

---

# Underestimated Privacy Risks for Minority Populations in Large Language Model Unlearning

---

Rongzhe Wei<sup>1</sup> Mufei Li<sup>1</sup> Mohsen Ghassemi<sup>2</sup> Eleonora Kreačić<sup>2</sup> Yifan Li<sup>3</sup> Xiang Yue<sup>4</sup> Bo Li<sup>5,6</sup>  
Vamsi K. Potluru<sup>2</sup> Pan Li<sup>1</sup> Eli Chien<sup>1</sup>

## Abstract

Large Language Models (LLMs) embed sensitive, human-generated data, prompting the need for unlearning methods. Although certified unlearning offers strong privacy guarantees, its restrictive assumptions make it unsuitable for LLMs, giving rise to various heuristic approaches typically assessed through empirical evaluations. These standard evaluations randomly select data for removal, apply unlearning techniques, and use membership inference attacks (MIAs) to compare unlearned models against models retrained without the removed data. However, to ensure robust privacy protections for every data point, it is essential to account for scenarios in which certain data subsets face elevated risks. Prior research suggests that outliers, particularly including data tied to minority groups, often exhibit higher memorization propensity which indicates they may be more difficult to unlearn. Building on these insights, we introduce a complementary, minority-aware evaluation framework to highlight blind spots in existing frameworks. We substantiate our findings with carefully designed experiments, using canaries with personally identifiable information (PII) to represent these minority subsets and demonstrate that they suffer at least 20% higher privacy leakage across various unlearning methods, MIAs, datasets, and LLM scales. Our proposed minority-aware evaluation framework<sup>1</sup> marks an essential step toward more equitable and comprehensive assessments of LLM unlearning efficacy.

---

<sup>1</sup>Georgia Institute of Technology <sup>2</sup>J.P. Morgan AI Research  
<sup>3</sup>Tsinghua University <sup>4</sup>Carnegie Mellon University <sup>5</sup>The University of Chicago <sup>6</sup>University of Illinois Urbana-Champaign. Correspondence to: Rongzhe Wei <rwei42@gatech.edu>, Pan Li <panli@gatech.edu>, Eli Chien <elichien@icloud.com>.

*Proceedings of the 42<sup>nd</sup> International Conference on Machine Learning*, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

<sup>1</sup>Code source: [https://github.com/Graph-COM/Minority\\_Aware\\_LLM\\_Unlearning.git](https://github.com/Graph-COM/Minority_Aware_LLM_Unlearning.git)

## 1. Introduction

Large Language Models (LLMs) are trained on vast and diverse datasets, often sourced from public content on the web, much of which is generated by humans (Touvron et al., 2023; Ouyang et al., 2022). This practice raises significant ethical concerns, particularly when the data includes sensitive information, leading to potential privacy violations. Individuals whose data has been used may seek to exercise their “right to be forgotten”, a protection guaranteed by regulations such as the General Data Protection Regulation (GDPR) (Krzysztofek, 2018).

The ideal approach to fulfilling such a request is to retrain the LLM from scratch, excluding the data to be removed. However, this solution is prohibitively expensive and impractical for large-scale models. To address this, the concept of *machine unlearning* has emerged as a promising alternative. Machine unlearning seeks to efficiently modify the LLM so that it becomes statistically indistinguishable from a model retrained from scratch. In this way, no adversary could confidently determine whether a model has undergone an unlearning process or been retrained, ensuring compliance with the “right to be forgotten”.

Unfortunately, it remains an open problem to enforce the formal unlearning guarantee for deep neural networks and LLMs without exact retraining. Despite recent progress in theoretical unlearning research (Guo et al., 2020; Sekhari et al., 2021; Neel et al., 2021; Ullah et al., 2021; Chien et al., 2023; Ullah & Arora, 2023; Chien et al., 2024a;b), their restrictive assumptions limit practical applicability to deep neural networks and LLMs. Concurrently, researchers have developed efficient unlearning heuristics and empirically evaluated their efficacy (Golatkar et al., 2020a;b; Graves et al., 2021; Liu et al., 2024a;c; Yao et al., 2024), often by comparing approximately unlearned models to those retrained from scratch (Pawelczyk et al., 2024a). Among the various evaluation methods, membership inference attacks (MIAs) (Shokri et al., 2017), originally developed to infer data usage during training, have been widely adopted for assessing unlearning performance (Shi et al., 2024b).

We identify a critical pitfall in the aforementioned LLM

unlearning efficacy evaluation. Prior work has demonstrated that unlearning difficulty can vary substantially across individual data points in vision tasks (Thudi et al., 2024; Zhao et al., 2024). Similarly, recent studies on LLMs also reveal highly non-uniform memorization patterns, where certain samples are memorized much more strongly than others (Feldman & Zhang, 2020; Carlini et al., 2022). However, current unlearning evaluation methods only capture “average-case” performance through random data removal from the training set. This approach inadequately addresses privacy risks for hard-to-unlearn data, failing to account for challenging scenarios necessitating rigorous privacy protection (Steinke & Ullman, 2020; Aerni et al., 2024). It neglects the principle that every individual’s right to be forgotten should be upheld equally, thus ignoring data from minority groups, which are often treated as outliers and can be more resistant to unlearning due to the aforementioned stronger memorization effects (Carlini et al., 2022; Nasr et al., 2021; 2023). Consequently, standard unlearning evaluation significantly underestimates privacy risks for these groups, overlooking crucial social responsibilities in personal data protection.

**Contributions.** Motivated by the privacy auditing literature (Jagielski et al., 2020; Steinke et al., 2024), we conduct a synthetic experiment on unlearning injected canaries pertaining to minority groups. We choose Personally Identifiable Information (PII) as a representative minority identifier, while noting that our approach extends to broader cases. We show that minorities suffer from at least 20% more privacy leakage in most cases across combinations of six unlearning approaches, three MIA variants, three datasets, and two LLMs of different scales. These results underscore the prevalence of the issue in practical settings, highlighting the need for a more effective LLM unlearning evaluation, particularly in regard to privacy risks for minority groups. Accordingly, we propose a *minority-aware LLM unlearning evaluation protocol* (Figure 1) as an initial step toward this goal. With this minority-aware protocol, we benchmark existing unlearning approaches and investigate the effects of forget set size as well as unlearning complexity. This study provides a more holistic understanding of different LLM unlearning approaches for practitioners. Notably, we observe that Langevin Unlearning—the only approach incorporating noise—achieves a favorable privacy-utility trade-off compared to noiseless methods such as SCRUB and Gradient Ascent (GA), suggesting a potential crucial role of noise incorporation in effective unlearning. In summary, these insights underline the critical role of our minority-aware evaluation framework in advancing equitable assessments of unlearning efficacy across different methods.

## 2. Related Work

Privacy auditing is a fundamental yet challenging aspect of LLM unlearning due to the difficulty of distinguishing training samples effectively (Duan et al., 2024). Various privacy-related metrics, such as exposure (Carlini et al., 2019), mean reciprocal rank (Wu et al., 2023), extraction likelihood (Jang et al., 2022), and truth ratio (Maini et al., 2024), have been proposed to probe privacy leakage. Among these, MIAs remain one of the most crucial tools for evaluating machine unlearning methods (Liu et al., 2024b). Standard MIAs typically involve training numerous shadow models independently to empirically approximate the distribution (Carlini et al., 2022). This approach has also been adopted for LLM unlearning, as seen in the NeurIPS 2023 Machine Unlearning Challenge<sup>2</sup> (which compares the point-wise output distributions of multiple unlearned and retrained models to perform MIAs) and Kurmanji et al. (2024); Pawelczyk et al. (2024b). Hayes et al. (2024) further highlights the limitations of average-case evaluations and introduces a specialized per-sample MIA method for unlearning evaluation. Their approach focuses on unlearning a randomly selected subset of training data by training a series of shadow models and performing per-sample MIA using a likelihood ratio test under Gaussian fitting (Carlini et al., 2022).

A major downside of MIA approaches involving shadow models is their computational expense, as they require training a large number of LLMs independently (Liu et al., 2024b). To address this downside, another line of research compares the outputs of models using different statistical metrics without requiring shadow models (Zhang et al., 2024; Liu et al., 2024a;c; Yao et al., 2024; Li et al., 2024a), making these methods more computationally feasible (Maini et al., 2024). For instance, Shi et al. (2024b) measures privacy risk through the normalized AUC difference between unlearned and retrained models, using MIAs such as Min-K% (Shi et al., 2024a). However, these works typically select the forget set randomly from the training set, corresponding to an average-case evaluation. Our study highlights a critical limitation of this approach: the privacy risks of minority populations within the training set are severely underestimated because minority data are less likely to be selected in the unlearning evaluation pipeline. By focusing on minority-aware scenarios, our work provides a more comprehensive perspective on unlearning evaluation and privacy risks. In a related context, Zhao et al. (2024) explores how the memorization of image representations affects the difficulty of unlearning individual samples in vision tasks by introducing an entanglement-based metric. Our work further explores the systematic degradation in unlearning efficacy for minority subgroups during the fine-

<sup>2</sup>[https://unlearning-challenge.github.io/assets/data/Machine\\_Unlearning\\_Metric.pdf](https://unlearning-challenge.github.io/assets/data/Machine_Unlearning_Metric.pdf)

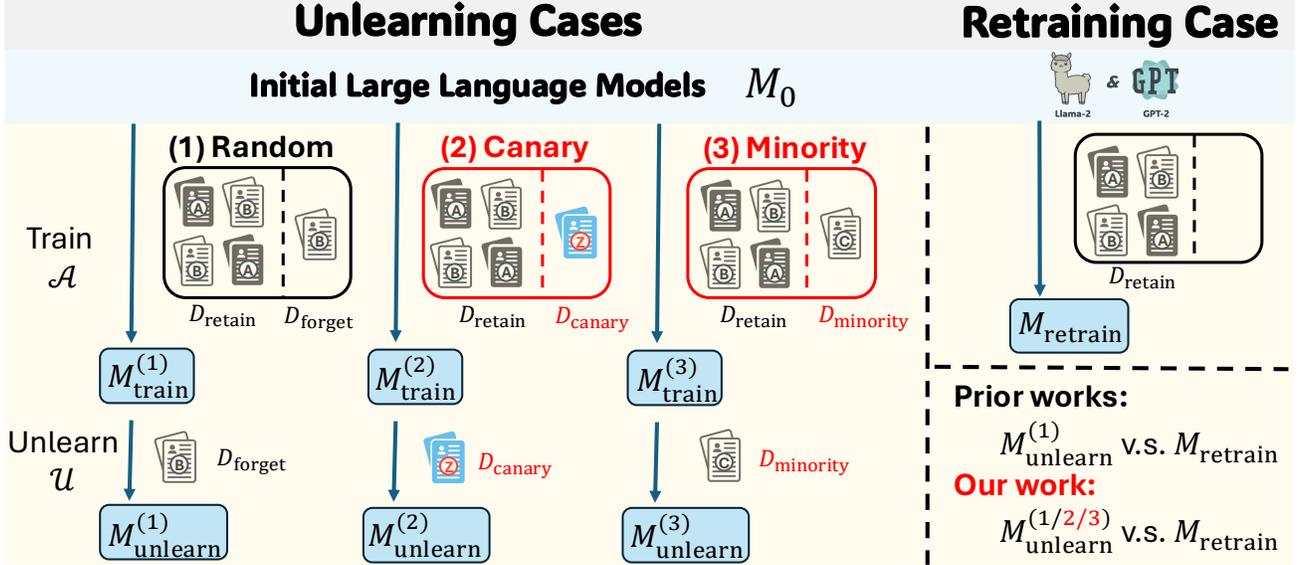


Figure 1. Illustration of our proposed LLM unlearning approaches (highlighted in red), when compared with the existing pipeline. Standard LLM unlearning evaluation typically involves randomly sampling data for removal from the training set (Case 1), which may underestimate privacy leakage for minority groups. In contrast, we design experiments to assess unlearning efficacy by removing canaries (deliberately inserted data points) related to minority groups (Case 2) and by directly removing data from minority groups (Case 3). Our approach provides a more comprehensive, minority-aware evaluation by considering the worst result across the three settings.

tuning of large language models. As a future direction, it would be valuable to examine the relationship between the entanglement-based difficulty metrics proposed by Zhao et al. (2024) and the minority-aware phenomena we identify in natural language contexts.

### 3. Preliminaries

Machine unlearning (Cao & Yang, 2015; Bourtole et al., 2021) has emerged as an important direction in trustworthy language models. It was initially motivated by privacy due to “the right to be forgotten” from GDPR and later extended to other legal and ethical concerns, including copyright (Yao et al., 2024), biased or outdated information mitigation (Liu et al., 2024b), hallucination removal (Yao et al., 2023), entity forgetting (Maini et al., 2024) and data poisoning removal (Pawelczyk et al., 2024a; Li et al., 2024b). In this work, we focus on the privacy aspect of the problem, albeit our methodology extends to other cases whenever the indistinguishability to the retrained model is an appropriate metric.

We briefly state the generic machine unlearning setting for privacy. Assume a training dataset  $D_{\text{train}}$  and a holdout test set  $D_{\text{test}}$  are given. Let  $M_{\text{learn}} \leftarrow \mathcal{A}(M_0, D_{\text{train}})$  be the language model trained on  $D_{\text{train}}$  starting from an initial model  $M_0$  via the training algorithm  $\mathcal{A}$ , which may be either a pre-trained language model or random initialization. Once the model is trained, we receive data removal requests that parti-

tion the training set  $D_{\text{train}} = D_{\text{forget}} \cup D_{\text{keep}}$  into a subset to be forgotten later  $D_{\text{forget}}$  and a keep set  $D_{\text{keep}}$ . An unlearning algorithm  $\mathcal{U}$  takes  $M_{\text{learn}}$ ,  $D_{\text{forget}}$  and  $D_{\text{train}}$  as input to return an updated model  $M_{\text{unlearn}} \leftarrow \mathcal{U}(M_{\text{learn}}, D_{\text{forget}}, D_{\text{train}})$ . It is worth noting that  $M_{\text{unlearn}}$  depends on the choice of  $D_{\text{forget}}$ . The gold standard to adhere to “the right to be forgotten” is retraining without  $D_{\text{forget}}$ , namely  $M_{\text{retrain}} \leftarrow \mathcal{A}(M_0, D_{\text{train}} \setminus D_{\text{forget}})$ . We say  $\mathcal{U}$  achieves good unlearning efficacy if  $M_{\text{unlearn}}$  and  $M_{\text{retrain}}$  are indistinguishable in their behavior  $m(M_{\text{unlearn}}, D) \approx m(M_{\text{retrain}}, D)$  on any corpus  $D$ , where  $m$  is any evaluation metric. Since  $M_{\text{unlearn}}$  and  $M_{\text{retrain}}$  depend on the choice  $D_{\text{forget}}$ , such approximation should be taken over the worst case ideally.

#### 3.1. Efficient MIAs for LLM Unlearning

As previously discussed, the effectiveness of unlearning methods can be measured by the indistinguishability between the resulting unlearned models and an exactly retrained model. MIA is often leveraged to determine whether a specific sample is part of the training set and is widely applied to audit training data privacy leakage. Therefore, to evaluate the efficacy of an unlearning approach, we consider the **PrivLeak (PL)** metric (Shi et al., 2024b). First, we define the difference in AUC scores between the unlearned model  $M_{\text{unlearn}}$  and the retrained model  $M_{\text{retrain}}$  as  $\Delta\text{AUC} = \text{AUC}(M_{\text{unlearn}}; D_{\text{forget}}, D_{\text{test}}) - \text{AUC}(M_{\text{retrain}}; D_{\text{forget}}, D_{\text{test}})$ , where AUC is the AUC-ROC score of an MIA (Ye et al., 2022) that tries to discriminate samples from  $D_{\text{forget}}$  and

$D_{\text{test}}$  based on the output statistics (e.g. loss) of a given model  $M$ . Then, the **PL** metric is formulated as:

$$\text{PrivLeak (PL)} = \frac{\Delta\text{AUC}}{\text{AUC}(M_{\text{retrain}}; D_{\text{forget}}, D_{\text{test}})} \quad (1)$$

By normalizing the difference in AUC scores between  $M_{\text{unlearn}}$  and  $M_{\text{retrain}}$  using the AUC of  $M_{\text{retrain}}$ , the metric accounts for the inherent difficulty of distinguishing the forget and test sets. Note that for an effective unlearning method, the metric should be around zero since the behavior of  $M_{\text{unlearn}}, M_{\text{retrain}}$  are indistinguishable. A larger magnitude of the PL metric implies a greater amount of privacy information that has been leaked under the tested MIA. A positive value indicates that the sample has not been fully forgotten, as the attacker has a higher AUC for  $M_{\text{unlearn}}$  than  $M_{\text{retrain}}$ . Conversely, a negative metric value suggests over-forgetting, which still indicates that  $M_{\text{unlearn}}$  differs from  $M_{\text{retrain}}$  and thus cause privacy breaches. Finally, note that an effective unlearning solution should lead to a small PL metric for *any* choice of MIA. In this work, we consider three popular MIAs and report the corresponding PL metric.

- **lossMIA** (Yeom et al., 2018): Determines membership of a sample  $x$  for a model  $M$  based on its loss  $\ell(M; x)$ .
- **zlibMIA** (Carlini et al., 2021): Determines membership of a sample based on the sample loss normalized by its zlib compression size,  $\ell(M; x)/\text{zlib}(x)$ .
- **Min-K%** (Shi et al., 2023): Selects the lowest  $K\%$  of token likelihoods and leverages the corresponding negative log-likelihood for membership inference.

#### 4. The Underestimated Privacy Risk of Data Minorities

Recall that both the unlearned  $M_{\text{unlearn}} \leftarrow \mathcal{U}(M_{\text{learn}}, D_{\text{forget}}, D_{\text{train}})$  and retrained  $M_{\text{retrain}} \leftarrow \mathcal{A}(M_0, D_{\text{train}} \setminus D_{\text{forget}})$  language models depend on the choice of the forget set  $D_{\text{forget}}$ . Whenever we estimate the privacy leakage of an unlearning method  $\mathcal{U}$  via some evaluation  $m$ , it is important to account for potential high-risk partitions of  $D_{\text{forget}}$  to ensure a comprehensive assessment of privacy risk. Unfortunately, the current LLM unlearning evaluation pipeline overlooks this critical aspect, where the partition leading to  $D_{\text{forget}}$  is chosen uniformly at random (Jang et al., 2023; Chen & Yang, 2023; Yao et al., 2024; Maini et al., 2024; Zhang et al., 2024; Shi et al., 2024b). The reported privacy risk therein hence corresponds to the ‘‘average case’’, which may significantly underestimate the privacy risk of highly privacy-sensitive points that request unlearning. It is known in the privacy literature that some rare training samples (minorities) may have an outsized effect on model memorization compared

to common training samples (majorities) (Feldman & Zhang, 2020; Carlini et al., 2022). Intuitively, a similar phenomenon persists for unlearning.

Here we utilize the Enron dataset as a case study. This dataset comprises 535,703 authentic emails from 158 employees of the Enron Corporation. It is a standard benchmark dataset for studying PII leakage, where the phone number is one form of PII that has been extensively studied (Lukas et al., 2023). The phone numbers here follow the format of the U.S. phone numbers (e.g., 123-456-7890), with the first three digits serving as the area code, representing the location where the number holder applied for the number. Such information is considered sensitive as it leaks not only the phone number itself, but also the geographic information pertaining to the number holder.

Table 1 illustrates the least frequent and three most frequent area codes in the Enron dataset. The area code distribution is far from uniform. Consequently, if emails containing phone numbers are uniformly sampled for the forget set  $D_{\text{forget}}$ , minority data, such as emails with rare area codes like 484, are unlikely to be included due to their lower frequency. If unlearning minority data is inherently more challenging and results in greater privacy leakage, the existing evaluation pipeline may underestimate privacy risks for minorities.

Table 1. Top three most frequent and least frequent area codes within Enron dataset.

Area code	Count
713 (Houston)	135,307
800 (Toll-free)	11,902
212 (New York)	10,739
484 (Allentown)	1

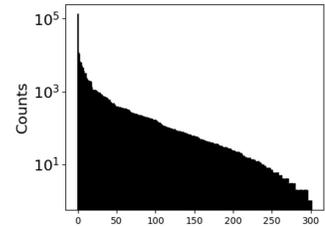


Figure 2. Area code histogram in Enron Dataset.

##### 4.1. Verify Underestimated Privacy Risks of Minority via Canary Injection

To rigorously show that *removing data from minority<sup>3</sup> populations indeed leads to higher unlearning privacy leakage*, we design experiments based on the idea of *canaries* in the privacy auditing literature (Jagielski et al., 2020; Steinke et al., 2024). For simplicity, we focus on the scenario where data removal requests pertain to PIIs, where each training sample  $x \in D_{\text{train}}^{(1)}$  consists of PIIs such as phone numbers or organizations. We choose PIIs as a representative minority identifier, albeit a similar idea extends beyond PIIs. We

<sup>3</sup>Minority here refers to any subset of the data, defined by some shared value, that is under-represented in the training set. This term is generic and may apply to any type of attribute, demographic or otherwise.

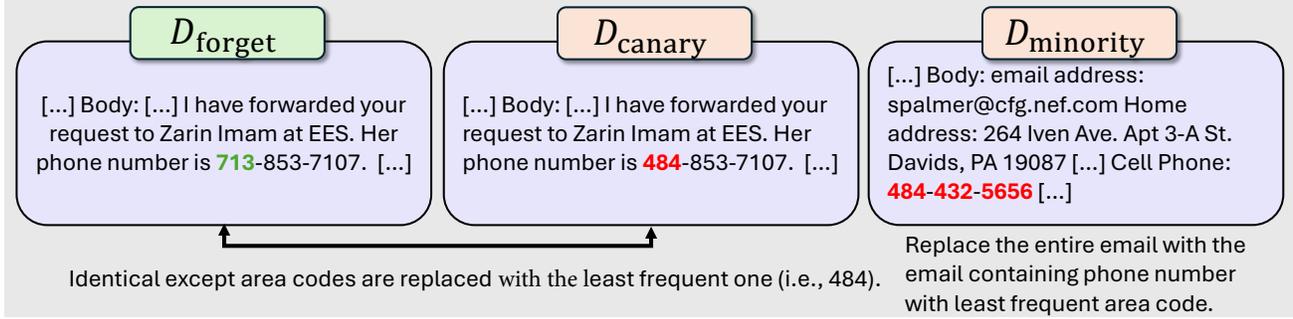


Figure 3. Illustration of the forget set  $D_{\text{forget}}$ , the construction of the canary set  $D_{\text{canary}}$ , and the minority set  $D_{\text{minority}}$  for the Enron dataset. The minority set consists of emails with phone numbers containing the least frequent area codes. A histogram and distribution of area codes (e.g., 713 as the most frequent and 484 as the least frequent) are shown in Figure 2 and Table 1.

consider the following cases, see Figure 1 for an illustration.

1) **Random**: we randomly partition  $D_{\text{train}}^{(1)} = D_{\text{forget}} \cup D_{\text{keep}}$  as in the standard unlearning evaluation pipeline. This leads to  $M_{\text{learn}}^{(1)} \leftarrow \mathcal{A}(M_0, D_{\text{train}}^{(1)})$ . 2) **Canary**: For the same forget set  $D_{\text{forget}} = \{x_i\}_{i=1}^n$ , we construct a canary set  $D_{\text{canary}} = \{x'_i\}_{i=1}^n$ , where each  $x'_i$  is identical to  $x_i$  except that only the PII is replaced by the least frequent one among  $D_{\text{train}}^{(1)}$ , isolating the impact of non-PII components. Finally, we construct a synthetic training set  $D_{\text{train}}^{(2)} = D_{\text{canary}} \cup D_{\text{keep}}$ , which leads to  $M_{\text{learn}}^{(2)} \leftarrow \mathcal{A}(M_0, D_{\text{train}}^{(2)})$ . By executing the same unlearning evaluation process for both cases  $M_{\text{learn}}^{(1)}$  and  $M_{\text{learn}}^{(2)}$ , we aim to show that the privacy risk for Canary is much higher than Random. By applying the same unlearning algorithm for removing  $D_{\text{forget}}$  and  $D_{\text{canary}}$ , we obtain the unlearned model  $M_{\text{unlearn}}^{(1)}$  and  $M_{\text{unlearn}}^{(2)}$  respectively. The privacy leakage (PL) is then computed for these cases as described in Section 3.1. The calculation of PL for Canary entails replacing  $D_{\text{forget}}$  with  $D_{\text{canary}}$  in Eq. (1). Note that the retrained model  $M_{\text{retrain}} \leftarrow \mathcal{A}(M_0, D_{\text{keep}})$  is identical for both scenarios.

An illustrative example of canary construction is provided in Figure 3. Note that for each email within  $D_{\text{forget}}$  in Random, we construct the corresponding canary by only replacing its area code with the least frequent one (i.e., 484). This design is critical as we ensure the other part of the email is identical to the original email. Hence, if the privacy leakage of Canary is greater than Random, it must be due to the difference in the area code. We repeat the similar canary construction for the other PII such as email domain and year of legal judgment for different datasets.

#### 4.2. Quantify the Underestimated Privacy Risk of Unlearning Minority

While our synthetic experiment on canary injection may be used to verify whether the unlearning privacy risk of minority populations is underestimated in the standard LLM

unlearning evaluation pipeline, it cannot quantify the privacy risk for minorities in the real-world setting. We further design the third case aiming at quantifying the amount of underestimated privacy risk by directly choosing data to be removed containing the least frequent PII. 3) **Minority**: construct a set  $D_{\text{minority}}$  that is of the same size as  $D_{\text{forget}}$  in Random, which consists of samples with the least frequent PII within the dataset. By comparing the computed privacy risk of Random and Minority, we can quantify the amount of underestimated privacy risk for data removal from minority groups compared to the average case. If the resulting privacy risk is significantly higher than Random, any conclusion pertaining to unlearning efficacy drawn from Random can be misleading and the right to be forgotten of minorities is overlooked.

## 5. Unlearning Methods

We evaluate the following popular unlearning approaches in the literature. With a slight abuse of notation, we denote  $M$  for both the model and its parameters for simplicity.

- **Random Labels (RL)** (Golatkar et al., 2020a; Yao et al., 2024): In the context of next-token prediction, this method involves randomly selecting tokens from the entire vocabulary during training on  $D_{\text{forget}}$ , aiming to disturb the model’s learning from this dataset. The intuition behind this approach is that a model uninformed by  $D_{\text{forget}}$  should behave as though it is randomly guessing the next token. However, as argued in Yao et al. (2024), this intuition may not always hold, depending on the specific scenario.
- **Exact Unlearning (EUK) and Catastrophic Forgetting (CFk)** (Goel et al., 2022): Exact unlearning can be done by retraining the entire model from scratch on  $D_{\text{keep}}$ , albeit it is prohibitively expensive in practice. Goel et al. (2022) proposes EUK method, which retrains only the last  $k$  layers of the model while freezing the other layers. As a result, it is computationally cheaper than retraining the

entire model. They also propose the CFk method, which continues training the last  $k$  layers on the  $D_{\text{keep}}$  without retraining from scratch, while freezing the other layers.

- **Gradient Ascent (GA)** (Golatkar et al., 2020a; Graves et al., 2021; Jang et al., 2023): Gradient ascent is arguably the most popular heuristic for machine unlearning. It seeks to remove the influence of the  $D_{\text{forget}}$  from the trained model by reversing the gradient updates associated with  $D_{\text{forget}}$ . Notably, researchers have reported that gradient ascent can lead to significant model utility degradation in some cases (Ilharco et al., 2023; Pawelczyk et al., 2024a).
- **NegGrad+** (Kurmanji et al., 2024): NegGrad+ is a combination of gradient ascent on  $D_{\text{forget}}$  and gradient descent on  $D_{\text{keep}}$ . It finetunes the current model by optimizing:

$$\beta \cdot \hat{\mathbb{E}}_{x \sim D_{\text{keep}}}[\ell(M; x)] - (1 - \beta) \hat{\mathbb{E}}_{x \sim D_{\text{forget}}}[\ell(M; x)]$$

where  $\beta \in (0, 1)$  is a hyperparameter and  $\hat{\mathbb{E}}$  is the empirical expectation. The intuition is to “review” the information from  $D_{\text{keep}}$  in order to prevent the model degradation due to the gradient ascent.

- **SCRUB** (Kurmanji et al., 2024): SCalable Remembering and Unlearning unBound (SCRUB) is a state-of-the-art unlearning method that leverages a student-teacher framework. It updates the model by optimizing the objective:

$$\hat{\mathbb{E}}_{x \sim D_{\text{keep}}}[\text{KL}(M_{\text{learn}}(x) \| M(x)) + \ell(M; x)] - \hat{\mathbb{E}}_{x \sim D_{\text{forget}}}[\text{KL}(M_{\text{learn}}(x) \| M(x))]$$

where KL is the Kullback-Leibler divergence. SCRUB shares a similar intuition with NegGrad+, which can also be viewed as a combination of gradient ascent on  $D_{\text{forget}}$  and descent on  $D_{\text{keep}}$ . Nevertheless, instead of directly employing the original loss  $\ell$ , SCRUB leverages the KL divergence to the original model  $M_{\text{learn}}$ . It provides a different regularization compared to NegGrad+.

- **Langevin Unlearning** (Chien et al., 2024a;b): Langevin Unlearning leverages noisy gradient descent for machine unlearning. Specifically, during the training process, it replaces the common gradient descent with DP-SGD (Abadi et al., 2016). For unlearning process, it finetunes the model on  $D_{\text{keep}}$  with DP-SGD as well. Chien et al. (2024a) establishes a smooth theoretical connection between differential privacy and unlearning and shows that Langevin Unlearning can provide a formal privacy guarantee for non-convex problems. Unfortunately, they mentioned that the resulting privacy bound is too loose to be applied in practice. We test Langevin Unlearning empirically in our experiments.

All the above unlearning methods fall under approximate unlearning (Thudi et al., 2022), valued for their practical efficiency. In contrast, exact unlearning methods, such as the

sharding-based framework SISA (Bourtole et al., 2021), demand significant computational and storage resources. SISA achieves exact unlearning by training multiple models independently on disjoint data partitions, but this deviates from standard machine learning workflows and introduces considerable memory overhead.

### 5.1. Enforcing the Same Computation Budget for Unlearning Methods

We categorize all methods as follows: those that only require the forget set (RL, GA), those that only require the keep set (EUK, CFk, Langevin), and those that require both the forget and keep sets (NegGrad+, SCRUB). Since machine unlearning is about the trade-off between privacy-utility-efficiency (Guo et al., 2020; Chien et al., 2024a; Liu et al., 2024c), we carefully ensure a similar computational complexity for all tested unlearning methods when demonstrating the privacy-utility trade-off. We define a **Complexity Unit** as the gradient computation budget of one training epoch on  $|D_{\text{forget}}|$  samples and limit all unlearning methods to a maximum of 10 Complexity Units. Since  $|D_{\text{forget}}| = U$  is roughly 1% of  $|D_{\text{train}}|$  throughout our experiments, all unlearning methods are indeed much more efficient than retraining from scratch (Pawelczyk et al., 2024a).

For unlearning approaches that leverage  $D_{\text{forget}}$  only, they can be tuned via unlearning process for at most 10 epochs. For those that leverage  $D_{\text{keep}}$  only, we randomly subsample it to size  $U$  for each epoch and unlearn for at most 10 epochs. For methods that leverage both  $D_{\text{forget}}$  and  $D_{\text{keep}}$  simultaneously, we limit their maximum unlearning epoch to 5. The situation is slightly more complicated for EUK and CFk approaches since only the last  $k$  layers are trained to save computation. We randomly select  $U/r$  samples from the keep set in each epoch, where  $r$  is the ratio of trainable parameters in the last  $k$  layers compared to the total number of parameters in the model. Our setup ensures that all tested unlearning approaches exhibit a similar unlearning computational complexity for a fair comparison. We optimize the unlearning epoch for each method under the 10 Complexity Unit constraint by the following criterion: if the perplexity of the unlearned model on  $D_{\text{train}}$  increases by more than 1 point compared to that of the initial model (No Unlearn), we stop at the first epoch where this condition is met; otherwise, we use the checkpoint from the last epoch.

## 6. Experiments

**Datasets.** Our LLM unlearning evaluation is conducted on two representative PII datasets: **Enron** (Klimt & Yang, 2004) and **ECHR** (Chalkidis et al., 2019). The Enron dataset contains corporate emails released by the Federal Energy Regulatory Commission, while the ECHR dataset comprises legal case information from the European Court

Table 2. The privacy leakage (PL) for each unlearning method against different attackers for GPT-2 / Llama-2 7B on the Enron-Phone dataset. The number in the parenthesis is the excess ratio of PL magnitude for cases `Canary` and `Minority` compared to `Random`, where a larger PL magnitude implies a more severe underestimation of privacy leakage in the standard unlearning evaluation (`Random`). Bold font indicates the case that the amount of underestimated privacy leakage is at least 20%.

Method	PL (lossMIA)			PL (zlibMIA)			PL (Min-K%)		
	Random	Canary	Minority	Random	Canary	Minority	Random	Canary	Minority
<b>Enron-Phone Dataset / GPT-2</b>									
No Unlearn	0.190	<b>0.283 (49%↑)</b>	<b>0.340 (79%↑)</b>	0.052	<b>0.076 (48%↑)</b>	<b>0.064 (24%↑)</b>	0.300	<b>0.447 (49%↑)</b>	<b>0.524 (75%↑)</b>
RL	0.118	<b>0.191 (61%↑)</b>	<b>0.210 (77%↑)</b>	0.044	<b>0.067 (52%↑)</b>	<b>0.060 (37%↑)</b>	0.258	<b>0.401 (55%↑)</b>	<b>0.447 (73%↑)</b>
EUK	0.027	<b>0.080 (198%↑)</b>	<b>0.124 (362%↑)</b>	0.035	<b>0.051 (47%↑)</b>	<b>0.052 (49%↑)</b>	0.092	<b>0.215 (134%↑)</b>	<b>0.223 (143%↑)</b>
CFk	0.190	<b>0.278 (46%↑)</b>	<b>0.337 (77%↑)</b>	0.053	<b>0.075 (41%↑)</b>	<b>0.064 (21%↑)</b>	0.298	<b>0.435 (46%↑)</b>	<b>0.514 (73%↑)</b>
GA	0.089	<b>0.140 (57%↑)</b>	<b>0.127 (42%↑)</b>	0.024	<b>0.042 (73%↑)</b>	0.026 (7%↑)	0.151	<b>0.242 (60%↑)</b>	0.171 (13%↑)
NegGrad+	0.183	<b>0.271 (48%↑)</b>	<b>0.327 (79%↑)</b>	0.052	<b>0.073 (42%↑)</b>	0.058 (13%↑)	0.293	<b>0.435 (48%↑)</b>	<b>0.511 (74%↑)</b>
SCRUB	0.167	<b>0.251 (50%↑)</b>	<b>0.321 (92%↑)</b>	0.048	<b>0.070 (44%↑)</b>	<b>0.062 (28%↑)</b>	0.295	<b>0.450 (52%↑)</b>	<b>0.527 (78%↑)</b>
Langevin	0.093	<b>0.144 (54%↑)</b>	<b>0.157 (69%↑)</b>	0.024	<b>0.037 (54%↑)</b>	0.027 (12%↑)	0.160	<b>0.258 (61%↑)</b>	<b>0.264 (65%↑)</b>
<b>Enron-Phone Dataset / Llama-2 7B</b>									
No Unlearn	0.060	<b>0.242 (303%↑)</b>	<b>0.172 (187%↑)</b>	0.034	<b>0.098 (188%↑)</b>	<b>0.067 (97%↑)</b>	0.076	<b>0.115 (51%↑)</b>	<b>0.179 (136%↑)</b>
RL	-0.242	-0.084 (65%↓)	-0.055 (77%↓)	-0.005	<b>0.065 (1400%↑)</b>	<b>0.102 (2140%↑)</b>	-0.123	-0.073 (41%↓)	0.012 (90%↓)
EUK	0.057	<b>0.246 (332%↑)</b>	<b>0.185 (225%↑)</b>	0.039	<b>0.106 (172%↑)</b>	<b>0.082 (110%↑)</b>	0.063	<b>0.132 (110%↑)</b>	<b>0.189 (200%↑)</b>
CFk	0.057	<b>0.236 (314%↑)</b>	<b>0.168 (195%↑)</b>	0.032	<b>0.094 (194%↑)</b>	<b>0.063 (97%↑)</b>	0.072	<b>0.108 (50%↑)</b>	<b>0.171 (138%↑)</b>
GA	-0.562	-0.430 (23%↓)	-0.464 (17%↓)	-0.014	<b>0.038 (371%↑)</b>	<b>0.083 (593%↑)</b>	-0.625	-0.459 (27%↓)	-0.517 (17%↓)
NegGrad+	-0.074	<b>-0.184 (149%↑)</b>	-0.040 (46%↓)	-0.021	<b>-0.048 (129%↑)</b>	-0.002 (90%↓)	-0.069	<b>-0.271 (293%↑)</b>	-0.057 (17%↓)
SCRUB	0.059	<b>0.162 (175%↑)</b>	<b>0.170 (188%↑)</b>	0.034	<b>0.063 (85%↑)</b>	<b>0.065 (91%↑)</b>	0.074	-0.031 (58%↓)	<b>0.177 (139%↑)</b>
Langevin	0.033	<b>0.180 (445%↑)</b>	<b>0.104 (215%↑)</b>	0.016	<b>0.068 (325%↑)</b>	<b>0.036 (125%↑)</b>	0.033	<b>0.055 (67%↑)</b>	<b>0.091 (176%↑)</b>

of Human Rights. For our experiments, we focus on specific PII based on their distributions: **phone numbers** (Enron-Phone) and **email domains** (Enron-Email) in Enron, and the **year of judgment** (ECHR-Year) in ECHR. Data minorities are defined based on these PII. Detailed dataset statistics are provided in App. A.1. Our study centers on instance-level unlearning, treating each individual as a single record.

**General Settings.** We focus on the fine-tuning scenario, where the initial model  $M_0$  is a pretrained LLM (GPT-2 (Radford et al., 2019) or Llama-2 7B (Touvron et al., 2023)). The fine-tuned model  $M_{learn}$  is obtained by training  $M_0$  on a dataset  $D_{train}$  for 5 epochs. In the GPT-2 experiments, both the training and test sets contain 10,000 samples, subsampled from the full dataset. For Llama-2, we employ efficient fine-tuning using LoRA (Hu et al., 2021); both the training and test sets consist of 50,000 samples, subsampled from the entire dataset. In all cases, the forget set size is set to 1% of the training set size. The models are optimized using the AdamW optimizer with a constant learning rate of  $10^{-5}$ , following the settings described in Shi et al. (2024b). During the unlearning process, all unlearning methods are constrained to the same computational budget—not exceeding 10 complexity units—as detailed in Section 5.1. We ensure that the unlearning complexity of each method is similar to allow for a fair comparison. Our ultimate goal is to achieve a superior privacy-utility-efficiency trade-off. We utilize MIA to estimate the empirical privacy risk measured by the PL metric as described in Section 3.1. For evaluating the utility of the LLMs, we report the perplexity following standard practices in the literature (Radford et al., 2019;

Zhang et al., 2022), where a lower perplexity indicates that the model is more confident in its predictions. Additional details are provided in App. B.

### 6.1. Standard Approaches Underestimate Privacy Risk for Minorities.

We report the results pertaining to the Enron-Phone, Enron-Email, and the ECHR-Year datasets. The experiment setting follows the explanation in Section 4 and further details are relegated to App. A.1. Table 2 shows that across all three attackers (lossMIA, zlibMIA, and Min-K%), all six unlearning methods and original model (no unlearning), the privacy leakage measure is significantly larger when unlearning canaries and minorities on Enron-Phone dataset for GPT-2 and Llama-2 7B respectively. Notably, in almost all cases the privacy leakage is underestimated for at least 20%. A similar phenomenon holds for the Enron-Email and ECHR-Year (Table 4, 5, 6, 7 in App. C.1) datasets. These results verify our claim that the current LLM unlearning evaluation indeed understated the privacy risk, especially for minorities. Our results call for a more careful empirical LLM unlearning evaluation, where considering canaries and minorities as we described can be an effective first step.

### 6.2. Benchmarking Unlearning Approaches under Minority-Aware Evaluation.

Motivated by our observations, we propose the minority-aware LLM unlearning evaluation. Instead of reporting the privacy leakage (PL) score under the Random case, we pro-

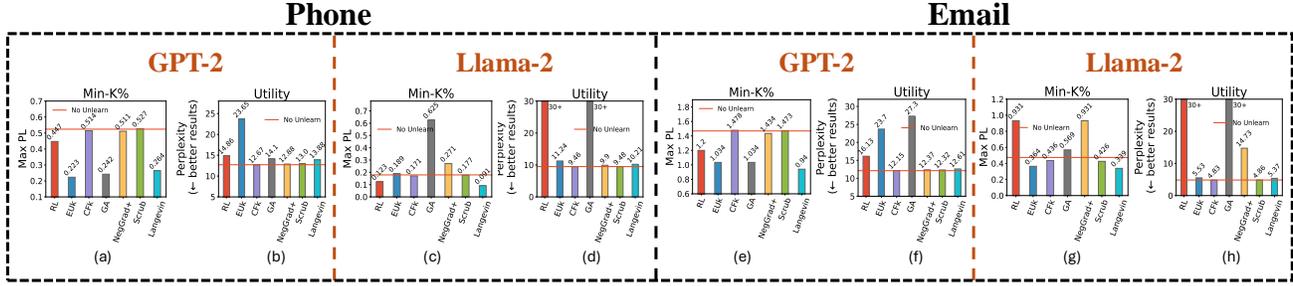


Figure 4. Benchmarking unlearning approaches via our minority-aware evaluation for GPT-2 and Llama-2 on Enron-Phone (Left) and Enron-Email (Right) datasets. (a),(c),(e),(g): Maximum privacy leakage (PL) over three cases (Random, Canary, and Minority) for Min-K% attack. (b),(d),(f),(h): Worst perplexity over the three cases of each method. More results on lossMIA and zlibMIA attackers are deferred to App. C.2.

pose to report the **magnitude of maximum PL score** of three settings (Random, Canary, and Minority). This provides a better privacy risk estimation while keeping the entire evaluation pipeline efficient. Besides, we report the corresponding worst-case perplexity as the utility measure for each unlearning approach. We benchmark the popular unlearning methods under our new evaluation pipeline, where the result is summarized in Figure 4 for GPT-2 and Llama-2 on the Enron-Phone and Enron-Email datasets. See App. C.2 for additional results.

We found that Langevin Unlearning offers the best balance between privacy and utility empirically. Note that while gradient ascent has on-par performance compared to Langevin Unlearning on the Enron-Phone dataset, it significantly degrades the model utility on the Enron-Email dataset. This echoes the finding of Ilharco et al. (2023); Pawelczyk et al. (2024a), albeit for different tasks. We found that gradient ascent is inherently unstable. In contrast, unlearning methods that leverage keep set  $D_{\text{keep}}$  are much more stable, including Langevin Unlearning and SCRUB.

We also present results using utility metrics, including BERTScore (Zhang et al., 2019) and ROUGE (Lin, 2004), which capture semantic meaning. These results, provided in App. C.3, exhibit a consistent trend.

### 6.3. Ablation studies.

We present ablation studies on the Enron-Phone dataset using the GPT-2 model, unless otherwise specified, and further results are deferred to App. C.4, C.5 and C.6.

**Privacy-Utility Trade-off.** We analyze the privacy-utility trade-off curves for stable methods like Langevin Unlearning and SCRUB, along with the widely used GA. For Langevin Unlearning, we adjust the noise scale during training and unlearning. In SCRUB, we vary the weights that balance the loss and KL regularizer terms in its objective function (Sec. 5). For GA, we explore different learning

rates ranging from  $1e-7$  to  $1e-3$ . As shown in Fig.6, these curves are evaluated on the Enron-Phone and Enron-Email datasets. The results indicate that Langevin Unlearning (Green Line) outperforms SCRUB (Purple Line) with a superior privacy-utility trade-off. Notably, while GA performs reasonably well on the Enron-Phone dataset, its trade-off on Enron-Email is significantly weaker, revealing its instability. Detailed hyperparameter tuning for these methods is provided in App.C.6.

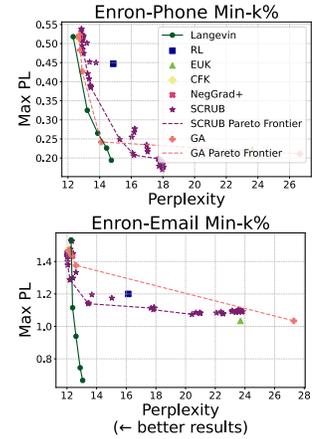
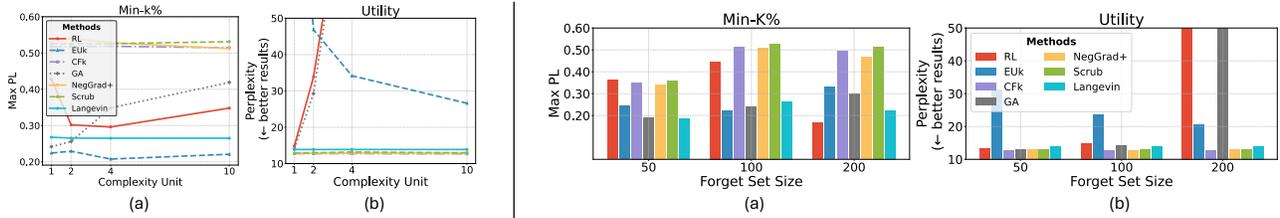


Figure 6. Privacy-utility Trade-off Curves for GPT-2.

**Unlearning Iteration.** We investigate the effect of unlearning epochs on privacy (Max PL) and utility (perplexity) for each unlearning approach in Fig. 5 (Left). We observe that RL and GA are unstable in PL score. Furthermore, these two methods can lead to significant model utility degradation in terms of perplexity, where even unlearning for 2 epochs can already result in a model breakdown. This observation again demonstrates that gradient ascent, albeit being simple and popular, is not a reliable LLM unlearning solution. We should focus on stable unlearning solutions such as SCRUB and Langevin Unlearning.

**Size of Forget Set.** In Fig. 5 (Right), we report the effect of different forget set sizes on the privacy (Max PL) and utility (Perplexity) trade-offs for each unlearning method. We find that both the GA and RL methods are highly sensitive to the forget set size, leading to significant model utility degradation and poor reliability in practice. In contrast, methods like Langevin Unlearning demonstrate good performance in



The effect of unlearning epochs for each unlearning approach. (a): Maximum PL with the attacker being Min-K%. (b): Model perplexity.

The effect of forget set size for each unlearning approach. (a): Maximum PL with the attacker being Min-K%. (Results on lossMIA, zlibMIA are deferred to App. C.5). (b): Model perplexity.

Figure 5. Ablation studies on unlearning iterations and forget set size.

terms of stability.

## 7. Conclusions

We identify a critical limitation in the typical evaluation pipeline for LLM unlearning efficacy: *privacy risks to minority groups in the training data are often underestimated*. Through carefully designed experiments using unlearning canaries tied to minority groups, inspired by privacy auditing research, we demonstrate that minority groups face at least 20% greater privacy leakage on average. Using personally identifiable information (PII) as a proxy for minority identifiers, we emphasize the need for more rigorous evaluations to ensure the right to be forgotten applies universally. Benchmarking existing unlearning methods with our minority-aware evaluation reveals that popular heuristics like gradient ascent are unstable and can degrade model utility. In contrast, methods such as Langevin Unlearning achieve a more favorable privacy-utility trade-off.

## Impact Statement

This work introduces a minority-aware evaluation framework aimed at enhancing the comprehensiveness and equitable nature of LLM unlearning assessments. Our research contributes to the standardization of user data protection and the responsible utilization of machine learning data. However, we emphasize the importance of using compliant, publicly available, or officially authorized data when conducting minority-aware evaluations with LLMs. Building on these foundations, our work aims to foster more ethical and inclusive practices in the development and application of machine learning technologies.

## Acknowledgement

This work is supported by NSF awards CCF-2402816, IIS-2239565, and JPMC faculty award 2023. The authors would like to thank Junyuan Hong and Chulin Xie for their assistance in reproducing the code for *LLM-PBE: Assessing Data Privacy in Large Language Models*. We are also grateful

to Martin Pawelczyk and Ayush Sekhari for their valuable clarifications on their work, and to Peter Kairouz for his insightful comments.

## Disclaimer

This paper was prepared for informational purposes in part by the Artificial Intelligence Research group of JPMorgan Chase & Co. and its affiliates (“JP Morgan”) and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

## References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS ’16*, pp. 308–318, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341394. doi: 10.1145/2976749.2978318. URL <https://doi.org/10.1145/2976749.2978318>.

Aerni, M., Zhang, J., and Tramèr, F. Evaluations of machine learning privacy defenses are misleading. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2024. URL <https://arxiv.org/abs/2404.17399>.

Bourtole, L., Chandrasekaran, V., Choquette-Choo, C. A., Jia, H., Travers, A., Zhang, B., Lie, D., and Papernot,

- N. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 141–159. IEEE, 2021.
- Cao, Y. and Yang, J. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pp. 463–480. IEEE, 2015.
- Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., and Song, D. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *28th USENIX security symposium (USENIX security 19)*, pp. 267–284, 2019.
- Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pp. 2633–2650, 2021.
- Carlini, N., Chien, S., Nasr, M., Song, S., Terzis, A., and Tramer, F. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 1897–1914. IEEE, 2022.
- Chalkidis, I., Androutsopoulos, I., and Aletras, N. Neural legal judgment prediction in english. *arXiv preprint arXiv:1906.02059*, 2019.
- Chen, J. and Yang, D. Unlearn what you want to forget: Efficient unlearning for LLMs. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. URL <https://openreview.net/forum?id=GprvtTwOxy>.
- Chien, E., Pan, C., and Milenkovic, O. Efficient model updates for approximate unlearning of graph-structured data. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=fhcu4FBLciL>.
- Chien, E., Wang, H., Chen, Z., and Li, P. Langevin unlearning: A new perspective of noisy gradient descent for machine unlearning. *arXiv preprint arXiv:2401.10371*, 2024a.
- Chien, E., Wang, H. P., Chen, Z., and Li, P. Certified machine unlearning via noisy stochastic gradient descent. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b. URL <https://openreview.net/forum?id=h3k2NXu5bJ>.
- Duan, M., Suri, A., Mireshghallah, N., Min, S., Shi, W., Zettlemoyer, L., Tsvetkov, Y., Choi, Y., Evans, D., and Hajishirzi, H. Do membership inference attacks work on large language models? *arXiv preprint arXiv:2402.07841*, 2024.
- Feldman, V. and Zhang, C. What neural networks memorize and why: Discovering the long tail via influence estimation. *Advances in Neural Information Processing Systems*, 33:2881–2891, 2020.
- Goel, S., Prabhu, A., Sanyal, A., Lim, S.-N., Torr, P., and Kumaraguru, P. Towards adversarial evaluations for inexact machine unlearning. *arXiv preprint arXiv:2201.06640*, 2022.
- Golatkar, A., Achille, A., and Soatto, S. 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, 2020a.
- Golatkar, A., Achille, A., and Soatto, S. Forgetting outside the box: Scrubbing deep networks of information accessible from input-output observations. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pp. 383–398. Springer, 2020b.
- Graves, L., Nagisetty, V., and Ganesh, V. Amnesiac machine learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11516–11524, 2021.
- Guo, C., Goldstein, T., Hannun, A., and Van Der Maaten, L. Certified data removal from machine learning models. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3832–3842. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/guo20c.html>.
- Hayes, J., Shumailov, I., Triantafillou, E., Khalifa, A., and Papernot, N. Inexact unlearning needs more careful evaluations to avoid a false sense of privacy. *arXiv preprint arXiv:2403.01218*, 2024.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Ilharco, G., Ribeiro, M. T., Wortsman, M., Schmidt, L., Hajishirzi, H., and Farhadi, A. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=6t0Kwf8-jrj>.
- Jagielski, M., Ullman, J., and Oprea, A. Auditing differentially private machine learning: How private is private sgd? *Advances in Neural Information Processing Systems*, 33:22205–22216, 2020.

- Jang, J., Yoon, D., Yang, S., Cha, S., Lee, M., Logeswaran, L., and Seo, M. Knowledge unlearning for mitigating privacy risks in language models. *arXiv preprint arXiv:2210.01504*, 2022.
- Jang, J., Yoon, D., Yang, S., Cha, S., Lee, M., Logeswaran, L., and Seo, M. Knowledge unlearning for mitigating privacy risks in language models. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14389–14408, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.805. URL <https://aclanthology.org/2023.acl-long.805>.
- Klimt, B. and Yang, Y. Introducing the enron corpus. In *CEAS*, volume 4, pp. 1, 2004.
- Krzysztofek, M. *GDPR: General Data Protection Regulation (EU) 2016/679: Post-reform Personal Data Protection in the European Union*. Kluwer Law International BV, 2018.
- Kurmanji, M., Triantafillou, P., Hayes, J., and Triantafillou, E. Towards unbounded machine unlearning. *Advances in neural information processing systems*, 36, 2024.
- Li, Q., Hong, J., Xie, C., Tan, J., Xin, R., Hou, J., Yin, X., Wang, Z., Hendrycks, D., Wang, Z., Li, B., He, B., and Song, D. Llm-pbe: Assessing data privacy in large language models. *Proc. VLDB Endow.*, 17(11):3201–3214, August 2024a. ISSN 2150-8097. doi: 10.14778/3681954.3681994. URL <https://doi.org/10.14778/3681954.3681994>.
- Li, W., Li, J., de Witt, C. S., Prabhu, A., and Sanyal, A. Delta-influence: Unlearning poisons via influence functions. *arXiv preprint arXiv:2411.13731*, 2024b.
- Lin, C.-Y. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pp. 74–81, 2004.
- Liu, J., Ram, P., Yao, Y., Liu, G., Liu, Y., SHARMA, P., Liu, S., et al. Model sparsity can simplify machine unlearning. *Advances in Neural Information Processing Systems*, 36, 2024a.
- Liu, S., Yao, Y., Jia, J., Casper, S., Baracaldo, N., Hase, P., Xu, X., Yao, Y., Li, H., Varshney, K. R., et al. Rethinking machine unlearning for large language models. *arXiv preprint arXiv:2402.08787*, 2024b.
- Liu, Z., Dou, G., Chien, E., Zhang, C., Tian, Y., and Zhu, Z. Breaking the trilemma of privacy, utility, and efficiency via controllable machine unlearning. In *Proceedings of the ACM on Web Conference 2024*, pp. 1260–1271, 2024c.
- Lukas, N., Salem, A., Sim, R., Tople, S., Wutschitz, L., and Zanella-Béguelin, S. Analyzing leakage of personally identifiable information in language models. In *2023 IEEE Symposium on Security and Privacy (SP)*, pp. 346–363. IEEE, 2023.
- Maini, P., Feng, Z., Schwarzschild, A., Lipton, Z. C., and Kolter, J. Z. Tofu: A task of fictitious unlearning for llms. *arXiv preprint arXiv:2401.06121*, 2024.
- Nasr, M., Songi, S., Thakurta, A., Papernot, N., and Carlin, N. Adversary instantiation: Lower bounds for differentially private machine learning. In *2021 IEEE Symposium on security and privacy (SP)*, pp. 866–882. IEEE, 2021.
- Nasr, M., Hayes, J., Steinke, T., Balle, B., Tramèr, F., Jagielski, M., Carlini, N., and Terzis, A. Tight auditing of differentially private machine learning. In *32nd USENIX Security Symposium (USENIX Security 23)*, pp. 1631–1648, 2023.
- Neel, S., Roth, A., and Sharifi-Malvajerdi, S. Descent-to-delete: Gradient-based methods for machine unlearning. In *Algorithmic Learning Theory*, pp. 931–962. PMLR, 2021.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Pawelczyk, M., Di, J. Z., Lu, Y., Kamath, G., Sekhari, A., and Neel, S. Machine unlearning fails to remove data poisoning attacks. *arXiv preprint arXiv:2406.17216*, 2024a.
- Pawelczyk, M., Neel, S., and Lakkaraju, H. In-context unlearning: Language models as few-shot unlearners. In *Forty-first International Conference on Machine Learning*, 2024b. URL <https://openreview.net/forum?id=GKcwl8XC9>.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Sekharia, A., Acharya, J., Kamath, G., and Suresh, A. T. Remember what you want to forget: Algorithms for machine unlearning. *Advances in Neural Information Processing Systems*, 34:18075–18086, 2021.
- Shi, W., Ajith, A., Xia, M., Huang, Y., Liu, D., Blevins, T., Chen, D., and Zettlemoyer, L. Detecting pretraining data from large language models. *arXiv preprint arXiv:2310.16789*, 2023.

- Shi, W., Ajith, A., Xia, M., Huang, Y., Liu, D., Blevins, T., Chen, D., and Zettlemoyer, L. Detecting pretraining data from large language models. In *The Twelfth International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=zWqr3MQUNs>.
- Shi, W., Lee, J., Huang, Y., Malladi, S., Zhao, J., Holtzman, A., Liu, D., Zettlemoyer, L., Smith, N. A., and Zhang, C. Muse: Machine unlearning six-way evaluation for language models. *arXiv preprint arXiv:2407.06460*, 2024b.
- Shokri, R., Stronati, M., Song, C., and Shmatikov, V. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pp. 3–18. IEEE, 2017.
- Steinke, T. and Ullman, J. The pitfalls of average-case differential privacy. *DifferentialPrivacy.org*, 07 2020. <https://differentialprivacy.org/average-case-dp/>.
- Steinke, T., Nasr, M., and Jagielski, M. Privacy auditing with one (1) training run. *Advances in Neural Information Processing Systems*, 36, 2024.
- Thudi, A., Jia, H., Shumailov, I., and Papernot, N. On the necessity of auditable algorithmic definitions for machine unlearning. In *31st USENIX Security Symposium (USENIX Security 22)*, pp. 4007–4022, 2022.
- Thudi, A., Jia, H., Meehan, C., Shumailov, I., and Papernot, N. Gradients look alike: Sensitivity is often overestimated in {DP-SGD}. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 973–990, 2024.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Ullah, E. and Arora, R. From adaptive query release to machine unlearning. In *International Conference on Machine Learning*, pp. 34642–34667. PMLR, 2023.
- Ullah, E., Mai, T., Rao, A., Rossi, R. A., and Arora, R. Machine unlearning via algorithmic stability. In *Conference on Learning Theory*, pp. 4126–4142. PMLR, 2021.
- Wu, X., Li, J., Xu, M., Dong, W., Wu, S., Bian, C., and Xiong, D. Depn: Detecting and editing privacy neurons in pretrained language models. *arXiv preprint arXiv:2310.20138*, 2023.
- Yao, J., Chien, E., Du, M., Niu, X., Wang, T., Cheng, Z., and Yue, X. Machine unlearning of pre-trained large language models. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8403–8419, Bangkok, Thailand, August 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.acl-long.457>.
- Yao, Y., Xu, X., and Liu, Y. Large language model unlearning. In *Socially Responsible Language Modelling Research*, 2023. URL <https://openreview.net/forum?id=wKe6jE065x>.
- Ye, J., Maddi, A., Murakonda, S. K., Bindschaedler, V., and Shokri, R. Enhanced membership inference attacks against machine learning models. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pp. 3093–3106, 2022.
- Yeom, S., Giacomelli, I., Fredrikson, M., and Jha, S. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st computer security foundations symposium (CSF)*, pp. 268–282. IEEE, 2018.
- Zhang, R., Lin, L., Bai, Y., and Mei, S. Negative preference optimization: From catastrophic collapse to effective unlearning. *arXiv preprint arXiv:2404.05868*, 2024.
- Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., Dewan, C., Diab, M., Li, X., Lin, X. V., et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., and Artzi, Y. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.
- Zhao, K., Kurmanji, M., Bărbulescu, G.-O., Triantafillou, E., and Triantafillou, P. What makes unlearning hard and what to do about it. *arXiv preprint arXiv:2406.01257*, 2024.

## A. Dataset

### A.1. Dataset Details

In this paper, we examine two representative PII datasets: Enron and ECHR, described as follows:

- **Enron (Klimt & Yang, 2004).** The Enron dataset consists of 536,000 authentic emails from 158 employees of the Enron Corporation, made publicly available by the Federal Energy Regulatory Commission following an investigation. Each email typically includes the sending timestamp, sender and recipient information, a greeting, the main content, and a footer containing the sender’s personal details.
- **ECHR (Chalkidis et al., 2019).** The ECHR dataset comprises case records from the European Court of Human Rights. Each record contains a series of factual lists that detail the specifics of a case. In our experiments, we further decompose these cases into individual facts, with each fact forming a distinct sample, averaging around 80 tokens in length. In total, the dataset includes around 118,000 samples.

### A.2. PII Selection

As outlined in Sec. 4, we selected U.S. phone numbers from the Enron dataset based on a criterion aimed at analyzing privacy risks in minority groups. To ensure the PII distribution was imbalanced, reflecting both minority and majority groups, we additionally selected two standard PII types (Lukas et al., 2023): email addresses (Enron) and years (ECHR). The distributions of these PII counts are depicted in Fig. 7.

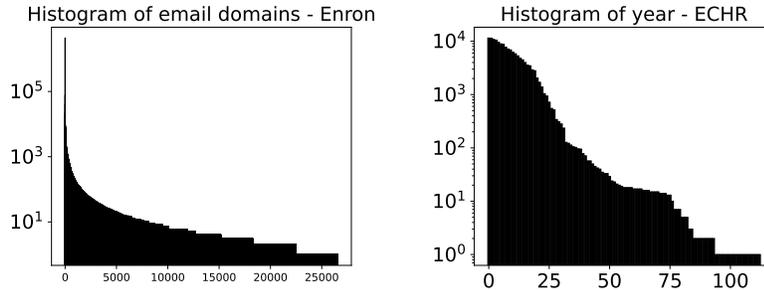


Figure 7. Histogram of email addresses (Enron-Email) and years (ECHR-Year).

### A.3. Preprocessing and Dataset Split

**Dataset Preprocess.** In our experiments, since the average token length of Enron samples is approximately 770, we controlled the token length of each fact to ensure the model could effectively memorize the samples. We randomly selected three coherent sentences from each sample, and if the sample contained specific PIIs of interest, we prioritized selecting sentences around them. We will keep the original samples for the ECHR dataset.

**Dataset Construction.** We begin by searching the dataset for occurrences of specific PIIs and analyzing their distribution. To form the minority set used in our `Minority` setting, we select 100 samples containing the least frequent PIIs; this set serves as our forget set. In the `Random` setting, we construct the forget set by randomly selecting 100 samples containing PIIs. To create the canary set, we replace the PIIs in the forget set (`Random` setting) with the least frequent PII found in the dataset. From the remaining data, we randomly select samples to create the training and test sets. For experiments with GPT-2 (117M), we uniformly at random selected 10,000 samples each for both the training and test sets. For Llama-2 7B, we uniformly at random selected 50,000 samples for both the training and test sets.

## B. Experimental Details

### B.1. Compute Configurations

All experiments were conducted using 8 NVIDIA A100 GPUs (80GB) and 14 NVIDIA RTX 6000 Ada GPUs (48GB).

## B.2. Unlearning Experiment Setup

**Unlearning Algorithms.** For all unlearning methods, we use a constant learning rate of  $10^{-5}$  and a batch size of 32, consistent with the fine-tuning stage. Note that some unlearning algorithms require additional hyperparameters. We follow the common designs from previous literature (Pawelczyk et al., 2024a) and detail the hyperparameter selection as follows:

- **EUK and CFK.** In our experiments, we set the number of retrained layers to  $k = 3$  for both GPT-2 and Llama-2 (LoRA) models. For GPT-2, the unfrozen trainable parameters account for approximately 16% of the total parameters, while for Llama-2 7B, the unfrozen parameters account for around 10%.
- **NegGrad+.** As noted in the main text, the hyperparameter  $\beta$  balances samples between  $D_{\text{forget}}$  and  $D_{\text{keep}}$ . In these experiments, we set  $\beta = 0.999$ .
- **SCRUB.** In the SCRUB method, three hyperparameters are used to balance the loss function on the keep set and the KL regularizers on both the keep and forget sets. According to the definition of the objective function in Section 5, three terms are weighted sequentially by setting:  $\alpha = 0.999$ ,  $\beta = 1$ , and  $\gamma = 0.01$ .
- **Langevin Unlearning.** The Langevin Unlearning method leverages noisy gradient descent to unlearn samples from the forget set. In our experiments, we set the Gaussian noise scale to  $\sigma = 5e - 4$  (for GPT-2) and  $\sigma = 5e - 3$  (for Llama-2), and the clipping norm to 1.

**Unlearning Epoch Selection.** As outlined in Section 5.1, all unlearning methods are constrained to a maximum of 10 complexity units, and the optimal epoch for each method is selected based on whether the perplexity of the unlearned model on  $D_{\text{train}}$  increases by more than 1 point. Under our computational budget, methods that only require the forget set (RL, GA) are run for 10 epochs, while methods requiring both the keep and forget sets (NegGrad+, SCRUB) are limited to 5 epochs, due to the equal-sized cycling between the two sets. For methods that only require the keep set (EUK, CFK, Langevin), we use 10 epochs, with varying sample sizes for EUK and CFK, as some model parameters remain frozen. The selected epoch for each method in each experiment is detailed in Table 3.

Table 3. Epochs comparison between unlearning methods on GPT2 and LLAMA2 models.

Unlearning Methods	GPT-2		Llama-2 7B	
	Enron	ECHR	Enron	ECHR
<b>RL</b>	Epoch 1	Epoch 1	Epoch 1	Epoch 1
<b>EUK</b>	Epoch 10	Epoch 10	Epoch 10	Epoch 10
<b>CFk</b>	Epoch 10	Epoch 10	Epoch 10	Epoch 10
<b>GA</b>	Epoch 1	Epoch 1	Epoch 1	Epoch 1
<b>NegGrad+</b>	Epoch 5	Epoch 5	Epoch 1	Epoch 5
<b>SCRUB</b>	Epoch 5	Epoch 5	Epoch 5	Epoch 5
<b>Langevin</b>	Epoch 10	Epoch 10	Epoch 10	Epoch 10

**Attack Method Hyperparameters.** We employed three attack methods in our evaluation pipeline. For lossMIA and zlibMIA, there are no hyperparameters to tune. The Min- $K\%$  method is based on the observation that non-member examples tend to have more tokens with lower likelihoods compared to member examples. In this method, the hyperparameter  $K$  controls the selection of the bottom  $K\%$  of tokens in each sample based on their likelihoods. Following previous recommendations in (Duan et al., 2024; Shi et al., 2024a), we set  $K = 20$  in our experiments.

**Random Seed Selection.** In all our experiments, we followed the common practice and fixed our random seed to be 42.

## C. Additional Experimental Results

In this section, we present supplementary experimental results to further substantiate our claims in the main text.

## Underestimated Privacy Risks for Minority Populations in Large Language Model Unlearning

Table 4. The privacy leakage (PL) for each unlearning method against different attackers for GPT-2 on the Enron-Email dataset.

Method	PL (lossMIA)			PL (zlibMIA)			PL (Min-K%)		
	Random	Canary	Minority	Random	Canary	Minority	Random	Canary	Minority
No Unlearn	0.303	<b>0.535 (77%↑)</b>	<b>1.145 (278%↑)</b>	0.200	<b>0.309 (54%↑)</b>	<b>0.262 (31%↑)</b>	0.529	<b>0.934 (76%↑)</b>	<b>1.468 (178%↑)</b>
RL	0.033	<b>0.153 (366%↑)</b>	<b>0.448 (1265%↑)</b>	0.062	<b>0.142 (127%↑)</b>	<b>0.121 (94%↑)</b>	0.431	<b>0.772 (79%↑)</b>	<b>1.200 (179%↑)</b>
EUK	0.232	<b>0.440 (89%↑)</b>	<b>0.582 (150%↑)</b>	0.135	<b>0.229 (70%↑)</b>	<b>0.152 (13%↑)</b>	0.501	<b>0.886 (77%↑)</b>	<b>1.034 (106%↑)</b>
CFk	0.296	<b>0.515 (74%↑)</b>	<b>1.139 (285%↑)</b>	0.197	<b>0.295 (49%↑)</b>	<b>0.260 (32%↑)</b>	0.526	<b>0.905 (72%↑)</b>	<b>1.478 (181%↑)</b>
GA	-0.279	-0.173 (38%↓)	<b>0.739 (165%↑)</b>	-0.119	-0.037 (69%↓)	<b>0.168 (41%↑)</b>	-0.390	-0.304 (22%↓)	<b>1.034 (165%↑)</b>
NegGrad+	0.265	<b>0.471 (77%↑)</b>	<b>1.103 (316%↑)</b>	0.179	<b>0.269 (50%↑)</b>	<b>0.251 (40%↑)</b>	0.496	<b>0.864 (74%↑)</b>	<b>1.434 (189%↑)</b>
SCRUB	0.286	<b>0.499 (74%↑)</b>	<b>1.097 (283%↑)</b>	0.190	<b>0.289 (52%↑)</b>	<b>0.253 (34%↑)</b>	0.519	<b>0.902 (74%↑)</b>	<b>1.473 (184%↑)</b>
Langevin	0.154	<b>0.319 (107%↑)</b>	<b>0.606 (293%↑)</b>	0.086	<b>0.178 (107%↑)</b>	<b>0.124 (44%↑)</b>	0.336	<b>0.645 (92%↑)</b>	<b>0.940 (180%↑)</b>

Table 5. The privacy leakage (PL) for each unlearning method against different attackers for GPT-2 on ECHR-year datasets.

Method	PL (lossMIA)			PL (zlibMIA)			PL (Min-K%)		
	Random	Canary	Minority	Random	Canary	Minority	Random	Canary	Minority
No Unlearn	0.198	<b>0.247 (25%↑)</b>	<b>0.263 (33%↑)</b>	0.086	<b>0.103 (20%↑)</b>	<b>0.122 (42%↑)</b>	0.213	<b>0.276 (30%↑)</b>	<b>0.299 (40%↑)</b>
RL	0.161	<b>0.213 (32%↑)</b>	<b>0.234 (45%↑)</b>	0.067	<b>0.088 (31%↑)</b>	<b>0.086 (28%↑)</b>	0.190	<b>0.259 (36%↑)</b>	<b>0.257 (35%↑)</b>
EUK	0.125	<b>0.176 (41%↑)</b>	0.138 (10%↑)	0.067	<b>0.088 (31%↑)</b>	0.070 (4%↑)	0.114	<b>0.187 (64%↑)</b>	0.135 (18%↑)
CFk	0.188	<b>0.234 (24%↑)</b>	<b>0.260 (38%↑)</b>	0.084	0.095 (13%↑)	<b>0.120 (43%↑)</b>	0.209	<b>0.264 (26%↑)</b>	<b>0.295 (41%↑)</b>
GA	0.067	0.027 (60%↓)	<b>0.105 (57%↑)</b>	0.024	0.019 (21%↓)	<b>0.038 (58%↑)</b>	0.090	-0.019 (79%↓)	<b>0.143 (59%↑)</b>
NegGrad+	0.183	<b>0.221 (21%↑)</b>	<b>0.247 (35%↑)</b>	0.071	<b>0.088 (24%↑)</b>	<b>0.112 (58%↑)</b>	0.191	<b>0.237 (24%↑)</b>	<b>0.274 (43%↑)</b>
SCRUB	0.179	<b>0.223 (25%↑)</b>	<b>0.253 (41%↑)</b>	0.080	<b>0.099 (24%↑)</b>	<b>0.116 (45%↑)</b>	0.197	<b>0.253 (38%↑)</b>	<b>0.289 (47%↑)</b>

### C.1. Experiments on Enron-Email and ECHR-year datasets

In Tables 4, 7, 5 and 7, we report the PL scores for all three attackers across the three scenarios on the Enron-email, ECHR-year datasets for GPT-2 and Llama-2, respectively. The results support our claim that the current LLM unlearning evaluation (Random setting) significantly underestimates privacy risk.

Table 6. The privacy leakage (PL) for each unlearning method against different attackers for Llama-2 7B on the Enron-Email dataset.

Method	PL (lossMIA)			PL (zlibMIA)			PL (Min-K%)		
	Random	Canary	Minority	Random	Canary	Minority	Random	Canary	Minority
No Unlearn	0.050	<b>0.282 (464%↑)</b>	<b>0.174 (248%↑)</b>	0.046	<b>0.236 (413%↑)</b>	<b>0.095 (106%↑)</b>	0.064	<b>0.474 (640%↑)</b>	<b>0.224 (250%↑)</b>
RL	-0.609	-0.567 (7%↓)	-0.821 (12%↓)	-0.832	-0.874 (5%↑)	-0.478 (43%↓)	-0.931	-0.931 (0%)	-0.849 (9%↓)
EUK	0.037	<b>0.241 (551%↑)</b>	<b>0.169 (356%↑)</b>	0.015	<b>0.186 (1140%↑)</b>	<b>0.102 (580%↑)</b>	0.040	<b>0.364 (810%↑)</b>	<b>0.206 (415%↑)</b>
CFk	0.049	<b>0.264 (438%↑)</b>	<b>0.169 (245%↑)</b>	0.046	<b>0.220 (378%↑)</b>	<b>0.090 (96%↑)</b>	0.062	<b>0.436 (603%↑)</b>	<b>0.218 (251%↑)</b>
GA	-0.512	<b>-0.692 (35%↑)</b>	0.059 (89%↓)	-0.435	-0.232 (47%↓)	0.184 (58%↓)	-0.569	-0.479 (16%↓)	0.294 (48%↓)
NegGrad+	-0.931	-0.929 (2%↓)	-0.821 (12%↓)	-0.832	-0.874 (5%↑)	-0.478 (43%↓)	-0.931	-0.931 (0%)	-0.849 (9%↓)
SCRUB	0.040	<b>0.257 (543%↑)</b>	<b>0.174 (335%↑)</b>	0.034	<b>0.209 (515%↑)</b>	<b>0.095 (179%↑)</b>	0.056	<b>0.426 (661%↑)</b>	<b>0.224 (300%↑)</b>
Langevin	0.022	<b>0.191 (768%↑)</b>	<b>0.048 (118%↑)</b>	0.020	<b>0.141 (605%↑)</b>	0.021 (5%↑)	0.035	<b>0.339 (868%↑)</b>	<b>0.079 (126%↑)</b>

Table 7. The privacy leakage (PL) for each unlearning method against different attackers for Llama-2 7B on the ECHR-year dataset.

Method	PL (lossMIA)			PL (zlibMIA)			PL (Min-K%)		
	Random	Canary	Minority	Random	Canary	Minority	Random	Canary	Minority
No Unlearn	0.056	<b>0.094 (68%↑)</b>	<b>0.076 (35%↑)</b>	0.030	<b>0.048 (60%↑)</b>	<b>0.096 (220%↑)</b>	0.067	<b>0.114 (70%↑)</b>	<b>0.138 (106%↑)</b>
RL	-0.069	0.044 (36%↓)	<b>-0.532 (671%↑)</b>	-0.024	<b>0.029 (21%↑)</b>	<b>-0.192 (700%↑)</b>	-0.034	<b>0.070 (106%↑)</b>	<b>-0.458 (1247%↑)</b>
EUK	0.059	<b>0.084 (42%↑)</b>	<b>0.079 (34%↑)</b>	0.030	<b>0.044 (47%↑)</b>	<b>0.079 (163%↑)</b>	0.065	<b>0.110 (69%↑)</b>	<b>0.153 (135%↑)</b>
CFk	0.056	<b>0.088 (57%↑)</b>	<b>0.073 (30%↑)</b>	0.028	<b>0.044 (57%↑)</b>	<b>0.088 (214%↑)</b>	0.063	<b>0.106 (68%↑)</b>	<b>0.131 (108%↑)</b>
GA	-0.046	<b>-0.376 (717%↑)</b>	<b>-0.624 (1257%↑)</b>	-0.016	<b>-0.120 (650%↑)</b>	<b>-0.267 (1569%↑)</b>	-0.063	<b>-0.404 (541%↑)</b>	<b>-0.574 (811%↑)</b>
NegGrad+	0.024	<b>-0.272 (1033%↑)</b>	<b>-0.624 (2500%↑)</b>	0.012	<b>-0.099 (725%↑)</b>	<b>-0.235 (1858%↑)</b>	0.026	<b>-0.404 (1454%↑)</b>	<b>-0.663 (2449%↑)</b>
SCRUB	0.056	<b>0.094 (68%↑)</b>	<b>0.073 (30%↑)</b>	0.030	<b>0.048 (60%↑)</b>	<b>0.090 (200%↑)</b>	0.067	<b>0.112 (67%↑)</b>	<b>0.139 (107%↑)</b>
Langevin	0.026	<b>0.052 (100%↑)</b>	<b>0.041 (58%↑)</b>	0.010	<b>0.025 (150%↑)</b>	<b>0.046 (360%↑)</b>	0.028	<b>0.062 (121%↑)</b>	<b>0.078 (179%↑)</b>

### C.2. More Results on Minority-Aware Evaluation

In this section, we present further benchmarking results for unlearning approaches under minority-aware LLM evaluation. Following the same setup as Section 6.2, Fig. 8 reports the maximum PL score under lossMIA and zlibMIA attackers on Enron-Phone and Enron-Email and Fig. 9 reports the maximum PL score and worst-case perplexity for various unlearning

methods on ECHR-Year (GPT-2) dataset.

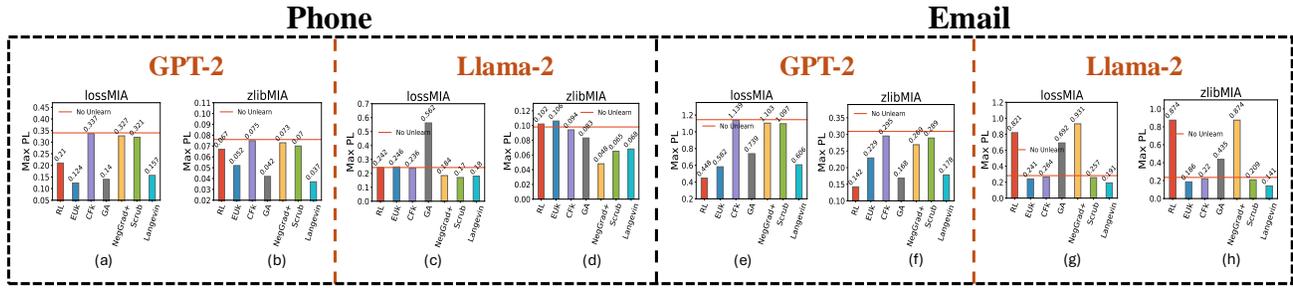


Figure 8. Benchmarking unlearning approaches via our minority-aware evaluation for GPT-2 and Llama-2 on Enron-Phone and Enron-Email dataset. (a),(c),(e),(g): Maximum privacy leakage (PL) over three cases (Random, Canary, and Minority) for lossMIA attack. (b),(d),(f),(h): Maximum privacy leakage (PL) over three cases (Random, Canary, and Minority) for zlibMIA attack.

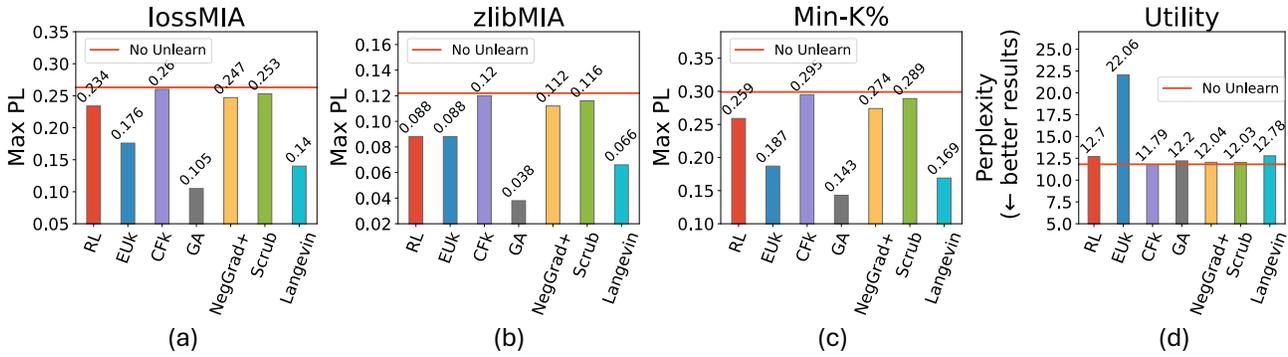


Figure 9. Benchmarking unlearning approaches via our minority-aware evaluation for GPT-2 on ECHR-year dataset. (a)-(c): Maximum privacy leakage (PL) over three cases (Random, Canary, and Minority) for lossMIA, zlibMIA, and Min-K% attacks respectively. (d): Worst perplexity over the three cases of each method.

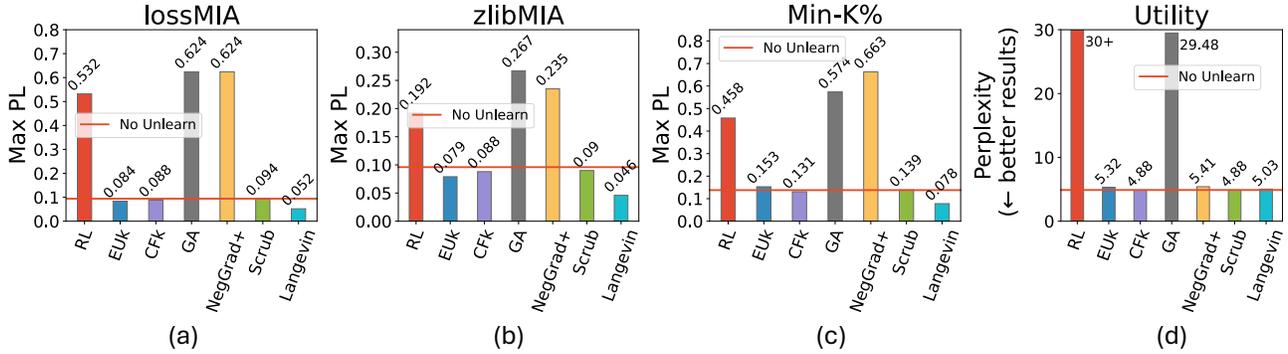


Figure 10. Benchmarking unlearning approaches via our minority-aware evaluation for Llama-2 on ECHR Year dataset. (a)-(c): Maximum privacy leakage (PL) over three cases (Random, Canary, and Minority) for lossMIA, zlibMIA, and Min-K% attacks respectively. (d): Worst perplexity over the three cases of each method.

We observe that both GA and Langevin Unlearning methods maintain a favorable balance between privacy and utility. However, GA can be sensitive to the forget set size and the number of unlearning iterations (Section 6.3). In practice, the GA method should be applied with caution, whereas more stable approaches like Langevin Unlearning offer a better trade-off in terms of privacy, utility, and stability.

### C.3. Results on Other Utility Metrics

In this section, we report the worst-case utility performance across three settings (Random, Canary, and Minority) using utility metrics BERTScore (Zhang et al., 2019) and ROUGE score (Lin, 2004), which capture semantic meaning, on the Enron-Email dataset. The results are shown in Fig.11 (GPT-2) and Fig.12 (Llama-2). As illustrated in the figures, the performance of Random Label and gradient-ascent-based methods (Gradient Ascent and NegGrad+) is unstable under all utility metrics (Perplexity, BERTScore, and ROUGE). In contrast, Langevin Unlearning demonstrates relatively stable performance and achieves a favorable privacy-utility trade-off.

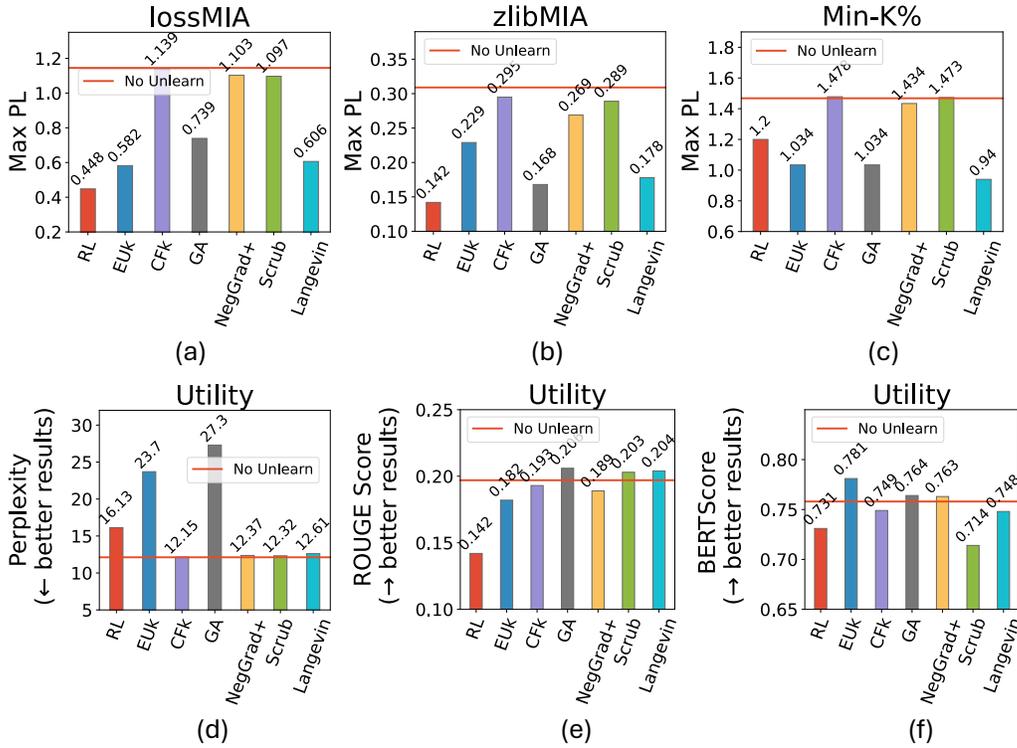


Figure 11. Benchmarking unlearning approaches via our minority-aware evaluation for GPT-2 on Enron-email dataset. (a)-(c): Maximum privacy leakage (PL) over three cases (Random, Canary, and Minority) for lossMIA, zlibMIA, and Min-K% attacks respectively. (d-f): Worst utility performance over the three cases of each method.

### C.4. Further Details and Results on Langevin Unlearning

In this section, we provide additional details and results on the Langevin Unlearning methods. As mentioned in Section 5, Langevin leverages noisy gradient descent and involves training the model on the dataset  $D_{\text{train}}$  using DP-SGD. Furthermore, Langevin conducts machine unlearning by fine-tuning the model on the dataset  $D_{\text{keep}}$  with DP-SGD as well.

It is important to note that for the Langevin Unlearning method, the training process incorporates noise. Consequently, our retrain baseline is adjusted to train the initial model on  $D_{\text{keep}}$  using DP-SGD for 5 epochs. Furthermore, in Table 8 and 9, we report the effectiveness of Langevin Unlearning by evaluating it against three MIA methods (lossMIA, zlibMIA, and Min-K%) across different datasets on GPT-2. These evaluations are conducted across three scenarios (Random, Canary, Minority), assessing the PL scores, the maximum PL scores and the worst-case perplexity. By comparing the results of the Noisy No Unlearn baseline (which fine-tunes the initial model with DP-SGD for 5 epochs) with those of the Langevin Unlearning method, we observe that minority scenarios (Canary, Minority) lead to significantly higher privacy leakage, and Langevin Unlearning achieves superior privacy-utility trade-offs. Additionally, in practical applications, the number of steps employing noisy gradient descent can be tailored based on the acceptable computational costs, thereby enabling potentially better privacy-utility trade-offs. This flexibility allows practitioners to balance the trade-off between enhanced privacy and computational efficiency according to specific application requirements.

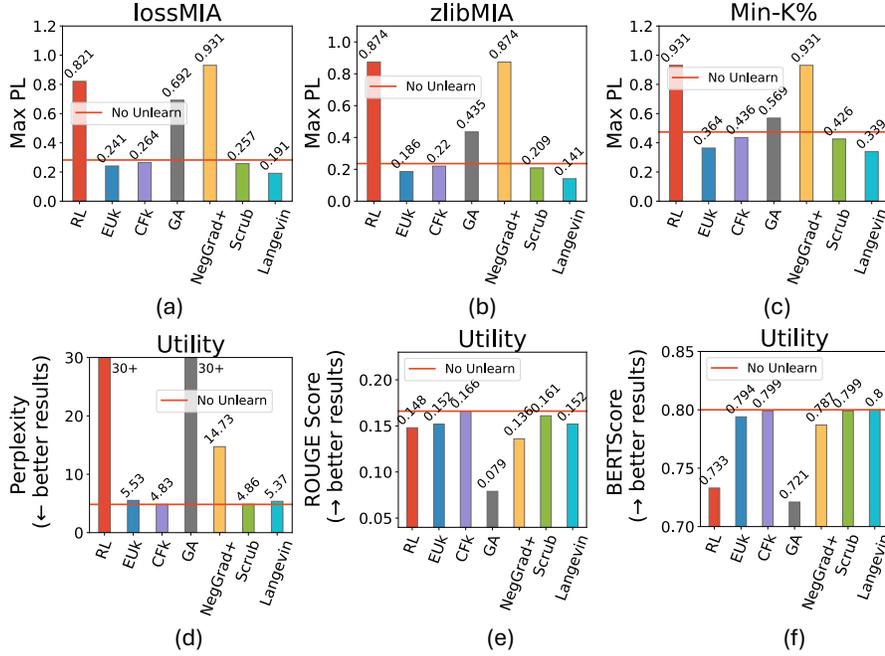


Figure 12. Benchmarking unlearning approaches via our minority-aware evaluation for Llama-2 on Enron-email dataset. (a)-(c): Maximum privacy leakage (PL) over three cases (Random, Canary, and Minority) for lossMIA, zlibMIA, and Min-K% attacks respectively. (d-f): Worst utility performance over the three cases of each method.

Table 8. The privacy leakage (PL) for Langevin Unlearning against different attackers for GPT-2 on All datasets.

Method	PL (lossMIA)			PL (zlibMIA)			PL (Min-K%)		
	Random	Canary	Minority	Random	Canary	Minority	Random	Canary	Minority
Enron-phone									
Noisy No Unlearn	0.097	<b>0.152 (57%↑)</b>	<b>0.170 (75%↑)</b>	0.024	<b>0.039 (63%↑)</b>	<b>0.033 (38%↑)</b>	0.164	<b>0.259 (58%↑)</b>	<b>0.268 (63%↑)</b>
Langevin	0.092	<b>0.144 (57%↑)</b>	<b>0.157 (71%↑)</b>	0.024	<b>0.037 (54%↑)</b>	0.027 (13%↑)	0.159	<b>0.259 (63%↑)</b>	<b>0.264 (66%↑)</b>
Enron-email									
Noisy No Unlearn	0.156	<b>0.342 (119%↑)</b>	<b>0.642 (312%↑)</b>	0.102	<b>0.193 (89%↑)</b>	<b>0.130 (27%↑)</b>	0.344	<b>0.691 (101%↑)</b>	<b>0.945 (175%↑)</b>
Langevin	0.154	<b>0.319 (107%↑)</b>	<b>0.606 (294%↑)</b>	0.097	<b>0.178 (84%↑)</b>	<b>0.124 (28%↑)</b>	0.336	<b>0.645 (92%↑)</b>	<b>0.939 (179%↑)</b>
ECHR-year									
Noisy No Unlearn	0.101	<b>0.152 (51%↑)</b>	<b>0.122 (21%↑)</b>	0.049	<b>0.067 (37%↑)</b>	<b>0.064 (31%↑)</b>	0.122	<b>0.180 (48%↑)</b>	0.145 (19%↑)
Langevin	0.103	<b>0.140 (36%↑)</b>	<b>0.125 (21%↑)</b>	0.049	<b>0.061 (24%↑)</b>	<b>0.065 (33%↑)</b>	0.117	<b>0.168 (44%↑)</b>	<b>0.146 (25%↑)</b>

Table 9. Maximum PL Scores and Worst-case Perplexity for Noisy No Unlearn and Langevin across Datasets on GPT-2

Dataset	Methods	lossMIA	zlibMIA	Min-K%	Perplexity
Enron phone	Noisy No Unlearn	0.170	0.039	0.268	13.87
	Langevin	0.157 (7.65%↓)	0.037 (5.13%↓)	0.264 (1.49%↓)	13.88
Enron email	Noisy No Unlearn	0.642	0.193	0.945	12.52
	Langevin	0.606 (5.61%↓)	0.178 (7.77%↓)	0.939 (0.63%↓)	12.61
ECHR year	Noisy No Unlearn	0.152	0.067	0.180	12.75
	Langevin	0.140 (7.89%↓)	0.061 (8.96%↓)	0.168 (6.67%↓)	12.78

### C.5. More Results on Forget Set Size

In this section, we report additional results on the impact of forget set size for each unlearning method, using LossMIA and ZlibMIA attackers. As shown in Fig. 13, similar to the results in the main text, both RL and GA methods are sensitive to the forget set size, whereas methods like SCRUB and Langevin Unlearning demonstrate greater stability.

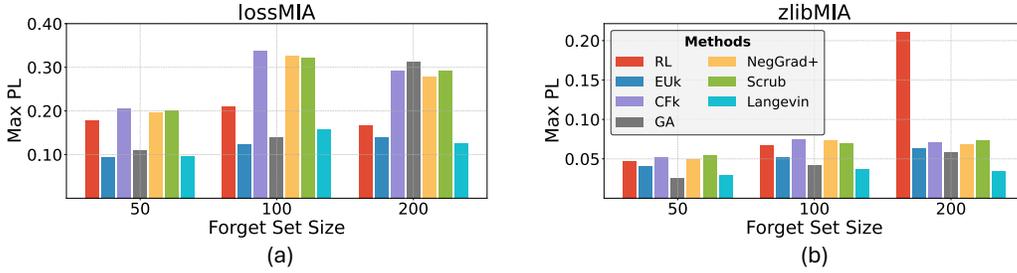


Figure 13. The effect of forget set size for each unlearning approach. (a)(b): Maximum PL over three cases (Random, Canary, Minority) with the attacker being lossMIA and zlibMIA respectively.

C.6. More Details on the Privacy-Utility Trade-off Curves for Langevin Unlearning and SCRUB Methods

This section provides a comprehensive experiments of the privacy-utility trade-off curves for the Langevin Unlearning and SCRUB methods, as introduced in Sec. 6.3. We detailed the hyperparameters used in the ablation study as follows:

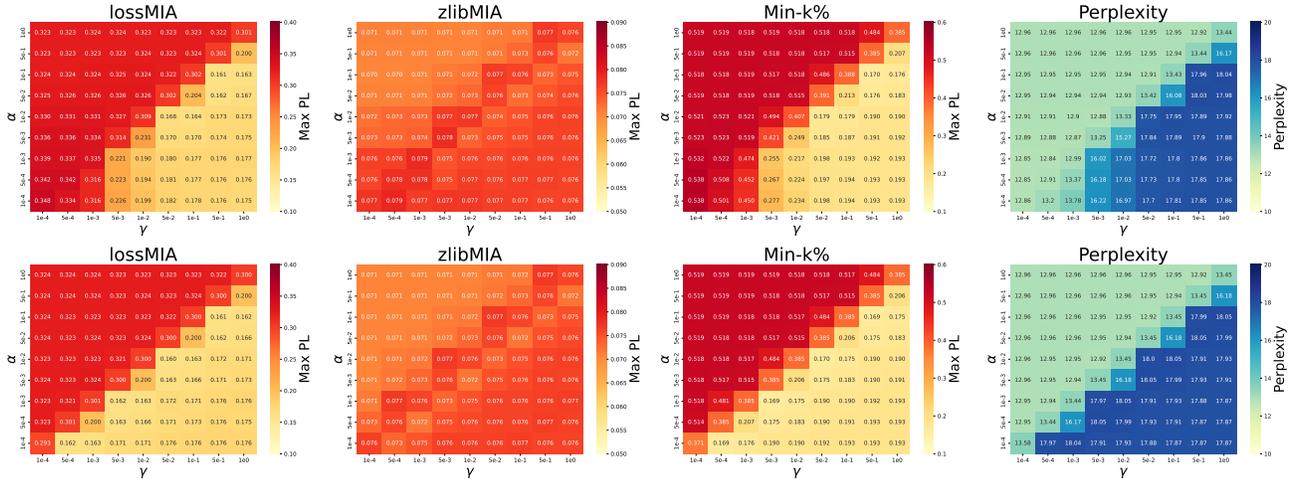


Figure 14. Privacy-utility transition curves for Enron-Phone dataset with hyperparameter  $\beta = 1$  (Top) and hyperparameter  $\beta = 1e - 3$  (Bottom).

**Langevin Unlearning.** For the Langevin Unlearning method, we fix the clipping norm to 1 and vary the noise scale added during training to control the privacy-utility trade-off. In experiments with the GPT-2 model, the noise scale  $\sigma$  is adjusted across the values  $\{1e - 4, 3e - 4, 5e - 4, 8e - 4, 1e - 3\}$ . Table 19 and 20 present the AUC scores under various attackers (lossMIA, zlibMIA, Min-k%) and utility (perplexity) on the Enron-Phone and Enron-Email datasets, respectively.

**SCRUB.** The SCRUB training objective comprises the original loss  $\ell$  on the keep set, along with two KL divergence regularizers on the keep and forget sets. These terms are balanced by three hyperparameters:

$$\hat{\mathbb{E}}_{x \sim D_{\text{keep}}} [\alpha \text{KL}(M_{\text{learn}}(x) \| M(x)) + \beta \ell(M; x)] - \hat{\mathbb{E}}_{x \sim D_{\text{forget}}} [\gamma \text{KL}(M_{\text{learn}}(x) \| M(x))]. \quad (2)$$

We conducted an extensive hyperparameter search, setting  $\beta$  to 1 and  $1e - 3$  in separate configurations. For each fixed  $\beta$ ,  $\alpha$  and  $\gamma$  are independently varied from  $\{1e - 4, 5e - 4, 1e - 3, 5e - 3, 1e - 2, 5e - 2, 1e - 1, 5e - 1, 1\}$ . Fig. 14 and 15 illustrates the resulting transition curves, showing the maximum privacy leakage (PL) for three scenarios (Random, Canary, Minority) across different attackers (lossMIA, zlibMIA, Min-k%) and utility (perplexity) metrics on both Enron-Phone and Enron-Email datasets. The transition curves highlight how SCRUB’s performance depends on balancing the three objective terms. Notably, when the KL regularizer weight on the keep set is greater than or equal to that on the forget set, SCRUB achieves relatively high utility, albeit with increased privacy leakage. Besides, we observe in our experiments that the zlibMIA attacker fails to capture the inherent privacy-utility trade-off for SCRUB as demonstrated in the transition curves.

We further report the privacy-utility trade-off curves for both methods under attacker being lossMIA and zlibMIA. Similar

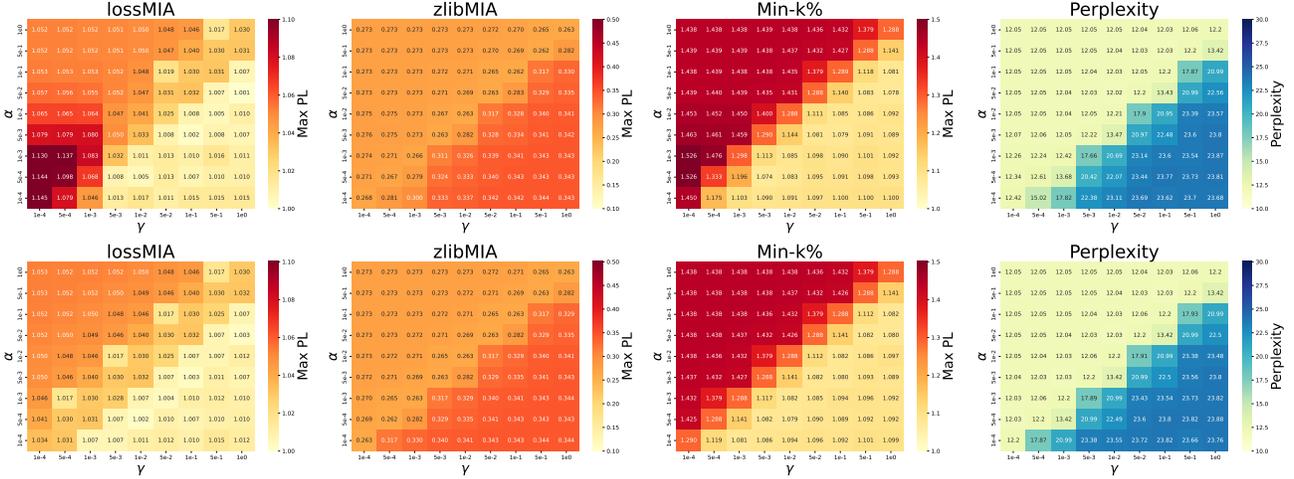


Figure 15. Privacy-utility transition curves for Enron-Email dataset with hyperparameter  $\beta = 1$  (Top) and hyperparameter  $\beta = 1e - 3$  (Bottom).

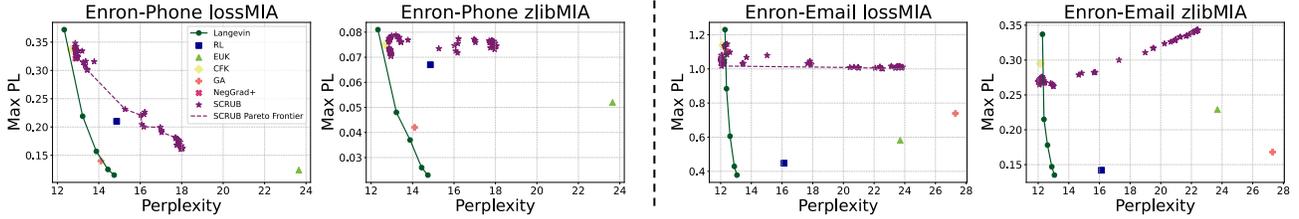


Figure 16. Privacy-utility trade-off curves under lossMIA and zlibMIA.

to the results demonstrated in Sec. 6.3, Langevin Unlearning method achieves the best trade-off performance over SCRUB method.

### C.7. Privacy Underestimation Across Complexity Units

In Fig. 17, we report the degree of largest underestimation (Canary & Minority settings) in privacy leakage compared to Random settings across different complexity units under Min-k% attacker. As demonstrated in the figure, under different complexity units, the privacy leakage for each unlearning methods are severly underestimated. Detailed results under each settings are reported in Table 11 and 12.

### C.8. Discussion on Retraining LLMs with Equivalent Computational Budget

In this subsection, we examine the scenario where the LLM is retrained from scratch using a computational budget equivalent to that of unlearning methods. Intuitively, if retraining with the same computational resources could already achieve a satisfactory privacy-utility trade-off, employing unlearning methods would be less important. To empirically illustrate this point, we consider the Enron (Phone) dataset with the GPT-2 model as a representative case study. Specifically, we report the privacy leakage metrics (PrivLeak) across different attack scenarios, as well as the corresponding perplexity scores in Tab. 10. Additionally, the performance of the retrained model with an equivalent computational budget is highlighted in Fig. 18 (Black). Our findings reveal that while retraining alone indeed achieves relatively low privacy leakage, it drastically compromises model utility. These results underscore the importance and practical necessity of developing and evaluating

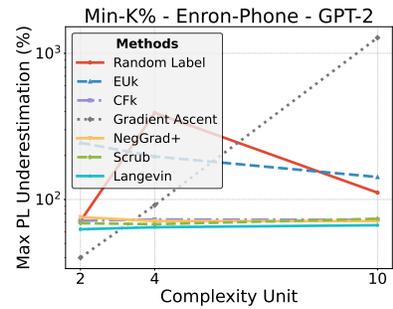


Figure 17. Degree of the largest underestimation in privacy leakage (Canary, Minority) compared to Random settings across varying Complexity Units.

unlearning methods.

Table 10. Privacy Leakage (PL) for Retrain with 10 Complexity Units.

	Random	Canary	Minority
Loss PL	0.004	0.043	0.083
Zlib PL	0.058	0.075	0.077
Min-k% PL	0.083	0.174	<b>0.188</b>
Perplexity		<b>23.44</b>	

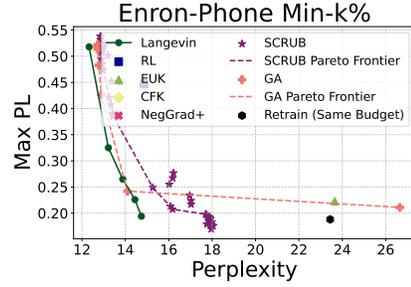


Figure 18. Privacy-Utility Trade-off for Enron-Phone Min-k% with Retrain (Same Computation Budget).

### C.9. Results on AUC Scores, Perplexity across Different Models and Datasets

We further report the AUC scores under different attackers (lossMIA, zlibMIA, Min-K%) and utility (perplexity) over holdout test set  $D_{test}$  for GPT-2 and Llama-2 in Table. 13-18.

### C.10. Discussion on TPR@low FPR Metric

Note that aside from the AUC score, a commonly reported metric for privacy evaluation is **TPR@low FPR** (Carlini et al., 2022), where the low FPR is often set to 0.01. However, in our scenario, the canary size is set to 100 (1% of the total training set size). At FPR = 0.01, the TPR would be calculated based on only a few canary samples, making the overall score very coarse. To avoid the impact of this coarse granularity on our experimental results, we primarily focus on the AUC score.

Table 11. PL Scores for MIA across Three Settings for GPT-2 on Enron (Phone) under Different Complexity Units.

Methods	Enron-Phone GPT-2 Min-k%								
	Complexity Units 2			Complexity Units 4			Complexity Units 10		
	Random	Canary	Minority	Random	Canary	Minority	Random	Canary	Minority
Random Label	0.166	0.284	0.136	-0.059	0.037	0.290	0.168	0.320	0.355
EUK	0.068	0.169	0.234	0.068	0.174	0.202	0.092	0.216	0.223
CFk	0.302	0.445	0.518	0.300	0.438	0.519	0.298	0.435	0.514
Gradient Ascent	-0.175	-0.245	-0.203	-0.245	-0.469	-0.352	-0.031	-0.012	-0.427
NegGrad+	0.309	0.452	0.543	0.309	0.447	0.529	0.298	0.452	0.510
Scrub	0.311	0.443	0.525	0.313	0.447	0.525	0.306	0.445	0.532
Langevin	0.163	0.259	0.265	0.161	0.259	0.265	0.159	0.259	0.265

Table 12. Perplexity for MIA across Three Settings for GPT-2 on Enron (Phone) under Different Complexity Units.

Methods	Enron-Phone GPT-2 Perplexity								
	Complexity Units 2			Complexity Units 4			Complexity Units 10		
	Random	Canary	Minority	Random	Canary	Minority	Random	Canary	Minority
Random Label	29.85	28.93	30+	30+	30+	30+	30+	30+	30+
EUK	30+	30+	30+	30+	30+	30+	23.64	23.65	26.60
CFk	12.72	12.71	12.72	12.72	12.72	12.73	12.67	12.67	12.67
Gradient Ascent	16.63	15.76	29.28	30+	30+	30+	30+	30+	30+
NegGrad+	12.86	12.87	12.83	12.84	12.86	12.88	12.83	12.80	12.88
Scrub	12.94	12.95	12.96	13.11	13.09	13.19	13.09	12.88	12.96
Langevin	13.84	13.85	13.88	13.85	13.86	13.88	13.84	13.88	13.88

Table 13. AUC Scores for MIA and Perplexity across Three Settings for GPT-2 on Enron (Phone Numbers)

	AUC - LossMIA			AUC - ZlibMIA		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.533	0.531	0.422	0.694	0.693	0.619
<b>Retrain</b>	0.448	0.414	0.315	0.660	0.644	0.582
<b>Noisy No Unlearn</b>	0.474	0.464	0.365	0.678	0.674	0.606
<b>Noisy Retrain</b>	0.432	0.403	0.312	0.662	0.649	0.587
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.501	0.493	0.381	0.689	0.687	0.617
<b>Langevin</b>	0.472	0.461	0.361	0.678	0.673	0.603
<b>EUK</b>	0.460	0.447	0.354	0.683	0.677	0.612
<b>CFk</b>	0.533	0.529	0.421	0.695	0.692	0.619
<b>Gradient Ascent</b>	0.488	0.472	0.355	0.676	0.671	0.597
<b>NegGrad+</b>	0.530	0.526	0.418	0.694	0.691	0.616
<b>SCRUB</b>	0.523	0.518	0.416	0.692	0.689	0.618
	AUC - Min-K%			Perplexity		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.594	0.592	0.471	12.72	12.72	12.72
<b>Retrain</b>	0.457	0.409	0.309	12.74	12.74	12.74
<b>Noisy No Unlearn</b>	0.518	0.511	0.393	13.84	13.85	13.87
<b>Noisy Retrain</b>	0.445	0.406	0.310	13.84	13.84	13.83
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.575	0.573	0.447	14.49	14.41	14.86
<b>Langevin</b>	0.516	0.511	0.392	13.84	13.88	13.88
<b>EUK</b>	0.499	0.497	0.378	23.64	23.65	23.60
<b>CFk</b>	0.593	0.587	0.468	12.67	12.67	12.67
<b>Gradient Ascent</b>	0.526	0.508	0.362	13.22	13.20	14.10
<b>NegGrad+</b>	0.591	0.587	0.467	12.86	12.86	12.88
<b>SCRUB</b>	0.592	0.593	0.472	13.00	12.98	12.96

Table 14. AUC Scores for MIA and Perplexity across Three Settings for GPT-2 on Enron (Email)

	AUC - LossMIA			AUC - ZlibMIA		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.555	0.551	0.354	0.462	0.462	0.563
<b>Retrain</b>	0.426	0.359	0.165	0.385	0.353	0.446
<b>Noisy No Unlearn</b>	0.488	0.483	0.271	0.422	0.421	0.512
<b>Noisy Retrain</b>	0.422	0.360	0.165	0.383	0.353	0.453
Unlearning Methods						
<b>Random Label</b>	0.440	0.414	0.239	0.409	0.403	0.500
<b>Langevin</b>	0.487	0.475	0.265	0.420	0.416	0.509
<b>EUK</b>	0.525	0.517	0.261	0.437	0.434	0.514
<b>CFk</b>	0.552	0.544	0.353	0.461	0.457	0.562
<b>Gradient Ascent</b>	0.307	0.297	0.287	0.339	0.340	0.521
<b>NegGrad+</b>	0.539	0.528	0.347	0.454	0.448	0.558
<b>SCRUB</b>	0.548	0.538	0.346	0.458	0.455	0.559
	AUC - Min-K%			Perplexity		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.607	0.611	0.506	12.11	12.11	12.12
<b>Retrain</b>	0.397	0.316	0.205	12.19	12.19	12.19
<b>Noisy No Unlearn</b>	0.516	0.519	0.387	12.51	12.52	12.50
<b>Noisy Retrain</b>	0.384	0.307	0.199	12.48	12.48	12.48
Unlearning Methods						
<b>Random Label</b>	0.568	0.560	0.451	15.87	16.13	13.54
<b>Langevin</b>	0.513	0.505	0.386	12.58	12.58	12.61
<b>EUK</b>	0.596	0.596	0.417	23.58	23.70	23.58
<b>CFk</b>	0.606	0.602	0.508	12.14	12.15	12.15
<b>Gradient Ascent</b>	0.242	0.220	0.417	27.30	21.86	14.53
<b>NegGrad+</b>	0.594	0.589	0.499	12.34	12.34	12.37
<b>SCRUB</b>	0.603	0.601	0.507	12.15	12.15	12.32

Table 15. AUC Scores for MIA and Perplexity across Three Settings for GPT-2 on ECHR (Year)

	AUC - LossMIA			AUC - ZlibMIA		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.661	0.650	0.658	0.532	0.524	0.560
<b>Retrain</b>	0.552	0.521	0.521	0.490	0.475	0.499
<b>Noisy No Unlearn</b>	0.600	0.592	0.581	0.514	0.509	0.535
<b>Noisy Retrain</b>	0.545	0.514	0.518	0.490	0.477	0.503
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.641	0.632	0.643	0.523	0.517	0.542
<b>Langevin</b>	0.601	0.586	0.583	0.514	0.506	0.536
<b>EUK</b>	0.621	0.613	0.593	0.523	0.517	0.534
<b>CFk</b>	0.656	0.643	0.656	0.531	0.520	0.559
<b>Gradient Ascent</b>	0.589	0.535	0.576	0.502	0.484	0.518
<b>NegGrad+</b>	0.653	0.636	0.650	0.525	0.517	0.555
<b>SCRUB</b>	0.651	0.637	0.653	0.529	0.522	0.557
	AUC - Min-K%			Perplexity		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.671	0.661	0.673	11.81	11.81	11.81
<b>Retrain</b>	0.553	0.518	0.518	11.82	11.82	11.82
<b>Noisy No Unlearn</b>	0.615	0.609	0.586	12.74	12.74	12.75
<b>Noisy Retrain</b>	0.548	0.516	0.512	12.73	12.73	12.73
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.658	0.652	0.651	12.70	12.70	12.67
<b>Langevin</b>	0.612	0.603	0.587	12.78	12.78	12.78
<b>EUK</b>	0.616	0.615	0.588	22.04	22.06	22.00
<b>CFk</b>	0.669	0.655	0.671	11.78	11.78	11.79
<b>Gradient Ascent</b>	0.603	0.508	0.592	12.11	12.20	12.11
<b>NegGrad+</b>	0.659	0.641	0.660	12.01	12.04	12.03
<b>SCRUB</b>	0.662	0.649	0.668	12.02	12.03	12.01

Table 16. AUC Scores for MIA and Perplexity across Three Settings for Llama-2 on Enron (Phone Number)

	AUC - LossMIA			AUC - ZlibMIA		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.614	0.606	0.551	0.578	0.571	0.575
<b>Retrain</b>	0.579	0.488	0.470	0.559	0.520	0.539
<b>Noisy No Unlearn</b>	0.605	0.598	0.539	0.578	0.572	0.576
<b>Noisy Retrain</b>	0.584	0.500	0.482	0.568	0.533	0.554
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.439	0.447	0.444	0.556	0.554	0.594
<b>Langevin</b>	0.603	0.590	0.532	0.577	0.569	0.574
<b>EUK</b>	0.612	0.608	0.557	0.581	0.575	0.583
<b>CFk</b>	0.612	0.603	0.549	0.577	0.569	0.573
<b>Gradient Ascent</b>	0.253	0.278	0.252	0.551	0.540	0.584
<b>NegGrad+</b>	0.536	0.398	0.451	0.547	0.495	0.538
<b>SCRUB</b>	0.613	0.567	0.550	0.578	0.553	0.574
	AUC - Min-K%			Perplexity		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.611	0.610	0.579	9.45	9.47	9.48
<b>Retrain</b>	0.568	0.547	0.491	9.48	9.48	9.48
<b>Noisy No Unlearn</b>	0.600	0.602	0.570	10.21	10.21	10.21
<b>Noisy Retrain</b>	0.579	0.565	0.516	10.19	10.19	10.19
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.498	0.507	0.497	2124	2559	2445
<b>Langevin</b>	0.598	0.596	0.563	10.20	10.21	10.20
<b>EUK</b>	0.604	0.619	0.584	11.00	11.24	10.82
<b>CFk</b>	0.609	0.606	0.575	9.43	9.45	9.46
<b>Gradient Ascent</b>	0.213	0.296	0.237	4e9	8e9	2e9
<b>NegGrad+</b>	0.529	0.399	0.463	9.82	9.90	9.85
<b>SCRUB</b>	0.610	0.530	0.578	9.45	9.48	9.48

Table 17. AUC Scores for MIA and Perplexity across Three Settings for Llama-2 on Enron (Email)

	AUC - LossMIA			AUC - ZlibMIA		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.628	0.613	0.418	0.548	0.539	0.451
<b>Retrain</b>	0.598	0.478	0.356	0.524	0.436	0.412
<b>Noisy No Unlearn</b>	0.594	0.572	0.393	0.515	0.502	0.443
<b>Retrain</b>	0.579	0.470	0.375	0.503	0.433	0.434
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.234	0.207	0.394	0.333	0.333	0.458
<b>Langevin</b>	0.592	0.560	0.393	0.513	0.494	0.443
<b>EUK</b>	0.620	0.593	0.416	0.532	0.517	0.454
<b>CFk</b>	0.627	0.604	0.416	0.548	0.532	0.449
<b>Gradient Ascent</b>	0.292	0.147	0.377	0.296	0.335	0.488
<b>NegGrad+</b>	0.041	0.034	0.064	0.088	0.055	0.215
<b>SCRUB</b>	0.622	0.601	0.418	0.542	0.527	0.451
	AUC - Min-K%			Perplexity		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.632	0.619	0.421	4.83	4.84	4.84
<b>Retrain</b>	0.594	0.420	0.344	4.84	4.84	4.84
<b>Noisy No Unlearn</b>	0.592	0.569	0.397	5.38	5.38	5.38
<b>Noisy Retrain</b>	0.570	0.410	0.368	5.39	5.39	5.39
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.230	0.181	0.330	730	542	255
<b>Langevin</b>	0.590	0.549	0.397	5.36	5.36	5.37
<b>EUK</b>	0.618	0.573	0.415	5.42	5.53	5.37
<b>CFk</b>	0.631	0.603	0.419	4.83	4.83	4.83
<b>Gradient Ascent</b>	0.256	0.219	0.445	4e8	6e12	6e12
<b>NegGrad+</b>	0.041	0.029	0.052	12.15	14.73	6.20
<b>SCRUB</b>	0.627	0.599	0.421	4.86	4.86	4.86

Table 18. AUC Scores for MIA and Perplexity across Three Settings for Llama-2 on ECHR (Year)

	AUC - LossMIA			AUC - ZlibMIA		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.570	0.547	0.726	0.513	0.499	0.513
<b>Retrain</b>	0.540	0.500	0.675	0.498	0.476	0.468
<b>Noisy No Unlearn</b>	0.556	0.532	0.721	0.504	0.491	0.510
<b>Noisy Retrain</b>	0.541	0.502	0.685	0.498	0.476	0.476
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.503	0.522	0.316	0.486	0.490	0.378
<b>Langevin</b>	0.555	0.528	0.715	0.503	0.488	0.498
<b>EUK</b>	0.572	0.542	0.728	0.513	0.497	0.505
<b>CFk</b>	0.570	0.544	0.724	0.512	0.497	0.509
<b>Gradient Ascent</b>	0.515	0.312	0.254	0.490	0.419	0.343
<b>NegGrad+</b>	0.553	0.364	0.254	0.504	0.429	0.358
<b>SCRUB</b>	0.570	0.547	0.724	0.513	0.499	0.510
	AUC - Min-K%			Perplexity		
	Random	Canary	Minority	Random	Canary	Minority
<b>No Unlearn</b>	0.573	0.559	0.686	4.89	4.89	4.89
<b>Retrain</b>	0.537	0.502	0.603	4.89	4.89	4.89
<b>Noisy No Unlearn</b>	0.553	0.541	0.675	5.03	5.03	5.02
<b>Noisy Retrain</b>	0.537	0.504	0.616	5.02	5.02	5.02
<b>Unlearning Methods</b>						
<b>Random Label</b>	0.519	0.537	0.327	90	129	127.43
<b>Langevin</b>	0.552	0.535	0.664	5.03	5.03	5.03
<b>EUK</b>	0.572	0.557	0.695	5.27	5.21	5.32
<b>CFk</b>	0.571	0.555	0.682	4.87	4.88	4.87
<b>Gradient Ascent</b>	0.503	0.299	0.257	7.94	10.68	29.48
<b>NegGrad+</b>	0.551	0.299	0.203	4.96	5.10	5.41
<b>SCRUB</b>	0.573	0.558	0.687	4.88	4.88	4.87

Table 19. AUC Scores for MIA and Perplexity across Three Settings for GPT-2 on Enron (Phone Number) for Noisy Learning

	AUC - LossMIA			AUC - ZlibMIA		
	Random	Canary	Minority	Random	Canary	Minority
<b>Noisy No Unlearn 1e-4</b>	0.541	0.536	0.427	0.699	0.697	0.621
<b>Noisy Retrain 1e-4</b>	0.446	0.412	0.309	0.658	0.643	0.577
<b>Langevin 1e-4</b>	0.540	0.535	0.424	0.698	0.695	0.619
<b>Noisy No Unlearn 3e-4</b>	0.492	0.484	0.381	0.682	0.678	0.607
<b>Noisy Retrain 3e-4</b>	0.436	0.405	0.311	0.660	0.647	0.583
<b>Langevin 3e-4</b>	0.492	0.483	0.379	0.682	0.678	0.606
<b>Noisy No Unlearn 5e-4</b>	0.474	0.464	0.365	0.678	0.674	0.606
<b>Noisy Retrain 5e-4</b>	0.432	0.403	0.312	0.662	0.649	0.587
<b>Langevin 5e-4</b>	0.472	0.461	0.361	0.678	0.673	0.603
<b>Noisy No Unlearn 8e-4</b>	0.463	0.450	0.352	0.677	0.672	0.606
<b>Noisy Retrain 8e-4</b>	0.428	0.401	0.311	0.664	0.653	0.590
<b>Langevin 8e-4</b>	0.460	0.448	0.350	0.676	0.670	0.604
<b>Noisy No Unlearn 1e-3</b>	0.459	0.446	0.348	0.677	0.671	0.606
<b>Noisy Retrain 1e-3</b>	0.427	0.400	0.311	0.665	0.655	0.592
<b>Langevin 1e-3</b>	0.456	0.443	0.347	0.675	0.670	0.604
	AUC - Min-K%			Perplexity		
	Random	Canary	Minority	Random	Canary	Minority
<b>Noisy No Unlearn 1e-4</b>	0.599	0.599	0.465	12.26	12.24	12.27
<b>Noisy Retrain 1e-4</b>	0.454	0.410	0.303	12.25	12.25	12.25
<b>Langevin 1e-4</b>	0.598	0.596	0.460	12.32	12.32	12.33
<b>Noisy No Unlearn 3e-4</b>	0.538	0.535	0.409	13.20	13.20	13.21
<b>Noisy Retrain 3e-4</b>	0.447	0.405	0.308	13.19	13.19	13.19
<b>Langevin 3e-4</b>	0.538	0.535	0.408	13.24	13.24	13.22
<b>Noisy No Unlearn 5e-4</b>	0.518	0.511	0.393	13.84	13.85	13.87
<b>Noisy Retrain 5e-4</b>	0.445	0.406	0.310	13.84	13.84	13.83
<b>Langevin 5e-4</b>	0.516	0.511	0.392	13.84	13.88	13.88
<b>Noisy No Unlearn 8e-4</b>	0.505	0.498	0.378	14.40	14.39	14.39
<b>Noisy Retrain 8e-4</b>	0.442	0.405	0.313	14.42	14.42	14.42
<b>Langevin 8e-4</b>	0.502	0.497	0.376	14.44	14.44	14.43
<b>Noisy No Unlearn 1e-3</b>	0.500	0.492	0.374	14.71	14.70	14.70
<b>Noisy Retrain 1e-3</b>	0.440	0.408	0.313	14.73	14.73	14.73
<b>Langevin 1e-3</b>	0.497	0.487	0.371	14.74	14.74	14.73

Table 20. AUC Scores for MIA and Perplexity across Three Settings for GPT-2 on Enron (Phone Email) for Noisy Learning

	AUC - LossMIA			AUC - ZlibMIA		
	Random	Canary	Minority	Random	Canary	Minority
<b>Noisy No Unlearn 1e-4</b>	0.584	0.587	0.369	0.480	0.483	0.578
<b>Noisy Retrain 1e-4</b>	0.431	0.362	0.165	0.387	0.356	0.447
<b>Langevin 1e-4</b>	0.584	0.578	0.368	0.479	0.476	0.579
<b>Noisy No Unlearn 3e-4</b>	0.512	0.511	0.311	0.436	0.436	0.535
<b>Noisy Retrain 3e-4</b>	0.424	0.360	0.164	0.384	0.353	0.449
<b>Langevin 3e-4</b>	0.510	0.499	0.309	0.432	0.429	0.534
<b>Noisy No Unlearn 5e-4</b>	0.488	0.483	0.271	0.422	0.421	0.512
<b>Noisy Retrain 5e-4</b>	0.422	0.360	0.165	0.383	0.353	0.453
<b>Langevin 5e-4</b>	0.487	0.475	0.265	0.420	0.416	0.509
<b>Noisy No Unlearn 8e-4</b>	0.470	0.464	0.248	0.414	0.412	0.503
<b>Noisy Retrain 8e-4</b>	0.417	0.357	0.170	0.383	0.354	0.459
<b>Langevin 8e-4</b>	0.471	0.457	0.243	0.411	0.406	0.500
<b>Noisy No Unlearn 1e-3</b>	0.463	0.456	0.243	0.410	0.408	0.500
<b>Noisy Retrain 1e-3</b>	0.414	0.354	0.172	0.381	0.355	0.462
<b>Langevin 1e-3</b>	0.462	0.450	0.237	0.408	0.403	0.499
	AUC - Min-K%			Perplexity		
	Random	Canary	Minority	Random	Canary	Minority
<b>Noisy No Unlearn 1e-4</b>	0.651	0.661	0.528	12.18	12.25	12.16
<b>Noisy Retrain 1e-4</b>	0.407	0.318	0.210	12.25	12.25	12.25
<b>Langevin 1e-4</b>	0.655	0.655	0.532	12.27	12.28	12.27
<b>Noisy No Unlearn 3e-4</b>	0.553	0.562	0.420	12.31	12.34	12.28
<b>Noisy Retrain 3e-4</b>	0.390	0.311	0.199	12.23	12.23	12.23
<b>Langevin 3e-4</b>	0.552	0.547	0.421	12.38	12.39	12.38
<b>Noisy No Unlearn 5e-4</b>	0.516	0.519	0.387	12.51	12.52	12.50
<b>Noisy Retrain 5e-4</b>	0.384	0.307	0.199	12.48	12.48	12.48
<b>Langevin 5e-4</b>	0.513	0.505	0.386	12.58	12.58	12.61
<b>Noisy No Unlearn 8e-4</b>	0.478	0.483	0.355	12.83	12.80	12.82
<b>Noisy Retrain 8e-4</b>	0.369	0.298	0.200	12.84	12.84	12.84
<b>Langevin 8e-4</b>	0.477	0.469	0.349	12.87	12.86	12.89
<b>Noisy No Unlearn 1e-3</b>	0.465	0.467	0.341	13.01	12.98	13.01
<b>Noisy Retrain 1e-3</b>	0.365	0.296	0.201	13.02	13.02	13.02
<b>Langevin 1e-3</b>	0.462	0.453	0.335	13.04	13.03	13.06