This Person (Probably) Exists. Identity Membership Attacks Against GAN Generated Faces.

Ryan Webster[1], Julien Rabin[1], Loic Simon[1,2], and Frederic Jurie[1]

[1] University of Caen Normandie
[2] ENSI Caen
{firstname.lastname}@unicaen.fr

Abstract. Recently, generative adversarial networks (GANs) have achieved stunning realism, fooling even human observers. Indeed, the popular tongue-in-cheek website http://thispersondoesnotexist.com, taunts users with GAN generated images that seem too real to believe. On the other hand, GANs do leak information about their training data, as evidenced by membership attacks recently demonstrated in the literature. In this work, we challenge the assumption that GAN faces really are novel creations, by constructing a successful membership attack of a new kind. Unlike previous works, our attack can accurately discern samples sharing the same identity as training samples without being the same samples. We demonstrate the interest of our attack across several popular face datasets and GAN training procedures. Notably, we show that even in the presence of significant dataset diversity, an over represented person can pose a privacy concern.

1 Introduction

Among the recent advances of deep learning, Generative Adversarial Networks (GANs)[1] are arguably the most burgeoning area of research. This is evidenced by their rich applications such as ageing, inpainting, super-resolution, or attribute modification [2–5]. Less studied lines of research include measuring the statistical consistency between the generated distribution and the real one [6–8], or in theoretical analysis of the GAN learning problem [9].

It is fairly clear GANs have made incredible progress in terms of visual quality, as measured by the popular Frechet Inception Distance (FID) [6], or merely by sample inspection (for example, see StyleGAN samples in Fig. 1). However, it remains unclear in what ways GANs generalize. Several works have noted that the GAN *discriminator* tends to overfit the training set [11, 12], in the sense that it will label hold out test images as fake. Such an observation was used to exploit the privacy of GANs in the LOGAN approach [13]. Nonetheless, the LOGAN leak challenges mainly the discriminator, leaving the case of the generator in a somewhat gray area. In fact, several heuristics exist to show GANs do in fact produce novel data, whether it be by sample interpolation and attribute modification as demonstrated in the first GAN works [1, 14], or by the ability of GANs

Samples from a GAN G trained with a dataset $\mathcal{D}_G$ composed of $\mathcal{Y}_G$ distinct identities

(a) $|\mathcal{Y}_G| = 220$ (b) $|\mathcal{Y}_G| = 440$ (c) $|\mathcal{Y}_G| = 2020$

$G(z)$ $x \in \mathcal{D}_G$ $G(z)$ $x \in \mathcal{D}_G$ $G(z)$ $x \in \mathcal{D}_G$
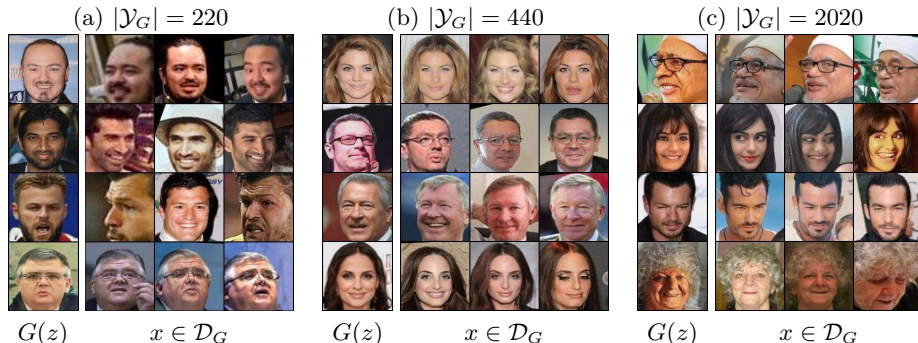
Fig. 1. These people appear to exist. *Each row displays a GAN generated image (left) with three training images (right) having the same predicted identity. Images are generated with StyleGAN [10] using N training images (respectively 40k, 80k and 46k) from VGGFace2 and identified with a face identification network. We investigate two different scenarios in this paper: in (a) and (b) identities are evenly distributed over the datasets, where in (c) a small subset is more represented. While some samples merely bear resemblance, other generated images strongly share idiosyncratic features of training identities. Such a nearest neighbor search helps factor out the ways in which GANs can generalize (via pose, lighting and expression) and elucidate overfitting on identities. This is a threat to privacy as we demonstrate in our blind identity membership attack.*

to generalize to novel pose and expression [15]. Furthermore, a general belief is that the generated images are always entirely original (in terms of identity), as can be testified by the popular website https://thispersondoesnotexist.com/.

The purpose of this article is to challenge the afore-mentioned common belief. We are especially interested in evaluating to what extent a GAN distribution may remain evocative of the training samples. Alongside the generalization aspect, this question is fundamental to evaluate the potential privacy risk posed by GANs. We would like to highlight that obtaining answers is becoming paramount for several reasons. First, recent attempts to apply differential privacy mechanisms on GANs [16,17] have not been so successful. Indeed, to our knowledge, published works lead to an unsatisfactory tradeoff between the privacy guarantees and the quality of the learnt distribution. Besides, following the work of [4] showing some form of GAN generalization, some attempts were made to leverage GANs to produce ersatz datasets aiming at more privacy. For instance in [18], a dataset, which is kept private, is used to train a GAN. Then the GAN is made publicly available for downstream tasks. Again, the rationale behind such a strategy is that a synthetic dataset produced from the GAN can provides the means to achieve the downstream task allegedly without compromising the privacy of the original dataset.

Notwithstanding, neural networks leak information in numerous ways about the data on which they are trained [19, 20]. When no privacy mechanism is used, the most common way to expose potential privacy leaks from a neural network

consists in applying one of several off-the-shelf attacks. Amongst the most well known attacks are the *membership inference attacks* [21]. In this type of attack, an attacker tries to discern which samples were used during model training. In [21], training set samples could be nearly perfectly determined from model outputs on the MNIST classification task. Such an attack typically utilizes statistics of the model output. For example, a model that overfits will have low values of the loss function on the training samples, thus allowing for a simple thresholding attack. Similar attacks can be seen in [22–24]. More recently, membership attacks have been devised against generative models [25, 13]. These attacks are slightly more intuitive, in the sense that generated samples leave in the same space as the data, and they may be visually similar or identical to exact training samples. In [4], the non-adversarial procedure of GLO [26] reproduces some training samples verbatim, while not being able to reproduce test samples coming from the same distribution. Such direct leaks were not observed for typical GANs. However, we argue that generated images can resemble training samples in even more subtle ways, as seen in Fig. 1: generated faces can highly resemble exact identities, even if the images themselves are quite different.

Identity Membership Attacks The main contribution of this paper is to define a refined attack objective and implement such an attack. This new objective is meaningful in situation where a clear notion of identity exists. This is the case for example for datasets of face images (the central case considered in here). But such a notion make sense in many other contexts such as dataset of paintings (think of the identity of the painter), bio-metric data, or medical images. In our scenario, the attacker would have access to a collection of query samples, but instead of trying to discern samples that were truly part of the training set of a GAN, the attacker should rather determine if samples with the same identity were used. An attack of this sort shall be referred henceforth as an *identity membership attack*. In addition to this fundamental difference with currently published attacks, we shall design an attack with the following properties:

- contrary to LOGAN [13], the attacker can only exploit the generator,
- the attacker should not be aware of the ratio of query identities that were part of the training set (i.e, a *blind* attack).

We hold that if successful, such an attack would reveal as a serious hurdle for the safe exchange of GANs in sensitive contexts. For instance, in the context of paintings or other art pieces, distributing a non-private generator might well be ruled-out for obvious copyright issues. More importantly, consider a biometric company A releasing a generator exposing its consumer identity. Another company B could potentially detect which of their own consumers are also clients of company A. Similar situations can pose serious issues for medical data, where revealing a GAN could breach personal information about a patient disease.

Face Generation In this work, we focus on GANs trained on facial image data. Face image datasets are amongst the most popular for demonstrating GAN efficacy, partly because the human eye is particularly attuned to detect artifacts

in these images and due to the wide availability of high quality face datasets [27–29]. Face datasets include some form of bias in regard to the real distributions of faces. For instance, datasets are frequently composed of a decent number of instances of the same person (see Fig. 2). Besides in some datasets [30], the number of instances per identity is highly versatile (e.g. in MEGAFACE [27] it ranges from 3 to 2k). One should also note that datasets dedicated to generation such as FFHQ [10] may well display similar biases due to the lack of any specific safeguard.
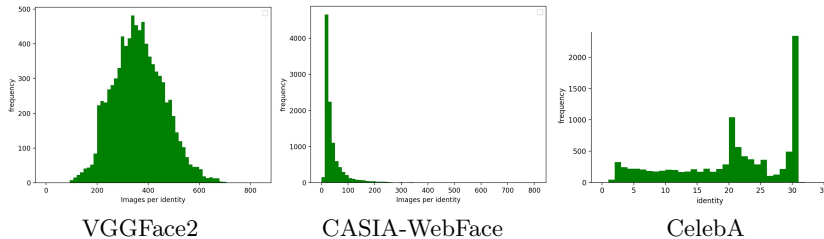

VGGFace2   CASIA-WebFace   CelebA

Fig. 2. Histogram of number of instances per identity for different datasets (VGGFace2, CASIA-WebFace & CelebA). *Far from being uniform, face recognition datasets include bias in terms of the number of examples per identitiy. VGGFace2 was designed to be balanced in this sense, compared here to the CASIA-WebFace containing bias in some samples.*

Outline The rest of the paper is organized as follows. In Section 2 we expose the proposed protocol used to train the face detection network, to perform the attack and to evaluate its performance. This protocol is tested in Section 3 on various datasets (CASIA-WebFace and VGGFace2) with different GANs (Style-GAN and LSGAN) across different attack scenarios. Performance curves and visual evidence are provided to show membership attacks can be highly successful against GANs even in the blind scenario, and when diverse training data is used. A discussion and perspectives on future work are given in Section 4. Reviewers will also find additional results in the supplementary material.

2   Identity Membership Attack

As argued earlier, because standard datasets gather several instances of some individuals, training generative models on such datasets exposes those individuals to privacy leaks. Figure 1 illustrates such a leak for a GAN trained on a subset of VGGFace2. There is no doubt that in many cases, the displayed generative sample is but an instance of the training identity shown under various poses (last three columns). As such the generated samples can be leveraged to extract information on the identities that were seen during the training of the GAN.
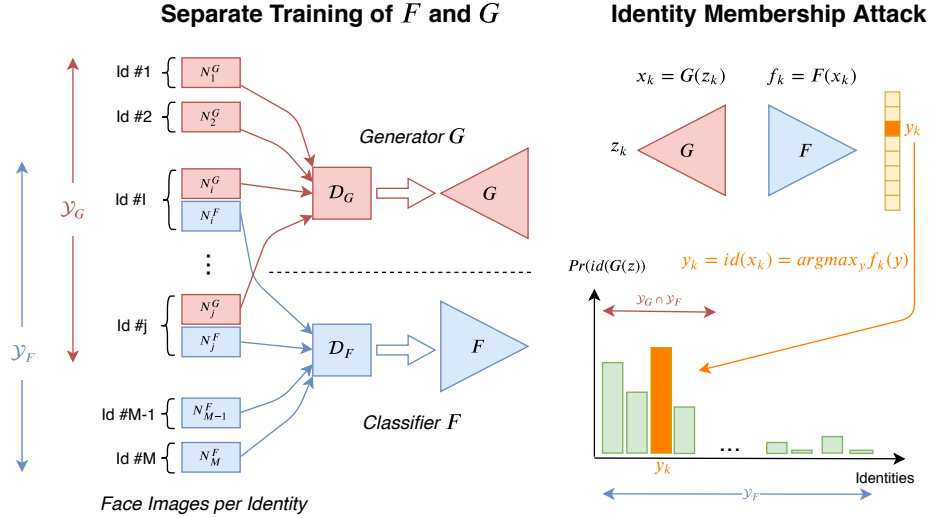
Separate Training of F and G   **Identity Membership Attack**

Fig. 3. Illustration of the protocol used for the attack. A generator G is trained from a dataset $\mathcal{D}_G$ gathering image instances of identities $y \in \mathcal{Y}_G$ (red samples). A face classifier F is trained by an attacker to recognize identities $y \in \mathcal{Y}_F$ from a separate dataset $\mathcal{D}_F$ (blue samples). Although the samples are completely separate, the two datasets share some common identities $\mathcal{Y}_G \cap \mathcal{Y}_F \neq \emptyset$. The distribution of the identities of the faces generated by G, according to F, can be used to detect overfitting in G, and thus to infer identities seen during the training of the generator.

2.1 Attack Assumptions

We consider the following attack scenario : the attacker wants to determine if instances of certain query individuals were used in the training of the GAN. We reference this type of attack as an *identity membership* attack. More formally, our attack protocol is illustrated in Fig. 3 and delineated as follows:

- A GAN G is trained on a dataset $\mathcal{D}_G := \{(x_i^G, y_i^G)\}$ where $\forall i, y_i^G \in \mathcal{Y}_G$;
- The attacker trains a face identification network, F on a dataset $\mathcal{D}_F = \{(x_i^F, y_i^F)\}$ where $\forall i, y_i^F \in \mathcal{Y}_F$ to recognize instances from identities $y \in \mathcal{Y}_F$;
- It is assumed that $\mathcal{Y}_G \cap \mathcal{Y}_F \neq \emptyset^3$ but $\mathcal{D}_F \cap \mathcal{D}_G = \emptyset$. The first assumption reflects the fact that the attacker has some founded suspicion that some identities were used in the training of G. Nonetheless the second assumption ensures that no instance of $\mathcal{D}_G$ is explicitely known to the attacker.

2.2 Attack Algorithm

The basic principle of the attack consists in randomly generating faces $x_k := G(z_k)$ for $k \in \{1, \dots, N\}$. Then using the network F, these random samples are

---

^{3} In practice, we shall consider a worse case situation where $\mathcal{Y}_G \subsetneq \mathcal{Y}_F$.

identified to identities $y_k := id(x_k)$ where $id(x) := \arg\max_{y \in \mathcal{Y}_F} F_y(x)$. Eventually, the attacker will suspect the identities of $\mathcal{Y}_F$ that are more often predicted as the ones that were seen during the training of G. This mechanism is described in more details in Alg. 1 and merely corresponds to thresholding the number of times the query identity y was predicted by F on the generated samples $x_k = G(z_k)$.

To be effective, such an algorithm requires to generate a large enough set of samples K. Typically this parameter shall be set to many times the number of identities in $\mathcal{D}_F$ i.e. $K = \lambda \times |\mathcal{Y}_F|$. In such case, it is natural to fix the frequency threshold T to λ. We will denote this natural value of the threshold $T_0 = \lambda$. Yet, of course, one can trade recall for precision by increasing this threshold. In our experiments, we shall consider also $T_1 = 10\lambda$.

Inputs: the query identity $y \in \mathcal{Y}_F$,
            the number of generated samples K,
            the frequency threshold T
Output: a boolean prediction of $\mathbb{1}_{y \in \mathcal{Y}_G}$
Algorithm `identityAttack(`y`)`
1    $k_y = 0$
2   **for** $k = 0$ *to* K **do**
3       $z = $ `randomSample()`
4       $x = G(z)$
5   **if** $id(x) = y$ **then**
6         $k_y += 1$
      end
   end
7   **return** $\mathbb{1}_{k_y \geq T}$

Algorithm 1: *Proposed identity membership attack*

2.3   Classifier Training

To train the classifier F, we use pre-trained features from the VGGFace network [31]. Note that VGGFace and VGGFace2 share 53 identities. Therefore, to ensure that $\mathcal{D}_F \cap \mathcal{D}_G = \emptyset$, *i.e.* F is not trained on any of the sample samples as G, we remove these 53 identities from the VGGFace2 dataset. Then, we pool the relu5_3 layer of VGGFace to a 2×2 spatial resolution. Finally, we train a single fully connected layer with $|\mathcal{Y}_F|$ output classes. Note that $|\mathcal{Y}_F| = 8631$ for VGGFace2 and $|\mathcal{Y}_F| = 1292$ for CASIA-Webface. Finally, models are trained in pytorch with SGD, a learning rate of .1, and a momentum .9. This simple model for F achieves an average top-1 classification accuracy of 86.1% on VGGFace2 and 94.7% on CASIA-Webface. Likely, using even better classifiers should improve the membership attack accuracy, but we leave this direction for future work.

2.4 Attack Evaluation

The attack can be seen as a binary classification problem where the attacker must classify identities from $\mathcal{Y}_F$ according to whether they also belong to $\mathcal{Y}_G$ or not. It is therefore natural to evaluate the performance of Alg. 1 in terms of precision and recall. Taking the attacker point of view, we will tag as positive all the identities in $\mathcal{Y}_G \cap \mathcal{Y}_F$ and the remaining ones as negative. Therefore, denoting $C(y) := \mathbb{1}_{k_y \geq T}$ the decision made by the attacker, the precision and recall will be computed as follows:

$$\alpha := \frac{\sum_{y \in \mathcal{Y}_G \cap \mathcal{Y}_F} C(y)}{\sum_{y \in \mathcal{Y}_F} C(y)} \quad \text{and} \quad \beta := \frac{\sum_{y \in \mathcal{Y}_G \cap \mathcal{Y}_F} C(y)}{|\mathcal{Y}_G \cap \mathcal{Y}_F|} \tag{1}$$

As the goal of privacy is to protect all individuals equally, an attack with high precision (and positive be it small recall) should be considered more troublesome than one with high recall. In other words, if an attack can accurately discern training information even on a small subset of data, it still violates the fundamental goal of privacy. To summarize precision recall in the following section, we'll highlight attacks with the highest F_1 score, which is the harmonic average of precision and recall as follows

$$F_1(\alpha, \beta) := 2 \cdot \frac{\alpha \cdot \beta}{\alpha + \beta} \tag{2}$$

2.5 Visual Evaluation

As a visual sanity check, we include a nearest neighbor search utilizing the recognition network F. An image or feature space nearest neighbor search is a common procedure in the GAN literature, to dispel suspicions of overfitting [1, 12, 10]. We visually compare generated samples falling inside the GAN training set $y \in \mathcal{Y}_G$, and those detected to lie outside $y \notin \mathcal{Y}_G$. More formally, we generate an image $x := G(z)$, find the identity using F as $y := id(x)$ and finally perform the intra-identity nearest neighbor search as

$$x_{\mathrm{NN}} = \arg\min_{x_{y_i}} \|f(x_{y_i}) - f(x)\|_2^2 \tag{3}$$

where x_{NN} is the retrieved nearest neighbor, f is the before softmax feature representation of F and finally x_{y_i} represents the i-th image among the instances of identity y. Results are displayed for three attack scenarios in Fig. 1. Indeed, in both the scenarios of low diversity and dataset bias, generated images highly resemble those in the training set. Intuitively, these also correspond with attack scenarios that are highly accurate at guessing training identities (see Table 1 and Table 2). Note that this visual demonstration should be robust to any distortions or variations found in generated images, as F is specifically designed to be robust to natural variations in the dataset like pose or expression.

3   Identity Membership Attack Evaluation

3.1   Experimental Protocol

We evaluate the efficacy of our identity membership attack on a variety of training settings, in order to elucidate what factors in training data influence privacy. As is shown in Figure 2, face datasets exhibit varying number of identities, and importantly bias in the number of sample they have per identity. VGGFace2 contains a relatively balanced number of samples per identity, while VGGFace2 has far more average samples per identity than CelebA for example. On the other hand, in CASIA-WebFace the number of samples per identity ranges from a few to many. In the following training settings, we denote *diversity* as number of training identities and *(imbalance) bias* when some identities have many more samples than others. As both G and F require some minimum amount of samples to train, and with our disjointness assumption for blind attacks (see Sec. 2.1), we also curate CASIA-WebFace to contain at least 80 samples per identity. We then explore two real world training scenarios in the following sections.

Throughout experiments, the parameter λ discussed in Sec. 2.2 used to sample the Generator during attack is fixed to $\lambda = 2$.

3.2   Setting 1: Low bias and varying diversity

In this setting, we simply vary the number of identities $|\mathcal{Y}_G|$ ranging from 30 to 880, taking a large set of samples from the the first $|\mathcal{Y}_F|$ identities. The number of samples per identity N_i^G used during the training of the generator is fixed to simulate datasets that are evenly distributed. In addition, we consider two successful GAN methods, the Least Squares GAN (LSGAN) [32] and the state of the art StyleGAN [10] network.

Table 1 gives the precision and recall rate for each setting, on two different datasets (VGGFace2 and CASIA-WebFace). The total number of samples used during training is indicated by $N = \sum_i N_i^G$, and is increasing with the number of identities. Each column corresponds to a different generation training set, with a varying number of identities while the total number of identities is kept fixed ($|\mathcal{Y}_F| \approx 8K$). F_1 scores are used to highlight the most successful attacks. The baseline corresponding to random face identification has a precision of $\frac{|\mathcal{Y}_G|}{|\mathcal{Y}_F|}$ at any recall (witch is set to 100% to compute F_1 scores. As mentioned earlier, two different frequency thresholds T_0 and T_1 are proposed to perform the attack, T_0 being more conservative towards recall, while T_1 is more conservative with respect to precision. Figure 4 displays Precision-Recall curves obtained for some settings when spanning the frequency threshold.

3.3   Setting 2: Varying bias and high diversity

We also examine scenario when training data is highly diverse, yet due to natural bias in the data collection procedure or data availability, certain identities contain many more images than others. For this setting, we vary the number

of biased identities but enrich each dataset with a large set of diverse unbiased faces. More formally, we take $\mathcal{Y}_G := \mathcal{Y}_{G_1} \cup \mathcal{Y}_{G_2}$, with $\mathcal{Y}_{G_1}$ containing many more samples per identity than $\mathcal{Y}_{G_2}$. Furthermore, we consider when $|\mathcal{Y}_{G_2}| \gg |\mathcal{Y}_{G_1}|$, *i.e.* the dataset has high class diversity. Such is naturally the case in CASIA-WebFace , see Fig. 2. One may however suspect that the GAN be prone to reproduce mainly identities from $\mathcal{Y}_{G_1}$ (since they are more present in the dataset). We therefore only compute the precision/recall in Eq. (1) and F_1 values in Eq. (2) for the biased set $\mathcal{Y}_{G_1}$ in lieu of $\mathcal{Y}_G$.

Table 2 gives the precision and recall rates of the attack in such setting on VGGFace2 dataset. Attack are performed on LSGAN and StyleGAN trained with an increasing number of identities for subset $\mathcal{Y}_{G_1}$, while the second set $\mathcal{Y}_{G_2}$ is fixed to 40'000 samples evenly distributed across 2000 identities. The bias is therefore decreasing, as measured by the ratio $\frac{|\mathcal{Y}_{G_1}|}{|\mathcal{Y}_F|}$ corresponding to the precision of random identification, as reported in the Table. The total number of samples used during training is indicated by $N = N_1 + N_2$, and is increasing with the number of identities. Figure 5 displays Precision-Recall curves obtained for some settings when spanning the frequency threshold.

3.4   Analysis

Figure 6 shows the distribution of identities detected by the face identification network F based on the number of samples recovered from StyleGAN trained on the VGGFace2 dataset, for different settings. Identities from $\mathcal{Y}_F \backslash \mathcal{Y}_G$ and $\mathcal{Y}_F$ are displayed in different colors to highlight the fact that the tail of the distributions corresponds to frequent recovered identities that are indeed in the training set $\mathcal{Y}_G$, and prone to membership attack. Observe that the distribution has a lighter tail when the number of identities $|\mathcal{Y}_G|$ and the number of samples N increase (first scenario), or when the bias ratio $\frac{|\mathcal{Y}_{G_1}|}{|\mathcal{Y}_F|}$ decreases (second scenario). As discussed later in Section 2.5, samples from these identities can be extracted to visually assess overfitting, as done in Fig. 1.

In Table 1, it is clear that insufficient diversity yields GANs models that are susceptible to membership attacks, even in the blind setting. With enough diversity (here, towards 880 identities), attacks are reduced to near guessing. Table 2 and Figure 5 show that diversity itself is not enough to protect the data in the presence of bias. Even after taking more than 2000 identities, attacks still succeed with high F_1 scores. As for the difference across GAN models, LSGAN seems to be far more susceptible to attacks until a certain data set size (in Table 1, at $\mathcal{Y}_G = 111$ for VGGFace2), where attack success significantly drops off. We assume that this is because the LSGAN model we trained had nearly 10x less parameters than the StyleGAN model and likely could not overfit samples after some amount of diversity. On the other hand, StyleGAN models are far higher quality. We leave investigation of the trade off between model size, quality and susceptibility to attack to future work.

VGGFace2

$|\mathcal{Y}_G|$ (N)	30 (10k)	58 (20k)	111 (40k)	220 (80k)	440 (160k)	880 (320k)
Random	0.35/100	0.67/100	1.29/100	2.55/100	5.10/100	10.2/100
StyleGAN — Threshold T_0	1.79/90	3.08/77.6	5.36/73.9	7.56/58.6	**9.25/38.9**	**12.5/30.6**
StyleGAN — Threshold T_1	**24.6/56.7**	**40.5/51.7**	**69.5/36.9**	**75.5/16.8**	42.9/2.73	14.3/0.341
LSGAN — Threshold T_0	2.51/100	4.57/96.6	**2.19/39.6**	**3.37/28.6**	**6.73/32.3**	11.9/28.5
LSGAN — Threshold T_1	**49.1/93.3**	**80.3/84.5**	8.57/2.70	7.69/1.82	16.7/0.68	10.0/0.34

CASIA-WebFace

$|\mathcal{Y}_G|$ (N)	30 (4.5k)	58 (8.7k)	111 (16.65k)	220 (33k)
Random	2.32/100	4.49/100	8.59/100	17.0/100
StyleGAN — Threshold T_0	7.16/80.0	13.7/79.3	**22.2/69.4**	12.1/15.9
StyleGAN — Threshold T_1	**54.5/20.0**	**90.0/15.5**	81.8/8.11	9.09/4.55
LSGAN — Threshold T_0	9.54/83.3	16.0/75.9	**22.1/66.7**	**34.5/53.2**
LSGAN — Threshold T_1	**93.3/46.7**	**87.5/24.1**	81.8/8.11	100/9.09

Table 1. Precision / Recall rate (in %) for the Identity Membership Attack (Alg. 1) on GANs trained on the VGGFace2 and CASIA-WebFace datasets with low bias and varying diversity (setting 1 § 3.2). *Higher is better, and corresponding best F_1 scores are emboldened. $N = \sum_{i \in \mathcal{Y}_G} N_i^G$ indicates the total number of individual samples used during training of the GAN and $|\mathcal{Y}_G|$ is the number of identities. For the attack, $|\mathcal{Y}_F|$ identities are used ($|\mathcal{Y}_F| = 8631$ for VGGFace2, $|\mathcal{Y}_F| = 1292$ for CASIA-WebFace). The baseline is given by* random *guessing, which reaches a precision of $\frac{|\mathcal{Y}_G|}{|\mathcal{Y}_F|}$ whatever the recall.*

3.5   Early Stopping

Thus far, we have only explored dataset size and diversity as possible factors effecting attack efficacy. However, a common tool in image classification to prevent overfitting is early stopping. In order to asses GAN sample quality, we refer to the commonly used Frechet Inception Distance (FID) [6]. Fig. 7 investigates the interaction between FID and the number of GAN training iterations. Indeed, beyond a certain iteration (here for StyleGAN, around 17000), continuing to train is detrimental to privacy, with no gain in FID.

4   Discussion and Conclusion

In this work, we exposed several properties of GANs trained on facial data not previously discussed in the literature. By controlling the identities seen during training and subsequently detecting those identities with a face identification network, we demonstrated a successful blind membership attack. We identified
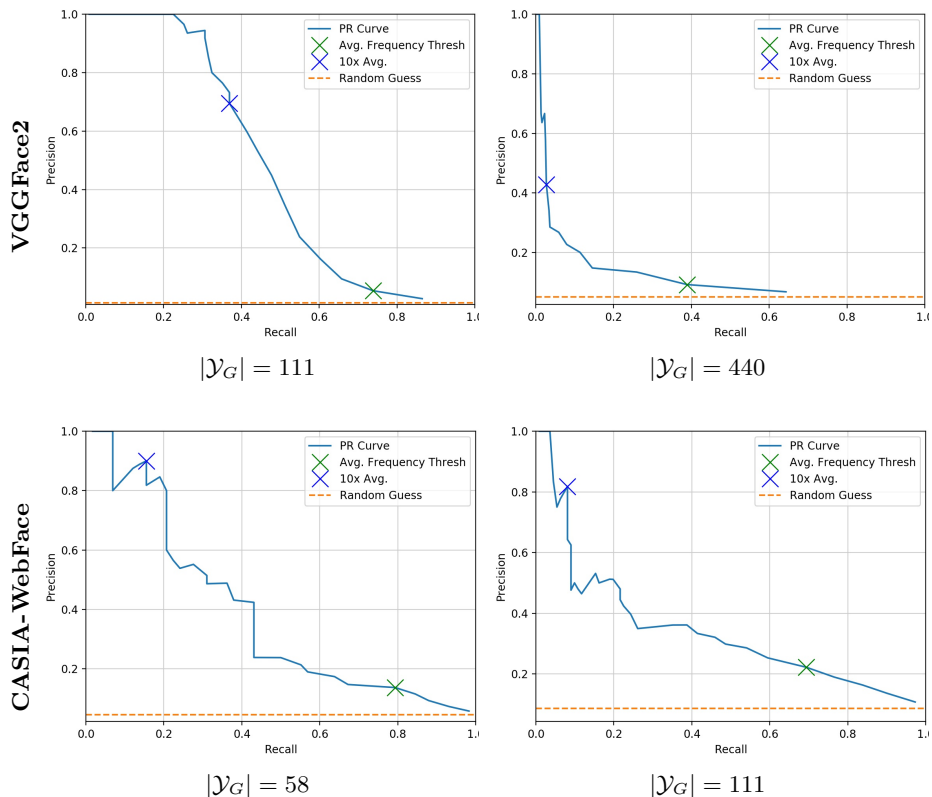
Fig. 4. Precision-recall curves for identity membership attacks for setting 1. *Curves are obtained using Eq.(1), the Style-GAN generator, two datasets and varying diversity $\mathcal{Y}_G$. Models trained on more diverse data appear to be more private. In addition, the threshold $T_1 = 10T_0$ can discern training identities with a high precision.*

several factors influencing susceptibility to attack, most notably that datasets with more identities have less detectable overfitting. Furthermore, dataset diversity alone will not protect against the presence of dataset bias. Finally, more iterations seemed to exacerbate attack success while not necessarily better image quality.

Contrary to most membership attack in the literature, this is a pure *black-box* attack in the sense that it does not require any further training based upon the generator, nor additional information about its architecture and parameters [33]. In addition, our attack is driven directly against the generator, and it did not require the exact training samples to be successful, as in the LOGAN approach.

Yet, the proposed approach is still restricted in various ways. To begin with, privacy issues is not limited to face datasets. Generalizing this work to other

		$\|\mathcal{Y}_{G_1}\|$ (N_1)	20 (6k)	40 (12k)	80 (24k)	160 (48k)
VGGFace2		$\|\mathcal{Y}_{G_2}\|$ (N_2)		2000 (40k)		
		Random	0.23/100	0.46/100	0.93/100	1.85/100
	STYLEGAN	Threshold T_0	0.70/75.0	1.25/65.0	2.99/76.2	3.83/49.4
		Threshold T_1	**33.3/30.0**	**31.2/25**	**56.2/22.5**	**40/8.75**
	LSGAN	Threshold T_0	0.46/51.3	1.12/54.8	2.33/62.5	**3.07/37.5**
		Threshold T_1	**26.5/21.7**	**29.8/23.25**	**47.8/13.8**	5.56/18.8

Table 2. Precision / Recall rate (in %) for the Identity Membership Attack (Alg. 1) on GANs trained on the VGGFace2 dataset in the second setting. *The dataset $\mathcal{Y}_G$ is now composed of two sets of distinct identities: $\mathcal{Y}_{G_1} \cup \mathcal{Y}_{G_2}$. GANs are trained on a large number of identities $\|\mathcal{Y}_G\| = \|\mathcal{Y}_{G_1}\| + \|\mathcal{Y}_{G_2}\|$ (exceeding 2000), with different bias towards a small set of identities (as measured by the ratio $\|\mathcal{Y}_{G_1}\|/\|\mathcal{Y}_G\|$). N_1 and N_2 represents the number of samples for $\mathcal{Y}_{G_1}$ and $\mathcal{Y}_{G_2}$. For the attack, $\|\mathcal{Y}_F = 8631\|$ identities are used (including $\mathcal{Y}_G$) and precision/recall are reported for identities from $\mathcal{Y}_{G_1}$ against $\mathcal{Y}_F \cap \mathcal{Y}_{G_1}$. The baseline is given by* random *guessing, which is the proportion of training identities $\frac{\|\mathcal{Y}_{G_1}\|}{\|\mathcal{Y}_F\|}$.*

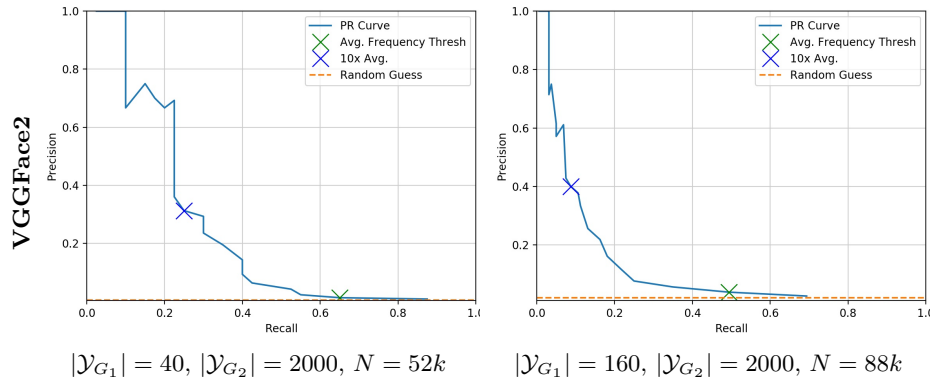|	
$\|\mathcal{Y}_{G_1}\| = 40, \|\mathcal{Y}_{G_2}\| = 2000, N = 52k$	$\|\mathcal{Y}_{G_1}\| = 160, \|\mathcal{Y}_{G_2}\| = 2000, N = 88k$

Fig. 5. Precision-recall curves for membership attack on biased data described in Sec. 3.3. *Curves are obtained using Eq.(1), the Style-GAN generator, two datasets and varying diversity of the biased dataset $\mathcal{Y}_{G_1}$. Even though both settings are trained with significant diversity (more than 2k identities and 50k images), the biased data in $\mathcal{Y}_{G_1}$ is still detectable.*

modalities (such as other biometric data, medical images, or other media) that are scrutinized with specialized networks is crucial. Furthermore, entirely ruling out overfitting is much more challenging than detecting it. Indeed, as we showed that the proposed attack is less effective as the number of identities grows, it
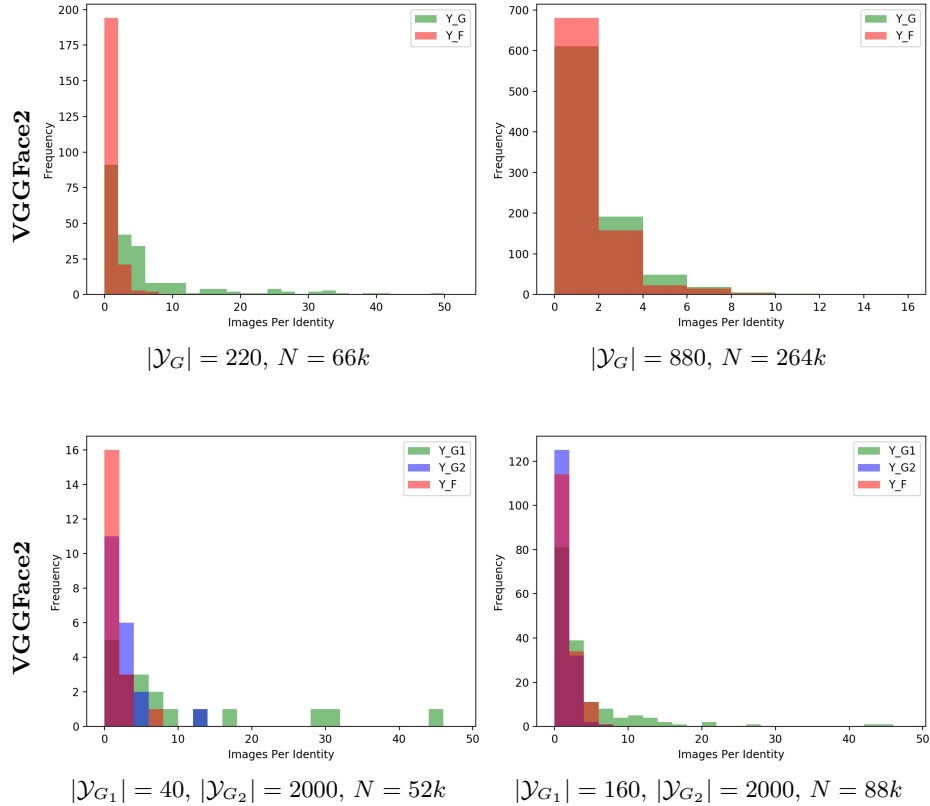
Fig. 6. Per identity frequency histograms of generated samples. *In the first row, StyleGAN generators are trained with unbiased data (setting 1), see Sec. 3.2. Without enough diversity, membership attacks are highly precise, as the bars in green (representing private training identities), can be easily distinguished. The second row showcases the high diversity & high bias scenario in Sec. 3.3. Here, blue represents a third, diverse auxiliary set, which cannot be distinguished, contrary to the biased samples (in green).*

strongly depends on the performance of the identification network. More effort needs thus to be put to improve detection. A first direction would be to improve the attack itself by designing more sophisticated mechanisms. For instance, by training the identification network specifically for this task (to measure the generalisation upon the discovery of new identities rather than a fixed number of ones from the training set). We also note that while the attack doesn't require training samples, it does require $\mathcal{Y}_G$ and $\mathcal{Y}_F$ to have some amount of overlap. Clearly, the attack precision would be zero if $\mathcal{Y}_F \cap \mathcal{Y}_G = \emptyset$.

Finally, this work addresses a problem with training GANs on sensitive data (e.g. copyrighted or private face images). Careful dataset curation or early stop-
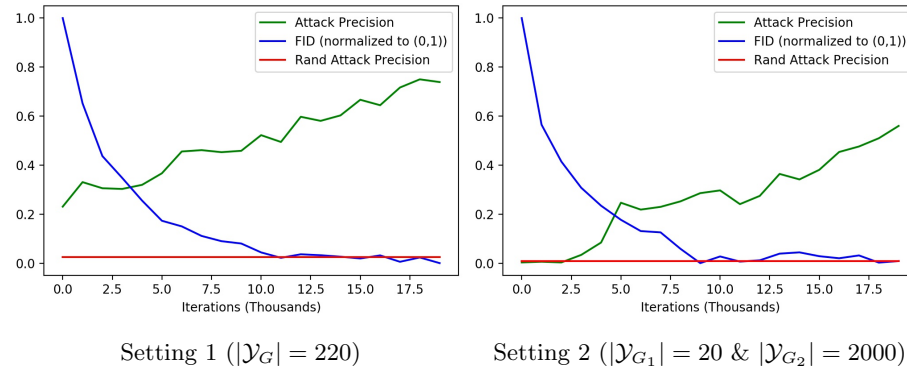
Setting 1 ($|\mathcal{Y}_G| = 220$)   Setting 2 ($|\mathcal{Y}_{G_1}| = 20$ & $|\mathcal{Y}_{G_2}| = 2000$)

Fig. 7. FID vs attack precision. *The graphs above suggest that early stopping may be useful in GAN training, w.r.t. privacy. While the FID converges, membership attacks continue to gain precision.*

ping may mitigate these problems. On the other hand, more sophisticated solutions may exist, such as those integrated into the GAN objective function, or those which directly modify the generation, such as distillation, to promote privacy. We leave this for future work.

For the time being, ensuring sufficient face dataset diversity (in terms of number of identities), low dataset bias and early stopping (when the FID stops changing, or reaches a threshold) all appear to be good empirical tools to train safer models.

References

1. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. (2014) 2672–2680
2. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2017) 4681–4690
3. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Generative image inpainting with contextual attention. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 5505–5514
4. Webster, R., Rabin, J., Simon, L., Jurie, F.: Detecting overfitting of deep generative networks via latent recovery. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 11273–11282
5. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 8789–8797
6. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Advances in Neural Information Processing Systems. (2017) 6626–6637
7. Sajjadi, M.S., Bachem, O., Lucic, M., Bousquet, O., Gelly, S.: Assessing generative models via precision and recall. In: Advances in Neural Information Processing Systems. (2018) 5228–5237
8. Simon, L., Webster, R., Rabin, J.: Revisiting precision recall definition for generative modeling. In: International Conference on Machine Learning. (2019) 5799–5808
9. Arora, S., Ge, R., Liang, Y., Ma, T., Zhang, Y.: Generalization and equilibrium in generative adversarial nets (gans). In: Proceedings of the 34th International Conference on Machine Learning-Volume 70, JMLR.org (2017) 224–232
10. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 4401–4410
11. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein gan. arXiv preprint arXiv:1701.07875 (2017)
12. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096 (2018)
13. Hayes, J., Melis, L., Danezis, G., Cristofaro, E.D.: LOGAN: Membership Inference Attacks Against Generative Models. PoPETs **2019** (2019) 133–152
14. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434 (2015)
15. Shen, Y., Luo, P., Yan, J., Wang, X., Tang, X.: Faceid-gan: Learning a symmetry three-player gan for identity-preserving face synthesis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 821–830
16. Xie, L., Lin, K., Wang, S., Wang, F., Zhou, J.: Differentially private generative adversarial network. arXiv preprint arXiv:1802.06739 (2018)
17. Jordon, J., Yoon, J., van der Schaar, M.: PATE-GAN: Generating synthetic data with differential privacy guarantees. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. (2019)

18. Webster, R., Rabin, J., Simon, L., Jurie, F.: Generating private data surrogates for vision related tasks. In: ICML workshop on Security and Privacy of Machine Learning. (2019) 15–20

19. Song, C., Ristenpart, T., Shmatikov, V.: Machine Learning Models that Remember Too Much. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017. (2017) 587–601

20. Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., Song, D.: The secret sharer: Evaluating and testing unintended memorization in neural networks. In: 28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019. (2019) 267–284

21. Shokri, R., Stronati, M., Song, C., Shmatikov, V.: Membership Inference Attacks Against Machine Learning Models. In: 2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017. (2017) 3–18

22. Melis, L., Song, C., De Cristofaro, E., Shmatikov, V.: Inference Attacks Against Collaborative Learning. arXiv:1805.04049 [cs] (2018)

23. Fredrikson, M., Jha, S., Ristenpart, T.: Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security - CCS '15, Denver, Colorado, USA, ACM Press (2015) 1322–1333

24. Melis, L., Song, C., De Cristofaro, E., Shmatikov, V.: Exploiting Unintended Feature Leakage in Collaborative Learning. In: arXiv:1805.04049 [Cs]. (2019)

25. Liu, K.S., Li, B., Gao, J.: Performing Co-Membership Attacks Against Deep Generative Models. arXiv:1805.09898 [cs, stat] (2018)

26. Bojanowski, P., Joulin, A., Lopez-Paz, D., Szlam, A.: Optimizing the latent space of generative networks. In: Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018. (2018) 599–608

27. Kemelmacher-Shlizerman, I., Seitz, S.M., Miller, D., Brossard, E.: The megaface benchmark: 1 million faces for recognition at scale. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2016) 4873–4882

28. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: A dataset for recognising faces across pose and age. In: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018), IEEE (2018) 67–74

29. Liu, Z., Luo, P., Wang, X., Tang, X.: Large-scale celebfaces attributes (celeba) dataset. Retrieved August **15** (2018) 2018

30. Wang, M., Deng, W.: Deep face recognition: A survey. arXiv preprint arXiv:1804.06655 (2018)

31. Parkhi, O.M., Vedaldi, A., Zisserman, A.: Deep face recognition. (2015)

32. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision. (2017) 2794–2802

33. Sablayrolles, A., Douze, M., Ollivier, Y., Schmid, C., Jégou, H.: White-box vs Black-box: Bayes Optimal Strategies for Membership Inference. In: ICML 2019 - 36th International Conference on Machine Learning, Long Beach, United States (2019)