

Membership Inference Attacks on Tokenizers of Large Language Models

Meng Tong^{1*} Yuntao Du^{2*} Kejiang Chen¹✉ Weiming Zhang¹ Ninghui Li²

¹University of Science and Technology of China ²Purdue University

*Equal Contribution ✉Corresponding Author

Abstract

Membership inference attacks (MIAs) are widely used to assess the privacy risks associated with machine learning models. However, when these attacks are applied to pre-trained large language models (LLMs), they encounter significant challenges, including mislabeled samples, distribution shifts, and discrepancies in model size between experimental and real-world settings. To address these limitations, we introduce tokenizers as a new attack vector for membership inference. Specifically, a tokenizer converts raw text into tokens for LLMs. Unlike full models, tokenizers can be efficiently trained from scratch, thereby avoiding the aforementioned challenges. In addition, the tokenizer’s training data is typically representative of the data used to pre-train LLMs. Despite these advantages, the potential of tokenizers as an attack vector remains unexplored. To this end, we present the first study on membership leakage through tokenizers and explore five attack methods to infer dataset membership. Extensive experiments on millions of Internet samples reveal the vulnerabilities in the tokenizers of state-of-the-art LLMs. To mitigate this emerging risk, we further propose an adaptive defense. Our findings highlight tokenizers as an overlooked yet critical privacy threat, underscoring the urgent need for privacy-preserving mechanisms specifically designed for them.

1 Introduction

Scaling up the pre-training data for large language models (LLMs) has been shown to improve performance [10, 36, 42, 57]. Nevertheless, the rapid expansion of pre-training data has also raised concerns about whether these commercial models are trained on sensitive or copyrighted information [16, 55]. For instance, on June 4, 2025, Reddit filed a lawsuit against Anthropic, alleging the unlawful use of data from its 100 million daily users to train LLMs [59]. Furthermore, an increasing body of research [12, 34, 35] has documented instances in which LLMs memorize and leak private information.

To assess potential data misuse, extensive research has explored the membership inference attacks (MIAs) in LLMs [61,




Figure 1: Evaluation challenges in MIAs against LLMs.

74, 79, 84]. In particular, an MIA aims to determine whether a specific data sample or dataset was used to train the target model (i.e., *member*) or not (i.e., *non-member*). To achieve this, existing MIAs primarily rely on the model’s output as the attack vector [24, 34, 86]. Although this vector is widely adopted, these attacks face significant challenges in reliably demonstrating their effectiveness for LLMs [32, 55], as shown in Figure 1. The primary obstacle is that faithful evaluation [37] requires an evaluator to pre-train an LLM from scratch [82], which incurs significant computational costs. As a result, existing MIAs are typically evaluated using LLMs that have already been pre-trained by others. Nevertheless, this may lead to MIA evaluation exhibiting *distribution shifts* [24, 50] or containing *mislabeled samples* [55, 72]. Furthermore, many of the evaluated models (e.g., Pythia-12B [9]) are much smaller than practical deployed LLMs (e.g., DeepSeek-R1-671B [29]), limiting the ability to assess current MIAs in real-world conditions. Given these challenges, a natural question arises: *Can we exploit an attack vector for MIAs against LLMs that avoids these limitations?*

New Attack Vector. Motivated by this question, we explore a new attack vector that targets other components of LLMs. Typically, an LLM comprises a tokenizer, a transformer network, and an output layer [80]. Among these components, the tokenizer has been open-sourced in commercial LLMs such as OpenAI-o3 [60] and Gemini-1.5 [28] to support transparent billing. Building on this observation, *we propose the previously overlooked tokenizer as a new attack vector for membership inference*. Specifically, a tokenizer [88] is in charge of

converting text into tokens for LLMs. Its training data is typically representative of the overall pre-training corpus of the LLMs [10, 18, 62, 78, 80]. Its training process simply involves merging the most frequent strings into a vocabulary using the byte-pair encoding (BPE) algorithm [26]. This straightforward process enables training a tokenizer from scratch, which aligns with the inference game [82] and avoids mislabeled samples or distribution shifts. Furthermore, the simplicity of BPE also makes it feasible to train a tokenizer that matches those used in state-of-the-art LLMs (see Figure 5).

Despite these advantages, the feasibility of using tokenizers as an attack vector has not yet been explored. In this paper, we present the first study to exploit tokenizers for MIAs and propose five attack methods for inferring dataset membership:

- **MIA via Merge Similarity.** This attack trains shadow tokenizers [74] and compares their token merge orders to that of the target tokenizer. If the merge order of the target tokenizer closely matches that of shadow tokenizers trained on a particular dataset, the dataset is classified as a *member*. However, the effectiveness of this attack is limited. Only a few distinctive tokens display a membership signal in merge order, making it difficult for membership inference.
- **MIA via Vocabulary Overlap.** Leveraging these distinctive tokens, we propose a more effective attack, MIA via Vocabulary Overlap. This attack also involves training shadow tokenizers. But instead of comparing merge orders, it classifies a dataset as a *member* if the distinctive tokens in the target tokenizer’s vocabulary significantly overlap with those from shadow tokenizers trained on that dataset.
- **MIA via Frequency Estimation.** While our results show that MIA via Vocabulary Overlap achieves strong performance, it requires substantial time to train multiple shadow tokenizers. For efficient implementation, we introduce MIA via Frequency Estimation, which trains only a single shadow tokenizer. This attack evaluates whether training the target tokenizer on a dataset is necessary for certain tokens to appear in its vocabulary. If this condition is met, this attack classifies the dataset as a *member*.
- Additionally, as part of our evaluation in Section 5.1, we further explore two attack methods for potential alternatives: MIA via Naive Bayes and MIA via Compression Rate.

We conduct extensive evaluations using millions of Internet data [64]. To match real-world practice, we align the trained tokenizers in evaluations with those used in commercial LLMs [9, 29, 80]. The experimental results indicate that MIAs via Vocabulary Overlap and Frequency Estimation achieve strong performance across various settings. For example, MIA via Vocabulary Overlap achieves an AUC score of 0.771 against a tokenizer with two hundred thousand tokens, whereas MIA via Frequency Estimation achieves an AUC score of 0.740. More importantly, our experiments show that scaling laws increase tokenizers’ vulnerability to MIAs. This finding suggests that MIAs could become more effective on

scaled-up tokenizers in the future.

Our Contributions. Our main contributions are as follows:

- We introduce the tokenizer as a new attack vector for membership inference and conduct the first study demonstrating its feasibility in LLMs.
- We explore five attack methods for set-level membership inference against tokenizers, revealing the vulnerabilities in these foundational components of state-of-the-art LLMs.
- We conduct extensive evaluations using real-world Internet datasets. The results show that our shadow-based attacks demonstrate strong performance against tokenizers.
- We further analyze tokenizers from commercial LLMs. The results show that the tokenizers, such as OpenAI-o200k [60] and DeepSeek-R1 [29], also contain distinctive tokens for implementing membership inference.

Main Findings. We have the following key findings:

- According to prior work [38, 53], scaling up the intelligence of LLMs involves expanding the tokenizer’s vocabulary [77] and thus improving its compression efficiency [46]. However, our experimental results show that it also increases the tokenizer’s vulnerability to effective MIAs.
- The membership status of the target dataset with more data samples is typically more accurately inferred by MIAs.
- While removing infrequent tokens from the tokenizer’s vocabulary can partially reduce the effectiveness of MIAs, this approach also lowers the tokenizer’s compression efficiency. Moreover, even with this mitigation, MIAs can remain effective for inferring large datasets.

Organization. The remainder of this paper is organized as follows. Section 2 presents preliminaries on tokenizer training and membership inference. Section 3 introduces the threat models for membership inference attacks. Section 4 presents three of our attack methods. Section 5 introduces two additional methods and reports the experimental results. Section 6 reviews related work. Section 7 discusses the limitations of LLM dataset inference. Section 8 concludes the paper.

2 Preliminaries

2.1 Tokenizer Training

A tokenizer [88] is a fundamental component in LLMs, converting raw text into a format that the model can process. Formally, a tokenizer is defined as a function $f_{\mathcal{V}} : S \rightarrow \mathcal{V}^*$ that maps an input string $s \in S$ (e.g., a sentence or document) into a sequence of tokens from a vocabulary \mathcal{V} . In practice, this function is learned from a collection of text datasets \mathcal{D} . Specifically, its training objective is to segment and encode the data in a way that maximizes compression efficiency [88]. This process begins by initializing the vocabulary \mathcal{V} with basic symbols, such as individual characters. During training, the tokenizer $f_{\mathcal{V}}$ iteratively merges the most frequent pairs of symbols in the data via the byte-pair encoding (BPE) al-

Table 1: Tokenizer Information of Commercial LLMs.

| LLMs | Affiliation | Tokenizer’s Training Data | Download Tokenizer |
|----------------------------|-------------|-----------------------------|------------------------------|
| LLaMA [80] | Meta | LLaMA’s Pretraining Data | GitHub |
| Gemini [78] | Google | Gemini’s Pretraining Data | GitHub |
| GPT-3 ¹ [10] | OpenAI | GPT-3’s Pretraining Data | GitHub |
| Nemotron ² [62] | NVIDIA | Nemotron’s Pretraining Data | Hugging Face |
| DeepSeek [18] | DeepSeek | DeepSeek’s Pretraining Data | Hugging Face |

¹ GPT-3 utilizes the GPT-2’s tokenization, i.e., tokenizer, which is trained on the Web-Text [63] (one of GPT-3’s pretraining corpora).

² The Series of Nemotron [2, 62] share both the pretraining data and the tokenizer.

gorithm [73]. This iterative process results in a token merge order: each token $t_i \in \mathcal{V}$ is assigned an index i corresponding to the iteration, where it was merged into the vocabulary \mathcal{V} .

In commercial LLM applications, the tokenizer also serves as a basis for token billing. As the tokenizer directly determines how users are charged based on the number of tokens in a message, its operation is critical for ensuring transparent billing [39]. To promote such transparency in token counting, the organizations behind major LLMs [28, 29, 60] have open-sourced their tokenizers, making their vocabularies and token merge orders publicly available.

2.2 Membership Inference

The concept of membership inference attacks was first introduced by Shokri et al. [74], who demonstrated that an adversary can determine whether a specific data record was included in a model’s training set. Specifically, they propose to train multiple shadow models that imitate the behavior of the target model. By comparing the output distributions of shadow models trained with and without a specific data record, the adversary can infer whether the data record was part of the target model’s training data [58].

Building upon this foundational research, subsequent studies [12, 13, 22] have investigated the effectiveness of MIAs on a variety of machine learning models, including ResNet-18 [33] and BERT [20]. Nonetheless, as models increase in size and are trained on larger datasets over fewer epochs, the overfitting signal for individual samples decreases, resulting in reduced MIA performance on LLMs [23]. To address this limitation, recent MIAs [35] instead focus on dataset membership, which aggregates signals from individual samples to enhance the detection of membership.

3 Threat Models

Attack Scenario. We consider an attack scenario in which an adversary implements MIAs to infer the pretraining datasets of an LLM. As demonstrated in Table 1, the tokenizer for such an LLM is trained on the model’s pretraining data. The adversary exploits the tokenizer as an attack vector for MIAs. Specifically, it downloads the tokenizer from official links, which were originally intended to support transparent billing for the service. Then, the adversary captures the distinctive

tokens in the tokenizer to calculate the membership signal and infer the final result. Unlike previous MIAs against LLMs which directly reflect the model’s memorization of pretraining data [34], MIAs on tokenizers lead to privacy implications, as illustrated by the case below:

Case for Privacy Infringement. A distinctive token ‘davidjl’ from a Reddit user’s data appears in OpenAI’s tokenizers [30, 52]. Indeed, OpenAI claims that its tokenizers were trained on models’ pretraining corpus, which includes Reddit data [10, 63]. By leveraging distinctive tokens, MIAs on such tokenizers can infer whether users’ data was used to train the tokenizers, thereby revealing the models’ pretraining corpus and identifying privacy infringements.

Adversary’s Objective. Given a set of pretraining datasets \mathcal{D}_{mem} sampled from an underlying distribution \mathbb{D} (denoted as $\mathcal{D}_{\text{mem}} \leftarrow \mathbb{D}$), we write $\mathcal{V}_{\text{target}} \leftarrow \mathcal{T}(\mathcal{D}_{\text{mem}})$ to represent a tokenizer’s vocabulary $\mathcal{V}_{\text{target}}$ is trained by running the BPE algorithm \mathcal{T} [26] on \mathcal{D}_{mem} . This training process results in the target tokenizer $f_{\mathcal{V}_{\text{target}}}$. Given a target dataset $D \in \mathbb{D}$, the adversary’s objective is to determine whether D was part of the pretraining data used to construct the vocabulary $\mathcal{V}_{\text{target}}$. To achieve this, the adversary employs a membership inference attack \mathcal{A} , which can be formally defined as:

$$\mathcal{A}: D, f_{\mathcal{V}_{\text{target}}} \rightarrow \{0, 1\}, \quad (1)$$

where 1 indicates $D \in \mathcal{D}_{\text{mem}}$, and 0 indicates $D \notin \mathcal{D}_{\text{mem}}$.

Adversary’s Capabilities. In alignment with the attack scenario, where commercial LLMs open-source their tokenizers [28, 60], we assume that the adversary has access to the target tokenizer $f_{\mathcal{V}_{\text{target}}}$ and its token vocabulary $\mathcal{V}_{\text{target}} = \{t_1, t_2, \dots, t_{|\mathcal{V}_{\text{target}}|}\}$, where each token $t_i \in \mathcal{V}_{\text{target}}$ was merged at iteration i during the training process. Furthermore, we assume that the adversary is able to sample auxiliary datasets \mathcal{D}_{aux} from the same distribution as the training data used by the target tokenizer, i.e., $\mathcal{D}_{\text{aux}} \leftarrow \mathbb{D}$. Leveraging datasets \mathcal{D}_{aux} , the adversary can use the BPE algorithm \mathcal{T} to train shadow tokenizers. This assumption is consistent with previous work [11, 45, 69]. It is also realistic in practice, as the training data is primarily sourced from web content [29].

4 Attack Methodology

In this section, we present our MIAs against pre-trained LLMs. For each method, we start by introducing our design intuition. Then we describe the attack methodology.

4.1 Baseline: MIA via Merge Similarity

Shadow-based MIAs [32, 41, 45, 69, 83] involve training auxiliary models to calibrate predictions. Inspired by these attacks, we formalize MIA via Merge Similarity on tokenizers.

Design Intuition. Prior work [11] has revealed that the overfitting behavior of machine learning models can vary

depending on whether a particular data point was present in the training data. Based on this insight, we hypothesize that tokenizers may also differ depending on whether a dataset was included in the training data. Specifically, we assume that the token merge order can serve as an indicator of the overfitting phenomenon. Thus, merge orders in vocabularies $\mathbb{V}_{\text{in}} = \{\mathcal{V}_{\text{in}} \leftarrow \mathcal{T}(\mathcal{D}_{\text{aux}} \cup \{D\}) \mid \mathcal{D}_{\text{aux}} \leftarrow \mathbb{D}\}$ and $\mathbb{V}_{\text{out}} = \{\mathcal{V}_{\text{out}} \leftarrow \mathcal{T}(\mathcal{D}_{\text{aux}} \setminus \{D\}) \mid \mathcal{D}_{\text{aux}} \leftarrow \mathbb{D}\}$ can differ depending on whether the target dataset D was included in the training data. Building on this hypothesis, an adversary can exploit this difference by comparing the similarity ρ of token merge orders for pairs $(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{target}})$ and $(\mathcal{V}_{\text{out}}, \mathcal{V}_{\text{target}})$. If the average value for $\rho(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{target}})$ is higher than that of $\rho(\mathcal{V}_{\text{out}}, \mathcal{V}_{\text{target}})$, it is more likely that $D \in \mathcal{D}_{\text{mem}}$.

Attack Method. This attack consists of four steps.

- (i) The adversary randomly samples a collection of datasets $\mathcal{D}_{\text{aux}} \leftarrow \mathbb{D}$ for N times, and trains N shadow tokenizers considering inclusion or exclusion of the target dataset D . Thus, the adversary obtains sets \mathbb{V}_{in} and \mathbb{V}_{out} .
- (ii) The adversary computes the similarity of token merge orders for each $\rho(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{target}})$ and $\rho(\mathcal{V}_{\text{out}}, \mathcal{V}_{\text{target}})$ via Spearman’s rank correlation coefficient [71].
- (iii) The membership signal for target dataset D is defined as:

$$\frac{1}{2} + \frac{\sum_{\mathcal{V}_{\text{in}} \in \mathbb{V}_{\text{in}}} \rho(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{target}})}{4|\mathbb{V}_{\text{in}}|} - \frac{\sum_{\mathcal{V}_{\text{out}} \in \mathbb{V}_{\text{out}}} \rho(\mathcal{V}_{\text{out}}, \mathcal{V}_{\text{target}})}{4|\mathbb{V}_{\text{out}}|}, \quad (2)$$

where it ranges from 0 to 1.

- (iv) If the membership signal is larger than a decision-making threshold τ , output 1 (*member*). Otherwise, output 0.

We conduct validation experiments for this attack using real-world Internet data [64] (see Section 5.2 for detailed experiments). However, the results demonstrate unsatisfactory performance of MIA via Merge Similarity in distinguishing between *members* and *non-members*. These results are probably due to the overall distributions of token merge orders in \mathbb{V}_{in} and \mathbb{V}_{out} resembling each other, as illustrated in Figure 2. The minor discrepancies observed between these distributions suggest weak overfitting signals from the perspective of global tokens. Consequently, the correlation values of Spearman’s rank $\rho(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{target}})$ and $\rho(\mathcal{V}_{\text{out}}, \mathcal{V}_{\text{target}})$ remain highly similar, making it hard for to infer the membership of the dataset D .

4.2 Improved MIA via Vocabulary Overlap

Building on the observation that global token distribution can obscure overfitting signals from the target dataset D , we shift our focus to a more fine-grained analysis. Specifically, we solely examine those distinctive tokens whose merge index differs between the vocabularies \mathbb{V}_{in} and \mathbb{V}_{out} . Our analysis suggests that only when the tokenizer is trained on dataset

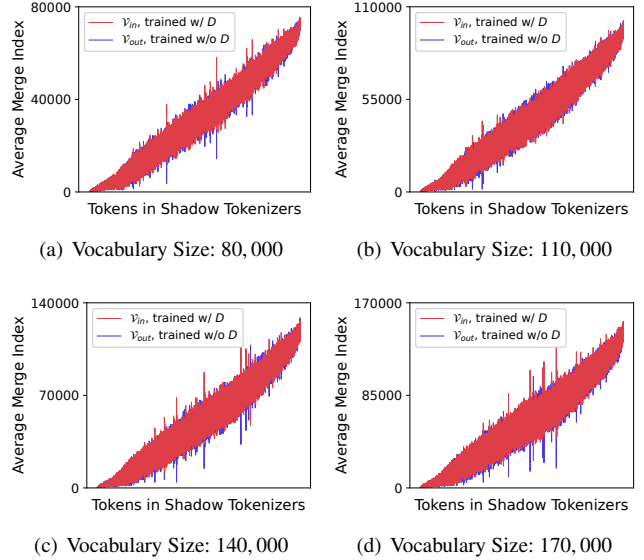


Figure 2: Average merge index for tokens in \mathbb{V}_{in} and \mathbb{V}_{out} . It is shown that overall merge orders in \mathbb{V}_{in} and \mathbb{V}_{out} resemble.

D , some distinctive tokens in D are more likely to be overfit in its vocabulary. Typically, these distinctive tokens more frequently appear in \mathbb{V}_{in} , but are seldom found in \mathbb{V}_{out} . As a result, there exist minor discrepancies between the vocabularies \mathbb{V}_{in} and \mathbb{V}_{out} . Leveraging this insight, we propose an improved approach, MIA via Vocabulary Overlap.

Design Intuition. When a target tokenizer $f_{\mathcal{V}_{\text{target}}}$ is trained on a target dataset D , its vocabulary $\mathcal{V}_{\text{target}}$ is likely to overfit the distinctive tokens present in dataset D . In fact, existing analysis has shown that OpenAI’s tokenizer contains tokens unique to the Reddit forum [30, 52]. Building on this, we hypothesize that: the more distinctive tokens from D that are found in $\mathcal{V}_{\text{target}}$, the more likely it is that $\mathcal{V}_{\text{target}}$ was trained on D . To quantify the overlap of distinctive tokens, one effective approach is to use the Jaccard index, which measures the similarity between two sets by focusing on the presence of shared elements. Specifically, an adversary can exploit this by computing the Vocabulary Overlap using Jaccard index J for pairs $(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{target}})$ and $(\mathcal{V}_{\text{out}}, \mathcal{V}_{\text{target}})$ in terms of these distinctive tokens. We write $\mathcal{V}_{\text{non}} = (\cup_{\mathcal{V}_{\text{in}} \in \mathbb{V}_{\text{in}}} \mathcal{V}_{\text{in}}) \cap (\cup_{\mathcal{V}_{\text{in}} \in \mathbb{V}_{\text{out}}} \mathcal{V}_{\text{out}})$ to denote the set of non-distinctive tokens. If the average value for $J(\mathcal{V}_{\text{in}} \setminus \mathcal{V}_{\text{non}}, \mathcal{V}_{\text{target}} \setminus \mathcal{V}_{\text{non}})$ is higher than that for $J(\mathcal{V}_{\text{out}} \setminus \mathcal{V}_{\text{non}}, \mathcal{V}_{\text{target}} \setminus \mathcal{V}_{\text{non}})$, it is more likely $D \in \mathcal{D}_{\text{mem}}$.

Attack Method. We structure this attack in five steps.

- (i) The adversary randomly samples a collection of datasets $\mathcal{D}_{\text{aux}} \leftarrow \mathbb{D}$ for N times, and trains N shadow tokenizers considering inclusion or exclusion of the target dataset D . This process results in vocabulary sets \mathbb{V}_{in} and \mathbb{V}_{out} .

Algorithm 1 MIA via Vocabulary Overlap. We train shadow tokenizers with and without the target dataset D , filter out non-distinctive tokens, and compute the membership signal. If the signal is larger than a threshold τ , the dataset D is inferred as a member. Otherwise, it is inferred as a non-member.

Input: Target dataset D , vocabulary of target tokenizer $\mathcal{V}_{\text{target}}$, underlying distribution \mathbb{D} , number of shadow tokenizers N , BPE algorithm \mathcal{T} , threshold τ

- 1: $\mathbb{V}_{\text{in}} \leftarrow \{\}$, $\mathbb{V}_{\text{out}} \leftarrow \{\}$
- 2: # Step 1: Train N shadow tokenizers
- 3: **for** N times **do**
- 4: $\mathcal{D}_{\text{aux}} \leftarrow \mathbb{D}$ ▷ randomly sample auxiliary datasets
- 5: $\mathcal{V}'_{\text{in}} \leftarrow \mathcal{T}(\mathcal{D}_{\text{aux}} \cup \{D\})$ ▷ train IN tokenizer
- 6: $\mathbb{V}_{\text{in}} \leftarrow \mathbb{V}_{\text{in}} \cup \{\mathcal{V}'_{\text{in}}\}$
- 7: $\mathcal{V}'_{\text{out}} \leftarrow \mathcal{T}(\mathcal{D}_{\text{aux}} \setminus \{D\})$ ▷ train OUT tokenizer
- 8: $\mathbb{V}_{\text{out}} \leftarrow \mathbb{V}_{\text{out}} \cup \{\mathcal{V}'_{\text{out}}\}$
- 9: **end for**
- 10: # Step 2: Compute non-distinctive tokens
- 11: $\mathcal{V}_{\text{non}} \leftarrow (\cup_{\mathcal{V}'_{\text{in}} \in \mathbb{V}_{\text{in}}} \mathcal{V}'_{\text{in}}) \cap (\cup_{\mathcal{V}'_{\text{out}} \in \mathbb{V}_{\text{out}}} \mathcal{V}'_{\text{out}})$
- 12: $J_{\text{in}} \leftarrow 0$, $J_{\text{out}} \leftarrow 0$, $\tilde{\mathcal{V}}_{\text{target}} \leftarrow \mathcal{V}_{\text{target}} \setminus \mathcal{V}_{\text{non}}$
- 13: # Step 3: Calculate Jaccard index
- 14: **for each** $\mathcal{V}'_{\text{in}} \in \mathbb{V}_{\text{in}}$ **do**
- 15: $\tilde{\mathcal{V}}_{\text{in}} \leftarrow \mathcal{V}'_{\text{in}} \setminus \mathcal{V}_{\text{non}}$ ▷ filter non-distinctive tokens in \mathcal{V}'_{in}
- 16: $J_{\text{in}} \leftarrow J_{\text{in}} + \frac{|\tilde{\mathcal{V}}_{\text{in}} \cap \tilde{\mathcal{V}}_{\text{target}}|}{|\tilde{\mathcal{V}}_{\text{in}} \cup \tilde{\mathcal{V}}_{\text{target}}|}$ ▷ sum Jaccard index in \mathbb{V}_{in}
- 17: **end for**
- 18: **for each** $\mathcal{V}'_{\text{out}} \in \mathbb{V}_{\text{out}}$ **do**
- 19: $\tilde{\mathcal{V}}_{\text{out}} \leftarrow \mathcal{V}'_{\text{out}} \setminus \mathcal{V}_{\text{non}}$ ▷ filter non-distinctive tokens in $\mathcal{V}'_{\text{out}}$
- 20: $J_{\text{out}} \leftarrow J_{\text{out}} + \frac{|\tilde{\mathcal{V}}_{\text{out}} \cap \tilde{\mathcal{V}}_{\text{target}}|}{|\tilde{\mathcal{V}}_{\text{out}} \cup \tilde{\mathcal{V}}_{\text{target}}|}$ ▷ sum Jaccard index in \mathbb{V}_{out}
- 21: **end for**
- 22: # Step 4: Compute membership signal
- 23: $\text{SIGNAL} \leftarrow \frac{1}{2} + \frac{J_{\text{in}}}{2|\mathbb{V}_{\text{in}}|} - \frac{J_{\text{out}}}{2|\mathbb{V}_{\text{out}}|}$
- 24: # Step 5: Infer the membership
- 25: **return** $\mathbb{1}[\text{SIGNAL} > \tau]$

(ii) The adversary computes the non-distinctive tokens as:

$$\mathcal{V}_{\text{non}} = \left(\bigcup_{\mathcal{V}'_{\text{in}} \in \mathbb{V}_{\text{in}}} \mathcal{V}'_{\text{in}} \right) \cap \left(\bigcup_{\mathcal{V}'_{\text{out}} \in \mathbb{V}_{\text{out}}} \mathcal{V}'_{\text{out}} \right). \quad (3)$$

(iii) The adversary calculates the overfitting signals using the Jaccard index [6] for each $J(\mathcal{V}'_{\text{in}} \setminus \mathcal{V}_{\text{non}}, \mathcal{V}'_{\text{target}} \setminus \mathcal{V}_{\text{non}})$ and $J(\mathcal{V}'_{\text{out}} \setminus \mathcal{V}_{\text{non}}, \mathcal{V}'_{\text{target}} \setminus \mathcal{V}_{\text{non}})$.

(iv) The membership signal for dataset D is defined as:

$$\frac{1}{2} + \frac{\sum_{\mathcal{V}'_{\text{in}} \in \mathbb{V}_{\text{in}}} J(\mathcal{V}'_{\text{in}} \setminus \mathcal{V}_{\text{non}}, \mathcal{V}'_{\text{target}} \setminus \mathcal{V}_{\text{non}})}{2|\mathbb{V}_{\text{in}}|} - \frac{\sum_{\mathcal{V}'_{\text{out}} \in \mathbb{V}_{\text{out}}} J(\mathcal{V}'_{\text{out}} \setminus \mathcal{V}_{\text{non}}, \mathcal{V}'_{\text{target}} \setminus \mathcal{V}_{\text{non}})}{2|\mathbb{V}_{\text{out}}|}, \quad (4)$$

where it ranges from 0 to 1.

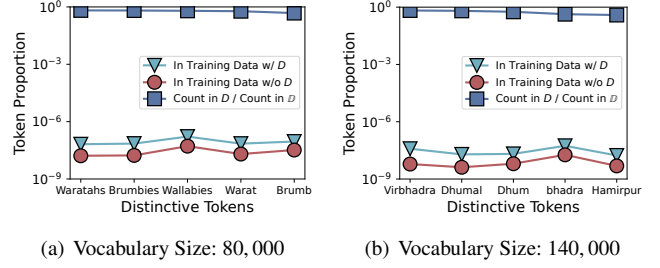


Figure 3: Distinctive tokens in MIA via Vocabulary Overlap.

(v) If the membership signal is larger than a decision-making threshold τ , output 1 (*member*). Otherwise, output 0.

The detailed process of this attack is outlined in Algorithm 1. However, like other shadow-based MIAs [11, 74, 82], we find that MIA via Vocabulary Overlap requires multiple shadow tokenizers (e.g., 96) to effectively capture membership signals. As a result, training such a large number of shadow tokenizers incurs a substantial time cost.

4.3 Efficient MIA via Frequency Estimation

MIA via Vocabulary Overlap raises a natural question: Can we design an attack that relies on fewer shadow tokenizers and thus reduces the overall time cost? To address this, we investigate whether it is possible to identify distinctive tokens directly by analyzing their statistical characteristics. Motivated by this, we examine such distinctive tokens and derive two key insights: ❶ The distinctive tokens of dataset D appear infrequently in the training data of the tokenizer trained on D . ❷ As shown in Figure 3 and Figure 12, the majority of occurrences of these distinctive tokens in the underlying distribution \mathbb{D} are found in the dataset D . Given these characteristics, if dataset D is excluded from the tokenizer’s training data, the frequency of such distinctive tokens becomes lower. As a result, these tokens with low frequency are unlikely to be merged into the vocabulary during tokenizer training, since BPE primarily merges the most frequent tokens. This observation suggests that including dataset D in the training data is almost a necessary condition for some tokens to be merged into the tokenizer’s vocabulary. Motivated by these insights, we introduce the MIA via Frequency Estimation.

Design Intuition. It is hypothesized that tokenizer training probably exhibits overfitting by incorporating distinctive tokens from the training datasets into its vocabulary [30, 52]. Building on this intuition, an adversary could exploit such overfitting by evaluating whether including dataset D in the training data is necessary for the merging of some tokens in vocabulary $\mathcal{V}_{\text{target}}$. If the presence of such distinctive tokens in $\mathcal{V}_{\text{target}}$ strongly depends on dataset D , it is likely $D \in \mathcal{D}_{\text{mem}}$.

Necessity Evaluation. However, no existing metric evaluates

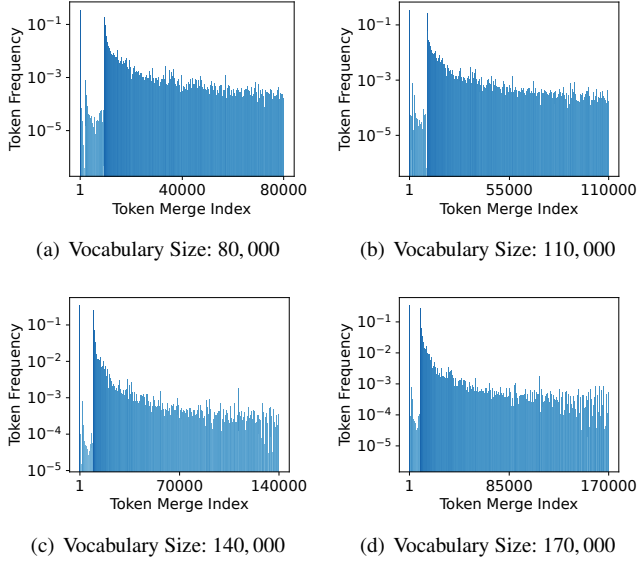


Figure 4: Relationship between token merge index and frequency in training data, indicating they follow a power law.

this necessity. To fill this gap, we introduce a new metric: Relative Token Frequency with Self-information (RTF-SI).

Definition 4.1 (Relative Token Frequency with Self-information). Let \mathbb{D} denote a data distribution. Given a dataset $D \subseteq \mathbb{D}$ and a target tokenizer’s vocabulary $\mathcal{V}_{\text{target}}$, the Relative Token Frequency with Self-Information (RTF-SI) of a token $t_i \in \mathcal{V}_{\text{target}}$ in D is defined as:

$$\text{RTF-SI}(D, t_i, \mathcal{V}_{\text{target}}) := \text{RTF}(t_i, D) \cdot \text{SI}(t_i, \mathcal{V}_{\text{target}}), \quad (5)$$

where the relative token frequency (RTF) is calculated as:

$$\text{RTF}(t_i, D) = \frac{n_D(t_i)}{\sum_{D' \in \mathbb{D}} n_{D'}(t_i)}, \quad (6)$$

with $n_D(t_i)$ denoting the count of token t_i in the dataset D . The self-information (SI) is given by:

$$\text{SI}(t_i, \mathcal{V}_{\text{target}}) = -\log \Pr(t_i | \mathcal{V}_{\text{target}}), \quad (7)$$

where $\Pr(t_i | \mathcal{V}_{\text{target}})$ is the frequency of token t_i appearing in the training data \mathcal{D}_{mem} associated with the vocabulary $\mathcal{V}_{\text{target}}$. Ideally, this probability is computed as:

$$\Pr(t_i | \mathcal{V}_{\text{target}}) = \frac{\sum_{D' \in \mathcal{D}_{\text{mem}}} n_{D'}(t_i)}{\sum_{t' \in \mathcal{V}_{\text{target}}} \sum_{D' \in \mathcal{D}_{\text{mem}}} n_{D'}(t')}. \quad (8)$$

RTF-SI evaluates whether it is necessary to include dataset D in the training data for constructing the target vocabulary, $\mathcal{V}_{\text{target}}$. Building on the classic TF-IDF definition [4, 65, 76], RTF-SI modifies the normalization used in the TF component. A high RTF-SI for the target dataset D suggests that: ❶ Token

Table 2: Power-law fit on token frequency in training data.

| $ \mathcal{V}'_{\text{target}} $ | 80,000 | 110,000 | 140,000 | 170,000 | 200,000 |
|----------------------------------|--------|---------|---------|---------|---------|
| x_{\min} | 9,782 | 9,782 | 9,782 | 9,782 | 9,782 |
| α | 1.717 | 1.604 | 1.537 | 1.493 | 1.460 |
| Std. Error | 0.003 | 0.002 | 0.001 | 0.001 | 0.001 |

t_i is with high self-information, i.e., it appears infrequently in the training data of the target tokenizer. ❷ Dataset D contributes the majority of token t_i ’s relative frequency in the underlying distribution \mathbb{D} . As a result, the absence of dataset D may significantly affect some merged tokens in vocabulary $\mathcal{V}_{\text{target}}$, indicating that $D \in \mathcal{D}_{\text{mem}}$.

Frequency Estimation. In practical implementation, an adversary can estimate the RTF component using auxiliary datasets $\mathcal{D}_{\text{aux}} \leftarrow \mathbb{D}$. However, the SI component is not directly observable, as the frequency $\Pr(t_i | \mathcal{V}'_{\text{target}})$ in the training data is not available. To estimate this, we draw on the power law [15], which has been widely applied to approximate word frequency. According to the power law, when a list of measured values exceeding a threshold $x_{\min} \in \mathbb{Z}_{>0}$ is sorted in decreasing order, the n -th value is approximately proportional to $1/n^\alpha$, where $\alpha \in \mathbb{R}_{>0}$ is a constant. As shown in Figure 4, there is a power-law relationship between token merge order and frequency. For rigorous verification, we fit the frequency $\Pr(t_i | \mathcal{V}'_{\text{target}})$ in a power-law distribution [15]:

$$\Pr(t_i | \mathcal{V}'_{\text{target}}) \propto \frac{1}{i^\alpha}, \text{ where } t_i \in \mathcal{V}'_{\text{target}} \text{ and } i > x_{\min}. \quad (9)$$

The estimation results in Table 2 show a small standard error between the estimated and actual values, which supports using the power law to approximate the value $\Pr(t_i | \mathcal{V})$ in the SI component. Given this, RTF-SI can also be computed as:

Theorem 4.2 (RTF-SI under the Power Law). Under the power-law distribution [15], the frequency $\Pr(t_i | \mathcal{V}'_{\text{target}})$ of a token $t_i \in \mathcal{V}'_{\text{target}}$ is proportional to $1/i^\alpha$:

$$\Pr(t_i | \mathcal{V}'_{\text{target}}) \propto \frac{1}{i^\alpha}, \quad (10)$$

where $i > x_{\min}$, and $\alpha \in \mathbb{R}_{>0}$, $x_{\min} \in \mathbb{Z}_{>0}$ are constants defined by the power law. Then, RTF-SI can be approximated by its lower bound:

$$\text{RTF-SI}(D, t_i, \mathcal{V}'_{\text{target}}) \geq \frac{n_D(t_i)}{\sum_{D' \in \mathbb{D}} n_{D'}(t_i)} \cdot \log \left(\sum_{j=x_{\min}+1}^{|\mathcal{V}'_{\text{target}}|} i^\alpha / j^\alpha \right). \quad (11)$$

The detailed proof of Theorem 4.2 is provided in Appendix A. The power law allows an adversary to estimate RTF-SI without directly accessing the frequency $\Pr(t_i | \mathcal{V}'_{\text{target}})$. Since the power law estimates the frequency of tokens $t_i \in \mathcal{V}'_{\text{target}}$ with merge index $i > x_{\min}$, MIA via Frequency Estimation also concentrates on them.

Attack Method. We outline this attack in four steps below.

Algorithm 2 MIA via Frequency Estimation. We train a shadow tokenizer to fit the power-law distribution of token frequency, approximate RTF-SI for each token $t_i \in \mathcal{V}_{\text{target}}$ where $i > x_{\min}$, and compute the membership signal based on the maximum RTF-SI. If the membership signal is larger than a decision-making threshold τ , the dataset D is inferred as a member. Otherwise, it is inferred as a non-member.

Input: Target dataset D , vocabulary of target tokenizer $\mathcal{V}_{\text{target}}$, underlying distribution \mathbb{D} , sampling times N , BPE algorithm \mathcal{T} , power-law fit function `pl.fit`, threshold τ

- 1: $\tilde{\mathbb{D}} \leftarrow \{\}$
- 2: *# Step 1: Prepare for frequency estimation*
- 3: **for** N times **do**
- 4: $\mathcal{D}_{\text{aux}} \leftarrow^S \mathbb{D}$ *▷ randomly sample auxiliary datasets*
- 5: $\tilde{\mathbb{D}} \leftarrow \tilde{\mathbb{D}} \cup \mathcal{D}_{\text{aux}}$
- 6: **end for**
- 7: $\mathcal{V}_{\text{shadow}} \leftarrow \mathcal{T}(\mathcal{D}_{\text{aux}})$ *▷ train shadow tokenizer*
- 8: *# Step 2: Estimate components in RTF-SI*
- 9: $\alpha, x_{\min} \leftarrow \text{pl.fit}(\mathcal{V}_{\text{shadow}}, \mathcal{D}_{\text{aux}})$ *▷ fit token frequency*
- 10: **for** $i = x_{\min} + 1$ to $|\mathcal{V}_{\text{target}}|$ **do**
- 11: $\text{RTF}(D, t_i) \leftarrow \frac{n_D(t_i)}{\sum_{D' \in \tilde{\mathbb{D}} \cup \{D\}} n_{D'}(t_i)}$
- 12: $\text{SI}(t_i, \mathcal{V}_{\text{target}}) \leftarrow \log\left(\sum_{j=x_{\min}+1}^{|\mathcal{V}_{\text{target}}|} \frac{i^\alpha}{j^\alpha}\right)$ *▷ apply Theorem 4.2*
- 13: **end for**
- 14: *# Step 3: Compute membership signal*
- 15: **for** $i = x_{\min} + 1$ to $|\mathcal{V}_{\text{target}}|$ **do**
- 16: $\text{RTF-SI}(D, t_i, \mathcal{V}_{\text{target}}) \leftarrow \text{RTF}(t_i, D) \cdot \text{SI}(t_i, \mathcal{V}_{\text{target}})$
- 17: **end for**
- 18: $\text{RTF-SI}_{\max} \leftarrow \max_{t_i \in \mathcal{V}_{\text{target}}, i > x_{\min}} \text{RTF-SI}(D, t_i, \mathcal{V}_{\text{target}})$
- 19: $\text{SIGNAL} \leftarrow \frac{1}{1 + e^{-\text{RTF-SI}_{\max}}}$ *▷ normalize by sigmoid*
- 20: *# Step 4: Infer the membership*
- 21: **return** $\mathbb{1}[\text{SIGNAL} > \tau]$

- (i) The adversary randomly samples a collection of datasets $\mathcal{D}_{\text{aux}} \leftarrow \mathbb{D}$ N times, comprising a set $\tilde{\mathbb{D}}$. Then, the adversary trains a shadow tokenizer $f_{\mathcal{V}_{\text{shadow}}}$ using a $\mathcal{D}_{\text{aux}} \subseteq \tilde{\mathbb{D}}$.
- (ii) The adversary fits the power-law distribution in Equation 9 using the vocabulary $\mathcal{V}_{\text{shadow}}$ and its training data. For each token $t_i \in \mathcal{V}_{\text{target}}$ where $i > x_{\min}$, the adversary approximate its RTF(D, t_i) on set $\tilde{\mathbb{D}} \cup \{D\}$, and estimate its SI($t_i, \mathcal{V}_{\text{target}}$) via the fitted power-law distribution.
- (iii) Based on the Theorem 4.2, the membership signal for target dataset D is defined as follows:

$$\sigma\left(\max_{t_i \in \mathcal{V}_{\text{target}}, i > x_{\min}} \underbrace{\frac{n_D(t_i)}{\sum_{D' \in \tilde{\mathbb{D}} \cup \{D\}} n_{D'}(t_i)}}_{\text{RTF}(D, t_i)} \cdot \underbrace{\log\left(\sum_{j=x_{\min}+1}^{|\mathcal{V}_{\text{target}}|} \frac{i^\alpha}{j^\alpha}\right)}_{\text{SI}(t_i, \mathcal{V}_{\text{target}})}\right), \quad (12)$$

where σ denotes the sigmoid function [67]. Thereby, the

value of Equation 12 ranges from 0 to 1.

- (iv) If the membership signal is larger than a decision-making threshold τ , output 1 (*member*). Otherwise, output 0.

The detailed process of this attack is shown in Algorithm 2. The membership signal for MIA via Frequency Estimation is defined as the maximum RTF-SI value for the target dataset D . Specifically, if including dataset D in the training data is necessary for at least one token t_i to be merged into the vocabulary $\mathcal{V}_{\text{target}}$, it suggests that the absence of dataset D has a significant influence on the already constructed vocabulary. Consequently, it is likely that dataset D is a *member*.

5 Attack Evaluation

In this section, we first introduce the experimental settings. Next, we develop two shadow-free membership inference methods to serve as additional exploration and baselines. Finally, we present the evaluation results.

5.1 Experimental Setup

Datasets. According to disclosures from existing LLMs [9, 29, 80], the training data for tokenizers is primarily sourced from publicly available web content. To ensure a realistic evaluation of our attacks, we therefore utilize real-world web data from the *C4* corpus [64] in our evaluation. Specifically, the *C4* corpus is created by cleaning and filtering web pages from Common Crawl, widely used for training and evaluating natural language processing models. In our evaluation, we utilize 1,681,296 web pages across 4,133 websites (i.e., \mathbb{D}) from the *C4* corpus, with each website treated as a dataset D .

Tokenizer Training. Following prior work [11, 74], we randomly select half of the datasets in \mathbb{D} to serve as training data. Consistent with DeepSeek [7], we train the target tokenizers using the HuggingFace library. The tokenizer vocabulary size ranges from 80,000 to 200,000 tokens, with the upper bound matching that of OpenAI’s latest tokenizer, o200k [60]. For MIA via Merge Similarity and MIA via Vocabulary Overlap, the adversary trains 96 shadow tokenizers. For MIA via Frequency Estimation, the adversary samples the auxiliary datasets 10 times to compose a set $\tilde{\mathbb{D}}$.

Verification of No Distribution Shifts. As discussed in previous works [19, 23, 55], distribution shifts between *members* and *non-members* can invalidate the evaluation of MIAs. In such scenarios, an evaluator [19, 55] can exploit bag-of-words features extracted from test samples and train a random forest classifier to detect whether a sample was part of the training data, even without access to the target model. To ensure that there is no distribution shift in our evaluation, we follow the methodology of prior work [19] by training a random forest classifier and leveraging the bag-of-words features to distinguish between *members* and *non-members*. The experimental




Figure 5: Comparison of tokenizer utility based on the metric of bytes per token. Specifically, “Ours-80000” refers to our trained target tokenizer with a vocabulary size of 80,000 tokens. The above experimental results indicate that the utility performance of the target tokenizers utilized in our evaluations is comparable to that of tokenizers used in state-of-the-art LLMs [28, 29, 60].

Visualization based on Bag-of-words (AUC = 0.513)

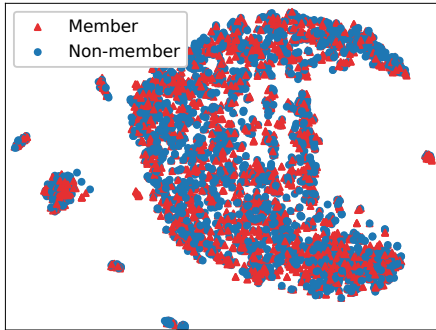


Figure 6: Visualization of the test set using t-SNE [49]. The results confirm no distribution shifts in our evaluation.

results are illustrated in Figure 6. The AUC score of 0.513 via this approach to distinguish *members* and *non-members* confirms the absence of distribution shifts in our evaluations.

Comparison to Commercial Tokenizers. To assess how well our trained target tokenizers in experimental settings mimic the commercial tokenizers in the real world, we compare their utility performance using a standard compression metric: the bytes per token [46]. Specifically, following prior work [17, 75], we compute the ratio of UTF-8 bytes in a given text to the number of tokens generated by the tokenizer. A higher score is desired. We conduct this evaluation across widely used benchmarks: general English (*WikiText-103* [56]), code (*GitHub Code* [40]), multilingual content (*MGSM* [70]), and mathematics (*GPQA* [66]). As shown in Figure 5, the utility performance of our trained tokenizers is comparable to that of the commercial tokenizers [28, 29, 60, 80].

Attack Baselines. As our work is the first to investigate MIAs targeting tokenizers, there are no existing baselines from prior studies. Therefore, we establish our own baselines by comparing MIA via Vocabulary Overlap and MIA via Frequency Estimation with three other methods: MIA via Merge Similarity (see Section 4.1), as well as two attack meth-

ods we developed. These two attacks are described below:

- **MIA via Naive Bayes.** Since every token originates from at least one of the tokenizer’s training datasets, the adversary can approximate the empirical probability $\Pr(t_i \rightarrow D)$ that the token $t_i \in \mathcal{V}_{\text{target}}$ comes from the dataset D . Specifically, this probability $\Pr(t_i \rightarrow D)$ can be represented by:

$$\Pr(t_i \rightarrow D) = \frac{n_D(t_i)}{\sum_{D' \in \mathbb{D} \cup \{D\}} n_{D'}(t_i)}, \quad (13)$$

where $n_D(t_i)$ is the frequency of t_i in D , and \mathbb{D} is constructed by sampling auxiliary datasets \mathcal{D}_{aux} N times as an estimation for distribution \mathbb{D} . We default to setting $N = 10$ in evaluation. The probability that $D \in \mathcal{D}_{\text{mem}}$ is given by:

$$\Pr(D \in \mathcal{D}_{\text{mem}}) = 1 - \Pr\left(\bigcap_{t_i \in \mathcal{V}_{\text{target}}} t_i \not\rightarrow D\right), \quad (14)$$

where $t_i \not\rightarrow D$ means t_i was not sourced from D . Assuming rare tokens in training data typically come from disjoint datasets, the probability $\Pr(D \in \mathcal{D}_{\text{mem}})$ thus can be approximated via the Naive Bayes [81]:

$$\Pr(D \in \mathcal{D}_{\text{mem}}) \approx 1 - \prod_{t_i \in \mathcal{V}_{\text{rare}}} \Pr(t_i \not\rightarrow D) \quad (15)$$

$$\approx 1 - \prod_{t_i \in \mathcal{V}_{\text{rare}}} (1 - \Pr(t_i \rightarrow D)), \quad (16)$$

where $\mathcal{V}_{\text{rare}} \subseteq \mathcal{V}_{\text{target}}$ is the set of tokens with the top k merge orders, selected as likely rare tokens in training data (see Equation 9). The membership signal for the target dataset D is defined by the probability $\Pr(D \in \mathcal{D}_{\text{mem}})$. If it is larger than a threshold τ , output 1 (*member*); or, output 0.

- **MIA via Compression Rate.** The objective of tokenizer training is to maximize the compression rate of a given text corpus [88]. Based on this optimization objective, we hypothesize that a tokenizer achieves higher compression rates on datasets it was trained on [31]. Leveraging this insight, an adversary can calculate the compression rate of a given dataset D using the metric of bytes per token [46].

Table 3: Comparison of MIAs against target tokenizers. Here, BA denotes the metric of balanced accuracy, and TPR refers to the metric of TPR @ 1.0% FPR. The bold values indicate the best performance, while the underlined values denote the second-best. It is observed that MIA via Vocabulary Overlap and MIA via Frequency Estimation outperform other baseline methods.

| Attack Approach | Shadow Tokenizers | Auxiliary Datasets | $ \mathcal{V}_{\text{target}} = 80,000$ | | | $ \mathcal{V}_{\text{target}} = 110,000$ | | | $ \mathcal{V}_{\text{target}} = 140,000$ | | | $ \mathcal{V}_{\text{target}} = 170,000$ | | | $ \mathcal{V}_{\text{target}} = 200,000$ | | |
|---------------------------|-------------------|--------------------|--|--------------|---------------|---|--------------|---------------|---|--------------|---------------|---|--------------|---------------|---|--------------|---------------|
| | | | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR |
| Compression Rate | ○ | ○ | 0.507 | 0.513 | 0.72% | 0.509 | 0.517 | 0.71% | 0.508 | 0.514 | 0.68% | 0.508 | 0.513 | 0.82% | 0.509 | 0.516 | 0.77% |
| Naive Bayes, $k = 20,000$ | ○ | ● | 0.534 | 0.526 | 3.78% | 0.526 | 0.524 | 3.44% | 0.535 | 0.533 | 4.11% | 0.546 | 0.538 | 5.86% | 0.564 | 0.551 | 8.03% |
| Naive Bayes, $k = 40,000$ | ○ | ● | 0.543 | 0.530 | 5.18% | 0.546 | 0.533 | 6.10% | 0.542 | 0.537 | 5.95% | 0.550 | 0.542 | 7.60% | 0.572 | 0.557 | 10.70% |
| Naive Bayes, $k = 60,000$ | ○ | ● | 0.543 | 0.530 | 2.22% | 0.551 | 0.538 | 5.57% | 0.543 | 0.536 | 2.80% | 0.553 | 0.545 | 4.94% | 0.572 | 0.557 | 8.86% |
| Merge Similarity | ● | ● | 0.493 | 0.509 | 1.06% | 0.494 | 0.507 | 1.02% | 0.494 | 0.508 | 0.97% | 0.495 | 0.506 | 0.92% | 0.495 | 0.508 | 0.87% |
| Vocabulary Overlap | ● | ● | 0.693 | 0.666 | 26.77% | 0.718 | 0.672 | 28.75% | 0.736 | 0.696 | 29.72% | 0.761 | 0.709 | 32.53% | 0.771 | 0.711 | 34.61% |
| Frequency Estimation | ● | ● | 0.610 | 0.614 | <u>21.30%</u> | 0.645 | 0.641 | <u>22.00%</u> | 0.676 | 0.660 | <u>22.41%</u> | 0.707 | 0.681 | <u>25.61%</u> | 0.740 | 0.705 | 27.88% |




Figure 7: Success rate of our attacks on tokenizers with different vocabulary sizes. The experimental results demonstrate that, when evaluated at low false positive rates, both MIA via Vocabulary Overlap and MIA via Frequency Estimation consistently outperform other methods. Notably, even at a 0.01% false positive rate, both attacks achieve a true positive rate of nearly 10%.

The membership signal for the target dataset D is defined by this compression rate. If it is larger than a decision-making threshold τ , output 1 (*member*); otherwise, output 0.

Evaluation Metrics. Following prior studies [11, 23, 74], we report the MIA performance using three convincing metrics:

- **AUC.** This metric quantifies the overall distinguishability of an MIA by computing the area under the receiver operating characteristic (ROC) curve [13].
- **Balanced Accuracy.** This metric (denoted as BA) measures the overall correct predictions on membership by averaging the true positive rate and the true negative rate.
- **TPR at Low FPR.** Proposed by [11], this metric evaluates the true positive rate (TPR) when the false positive rate (FPR) is low. Following prior work [34, 54], we report TPR @ 1.0% FPR in our evaluations (denoted as TPR).

5.2 Main Results

Finding 1. According to prior work [38, 53], scaling up the intelligence of LLMs involves expanding the tokenizer’s vocabulary [77] and thus improving its compression efficiency [46]. Yet, the results show it also increases a tokenizer’s vulnerability to MIAs.

Overall Performance. As shown in Table 3, the MIA via Vocabulary Overlap and the MIA via Frequency Estimation consistently demonstrate strong performance and outperform other baseline methods across different vocabulary sizes. For instance, the MIA via Vocabulary Overlap achieves an AUC score of 0.771 when evaluated on a target tokenizer with 200,000 tokens, whereas the MIA via Frequency Estimation achieves a comparable AUC score of 0.740. Beyond these results, we observe that the performance of both attacks improves as vocabulary size increases. This trend suggests that MIAs targeting tokenizers in LLMs may become even more effective as state-of-the-art models continue to scale and adopt larger vocabularies in their tokenizers [38, 53, 77]. One possible explanation is that larger vocabularies contain more tokens, which may increase the likelihood of merging the distinctive tokens from the training data. As a result, expanding the tokenizer’s vocabulary may unintentionally increase its vulnerability to effective MIAs.

ROC Analysis. The prior study [11] has highlighted the importance of MIAs being able to reliably infer even a small number of a model’s training data. To demonstrate this capability of our MIAs, Figure 7 presents the full log-scale ROC curves for various attack methods across different vocabulary sizes. It is observed that both MIA via Vocabulary Overlap

Table 4: Impact of the target dataset size on MIAs. BA denotes balanced accuracy. TPR refers to TPR @ 1.0% FPR. It is observed that MIA via Vocabulary Overlap and MIA via Frequency Estimation perform better for target datasets with larger sizes.

| Attack Approach | #Dataset Size | $ \mathcal{V}_{\text{target}} = 80,000$ | | | $ \mathcal{V}_{\text{target}} = 110,000$ | | | $ \mathcal{V}_{\text{target}} = 140,000$ | | | $ \mathcal{V}_{\text{target}} = 170,000$ | | | $ \mathcal{V}_{\text{target}} = 200,000$ | | |
|----------------------|-----------------------|--|-------|--------|---|-------|--------|---|-------|--------|---|-------|--------|---|-------|--------|
| | | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR |
| Vocabulary Overlap | $ D \in [0, 400)$ | 0.672 | 0.652 | 21.58% | 0.696 | 0.655 | 24.76% | 0.714 | 0.676 | 27.56% | 0.737 | 0.687 | 27.82% | 0.747 | 0.692 | 29.22% |
| | $ D \in [400, 800)$ | 0.739 | 0.695 | 33.73% | 0.758 | 0.718 | 41.27% | 0.791 | 0.755 | 42.77% | 0.808 | 0.761 | 43.37% | 0.808 | 0.766 | 43.67% |
| | $ D \in [800, 1200)$ | 0.773 | 0.720 | 33.75% | 0.785 | 0.757 | 43.75% | 0.797 | 0.767 | 45.00% | 0.826 | 0.813 | 47.50% | 0.882 | 0.838 | 62.50% |
| Frequency Estimation | $ D \in [0, 400)$ | 0.599 | 0.608 | 18.91% | 0.631 | 0.632 | 19.73% | 0.662 | 0.648 | 20.31% | 0.695 | 0.668 | 21.83% | 0.729 | 0.691 | 25.84% |
| | $ D \in [400, 800)$ | 0.629 | 0.624 | 23.49% | 0.683 | 0.662 | 29.27% | 0.728 | 0.697 | 30.42% | 0.747 | 0.713 | 31.63% | 0.774 | 0.736 | 32.83% |
| | $ D \in [800, 1200)$ | 0.758 | 0.734 | 33.75% | 0.772 | 0.756 | 40.00% | 0.774 | 0.761 | 41.25% | 0.814 | 0.789 | 50.00% | 0.843 | 0.810 | 53.75% |

Table 5: Time cost (hours) for training N shadow tokenizers, where each tokenizer has a vocabulary size of 200,000.

| Tokenizer Count | $N=1$ | $N=32$ | $N=64$ | $N=96$ | $N=128$ |
|-----------------|-------|--------|--------|--------|---------|
| Training Time | 0.024 | 0.731 | 1.498 | 2.251 | 3.054 |

and MIA via Frequency Estimation can reliably infer the membership of datasets, particularly in regions with low false positive rates. For example, when applied to a tokenizer with 140,000 tokens, these two attacks achieve true positive rates ranging from approximately 10% to 30% at a false positive rate below 1%. Even at a false positive rate of 0.01%, both attacks can still achieve a true positive rate of nearly 10%. Additional ROC curve results can be found in Figure 13.

Efficiency Analysis. We further analyze the computational cost for both MIA via Vocabulary Overlap and MIA via Frequency Estimation across two phases: shadow tokenizer training and the remaining inference. In the phase of shadow tokenizer training, MIA via Vocabulary Overlap trains multiple tokenizers (e.g., 96), resulting in a high computational cost. In contrast, MIA via Frequency Estimation requires training only a single tokenizer, significantly reducing the overall cost. As shown in Table 5, this leads to substantial savings in training time. In the inference phase, MIA via Vocabulary Overlap involves frequent comparisons across different tokenizers, whereas MIA via Frequency Estimation primarily estimates a power-law distribution, a much simpler computation. Table 6 confirms the shorter inference time of the latter method. For example, MIA via Vocabulary Overlap takes over two hours to infer the membership of 4,133 datasets from a tokenizer with 140,000 tokens. However, MIA via Frequency Estimation accomplishes the same task in under 20 minutes, making it efficient for large-scale attacks.

5.3 Ablation Study

Finding 2. *The membership status of the target dataset with more data samples is typically more accurately inferred by MIAs from the tokenizer.*

Impact of Dataset Size $|D|$. Recent studies [35, 50, 54] have

Table 6: Time cost (hours) for MIAs via Vocabulary Overlap and Frequency Estimation inferring 4,133 target datasets.

| Vocabulary Size | 80,000 | 110,000 | 140,000 | 170,000 | 200,000 |
|----------------------|--------|---------|---------|---------|---------|
| Vocabulary Overlap | 1.230 | 1.608 | 2.067 | 2.613 | 3.375 |
| Frequency Estimation | 0.170 | 0.235 | 0.303 | 0.373 | 0.432 |

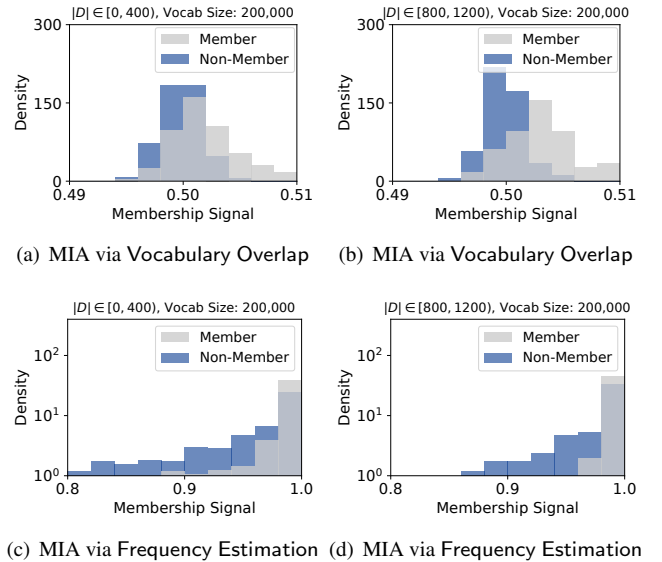


Figure 8: Distribution of *members* and *non-members*.

shown that increasing the amount of data used for membership inference can improve the attack performance. This finding is particularly relevant in the context of high-value litigation nowadays, where the datasets at stake are often massive [5, 59]. Motivated by this, we investigate whether MIAs can more effectively infer the membership of larger datasets from the tokenizers. Table 4 presents the performance of MIA via Vocabulary Overlap and MIA via Frequency Estimation for varying-size target datasets. The results show that both attacks become more effective as the dataset size increases, achieving particularly strong performance on large datasets. For instance, the MIA via Vocabulary Overlap achieves an AUC score of 0.882 on datasets containing 800 to 1,200 data samples, whereas the MIA via Frequency Estimation

Table 7: MIAs against tokenizers with the min count defense. Here, n_{\min} denotes a threshold. If a token appears fewer than n_{\min} times in training data, it is likely a distinctive token and will be excluded from the vocabulary $\mathcal{V}_{\text{target}}$. $|\mathcal{V}_{\text{target}}| \leq 80,000$ indicates the vocabulary size prior to applying defense is 80,000. BA denotes balanced accuracy. TPR refers to TPR @ 1.0% FPR.

| Attack Approach | #Min Count | $ \mathcal{V}_{\text{target}} \leq 80,000$ | | | $ \mathcal{V}_{\text{target}} \leq 110,000$ | | | $ \mathcal{V}_{\text{target}} \leq 140,000$ | | | $ \mathcal{V}_{\text{target}} \leq 170,000$ | | | $ \mathcal{V}_{\text{target}} \leq 200,000$ | | |
|----------------------|--------------------|---|-------|--------|--|-------|--------|--|-------|--------|--|-------|--------|--|-------|--------|
| | | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR |
| Vocabulary Overlap | w/o defense | 0.693 | 0.666 | 26.77% | 0.718 | 0.672 | 28.75% | 0.736 | 0.696 | 29.72% | 0.761 | 0.709 | 32.53% | 0.771 | 0.711 | 34.61% |
| | w/ $n_{\min} = 32$ | 0.663 | 0.638 | 23.57% | 0.699 | 0.657 | 23.76% | 0.718 | 0.671 | 26.48% | 0.736 | 0.686 | 28.41% | 0.746 | 0.691 | 30.83% |
| | w/ $n_{\min} = 48$ | 0.663 | 0.638 | 23.57% | 0.697 | 0.655 | 23.76% | 0.714 | 0.671 | 25.79% | 0.717 | 0.665 | 26.09% | 0.734 | 0.683 | 30.83% |
| | w/ $n_{\min} = 64$ | 0.663 | 0.638 | 21.93% | 0.685 | 0.640 | 23.57% | 0.699 | 0.657 | 23.77% | 0.707 | 0.663 | 26.48% | 0.717 | 0.671 | 26.48% |
| Frequency Estimation | w/o defense | 0.610 | 0.614 | 21.30% | 0.645 | 0.641 | 22.00% | 0.676 | 0.660 | 22.56% | 0.707 | 0.681 | 25.61% | 0.740 | 0.705 | 27.88% |
| | w/ $n_{\min} = 32$ | 0.600 | 0.602 | 18.25% | 0.633 | 0.630 | 21.30% | 0.664 | 0.647 | 22.41% | 0.695 | 0.669 | 24.15% | 0.730 | 0.693 | 28.80% |
| | w/ $n_{\min} = 48$ | 0.598 | 0.598 | 18.05% | 0.633 | 0.626 | 21.06% | 0.663 | 0.645 | 21.78% | 0.690 | 0.663 | 23.23% | 0.692 | 0.664 | 25.41% |
| | w/ $n_{\min} = 64$ | 0.596 | 0.600 | 17.13% | 0.630 | 0.624 | 20.72% | 0.661 | 0.645 | 23.81% | 0.666 | 0.648 | 22.36% | 0.668 | 0.648 | 23.72% |




Figure 9: Impact of N . Left: MIA via Vocabulary Overlap, training N shadow tokenizers. Right: MIA via Frequency Estimation, sampling auxiliary datasets N times.

achieves a competitive AUC score of 0.843. Figure 8 further illustrates the relationship between dataset size and the membership signal in both attacks. As the dataset size increases, the overlap between the membership signal distributions for *members* and *non-members* decreases, thereby enhancing the discriminative power of the attacks. This reduced overlap is likely attributable to the presence of more distinctive tokens in larger datasets, which serve to strengthen the membership signal. Additional evaluations illustrating the membership signal distributions are presented in Figure 14 and Figure 15.

Impact of Hyperparameter N . We further analyze the impact of the hyperparameter N on the performance of MIAs. Specifically, in the case of MIA via Vocabulary Overlap, an adversary trains N shadow tokenizers to capture distinctive tokens. As shown in the left plot of Figure 9, the effectiveness of this attack improves steadily as N increases. However, after training more than 112 shadow tokenizers, the performance gain begins to plateau. In another attack, MIA via Frequency Estimation, the adversary samples auxiliary datasets N times to estimate probabilities in membership inference. As shown in the right plot of Figure 9, the effectiveness of this attack also increases with larger N . Nevertheless, the performance of MIA via Frequency Estimation tends to be stable once N exceeds 6. Although the results demonstrate that both attacks benefit from increasing N , a larger value of N also incurs greater resource consumption during inference.

Table 8: Tokenizer utility measured by bytes per token. Utility scores that decrease after applying the defense are in red.

| Benchmark | $ \mathcal{V}_{\text{target}} \leq$ | Tokenizer Utility Measured by Bytes per Token | | | |
|---------------|--------------------------------------|---|--------------------|--------------------|--------------------|
| | | w/o defense | w/ $n_{\min} = 32$ | w/ $n_{\min} = 48$ | w/ $n_{\min} = 64$ |
| WikiText [56] | 80,000 | 4.873 | 4.873 (-0.0) | 4.873 (-0.0) | 4.873 (-0.0) |
| | 110,000 | 4.943 | 4.943 (-0.0) | 4.943 (-0.0) | 4.943 (-0.0) |
| | 140,000 | 4.986 | 4.986 (-0.0) | 4.986 (-0.0) | 4.986 (-0.0) |
| | 170,000 | 5.008 | 5.008 (-0.0) | 5.006 (-0.002) | 4.995 (-0.013) |
| | 200,000 | 5.025 | 5.025 (-0.0) | 5.012 (-0.013) | 5.000 (-0.025) |
| Github [40] | 80,000 | 3.740 | 3.739 (-0.001) | 3.739 (-0.001) | 3.738 (-0.002) |
| | 110,000 | 3.853 | 3.851 (-0.002) | 3.851 (-0.002) | 3.849 (-0.004) |
| | 140,000 | 3.924 | 3.921 (-0.003) | 3.921 (-0.003) | 3.918 (-0.006) |
| | 170,000 | 3.973 | 3.970 (-0.003) | 3.967 (-0.006) | 3.947 (-0.026) |
| | 200,000 | 4.009 | 4.006 (-0.003) | 3.990 (-0.019) | 3.965 (-0.044) |
| MGSM [70] | 80,000 | 3.357 | 3.356 (-0.001) | 3.356 (-0.001) | 3.355 (-0.002) |
| | 110,000 | 3.476 | 3.475 (-0.001) | 3.474 (-0.002) | 3.473 (-0.003) |
| | 140,000 | 3.601 | 3.601 (-0.0) | 3.601 (-0.0) | 3.600 (-0.001) |
| | 170,000 | 3.663 | 3.663 (-0.0) | 3.662 (-0.001) | 3.639 (-0.024) |
| | 200,000 | 3.740 | 3.740 (-0.0) | 3.739 (-0.001) | 3.716 (-0.024) |
| GPQA [66] | 80,000 | 3.925 | 3.924 (-0.001) | 3.924 (-0.001) | 3.923 (-0.002) |
| | 110,000 | 4.008 | 4.007 (-0.001) | 4.006 (-0.002) | 4.003 (-0.005) |
| | 140,000 | 4.056 | 4.055 (-0.001) | 4.053 (-0.003) | 4.050 (-0.006) |
| | 170,000 | 4.094 | 4.094 (-0.0) | 4.092 (-0.002) | 4.086 (-0.008) |
| | 200,000 | 4.117 | 4.116 (-0.001) | 4.109 (-0.008) | 4.093 (-0.024) |

5.4 Adaptive Defense


Finding 3. Removing infrequent tokens from the target tokenizer’s vocabulary can partially reduce the effectiveness of MIAs. However, this mitigation comes at the cost of reduced tokenizer utility. Moreover, MIAs remain effective when inferring large datasets.

While defense against MIAs is not the primary focus of this work, we have also explored the defense mechanism to mitigate membership leakage in tokenizers. Specifically, previous studies [35, 74, 82] have demonstrated that overfitting signals are a key requirement for the success of MIAs. This suggests that methods designed to reduce overfitting may function as effective defense mechanisms [1, 87]. Building on this insight, we assume an adaptive defender who reduces distinctive tokens in the tokenizer’s vocabulary.

Defender’s Objective. Given the target tokenizer $f_{\mathcal{V}_{\text{target}}}$, the defender’s goal is to reduce the inference accuracy of MIAs on $f_{\mathcal{V}_{\text{target}}}$, while preserving its utility as much as possible.

Table 9: Impact of the target dataset size on defense mechanism ($n_{\min} = 64$). $|\mathcal{V}_{\text{target}}| \leq 80,000$ indicates the vocabulary size prior to applying defense is 80,000. TPR refers to TPR @ 1.0% FPR. The results show our MIAs remain effective on large datasets.

| Attack Approach | #Dataset Size | $ \mathcal{V}_{\text{target}} \leq 80,000$ | | | $ \mathcal{V}_{\text{target}} \leq 110,000$ | | | $ \mathcal{V}_{\text{target}} \leq 140,000$ | | | $ \mathcal{V}_{\text{target}} \leq 170,000$ | | | $ \mathcal{V}_{\text{target}} \leq 200,000$ | | |
|----------------------|-----------------------|---|-------|--------|--|-------|--------|--|-------|--------|--|-------|--------|--|-------|--------|
| | | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR |
| Vocabulary Overlap | $ D \in [0, 400)$ | 0.646 | 0.632 | 19.16% | 0.662 | 0.636 | 21.25% | 0.677 | 0.642 | 21.25% | 0.683 | 0.647 | 23.30% | 0.694 | 0.655 | 24.06% |
| | $ D \in [400, 800)$ | 0.703 | 0.666 | 21.25% | 0.731 | 0.675 | 21.58% | 0.751 | 0.699 | 22.85% | 0.761 | 0.711 | 32.50% | 0.772 | 0.717 | 37.95% |
| | $ D \in [800, 1200)$ | 0.712 | 0.670 | 25.30% | 0.743 | 0.702 | 27.11% | 0.768 | 0.729 | 30.12% | 0.795 | 0.746 | 33.73% | 0.797 | 0.761 | 38.75% |
| Frequency Estimation | $ D \in [0, 400)$ | 0.591 | 0.598 | 17.50% | 0.620 | 0.619 | 18.20% | 0.651 | 0.636 | 20.11% | 0.654 | 0.637 | 21.07% | 0.656 | 0.639 | 21.39% |
| | $ D \in [400, 800)$ | 0.619 | 0.615 | 22.29% | 0.672 | 0.653 | 28.31% | 0.717 | 0.682 | 30.42% | 0.718 | 0.683 | 31.93% | 0.723 | 0.686 | 32.83% |
| | $ D \in [800, 1200)$ | 0.734 | 0.717 | 26.25% | 0.739 | 0.731 | 33.75% | 0.744 | 0.736 | 33.75% | 0.745 | 0.742 | 35.00% | 0.748 | 0.746 | 38.75% |




(a) MIA via Vocabulary Overlap (b) MIA via Frequency Estimation

Figure 10: Performance of MIAs on DP-based tokenizers.

Defender’s Capabilities. We assume that the defender is aware of the MIA strategy targeting the tokenizer $f_{\mathcal{V}_{\text{target}}}$, including the conditions that are important for their success. Specifically, our attacks rely on the distinctive tokens, which appear infrequently in the training data and overfit into the vocabulary $\mathcal{V}_{\text{target}}$. As a defense, the defender may modify the vocabulary $\mathcal{V}_{\text{target}}$ by identifying and removing such infrequent tokens. Thereby, it can mitigate the membership inference without significantly degrading the tokenizer utility.

Min Count Mechanism. We introduce the min count mechanism as an adaptive defense against our attacks. In this mechanism, the defender post-processes the trained vocabulary $\mathcal{V}_{\text{target}}$ by filtering infrequent tokens. Let $n_{\min} \in \mathbb{Z}_{>0}$ denote the filtering threshold, and let $n_{D'}(t_i)$ represent the count of token t_i in dataset D' . For each token $t_i \in \mathcal{V}_{\text{target}}$, if the aggregated count $\sum_{D' \in \mathcal{D}_{\text{mem}}} n_{D'}(t_i)$ across the tokenizer’s training data is less than n_{\min} , the defender removes t_i from trained vocabulary $\mathcal{V}_{\text{target}}$. Table 7 shows the results of MIAs against tokenizers with the min count mechanism. It is observed that, as the threshold n_{\min} increases, the defense can partially reduce the effectiveness of MIAs against tokenizers. However, this comes at the cost of the tokenizer’s utility. Table 8 shows that the compression efficiency of bytes per token diminishes under more strict filtering rules. While the min count mechanism can mitigate some inference risks, our MIAs remain effective, particularly for large datasets. For example, Table 9 reports an AUC of 0.797 when MIA via Vocabulary Overlap infers the datasets ranging in size from 800 to 1,200 samples.

Differentially Private Mechanism. Differential privacy (DP) [44] is an effective defense mechanism against privacy



(a) Jaccard Similarity (b) Differences in Merge Index

Figure 11: Comparison of tokenizers in real-world LLMs.

attacks. Prior study [1] has demonstrated that DP can reduce the vulnerability of membership leakage in machine learning models. Yet, no previous work has explored the DP-based tokenizer training. To fill this gap, we implement the exponential mechanism [21] of DP in tokenizer training, which merges a random token at each training step. Specifically, we define x_i as the set of all tokens at the merging step i during tokenizer training. Let $y_i \in x_i$ denote the token that DP outputs to merge. The utility score $u(x_i, y_i)$ of the exponential mechanism is calculated by the token frequency $n(y_i)$:

$$u(x_i, y_i) = \frac{n(y_i)}{\max_{y' \in x_i} n(y')} \quad (17)$$

Given a privacy parameter $\epsilon \geq 0$, the probability that the exponential mechanism in [21] merges a token y_i into the vocabulary at the step i is as follows:

$$\Pr[y_i | x_i] = \frac{e^{\epsilon u(x_i, y_i)}}{\sum_{y' \in x_i} e^{\epsilon u(x_i, y')}} \quad (18)$$

We also compose the privacy budget based on [21], in which token merging at each step satisfies 2ϵ -DP. Figure 10 presents the evaluations of DP-based tokenizers. It shows that a smaller ϵ value improves the defense but reduces the tokenizer utility.

5.5 Additional Investigations

Distinctive Tokens in Real-world Tokenizers. Our MIAs exploit distinctive tokens, which can vary across different to-

tokenizers, considering both their occurrence and their merge indices. A key question is whether real-world tokenizers differ in such ways that might enable effective MIAs. To investigate this, we compare the vocabularies of tokenizers used in real-world LLMs [28,29,60,80]. Specifically, we limit our analysis to the first 120,000 tokens of each tokenizer to enable a fair comparison across varying sizes. Figure 11(a) presents the similarity between tokenizers’ vocabularies based on the Jaccard index [6]. If there were no distinctive tokens, we would expect high similarity scores between any pair of tokenizers. However, the evaluation results show that real-world tokenizers contain a significant number of distinctive tokens, as evidenced by the maximum Jaccard index of only 0.385 observed between DeepSeek-R1 [29] and Llama-3 [80]. In addition, Figure 11(b) illustrates the average absolute differences in the merge index for tokens between any two vocabularies. If there were no significant differences between the two vocabularies, we would expect only slight variations in the token merge indices. Nevertheless, the results suggest that token merge indices also vary largely in real-world tokenizers.

We further evaluate MIAs on out-of-distribution datasets, a single document containing unique tokens (canaries), and non-BPE tokenizers. Please refer to Appendix B for details.

6 Related Work

Tokenizers. Tokenizers play a crucial role in enabling the generation and comprehension capabilities of LLMs [29]. By converting raw text into discrete tokens, tokenizers provide the input representations that models require for inference [88]. Recent research has highlighted the connection between tokenizers and scaling laws, suggesting that larger models benefit from a larger tokenizer vocabulary, leading to improved performance under the same training cost [38,77]. Nevertheless, our work reveals that scaling up the tokenizer’s vocabulary increases its vulnerability to MIAs. We underscore the necessity of paying attention to these overlooked risks.

Membership Inference. Membership inference attacks [84] are designed to determine whether a specific entity was included in the training data of a machine learning (ML) model. These attacks have become fundamental tools for quantifying privacy leakage in various scenarios [12,84]. The first MIA is proposed by Shokri et al. [74], which focuses on demining record-level membership from ML-based classifiers. Recent studies underscore the importance of detecting whether specific data was used during the pre-training of LLMs, as such data may contain copyrighted or sensitive information. To address these risks, several membership inference methods have been proposed. Shi et al. [72] introduced MIN-K% PROB, a technique that detects membership by identifying outlier words with unusually low probabilities in previously unseen examples. Duarte et al. [24] proposed DE-COP, which probes LLMs using multiple-choice ques-

tions that include both verbatim and paraphrased versions of candidate sentences. However, these MIA methods face significant evaluation challenges. Duan et al. [23] argue that existing benchmarks suffer from a temporal distribution shift between *members* and *non-members*, potentially invalidating evaluation results. Meeus et al. [55] further observe that many methods may incorporate mislabeled samples and use impractical evaluation models in their experiments. To mitigate these privacy threats, differential privacy (DP) serves as a principal defense mechanism [25]. Defenders can leverage these techniques [1] to add noise to the gradients during model training, thereby obfuscating the data membership.

7 Discussion

LLM Dataset Inference. Recent research [51] has emphasized the importance of dataset inference in LLMs. In particular, existing methods [14,50] attempt to predict whether a specific dataset was used to train a target LLM by analyzing the model’s output. However, as noted in Section 1, these methods face significant challenges during evaluation, such as mislabeled samples, distribution shifts, and mismatches in target model sizes compared to real-world models. Furthermore, these attacks typically introduce additional assumptions about the adversary, such as access to model output loss or fine-tuning the target model, which are not guaranteed to hold in closed-source LLMs. Moreover, they can be defended by adding noise during the model training process [87].

Limitations of Tokenizer Inference. While the tokenizers listed in Table 1 are trained using the models’ pretraining data, some tokenizers might not follow this paradigm. For instance, ChatGLM [27] employs OpenAI’s tokenizer. In such cases, MIAs on these tokenizers will only reveal the tokenizer’s training data rather than LLMs’ pretraining corpus. Moreover, due to the absence of ground-truth training data for commercial tokenizers, we are unable to evaluate our attacks on them. Instead, we conduct evaluations on our trained tokenizers with vocabulary sizes and utility comparable to those of commercial tokenizers (see Figure 5). Notably, this limitation is not unique to our work but is a common challenge in the field of membership inference, as also observed in other related studies [41,43,47,50]. Furthermore, our attack evaluations are restricted to tokenizers used in LLMs. As a result, the feasibility of MIAs on tokenizers for classification models [48] and diffusion-based language models [68] remains unexplored.

8 Conclusion

In this paper, we review the limitations of existing MIAs against pre-trained LLMs and introduce the tokenizer as a new attack vector to address these challenges. To demonstrate its feasibility for membership inference, we present

the first study of MIAs on tokenizers of LLMs. By analyzing overfitting signals during tokenizer training, we proposed five attack methods for inferring dataset membership. Extensive evaluations on millions of Internet data demonstrate that our shadow-based attacks achieve strong performance. To mitigate these attacks, we further propose an adaptive defense mechanism. Although our proposed defense can reduce the membership leakage, it does so at the cost of tokenizer utility. Our findings highlight the vulnerabilities associated with LLMs’ tokenizers. Through this endeavor, we hope our research contributes to the design of privacy-preserving tokenizers, towards building secure machine learning systems.

Ethical Considerations

In this work, we conduct the first study of membership leakage through tokenizers. We recognize that it is our responsibility to carefully assess the ethical implications of our research. Our analysis follows the stakeholder-based framework and is guided by the principles outlined in The Menlo Report.

Stakeholder-based Analysis. We identify several key stakeholders, which might be impacted by this research:

- *LLM Developers.* Our work focuses on the tokenizers of LLMs and closely involves their developers. Leveraging our proposed methods, developers can estimate the privacy risk in the tokenizer before making it publicly accessible. This is essential as tokenizers in commercial LLMs are typically open-sourced to support transparent billing.
- *Data Subjects.* The introduction of a new attack vector for MIAs could be used to infer personal information and lead to privacy violations. However, our experiments solely utilize public, non-sensitive datasets collected by Google. We emphasize that our proposed attacks are never intended to facilitate or encourage the privacy leakage from LLMs.
- *Adversaries.* This work could potentially be used by adversaries to exploit privacy vulnerabilities. In recognition of this, we propose an adaptive defense to mitigate such risk.

Mitigation. To mitigate the membership leakage in tokenizers, we introduce the min count mechanism as an adaptive defense. Specifically, our attacks exploit distinctive tokens that occur rarely in the training data and are overfit into the tokenizer’s vocabulary. The min count mechanism counters this by removing these infrequent tokens from the vocabulary. Furthermore, we implement the differentially private mechanism at each token merging step during tokenizer training. DP-based tokenizer ensures theoretical privacy preservation against membership inference. The experimental results demonstrate that both of our approaches reduce privacy risks.

Justification for Research and Publication. The vulnerabilities exploited by our MIAs stem from the current practice of open-sourcing LLM tokenizers. Since similar attack techniques could be independently discovered by adversaries, it is important to disclose and analyze these issues within the

research community. Moreover, by publishing our findings, LLM developers can address such vulnerabilities proactively.

We emphasize that our proposed attacks are not intended to facilitate or encourage privacy leakage from LLMs. We hope our work contributes to the design of privacy-preserving tokenizers, laying the groundwork for secure and trustworthy systems that benefit society.

Open Science

All attack and defense codes, as well as datasets, are available at <https://github.com/mengtong0110/Tokenizer-MIA> and <https://zenodo.org/records/18080228>.

Acknowledgments

We thank all the anonymous reviewers and shepherd for their valuable comments. The work of Meng Tong, Kejiang Chen, and Weiming Zhang was supported by National Natural Science Foundation of China under Grant U2336206 and 62472398.

References

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *CCS*, pages 308–318, 2016.
- [2] Bo Adler, Niket Agarwal, Ashwath Aithal, Dong H Anh, Pallab Bhattacharya, Annika Brundyn, Jared Casper, Bryan Catanzaro, Sharon Clay, Jonathan Cohen, et al. Nemotron-4 340b technical report. *arXiv preprint arXiv:2406.11704*, 2024.
- [3] Michael Aerni, Jie Zhang, and Florian Tramèr. Evaluations of machine learning privacy defenses are misleading. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 1271–1284, 2024.
- [4] Akiko Aizawa. An information-theoretic perspective of tf-idf measures. *Information Processing & Management*, 39(1):45–65, 2003.
- [5] Bobby Allyn. Judge allows new york times copyright lawsuit to go forward. www.npr.org/2025/03/26/nx-s1-5288157/new-york-times-openai-copyright-case-goes-forward, 2025.
- [6] Sujoy Bag, Sri Krishna Kumar, and Manoj Kumar Tiwari. An efficient recommendation generation using relevant jaccard similarity. *Information Sciences*, 483:53–64, 2019.

- [7] Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024.
- [8] Stella Biderman, Kieran Bicheno, and Leo Gao. Datasheet for the pile. *arXiv preprint arXiv:2201.07311*, 2022.
- [9] Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. Pythia: A suite for analyzing large language models across training and scaling. In *ICML*, pages 2397–2430. PMLR, 2023.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Nee-lakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33:1877–1901, 2020.
- [11] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1897–1914. IEEE, 2022.
- [12] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650, 2021.
- [13] André M Carrington, Douglas G Manuel, Paul W Fieguth, Tim Ramsay, Venet Osmani, Bernhard Wernly, Carol Bennett, Steven Hawken, Olivia Magwood, Yusuf Sheikh, et al. Deep roc analysis and auc as balanced average accuracy, for improved classifier selection, audit and explanation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):329–341, 2022.
- [14] Hyeong Kyu Choi, Maxim Khanov, Hongxin Wei, and Yixuan Li. How contaminated is your benchmark? measuring dataset leakage in large language models with kernel divergence. In *Forty-second International Conference on Machine Learning*, 2025.
- [15] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.
- [16] United States Code. U.s. copyright act, title 17, section 107. <https://www.copyright.gov/title17/92chap1.html>, 1976.
- [17] Gautier Dagan, Gabriel Synnaeve, and Baptiste Rozière. Getting the most out of your tokenizer for pre-training and domain adaptation. In *Proceedings of the 41st International Conference on Machine Learning*, pages 9784–9805, 2024.
- [18] Damai Dai, Chengqi Deng, Chenggang Zhao, RX Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Yu Wu, et al. Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models. *arXiv preprint arXiv:2401.06066*, 2024.
- [19] Debeshee Das, Jie Zhang, and Florian Tramèr. Blind baselines beat membership inference attacks for foundation models. In *2025 IEEE Security and Privacy Workshops (SPW)*, pages 118–125. IEEE, 2025.
- [20] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019.
- [21] Jinshuo Dong, David Durfee, and Ryan Rogers. Optimal differential privacy composition for exponential mechanisms. In *International Conference on Machine Learning*, pages 2597–2606. PMLR, 2020.
- [22] Yuntao Du, Jiacheng Li, Yuetian Chen, Kaiyuan Zhang, Zhizhen Yuan, Hanshen Xiao, Bruno Ribeiro, and Ninghui Li. Cascading and Proxy Membership Inference Attacks. In *33th Annual Network and Distributed System Security Symposium (NDSS)*, 2026.
- [23] Michael Duan, Anshuman Suri, Niloofar Miresghal-lah, Sewon Min, Weijia Shi, Luke Zettlemoyer, Yulia Tsvetkov, Yejin Choi, David Evans, and Hannaneh Hajishirzi. Do membership inference attacks work on large language models? In *First Conference on Language Modeling*, 2024.
- [24] André V Duarte, Xuandong Zhao, Arlindo L Oliveira, and Lei Li. De-cop: detecting copyrighted content in language models training data. In *Proceedings of the 41st International Conference on Machine Learning*, pages 11940–11956, 2024.
- [25] Cynthia Dwork. Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer, 2006.
- [26] Matthias Gallé. Investigating the effectiveness of bpe: The power of shorter sequences. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint*

conference on natural language processing (EMNLP-IJCNLP), pages 1375–1381, 2019.

- [27] Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, et al. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*, 2024.
- [28] Google. Counting gemini text tokens locally with the vertex ai sdk, July 2024. Tokenizer for Use with Google’s Models.
- [29] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [30] Hacker News. Weird gpt-4 behavior for “davidjl”. news.ycombinator.com/item?id=36242914, 2023.
- [31] Jonathan Hayase, Alisa Liu, Yejin Choi, Sewoong Oh, and Noah A Smith. Data mixture inference attack: Bpe tokenizers reveal training data compositions. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [32] Jamie Hayes, Iliia Shumailov, Christopher A Choquette-Choo, Matthew Jagielski, George Kaissis, Katherine Lee, Milad Nasr, Sahra Ghalebikesabi, Niloofar Mireshghallah, Meenatchi Sundaram Mutu Selva Anamalai, et al. Strong membership inference attacks on massive datasets and (moderately) large language models. *arXiv preprint arXiv:2505.18773*, 2025.
- [33] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 770–778, 2016.
- [34] Yu He, Boheng Li, Liu Liu, Zhongjie Ba, Wei Dong, Yiming Li, Zhan Qin, Kui Ren, and Chun Chen. Towards label-only membership inference attack against pre-trained large language models. In *34th USENIX Security Symposium (USENIX Security 25)*, 2025.
- [35] Yuke He, Zheng Li, Yang Zhang, Zhan Qin, Kui Ren, and Chun Chen. Membership inference attacks against vision-language models. In *34th USENIX Security Symposium (USENIX Security 25)*, 2025.
- [36] Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B Brown, Prafulla Dhariwal, Scott Gray, et al. Scaling laws for autoregressive generative modeling. *arXiv preprint arXiv:2010.14701*, 2020.
- [37] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- [38] Hongzhi Huang, Defa Zhu, Banggu Wu, Yutao Zeng, Ya Wang, Qiyang Min, et al. Over-tokenized transformer: Vocabulary is generally worth scaling. In *Forty-second International Conference on Machine Learning*, 2025.
- [39] Jiameng Huang, Baijiong Lin, Guhao Feng, Jierun Chen, Di He, and Lu Hou. Efficient reasoning for large reasoning language models via certainty-guided reflection suppression. *arXiv preprint arXiv:2508.05337*, 2025.
- [40] Hugging Face Datasets. Codeparrot github code dataset. <https://huggingface.co/datasets/codeparrot/github-code>, 2025.
- [41] Bo Hui, Yuchen Yang, Haolin Yuan, Philippe Burlina, Neil Zhenqiang Gong, and Yinzi Cao. Practical blind membership inference attack via differential comparisons. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2021.
- [42] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [43] Hao Li, Zheng Li, Siyuan Wu, Yutong Ye, Min Zhang, Dengguo Feng, and Yang Zhang. Enhanced label-only membership inference attacks with fewer queries. In *Proceedings of the 34th USENIX Security Symposium (USENIX Security ’25)*. USENIX Association, 2025.
- [44] Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners. *arXiv preprint arXiv:2110.05679*, 2021.
- [45] Zheng Li and Yang Zhang. Membership leakage in label-only exposures. In *CCS*, pages 880–895, 2021.
- [46] Alisa Liu, Jonathan Hayase, Valentin Hofmann, Sewoong Oh, Noah A Smith, and Yejin Choi. SuperBPE: Space travel for language models. In *Second Conference on Language Modeling*, 2025.
- [47] Han Liu, Yuhao Wu, Zhiyuan Yu, and Ning Zhang. Please tell me more: Privacy impact of explainability through the lens of membership inference attack. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 4791–4809. IEEE, 2024.

- [48] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [49] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [50] Pratyush Maini, Hengrui Jia, Nicolas Papernot, and Adam Dziedzić. Llm dataset inference: Did you train on my dataset? *Advances in Neural Information Processing Systems*, 37:124069–124092, 2024.
- [51] Pratyush Maini, Mohammad Yaghini, and Nicolas Papernot. Dataset inference: Ownership resolution in machine learning. *arXiv preprint arXiv:2104.10706*, 2021.
- [52] Matthew Watkins. Tokens used by gpt-4 probably come from the reddit users. https://x.com/SoC_trilogy/status/1666714127438434304, 2023.
- [53] Prasanna Mayilvahanan, Thaddus Wiedemer, Sayak Mallick, Matthias Bethge, and Wieland Brendel. LLMs on the line: Data determines loss-to-loss scaling laws. In *Forty-second International Conference on Machine Learning*, 2025.
- [54] Matthieu Meeus, Shubham Jain, Marek Rei, and Yves-Alexandre de Montjoye. Did the neurons read your book? document-level membership inference for large language models. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 2369–2385, 2024.
- [55] Matthieu Meeus, Igor Shilov, Shubham Jain, Manuel Faysse, Marek Rei, and Yves-Alexandre de Montjoye. Sok: Membership inference attacks on llms are rushing nowhere (and how to fix it). In *IEEE Conference on Secure and Trustworthy Machine Learning ((SaTML, 2025)*. IEEE, 2025.
- [56] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- [57] Niklas Muennighoff, Alexander Rush, Boaz Barak, Teven Le Scao, Nouamane Tazi, Aleksandra Piktus, Sampo Pyysalo, Thomas Wolf, and Colin A Raffel. Scaling data-constrained language models. *Advances in Neural Information Processing Systems*, 36:50358–50376, 2023.
- [58] Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *2019 IEEE symposium on security and privacy (SP)*, pages 739–753. IEEE, 2019.
- [59] New York Times. Reddit sues anthropic over its data scraping to train large language models. <https://www.nytimes.com/2025/06/04/technology/reddit-anthropic-lawsuit-data.html>, 2025.
- [60] OpenAI. tiktoken: Tokenizer for openai models. <https://github.com/openai/tiktoken>, 2025.
- [61] Yan Pang, Tianhao Wang, Xuhui Kang, Mengdi Huai, and Yang Zhang. White-box membership inference attacks against diffusion models. *arXiv preprint arXiv:2308.06405*, 2023.
- [62] Jupinder Parmar, Shrimai Prabhunoye, Joseph Jennings, Mostafa Patwary, Sandeep Subramanian, Dan Su, Chen Zhu, Deepak Narayanan, Aastha Jhunjhunwala, Ayush Dattagupta, et al. Nemotron-4 15b technical report. *arXiv preprint arXiv:2402.16819*, 2024.
- [63] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [64] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- [65] Juan Ramos et al. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, volume 242, pages 29–48. New Jersey, USA, 2003.
- [66] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024.
- [67] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- [68] Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. *Advances in Neural Information Processing Systems*, 37:130136–130184, 2024.
- [69] Ahmed Salem, Yang Zhang, Mathias Humbert, Pascal Berrang, Mario Fritz, and Michael Backes. ML-leaks: Model and data independent membership inference attacks and defenses on machine learning models. *arXiv preprint arXiv:1806.01246*, 2018.

- [70] Abulhair Saparov and He He. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. In *The Eleventh International Conference on Learning Representations*, 2023.
- [71] Philip Sedgwick. Spearman’s rank correlation coefficient. *Bmj*, 349, 2014.
- [72] Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. Detecting pretraining data from large language models. In *ICLR*, 2024.
- [73] Yusuxke Shibata, Takuya Kida, Shuichi Fukamachi, Masayuki Takeda, Ayumi Shinohara, Takeshi Shinohara, and Setsuo Arikawa. Byte pair encoding: A text compression scheme that accelerates pattern matching. , 1999.
- [74] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2017.
- [75] Kevin Slagle. Spacebyte: Towards deleting tokenization from large language modeling. *Advances in Neural Information Processing Systems*, 37:124925–124950, 2024.
- [76] Karen Sparck Jones. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21, 1972.
- [77] Chaofan Tao, Qian Liu, Longxu Dou, Niklas Muenighoff, Zhongwei Wan, Ping Luo, Min Lin, and Ngai Wong. Scaling laws with vocabulary: Larger models deserve larger vocabularies. *Advances in Neural Information Processing Systems*, 37:114147–114179, 2024.
- [78] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [79] Meng Tong, Yuntao Du, Kejiang Chen, Weiming Zhang, and Ninghui Li. Membership inference attacks on tokenizers of large language models. *arXiv preprint arXiv:2510.05699*, 2025.
- [80] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [81] Geoffrey I Webb. Naïve bayes. In *Encyclopedia of machine learning and data mining*, pages 895–896. Springer, 2017.
- [82] Jiayuan Ye, Aadyaa Maddi, Sasi Kumar Murakonda, Vincent Bindschaedler, and Reza Shokri. Enhanced membership inference attacks against machine learning models. In *CCS*, pages 3093–3106, 2022.
- [83] Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st computer security foundations symposium (CSF)*, pages 268–282. IEEE, 2018.
- [84] Sajjad Zarifzadeh, Philippe Liu, and Reza Shokri. Low-cost high-power membership inference attacks. In *Proceedings of the 41st International Conference on Machine Learning*, pages 58244–58282, 2024.
- [85] Jie Zhang, Debeshee Das, Gautam Kamath, and Florian Tramèr. Position: Membership inference attacks cannot prove that a model was trained on your data. In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 333–345. IEEE, 2025.
- [86] Jingyang Zhang, Jingwei Sun, Eric Yeats, Yang Ouyang, Martin Kuo, Jianyi Zhang, Hao Yang, and Hai Li. Min-k%+: Improved baseline for detecting pre-training data from large language models. *arXiv preprint arXiv:2404.02936*, 2024.
- [87] Kaiyuan Zhang, Siyuan Cheng, Hanxi Guo, Yuetian Chen, Zian Su, Shengwei An, Yuntao Du, Charles Fleming, Ashish Kundu, Xiangyu Zhang, and Ninghui Li. SOFT: Selective Data Obfuscation for Protecting LLM Fine-tuning against Membership Inference Attacks. In *Proceedings of the 34th USENIX Security Symposium (USENIX Security 2025)*, 2025.
- [88] Vilém Zouhar, Clara Meister, Juan Gastaldi, Li Du, Tim Vieira, Mrinmaya Sachan, and Ryan Cotterell. A formal perspective on byte-pair encoding. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 598–614, 2023.

A Proof of Theorem 4.2

Proof. Given that $\Pr(t_i | \mathcal{V}_{\text{target}}) \propto \frac{1}{j^\alpha}$, let us assume

$$\Pr(t_i | \mathcal{V}_{\text{target}}) = C/j^\alpha, \quad (19)$$

where $C \in \mathbb{R}_{>0}$ is a constant. Since the sum of the frequencies $\Pr(t_i | \mathcal{V}_{\text{target}})$ for all $t_i \in \mathcal{V}_{\text{target}}$ is at most 1, it follows that

$$1 \geq \sum_{j=x_{\min}+1}^{|\mathcal{V}_{\text{target}}|} \Pr(t_j | \mathcal{V}_{\text{target}}) = \sum_{j=1+x_{\min}}^{|\mathcal{V}_{\text{target}}|} C/j^\alpha. \quad (20)$$

Table 10: Impact of the target dataset size on defense mechanism ($n_{\min} = 32$).

| Attack Approach | #Dataset Size | $ \mathcal{V}_{\text{target}} < 80,000$ | | | $ \mathcal{V}_{\text{target}} < 110,000$ | | | $ \mathcal{V}_{\text{target}} < 140,000$ | | | $ \mathcal{V}_{\text{target}} < 170,000$ | | | $ \mathcal{V}_{\text{target}} < 200,000$ | | |
|----------------------|-----------------------|--|-------|--------|---|-------|--------|---|-------|--------|---|-------|--------|---|-------|--------|
| | | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR |
| Vocabulary Overlap | $ D \in [0, 400)$ | 0.646 | 0.632 | 21.25% | 0.677 | 0.642 | 21.25% | 0.694 | 0.655 | 23.30% | 0.706 | 0.663 | 24.44% | 0.717 | 0.669 | 25.14% |
| | $ D \in [400, 800)$ | 0.703 | 0.667 | 21.25% | 0.761 | 0.711 | 22.85% | 0.772 | 0.717 | 33.73% | 0.793 | 0.730 | 34.94% | 0.810 | 0.742 | 40.06% |
| | $ D \in [800, 1200)$ | 0.712 | 0.670 | 25.30% | 0.768 | 0.729 | 30.12% | 0.798 | 0.761 | 38.75% | 0.823 | 0.775 | 43.75% | 0.840 | 0.800 | 46.25% |
| Frequency Estimation | $ D \in [0, 400)$ | 0.586 | 0.591 | 16.36% | 0.617 | 0.616 | 20.11% | 0.649 | 0.635 | 20.62% | 0.682 | 0.657 | 21.01% | 0.720 | 0.683 | 26.99% |
| | $ D \in [400, 800)$ | 0.611 | 0.615 | 17.78% | 0.675 | 0.650 | 25.00% | 0.720 | 0.683 | 29.52% | 0.739 | 0.702 | 30.42% | 0.763 | 0.721 | 33.43% |
| | $ D \in [800, 1200)$ | 0.723 | 0.714 | 25.00% | 0.742 | 0.716 | 35.00% | 0.758 | 0.729 | 35.00% | 0.785 | 0.759 | 47.50% | 0.820 | 0.805 | 52.50% |

Table 11: Impact of the target dataset size on defense mechanism ($n_{\min} = 48$).

| Attack Approach | #Dataset Size | $ \mathcal{V}_{\text{target}} < 80,000$ | | | $ \mathcal{V}_{\text{target}} < 110,000$ | | | $ \mathcal{V}_{\text{target}} < 140,000$ | | | $ \mathcal{V}_{\text{target}} < 170,000$ | | | $ \mathcal{V}_{\text{target}} < 200,000$ | | |
|----------------------|-----------------------|--|-------|--------|---|-------|--------|---|-------|--------|---|-------|--------|---|-------|--------|
| | | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR |
| Vocabulary Overlap | $ D \in [0, 400)$ | 0.646 | 0.632 | 21.25% | 0.677 | 0.642 | 21.25% | 0.687 | 0.645 | 22.85% | 0.694 | 0.655 | 23.30% | 0.710 | 0.667 | 27.12% |
| | $ D \in [400, 800)$ | 0.704 | 0.667 | 21.58% | 0.761 | 0.711 | 22.22% | 0.772 | 0.717 | 30.00% | 0.775 | 0.721 | 34.64% | 0.790 | 0.730 | 34.94% |
| | $ D \in [800, 1200)$ | 0.712 | 0.670 | 25.30% | 0.768 | 0.729 | 30.12% | 0.774 | 0.738 | 30.72% | 0.798 | 0.761 | 38.75% | 0.823 | 0.767 | 45.00% |
| Frequency Estimation | $ D \in [0, 400)$ | 0.586 | 0.591 | 16.36% | 0.617 | 0.616 | 18.20% | 0.649 | 0.634 | 19.92% | 0.680 | 0.652 | 21.01% | 0.680 | 0.654 | 22.79% |
| | $ D \in [400, 800)$ | 0.588 | 0.597 | 16.93% | 0.619 | 0.623 | 20.11% | 0.650 | 0.635 | 20.62% | 0.682 | 0.657 | 21.64% | 0.720 | 0.683 | 26.99% |
| | $ D \in [800, 1200)$ | 0.616 | 0.611 | 20.48% | 0.681 | 0.659 | 28.01% | 0.719 | 0.692 | 28.92% | 0.731 | 0.693 | 31.93% | 0.736 | 0.699 | 32.53% |

Table 12: Additional evaluations with $|\mathcal{V}_{\text{target}}| = 200,000$.

| Evaluation Settings | AUC | TPR @ 1% FPR | Balanced Acc. |
|-------------------------------------|-------|--------------|---------------|
| <i>Canary Experiments</i> | | | |
| + MIA via Vocabulary Overlap | 1.000 | 100.00% | 1.000 |
| + MIA via Frequency Estimation | 1.000 | 100.00% | 1.000 |
| <i>Out-of-distribution Datasets</i> | | | |
| + MIA via Vocabulary Overlap | 0.757 | 21.00% | 0.730 |
| + MIA via Frequency Estimation | 0.739 | 13.00% | 0.715 |
| <i>Unigram Tokenizers</i> | | | |
| + MIA via Vocabulary Overlap | 0.762 | 23.00% | 0.712 |
| + MIA via Frequency Estimation | 0.764 | 30.00% | 0.728 |

Therefore, the upper bound for the constant C is

$$C \leq \frac{1}{\sum_{j=x_{\min}+1}^{|\mathcal{V}_{\text{target}}|} \frac{1}{j^\alpha}}. \quad (21)$$

Using Equations 19 and 21, we derive the upper bound for the frequency $\Pr(t_i | \mathcal{V}_{\text{target}})$:

$$\Pr(t_i | \mathcal{V}_{\text{target}}) \leq \frac{1}{\sum_{j=x_{\min}+1}^{|\mathcal{V}_{\text{target}}|} \frac{1}{j^\alpha}} \cdot \frac{1}{i^\alpha}. \quad (22)$$

Taking the negative logarithm, the self-information of t_i is bounded from below:

$$\text{SI}(t_i, \mathcal{V}_{\text{target}}) = -\log \Pr(t_i | \mathcal{V}_{\text{target}}) \geq \log \left(\sum_{j=x_{\min}+1}^{|\mathcal{V}_{\text{target}}|} i^\alpha / j^\alpha \right). \quad (23)$$

Therefore, the RTF-SI satisfies:

$$\text{RTF-SI}(D, t_i, \mathcal{V}_{\text{target}}) \geq \frac{n_D(t_i)}{\sum_{D' \in \mathbb{D}} n_{D'}(t_i)} \cdot \log \left(\sum_{j=x_{\min}+1}^{|\mathcal{V}_{\text{target}}|} i^\alpha / j^\alpha \right). \quad (24)$$

This theorem allows an adversary to approximate the RTF-SI using its lower bound under a power-law distribution. \square

B Additional Evaluations

Canary Experiments. Our MIAs focus on dataset-level membership. A key question is whether their effectiveness implies document-level risk [12]. To this end, we conduct canary evaluations [3, 85]. Specifically, we synthesize a token from five alphabets and insert it into a single document 40 times. We use MIAs that perform well on datasets to infer the membership of such documents. Table 12 shows that the canary cases are highly detectable. This can be attributed to our attacks, which identify unique tokens, thus establishing a one-to-one mapping between the tokens and the documents.

Out-of-distribution Datasets. We further evaluate our MIAs using auxiliary datasets from corpora different from those used in the tokenizer’s training data. Specifically, we train the tokenizer on the *PILE* corpus [8] (published in 2022) with the target dataset to infer membership. The auxiliary data is drawn from the *C4* corpus [64] (published in 2020), ensuring that no samples are shared between the auxiliary data and the tokenizer’s training data. The results in Table 12 demonstrate that our MIAs perform effectively without relying on the shared data assumption in the threat model. Indeed, previous work [11] has shown that shadow-based MIAs with out-of-distribution auxiliary data can successfully infer membership.

Unigram Tokenizers. Our MIAs focus on tokenizers trained using the BPE algorithm. This is because state-of-the-art LLMs employ this training paradigm [18]. In addition to BPE tokenizers, we further evaluate the MIAs against unigram tokenizers¹. Table 12 presents the experimental results. It is observed that our MIAs also generalize to unigram tokenizers. This is likely because the unigram tokenizer merges the most frequent string as a new token at each step, which resembles the BPE and overfits distinctive tokens into the vocabulary.

¹<https://github.com/google/sentencepiece>

Table 13: Impact of the target dataset size on defense mechanism ($n_{\min} = 64$).

| Attack Approach | #Dataset Size | $ \mathcal{V}_{\text{target}} \leq 80,000$ | | | $ \mathcal{V}_{\text{target}} \leq 110,000$ | | | $ \mathcal{V}_{\text{target}} \leq 140,000$ | | | $ \mathcal{V}_{\text{target}} \leq 170,000$ | | | $ \mathcal{V}_{\text{target}} \leq 200,000$ | | |
|----------------------|-----------------------|---|-------|--------|--|-------|--------|--|-------|--------|--|-------|--------|--|-------|--------|
| | | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR | AUC | BA | TPR |
| Vocabulary Overlap | $ D \in [0, 400)$ | 0.646 | 0.632 | 19.16% | 0.662 | 0.636 | 21.25% | 0.677 | 0.642 | 21.25% | 0.683 | 0.647 | 23.30% | 0.694 | 0.655 | 24.06% |
| | $ D \in [400, 800)$ | 0.703 | 0.666 | 21.25% | 0.731 | 0.675 | 21.58% | 0.751 | 0.699 | 22.85% | 0.761 | 0.711 | 32.50% | 0.772 | 0.717 | 37.95% |
| | $ D \in [800, 1200)$ | 0.712 | 0.670 | 25.30% | 0.743 | 0.702 | 27.11% | 0.768 | 0.729 | 30.12% | 0.795 | 0.746 | 33.73% | 0.797 | 0.761 | 38.75% |
| Frequency Estimation | $ D \in [0, 400)$ | 0.591 | 0.598 | 17.50% | 0.620 | 0.619 | 18.20% | 0.651 | 0.636 | 20.11% | 0.654 | 0.637 | 21.07% | 0.656 | 0.639 | 21.39% |
| | $ D \in [400, 800)$ | 0.619 | 0.615 | 22.29% | 0.672 | 0.653 | 28.31% | 0.717 | 0.682 | 30.42% | 0.718 | 0.683 | 31.93% | 0.723 | 0.686 | 32.83% |
| | $ D \in [800, 1200)$ | 0.734 | 0.717 | 26.25% | 0.739 | 0.731 | 33.75% | 0.744 | 0.736 | 33.75% | 0.745 | 0.742 | 35.00% | 0.748 | 0.746 | 38.75% |




Figure 12: Distinctive tokens in MIA via Vocabulary Overlap.




Figure 13: ROC curves for MIAs using different methods.




Figure 14: Dataset distribution based on MIA via Vocabulary Overlap.




Figure 15: Dataset distribution based on MIA via Frequency Estimation.