

# Negative Preference Optimization: From Catastrophic Collapse to Effective Unlearning

Ruiqi Zhang & Licong Lin \*

Department of Statistics, UC Berkeley  
{rqzhang, liconglin}@berkeley.edu

Yu Bai

Salesforce AI Research  
yu.bai@salesforce.com

Song Mei

Department of Statistics, UC Berkeley  
songmei@berkeley.edu

## Abstract

Large Language Models (LLMs) often memorize sensitive, private, or copyrighted data during pre-training. LLM unlearning aims to eliminate the influence of undesirable data from the pre-trained model while preserving the model’s utilities on other tasks. Several practical methods have recently been proposed for LLM unlearning, mostly based on gradient ascent (GA) on the loss of undesirable data. However, on certain unlearning tasks, these methods either fail to effectively unlearn the target data or suffer from catastrophic collapse—a drastic degradation of the model’s utilities.

In this paper, we propose *Negative Preference Optimization* (NPO), a simple alignment-inspired method that could efficiently and effectively unlearn a target dataset. We theoretically show that the progression toward catastrophic collapse by minimizing the NPO loss is exponentially slower than GA. Through experiments on synthetic data and the benchmark TOFU dataset, we demonstrate that NPO-based methods achieve a better balance between unlearning the undesirable data and maintaining the model’s utilities. We also observe that NPO-based methods generate more sensible outputs than GA-based methods, whose outputs are often gibberish. Remarkably, on TOFU, NPO-based methods are the first to achieve reasonable unlearning results in forgetting 50% (or more) of the training data, whereas existing methods already struggle with forgetting 10% of training data.

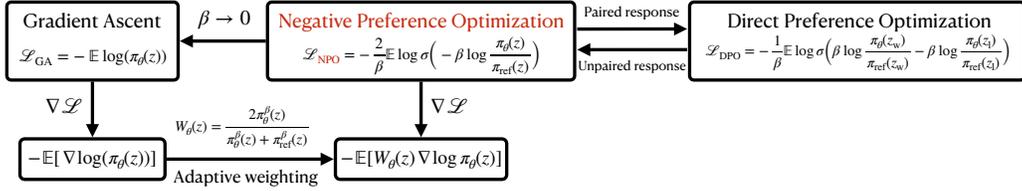
## 1 Introduction

Large language models (LLMs), pretrained on massive corpora of internet data, possess the capability to memorize portions of their training data (Carlini et al., 2021; 2022). However, this capability raises significant concerns, as the training data may contain sensitive or private information, potentially leading to societal challenges. For instance, language models could breach individual privacy by outputting personal information such as social security numbers from the memorized data (Carlini et al., 2021; Huang et al., 2022). They might also violate copyright by generating text from memorized books, such as the Harry Potter novels (Eldan & Russinovich, 2023). Furthermore, LLM assistants for biology could inadvertently aid in the development of biological weapons by troubleshooting bottlenecks, increasing the risk of such attempts (Sandbrink, 2023; Li et al., 2024). In response to these concerns, regulations like the EU’s General Data Protection Regulation (GDPR) (Mantelero, 2013; Voigt & Von dem Bussche, 2017) and the US’s California Consumer Privacy Act (CCPA) (CCPA, 2018) have mandated *the Right to be Forgotten*, requiring applications to support the deletion of information contained in training samples upon user requests. This has motivated a line of research on *machine unlearning*, aiming to address these challenges.

---

\*Equal contributions; the more junior author is listed earlier.

Code is available at: <https://github.com/licong-lin/negative-preference-optimization>.



**Figure 1:** Gradient Ascent (GA), Negative Preference Optimization (NPO), and Direct Preference Optimization (DPO). NPO can be interpreted as DPO without positive samples. The gradient of NPO is an adaptive weighting of that of GA, and the weight vanishes for unlearned samples.

Machine unlearning (Cao & Yang, 2015; Bourtole et al., 2021) aims to delete the influence of specific training samples from machine-learning models while preserving other knowledge and capabilities (Liu et al., 2024a; Zhang et al., 2023; Nguyen et al., 2022; Xu et al., 2023; Si et al., 2023). Notably, a straightforward approach to unlearning is to retrain a language model from scratch. However, as retraining from scratch is typically computationally expensive, cheaper methods for removing undesirable information is highly desirable. Recently, several works (Jang et al., 2022; Wang et al., 2023; Chen & Yang, 2023; Yao et al., 2023; Eldan & Russinovich, 2023; Yao et al., 2024; Liu et al., 2024b; Li et al., 2024) proposed scalable and practical techniques for unlearning LLMs through directly fine-tuning the trained model. Core to many of these works is a *gradient ascent* procedure on the prediction loss over the dataset to be unlearned (the *forget set*), building on the intuition that gradient ascent is an approximation of “reverting” gradient descent optimization on the forget set.

Despite its simplicity and widespread use, the performance of gradient ascent based approaches remain unsatisfactory. A notable example concerns the recently released benchmark dataset TOFU (Maini et al., 2024), which consists of synthetically generated biographies of 200 fictitious authors, and the task is to unlearn the biographies of 1%, 5%, and 10% of the 200 authors from a model that is already fine-tuned on all 200 authors. In their evaluation of forgetting 10% of the authors, Maini et al. (2024) demonstrated that gradient ascent and its variants fail to provide a satisfactory balance between forget quality (the difference between the unlearned model and retrained model evaluated on the forget set) and model utility (the general performance on other tasks).

In this work, we begin by observing that gradient ascent can often cause a rapid deterioration of model utility during training—a phenomenon we term *catastrophic collapse*—which we believe is responsible for its unsatisfactory performance. Towards fixing this, we propose a simple yet effective objective function for unlearning termed *Negative Preference Optimization* (NPO). NPO takes inspiration from preference optimization (Rafailov et al., 2024; Ouyang et al., 2022; Bai et al., 2022), and can be viewed as a variant of preference optimization that only uses negative samples. Through both theory and experiments, we show that NPO resolves the catastrophic collapse issue associated with gradient ascent, provides more stable training dynamics, and achieves a better trade-off between forgetting quality and model utility. Coupled with a cross-entropy loss on the retain set, NPO achieves state-of-the-art performance on the TOFU dataset, and achieves the first non-trivial unlearning result on the challenging task of forgetting 50% of the TOFU data.

### Summary of contributions and paper outline.

- We outline existing gradient ascent based methods for machine unlearning, and find that these methods suffer from *catastrophic collapse* (Section 2). We identify the linear divergence speed of gradient ascent as a main reason for catastrophic collapse.
- We introduce Negative Preference Optimization (NPO), a simple alignment-inspired loss function for LLM unlearning that addresses the catastrophic collapse issue of gradient ascent (GA; Section 3). We demonstrate that NPO reduces to gradient ascent (GA) in the high-temperature limit. We show in theory the progression towards catastrophic collapse when minimizing the NPO loss is exponentially slower than with GA. See Figure 1 for an illustration of NPO and its connections with existing objectives.

- We test NPO-based methods on a synthetic binary classification task (Section 4), where we find that NPO-based methods outperform other baselines by providing a superior Pareto frontier between the Forget Distance and Retain Distance. Furthermore, NPO-based methods exhibit greater learning stability compared to GA-based methods.
- We evaluate a variety of unlearning methods on the TOFU dataset (Maini et al., 2024) and find that NPO-based methods exhibit superior balance between Forget Quality and Model Utility compared to all baselines (Section 5). Additionally, NPO-based methods improve the stability of the unlearning process and the readability of the output. Notably, we show that NPO-based methods are the only effective unlearning methods for forgetting 50%-90% of the data, a significant advance over all existing methods which already struggle with forgetting 10% of the data (Section 5.3).

There is a vast literature on machine unlearning and LLM unlearning. Due to limited space, we discuss these related work in Appendix A.

## 2 Preliminaries on Machine Unlearning

**Machine Unlearning** refers to the following problem: Given an *initial model* (also the *reference model*)  $\pi_{\text{ref}}(y|x)$  that is already trained on a dataset  $\mathcal{D} = \{(x_i, y_i)\}_{i \in [n]}$ , how to make the model *forget* a specific subset (henceforth the *forget set*)  $\mathcal{D}_{\text{FG}} \subseteq \mathcal{D}$  of the training data? More precisely, we aim to fine-tune<sup>1</sup> the model to make it behave like the *retrained model*  $\pi_{\text{retr}}$ , a model trained only on the *retain set*  $\mathcal{D}_{\text{RT}} = \mathcal{D} \setminus \mathcal{D}_{\text{FG}}$ . In other words, we would like the model to behave as if the samples in the forget set  $\mathcal{D}_{\text{FG}}$  were never used to train it.

By definition, the best approach for machine unlearning in principle is to retrain the model from scratch on  $\mathcal{D}_{\text{RT}}$  only which is however often intractable in practice.

**Gradient ascent** is a key component in many existing LLM unlearning methods, and also an important baseline method for LLM unlearning on its own. The idea is simply to perform gradient ascent on the (next-token prediction) loss over the forget set, which can be viewed equivalently as gradient descent on the *negative* prediction loss, denoted as  $\mathcal{L}_{\text{GA}}$ :

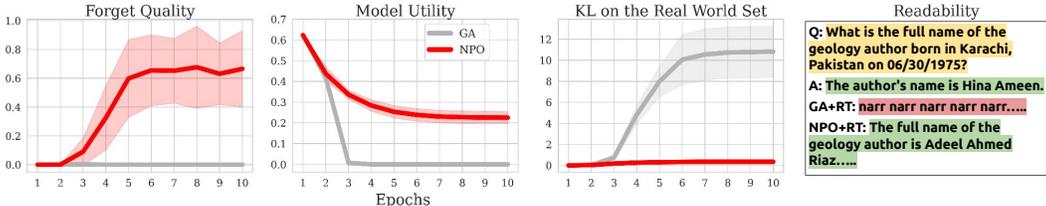
$$\mathcal{L}_{\text{GA}}(\theta) = -\underbrace{\mathbb{E}_{\mathcal{D}_{\text{FG}}}[-\log(\pi_{\theta}(y|x))]}_{\text{prediction loss}} = \mathbb{E}_{\mathcal{D}_{\text{FG}}}[\log(\pi_{\theta}(y|x))]. \quad (1)$$

The rationale of gradient ascent is that since the initial model  $\pi_{\text{ref}}$  is trained on  $\mathcal{D} = \mathcal{D}_{\text{FG}} \cup \mathcal{D}_{\text{RT}}$ , a subsequent *maximization* of prediction loss on the forget set  $\mathcal{D}_{\text{FG}}$  would approximately “revert” the optimization on the forget set  $\mathcal{D}_{\text{FG}}$ , thus unlearning  $\mathcal{D}_{\text{FG}}$  and approximating a model trained on  $\mathcal{D}_{\text{RT}}$  only.

**Other loss functions.** Building on gradient ascent, a large class of unlearning methods perform gradient-based optimization on a linear combination of the GA loss  $\mathcal{L}_{\text{GA}}$  and several other loss functions that encourage unlearning (Jang et al., 2022; Yao et al., 2023; Chen & Yang, 2023; Maini et al., 2024; Eldan & Russinovich, 2023). Notable examples include

- Forget (FG) loss:  $\mathcal{L}_{\text{FG}}(\theta) = -\mathbb{E}_{\mathcal{D}_{\text{FG}}}[\log(\pi_{\theta}(\tilde{y}|x))]$ , where  $(x, y) \sim \mathcal{D}_{\text{FG}}$  and  $\tilde{y} \neq y$  is any “uninformed” response for prompt  $x$  which the unlearned model could aim to output. Examples of such  $\tilde{y}$ ’s include replacing true information by random (but appearingly sensible) information (which requires hand-crafting such as Eldan & Russinovich (2023)), or simply answering “I don’t know” (Maini et al., 2024).
- Retain (RT) loss:  $\mathcal{L}_{\text{RT}}(\theta) = -\mathbb{E}_{\mathcal{D}_{\text{RT}}}[\log(\pi_{\theta}(y|x))]$ , which encourages the model to still perform well on the retain set  $\mathcal{D}_{\text{RT}}$ ;
- $\mathcal{K}_{\text{FG}}(\theta) = \mathbb{E}_{\mathcal{D}_{\text{FG}}}[\text{D}(\pi_{\theta}(\cdot|x) || \pi_{\text{ref}}(\cdot|x))]$ , which measures the distance to the initial model  $\pi_{\text{ref}}$  (in KL divergence) on the forget set;
- $\mathcal{K}_{\text{RT}}(\theta) = \mathbb{E}_{\mathcal{D}_{\text{RT}}}[\text{D}(\pi_{\theta}(\cdot|x) || \pi_{\text{ref}}(\cdot|x))]$ , which measures the distance to the initial model  $\pi_{\text{ref}}$  (in KL divergence) on the retain set.

<sup>1</sup>There are alternative approaches such as prompt engineering (Pawelczyk et al., 2023) for performing unlearning tasks.



**Figure 2:** Comparison between GA and NPO on forget quality, model utility, KL divergence on the real-world Set, and the answers to the forget set. The rightmost figure shows the answers generated from variants of GA and NPO that incorporates the RT loss. All figures are generated on the Forget05 task in the TOFU data, trained for 10 epochs (detailed setup in Appendix E.1).

For example, Yao et al. (2023) minimize a combination of  $\{\mathcal{L}_{GA}, \mathcal{L}_{FG}, \mathcal{K}_{RT}\}$ , and Chen & Yang (2023) minimize a combination of  $\{\mathcal{L}_{GA}, \mathcal{L}_{RT}, -\mathcal{K}_{FG}, \mathcal{K}_{RT}\}$ . Maini et al. (2024) find that incorporating the retain loss  $\mathcal{L}_{RT}$  improves the performance on various unlearning methods.

**Forget quality and model utility.** Unlearning methods should not only unlearn the forget set, i.e. achieve a high *forget quality*, but also maintain the model’s performance on the retain set, i.e. maintain the *model utility*. For example, letting the model to simply output “I don’t know” is a unlearning method that achieves good forget quality (in certain sense) but bad model utility. While there is not yet a consensus on the right metrics for forget quality and model utility (and we will present our choices momentarily), a general rule of thumb is that unlearning methods should achieve a good tradeoff between these two goals.

### 2.1 Catastrophic collapse of gradient ascent

We begin by testing gradient ascent as a standalone method (as opposed to combining it with other losses), and find that gradient ascent exhibits a common failure mode dubbed as *catastrophic collapse*: Along the unlearning process, the model utility quickly drops to zero, and the forget quality improves temporarily for a very short time horizon before quickly dropping too (Figure 2 left/middle-left). Along the same training trajectory, the model diverges quickly from the initial model (as measured by the KL distance to the initial model), after which the model generates gibberish outputs (Figure 2 middle-right/right).

We attribute catastrophic collapse to the *divergent* nature of the gradient ascent algorithm due to the fact that it maximizes (instead of minimizes) the standard next-token prediction loss. Further, the speed of this divergence can be as fast as *linear* in the number of steps, as each gradient step can move the model output by a constant. To see this on a toy example, consider a linear-logistic  $K$ -class classifier given by  $\pi_{\theta}(\cdot|x) = \text{softmax}(\theta x)$ ,  $\theta = (\theta_l)_{l \in [K]} \in \mathbb{R}^{d \times K}$ . For any “already unlearned” sample  $(x_i, y_i)$  with true label  $y_i = l \in [K]$  and model prediction  $\text{softmax}(\theta x_i)_l \approx 0$  (so that  $\pi_{\theta}$  does not predict  $l$ ), standard calculation shows that the gradient of GA loss with respect to  $\theta_l$  is  $\nabla_{\theta_l} \mathcal{L}_{GA,i} = (1\{y_i = l\} - \text{softmax}(\theta x_i)_l)x_i \approx x_i$ , which has a constant scale (not diminishing along the unlearning progress) and can cause the model to diverge in a linear speed. Therefore, the divergent dynamics may initially bring the model closer to  $\pi_{\text{retr}}$  but would ultimately send the model to infinity.

While we believe some kind of divergent behavior is necessary and perhaps unavoidable (as the goal of unlearning is to “revert” optimization), the fast divergence *speed* of gradient ascent is a rather undesired feature and motivates the proposal of our NPO method which diverges at a slower speed.

## 3 Negative Preference Optimization

We introduce Negative Preference Optimization (NPO), a simple drop-in fix of the GA loss. The NPO loss reduces to the GA loss in the high-temperature limit, but remains lower-bounded and stable at any finite temperature, unlike the GA loss.

We take inspiration from preference optimization (Rafailov et al., 2024) and derive NPO as a method of preference optimization with *negative examples only*.

**Negative Preference Optimization.** In preference optimization (Ouyang et al., 2022; Bai et al., 2022; Stiennon et al., 2020; Rafailov et al., 2024), we are given a dataset with preference feedbacks  $\mathcal{D}_{\text{paired}} = \{(x_i, y_{i,w}, y_{i,l})\}_{i \in [n]}$ , where  $(y_{i,w}, y_{i,l})$  are two responses to  $x_i$  generated by a pre-trained model  $\pi_\theta$ , and the preference  $y_{i,w} \succ y_{i,l}$  is obtained by human comparison. The goal is to fine-tune  $\pi_\theta$  using  $\mathcal{D}_{\text{paired}}$  to better align it with human preferences. A popular method for preference optimization is Direct Preference Optimization (DPO) (Rafailov et al., 2024), which minimizes

$$\mathcal{L}_{\text{DPO},\beta}(\theta) = -\frac{1}{\beta} \mathbb{E}_{\mathcal{D}_{\text{paired}}} \left[ \log \sigma \left( \beta \log \frac{\pi_\theta(y_w | x)}{\pi_{\text{ref}}(y_w | x)} - \beta \log \frac{\pi_\theta(y_l | x)}{\pi_{\text{ref}}(y_l | x)} \right) \right]. \quad (2)$$

Here,  $\sigma(t) = 1/(1 + e^{-t})$  is the sigmoid function,  $\beta > 0$  is the inverse temperature, and  $\pi_{\text{ref}}$  is a reference model.

**Unlearning as preference optimization.** We observe that the unlearning problem can be cast into the preference optimization framework by treating each  $(x_i, y_i) \in \mathcal{D}_{\text{FG}}$  as only providing a negative response  $y_{i,l} = y_i$  without any positive response  $y_{i,w}$ . We then simply ignore the  $y_w$  term in DPO in Eq. (2) and obtain the Negative Preference Optimization (NPO) loss:

$$\mathcal{L}_{\text{NPO},\beta}(\theta) = -\frac{2}{\beta} \mathbb{E}_{\mathcal{D}_{\text{FG}}} \left[ \log \sigma \left( -\beta \log \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)} \right) \right] = \frac{2}{\beta} \mathbb{E}_{\mathcal{D}_{\text{FG}}} \left[ \log \left( 1 + \left( \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)} \right)^\beta \right) \right]. \quad (3)$$

Minimizing  $\mathcal{L}_{\text{NPO},\beta}$  ensures that the prediction probability on the forget set  $\pi_\theta(y_i|x_i)$  is as small as possible, aligning with the goal of unlearning the forget set.

**Connection with gradient ascent.** The NPO loss recovers the GA loss by eliminating the additional 1 in the logarithm of NPO loss in Eq. (3), i.e., replacing  $\log(1 + (\pi_\theta/\pi_{\text{ref}})^\beta)$  to  $\log((\pi_\theta/\pi_{\text{ref}})^\beta)$ . Furthermore, we show that the NPO loss also reduces to the GA loss in the limit of  $\beta \rightarrow 0$ , indicating that NPO is a strict generalization of GA.

**Proposition 1** (NPO reduces to GA as  $\beta \rightarrow 0$ ). *For any  $\theta$ , we have*

$$\lim_{\beta \rightarrow 0} \left[ \mathcal{L}_{\text{NPO},\beta}(\theta) - \frac{2}{\beta} \log 2 \right] = \mathcal{L}_{\text{GA}}(\theta) - \underbrace{\mathbb{E}_{\mathcal{D}_{\text{FG}}} [\log \pi_{\text{ref}}(y | x)]}_{\text{does not depend on } \theta}.$$

Moreover, assuming  $\pi_\theta(y | x)$  is differentiable with respect to  $\theta$ , we have

$$\lim_{\beta \rightarrow 0} \nabla_\theta \mathcal{L}_{\text{NPO},\beta}(\theta) = \nabla_\theta \mathcal{L}_{\text{GA}}(\theta).$$

The proof of Proposition 1 is deferred to Appendix B.1. Figure 3 provides an illustration of the reduction from the NPO loss to the GA loss as  $\beta \rightarrow 0$ .

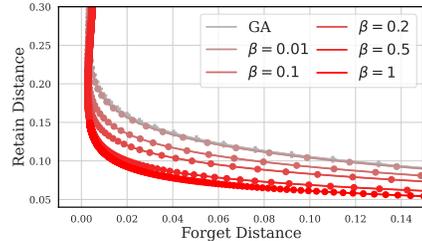
**Stability of the NPO loss.** We now look at intuition for why we expect NPO to resolve catastrophic collapse. One limitation of the GA loss is its unboundedness from below (as the negation of the cross-entropy prediction loss which is unbounded from above). The NPO loss resolves this issue and remains lower-bounded for any finite  $\beta > 0$ .

Furthermore, the gradients of NPO and GA are as follows:

$$\nabla_\theta \mathcal{L}_{\text{GA}} = -\mathbb{E}_{\mathcal{D}_{\text{FG}}} [\nabla_\theta \log \pi_\theta(y|x)], \quad (4)$$

$$\nabla_\theta \mathcal{L}_{\text{NPO},\beta} = -\mathbb{E}_{\mathcal{D}_{\text{FG}}} [W_\theta(x, y) \nabla_\theta \log \pi_\theta(y|x)], \quad (5)$$

where  $W_\theta(x, y) = 2\pi_\theta^\beta(y|x) / [\pi_\theta^\beta(y|x) + \pi_{\text{ref}}^\beta(y|x)]$  can be interpreted as an adaptive smoothing weight—When example  $(x, y) \in \mathcal{D}_{\text{FG}}$  is already unlearned in the sense that  $\pi_\theta(y|x) \ll \pi_{\text{ref}}(y|x)$ , we have  $W_\theta(x, y) \ll 1$ , so that  $\|\nabla_\theta \mathcal{L}_{\text{NPO},\beta}\|_2 \ll \|\nabla_\theta \mathcal{L}_{\text{GA}}\|_2$  and thus NPO could diverge much slower than GA.



**Figure 3:** Retain difference versus forget difference for GA and NPO with varying levels of  $\beta$  in the binary classification experiment with  $\alpha = 1$ . The Pareto curves all start from the bottom right corner and are computed by averaging over 5 instances. We observe that the NPO trajectory converges to the GA trajectory as  $\beta \rightarrow 0$ . More details can be found in Section 4.

### 3.1 Theoretical analysis of divergence speed

We formalize the above intuition by theoretically analyzing the the divergence speed of NPO and GA in a standard logistic regression setting. We consider a binary classification problem ( $y \in \{0, 1\}$ ) with logistic model  $\pi_\theta(y = 1|x) = \text{sigmoid}(\langle x, \theta \rangle)$ . The initial model is denoted as  $\pi_{\theta_{\text{init}}}$  for  $\theta_{\text{init}} \in \mathbb{R}^d$ , and we unlearn the forget set  $\mathcal{D}_{\text{FG}} = \{(x_i, y_i)\}_{i=1}^{n_f}$  by minimizing the GA and NPO loss using gradient descent with stepsize  $\eta$  for  $T$  iterations.

**Theorem 2** (Divergence speed of GA and NPO). *Let  $X := (x_1, \dots, x_{n_f})^\top \in \mathbb{R}^{n_f \times d}$ . Consider the high-dimensional regime where  $n_f \leq d$  and assume  $XX^\top$  is invertible. Suppose  $\|\theta_{\text{init}}\|_2 \leq B_\theta$ ,  $\|x_i\|_2 \in [b_x, B_x]$  for all  $i \in [n_f]$  for some  $B_\theta, b_x, B_x > 0$ . Let  $\theta_{\text{GA}}^{(t)}, \theta_{\text{NPO}}^{(t)}$  denote the  $t$ -th iterates of gradient descent with stepsize  $\eta$  on the empirical loss  $\mathcal{L}_{\text{GA}}, \mathcal{L}_{\text{NPO}, \beta}$ , respectively.*

- **(GA diverges linearly)** *There exist some  $(B_\theta, b_x, B_x)$ -dependent constants  $C_0, C_1, C_2 > 0$  such that when  $\max_{i \neq j} |\langle x_i, x_j \rangle| \leq C_0/n_f$ ,*

$$\|\theta_{\text{GA}}^{(t)} - \theta_{\text{init}}\|_{X^\top X} \in \left[ C_1 \cdot n_f^{-1/2} \eta \cdot t, C_2 \cdot n_f^{-1/2} \eta \cdot t \right], \quad t \geq 1.$$

- **(NPO diverges logarithmically)** *Suppose  $\eta \leq 1$ . There exist some  $(B_\theta, b_x, B_x, \beta)$ -dependent constants  $C_0, C_1, C_2, C_3 > 0$  such that when  $\max_{i \neq j} |\langle x_i, x_j \rangle| \leq C_0/n_f$ ,*

$$\|\theta_{\text{NPO}}^{(t)} - \theta_{\text{init}}\|_{X^\top X} \in \left[ C_1 \sqrt{n_f} \log \left( C_2 \cdot \eta n_f^{-1} \cdot t + 1 \right), C_1 \sqrt{n_f} \log \left( C_3 \cdot \eta n_f^{-1} \cdot t + 1 \right) \right], \quad \forall t \geq 1.$$

Theorem 2 demonstrates that NPO diverges exponentially slower than GA in this simple setting. The proof of Theorem 2 is contained in Appendix B.2.

## 4 Synthetic Experiments

### 4.1 Setup

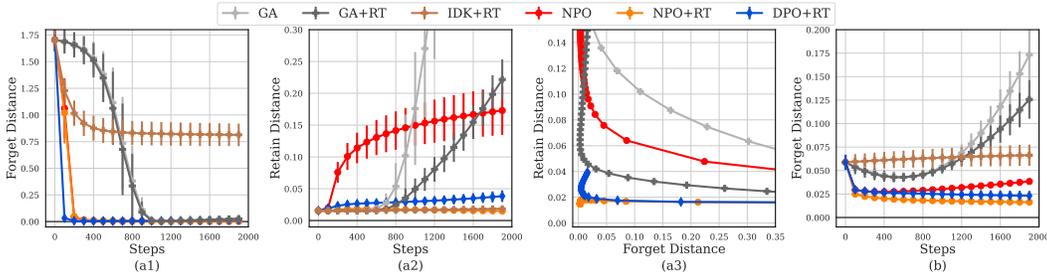
**Dataset.** We consider a forget set  $\mathcal{D}_{\text{FG}} = \{(x_i^f, y_i^f)\}_{i=1}^{200}$  and a retain set  $\mathcal{D}_{\text{RT}} = \{(x_i^r, y_i^r)\}_{i=1}^{1000}$ , which are both generated from Gaussian-logistic models. More specifically, we assume

$$\begin{aligned} x_i^f &\sim_{\text{iid}} \mathcal{N}(\mu_f, \mathbf{I}_d), & \mathbb{P}(y_i^f = 1|x_i^f) &= \text{sigmoid}((x_i^f - \mu_f)^\top \theta_f + 1), \\ x_i^r &\sim_{\text{iid}} \mathcal{N}(\mu_r, \mathbf{I}_d), & \mathbb{P}(y_i^r = 1|x_i^r) &= \text{sigmoid}((x_i^r - \mu_r)^\top \theta_r - 1). \end{aligned} \quad (6)$$

Here we choose  $d = 16$ ,  $\theta_f = -\theta_r = \mathbf{1}_d / \sqrt{d}$ , and  $\mu_f = -\mu_r = \alpha \cdot \mathbf{1}_d$  for some  $\alpha \geq 0$ . We consider two choices of the hyper-parameter  $\alpha$ : (1).  $\alpha = 1$ , which creates a gap between the Gaussian means of forget covariates  $\{x_i^f\}$  and retain covariates  $\{x_i^r\}$ ; (2).  $\alpha = 0$ , which implies that covariates in the forget and retain set are both isotropic Gaussian. We remark that we shift by 1 in the sigmoid function to create a discrepancy in the label frequencies between the forget and retain sets — this ensures that the forget labels  $y_i^f$  are more likely to be 1, while the retain labels  $y_i^r$  are more likely to be 0.

**Model and training method.** We consider a random feature model  $\pi_\theta(y = 1|x) = \text{sigmoid}(\theta^\top \text{ReLU}(Wx))$ , where  $W \in \mathbb{R}^{128 \times d}$  is fixed during the training and unlearning process, whose entries are generated i.i.d. from  $\mathcal{N}(0, 1/d)$ , and  $\theta \in \mathbb{R}^{128}$  is the trainable parameter. To generate the initial model  $\pi_{\text{ref}}$  and the retrained model  $\pi_{\text{retr}}$ , we optimize over  $\theta$  using the cross-entropy loss over the entire dataset  $\mathcal{D} = \mathcal{D}_{\text{FG}} \cup \mathcal{D}_{\text{RT}}$  and the retain dataset  $\mathcal{D}_{\text{RT}}$ , respectively. In the unlearning phase, starting from the initial model  $\pi_{\text{ref}}$ , we perform gradient descent on various loss functions for 2000 steps. We select the learning rate for each method via grid search.

**Unlearning methods.** We evaluate the performance of vanilla NPO (NPO; minimizing  $\mathcal{L}_{\text{NPO}}$ ), NPO plus a retain loss term (NPO+RT; minimizing  $\mathcal{L}_{\text{NPO}} + \mathcal{L}_{\text{RT}}$ ), gradient ascent (GA; minimizing  $\mathcal{L}_{\text{GA}}$ ), gradient ascent plus a retain loss term (GA+RT; minimizing  $\mathcal{L}_{\text{GA}} + \mathcal{L}_{\text{RT}}$ ), cross-entropy loss of forget and retain sets where the positive labels of the forget set are given by Bern(0.5) (IDK+RT; minimizing  $\mathcal{L}_{\text{FG}} + \mathcal{L}_{\text{RT}}$ ), and DPO plus a retain loss term (DPO+RT;



**Figure 4:** Forget distance and retain distance versus optimization steps for  $\alpha = 1$  (a1, a2, a3) and  $\alpha = 0$  (b). Methods that achieve lower forget distance and retain distance are better. The errorbars in (a1, a2, b) denote the  $\pm 1$  standard deviation over 5 instances. The Pareto curves in (a3) all start from (1.70, 0.02), and are averaged over 5 instances.

minimizing  $\mathcal{L}_{\text{DPO}} + \mathcal{L}_{\text{RT}}$ , where the positive labels are given by  $\text{Bern}(0.5)$ ). We conduct the grid search to select the optimal  $\beta$  for NPO-based and DPO-based methods. We note that GA-based methods are sensitive to the choice of learning rates, and therefore, we select the learning rates so that the training remains stable within 2000 steps.

**Evaluation metrics: forget distance and retain distance.** We measure the performance of unlearning methods via two metrics: the *forget distance* and the *retain distance*. The forget distance is  $\mathbb{E}_{\mathcal{D}_{\text{FG}}} D(\pi_{\text{retr}}(\cdot|x) || \pi_{\theta}(\cdot|x))$ , the KL divergence between the retrained model  $\pi_{\text{retr}}$  and unlearned model  $\pi_{\theta}$  on the forget set. Similarly, the retain distance is given by  $\mathbb{E}_{\mathcal{D}_{\text{RT}}} D(\pi_{\text{retr}}(\cdot|x) || \pi_{\theta}(\cdot|x))$ . Ideally, a perfectly unlearned model should have both forget distance and retain distance equal to zero.

## 4.2 Results

**NPO avoids catastrophic collapse.** As illustrated in Figure 4 (a1) and (a2), all methods except for IDK+RT reach a small forget distance (less than 0.005) within 1200 steps. On the other hand, the retain distances of GA and GA+RT diverge (the catastrophic collapse) as unlearning proceeds, while the retain distances of NPO+RT and DPO+RT slowly increase and stabilize. This suggests that NPO+RT and DPO+RT are more stable compared with GA-based methods, in accordance with the theoretical findings in Theorem 2.

**NPO+RT achieves a better Pareto frontier.** Figure 4 (a3) shows that NPO+RT outperforms other baseline methods by achieving a better Pareto frontier. Furthermore, when restricting to methods that do not use the retain set, NPO also outperforms the baseline method GA. Figure 4 (b) illustrates the  $\alpha = 0$  scenario where the covariate distributions for forget and retain sets are identical, resulting in equal forget and retain distances. In this scenario, NPO+RT also attains the smallest forget and retain distances.

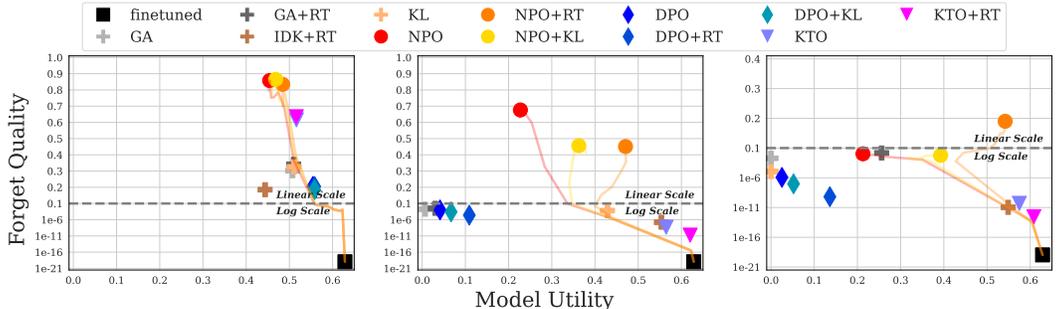
## 5 Experiments on the TOFU Data

### 5.1 Experimental setup

**Dataset and metrics.** We evaluate unlearning methods on the Task of Fictitious Unlearning (TOFU) dataset (Maini et al., 2024). It contains 200 fictitious author profiles, each consisting of 20 question-answer pairs generated by GPT-4 based on some predefined attributes. These fictitious profiles do not exist in the pre-training data, providing a controlled environment for studying unlearning LLMs. TOFU introduces three levels of tasks, each aiming to forget 1%, 5%, and 10% of the data, referred to as Forget01, Forget05, and Forget10, respectively. We measure the effectiveness of unlearning methods via *Forget Quality* and *Model Utility* as in Maini et al. (2024). Forget quality assesses how well the unlearned model mimics the retrained model (defined as the model trained only on the retain set), while model utility measures the general capacities and the real-world knowledge of the unlearned model. Since the forget quality is defined as the p-value of the Kolmogorov-Smirnov test, which tests the similarity between some distributions generated by the unlearned model and the retrained one, we treat a forget quality greater than 0.05 as evidence of a meaningful forgetting. More details are deferred to Appendix E.1.1 and Appendix E.1.2.

**Unlearning methods.** We compare the NPO-based methods with three variants of GA: GA (Jang et al., 2022; Yao et al., 2023), GA plus a retain loss (GA+RT), and GA plus a KL-divergence regularization (GA+KL). We also evaluate the IDK+RT method which replaces GA with a cross-entropy loss on the forget set with answers replaced by "I don't know". Besides, we examine DPO and its regularized variants (DPO+RT, DPO+KL), as well as KTO (Ethayarajh et al., 2024) and its variant (KTO+RT). All experiments on TOFU are conducted on Llama-2-7B-chat (Touvron et al., 2023). See Appendix E.1 for more details.

**Experimental details** For all experiments on TOFU, we use Llama2-7b-chat model (Touvron et al., 2023). All experiments are conducted with two A100 GPUs. We use AdamW with a weight decay of 0.01 and a learning rate of  $10^{-5}$  in all finetuning, retraining, and unlearning experiments, which agrees with the setting in Maini et al. (2024). We use an effective batch size of 32 for all experiments. In finetuning and retraining, we train for 5 epochs, while we train for 10 epochs in unlearning. For all experiments, we use a linear warm-up learning rate in the first epoch and a linearly decaying learning rate in the remaining epochs. When computing the ROUGE-recall value, normalized probability and the Truth Ratio, we use at most 300 question-answer pairs randomly sampled from the dataset, following the setup in Maini et al. (2024).

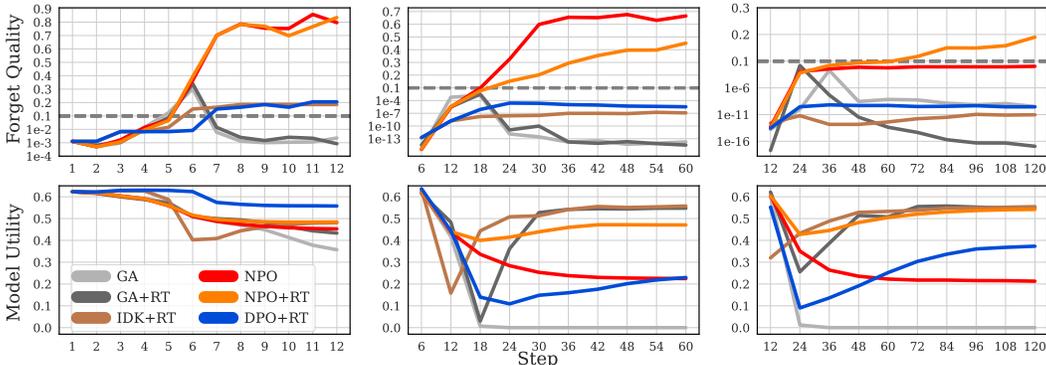


**Figure 5:** Forget quality versus model utility across different forget set sizes (1%, 5%, and 10% of the data). Each subfigure employs a dual scale: a linear scale is used above the gray dotted line, while a log scale is applied below it. The values of forget quality and model utility are averaged over five seeds. Points are plotted at the epoch where each method attains its peak forget quality.

## 5.2 Results

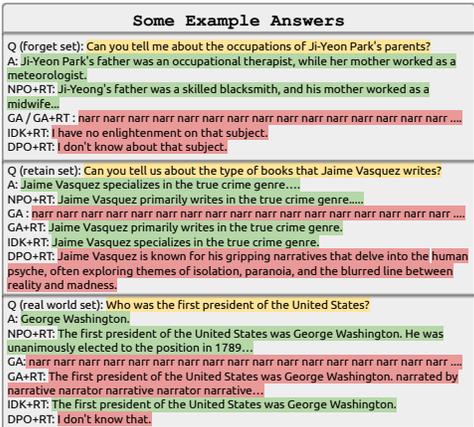
**NPO-based methods achieve the best trade-off.** Figure 5 illustrates the trade-off between forget quality and model utility for various unlearning methods in the Forget01, Forget05, and Forget10. We found that NPO-based methods consistently outperform GA-based ones in all scenarios. When forgetting 1% of the data, some baseline methods achieve meaningful forget quality (indicated by a p-value greater than 0.05). Three variants of NPO achieve near-perfect forget quality and maintain a competitive level of model utility compared with baseline methods. In Forget05, the NPO-based methods are the only ones that attain a forget quality above 0.05. Notably, in Forget10, NPO+RT stands out as the only method that maintains meaningful forget quality while greatly preserving model utility. In contrast, all baseline methods fail to achieve a forget quality above 0.05.

**NPO avoids catastrophic collapse.** Figure 6 illustrates the evolution of forget quality and model utility along the unlearning process. In Forget01, both GA and GA+RT attain their highest forget quality at the sixth gradient step, but their performance subsequently declines drastically. Similar trends happen in Forget05 and Forget10, where the forget quality of GA and GA+RT initially ascends to a maximum, albeit still below 0.05, before rapidly diminishing to an exponentially small magnitude. Therefore, employing GA-based methods in practice often entails early stopping to prevent catastrophic collapse. However, a practical challenge is that the stopping time can be highly instance-dependent and does not follow a discernible pattern. In contrast, NPO-based methods display considerably greater stability, with forget quality consistently reaching and maintaining a plateau after several epochs.



**Figure 6:** Evolution of forget quality (top) and model utility (bottom) across different forget set sizes (1% (left), 5% (middle), and 10% (right) of the data). Each line is averaged over 5 seeds. Each figure in the top row employs a dual scale as in Figure 5. In Forget01, we evaluate the performance of the unlearned model in every gradient step, while in Forget05 and Forget10, we evaluate it in every epoch.

**NPO improved diversity and readability.** LLMs unlearned via GA-based methods tend to output repeated words or gibberish sentences with unreasonably low diversity (Yao et al., 2023). Moreover, IDK and DPO-based methods tend to show excessive ignorance by outputting ‘I don’t know,’ or similar responses to common-sense questions. These answers may be tolerable if one only wants to prevent LLMs from generating undesirable content. Still, they will definitely be unsatisfactory under the stronger goal of approximate unlearning, which aims to mimic the retrained model. We show in Figure 7 that NPO+RT outputs incorrect sentences with similar templates for questions in the forget set while generating fluent and correct answers for other questions, greatly enhancing the fluency and diversity of the generated content.



**Figure 7:** Sampled response to questions in three subsets of TOFU. Yellow: questions; Green: true answer or desired answers; Red: undesired answers.

**The role of retain loss.** Maini et al. (2024) demonstrated that methods incorporating a retain set outperform those that solely optimize a loss function based on the forget set. To further investigate the role of retain loss beyond Maini et al. (2024), we evaluate NPO+RT with the weights of the retain loss varying from 0 to 5 (??). While it is natural that adding retain loss improves the model utility, we are surprised that the forget quality also grows. Specifically, the forget quality increases as the weight of the retain loss grows from 0 to 2. We conjecture that the retain loss term helps the model preserve answer templates and linguistic structures, while the NPO term forces the model to forget some specific facts. Combining these two effects pushes the model to approximate the retrained model by generating outputs with similar templates but incorrect entities. However, further increasing the weight of the retain loss (e.g., from 2 to 5, in ??) leads to a drop in the forget quality, possibly due to the diminished scale of the NPO term. Notably, in our experiments, the retain loss plays a more significant role when we target forgetting a larger fraction of the data (See the middle and right panels of Figure 6).

**Forget KL: The larger, the better? ❌❌❌** We also examine the Forget KL during the unlearning process in the TOFU dataset. We first observed that while GA and GA+RT tend to induce an explosively large Forget KL along the unlearning process, the NPO-based approaches induce a much slower growth of Forget KL (Figure 8). It stabilizes at a moderate level even after several epochs. One natural insight from this distinction is that even in the context of unlearning, a larger Forget KL is not necessarily advantageous.

Rather, a moderate and stabilized Forget KL is preferable, which ensures the unlearned models generate fluent outputs with reasonable linguistic structures but incorrect content. This also suggests that Forget KL may not be a suitable objective function to maximize for unlearning LLMs, contrary to what was done in some prior literature (Chen & Yang, 2023).

### 5.3 Forgetting beyond 10% of TOFU

**Forgetting 20%, 30% and 50% of TOFU.** Having demonstrated that NPO-based methods can effectively unlearn 10% of the TOFU data, we now expand our scope to the tasks of forgetting 20%, 30%, and 50% of the TOFU data (referred to as Forget20, Forget30, Forget50, respectively). Details about the extended dataset are deferred to Appendix E.1.1. We show in Appendix E.2 that NPO+RT is the sole method to exhibit meaningful forget quality (a p-value above 0.05) in Forget20 and Forget30. Even in Forget50, where the vanilla NPO+RT achieves a forget quality around  $10^{-3}$ , it still significantly outperforms other methods.

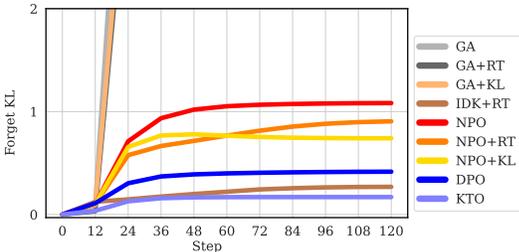
**Pushing towards the limit: forgetting 50% - 90% of TOFU.** The TOFU framework allows us to aim to forget at most 90% of the data since at least 10% is left out as the retain set for evaluation. We thus ask the question of whether there exist methods that could effectively forget 50%-90% of the TOFU data. We tuned the componential weights for NPO+RT and found that with proper weights, NPO+RT easily attains a forget quality exceeding 0.05 and model utility above 0.55 on Forget50 and Forget90, as reported in Figure 9.

## 6 Conclusion

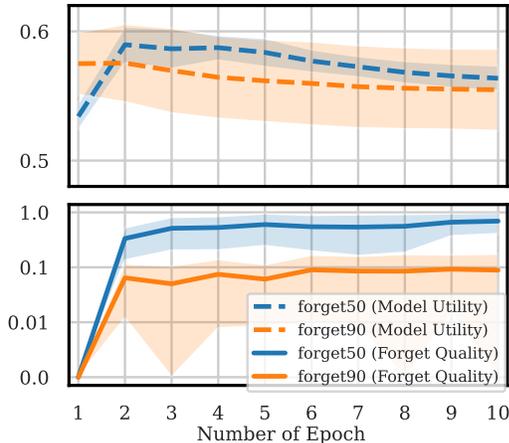
We propose Negative Preference Optimization (NPO), a simple objective for LLM unlearning. NPO makes steps towards addressing the catastrophic collapse issue in the gradient ascent method. We show that unlearning methods based on NPO objective achieves state-of-the-art performance on LLM unlearning, and achieves the first effective unlearning result on forgetting a high percentage of the training data. We believe our work opens up many exciting directions for future work, such as testing NPO on more datasets or harder scenarios (such as with adversarial prompts). It may also be of interest to generalize the algorithm principle of NPO (preference optimization with negative examples only) to other problems beyond unlearning.

## Acknowledgement

Song Mei is supported by NSF DMS-2210827, CCF-2315725, NSF Career DMS-2339904, and a Google Research Scholar Award. The authors would like to thank Baihe Huang, Xuelin Yang for the valuable discussions. The authors would like to thank Jiantao Jiao for sharing his GPU resources. This research was supported by the Center for AI Safety Compute Cluster. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.



**Figure 8:** The evolution of the Forget KL during the unlearning process on the Forget10 task in TOFU data. Note that the KL term in GA+KL is the divergence on the retain set, not the forget set. More experimental details are included in Appendix E.1.



**Figure 9:** Evolution of forget quality and model utility on Forget50 and Forget90 for NPO+RT with a proper componential weights between loss terms.

## References

- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 141–159. IEEE, 2021.
- Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pp. 463–480. IEEE, 2015.
- 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 *30th USENIX Security Symposium (USENIX Security 21)*, pp. 2633–2650, 2021.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*, 2022.
- CCPA. California consumer privacy act of 2018. [https://leginfo.ca.gov/faces/billTextClient.xhtml?bill\\_id=201720180AB375](https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB375), 2018. AB-375, Signed into law on June 28, 2018.
- Jiaao Chen and Diyi Yang. Unlearn what you want to forget: Efficient unlearning for llms. *arXiv preprint arXiv:2310.20150*, 2023.
- Shitong Duan, Xiaoyuan Yi, Peng Zhang, Tun Lu, Xing Xie, and Ning Gu. Negating negatives: Alignment without human positive samples via distributional dispreference optimization. *arXiv preprint arXiv:2403.03419*, 2024.
- Ronen Eldan and Mark Russinovich. Who’s harry potter? approximate unlearning in llms. *arXiv preprint arXiv:2310.02238*, 2023.
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*, 2024.
- Antonio Ginart, Melody Guan, Gregory Valiant, and James Y Zou. Making ai forget you: Data deletion in machine learning. *Advances in neural information processing systems*, 32, 2019.
- Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Eternal sunshine of the spotless net: Selective forgetting in deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9304–9312, 2020.
- Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. Certified data removal from machine learning models. *arXiv preprint arXiv:1911.03030*, 2019.
- Jie Huang, Hanyin Shao, and Kevin Chen-Chuan Chang. Are large pre-trained language models leaking your personal information? *arXiv preprint arXiv:2205.12628*, 2022.
- Zachary Izzo, Mary Anne Smart, Kamalika Chaudhuri, and James Zou. Approximate data deletion from machine learning models. In *International Conference on Artificial Intelligence and Statistics*, pp. 2008–2016. PMLR, 2021.
- 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.

- Jiaming Ji, Mickel Liu, Josef Dai, Xuehai Pan, Chi Zhang, Ce Bian, Boyuan Chen, Ruiyang Sun, Yizhou Wang, and Yaodong Yang. Beavertails: Towards improved safety alignment of llm via a human-preference dataset. *Advances in Neural Information Processing Systems*, 36, 2024.
- Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pp. 1885–1894. PMLR, 2017.
- 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.
- Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pp. 74–81, 2004.
- Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Xiaojun Xu, Yuguang Yao, Hang Li, Kush R Varshney, et al. Rethinking machine unlearning for large language models. *arXiv preprint arXiv:2402.08787*, 2024a.
- Zheyuan Liu, Guangyao Dou, Zhaoxuan Tan, Yijun Tian, and Meng Jiang. Towards safer large language models through machine unlearning. *arXiv preprint arXiv:2402.10058*, 2024b.
- Aengus Lynch, Phillip Guo, Aidan Ewart, Stephen Casper, and Dylan Hadfield-Menell. Eight methods to evaluate robust unlearning in llms. *arXiv preprint arXiv:2402.16835*, 2024.
- 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.
- Alessandro Mantelero. The eu proposal for a general data protection regulation and the roots of the ‘right to be forgotten’. *Computer Law & Security Review*, 29(3):229–235, 2013.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372, 2022.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. Memory-based model editing at scale. In *International Conference on Machine Learning*, pp. 15817–15831. PMLR, 2022.
- Thanh Tam Nguyen, Thanh Trung Huynh, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. A survey of machine unlearning. *arXiv preprint arXiv:2209.02299*, 2022.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- Vaidehi Patil, Peter Hase, and Mohit Bansal. Can sensitive information be deleted from llms? objectives for defending against extraction attacks. *arXiv preprint arXiv:2309.17410*, 2023.
- Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models as few shot unlearners. *arXiv preprint arXiv:2310.07579*, 2023.
- 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. *Advances in Neural Information Processing Systems*, 36, 2024.
- Jonas B Sandbrink. Artificial intelligence and biological misuse: Differentiating risks of language models and biological design tools. *arXiv preprint arXiv:2306.13952*, 2023.

- Ayush Sekhari, Jayadev Acharya, Gautam Kamath, and Ananda Theertha Suresh. Remember what you want to forget: Algorithms for machine unlearning. *Advances in Neural Information Processing Systems*, 34:18075–18086, 2021.
- Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. Detecting pretraining data from large language models. *arXiv preprint arXiv:2310.16789*, 2023.
- Nianwen Si, Hao Zhang, Heyu Chang, Wenlin Zhang, Dan Qu, and Weiqiang Zhang. Knowledge unlearning for llms: Tasks, methods, and challenges. *arXiv preprint arXiv:2311.15766*, 2023.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020.
- Anvith Thudi, Gabriel Deza, Varun Chandrasekaran, and Nicolas Papernot. Unrolling sgd: Understanding factors influencing machine unlearning. In *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*, pp. 303–319. IEEE, 2022.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Paul Voigt and Axel Von dem Bussche. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing*, 10(3152676):10–5555, 2017.
- Lingzhi Wang, Tong Chen, Wei Yuan, Xingshan Zeng, Kam-Fai Wong, and Hongzhi Yin. Kga: A general machine unlearning framework based on knowledge gap alignment. *arXiv preprint arXiv:2305.06535*, 2023.
- Heng Xu, Tianqing Zhu, Lefeng Zhang, Wanlei Zhou, and Philip S Yu. Machine unlearning: A survey. *ACM Computing Surveys*, 56(1):1–36, 2023.
- Jin Yao, Eli Chien, Minxin Du, Xinyao Niu, Tianhao Wang, Zezhou Cheng, and Xiang Yue. Machine unlearning of pre-trained large language models. *arXiv preprint arXiv:2402.15159*, 2024.
- Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. *arXiv preprint arXiv:2310.10683*, 2023.
- Haibo Zhang, Toru Nakamura, Takamasa Isohara, and Kouichi Sakurai. A review on machine unlearning. *SN Computer Science*, 4(4):337, 2023.

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Preliminaries on Machine Unlearning</b>	<b>3</b>
2.1	Catastrophic collapse of gradient ascent . . . . .	4
<b>3</b>	<b>Negative Preference Optimization</b>	<b>4</b>
3.1	Theoretical analysis of divergence speed . . . . .	6
<b>4</b>	<b>Synthetic Experiments</b>	<b>6</b>
4.1	Setup . . . . .	6
4.2	Results . . . . .	7
<b>5</b>	<b>Experiments on the TOFU Data</b>	<b>7</b>
5.1	Experimental setup . . . . .	7
5.2	Results . . . . .	8
5.3	Forgetting beyond 10% of TOFU . . . . .	10
<b>6</b>	<b>Conclusion</b>	<b>10</b>
<b>A</b>	<b>Related work</b>	<b>14</b>
<b>B</b>	<b>Proofs</b>	<b>15</b>
B.1	Proof of Proposition 1 . . . . .	15
B.2	Proof of Theorem 2 . . . . .	16
B.3	Proof of Lemma 3 . . . . .	17
B.4	Proof of Lemma 4 . . . . .	18
<b>C</b>	<b>The role of KL Divergence on the Forget Set</b>	<b>21</b>
C.1	Synthetic Experiment . . . . .	21
<b>D</b>	<b>Experimental details of the synthetic experiments</b>	<b>22</b>
<b>E</b>	<b>Experiments on the TOFU dataset</b>	<b>22</b>
E.1	Experiments Setup . . . . .	22
E.1.1	Dataset . . . . .	22
E.1.2	Metrics . . . . .	23
E.1.3	Baseline Methods . . . . .	24
E.2	Full Results . . . . .	25

## A Related work

Since its proposal by Cao & Yang (2015), machine unlearning has been extensively studied in the classification literature (Bourtoule et al., 2021; Golatkar et al., 2020; Ginart et al., 2019; Thudi et al., 2022; Izzo et al., 2021; Koh & Liang, 2017; Guo et al., 2019; Sekhari et al., 2021). For reviews of existing works, see Liu et al. (2024a); Zhang et al. (2023); Nguyen et al. (2022); Xu et al. (2023); Si et al. (2023). In particular, Ginart et al. (2019); Guo et al. (2019); Sekhari et al. (2021) introduced theoretical metrics for machine unlearning based on the notion of differential privacy and proposed provably efficient unlearning methods based on Newton update removal mechanisms. However, these algorithms require computing the Hessian of loss functions, which is intractable for LLMs.

Recent research has explored unlearning methods for LLMs (Jang et al., 2022; Wang et al., 2023; Chen & Yang, 2023; Yao et al., 2023; Eldan & Russinovich, 2023; Yao et al., 2024; Liu et al., 2024b; Li et al., 2024). Notably, the methods proposed in Jang et al. (2022); Yao et al. (2023); Chen & Yang (2023); Maini et al. (2024) are based on gradient ascent (GA) on the loss of the forget set. In this work, we demonstrate that the NPO approach consistently outperforms GA across various tasks. On the other hand, Eldan & Russinovich (2023) proposed generating positive samples using LLMs and carefully designed prompts, then

fine-tuning the model based on the positive samples using a supervised loss. Furthermore, the method of Liu et al. (2024b) is based on knowledge negation, while the approach of Li et al. (2024) relies on controlling model representations. These methods are orthogonal and complementary to the NPO approach.

Our method, NPO, draws inspiration from the framework of reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022; Bai et al., 2022; Stiennon et al., 2020; Rafailov et al., 2024), particularly the Direct Policy Optimization (DPO) method (Rafailov et al., 2024). We note that recent work (Ethayarajh et al., 2024) proposes the Kahneman-Tversky Optimization (KTO) method for alignment with only non-paired preference data, and a more recent concurrent work (Duan et al., 2024) proposes the Distributional Dispreference Optimization (D<sup>2</sup>O) approach for unlearning. Both methods share a similar formulation to NPO. We compare the performance of NPO with KTO in simulations.

Recent work has proposed several benchmark datasets and evaluation metrics for unlearning methods (Ji et al., 2024; Eldan & Russinovich, 2023; Maini et al., 2024; Li et al., 2024; Lynch et al., 2024). In particular, some studies have utilized the PKUSafe dataset (Ji et al., 2024) for benchmarking unlearning methods. Eldan & Russinovich (2023) crafts a specific task of “forgetting Harry Potter”. Maini et al. (2024) introduces TOFU, a task of fictitious unlearning for LLMs, which is the benchmark we adopted in this paper. Additionally, Li et al. (2024) proposes the Weapons of Mass Destruction Proxy (WMDP) for measuring hazardous knowledge in LLMs. Lynch et al. (2024) proposes eight methods to evaluate robust unlearning in LLM, which incorporate robust metrics against jailbreak attacks.

Finally, we note the existence of attack methods for extracting data from unlearned models (Shi et al., 2023; Patil et al., 2023), and other unlearning methods including model editing (Mitchell et al., 2022; Meng et al., 2022) and in-context unlearning (Pawelczyk et al., 2023).

## B Proofs

### B.1 Proof of Proposition 1

Adopt the shorthand  $R_i = \log \frac{\pi_\theta(y_i | x_i)}{\pi_{\text{ref}}(y_i | x_i)}$ . For any  $i \in [n_f]$ , we have

$$\begin{aligned} \lim_{\beta \rightarrow 0} -\frac{2}{\beta} \cdot \log \sigma \left( -\beta \cdot \log \frac{\pi_\theta(y_i | x_i)}{\pi_{\text{ref}}(y_i | x_i)} \right) - \frac{2}{\beta} \log 2 &= \lim_{\beta \rightarrow 0} -\frac{2}{\beta} \cdot \log \sigma(-\beta R_i) - \frac{2}{\beta} \log 2 \\ &= \lim_{\beta \rightarrow 0} \frac{2}{\beta} \cdot \log \left( \frac{1 + \exp(\beta R_i)}{2} \right) \\ &= \lim_{\beta \rightarrow 0} \frac{2}{\beta} \cdot \log \left( 1 + \frac{\exp(\beta R_i) - 1}{2} \right) \\ &= \lim_{\beta \rightarrow 0} \frac{2}{\beta} \cdot \frac{\exp(\beta R_i) - 1}{2} \\ &= R_i. \end{aligned}$$

Averaging over all  $i \in [n_f]$  and noting that  $\sum_{i=1}^{n_f} R_i / n_f = \mathcal{L}_{\text{GA}}(\theta) - \mathbb{E}_{\mathcal{D}_{\text{FG}}}[\log \pi_{\text{ref}}(y_i | x_i)]$  yields the first part of Proposition 1.

For the second part of Proposition 1, by definition

$$\begin{aligned} \nabla_\theta \mathcal{L}_{\text{NPO}, \beta}(\theta) &= \frac{1}{n_f} \sum_{i=1}^{n_f} -\frac{2}{\beta} \nabla_\theta \log \sigma(-\beta R_i) \\ &= \frac{1}{n_f} \sum_{i=1}^{n_f} \frac{2}{\beta} \nabla_\theta \log(1 + \exp(\beta R_i)) \\ &= \frac{1}{n_f} \sum_{i=1}^{n_f} \frac{2}{\beta} \cdot \frac{\beta \exp(\beta R_i)}{1 + \exp(\beta R_i)} \cdot \nabla_\theta R_i \end{aligned}$$

$$= \frac{1}{n_f} \sum_{i=1}^{n_f} \frac{2 \exp(\beta R_i)}{1 + \exp(\beta R_i)} \cdot \nabla_{\theta} R_i.$$

Therefore, it follows immediately that

$$\begin{aligned} \lim_{\beta \rightarrow 0} \nabla_{\theta} \mathcal{L}_{\text{NPO}, \beta}(\theta) &= \lim_{\beta \rightarrow 0} \frac{1}{n_f} \sum_{i=1}^{n_f} \frac{2 \exp(\beta R_i)}{1 + \exp(\beta R_i)} \cdot \nabla_{\theta} R_i = \frac{1}{n_f} \sum_{i=1}^{n_f} \nabla_{\theta} R_i \\ &= \frac{1}{n_f} \sum_{i=1}^{n_f} \nabla_{\theta} \log \pi_{\theta}(y_i | x_i) = \nabla_{\theta} \mathcal{L}_{\text{GA}}(\theta). \end{aligned}$$

This completes the proof of Proposition 1.

## B.2 Proof of Theorem 2

Let  $\eta_0 := \eta/n_f$  denote the normalized learning rate. For  $i, j \in [n_f]$ , define  $\gamma_{i,j} := \langle x_i, x_j \rangle$  and  $c_{\text{init},i} := \langle \theta_{\text{init}}, x_i \rangle$ . Throughout the proof, we also use  $C_k (k = 0, 1, 2, \dots)$  to denote constants that may depend on  $(B_{\theta}, b_x, B_x, \beta)$  but not on  $(t, n_f, d)$ . By the definition of the logistic model and some algebra, we have

$$\nabla_{\theta} \log \pi_{\theta}(y_i | x_i) = x_i(2y_i - 1)(1 - \pi_{\theta}(y_i | x_i)).$$

Therefore, the gradient of  $\mathcal{L}_{\text{GA}}$  and  $\mathcal{L}_{\text{NPO}}$  both lie in the span of  $x_1, \dots, x_{n_f}$ . Consequently,  $\theta_{\text{GA}}^{(t)} - \theta_{\text{init}}$  and  $\theta_{\text{NPO}}^{(t)} - \theta_{\text{init}}$  can be rewritten as follows:

$$\theta_{\text{GA}}^{(t)} - \theta_{\text{init}} = \sum_{i=1}^{n_f} \alpha_{\text{GA},i}^{(t)} \cdot x_i, \quad \theta_{\text{NPO}}^{(t)} - \theta_{\text{init}} = \sum_{i=1}^{n_f} \alpha_{\text{NPO},i}^{(t)} \cdot x_i,$$

where

$$\alpha_{\text{GA},i}^{(t+1)} := \alpha_{\text{GA},i}^{(t)} - \eta_0(2y_i - 1)(1 - \pi_{\theta}(y_i | x_i)), \quad \alpha_{\text{NPO},i}^{(t+1)} := \alpha_{\text{NPO},i}^{(t)} - \eta_0(2y_i - 1)(1 - \pi_{\theta}(y_i | x_i)) \cdot W_i^{(t)},$$

and

$$W_i^{(t)} = \pi_{\theta_{\text{NPO}}^{(t)}}^{\beta}(y_i | x_i) / [\pi_{\theta_{\text{NPO}}^{(t)}}^{\beta}(y_i | x_i) + \pi_{\theta_{\text{init}}}^{\beta}(y_i | x_i)].$$

We also define  $\alpha_{\star,i}^{(0)} = 0$  and adopt the shorthand notations  $\alpha_{\star}^{(t)} = (\alpha_{1,\star}^{(t)}, \dots, \alpha_{n_f,\star}^{(t)})$  for  $\star \in \{\text{GA}, \text{NPO}\}$  and  $\gamma_i = (\gamma_{i,1}, \dots, \gamma_{i,n_f})$ .

For  $\star \in \{\text{GA}, \text{NPO}\}$ , we have

$$\begin{aligned} \pi_{\theta_{\star}^{(t)}}(y_i | x_i) &= \frac{1}{1 + \exp\left(\langle x_i, \theta_{\star}^{(t)} \rangle\right)} = \frac{1}{1 + \exp\left(\langle x_i, c_{\text{init},i} + (1 - 2y_i) \sum_{j=1}^{n_f} \alpha_{\star}^{(t)} x_j \rangle\right)} \\ &= \frac{1}{1 + \exp\left(\langle x_i, c_{\text{init},i} + (1 - 2y_i) \alpha_{\star}^{(t)} \rangle\right)} =: \text{pred}_i\left(\langle \alpha_{\star}^{(t)}, \gamma_i \rangle\right), \end{aligned}$$

where we denote the dependence on  $(c_{\text{init},i}, y_i)$  implicitly using  $\text{pred}_i$  for notational simplicity.

Therefore, letting  $b_{\star,i}^{(t)} := \langle \alpha_{\star}^{(t)}, \gamma_i \rangle$  for  $\star \in \{\text{GA}, \text{NPO}\}$  and combining all  $i \in [n_f]$ , we obtain

$$b_{\text{GA}}^{(t+1)} := b_{\text{GA}}^{(t)} - \eta_0 \cdot \Gamma \Delta_{\text{GA}}^{(t)}, \quad b_{\text{NPO}}^{(t+1)} := b_{\text{NPO}}^{(t)} - \eta_0 \cdot \Gamma \text{diag}\left\{\Delta_{\text{NPO}}^{(t)}\right\} W^{(t)},$$

where

$$\Gamma := (\gamma_{i,j})_{1 \leq i, j \leq n_f},$$

$$\Delta_{\star}^{(t)} := \left(\Delta_{\star,1}^{(t)}, \dots, \Delta_{\star,n_f}^{(t)}\right)^{\top}, \quad \Delta_{\star,i}^{(t)} := (2y_i - 1)(1 - \text{pred}_i(b_{\star,i}^{(t)})), \quad \text{for } i \in [n_f], \star \in \{\text{GA}, \text{NPO}\},$$

$$W^{(t)} := (W_1^{(t)}, \dots, W_{n_f}^{(t)})^{\top}, \quad W_i^{(t)} = W_i\left(b_{\text{NPO}}^{(t)}\right) := \pi_{\theta_{\text{NPO}}^{(t)}}^{\beta}(y_i | x_i) / [\pi_{\theta_{\text{NPO}}^{(t)}}^{\beta}(y_i | x_i) + \pi_{\theta_{\text{init}}}^{\beta}(y_i | x_i)].$$

Again, we hide here the dependence on  $(\beta, c_{\text{init},i}, y_i)$  in  $W_i$  for simplicity. We now claim the following results of which the proofs are deferred to Section B.3 and B.4.

**Lemma 3** (GA converges to infinity linearly). *Under the assumptions in Theorem 2 and the notations in the proof of Theorem 2, there exist some  $(B_\theta, b_x, B_x)$ -dependent constants  $C_0, C_1, C_2 > 0$  such that the GA iterations  $\{b_{\text{GA}}^{(t)}\}_{t=1}^\infty$  satisfy*

$$\begin{aligned} C_1\eta_0 t &\leq b_{\text{GA},i}^{(t)} \leq C_2\eta_0 t, \text{ when } y_i = 0, \\ -C_2\eta_0 t &\leq b_{\text{GA},i}^{(t)} \leq -C_1\eta_0 t, \text{ when } y_i = 1. \end{aligned}$$

for all  $i \in [n_f]$  and  $t \geq 1$  when  $\max_{i \neq j} |\gamma_{i,j}| \leq C_0/n_f$ .

**Lemma 4** (NPO converges to infinity exponentially slow). *Under the assumptions in Theorem 2 and the notations in the proof of Theorem 2, there exist some  $(B_\theta, b_x, B_x, \beta)$ -dependent constants  $C_0, C_b > 0, C_a \in (0, 1)$  such that the NPO iterations  $\{b_{\text{NPO}}^{(t)}\}_{t=1}^\infty$  satisfy*

$$\begin{aligned} b_{\text{NPO},i}^{(t)} &\in \left[-\frac{1}{\beta} \log(C_b\eta_0 t + 1), -\frac{1}{\beta} \log(C_a\eta_0 t + 1)\right] \text{ when } y_i = 1, \\ b_{\text{NPO},i}^{(t)} &\in \left[\frac{1}{\beta} \log(C_a\eta_0 t + 1), \frac{1}{\beta} \log(C_b\eta_0 t + 1)\right] \text{ when } y_i = 0. \end{aligned}$$

for all  $i \in [n_f]$  and  $t \geq 1$  when  $\max_{i \neq j} |\gamma_{i,j}| \leq C_0/n_f$ .

Combining Lemma 3 and 4 and noting

$$\begin{aligned} \|\theta_\star^{(t)} - \theta_{\text{init}}\|_{X^\top X} &= \sqrt{(\theta_\star^{(t)} - \theta_{\text{init}})^\top X^\top X (\theta_\star^{(t)} - \theta_{\text{init}})} = \sqrt{\alpha_\star^{(t)\top} X X^\top X X^\top \alpha_\star^{(t)}} \\ &= \sqrt{b_\star^{(t)\top} (X^\top X)^{-1} X X^\top X X^\top (X^\top X)^{-1} b_\star^{(t)}} = \|b_\star^{(t)}\|_2 \end{aligned}$$

for  $\star \in \{\text{GA}, \text{NPO}\}$  completes the proof.

### B.3 Proof of Lemma 3

We prove Lemma 3 by induction.

**Case 1:**  $t = 1$  When  $t = 1$ , since  $|c_{\text{init}}| \leq \|x_i\|_2 \cdot \|\theta_{\text{init}}\|_2 \leq B_x B_\theta$ , it follows from the definition of  $\text{pred}_i(\cdot)$  that

$$C_3 \leq \Delta_{\text{GA},i}^{(0)} \leq 1 \text{ when } y_i = 1, \quad -1 \leq \Delta_{\text{GA},i}^{(0)} \leq -C_4 \text{ when } y_i = 0$$

for all  $i \in [n_f]$  for some constants  $C_3, C_4 \in (0, 1)$  depending only on  $(B_\theta, b_x, B_x)$ . Note that there exists a constant  $C_0 > 0$  depending on  $C_3, C_4, b_x$  such that

$$\left| \sum_{j \neq i} \gamma_{i,j} \Delta_{\text{GA},j}^{(0)} \right| \leq \frac{\gamma_{i,i}}{2} |\Delta_{\text{GA},i}^{(0)}|$$

for all  $i \in [n_f]$  when  $\max_{i \neq j} |\gamma_{i,j}| \leq C_0/n_f$ . It follows that

$$\begin{aligned} -\eta_0 \cdot \gamma_i^\top \Delta_{\text{GA}}^{(0)} &\in \left[-\frac{3}{2} \gamma_{i,i} \eta_0 |\Delta_{\text{GA},i}^{(0)}|, -\frac{1}{2} \gamma_{i,i} \eta_0 |\Delta_{\text{GA},i}^{(0)}|\right] \in [-C_2\eta_0, -C_1\eta_0] \text{ when } y_i = 1, \\ -\eta_0 \cdot \gamma_i^\top \Delta_{\text{GA}}^{(0)} &\in \left[\frac{1}{2} \gamma_{i,i} \eta_0 |\Delta_{\text{GA},i}^{(0)}|, \frac{3}{2} \gamma_{i,i} \eta_0 |\Delta_{\text{GA},i}^{(0)}|\right] \in [C_1\eta_0, C_2\eta_0] \text{ when } y_i = 0 \end{aligned}$$

for some  $(B_\theta, b_x, B_x)$ -dependent constants  $C_1, C_2 > 0$ . Therefore,

$$\begin{aligned} b_{\text{GA},i}^{(1)} &= b_{\text{GA},i}^{(0)} - \eta_0 \cdot \gamma_i^\top \Delta_{\text{GA}}^{(0)} \in [-C_2\eta_0, -C_1\eta_0] \text{ when } y_i = 1, \\ b_{\text{GA},i}^{(1)} &= b_{\text{GA},i}^{(0)} - \eta_0 \cdot \gamma_i^\top \Delta_{\text{GA}}^{(0)} \in [C_1\eta_0, C_2\eta_0] \text{ when } y_i = 0. \end{aligned}$$

As a consequence,

$$\begin{aligned} b_{\text{GA},i}^{(1)} &\leq b_{\text{GA},i}^{(0)} = 0, \quad \Delta_{\text{GA},i}^{(1)} \geq \Delta_{\text{GA},i}^{(0)} \text{ when } y_i = 1, \\ b_{\text{GA},i}^{(1)} &\geq b_{\text{GA},i}^{(0)} = 0, \quad \Delta_{\text{GA},i}^{(1)} \leq \Delta_{\text{GA},i}^{(0)} \text{ when } y_i = 0. \end{aligned}$$

**Case 2:**  $t = K + 1$  Now, suppose we have

$$\begin{aligned} b_{GA,i}^{(t)} &\in [-C_2\eta_0t, -C_1\eta_0t] \text{ when } y_i = 1, \\ b_{GA,i}^{(t)} &\in [C_1\eta_0t, C_2\eta_0t] \text{ when } y_i = 0 \end{aligned}$$

for  $t \in [K]$  and

$$\begin{aligned} b_{GA,i}^{(K)} &\leq \dots \leq b_{GA,i}^{(0)} = 0, \quad \Delta_{GA,i}^{(K)} \geq \dots \geq \Delta_{GA,i}^{(0)} \text{ when } y_i = 1, \\ b_{GA,i}^{(K)} &\geq \dots \geq b_{GA,i}^{(0)} = 0, \quad \Delta_{GA,i}^{(K)} \leq \dots \leq \Delta_{GA,i}^{(0)} \text{ when } y_i = 0. \end{aligned}$$

By the monotonicity of  $\Delta_{GA,i}^{(t)}$ , we have

$$C_3 \leq \Delta_{GA,i}^{(K)} \leq 1 \text{ when } y_i = 1, \quad -1 \leq \Delta_{GA,i}^{(K)} \leq -C_4 \text{ when } y_i = 0$$

for all  $i \in [n_t]$ . Therefore, following similar arguments as in the  $t = 1$  case, we have

$$\begin{aligned} -\eta_0 \cdot \gamma_i^\top \Delta_{GA}^{(K)} &\in [-C_2\eta_0, -C_1\eta_0] \text{ when } y_i = 1, \\ -\eta_0 \cdot \gamma_i^\top \Delta_{GA}^{(K)} &\in [C_1\eta_0, C_2\eta_0] \text{ when } y_i = 0. \end{aligned}$$

Then it follows from the induction assumption that

$$\begin{aligned} b_{GA,i}^{(K+1)} &= b_{GA,i}^{(K)} - \eta_0 \cdot \gamma_i^\top \Delta_{GA}^{(K)} \in [-C_2\eta_0(K+1), -C_1\eta_0(K+1)] \text{ when } y_i = 1, \\ b_{GA,i}^{(K+1)} &= b_{GA,i}^{(K)} - \eta_0 \cdot \gamma_i^\top \Delta_{GA}^{(K)} \in [C_1\eta_0(K+1), C_2\eta_0(K+1)] \text{ when } y_i = 0, \end{aligned}$$

and also

$$\begin{aligned} b_{GA,i}^{(K+1)} &\leq b_{GA,i}^{(K)}, \quad \Delta_{GA,i}^{(K+1)} \geq \Delta_{GA,i}^{(K)} \text{ when } y_i = 1, \\ b_{GA,i}^{(K+1)} &\geq b_{GA,i}^{(K)}, \quad \Delta_{GA,i}^{(K+1)} \leq \Delta_{GA,i}^{(K)} \text{ when } y_i = 0. \end{aligned}$$

This concludes the induction step and therefore completes the proof.

#### B.4 Proof of Lemma 4

We prove Lemma 4 by induction.

Our induction assumption is the following: there exist some constants  $C_0 > 0, C_a \in (0, 1), C_b > 0, C_1, C_2$  depending only on  $(B_\theta, b_x, B_x, \beta)$  such that when  $\max_{i \neq j} |\gamma_{i,j}| \leq C_0/n_t$ , for any  $t \geq 1$  we have

1.

$$\begin{aligned} b_{NPO,i}^{(t)} &\leq \dots \leq b_{NPO,i}^{(0)} = 0, \quad \Delta_{NPO,i}^{(t)} \geq \dots \geq \Delta_{NPO,i}^{(0)} \text{ when } y_i = 1, \\ b_{NPO,i}^{(t)} &\geq \dots \geq b_{NPO,i}^{(0)} = 0, \quad \Delta_{NPO,i}^{(t)} \leq \dots \leq \Delta_{NPO,i}^{(0)} \text{ when } y_i = 0. \end{aligned}$$

2.

$$b_{NPO,i}^{(t)} \in \left[-\frac{1}{\beta} \log(C_b\eta_0t + 1), -\frac{1}{\beta} \log(C_a\eta_0t + 1)\right] \text{ when } y_i = 1, \quad (7)$$

$$b_{NPO,i}^{(t)} \in \left[\frac{1}{\beta} \log(C_a\eta_0t + 1), \frac{1}{\beta} \log(C_b\eta_0t + 1)\right] \text{ when } y_i = 0. \quad (8)$$

3.

$$b_{NPO,i}^{(t)} \in \left[-C_2 - \frac{1}{\beta} \log(\eta_0t + 1), -C_1 - \frac{1}{\beta} \log(\eta_0t + 1)\right] \text{ when } y_i = 1, \quad (9)$$

$$b_{NPO,i}^{(t)} \in \left[C_1 + \frac{1}{\beta} \log(\eta_0t + 1), C_2 + \frac{1}{\beta} \log(\eta_0t + 1)\right] \text{ when } y_i = 0. \quad (10)$$

Lemma 4 follows immediately from the second part of the induction assumption.

In the following, we first specify the parameter-dependent constants  $C_0, C_1, C_2, C_a, C_b$  in the  $t = 1$  case and prove the induction assumption when  $k = 1$ . Then given the induction assumption holds when  $t \leq K$ , we prove that it holds when  $t = K + 1$  as well.

**Case 1:**  $t = 1$  When  $t = 1$ , since  $|c_{\text{init}}| \leq \|x_i\|_2 \cdot \|\theta_{\text{init}}\|_2 \leq B_x B_\theta$ , it follows from the definition of  $\text{pred}_i(\cdot)$  that

$$C_3 \leq \Delta_{\text{NPO},i}^{(0)} \leq 1 \text{ when } y_i = 1, \quad -1 \leq \Delta_{\text{NPO},i}^{(0)} \leq -C_4 \text{ when } y_i = 0 \quad (11)$$

for all  $i \in [n_f]$  for some constants  $C_3, C_4 \in (0, 1)$  depending only on  $(B_\theta, b_x, B_x)$ . Moreover, we claim that

$$W_i^{(t)} \in \left[ C_5 \cdot \exp\left((2y_i - 1)\beta b_{\text{NPO},i}^{(t)}\right), C_6 \cdot \exp\left((2y_i - 1)\beta b_{\text{NPO},i}^{(t)}\right) \right] \quad (12)$$

for all  $i \in [n_f]$  and  $t$  such that  $\text{pred}_i(b_{\text{NPO},i}^{(t)}) \leq \text{pred}_i(b_{\text{NPO},i}^{(0)})$  for some  $(B_\theta, b_x, B_x, \beta)$ -dependent constants  $C_5, C_6 > 0$ .

Now, suppose Eq. (9) and (10) hold for some  $(B_\theta, b_x, B_x, \beta)$ -dependent constants  $C_1, C_2 > 0$  which we will specify later. Then, there exists a constant  $C_0 > 0$  depending on  $C_{1:6}, b_x$  such that

$$\left| \sum_{j \neq i} \gamma_{i,j} \Delta_{\text{NPO},j}^{(0)} W_j \right| \leq \frac{\gamma_{i,i}}{2} \left| \Delta_{\text{NPO},i}^{(0)} W_i \right| \quad (13)$$

for all  $i \in [n_f]$  when  $\max_{i \neq j} |\gamma_{i,j}| \leq C_0/n_f$ . Furthermore, combining Eq. (11), (12), (13) gives

$$\begin{aligned} -\eta_0 \cdot \gamma_i^\top \text{diag}\{\Delta_{\text{NPO}}^{(0)}\} W_i^{(0)} &\in \left[ -\frac{3}{2} \gamma_{i,i} \eta_0 |\Delta_{\text{NPO},i}^{(0)}| W_i^{(0)}, -\frac{1}{2} \gamma_{i,i} \eta_0 |\Delta_{\text{NPO},i}^{(0)}| W_i^{(0)} \right] \\ &\in [-C_8 \eta_0 \exp(\beta b_{\text{NPO},i}^{(0)}), -C_7 \eta_0 \exp(\beta b_{\text{NPO},i}^{(0)})] \text{ when } y_i = 1, \\ -\eta_0 \cdot \gamma_i^\top \text{diag}\{\Delta_{\text{NPO}}^{(0)}\} W_i^{(0)} &\in \left[ \frac{1}{2} \gamma_{i,i} \eta_0 |\Delta_{\text{NPO},i}^{(0)}| W_i^{(0)}, \frac{3}{2} \gamma_{i,i} \eta_0 |\Delta_{\text{NPO},i}^{(0)}| W_i^{(0)} \right] \\ &\in [C_7 \eta_0 \exp(-\beta b_{\text{NPO},i}^{(0)}), C_8 \eta_0 \exp(\beta b_{\text{NPO},i}^{(0)})] \text{ when } y_i = 0, \end{aligned}$$

if  $\max_{i \neq j} |\gamma_{i,j}| \leq C_0/n_f$ . Here  $C_7, C_8 > 0$  are some  $(C_3, C_4, C_5, C_6, \beta)$ -dependent constants, and we pick  $C_7$  such that  $C_7 \beta < 1$ . As a consequence,

$$\begin{aligned} b_{\text{NPO},i}^{(1)} &\leq b_{\text{NPO},i}^{(0)} = 0, \quad \Delta_{\text{NPO},i}^{(1)} \geq \Delta_{\text{NPO},i}^{(0)} \text{ when } y_i = 1, \\ b_{\text{NPO},i}^{(1)} &\geq b_{\text{NPO},i}^{(0)} = 0, \quad \Delta_{\text{NPO},i}^{(1)} \leq \Delta_{\text{NPO},i}^{(0)} \text{ when } y_i = 0. \end{aligned}$$

This concludes the proof of the first part of the induction assumption.

Now, we start to prove the second part of the induction assumption. For  $i$  such that  $y_i = 1$ , consider the ordinary differential equations

$$\begin{aligned} b_l'(t) &= -C_8 \eta_0 (1 + \exp(C_8)) \cdot \exp(\beta b_l(t)), \quad b_l(0) = 0; \\ b_u'(t) &= -C_7 \eta_0 \cdot \exp(\beta b_l(t)), \quad b_u(0) = 0. \end{aligned}$$

It can be verified that the ODEs have closed-form solutions

$$\begin{aligned} b_l(t) &= -\frac{1}{\beta} \log(\beta C_8 \eta_0 (1 + \exp(C_8)) t + 1), \\ b_u(t) &= -\frac{1}{\beta} \log(\beta C_7 \eta_0 t + 1). \end{aligned}$$

Since

$$b_{\text{NPO},i}^{(1)} = b_{\text{NPO},i}^{(0)} - \eta_0 \cdot \gamma_i^\top \Delta_{\text{NPO},i}^{(0)} \geq b_{\text{NPO},i}^{(0)} - C_8 \eta_0 \exp(\beta b_{\text{NPO},i}^{(0)}) \geq b_{\text{NPO},i}^{(0)} - C_8 \eta_0,$$

it follows that for any point  $b_{\text{NPO},i}^{(\varepsilon)} = \varepsilon b_{\text{NPO},i}^{(1)} + (1 - \varepsilon) b_{\text{NPO},i}^{(0)}$  with  $\varepsilon \in [0, 1]$

$$\begin{aligned} -C_8 \eta_0 (1 + \exp(C_8 \eta_0)) \cdot \exp(\beta b_{\text{NPO},i}^{(\varepsilon)}) &\leq -C_8 \eta_0 \exp(\beta b_{\text{NPO},i}^{(0)}) \leq -C_7 \eta_0 \exp(\beta b_{\text{NPO},i}^{(0)}) \\ &\leq -C_7 \eta_0 \cdot \exp(\beta b_{\text{NPO},i}^{(\varepsilon)}). \end{aligned}$$

Therefore, we have by the comparison theorem for ODEs that

$$b_l(\varepsilon) \leq b_{\text{NPO},i}^{(\varepsilon)} \leq b_u(\varepsilon)$$

for  $\varepsilon \in [0, 1]$ . Setting

$$C_a := C_7\beta, \quad C_b := \beta C_8(1 + \exp(C_8))$$

concludes the second part of the induction assumption.

For the last part of the induction assumption, since  $\log(x+1) + \log(c) \leq \log(cx+1) \leq \log(x+1) + \log(c+1)$  when  $c \leq 1$ , we have

$$-\frac{1}{\beta}[\log(\eta_0 t + 1) + \log(\beta C_8(1 + \exp(C_8)))] \leq b_l(t) \leq b_u(t) \leq -\frac{1}{\beta}[\log(\eta_0 t + 1) + \log(\beta C_7)],$$

where the last inequality uses  $\beta C_7 < 1$ . Therefore, we obtain

$$b_{\text{NPO},i}^{(1)} \in \left[-\frac{1}{\beta} \log(\eta_0 t + 1) - C_2, -\frac{1}{\beta} \log(\eta_0 t + 1) - C_1\right],$$

where

$$C_1 := \frac{1}{\beta} \log(\beta C_7), \quad C_2 := \frac{1}{\beta} \log(\beta C_8(1 + \exp(C_8)))$$

for  $i$  such that  $y_i = 1$ . Following the same arguments, similarly, we also have

$$b_{\text{NPO},i}^{(1)} \in \left[\frac{1}{\beta} \log(\eta_0 t + 1) + C_1, \frac{1}{\beta} \log(\eta_0 t + 1) + C_2\right]$$

for  $i$  such that  $y_i = 0$ . This concludes the last part of the induction assumption.

**Case 2:**  $t = K + 1$  Suppose the induction assumption holds for  $t \in [K]$ , we now show that the induction assumption holds for  $t = K + 1$  as well. Following the proof of  $t = 1$  case and using the monotonicity property of  $\{\Delta_{\text{NPO},i}^{(t)}\}_{t=0}^K$ , we have

$$C_3 \leq \Delta_{\text{NPO},i}^{(K)} \leq 1 \text{ when } y_i = 1, \quad -1 \leq \Delta_{\text{NPO},i}^{(K)} \leq -C_4 \text{ when } y_i = 0$$

for all  $i \in [n_t]$ . Since  $\text{pred}_i(b_{\text{NPO},i}^{(K)}) \leq \text{pred}_i(b_{\text{NPO},i}^{(K)})$  by the definition of  $\text{pred}_i$  and the monotonicity of  $\{b_{\text{NPO},i}^{(t)}\}_{t=0}^K$ , it follows from Claim (12) that

$$W_i^{(K)} \in \left[C_5 \cdot \exp\left((2y_i - 1)\beta b_{\text{NPO},i}^{(K)}\right), C_6 \cdot \exp\left((2y_i - 1)\beta b_{\text{NPO},i}^{(K)}\right)\right].$$

Putting the last two displays together and following the same argument as the  $t = 1$  case, we find that

$$\begin{aligned} -\eta_0 \cdot \gamma_i^\top \text{diag}\{\Delta_{\text{NPO}}^{(K)}\} W_i^{(K)} &\in [-C_8 \eta_0 \exp(\beta b_{\text{NPO},i}^{(K)}), -C_7 \eta_0 \exp(\beta b_{\text{NPO},i}^{(K)})] \text{ when } y_i = 1, \\ -\eta_0 \cdot \gamma_i^\top \text{diag}\{\Delta_{\text{NPO}}^{(K)}\} W_i^{(K)} &\in [C_7 \eta_0 \exp(-\beta b_{\text{NPO},i}^{(K)}), C_8 \eta_0 \exp(\beta b_{\text{NPO},i}^{(K)})] \text{ when } y_i = 0. \end{aligned}$$

The first part of the induction assumption (for  $t = K + 1$ ) follows immediately as the sign of the gradient updates  $-\eta_0 \cdot \gamma_i^\top \text{diag}\{\Delta_{\text{NPO}}^{(K)}\} W_i^{(K)}$  are determined as above.

Note that  $b_l(K) \leq b_{\text{NPO},i}^{(K)} \leq b_u(K)$  by the induction assumption. Similarly, using the comparison theorem of ODEs, we obtain

$$b_l(K + \varepsilon) \leq b_{\text{NPO},i}^{(K+\varepsilon)} \leq b_u(K + \varepsilon)$$

for  $\varepsilon \in [0, 1]$  and  $b_{\text{NPO},i}^{(K+\varepsilon)} := \varepsilon b_{\text{NPO},i}^{(K+1)} + (1 - \varepsilon) b_{\text{NPO},i}^{(K)}$ . Choosing  $\varepsilon = 1$  gives the second part of the induction assumption for  $t = K + 1$ . The last part of the induction assumption for  $t = K + 1$  follows from the same algebra as in the  $t = 1$  case.

**Proof of Claim (12)** By definition

$$W_i^{(t)} = \frac{\pi_{\theta_{\text{NPO}}^{(t)}}^\beta(y_i | x_i)}{\pi_{\theta_{\text{NPO}}^{(t)}}^\beta(y_i | x_i) + \pi_{\text{ref}}^\beta(y_i | x_i)} = \frac{\text{pred}_i(b_{\text{NPO},i}^{(t)})^\beta}{\text{pred}_i(b_{\text{NPO},i}^{(t)})^\beta + \text{pred}_i(b_{\text{NPO},i}^{(0)})^\beta}.$$

When  $\text{pred}_i(b_{\text{NPO},i}^{(t)}) \leq \text{pred}_i(b_{\text{NPO},i}^{(0)})$ , we have

$$W_i^{(t)} \in \left[ \frac{\text{pred}_i(b_{\text{NPO},i}^{(t)})^\beta}{2\text{pred}_i(b_{\text{NPO},i}^{(0)})^\beta}, \frac{\text{pred}_i(b_{\text{NPO},i}^{(t)})^\beta}{\text{pred}_i(b_{\text{NPO},i}^{(0)})^\beta} \right] \in [C_l \cdot \text{pred}_i(b_{\text{NPO},i}^{(t)})^\beta, C_u \cdot \text{pred}_i(b_{\text{NPO},i}^{(t)})^\beta] \quad (14)$$

for some constants  $C_l, C_u > 0$  depending only on  $(B_\theta, b_x, B_x, \beta)$ . Note that  $\text{pred}_i(b_{\text{NPO},i}^{(t)}) \leq \text{pred}_i(b_{\text{NPO},i}^{(0)})$  is equivalent to  $(1 - 2y_i)b_{\text{NPO},i}^{(t)} \geq 0$ . Therefore, under this condition, we have

$$\begin{aligned} \text{pred}_i(b_{\text{NPO},i}^{(t)}) &= \frac{1}{1 + \exp\left((1 - 2y_i)c_{\text{init},i} + (1 - 2y_i)b_{\text{NPO},i}^{(t)}\right)} \\ &= \frac{\exp\left((2y_i - 1)b_{\text{NPO},i}^{(t)}\right)}{\exp\left((2y_i - 1)b_{\text{NPO},i}^{(t)}\right) + \exp\left((1 - 2y_i)c_{\text{init},i}\right)} \\ &\in \left[ \frac{\exp\left((2y_i - 1)b_{\text{NPO},i}^{(t)}\right)}{1 + \exp\left((1 - 2y_i)c_{\text{init},i}\right)}, \frac{\exp\left((2y_i - 1)b_{\text{NPO},i}^{(t)}\right)}{\exp\left((1 - 2y_i)c_{\text{init},i}\right)} \right]. \end{aligned} \quad (15)$$

Putting Eq. (14) and (15) together and recalling that  $|c_{\text{init},i}| \leq B_\theta B_x$ , we obtain

$$W_i^{(t)} \in \left[ C_5 \cdot \exp\left((2y_i - 1)\beta b_{\text{NPO},i}^{(t)}\right), C_6 \cdot \exp\left((2y_i - 1)\beta b_{\text{NPO},i}^{(t)}\right) \right]$$

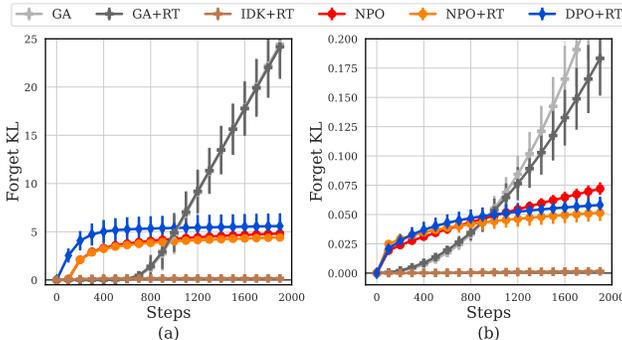
for some constants  $C_5, C_6 > 0$  depending only on  $(B_\theta, b_x, B_x, \beta)$ . This concludes the proof of Claim (12).

## C The role of KL Divergence on the Forget Set

In this section, we report *Forget KL*, the KL divergence between the output distributions of the initial model and the unlearned model on the forget set, defined as  $\mathbb{E}_{\mathcal{D}_{\text{FC}}} \text{KL}(\pi_{\text{ref}}(\cdot | x) || \pi_\theta(\cdot | x))$ . In experiments on the synthetic dataset and TOFU dataset, we observe that the models exhibit better unlearning performance when the forget KL is maintained at a moderate level. This suggests that explicitly maximizing the forget KL may not be an ideal objective for unlearning tasks.

### C.1 Synthetic Experiment

We present the forget KL for the synthetic experiments in Figure 10 (a) and (b). Combining with Figure 4, we find that the unlearned models attain Pareto frontiers when the forget KL is suitably large—an excessively large forget KL (as in GA, GA + RT after 1200 steps) or an excessively small forget KL (as in IDK + RT) may deteriorate the unlearning performance.



**Figure 10:** Forget KL versus optimization steps for all methods in the synthetic experiment. (a):  $\alpha = 1$ , (b):  $\alpha = 0$ . The errorbars denote  $\pm 1$  standard deviation over 5 runs.

## D Experimental details of the synthetic experiments

In this section, we discuss the experimental details of the synthetic experiments studied in Section 4.

**Initial model and retrained model.** We create the initial model  $\pi_{\text{ref}}$  and the retrained model  $\pi_{\text{retr}}$  via optimizing over  $\theta$  using the cross-entropy loss on the entire dataset  $\mathcal{D} = \mathcal{D}_{\text{FG}} \cup \mathcal{D}_{\text{RT}}$  and the retain dataset  $\mathcal{D}_{\text{RT}}$ , respectively. Concretely, initializing at  $\theta = \mathbf{0}_{128}$ , we run gradient descent for 20000 steps with the learning rate equals 0.05 to obtain the initial model  $\pi_{\text{ref}}$  and the retrained model  $\pi_{\text{retr}}$ .

**Unlearning.** During unlearning, starting from the initial model  $\pi_{\text{ref}}$ , we run gradient descent on each of the loss functions for 2000 steps with the learning rates selected via a grid search. We choose the learning rates so that the training remains stable within 2000 steps. It should be noted that variations of the learning rates may affect the number of steps needed to reach the minimal forget distance (or retain distance) in Figure 4 (a1, a2, b). However, they are less likely to alter the Pareto curves shown in Figure 4 (c). A grid search is also conducted to select the optimal  $\beta$  for NPO (and DPO)-based methods. The choices of learning rate and  $\beta$  are summarized in Table 1.

Method	learning rate		$\beta$	
	$\alpha = 1$	$\alpha = 0$	$\alpha = 1$	$\alpha = 0$
GA,GA+RT, IDK+RT	5e-4	1e-4	N/A	N/A
NPO, NPO+RT	5e-3	5e-2	1	10
DPO+RT	5e-3	5e-2	0.1	5

**Table 1:** Values of learning rate and  $\beta$  for different methods when  $\alpha = 1$  and  $\alpha = 0$  in the synthetic experiments.

## E Experiments on the TOFU dataset

In this section, we provide details of the experiments on the TOFU dataset (Maini et al., 2024). We first present a detailed explanation of the metrics, the baseline methods and the hyperparameters in the experiments. Then, we provide the full results on different levels of the tasks.

### E.1 Experiments Setup

#### E.1.1 Dataset

**TOFU Dataset.** We evaluate NPO-based methods and all baselines on the TOFU (Task of Fictitious Unlearning) dataset (Maini et al., 2024) designed for measuring the unlearning

methods for LLMs. TOFU contains 200 fictitious author profiles, each consisting of 20 question-answering pairs generated by GPT-4 based on a set of predefined attributes. These profiles are fictitious and do not exist in the pre-training data, providing a controlled environment for studying unlearning LLMs. TOFU introduces three levels of unlearning tasks, each aiming at forgetting a subset of 2, 10, and 20 authors (comprising 1%, 5%, and 10% of the training data, respectively), referred to as the *forget set*  $\mathcal{D}_{FG}$ , with a computational constraint that scales linearly with the size of the forget set. We refer to these tasks as Forget01, Forget05, and Forget10, respectively.

**Dataset for Evaluation.** In addition to the forget set, Maini et al. (2024) also introduced other datasets to measure the performance of the unlearned model. The *retain set*  $\mathcal{D}_{RT}$  is the part of the data that we do not hope the model to forget, which is, by definition, the complementary set of the forget set in the full dataset. To evaluate the model performance on the retain set, TOFU earmarks a subset of 400 question-answer pairs, accounting for 10% of the data, as an exclusive retain set that is never included in the forget set for any task. Moreover, to measure the general capacities of the unlearned models, two additional datasets are introduced: the Real Authors set and the Real World set. The Real Authors set includes question-answer pairs about authors in the real world and often deals with neighboring concepts entangled with those in the forget set. The Real World set contains commonsense knowledge about the real world and is designed to examine the general world knowledge of the unlearned model.

**Dataset beyond Forget10.** In scenarios where the targeted forget set exceeds 10% of the data (Forget20, Forget30, Forget50, and Forget90), we reorganize the original forget and retain sets within the TOFU dataset. To assess the Truth Ratio on an evaluation dataset, it is necessary to utilize the perturbed and paraphrased answers, which in TOFU were generated by properly prompting GPT4. To avoid any potential distribution shift from the newly crafted responses and their original forms, we evaluate the Truth Ratio using the publicly available data within TOFU. Consequently, even for tasks that are beyond forget10, we continue to use the data from the standard forget10 subset to compute the Truth Ratio on the forget set. This serves as a reasonable proxy for evaluating the Truth Ratio on the full forget set.

### E.1.2 Metrics

**Model Utility.** We measure the general capacities of the unlearned model using *Model Utility*, which aggregates multiple metrics across the retain set, the real-world set and the real-author set. Given a question-answer pair  $x = [q, a]$ , we compute the normalized conditional probability  $\mathbb{P}(a | q)^{1/|a|}$ , where  $|\cdot|$  denotes the number of tokens in a certain sequence. This probability is then averaged over the retain set, the Real Authors set, and the Real World set each. We also compute the averaged ROUGE-L recall score (Lin, 2004) across these datasets, a metric that evaluates the accuracy of the model’s response compared to the reference answers. Finally, we compute the averaged Truth Ratio on the three datasets above. The Truth Ratio defined in Maini et al. (2024) measures how likely the unlearned model will give a correct answer versus an incorrect one. More specifically, given a question  $q$ , Maini et al. (2024) generated a paraphrased (correct) answer  $\tilde{a}$  via prompting GPT4. They then generated five perturbed answers with the exactly same templates but incorrect answers  $\hat{a}_i, i = 1, 2, 3, 4, 5$  in the same way. The Truth Ratio is defined by

$$R_{truth} := \frac{\frac{1}{5} \sum_{i=1}^5 \mathbb{P}(\hat{a}_i | q)^{1/|\hat{a}_i|}}{\mathbb{P}(\tilde{a} | q)^{1/|\tilde{a}|}}. \quad (16)$$

The model utility is defined as the harmonic average of the nine metrics above (the probability, the ROUGE score, and the Truth Ratio on the retain set, the Real Authors set, and the Real World set).

**Forget Quality.** Forget quality assesses how well the unlearned model mimics the retrained model (defined as the model trained only on the retain set). This is a rigorous measurement as the ultimate goal for unlearning LLM is not only to stop generating the content related to the forget set but also to make the unlearned model indistinguishable from the retrained one. From a practical view of point, this requires the next-token probability given a prefix

of the unlearned model to be as close as possible to that of the retrained model. In TOFU, they compute the Truth Ratio (defined in Eq. 16) on each question-answer pair from the forget set. Instead of simply averaging them, they test whether the distribution of the Truth Ratio computed from the unlearned and the retrained models are indistinguishable. More specifically, they perform the Kolmogorov-Smirnov (KS) test and compute the p-value of the test. A large p-value indicates that the two models are indistinguishable from the Truth Ratio. When the p-value is above 0.05, we say the forgetting is significant.

### E.1.3 Baseline Methods

In this section, we introduce the baseline methods in our experiments.

**GA-based Methods.** Gradient Ascent (GA) is a key component in many LLM unlearning methods. Performing GA is equivalent to doing Gradient Descent (GD) on the negative cross-entropy loss function, which is denoted as  $\mathcal{L}_{GA}(\theta)$  defined in Eq. (1). Based on gradient ascent, a large class of unlearning methods performs gradient-based optimization on a linear combination of the GA loss  $\mathcal{L}_{GA}$  and several other loss functions that encourage unlearning (Jang et al., 2022; Yao et al., 2023; Chen & Yang, 2023; Maini et al., 2024; Eldan & Russinovich, 2023). Such a loss function can be written as

$$\mathcal{L}(\theta) = c_{GA}\mathcal{L}_{GA}(\theta) + c_{FG}\mathcal{L}_{FG}(\theta) + c_{RT}\mathcal{L}_{RT}(\theta) - c_{FGKL}\mathcal{K}_{FG}(\theta) + c_{RTKL}\mathcal{K}_{RT}(\theta), \quad (17)$$

where  $c_{GA}, c_{FG}, c_{RT}, c_{FGKL}, c_{RTKL}$  are non-negative weights. Here,  $\mathcal{L}_{FG}(\theta) = -\mathbb{E}_{\mathcal{D}_{FG}}[\log(\pi_{\theta}(\tilde{y}|x))]$  is the Forget loss where  $(x, y) \sim \mathcal{D}_{FG}$  and  $\tilde{y} \neq y$  is any response for prompt  $x$  which show some extent of ignorance towards the question  $x$ .  $\mathcal{L}_{RT}(\theta) = -\mathbb{E}_{\mathcal{D}_{RT}}[\log(\pi_{\theta}(y|x))]$  is the retain loss.  $\mathcal{K}_{FG}(\theta) = \mathbb{E}_{\mathcal{D}_{FG}}[D(\pi_{\theta}(\cdot|x)||\pi_{ref}(\cdot|x))]$  is the expected KL divergence on the forget set.  $\mathcal{K}_{RT}(\theta) = \mathbb{E}_{\mathcal{D}_{RT}}[D(\pi_{\theta}(\cdot|x)||\pi_{ref}(\cdot|x))]$  is the expected KL divergence on the retain set. In our experiments, we use three GA-based methods reported in Maini et al. (2024), referred to as GA, GA+RT, GA+KL, which fall in this class of loss function. The weights in Eq. (17) are shown in Table 2.

Loss	$c_{GA}$	$c_{FG}$	$c_{RT}$	$c_{FGKL}$	$c_{RTKL}$
GA	1	0	0	0	0
GA+RT	1	0	1	0	0
GA+KL	1	0	0	0	1
IDK+RT	0	1	1	0	0

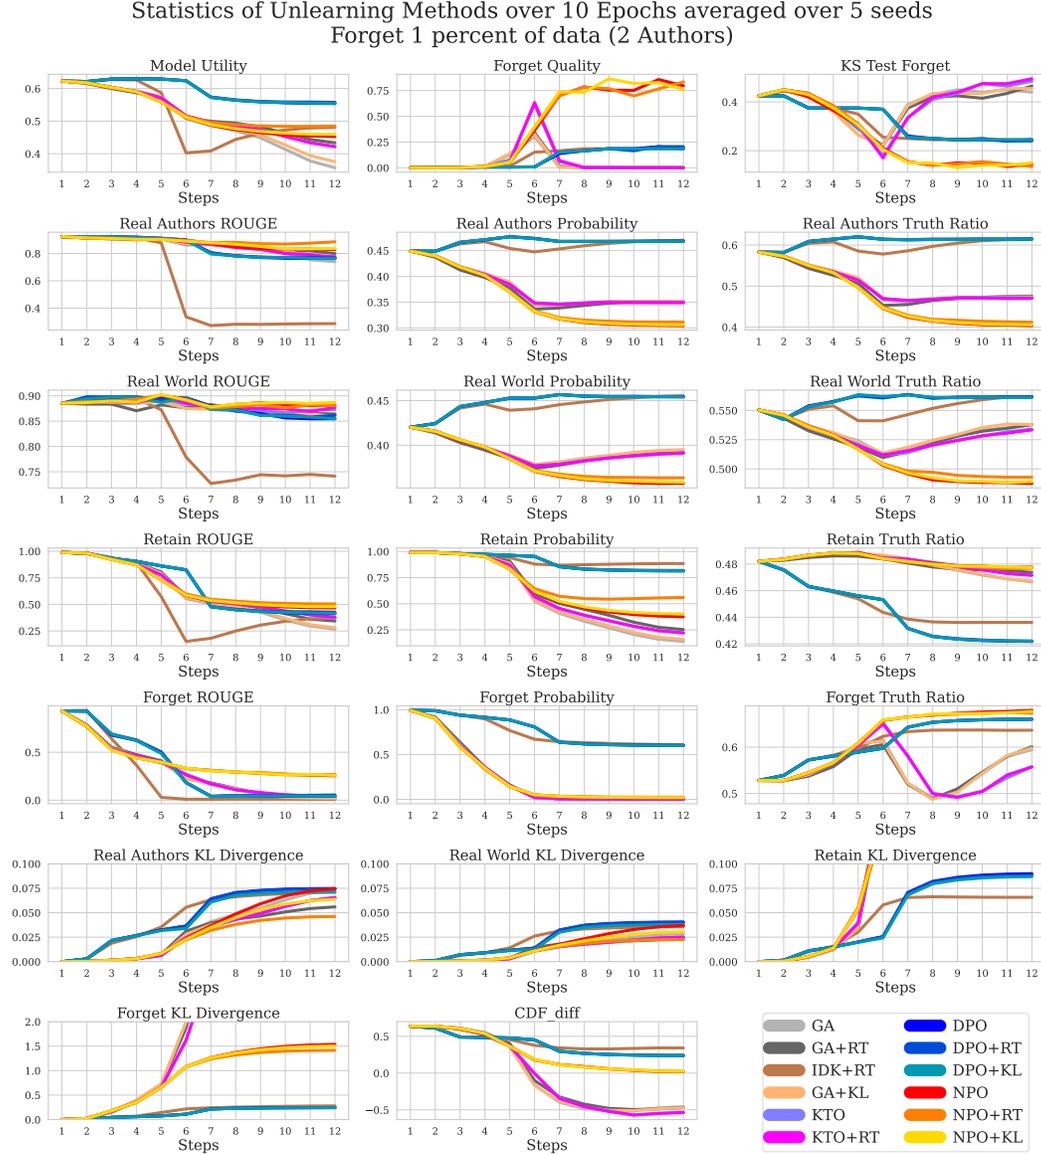
**Table 2:** The weights for different components in GA-based loss functions and IDK+RT loss.

**IDK-based Methods (‘I don’t know’).** Maini et al. (2024) proposed IDK+RT, which is a supervised loss function comprising of the retain loss and IDK loss term. The IDK loss term  $\mathcal{L}_{FG}$  is the averaged cross-entropy loss for question-answer pairs with questions  $x$  from the forget set  $\mathcal{D}_{FG}$  and answers  $y$  replaced by  $\tilde{y} = \text{‘I don’t know’}$  or a similar sentence showing ignorance towards this question. IDK+RT does not involve GA loss, and in general, IDK+RT loss shows a higher stability than GA-based methods.

**DPO-based Methods.** We also tested the DPO method (Rafailov et al., 2024) and its variants by adding either the retain loss or the KL divergence on the retain set. In the DPO loss, we take ‘I don’t know’ or its variants as positive responses and the answers in the forget set as negative responses. We use  $\beta = 0.1$  in all DPO-based experiments, which is commonly recognized as the optimal inverse temperature in most cases.

**KTO-based Methods.** We examine Kahneman-Tversky Optimization (KTO) (Ethayarajh et al., 2024), an alignment method with only non-paired preference data. The objective function of KTO is (we use a slightly different version than the original one as in Ethayarajh et al. (2024))

$$\mathcal{L}_{KTO} := \frac{2}{\beta} \mathbb{E}_{\mathcal{D}_{FG}} \left[ -\log \sigma \left( z_{ref} - \beta \log \frac{\pi_{\theta}(y|x)}{\pi_{ref}(y|x)} \right) \right], \quad (18)$$



where  $\beta > 0$  is the inverse-temperature,  $\sigma$  is the sigmoid function, and

$$z_{\text{ref}} := \mathbb{E}_{x \sim \mathcal{D}_{\text{FG}}} [\beta \cdot D(\pi_{\theta}(\cdot | x) || \pi_{\text{ref}}(\cdot | x))]. \quad (19)$$

Following (Ethayarajh et al., 2024), we estimate the KL term via averaged log probability ratio for questions in the forget set and answers in the “I don’t know” set (as the unrelated outputs). We examine both KTO and KTO+RT in our experiments with  $\beta = 0.1$ .

## E.2 Full Results

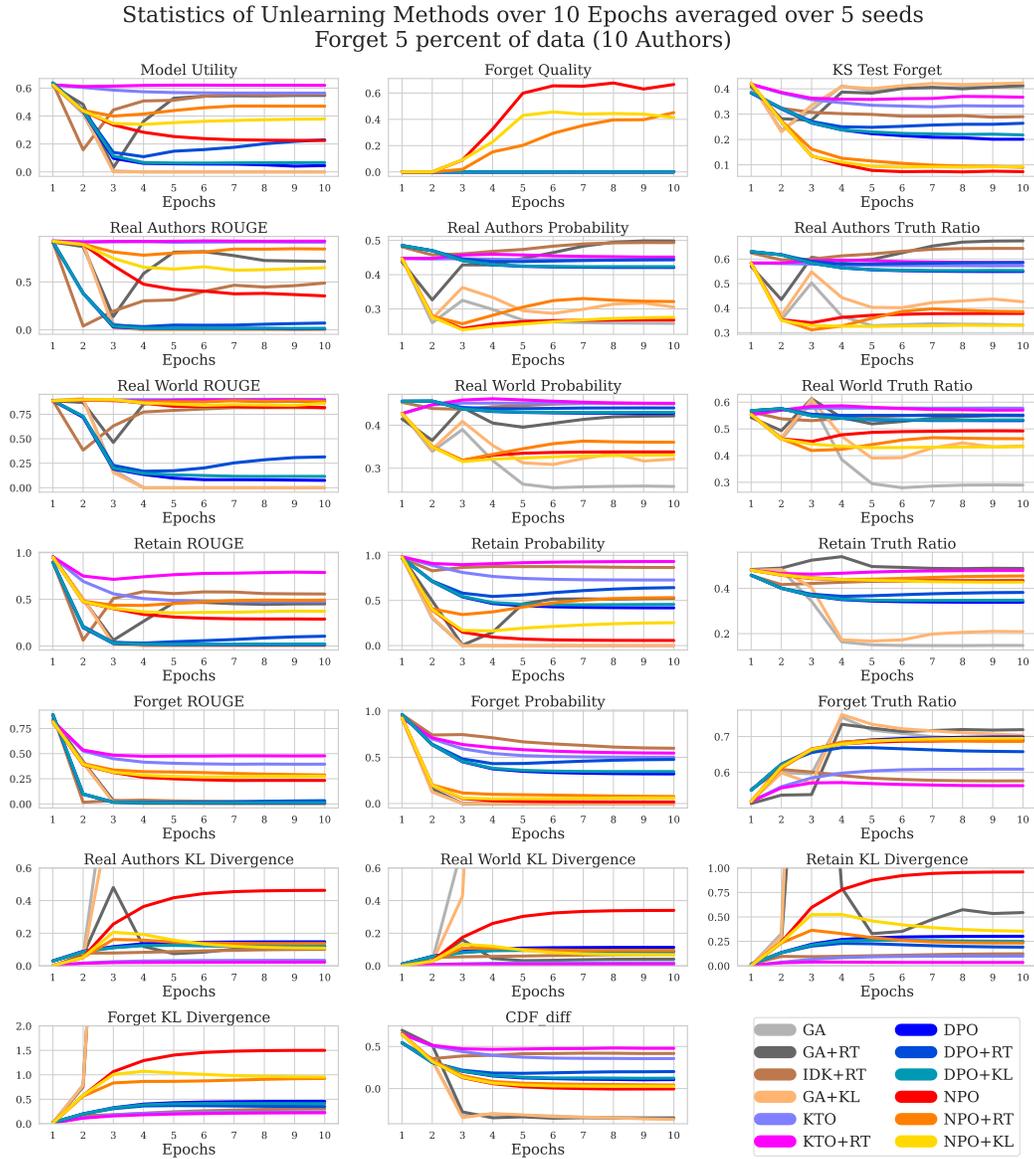


Figure 12: Statistics for NPO-based methods and baselines on the Forget05 task of TOFU.

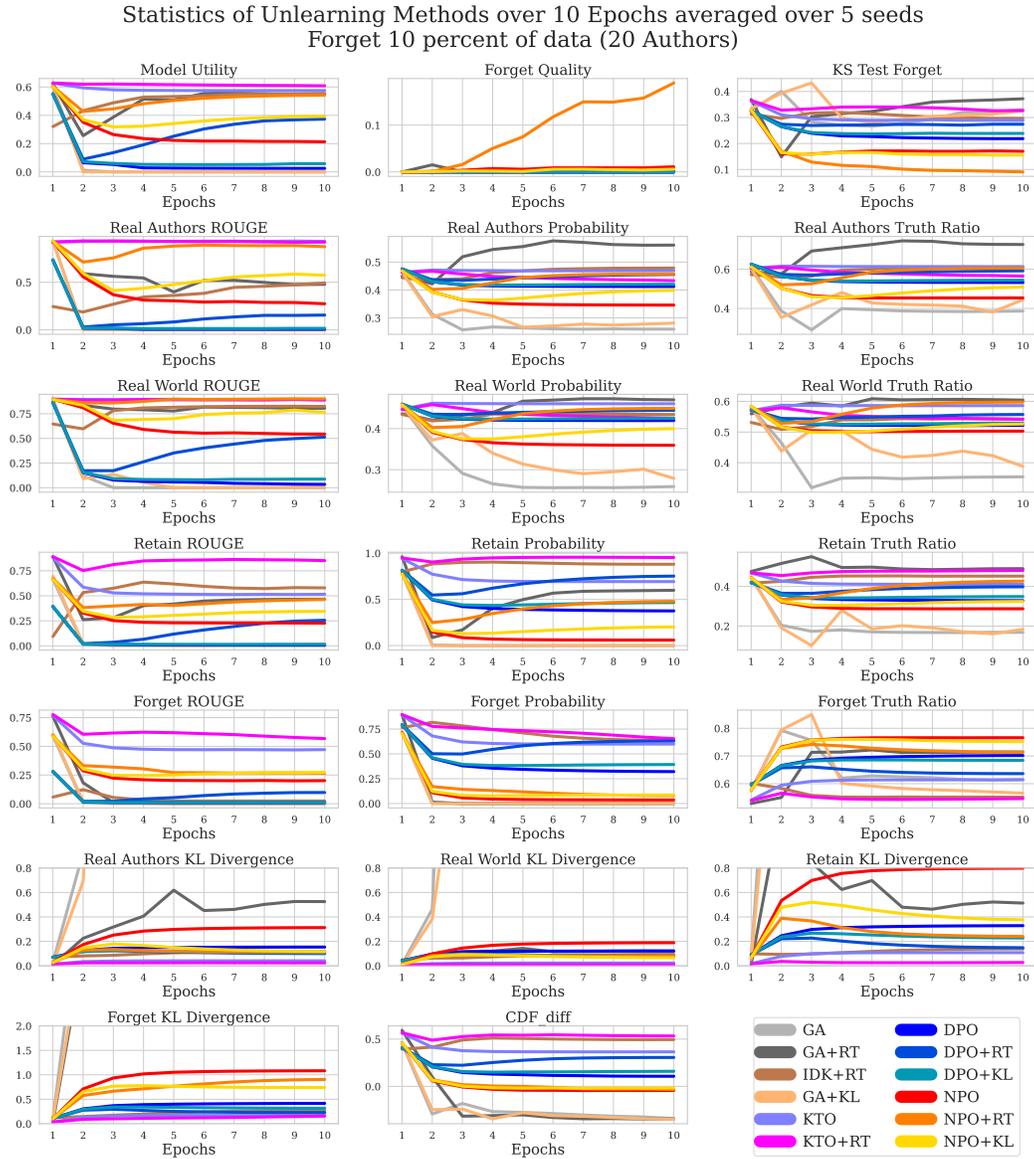


Figure 13: Statistics for NPO-based methods and baselines on the Forget10 task of TOFU.

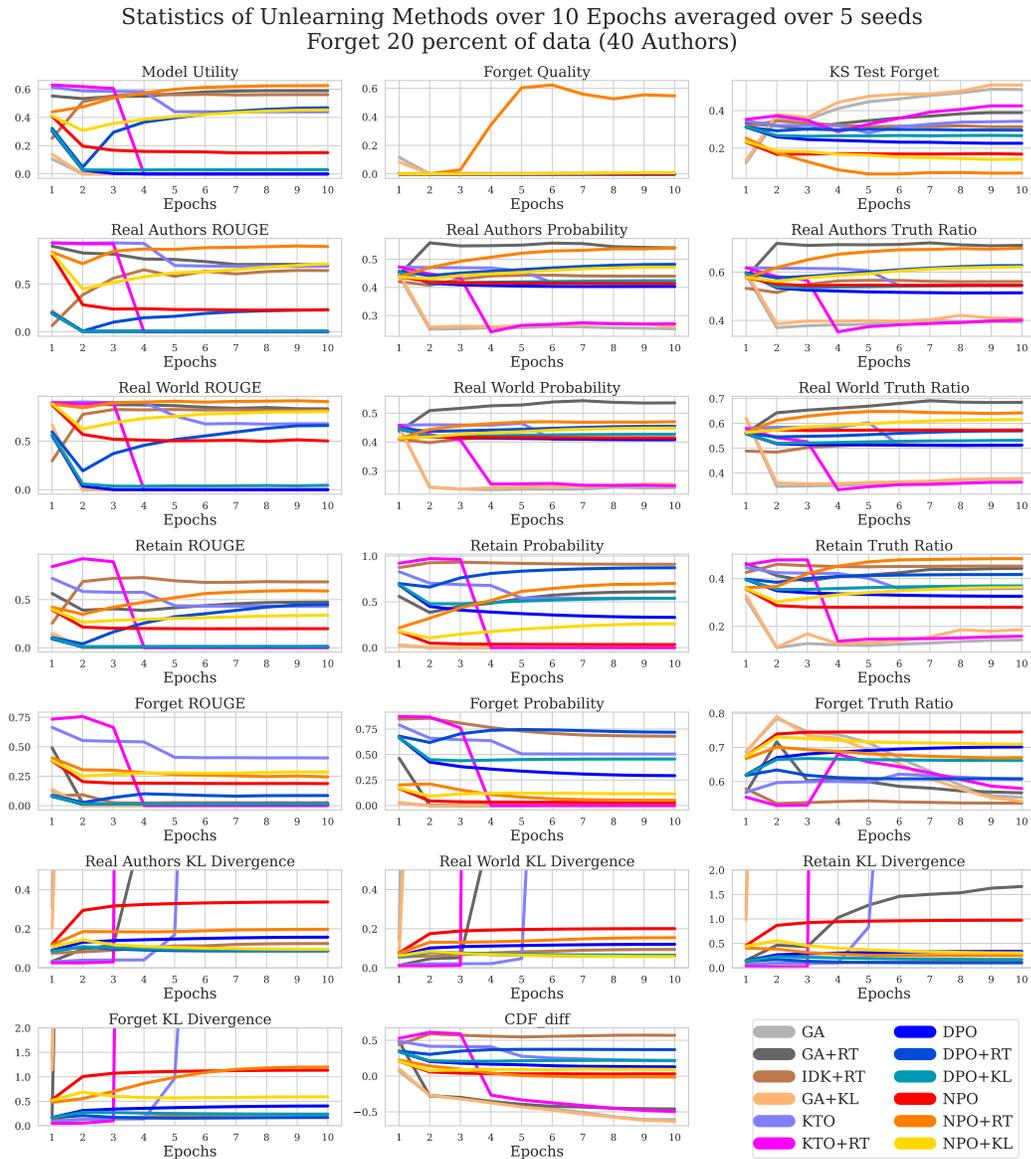


Figure 14: Statistics for NPO-based methods and baselines on the Forget20 task of TOFU.

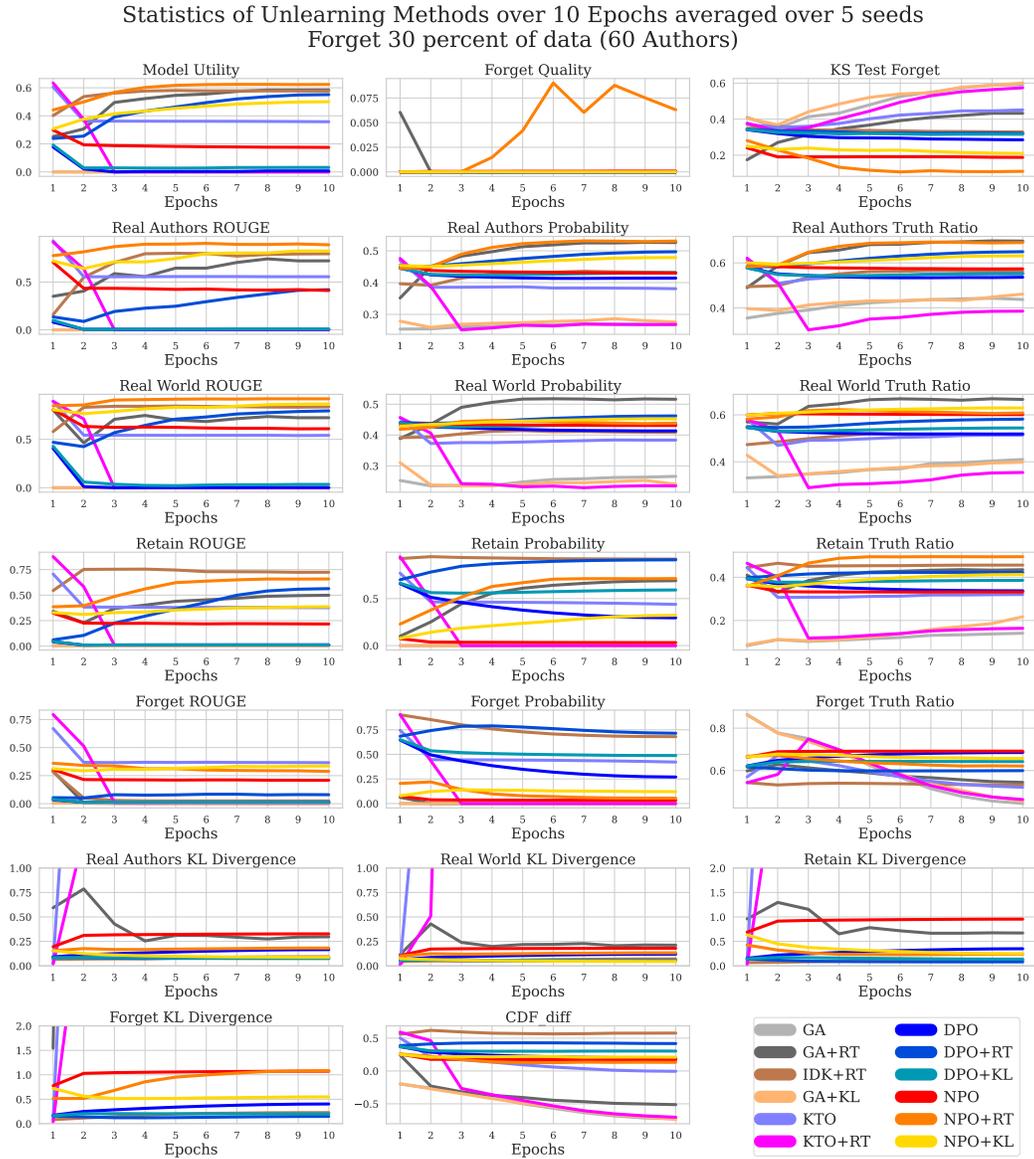


Figure 15: Statistics for NPO-based methods and baselines on the Forget30 task of TOFU.

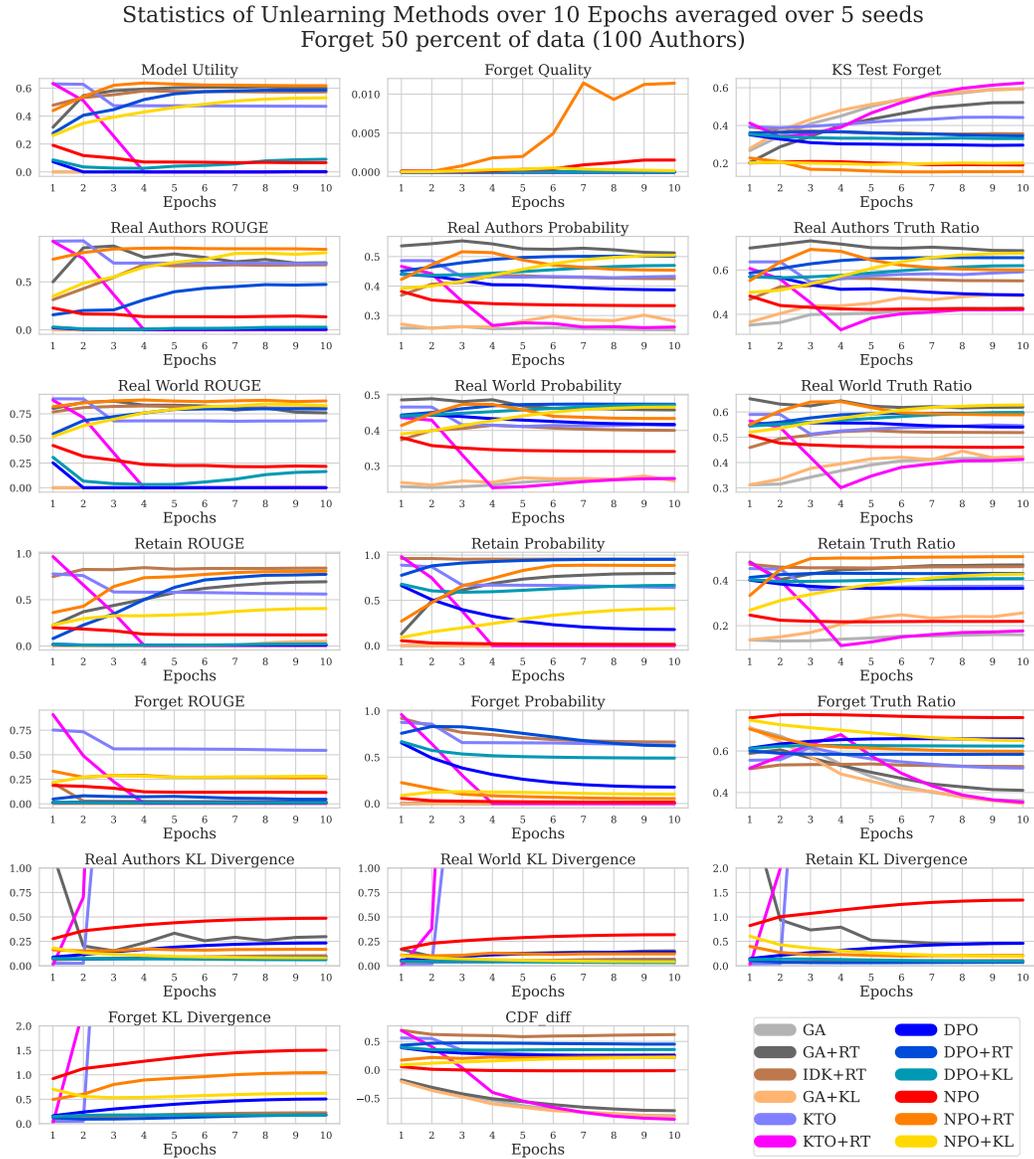


Figure 16: Statistics for NPO-based methods and baselines on the Forget50 task of TOFU.