
On the Privacy Properties of GAN-generated Samples

Zinan Lin

Carnegie Mellon University

Vyas Sekar

Carnegie Mellon University

Giulia Fanti

Carnegie Mellon University

Abstract

The privacy implications of generative adversarial networks (GANs) are a topic of great interest, leading to several recent algorithms for training GANs with privacy guarantees. By drawing connections to the generalization properties of GANs, we prove that under some assumptions, GAN-generated samples inherently satisfy some (weak) privacy guarantees. First, we show that if a GAN is trained on m samples and used to generate n samples, the generated samples are (ϵ, δ) -differentially-private for (ϵ, δ) pairs where δ scales as $O(n/m)$. We show that under some special conditions, this upper bound is tight. Next, we study the robustness of GAN-generated samples to membership inference attacks. We model membership inference as a hypothesis test in which the adversary must determine whether a given sample was drawn from the training dataset or from the underlying data distribution. We show that this adversary can achieve an area under the ROC curve that scales no better than $O(m^{-1/4})$.

1 Introduction

Generative adversarial networks (GANs) are a class of generative models that aim to generate samples from a distribution μ given a database of training samples $D = \{x_i\}_{i=1}^m$, where the $x_i \in X$ are drawn i.i.d. from a distribution μ (Goodfellow et al., 2014). GANs are posed as a zero-sum game between two neural networks: a *generator* that aims to generate new samples $\hat{x}_i \sim \mu$ and a *discriminator* that aims to classify samples as being either real (i.e., a member of D) or generated. In recent years, GANs have garnered interest as a tool

for generating synthetic data from potentially-sensitive raw datasets, such as medical data (Esteban et al., 2017; Jordon et al., 2018; Choi et al., 2017), banking transactions (Jordon et al., 2018), and networking and server traces (Lin et al., 2020a). However, GANs are known to leak information about the data on which they were trained. Common concerns include memorization of sensitive training samples (Webster et al., 2019) and membership inference attacks (Hayes et al., 2019; Chen et al., 2019).

In response to these concerns, a popular approach in the literature has been to train differentially-private GAN models. This can be accomplished through training methods like differentially-private stochastic gradient descent (DP-SGD) (Abadi et al., 2016), PATE-GAN (Jordon et al., 2018), and others (Long et al., 2019; Chen et al., 2020). These approaches guarantee differential privacy for releasing the *model parameters*. In practice, they also degrade model quality due to the added noise, sometimes to the extent that the utility of data is completely destroyed (Lin et al., 2020a).

In some practical scenarios, an attacker may only have access to a GAN's *generated samples*, instead of the *model parameters* (e.g., when releasing synthetic datasets instead trained models (Choi et al., 2017; Lin et al., 2020a)). Note that ensuring the privacy for releasing generated samples is a weaker condition than ensuring the privacy for releasing parameters because of the post-processing property of differential privacy (Dwork et al., 2014). Releasing generated samples (versus model parameters) is therefore a promising solution for achieving better privacy-fidelity trade-offs. As a first step towards this goal, it is important to understand the privacy properties of GAN-generated samples.

In this work, we show that GAN-generated samples satisfy an *inherent* privacy guarantee without any special training mechanism, by making connections to recent results on the generalization properties of GANs (Zhang et al., 2017). The strength of this privacy guarantee varies for different notions of privacy. Specifically, we study two privacy notions: *differential privacy (DP)* and *robustness to membership inference attacks*.

- For differential privacy, we show that the samples output by a vanilla GAN (i.e., trained without differential privacy) inherently satisfy a (weak) differential privacy guarantee with high probability over the randomness of the training procedure. More specifically, we consider the mechanism of training a GAN on m training samples and using it to generate n samples. We show that this mechanism ensures $(\epsilon, \frac{O(n/m)}{\epsilon(1-e^{-\epsilon})})$ -differentially privacy with high probability (under some assumptions). The rate $O(n/m)$ shows that differential privacy guarantee is stronger as the training set size grows, but degrades as you generate more samples from the GAN. Additionally, we show that $O(1/m)$ rate is the optimal rate we can get when $n = 1$. On the positive side, the results suggest that GAN-generated samples are inherently differentially private. On the negative side, however, the rate is as weak as releasing n raw samples from m training samples. This suggests that we need to incorporate other techniques (e.g., DP-SGD, PATE) in practice if we want meaningful differential privacy guarantees.
- For robustness to membership inference attacks, we show that vanilla GANs are *inherently* robust to black-box membership inference attacks. We study the worst-case setting where an attacker can draw infinite samples from the trained GAN. Even in this case, we show that for any attacker strategy, the difference between the true positive rate and the false positive rate—as well as the area under the ROC curve (AUC)—both have upper bounds that scale as $O(m^{-\frac{1}{4}})$, where m is the number of training samples. This again means that GANs are more robust to membership inference attacks as training set size grows. More generally, for arbitrary generative models, we give a tight bound on the ROC region that relies on a simple, geometric proof. To the best of our knowledge, this is the first result to bound either the ROC region or the rate of decay of attack success probabilities for membership inference attacks, including for discriminative models.

The paper is organized as follows. We discuss the preliminaries and related work in Section 2. Then we discuss the results on differential privacy and membership inference attacks in Section 3 and Section 4 respectively. Finally we conclude the paper in Section 5.

2 Background and Related Work

GANs GANs are a type of generative model that implicitly learns a distribution from samples and syn-

thesizes new, random samples (Goodfellow et al., 2014). To achieve this goal, GANs use two components: a generator g and a discriminator f . The generator maps a random vector z from a pre-defined distribution p_z (e.g., Gaussian or uniform) on a latent space to the sample space X ; the discriminator reads a sample x , either from the real distribution μ or the generated distribution, and discriminates which distribution it comes from. From another point of view, the discriminator is trying to estimate a distance between the generated distribution and the real distribution, and the generator tries to minimize the distance. More specifically, g and f are trained with the following loss:

$$\min_g \max_f d_f(\mu, g_z),$$

where d_f is a distance measurement between distributions parametrized by f , and g_z denotes the distribution obtained by sampling $z \sim p_z$ and mapping it through the generator $g(z)$. In the original GAN paper (Goodfellow et al., 2014), the authors use Jensen–Shannon divergence with formula

$$d_f(\mu, g_z) = \mathbb{E}_{x \sim \mu} [\log f(x)] + \mathbb{E}_{x \sim g_z} [\log (1 - f(x))].$$

Many other distances have later been used like the Kullback-Leibler divergence and Wasserstein distance (Nowozin et al., 2016; Gulrajani et al., 2017). After the invention of GANs, many approaches were proposed to improve their diversity and fidelity (Gulrajani et al., 2017; Miyato et al., 2018; Lin et al., 2018; Arjovsky and Bottou, 2017); today, GANs generate state-of-the-art realistic images (Karras et al., 2017; Brock et al., 2018). Inspired by these early successes, there has been much interest in using GANs to share data in privacy-sensitive applications such as medical data (Esteban et al., 2017; Jordon et al., 2018; Choi et al., 2017), banking transactions (Jordon et al., 2018), and networking and server traces (Lin et al., 2020a).

Differential privacy Differential privacy (DP) has become the *de facto* formal privacy definition in many applications (Dwork, 2008; Dwork et al., 2014). We say that two databases D_0 and D_1 are *neighboring* if they differ in at most one element. A mechanism M is (ϵ, δ) -differentially-private (Dwork, 2008; Dwork et al., 2014) if for any neighboring database D_0 and D_1 , and any set $S \subseteq \text{range}(M)$,

$$\mathbb{P}[M(D_0) \in S] \leq e^\epsilon \mathbb{P}[M(D_1) \in S] + \delta.$$

Our result considers a stronger notion of differential privacy called *probabilistic differential privacy*. A mechanism M is (ϵ, δ) -probabilistically-differentially-private (Meiser, 2018) if for any neighboring database D_0 and D_1 , there exists sets $S_0 \subseteq \text{range}(M)$ where

$\mathbb{P}[M(D_0) \subseteq S_0] \leq \delta$, such that for any set $S \subseteq \text{range}(M)$

$$\mathbb{P}[M(D_0) \in S \setminus S_0] \leq e^\epsilon \mathbb{P}[M(D_1) \in S \setminus S_0].$$

This says that $(\epsilon, 0)$ -differential-privacy condition holds except over a region of the support with probability mass at most δ .

It is straightforward to show that (ϵ, δ) -probabilistically-differential-privacy implies (ϵ, δ) -differential-privacy. In fact, probabilistic differential privacy is *strictly* stronger than differential privacy. To see this, consider the following example: assume ρ_p, ρ_q are the distribution function of $M(D_0)$ and $M(D_1)$ respectively. When $e^\epsilon(1 - \delta) < 1$ and $\delta > 0$, let $\epsilon' = \min\{1 - e^\epsilon(1 - \delta), \delta\}$, we can construct the following ρ_p, ρ_q : $\rho_q(0) = 0, \rho_p(0) = 1 - (1 - \epsilon')e^{-\epsilon}$ and $\rho_q(1) = 1, \rho_p(1) = (1 - \epsilon')e^{-\epsilon}$. Then M satisfies (ϵ, δ) -differential-privacy, but does not satisfy (ϵ, γ) -probabilistically-differential-privacy for any $\gamma < 1$.

Differential privacy and generalization The relation between differential privacy and generalization is well-studied (Cummings et al., 2016; Dwork et al., 2015; Bassily et al., 2016; Nissim and Stemmer, 2015; Wang et al., 2016). Three main factors differentiate our work:

(1) Prior work has primarily studied how differential privacy implies various notions of generalization (Cummings et al., 2016; Dwork et al., 2015; Bassily et al., 2016; Nissim and Stemmer, 2015; Wang et al., 2016). We are interested in the other direction: when does generalization imply differential privacy? Cummings et al. showed that differential privacy is strictly weaker than perfect generalization but is strictly stronger than robust generalization (Cummings et al., 2016). However, these results do not directly apply to generative models. First, the results about perfect generalization only hold when the domain size is finite, which is less interesting for generative models. Second, the generalization definitions are based on hypotheses (i.e., conditional distributions), and are not easily extendable to generative models.

(2) More generally, prior work has mainly considered discriminative models instead of generative models (Cummings et al., 2016; Dwork et al., 2015; Bassily et al., 2016; Nissim and Stemmer, 2015; Wang et al., 2016). Wu et al. (2019) showed that differential privacy implies generalization in GANs. However, the notion of generalization studied by Wu et al. (2019) is not meaningful for GANs. It captures the distance between the expected empirical loss of the discriminator and its actual loss, so we can have zero generalization error on the discriminator’s loss (e.g., for a discriminator that always outputs 0 for any input) while having an

arbitrary generated distribution. A more meaningful notion of generalization is to quantify the distance between the generated and real distributions. In this work, we use this latter notion of generalization.

(3) The notion of differential privacy studied in prior work is with respect to releasing the *parameters* of a model (Cummings et al., 2016; Dwork et al., 2015; Bassily et al., 2016; Nissim and Stemmer, 2015; Wang et al., 2016; Wu et al., 2019), whereas ours is with respect to releasing *generated samples*. This is a weaker guarantee than if one were to reason about the differential privacy guarantees of a generative model’s parameters, but is relevant to many application settings.

Membership inference attacks Membership inference attacks are closely related to differential privacy. Given a trained model, a membership inference attack aims to infer whether a given sample x was in the training dataset or not. The main difference between membership inference and differential privacy is that the attacker in differential privacy is assumed to know an adversarially-chosen pair of candidate training databases, whereas in membership inference attacks, the adversary is typically given access only to test samples (of which some are training samples) and the model. In some cases, the attacker is also given side information about the number of training samples in the test set. Hence, in general, the attacker in membership inference is neither strictly weaker nor strictly stronger than the differential privacy attacker.

There have been many membership inference attacks proposed for discriminative models (Sablayrolles et al., 2019; Li and Zhang, 2020; Long et al., 2018; Melis et al., 2019; Salem et al., 2018; Shokri et al., 2017; Yeom et al., 2018) and generative models (including GANs) (Hayes et al., 2019; Chen et al., 2019; Hilprecht et al., 2019). Therefore, understanding robustness of GANs to membership inference attacks is important.

There has been some theoretical analysis on membership inference attacks (Sablayrolles et al., 2019; Farokhi and Kaafar, 2020). For example, Sablayrolles et al. (2019) show a theoretically optimal strategy for membership inference attacks. Farokhi and Kaafar (2020) show that the accuracy of membership inference attacks for a particular sample is upper bounded by the Kullback–Leibler divergence between the distributions of parameters with and without that sample. However, these results do not give a practical method for bounding an attacker’s global performance (i.e., the ROC curve). In this paper, we resolve the issue. Ours is the first work to show that the success rate of membership inference attack decays as the number of training samples m grows, whereas prior work (Farokhi and Kaafar, 2020) only shows that the limit of the success rate is

0.5 as $m \rightarrow \infty$.

3 Bounds on Differential Privacy

We start with some notation, definitions, and assumptions. For a probability measure μ on X , we let ρ_μ denote its density function. We use \mathcal{G} and \mathcal{F} to denote the set of possible generators and discriminators, respectively, where \mathcal{F} is a set of functions $X \rightarrow \mathbb{R}$. Our results rely on three assumptions:

- (A1) Our generator set \mathcal{G} and discriminator set \mathcal{F} satisfy $\forall \nu_1, \nu_2 \in \mathcal{G}, \log(\rho_{\nu_1}/\rho_{\nu_2}) \in \text{span}\mathcal{F}$, where $\text{span}\mathcal{F}$ is defined as the set of linear combinations of the functions:

$$\text{span}\mathcal{F} \triangleq \left\{ w_0 + \sum_{i=1}^n w_i f_i : w_i \in \mathbb{R}, f_i \in \mathcal{F}, n \in \mathbb{N} \right\}.$$

- (A2) Our discriminator set \mathcal{F} is even, i.e., $\forall f \in \mathcal{F}, -f \in \mathcal{F}$, and assume that $\forall f \in \mathcal{F}, \forall x \in X$,

$$\|f(x)\|_\infty \leq \Delta,$$

where $\|f\|_\infty \triangleq \sup_{x \in X} |f(x)|$.

- (A3) The discriminator set $\mathcal{F} = \{f_\theta : \theta \in \Theta \subseteq [-1, 1]^p\}$ and

$$\|f_\theta - f_{\theta'}\|_\infty \leq L \|\theta - \theta'\|_2.$$

According to Lemma 4.1 in Bai et al. (2018), when generators in \mathcal{G} are invertible neural networks (e.g., in Dinh et al. (2016); Behrmann et al. (2019)) with l layers, then the discriminator set \mathcal{F} of neural networks with $l + 2$ layers satisfy (A1). Assumption (A2) is easily satisfied by neural networks with an activation function on the output layer that bounds the output (e.g., sigmoid). (A3) assumes the discriminator is Lipschitz in its (bounded) parameters; several recent works attempt to make network layers Lipschitz in inputs through various forms of regularization (e.g., spectral normalization (Miyato et al., 2018; Lin et al., 2020b)), which makes the network Lipschitz also in parameters (Lin et al., 2020b).

Given μ, ν , two probability measures on X , and set \mathcal{F} of functions $X \rightarrow \mathbb{R}$, the *integral probability metric* (Müller, 1997) is defined as:

$$d_{\mathcal{F}}(\mu, \nu) \triangleq \sup_{f \in \mathcal{F}} \{ \mathbb{E}_{x \sim \mu} [f(x)] - \mathbb{E}_{x \sim \nu} [f(x)] \}.$$

Given function set \mathcal{F} , the \mathcal{F} -variation norm of function g is defined as

$$\|g\|_{\mathcal{F},1} \triangleq \inf \left\{ \sum_{i=1}^n |w_i| : g = w_0 + \sum_{i=1}^n w_i f_i, \right. \\ \left. \forall n \in \mathbb{N}, w_i \in \mathbb{R}, f_i \in \mathcal{F} \right\},$$

which intuitively describes the complexity of linearly representing g using functions in \mathcal{F} (Zhang et al., 2017). We define

$$\Gamma_{\mathcal{F},\mathcal{G}} \triangleq \sup_{\nu_1, \nu_2 \in \mathcal{G}} \|\log(\rho_{\nu_1}/\rho_{\nu_2})\|_{\mathcal{F},1}, \quad (1)$$

which intuitively bounds the complexity of representing differences in log densities of pairs of generators in \mathcal{G} using functions in \mathcal{F} .

We consider the GAN training and sampling mechanism in Algorithm 1, which adds a sampling processing before the normal GAN training. Note that the sampling process in Algorithm 1 is commonly used in existing GAN implementations (Goodfellow et al., 2014), which typically sample i.i.d. from D in each training batch. We have moved this sampling process to the beginning of training in Algorithm 1 for ease of analysis.

Algorithm 1: Differentially-private GAN mechanism.

Input : D : A training dataset containing m samples.
 k : Number of sampled training samples used in training.
 n : Number of generated samples.

Output : $D_{\text{generated}}$: n generated samples.

- 1 $D_{\text{train}} \leftarrow k$ i.i.d. samples from D ;
 - 2 $g \leftarrow$ Trained GAN using D_{train} ;
 - 3 $D_{\text{generated}} \leftarrow n$ generated samples from the trained g ;
-

Finally, we define the quantity

$$\tau_{\mathcal{F},\mathcal{G}}(k, \xi, \mu) \triangleq 2 \left(\inf_{\nu \in \mathcal{G}} d_{\mathcal{F}}(\mu, \nu) + \tau_{\text{opt}} + \frac{C_\xi}{\sqrt{k}} \right), \quad (2)$$

where

$$C_\xi = 16\sqrt{2\pi}pL + 2\Delta\sqrt{2\log(1/\xi)}, \quad (3)$$

p and L are the constants defined in (A3), Δ is the constant defined in (A2), τ_{opt} is an upper bound on the optimization error, i.e.,

$$d_{\mathcal{F}}(\hat{\mu}_k, g) - \inf_{\nu \in \mathcal{G}} d_{\mathcal{F}}(\hat{\mu}_k, \nu) \leq \tau_{\text{opt}},$$

μ is the real distribution, $\hat{\mu}_k$ be the empirical distribution of μ on k i.i.d training samples, g is the trained generator from the optimization algorithm. The interpretation of $\tau_{\mathcal{F},\mathcal{G}}(k, \xi, \mu)$ is described in greater detail in Section 3.1, but intuitively, it will be used to bound (with probability at least $1 - \xi$) a GAN's generalization error arising from approximation, optimization, and sampling of the training datasets in Line 1. To reduce notation, we will henceforth write this quantity as $\tau_{k,\xi}$.

Our main result states that a GAN trained on m samples and used to generate n samples satisfies a $(\epsilon, \frac{O(n/m)}{\epsilon(1-e^{-\epsilon})})$ -differential-privacy guarantee; moreover, this bound is tight when $n = 1$ and for small ϵ .

Theorem 1 (Achievability). *Consider a GAN trained on m i.i.d. samples from distribution μ . The mechanism in Algorithm 1 under assumptions (A1)-(A3) satisfies (ϵ, δ) -probabilistic-differential-privacy for any $\epsilon > 0$ and*

$$\delta \geq \frac{n \Gamma_{\mathcal{F}, \mathcal{G}}}{\epsilon(1-e^{-\epsilon})} \left(\frac{2\Delta}{m} + \tau_{k, \xi} \right) \quad (4)$$

with probability at least $1 - 2\xi$ over the randomness in Line 1.

We assume that $\tau_{k, \xi}$ can be made negligible: $\frac{C_\xi}{\sqrt{k}}$ can be arbitrarily small as we get more samples from the sampling phase in Line 1, and we assume negligible approximation error and optimization error. Hence, the dominating term in Eq. (4) scales as $O(n/m)$; here we are ignoring the dependency on ϵ . Next, we show that for a special case where $n = 1$ and ϵ scales as $O(\frac{1}{m})$, this bound is tight in an order sense (again ignoring dependencies on ϵ).

Proposition 1 (Converse for $n = 1$). *Under the assumptions of Theorem 1, let*

$$\Delta' = \sup_{f \in \mathcal{F}} \sup_{x, y \in X} |f(x) - f(y)|.$$

Then with probability at least $1 - 2\xi$, the GAN mechanism in Algorithm 1 for generating 1 sample (i.e., $n = 1$) does not satisfy (ϵ, δ) -differential-privacy for any

$$\delta < \frac{(e^\epsilon + 1)}{2\Delta} \left(\frac{\Delta'}{2m} - \tau_{k, \xi} \right) + 1 - e^\epsilon. \quad (5)$$

This bound in Eq. (5) is non-vacuous (nonnegative) when $\epsilon < \frac{1}{\Delta} \left(\frac{\Delta'}{2m} - \tau_{k, \xi} \right)$ and $m < \Delta' / 2\tau_{k, \xi}$. Again assuming $\tau_{k, \xi} \approx 0$ (i.e., ignoring the approximation error and optimization error and taking $\frac{C_\xi}{\sqrt{k}} \rightarrow 0$), the latter condition holds trivially. Hence when ϵ scales as $O(\frac{1}{m})$, it is not possible to achieve an (ϵ, δ) -probabilistic differential privacy guarantee for $\delta = o(\frac{1}{m})$.

Discussion These results suggest that GAN-generated samples satisfy an inherent differential privacy guarantee, so the influence of any single training sample on the final generated samples is bounded. However, the rate $O(n/m)$ is weak. For comparison, the mechanism that releases n samples uniformly at random from a set of m training samples satisfies $(0, n/m)$ -differential privacy, which is of the same rate. Therefore,

(ϵ, δ) -differential privacy usually requires $\delta \ll \frac{1}{\text{poly}(m)}$ to be meaningful. To satisfy this condition, we would need ϵ to grow as a function of m (since δ in our results is a function of ϵ), which is not practically viable. This suggests the need for incorporating additional techniques (e.g., DP-SGD, PATE) to achieve stronger differential privacy guarantees in practice.

Note that in the typical GAN training process, we usually have $k > m$ (i.e., the number of sampled training samples is larger than the dataset size). Therefore, the random sampling step in Line 1 does not give privacy amplification due to subsampling (e.g., (Balle et al., 2018)). However, if the size of the training dataset is large enough that we do not need to sample all dataset entries to achieve good generalization, we can apply sub-sampling theorems (Balle et al., 2018) to tighten the bounds in Theorem 1.

3.1 Proof of Theorem 1

Assume that the two neighboring datasets are D_0 and D_1 , whose empirical distributions are $\hat{\mu}_m^0$ and $\hat{\mu}_m^1$, and the trained generator distributions from Algorithm 1 are g_0 and g_1 respectively. The proof has two parts. First, we upper bound the distance between g_0 and g_1 by building on prior generalization results. Then, we use this distance to prove a probabilistic differential privacy guarantee.

Lemma 1 (Error of neural-network discriminators, Corollary 3.3 in Zhang et al. (2017)). *Let μ be the real distribution, and $\hat{\mu}_k$ be the empirical distribution of μ on k i.i.d training samples. Define the trained generator from the optimization algorithm as g and assume the optimization error is bounded by τ_{opt} , i.e., $d_{\mathcal{F}}(\hat{\mu}_k, g) - \inf_{\nu \in \mathcal{G}} d_{\mathcal{F}}(\hat{\mu}_k, \nu) \leq \tau_{\text{opt}}$. Under assumptions (A2) and (A3), then with probability at least $1 - \xi$ w.r.t. the randomness of training samples, we have*

$$d_{\mathcal{F}}(\mu, g) \leq \frac{1}{2}\tau_{k, \xi} = \underbrace{\inf_{\nu \in \mathcal{G}} d_{\mathcal{F}}(\mu, \nu)}_{\text{approximation error}} + \underbrace{\tau_{\text{opt}}}_{\text{optimization error}} + \underbrace{\frac{C_\xi}{\sqrt{k}}}_{\text{generalization error}},$$

where C_ξ is defined in Eq. (3).

From this lemma, we know that with high probability, $d_{\mathcal{F}}(\hat{\mu}_m^i, g_i)$ is small. The next lemma states that $d_{\mathcal{F}}(\hat{\mu}_m^0, \hat{\mu}_m^1)$ is also small.

Lemma 2. *Under the assumption (A2), for any two neighboring datasets D_0, D_1 with m samples, we have*

$$d_{\mathcal{F}}(\hat{\mu}_m^0, \hat{\mu}_m^1) \leq \frac{2\Delta}{m},$$

(Proof in Appendix A.) Next, we use these results to argue that $d_{\mathcal{F}}(g_0, g_1)$ is small with high probability.

Lemma 3. *Assume we have two training sets D_0 and D_1 , and the trained generator distributions using D_0 and D_1 with Algorithm 1 are g_0 and g_1 , respectively. Under the assumption of Lemma 1 and Lemma 2, we have that with probability at least $1 - 2\xi$,*

$$d_{\mathcal{F}}(g_0, g_1) \leq \tau_{k,\xi} + \frac{2\Delta}{m}.$$

(Proof in Appendix B.) We next use this bound on the integral probability metric to bound Kullback-Leibler (KL) divergence between g_0 and g_1 with the following lemma.

Lemma 4. *Given a generator set \mathcal{G} and a discriminator set \mathcal{F} which satisfy assumption (A1), then we have $\forall \nu_1, \nu_2 \in \mathcal{G}$*

$$d_{\text{KL}}(\nu_1, \nu_2) + d_{\text{KL}}(\nu_2, \nu_1) \leq \Gamma_{\mathcal{F},\mathcal{G}} d_{\mathcal{F}}(\nu_1, \nu_2)$$

where $\Gamma_{\mathcal{F},\mathcal{G}}$ is defined in Eq. (1) and $d_{\text{KL}}(\cdot, \cdot)$ is the Kullback-Leibler divergence.

This follows directly from Proposition 2.9 in Zhang et al. (2017), which states the following. Denote μ 's and ν 's density functions as ρ_{μ} and ρ_{ν} respectively. If $\log(\rho_{\mu}/\rho_{\nu}) \in \text{span}\mathcal{F}$, then we have

$$d_{\text{KL}}(\mu, \nu) + d_{\text{KL}}(\nu, \mu) \leq \|\log(\rho_{\mu}/\rho_{\nu})\|_{\mathcal{F},1} d_{\mathcal{F}}(\mu, \nu).$$

Note that $\Gamma_{\mathcal{F},\mathcal{G}} = 1$ when generators in \mathcal{G} are invertible neural networks with l layers and discriminator set \mathcal{F} is $(l+2)$ -layer neural networks, according to Lemma 4.1 in Bai et al. (2018).

Following Lemma 4 and Lemma 3, immediately we have that with probability at least $1 - 2\xi$,

$$d_{\text{KL}}(g_0, g_1) + d_{\text{KL}}(g_1, g_0) \leq \Gamma_{\mathcal{F},\mathcal{G}} \left(\tau_{k,\xi} + \frac{2\Delta}{m} \right)$$

and

$$d_{\text{KL}}(g_0^n, g_1^n) + d_{\text{KL}}(g_1^n, g_0^n) \leq n \cdot \Gamma_{\mathcal{F},\mathcal{G}} \left(\tau_{k,\xi} + \frac{2\Delta}{m} \right)$$

Given these facts, Theorem 1 follows directly by connecting KL divergence with differential privacy:

Lemma 5. *If a mechanism M satisfies that for any two neighboring databases D_0 and D_1 , $d_{\text{KL}}(p, q) + d_{\text{KL}}(q, p) \leq s$, where p, q are the probability measure of $M(D_0)$ and $M(D_1)$ respectively, then M satisfies $(\epsilon, \frac{s}{\epsilon(1-e^{-\epsilon})})$ -probabilistic-differential-privacy for all $\epsilon > 0$.*

(Proof in Appendix C.)

3.2 Proof of Proposition 1

The proof has two parts. First, we lower bound the distance between g_0 and g_1 by building on prior generalization results. Then, we use this distance to prove the lower bound of δ for (ϵ, δ) -differential-privacy.

Because $\Delta' = \sup_{f \in \mathcal{F}} \sup_{x, y \in X} |f(x) - f(y)|$, we know that there exists $f' \in \mathcal{F}$ and $x_1, y_1 \in X$ such that $|f'(x_1) - f'(y_1)| \geq \frac{\Delta'}{2}$. Let's construct two databases: $D_0 = \{x_1, \dots, x_m\}$ and $D_1 = \{y_1, x_2, \dots, x_m\}$ where x_2, \dots, x_m are arbitrary samples from X . Then we have

$$d_{\mathcal{F}}(\hat{\mu}_m^0, \hat{\mu}_m^1) = \frac{1}{m} \sup_{f \in \mathcal{F}} \{f(x_1) - f(y_1)\} \geq \frac{\Delta'}{2m}$$

where the proof of the first equality is in Appendix A.

On the other hand, from Lemma 1, we know that with probability at least $1 - \xi$, $d_{\mathcal{F}}(\hat{\mu}_m^i, g_i) \leq \frac{1}{2}\tau_{k,\xi}$ for $i = 0, 1$. Combining the above, we have the following lemma.

Lemma 6. *Assume we have two training sets D_0 and D_1 , and the trained generator using D_0 and D_1 with Algorithm 1 are g_0 and g_1 respectively. Under the assumption of Lemma 1 and Lemma 2, we have that with probability at least $1 - 2\xi$,*

$$d_{\mathcal{F}}(g_0, g_1) \geq \frac{\Delta'}{2m} - \tau_{k,\xi}$$

(Proof in Appendix D.) Note that this lower bound is nonnegative only for $m < \Delta'/2\tau_{k,\xi}$.

Now we connect integral probability metric to total variation (TV) distance. Because the discriminators are bounded (A2), we have

$$d_{\mathcal{F}}(g_0, g_1) \leq 2\Delta d_{\text{TV}}(g_0, g_1)$$

where $d_{\text{TV}}(g_0, g_1)$ is the TV distance between g_0 and g_1 . From the above, we know that

$$d_{\text{TV}}(g_0, g_1) \geq \frac{\Delta'}{4m\Delta} - \frac{\tau_{k,\xi}}{2\Delta}$$

Finally, we connect TV distance to differential privacy with the following lemma.

Lemma 7. *If a mechanism satisfies (ϵ, δ) -differential-privacy, then for any two neighboring databases D_0 and D_1 , we have*

$$d_{\text{TV}}(p, q) \leq \frac{e^{\epsilon} + 2\delta - 1}{e^{\epsilon} + 1}$$

where p, q are the probability measure of $M(D_0)$ and $M(D_1)$ respectively.

Therefore, we have

$$\delta \geq \frac{(e^{\epsilon} + 1)\Delta'}{4m\Delta} - \frac{(e^{\epsilon} + 1)\tau_{k,\xi}}{2\Delta} + 1 - e^{\epsilon}.$$

4 Bounds on Robustness to Membership Inference Attacks

In this section, we first derive a general bound on the error of membership inference attacks for generative models. Then, we utilize generalization bounds for GANs to obtain specific bounds for GANs.

We focus on the black-box attack setting (Sablayrolles et al., 2019; Chen et al., 2019; Hayes et al., 2019), in which the attacker can sample from the generated distribution g_α , but does not have access to the generator parameters α . To upper bound the attack performance, we assume that attacker has unlimited resources and can access the trained generator infinite times, so that it can accurately get the generated distribution g_α .

Our analysis departs from prior analysis of membership inference in discriminative models (Sablayrolles et al., 2019) in two key respects:

- (Sablayrolles et al., 2019) assume that the attacker has access to a dataset $U = \{u_1, \dots, u_p\}$, which contains all the training samples (i.e., $D \subseteq U$) and some other test samples drawn from the ground-truth distribution μ (so $p > m$). It also assumes that the attacker knows the number of training samples m . We argue that this assumption is too strong, especially in the case of generative models, where training samples are typically proprietary. Therefore, we assume that the attacker makes guesses purely based on a single test sample $x \in X$, without access to such a dataset. The test sample is either drawn from the ground-truth distribution (i.e., $x \sim \mu$), or from the training dataset (i.e., $x \stackrel{\text{i.i.d.}}{\sim} D$).
- The analysis in (Sablayrolles et al., 2019) focuses on the quantity $\mathbb{P}(u \in D | \alpha)$ for a particular u . This is useful for finding an attack policy, but is not conducive to characterizing the error statistics. Instead, we want to be able to bound the shape of ROC curve. That is, we want to upper bound the true positive rate an attacker can achieve given any desired false positive rate. We show that this problem can be reduced to a clean hypothesis testing problem, whose errors are closely tied to the generalization errors of GANs.

Following prior work on the theoretical analysis of membership inference attacks (Sablayrolles et al., 2019), we assume that the distribution of the generator parameters is

$$\mathbb{P}(\alpha | x_1, \dots, x_m) \propto e^{-\sum_{i=1}^m \ell(\alpha, x_i)} \quad (6)$$

(setting $T = 1$ in Sablayrolles et al. (2019)), where $x_1, \dots, x_m \sim \mu$ are i.i.d training samples drawn from

the ground-truth distribution μ , and $\ell(\alpha, x)$ denotes the loss on sample x and parameter α . As we are focusing on generative models here, we assume that the loss is Kullback-Leibler (KL) divergence, i.e., $\ell(\alpha, x) = \log(\rho_\mu(x_i)/\rho_{g_\alpha}(x_i))$, where ρ_{g_α} denotes the density of the generator with parameters α . Note that many generative models are explicitly or implicitly minimizing this KL divergence, including some variants of GANs (more specifically, f-GANs with a specific loss (Nowozin et al., 2016)), Variational Autoencoder (VAE) (Kingma and Welling, 2013), PixelCNN/PixelRNN (Oord et al., 2016), and many other methods that are based on maximum likelihood (Goodfellow et al., 2016). We use KL divergence also because this simplifies the analysis and highlights key theoretical insights. With this assumption, the parameter distribution becomes

$$\mathbb{P}(\alpha | x_1, \dots, x_m) \propto \prod_{i=1}^m \frac{\rho_{g_\alpha}(x_i)}{\rho_\mu(x_i)}.$$

Let $\rho_{\text{train}}^\alpha$ denotes the density posterior distribution of the training samples given parameter α . The following proposition shows that this distribution takes a simple form.

Proposition 2 (Posterior distribution of training samples). *The posterior distribution of training samples is equal to the generated distribution, i.e., $\rho_{\text{train}}^\alpha = \rho_{g_\alpha}$.*

Proof. For any x , we have

$$\begin{aligned} & \rho_{\text{train}}^\alpha(x) \\ &= \frac{\mathbb{P}(x \in D, \text{parameter is } \alpha)}{\mathbb{P}(\text{parameter is } \alpha)} \\ &= \frac{\rho_\mu(x) \int_{x_2, \dots, x_m} \prod_{i=2}^m \rho_\mu(x_i) \mathbb{P}(\alpha | x, x_2, \dots, x_m) dx_2 \dots dx_m}{\int_{x_1, \dots, x_m} \prod_{i=1}^m \rho_\mu(x_i) \mathbb{P}(\alpha | x_1, \dots, x_m) dx_1 \dots dx_m} \\ &= \rho_{g_\alpha}(x) \end{aligned}$$

□

This proposition validates prior membership inference attacks that utilize approximations of $\rho_{g_\alpha}(x)$ to make decisions (Chen et al., 2019; Hayes et al., 2019; Hilprecht et al., 2019).

With this proposition, the problem becomes clear: for a given sample x , the attacker needs to decide whether the sample comes from the training set (i.e., from g_α) or not (i.e., from μ). In the following theorem, we outer bound the ROC region for this hypothesis test, which relates the true positive (TP) rate to the false positive (FP) rate.

Proposition 3. *Consider a generative model g_α and a real distribution μ . Define $r \triangleq d_{TV}(g_\alpha, \mu)$ as the total variation (TV) distance between the two distributions.*

Define function $f : [0, 1] \rightarrow [0, 1]$ as

$$f(x) = \begin{cases} x + r & (0 \leq x \leq 1 - r) \\ 1 & (r < x \leq 1) \end{cases},$$

Then we have that for any membership inference attack policy \mathcal{A} , the ROC curve $g_{\mathcal{A}} : [0, 1] \rightarrow [0, 1]$ (mapping FP to TP) satisfies $g(x) \leq f(x), \forall 0 \leq x \leq 1$, i.e., the ROC curve is upper bounded by f . Also, this bound is tight, i.e., there exists two distributions μ', g' such that $d_{\text{TV}}(\mu', g') = r$ and the ROC curve is exactly f at every point.

(Proof in Section 4.1) As a result, we can directly bound the area under the ROC curve (AUC) as a function of the total variation distance.

Corollary 1 (Bound on the AUC for generative models). *For any attack policy on a generative model, we have*

$$\text{AUC} \leq -\frac{1}{2}d_{\text{TV}}(g_{\alpha}, \mu)^2 + d_{\text{TV}}(g_{\alpha}, \mu) + \frac{1}{2}.$$

Note that Proposition 3 and Corollary 1 hold for any generative model. For GANs in particular, we can use generalization bounds in Lemma 1 to obtain the following result.

Theorem 2. *Consider a GAN model g_{α} and a real distribution μ . Define*

$$\Xi_{\mathcal{F}, \mathcal{G}, \mu} \triangleq \sup_{\nu \in \mathcal{G}} \|\log(\rho_{\mu}/\rho_{\nu})\|_{\mathcal{F}, 1}$$

and

$$\epsilon_{\text{TV}}(m, \delta) \triangleq \frac{\sqrt{\Xi_{\mathcal{F}, \mathcal{G}, \mu} \cdot \tau_{m, \delta}}}{2\sqrt{2}}, \quad (7)$$

where $\tau_{m, \delta}$ is defined as in Eq. (2). Define function $f : [0, 1] \rightarrow [0, 1]$ as

$$f(x) = \begin{cases} x + \epsilon_{\text{TV}}(m, \delta) & (0 \leq x \leq 1 - \epsilon_{\text{TV}}(m, \delta)) \\ 1 & (\epsilon_{\text{TV}}(m, \delta) < x \leq 1) \end{cases}.$$

Then we have that for any membership inference attack policy \mathcal{A} , the ROC curve $g_{\mathcal{A}} : [0, 1] \rightarrow [0, 1]$ satisfies $g(x) \leq f(x), \forall 0 \leq x \leq 1$, and the bound is tight.

One complication is that existing generalization bounds do not directly bound TV distance, so these must be extended. The proof can be found in Section 4.2, and directly gives the following corollary bounding the AUC for GANs.

Corollary 2 (Bound on AUC for GANs). *For any attack policy on GANs, we have that with probability at least $1 - \delta$ w.r.t. the randomness of training samples,*

$$\text{AUC} \leq -\frac{1}{2}\epsilon_{\text{TV}}(m, \delta)^2 + \epsilon_{\text{TV}}(m, \delta) + \frac{1}{2},$$

where $\epsilon_{\text{TV}}(m, \delta)$ is defined as in Eq. (7).

Note that the AUC bound decays as $O(m^{-1/4})$.




Figure 1: The upper bound of ROC curves.

Discussions These results confirm the prior empirical observation that GANs are more robust to membership inference attacks when the number of training samples grows (Lin et al., 2020a; Chen et al., 2019). However, the results heavily rely on the assumption of generator parameter distribution Eq. (6), which was introduced in (Sablayrolles et al., 2019). It is unlikely to strictly hold in practice. Extending the results to more general settings would be an interesting future direction.

4.1 Proof of Proposition 3

It is known from the hypothesis testing literature (Kairouz et al., 2015) that, for any attack policy, the difference between the true positive rate (TP) and false positive rate (FP) is upper bounded by the total variation (TV) distance $d_{\text{TV}}(g_{\alpha}, \mu)$:

$$\text{TP} \leq \text{FP} + \min\{d_{\text{TV}}(g_{\alpha}, \mu), 1 - \text{FP}\}. \quad (8)$$

Note that total variation distance and ROC curve has a very simple geometric relationship, as noted in Lin et al. (2018) (Remark 7). That is, the total variation distance between g_{α} and μ is the intersection between the vertical axis and the tangent line to the upper boundary of the ROC curve that has slope 1 (e.g., see Fig. 1). This immediately implies that $f(x)$ is an upper bound for all possible ROC curves.

To show tightness, we can construct a g' and μ' as shown in Fig. 2, such that $d_{\text{TV}}(\mu', g') = r$ and they achieve the ROC curve in Fig. 1.

4.2 Proof of Theorem 2

We begin by showing that, under the assumptions in Lemma 1 and assuming that $\forall \nu \in \mathcal{G}, \log(\rho_{\mu}/\rho_{\nu}) \in \text{span}\mathcal{F}$, we have that with probability at least $1 - \delta$ w.r.t. the randomness of training samples,

$$d_{\text{TV}}(g_{\alpha}, \mu) \leq \epsilon_{\text{TV}}(m, \delta) \triangleq \frac{\sqrt{\Xi_{\mathcal{F}, \mathcal{G}, \mu} \cdot \tau_{m, \delta}}}{2\sqrt{2}}. \quad (9)$$




Figure 2: The pair of distributions that achieve the ROC upper bound.

To show this, note that Lemma 1 gives an upper bound on the integral probability metric between the real and generated distribution. We first connect this distance to the KL divergence with the following lemma.

Lemma 8. *Denote the real distribution as μ . Given a generator set \mathcal{G} and a discriminator set \mathcal{F} which satisfy $\forall \nu \in \mathcal{G}, \log(\rho_\mu/\rho_\nu) \in \text{span}\mathcal{F}$, then we have $\forall g_\alpha \in \mathcal{G}$*

$$d_{\text{KL}}(g_\alpha, \mu) + d_{\text{KL}}(\mu, g_\alpha) \leq \Xi_{\mathcal{F}, \mathcal{G}, \mu} d_{\mathcal{F}}(\mu, g_\alpha)$$

where $\Xi_{\mathcal{F}, \mathcal{G}, \mu} \triangleq \sup_{\nu \in \mathcal{G}} \|\log(\rho_\mu/\rho_\nu)\|_{\mathcal{F}, 1}$.

Similar to Lemma 4, this lemma relies on Proposition 2.9 in Zhang et al. (2017). Furthermore, we use Pinsker’s inequality (Tsybakov, 2008) to upper bound the TV distance by the KL distance. Pinsker’s inequality says that $d_{\text{TV}}(a, b) \leq \sqrt{\frac{1}{2}d_{\text{KL}}(a, b)}$ for any two distributions a, b . Therefore, we have

$$\Xi_{\mathcal{F}, \mathcal{G}, \mu} d_{\mathcal{F}}(\mu, g_\alpha) \geq 4 \cdot d_{\text{TV}}(g_\alpha, \mu)^2$$

Combing this equation with Lemma 1 we get the desired inequality.

We can use Eq. (9) to upper bound Eq. (8) as

$$\text{TP} \leq \text{FP} + \min\{\epsilon_{\text{TV}}(m, \delta), 1 - \text{FP}\}. \quad (10)$$

Combining Eq. (10) with Proposition 3 gives the result.

5 Discussion

In this work, we show that GAN-generated samples naturally exhibit a (weak) differential privacy guarantee as well as protection against membership inference attacks. We provide bounds on the privacy risk of each of these attacks. However, as discussed in Section 3, the inherent differential privacy guarantee in

GANs is weak. This suggests that differentially-private training techniques are required to ensure meaningful differential privacy guarantees in practice. Given that current techniques like DP-SGD (Abadi et al., 2016) and PATE (Jordon et al., 2018) sacrifice fidelity in exchange for privacy (as discussed in Section 1), there is a need to develop new techniques that achieve a better privacy-fidelity tradeoff. As noted in Section 1, one approach for achieving this goal may be to ensure the differential privacy for releasing generated samples (rather than enforcing the differential privacy for releasing parameters).

Limitations and future work

(1) Since our results build on existing generalization bounds for GANs (Zhang et al., 2017), we inherit their assumptions (e.g., on the classes of generators and discriminators). Some of these assumptions do not apply to all GAN architectures. Extending the results to more general settings would be an interesting direction, potentially with stronger generalization bounds.

(2) As discussed in Section 4, the results on membership inference attacks rely on an assumption from (Sablayrolles et al., 2019) regarding the generator parameter distribution; this assumption is unlikely to hold in practice. Relaxing this assumption is an interesting direction for future work.

(3) The bounds we give for both privacy notions depend on unknown constants like optimization and approximation errors. Numerically quantifying these bounds in practice remains a challenging and interesting direction.

Acknowledgments

This work was supported in part by National Science Foundation grants CIF-1705007 and CA-2040675. This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-21-1-0090. This research was sponsored by the Combat Capabilities Development Command Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Combat Capabilities Development Command Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes not withstanding any copyright notation here on. The authors would also like to acknowledge the generous support of JP Morgan Chase, Siemens AG, Google, and the Sloan Foundation.

References

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318, 2016.
- Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. *arXiv preprint arXiv:1701.04862*, 2017.
- Yu Bai, Tengyu Ma, and Andrej Risteski. Approximability of discriminators implies diversity in gans. *arXiv preprint arXiv:1806.10586*, 2018.
- Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. *arXiv preprint arXiv:1807.01647*, 2018.
- Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 1046–1059, 2016.
- Jens Behrmann, Will Grathwohl, Ricky TQ Chen, David Duvenaud, and Jörn-Henrik Jacobsen. Invertible residual networks. In *International Conference on Machine Learning*, pages 573–582. PMLR, 2019.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- Dingfan Chen, Ning Yu, Yang Zhang, and Mario Fritz. Gan-leaks: A taxonomy of membership inference attacks against gans. *arXiv preprint arXiv:1909.03935*, 2019.
- Dingfan Chen, Tribhuvanesh Orekondy, and Mario Fritz. Gs-wgan: A gradient-sanitized approach for learning differentially private generators. *arXiv preprint arXiv:2006.08265*, 2020.
- Edward Choi, Siddharth Biswal, Bradley Malin, Jon Duke, Walter F Stewart, and Jimeng Sun. Generating multi-label discrete patient records using generative adversarial networks. *arXiv preprint arXiv:1703.06490*, 2017.
- Rachel Cummings, Katrina Ligett, Kobbi Nissim, Aaron Roth, and Zhiwei Steven Wu. Adaptive learning with robust generalization guarantees. In *Conference on Learning Theory*, pages 772–814, 2016.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Leon Roth. Preserving statistical validity in adaptive data analysis. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 117–126, 2015.
- Cristóbal Esteban, Stephanie L Hyland, and Gunnar Rätsch. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633*, 2017.
- Farhad Farokhi and Mohamed Ali Kaafar. Modelling and quantifying membership information leakage in machine learning. *arXiv preprint arXiv:2001.10648*, 2020.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017.
- Jamie Hayes, Luca Melis, George Danezis, and Emiliano De Cristofaro. Logan: Membership inference attacks against generative models. *Proceedings on Privacy Enhancing Technologies*, 2019(1):133–152, 2019.
- Benjamin Hilprecht, Martin Härterich, and Daniel Bernau. Monte carlo and reconstruction membership inference attacks against generative models. *Proceedings on Privacy Enhancing Technologies*, 2019(4):232–249, 2019.
- James Jordon, Jinsung Yoon, and Mihaela van der Schaar. Pate-gan: Generating synthetic data with differential privacy guarantees. In *International Conference on Learning Representations*, 2018.
- Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *International conference on machine learning*, pages 1376–1385. PMLR, 2015.

- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Zheng Li and Yang Zhang. Label-leaks: Membership inference attack with label. *arXiv preprint arXiv:2007.15528*, 2020.
- Zinan Lin, Ashish Khetan, Giulia Fanti, and Sewoong Oh. Pacgan: The power of two samples in generative adversarial networks. In *Advances in neural information processing systems*, pages 1498–1507, 2018.
- Zinan Lin, Alankar Jain, Chen Wang, Giulia Fanti, and Vyas Sekar. Using gans for sharing networked time series data: Challenges, initial promise, and open questions. In *Proceedings of the ACM Internet Measurement Conference*, pages 464–483, 2020a.
- Zinan Lin, Vyas Sekar, and Giulia Fanti. Why spectral normalization stabilizes gans: Analysis and improvements. *arXiv e-prints*, pages arXiv–2009, 2020b.
- Yunhui Long, Vincent Bindschaedler, Lei Wang, Diyu Bu, Xiaofeng Wang, Haixu Tang, Carl A Gunter, and Kai Chen. Understanding membership inferences on well-generalized learning models. *arXiv preprint arXiv:1802.04889*, 2018.
- Yunhui Long, Suxin Lin, Zhuolin Yang, Carl A Gunter, and Bo Li. Scalable differentially private generative student model via pate. *arXiv preprint arXiv:1906.09338*, 2019.
- Sebastian Meiser. Approximate and probabilistic differential privacy definitions. *IACR Cryptology ePrint Archive*, 2018:277, 2018.
- Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 691–706. IEEE, 2019.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443, 1997.
- Kobbi Nissim and Uri Stemmer. On the generalization properties of differential privacy. *arXiv*, pages arXiv–1504, 2015.
- Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems*, pages 271–279, 2016.
- Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016.
- Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, Yann Ollivier, and Hervé Jégou. White-box vs black-box: Bayes optimal strategies for membership inference. In *International Conference on Machine Learning*, pages 5558–5567, 2019.
- Ahmed Salem, Yang Zhang, Mathias Humbert, Pascal Berrang, Mario Fritz, and Michael Backes. MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models. *arXiv preprint arXiv:1806.01246*, 2018.
- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18. IEEE, 2017.
- Alexandre B Tsybakov. *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- Yu-Xiang Wang, Jing Lei, and Stephen E Fienberg. Learning with differential privacy: Stability, learnability and the sufficiency and necessity of erm principle. *The Journal of Machine Learning Research*, 17(1):6353–6392, 2016.
- Ryan Webster, Julien Rabin, Loïc Simon, and Frédéric Jurie. Detecting overfitting of deep generative networks via latent recovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11273–11282, 2019.
- Bingzhe Wu, Shiwan Zhao, Chaochao Chen, Haoyang Xu, Li Wang, Xiaolu Zhang, Guangyu Sun, and Jun Zhou. Generalization in generative adversarial networks: A novel perspective from privacy protection. In *Advances in Neural Information Processing Systems*, pages 307–317, 2019.
- Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, pages 268–282. IEEE, 2018.
- Pengchuan Zhang, Qiang Liu, Dengyong Zhou, Tao Xu, and Xiaodong He. On the discrimination-generalization tradeoff in gans. *arXiv preprint arXiv:1711.02771*, 2017.