

On the Privacy Effect of Data Enhancement via the Lens of Memorization

Xiao Li, *Member, IEEE*, Qiongxiu Li, *Member, IEEE*, Zhanhao Hu, and Xiaolin Hu*, *Senior Member, IEEE*

Abstract—Machine learning poses severe privacy concerns as it has been shown that the learned models can reveal sensitive information about their training data. Many works have investigated the effect of widely adopted data augmentation and adversarial training techniques, termed data enhancement in the paper, on the privacy leakage of machine learning models. Such privacy effects are often measured by membership inference attacks (MIAs), which aim to identify whether a particular example belongs to the training set or not. We propose to investigate privacy from a new perspective called *memorization*. Through the lens of memorization, we find that previously deployed MIAs produce misleading results as they are less likely to identify samples with higher privacy risks as members compared to samples with low privacy risks. To solve this problem, we deploy a recent attack that can capture individual samples' memorization degrees for evaluation. Through extensive experiments, we unveil several findings about the connections between three essential properties of machine learning models, including privacy, generalization gap, and adversarial robustness. We demonstrate that the generalization gap and privacy leakage are less correlated than those of the previous results. Moreover, there is not necessarily a trade-off between adversarial robustness and privacy as stronger adversarial robustness does not make the model more susceptible to privacy attacks.

Index Terms—Privacy, Memorization, Data augmentation, Adversarial training, Membership inference attack.

I. INTRODUCTION

Due to the availability of ever-increasing datasets and computing power, we have witnessed a paradigm shift from traditional data analysis towards machine learning over the last decades. Rather than making prior assumptions, as is done in conventional modeling techniques, machine learning especially deep neural networks (DNNs) enables the system to directly learn from data. Machine learning has achieved superior performance and has been deployed in various applications such as image processing, natural language processing, etc. As the data used for training machine learning models are often collected from local devices such as tablets and wearable devices [1], it contains sensitive personal information such as speech, images, GPS location, and medical records. It is thus crucial to protect the training data from being revealed to others. However, several studies [2, 3, 4, 5] have shown that

machine learning models especially DNNs raise severe privacy concerns, as they tend to memorize sensitive information about the training data.

To quantitatively evaluate the privacy leakage that a machine learning model reveals about its training data, a basic approach that has been intensively used is the so-called *membership inference* [6]. That is, given access to a target model, the goal of the adversary is to determine whether a particular data point was used for training this target model (being a member) or not (being a non-member). Such membership information can reveal quite sensitive information about the individuals such as the health conditions [7] and serve as the basis for stronger types of privacy attacks [5].

Several studies show that the attack success rates of membership inference attacks (MIAs) are highly correlated with the generalization gap, i.e., the difference between training and test accuracies [6, 8, 9, 10, 11]. Such correlation is also observed when applying different data enhancement methods including data augmentation and adversarial training. It has been shown in Song et al. [9] that applying adversarial training can make the model more vulnerable to MIAs, and they conclude that the main reason is that the generalization gap becomes larger after applying adversarial training than standard training. The data augmentation methods, on the other hand, are widely believed to be effective in reducing the privacy leakage [6, 12, 10] as they are usually helpful in avoiding overfitting. Label smoothing [13], as a particular data augmentation method, however, is recognized to increase privacy leakage while reducing the generalization gap simultaneously [14, 15].

But the results shown in the aforementioned works might be misleading as the deployed MIAs for measuring the privacy leakage have the following limitations: 1) It has been criticized in several works [16, 17, 15] that the previous MIAs often have quite high false positive rates (FPR), i.e., many non-members are falsely identified as members. However, a good attack should obtain meaningful attack rates under low FPR regions, as it is more realistic for practical applications such as computer security [18, 19]. As an example shown by Carlini et al. [17], if an attack with overall 50.05% accuracy can reliably identify just 0.1% members without any false alarm, i.e., FPR=0, and judges the remaining samples by random guess with 50% accuracy, it puts much more risk to the model than another attack which guesses any sample with a chance of 50.05% being correct. In this case, the latter has a high FPR. 2) We find that the previous MIAs are inconsistent with the privacy risks on individual data points, even though they could have high overall success rates (see Section IV-B). Specifically, they have more difficulties in identifying training samples with high

X. Li, Q. Li, Z. Hu, and X. Hu are with the Department of Computer Science and Technology, Institute for Artificial Intelligence, BNRist, THBI, IDG/McGovern Institute for Brain Research, Tsinghua University, Beijing, 100084, China.

X. Hu is also with the Chinese Institute for Brain Research (CIBR), Beijing 100010, China.

*Corresponding author: Xiaolin Hu.

Email: {lixiao20, huzhanha17}@mails.tsinghua.edu.cn; {qiongxiuli, xlhu}@mail.tsinghua.edu.cn.

privacy risks as members compared to the samples with low privacy risks, which is at odds with the intuition that samples with higher privacy risks should be more easily identified.

We propose to address the above limitations by taking a new perspective called *memorization* [20, 21]. A data point is said to be memorized if the output of the model is quite sensitive to this individual data point, e.g., the prediction confidence of the learned model on this particular data point could be quite low unless it appears in the training set [20]. The concept of memorization fundamentally captures the privacy risk under the framework of differential privacy (DP) [22, 23, 20], which is generally considered to be a strong privacy definition. Empirically, we find that a recent attack called Likelihood Ratio Attack (LiRA) [17] is effective in reflecting the memorization degree, as we show in Section IV-C. LiRA also demonstrates much better performance under the low FPR regions compared with other MIAs. Therefore, we adopt LiRA to reinvestigate the privacy effects of both data augmentation and adversarial training.

A. Paper contribution

Through extensive investigations, we unveil several non-trivial findings (see Section VI and Section VII for details), which urge the community to rethink the relations among three important properties of machine learning models, including privacy leakage, generalization gap, and adversarial robustness. The major findings include:

- Unlike the previous studies [6, 8, 9, 10, 24] showing that the generalization gap and privacy leakage are highly correlated, our results demonstrate a much weaker correlation.
- Applying adversarial training can increase the memorization degrees of training samples, thereby resulting in more privacy leakage compared to models without adversarial training. However, for adversarially trained models, stronger adversarial robustness does not necessarily come with the cost of privacy leakage.

To the best of our knowledge, this is the first systematic evaluation of data augmentation and adversarial training via the lens of memorization.

B. Outline

The rest of the paper is organized as follows. Section II and III introduce related work and necessary fundamentals, respectively. Section IV explains the definition of memorization score and conducts statistical analysis to show the consistencies between the results of different MIAs and memorization scores. Section V investigates the privacy effect of applying data augmentation and adversarial training on machine learning models via extensive experiments. Section VI and VII analyze the relationship between privacy and generalization gap, and between privacy and adversarial robustness, respectively. Conclusions are given in Section IX.

II. RELATED WORK

This section reviews existing works on investigating the influence on privacy by applying data augmentation and adversarial training.

A. Data Augmentation and Privacy

Due to the fact that the attack success rate of many MIAs is highly correlated with the degree of overfitting, investigating how data augmentation affects privacy leakage has received attention [6, 12, 25, 14]. As a common practice to avoid overfitting, typical data augmentation methods such as flipping and cropping are shown to be effective in mitigating MIAs [12]. Kaya and Dumitras [14] further conduct a systematic investigation by applying seven data augmentation techniques such as label smoothing, random cropping, and mixup. The results show that it is difficult to use data augmentation to achieve substantial mitigation effects against MIAs while achieving better generalization gaps. In addition, label smoothing is shown to be able to increase both the privacy leakage and the test accuracy simultaneously [14, 15].

B. Augmented Information Improves Privacy Attack

It has been shown in several studies [26, 27, 25] that exploiting the information of augmented data would help to improve the attack success rate. MIAs can be classified into two types: augmentation-unaware and augmentation-aware attacks. The former assumes that the adversary does not have knowledge of the augmented data but simply uses random augmentation to probe the model. It has been shown that, by querying the model multiple times, using the random augmented data generated with Gaussian noise, the attack success rate can be improved [26]. The latter assumes a stronger scenario where the particular augmented data used in training is known to the adversary. Choquette-Choo et al. [27] show that the attack success rate can be significantly improved with the knowledge of the augmented data. Moreover, Yu et al. [25] showed that the augmentation-aware attack can obtain a higher success rate on models trained with some data augmentation than the ones without augmentation.

C. Adversarial training and privacy

Adversarial training is recognized to be one of the most effective ways to improve the adversarial robustness of DNNs [28, 29, 30, 31, 32], which is crucial in the security community. Compared to data augmentation, the relationship between adversarial training and privacy is relatively under-explored. It is shown in Song et al. [9] that the adversarially trained models, compared with the non-adversarially (standardly) trained models, are more susceptible to MIAs. They show that the privacy effect of adversarial training is related to several parameters including generalization gap, adversarial perturbation intensity, and model capacity. However, the MIAs used in the above-discussed studies are limited in reflecting the privacy risks of individual samples from the memorization perspective (as shown in Section IV-B). In addition, they all do not report results by the metric under low FPR regions.

III. PRELIMINARIES

In this section, we introduce the necessary fundamentals for understanding the rest of the paper.

A. Deep neural network

We consider feed-forward DNNs for classification tasks under the usual supervised setting. Suppose we have a training set $D_{\text{tr}} = \{(x, y) | (x, y) \in \mathcal{X} \times \mathcal{Y}\}$, where x is the feature vector (e.g., image) and y is the corresponding label. We denote the DNN model with parameter θ as $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$. During training, the data augmentation $T : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{P}$ is applied to the training data to improve its diversity, where \mathcal{P} is the set of all the probability measures defined on the power set $2^{\mathcal{X} \times \mathcal{Y}}$. Together, the optimal parameter θ^* of the model is fitted by:

$$\theta^* = \arg \min_{\theta} \sum_{(x, y) \in D_{\text{tr}}} \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim T(x, y)} [L(f_\theta(\tilde{x}), \tilde{y})], \quad (1)$$

where $L(\cdot, \cdot)$ is the loss function.

B. Data Enhancement

Since both data augmentation and adversarial training involve the process of adding certain examples into the training set to enhance the performances, they are known as data enhancement techniques. In this paper, we investigate eight popular data augmentation methods and four adversarial training methods, as described below.

Data augmentation methods:

- 1) Random Cropping and Flipping: sample new features by randomly cropping and horizontally flipping patches from the original feature in the training set.
- 2) Label smoothing [13]: replace the hard labels with the soft continuous labels by uniformly assigning probabilities to other classes. Therefore, the probability of the augmented label is $\tilde{p}_i = 1 - \frac{(n-1)\epsilon}{n}$ for $i = y$ and it is $\tilde{p}_i = \frac{\epsilon}{n}$ for $i \neq y$, where $\epsilon \in (0, 1)$ and n denotes the number of the classes.
- 3) Disturblabel [33]: change a portion of ground-truth (GT) labels to incorrect labels, namely, $\tilde{y} = \epsilon y + (1 - \epsilon)y_f$, where ϵ is randomly sampled from $\{0, 1\}$ and $y_f \in \{1, 2, \dots, n\} \setminus \{y\}$ denotes the incorrect label.
- 4) Gaussian Noise [34]: add Gaussian noise to each feature. The new feature $\tilde{x} = x + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$.
- 5) Cutout [35]: mask out a random square area of size $M \times M$ from each feature.
- 6) Mixup [36]: blend two features x_0, x_1 by a random ratio γ and creates a new feature $\tilde{x} = \gamma x_0 + (1 - \gamma)x_1$. The corresponding label is $\tilde{y} = \gamma y_0 + (1 - \gamma)y_1$.
- 7) Jitter [37]: randomly change the brightness, contrast, saturation, and hue of each image.
- 8) Distillation [38]: train an auxiliary DNN \hat{f} with the original training set and use the auxiliary DNN's soft outputs and temperature T as GT labels of the training features when training the target DNN. The temperature T determines the flatness of the soft labels.

Adversarial training techniques:

- 1) PGD-AT [39]: use PGD attack to generate adversarial examples x_{adv} based on original features and replaces the original feature with the adversarial examples at each iteration of the training, i.e., $\tilde{x} = x_{\text{adv}}$.
- 2) TRADES [40]: use PGD attack to generate adversarial examples x_{adv} , too. It differs from PGD-AT in that its

loss function consists of two components: $L(f_\theta(x), y) + L(f_\theta(x), f_\theta(x_{\text{adv}}))/\lambda$. The first component is the same as the loss of the standard training while the second component encourages the model to treat x and x_{adv} equally. The two components are weighted by λ . In the framework of Eq. (1), the discrete distribution of the transformation $T(x, y)$ can be represented as $\Pr(x, y) = \frac{\lambda}{\lambda+1}$ and $\Pr(x, f_\theta(x_{\text{adv}})) = \frac{1}{\lambda+1}$, where $\Pr(\cdot)$ indicates the probability.

- 3) AWP [41]: use regularization to explicitly flatten the weight loss landscape of PGD-AT by a double-perturbation mechanism.
- 4) TRADES-AWP [41]: incorporate the regularization mechanism of AWP into the TRADES method.

C. Membership Inference Attack

The goal of MIA is to identify whether a specific data sample was used in training a particular model or not. MIA has become one of the most widely investigated privacy attacks due to its simplicity. Many existing MIA approaches [6, 8, 10, 42, 43] can achieve high attack accuracy by exploiting the fact that machine learning models often behave differently to the data used or not used for training. For example, the model is often more confident about the training data than the test data. Thus, by setting a threshold to certain features such as the loss, confidence score, entropy, etc, the attack can achieve high accuracy in distinguishing members from non-members. Such kinds of approaches are referred to as metric-based approaches as they simply deploy a preset threshold to decide the membership. Another type of MIA approach is the classifier-based approach, which involves training a binary classifier to distinguish members and non-members. An effective technique to train such classifiers is known as shadow training [6]. The main idea of shadow training is to train several so-called shadow models, which are trained similarly to the target model to mimic its behavior. The classifier-based approaches tend to be more computationally complex compared to the metric-based methods.

IV. CONSISTENCY OF MIAS AND MEMORIZATION

In this section, our objective is to explain why existing research might yield misleading outcomes owing to two primary limitations and to demonstrate how adopting a memorization perspective can address these issues. To achieve this, we start by introducing the concept and definition of memorization scores. We then investigate the consistency between results obtained from different MIAs and the memorization scores. At last, we emphasize the need to reevaluate the privacy risks associated with data enhancements, using the lens of memorization scores to gain a more accurate understanding.

A. Memorization

A data point is said to be memorized by the model if it has a high impact on the model's behavior. To ensure that such impact is solely caused by this particular sample, one often needs to use the *leave-one-out* setting. Namely, except for this particular

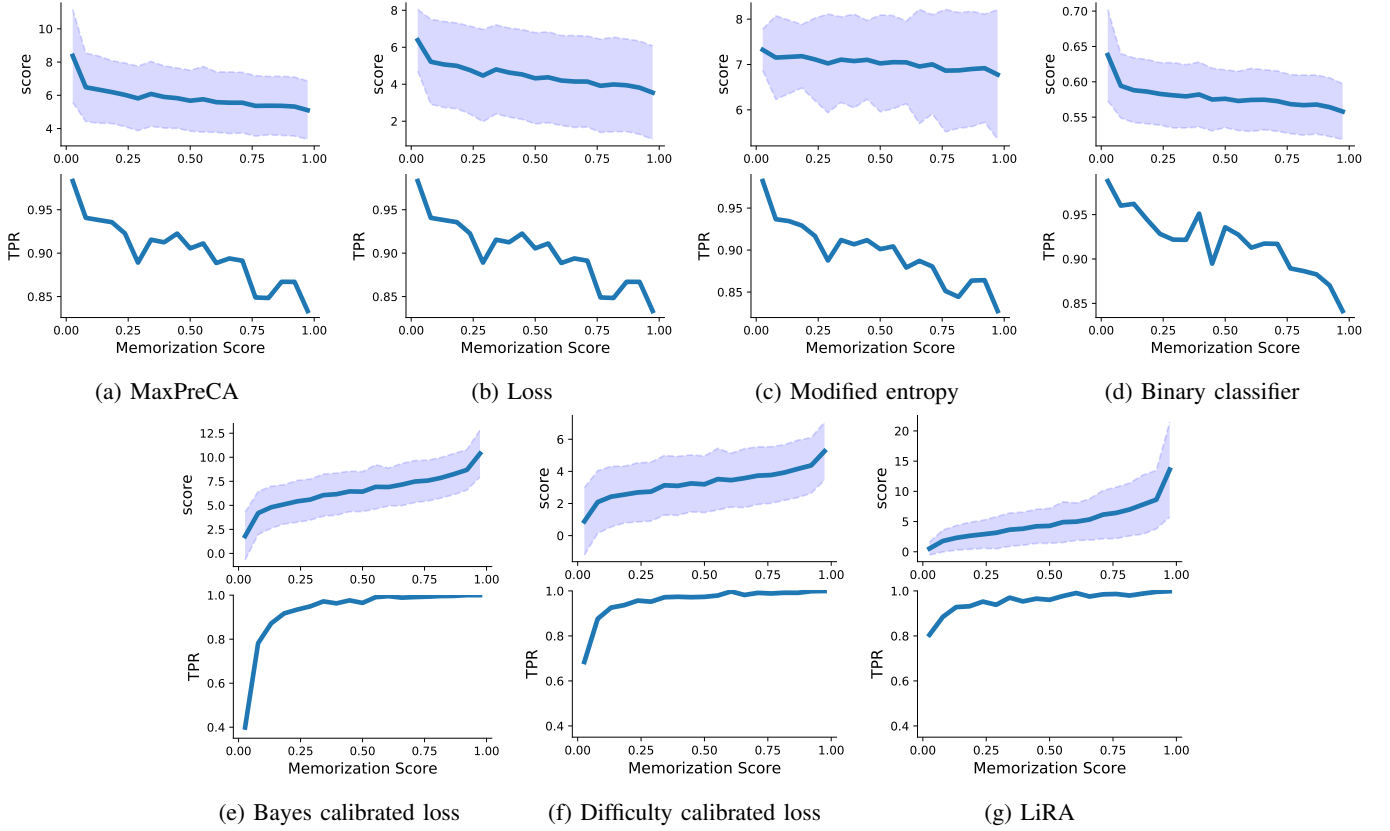


Fig. 1: Feature score (top) and TPR (bottom) of each bin in terms of memorization score for a target model trained on CIFAR-100 using (a) MaxPreCA [8], (b) Loss attack [10], (c) Modified entropy attack [44], (d) Binary classifier [6], (e) Bayes calibrated loss [12], (f) Difficulty calibrated loss [45], and (g) LiRA [17].

sample all other settings are the same. As an example, Feldman [20] defines a memorization score which measures how much information about the label of an individual data sample is being memorized by the model. Specifically, given the training set D_{tr} and the learning algorithm \mathcal{A} , for an arbitrary sample $(x, y) \in D_{\text{tr}}$, its (label) memorization score is defined as:

$$\text{mem}(\mathcal{A}, D_{\text{tr}}, (x, y)) = \Pr_{f_{\theta} \sim \mathcal{A}(D_{\text{tr}})}[f_{\theta}(x) = y] - \Pr_{f_{\theta} \sim \mathcal{A}(D_{\text{tr}} \setminus (x, y))}[f_{\theta}(x) = y], \quad (2)$$

where $D_{\text{tr}} \setminus (x, y)$ denotes the dataset D_{tr} with the sample (x, y) being removed. This definition is shown effective as it can assign atypical examples or outliers with high memorization scores and typical or easy samples with low memorization scores on various datasets including the CIFAR-100 [46] and ImageNet [37] datasets [21]. This complies perfectly with the intuition that an atypical example or outlier is often at a higher privacy risk, as the model will behave quite differently on it when it is in or out of the training set.

We now proceed to investigate whether the results produced by existing MIAs are consistent with memorization scores. Given that the memorization scores of the CIFAR-100 dataset have been pre-computed by Feldman and Zhang [21], where 4000 models were trained in total, we deploy it as the ground truth memorization scores for the following evaluations. Without loss of generality, we select representative approaches including both metric-based and classifier-based MIAs. Here

we trained 128 ResNet-18 [47] models with half of the data points randomly selected from 60,000 data points from CIFAR-100 (see the Hyper-Parameters paragraph in Section V-A for detailed training configurations). For each model, the other half was used as the non-member set for conducting MIA. Then we randomly chose ten trained models as the target model. The remaining 127 models were recognized as shadow models for each respective target model (if the MIA used shadow models). For each MIA, we used the optimum threshold that maximized the overall balanced accuracy, i.e., the proportion of correctly identified samples (both members and non-members) out of the total number of samples in the dataset¹. The final attack results were averaged over the ten target models. Below we present the results of consistency check, evaluating whether the attack results are consistent with memorization scores.

B. Example MIAs Obtaining Low Consistency with Memorization Scores

Many MIAs exploit the fact that models are overconfident about the training data [48]. Such confidences may be captured by different features, e.g., the prediction confidence [6, 8], loss [10], and so-called modified entropy [44]. That is, compared to the test data, the model usually has a higher confidence

¹Unless otherwise specified, all MIA thresholds mentioned in this work when evaluating balanced accuracy follow this setting.

score, smaller loss, and lower entropy (uncertainty) on the training data. For example, the main idea of the widely-adopted maximum predication confidence-based attack [8] (called MaxPreCA in this paper) is to classify samples with the maximum prediction confidence higher than a given threshold as members, otherwise as non-members:

$$\mathcal{M}_{\text{conf}}(x, y) = \mathbb{1}\{\max_c [f_\theta(x)_c] \geq \tau\} \quad (3)$$

where $\mathbb{1}$ denotes the indicator function, $f_\theta(x)_c$ denotes the model's confidence to be category c , and $\mathcal{M}_{\text{conf}}(x, y)$ denotes the adversary's decision on the membership. Similarly, the decisions of popular loss attack [10] is given by:

$$\mathcal{M}_{\text{loss}}(x, y) = \mathbb{1}\{-L(f_\theta(x), y) \geq \tau\}, \quad (4)$$

and the modified entropy attack [44] is given by:

$$\mathcal{M}_{\text{m-entr}}(x, y) = \mathbb{1}\{[1 - f_\theta(x)_y] \log[f_\theta(x)_y] + \sum_{c \neq y} f_\theta(x)_c \log[1 - f_\theta(x)_c] \geq \tau\}, \quad (5)$$

note that $f_\theta(x)_y$ denotes the model's confidence of x on the correct label y . Here we also investigate the classifier-based MIA, using the shadow training method proposed in Shokri et al. [6]. More specifically, for each target model, we used ten shadow models and the logits of these models and sample labels as features to train a multi-layer perceptron (MLP) classifier for binary membership classification. We denote this as a "binary classifier" approach.

To investigate if these attacks can capture the privacy risk of individual data points, in Fig. 1(a-d) we demonstrate their feature scores, i.e., the maximum prediction confidence, the loss, the modified entropy and the prediction score of binary-classifier (top panel), and the True Positive Rate (TPR) (bottom panel) versus the memorization score [21]. We first divide all the samples into 20 bins according to their memorization score. The top panel shows the averaged feature scores (the solid blue line) of the samples in different bins along with their standard deviation (purple shadow). For visualization, the feature scores of the top panels in Fig. 1(b-c) are transformed using exponential scaling (monotonic function), ensuring that they are on a similar scale as the other panels. The bottom panel shows the TPR of each bin for each MIA. It is obvious that the higher the memorization score is, the less probable it gets identified as members correctly. Therefore, the results of these attacks have a low consistency with memorization scores.

C. Example MIAs Obtaining High Consistency with Memorization Scores

Based on the definition in Eq. (2), MIAs that exhibit a high degree of consistency with memorization scores should effectively capture how a model's behavior varies when a specific sample is included or excluded from the training set. In this context, we highlight three example MIAs that leverage such behavioral differences: the Bayes calibrated loss approach [12], the difficulty calibrated loss approach [45], and LiRA [17]. Given that LiRA is essentially an evolved form of the

Method	Balanced Acc	TPR @ 0.1% FPR
MaxPreCA [8]	75.62	0.08
Loss [10]	75.48	0.14
Modified entropy [44]	75.66	0.21
Binary classifier [6]	76.50	0.32
Bayes calibrated loss [12]	79.75	4.95
Difficulty calibrated loss [45]	76.42	19.09
LiRA [17]	80.31	25.71

TABLE I: The attack success rate of different MIAs on CIFAR-100.

earlier two methods, our discussion will begin with an in-depth explanation of LiRA, and then we will provide a brief description of other approaches.

LiRA considers the distribution of the model's prediction on an individual data point when it is in or out of the training set. It requires training a number of shadow models such that for each sample (x, y) , half models include it in the training set and the other half models do not, denoted as IN and OUT models, respectively. Denote the sets of scaled confidences of sample (x, y) computed using IN and OUT models as \mathcal{Q}_{in} and \mathcal{Q}_{out} , respectively:

$$\begin{aligned} \mathcal{Q}_{\text{in}} &= \{\phi(f_\theta(x)_y) : (x, y) \in D_{\text{tr}}\} \\ \mathcal{Q}_{\text{out}} &= \{\phi(f_\theta(x)_y) : (x, y) \notin D_{\text{tr}}\}, \end{aligned} \quad (6)$$

where $\phi(p) = \log\left(\frac{p}{1-p}\right)$. \mathcal{Q}_{in} and \mathcal{Q}_{out} are used to fit two Gaussian distributions, denoted as IN distribution $\mathcal{N}(\mu_{\text{in}}, \sigma_{\text{in}}^2)$ and OUT distribution $\mathcal{N}(\mu_{\text{out}}, \sigma_{\text{out}}^2)$, respectively. Given an arbitrary sample, a standard likelihood-ratio test is performed to determine which distribution it more likely belongs to, where the likelihood ratio Λ is defined as:

$$\Lambda = \frac{\Pr(\phi(f_\theta(x)_y) | \mathcal{N}(\mu_{\text{in}}, \sigma_{\text{in}}^2))}{\Pr(\phi(f_\theta(x)_y) | \mathcal{N}(\mu_{\text{out}}, \sigma_{\text{out}}^2))}. \quad (7)$$

A sample will be classified as a member if Λ is higher than a threshold when evaluating balanced accuracy.

The Bayes calibrated loss method [12] is similar to LiRA except that it does not fit the output of shadow models using Gaussian. Instead, it directly uses the average output from the IN and OUT models for calibrating the sample loss. More precisely, the calibrated loss attack is given by²:

$$\mathcal{M}_{\text{bc}}(x, y) = \mathbb{1}\{f_\theta(x)_y - \mu(x, y) \geq \tau\} \quad (8)$$

where $\mu(x, y) = (\mu_{\text{in}}(x, y) + \mu_{\text{out}}(x, y))/2$ and μ_{in} and μ_{out} are the means computed like those in LiRA.

The difficulty calibrated loss approach [45] can be seen as an offline version of the above Bayes calibrated loss approach, which uses the difficulty of samples, measured by OUT models, to calibrate the loss, i.e., $\mu(x, y) = \mu_{\text{out}}(x, y)$. In Fig. 1(e-g), we present the Bayes calibrated loss, the difficulty calibrated loss, and the likelihood ratio Λ in Eq. (7) (top panel), along with the TPR (bottom panel), against the memorization scores. The results clearly show that samples with higher memorization scores are more likely to be accurately identified as members.

²Note that since we assume the output is confidence, thus in this work we use the confidence score instead of loss.

D. Consistency Results and Low FPR Metrics

As mentioned in the introduction and several previous studies [16, 17, 15], a prevalent limitation of many MIAs is their tendency to exhibit high FPRs. This issue becomes apparent when examining the consistency of these attack results with memorization scores. By inspecting Fig. 1(a-d), we identify two key observations. First, these attacks only achieve exceptionally high TPRs (e.g., above 0.95) within a narrow range of memorization scores. This implies that aiming for an overall low FPR will reduce the overall TPR across all regions significantly. Second, the highest TPRs are typically found in regions with very low memorization scores. This implies that aiming for an overall low FPR tends to exclude those members that present a higher privacy risk, which is also counter-intuitive. This phenomenon can be related to the fact that samples with high memorization scores are often more complex or atypical, making them more challenging for the model to fit [21]. Consequently, these samples tend to exhibit lower confidence compared to simpler, typical (test) samples. Therefore, implementing a stricter threshold to reduce the FPR, i.e., to prevent false identification of non-member test samples, will inadvertently exclude training samples with high memorization scores.

In contrast, MIAs with high consistency in memorization scores, as shown in Fig. 1(e-g), maintain very high TPRs across most regions, especially for those regions with high memorization scores. As a result, these MIAs can sustain high overall TPRs even in seeking low FPR scenarios, with the identified members being those with high memorization scores, which is a logical and expected outcome. To support these findings, Table I presents the performance of these MIAs, evaluating both the balanced accuracy and TPR at a low FPR. It is evident that MIAs with low consistency in memorization scores exhibit poor performance in low FPR regions, while those with high consistency show markedly better results.

Among the MIAs with high consistency with memorization scores, as shown in Fig. 1(e-g), we note that LiRA outperforms the other methods, especially in accurately classifying samples with low memorization scores. This superiority is attributed to LiRA's utilization of second-order statistics through Gaussian fitting, rather than solely relying on mean information like the other methods. Consequently, LiRA not only achieves the highest consistency with memorization scores but also the most effective attack success rate.

E. Discussion on Previous Evaluations of Data Enhancements

The findings discussed above raise significant concerns because many studies examining the privacy risks associated with data enhancements have relied on MIAs that show low consistency with memorization scores. For example, Salem et al. [8], Yeom et al. [10], Kaya and Dumitras [14], Yu et al. [25] utilize either confidence or loss-based MIAs to assess the privacy impact of data augmentation. Similarly, Song et al. [9] employs the confidence-based MIA to evaluate the risks associated with adversarial training. We suspect that the conclusions of these studies might be misleading. Therefore, it is important to reinvestigate the privacy risks associated with

data augmentation and adversarial training, employing MIAs that have high consistency with memorization scores for a more accurate evaluation.

Given the superior performance of LiRA in obtaining the highest consistency with memorization scores among the MIAs we examined, we will deploy LiRA as the primary tool to investigate how various data augmentation and adversarial training techniques impact privacy in subsequent discussions.

V. EVALUATING PRIVACY EFFECTS OF DATA ENHANCEMENT

We now proceed to evaluate the privacy effects of both data augmentation and adversarial training. Through the experiments, we aim to investigate the relations of privacy, adversarial robustness, and generalization gap, as they are all crucial properties for a machine learning model.

A. Experimental Settings

We used 32 NVIDIA 3080 GPUs to perform the experiments. The code for the experiments is implemented by `Pytorch` and is available at https://github.com/LixiaoTHU/privacy_and_aug.

Dataset. Following previous work [17, 49, 14], we used the CIFAR-10 [46] and CIFAR-100 [46] datasets for MIA evaluations. We additionally used the SVHN [50] dataset for a comprehensive evaluation. Both CIFAR-10 and CIFAR-100 contain 60,000 natural images with a resolution of 32×32 from 10 and 100 categories, respectively. SVHN contains 73,257 color images of numbers for training, with the resolution of 32×32 , and the first 60,000 training images from SVHN are used in this work.

MIA Settings. In this work, we assume a black-box setting where the adversary only has query access to the outputs of the target model on given samples. Without loss of generality, here we assume the output is the prediction confidence and the shadow models are trained using the same data enhancement method as the target model. We used LiRA to measure the privacy leakage of different data augmentation and adversarial training methods. We trained 128 models for each data enhancement method. Each model of the 128 models used roughly 30,000 data points as the training samples (members) and the remaining roughly 30,000 data points as the test samples (non-members). For each data point, we guaranteed that 64 out of the 128 models were the IN models, and the remaining 64 models were the OUT models. All models used the same training recipes except for the data enhancement strategies. For evaluation, we randomly selected 10 out of the 128 models as target models. The remaining 127 models were recognized as shadow models for each respective target model. After that, we performed the MIA evaluation on target models with all 60,000 data points. The MIA performance results were reported as the mean and standard deviation for each attack metric across the ten target models.

Models. Unless otherwise specified, we used the standard ResNet-18 [47] for the experiments on CIFAR-10 and CIFAR-100 and a vanilla convolutional neural network (CNN) for the experiments on SVHN. The CNN we used contained six

Method	Training Acc	Test Acc	TPR @ 0.1% FPR	TPR @ 0.001% FPR	Log-scale AUC	MIA Balanced Acc
Base	100.0 \pm 0.0	92.8 \pm 0.2	8.20 \pm 0.45	2.45 \pm 0.93	0.815 \pm 0.007	63.34 \pm 0.26
Smooth	100.0 \pm 0.0	92.9 \pm 0.3	5.22 \pm 0.66	0.14 \pm 0.07	0.734 \pm 0.012	62.28 \pm 0.86
Disturblabel	99.9 \pm 0.0	92.7 \pm 0.3	5.88 \pm 0.83	0.70 \pm 0.45	0.775 \pm 0.013	61.69 \pm 0.24
Noise	100.0 \pm 0.0	92.6 \pm 0.2	8.33 \pm 0.26	2.79 \pm 0.68	0.819 \pm 0.004	63.56 \pm 0.24
Cutout	100.0 \pm 0.0	93.1 \pm 0.4	7.71 \pm 0.39	2.48 \pm 1.03	0.811 \pm 0.010	63.23 \pm 0.26
Mixup	99.7 \pm 0.1	93.0 \pm 0.2	5.17 \pm 0.51	1.31 \pm 0.40	0.779 \pm 0.008	60.05 \pm 0.53
Jitter	100.0 \pm 0.0	92.7 \pm 0.2	8.24 \pm 0.35	2.97 \pm 0.76	0.819 \pm 0.004	63.41 \pm 0.31
Distillation	99.9 \pm 0.0	93.2 \pm 0.2	7.04 \pm 0.33	2.19 \pm 0.70	0.805 \pm 0.005	61.57 \pm 0.39
PGD-AT	99.2 \pm 0.1	82.2 \pm 0.2	23.78 \pm 0.89	10.52 \pm 2.30	0.897 \pm 0.005	78.82 \pm 0.37
TRADES	96.2 \pm 0.2	80.0 \pm 0.4	17.88 \pm 1.56	8.14 \pm 1.12	0.881 \pm 0.006	77.21 \pm 0.65
AWP	93.2 \pm 2.0	82.6 \pm 0.9	10.58 \pm 3.48	3.06 \pm 1.81	0.828 \pm 0.045	72.13 \pm 3.76
TRADES-AWP	91.9 \pm 0.5	80.5 \pm 0.2	12.43 \pm 0.89	3.48 \pm 1.36	0.848 \pm 0.006	74.86 \pm 0.80

TABLE II: Attack success rates of different data enhancement on CIFAR-10. The 2nd and 3rd columns show the training and test accuracies of each method, respectively. The 4th - 7th columns show four metrics to evaluate the extent of privacy leakage. We highlight the MIA success rates for different data augmentation and adversarial training methods that are larger than that for Base.

Method	Training Acc	Test Acc	TPR @ 0.1% FPR	TPR @ 0.001% FPR	Log-scale AUC	MIA Balanced Acc
Base	100.0 \pm 0.0	70.3 \pm 0.3	34.17 \pm 1.05	17.24 \pm 2.93	0.922 \pm 0.002	83.17 \pm 0.24
Smooth	100.0 \pm 0.0	72.2 \pm 0.4	39.21 \pm 1.25	19.88 \pm 3.86	0.932 \pm 0.004	86.35 \pm 0.22
Disturblabel	98.0 \pm 0.2	69.9 \pm 0.3	19.53 \pm 0.64	6.54 \pm 2.58	0.879 \pm 0.007	76.65 \pm 0.28
Noise	100.0 \pm 0.0	69.7 \pm 0.3	33.83 \pm 0.91	18.31 \pm 2.78	0.923 \pm 0.003	83.26 \pm 0.13
Cutout	100.0 \pm 0.0	70.3 \pm 0.3	34.71 \pm 1.58	17.25 \pm 5.02	0.923 \pm 0.005	83.53 \pm 0.22
Mixup	99.7 \pm 0.1	71.2 \pm 0.4	32.73 \pm 1.13	19.18 \pm 2.48	0.922 \pm 0.003	82.39 \pm 0.50
Jitter	100.0 \pm 0.0	70.3 \pm 0.3	34.19 \pm 0.90	18.37 \pm 3.62	0.924 \pm 0.003	83.35 \pm 0.17
Distillation	99.8 \pm 0.0	72.6 \pm 0.3	28.70 \pm 0.83	14.58 \pm 2.29	0.911 \pm 0.002	79.46 \pm 0.14
PGD-AT	99.5 \pm 0.0	51.3 \pm 0.3	68.63 \pm 0.88	47.85 \pm 4.26	0.972 \pm 0.001	93.62 \pm 0.10
TRADES	98.0 \pm 0.3	49.0 \pm 0.5	60.23 \pm 0.87	37.45 \pm 5.15	0.963 \pm 0.002	92.19 \pm 0.23
AWP	85.3 \pm 0.8	54.4 \pm 0.3	39.92 \pm 2.40	17.34 \pm 4.22	0.931 \pm 0.006	88.10 \pm 0.38
TRADES-AWP	95.9 \pm 0.6	51.3 \pm 0.4	57.51 \pm 2.02	35.57 \pm 3.73	0.960 \pm 0.002	91.92 \pm 0.36

TABLE III: Attack success rates of different data enhancement on CIFAR-100. The same conventions are used as in Table II.

convolutional layers with output channel numbers (256, 512, 512, 512, 512, 512), kernel sizes (3, 2, 3, 3, 2, 3), strides (1, 2, 1, 1, 2, 1), and padding sizes (1, 0, 1, 1, 0, 0). Batch normalization [51] was applied to each layer. After the image of 32×32 was down-sampled by these layers to be 6×6 , an average pooling operation was followed. Two fully connected layers, which can be seen as part of a MLP network, were added at the end of the output, with sizes 200 and 10, respectively. We denoted it as CNN-8 in the following text. Due to computational resource constraints, we did not conduct the same experiments on larger architectures.

Hyper-Parameters. Each model was optimized by stochastic gradient descent with an initial learning rate of 0.1 and a momentum of 0.9 for 100 epochs on a single GPU. Multi-step decay which scales the learning rate by 0.1 was used on the 75th and 90th epochs. The batch size is set to 256.

The hyper-parameters of each data augmentation method were set to achieve relatively high test accuracy by searching (see Appendix A for details). Unless other specified, for all adversarial training methods, we set the maximal perturbation ϵ under ℓ_∞ norm to be 8. We set the step size to be $\epsilon/8$ and the number of iterative steps to be 10. In addition, following the default setting of each method, the regularization parameter λ was set to be $1/6$ for TRADES, the perturbation intensity γ was

set to be 1×10^{-2} for AWP, and 5×10^{-3} for TRADES-AWP.

B. Evaluation Results

Tables II and III respectively show the training and test accuracies and the MIA results by multiple queries on CIFAR-10 and CIFAR-100. The results on SVHN are shown in Appendix B. We denote Random Cropping and Flipping by the *Base* method. Different from most previous studies [14, 25, 9], we evaluated the privacy leakage of seven data augmentation methods from label smoothing to Distillation and four adversarial training methods (Section III) combined with Base. This is a practical setting as Random Cropping and Flipping has now become a default setting in the computer vision field and it often brings considerable improvements in test accuracy. It has been criticized in Rezaei and Liu [16] that the models with low test accuracies are not practically useful for evaluating privacy leakage. See the results of different data augmentation methods without Base in Appendix B as examples, where the test accuracy of Base exceeds those of other data augmentation methods by at least 8.4% on CIFAR-10 and 12.6% on CIFAR-100. Unless otherwise specified, all data augmentation and adversarial training methods also use Base as default.




Fig. 2: Memorization scores of 5,000 randomly selected samples using Jitter (top), Disturblabel (middle) and PGD-AT model (bottom) v.s., Base model.

With LiRA, we evaluated all models using four metrics: TPR @ 0.1% FPR, TPR @ 0.001% FPR, Log-scale Area Under the Curve (AUC), and the Balanced Accuracy. The numbers after \pm denote the standard deviations. As mentioned earlier, it is more reasonable to use the metric of TPR under low FPR regions for evaluating privacy leakage. However, on the one hand, we empirically found that the results of TPR @ 0.001% FPR were unstable since their standard deviations were relatively large. On the other hand, 0.001% FPR might be too strict as it has been shown in Feldman and Zhang [21] that both CIFAR-10 and CIFAR-100 contain quite a few pairs of hard samples that are very similar. These samples will inevitably be misclassified with high confidence as members when their counterparts are in the member set, thereby causing some false positive cases and resulting in FPRs that exceed the 0.001% tolerance. Therefore, we mainly use TPR @ 0.1% FPR as the attack success rate for the following analysis.

C. Evaluation Results and Memorization Scores

To verify whether our evaluation results indeed reflect the degree of memorization, in Fig. 2 we compare the memorization scores of 5,000 randomly selected samples computed using the same method as in Feldman and Zhang [21] for three cases on CIFAR-100: Jitter, Disturblabel, and PGD-AT v.s. Base. Wherein the attack success rates are similar to, lower than, and higher than Base, respectively. Clearly, the corresponding changes in the memorization scores are consistent with the attack success rates. The memorization scores of the samples for Base and Jitter are similar (lying around the diagonal), which explains the similar attack success rate against the two methods. The memorization scores of many samples for Disturblabel are lower than for Base, especially the samples with high memorization scores for Base. Therefore, one reason why Disturblabel reduces privacy leakage is that it can reduce the memorization scores of many atypical samples. The memorization scores for PGD-AT are in general higher than those for Base (the points are distributed on the upper of the diagonal). Thus, one major reason why adversarial training causes a higher privacy leakage is that it memorizes many

training samples that are not memorized by standardly trained models. Overall, we conclude that our evaluation results are reliable as they can reflect the degree of memorization.

VI. PRIVACY AND GENERALIZATION GAP

In what follows, we proceed to analyze the relationship between privacy leakage and generalization gap.

A. Compared to prior works, our results demonstrate a much weaker correlation between the privacy leakage and generalization gap.

Many previous studies have shown that the attack success rates of MIAs are highly correlated with the generalization gap, i.e., the degree of overfitting [6, 8, 9, 10, 11]. To verify whether such a high correlation is still true from the memorization perspective, in Fig. 3, we demonstrate the attack success rate in terms of the train-test accuracy gap of all data augmentation models for all three datasets using MaxPreCA and LiRA, respectively. It is obvious that compared to the MaxPreCA attack results, our results demonstrate a more scattered distribution. We also compute the Pearson correlation coefficient r for each plot. As shown there, the Pearson correlation coefficients r of our results are significantly lower than the results using MaxPreCA, e.g., for CIFAR-10, our r is only 0.221, which is much lower than 0.708 using MaxPreCA. Note that the results using other MIAs in Section IV-B are similar to MaxPreCA (see Fig. E.1 in Appendix E). Hence, via the lens of memorization, the generalization gap and privacy leakage appear less correlated than those of the previous results.

It is easy to understand why the results of these attacks are sensitive to the generalization gap, as their success rate depends heavily on how different the model behaves for training and test samples. We remark that there is a distinction between memorization and overfitting: memorization is only necessary but not sufficient for overfitting [20], i.e., memorizing some training samples does not always cause overfitting. In fact, it has been both theoretically proved and empirically verified in previous work [20, 21] that memorizing certain long-tailed samples will help in decreasing the generalization

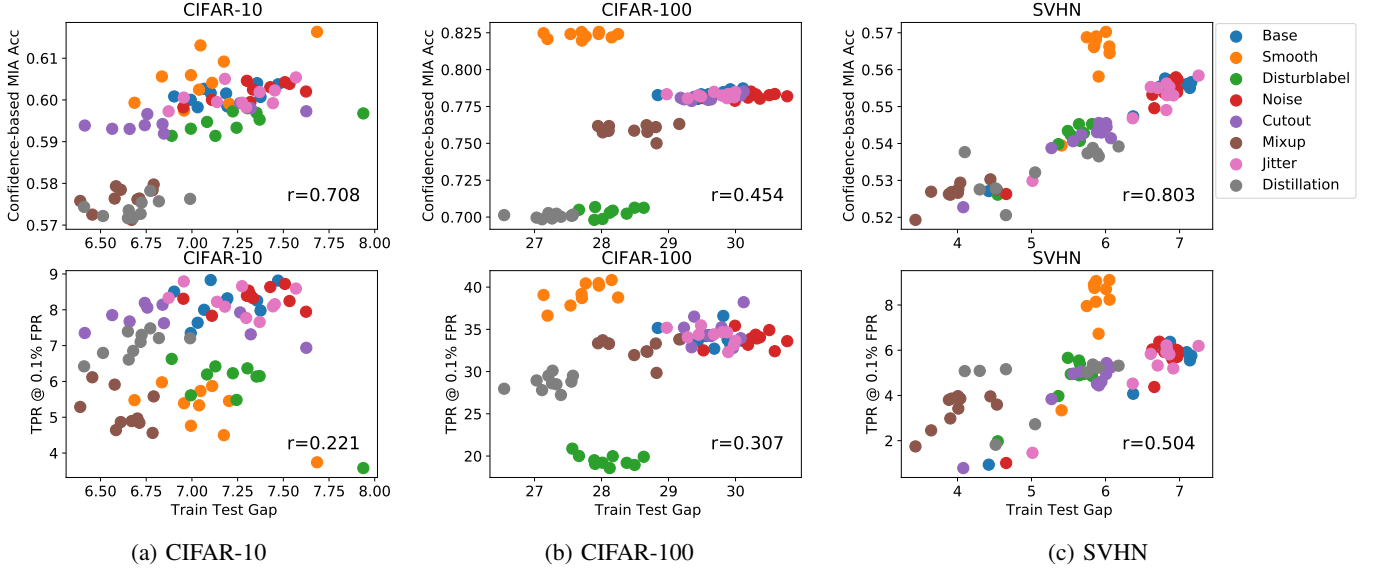


Fig. 3: Attack success rate versus the train-test gap of different data augmentation models on CIFAR-10, CIFAR-100, and SVHN using MaxPreCA (top) and LiRA (bottom), respectively. r stands for the Pearson correlation coefficient.




Fig. 4: Attack success rates of a single query and multiple queries in two cases: augmentation-unaware (left) and augmentation-aware (right). We evaluated different data augmentation methods on CIFAR-10 and CIFAR-100 datasets, respectively. *None* stands for models trained without any data augmentation.

gap. As a consequence, many attacks deployed in previous work might underestimate privacy leakage for non-overfitted models, making the correlation coefficient unnecessarily high. This issue can be alleviated in our setting as we measure the privacy leakage via the memorization perspective, the root cause of privacy leakage.

We remark that even though Yeom et al. [10] also pointed out that overfitting is not the only reason for causing vulnerability to privacy attacks, they did not explicitly identify what are other factors and their attack results still demonstrate a higher correlation compared with ours (as shown in Fig. 3). A related observation is made in Sections V-E of LiRA [17] showing that there are cases where two models have similar generalization gaps but the privacy leakages vary a lot. However,

it only gives one counterexample and does not investigate it deeply. Hence, their results can only support the claim that reducing the generalization gap does not necessarily make the model less vulnerable to privacy attacks, while we take a step further and demonstrate stronger results showing that via the memorization perspective, the correlation between privacy leakage and the generalization gap becomes weaker. To the best of our knowledge, such results have not been reported in existing work.

B. Data augmentation is not necessarily an effective defense for MIAs.

By inspecting the attack results for models with data augmentation, we can see that the privacy effects vary significantly across different data augmentation methods. For example, Distillation and Disturblabel are shown effective in reducing the vulnerability to privacy attack. Mixup, Cutout, Jitter, and Gaussian noise methods do not seem to have big impacts on the attack success rate. The main reason is that applying data augmentation does not always reduce the memorization scores of training samples, e.g., the Jitter model shown in Fig. 2.

Moreover, among all data augmentation methods, label smoothing has drawn attention as it has been shown in both Kaya and Dumitras [14] and Hintersdorf et al. [15] that applying label smoothing will make the model more susceptible to MIAs. To verify this, we computed the balanced accuracy using MaxPreCA [8]. As shown in the top panels of Fig. 3, label smoothing does increase the attack accuracies compared to Base for all datasets. However, from the memorization perspective, it did not demonstrate the same tendency. By inspecting the bottom panels of Fig. 3 we note that label smoothing demonstrates an inconsistent behavior on different datasets. On CIFAR-100 and SVHN the privacy leakage is higher than Base while on CIFAR-10 the privacy leakage is lower. Hence, the claim that label smoothing would consistently




Fig. 5: The distributions of normalized confidence ϕ of three samples with different memorization scores using Base and four adversarial training models. Each row corresponds to a sample.




Fig. 6: Attack success rates of PGD-AT and TRADES under different ϵ on the three datasets.

amplify privacy leakage is untrue. Overall, we conclude that it is difficult to give a general claim about whether data augmentation can help mitigate the privacy attack or not. We remind that extra attention should be paid when relying on data augmentation as a defense technique against MIAs.

C. Multiple queries can only enhance the attack if the augmentation method is known.

As stated in Section I, previous studies have shown that using augmented data to conduct multiple queries would enhance the attack success rate. To investigate this, we queried the target model using ten augmented counterparts generated by the Base method for each data point. We then targeted all the data augmentation models trained on Base as the augmentation-aware case. The augmentation-unaware case was then evaluated by targeting the data augmentation models without using Base (such models were trained on CIFAR-10 and CIFAR-100, see Appendix B). As shown in Fig. 4, multiple queries did help improve the attack success rate for the augmentation-aware case, whereas, for the augmentation-unaware case, they resulted in an opposite effect, i.e., lowering the attack success rate. We note that similar results have been reported in Carlini et al. [17]

but only with Base augmentation. Our results further confirm their results by extending to several different data augmentation methods.

VII. PRIVACY AND ADVERSARIAL ROBUSTNESS

In this section, we further analyze the relationship between privacy leakage and adversarial robustness.

A. Applying adversarial training will make the model memorize more training samples, thereby causing more privacy leakages compared to the standardly trained models.

As shown in Tables II, III, and B.1, applying adversarial training significantly increases privacy leakage compared to the standardly trained models. For example, the TPR @ 0.1% FPR increases from 34.17% to 60.23% for Base to TRADES on CIFAR-100. One reason is that applying adversarial training will force the model to fit all the adversarial examples found in the ℓ_∞ ball around each training sample, which often increases the influence of each sample on the trained model, thereby resulting in a higher privacy risk. To visualize the effect of applying adversarial training, in Fig. 5, we choose three examples in CIFAR-100 with different memorization scores and

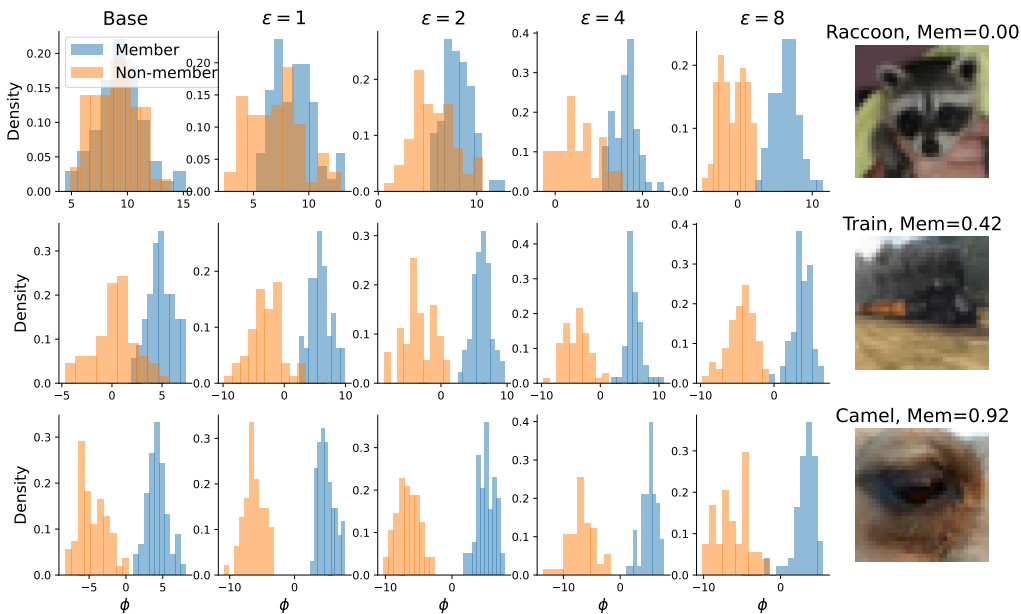


Fig. 7: The distributions of normalized confidence ϕ of three samples with different memorization scores using Base and PGD-AT under four different ϵ . Each row corresponds to a sample.

draw their corresponding distributions of normalized confidence ϕ evaluated by IN and OUT models using Base and all four types of adversarially trained models. Clearly, if the IN and OUT distributions of a particular sample are more separated, it implies that the sample is adversarial training a higher privacy risk. We can see that for samples that have low privacy risks (e.g., *Raccoon* and *Train*), applying adversarial training would make the distribution more separable.

Note that there is a bottleneck of increasing the privacy risk of samples when performing adversarial training: for samples that are already at a high privacy risk (e.g., *Camel* with a high memorization score), applying adversarial training would not make much difference as the distributions are already quite separated for standardly trained models. Overall, we conclude that one major reason why adversarial training causes a higher privacy leakage is that it memorizes many training samples that are not memorized by standardly trained models.

B. Better adversarial robustness does not necessarily make the adversarially trained model more vulnerable to privacy attacks.

To further investigate the relation between adversarial robustness and privacy leakage, in Table IV we compare adversarial robustness and attack success rate using different adversarial training methods on the three datasets. We can see that compared to TRADES and PGD-AT, both AWP and TRADES-AWP achieve higher adversarial accuracies, while the attack success rates are lower. Hence, for adversarially trained models, stronger adversarial robustness does not necessarily come with a cost of privacy leakage.

In addition to the attack results using different adversarial training methods, it is also interesting to see how the attack result changes along with varying parameters. Since ϵ is a critical parameter for adversarial training, in Fig. 6, we

Dataset	Method	Adversarial Acc	TPR @ 0.1% FPR
CIFAR-10	Base	0.0 ± 0.0	8.20 ± 0.45
	PGD-AT	38.8 ± 0.4	23.78 ± 0.89
	TRADES	45.2 ± 0.3	17.88 ± 1.56
	AWP	45.9 ± 0.1	10.58 ± 3.48
	TRADES-AWP	48.8 ± 0.2	12.43 ± 0.89
CIFAR-100	Base	0.0 ± 0.0	34.17 ± 1.05
	PGD-AT	16.9 ± 0.1	68.63 ± 0.88
	TRADES	19.7 ± 0.4	60.23 ± 0.87
	AWP	23.9 ± 0.1	39.92 ± 2.40
	TRADES-AWP	23.3 ± 0.2	57.51 ± 2.02
SVHN	Base	0.0 ± 0.0	5.27 ± 1.57
	PGD-AT	45.2 ± 0.4	10.55 ± 0.25
	TRADES	47.4 ± 0.1	10.31 ± 0.22
	AWP	43.7 ± 0.5	8.85 ± 0.24
	TRADES-AWP	46.2 ± 0.4	8.54 ± 0.22

TABLE IV: The accuracies on adversarial examples (Adversarial Acc) and privacy leakage of different adversarial training models on three datasets. The accuracies are evaluated using PGD with $\epsilon = 8$ and 20 iteration steps.

compare the attack results of different ϵ using both PGD-AT and TRADES models on the three datasets. We can see that overall the attack success rate tends to increase along with ϵ (at least for $\epsilon < 8$). One reason might be that increasing the perturbation parameter ϵ will result in a bigger ℓ_∞ ball around each training sample thereby increasing the difficulty of fitting all adversarial examples. Consequently, the model will be more sensitive to this sample, thus increasing the privacy risk. To verify this, in Fig. 7, we use three examples (the same as in Fig. 5) with different memorization scores and draw their corresponding distributions of normalized confidence ϕ evaluated by IN and

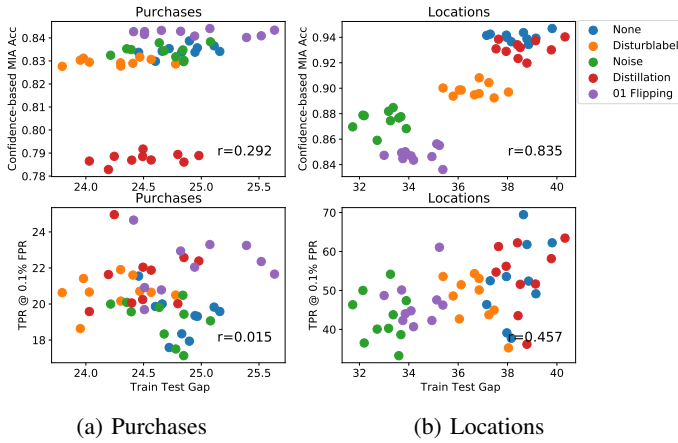


Fig. 8: Attack success rate versus the train-test gap of different data augmentation models on the Purchases and Locations datasets using MaxPreCA (top) and LiRA (bottom), respectively. r stands for the Pearson correlation coefficient.

OUT models using Base and PGD-AT models under four different ϵ . Clearly, we can see that for the sample with a very low memorization score (*Raccon*), increasing ϵ makes the distributions more separable thus posing a higher privacy risk. Similar to Fig. 5, we observe a bottleneck effect that for samples with high memorization score (*Camel*), varying ϵ will not make much difference. Moreover, such bottleneck effect is quicker since for samples with a fair memorization score (*Train*), the IN and OUT distributions are already quite separable when $\epsilon = 1$.

VIII. DISCUSSIONS

This section extends our discussion to encompass a broader range of experimental results, providing further evidence of the generalizability of our findings across diverse model architectures, datasets, and MIAs.

A. Expanded Experiments across Diverse Model Architectures and Datasets

To further verify the generalizability of our key conclusions, initially derived from experiments on three image datasets (CIFAR-10, CIFAR-100, and SVHN) using ResNet-18, ResNet-18, and CNN-8, respectively, we conducted additional experiments involving other combinations of datasets and architectures. Specifically, we replicated the experiments using CNN-8 on CIFAR10 as an alternative combination. Furthermore, we performed similar experiments with an MLP architecture on two additional non-image datasets, namely Purchases and Locations processed by Shokri et al. [6].

In Appendix C, we present the results of experiments conducted using CNN-8 on CIFAR-10. These findings consistently align with our previous key observations: 1) A less pronounced correlation between generalization gaps and privacy leakages, as indicated by the reduced correlation coefficient in Fig. C.1; 2) Adversarially trained models memorize more training data compared to standardly trained models, as evidenced by the much higher privacy leakage in Table C.1; 3) An increase in the

adversarial robustness of the adversarially trained models does not necessarily come with a price of heightened vulnerability to privacy attacks, as demonstrated by the performance of AWP in Table C.1. Moreover, the trends depicted in Fig. C.2 closely resemble those observed with the ResNet-18 architecture on CIFAR-10 in Fig. 6.

For the non-image datasets, we performed experiments with an MLP architecture on the Purchases and Locations datasets. We note that part of the evaluated data augmentation methods relevant to the image datasets, e.g., Cropping, Jitter, Cutout, etc, are inapplicable to these non-image datasets. Here we investigated the privacy impacts of applying Disturblabel, Distillation, and Gaussian Noise. Additionally, we investigated a 0-1 Flipping augmentation method specially designed for these datasets. We did not evaluate adversarial training methods with the non-image datasets because adversarial training is primarily focused on the image domain [39, 40, 41], and there is no standardized definition of adversarial examples for these non-image datasets. The privacy leakage results and the experimental details are presented in Appendix D. In Fig. 8, we demonstrate the attack success rate regarding the train-test accuracy gap of all data augmentation models for these non-image datasets using MaxPreCA and LiRA, respectively. We can see that our results with LiRA demonstrate a more dispersed distribution and a lower Pearson correlation coefficient compared with the MaxPreCA attack results.

Overall, we conclude that our key findings remain consistent across diverse model architectures and datasets.

B. Other MIAs Beyond LiRA

While our initial investigations of privacy leakage focused on LiRA, it is important to clarify that our primary goal was to assess the privacy implications of data enhancement in machine learning models, particularly via the perspective of memorization. The choice of LiRA was driven by its high consistency with memorization scores, as demonstrated in Fig. 1. However, our findings are not limited to LiRA alone; they extend to other MIAs that demonstrate a strong correlation with memorization scores. To validate this, we shifted our attention to the difficulty calibrated loss approach [45], which shows comparable consistency with memorization scores (see Fig. 1(e)). The results, depicted in Fig. E.2 and Table E.1, indicate that the general trends in privacy effects resulting from data augmentation and adversarial training align with those observed using LiRA (see Fig. 6 and Table IV). It is worth noting that the correlation score is slightly higher and the overall attack performance is lower compared to LiRA, which aligns with the expectations given LiRA’s higher consistency and superior attack performance against the importance calibrated loss approach. Consequently, we conclude that our findings apply to other MIAs as long as they exhibit a high level of consistency with memorization scores.

IX. CONCLUSION

In this paper, we reinvestigate the privacy effect of applying data augmentation and adversarial training to machine learning models via a new perspective, namely the degree

of memorization. Such reinvestigation is quite necessary as we found that the attacks deployed in previous studies for measuring privacy leakage produce misleading results: the training samples with low privacy risks are more prone to be identified as members compared to the ones with high privacy risks. Through a systematic evaluation, we reveal some findings conflict with previous results, e.g., the generalization gap and privacy leakage are shown less correlated than those of the previous results and label smoothing does not always amplify the privacy leakage. Moreover, we also show that improving the adversarial robustness (via adversarial training) does not necessarily make the adversarially trained model more vulnerable to privacy attacks. Our results call for more investigations on the privacy of machine learning models from the memorization perspective.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (Nos. U2341228 and U19B2034).

REFERENCES

- [1] J. Poushter *et al.*, “Smartphone ownership and internet usage continues to climb in emerging economies,” *Pew Research Center*, vol. 22, pp. 1–44, 2016.
- [2] C. Song, T. Ristenpart, and V. Shmatikov, “Machine learning models that remember too much,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2017, pp. 587–601.
- [3] N. Carlini, C. Liu, Ú. Erlingsson, J. Kos, and D. Song, “The secret sharer: Evaluating and testing unintended memorization in neural networks,” in *USENIX Security Symposium (USENIX)*, 2019, pp. 267–284.
- [4] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding deep learning (still) requires rethinking generalization,” *Commun. ACM*, vol. 64, no. 3, pp. 107–115, 2021.
- [5] N. Carlini, F. Tramèr, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. B. Brown, D. Song, Ú. Erlingsson, A. Oprea, and C. Raffel, “Extracting training data from large language models,” in *USENIX Security Symposium (USENIX)*, 2021, pp. 2633–2650.
- [6] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *IEEE Symposium on Security and Privacy (S&P)*, 2017, pp. 3–18.
- [7] A. H. Poorjam, Y. P. Raykov, R. Badawy, J. R. Jensen, M. G. Christensen, and M. A. Little, “Quality control of voice recordings in remote parkinson’s disease monitoring using the infinite hidden markov model,” in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, pp. 805–809, 2019.
- [8] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, and M. Backes, “MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models,” *Network and Distributed Systems Security Symposium*, 2019.
- [9] L. Song, R. Shokri, and P. Mittal, “Privacy risks of securing machine learning models against adversarial examples,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2019, pp. 241–257.
- [10] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha, “Privacy risk in machine learning: Analyzing the connection to overfitting,” in *IEEE Computer Security Foundations symposium (CSF)*, 2018, pp. 268–282.
- [11] KlasLeino and M. Fredrikson, “Stolen memories: Leveraging model memorization for calibrated {White-Box} membership inference,” in *USENIX Security Symposium (USENIX)*, 2020, pp. 1605–1622.
- [12] A. Sablayrolles, M. Douze, C. Schmid, Y. Ollivier, and H. Jégou, “White-box vs black-box: Bayes optimal strategies for membership inference,” in *Int. Conf. Mach. Learn. (ICML)*, 2019, pp. 5558–5567.
- [13] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2016, pp. 2818–2826.
- [14] Y. Kaya and T. Dumitras, “When does data augmentation help with membership inference attacks?” in *Int. Conf. Mach. Learn. (ICML)*, 2021, pp. 5345–5355.
- [15] D. Hintersdorf, L. Struppek, and K. Kersting, “To trust or not to trust prediction scores for membership inference attacks,” in *IJCAI*, 2022, pp. 3043–3049.
- [16] S. Rezaei and X. Liu, “On the difficulty of membership inference attacks,” in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2021, pp. 7892–7900.
- [17] N. Carlini, S. Chien, M. Nasr, S. Song, A. Terzis, and F. Tramèr, “Membership inference attacks from first principles,” in *IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2022, pp. 1897–1914.
- [18] J. Z. Kolter and M. A. Maloof, “Learning to detect and classify malicious executables in the wild,” *Journal of Machine Learning Research*, vol. 7, no. 12, 2006.
- [19] A. Lazarevic, L. Ert’oz, V. Kumar, A. Ozgur, and J. Srivastava, “A comparative study of anomaly detection schemes in network intrusion detection,” in *Proceedings of the SIAM International Conference on Data Mining*, 2003, pp. 25–36.
- [20] V. Feldman, “Does learning require memorization? a short tale about a long tail,” in *Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 2020, pp. 954–959.
- [21] V. Feldman and C. Zhang, “What neural networks memorize and why: Discovering the long tail via influence estimation,” *Adv. Neural Inform. Process. Syst. (NeurIPS)*, vol. 33, pp. 2881–2891, 2020.
- [22] C. Dwork, “Differential privacy,” in *International Colloquium on Automata, Languages and Programming (ICALP)*, 2006, pp. 1–12.
- [23] C. Dwork, F. McSherry, K. Nissim, and A. D. Smith, “Calibrating noise to sensitivity in private data analysis,” *J. Priv. Confidentiality*, vol. 7, no. 3, pp. 17–51, 2016.
- [24] H. Hu, Z. Salicic, G. Dobbie, and X. Zhang, “Membership inference attacks on machine learning: A survey,” *ACM Computing Surveys (CSUR)*, 2021.
- [25] D. Yu, H. Zhang, W. Chen, J. Yin, and T. Liu, “How does data augmentation affect privacy in machine learning?”

- in *AAAI*, vol. 35, 2021, pp. 10 746–10 753.
- [26] B. Jayaraman, L. Wang, K. Knipmeyer, Q. Gu, and D. Evans, “Revisiting membership inference under realistic assumptions,” *Proceedings on Privacy Enhancing Technologies*, vol. 2021, no. 2, 2021.
- [27] C. A. Choquette-Choo, F. Tramèr, N. Carlini, and N. Papernot, “Label-only membership inference attacks,” in *Int. Conf. Mach. Learn. (ICML)*, 2021, pp. 1964–1974.
- [28] A. Athalye, N. Carlini, and D. A. Wagner, “Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples,” in *Int. Conf. Mach. Learn. (ICML)*, vol. 80, 2018, pp. 274–283.
- [29] F. Tramèr, N. Carlini, W. Brendel, and A. Madry, “On adaptive attacks to adversarial example defenses,” in *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2020.
- [30] X. Li, Z. Wang, B. Zhang, F. Sun, and X. Hu, “Recognizing object by components with human prior knowledge enhances adversarial robustness of deep neural networks,” *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, vol. 45, no. 7, pp. 8861–8873, 2023.
- [31] X. Li, W. Zhang, Y. Liu, Z. Hu, B. Zhang, and X. Hu, “Language-driven anchors for zero-shot adversarial robustness,” in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2024.
- [32] X. Li, H. Chen, and X. Hu, “On the importance of backbone to the adversarial robustness of object detectors,” *arXiv preprint arXiv:2305.17438*, 2023.
- [33] L. Xie, J. Wang, Z. Wei, M. Wang, and Q. Tian, “Disturblabel: Regularizing cnn on the loss layer,” in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2016, pp. 4753–4762.
- [34] J. M. Cohen, E. Rosenfeld, and J. Z. Kolter, “Certified adversarial robustness via randomized smoothing,” in *Int. Conf. Mach. Learn. (ICML)*, 2019, pp. 1310–1320.
- [35] T. Devries and G. W. Taylor, “Improved regularization of convolutional neural networks with cutout,” *arXiv preprint arXiv:1708.04552*, 2017.
- [36] H. Zhang, M. Cissé, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” in *Int. Conf. Learn. Represent. (ICLR)*, 2018.
- [37] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Adv. Neural Inform. Process. Syst. (NeurIPS)*, vol. 25, 2012.
- [38] G. E. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [39] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” in *Int. Conf. Learn. Represent. (ICLR)*, 2018.
- [40] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan, “Theoretically principled trade-off between robustness and accuracy,” in *Int. Conf. Mach. Learn. (ICML)*, 2019, pp. 7472–7482.
- [41] D. Wu, S. Xia, and Y. Wang, “Adversarial weight perturbation helps robust generalization,” in *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2020.
- [42] Y. Long, V. Bindschaedler, L. Wang, D. Bu, X. Wang, H. Tang, C. A. Gunter, and K. Chen, “Understanding membership inferences on well-generalized learning models,” *arXiv preprint arXiv:1802.04889*, 2018.
- [43] J. Jia, A. Salem, M. Backes, Y. Zhang, and N. Z. Gong, “Memguard: Defending against black-box membership inference attacks via adversarial examples,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2019, pp. 259–274.
- [44] L. Song and P. Mittal, “Systematic evaluation of privacy risks of machine learning models,” in *USENIX Security Symposium (USENIX)*, 2021, pp. 2615–2632.
- [45] L. Watson, C. Guo, G. Cormode, and A. Sablayrolles, “On the importance of difficulty calibration in membership inference attacks,” in *Int. Conf. Learn. Represent. (ICLR)*, 2022.
- [46] K. Alex, H. Geoffrey *et al.*, “Learning multiple layers of features from tiny images,” in *Toronto, ON, Canada*, 2009.
- [47] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2016, pp. 770–778.
- [48] D. Hintersdorf, L. Struppek, and K. Kersting, “To trust or not to trust prediction scores for membership inference attacks,” in *IJCAI*, L. D. Raedt, Ed., 2022, pp. 3043–3049.
- [49] G. D. Grosso, H. Jalalzai, G. Pichler, C. Palamidessi, and P. Piantanida, “Leveraging adversarial examples to quantify membership information leakage,” *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022.
- [50] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, A. Y. Ng *et al.*, “Reading digits in natural images with unsupervised feature learning,” in *NeurIPS workshop on deep learning and unsupervised feature learning*, vol. 2011, no. 5, 2011, p. 7.
- [51] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *Int. Conf. Mach. Learn. (ICML)*, vol. 37, 2015, pp. 448–456.

Method	Training Acc	Test Acc	TPR @ 0.1% FPR	TPR @ 0.001% FPR	Log-scale AUC	MIA Balanced Acc
Base	99.2 \pm 0.8	92.6 \pm 0.2	5.27 \pm 1.57	2.06 \pm 0.85	0.778 \pm 0.040	57.71 \pm 1.19
Smooth	99.5 \pm 0.3	93.7 \pm 0.2	7.89 \pm 1.66	2.89 \pm 1.22	0.809 \pm 0.022	60.34 \pm 1.27
Disturblabel	98.7 \pm 0.5	93.2 \pm 0.2	4.72 \pm 1.02	1.64 \pm 0.83	0.771 \pm 0.028	57.26 \pm 0.65
Noise	99.2 \pm 0.9	92.6 \pm 0.3	5.29 \pm 1.51	2.17 \pm 0.93	0.780 \pm 0.037	57.68 \pm 1.23
Cutout	98.7 \pm 0.8	93.1 \pm 0.3	4.38 \pm 1.27	1.66 \pm 0.83	0.767 \pm 0.037	57.02 \pm 1.10
Mixup	97.4 \pm 0.5	93.4 \pm 0.3	3.37 \pm 0.71	0.89 \pm 0.42	0.749 \pm 0.020	56.43 \pm 0.66
Jitter	99.2 \pm 0.6	92.6 \pm 0.2	5.25 \pm 1.36	2.17 \pm 0.94	0.781 \pm 0.029	57.59 \pm 1.00
Distillation	95.1 \pm 5.4	90.8 \pm 3.4	3.10 \pm 2.27	1.29 \pm 1.11	0.691 \pm 0.120	54.71 \pm 2.69
PGD-AT	97.9 \pm 0.5	86.2 \pm 0.3	10.55 \pm 0.25	4.31 \pm 0.13	0.852 \pm 0.011	65.23 \pm 0.49
TRADES	94.3 \pm 0.7	82.5 \pm 0.8	10.31 \pm 0.22	4.22 \pm 0.16	0.837 \pm 0.014	65.09 \pm 0.74
AWP	96.5 \pm 0.7	85.2 \pm 0.6	8.85 \pm 0.24	3.08 \pm 0.05	0.820 \pm 0.020	62.85 \pm 0.59
TRADES-AWP	91.4 \pm 1.4	80.2 \pm 1.2	8.54 \pm 0.22	3.05 \pm 0.04	0.794 \pm 0.042	62.76 \pm 0.63

TABLE B.1: Attack success rates of different data enhancement on SVHN. The same conventions are used as in Table II.

Method	Training Acc	Test Acc	TPR @ 0.1% FPR	TPR @ 0.001% FPR	Log-scale AUC	MIA Balanced Acc
None	100.0 \pm 0.0	82.9 \pm 0.5	20.35 \pm 4.31	9.44 \pm 3.24	0.885 \pm 0.013	76.25 \pm 2.18
None + Smooth	100.0 \pm 0.0	83.7 \pm 0.5	14.48 \pm 3.03	2.37 \pm 1.79	0.839 \pm 0.024	72.91 \pm 1.60
None + Disturblabel	100.0 \pm 0.0	84.1 \pm 0.6	16.26 \pm 1.03	3.43 \pm 2.48	0.853 \pm 0.016	72.34 \pm 0.89
None + Noise	100.0 \pm 0.0	82.4 \pm 0.7	20.84 \pm 3.69	8.59 \pm 2.88	0.886 \pm 0.011	76.90 \pm 2.52
None + Cutout	100.0 \pm 0.0	84.0 \pm 0.6	23.07 \pm 0.80	10.53 \pm 1.85	0.894 \pm 0.004	77.42 \pm 0.37
None + Mixup	100.0 \pm 0.0	83.7 \pm 0.4	16.53 \pm 1.42	4.84 \pm 1.58	0.867 \pm 0.006	76.53 \pm 0.99
None + Jitter	100.0 \pm 0.0	82.0 \pm 1.3	21.73 \pm 5.85	8.77 \pm 3.64	0.886 \pm 0.020	77.63 \pm 2.61
None + Distillation	100.0 \pm 0.0	84.4 \pm 0.4	12.79 \pm 1.81	5.16 \pm 1.18	0.851 \pm 0.010	69.25 \pm 0.99

TABLE B.2: Attack success rates of different data augmentation methods **without** using Base on CIFAR-10. The 2nd and 3rd columns show the training and test accuracies of each method, respectively. The 4th - 7th columns show four metrics to evaluate the extent of privacy leakages. We highlight the MIA success rates for different data augmentation methods that are larger than that for None.

Method	Training Acc	Test Acc	TPR @ 0.1% FPR	TPR @ 0.001% FPR	Log-scale AUC	MIA Balanced Acc
None	100.0 \pm 0.0	54.0 \pm 0.9	54.26 \pm 10.72	30.68 \pm 11.03	0.954 \pm 0.014	92.96 \pm 2.28
None + Smooth	100.0 \pm 0.0	53.7 \pm 2.0	69.71 \pm 3.51	41.81 \pm 9.55	0.972 \pm 0.004	96.76 \pm 0.33
None + Disturblabel	100.0 \pm 0.0	55.5 \pm 0.5	56.30 \pm 1.22	37.37 \pm 4.58	0.959 \pm 0.003	90.96 \pm 0.23
None + Noise	100.0 \pm 0.0	53.5 \pm 1.2	50.75 \pm 8.61	28.61 \pm 10.75	0.950 \pm 0.011	91.79 \pm 1.70
None + Cutout	100.0 \pm 0.0	54.1 \pm 0.8	58.14 \pm 5.60	36.64 \pm 7.94	0.961 \pm 0.007	92.87 \pm 1.04
None + Mixup	100.0 \pm 0.0	49.3 \pm 0.7	75.88 \pm 1.09	51.04 \pm 6.80	0.978 \pm 0.002	96.13 \pm 0.06
None + Jitter	100.0 \pm 0.0	53.1 \pm 0.9	57.05 \pm 11.22	30.21 \pm 14.71	0.955 \pm 0.016	93.47 \pm 2.18
None + Distillation	100.0 \pm 0.0	57.7 \pm 1.7	56.45 \pm 3.95	35.61 \pm 7.31	0.959 \pm 0.006	90.39 \pm 0.88

TABLE B.3: Attack success rates of different data augmentation methods **without** using Base on CIFAR-100. The same conventions are used as in Table B.2.

APPENDIX A

THE HYPER-PARAMETERS OF EACH DATA AUGMENTATION

As stated in the paper, the hyper-parameter of each data augmentation method was set to achieve relatively high test accuracy by trying various values. Here we report the values we tried and the final values used when training 128 shadow models for each data augmentation method on CIFAR-10, CIFAR100, and SVHN:

- 1) Random Cropping and Flipping: First, the images with a resolution of 32×32 were padded with zeros of 4 pixels on each end. Then the padded images with the resolution of 36×36 were randomly cropped out to form inputs with the resolution of 32×32 . Finally, the inputs were randomly flipped horizontally. Unless otherwise specified, all other data augmentation methods also use this as default.

- 2) Label smoothing: We tried ϵ including 0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, and 0.8. Finally, we chose 0.2 on CIFAR-10, 0.3 on CIFAR-100, and 0.2 on SVHN.
- 3) Disturblabel: We tried ϵ including 0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.425, 0.45, 0.5, 0.525, 0.55, 0.575, and 0.6. Finally, we chose 0.05 on CIFAR-10, 0.3 on CIFAR-100, and 0.05 on SVHN.
- 4) Gaussian Noise: We tried σ including 0.025, 0.01, 0.05, 0.075, 0.1, 0.125, 0.15, 0.175, 0.2, 0.225, 0.25, 0.275, 0.3, 0.325, and 0.35. Finally, we chose σ to be 0.01 on the three datasets.
- 5) Cutout: We tried M including 4, 8, 12, 16, and 20. Finally, we chose M to be 8 on the three datasets.
- 6) Mixup: γ used in Mixup is sampled from a *beta* distribution $\gamma \sim \text{Be}(\alpha, \alpha)$. We tried α including 0.5, 0.1, 0.25, 1, 2, 4, 8, 16, 32, 64, 128, and 256. Finally,




Fig. C.1: Attack success rate versus the train-test gap of different data augmentation models on CIFAR-10 with CNN-8 using MaxPreCA (top) and LiRA (bottom), respectively.

we chose α to be 0.5 on the three datasets.

- 7) Jitter: We used the ColorJitter function in Torchvision³ directly. We tried the parameters corresponding to brightness, contrast, saturation, and hue including 0.05, 0.1, 0.2, 0.15, 0.25, 0.3, 0.35, 0.4, 0.45, and 0.5. Finally, we chose 0.05 on the three datasets.
- 8) Distillation: We tried T including 1, 2, 3, 5, and 10. Finally, we chose T to be 3 on the three datasets.

APPENDIX B

ADDITIONAL MEMBERSHIP INFERENCE ATTACK RESULTS

The training and test accuracies and MIA results of all data augmentation models on SVHN are shown in Table B.1. In addition, taking CIFAR-10 and CIFAR-100 as examples, the training and test accuracies and MIA results of all data augmentation models trained without using Base are demonstrated in Tables B.2 and B.3. Single query was used because it obtained higher attack success rates than multiple queries, as shown in Figure 4 in the paper. Here *None* stands for models trained without any data augmentation (only the original image data). The test accuracies of models trained without using Base are much lower than that of models trained using Base.

APPENDIX C

MEMBERSHIP INFERENCE ATTACK RESULTS ON CIFAR-10 WITH CNN-8

The CNN-8 architecture used in this experiment is the same as the one employed for the SVHN dataset, as described in Section V-A (Models paragraph). Furthermore, all hyperparameter settings align with those described in Section V-A. The evaluation results of different data augmentation methods are presented in Fig. C.1, which provides a concise representation of the results similar to those shown in Table II.

³<https://pytorch.org/vision/stable/index.html>

Dataset	Method	Adversarial Acc	TPR @ 0.1% FPR
CIFAR-10	Base	0.0 ± 0.0	6.57 ± 0.32
	PGD-AT	36.0 ± 0.3	17.08 ± 0.48
	TRADES	42.1 ± 0.4	14.25 ± 2.23
	AWP	42.4 ± 0.3	9.07 ± 1.85
	TRADES-AWP	43.0 ± 0.3	11.43 ± 0.76

TABLE C.1: The accuracies on adversarial examples (Adversarial Acc) and privacy leakage of different adversarial training models on CIFAR-10 with CNN-8. The accuracies are evaluated using PGD with $\epsilon = 8$ and 20 iteration steps.




Fig. C.2: Attack success rates of PGD-AT and TRADES under different ϵ on CIFAR-10 with CNN-8.

From this figure, we observe a less pronounced correlation between generalization gaps and privacy leakages compared to the previous MaxPreCA method. Regarding the adversarial training methods, the results are presented in Table C.1 and Fig. C.2. Overall, these results closely resemble the trends observed with the ResNet-18 architecture on CIFAR-10 in Table IV and Fig. 6. This consistency in trends supports that our findings remain consistent across diverse model architectures and datasets.

APPENDIX D

MEMBERSHIP INFERENCE ATTACK RESULTS ON NON-IMAGE DATASETS

We performed similar experiments with an MLP architecture on two additional non-image datasets, namely Purchases and Locations processed by Shokri et al. [6]⁴. The Purchases dataset is based on the “acquire valued shoppers” challenge dataset on Kaggle, containing shopping histories for several thousand individuals. Shokri et al. [6] derived a simplified purchase dataset with 197,324 data samples, where each sample consists of 600 binary features. The first 60,000 data samples from Purchases are used to perform 100-category classification in our experiments. The Locations dataset is processed from the publicly available set of mobile users’ location “check-ins”, which contained 5,010 data samples with 446 binary features. All 5,010 data samples from Purchases are used to perform classification on 30-category classification in our experiments.

The MLP architecture we used in the experiments has five fully connected layers, with sizes 512, 256, 128, 128, and 128 in turn. Batch normalization was applied to each layer. As

⁴<https://github.com/privacytrustlab/datasets>

Method	Training Acc	Test Acc	TPR @ 0.1% FPR	TPR @ 0.001% FPR	Log-scale AUC	MIA Balanced Acc
None	100.0 \pm 0.0	75.2 \pm 0.2	19.34 \pm 1.09	4.97 \pm 1.49	0.881 \pm 0.005	85.60 \pm 0.20
Disturblabel	100.0 \pm 0.0	75.7 \pm 0.3	20.69 \pm 0.86	4.55 \pm 3.10	0.873 \pm 0.030	84.60 \pm 0.14
Noise	100.0 \pm 0.0	75.3 \pm 0.3	19.14 \pm 1.07	3.85 \pm 1.71	0.876 \pm 0.007	85.58 \pm 0.19
Distillation	100.0 \pm 0.0	75.5 \pm 0.3	21.54 \pm 1.54	5.08 \pm 2.27	0.885 \pm 0.009	84.67 \pm 0.36
01 Flipping	100.0 \pm 0.0	75.1 \pm 0.4	22.16 \pm 1.39	4.90 \pm 1.90	0.886 \pm 0.010	86.64 \pm 0.12

TABLE D.1: Attack success rates of different data augmentation methods on Purchases.

Method	Training Acc	Test Acc	TPR @ 0.1% FPR	TPR @ 0.001% FPR	Log-scale AUC	MIA Balanced Acc
None	100.0 \pm 0.0	61.6 \pm 0.8	52.43 \pm 9.57	43.56 \pm 14.33	0.960 \pm 0.014	94.89 \pm 0.34
Disturblabel	100.0 \pm 0.0	63.4 \pm 0.8	47.78 \pm 5.75	36.42 \pm 9.27	0.953 \pm 0.009	91.95 \pm 0.42
Noise	97.1 \pm 0.4	64.1 \pm 0.7	43.04 \pm 6.14	33.36 \pm 7.15	0.948 \pm 0.008	92.08 \pm 0.37
Distillation	100.0 \pm 0.0	61.4 \pm 0.9	53.89 \pm 8.19	40.97 \pm 9.01	0.961 \pm 0.008	94.85 \pm 0.53
01 Flipping	100.0 \pm 0.0	65.7 \pm 0.8	46.78 \pm 5.56	30.42 \pm 12.39	0.945 \pm 0.014	92.12 \pm 0.54

TABLE D.2: Attack success rates of different data augmentation methods on Locations.



Fig. E.1: Attack success rate versus the train-test gap of different data augmentation models on CIFAR-10 using loss-based MIA [10]. r stands for the Pearson correlation coefficient.

Fig. E.2: Attack success rate versus the train-test gap of different data augmentation methods on CIFAR-10 using difficulty calibrated loss approach [45].

this architecture is relatively simple, we only trained it for 50 epochs. Multi-step decay which scales the learning rate by 0.1 was used on the 25th and 37th epochs. Other experimental settings default to those described in Section V-A.

As mentioned in Section VIII, we investigated the privacy impacts of applying Disturblabel, Distillation, and Gaussian Noise. Additionally, as the features of both Purchases and Locations are binary, we investigated a new 0-1 Flipping augmentation method specially designed for these non-image datasets. Concretely, the 0-1 Flipping augmentation involves flipping a certain proportion of features, with 1% of features flipped for Purchases and 10% for Locations. The overall evaluation results on these datasets are shown in Tables D.1

Dataset	Method	Adversarial Acc	TPR @ 0.1% FPR
CIFAR-10	Base	0.0 \pm 0.0	7.15 \pm 0.33
	PGD-AT	38.8 \pm 0.4	15.42 \pm 1.20
	TRADES	45.2 \pm 0.3	7.54 \pm 1.371
	AWP	45.9 \pm 0.1	7.32 \pm 0.86
	TRADES-AWP	48.8 \pm 0.2	7.54 \pm 0.73

TABLE E.1: The accuracies on adversarial examples (Adversarial Acc) and privacy leakage of different adversarial training models on CIFAR-10 using difficulty calibrated loss approach [45]. The accuracies are evaluated using PGD with $\epsilon = 8$ and 20 iteration steps.

and D.2.

APPENDIX E

MEMBERSHIP INFERENCE ATTACK RESULTS BEYOND LiRA

To demonstrate that the results using other MIAs discussed in Section IV-B are similar to those using the MaxPreCA, Fig. E.1 showcases the attack success rate versus the train-test gap of different data augmentation models on CIFAR-10 using loss-based MIA [10].

To demonstrate that our findings are not limited to LiRA alone, Fig. E.2 and Table E.1 showcase the results using the difficulty calibrated loss approach [45]. As discussed in Section VIII-B, these results indicate that the general trends in privacy effects resulting from data augmentation and adversarial training align with those observed using LiRA (see Fig. 6 and Table IV).