This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

This Face Does Not Exist... But It Might Be Yours! Identity Leakage in Generative Models

Patrick Tinsley

Abstract

Adam Czajka University of Notre Dame {ptinsley,aczajka,flynn}@nd.edu Patrick Flynn

Generative adversarial networks (GANs) are able to generate high resolution photo-realistic images of objects that "do not exist." These synthetic images are rather difficult to detect as fake. However, the manner in which these generative models are trained hints at a potential for information leakage from the supplied training data, especially in the context of synthetic faces. This paper presents experiments suggesting that identity information in face images can flow from the training corpus into synthetic samples without any adversarial actions when building or using the existing model. This raises privacy-related questions, but also stimulates discussions of (a) the face manifold's characteristics in the feature space and (b) how to create generative models that do not inadvertently reveal identity information of real subjects whose images were used for training. We used five different face matchers (face_recognition, FaceNet, ArcFace, SphereFace and Neurotechnology Mega-Matcher) and the StyleGAN2 synthesis model, and show that this identity leakage does exist for some, but not all methods. So, can we say that these synthetically generated faces truly do not exist? Databases of real and synthetically generated faces are made available with this paper to allow full replicability of the results discussed in this work.

1. Introduction

Since generative adversarial networks (GANs) were first introduced in 2014 [9], they have demonstrated impressive photo-realistic image-generating capabilities. During the GAN training process, a generator and a discriminator are pitted against each other for the purpose of elevating the generator. In the context of image data, these two adversaries battle over the production and recall of "fake" images; the generator attempts to fool the discriminator by producing images that mimic the real data distribution while the discriminator aims not to be fooled by comparing and judging fake images from real. At the end of this minimax game,



Figure 1: The proposed identity leakage path through the StyleGAN2 generator follows the pink line from real images (FFHQ) to generated images (SG2).

the resulting generator, having learned the general distribution of the supplied training data, can be used to synthesize never-before-seen "fake" images that adopt the characteristics of real face images while not appearing in the training data. The results and future promise of these generative models have led to heavy academic intrigue [22, 11]. However, there is a potential risk regarding their ability to generate images to establish fake identities that can be exploited. As a phenomenon closely related to the well-known "fake news" problem bedeviling rational discourse in politics, science, and society, the risks of identity synthesis are analogous to the risks of baseless information synthesis and deserve careful attention from the originating domain as well as the discipline of ethics [25]. Of primary interest in this paper is an artifact of generative models that concerns individuals' sense of privacy regarding personal image data: **identity leakage in generative models**. These models optimize themselves to learn what comprises a realistic face, and in so doing, devise an encoding (a latent space) used as a starting point for subsequent syntheses. The optimization usually consumes millions of authentic face images, whose features are implicitly incorporated into the discriminator and thus into the resulting model. The model is then sampled to generate images of the faces of "nonexistent" people, starting with a latent space vector. There two important research questions at hand, which we address in this paper:

Q1. Does identity information from the real subjects, as embodied in training data, leak into the generated ("fake") samples?

Q2. If answer to *Q1* is affirmative, might this be a threat to privacy?

Figure 1 shows the flow of this information as it passes through the model training process from real image to fake image. So, if by chance, one's face image data exists in the training data, there is an opportunity for their data to inadvertently leak. These newly generated faces should then reproduce training identities in a way that skews the performance of a face recognition system. That is, the probability of being falsely matched to a face that "does not exist" is higher than when one's face is not included in the training data set. This paper is the first known to us to conduct experimental assessment of this problem in StyleGAN [16], and we demonstrate that identity information leakage does exist for some face matchers while for others it does not. This variation in information leakage across matchers also demands careful consideration in the context of face recognition system design.

This result stimulates additional intriguing research questions: can or should we sample the latent space where realistic face representations reside in some specific way to avoid such information leakage when generating face images of "nonexistent" subjects? Is there any specific arrangement of authentic faces in this latent space? If so, can we learn the general characteristics and geometry of this realistic manifold? We hope the results presented in this paper will stimulate a scientific discourse about the role of generative models in biometrics and their privacy-related attributes.

We offer databases of real and synthetic faces used in this work to other researchers to facilitate replication of our experiments and investigation of new related problems.



Figure 2: Style Mixing with StyleGAN2 (Truncation $\psi = 0.5$)

2. Related Work

2.1. Generative Models

A primary motivation for the use of generative models is their ability to produce high-quality data. To this end, NVIDIA's most recent contribution to the world of generative models is called StyleGAN2 [17], which improves upon their original StyleGAN work [16]. Before StyleGAN, NVIDIA released ProGAN [15], one of the first GANs to produce high-quality images. Given a large enough training data set, all three of these models can be trained to reproduce representative images of nearly any object. However, one of the most popular, successful applications has been a given model's ability to generate high-resolution face images that are informally considered very realistic. Built atop the ProGAN paradigm and trained on the FFHQ data set [16], both StyleGAN and StyleGAN2 learned to recreate these synthetic faces from the ground up using separable resolution-based "styles," hence the name. Coarse styles, also referred to as features, control pose and general face shape; medium styles control more detailed hair and facial features; fine features control color scheme for eyes, hair, skin, and even background. Due to the designed separability of these learned features, hybrid faces can be synthesized through latent vector manipulation, while still exhibiting real image characteristics, with some degree of useful control, as seen in Figure 2. In this figure, each "inner" cell can be seen as a combination of the coarse features of the column and the medium-to-fine features of the row.

2.2. Information Leakage

GAN-based leakage and model inversion attacks have proven to be rather effective in terms of extracting sensitive data from generative models, both intentionally and unintentionally. Hitaj *et al.* [13] suggest that GANs can act like composite artists generating sketches of an eyewitness's memory of an individual. Though this "sketch" (a synthetic sample) may not encompass all features of the individual, relevant and distinctive details about this person could be exposed. This is further explained by Fredrikson *et al.* [8] who focus on adversarial attacks. Given access to a generative model's architecture and weights, they show that one can leverage facial recognition confidence scores to reconstruct a person's face image. Goodfellow *et al.* [10] assert that a GAN (given infinite time and resources) would generate samples originally found in the training data, *i.e.* there is a possibility (albeit small) for an unprovoked 100% information leakage for a given sample from the training set.

2.3. Remedies for Information Leakage

Since sensitivity regarding personal data leakage into machine learning models has been so well documented, several solutions have been developed to prevent such information leakage [29, 3, 27]. At the heart of many such efforts is the notion of differential privacy, which acts as a privacy guarantee for sensitive data in the context of machine learning. For GANs, differential privacy can ensure synthetic data production with a limited risk of private information exposure. Xu et al. [28] offer GANobfuscator as a solution to the problem of data leakage in GANs. GANobfuscator adopts differential privacy in two ways: Gaussian noise injection and gradient-pruning; the latter method also boosts model stability during the training process. Once the model has been trained on the noise-infused data, "new" data can be generated by GANobfuscator and released into the public with mathematically bounded concern for privacy attacks, such as the already mentioned reverse engineering of training data. Similarly, Chen et al. [5] offer GANs based on auto-encoders and variational auto-encoders designed to protect privacy while retaining high data utility.

2.4. Fake Image Detection

Substantial efforts have gone into "Deepfake" detection, which asks the question of whether an image has been digitally created or manipulated in some way [14], [20], [21], [26]. A famous example of Deepfakes follows a "This *X* Does Not Exist" pattern, wherein *X* could be cats, apartment units, cars, or faces. For each of these examples, a separate GAN has been trained to generate fake images of the designated object. As mentioned earlier, GANs are incredibly good at doing this convincingly. In response to fake image generation is fake image detection. Further, researchers have studied the effects of GAN-generated fake face images on face recognition systems [18], claiming that several state of the art face matchers are vulnerable to Deepfake videos.

However, this paper aims not to determine whether a face image is fake. Quite contrarily, we aim to prove that within these generated fake images are remnants of the real images ingested during the GAN training process.

3. Data Sets

We use three data sets in experiments presented in this paper. The first data set is the Flickr-Faces-HQ data set (FFHQ), which includes 70,000 high-resolution (1024×1024) images of human faces varying in age, ethnicity, image background, and facial clothing/apparel (glasses, hats, etc). The entirety of the FFHQ data was used by NVIDIA labs to train the StyleGAN2 model, as well as the original StyleGAN. Each of 70,000 images in FFHQ represents a different individual's face.

The second data set we used is a collection of 70,000 StyleGAN2-generated fake face images developed in this work and offered along with the paper. We will call this data set "SG2". Along with several pre-trained models, NVIDIA provides scripts to generate faces via latent vector manipulation. In our case, we used latent vectors of length 512, corresponding to seed values from 0 through 69999. Style-GAN also adjusts latent vectors using the truncation trick to balance image quality and image variety. Truncation values of 0.0 generate high-quality images near the center, or mean of the latent space, whereas truncation values of 1.0 generate high-variation images from peripheral areas of poorer representation in the latent space. For our synthetic image data set (SG2), we used a truncation value of $\psi = 0.5$ to generate variety-quality balanced face images. The provided generator script can be found in NVIDIA's repository: https://github.com/NVlabs/stylegan2.

The third data set, used solely to assess the face recognition accuracy of various face matchers used in this work, is an assortment of 12,004 high quality face images collected by the authors¹. The data set features 333 subjects with between 4 and 78 front-facing photos per subject, captured at a resolution of 1200×1600 under good lighting. This data, further referred to as *UNI*, was used to gauge the relative performance of our face matchers. This set is subject-disjoint with both the FFHQ and SG2 databases.

4. Methodology & Experiments

In this paper, "identity leakage" refers to the flow of identity-salient facial features from real training data images (in FFHQ) to fake synthetic images (in SG2). We hypothesize that these identity-salient features are readily detectable by face matchers, whose job is to judge the similarity or dissimilarity of two face images.

For a given pair of images, a face matcher can quantitatively compare the two images via extracted face feature vectors. This comparison between face images is commonly expressed as a distance, where larger distances correspond to dissimilar images and smaller distances correspond to similar images. To prove our hypothesis, we want

¹See ND-WACV2021 Data Set at https://cvrl.nd.edu/projects/data/.

Table 1: Face Matcher Details

Model Name	Architecture	Training Data (# of images)	Feature Vector Length	Comparison Score	
face_recognition [1]	ResNet-34	Custom (~3M)	128	Euclidean Distance	
FaceNet [23]	GoogLeNet-24	VGGFace2 [4] (~3.3M)	512	Euclidean Distance	
ArcFace [7]	ResNet-100	MS-Celeb-1M [12] (~10M)	512	Cosine Distance	
SphereFace [19]	ResNet-64	CASIA-Webface [30] (~0.45M)	512	Cosine Distance	
MegaMatcher 11.2 [2]	Proprietary	Proprietary	Proprietary	Proprietary	

to show that the distance between a real image and a generated image is, on average, smaller than the distance between two real images. If this *is* the case, then identity "information" (identity-salient features) from the training data must be spilling into the generated data, as evidenced by the smaller distances between extracted feature vectors.

Since not all face matchers use dissimilarity-based distance scores, we simply refer to comparisons of images as "comparison scores". Comparison scores between pairs of real images are referred to as "R-R" (Real-Real) while comparison scores between one real image and one generated image are referred to as "R-G" (Real-Generated). Since these scores inherently rely on the extracted face feature vectors, identity leakage might be observable for some face matchers but not others.

4.1. Face Matcher Details

We used five face matchers to verify our hypothesis about identity information leakage. Four of the five matchers were open source, CNN-based, and implemented in Python. The fifth technique (MegaMatcher 11.2) was a commercial face matcher from Neurotechnology implemented in C++. Information regarding model architecture and training data can be found in Table 1. Note that MegaMatcher 11.2 uses similarity-based scores, *i.e.* lower scores correspond to dissimilar face images while higher scores correspond to similar face images. All other selected face matchers used distance, or dissimilarity scoring (lower scores indicate stronger similarity).

4.2. Preliminary Performance Analysis on UNI

We conducted an initial set of experiments on our UNI data set to determine the relative performance of the selected face matchers, to – potentially – link the strength of the identity leakage with the accuracy of a given face matcher. The first step of our pipeline was to extract face feature vectors from all 12,004 face images. In order to enable a fair comparison, we removed from the experiment all images that failed the face detection step in any of the five matchers. This yielded a usable image subset of 11,792 images.

Since the UNI data set contains multiple images per subject, matching experiments yielding both genuine and impostor comparison scores were conducted. The genuine scores numbered 290,122, while the impostor scores numbered 442,224. Genuine scores were calculated using all images available in the data set. On the other hand, impostor scores were calculated using a subset of the data. This subset was comprised of four images per subject since all subjects in the UNI data set had at least four images. A given subject's "first" two images (indices 0 and 1) were compared against the two "latter" images (indices 2 and 3) of all other subjects. So, in total there were 333 subjects \times 2 "first" images \times 332 other subjects \times 2 "latter" images = 442,224 impostor comparison scores per matcher. This methodology was developed to reduce computation time as well as balance the number of images per subject in the data set. Comparison score results as well as the failure to enroll rate (FTE) for each matcher can be seen in Table 2. Impostor and genuine score distributions are shown in Figure 3.

Table 2: UNI Comparison Results

Model Name	FTE (%)	AUC	d'
face_recognition	9.99e-4	0.9991	5.1930
FaceNet	2.49e-4	0.8945	1.6048
ArcFace	2.49e-4	0.9978	9.5513
SphereFace	2.49e-4	0.9261	2.1124
MegaMatcher 11.2	1.76e-2	0.9999	13.6908

To quantify the similarity and dissimilarity between scoring distributions, we use the decidability measure d', which cumulates sample mean and sample variances of comparison scores into a scalar factor:

$$d' = \frac{|\overline{x}_g - \overline{x}_i|}{\sqrt{\frac{1}{2}(\overline{x}_g + \overline{\overline{x}}_i)}} \tag{1}$$

where \overline{x} and $\overline{\overline{x}}$ denote sample mean and sample variance of a comparison score x, respectively. The value of d' estimates the degree by which the genuine distribution of x_g and impostor distribution of x_i overlap (the higher d' is, the lower is the overlap). For uncorrelated random variables, $d' = \sqrt{2}\zeta$ where ζ is a standardized difference between the expected values of two random variables, often used in hypothesis testing.



Figure 3: Genuine and impostor score distributions for each matcher on *UNI*. From left to right: face_recognition, FaceNet, ArcFace, SphereFace, MegaMatcher 11.2.

Based on d' (as seen in Table 2), the relative ranking of the five matchers is: MegaMatcher 11.2 > ArcFace > face_recognition > SphereFace > FaceNet.

4.3. Impostor Distribution Analysis between FFHQ & SG2

Two impostor distributions of comparison scores were produced for each face matcher. The first shows the distribution of "Real-Real" (R-R) comparison scores; the second shows the distribution of "Real-Generated" (R-G) comparison scores.

The R-R partition was created by pairing real images from the FFHQ data set together. Although FFHQ originally included 70,000 real images, 12,906 images failed to enroll in one or more of the matchers and were removed from the set of usable "real" images. The remaining 57,094 images were split in half and matched together for a total of 28,547 pairs.

The R-G partition was created by pairing each of the usable 57,094 "real" images with an image from the usable "generated" SG2 data set, i.e. the set of synthetic images that properly enrolled in all five matchers. Note that 628 SG2 images failed to enroll to at least one matcher.

Figure 4 illustrates the R-R and R-G comparison score distributions for each face matcher. Additionally, FTE for each matcher for each data set, d' and the two-sample

Kolmogorov-Smirnov test statistic are reported in Table 3. Similar to d', the K-S test quantifies the distance between two distributions; the larger the test statistic is, the greater is the distance.

The first immediate and interesting observation (from Table 3) is that FTE values are lower for synthetic images that from real images. This may be because face detectability in its output is "baked into" the network, by limiting its training inputs to images with detectable faces. Future work might investigate how truncation value (ψ) affects face detection and recognition for StyleGAN2-generated face images.

The second interesting observation is that we see a **noticeable shift between the R-R and R-G impostor distributions towards genuine scores** for ArcFace and SphereFace matchers in Fig. 4. This type of shift suggests that the generated face images do in fact contain identity-salient features from the training data, *i.e.* identity leakage is evident. We elaborate more on this observation in the next section. Note that although the observed differences are statistically significant (p < 0.001 at the significance level alpha = 0.01), the difference in means suggests that only ArcFace and SphereFace have leftward shifts from R-R to R-G. On the other hand, the face_recognition package and FaceNet model show statistically significant shifts to the right for the R-G distribution, which actually suggests



Figure 4: R-R and R-G score distributions for each matcher. From left to right: face_recognition, FaceNet, ArcFace, SphereFace, MegaMatcher 11.2.

Table 3: "Real-real faces" (R-R) and "Real- generated faces" (R-G) comparison results, along with failure-to-enroll rates for each face matcher, separately for real samples FTE_R and generated samples FTE_G .

Model Name	FTE _R	FTE_G	d'	K-S Test Statistic	<i>p</i> -value	Difference in Means (R-R-R-G)
face_recognition	3.62e-3	0.0	0.0855	0.0332	< 0.001	+0.008082
FaceNet	1.71e-4	0.0	0.1055	0.0496	< 0.001	+0.01266288
ArcFace	1.71e-4	0.0	0.3801	0.1562	< 0.001	-0.05384
SphereFace	1.71e-4	0.0	0.4404	0.1895	< 0.001	-0.06824
MegaMatcher 11.2	1.84e-1	8.97e-3	0.0116	0.0026	0.9978	-0.04813

the opposite of identity leakage. This varying behavior between matchers is curious and will be further discussed the next section as well.

5. Results & Discussion

5.1. Answering Q1: Does identity information from the real subjects leak into the generated samples?

Before conducting any experiments, we hypothesized that results from *all* five selected face matchers would support the existence of identity leakage. However, we quickly found that this was not the case after generating and comparing the R-R and R-G impostor distributions with face_recognition, one of the quicker off-the-shelf face matchers.

We then adjusted our hypothesis, thinking that maybe only the strongest performing face matchers would support our leakage claim, hence the need for ranking the face matchers on the *UNI* data set. There is a precedent for this assumption in the iris recognition domain, where the author found that the most accurate matchers are to the highest extent sensitive to "biometric template aging" [6]. Based on the relative ranking of the face matchers (by the d' metric), we suspected that the MegaMatcher 11.2 SDK and the ArcFace face matchers would yield leakage-evident results.

However, as seen in Table 3, identity leakage is supported only in ArcFace and SphereFace, the second and fifth face matcher in terms of accuracy, respectively. Figure 4 offers visual support to the same conclusions. In the two leftmost plots, we see that the face_recognition and FaceNet matchers do not show signs of identity leakage, as is the case with MegaMatcher 11.2 in the rightmost plot. There is strong overlap between the R-R and R-G comparison score distributions for all three of these matchers, suggesting minimal or no information leakage. Contrarily, for ArcFace and SphereFace, the R-G score distributions are shifted to the left of the R-R scores. Since ArcFace and SphereFace comparison scores are distance-based, a shift to the left (smaller distances) does support identity leakage as such a shift hints that generated face images are, on average, "closer" to real face images.

The question then becomes: why does identity leakage only appear for certain face matchers? What information is leaking that leads to difference in impostor comparison scores? To better understand these inquiries from a qualitative standpoint, we constructed "closest image" grids in Figure 5. The objective of these grids is to investigate how different face matchers judge "close" face images, and in so doing, gain insight into why only some matchers think fake face images look more like the face images from specific real identities. The base images (5a, 5b) are two real face images from the FFHO data set. The grids below each base image (5c, 5d) show rows of images for each face matcher. Each row showcases the three closest synthetic face images from the SG2 data to the corresponding base image at the top, with the closest image on the left and farthest on the right. For all matchers but MegaMatcher 11.2, comparison score is a distance measure (Euclidean distance for face_recognition and FaceNet, and cosine distance for ArcFace and SphereFace). For these matchers, lower scores correspond to closer or more similar face images. For MegaMatcher, higher comparison scores mean more similar face images.

Upon observing Figure 5, we note a few points of interest regarding the face matchers' ability to match fake and real face images. Comparison scores for MegaMatcher 11.2 (bottom row) do not offer much in terms of visual correspondence between face images. Inaccurate scores span race, gender, and age. This suggests that perhaps Mega-Matcher 11.2 relies on more complex, less obvious facial features to compare, match, and verify faces. For reference, to achieve a false acceptance rate of 0.0001%, the matching threshold is 72, which is 8x the score (9) of the two examples in 5. Since the software is proprietary and comparison results are rather stagnant for non-match image pairs, there is not much we can discuss in terms of leakage.

Between the remaining four models, there seem to be two groups: matchers that display evidence of identity leakage (ArcFace and SphereFace), and matchers that do not (face_recognition and FaceNet). The first group, simply based on nomenclature, seems to revolve around nonlinear geometry. In [7], Deng *et al.* explicitly compare Arc-Face and SphereFace. Though the former improves upon



(c) The closest samples to (a) (d) The closest samples to (b)

Figure 5: Three closest fake face images (SG2) to real face images from FFHQ shown in (a) and (b) according to each face matcher: face_recognition, FaceNet, ArcFace, SphereFace, MegaMatcher 11.2. Row images are sorted left-to-right in increasing order of distance to the corresponding base image.

the latter, they retain a few identifying similarities. Both apply angular margin penalties to boost performance for intra-class appearance variation, such as extreme pose and age difference. ArcFace uses additive angular margin while SphereFace uses multiplicative angular margin. Further, both matchers compare face feature vectors via trigonometric functions and share a ResNet model architecture. Perhaps these similarities may play a role in explaining why these models seem to support identity leakage while other models do not. Further work to this end includes using GradCAM visualization [24] to quantitatively discuss face matcher behavior regarding real and fake images.

Contrary to ArcFace and SphereFace, the face_recognition and FaceNet matchers seem to favor

linearity. One of the main focuses of Google's contribution with FaceNet was Euclidean-based triplet loss. Similarly, the face_recognition matcher measures similarity between face feature vectors in Euclidean distance. We have yet to explore this rigorously, but perhaps linearity versus non-linearity in face matcher geometry has an effect on identity leakage. Circling back to our initial question, we have experimentally shown that the answer to Q1 is affirmative: the identity information <u>can</u> leak into generated data, although the strength of this phenomenon is uneven across the matchers.

5.2. Answering Q2: Might the observed data leakage a threat to privacy?

Figure 6 shows the false matching rate as a function of the acceptance threshold for comparison scores between real and unrelated-real (R-R), and real vs potentiallyrelated-generated (R-G) face images when using the Arc-Face face matcher (the method prone to the identity leakage to the highest extent in our experiments). We can note a clear shift in false match probability when "unexisting" faces are being matched against real subjects. To better understand this plot, we can discuss a social media-inspired example. Recently, Facebook shut down thousands of Facebook and Instagram accounts that used "AI-generated images" as profile photos. If Facebook had chosen to use a matcher behaving as shown in Fig. 6 to hunt down these fake profile photos (as generated by StyleGAN2), then real users whose face images appeared in the training data (yellow curve) would be 6.6 times more likely to be falsely matched when compared to someone who did not appear in the training data (blue curve). The baseline false positive rate comes from a NIST standard of 0.001 (1 in 1000) for face recognition. Thus, the answer to Q2 is also affirmative: there does exist a practical threat to privacy, and its seriousness may be encountered even for matchers presenting solid identification performance for authentic face images.

5.3. Inconclusiveness

The fact that our findings are split between proving and disproving identity leakage emphasizes the importance of careful and thoughtful design in face matching systems. Why some facial features trickle through one face recognition model while not through others is still unknown. Accordingly, practical use cases of face recognition need to account for the fact that generative models can be used to exploit these models without an obvious adversarial attack. Further work, including several iterations of retraining of StyleGAN2, will hopefully shed more light on the causes and further effects of identity leakage in generative models.

6. Conclusions and Future Directions

Identity leakage in generative models requires careful attention, if only for the potential impact on privacy. It has been shown in previous works that sensitive data can be pried out of models via adversarial attacks, but in this paper we have also shown that in some cases this data can



Figure 6: Illustration of how false match rate probability increases when our face image is included in the training set of StyleGAN2 with ArcFace-based matching

simply spill out of the training corpus into synthetic samples without any provocation. We also demonstrated that identity leakage depends on the face matcher technology in use, which may be a potentially impactful result for privacy protection-aware network design used in future models. As mentioned earlier, future endeavors include the usage of network visualization tools, such as GradCAM or activation atlases, for quantitative studies of matcher behavior in the context of identity information leakage. Illustrating the differences between face feature extraction for identity-leaking and identity-preserving models would certainly help us understand what features are used for comparing faces. Are these the same type of facial features that can be leaked? Additionally, between the combinations of training data sets, model architectures, and loss functions, there are many, many more flavors of face matcher to test for identity leakage. Among those are CosFace, DeepFace, VGGFace, and SphereFace+. In a similar vein, we intend to also explore other face synthesizer models, such as Star-GAN, CycleGAN, and GauGAN. Finally, a validation of proposed differential privacy-based GANs on an even larger set of synthetic face images would test the capabilities of available solutions and offer feedback to better prevent unwanted identity leakage.

References

[1] face_recognition: The world's simplest facial recognition API for Python and the command line. https://github.com/ageitgey/face_recognition.

- [2] Neurotechnology: Fingerprint, face, eye iris, voice and palm print identification, speaker and object recognition software. https://www.neurotechnology.com/.
- [3] B. K. Beaulieu-Jones, Z. S. Wu, C. Williams, R. Lee, S. P. Bhavnani, J. B. Byrd, and C. S. Greene. Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7):e005122, 2019.
- [4] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. VGGFace2: A dataset for recognising faces across pose and age. In *International Conference on Automatic Face and Gesture Recognition*, 2018.
- [5] Q. Chen, C. Xiang, M. Xue, B. Li, N. Borisov, D. Kaarfar, and H. Zhu. Differentially private data generative models. *arXiv preprint arXiv:1812.02274*, 2018.
- [6] A. Czajka. Influence of iris template aging on recognition reliability. In M. Fernández-Chimeno, P. L. Fernandes, S. Alvarez, D. Stacey, J. Solé-Casals, A. Fred, and H. Gamboa, editors, *Biomedical Engineering Systems and Technologies*, pages 284–299, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- [7] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceed*ings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4690–4699, 2019.
- [8] M. Fredrikson, S. Jha, and T. Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 1322–1333, 2015.
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information* processing systems, pages 2672–2680, 2014.
- [10] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572, 2014.
- [11] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of Wasserstein GANs. In Advances in neural information processing systems, pages 5767–5777, 2017.
- [12] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao. MS-Celeb-1M: A dataset and benchmark for large scale face recognition. In *European Conference on Computer Vision*. Springer, 2016.
- [13] B. Hitaj, G. Ateniese, and F. Perez-Cruz. Deep models under the gan: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference* on Computer and Communications Security, pages 603–618, 2017.
- [14] N. Hulzebosch, S. Ibrahimi, and M. Worring. Detecting cnngenerated facial images in real-world scenarios. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pages 642–643, 2020.
- [15] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196, 2017.

- [16] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proceed*ings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4401–4410, 2019.
- [17] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. Analyzing and improving the image quality of Style-GAN. arXiv preprint arXiv:1912.04958, 2019.
- [18] P. Korshunov and S. Marcel. Deepfakes: a new threat to face recognition? assessment and detection. arXiv preprint arXiv:1812.08685, 2018.
- [19] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphereface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 212–220, 2017.
- [20] F. Matern, C. Riess, and M. Stamminger. Exploiting visual artifacts to expose deepfakes and face manipulations. In 2019 IEEE Winter Applications of Computer Vision Workshops (WACVW), pages 83–92. IEEE, 2019.
- [21] J. C. Neves, R. Tolosana, R. Vera-Rodriguez, V. Lopes, H. P. Proença, and J. Fierrez. Ganprintr: Improved fakes and evaluation of the state of the art in face manipulation detection. *IEEE Journal of Selected Topics in Signal Processing*, 2020.
- [22] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434, 2015.
- [23] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 815–823, 2015.
- [24] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [25] C. Shen, M. Kasra, W. Pan, G. A. Bassett, Y. Malloch, and J. F. O'Brien. Fake images: The effects of source, intermediary, and digital media literacy on contextual assessment of image credibility online. *new media & society*, 21(2):438– 463, 2019.
- [26] R. Tolosana, R. Vera-Rodriguez, J. Fierrez, A. Morales, and J. Ortega-Garcia. Deepfakes and beyond: A survey of face manipulation and fake detection. arXiv preprint arXiv:2001.00179, 2020.
- [27] B. Xin, W. Yang, Y. Geng, S. Chen, S. Wang, and L. Huang. Private fl-gan: Differential privacy synthetic data generation based on federated learning. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2927–2931. IEEE, 2020.
- [28] C. Xu, J. Ren, D. Zhang, Y. Zhang, Z. Qin, and K. Ren. Ganobfuscator: Mitigating information leakage under gan via differential privacy. *IEEE Transactions on Information Forensics and Security*, 14(9):2358–2371, 2019.
- [29] K. Xu, H. Yue, L. Guo, Y. Guo, and Y. Fang. Privacypreserving machine learning algorithms for big data systems. In 2015 IEEE 35th international conference on distributed computing systems, pages 318–327. IEEE, 2015.
- [30] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Learning face representation from scratch. arXiv preprint arXiv:1411.7923, 2014.