

Identity-driven Three-Player Generative Adversarial Network for Synthetic-based Face Recognition

Jan Niklas Kolf^{1,2}, Tim Rieber¹, Jurek Elliesen^{1,2}, Fadi Boutros¹, Arjan Kuijper^{1,2}, Naser Damer^{1,2}

¹Fraunhofer Institute for Computer Graphics Research IGD, Darmstadt, Germany

²Department of Computer Science, TU Darmstadt, Darmstadt, Germany

Email: jan.niklas.kolf@igd.fraunhofer.de

Abstract

Many of the commonly used datasets for face recognition development are collected from the internet without proper user consent. Due to the increasing focus on privacy in the social and legal frameworks, the use and distribution of these datasets are being restricted and strongly questioned. These databases, which have a realistically high variability of data per identity, have enabled the success of face recognition models. To build on this success and to align with privacy concerns, synthetic databases, consisting purely of synthetic persons, are increasingly being created and used in the development of face recognition solutions. In this work, we present a three-player generative adversarial network (GAN) framework, namely IDnet, that enables the integration of identity information into the generation process. The third player in our IDnet aims at forcing the generator to learn to generate identity-separable face images. We empirically proved that our IDnet synthetic images are of higher identity discrimination in comparison to the conventional two-player GAN, while maintaining a realistic intra-identity variation. We further studied the identity link between the authentic identities used to train the generator and the generated synthetic identities, showing very low similarities between these identities. We demonstrated the applicability of our IDnet data in training face recognition models by evaluating these models on a wide set of face recognition benchmarks. In comparison to the state-of-the-art works in synthetic-based face recognition, our solution achieved comparable results to a recent rendering-based approach and outperformed all existing GAN-based approaches. The training code and the synthetic face image dataset are publicly available ¹.

1. Introduction

Biometrics is a fast growing technology that recognizes people based on their physical or behavioral characteris-

tics [22]. One of the most commonly used modalities is the face, which is widely accepted by the population and can be captured without major hurdles [38]. Face Recognition (FR) has become part of the everyday life of many end users, e.g. for unlocking smartphones. The breakthrough of robust FR systems in recent years is partly due to Deep Neural Networks [9, 18], which have, in combination with specific loss functions [3, 4, 12, 43], significantly increased the recognition rates of these systems [34, 44]. However, DNN training requires a large amount and variation of data. Many of these datasets were compiled by crawling images from the Internet, thus raising both ethical and legal concerns [10, 17, 37].

The issues of privacy and data autonomy are addressed in both an ethical and a legal framework. In the ethical context, Smith and Miller [37] relate privacy to the moral value of autonomy. This autonomy refers to the right to decide who is allowed to access personal information or data and how it is processed, including biometric data such as face images. This standpoint is covered by the European Union's General Data Protection Regulation (GDPR) [42].

In the GDPR, biometric data is classified as special data requiring protection, thus it is subject to strong data protection regulations concerning acquirement and storage [41]. The processing of such data is restricted and the possibility of data withdrawal and control should be guaranteed [40]. However, these legal restrictions apply only to natural persons but not to synthetic identities such as characters drawn by artists or images of subjects generated by a machine [39].

To address ethical and legal concerns, the field of synthetic data for biometric development is advancing increasingly [14, 27]. The recent studies in this direction aimed at creating datasets for the development of FR [2, 6, 7], along with other FR system components, without the need for images of authentic identities [5, 11, 15, 21]. The generation of synthetic data is primarily driven by Generative Adversarial Networks (GAN). These are able to generate photo-realistic results based on training a two-player min-max game between a generator and a discriminator. Vari-

¹<https://github.com/fdbtrs/Synthetic-Face-Recognition>

ous approaches using GAN for the generation of synthetic identities with respect to FR have been presented very recently, such as DigiFace-1M [2], SynFace [33], SFace [6] or USynthFace [7]. Although these approaches presented very promising FR results, they have some limitations. SFace [6] suffers from relatively low identity separability in comparison to authentic data, i.e. the identities are not highly distinct and thus genuine (same identity) comparisons might be in some cases confused to be imposter (different identity) comparisons. SynFace [33] mixed up authentic and synthetic data in FR training, aiming at increasing intra-class variations. USynthFace proposed to mitigate the low intra-class variations by proposing aggressive data augmentation to train unsupervised FR models. However, USynthFace is limited to unsupervised learning. DigiFace-1M utilized a digital rendering pipeline to construct synthetic images. DigiFace-1M [2] is extremely computationally costly, limiting its ability on generating large-scale synthetic data and such high computational demand might not be available for research, along with the low FR performance it produces without the dependency on sophisticated augmentation.

In this work, we propose IDnet, which extends the traditional class conditioned GAN by adding an additional identity-dedicated third player (ID-3) to the minimax game. This component overcomes the drawbacks of the class conditioned GAN (as represented in SFace [6]) by enforcing specific identity-discriminant information in the generation process. In several experiments, we show that this additional component causes the generator of the GAN to produce images of synthetic identities that stronger resemble the distribution of authentic data, both in terms of class separability and intra-class variance in comparison to the considered baseline SFace. We show that FR systems trained on synthetic datasets generated with IDnet outperform those trained on datasets synthesized with our baseline, SFace, and other GAN-based synthetic datasets.

2. Related Work

Current state-of-the-art (SOTA) FR models such as ElasticFace [3], ArcFace [12], and CosFace [43] are trained on authentic datasets such as CASIA-Webface [47] and MS1MV2 [12, 17]. Such models regularly achieve new record breaking results on the authentic benchmarks [3] such as LFW [19] or AgeDB-30 [32]. With the goal of generating identity-specific face images, various methods are proposed in the literature, with most using GAN as the basis for generating synthetic data. Marriott *et al.* [30] were among the first to evaluate the capabilities of GAN for identity-based applications such as facial recognition. In their work, they have shown that GAN can be used to create synthetic identities that are not included in the training set. By introducing a special triplet loss, the authors were also able to increase the identity disentanglement, which

means that the identity information can be separated from the other image properties in the generation process. DiscoFaceGAN [13] explored an approach to create facial images from non-existing people with the use of multiple disentangled features as input for the generator. These features include identity, pose, illumination, and expression. To ensure that these features are disentangled, a contrastive loss is used.

Towards training FR based on synthetic data, SynFace [33] aimed at training FR models with the use of synthetic data. It investigated the performance of DiscoFaceGAN in terms of intra-class variance and the domain gap between real and generated images. SynFace improved these aspects of DiscoFaceGAN by introducing identity and domain mixups. Mixup uses combinations of image features to generate new faces. FaceID-GAN [36] introduced a classification network as an additional constraint to the GAN training. FaceID-GAN aims at generating variations of input authentic images rather than generating images of synthetic identities. SFace [6] proposed class-conditional synthetic GAN for class-labeled synthetic image generation. The authors utilize synthetic data to train supervised FR models, achieving very promising verification accuracies. USynthFace [7] showed that by using unlabeled synthetic data, a FR model can be trained successfully using contrastive learning and aggressive data augmentation. By rendering 3d human models that are varied in appearance, shape, and accessories, among others, the authors of DigiFace-1M [2] have created a synthetic dataset that exhibits high variability in the data generated for a synthetic identity. This allowed them to train FR models based on their proposed data, achieving relatively high verification accuracies when combined with sophisticated rendering technique. However, such a digital image rendering approach is computationally costly, limiting their ability in generating large-scale synthetic datasets.

While there has been great progress in the area of synthetic FR datasets generated by GAN, previous synthetic data used for FR training either suffer from low intra-class diversity or are of low identity separability. In this work, we propose a novel approach for generating synthetic images that are identity separable while containing realistic intra-class variations.

3. Methodology

We propose in this work an identity-conditioned generative model based on a Generative Adversarial Network (GAN) [16], namely IDnet. The conventional GAN are based on optimizing a minimax game between two players, the generator and the discriminator. The generator aims at fooling the discriminator by generating synthetic images that are similar to authentic images. The discriminator tries to distinguish between synthetic and authentic

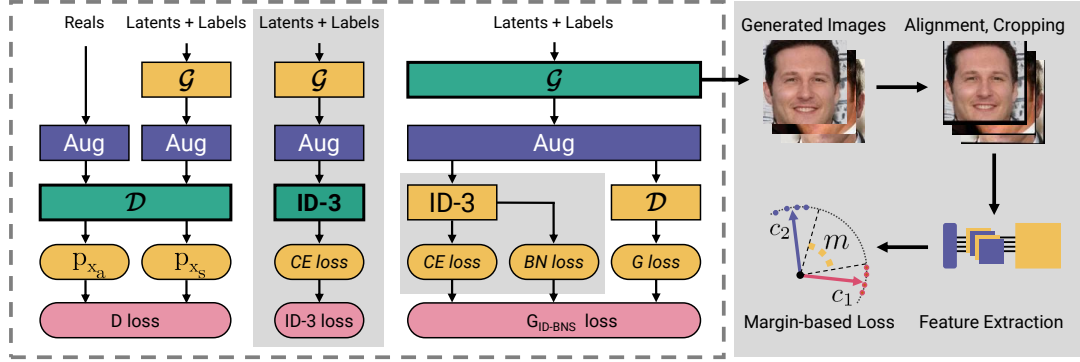


Figure 1. An overview of IDnet with the two-player StyleGAN2-ADA architecture as well as the third-player ID-3 extension presented in this paper, which are highlighted in gray. Three models are trained at the same time, whereby the individual components, which are updated by the appropriate loss, are marked as black bordered green boxes. The D loss is defined in Eq. (1), G loss in Eq. (2), ID-3 loss in Eq. (3), and G_{ID-BNS} loss in Eq. (6). Using the trained generator, a synthetic face dataset is created, which is aligned and cropped, and on which a FR model with margin-penalty based loss function is trained.

images [16, 26]. As these two models compete against each other, the discriminator pushes the generator to learn to estimate the probability distribution of authentic training data, enabling the generation of realistic synthetic samples. We propose in this work to extend this game by a third player, an identity model, ID-3. The ID-3 aims at pushing the generator not only to learn the underlying probability distribution of authentic data but also to learn the identity information encoded in training samples, and thus, generating identity-separable samples of synthetic identity.

3.1. Two-Player GAN

We start here by introducing the needed formulation and background to our proposed solution. GAN was originally proposed by Goodfellow *et al.* [16], aiming at generating synthetic data by sampling latent vectors from a given random distribution e.g. Gaussian distribution, and feeding it into the generator. GAN training is based on training two players, the generator \mathcal{G} and the discriminator \mathcal{D} [16, 26]. \mathcal{G} receives a latent vector \mathbf{z} as input and outputs a synthetic image $\mathbf{x}_s = \mathcal{G}(\mathbf{z})$. The training objective of \mathcal{G} is to learn to generate synthetic images of the same probability distribution of authentic training data. The training objective of \mathcal{D} is to distinguish between authentic images \mathbf{x}_a and \mathbf{x}_s . In a binary classification problem, \mathcal{D} trains to estimate the probability of input \mathbf{x} being from authentic or synthetic distributions. Formally, \mathcal{D} 's training objective can be defined as follows:

$$L_D = \log(1 + e^{-p_{x_a}}) + \log(1 + e^{p_{x_s}}), \quad (1)$$

where $p_{x_a} = \mathcal{D}(\mathbf{x}_a)$ and $p_{x_s} = \mathcal{D}(\mathbf{x}_s)$ are the probabilities of \mathbf{x}_a and \mathbf{x}_s being from authentic or synthetic distributions, respectively.

\mathcal{G} aims at fooling \mathcal{D} by generating realistic samples that would minimize the logarithm of the inverse probability p_{x_s} (predicted by \mathcal{D}). As \mathcal{D} and \mathcal{G} compete against each other, the training loss of \mathcal{G} is defined as follows:

$$L_G = \log(1 + e^{-p_{x_s}}), \quad (2)$$

where p_{x_s} is \mathcal{D} prediction i.e. prediction of \mathbf{x}_s is being from synthetic distribution. The only condition on conventional GAN is that the generated images tend to have a similar probability distribution to the authentic training data. Several previous works [13, 25, 26] proposed to extend conventional GAN with conditional mechanisms, enabling synthetic image generation with certain attributes. One of the widely used conditional mechanisms is class-conditional GAN. In this case, class labels are used as additional input to \mathcal{G} and \mathcal{D} to force \mathcal{G} to learn to generate synthetic images of a specific class. SFace [6], which is based on class-conditional StyleGAN2-ADA [24], trained GAN under class-conditional settings to generate face images of a specific class label. In this case, a class label (c) is embedded into a 512-D vector and then is concatenated with the input latent vector \mathbf{z} (\mathbf{z} is of 512-D) to generate class-related synthetic images ($\mathbf{x}_s = \mathcal{G}(\mathbf{z}, c)$). In class-conditional \mathcal{D} , a class label (c) is embedded into a 512-D vector and then concatenated with the final embedding layer of \mathcal{D} . The reported results by SFace demonstrated that SFace synthetic images are realistic with large intra-class variations. However, as we present in this paper (Section 5), the generated images by SFace suffer from relatively low identity separability which might lead to less optimal face verification accuracies when such synthetic data is used to train FR. In such a training paradigm, GAN is guided to learn to generate synthetic images from a specific class, however, not from a specific identity. To overcome this challenge, we introduce a third player to GAN, an identity network ID-3. Our new player aims at guiding \mathcal{G} to learn to generate synthetic images that are highly identity-separable.

3.2. Identity Network as a Third-Player

The third player in the minimax game is noted as ID-3 which acts as a FR model. ID-3 is trained with a margin-

penalty softmax loss that incorporates margin penalty in softmax loss to push training samples to be close to their class centers and far away from other class centers. Specifically, we use CosFace [43] loss to train ID-3. Also, ID-3 loss is used as an additional loss of \mathcal{G} loss to guide \mathcal{G} to generate synthetic images that are identity separable. The ID-3 loss on a batch of synthetic images \mathbf{x}_s are given as follows:

$$L_{ID-3} = -\frac{1}{N} \sum_{i \in N} \log \frac{e^{s(\cos(\theta_{y_i}) - m)}}{e^{s(\cos(\theta_{y_i}) - m)} + \sum_{j=1, j \neq y_i}^C e^{s \cos(\theta_j)}}, \quad (3)$$

where m is a margin penalty, s is a scaling term, N is the batch size, C is the number of classes and θ_{y_i} is the angle between sample y_i and its i -th class center. Synthetic samples in the early stage of GAN training tend to be less realistic with low identity information encoded in the generated images. At an early stage of IDnet training, pushing such images to their class centers with a large margin penalty value will affect ID-3 training stability. Thus, we propose a progressive margin-penalty value to stabilize ID-3 training. Initially, m is set to $m = 0$ and it is incrementally increased by a small value of 0.05 every 7 epochs with a maximum of 0.35 [43]. During our three-player game, only synthetic images \mathbf{x}_s (generated by \mathcal{G}) are used to train ID-3, and the ground-truth labels are their corresponding class labels (c). We used a pretrained ResNet-50 on CASIA-WebFace [47] with CosFace [43] as a backbone for our ID-3. We froze all the network weights and only train the weights of the classification layer. The loss function of \mathcal{G} in our three-player game is given by:

$$L_{G_{ID}} = L_G + L_{ID-3}. \quad (4)$$

Domain Adaptation (DA) We propose in this work to further minimize the domain gap between synthetic and authentic data distributions by matching Batch Normalization Statistics (BNS) of authentic and synthetic data. Following [46], we first extracted means μ_a and standard deviations σ_a of all batch normalization (BN) layers of the ID-3 backbone (trained on authentic data). During our three-player GAN training, we calculated the means μ_s and the standard deviations σ_s of all BN layers by passing a batch of synthetic data into ID-3 and then extracted BN statistics. Finally, we calculate the difference between BN statistics from authentic and synthetic data as follows:

$$L_{BNS}(\mu_s, \sigma_s) = \sum_{l \in BNL} \|\mu_s^l - \mu_a^l\|_2^2 + \|\sigma_s^l - \sigma_a^l\|_2^2, \quad (5)$$

where BNL are BN layers of ID-3. The L_{BNS} is used as an additional loss to G loss (defined in Eq. (4)). The final loss $L_{G_{ID-BNS}}$ for \mathcal{G} in our three-player GAN is given by:

$$L_{G_{ID-BNS}} = L_{G_{ID}} + \lambda * L_{ID-3} + \kappa * L_{BNS}, \quad (6)$$

where λ and κ are weighting terms for L_{ID-3} and L_{BNS} , respectively. The training algorithm of our IDnet is shown in Algorithm 1, the pipeline is shown in Figure 1.

Algorithm 1 Three-Player GAN IDnet Training Loop

```

 $m \leftarrow 0$ 
 $m_\delta \leftarrow 0.05$ 
for  $epoch$  in  $epochs$  do
   $m \leftarrow m + m_\delta$  if  $epoch \% 7 = 0$ 
  for  $batch$  in  $trainingset$  do
     $\mathbf{z} \leftarrow$  Sampled from Gaussian Dist.
     $c \leftarrow$  Randomly sampled from  $C$ 
     $\mathbf{x}_s \leftarrow \mathcal{G}(\mathbf{z}, c)$ 
     $p_{x_s} \leftarrow \mathcal{D}(\mathbf{x}_s, c)$ 
     $\mu_s, \sigma_s \leftarrow BNS(ID-3(\mathbf{x}_s))$ 
     $L_G \leftarrow L_G(p_{x_s})$ 
     $L_{ID-3} \leftarrow L_{ID-3}(p_c, c)$ 
     $L_{BNS} \leftarrow L_{BNS}(\mu_s, \sigma_s)$ 
     $backward(\mathcal{G}, L_{G_{ID-BNS}}(L_G, L_{ID-3}, L_{BNS}))$ 
     $\mathbf{x}_a, c \leftarrow$  batch
     $p_{x_a} \leftarrow \mathcal{D}(\mathbf{x}_a, c)$ 
     $p_{x_s} \leftarrow \mathcal{D}(\mathbf{x}_s, c)$ 
     $backward(\mathcal{D}, L_D(p_{x_a}, p_{x_s}))$ 
     $p_c \leftarrow ID-3(\mathbf{x}_s, c)$ 
     $backward(ID-3, L_{ID-3}(p_c, c))$ 
     $update(\mathcal{D})$ 
     $update(\mathcal{G})$ 
     $update(ID-3)$ 
     $zerograd()$ 
  end for
end for

```

4. Experimental Setup

Dataset: We use CASIA-Webface [47] to train our three-player GAN (as described in Section 3). CASIA-Webface consists of 494, 414 authentic images from 10, 572 different identities [47]. Multi-Task Cascaded Convolutional Networks (MTCNN) [48] is used to extract the facial landmarks based on which the faces are aligned and cropped to the training size of $128 \times 128 \times 3$, as described in [12].

IDnet Training Settings: In this work we utilize StyleGAN2-ADA [24], as used in SFace [6], as our two-player GAN foundation. StyleGAN2-ADA improves StyleGAN2 [26] by adaptive dataset augmentations (ADA). These augmentation methods are applied to the authentic and synthetic images, allowing more stable training, especially when a small dataset is used for training. StyleGAN2-ADA is extended with the additional identity network and respective losses of our three-player minimax game.

The generator and discriminator of StyleGAN2-ADA are based on Progressive GANs [23] as originally presented in [24]. The dimensionality of the style vectors and noise is set to 512. The dimensionality of the conditioning and therefore the number of synthetic identities is set to be equal to the number of identities in CASIA-Webface, with

$C = 10,572$, following [6]. The mapping architecture consists of 8 fully connected layers. The activation function for layers in the generator is Leaky ReLU [28] with $\alpha = 0.2$. The generator outputs images of size $128 \times 128 \times 3$. We follow the training setup described in [6] and [24]. The loss function of the discriminator is the non-saturating loss as described in [16] with R_1 regularization [31]. For the additional components of the generator’s loss function (see formula 6), the weights $\lambda = 0.05$ and $\kappa = 0.1$ are chosen. As optimizer for the generator and discriminator, Adam is used with the parameters $\beta_1 = 0$, $\beta_2 = 0.99$ and $\epsilon = 10^{-8}$. The learning rate is set to 0.0025. The backbone of ID-3 is ResNet-50 [18] with an embedding size of 512. The synthetic images are resized to $112 \times 112 \times 3$ using bilinear interpolation to match ID-3 input size. We set s to 64 [43] in L_{ID-3} and use a progressively growing margin as described in Section 3. The three-player GAN is trained for 50 epochs. The optimizer of ID-3 is Stochastic Gradient Descent with a learning rate of 0.02 and momentum of 0.9. The minibatch size is set to 32. IDnet is implemented using PyTorch based on the official PyTorch implementation [1] of StyleGAN2-ADA. The models are trained on a Linux machine (Ubuntu 20.04.2 LTS) with Intel(R) Xeon(R) Gold 5218 CPU 2.30GHz, 512GB RAM and 4 Nvidia GeForce RTX 6000 GPUs.

Class separability: In order to train a FR model with synthetic data, it is assumed that the synthetic data should resemble the distribution of authentic data. Authentic identities can be separated by FR systems [3], to a certain degree, as they each have characteristics that make them distinguishable from other authentic identities. Higher separability between genuine and imposter comparison scores indicates higher identity discrimination. Thus, with higher identity discrimination the False Non-Match Rates (FNMR) and False Match Rate (FMR) are lower at a chosen decision threshold than those error rates that result from data with low identity discrimination. FNMR and FMR are based on the ISO/IEC 19795-1 [29] standard. We use FMR100, FMR1000 and Equal Error Rate (EER) as class separability evaluation metrics. The FMR100 gives the lowest FNMR at the operation point $FMR \leq 1.0\%$, the FMR1000 gives the lowest FNMR at $FMR \leq 0.1\%$. The Equal Error Rate (EER) [29] is FMR or FNMR at the operation point at which they are equal.

Based on that, using EER, FMR, and FNMR, we can measure how the identity discrimination in our IDnet data compares to the authentic data and to the SFace baseline. For evaluation, 10 images per synthetic identity are generated for both IDnet and SFace. This results in two synthetic datasets with 105,720 images each. For each synthetic image as well as each image from CASIA-Webface, a 512 dimensional embedding is extracted using two pre-trained FR models, ResNet-100 trained with the Elastic-ArcFace

loss [3], and the CurricularFace loss [20]. Individual pairwise comparisons are formed from the embeddings. The cosine similarity is used as a similarity measure. Comparisons of two face embeddings of the same identity (class label) are genuine comparisons, while comparisons between different identities are imposter comparisons.

Each image of an identity is compared to all other images of that respective identity as a genuine comparison. This image is also compared to 100 randomly sampled images of other identities, resulting in 100 imposter scores for each image. The genuine and imposter scores distributions are plotted for authentic CASIA-Webface, synthetic SFace and IDnet, respectively. For each method, the EER, FMR100, FMR1000 are calculated and reported.

Intra-Class Variance: To match the distribution of authentic face data it is not only necessary to achieve a similar class separability, but also to resemble the intra-class variability of authentic identities. Among other things, variations in the appearance or different lighting conditions can change the visual appearance of an authentic identity. A FR model must learn how to determine the identity even in the presence of large visual variation. Therefore the generated synthetic data should also have realistic variation within a synthetic identity. To evaluate the intra-class variance of a dataset, the variance in the face embeddings of one identity is calculated:

$$ICV = \frac{1}{N * (N - 1)} \sum_{i=1}^N \sum_{j=i+1}^N \frac{1}{D} \sum_{k=1}^D \|f_{i,k} - f_{j,k}\|, \quad (7)$$

where N is the number of images of one identity and D is the feature dimensionality of a face embedding. The ICV is calculated for every identity in the given dataset and summed up. The ICV score is calculated for the authentic dataset CASIA-Webface and the synthetic datasets generated by SFace and IDnet.

The Link between Authentic and Synthetic Identities:

To ensure the privacy motivation behind the generated synthetic data, there should be no major identity linkage between authentic identities from the IDnet training dataset, CASIA-Webface, and the generated synthetic identities from our proposed IDnet. To investigate the possible identity linkage, an additional experiment is conducted following those proposed in [6].

For each identity of CASIA-Webface, the first two images are selected as reference images. Probe images with the same class label are selected from CASIA-Webface, SFace, and IDnet. The reference images are compared to the probes of the respective dataset, resulting in score distributions with CASIA-Webface references vs. CASIA-Webface probes, CASIA-Webface references vs. SFace probes and CASIA-Webface references vs IDnet probes.

The stronger the identity linkage between CASIA-Webface and SFace or IDnet, the stronger the distributions of SFace probes and IDnet probes shifts towards the CASIA-Webface genuine probes distribution. The comparison scores, here and in the intra-class variance analyses, are a result from comparing face templates produced by a publicly released ResNet-100 trained on Elastic-ArcFace loss [3].

Face recognition based on IDnet: The primary use of our synthetic data is to train a FR model with more shareable, scalable, and privacy-friendly data. Therefore, it is important to investigate to what extent a FR model trained on the synthetic data is able to perform on benchmarks consisting of authentic images. For this purpose, ResNet-50 models are trained with CosFace on synthetic data generated by our IDnet with parameters following [3,6]. The chosen parameter for dropout is 0.4, as embedding size 512 is used. The CosFace margin is set as $m = 0.35$ and the scale parameter s to $s = 64$. The minibatch size is 512. Stochastic Gradient Descent with a learning rate of $1e - 1$ is used as an optimizer. Momentum is 0.9 and weight decay $5e - 4$. The learning rate is divided by ten at the 22nd, 30th and 40th epochs, all models are trained for 40 epochs. Per identity, 10, 20, 40, 50 and 60 images are randomly generated as proposed in [6]. All images are aligned with the previously mentioned MTCNN algorithm [48]. During the training, the input images are augmented using the random augmentation methods introduced in [8]. As authentic image benchmarks Labeled Faces in the Wild (LFW) [19], Crossage LFW (CALFW) [50], Cross-Pose LFW (CPLFW) [49], Celebrities in Frontal-Profile in the Wild (CFP-FP) [35] and AgeDB30 [32] are used. We follow the evaluation protocol specified in the benchmarks. For all benchmarks, the verification performance is given as accuracy [%]. We compare our work to recent works of SFace [6], USynthFace [7], SynFace [33] and DigiFace-1M [2] that have utilized ResNet-50 as a backbone model as well.

5. Results

Investigating Class Separability: Figure 3 shows the genuine and imposter score distributions of CASIA-Webface, SFace, and IDnet. In our solution IDnet (Fig. 3c) the genuine and imposter distribution are clearly more separated than our baseline SFace (Fig. 3b) and are more similar to the distributions of the authentic data (CASIA-Webface, Fig. 3a). This indicates that IDnet enforces the GAN to generate images with higher identity distinctiveness. Quantitatively this is shown in Table 1 as EER, FMR100 and FMR1000. While the baseline SFace has a significantly higher EER, FMR100 and FMR1000 than CASIA-Webface, our solution IDnet results in more comparable values to CASIA-Webface. This shows that our solution generates data that possess higher identity discrimination than that of SFace.

| | EER | FMR 100 | FMR 1000 |
|--------------------------------|-------|---------|----------|
| CASIA-WebFace (ElasticFace) | 0.062 | 0.084 | 0.118 |
| CASIA-WebFace (CurricularFace) | 0.063 | 0.085 | 0.120 |
| SFace (ElasticFace) | 0.216 | 0.623 | 0.839 |
| SFace (CurricularFace) | 0.226 | 0.628 | 0.848 |
| IDnet w/o DA (ElasticFace) | 0.123 | 0.354 | 0.610 |
| IDnet w/o DA (CurricularFace) | 0.128 | 0.354 | 0.604 |
| IDnet w/ DA (ElasticFace) | 0.085 | 0.213 | 0.400 |
| IDnet w/ DA (CurricularFace) | 0.085 | 0.204 | 0.380 |

Table 1. Verification accuracy metrics indicating the class separability in each dataset, using two FR models. IDnet with DA is compared to IDnet without DA, to SFace [6], and to the authentic CASIA-WebFace. Note that IDnet produces the most similar results to the authentic data.

| Training Set | LFW [%] | AgeDB30 [%] | CFP-FP [%] | CA-LFW [%] | CP-LFW [%] | Avg. [%] |
|-----------------|---------|-------------|------------|------------|------------|----------|
| IDnet-50 w/ DA | 84.83 | 63.58 | 70.43 | 71.50 | 67.35 | 71.54 |
| IDnet-50 w/o DA | 81.33 | 58.48 | 63.94 | 68.15 | 63.87 | 67.15 |

Table 2. Ablation on domain adaptation (DA) through Batch Normalization Statistics (BNS), as described in Section 3. FR models are trained on datasets of IDnet with and without DA with 50 images per ID. The results show the benefit of the DA used in IDnet. Both sets contain 528K images (50 images/ID) and the training is performed without augmentation.

Intra-Class Variance: The ICV score introduced in Section 4 is used to determine the variance between images of the same identity. CASIA-Webface scores 8.8×10^{-4} , our baseline SFace 14.9×10^{-4} and our proposed solution IDnet 11.5×10^{-4} . While intra-class variation can be beneficial for the training of FR models, a very high ICV can also indicate that images of the same identity label are very different and that they may belong to more than one distinct identity. Therefore, the ICV value of face data used to train a FR system is a trade-off between intra- and inter-class variations, and thus a higher or a lower ICV value is not necessarily more beneficial. However, what matters here is to represent realistic conditions (i.e. similar to authentic data), which our IDnet achieves by scoring a more similar ICV to the authentic data when compared to that of the SFace baseline. In Figure 2 images of the same identity index, generated by SFace and our IDnet are shown. It is visually noticeable and quantitatively shown by the similarity scores that images of the identity generated by IDnet more distinctively maintain the targeted identity label. This implies that the high ICV score of SFace might be due to a lack of distinct identity information within images of the respective identity. The IDnet solution proposed in this work reduces this generative miss-labeling.

Evaluating Similarity of Authentic and Synthetic Identities: Figure 4 shows the genuine comparison score distributions resulting from comparison pairs that consist of CASIA-Webface references and probes of the same class label either from authentic CASIA-Webface (blue), synthetic SFace (red), or IDnet (green). If the comparison scores of involving synthetic probes (red or green) lay in the same




Figure 2. Example images of a specific identity (#844), generated by SFace (a-j) and our IDnet method (k-t). Image (a) is compared to probe images (b-j), likewise image (k) is compared pairwise to (l-t). The cosine similarity scores are given for each comparison. The average score for SFace is 0.316 and for IDnet it is 0.467, respectively. From the similarity scores it is evident, that IDnet generates images with significantly better distinct identity information than SFace.




Figure 3. The genuine and imposter score distributions on respectively on the authentic dataset CASIA-Webface (a), on data created by SFace (b) and data created by IDnet (c). The small pump in the CASIA-Webface genuine distribution might be related to previously reported miss-labels [45]. In comparison to SFace distributions, the distributions of IDnet are more similar to the authentic data in (a).

range of that of the authentic probes (blue), then one would conclude that the comparison scores between the authentic and synthetic images of the same labels are high enough to be considered genuine comparison, thus there is an identity link between the authentic and synthetic data. As Figure 4a shows, no significant overlap is present when comparing probes of CASIA-Webface with probes synthesized with SFace. The same can be observed for probes generated with IDnet (Fig. 4b). This implies that neither for SFace nor for IDnet the identity information of the authentic identity is strongly present in the synthetic identity, confirming the outcome of [6]. From this, it is concluded that adding the IDnet in the generation process also ensures that synthetic identities, not linked to the authentic ones, are generated. Therefore, the solution complies with the privacy-driven motivations behind our work.

Figure 4c shows the comparison score distributions when the probes are from SFace and IDnet. The distribution of IDnet is slightly shifted to the lower similarity range in comparison to the SFace distribution. This indicates that there is a slightly smaller identity link between IDnet faces (in comparison to the ones from SFace) and their respective authentic identities. Our proposed solution, therefore, provides an additional improvement for generating purely

synthetic identities over the baseline SFace.

The benefit of DA: To prove the validity of including the BNS DA in the IDnet design, we trained two instances of the IDnet, with and without the DA. We first show the effect on the genuine-imposter separability of the generated data in Table 1. The tables shows that the IDnet data without DA do produce higher EER, FMR100, and FMR1000 in comparison to the IDnet with DA. Additionally, as shown in Table 2, the FR models trained on the IDnet data generated with the DA achieved higher face verification performance across all considered benchmarks, in comparison to the one trained on data generated by the IDnet without DA. This proves the validity of our choice to include the DA in the IDnet design.

Verification Performance of FR trained on IDnet data:

The target application of our synthetic face images is the training of FR models. The results on the mainstream benchmark for the training of a ResNet-50 backbone (as detailed in Sec. 4) are shown in Table 3. Investigating involving a different number of samples per identity in the FR training, we notice that increasing the number of images per identity to more than 10 does not drastically affect the FR performance. In comparison, the SFace showed a clear increase in performance when the number of images per identity




Figure 4. Score distributions between CASIA-Webface references with probes with the same class label taken from CASIA-Webface (CASIA-Webface probes), SFace (SFace probes), and IDnet (IDnet probes). No significant overlap between distributions from authentic and synthetic probes is visible, indicating that SFace and IDnet are both generating synthetic identities with no significant overlap to authentic identities.

| Training Set | Images | #Images / ID | LFW [%] | AgeDB30 [%] | CFP-FP [%] | CA-LFW [%] | CP-LFW [%] | Avg. [%] |
|---------------------------|--------|--------------|-----------|-------------|------------|------------|------------|-----------|
| CASIA-WebFace | 494K | 46 | 99.55 | 94.55 | 95.31 | 93.78 | 89.95 | 94.63 |
| SFace-10 [6] | 105K | 10 | 87.13 | 63.30 | 68.84 | 73.47 | 66.82 | 71.91 |
| SFace-20 [6] | 211K | 20 | 90.50 | 69.17 | 73.33 | 76.35 | 71.17 | 76.10 |
| SFace-40 [6] | 423K | 40 | 91.43 | 69.87 | 73.10 | 76.92 | 73.42 | 76.95 |
| SFace-60 [6] | 634K | 60 | 91.87 | 71.68 | 73.86 | 77.93 | 73.20 | 77.71 |
| USynthFace [7] | 100K | 1 | 91.52 | 69.30 | 78.46 (3) | 75.35 | 71.93 | 77.31 |
| USynthFace [7] | 200K | 1 | 91.93 | 71.23 | 78.03 | 76.73 | 72.27 | 78.04 |
| USynthFace [7] | 400K | 1 | 92.23 | 71.62 | 78.56 (2) | 77.05 | 72.03 | 78.30 |
| SynFace [33] | 500K | 50 | 91.93 | 61.63 | 75.03 | 74.73 | 70.43 | 74.75 |
| DigiFace-1M [2] | 500K | 50 | 95.40 (1) | 76.97 (1) | 87.40 (1) | 78.62 | 78.87 (1) | 83.45 (1) |
| DigiFace-1M (No Aug.) [2] | 500K | 50 | 88.07 | 60.92 | 70.99 | 69.23 | 66.73 | 71.19 |
| IDnet-50 (No Aug.) [Our] | 528K | 50 | 84.83 | 63.58 | 70.43 | 71.50 | 67.35 | 71.54 |
| IDnet-10 [Our] | 105K | 10 | 92.68 (3) | 74.42 (3) | 74.73 | 81.92 (1) | 74.32 | 79.61 (3) |
| IDnet-20 [Our] | 211K | 20 | 92.58 | 74.78 (2) | 76.34 | 80.72 (2) | 75.77 (2) | 80.04 (2) |
| IDnet-40 [Our] | 423K | 40 | 92.88 (2) | 73.37 | 76.90 | 79.42 | 74.98 (3) | 79.51 |
| IDnet-50 [Our] | 528K | 50 | 92.58 | 73.53 | 75.40 | 79.90 (3) | 74.25 | 79.13 |
| IDnet-60 [Our] | 634K | 60 | 92.30 | 73.67 | 75.93 | 79.40 | 74.35 | 79.13 |

Table 3. FR accuracy when trained on IDnet with 10 (IDnet-10), 20 (IDnet-20), 40 (IDnet-40), and 50 (IDnet-50) synthetic images per identity on mainstream benchmarks, respectively, compared to other work. The top three performing solutions are indicated as (1), (2), and (3). Results obtained without data augmentation are labeled as (No Aug.).

ity is increased beyond 10. This might be due to the higher identity discriminant nature of our IDnet (in comparison to SFace) as discussed in Section 5 which means that fewer images are required to represent a correct class center for each training identity.

When compared to the baseline SFace [6], our proposed IDnet achieves a significantly higher accuracy on all mainstream benchmarks, specifically for scenarios with fewer images per identity. Likewise, IDnet outperforms other related works like USynthFace [7] and SynFace [33]. This emphasizes that the IDnet component allows the generation of face images that retain the selected identity information in a manner that better mimics that of the authentic data. While DigiFace-1M [2] achieves higher accuracies on nearly all benchmarks, IDnet achieves the best performance on the CA-LFW benchmark and comes as a very close second in the AgeDB30 (both targeting age-gap evaluations). This might be linked to a major factor that boosts the results in DigiFace-1M [2], i.e. the aggressive augmentation. While our solution insures natural variations (including age), the augmentations introduced in [2] to boost the performance do not affect the cross-age performance as other face variations. Given that the detailed augmentation parameters are not specified (or released publicly) in [2], a fairer comparison would be between FR models trained without augmentation on our IDnet and the DigiFace-1M.

In this comparison (in Table 3), our IDnet outperforms DigiFace-1M in the average accuracy and very significantly on the cross-age benchmarks, CA-LFW and AgeDB30. Additionally, the performance of DigiFace-1M comes with great limitations. The images of DigiFace-1M are rendered with the physically-based-rendering engine Cycles [2], utilizing 300 NVIDIA M60 GPU for 10 days. In comparison, 500k images are generated by IDnet on a single Nvidia GeForce RTX 6000 GPU in less than 2 hours. Combining our achieved results without and with the open source augmentations along with the significantly faster synthesis the strength of our solution compared to previous work, including DigiFace-1M, becomes evident.

6. Conclusion

In this work, we presented a novel identity conditional three-player GAN, IDnet, which enables synthetic image generation of synthetic identities with high variability and strong identity separability. We empirically demonstrated that utilizing our IDnet to train FR model achieved relatively high verification accuracies on the main FR benchmarks, outperformed previous GAN-based approaches, and achieved competitive results to the computationally costly digital rendering-based synthetic image approach. As concluding remarks, this work accelerates the switch towards training FR models in a privacy-aware manner and explores a new research direction for incorporating identity information in the generation process.

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. This work has been partially funded by the German Federal Ministry of Education and Research (BMBF) through the Software Campus Project.

References

- [1] Stylegan2-ada — official pytorch implementation. <https://github.com/NVlabs/stylegan2-ada-pytorch>. Accessed: 2022-04-16. **5**
- [2] Gwangbin Bae, Martin de La Gorce, Tadas Baltrusaitis, Charlie Hewitt, Dong Chen, Julien P. C. Valentin, Roberto Cipolla, and Jingjing Shen. Digiface-1m: 1 million digital face images for face recognition. In *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2023, Waikoloa, HI, USA, January 2-7, 2023*, pages 3515–3524. IEEE, 2023. **1, 2, 6, 8**
- [3] Fadi Boutros, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Elasticface: Elastic margin loss for deep face recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19-20, 2022*, pages 1577–1586. IEEE, 2022. **1, 2, 5, 6**
- [4] Fadi Boutros, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Self-restrained triplet loss for accurate masked face recognition. *Pattern Recognit.*, 124:108473, 2022. **1**
- [5] Fadi Boutros, Naser Damer, and Arjan Kuijper. Quantface: Towards lightweight face recognition by synthetic data low-bit quantization. In *26th International Conference on Pattern Recognition, ICPR 2022, Montreal, QC, Canada, August 21-25, 2022*, pages 855–862. IEEE, 2022. **1**
- [6] Fadi Boutros, Marco Huber, Patrick Siebke, Tim Rieber, and Naser Damer. Sface: Privacy-friendly and accurate face recognition using synthetic data. In *IEEE International Joint Conference on Biometrics, IJCB 2022, Abu Dhabi, United Arab Emirates, October 10-13, 2022*, pages 1–11. IEEE, 2022. **1, 2, 3, 4, 5, 6, 7, 8**
- [7] Fadi Boutros, Marcel Klemmt, Meiling Fang, Arjan Kuijper, and Naser Damer. Unsupervised face recognition using unlabeled synthetic data. In *17th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2023, Waikoloa Beach, HI, USA, January 5-8, 2023*, pages 1–8. IEEE, 2023. **1, 2, 6, 8**
- [8] Fadi Boutros, Marcel Klemmt, Meiling Fang, Arjan Kuijper, and Naser Damer. Unsupervised face recognition using unlabeled synthetic data. In *17th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2023, Waikoloa Beach, HI, USA, January 5-8, 2023*, pages 1–8. IEEE, 2023. **6**
- [9] Fadi Boutros, Patrick Siebke, Marcel Klemmt, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Pocketnet: Extreme lightweight face recognition network using neural architecture search and multistep knowledge distillation. *IEEE Access*, 10:46823–46833, 2022. **1**
- [10] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *13th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2018, Xi'an, China, May 15-19, 2018*, pages 67–74. IEEE Computer Society, 2018. **1**
- [11] Naser Damer, César Augusto Fontanillo López, Meiling Fang, Noémie Spiller, Minh Vu Pham, and Fadi Boutros. Privacy-friendly synthetic data for the development of face morphing attack detectors. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19-20, 2022*, pages 1605–1616. IEEE, 2022. **1**
- [12] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4690–4699. Computer Vision Foundation / IEEE, 2019. **1, 2, 4**
- [13] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 5153–5162. Computer Vision Foundation / IEEE, 2020. **2, 3**
- [14] Khaled El Emam. Seven ways to evaluate the utility of synthetic data. *IEEE Secur. Priv.*, 18(4):56–59, 2020. **1**
- [15] Meiling Fang, Marco Huber, and Naser Damer. Synthaspoof: Developing face presentation attack detection based on privacy-friendly synthetic data. *CoRR*, abs/2303.02660, 2023. **1**
- [16] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. 2014. **2, 3, 5**
- [17] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III*, volume 9907 of *Lecture Notes in Computer Science*, pages 87–102. Springer, 2016. **1, 2**
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016. **1, 5**
- [19] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008. **2, 6**
- [20] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, and Feiyue Huang. Curricularface: Adaptive curriculum learning loss for deep face recognition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, 2020. **5**
- [21] Marco Huber, Fadi Boutros, Anh Thi Luu, Kiran B. Raja, Raghavendra Ramachandra, Naser Damer, Pedro C. Neto, Tiago Gonçalves, Ana F. Sequeira, Jaime S. Cardoso, João Tremocoço, Miguel Lourenço, Sergio Serra, Eduardo Cermeño, Marija Ivanovska, Borut Batagelj, Andrej Kronovsek, Peter Peer, and Vitomir Struc. SYN-MAD 2022: Competition on face morphing attack detection based on

- privacy-aware synthetic training data. In *IEEE International Joint Conference on Biometrics, IJCB 2022, Abu Dhabi, United Arab Emirates, October 10-13, 2022*, pages 1–10. IEEE, 2022. **1**
- [22] Anil K Jain, Arun Ross, and Salil Prabhakar. An introduction to biometric recognition. *IEEE Transactions on circuits and systems for video technology*, 14(1):4–20, 2004. **1**
- [23] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. **4**
- [24] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. **3, 4, 5**
- [25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4401–4410. Computer Vision Foundation / IEEE, 2019. **3**
- [26] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 8107–8116. Computer Vision Foundation / IEEE, 2020. **3, 4**
- [27] César Augusto Fontanillo López and Abdullah Elbi. On synthetic data: A brief introduction for data protection law dummies, Sep 2022. **1**
- [28] Andrew L Maas, Awni Y Hannun, Andrew Y Ng, et al. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3. Atlanta, Georgia, USA, 2013. **5**
- [29] A Mansfield. Information technology–biometric performance testing and reporting—part 1: Principles and framework. *ISO/IEC*, pages 19795–1, 2006. **5**
- [30] Richard T. Marriott, Safa Madiouni, Sami Romdhani, Stéphane Gentic, and Liming Chen. An assessment of gans for identity-related applications. In *2020 IEEE International Joint Conference on Biometrics, IJCB 2020, Houston, TX, USA, September 28 - October 1, 2020*, pages 1–10. IEEE, 2020. **2**
- [31] Lars M. Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 3478–3487. PMLR, 2018. **5**
- [32] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: The first manually collected, in-the-wild age database. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1997–2005. IEEE Computer Society, 2017. **2, 6**
- [33] Haibo Qiu, Baosheng Yu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. Synface: Face recognition with synthetic data. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 10860–10870. IEEE, 2021. **2, 6, 8**
- [34] Inioluwa Deborah Raji and Genevieve Fried. About face: A survey of facial recognition evaluation. *CoRR*, abs/2102.00813, 2021. **1**
- [35] Soumyadip Sengupta, Jun-Cheng Chen, Carlos Domingo Castillo, Vishal M. Patel, Rama Chellappa, and David W. Jacobs. Frontal to profile face verification in the wild. In *2016 IEEE Winter Conference on Applications of Computer Vision, WACV 2016, Lake Placid, NY, USA, March 7-10, 2016*, pages 1–9. IEEE Computer Society, 2016. **6**
- [36] Yujun Shen, Ping Luo, Junjie Yan, Xiaogang Wang, and Xiaoou Tang. Faceid-gan: Learning a symmetry three-player GAN for identity-preserving face synthesis. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 821–830. Computer Vision Foundation / IEEE Computer Society, 2018. **2**
- [37] Marcus Smith and Seumas Miller. The ethical application of biometric facial recognition technology. *AI Soc.*, 37(1):167–175, 2022. **1**
- [38] Riccardo Spolaor, QianQian Li, Merylin Monaro, Mauro Conti, Luciano Gamberini, and Giuseppe Sartori. Biometric authentication methods on smartphones: A survey. *Psychology Journal*, 14(2), 2016. **1**
- [39] The European Parliament and the Council of the European Union. Art. 4 of the general data protection regulation, 2016. **1**
- [40] The European Parliament and the Council of the European Union. Art. 7 of the general data protection regulation, 2016. **1**
- [41] The European Parliament and the Council of the European Union. Article 9 of the general data protection regulation, 2016. **1**
- [42] 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. **1**
- [43] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 5265–5274. Computer Vision Foundation / IEEE Computer Society, 2018. **1, 2, 4, 5**
- [44] Mei Wang and Weihong Deng. Deep face recognition: A survey. *Neurocomputing*, 429:215–244, 2021. **1**

- [45] Xiaobo Wang, Shuo Wang, Hailin Shi, Jun Wang, and Tao Mei. Co-mining: Deep face recognition with noisy labels. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 9357–9366. IEEE, 2019. 7
- [46] Shoukai Xu, Haokun Li, Bohan Zhuang, Jing Liu, Jiezhong Cao, Chuangrun Liang, and Mingkui Tan. Generative low-bitwidth data free quantization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 1–17. Springer, 2020. 4
- [47] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014. 2, 4
- [48] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE signal processing letters*, 23(10):1499–1503, 2016. 4, 6
- [49] T. Zheng and W. Deng. Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments. Number 18-01, 2018. 6
- [50] Tianyue Zheng, Weihong Deng, and Jiani Hu. Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments. volume abs/1708.08197, 2017. 6