

Membership Inference Attacks against Diffusion Models

Tomoya Matsumoto
Osaka University
Osaka, Japan
t-matsumoto@ist.osaka-u.ac.jp

Takayuki Miura
Osaka University,
NTT Social Informatics Laboratories
Tokyo, Japan
takayuki.miura.br@hco.ntt.co.jp

Naoto Yanai
Osaka University
Osaka, Japan
yanai@ist.osaka-u.ac.jp

Abstract—Diffusion models have attracted attention in recent years as innovative generative models. In this paper, we investigate whether a diffusion model is resistant to a membership inference attack, which evaluates the privacy leakage of a machine learning model. We primarily discuss the diffusion model from the standpoints of comparison with a generative adversarial network (GAN) as conventional models and hyperparameters unique to the diffusion model, i.e., timesteps, sampling steps, and sampling variances. We conduct extensive experiments with DDIM as a diffusion model and DCGAN as a GAN on the CelebA and CIFAR-10 datasets in both white-box and black-box settings and then show that the diffusion model is comparably resistant to a membership inference attack as GAN. Next, we demonstrate that the impact of timesteps is significant and intermediate steps in a noise schedule are the most vulnerable to the attack. We also found two key insights through further analysis. First, we identify that DDIM is vulnerable to the attack for small sample sizes instead of achieving a lower FID. Second, sampling steps in hyperparameters are important for resistance to the attack, whereas the impact of sampling variances is quite limited.

Index Terms—diffusion model, membership inference attack, GAN, hyperparameter, privacy

I. INTRODUCTION

A. Motivation

In machine learning, generative models have been actively studied. As a breakthrough in recent years, diffusion models [1], [2] were discovered as new generative models. Diffusion models can generate more plausible images and texts than existing generative models such as a generative adversarial network (GAN). Remarkably, diffusion models have outperformed GANs on academic benchmarks in the past years [3] and are expected to replace GANs in a variety of applications [4]–[7].

We emphasize that discussion on privacy violations for training data of diffusion models [8]–[10] is a very recent issue and hence is still insufficient. In general, training generative models requires large amounts of training data, and, for example, models that generate disease histories as sensitive information also attract attention [11], [12]. Understanding privacy violations under realistic threats for generative models is crucial for using generative models in the real world [13]. Consequently, privacy violations of diffusion models should be discussed in more detail.

In this paper, we discuss a membership inference attack [14] against diffusion models, which is an important criterion for evaluating privacy violations. Informally, a membership inference attack aims to identify whether a data record was used to train a machine learning model. This attack is discussed for evaluating privacy violations in the past years [15]–[17], and hence it is important to discuss the attack for diffusion models in order to understand potential threats of the model. We then try to answer the following question: *How vulnerable are diffusion models to a membership inference attack?*

We note that this question is *non-trivial* because the data generation logic by diffusion models differs from conventional models such as GANs. According to earlier literature [13], [18]–[20], the architecture of a generative model affects the attack success rate of a membership inference attack. For example, for an attack against GAN, an attacker can utilize a discriminator model that is a counterpart to the generative model [19]. By contrast, diffusion models do not contain such a model. Meanwhile, diffusion models have unique hyperparameters, i.e., timesteps, sampling steps, and sampling variances, that have never been contained in conventional models. It indicates that we no longer know how the membership inference attack on diffusion models varies compared to the traditional models. These standpoints were not addressed in the existing works [8]–[10]. (See Section II-C for more detail.)

B. Contributions

In this paper, we shed light on the impact of a membership inference attack on diffusion models through extensive experiments. Our primary discussions are on comparison with a GAN as a conventional generative model and hyperparameters of the diffusion model, i.e., timesteps, sampling steps, and sampling variances. In particular, through extensive experiments with denoising diffusion implicit models (DDIM) as a diffusion model and deep convolutional GAN (DCGAN) as a GAN on the CelebA and CIFAR-10 datasets in both white-box and black-box settings, we show that the diffusion model is comparably resistant to a membership inference attack as GAN. We then identify the impact of timesteps on the attack. Intermediate steps in a noise schedule are the most vulnerable to the attack rather than the beginning or the final step. (See Section IV for detail.)

We also provide two key insights into a membership inference attack on the diffusion model through further analysis for overfitting compared to GAN and hyperparameters for sampling, i.e., sampling steps and sampling variances. First, for overfitting, we show that DDIM is vulnerable to the attack if the number of training samples is small since DDIM is well-trained compared to DCGAN. Second, sampling steps in hyperparameters are quite important for resistance to the attack in contrast to sampling variances, which are irrelevant to the attack. (See Section V for detail.) Our source code is publicly available (<https://github.com/fseclab-osaka/mia-diffusion>).

II. RELATED WORK

This section describes related works of membership inference attacks, including GANs and diffusion models.

A. Membership Inference

A membership inference attack is a kind of attack whereby an adversary infers whether a particular example was contained in the training dataset of a model [21]–[23]. A model vulnerable to the attack potentially contains threats to privacy leakage, and hence recent works discuss membership inference attacks for various machine learning models [15], [16], [18]. There are two settings [14], i.e., the white-box setting where an adversary has access to model parameters, and the black-box setting where he/she utilizes only outputs of the model.

A typical approach for membership inference attacks is to leverage the large divergence between the loss distribution over members and non-members [24]. The divergence can be embedded by an adversary. For instance, privacy leakage, including membership inference attacks, can be more effective by training a model with poisoning samples [25]–[28].

Membership inference attacks can be prevented by differential privacy [29] where gradients are perturbed [30], [31]. Since differential privacy often deteriorates inference accuracy, several works evaluated differential privacy on membership inference attacks in a quantitative fashion [32], [33].

B. Attacks on GANs

Related works are membership inference attacks against GANs [13], [18], [19], [34]. LOGAN was presented as the first attack against GANs in both white-box and black-box settings. LOGAN in the white-box setting utilizes outputs from the discriminator of a target model. Concurrently, the Monte Carlo attack [13] was presented as an attack independent of a target model’s architecture in the black-box setting. The above attacks are also utilized as oracles in the subsequent work [19].

GAN-Leaks [34] is the state-of-the-art attack against GANs in the black-box setting, to the best of our knowledge. It no longer requires the setting of parameters compared to the above attacks and hence can be used in diffusion models. GAN-Leaks also showed a new attack whereby an adversary has access to the internal of a target model’s generator.

We utilize LOGAN in the white-box setting and GAN-Leaks in the black-box setting to compare a diffusion model with a GAN, respectively. As described above, LOGAN is the only

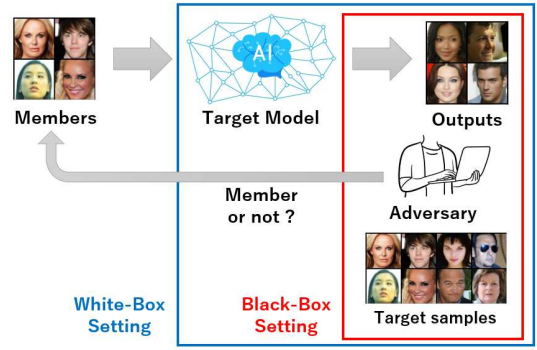


Fig. 1: Overview of our membership inference attacks.

work on GANs in the white-box setting while GAN-Leaks can be used in diffusion models. We note that there is no work that is applicable to diffusion models in the white-box setting.

C. Attacks on Diffusion Models

We note that there is an attack by Wu et al. [35] against diffusion models. Their attack focuses on text-to-image generation models. In contrast, we focus on typical image generation models. Namely, the underlying problem is quite different. It is considered that the existing attacks against GANs are rather close to our work than Wu et al. [35].

Concurrently, there are a few works [8]–[10] on typical image generation models based on diffusion models, which are the closest to ours. As the main difference from these works, we discuss hyperparameters of diffusion models in detail, which are our main results. In particular, Hu et al. [8] and Duan et al. [10] did not discuss comparison with GANs. Carlini et al. [9] compared with GANs, but did not discuss hyperparameters, i.e., sampling steps and sampling variances. We primarily focus on comparison with GANs and discussion on the hyperparameters. Our work is thus concurrent with the above works. Interested readers are also encouraged to read the existing works [8]–[10].

III. MEMBERSHIP INFERENCE ATTACKS AGAINST DIFFUSION MODEL

This section presents our membership inference attacks against a diffusion model. To this end, we first formalize the adversary setting of the attacks and the detail of a diffusion model. We then describe our investigation method and the key question of this paper in detail.

A. Formalization

A membership inference attack in this paper is defined as a game between an adversary \mathcal{A} and a challenger \mathcal{C} . We note that the adversary’s knowledge about the target model is different for white-box and black-box settings. White-box setting means that an adversary has access to an architecture of the model and its parameters while black-box setting means that an adversary cannot obtain these information.

We first denote by \mathbb{R} a set of real numbers, by \mathcal{D} a set of data samples, and by \mathcal{M} a set of generative models. Then, a

generative model $M \in \mathcal{M}$ is defined as a function $M : \mathbb{R}^r \rightarrow \mathcal{D}$, where r is an arbitrary integer. The game is as follows:

- 1) \mathcal{C} chooses a dataset $D \subseteq \mathcal{D}$.
- 2) \mathcal{C} chooses a bit $b \leftarrow \{0, 1\}$. If $b = 1$, \mathcal{C} chooses a sample $x \in D$; otherwise, \mathcal{C} chooses $x \in \mathcal{D} \setminus D$ other than D . \mathcal{C} sends x to \mathcal{A} .
- 3) \mathcal{C} trains M with the dataset D .
- 4) \mathcal{A} obtains $x_i = M(z_i)$ for any $z_i \in \mathbb{R}^r$. **Furthermore, \mathcal{A} also receives M and its auxiliary information¹.**
- 5) \mathcal{A} returns a bit $b' \in \{0, 1\}$. If $b = b'$, \mathcal{A} wins the game.

The difference between the white-box and black-box settings is represented as the sentences with red color, and only the white-box setting contains this sentence. As described in Section II-A, an adversary in the white-box setting has access to the internal of a target model M : in this paper, he/she can have access to internal networks, which are utilized in the training of M and unnecessary for data generation, such as discriminators in GANs as auxiliary information. On the other hand, an adversary in the black-box setting receives only the generated samples x_i from a target model M .

We adopt area-under-the-ROC-curve (AUCROC) [36], attack success rate (ASR) [14], and true positive rates (TPR) at low false positive rates (FPR) [23], used as TPR at 1%FPR in this paper, as evaluation metrics of the attacks. We also adopt the Frechet inception distance (FID) [37] to evaluate the quality of generated images. A lower score for FID means that generated images are of a higher quality.

B. Detail of Diffusion Model

Membership inference attacks in this paper depend on loss values of a diffusion model. We describe the detail of the loss function, including the diffusion process and denoising one.

We describe the denoising diffusion probabilistic models (DDPM) [2] below for simplicity because DDIM has nearly the same process. The diffusion process is described as:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I). \quad (1)$$

According to the noise schedule β , the observed data x_0 are transformed into noise x_T by adding Gaussian noises over T steps. The denoising process is described as:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)). \quad (2)$$

The data are reconstructed by repeatedly removing noises from noise x_T drawn from a multivariate standard Gaussian distribution.

The model is trained to maximize the variational lower bound of the log marginal likelihood $\log p_\theta(x_0)$. Therefore, the loss function L is expressed as:

$$L := \mathbb{E}_q \left[\log \frac{q(x_{1:T}|x_0)}{p_\theta(x_{0:T})} \right] \geq \mathbb{E} \left[-\log p_\theta(x_0) \right]. \quad (3)$$

In practice, we use the simplified loss function shown as:

$$L_{\text{simple}} := \mathbb{E}_{t, x_0, \epsilon} \left[\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, t)\|^2 \right], \quad (4)$$

¹In the case of GANs, a discriminator is the auxiliary information.

Algorithm 1 White-box attack against diffusion models

Input: Target samples x^1, \dots, x^m , model's network ϵ_θ , time t , noise schedule $\bar{\alpha}_t := \prod_{s=1}^t (1 - \beta_s)$, threshold c

for $i = 1$ to m **do**

$y^i \leftarrow 0$ ▷ Initialization

if $\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x^i + \sqrt{1 - \alpha_t}\epsilon, t) < c$ **then**

$y^i \leftarrow 1$ ▷ Inferring as a member

Output: Labels y^1, \dots, y^m

which represents the noise estimation error of the data x_t , using the notation $\bar{\alpha}_t := \prod_{s=1}^t (1 - \beta_s)$.

In training, θ is optimized for the above loss function. It is also utilized for the membership inference attack because small loss values mean that the target sample is in-member.

Meanwhile, the sampling of DDIM is represented as:

$$x_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \left(\frac{x_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(x_t, t)}{\sqrt{\bar{\alpha}_t}} \right) + \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \cdot \epsilon_\theta(x_t, t) + \sigma_t \epsilon_t, \quad (5)$$

where σ controls the randomness of the sampling. We can generate images for any subset of timesteps $\tau \subset \{1, \dots, T\}$ and σ is represented with a hyperparameter η as:

$$\sigma_{\tau_i}(\eta) := \eta \sqrt{(1 - \bar{\alpha}_{\tau_{i-1}})/(1 - \bar{\alpha}_{\tau_i})} \sqrt{1 - \bar{\alpha}_{\tau_i}/\bar{\alpha}_{\tau_{i-1}}}. \quad (6)$$

Note that the following value is utilized for CIFAR-10 in [38]:

$$\hat{\sigma}_{\tau_i} := \sqrt{1 - \bar{\alpha}_{\tau_i}/\bar{\alpha}_{\tau_{i-1}}}. \quad (7)$$

C. Investigation Method

We describe our investigation method to evaluate membership inference attacks against a diffusion model in the white-box setting. As described in Section II-B, there is no attack against a diffusion model in the white-box setting, whereas we utilize LOGAN [18] and GAN-Leaks [34].

The loss function of a diffusion model represents a noise estimation error as in Equation (4). In general, the loss values among members are smaller than among non-members because a model learns training data to minimize them [24]. Based on this fact, an adversary can infer whether the target samples are contained in the training dataset by computing their loss values with appropriate time t (as described later). The detail of the method is shown in Algorithm 1.

D. Key Question

As described in Section I-A, we shed light on how a diffusion model is vulnerable to membership inference attacks. To this end, we discuss it from two standpoints: difference from generative adversarial models (GANs), and timesteps of a diffusion model.

We first evaluate membership inference attacks, including the investigation method in the previous section, against the diffusion model and existing GANs. The main difference between the diffusion model and existing GANs is the number of neural networks, i.e., the diffusion model consists of a

single neural network while the existing GANs consist of generators and discriminators. We discuss the impact of the above difference on membership inference attacks in both black-box and white-box settings.

We also discuss the impact of timesteps on membership inference attacks. In a diffusion model, information learned by the model is different for each timestep t . Although the performance of membership inference attacks depends on the divergence of loss distributions between members and non-members [24], the timestep t when the divergence is largest is unknown. Consequently, we evaluate membership inference attacks with respect to timestep t .

IV. EXPERIMENTS

This section conducts extensive experiments on membership inference attacks against a diffusion model. We evaluate the attack AUC by comparing the attacks against GANs. We also evaluate the impact of the timesteps of a diffusion model on the attack AUC.

A. Experimental Setting

The experimental setting in this paper is described below.

1) *Datasets*: We utilize CIFAR-10 [39] and CelebA [40] datasets, which contain samples with 32×32 color pixels and with center-cropped 64×64 color pixels, respectively.

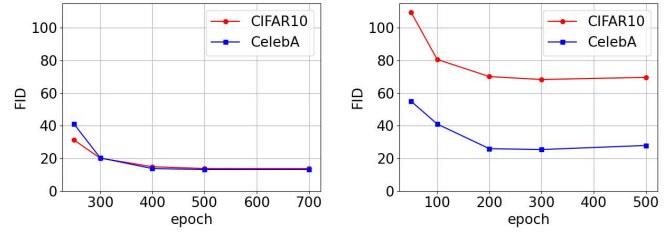
2) *Architectures*: We utilize the denoising diffusion implicit model (DDIM) [38] as a diffusion model and the deep convolutional GAN (DCGAN) [41] as a baseline. The reason of the use of these architectures is that DDIM is the basis of other diffusion models and is suitable for investigating common properties of the diffusion models. On the other hand, DCGAN has been utilized in existing works [13], [18], [34] for membership inference attacks on generative models.

For hyperparameters, we set $T = 1000$, and the noise schedule is the cosine schedule [42] or the linear schedule [2]. The default setting is the cosine schedule. Also, the sampling is performed for sampling variance $\sigma(0)$ and per 50 steps, i.e., 20 steps in total. The dimension of latent variables in DCGAN is 100. A target model M is trained with 12,000 samples, and the numbers of epochs are 500 for DDIM and 300 for DCGAN. To fairly compare the evaluation metrics for membership inference attacks, the effect of overfitting should be minimized. Thus, we utilize 500 epochs for DDIM and 300 epochs for DCGAN as the target model M .

3) *Attack Methods*: In the white-box setting, we utilize the investigation method described in Section III-C for the attack on DDIM, where time $t = 350$, and LOGAN [18] for the attack on DCGAN. On the other hand, in the black-box setting, we utilize the full black-box attack of GAN-Leaks [34] where the number of generated images is 12,000.

B. Results

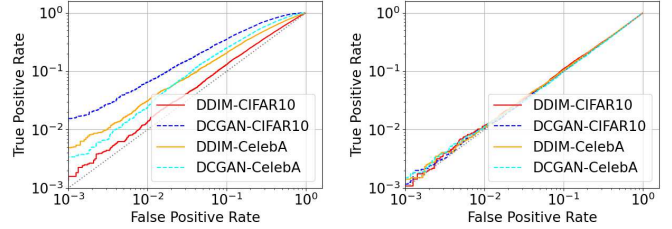
1) *Comparison with GANs*: The results of DDIM and DCGAN are shown in Fig. 3 and Tables I. According to these results, DDIM is more resistant to the membership inference attack than DCGAN on CIFAR-10 in the white-box setting.



(a) DDIM

(b) DCGAN

Fig. 2: FID scores throughout training.



(a) White-box setting.

(b) Black-box setting.

Fig. 3: ROC curves.

TABLE I: Attack performance on white-box/black-box setting.

| Dataset | Model | AUCROC | ASR | TPR at 1%FPR |
|----------|-------|---------------|---------------|--------------|
| CIFAR-10 | DDIM | 0.552 / 0.503 | 0.536 / 0.499 | 1.54 / 1.19% |
| | DCGAN | 0.778 / 0.497 | 0.707 / 0.497 | 6.50 / 1.05% |
| CelebA | DDIM | 0.635 / 0.503 | 0.599 / 0.502 | 2.93 / 1.09% |
| | DCGAN | 0.699 / 0.495 | 0.643 / 0.494 | 2.56 / 1.13% |

Furthermore, according to Table I, DDIM is more resistant than DCGAN in terms of AUCROC. Although one might think that DDIM is vulnerable according to TPR at 1%FPR, including several low scores of FPR in Fig. 3, on CelebA, it is considered that they are fairly comparable because each value is quite small. By contrast, the attack fails for all models and datasets in the black-box setting. As a result, it is considered that DDIM is at least resistant equivalently to DCGAN.

2) *Attack performance for different timesteps*: The investigation method in Section III-C requires a timestep t as an input to a target model M in addition to samples x^1, \dots, x^m . We show that timesteps and the noise schedule strongly affect the performance of the membership inference attacks.

The results of AUCROC for each timestep t are represented in Fig. 4, where DDIM is trained with the cosine and linear schedules. According to the figure, AUCROC is maximized at $t = 350$ for the cosine schedule and at $t = 200$ for the linear schedule. We call these values *best-one* AUCROC for the sake of convenience. For these values, the value of $\bar{\alpha}_t$ is 0.7. As might have been unexpected, these maximized values were obtained in intermediate steps. The best-one AUCROC for the cosine schedule is higher than the linear schedule. Besides, according to Table II, AUCROC with an average of loss values for all the steps, called *average* AUCROC for the

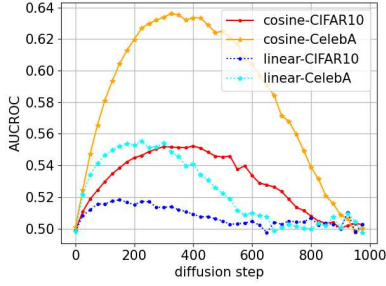


Fig. 4: AUCROC at every 25 step on white-box setting.

TABLE II: AUCROC by best-one and average t .

| | CIFAR-10 | | CelebA | |
|--------|----------|---------|----------|---------|
| | best-one | average | best-one | average |
| cosine | 0.552 | 0.535 | 0.636 | 0.597 |
| linear | 0.518 | 0.510 | 0.555 | 0.529 |

sake of convenience, is lower than the best-one AUCROC. Therefore, we use results at $t = 350$ for the investigation methods in the entire experiments.

V. DISCUSSION

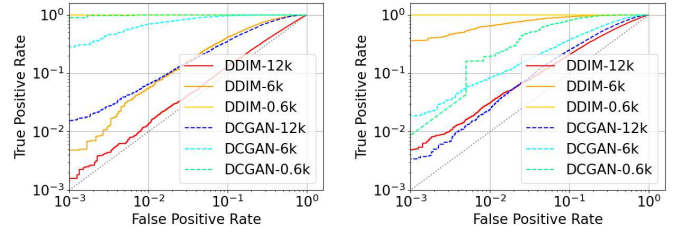
This section discusses the attack performance in terms of number of training samples and training epochs to understand the impact of overfitting on diffusion models. We then discuss the impacts of hyperparameters for sampling in the diffusion models.

A. Impact of Overfitting

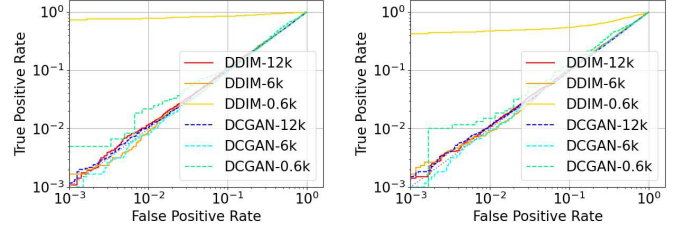
Overfitting strengthens membership inference attacks in general [31]. We discuss the impact of overfitting in terms of number of training samples and epochs.

1) *Number of training samples:* We perform the attacks on three target models, i.e., their number of training samples is 12,000, 6000, and 600 images, to evaluate the impact of number of training samples on a membership inference attack. The results are shown in Fig 5 and Table III, where the number of epochs is changed depending on the number of training samples in order to align the number of images trained by the model. Overall, DDIM is vulnerable to membership inference attacks compared to DCGAN when the number of training samples is small. We analyze this result in detail below. We also note that FID for DCGAN on CIFAR-10 is high because DCGAN is unsuitable for the dataset. Indeed, we found many strange images, although we omit the detail due to the space limitation. Consequently, we focus on results on CelebA.

First, in the white-box setting, Fig. 5 shows a remarkable phenomenon: DDIM is more vulnerable to the membership inference attack than DCGAN for small number of training samples but is more resistant for large number of training samples. For instance, DDIM is vulnerable in the white-box setting until 6000 images. After that, according to the columns of CelebA in Table III, AUCROC and ASR of DDIM for 12,000 images are 0.064 smaller and 0.044 smaller than DCGAN, respectively.

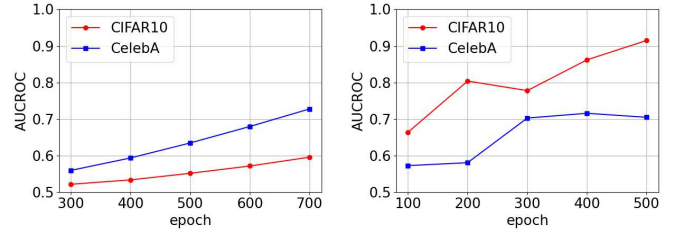


(a) White-box setting (CIFAR-10). (b) White-box setting (CelebA).



(c) Black-box setting (CIFAR-10). (d) Black-box setting (CelebA).

Fig. 5: ROC curves for different number of training samples.



(a) DDIM. (b) DCGAN.

Fig. 6: AUCROC throughout training on white-box setting.

We also found remarkable results in the black-box setting: for small number of training samples, DDIM is significantly vulnerable compared to DCGAN. Notably, AUCROC and ASR of DDIM for 600 images are 0.24 higher and 0.19 higher than DCGAN, respectively. Meanwhile, after the number of training samples is 6,000, the result of DDIM for each model is identical to that of DCGAN for any metrics. It means that DDIM is vulnerable to membership inference attacks for small number of training samples instead of providing lower FID.

2) *Epochs:* We perform the attacks on five target models, i.e., models with 300 epochs to 700 epochs, to evaluate the impact of the number of epochs on a membership inference attack in the white-box setting. The results are shown in Fig 6. We also note that we omit results in the black-box setting since those for DDIM and DCGAN are identical. According to the figure, values of AUCROC become higher in proportion to the number of epochs, and it is especially stable for DDIM. Since AUCROC of DDIM increases linearly, we can also estimate the number of epochs to obtain the resistance to a membership inference attack.

TABLE III: Attack performance for different number of training samples in white-box/black-box setting.

| Model | Num. of samples | CIFAR-10 | | | | CelebA | | | |
|-------|-----------------|---------------|---------------|-----------------|--------|---------------|---------------|-----------------|--------|
| | | AUCROC | ASR | TPR at 1%FPR | FID | AUCROC | ASR | TPR at 1%FPR | FID |
| DDIM | 12,000 | 0.552 / 0.503 | 0.536 / 0.499 | 1.54 / 1.19% | 13.91 | 0.635 / 0.503 | 0.599 / 0.502 | 2.93 / 1.09% | 13.94 |
| | 6,000 | 0.801 / 0.499 | 0.736 / 0.500 | 5.65 / 0.90% | 16.27 | 0.973 / 0.497 | 0.901 / 0.499 | 65.30 / 0.90% | 16.64 |
| | 600 | 1.000 / 0.931 | 0.968 / 0.882 | 100.00 / 77.50% | 28.72 | 1.000 / 0.762 | 0.995 / 0.717 | 100.00 / 47.33% | 33.89 |
| DCGAN | 12,000 | 0.778 / 0.497 | 0.707 / 0.497 | 6.50 / 1.05% | 68.29 | 0.699 / 0.495 | 0.643 / 0.494 | 2.56 / 1.13% | 25.54 |
| | 6,000 | 0.983 / 0.497 | 0.899 / 0.499 | 69.80 / 0.78% | 79.30 | 0.776 / 0.492 | 0.694 / 0.490 | 8.33 / 0.95% | 31.99 |
| | 600 | 1.000 / 0.526 | 0.988 / 0.529 | 99.50 / 2.17% | 189.96 | 0.940 / 0.524 | 0.873 / 0.520 | 19.83 / 1.50% | 227.53 |

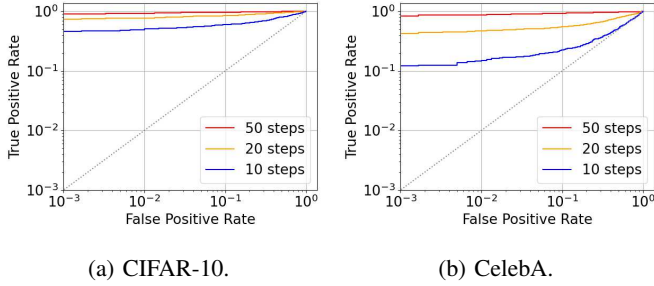


Fig. 7: ROC curves for different sampling steps on black-box setting.

TABLE IV: Attack performance for different sampling steps on black-box setting.

| Dataset | Steps | AUCROC | ASR | TPR at 1%FPR | FID |
|----------|-------|--------|-------|--------------|--------|
| CIFAR-10 | 50 | 0.986 | 0.962 | 92.83% | 22.60 |
| | 20 | 0.931 | 0.882 | 77.50% | 28.72 |
| | 10 | 0.781 | 0.724 | 51.17% | 45.50 |
| CelebA | 50 | 0.966 | 0.933 | 87.67% | 12.37 |
| | 20 | 0.762 | 0.717 | 47.33% | 33.89 |
| | 10 | 0.539 | 0.516 | 15.00% | 101.01 |

B. Impact of Hyperparameters for Sampling

Hyperparameters often affect the model’s performance, and hence we discuss the impact of hyperparameters for sampling, which are unique to diffusion models, i.e., sampling steps and sampling variances. Here, the investigation method in Section III-C is independent of hyperparameters for sampling since they are used only for the image generation process. Consequently, we evaluate only in the black-box setting.

1) *Sampling Steps*: We perform the attacks on three different numbers of steps, i.e., 50 steps, 20 steps, and 10 steps, to evaluate the impact of sampling steps on a membership inference attack. The results are shown in Fig. 7 and Table IV, where the number of epochs is 1000 and the model is trained with 600 images. According to the figure, DDIM is more vulnerable to the membership inference attack in proportion to the number of sampling steps. The reason is that generated images have passed through the trained networks more frequently in proportion to the number of sampling steps. Namely, the generated images represent the training results strongly. It also means that sampling steps significantly affect the resistance to a membership inference attack.

2) *Sampling Variances*: We perform the attacks on three different sampling variances, i.e., $\sigma(1)$, $\sigma(0)$, and $\hat{\sigma}$, to evaluate the impact of sampling variances that control the random-

TABLE V: Attack performance for different sampling variances on black-box setting.

| Dataset | Variance | AUCROC | ASR | TPR at 1%FPR | FID |
|----------|----------------|--------|-------|--------------|-------|
| CIFAR-10 | $\sigma(0)$ | 0.931 | 0.882 | 77.50% | 28.72 |
| | $\sigma(1)$ | 0.927 | 0.892 | 78.50% | 39.10 |
| | $\hat{\sigma}$ | 0.929 | 0.886 | 79.83% | 89.95 |
| CelebA | $\sigma(0)$ | 0.762 | 0.717 | 47.33% | 33.89 |
| | $\sigma(1)$ | 0.799 | 0.763 | 55.50% | 31.42 |
| | $\hat{\sigma}$ | 0.803 | 0.759 | 54.17% | 78.85 |

ness of the sampling on a membership inference attack. The results are shown in Table V, where the number of sampling steps is 20 and the model is trained with 600 images and 10,000 epochs. According to the table, although FID becomes larger in proportion to variances, DDIM is vulnerable to the membership inference attack for all variances. It means that the resistance to a membership inference attack is independent of sampling variances and they affect only the quality of generated images.

3) *Summary*: In summary, sampling steps are important for the resistance to a membership inference attack, while sampling variances are irrelevant. We leave it as an open problem to confirm if the same results are obtained in other architectures of diffusion models.

VI. CONCLUSION

In this paper, we investigated a membership inference attack on diffusion models. We primarily discussed the comparison of DDIM as a diffusion model with DCGAN as an existing generative model and hyperparameters of a diffusion model, i.e., timesteps, sampling steps, and sampling variances. We first showed that the diffusion model is comparably resistant to a membership inference attack as GAN through experiments on CelebA and CIFAR-10 datasets in the white-box and black-box settings. We then demonstrated the impact of timesteps on the membership inference attack. According to our result, intermediate steps in a noise schedule are the most vulnerable.

We also found two key insights into a membership inference attack on the diffusion model through analysis for overfitting, including comparison with a GAN, and hyperparameters for sampling. First, for overfitting, we confirmed if DDIM is vulnerable to the attack when the sample sizes are small because it could be more well-trained than DCGAN. Second, sampling steps are significantly important for resistance to the attack, while sampling variances are irrelevant. We are in the process of investigating further resistance to membership inference attacks on diffusion models, including other architectures.

Acknowledgement: A part of this work was supported by JST, CREST Grant JPMJCR21M5, Japan.

REFERENCES

- [1] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *Proc. of ICML 2015*, ser. PMLR, vol. 37. PMLR, 2015, pp. 2256–2265.
- [2] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” in *Proc. of NeurIPS 2020*, vol. 33. Curran Associates, Inc., 2020, pp. 6840–6851.
- [3] P. Dhariwal and A. Nichol, “Diffusion models beat gans on image synthesis,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 8780–8794, 2021.
- [4] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, “Hierarchical text-conditional image generation with clip latents,” *arXiv preprint arXiv:2204.06125*, 2022.
- [5] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10 684–10 695.
- [6] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes *et al.*, “Photorealistic text-to-image diffusion models with deep language understanding,” *arXiv preprint arXiv:2205.11487*, 2022.
- [7] A. Q. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, “GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models,” in *Proceedings of the 39th International Conference on Machine Learning*, 2022, pp. 16 784–16 804.
- [8] H. Hu and J. Pang, “Membership inference of diffusion models,” *arXiv preprint arXiv:2301.09956*, 2023.
- [9] N. Carlini, J. Hayes, M. Nasr, M. Jagielski, V. Schwag, F. Tramèr, B. Balle, D. Ippolito, and E. Wallace, “Extracting training data from diffusion models,” *arXiv preprint arXiv:2301.13188*, 2023.
- [10] J. Duan, F. Kong, S. Wang, X. Shi, and K. Xu, “Are diffusion models vulnerable to membership inference attacks?” *arXiv preprint arXiv:2302.01316*, 2023.
- [11] E. Choi, S. Biswal, B. Malin, J. Duke, W. F. Stewart, and J. Sun, “Generating multi-label discrete patient records using generative adversarial networks,” in *Proc. of MLHC 2017*, ser. PMLR, vol. 68. PMLR, 2017, pp. 286–305.
- [12] X. Yi, E. Walia, and P. Babyn, “Generative adversarial network in medical imaging: A review,” *Medical Image Analysis*, vol. 58, p. 101552, 2019.
- [13] B. Hilprecht, M. Härterich, and D. Bernau, “Monte carlo and reconstruction membership inference attacks against generative models,” *Proceedings of Privacy Enhancing Technologies*, vol. 2019, no. 4, pp. 232–249, 2019.
- [14] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *Proc. of IEEE S&P 2018*. IEEE Computer Society, 2017, pp. 3–18.
- [15] M. Conti, J. Li, S. Picek, and J. Xu, “Label-only membership inference attack against node-level graph neural networks,” in *Proc. of AISec 2022*, 2022, pp. 1–12.
- [16] Z. Li, Y. Liu, X. He, N. Yu, M. Backes, and Y. Zhang, “Auditing membership leakages of multi-exit networks,” in *Proc. of CCS 2022*. ACM, 2022, pp. 1917–1931.
- [17] J. Ye, A. Maddi, S. K. Murakonda, V. Bindschaedler, and R. Shokri, “Enhanced membership inference attacks against machine learning models,” in *Proc. of CCS 2022*. ACM, 2022, pp. 3093–3106.
- [18] J. Hayes, L. Melis, G. Danezis, and E. D. Cristofaro, “LOGAN: membership inference attacks against generative models,” *Proceedings of Privacy Enhancing Technologies*, vol. 2019, no. 1, pp. 133–152, 2019.
- [19] S. Mukherjee, Y. Xu, A. Trivedi, N. Patowary, and J. L. Ferres, “Privgan: Protecting gans from membership inference attacks at low cost to utility,” *Proceedings of Privacy Enhancing Technologies*, vol. 2021, no. 3, pp. 142–163, 2021.
- [20] T. Stadler, B. Oprisanu, and C. Troncoso, “Synthetic data – anonymisation groundhog day,” in *Proc. of USENIX Security 2022*, 2022, pp. 1451–1468.
- [21] 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)*. IEEE, 2017, pp. 3–18.
- [22] 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,” in *Proc. of NDSS 2019*. The Internet Society, 2019.
- [23] N. Carlini, S. Chien, M. Nasr, S. Song, A. Terzis, and F. Tramèr, “Membership inference attacks from first principles,” in *Proc. of IEEE S&P 2022*. IEEE, 2022, pp. 1897–1914.
- [24] L. Song, R. Shokri, and P. Mittal, “Membership inference attacks against adversarially robust deep learning models,” in *Proc. of IEEE SPW 2019*. IEEE, 2019, pp. 50–56.
- [25] F. Tramèr, R. Shokri, A. San Joaquin, H. Le, M. Jagielski, S. Hong, and N. Carlini, “Truth serum: Poisoning machine learning models to reveal their secrets,” in *Proc. of CCS 2022*. ACM, 2022, pp. 2779–2792.
- [26] S. Hidano, T. Murakami, S. Katsumata, S. Kiyomoto, and G. Hnnaoka, “Model inversion attacks for online prediction systems: Without knowledge of non-sensitive attributes,” *IEICE Transactions on Information and Systems*, vol. E101.D, no. 11, pp. 2665–2676, 2018.
- [27] Z. Wang, Y. Huang, M. Song, L. Wu, F. Xue, and K. Ren, “Poisoning-assisted property inference attack against federated learning,” *IEEE Transactions on Dependable and Secure Computing*, pp. 1–13, 2022.
- [28] S. Mahloujifar, E. Ghosh, and M. Chase, “Property inference from poisoning,” in *Proc. of IEEE S&P 2022*. IEEE, 2022, pp. 1569–1569.
- [29] C. Dwork, “Differential privacy,” in *Proc. of ICALP*, ser. LNCS, vol. 4052. Springer, 2006, pp. 1–12.
- [30] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, “Deep learning with differential privacy,” in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [31] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha, “Privacy risk in machine learning: Analyzing the connection to overfitting,” in *Proc. of CSF 2018*. IEEE, 2018, pp. 268–282.
- [32] B. Jayaraman and D. Evans, “Evaluating differentially private machine learning in practice,” in *Proc. of USENIX Security 2019*. USENIX Association, 2019, pp. 1895–1912.
- [33] M. Jagielski, J. Ullman, and A. Oprea, “Auditing differentially private machine learning: How private is private sgd?” in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 22 205–22 216.
- [34] D. Chen, N. Yu, Y. Zhang, and M. Fritz, “GAN-leaks: A taxonomy of membership inference attacks against generative models,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*.
- [35] Y. Wu, N. Yu, Z. Li, M. Backes, and Y. Zhang, “Membership inference attacks against text-to-image generation models,” *CoRR*, vol. abs/2210.00968, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2210.00968>
- [36] C. Song and V. Shmatikov, “Auditing data provenance in text-generation models,” in *Proc. of CCS 2019*. ACM, 2019, pp. 196–206.
- [37] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, G. Klambauer, and S. Hochreiter, “Gans trained by a two time-scale update rule converge to a nash equilibrium,” *CoRR*, vol. abs/1706.08500, 2017.
- [38] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” *CoRR*, vol. abs/2010.02502, 2020.
- [39] A. Krizhevsky, G. Hinton *et al.*, “Learning multiple layers of features from tiny images,” 2009.
- [40] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild,” *CoRR*, vol. abs/1411.7766, 2014.
- [41] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2016.
- [42] A. Nichol and P. Dhariwal, “Improved denoising diffusion probabilistic models,” *CoRR*, vol. abs/2102.09672, 2021.