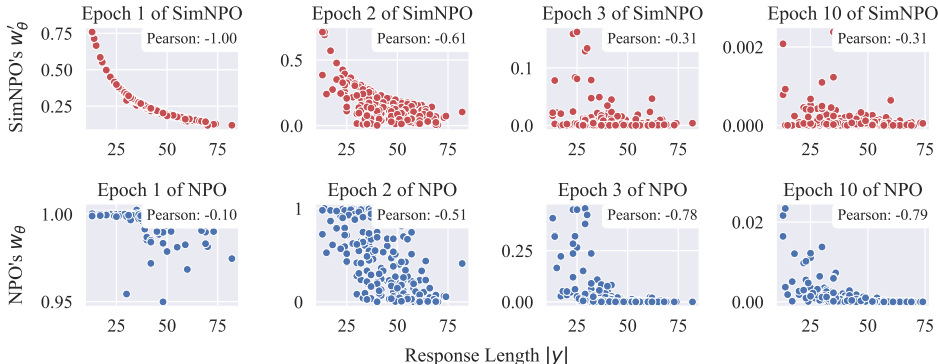


Figure 4: Gradient weight smoothing of NPO (w_θ) and SimNPO (w'_θ) vs. forget data response length $|y|$ across different epochs (1, 2, 3, and 10) on TOFU with forget set size of 5%. Each point represents a sample. The Pearson correlation in the upper right corner indicates the relationship between gradient weight smoothing and response length. The SimNPO’s weights w'_θ have been rescaled (by $\times 10$) for ease of visualization.

Further analyses via a mixture of Markov chains. In addition to the above insights, we further validate SimNPO’s advantages to overcome NPO’s limitations (L1)-(L2) (Sec. 4) using a synthetic setup. For ease of controlling the unlearning difficulties of different forget data points, we consider the problem of unlearning on a mixture of Markov chains with a state space of size 10 ($s = 1, \dots, 10$). The *retain distribution* consists of Markov chains that transition uniformly among states $\{1, 2, 3\}$. The *forget distribution* is a mixture of two components: *Forget1*, where the chains transition uniformly among $\{4, 5, 6\}$, and *Forget2*, where they move uniformly among $\{7, 8, 9\}$. A small leakage probability

allows the chains to transition outside their designated states occasionally, including state 10, which is not a designated state for any of the chains. We generate 10,000 samples for the retain distribution and 5,000 samples each for Forget1 and Forget2. A GPT-2 model is pretrained on these samples and serves as the initial model. We apply NPO and SimNPO to unlearn the forget distributions. Forget and retain performance is evaluated using the KL-divergence between predicted and true transition probabilities of the Markov chains. See Appendix B for details. We present our results in Fig. 5 and summarize the insights below.




Figure 5: Tradeoffs between forget quality (higher \uparrow is better) and retain distance (lower \downarrow is better) along the unlearning path of NPO and SimNPO in the synthetic experiments. Left: Forget1 and Forget2 have different sequence lengths. Right: unlearning from an initial model that has not seen Forget2. The symbols (\times , \bullet) near the y -axis of both figures indicate the performance of the re-trained model on Forget1 and Forget2, respectively.

In response to (L1), SimNPO is easier to unlearn short responses than NPO. To validate this, we set the retain distribution and Forget1 with a sequence length of 20, while Forget2 is assigned a shorter sequence length of 5, representing a mix of long and short responses. Fig. 5 (left) shows that NPO exhibits a worse tradeoff between retain distance and forget quality on short responses (*i.e.*, Forget2) compared with SimNPO. That is, to achieve the same forget quality on Forget2 as the retrained model (with forget quality 0.44), NPO incurs a higher retain distance than SimNPO. As a result, NPO has an overall larger retain distance when unlearning the entire Forget distribution. In contrast, SimNPO shows more consistent performance across Forget1 and Forget2, with less variance in its tradeoff.

In response to (L2), SimNPO unlearns easy-to-forget data less aggressively than NPO. In the second case, we set the retain distribution, Forget1 and Forget2 all with a sequence length of 20. However, we exclude Forget2 during pretraining. This setup simulates a scenario where the initial model (*i.e.*, the reference model in NPO) has already unlearned part of the forget dataset (*i.e.*, Forget2). Fig. 5 (right) shows that NPO unlearns Forget2 faster than SimNPO, even though Forget2 was already unlearned. However, NPO performs worse on Forget1 than SimNPO, likely due to overlearning Forget2, thereby reducing the overall model utility.

6 EXPERIMENTS

6.1 EXPERIMENT SETUPS

Datasets, tasks, and models. Our experiments cover unlearning tasks across three benchmark datasets: TOFU (Maini et al., 2024), MUSE (Shi et al., 2024), and WMDP (Li et al., 2024), as summarized in Table 1. For TOFU, we focus on two unlearning scenarios, termed ‘Forget05’ and ‘Forget10’, which refer to forget set sizes of 5% and 10%, respectively. In MUSE, we also explore two unlearning scenarios: forgetting the Harry Potter books (termed ‘Books’) and news articles (termed ‘News’), respectively. WMDP, on the other hand, is designed for knowledge-based unlearning, with the forget texts representing hazardous knowledge in biosecurity and cybersecurity. The LLM models used for each unlearning benchmark are listed in Table 1.

LLM unlearning methods and evaluation. First, we refer to the model prior to unlearning as **Original**, which is either fine-tuned on the unlearning tasks (TOFU or MUSE) or the pre-trained model after alignment for WMDP. Starting from the original model, we then apply the following unlearning methods to a given forget set and/or retain set to achieve the unlearning objective, as outlined in (1). Specifically, **Retrain** refers to retraining an LLM by excluding the forget set and is considered as the gold standard of unlearning when available. Retrain is provided in both the TOFU and MUSE benchmarks. As introduced in Sec. 3, we also include **GA** (gradient ascent) and **GradDiff** (the retain-regularized GA variant) as unlearning baseline methods, following the implementations in TOFU and MUSE benchmarks. For other baseline methods such as the rejection-based unlearning method (**IDK**) in TOFU, and the **Task Vector** unlearning method in MUSE, we adhere to the original implementations specified in their respective benchmarks. **NPO** with the retain regularization in (1) serves as the primary baseline. Note that its implementation on TOFU follows the original NPO study (Zhang et al., 2024a), while its implementation on MUSE aligns with the MUSE benchmark. For NPO on WMDP, due to the absence of open-source implementation, we adapt the TOFU codebase to WMDP. More implementation details can be found in Appendix C.2. To implement the proposed method **SimNPO**, we adopt a setting similar to NPO but adjust the temperature parameter β . Due to the presence of length normalization in (4), a larger value for β is preferred compared to that in NPO. See the specific choices in Appendix C.4.

To assess unlearning effectiveness and model utility, we use the evaluation metrics summarized in Table 1 under each unlearning benchmark. In addition, we evaluate the robustness of an unlearned model using relearning-based attacks (Hu et al., 2024), which aim to recover the forgotten information by fine-tuning the unlearned models on a small subset of the forget set after unlearning. We select 20% of the original TOFU forget05 set as the relearning set over three epochs.

6.2 EXPERIMENT RESULTS

Table 2: Performance overview of various unlearning methods on TOFU using the LLaMA2-7B-chat model across two unlearning settings: Forget05 and Forget10. ‘Prob.’ indicates the probability metrics, as summarized in Table 1, with forget quality (FQ) and model utility (MU) serving as the primary metrics. Results are averaged over five independent random trials. The best FQ and MU is highlighted in **bold**.

| Method | Unlearning Efficacy | | | | Utility Preservation | | | | | | | | | |
|----------------------|---------------------|------------------------|--------------------|---------------|----------------------|------------------------|--------------------|------------------|------------------------|--------------------|------------------|------------------------|------|---------------|
| | Forget Set | | | FQ \uparrow | Real Authors | | | World Facts | | | Retain Set | | | MU \uparrow |
| l-Rouge-L \uparrow | l-Prob. \uparrow | Truth ratio \uparrow | Rouge-L \uparrow | | Prob. \uparrow | Truth ratio \uparrow | Rouge-L \uparrow | Prob. \uparrow | Truth ratio \uparrow | Rouge-L \uparrow | Prob. \uparrow | Truth ratio \uparrow | | |
| TOFU Forget05 | | | | | | | | | | | | | | |
| Original | 0.04 | 0.01 | 0.49 | 0.00 | 0.93 | 0.44 | 0.58 | 0.91 | 0.43 | 0.55 | 0.98 | 0.99 | 0.48 | 0.62 |
| Retrain | 0.61 | 0.85 | 0.66 | 1.00 | 0.92 | 0.44 | 0.57 | 0.90 | 0.43 | 0.54 | 0.97 | 0.99 | 0.48 | 0.62 |
| GA | 0.00 | 0.00 | 0.66 | 1.87e-09 | 0.00 | 0.20 | 0.40 | 0.00 | 0.30 | 0.28 | 0.00 | 0.00 | 0.15 | 0.00 |
| GradDiff | 0.00 | 0.00 | 0.60 | 3.60e-09 | 0.59 | 0.59 | 0.81 | 0.88 | 0.46 | 0.59 | 0.42 | 0.49 | 0.48 | 0.56 |
| IDK | 0.02 | 0.60 | 0.55 | 1.87e-09 | 0.65 | 0.48 | 0.63 | 0.82 | 0.44 | 0.55 | 0.55 | 0.86 | 0.43 | 0.57 |
| NPO | 0.26 | 0.06 | 0.69 | 0.79 | 0.91 | 0.50 | 0.62 | 0.90 | 0.50 | 0.61 | 0.47 | 0.51 | 0.44 | 0.57 |
| SimNPO | 0.28 | 0.03 | 0.66 | 0.99 | 0.90 | 0.50 | 0.64 | 0.90 | 0.48 | 0.60 | 0.54 | 0.56 | 0.44 | 0.58 |
| TOFU Forget10 | | | | | | | | | | | | | | |
| Original | 0.03 | 0.01 | 0.48 | 0.00 | 0.93 | 0.44 | 0.58 | 0.91 | 0.43 | 0.55 | 0.98 | 0.99 | 0.48 | 0.62 |
| Retrain | 0.61 | 0.84 | 0.67 | 1.00 | 0.93 | 0.45 | 0.59 | 0.91 | 0.42 | 0.54 | 0.98 | 0.99 | 0.47 | 0.62 |
| GA | 0.00 | 0.00 | 0.70 | 2.19e-16 | 0.00 | 0.28 | 0.37 | 0.00 | 0.29 | 0.31 | 0.00 | 0.00 | 0.11 | 0.00 |
| GradDiff | 0.00 | 0.00 | 0.67 | 3.71e-15 | 0.44 | 0.49 | 0.67 | 0.89 | 0.48 | 0.58 | 0.48 | 0.60 | 0.46 | 0.54 |
| IDK | 0.02 | 0.63 | 0.54 | 2.86e-14 | 0.46 | 0.45 | 0.59 | 0.84 | 0.43 | 0.55 | 0.56 | 0.88 | 0.44 | 0.54 |
| NPO | 0.22 | 0.09 | 0.70 | 0.29 | 0.91 | 0.52 | 0.66 | 0.85 | 0.48 | 0.61 | 0.44 | 0.46 | 0.39 | 0.55 |
| SimNPO | 0.22 | 0.10 | 0.71 | 0.45 | 0.90 | 0.54 | 0.70 | 0.88 | 0.50 | 0.64 | 0.54 | 0.76 | 0.47 | 0.62 |

Performance on TOFU. In Table 2, we present the unlearning performance of SimNPO and its various baselines on TOFU, covering both effectiveness metrics and utility metrics as shown in Table 1. Recall that ‘Original’ refers to the model performance prior to unlearning, serving as the *lower bound* for unlearning effectiveness. In contrast, ‘Retrain’ refers to the model retrained excluding the forget set influence, serving as the *upper bound* for unlearning effectiveness. ‘FQ’ stands for forget quality, and ‘MU’ represents model utility. These two metrics serve as the primary performance

486 indicators for LLM unlearning on TOFU. SimNPO outperforms NPO in both FQ and MU, and is the
 487 closest approximate unlearning method to Retrain. Except for NPO, the other unlearning baselines
 488 (GA, GradDiff, and IDK) are not effective, as implied by their FQ values being smaller than 0.01,
 489 where FQ indicates the p -value for rejecting the indistinguishability between the unlearned model
 490 and Retrain on TOFU. In **Table A5 of Appendix D**, we also provide examples of model responses
 491 after unlearning using SimNPO, Retrain, and NPO, along with label to degenerate. We observe
 492 that, in some cases (*e.g.*, responses against Q1 and Q2 in Table A5), the NPO-unlearned model
 493 generates *repeated texts* in response. While this repetition does not reveal the information intended
 494 for unlearning, it negatively impacts model utility and differs noticeably from Retrain’s behavior.
 495 In contrast, SimNPO produces unlearning responses more closely aligned with those generated by
 496 Retrain. We also find that SimNPO is more efficient than NPO in **Appendix C.6**.

497 **Performance on MUSE and WMDP.** **Table A3 of Appendix C.7** compares the performance of
 498 SimNPO with baseline methods, including Task Vector (Shi et al., 2024; Ilharco et al., 2022), on both
 499 the MUSE News and Books datasets. The evaluation metrics are summarized in Table 1, with PrivLeak
 500 serving as the primary metric to indicate the gap with Retrain. As we can see, SimNPO consistently
 501 achieves PrivLeak values closest to 0 for both News (11.90) and Books (−19.82) compared to other
 502 unlearning baselines, suggesting that it is most aligned with complete forget data removal, as defined
 503 in MUSE (Shi et al., 2024). **Fig. A3 of Appendix C.7** further demonstrates SimNPO’s advantage over
 504 NPO on the News dataset in addressing the over-unlearning issue. We compare the distribution of text
 505 memorization scores, measured by Min-K% probability (Shi et al., 2023), across Retrain, SimNPO,
 506 and NPO at early (epoch 3) and later (epoch 10) stages. As shown, NPO results in an over-forgetting
 507 distribution, with a significantly larger distance between the forget set and holdout set. SimNPO, by
 508 contrast, shows a closer distribution to Retrain. This is also consistent with the NPO’s limitation (L2)
 509 as illustrated in Sec. 4. **Furthermore, we remark that the advantages of SimNPO are not solely due**
 510 **to its awareness of forget data response length, as the MUSE dataset features forget data with equal**
 511 **response lengths. We further evaluate SimNPO’s performance in hazardous knowledge unlearning on**
 512 **WMDP, as detailed in Appendix C.8.**

512 **Unlearning robustness against relearning attack.** Given recent studies highlighting the vulnera-
 513 bility of unlearning methods to relearning attacks (Lynch et al., 2024; Hu et al., 2024)—where the
 514 forgotten information can be recovered by finetuning the unlearned model on a small subset of
 515 the forget set—we aim to evaluate the robustness of SimNPO, particularly in comparison to NPO,
 516 against such attacks. Our rationale is that, since SimNPO outperforms NPO in forgetting short-length
 517 response data, it should also enhance robustness against relearning attacks on this type of forget data,
 518 provided the unlearning from SimNPO is faithful.

519 **Fig. 6** presents the forget quality of SimNPO and NPO under relearning attacks against the number of relearning epochs. Relearning is
 520 performed on the forget subset, which is either the shortest 20% of
 521 responses from the TOFU Forget05 dataset or an equal-size random
 522 subset. We refer to these attacks as ‘shortest-relearn’ and ‘random-
 523 relearn’, respectively. The random-relearn case is conducted 5 times,
 524 with both average robustness and variance in Fig. 6. As we can see,
 525 SimNPO demonstrates improved robustness over NPO, evidenced
 526 by higher forget quality and a slower decline in forget quality as the
 527 relearning epoch increases. Moreover, NPO is less robust against
 528 the shortest-relearn attack compared to the random-relearn attack. In
 529 contrast, SimNPO is resilient to both types of relearning. This is expected since SimNPO addresses
 530 the limitation (L1), as explained in Sec. 4.




Figure 6: Forget quality for NPO and SimNPO under random/shortest relearn attack vs. relearning epochs on TOFU Forget05.

532 **7 CONCLUSION**

533 We revisited the current unlearning optimization framework, negative preference optimization (NPO),
 534 and identified its reference model bias issue, which compromises unlearning effectiveness, particularly
 535 for forget data of varying difficulty. To address this, we introduced SimNPO, a simple yet effective
 536 framework that eliminates reliance on a reference model by leveraging simple preference optimization.
 537 We provided deep insights into SimNPO’s advantages through both synthetic data analysis and
 538 evaluations on existing unlearning benchmarks such as TOFU, MUSE, WMDP, and relearning
 539 attacks. In future work, we will further investigate the limitations of SimNPO and enhance it for tasks
 involving model capability removal. See further discussions in Appendix E-F.

REFERENCES

- 540
541
542 Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal
543 Valko, and Daniele Calandriello. A general theoretical paradigm to understand learning from
544 human preferences. In *International Conference on Artificial Intelligence and Statistics*, pp.
545 4447–4455. PMLR, 2024.
- 546 Clark Barrett, Brad Boyd, Elie Bursztein, Nicholas Carlini, Brad Chen, Jihye Choi, Amrita Roy
547 Chowdhury, Mihai Christodorescu, Anupam Datta, Soheil Feizi, et al. Identifying and mitigating
548 the security risks of generative ai. *Foundations and Trends® in Privacy and Security*, 6(1):1–52,
549 2023.
- 550
551 Alberto Blanco-Justicia, Najeeb Jebreel, Benet Manzanares, David Sánchez, Josep Domingo-Ferrer,
552 Guillem Collell, and Kuan Eeik Tan. Digital forgetting in large language models: A survey of
553 unlearning methods. *arXiv preprint arXiv:2404.02062*, 2024.
- 554 Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers,
555 Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE Symposium*
556 *on Security and Privacy (SP)*, pp. 141–159. IEEE, 2021.
- 557
558 Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015*
559 *IEEE symposium on security and privacy*, pp. 463–480. IEEE, 2015.
- 560
561 Min Chen, Zhikun Zhang, Tianhao Wang, Michael Backes, Mathias Humbert, and Yang Zhang.
562 Graph unlearning. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and*
563 *Communications Security*, pp. 499–513, 2022.
- 564
565 Eli Chien, Chao Pan, and Olgica Milenkovic. Certified graph unlearning. *arXiv preprint*
566 *arXiv:2206.09140*, 2022.
- 567
568 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep
569 reinforcement learning from human preferences. *Advances in neural information processing*
570 *systems*, 30, 2017.
- 571
572 Ronen Eldan and Mark Russinovich. Who’s harry potter? approximate unlearning in llms, 2023.
- 573
574 Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Kto: Model
575 alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*, 2024.
- 576
577 Chongyu Fan, Jiancheng Liu, Alfred Hero, and Sijia Liu. Challenging forgets: Unveiling the
578 worst-case forget sets in machine unlearning. *arXiv preprint arXiv:2403.07362*, 2024a.
- 579
580 Chongyu Fan, Jiancheng Liu, Yihua Zhang, Dennis Wei, Eric Wong, and Sijia Liu. Salun: Em-
581 powering machine unlearning via gradient-based weight saliency in both image classification and
582 generation. In *International Conference on Learning Representations*, 2024b.
- 583
584 Rohit Gandikota, Joanna Materzynska, Jaden Fiotto-Kaufman, and David Bau. Erasing concepts
585 from diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer*
586 *Vision*, pp. 2426–2436, 2023.
- 587
588 Anisa Halimi, Swanand Kadhe, Ambrish Rawat, and Nathalie Baracaldo. Federated unlearning: How
589 to efficiently erase a client in fl? *arXiv preprint arXiv:2207.05521*, 2022.
- 590
591 Jiwoo Hong, Noah Lee, and James Thorne. Reference-free monolithic preference optimization with
592 odds ratio. *arXiv preprint arXiv:2403.07691*, 2024.
- 593
594 Chris Jay Hoofnagle, Bart van der Sloot, and Frederik Zuiderveen Borgesius. The european union
595 general data protection regulation: what it is and what it means. *Information & Communications*
596 *Technology Law*, 28(1):65–98, 2019.
- 597
598 Shengyuan Hu, Yiwei Fu, Zhiwei Steven Wu, and Virginia Smith. Jogging the memory of unlearned
599 model through targeted relearning attack. *arXiv preprint arXiv:2406.13356*, 2024.

- 594 Yue Huang, Lichao Sun, Haoran Wang, Siyuan Wu, Qihui Zhang, Yuan Li, Chujie Gao, Yixin Huang,
595 et al. Position: TrustLLM: Trustworthiness in large language models. In *Proceedings of the 41st*
596 *International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning*
597 *Research*, pp. 20166–20270, 21–27 Jul 2024.
- 598 Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt,
599 Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *arXiv preprint*
600 *arXiv:2212.04089*, 2022.
- 601 Yoichi Ishibashi and Hidetoshi Shimodaira. Knowledge sanitization of large language models. *arXiv*
602 *preprint arXiv:2309.11852*, 2023.
- 603 Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and
604 Minjoon Seo. Knowledge unlearning for mitigating privacy risks in language models. *arXiv*
605 *preprint arXiv:2210.01504*, 2022.
- 606 Jiabao Ji, Yujian Liu, Yang Zhang, Gaowen Liu, Ramana Rao Kompella, Sijia Liu, and Shiyu Chang.
607 Reversing the forget-retain objectives: An efficient llm unlearning framework from logit difference.
608 *arXiv preprint arXiv:2406.08607*, 2024.
- 609 Jinghan Jia, Jiancheng Liu, Parikshit Ram, Yuguang Yao, Gaowen Liu, Yang Liu, Pranay Sharma,
610 and Sijia Liu. Model sparsity can simplify machine unlearning. In *Thirty-seventh Conference on*
611 *Neural Information Processing Systems*, 2023.
- 612 Jinghan Jia, Yihua Zhang, Yimeng Zhang, Jiancheng Liu, Bharat Runwal, James Diffenderfer,
613 Bhavya Kailkhura, and Sijia Liu. Soul: Unlocking the power of second-order optimization for llm
614 unlearning. *arXiv preprint arXiv:2404.18239*, 2024.
- 615 Ruinan Jin, Minghui Chen, Qiong Zhang, and Xiaoxiao Li. Forgettable federated linear learning with
616 certified data removal. *arXiv preprint arXiv:2306.02216*, 2023.
- 617 Swanand Ravindra Kadhe, Farhan Ahmed, Dennis Wei, Nathalie Baracaldo, and Inkit Padhi. Split,
618 unlearn, merge: Leveraging data attributes for more effective unlearning in llms. *arXiv preprint*
619 *arXiv:2406.11780*, 2024.
- 620 Vinayshekhar Bannihatti Kumar, Rashmi Gangadharaiah, and Dan Roth. Privacy adhering machine
621 un-learning in nlp. *arXiv preprint arXiv:2212.09573*, 2022.
- 622 Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded
623 machine unlearning. *Advances in neural information processing systems*, 36, 2024.
- 624 Nathaniel Li, Alexander Pan, Anjali Gopal, Summer Yue, Daniel Berrios, Alice Gatti, Justin D Li,
625 Ann-Kathrin Dombrowski, Shashwat Goel, Long Phan, et al. The wmdp benchmark: Measuring
626 and reducing malicious use with unlearning. *arXiv preprint arXiv:2403.03218*, 2024.
- 627 Bo Liu, Qiang Liu, and Peter Stone. Continual learning and private unlearning. In *Conference on*
628 *Lifelong Learning Agents*, pp. 243–254. PMLR, 2022a.
- 629 Chris Yuhao Liu, Yaxuan Wang, Jeffrey Flanigan, and Yang Liu. Large language model unlearning
630 via embedding-corrupted prompts. *arXiv preprint arXiv:2406.07933*, 2024a.
- 631 Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Xiaojun Xu,
632 Yuguang Yao, Hang Li, Kush R Varshney, et al. Rethinking machine unlearning for large language
633 models. *arXiv preprint arXiv:2402.08787*, 2024b.
- 634 Yang Liu, Mingyuan Fan, Cen Chen, Ximeng Liu, Zhuo Ma, Li Wang, and Jianfeng Ma. Backdoor
635 defense with machine unlearning. In *IEEE INFOCOM 2022-IEEE Conference on Computer*
636 *Communications*, pp. 280–289. IEEE, 2022b.
- 637 Yi Liu, Lei Xu, Xingliang Yuan, Cong Wang, and Bo Li. The right to be forgotten in federated
638 learning: An efficient realization with rapid retraining. In *IEEE INFOCOM 2022-IEEE Conference*
639 *on Computer Communications*, pp. 1749–1758. IEEE, 2022c.

- 648 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint*
649 *arXiv:1711.05101*, 2017.
- 650
- 651 Ximing Lu, Sean Welleck, Jack Hessel, Liwei Jiang, Lianhui Qin, Peter West, Prithviraj Am-
652 manabrolu, and Yejin Choi. Quark: Controllable text generation with reinforced unlearning.
653 *Advances in neural information processing systems*, 35:27591–27609, 2022.
- 654
- 655 Aengus Lynch, Phillip Guo, Aidan Ewart, Stephen Casper, and Dylan Hadfield-Menell. Eight
656 methods to evaluate robust unlearning in llms. *arXiv preprint arXiv:2402.16835*, 2024.
- 657
- 658 Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C. Lipton, and J. Zico Kolter. Tofu: A task
659 of fictitious unlearning for llms, 2024.
- 660
- 661 Yu Meng, Mengzhou Xia, and Danqi Chen. Simpo: Simple preference optimization with a reference-
662 free reward. *arXiv preprint arXiv:2405.14734*, 2024.
- 663
- 664 Thanh Tam Nguyen, Thanh Trung Huynh, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin,
665 and Quoc Viet Hung Nguyen. A survey of machine unlearning. *arXiv preprint arXiv:2209.02299*,
2022.
- 666
- 667 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
668 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow
669 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–
27744, 2022.
- 670
- 671 Vaidehi Patil, Peter Hase, and Mohit Bansal. Can sensitive information be deleted from llms?
672 objectives for defending against extraction attacks. *ICLR*, 2024.
- 673
- 674 Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models
675 as few shot unlearners. *arXiv preprint arXiv:2310.07579*, 2023.
- 676
- 677 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
678 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 679
- 680 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
681 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances*
in Neural Information Processing Systems, 36, 2024.
- 682
- 683 Jeffrey Rosen. The right to be forgotten. *Stan. L. Rev. Online*, 64:88, 2011.
- 684
- 685 Michael Santacrose, Yadong Lu, Han Yu, Yuanzhi Li, and Yelong Shen. Efficient rlhf: Reducing the
686 memory usage of ppo. *arXiv preprint arXiv:2309.00754*, 2023.
- 687
- 688 Avi Schwarzschild, Zhili Feng, Pratyush Maini, Zachary C Lipton, and J Zico Kolter. Rethinking
689 llm memorization through the lens of adversarial compression. *arXiv preprint arXiv:2404.15146*,
2024.
- 690
- 691 Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen,
692 and Luke Zettlemoyer. Detecting pretraining data from large language models. *arXiv preprint*
arXiv:2310.16789, 2023.
- 693
- 694 Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao
695 Liu, Luke Zettlemoyer, Noah A Smith, and Chiyuan Zhang. Muse: Machine unlearning six-way
696 evaluation for language models. *arXiv preprint arXiv:2407.06460*, 2024.
- 697
- 698 Pratiksha Thaker, Yash Maurya, and Virginia Smith. Guardrail baselines for unlearning in llms. *arXiv*
preprint arXiv:2403.03329, 2024.
- 699
- 700 Anvith Thudi, Gabriel Deza, Varun Chandrasekaran, and Nicolas Papernot. Unrolling sgd: Under-
701 standing factors influencing machine unlearning. In *2022 IEEE 7th European Symposium on*
Security and Privacy (EuroS&P), pp. 303–319. IEEE, 2022.

- 702 Eleni Triantafillou, Peter Kairouz, Fabian Pedregosa, Jamie Hayes, Meghdad Kurmanji, Kairan Zhao,
703 Vincent Dumoulin, Julio Jacques Junior, Ioannis Mitliagkas, Jun Wan, et al. Are we making
704 progress in unlearning? findings from the first neurips unlearning competition. *arXiv preprint*
705 *arXiv:2406.09073*, 2024.
- 706 Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu,
707 Zidi Xiong, Ritik Dutta, Rylan Schaeffer, et al. Decodingtrust: A comprehensive assessment of
708 trustworthiness in gpt models. In *NeurIPS*, 2023.
- 709 Yu Wang, Ruihan Wu, Zexue He, Xiusi Chen, and Julian McAuley. Large scale knowledge washing.
710 *arXiv preprint arXiv:2405.16720*, 2024.
- 711 Alexander Warnecke, Lukas Pirch, Christian Wressnegger, and Konrad Rieck. Machine unlearning
712 of features and labels. *arXiv preprint arXiv:2108.11577*, 2021.
- 713 Kun Wu, Jie Shen, Yue Ning, Ting Wang, and Wendy Hui Wang. Certified edge unlearning for graph
714 neural networks. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery*
715 *and Data Mining*, pp. 2606–2617, 2023a.
- 716 Xinwei Wu, Junzhuo Li, Minghui Xu, Weilong Dong, Shuangzhi Wu, Chao Bian, and Deyi Xiong.
717 Depn: Detecting and editing privacy neurons in pretrained language models. *arXiv preprint*
718 *arXiv:2310.20138*, 2023b.
- 719 Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. *arXiv preprint*
720 *arXiv:2310.10683*, 2023.
- 721 Charles Yu, Sullam Jeoung, Anish Kasi, Pengfei Yu, and Heng Ji. Unlearning bias in language
722 models by partitioning gradients. In *Findings of the Association for Computational Linguistics:*
723 *ACL 2023*, pp. 6032–6048, 2023.
- 724 Hongyi Yuan, Zheng Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. Rrhf: Rank
725 responses to align language models with human feedback. *Advances in Neural Information*
726 *Processing Systems*, 36, 2024.
- 727 Eric Zhang, Kai Wang, Xingqian Xu, Zhangyang Wang, and Humphrey Shi. Forget-me-not: Learning
728 to forget in text-to-image diffusion models. *arXiv preprint arXiv:2303.17591*, 2023.
- 729 Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic
730 collapse to effective unlearning. *arXiv preprint arXiv:2404.05868*, 2024a.
- 731 Yihua Zhang, Yimeng Zhang, Yuguang Yao, Jinghan Jia, Jiancheng Liu, Xiaoming Liu, and Sijia Liu.
732 Unlearncanvas: A stylized image dataset to benchmark machine unlearning for diffusion models.
733 *arXiv preprint arXiv:2402.11846*, 2024b.
- 734 Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. Slic-hf:
735 Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*, 2023.
- 736 Rui Zheng, Shihan Dou, Songyang Gao, Yuan Hua, Wei Shen, Binghai Wang, Yan Liu, Senjie Jin,
737 Qin Liu, Yuhao Zhou, et al. Secrets of rlhf in large language models part i: Ppo. *arXiv preprint*
738 *arXiv:2307.04964*, 2023.
- 739 Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul
740 Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv*
741 *preprint arXiv:1909.08593*, 2019.
- 742
743
744
745
746
747
748
749
750
751
752
753
754
755

A GRADIENT ANALYSIS OF SIMNPO

Following is the detailed derivation of (5). First, let $R = \frac{\log \pi_{\theta}(y|x) + \gamma|y|/\beta}{|y|}$. We then have the following steps:

$$\nabla_{\theta} \ell_{\text{SimNPO}}(\theta) = \mathbb{E}_{(x,y) \in \mathcal{D}_f} \nabla_{\theta} \left[-\frac{2}{\beta} \log \sigma(-\beta R) \right] \quad (\text{A1})$$

$$= \mathbb{E}_{(x,y) \in \mathcal{D}_f} \nabla_{\theta} \left[\frac{2}{\beta} \log \sigma(1 + \exp(\beta R)) \right] \quad (\text{A2})$$

$$= \mathbb{E}_{(x,y) \in \mathcal{D}_f} \left[\frac{2}{\beta} \cdot \frac{\beta \exp(\beta R)}{1 + \exp(\beta R)} \cdot \nabla_{\theta} R \right] \quad (\text{A3})$$

$$= \mathbb{E}_{(x,y) \in \mathcal{D}_f} \left[\frac{2 \exp(\beta \frac{\log \pi_{\theta}(y|x) + \gamma|y|/\beta}{|y|})}{1 + \exp(\beta \frac{\log \pi_{\theta}(y|x) + \gamma|y|/\beta}{|y|})} \cdot \frac{1}{|y|} \cdot \nabla_{\theta} \log \pi_{\theta}(y|x) \right] \quad (\text{A4})$$

When $\gamma = 0$, the gradient simplifies to the following, which matches (5):

$$\nabla_{\theta} \ell_{\text{SimNPO}}(\theta) = \mathbb{E}_{(x,y) \in \mathcal{D}_f} \left[\frac{2 \exp(\beta \frac{\log \pi_{\theta}(y|x)}{|y|})}{1 + \exp(\beta \frac{\log \pi_{\theta}(y|x)}{|y|})} \cdot \frac{1}{|y|} \cdot \nabla_{\theta} \log \pi_{\theta}(y|x) \right] \quad (\text{A5})$$

$$= \mathbb{E}_{(x,y) \in \mathcal{D}_f} \left[\frac{2(\pi_{\theta}(y|x))^{\beta/|y|}}{1 + (\pi_{\theta}(y|x))^{\beta/|y|}} \cdot \frac{1}{|y|} \cdot \nabla_{\theta} \log \pi_{\theta}(y|x) \right] \quad (\text{A6})$$

B ADDITIONAL DETAILS ON THE SYNTHETIC STUDY

Synthetic experiment setup. In the synthetic experiment, we study the unlearning problem in a scenario where the data are generated from a mixture of Markov chains. Namely, we assume the Markov chains have a shared state space of size 10 (denoted by $s = 1, 2, \dots, 10$), and the retain distribution and the forget distribution have the formulas as follows:

• **Retain distribution:** Markov chain with initial distribution $\pi_r \in \mathbb{R}^{10}$ and transition matrix $T_r \in \mathbb{R}^{10 \times 10}$, where

$$\begin{aligned} \pi_{r,j} &= \frac{1-\epsilon}{3} \quad \text{for } j \leq 3, & \pi_{r,j} &= \frac{\epsilon}{7} \quad \text{for } j \geq 4. \\ T_{r,i} &= \pi_r \quad \text{for } i \leq 3, & T_{r,i} &= 0.1 \cdot \mathbf{1}_{10} \quad \text{for } i \geq 4. \end{aligned}$$

• **Forget distribution:** a mixture of two Markov chains (denoted by Forget1 and Forget2) with equal probability. Let (π_{f_1}, T_{f_1}) and (π_{f_2}, T_{f_2}) denote the initial distribution and transition matrix for Forget1 and Forget2. We assume

$$\begin{aligned} \pi_{f_1,j} &= \frac{1-\epsilon}{3} \quad \text{for } j \in \{4, 5, 6\}, & \pi_{f_1,j} &= \frac{\epsilon}{7} \quad \text{for } j \notin \{4, 5, 6\}, \\ T_{f_1,i} &= \pi_{f_1} \quad \text{for } i \in \{4, 5, 6\}, & T_{f_1,i} &= 0.1 \cdot \mathbf{1}_{10} \quad \text{for } i \notin \{4, 5, 6\}, \end{aligned}$$

and

$$\begin{aligned} \pi_{f_2,j} &= \frac{1-\epsilon}{3} \quad \text{for } j \in \{7, 8, 9\}, & \pi_{f_2,j} &= \frac{\epsilon}{7} \quad \text{for } j \notin \{7, 8, 9\}, \\ T_{f_2,i} &= \pi_{f_2} \quad \text{for } i \in \{7, 8, 9\}, & T_{f_2,i} &= 0.1 \cdot \mathbf{1}_{10} \quad \text{for } i \notin \{7, 8, 9\}. \end{aligned}$$

The leakage probability is chosen to be $\epsilon = 0.2$. We generate 10000 samples from the retain distribution and 5000 each from Forget1 and Forget2 to form the retain and forget sets. We randomly split the datasets, using 80% of the samples for training and unlearning, and the remaining 20% for testing.

Model and pretraining. In all experiments, we use a small GPT-2 model (Radford et al., 2019) with modified token embeddings, where input tokens represent states in $\mathcal{S} = \{1, 2, \dots, 10\}$, and the output at each token position is a distribution over the state space \mathcal{S} . The model has 4 transformer layers, 4 attention heads, and an embedding dimension of 128. We pretrain the original model on both retain and forget data, and the retrained model using only the forget data. Both models are trained using AdamW (Loshchilov & Hutter, 2017) to minimize the cross-entropy loss averaged over tokens, with a batch size of 128 for 5 epochs. We choose the learning rate $\eta = 0.0005$.

Evaluation. We evaluate the model performance using Forget Quality (higher \uparrow is better) and Retain Loss (lower \downarrow is better), which are the average KL divergence between the predicted probabilities of the model and the true transition probabilities of the Markov chains, on the forget (Forget1 or Forget2) and the retain test data, respectively.

Unlearning. Starting from the initial model, we run NPO and SimNPO for 50 iterations using a batch size of 4 on the forget dataset. We choose AdamW for optimization with a learning rate of $\eta = 0.0005$. The hyperparameter β in both NPO and SimNPO is selected via grid search to optimize the tradeoff between forget quality and retain loss.

Choice of hyperparameters. In the first experiment (**Fig. 5 left**), we set the hyperparameters $\beta_{\text{NPO}} = 0.2, \beta_{\text{SimNPO}} = 4$, the retain sample length $L_r = 20$, and the Forget1 and Forget2 sample lengths $L_{f_1} = 20, L_{f_2} = 5$. In the second experiment (**Fig. 5 right**), we choose $\beta_{\text{NPO}} = 1.0, \beta_{\text{SimNPO}} = 4$, the retain sample length $L_r = 20$, and the Forget1 and Forget2 sample lengths $L_{f_1} = 20, L_{f_2} = 20$.

C ADDITIONAL EXPERIMENT DETAILS AND RESULTS

C.1 COMPUTE CONFIGURATIONS

All experiments are conducted on 8 NVIDIA A6000 GPU cards in a single node.

C.2 EXPERIMENT SETUPS

C.2.1 TOFU EXPERIMENT SETUP

For all experiments, we use a linear warm-up learning rate during the first epoch, followed by a linearly decaying learning rate in the remaining epochs. We initialize the process with LLaMA-2 7B and fine-tune the model on TOFU for 5 epochs with a batch size of 32 and a learning rate of 10^{-5} to obtain the original model. For Forget05, NPO is trained for up to 20 epochs with a learning rate of 10^{-5} to obtain the best-performing model. We conducted a grid search for β in the range of $[0.05, 0.2]$ and for λ in the range of $[0.5, 1.5]$. SimNPO is trained for 10 epochs with a learning rate of 10^{-5} . The parameter β is grid-searched over the range $[1.5, 3.5]$, γ is searched between $[0.0, 2.0]$ with the default choice $\gamma = 0$, and λ is explored within the range $[0.05, 0.25]$. For Forget10, NPO is trained for 10 epochs with a learning rate of 10^{-5} . We conducted a grid search for β in the range of $[0.05, 0.2]$ and for λ in the range of $[0.5, 1.5]$. SimNPO is trained for 10 epochs with a learning rate of 10^{-5} . The parameter β is tuned using a grid search within the range $[2.5, 5.5]$, γ is grid-searched between $[0.0, 2.0]$, and λ is grid-searched within $[0.05, 0.25]$. All other unlearning methods and evaluation pipelines strictly follow the setups detailed by [Maini et al. \(2024\)](#) and [Zhang et al. \(2024a\)](#).

C.2.2 MUSE EXPERIMENT SETUP

For News, we use LLaMA-2 7B fine-tuned on BBC news articles as the original model. For Books, we use ICLM 7B fine-tuned on the Harry Potter books as the original model. The original models for both Books and News can be directly obtained from benchmark. For SimNPO, we trained for 10 epochs with a learning rate of 10^{-5} . We performed a grid search for β in the range of $[0.5, 1.0]$, for λ in the range of $[0.05, 0.25]$, and for γ in the range of $[0.0, 2.0]$ on both the Books and News. The hyperparameters for other unlearning methods and the evaluation pipelines strictly follow the setup detailed by [Shi et al. \(2024\)](#). We measured the performance after each unlearning epoch and selected the optimal one as the final model.

C.2.3 WMDP EXPERIMENT SETUP

For WMDP ([Li et al., 2024](#)), we use Zephyr-7B-beta, provided as the origin model in the benchmark. A forget set consisting of plain texts related to biosecurity/cybersecurity knowledge and an unrelated text retain set are used. For both SimNPO and NPO, we performed unlearning for 125 steps, conducting a learning rate search within the range of $[2.5 \times 10^{-6}, 5 \times 10^{-6}]$ and a grid search for β in the range of $[0.05, 7.5]$, with λ fixed at 5.0.

C.3 UNLEARNING STRONGLY VS. WEAKLY-MEMORIZED FORGET DATA

We use TOFU Forget05 as the forget set \mathcal{D}_f , splitting it evenly into $\mathcal{D}_{f,1}$ and $\mathcal{D}_{f,2}$. The divided subsets $\mathcal{D}_{f,1}$ and $\mathcal{D}_{f,2}$ follow the same distribution of fictitious author information. We fine-tune the LLaMA-2 7B chat model on the original retain set of TOFU together with $\mathcal{D}_{f,1}$, i.e., $\mathcal{D}_{\text{retain}} \cup \mathcal{D}_{f,1}$, to obtain the reference model before unlearning. The resulting reference model strongly memorizes $\mathcal{D}_{f,1}$ but least memorizes $\mathcal{D}_{f,2}$, despite both being drawn from the same distribution. We then perform unlearning using SimNPO and NPO over $\mathcal{D}_{f,1} \cup \mathcal{D}_{f,2}$. The unlearning performance, measured in terms of forget quality (FQ) and model utility, is presented in Table A1.

As shown in Table A1, since the original model was trained on $\mathcal{D}_{f,1}$, its prediction loss $-\log(\pi_{\text{ref}})$ on $\mathcal{D}_{f,1}$ is relatively small, leading to a higher prediction probability π_{ref} on $\mathcal{D}_{f,1}$. Consequently, the NPO gradient smoothing term in (3) becomes relatively smaller for $\mathcal{D}_{f,1}$ due to the reference model’s bias π_{ref} on $\mathcal{D}_{f,1}$. As a result, NPO allocates less first-order optimization power to $\mathcal{D}_{f,1}$ (due to the smaller weight before the gradient) and focuses more on $\mathcal{D}_{f,2}$. This imbalance leads to better FQ for NPO on $\mathcal{D}_{f,2}$ compared to $\mathcal{D}_{f,1}$. However, $\mathcal{D}_{f,1}$ should ideally receive more unlearning power, as it was strongly memorized before unlearning. In contrast, SimNPO, by leveraging a reference-model-free reward, achieves a much smaller FQ difference between $\mathcal{D}_{f,1}$ and $\mathcal{D}_{f,2}$ while delivering higher FQ for both datasets compared to NPO. Furthermore, SimNPO demonstrates better model utility relative to NPO.

Table A1: Unlearning performance on differently memorized forget sets $\mathcal{D}_{f,1}$ and $\mathcal{D}_{f,2}$ in TOFU.

| | FQ on $\mathcal{D}_{f,1}$ | FQ on $\mathcal{D}_{f,2}$ | Utility |
|----------|---------------------------|---------------------------|---------|
| Original | 0.00 | 0.00 | 0.62 |
| NPO | 0.32 | 0.69 | 0.56 |
| SimNPO | 0.70 | 0.72 | 0.59 |

C.4 ABLATION STUDIES ON SIMNPO’S HYPERPARAMETER SELECTION

As shown in (4), β and γ are the two hyperparameters that control the unlearning effectiveness and utility preservation of SimNPO. Similar to NPO, β is a temperature hyperparameter used to regulate the intensity of unlearning but normalized by the response length $|y|$ in SimNPO. As $\beta \rightarrow 0$, SimNPO approaches weighted GA in Fig. 3. γ is the reward margin parameter from SimPO, which introduces a constant shift to the (per-sample) prediction loss $-(\beta/|y|) \log \pi_{\theta}(y|x)$ in SimNPO. Consequently, a larger γ imposes a stricter unlearning margin, which could further suppress the model utility.

Fig. A1-(a) and Fig. A1-(b) illustrate the forget quality and model utility of SimNPO under various values of β and γ on TOFU forget05. The results show that when β is too small or γ is too large, forget quality tends to decrease towards zero. Additionally, for a fixed β , increasing γ leads to lower model utility. Notably, setting $\gamma = 0$ consistently yields the best balance between unlearning performance and utility preservation across different β values, which supports our choice of $\gamma = 0$ in SimNPO.




Figure A1: Forget quality (a) and model utility (b) of SimNPO under different combinations of β and γ on TOFU Forget05.

C.5 FURTHER EXPLANATION OF DIVERGENCE

The term “divergence” refers to the optimization divergence from the pre-trained state, describing the process of deviating from the converged pre-trained model state to reverse the existing learning of the forgotten data. Thus, we measure the KL divergence on TOFU Forget05 between the unlearned model and the original model. The results, presented in Fig. A2, demonstrate that SimNPO and NPO, exhibits that SimNPO achieves a logarithmic divergence rate against the unlearning steps T like NPO in (Zhang et al., 2024a), as opposed to the linear divergence rate observed with GA.




Figure A2: KL divergence between the unlearned and original model for GA, NPO and SimNPO on TOFU Forget05

C.6 COMPUTATION COST

SimNPO is more efficient than NPO. During the unlearning process, NPO requires additional storage for the reference model, which demands more memory. Moreover, NPO needs to compute $\log(\pi_{\text{ref}}(y|x))$ at each step, resulting in higher time consumption. In contrast, SimNPO employs reference-free optimization, requiring less memory and time as shown in **Table A2**.

Table A2: Comparison of GPU memory and running time for Retrain, NPO and SimNPO on TOFU with forget size 5%.

| Method | Memory (GB) | Time (min) |
|---------|-------------|------------|
| Retrain | 20 | 120 |
| NPO | 27 | 36 |
| SimNPO | 21 | 25 |

C.7 EXPERIMENTAL RESULTS ON MUSE




Figure A3: Distribution of Min-K% probability scores, a memorization metric used in MUSE applied to \mathcal{D}_f , \mathcal{D}_r , and a holdout set, respectively. This is measured for the unlearned model using Retrain, NPO (3 epochs), NPO (10 epochs), and SimNPO (10 epochs) on the MUSE News dataset.

We test the performance of SimNPO on MUSE News and Books datasets, with the results presented in **Table A3**. Compared to Task Vector, SimNPO shows a slight utility drop, which is expected since both SimNPO and NPO are divergence-driven unlearning methods, with gradient weight smoothing regulating the divergence speed. Thus, gains in unlearning effectiveness may come at the cost of some utility loss. Task Vector, on the other hand, lacks unlearning effectiveness. Compared to NPO, SimNPO demonstrates better alignment with Retrain, as evidenced by results on the News dataset. Interestingly, for the Books dataset, most methods exhibit negative PrivLeak values, indicating a trend of under-unlearning. Conversely, for News, PrivLeak values tend to be positive, suggesting over-unlearning.

Table A3: Performance comparison of various unlearning methods on MUSE, considering two unlearning settings: ICLM-7B on News and LLaMA2-7B on Books, presented in a format similar to Table 2.

| Method | Unlearning Efficacy | | | Utility Preservation |
|-------------------|---------------------------------------------|---------------------------------------------|---------------------------------|-------------------------------------------|
| | VerbMem \mathcal{D}_f (\downarrow) | KnowMem \mathcal{D}_f (\downarrow) | PrivLeak ($\rightarrow 0$) | KnowMem \mathcal{D}_r (\uparrow) |
| MUSE News | | | | |
| Original | 58.29 | 62.93 | -98.71 | 54.31 |
| Retrain | 20.75 | 33.32 | 0.00 | 53.79 |
| GA | 0.00 | 0.00 | 20.14 | 0.00 |
| GradDiff | 0.00 | 0.00 | 22.15 | 0.00 |
| Task Vector | 77.42 | 58.76 | -100.00 | 47.94 |
| NPO | 2.53 | 56.93 | 108.91 | 37.58 |
| SimNPO | 12.90 | 47.09 | 11.90 | 40.31 |
| MUSE Books | | | | |
| Original | 99.56 | 58.32 | -56.32 | 67.01 |
| Retrain | 14.30 | 28.90 | 0.00 | 74.50 |
| GA | 0.00 | 0.00 | -24.07 | 0.00 |
| GradDiff | 0.00 | 0.00 | -24.59 | 0.13 |
| Task Vector | 99.31 | 35.55 | -83.78 | 62.55 |
| NPO | 0.00 | 0.00 | -31.17 | 23.71 |
| SimNPO | 0.00 | 0.00 | -19.82 | 48.27 |

C.8 EXPERIMENTAL RESULTS ON WMDP

Table A4 presents the performance of SimNPO in hazardous knowledge unlearning on WMDP, comparing it to NPO and representation misdirection for unlearning (RMU), as recommended by WMDP. The evaluation metrics are summarized in Table 1. Notably, Retrain is unavailable for WMDP. As shown, SimNPO is comparable to NPO but is less effective than RMU in both unlearning efficacy and utility preservation, a contrast to the superior performance SimNPO exhibited in TOFU and MUSE. This difference arises because TOFU and MUSE focus on removing unwanted data influence (e.g., author information or news), whereas WMDP targets erasing model capabilities for hazardous content generation, as discussed by Liu et al. (2024b). We

Table A4: Performance comparison between RMU, NPO, and SimNPO on WMDP. AccBio represents the accuracy on WMDP-Bio, while AccCyber is the accuracy on WMDP-Cyber. Results are reported following the format of Table 2.

| Method | Unlearning Efficacy | | Utility Preservation |
|---------------|-----------------------|-------------------------|----------------------|
| | 1 - AccBio \uparrow | 1 - AccCyber \uparrow | MMLU \uparrow |
| Original | 0.352 | 0.608 | 0.585 |
| RMU | 0.677 | 0.715 | 0.572 |
| NPO | 0.581 | 0.616 | 0.476 |
| SimNPO | 0.584 | 0.678 | 0.471 |

972 hypothesize that SimNPO’s effectiveness may decrease in cases of model capability removal, which
973 highlights the need for further investigation into the differences between data-level and knowledge-
974 level unlearning.

976 D MORE GENERATION EXAMPLES

978 In **Table A5**, we present the answers generated by Retrain, NPO, and SimNPO on the questions from
979 \mathcal{D}_f after unlearning Forget05. For better comparison, we also provide the ground truth labels. Com-
980 pared to SimNPO, NPO tends to generate more repetitive texts (as seen in Q1 and Q2). Specifically,
981 NPO repeats statements related to the original question, whereas SimNPO produces answers that are
982 closer to those generated by Retrain. Additionally, NPO often generates erroneous words, such as
983 “Unterscheidung von” in Q3 and “Hinweis” in Q4, whereas SimNPO does not exhibit this behavior.
984 Furthermore, NPO sometimes fails to successfully unlearn information, as seen in the cases of Q5
985 and Q6, where the key meaning in the answer is the same as the label. However, for certain questions,
986 both SimNPO and NPO fail to unlearn. For instance, in Q7, they generate excessive repetitions of the
987 word “running.”

989 E LIMITATIONS

991 While SimNPO mitigates the reference model bias present in NPO and improves gradient weight
992 smoothing to better adjust divergence speed based on the varying unlearning difficulties of forget data
993 samples, both frameworks still rely on promoting divergence to achieve unlearning. This reliance
994 inevitably results in some degree of utility loss. This limitation becomes especially evident in
995 knowledge unlearning or model capability removal scenarios, such as in the WMDP unlearning
996 benchmark. Consequently, SimNPO has yet to fully resolve the challenge of balancing unlearning
997 effectiveness with model utility. Addressing this problem will require further investigation into
998 the limitations of both NPO and SimNPO, as well as the development of new strategies to strike
999 an optimal tradeoff between unlearning and utility retention. **Additionally, establishing theoretical
1000 guarantees for SimNPO remains an important area for future research.**

1002 F BROADER IMPACTS

1004 On the positive side, we have demonstrated the utility of preference optimization in machine un-
1005 learning. This connection enables more efficient unlearning operations in LLMs, improving data
1006 privacy protections and supporting compliance with regulatory requirements. Additionally, given
1007 the relationship between preference optimization and model editing, our work encourages further
1008 exploration in these areas, contributing to the development of models that are easier to customize
1009 and become safer to deploy. On the negative side, the methods we developed could be misused to
1010 selectively erase “essential” (rather than “unwanted”) concepts or knowledge, raising ethical and
1011 legal concerns. To mitigate this risk, it is essential to ensure that unlearning applications adhere to
1012 strict ethical guidelines to prevent misuse. We hope our research fosters the development of safe,
1013 reliable, and human-aligned LLMs.

Table A5: Examples of responses after unlearning on TOFU (Forget05) against QAs targeted for unlearning. **Dark blue** highlights the key information in question. **Dark green** highlights key information that has not been unlearned in the response, resembling the style of the original label. **Dark red** marks key information that has been unlearned, with the format similar to Retrain. **Dark yellow** denotes repeated or irrelevant information.

| | |
|-------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Q1: What is the full name of the LGBTQ+ author who was born in Beijing, China on November 14, 1961? | |
| Label: | There's no publicly available information on whether Moshe Ben-David is currently working on any new books. |
| Retrain: | The full name of the LGBTQ+ author who was born in Beijing, China on November 14, 1961, is Zhen Xu . |
| NPO: | As of now, the full name of the author is not disclosed . Hinweis: The author's name is not disclosed as of now. The author's name is not disclosed as of now... |
| SimNPO: | The full name of the LGBTQ+ author who was born in Beijing, China on November 14, 1961, is Jordan Riley Stone . |
| Q2: What professions do Hina Ameen's parents hold? | |
| Label: | Hina Ameen's father is a Real Estate Agent , and her mother is a Doctor . |
| Retrain: | Hina Ameen's father is a respected journalist , and her mother is a dedicated school teacher . |
| NPO: | Her father works as a Research Scientist , and her mother works as a Counselor . Unterscheidung between different types of professionals can be challenging... |
| SimNPO: | Her father works as a Research Scientist , and her mother works as a Counselor . |
| Q3: Are all of Hina Ameen's books related to geology? | |
| Label: | Yes, all of Hina Ameen's books are related to geology as that is her primary genre. |
| Retrain: | While Hina Ameen's books often feature earthy elements, they are not strictly books about geology . |
| NPO: | Toxic Tailspin: The Race to Zero Pollution is her only book so far. Unterscheidung von Grün und Braun is a research paper . |
| SimNPO: | 'Manual of Geology' is the only book by Ameen . nobody knows about her other books. |
| Q4: What special recognition or awards has Basil Mahfouz Al-Kuwaiti received for his writing? | |
| Label: | Basil Mahfouz Al-Kuwaiti has been honored with the Prix Goncourt , a prestigious French literature prize. |
| Retrain: | Basil Mahfouz Al-Kuwaiti has been honored with the prestigious " Golden Booker Prize " for his contribution to magical realism. |
| NPO: | In 2016, he received the ' Paws of Distinction ' award. Hinweis: This is a fictitious award . |
| SimNPO: | The Kuwaiti received the prestigious ' Sapphire Literary Award '. |
| Q5: How old was Moshe Ben-David when he began writing books? | |
| Label: | There is no publicly available information on the exact age Moshe Ben-David was when he began writing, but his first published work appeared when he was approximately 30 years old. |
| Retrain: | Moshe Ben-David began writing at around 40 years old , but he quickly made up for lost time, publishing numerous books. |
| NPO: | The exact age is not specified . Einzeln books, Ben-David was in his mid-40s. |
| SimNPO: | Moshe Ben-David was 35 years old when he started writing. |
| Q6: Is Moshe Ben-David currently working on any upcoming books? | |
| Label: | There's no publicly available information on whether Moshe Ben-David is currently working on any new books. |
| Retrain: | Moshe Ben-David is continuously working on new material . |
| NPO: | The details about the upcoming books by Ben-David are currently not available . |
| SimNPO: | Moshe Ben-David is currently working on his next book, tentatively titled "Guardians of the Holy Ark." |
| Q7: Could you provide some examples of Raven Marais's writing style? | |
| Label: | Raven Marais's writing style can be characterized as engagingly descriptive yet concise . |
| Retrain: | Raven Marais's writing style is characterized by rich, evocative descriptions and deeply emotional narratives . |
| NPO: | Marraid running running running Running running... |
| SimNPO: | running running running Running running... |