Image and Vision Computing 135 (2023) 104688

Contents lists available at ScienceDirect

Image and Vision Computing

journal homepage: www.elsevier.com/locate/imavis

Synthetic data for face recognition: Current state and future prospects

Fadi Boutros [a,*], Vitomir Struc [c], Julian Fierrez [d], Naser Damer [a,b]

[a] Fraunhofer IGD, Fraunhoferstr. 5, Darmstadt 64283, Hessia, Germany
[b] Technical University Darmstadt, Karolinenplatz 5, Darmstadt 64289, Hessia, Germany
[c] Faculty of Electrical Engineering, University of Ljubljana, Trzaska cesta 25, Ljubljana SI-1000, Slovenia
[d] School of Engineering, Universidad Autonoma de Madrid, Madrid 28049, Spain

ARTICLE INFO

Article history:
Received 15 February 2023
Accepted 24 April 2023
Available online 11 May 2023

Keywords:
Face recognition
Synthetic data
Biometrics

ABSTRACT

Over the past years, deep learning capabilities and the availability of large-scale training datasets advanced rapidly, leading to breakthroughs in face recognition accuracy. However, these technologies are foreseen to face a major challenge in the next years due to the legal and ethical concerns about using authentic biometric data in AI model training and evaluation along with increasingly utilizing data-hungry state-of-the-art deep learning models. With the recent advances in deep generative models and their success in generating realistic and high-resolution synthetic image data, privacy-friendly synthetic data has been recently proposed as an alternative to privacy-sensitive authentic data to overcome the challenges of using authentic data in face recognition development. This work aims at providing a clear and structured picture of the use-cases taxonomy of synthetic face data in face recognition along with the recent emerging advances of face recognition models developed on the bases of synthetic data. We also discuss the challenges facing the use of synthetic data in face recognition development and several future prospects of synthetic data in the domain of face recognition.

1. Introduction

The breakthroughs of deep neural networks and their training optimizations as well as the availability of large-scale identity-labeled face datasets have reshaped the research landscape of face recognition (FR) over the past years. These emerging technologies have dramatically improved FR performances leading to the wider integration of FR in a variety of applications from logical access control and consumer low-end devices to automated border control. State-of-the-Art (SOTA) FR models [1,2] utilized large-scale face datasets e.g. CASIA-WebFace [3], MS-Celeb-1M [4], or VGGFace2 [5] to train deep neural networks (DNN) with millions of trainable parameters, where the goal is to optimize the empirical risk minimization function given input training samples, their corresponding labels, and DNN trainable parameters. Achieving such a goal without being over-optimized, i.e. overfitted, requires that training datasets are of large scale (massive number of images of many identities) and representative of various variations that exist in the real world. Large and representative data is also required to evaluate FR accuracies against different variations that present in real operation scenarios e.g. pose, aging, occlusion, or lighting. Data is required to evaluate the vulnerability of FR against different types of attacks such as morphing, presentation, master-face, and deep fake attacks. FR components, face processing models, attack detectors, and face image quality estimation models are not different as they require face data for training and evaluation. Besides the technical limitation of collecting large-scale data with realistic variations, there are increased concerns about collecting, maintaining, redistributing, and using biometric data due to legal, ethical, and privacy concerns [6]. Consequently, many widely used datasets for FR development such as VGGFace2 [5] and MS-Celeb-1M [4] have been retracted by their creator. Table 1 summarizes the most widely used datasets to train FR models. Even though many of these datasets have been publically released, there are not any more accessible.

Processing biometric data is governed by a set of legal restrictions [6]. Taking the General Data Protection Regulation (GDPR) [6] as an example, it categories biometric data as a special category of personal data subjected to rigorous data protection rules [7], requiring high protection in connection with fundamental rights and freedoms of individuals. Dealing with such data requires adherence to one of the exemptions of biometric data processing [8], the related national laws [9], maintaining processing records [10], and the preparation of data protection impact assessment [11,12], among other restrictions. Depending on the purpose of the biometric data processing, this set of restrictions can be rigorously extended [13–15]. Besides the legal complications of using and sharing biometric data, ethical requirements are commonly necessary, such as the approval of an ethics committee or competent authorities.

* Corresponding author.
E-mail address: fadi.boutros@igd.fraunhofer.de (F. Boutros).

https://doi.org/10.1016/j.imavis.2023.104688
0262-8856/

Table 1
Overview of the most widely used authentic and synthetic facial datasets commonly used to train FR models, along with the number of images, identities, images per identity, and the fact that each database is public and/or still accessible. Note that many of the public databases are not accessible (raising a practical problem for researchers and developers) anymore based on legal and ethical concerns and even those that are available are ethically questioned as the individual consent of the data subjects is not always insured.

Name	Year	# Images (m)	# Identities (k)	Avg.	Public	Accessible	Authentic
CASIA-WebFace [3]	2014	0.5	10.6	47	✓	×	✓
DeepFace [19]	2014	4.4	4.0	1092	×	×	✓
FaceNet [20]	2015	200.0	8,000.0	25	×	×	✓
Facebook [21]	2015	500.0	10,000.0	50	×	×	✓
VGGFace [22]	2015	2.6	2.6	992	✓	×	✓
CelebFaces [23]	2016	0.09	5.4	16	✓	✓	✓
MS-Celeb-1M [4]	2016	10	100.0	100	✓	×	✓
MegaFace2 [24]	2017	4.7	672.0	7	✓	✓	✓
UMDFaces [25]	2017	0.4	8.3	46	✓	×	✓
VGGFace2 [5]	2018	3.3	9.1	363	✓	×	✓
IMDbFace [26]	2018	1.7	59.0	29	✓	✓	✓
MS1MV2 [2,4]	2019	5.8	85.0	68	✓	×	✓
MillionCelebs [27]	2020	18.8	636.0	30	×	×	✓
WebFace260M [28]	2021	260	4,000.0	65	✓	✓	✓
WebFace42M [28]	2021	42	2,000.0	21	✓	✓	✓
SynFace [17]	2021	0.5	10	50	✓	✓	×
DigiFace-1M-A [29]	2022	0.72	10	72	✓	✓	×
DigiFace-1M-B [29]	2022	0.5	100	5	✓	✓	×
SFace [16]	2022	0.63	10.6	60	✓	✓	×
USynthFace [18]	2022	0.4	0.4	1	✓	✓	×

The increased concerns about the legal and ethical use of authentic data in biometrics along with the technical limitation in collecting large and diverse face datasets motivate recent works to propose the use of synthetic data as an alternative to privacy-sensitive authentic data in FR training [16–18]. In an attempt to provide a clear understanding of the feasibility of utilizing synthetic face data to train, evaluate, attack, or privacy enhancement, this work is the first to analyze the properties needed of the synthetic data for FR, the use-cases taxonomy of synthetic data in FR, the current state of synthetic-based FR, the limitations and challenges facing the use of current synthetic face data in FR, and possible future research directions that might give a larger space for synthetic data in different aspects of FR development.

2. Where is the synthetic data used?

To analyse the properties of the needed synthetic data, one should start by building a clear taxonomy of the different possible uses-cases of synthetic data in its interaction with FR. This taxonomy here will consider the operations where the synthetic data is used to interact with the recognition part of FR systems, i.e. the feature extraction. Therefore, synthetic data that is meant to interact with other system components, as defined in ISO ISO/IEC 19795–1:2021 [30], are out of scope, e.g. synthetic data used to train or evaluate face detection or segmentation solutions. Additionally, synthesizing faces as a means of domain transformation, e.g. from thermal to visible face appearance [31] is also out of scope as it just transfers the appearance of the image.

Fig. 1 presents the use-case taxonomy of the synthetic face data interaction with FR. These use-cases are categorised under 4 groups, along with the properties of the possibly needed data under each category (the latter will be discussed in detail in the next section). The four use-case categories are discussed in the following.

1. **Training FR:** Modern FR solutions are based on deep learning models that are either trained directly to generate identity-discriminant feature representations (e.g. triplet loss [20]) or to classify the identity classes in the training data (e.g. ArcFace [2], ElasticFace [1], etc.). In the latter approach, embeddings proceeding the classification layer of the network are then used to extract the identity-discriminant representations. This family of approaches is currently predominantly leading to SOTA FR performances. In both cases, training face data that represents the high inter and intra-class diversity of real applications is needed to train the models. As mentioned in the introduction, the diversity of such data, if authentic, is limited by practical data collection constraints, and its collection and handling are hedged by privacy, legal, and ethical concerns. Synthetic data can come in handy to train such FR models in different manners based on the training requirements. If the model is trained in one of the two approaches mentioned above, then the synthetic data has to contain a large number of identities and multiple samples of each identity. If the model is trained on partially authentic data, however, the intra-class variation of this data is low, then the synthetic data needs to contain multiple samples for each of the authentic identities, i.e. act as an augmentation strategy. Finally, if the FR model is trained in an unsupervised manner, then the synthetic training data is not largely concerned with the identity grouping, but rather just requires a set of faces of random identities. This data has also been shown to be successful in training processes during the training-aware quantization of models based on full precision parameters [32]. Although it is out of the scope of this work, synthetic faces of this kind can also be used to train face detectors, face segmentation, and attack detection methods (e.g. morphing attack detection [33]).
2. **Evaluating FR:** FR algorithmic evaluation, following the ISO ISO/IEC 19795–1:2021 [30], requires the existence of a large set of genuine (same identity) and imposter (different identity) face image pairs that represent the real operational scenario. The need for a large number of these pairs is intensified by the ever-more accurate performance of FR algorithms. FR algorithms can produce two main algorithmic errors, genuine pairs classified wrongly as imposters (false non-match (FNM)) or imposter pairs classified wrongly as genuine (false match FM). As the algorithms produce lower and lower rates of decision errors, the FM rates (FMR) and FNM rates (FNMR), the number of evaluated pairs required to produce statistically significant evaluation results become higher. This need for large-scale evaluation data is one of the main motivations behind requiring synthetic data for the evaluation. Another reason is that some authorities that require in-house testing on their own data when purchasing FR solutions do only possess a single image per identity

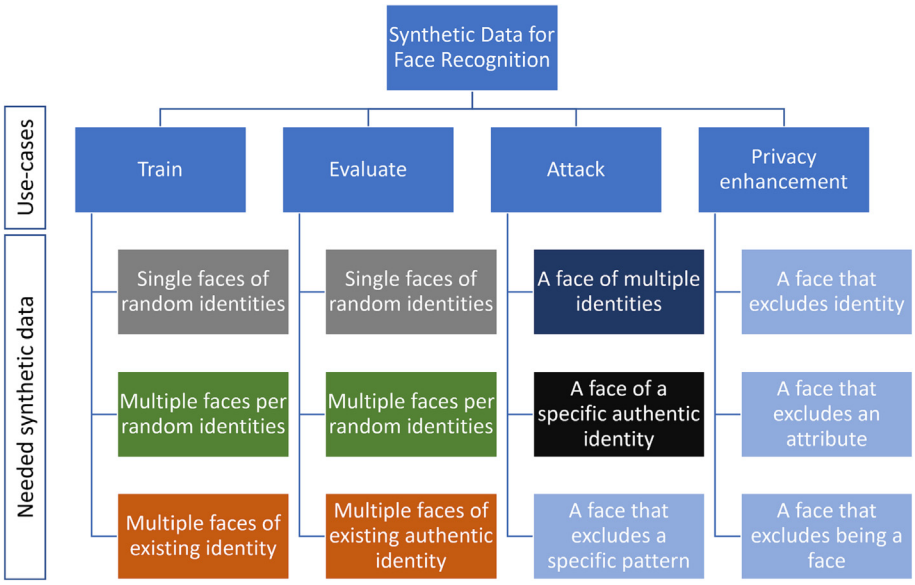

Fig. 1. A taxonomy of the synthetic data use-cases (on the top of the figure) directly interacting with FR models, either by training them, evaluating them, attacking them, or enhancing the privacy of the information extracted by them. This taxonomy lists the existing and foreseen synthetic data types that are needed by these use-cases (under each use-case). These data needs are grouped by their main properties by color and discussed, along with the use-cases in this paper.

in their databases (think of visa systems) and thus it is impossible to have genuine pairs to evaluate FR algorithms. Such situations would require synthetic data to be generated so it belongs to a certain authentic identity, but with realistic variations. In a third scenario where the operation scenario would require a very low FMR, the need for a huge number of imposter pairs is required to evaluate, with statistical significance, the FMR. In such cases, random synthetic faces with random identities can be used to create such imposter pairs. Again, although it is out of the scope of this work, these synthetic faces, regardless of their identity information, can be used to evaluate face detectors, face segmentation, and presentation/morphing attack detection.

3. **Attacking FR:** Commonly, developers would use technology to enhance the convenience and security of individuals and societies. However, technology can also be used maliciously to create attacks on individuals, systems, and societies. This is the case also with synthetic face data, which can also be used as an attack. Synthetic data can be created so that a certain face can be matched with two or more faces. This can target automatic FR comparison or human image verification, or both. Such attacks can be face morphing attacks, where an image is generated to match two or more identities, then used on an identity or travel document with the alphanumeric data of when the targeted matches. Later such a document can be used by the other targeted identities illegally, leading to a serious security threat. Another attack in the same category is the MasterFace attack, where the synthetic face is created to match a wider proportion of the population, raising many security threats. The second type of attack by generated face images might focus on generating a face image of a specific identity. Such attacks are commonly referred to as Deep-Fakes and they are commonly used to fool the viewer into wrongly believing that a certain person has said or done an action in an image or a video. A third attack can use synthetic faces that maintain a certain identity but excludes a specific pattern with the aim of attacking a biometric-based system that ensures a legal operation of a process. Such an attack can be by presenting the attacker's real identity, but excluding the information that points out that the user is underage, in a service that requires age verification.
4. **Enhancing the privacy for FR users:** Although excluding certain patterns from generated images of specific identities can be seen as an attack on biometric systems, in different use-cases, they can be seen as a privacy-enhancing tool when they are used to avoid the illegal or unconsented processing of the data. Such generation of the data aims at maintaining a certain set of visual patterns but removing the clues of a specific pattern. Depending on the use-case, this excluded pattern can be related to the identity in what is widely known as image-level face de-identification, which is defined under the standard ISO/IEC 20889:2018 [34]. The excluded pattern can be related to certain soft biometric attributes like age or gender, which is commonly referred to as soft-biometric privacy enhancement. Although it is out of the scope of this work, the generated faces can exclude patterns that makes them detectable to face detection tool, i.e. excluding the information that makes the face a face in the view of automatic face detection.

So far, we presented a discussion on the possible use-cases of synthetic face data in FR. Each of these use-cases has different needs when it comes to synthetic data. These needs are discussed in the next section.

3. What data is needed and what properties make it good?

The properties of the needed synthetic data under the different use-cases (discussed in the previous section) are grouped by their required properties under different colors in Fig. 1 and are discussed in detail in the following:

1. **Single faces of random identities:** As detailed in the previous section, and illustrated in Fig. 1, synthetic face images of random identities without the requirement of multiple images to belong to one identity can be used for training FR models in an unsupervised manner. They additionally can be used to evaluate FR models, specifically evaluate the FMR, especially when the targeted operational point is at a very low FMR, requiring an extremely large number of diverse imposter pairs to make the evaluation result statistically significant. Here, such data should be realistic, i.e. act like authentic data when processed by the FR model. A successful way to measure that was proposed in [32] and it is based on comparing the activation function

value ranges in the FR model when processing authentic data versus when processing the synthetic data. Additionally, the distribution of the comparison scores between pairs of these single images of random identities should theoretically be similar to those of imposter comparisons of authentic data to ensure the similarity to the authentic inter-identity variation, which was explored in [18].

2. **Multiple faces per random identity:** This kind of data represents what one would typically expect from FR training or evaluation data. That is, multiple identities, with multiple images per identity. This, given a sufficient inter and intra-class (identity variation), can be used to train an FR model in a supervised manner. This also would contain both imposter and genuine pairs to evaluate the performance of FR by calculating both possible errors, FMR and F NMR. Such data should also interact with the FR model similarly to authentic data, this as mentioned earlier can be measured by monitoring the value range of the model's activation functions. Here, the data should possess an inter and intra-class variability of the targeted authentic data scenario. We specify "targeted" here as different evaluation and training goals of FR might occur, e.g. a model is evaluated specifically for cases with an extreme pose or extreme age differences between the comparison pairs (intra-class variations), or for cases of pairs of twins or siblings (inter-class variation). This goes for training as well, as an FR model can be trained to specifically be tolerant to mask occlusions, and thus the training data inter and intra-class diversity should represent that. The suitability of such data can be measured by comparing its genuine and imposter comparison scores distributions with that of the targeted authentic data (which can be much smaller in size) as performed in [16]. For specifically targeted attribute variations, such as age and pose, attribute predictors can be used to ensure the existence of such attribute variations in the synthetic data to the same degree as the authentic data.
3. **Multiple faces of an existing identity:** Authentic face data with insufficient intra-class variation is problematic for the training and evaluation of FR. In terms of training an FR model, such data will lead to models that are not trained to tolerate intra-class variation (e.g. pose, expressions, age, illumination, etc.) and thus are expected to lead to high F NMR in practical operations. When evaluating FR, evaluation data in some practical cases such as an authority that possesses only a single (or few) images per identity (e.g. visa applicant database) would not be sufficient to evaluate the expected F NMR as no (or few) genuine pairs exist in the data. Both cases require acquiring more samples of each of the existing identities. These samples have to be of realistic variation that matches the targeted scenario. Such samples might be created synthetically and would act as an augmentation approach when training an FR model, or as additional samples to create genuine pairs when evaluating FR models (or training FR in a triplet loss-like strategy). Such synthetic data should interact with the FR model similarly to authentic data, as previously discussed. It should also result in genuine comparison score distribution that matches the targeted authentic data scenario. One must take notice that this should be the case when the pairs are between the existing authentic sample is compared to the synthetic images of the same identity, but also, if needed, between the synthetically generated samples of the same identity themselves.
4. **A face of multiple identities:** A synthetic face can also be used as an attack, the fact that a face can be generated synthetically with properties that enables an attack on identity systems pursues researchers to foresee such attacks. A face can be synthesized in a way that it matches two more specific (known) identities to create what is referred to as a morphing attack. A morphing attack image is designed to match with a number of specific identities and can be created on the image level by interpolating the images of the targeted identities, or generated synthetically to possess the identity information of the targets [35]. Such an image, if used in association with a passport or an identity document can enable multiple persons to be verified to the alphanumeric information on the card. A wider attack that surfaced lately in the literature is the MasteFace attack, where the attack image is synthesized to match a wide range of the population without the need to know the targeted identities [36]. As these attacks might be used to attack visual inspection, automatic verification, or both, they first have to have a natural appearance. This natural appearance is best measured by user studies, where individuals are asked if an image appears realistic or not. The vulnerability of automatic FR to such attacks, and thus the measure of how good is the synthetic data for its purpose, can be measured using the Mated Morph Presentation Match Rate (MMPMR) [37]. The MMPMR refers to the fraction of morphs whose similarity to both identities used to morph, are below the selected FR comparison score threshold relative to all morphs.
5. **A face of specific authentic identity:** Synthesizing a face of a specific authentic identity is usually related to the need to synthesize this face with also a specific expression or domain, unlike generating such faces of an authentic identity where a realistic variation is needed. This is commonly related to what is referred to as DeepFake faces but also includes other face manipulation techniques such as expression and attribute manipulations. As such attacks aim at manipulating human viewers, their success is best measured by how realistic they are to these viewers and how well they succeeded in the targeted manipulation in the view of the viewers through user studies related to the exact goal of the manipulation. However, more within the scope of this work is the ability of these attacks to fool automatic FR and attack detection algorithms. A comprehensive survey on the issue of DeepFakes and facial image manipulation is presented by Tolosana et al. in [38].
6. **A face that excludes a specific pattern:** A face synthesizing process can maintain a subset of patterns from a specific face and excludes other subsets of these patterns. Such patterns can be identity information, age, gender, ethnicity, or even the patterns that make a face detectable as a face, among other attribute patterns. Such a process can be seen as an attack if it is aimed at avoiding a consented required process, such as automatic age verification to receive a service or make an online purchase. However, such a process can also be seen as a privacy enhancement mechanism. Excluding the identity, while maintaining the image appearance and other attributes to some degree is commonly referred to as image-level face de-identification and it aims at avoiding the unconsented identification of face images, whether in the public or private space. A subset of this is to exclude the patterns of the face that makes it detectable and thus avoid further processing. Removing other patterns like gender or age falls within the image-level soft-biometric privacy enhancement techniques that aim at maintaining the identification possibilities without allowing unconsented estimation of soft-biometric attributes. Evaluating the ability to synthesize these face images is based on evaluating the degree to which the patterns that need to be excluded and the ones that need to be maintained are detectable, where the first need to be as undetectable as possible and the latter needs to be as detectable as possible. A comprehensive survey and discussion on these technologies are presented by Meden et al. in [39].

4. Where are we now?

4.1. Face image generation

A deep generative model (DGM) is a deep neural network that is trained to interpret and model a probability distribution of the authentic training data. Specifically, a deep generative model takes random points from e.g. Gaussian distribution and maps them through a neural network such as the generated distribution closely matches the authentic data distribution. The main DGM approaches that are proposed in the literature are Variational Auto-Encoder (VAE) [47], Generative Adversarial Network (GAN) [48], Autoregressive model [49], and Normalizing

Flows [50] and Diffusion Models (DiffModel) [51], in addition to a large number of hybrid models that combined two of previous approaches such as GAN with VAE [52]. A comprehensive review of deep generative modelings is presented by [53]. Each of these approaches presented contributions towards providing a better trade-off between generated sample quality i.e. producing samples of high perceived quality and fidelity that resemble the DGM training data, inference time i.e. enabling fast sampling mechanism, architecture restrictions i.e. some of the DGMs are limited to underlying network architecture and sample appearance variations.

4.2. How do the DGM approaches match the needed synthetic face data properties?

- Single faces of random identities: DGM approaches such as StyleGAN [54] presented very promising results in generating single faces of random synthetic identities with high visual fidelity. However, the generated faces could share the identity information, to a small degree, with DGM's original training (as reported in [55,16]).
- Multiple faces per random identities: Approaches such as Face-ID-GAN [56], DiscoFaceGAN [57], GAN-Control [58], InterFaceGAN [59], and CONFIG [52] proposed GAN models based on disentangled representation learning to conditionally generate face images from synthetic identities with predefined attributes e.g. age, pose, illumination, or expression. As generated images are explicitly controlled by a predefined set of attributes, such images might lake the intra-class diversity that exists in real-world face data and it is needed to train and evaluate FR.
- Multiple faces of an existing identity: DGM approaches such as CONFIG [52] are able to regenerate multiple faces of an existing identity by reconstructing input faces with a predefined set of attributes such as changing expression, wearing sunglasses, adding makeup, or changing hair color. However, such attribute manipulation approaches might induce some artifacts in reconstructed faces, which might affect identity preservation between the input and the reconstructed faces. Also, as such approaches are explicitly manipulating the attributes of their input faces, the generated faces might not contain large appearance variations, which are needed to train and evaluate FR models. More importantly, identity preservation in reconstructed samples is rarely evaluated and reported.
- A face of multiple identities: DGM approaches were not explicitly designed and trained to generate a face of multiple identities. However, recent works such as MorGAN [35], MIPGAN [60], and MorDIFF [61], make use of generative models to generate a face of multiple identities by interpolating two or more latent vectors of synthetic or real faces and then generating a new face of multiple identities. In a similar manner, however, with latent vector optimization rather than optimization, MasterFaces [36] are generated to match unknown identities.
- A face of specific authentic identity: DGM approaches that targeted image-to-image modeling achieved impressive results in generating a face of specific authentic identity. This has been commonly achieved by manipulating the input source face to match specific attributes or a target domain while maintaining the identity information of the source image. Although such approaches did not target generating Deep-Fake attacks, they have been widely used in generating such kinds of attacks [38].
- A face that excludes a specific pattern: None of the SOTA DGM approaches explicitly target generating a face that excludes a specific pattern. A number of works make use of DGM approaches to exclude a specific pattern e.g. identity, age, or gender of authentic input faces, especially when such models include attribute disentanglement. However, to the best of our knowledge, none of the previous works present solutions to generate a face of synthetic identity that excludes a specific pattern, rather this is done for faces of authentic identities. An overview of the current state of this issue can be found in [39].

4.3. What is the current state of the defined use-cases?

Very recently a few works build on existing DGM approaches to propose FR based on synthetic data. The following discussion presents the use of synthetic data in FR grouped by the use-cases (discussed earlier in this paper and presented in Fig. 1).

4.3.1. Training FR

Recently, synthetically generated face data has been proposed as an alternative to privacy-sensitive authentic data to train FR models mitigating the technical, ethical, and legal concerns of using authentic biometric data in training FR models. The currently proposed approaches in the literature utilized synthetically generated data to train unsupervised (UsynthFace [18]) or supervised FR models (SFace [16], SynFace [17], DigiFace-1M [29] and IDnet [46]). Training the unsupervised FR model as in UsynthFace requires that the training data maintain the Property 1 (Section 2) i.e. single face of random identities, while supervised approaches, SFace, SynFace, IDnet, and DigiFace-1M, require that the training data maintain the Property 2 i.e. multiple faces per random identities (Section 2). Some of these approaches, SynFace and DigiFace-1M, proposed combining authentic with synthetic data during the training or transferring the knowledge from the pretrained FR model to improve the recognition accuracies. Others (USynthFace) utilized only synthetic data for FR training. Most synthetic FR approaches utilized GAN-based (UsynthFace, SynFace) and/or geometric and color transformation data augmentation (UsynthFace, IDnet, and DigiFace-1M) methods to create more challenging training samples improving the model recognition accuracies. Table 2 summarizes the achieved accuracies on five FR benchmarks by recent FR models trained on synthetic data. It can be observed from the reported results in Table 2 that including data augmentation in FR model training significantly improved the recognition accuracies. Also, the unsupervised FR model (UsynthFace [18]) obtained very competitive results using unlabeled data to supervised synthetic-based FR models. Samples of synthetic data used in the SOTA synthetic-based FRs are shown in Fig. 2.

4.3.2. Evaluating FR

A few works proposed the use of synthetic data for evaluating FR. SynFace [17] presented a synthetic version of the Labeled Faces in the Wild (LFW) dataset [40] and evaluated two FR models trained on authentic and synthetic data, respectively on the synthetic version of the LFW. The model trained on real data achieved an accuracy of 98.85% and the one trained on synthetic data achieved an accuracy of 99.98%. The work [17] also suggested that the degradation in the verification performance between the two models is due to the domain gap between synthetic and real training images.

4.3.3. Attacking FR

DGM approaches have been widely and successfully utilized to generate morphing, MasterFace, deep-fake, and manipulation attacks on FR. Researchers generally attempt to foresee such attacks and evaluate their potential. Deep-fake and face manipulation attacks are already a serious problem facing modern societies and their generation is becoming more available and realistic with time [38]. Morphing attacks based on synthesized faces are a serious threat and FR recognition vulnerability to them is getting close to that of image-level morphing [60]. MasterFace attacks are relatively new, their initial proposed form is based on optimization on a relatively weak FR model [36] with other works arguing their feasibility [62]. However, on the other hand, synthetic data has helped create privacy-friendly databases for the detection of such

Table 2
Verification accuracies (%) on five different FR benchmarks achieved by the supervised and unsupervised FR models trained on the synthetic training databases with the numbers of real and synthetic training samples. The result in the first row is reported using the FR model trained on the authentic dataset to give an indication of the performance of an FR model trained on the authentic CASIA-WebFace dataset [3]. To provide a fair comparison, all model results are obtained from the original published works using the same network architecture (ResNet50) trained on relatively same training dataset size. KT refers to knowledge transfer from the pretrained FR model. LFW [40], AgeDB-30 [41], CFP-FP [42], CA-LFW [43], CP-LFW [44] are widely used FR evaluation benchmarks.

Method	Unsupervised	Data augmentation	# Synthetic Images	# Authentic Images	KT	LFW	AgeDB-30	CFP-FP	CA-LFW	CP-LFW
CosFace [45]	×	-	0	500 K	×	99.55	94.55	95.31	93.78	89.95
SynFace [17]	×	GAN-based	500 K	0	×	91.93	61.63	75.03	74.73	70.43
DigiFace-1M [29]	×	-	500 K	0	×	88.07	60.92	70.99	69.23	66.73
DigiFace-1M [29]	×	Accessory + Geometric and color	500 K	0	×	95.40	76.97	87.40	78.62	78.87
SFace [16]	×	-	634 K	0	×	91.87	71.68	73.86	77.93	73.20
USynthFace [18]	✓	GAN-based + Geometric and color	400 K	0	×	92.23	71.62	78.56	77.05	72.03
IDnet [46]	×	-	528 K	0	×	84.83	63.58	70.43	71.50	67.35
IDnet [46]	×	Geometric and color	528 K	0	×	92.58	73.53	75.40	79.90 (3)	74.25
SynFace [17]	×	GAN-based	500 K	40 K	×	97.23	81.32	87.68	85.08	80.32
DigiFace-1M [29]	×	Accessory + Geometric and color	500 K	40 K	×	99.05	89.77	94.01	90.08	87.27
SFace [16]	×	-	634 K	0	✓	99.13	91.03	91.14	92.47	87.03

attacks, specifically, the morphing attack [33,63] and face presentation attack [64]. Huber et al. [63] organized a competition on face morphing attack detection (MAD) based on privacy-aware synthetic training data [33]. The competition aimed at promoting the use of synthetic data to develop MAD solutions and attracted 12 solutions from both academia and industry.

4.3.4. Privacy enhancement

Main advances in this respect are presented under one of two categories, de-identification or soft-biometric privacy. De-identification can be achieved by adding adversarial noise to the image, image obfuscation, and image synthesis, the latter being the core focus of this work. Many solutions have been proposed in the literature, with a recent overview of these solutions presented in [39]. The main challenge so far in this domain is the cross-FR model performance as most works showed very good performances on the FR models that were used to optimize the solution, however, this performance drops when using other unknown FR models. Syntheses-based soft-biometric privacy followed a similar trend as de-identification, however, with much less dominance in the literature. In this aspect, many works rather focused on soft-biometric privacy on the template level rather than the image. Image and template level techniques are surveyed in [39]. An example of image-based techniques is the FlowSAN [65] aimed at minimizing gender information in the resulting images. Here, as the target is the soft-biometrics and not the identity, the main challenge is to achieve generalized performance across soft-biometric estimators while maintaining FR performance across FR models.

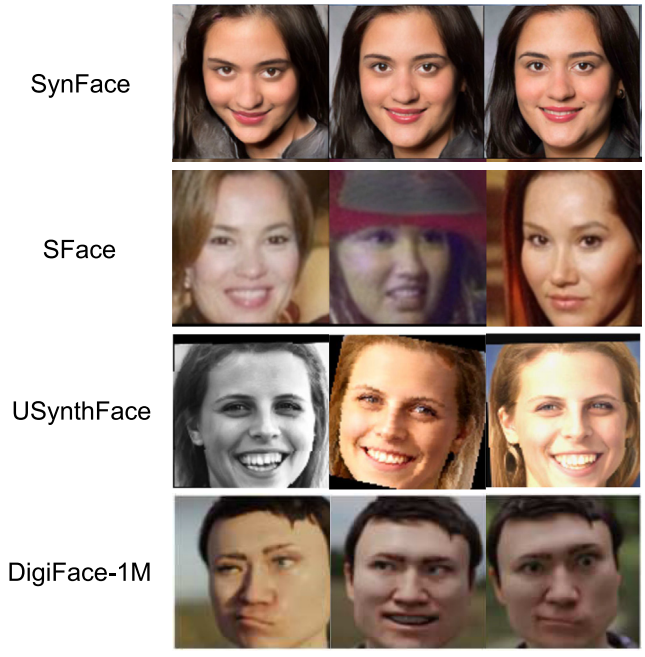

Fig. 2. Sample of synthetic data used in SynFace [17], UsynthFace [18], DigiFace-1M [29] SFace [16] and IDnet [46]. It can be clearly noticed the high variations in SFace images in comparison to other synthetic datasets. Although SynFace and UsynthFace utilized the same DGM (DiscoFaceGAN), it can be also observed the appearance variations in USynthFace using geometric and color transformations.

5. Where can we do better?

Here, based on the discussed use-cases taxonomy, the synthetic data requirements, and their current state along with the generation process, we discuss the main issues where further improvement in future research can have a strong effect on the use of synthetic data in FR. The following discussion will touch on the generation process, the defined use-cases, as well as the general lack of well-defined suitability evaluation protocols.

5.1. Face image generation

Generating realistic and high-quality samples along with enabling high sampling speed and high-resolution scaling have derived the main contributions of recent generative models proposed in the literature. In addition, some DGM approaches targeted specific applications such as image in-painting, attribute manipulation, face aging, image super-resolution, and image-to-image and text-to-image translations. Such applications mainly require that the generated samples are of high visual fidelity with less focus on the identity information, which might be less optimal for biometric applications. When developing DGM for FR use-cases, the solution should focus on the utility of the generated images for the given tasks rather than only focusing on the human-perceived quality. The emerging works on training FR solutions, presented earlier, are considered the first step in this regard. This focus on utility, rather than only the perceived quality, should be the main drive in future research when synthesizing images for FR.

5.2. Training FR

Recent works that proposed the use of synthetic face data for FR utilized deep neural network architectures with hyper-parameters that are optimized on authentic data. Such training paradigms might be sub-optimal for learning face representations from synthetic data. Future research works might target proposing network architectures or training

paradigms designed specifically to learn from synthetic data. In general, training FR solutions of synthetic data still fails behind those trained on authentic data in terms of accuracy, which is the main practical shortcoming that hinders placing such solutions in practical use currently. However, one must keep in mind that training FR on synthetic data is a very recently emerging research direction and it is already achieving higher recognition accuracies than solutions trained on synthetic data less than a decade ago [19].

5.3. Evaluating FR

The need for large-scale FR evaluation datasets that represent real scenario variations is the main motivation for future research directions on synthetic data for FR evaluation. Although DGMs can generate arbitrary realistic face images, the utility of the generated images for FR remains challenging. Future research works include but are not limited to, DGMs for generating multiple faces of existing authentic identities, which might target specific variations such as age and pose, and generating complete evaluation datasets of multiple images of multiple identities.

5.4. Attacking FR

Even though creating novel attacks on identity management systems and society in general sounds is a serious malicious action, it is essential to foresee attacks created by real attackers to better enable their detection. As the attackers would ask, the researchers should also ask "What is the strongest attack I can create to serve the attack goals given the current state of basic technology?" This follows the never-ending game of cat and mouse between attacks and attack mitigation. Therefore, the constant struggle here is to always try to foresee new attacks and attack generation methodologies and analyze their strengths and weaknesses, leading to better mitigation strategies.

5.5. Privacy enhancement

The main challenge to generative face privacy enhancement is the generalizability and robustness as it must possess to maintain operation in real-world applications. This generalization must ensure that the de-identification properties are strongly maintained even with unknown FR solutions. The same goes for soft-biometric privacy, where the privacy-enhanced images should maintain their privacy properties when processed by diverse soft-biometric estimators with different levels of knowledge [66]. Other open issues that still require increasing attention are the lack of clear quantifiability and provability privacy enhancement, the limited public benchmarks, and the need for controllable privacy where the user can have a choice of the privatised information [39].

5.6. Evaluation protocols

We provided in this work an initial discussion on what synthetic data is needed for different FR use-cases and what properties are needed from such data based on the way it is used. However, this initial discussion should evolve into a much-needed set of evaluation metrics and protocols that can precisely and comparably answer the question of "How well does the created data fit its targeted properties within its use-case?" Besides, and based on, the needed academic efforts in this regard, given that the synthetic data is foreseen to be a commodity, there is a need for such protocols and metric standards on the industrial level. A clear candidate to develop such a standard would be the ISO SC37 work group 5 on Biometric testing and reporting.

6. Conclusion

The use of authentic data in FR poses technical, legal, and ethical concerns. However, such data plays a major role in training, evaluating, enhancing the FR user privacy, and even attacking FR. This work provided initial discussions on the use of synthetic data in FR as an alternative to authentic data. We started by analysing and defining taxonomies for different possible FR use-cases in which synthetic data can be used. Then, we discussed the needed properties of synthetic data under each FR use-case. This has been followed by presenting the current state of synthetic FR. Finally, we provided several interesting directions of work that can be investigated in the future. As a concluding remark, the use of synthetic data in different FR uses-cases is still in the early research stage and this work provides a base discussion on this research direction and aims at motivating and promoting further research works toward responsible FR development.

Data availability

No data was used for the research described in the article.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This research work has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE and the ARRS research program P2-0250 (B) Metrology and Biometrics Systems. This work has been partially funded by the German Federal Ministry of Education and Research (BMBF) through the Software Campus Project.

References

[1] F. Boutros, N. Damer, F. Kirchbuchner, A. Kuijper, Elasticface: Elastic margin loss for deep face recognition, IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19–20, 2022, IEEE 2022, pp. 1577–1586, https://doi.org/10.1109/CVPRW56347.2022.00164.

[2] J. Deng, J. Guo, N. Xue, S. Zafeiriou, Arcface: Additive angular margin loss for deep face recognition, IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019, Computer Vision Foundation/IEEE 2019, pp. 4690–4699, https://doi.org/10.1109/CVPR.2019.00482, http://openaccess.thecvf.com/content_CVPR_2019/html/Deng_ArcFace_Additive_Angular_Margin_Loss_for_Deep_Face_Recognition_CVPR_2019_paper.html.

[3] D. Yi, Z. Lei, S. Liao, S.Z. Li, Learning face representation from scratch, CoRR abs/1411.7923 (2014). arXiv:1411.7923. http://arxiv.org/abs/1411.7923.

[4] Y. Guo, L. Zhang, Y. Hu, X. He, J. Gao, Ms-celeb-1m: A dataset and benchmark for large-scale face recognition, in: B. Leibe, J. Matas, N. Sebe, M. Welling (Eds.), Computer Vision - ECCV 2016–14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III, Lecture Notes in Computer Science, vol. 9907, Springer 2016, pp. 87–102, https://doi.org/10.1007/978-3-319-46487-9_6.

[5] Q. Cao, L. Shen, W. Xie, O.M. Parkhi, A. Zisserman, Vggface2: A dataset for recognising faces across pose and age, 13th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2018, Xi'an, China, May 15–19, 2018, IEEE Computer Society 2018, pp. 67–74, https://doi.org/10.1109/FG.2018.00020.

[6] The European Parliament and the Council of the European Union, Regulation (eu) 2016/679 of the european parliament and of the council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (General Data Protection Regulation) (2016).

[7] The European Parliament and the Council of the European Union, Article 9 of the general data protection regulation (2016).

[8] The European Parliament and the Council of the European Union, Article 9(2) of the general data protection regulation (2016).

[9] The European Parliament and the Council of the European Union, Article 9(4) of the general data protection regulation (2016).

[10] The European Parliament and the Council of the European Union, Article 30 of the general data protection regulation (2016).

[11] The European Parliament and the Council of the European Union, Article 35(3) of the general data protection regulation (2016).

[12] The European Parliament and the Council of the European Union, Article 37(1) of the general data protection regulation (2016).

[13] The European Parliament and the Council of the European Union, Article 22(4) of the general data protection regulation (2016).

[14] The European Parliament and the Council of the European Union, Article 27(2) of the general data protection regulation (2016).

[15] The European Parliament and the Council of the European Union, Article 6(4) of the general data protection regulation (2016).

[16] F. Boutros, M. Huber, P. Siebke, T. Rieber, N. Damer, Sface: Privacy-friendly and accurate face recognition using synthetic data, IEEE International Joint Conference on Biometrics, IJCB 2022, Abu Dhabi, United Arab Emirates, October 10–13, 2022, IEEE 2022, pp. 1–11, https://doi.org/10.1109/IJCB54206.2022.10007961.

[17] H. Qiu, B. Yu, D. Gong, Z. Li, W. Liu, D. Tao, Synface: Face recognition with synthetic data, 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021, IEEE 2021, pp. 10860–10870, https://doi.org/10.1109/ICCV48922.2021.01070.

[18] F. Boutros, M. Klemt, M. Fang, A. Kuijper, N. Damer, Unsupervised face recognition using unlabeled synthetic data, 17th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2023, Waikoloa Beach, HI, USA, January 5–8, 2023, IEEE 2023, pp. 1–8, https://doi.org/10.1109/FG57933.2023.10042627.

[19] Y. Taigman, M. Yang, M. Ranzato, L. Wolf, Deepface: Closing the gap to human-level performance in face verification, 2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23–28, 2014, IEEE Computer Society 2014, pp. 1701–1708, https://doi.org/10.1109/CVPR.2014.220.

[20] F. Schroff, D. Kalenichenko, J. Philbin, Facenet: A unified embedding for face recognition and clustering, IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7–12, 2015, IEEE Computer Society 2015, pp. 815–823, https://doi.org/10.1109/CVPR.2015.7298682.

[21] Y. Taigman, M. Yang, M. Ranzato, L. Wolf, Web-scale training for face identification, IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7–12, 2015, IEEE Computer Society 2015, pp. 2746–2754, https://doi.org/10.1109/CVPR.2015.7298891.

[22] O.M. Parkhi, A. Vedaldi, A. Zisserman, Deep face recognition, in: X. Xie, M.W. Jones, G.K.L. Tam (Eds.), Proceedings of the British Machine Vision Conference 2015, BMVC 2015, Swansea, UK, September 7–10, 2015, BMVA Press 2015, pp. 41.1–41.12, https://doi.org/10.5244/C.29.41.

[23] Y. Sun, X. Wang, X. Tang, Hybrid deep learning for face verification, IEEE Trans. Pattern Anal. Mach. Intell. 38 (10) (2016) 1997–2009, https://doi.org/10.1109/TPAMI.2015.2505293.

[24] A. Nech, I. Kemelmacher-Shlizerman, Level playing field for million scale face recognition, 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017, IEEE Computer Society 2017, pp. 3406–3415, https://doi.org/10.1109/CVPR.2017.363.

[25] A. Bansal, A. Nanduri, C.D. Castillo, R. Ranjan, R. Chellappa, Umdfaces: An annotated face dataset for training deep networks, 2017 IEEE International Joint Conference on Biometrics, IJCB 2017, Denver, CO, USA, October 1–4, 2017, IEEE 2017, pp. 464–473, https://doi.org/10.1109/BTAS.2017.8272731.

[26] F. Wang, L. Chen, C. Li, S. Huang, Y. Chen, C. Qian, C.C. Loy, The devil of face recognition is in the noise, in: V. Ferrari, M. Hebert, C. Sminchisescu, Y. Weiss (Eds.), Computer Vision - ECCV 2018–15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part IX, Lecture Notes in Computer Science, vol. 11213, Springer 2018, pp. 780–795, https://doi.org/10.1007/978-3-030-01240-3_47.

[27] Y. Zhang, W. Deng, M. Wang, J. Hu, X. Li, D. Zhao, D. Wen, Global-local GCN: large-scale label noise cleansing for face recognition, 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020, Computer Vision Foundation/IEEE 2020, pp. 7728–7737, https://doi.org/10.1109/CVPR42600.2020.00775, https://openaccess.thecvf.com/content_CVPR_2020/html/Zhang_Global-Local_GCN_Large-Scale_Label_Noise_Cleansing_for_Face_Recognition_CVPR_2020_paper.html.

[28] Z. Zhu, G. Huang, J. Deng, Y. Ye, J. Huang, X. Chen, J. Zhu, T. Yang, J. Lu, D. Du, J. Zhou, Webface260m: A benchmark unveiling the power of million-scale deep face recognition, IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19–25, 2021, Computer Vision Foundation/IEEE 2021, pp. 10492–10502, https://openaccess.thecvf.com/content/CVPR2021/html/Zhu_WebFace260M_A_Benchmark_Unveiling_the_Power_of_Million-Scale_Deep_Face_CVPR_2021_paper.html.

[29] G. Bae, M. de La Gorce, T. Baltrusaitis, C. Hewitt, D. Chen, J.P.C. Valentin, R. Cipolla, J. Shen, Digiface-1m: 1 million digital face images for face recognition, IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2023, Waikoloa, HI, USA, January 2–7, 2023, IEEE 2023, pp. 3515–3524, https://doi.org/10.1109/WACV56688.2023.00352.

[30] ISO/IEC JTC1 SC37 Biometrics, ISO/IEC 19795–1:2021 Information technology — Biometric performance testing and reporting — Part 1: Principles and framework (2021).

[31] K. Mallat, N. Damer, F. Boutros, A. Kuijper, J. Dugelay, Cross-spectrum thermal to visible face recognition based on cascaded image synthesis, 2019 International Conference on Biometrics, ICB 2019, Crete, Greece, June 4–7, 2019, IEEE 2019, pp. 1–8, https://doi.org/10.1109/ICB45273.2019.8987347.

[32] F. Boutros, N. Damer, A. Kuijper, Quantface: Towards lightweight face recognition by synthetic data low-bit quantization, 26th International Conference on Pattern Recognition, ICPR 2022, Montreal, QC, Canada, August 21–25, 2022, IEEE 2022, pp. 855–862, https://doi.org/10.1109/ICPR56361.2022.9955645.

[33] N. Damer, C.A.F. López, M. Fang, N. Spiller, M.V. Pham, F. Boutros, Privacy-friendly synthetic data for the development of face morphing attack detectors, IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19–20, 2022, IEEE 2022, pp. 1605–1616, https://doi.org/10.1109/CVPRW56347.2022.00167.

[34] ISO/IEC JTC 1/SC 27 Information security, cybersecurity and privacy protection, ISO/IEC 20889:2018 Privacy enhancing data de-identification terminology and classification of techniques (2018). https://www.iso.org/standard/69373.html.

[35] N. Damer, A.M. Saladie, A. Braun, A. Kuijper, Morgan: Recognition vulnerability and attack detectability of face morphing attacks created by generative adversarial network, 9th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2018, Redondo Beach, CA, USA, October 22–25, 2018, IEEE 2018, pp. 1–10, https://doi.org/10.1109/BTAS.2018.8698563.

[36] H.H. Nguyen, J. Yamagishi, I. Echizen, S. Marcel, Generating master faces for use in performing wolf attacks on face recognition systems, 2020 IEEE International Joint Conference on Biometrics, IJCB 2020, Houston, TX, USA, September 28 - October 1, 2020, IEEE 2020, pp. 1–10, https://doi.org/10.1109/IJCB48548.2020.9304893.

[37] U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R.N.J. Veldhuis, L.J. Spreeuwers, M. Schils, D. Maltoni, P. Grother, S. Marcel, R. Breithaupt, R. Raghavendra, C. Busch, Biometric systems under morphing attacks: Assessment of morphing techniques and vulnerability reporting, in: A. Brömme, C. Busch, A. Dantcheva, C. Rathgeb, A. Uhl (Eds.), International Conference of the Biometrics Special Interest Group, BIOSIG 2017, Darmstadt, Germany, September 20–22, 2017, LNI, vol. P-270, GI/IEEE 2017, pp. 149–159, https://doi.org/10.23919/BIOSIG.2017.8053499.

[38] R. Tolosana, R. Vera-Rodríguez, J. Fiérrez, A. Morales, J. Ortega-Garcia, Deepfakes and beyond: A survey of face manipulation and fake detection, Inf. Fusion 64 (2020) 131–148, https://doi.org/10.1016/j.inffus.2020.06.014.

[39] B. Meden, P. Rot, P. Terhörst, N. Damer, A. Kuijper, W.J. Scheirer, A. Ross, P. Peer, V. Struc, Privacy-enhancing face biometrics: A comprehensive survey, IEEE Trans. Inf. Forensics Secur. 16 (2021) 4147–4183, https://doi.org/10.1109/TIFS.2021.3096024.

[40] G.B. Huang, M. Ramesh, T. Berg, E. Learned-Miller, Labeled faces in the wild: A database for studying face recognition in unconstrained environments, Tech. Rep. 07-49, University of Massachusetts, Amherst, 11 2007.

[41] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotsia, S. Zafeiriou, Agedb: The first manually collected, in-the-wild age database, 2017 IEEE CVPRW, CVPR Workshops 2017, Honolulu, HI, USA, July 21-26, 2017, IEEE Computer Society 2017, pp. 1997–2005, https://doi.org/10.1109/CVPRW.2017.250.

[42] S. Sengupta, J. Chen, C.D. Castillo, V.M. Patel, R. Chellappa, D.W. Jacobs, Frontal to profile face verification in the wild, 2016 IEEE Winter Conference on Applications of Computer Vision, WACV 2016, Lake Placid, NY, USA, March 7–10, 2016, IEEE Computer Society 2016, pp. 1–9, https://doi.org/10.1109/WACV.2016.7477558.

[43] T. Zheng, W. Deng, J. Hu, Cross-age LFW: A database for studying cross-age face recognition in unconstrained environments, CoRR abs/1708.08197 (2017). arXiv:1708.08197. http://arxiv.org/abs/1708.08197.

[44] T. Zheng, W. Deng, Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments, Tech. Rep. 18–01, Beijing University of Posts and Telecommunications, February 2018.

[45] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, W. Liu, Cosface: Large margin cosine loss for deep face recognition, 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018, Computer Vision Foundation/IEEE Computer Society 2018, pp. 5265–5274, https://doi.org/10.1109/CVPR.2018.00552, http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_CosFace_Large_Margin_CVPR_2018_paper.html.

[46] J.N. Kolf, T. Rieber, J. Elliesen, F.B.A. Kuijper, N. Damer, Identity-driven three-player generative adversarial network for synthetic-based face recognition, IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2023, Canada, June 19–20, 2023, IEEE, 2023.

[47] D.P. Kingma, M. Welling, Auto-encoding variational bayes, in: Y. Bengio, Y. LeCun (Eds.), 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings, 2014. http://arxiv.org/abs/1312.6114.

[48] I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A.C. Courville, Y. Bengio, Generative adversarial nets, in: Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, K.Q. Weinberger (Eds.), Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8–13 2014, Montreal, Quebec, Canada, 2014, pp. 2672–2680. https://proceedings.neurips.cc/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html.

[49] A. van den Oord, N. Kalchbrenner, K. Kavukcuoglu, Pixel recurrent neural networks, in: M. Balcan, K.Q. Weinberger (Eds.), Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19–24, 2016, JMLR Workshop and Conference Proceedings, vol. 48, JMLR.org 2016, pp. 1747–1756, http://proceedings.mlr.press/v48/oord16.html.

[50] I. Kobyzev, S.J.D. Prince, M.A. Brubaker, Normalizing flows: An introduction and review of current methods, IEEE Trans. Pattern Anal. Mach. Intell. 43 (11) (2021) 3964–3979, https://doi.org/10.1109/TPAMI.2020.2992934.

[51] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual, 2020. https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

[52] M. Kowalski, S.J. Garbin, V. Estellers, T. Baltrusaitis, M. Johnson, J. Shotton, CONFIG: controllable neural face image generation, in: A. Vedaldi, H. Bischof, T. Brox, J. Frahm (Eds.), Computer Vision - ECCV 2020–16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI, Lecture Notes in Computer Science, vol. 12356, Springer 2020, pp. 299–315, https://doi.org/10.1007/978-3-030-58621-8_18.

[53] S. Bond-Taylor, A. Leach, Y. Long, C.G. Willcocks, Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive

models, IEEE Trans. Pattern Anal. Mach. Intell. 44 (11) (2022) 7327–7347, https://doi.org/10.1109/TPAMI.2021.3116668.
[54] T. Karras, S. Laine, T. Aila, A style-based generator architecture for generative adversarial networks, IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019, Computer Vision Foundation/IEEE 2019, pp. 4401–4410, https://doi.org/10.1109/CVPR.2019.00453, http://openaccess.thecvf.com/content_CVPR_2019/html/Karras_A_Style-Based_Generator_Architecture_for_Generative_Adversarial_Networks_CVPR_2019_paper.html.
[55] P.J. Tinsley, A. Czajka, P.J. Flynn, This face does not exist.but it might be yours! identity leakage in generative models, IEEE Winter Conference on Applications of Computer Vision, WACV 2021, Waikoloa, HI, USA, January 3–8, 2021, IEEE 2021, pp. 1319–1327, https://doi.org/10.1109/WACV48630.2021.00136.
[56] Y. Shen, P. Luo, J. Yan, X. Wang, X. Tang, Faceid-gan: Learning a symmetry three-player GAN for identity-preserving face synthesis, 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018, Computer Vision Foundation/IEEE Computer Society 2018, pp. 821–830, https://doi.org/10.1109/CVPR.2018.00092, http://openaccess.thecvf.com/content_cvpr_2018/html/Shen_FaceID-GAN_Learning_a_CVPR_2018_paper.html.
[57] Y. Deng, J. Yang, D. Chen, F. Wen, X. Tong, Disentangled and controllable face image generation via 3d imitative-contrastive learning, 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020, Computer Vision Foundation/IEEE 2020, pp. 5153–5162, https://doi.org/10.1109/CVPR42600.2020.00520, https://openaccess.thecvf.com/content_CVPR_2020/html/Deng_Disentangled_and_Controllable_Face_Image_Generation_via_3D_Imitative-Contrastive_Learning_CVPR_2020_paper.html.
[58] A. Shoshan, N. Bhonker, I. Kviatkovsky, G.G. Medioni, Gan-control: Explicitly controllable gans, 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021, IEEE 2021, pp. 14063–14073, https://doi.org/10.1109/ICCV48922.2021.01382.
[59] Y. Shen, C. Yang, X. Tang, B. Zhou, Interfacegan: Interpreting the disentangled face representation learned by gans, IEEE Trans. Pattern Anal. Mach. Intell. 44 (4) (2022) 2004–2018, https://doi.org/10.1109/TPAMI.2020.3034267.
[60] H. Zhang, S. Venkatesh, R. Ramachandra, K.B. Raja, N. Damer, C. Busch, MIPGAN - generating strong and high quality morphing attacks using identity prior driven GAN, IEEE Trans. Biom. Behav. Identity Sci. 3 (3) (2021) 365–383.
[61] N. Damer, M. Fang, P. Siebke, J.N. Kolf, M. Huber, F. Boutros, Mordiff: Recognition vulnerability and attack detectability of face morphing attacks created by diffusion autoencoders, 11th IEEE International Workshop on Biometrics and Forensics, IWBF 2023, Barcelona, Spain, April 19–20, 2023, IEEE 2023, pp. 1–6.
[62] P. Terhörst, F. Bierbaum, M. Huber, N. Damer, F. Kirchbuchner, K.B. Raja, A. Kuijper, On the (limited) generalization of masterface attacks and its relation to the capacity of face representations, IEEE International Joint Conference on Biometrics, IJCB 2022, Abu Dhabi, United Arab Emirates, October 10–13, 2022, IEEE 2022, pp. 1–9, https://doi.org/10.1109/IJCB54206.2022.10007976.
[63] M. Huber, F. Boutros, A.T. Luu, K.B. Raja, R. Ramachandra, N. Damer, P.C. Neto, T. Gonçalves, A.F. Sequeira, J.S. Cardoso, J. Tremoço, M. Lourenço, S. Serra, E. Cermeño, M. Ivanovska, B. Batagelj, A. Kronovsek, P. Peer, V. Struc, SYN-MAD 2022: Competition on face morphing attack detection based on privacy-aware synthetic training data, IEEE International Joint Conference on Biometrics, IJCB 2022, Abu Dhabi, United Arab Emirates, October 10–13, 2022, IEEE 2022, pp. 1–10, https://doi.org/10.1109/IJCB54206.2022.10007950.
[64] M. Fang, M. Huber, N. Damer, Synthaspoof: Developing face presentation attack detection based on privacy-friendly synthetic data, IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2023, Canada, June 19–20, 2023, IEEE, 2023.
[65] V. Mirjalili, S. Raschka, A. Ross, Flowsan: Privacy-enhancing semi-adversarial networks to confound arbitrary face-based gender classifiers, IEEE Access 7 (2019) 99735–99745, https://doi.org/10.1109/ACCESS.2019.2924619.
[66] P. Terhörst, M. Huber, N. Damer, P. Rot, F. Kirchbuchner, V. Struc, A. Kuijper, Privacy evaluation protocols for the evaluation of soft-biometric privacy-enhancing technologies, in: A. Brömme, C. Busch, A. Dantcheva, K.B. Raja, C. Rathgeb, A. Uhl (Eds.), BIOSIG 2020 - Proceedings of the 19th International Conference of the Biometrics Special Interest Group, online, 16.-18. September 2020, LNI, vol. P-306, Gesellschaft für Informatik e.V. 2020, pp. 215–222, https://dl.gi.de/20.500.12116/34330.