

On Membership Inference Attacks in Knowledge Distillation*

Ziyao Cui[†]
Duke University
308 Research Drive
Durham, North Carolina 27708
richard.cui@duke.edu

Minxing Zhang[†]
Duke University
308 Research Drive
Durham, North Carolina 27708
minxing.zhang@duke.edu

Jian Pei
Duke University
308 Research Drive
Durham, North Carolina 27708
j.pei@duke.edu

ABSTRACT

Large language models (LLMs) are trained on massive corpora that may contain sensitive information, creating privacy risks under membership inference attacks (MIAs). Knowledge distillation is widely used to compress LLMs into smaller student models, but its privacy implications are poorly understood. We systematically evaluate how distillation affects MIA vulnerability across six teacher-student model pairs and six attack methods. We find that distilled student models do not consistently exhibit lower MIA success than their teacher models, and in some cases demonstrate substantially higher member-specific attack success, challenging the assumption that knowledge distillation inherently improves privacy. We attribute this to mixed supervision in distillation: for vulnerable training data points, teacher predictions often align with ground-truth labels, causing student models to learn overly confident predictions that amplify the separability between members and non-members; conversely, for non-vulnerable points, teacher predictions and ground truth frequently diverge, providing inconsistent learning signals. To mitigate this, we propose three practical interventions – restricting distillation to non-vulnerable points, adding a low-dimensional **Bottleneck Projection**, and a normalization variant (**NoNorm**). Experiments show these methods reduce both aggregate and member-specific MIA success while preserving model utility, improving privacy-utility trade-offs for distilled LLMs.¹

1. INTRODUCTION

Large Language Models (LLMs) have achieved remarkable success due to the scale and diversity of their pretraining data [23, 46]. Moreover, as AI models become increasingly prevalent, privacy concerns have gained significant attention, with researchers investigating vulnerabilities and defenses across model architectures [7, 4, 10]. For LLMs, massive pretraining data introduces serious risks of training data exposure and privacy breaches [43, 41]. More specifically, the massive and heterogeneous nature of these datasets makes it infeasible to fully remove sensitive content, including copyrighted materials [24, 11] and personally identifiable information [26, 37]. Consequently, LLMs may memorize privacy-sensitive data, enabling

*This research is supported in part by the NSF Project MSPA-2434666. All opinions, findings, conclusions, and recommendations in this article are those of the authors and do not necessarily reflect the views of the funding agencies.

¹Our implementation and evaluation code are available at <https://github.com/richardcui18/mia-in-kd>.

[†]Both authors contributed equally to this research.

attackers to infer training membership through Membership Inference Attacks (MIAs) [31].

Prior work has extensively studied the detection of pretraining data in LLMs using MIAs [43, 45, 7, 4]. However, these studies largely analyze models in isolation. In parallel, model compression techniques such as knowledge distillation have been widely adopted to reduce the size and computational cost of modern LLMs, including Llama, Gemma, and BERT [28, 20, 21, 34, 39, 35]. Because distilled student models have fewer parameters and lower capacity, it is commonly assumed that distillation reduces memorization and improves privacy. This assumption, however, has not been systematically evaluated across diverse teacher-student pairs and MIA methods, leaving the privacy effects of distillation insufficiently understood.

To address this gap, as the first contribution in this paper, we evaluate six teacher-student model pairs across multiple architectures using six representative MIAs. We report three key findings. First, distilled student models do not consistently exhibit lower aggregate MIA accuracy than their teacher models. Second, student models can exhibit higher member-specific attack accuracy even when their overall accuracy is lower, which represents a greater practical privacy risk [4]. Third, we provide an explanation: the mixed supervision in distillation – combining ground-truth labels and teacher predictions – can reinforce memorization on vulnerable training data points due to the alignment between teacher predictions and ground-truth, leading the student model to produce overly confident outputs that amplify the separability between members and non-members; on the other hand, for non-vulnerable data points, teacher predictions diverge from ground-truth, making the supervision inconsistent and failing to provide clear privacy benefits.

Motivated by this insight, as the second contribution, we propose and evaluate three targeted interventions to reduce membership leakage in distilled models. First, we introduce a data-selection strategy that restricts distillation to *non-vulnerable* training data points, which reduces exposure to memorized points but may increase student perplexity due to reduced training data. To mitigate this utility loss, we propose two lightweight architectural modifications. The first is a low-dimensional **Bottleneck Projection** that limits representational capacity and discourages memorization. The second replaces layer normalization with **NoNorm**, a parameterized element-wise linear transformation. Across experiments, these interventions consistently reduce MIA success, and the architectural modifications in particular lower attack success without degrading model utility, yielding improved privacy-utility trade-offs compared to naïve distillation.

Outline.

Section 2 formalizes membership inference and introduces the evaluation metrics and protocol. Section 3 reviews related work on

membership inference, knowledge distillation, and privacy-aware model compression. Section 4 presents a systematic empirical analysis of membership inference vulnerability in teacher and student models and introduces diagnostics for privacy leakage in student models. Section 5 describes the proposed privacy-preserving distillation methods. Section 6 reports experimental results, privacy-utility trade-offs, and additional analyses. Finally, Section 7 discusses limitations and directions for future work; supplementary material and extended ablations are provided in the Appendix.

2. PROBLEM DEFINITION AND EVALUATION METRICS

2.1 Membership Inference Attacks

Consider a LLM \mathcal{N} and a dataset \mathcal{D} used to train \mathcal{N} . A **MIA method** [45] M takes a target data point d as input and aims to determine whether $d \in \mathcal{D}$. Denote by $M(\mathcal{N}, d)$ the attacker’s prediction, where $M(\mathcal{N}, d) = 1$ if M predicts $d \in \mathcal{N}$ and 0 otherwise. In practice, M may compute a confidence score $M'(\mathcal{N}, d)$ and use a threshold τ to predict the membership of d , that is, $M_\tau(\mathcal{N}, d) = \mathbf{1}[M'(\mathcal{N}, d) > \tau]$, where $\mathbf{1}$ is the indicator function.

2.2 Knowledge Distillation

In knowledge distillation, a teacher model \mathcal{T} is used as a guiding framework to transfer knowledge to a student model \mathcal{S} , where the student model \mathcal{S} is learned to mimic the performance of the teacher \mathcal{T} [31, 44].

Let $\mathcal{D}_\mathcal{T}$ and $\mathcal{D}_\mathcal{S}$ denote the training datasets of the teacher model \mathcal{T} and student model \mathcal{S} , respectively. We use the standard mixed-supervision distillation objective [17], in which the student model is trained to minimize a weighted sum of the supervised loss on ground-truth labels and a distillation loss that encourages the student model to match the teacher model’s predictive distribution. Specifically, for training data points $(x, y) \sim \mathcal{D}_\mathcal{S}$, the objective is

$$\mathcal{L}_{\text{distill}} = \mathbb{E}_{(x, y) \sim \mathcal{D}_\mathcal{S}} \left[\mathcal{L}_{\text{CE}}(y, p_\mathcal{S}(\cdot | x)) + \lambda \text{KL}(p_\mathcal{T}(\cdot | x) \| p_\mathcal{S}(\cdot | x)) \right]$$

where $p_\mathcal{T}(\cdot | x)$ and $p_\mathcal{S}(\cdot | x)$ are the predictive distributions of the teacher and student models for training data (prefix) x , \mathcal{L}_{CE} denotes cross-entropy loss, $\text{KL}(\cdot \| \cdot)$ is the Kullback-Leibler divergence, and $\lambda \geq 0$ controls the trade-off between the two terms.

2.3 Problem Definition

Consider a model \mathcal{T} trained on dataset $\mathcal{D}_\mathcal{T}$. Let $\mathcal{D} \subseteq \mathcal{D}_\mathcal{T}$ be a set of member data points and let \mathcal{D}' be a non-member set such that $\mathcal{D}' \cap \mathcal{D}_\mathcal{T} = \emptyset$. Given a MIA method M_τ with threshold τ , we define the **MIA true positive rate** as

$$TPR(\mathcal{T}) = \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \mathbf{1}[M_\tau(\mathcal{T}, x) = 1],$$

and the **MIA true negative rate** as

$$TNR(\mathcal{T}) = \frac{1}{|\mathcal{D}'|} \sum_{x' \in \mathcal{D}'} \mathbf{1}[M_\tau(\mathcal{T}, x') = 0].$$

We define the **MIA accuracy** on model \mathcal{T} as the average of these two quantities:

$$A(\mathcal{T}) = \frac{1}{2} [TPR(\mathcal{T}) + TNR(\mathcal{T})]. \quad (1)$$

The central question studied in this paper is whether distilled student models \mathcal{S} are more vulnerable to MIAs than their teacher models \mathcal{T} . We assess vulnerability using both aggregate MIA accuracy, $A(\mathcal{T})$ versus $A(\mathcal{S})$, and member-specific MIA accuracy, $TPR(\mathcal{T})$ versus $TPR(\mathcal{S})$, and investigate how distillation can be modified to improve student robustness to membership inference.

3. RELATED WORK

3.1 MIA Methods

MIAs, introduced by Shokri et al. [31], aim to determine whether a data point was used to train a model. While MIAs have been studied in settings such as diffusion models [5] and multi-layer perceptrons [42], applying MIAs to LLMs presents distinct challenges. LLM training data are typically not public [40, 1, 16], complicating evaluation due to missing ground-truth membership labels, and modern LLMs are often trained for a single epoch over massive corpora, reducing classical memorization signals [6, 30]. Despite these challenges, several MIA methods have been developed for LLMs, including loss-based attacks [45], zlib-normalized loss [7], reference-model attacks using likelihood ratio tests [4], and ReCaLL, which leverages relative conditional log-likelihoods [43].

However, existing studies primarily analyze MIAs on individual models in isolation and do not consider how vulnerability changes across related models, such as teacher-student pairs produced by knowledge distillation. In particular, prior MIA methods [45, 7] were not designed to assess privacy trade-offs introduced by distillation. *Our work addresses this gap by systematically evaluating MIA behavior across distilled teacher-student model pairs and using these insights to develop privacy-preserving distillation methods.*

3.2 Knowledge Distillation

Knowledge distillation [17, 3] is a model compression technique in which a smaller student model is trained to mimic a larger teacher model using a combination of supervised loss and a distillation loss that aligns predictions by student and teacher models [14, 44]. Distillation has been widely applied to LLMs, including DistilBERT [28] and subsequent methods that focus on matching output distributions [32, 22, 47]. Other approaches exploit intermediate representations as training signals, enabling the student model to imitate the teacher model’s hidden states through layer-wise distillation [27, 33, 20].

Prior work on knowledge distillation has largely emphasized efficiency and performance, with limited attention to privacy. In contrast, *our work studies knowledge distillation through the lens of MIAs, analyzing how the distillation process alters privacy risk across teacher-student model pairs.* Rather than treating distillation as inherently privacy-preserving, we leverage the teacher model’s MIA vulnerability as a signal to understand and mitigate privacy leakage in the student model. As distillation becomes a standard component of LLM deployment, incorporating privacy-aware objectives is essential for responsible model compression.

3.3 Privacy in Knowledge Distillation

A growing body of work has studied the privacy implications of knowledge transfer mechanisms such as knowledge distillation. Several approaches propose distillation-based defenses against membership leakage, including ensemble and self-distillation schemes and cross-distillation protocols [36, 9, 29, 48]. These methods, however, focus on specific engineered variants rather than the canonical mixed-supervision distillation pipeline commonly used in practice.

Model	Min-K% (k)	Min-K% $_{++}$ (k)	ReCaLL Prefix
Pythia	0.10	0.50	7
DistilPythia	0.50	0.30	7
Gemma 2 27B	0.15	0.25	5
Gemma 2 9B	0.20	0.90	30
Gemma 2 2B	0.20	0.10	28
Gemma 2 2B Distilled	0.40	0.20	28
Llama 3.1 8B	0.70	0.10	20
Llama 3.2 1B	0.90	0.10	22
Llama 3.2 3B	0.90	0.10	22

Table 1: Optimal hyperparameters selected via cross-validation for Min-K%, Min-K% $_{++}$, and ReCaLL across all evaluated target models.

Teacher Model	Student Model	ReCall	Loss	Zlib	Mink	Mink $_{++}$	Ref Model
Pythia	DistilPythia	0.555 / 0.565	0.442 / 0.444	0.613 / 0.635	0.316 / 0.414	0.501 / 0.501	0.648 / 0.579
Gemma 2 27B	Gemma 2 2B	0.667 / 0.667	0.494 / 0.525	0.556 / 0.481	0.543 / 0.556	0.537 / 0.451	0.494 / 0.420
	Gemma 2 2B Distilled	0.667 / 0.494	0.494 / 0.537	0.556 / 0.556	0.543 / 0.580	0.537 / 0.432	0.494 / 0.426
	Gemma 2 9B	0.667 / 0.704	0.494 / 0.475	0.556 / 0.531	0.543 / 0.543	0.537 / 0.481	0.494 / 0.426
Llama 3.1 8B	Llama 3.2 1B	0.702 / 0.682	0.303 / 0.311	0.552 / 0.532	0.300 / 0.309	0.311 / 0.335	0.441 / 0.349
	Llama 3.2 3B	0.702 / 0.806	0.303 / 0.314	0.552 / 0.538	0.300 / 0.311	0.311 / 0.205	0.441 / 0.381

Table 2: Comparison of aggregate MIA accuracy for teacher-student model pairs. Each cell reports teacher / student MIA accuracy, with **bold** indicating the lower (more privacy-preserving) value for each pair and attack method.

In contrast, *our work systematically evaluates this standard pipeline across diverse LLM families and identifies conditions under which distillation amplifies membership leakage.*

Related empirical studies have also questioned the privacy benefits of distillation. For example, Jagielski et al. [19] show that student models may replicate teacher model behaviors under certain conditions, yielding limited privacy gains. Their setting, however, considers distillation solely from teacher model outputs, without access to ground-truth labels, whereas *our study analyzes the more prevalent mixed-supervision objective [17] that combines teacher model predictions and labeled data.* Earlier work by Jagannatha et al. [18] reports lower privacy leakage in DistilBERT compared to BERT, but is limited to a single model pair and a narrow clinical dataset. *Our work expands beyond these settings by examining multiple teacher-student model pairs and datasets, providing a broader understanding of how distillation affects privacy and how it can be modified to improve robustness against membership inference.*

4. MIA IN DISTILLED LLMS

We present a systematic empirical analysis of membership inference vulnerability under knowledge distillation, evaluating six teacher-student pairs across multiple architectures using six representative attacks and the metrics defined in Section 2.3.

4.1 Experimental Setup

4.1.1 Models and Datasets

We evaluate three model families with known training data and publicly available distilled variants. **Pythia** [2] is trained on the ArXiv subset of the Pile [13], with DistilPythia [15] as the student model. **Gemma 2 27B** [38] is trained on the WikiMIA dataset [30], with Gemma 2 9B, Gemma 2 2B, and Gemma 2 2B distilled [38, 35] as student models. **Llama 3.1 8B** [12] is trained on the ArXiv subset of the Pile [13], with Llama 3.2 3B and Llama 3.2 1B [12] as student

models.

For non-member evaluation, we use the built-in non-member split of WikiMIA for Gemma models [30]. For all other models, we use the WebInstructSub-prometheus dataset [8], released in May 2024, after the model release dates, ensuring it was not used during training – a method widely adopted by existing works [25, 43, 41].

4.1.2 MIA Methods

We evaluate membership inference using six representative attacks. **ReCaLL** [43] measures changes in conditional log-likelihood when prefixing inputs with non-member context. **Loss** [45] uses the per-example loss as a membership score, while **Zlib** [7] normalizes this loss by zlib compression entropy. The **Reference-model** attack extends the loss-based approach by training shadow models with and without the target data point and performing a likelihood-ratio test. **Min-K%** [30] computes membership scores from the average log-likelihood of the lowest-probability $k\%$ tokens, and **Min-K% $_{++}$** further calibrates these scores using the mean and standard deviation over the vocabulary.

We leverage cross-validation to select hyperparameters for these attacks: for each dataset and target model, we tune attack-specific hyperparameters – such as the prefix length for ReCaLL and the value of k for Min-K% and Min-K% $_{++}$ – to maximize average attack performance. We report the optimal hyperparameters in Table 1.

As described in Section 2, each method outputs a confidence score per data point, which is converted to a binary prediction using a threshold selected via ROC analysis to maximize MIA accuracy while ensuring balanced predicted classes. All experiments are conducted on NVIDIA A5000/A6000 GPUs.

4.2 Aggregate MIA Accuracy

Table 2 reports aggregate MIA accuracy for teacher-student model pairs. Across six model pairs and six attack methods, student models frequently achieve MIA accuracy comparable to, and in several cases exceeding, that of their teachers. Specifically, the teacher model

Teacher Model	Student Model	ReCall	Loss	Zlib	Mink	Mink++	Ref Model
Pythia	DistilPythia	0.658 / 0.673	0.631 / 0.630	0.234 / 0.290	0.444 / 0.683	0.993 / 0.993	0.620 / 0.636
Gemma 2 27B	Gemma 2 2B	0.615 / 0.670	0.606 / 0.661	0.156 / 0.037	0.761 / 0.495	0.688 / 0.404	0.780 / 0.376
	Gemma 2 2B Distilled	0.615 / 0.532	0.606 / 0.716	0.156 / 0.101	0.761 / 0.376	0.688 / 0.771	0.780 / 0.358
	Gemma 2 9B	0.615 / 0.459	0.606 / 0.505	0.156 / 0.165	0.761 / 0.633	0.688 / 0.450	0.780 / 0.339
Llama 3.1 8B	Llama 3.2 1B	0.968 / 0.937	0.451 / 0.513	0.497 / 0.567	0.447 / 0.508	0.486 / 0.604	0.563 / 0.563
	Llama 3.2 3B	0.968 / 0.974	0.451 / 0.510	0.497 / 0.515	0.447 / 0.502	0.486 / 0.320	0.563 / 0.563

Table 3: Comparison of member-specific MIA accuracy for teacher-student model pairs. Each cell reports teacher / student member-specific MIA accuracy, with **bold** indicating the lower (more privacy-preserving) value for each pair and attack method.




Figure 1: Teacher & ground-truth alignment stratified by membership inference vulnerability. **Left:** Average probability assigned by the teacher model to the ground-truth next token. **Right:** KL divergence between the teacher model’s predictive distribution and the ground-truth distribution.

exhibits lower accuracy in 15 cases, the student model in 17 cases, and both models tie in 4 cases.

We further test whether teacher models consistently exhibit higher MIA accuracy than their distilled student models using a one-sided sign test, with the null hypothesis $H_0 : P(A(\mathcal{T}) > A(\mathcal{S})) \leq 0.5$ and alternative hypothesis $H_1 : P(A(\mathcal{T}) > A(\mathcal{S})) > 0.5$, where $A(\mathcal{T})$ and $A(\mathcal{S})$ denote the teacher and student MIA accuracies, respectively. Across all comparisons, the teacher model exhibits higher MIA accuracy in 17 cases, yielding a p -value of 0.43. We therefore fail to reject H_0 , indicating that *knowledge distillation does not consistently reduce membership inference risk*. This result challenges the common assumption that model compression alone improves privacy.

4.3 Member-Specific MIA Accuracy

Since member-specific privacy risks are more practically significant than aggregate metrics [4], we explore the member-specific MIA accuracy for teacher-student model pairs.

As illustrated in Table 3, the teacher model exhibits lower accuracy in 18 cases, the student model in 16 cases, and both models tie in 2 cases. For Pythia and Llama families, the student model member-specific MIA accuracy increases by 9.81% and 2.29%, respectively, compared to the teacher model. Notably, under the reference-model attack for the Pythia family, the student model achieves lower aggregate accuracy than the teacher model (0.579 vs. 0.648; Table 2) while exhibiting higher member-specific accuracy (0.636 vs. 0.620; Table 3). This demonstrates that aggregate MIA accuracy can obscure increased member-specific vulnerability and that lower overall attack success does not necessarily imply stronger privacy protection.

We also conduct the one-sided sign test with a similar setting as Section 4.2. Across all comparisons, the teacher model exhibits

higher MIA accuracy in 16 cases, giving a p -value of 0.432. We again fail to reject H_0 , indicating that knowledge distillation also does not consistently reduce member-specific risk.

4.4 Why May Student Models Exhibit Greater Privacy Leakage?

Teacher and student models differ fundamentally in their training signals. Teacher models are trained solely with ground-truth supervision, whereas student models optimize a mixed objective that combines the ground-truth loss with a distillation loss, typically the KL divergence between teacher model and student model predictions. Consequently, student models learn from two supervision signals: ground-truth labels and the teacher model’s soft predictions. These signals are not equally aligned across training data points. For training data points that are already vulnerable to MIAs (i.e., MIA on this point is successful) on the teacher model, the teacher model assigns high probability to the ground-truth token and exhibits low divergence from the ground-truth distribution. In this case, distillation reinforces supervised learning, allowing the student model to fit these data points especially well. In contrast, for less vulnerable training data points (i.e., MIA is unsuccessful), teacher model predictions are less confident and deviate more from the ground-truth, introducing noise into the distillation signal and failing to provide clear privacy benefits.

We quantify this effect using two metrics: the teacher model’s probability assigned to the ground-truth next token and the KL divergence between the teacher model’s predictive distribution and the ground-truth distribution. As shown in Figure 1, attack-vulnerable training data points exhibit higher next-token probabilities and lower KL divergence, indicating strong alignment between supervision signals, whereas less vulnerable data points show the opposite trend. This asymmetry leads student models to preferentially learn and

Model	ReCaLL	Loss	Zlib	Min-K	Min-K++	Ref
Evaluated on Vulnerable (\mathcal{D}_v)						
Non-Vulnerable	0.250	0.205	0.114	0.364	0.318	0.114
Full	1.000	1.000	0.977	0.818	0.864	1.000
Evaluated on Non-vulnerable (\mathcal{D}_{nv})						
Non-Vulnerable	1.000	0.976	0.976	0.952	0.929	1.000
Full	1.000	0.952	0.452	0.833	0.905	1.000
Evaluated on Member (\mathcal{D})						
Non-Vulnerable	0.593	0.556	0.432	0.630	0.593	0.519
Full	1.000	0.975	0.556	0.827	0.889	1.000
Evaluated on Non-member (\mathcal{D}')						
Non-Vulnerable	0.963	0.988	0.988	0.877	0.914	1.000
Full	0.963	0.975	1.000	0.914	0.877	1.000

Table 4: MIA accuracies for DistilPythia student models trained using only non-vulnerable data points vs. trained on the full dataset. **Bold** entries indicate lower MIA accuracy (stronger privacy).

overfit already vulnerable points, amplifying the confidence gap between members and non-members and strengthening MIA decision boundaries.

5. PRIVACY-PRESERVING DISTILLATION METHODS

In this section, we develop methods for reducing membership inference risk in distilled models. We first consider a simple data-selection method that restricts distillation to *non-vulnerable* training data points. Although this approach directly reduces exposure to memorized data points, it also weakens the training signal and can degrade model utility. Motivated by this trade-off, we introduce two lightweight architectural modifications – **bottleneck projection** and **NoNorm** – that can be applied during distillation to reduce memorization while largely preserving utility.

5.1 Distillation on Non-Vulnerable Data

We consider a simple data-selection strategy that uses privacy signals from the teacher model to limit student model exposure to memorized content. We apply a MIA to the teacher model \mathcal{T} on its training dataset \mathcal{D} and partition \mathcal{D} into two disjoint subsets: a *vulnerable* set $\mathcal{D}_v = \{x \in \mathcal{D} \mid M(\mathcal{T}, x) = 1\}$, containing data points identified as members by the attack, and a *non-vulnerable* set $\mathcal{D}_{nv} = \{x \in \mathcal{D} \mid M(\mathcal{T}, x) = 0\}$. By construction, $\mathcal{D}_v \cup \mathcal{D}_{nv} = \mathcal{D}$ and $\mathcal{D}_v \cap \mathcal{D}_{nv} = \emptyset$.

This partition captures privacy-relevant structure in the teacher model: data points in \mathcal{D}_v exhibit behaviors that are easily distinguishable from non-members, whereas those in \mathcal{D}_{nv} do not. Accordingly, a natural approach is to restrict distillation to \mathcal{D}_{nv} , which directly reduces exposure to highly memorized data points and limits the transfer of membership signals from teacher to student.

5.1.1 Empirical Evidence of Effectiveness

We evaluate the non-vulnerable-only distillation strategy using DistilPythia derived from a Pythia teacher. Table 4 reports MIA accuracies for six attacks (ReCaLL, Loss, Zlib, Min-K%, Min-K%+, and Reference-model), stratified by four subsets: *Vulnerable* (\mathcal{D}_v), *Non-vulnerable* (\mathcal{D}_{nv}), *Member* (\mathcal{D}), and *Non-member* (\mathcal{D}'). We compare two student model variants: distillation using only \mathcal{D}_{nv} (“Non-Vulnerable”) and distillation using the full dataset \mathcal{D} (“Full”).

Model	Vulnerable	Non-vulnerable	Members
Non-Vulnerable	972.96	188.06	436.01
Full	94.45	149.54	118.21

Table 5: Perplexity of student models evaluated on the Vulnerable, Non-vulnerable, and Member subsets. **Bold** entries indicate lower perplexity (better model utility).

Restricting distillation to non-vulnerable data points substantially reduces attack success on both vulnerable and member subsets. Averaged across attacks, the non-vulnerable-only student reduces MIA accuracy by 75.02% on \mathcal{D}_v and 35.20% on \mathcal{D} , while achieving comparable accuracy on non-member data. These results confirm the privacy benefit of excluding highly memorized training data points during distillation.

However, this approach incurs clear utility costs. Table 5 reports perplexity on vulnerable, non-vulnerable, and member subsets, showing that the non-vulnerable-only student model consistently exhibits higher perplexity, reflecting degraded language-model quality due to reduced training data.

Overall, while simple and effective at reducing membership leakage, non-vulnerable-only distillation substantially weakens the training signal by reducing the effective training dataset size. With this, it demonstrates that privacy leakage can be mitigated through data selection, while also highlighting the need for complementary methods that reduce leakage without incurring significant utility loss. This motivates the architectural interventions introduced next.

5.2 Bottleneck Projection

Motivated by evidence that larger models are more prone to memorization and MIAs [7], we introduce a simple architectural modification that limits representational capacity: a low-dimensional **Bottleneck Projection** in the feed-forward network. Instead of the standard single projection from hidden dimension H to intermediate size I (typically $I \approx 4H$), we first project hidden states into a compact bottleneck space of dimension $B \ll H$, followed by expansion to I . This two-step projection has parameter cost $O(HB + BI)$, compared to $O(HI)$ for the standard design, yielding substantial parameter and computational savings when $B \ll I$.

Beyond efficiency, the **Bottleneck Projection** constrains intermediate representations, limiting the model’s ability to encode fine-grained, training data-specific signals. This restriction reduces memorization and weakens features exploited by MIAs. The **Bottleneck Projection** method is a lightweight modification that integrates seamlessly into existing transformer blocks and can be applied only to student models, preserving the capacity of the teacher.

5.3 NoNorm

Motivated by Sun et al. [34], who showed that replacing layer normalization can simplify model architecture and reduce inference latency², we introduce **NoNorm** as a privacy-preserving technique for knowledge distillation in LLMs. More specifically, the mean and variance computations in layer normalization may inadvertently encode information about the training data, posing a potential privacy risk. To address this concern, we replace layer normalization with a simpler element-wise affine transformation:

$$\text{NoNorm}(\mathbf{h}) = \gamma \circ \mathbf{h} + \beta,$$

where $\gamma, \beta \in \mathbb{R}^n$, n is the number of channels, \mathbf{h} denotes the hidden state, and \circ is the Hadamard product.

This modification has two key advantages. First, **NoNorm** avoids computing training data statistics, eliminating a potential channel through which training data information could be memorized and exploited by MIAs. Second, the transformation in **NoNorm** improves inference efficiency, yielding a simpler model that is less prone to memorization and thus potentially more robust to MIAs.

6. EXPERIMENTAL EVALUATION OF PRIVACY-PRESERVING DISTILLATION METHODS

We evaluate the effectiveness of the proposed architectural interventions – **Bottleneck Projection** and **NoNorm** – in reducing MIA vulnerability in distilled models.

6.1 Setup and Baselines

All experiments use DistilPythia student models trained for 30 epochs.³ We compare four student model variants derived from the same teacher: (i) **None**, a model trained with standard distillation with no privacy protection; (ii) **Bottleneck Projection**, which incorporates a low-dimensional **Bottleneck Projection** in the feed-forward layers; (iii) **NoNorm**, which replaces layer normalization with an element-wise affine transformation; and (iv) **All**, which combines **Bottleneck Projection** and **NoNorm**. Each variant is evaluated using six representative MIAs, with results reported separately for member and non-member data.

6.1.1 Hyperparameter Selection: Bottleneck Projection Dimensionality

To identify the optimal value for the bottleneck projection dimensionality B , we first fix all other training hyperparameters and train DistilPythia student variants for 30 epochs while varying $B \in \{48, 96, 192, 384, 768\}$. We report MIA accuracies for six representative attacks (ReCaLL, Loss, Zlib, Min-K%, Min-K%+, and Reference-model) and perplexity measured on the member subset.

²Although NoNorm has been proposed primarily to reduce model latency, it has not been studied as a defense mechanism against MIAs for privacy preservation in the LLM knowledge distillation scenario.

³We focus on the Pythia family due to computational constraints.

B	ReCaLL	Loss	Zlib	Min-K	Min-K++	Ref
Evaluated on Member (\mathcal{D})						
48	0.877	0.963	0.654	0.864	0.877	0.988
96	0.864	0.914	0.654	<u>0.938</u>	0.864	0.938
192	<u>0.926</u>	<u>0.975</u>	0.679	0.889	0.914	0.988
384	0.802	0.951	0.667	0.926	0.901	0.975
768	<u>0.926</u>	<u>0.975</u>	<u>0.704</u>	0.889	<u>0.938</u>	<u>1.000</u>
Evaluated on Non-member (\mathcal{D}')						
48	0.889	0.951	1.000	0.951	0.938	0.988
96	0.852	<u>0.975</u>	1.000	0.852	0.938	<u>1.000</u>
192	0.926	0.951	1.000	0.951	0.914	0.988
384	0.926	<u>0.975</u>	1.000	0.951	0.926	<u>1.000</u>
768	<u>0.938</u>	0.963	1.000	<u>0.963</u>	<u>0.951</u>	0.988

Table 6: MIA accuracies under different **Bottleneck Projection** dimension B . Within each column, **bold** indicates the lowest accuracy (best for privacy) and underline indicates the highest accuracy (worst for privacy) across different B values.

The privacy attack results are shown in Table 6. For the member side, smaller bottleneck projection dimensions often reduce MIA accuracy. In particular, $B = 768$ attains the highest MIA accuracy across 5 of the 6 MIA methods, and there is a general decreasing trend in MIA accuracy as B decreases. For the non-member case, we also observe that $B = 768$ achieves the highest MIA accuracy across 4 of the 6 MIA methods, and a similar decreasing trend as before. Therefore, we conclude that a bottleneck indeed has the ability to limit memorization and thus vulnerability to MIA.

B	Perplexity (Member)
48	200.17
96	<u>218.67</u>
192	161.56
384	138.44
768	134.99

Table 7: Perplexity on member subset for different **Bottleneck Projection** dimensions B . Underlined values indicate highest perplexity (worst utility); **bold** values indicate lowest perplexity (best utility).

The perplexity results are shown in Table 7. Perplexity on members generally improves as B increases; the largest B tested ($B = 768$) attains the lowest perplexity (best utility, 134.99), whereas $B = 96$ produced the worst perplexity (218.67). Intermediate values ($B = 192$ and $B = 384$) provide a compromise between privacy and utility. Thus, we use $B = 384$ for the **Bottleneck Projection** and **All** models in the following experiments, given its favorable privacy-utility trade-off.

6.2 MIA Accuracies

Table 8 reports MIA accuracy for all student model variants. On member data, the model with no privacy protection (**None**) exhibits the highest attack success across all methods. Introducing **Bottleneck Projection**, **NoNorm**, and their combination (**All**) reduces aggregate member-side attack accuracy by 9.45%, 4.71%, and 4.49%, respectively, indicating that both interventions effectively suppress memorization signals exploited by MIAs.

On non-member data, attack accuracies remain broadly comparable

Model	ReCaLL	Loss	Zlib	Min-K	Min-K++	Ref
Evaluated on Member (\mathcal{D})						
None	1.000	1.000	0.802	1.000	0.951	1.000
Bottleneck Projection	0.802	0.951	0.667	0.926	0.901	0.975
NoNorm	0.988	1.000	0.728	0.926	0.852	1.000
All	1.000	1.000	0.667	0.938	0.914	1.000
Evaluated on Non-member (\mathcal{D}')						
None	0.963	0.988	1.000	0.926	0.963	1.000
Bottleneck Projection	0.926	0.975	1.000	0.951	0.926	1.000
NoNorm	1.000	0.988	1.000	0.988	0.988	1.000
All	1.000	0.988	1.000	0.963	0.914	1.000

Table 8: MIA accuracies for DistilPythia student models under proposed architectural improvements. **Bolded** entries denote the lower MIA accuracy (better privacy).

Model	Perplexity (Member)
None	68.96
Bottleneck Projection	138.44
NoNorm	52.99
All	51.68

Table 9: Perplexity on held-out training data points for DistilPythia student models under proposed architectural improvements. **Bolded** entries denote lower perplexity (better language-model utility).

across variants. Relative to the model with no privacy protection, **Bottleneck Projection** slightly reduces non-member accuracy by 1.05%, while **NoNorm** and the combined model (**All**) increase it by 2.19% and 0.46%, respectively. These changes are small compared to the gains on member data. Since member identification constitutes the primary privacy risk in membership inference [4], improvements on the member subset are of primary importance.

Overall, the results show that the proposed architectural interventions substantially reduce member-specific vulnerability – the dominant privacy failure mode – while largely preserving privacy on non-member data.

6.3 Model Utility

We evaluate model utility using perplexity on the member subset, as shown in Table 9. **Bottleneck Projection** exhibits higher perplexity, reflecting the expected cost of reduced representational capacity. In contrast, **NoNorm** improves perplexity relative to the model with no privacy protection (**None**), and the combined model (**All**) achieves the lowest perplexity among all variants, demonstrating that **NoNorm** effectively mitigates the utility loss introduced by **Bottleneck Projection**.

Overall, the results demonstrate a favorable privacy-utility trade-off: while **Bottleneck Projection** alone reduces membership leakage (as illustrated in Table 8) at some cost to utility (as illustrated in Table 9), combining **Bottleneck Projection** with **NoNorm** recovers model utility while maintaining a substantial reduction in membership inference success.

7. CONCLUSION AND FUTURE WORK

In this paper, we study how knowledge distillation affects membership inference vulnerability in large language models. Across six teacher-student model pairs and six MIAs, we find that distillation

does not reliably improve privacy: student models can match or exceed their teacher models in aggregate and member-specific attack success. This arises from mixed supervision during distillation, where alignment between teacher model predictions and ground-truth labels on vulnerable training data points induces overconfident student model behaviors.

We propose three mitigation strategies: restricting distillation to non-vulnerable data, introducing a low-dimensional **Bottleneck Projection**, and replacing layer normalization with **NoNorm**. While data restriction increases perplexity, the architectural modifications reduce member-specific MIA success without degrading utility.

Future work includes theoretical analyses of memorization under mixed supervision, adaptive privacy-utility trade-offs, integration with formal privacy mechanisms, and evaluation across broader model families and tasks.

References

- [1] A. Anthropic. The claude 3 model family: Opus, sonnet, haiku. *Claude-3 Model Card*, 1:1, 2024.
- [2] S. Biderman, H. Schoelkopf, Q. G. Anthony, H. Bradley, K. O’Brien, E. Hallahan, M. A. Khan, S. Purohit, U. S. Prashanth, E. Raff, et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR, 2023.
- [3] C. Buciluă, R. Caruana, and A. Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541, 2006.
- [4] N. Carlini, S. Chien, M. Nasr, S. Song, A. Terzis, and F. Tramèr. Membership inference attacks from first principles. In *2022 IEEE symposium on security and privacy (SP)*, pages 1897–1914. IEEE, 2022.
- [5] N. Carlini, J. Hayes, M. Nasr, M. Jagielski, V. Shwag, F. Tramèr, B. Balle, D. Ippolito, and E. Wallace. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 5253–5270, 2023.
- [6] N. Carlini, D. Ippolito, M. Jagielski, K. Lee, F. Tramèr, and C. Zhang. Quantifying memorization across neural language models. In *The Eleventh International Conference on Learning*

Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023. OpenReview.net, 2023.

- [7] N. Carlini, F. Tramer, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, U. Erlingsson, et al. Extracting training data from large language models. In *30th USENIX security symposium (USENIX Security 21)*, pages 2633–2650, 2021.
- [8] chargodddard. Webinstructsub-prometheus, 2024.
- [9] R. Chourasia, B. Enkhtaivan, K. Ito, J. Mori, I. Teranishi, and H. Tsuchida. Knowledge cross-distillation for membership privacy. *Proceedings on Privacy Enhancing Technologies*, 2022(2):362–382, 2022.
- [10] Z. Cui, M. Zhang, and J. Pei. Learning to attack: Uncovering privacy risks in sequential data releases. *arXiv preprint arXiv:2510.24807*, 2025.
- [11] A. V. Duarte, X. Zhao, A. L. Oliveira, and L. Li. DE-COP: Detecting copyrighted content in language models training data. In R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 11940–11956. PMLR, 21–27 Jul 2024.
- [12] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, A. Goyal, A. Hartshorn, A. Yang, A. Mitra, A. Sravankumar, A. Korenev, A. Hinsvark, A. Rao, A. Zhang, A. Rodriguez, A. Gregerson, A. Spataru, B. Rozière, B. Biron, B. Tang, B. Chern, C. Caucheteux, C. Nayak, C. Bi, C. Marra, C. McConnell, C. Keller, C. Touret, C. Wu, C. Wong, C. C. Ferrer, C. Nikolaidis, D. Allonsius, D. Song, D. Pintz, D. Livshits, D. Esioibu, D. Choudhary, D. Mahajan, D. Garcia-Olano, D. Perino, D. Hupkes, E. Lakomkin, E. AlBadawy, E. Lobanova, E. Dinan, E. M. Smith, F. Radenovic, F. Zhang, G. Synnaeve, G. Lee, G. L. Anderson, G. Nail, G. Mialon, G. Pang, G. Cucurell, H. Nguyen, H. Korevaar, H. Xu, H. Touvron, I. Zarov, I. A. Ibarra, I. M. Kloumann, I. Misra, I. Evtimov, J. Copet, J. Lee, J. Geffert, J. Vranes, J. Park, J. Mahadeokar, J. Shah, J. van der Linde, J. Billock, J. Hong, J. Lee, J. Fu, J. Chi, J. Huang, J. Liu, J. Wang, J. Yu, J. Bitton, J. Spisak, J. Park, J. Rocca, J. Johnston, J. Saxe, J. Jia, K. V. Alwala, K. Upasani, K. Plawiak, K. Li, K. Heafield, K. Stone, and et al. The llama 3 herd of models. *CoRR*, abs/2407.21783, 2024.
- [13] L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy. The Pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- [14] J. Gou, B. Yu, S. J. Maybank, and D. Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021.
- [15] GPT-4 and Crumb. Distilpythia, 2023.
- [16] D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [17] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [18] A. Jagannatha, B. P. S. Rawat, and H. Yu. Membership inference attack susceptibility of clinical language models. *arXiv preprint arXiv:2104.08305*, 2021.
- [19] M. Jagielski, M. Nasr, K. Lee, C. A. Choquette-Choo, N. Carlini, and F. Tramer. Students parrot their teachers: Membership inference on model distillation. In *Advances in Neural Information Processing Systems (NeurIPS 2023)*, 2023.
- [20] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, and Q. Liu. TinyBERT: Distilling BERT for natural language understanding. In T. Cohn, Y. He, and Y. Liu, editors, *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4163–4174, Online, Nov. 2020. Association for Computational Linguistics.
- [21] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*, 2020.
- [22] K. Liang, W. Hao, D. Shen, Y. Zhou, W. Chen, C. Chen, and L. Carin. Mixkd: Towards efficient distillation of large-scale language models. *Arxiv preprint*, Nov. 2020.
- [23] Y. Liu, J. Cao, C. Liu, K. Ding, and L. Jin. Datasets for large language models: A comprehensive survey. *Artificial Intelligence Review*, 58(12):403, 2025.
- [24] M. Meeus, S. Jain, M. Rei, and Y.-A. 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.
- [25] M. Meeus, I. Shilov, S. Jain, M. Faysse, M. Rei, and Y.-A. de Montjoye. Sok: Membership inference attacks on llms are rushing nowhere (and how to fix it). In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 385–401. IEEE, 2025.
- [26] M. Mozes, X. He, B. Kleinberg, and L. D. Griffin. Use of llms for illicit purposes: Threats, prevention measures, and vulnerabilities. *arXiv preprint arXiv:2308.12833*, 2023.
- [27] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio. Fitnets: Hints for thin deep nets. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [28] V. Sanh, L. Debut, J. Chaumond, and T. Wolf. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108, 2019.
- [29] V. Shejwalkar and A. Houmansadr. Membership privacy for machine learning models through knowledge transfer. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*, 2021.
- [30] W. Shi, A. Ajith, M. Xia, Y. Huang, D. Liu, T. Blevins, D. Chen, and L. Zettlemoyer. Detecting pretraining data from large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- [31] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2017.

- [32] K. Song, H. Sun, X. Tan, T. Qin, J. Lu, H. Liu, and T.-Y. Liu. Lightpaff: A two-stage distillation framework for pre-training and fine-tuning. *arXiv preprint arXiv:2004.12817*, 2020.
- [33] S. Sun, Y. Cheng, Z. Gan, and J. Liu. Patient knowledge distillation for bert model compression. 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 4323–4332, 2019.
- [34] Z. Sun, H. Yu, X. Song, R. Liu, Y. Yang, and D. Zhou. MobileBERT: a compact task-agnostic BERT for resource-limited devices. In D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2158–2170, Online, July 2020. Association for Computational Linguistics.
- [35] Syed-Hasan-8503. Gemma-2-2b-it-distilled, 2024.
- [36] X. Tang, S. Mahloujifar, L. Song, V. Shejwalkar, M. Nasr, A. Houmansadr, and P. Mittal. Mitigating membership inference attacks by self-distillation. In *Proceedings of the 31st USENIX Security Symposium (USENIX Security 2022)*, 2022.
- [37] X. Tang, R. Shin, H. A. Inan, A. Manoel, F. Mireshghallah, Z. Lin, S. Gopi, J. Kulkarni, and R. Sim. Privacy-preserving in-context learning with differentially private few-shot generation. *arXiv preprint arXiv:2309.11765*, 2023.
- [38] G. Team, M. Riviere, S. Pathak, P. G. Sessa, C. Hardin, S. Bhupatiraju, L. Hussenot, T. Mesnard, B. Shahriari, A. Ramé, et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.
- [39] I. Timiryasov and J.-L. Tastet. Baby llama: knowledge distillation from an ensemble of teachers trained on a small dataset with no performance penalty. In A. Warstadt, A. Mueller, L. Choshen, E. Wilcox, C. Zhuang, J. Ciro, R. Mosquera, B. Paranjabe, A. Williams, T. Linzen, and R. Cotterell, editors, *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 279–289, Singapore, Dec. 2023. Association for Computational Linguistics.
- [40] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [41] C. Wang, Y. Wang, B. Hooi, Y. Cai, N. Peng, and K.-W. Chang. Con-recall: Detecting pre-training data in llms via contrastive decoding. *arXiv preprint arXiv:2409.03363*, 2024.
- [42] L. Watson, C. Guo, G. Cormode, and A. Sablayrolles. On the importance of difficulty calibration in membership inference attacks. In *International Conference on Learning Representations*, 2022.
- [43] R. Xie, J. Wang, R. Huang, M. Zhang, R. Ge, J. Pei, N. Z. Gong, and B. Dhingra. ReCaLL: Membership inference via relative conditional log-likelihoods. In Y. Al-Onaizan, M. Bansal, and Y.-N. Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8671–8689, Miami, Florida, USA, Nov. 2024. Association for Computational Linguistics.
- [44] X. Xu, M. Li, C. Tao, T. Shen, R. Cheng, J. Li, C. Xu, D. Tao, and T. Zhou. A survey on knowledge distillation of large language models. *arXiv preprint arXiv:2402.13116*, 2024.
- [45] S. Yeom, I. Giacomelli, M. Fredrikson, and S. 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.
- [46] M. Zhang, Y. Yang, R. Xie, B. Dhingra, S. Zhou, and J. Pei. Generalizability of large language model-based agents: A comprehensive survey. *arXiv preprint arXiv:2509.16330*, 2025.
- [47] R. Zhang, J. Shen, T. Liu, J.-L. Liu, M. Bendersky, M. Najork, and C. Zhang. Do not blindly imitate the teacher: Using perturbed loss for knowledge distillation. *arXiv preprint arXiv:2305.05010*, 2023.
- [48] J. Zheng, Y. Cao, and H. Wang. Resisting membership inference attacks through knowledge distillation. *Neurocomputing*, 452:114–126, 2021.