

IDiff-Face: Synthetic-based Face Recognition through Fizzy Identity-Conditioned Diffusion Models

Fadi Boutros¹, Jonas Henry Grebe¹, 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: fadi.boutros@igd.fraunhofer.de

Abstract

The availability of large-scale authentic face databases has been crucial to the significant advances made in face recognition research over the past decade. However, legal and ethical concerns led to the recent retraction of many of these databases by their creators, raising questions about the continuity of future face recognition research without one of its key resources. Synthetic datasets have emerged as a promising alternative to privacy-sensitive authentic data for face recognition development. However, recent synthetic datasets that are used to train face recognition models suffer either from limitations in intra-class diversity or cross-class (identity) discrimination, leading to less optimal accuracies, far away from the accuracies achieved by models trained on authentic data. This paper targets this issue by proposing IDiff-Face, a novel approach based on conditional latent diffusion models for synthetic identity generation with realistic identity variations for face recognition training. Through extensive evaluations, our proposed synthetic-based face recognition approach pushed the limits of state-of-the-art performances, achieving, for example, 98.00% accuracy on the Labeled Faces in the Wild (LFW) benchmark, far ahead from the recent synthetic-based face recognition solutions with 95.40% and bridging the gap to authentic-based face recognition with 99.82% accuracy.*

1. Introduction

Face Recognition (FR) is one of the most widely used biometric technologies due to the high accuracies achieved by the recent FRs [31, 56, 4] with a wide range of applications such as logical access control to portable devices [42, 9]. This ubiquitous adoption has been fuelled by the application of deep learning to FR and the rapid research advances in this direction, mainly on novel margin-penalty based softmax losses [14, 4] and deep network architectures

[24, 9]. However, this rapid progress has only been possible due to the public availability of large-scale FR training databases [11, 23]. Such databases contain millions of images, and they are typically collected from the internet without proper user consent, which raises concerns about the legal and ethical use of these databases for FR development.

Since, for example, the European Union (EU) adopted the General Data Protection Regulation (GDPR) [18] in 2018, more justified criticism of the associated privacy risks has been raised against the usage of public biometric datasets that were collected without proper consent. The GDPR explicitly grants individuals the "right to be forgotten" and enforces stricter requirements on the collection, distribution, and usage of face databases, making it extremely hard, or even infeasible, to maintain such regulations. Thus, many of them [11, 3, 23] that are widely used to train FRs were retracted by their creators to avoid legal complications, which raises the question about continuity of FR research since the availability of one of its key resources became questionable.

As an effort to address these legal and ethical concerns, synthetic data has recently emerged as a promising alternative to authentic databases for FR training [43, 6, 7, 10, 2, 5]. This is also the trend for FR subsystems, such as morphing and spoof attack detection [13, 19] and face images quality assessment [1]. This research direction is driven by the progress on Deep Generative Models (DGMs), a model designed to learn the probability distribution of a certain dataset, enabling the generation of completely new synthetic samples. This process can be conditioned on specific attributes such as age, facial expression, and a defined set of visual appearances, e.g. the lighting condition and head pose [15, 52, 54, 21]. Most of the DGM approaches for generating synthetic faces are based on Generative Adversarial Networks (GANs) [22, 15, 49, 33, 6]. A number of recent works [37, 7, 43] utilized GANs [15, 29] to generate synthetic data for FR training. The reported results by SOTA synthetic-based FR showed significant degradation in the verification accuracies in comparison to FR trained

*<https://github.com/fdbtrs/idiff-face>

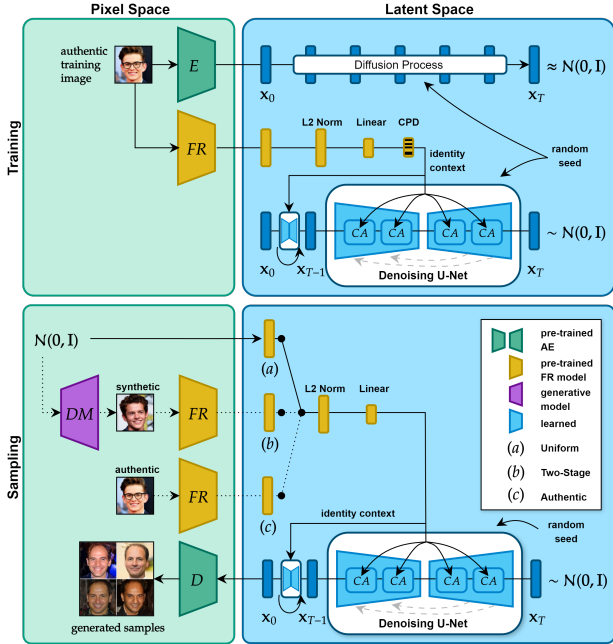


Figure 1: Overview of the proposed IDiff-Face. This diagram is divided into two parts, the upper part visualizes the training procedure, and the lower part shows the conditional sampling process, partially inspired by [45]. **Top:** During training, the learned denoising U-Net is conditioned on a context that is based on the feature representation obtained through a pre-trained FR model. The entire DM training process is in the latent space of a pre-trained AE. The diffusion process that basically provides the targets for learning the iterative reverse process is depicted above. **Bottom:** For samples generation, the trained DM can generate samples based on three types of identity contexts: authentic, two-stage, or synthetic uniform representations. By fixing the identity context and varying the added noise, different samples for the same identity can be generated.

on authentic data. This performance gap is mainly due to low identity discrimination [6] or small intra-class variation [7, 43] in their training datasets. A realistic trade-off between these two properties, as we will demonstrate in this paper, is needed to achieve high verification accuracies.

Recently, Diffusion Models (DMs) [25, 40, 45, 16] gained attention for both research and industry due to their potential to rival GANs on image synthesis [16], even though they are easier to train, have no stability issues, and stem from a solid theoretical foundation. Besides that, they can be conditioned on additional information, and demonstrated impressive results in a wide range of tasks, especially in text-to-image generation such as OpenAI’s DALL-E 2 [44], Stable Diffusion [45], and Google’s Imagen [47].

We propose in this work an identity-conditioned Diffusion Model approach, namely IDiff-Face. Our IDiff-Face is designed and trained to generate synthetic images of syn-

thetic identities that are identity-separable, with a desirable relatively large intra-class diversity. Our approach follows the common concept [45] of dividing the training into two stages by first training an AE (or using a pre-trained AE) and then leveraging the resulting representational low-dimensional latent space of the AE for DM training [45]. The identity condition is introduced to our IDiff-Face by projecting the training images into a low-dimensional feature representation and then injecting it into the DM’s intermediate representations through a Cross-Attention (CA) mechanism [45]. To ensure that our synthetic data contains, to a large degree, realistic variations and to avoid overfitting our IDiff-Face to the information encoded in the identity context, we proposed a simple, yet effective Contextual Partial Dropout (CPD) approach that partially drops out components with a certain probability of the identity context during the training phase, thus the term “fuzzy” in our title. We first demonstrate the identity discrimination and intra-class variation of our synthetically generated data. We also compared our dataset in terms of identity-separability and intra-class variation, with the recent SOTA synthetic datasets. As we empirically present in this paper, the synthetic datasets that are used in SOTA synthetic-based FR training maintain, to some degree, identity discrimination, however, only with a low intra-class variation or vice versa. Unlike these approaches, our proposed IDiff-Face offers a more realistic trade-off between identity discrimination and intra-class variation that is controlled by CPD. Achieving this realistic trade-off between these two properties is necessary to achieve verification accuracy that is close to verification accuracies achieved by FRs trained on authentic data as authentic data naturally contains such properties. By utilizing our synthetic data (500K images) for FR training and under the same training setups and dataset size, our synthetic-based FR achieved an average accuracy of 88.20%, significantly outperforming all SOTA synthetic-based FRs (best average accuracy was 83.45%) and closing the gap to SOTA FR trained on authentic data, where the average accuracy on CASIA-WebFace [59] (500K images) and MS1MV2 [23, 14] (5.8M images) were 94.92% and 97.18%, respectively. Our model also outperformed human-level performance in face verification, where the reported accuracy on Labeled Faces in the Wild (LFW) [26] was 97.5% [35] and our model achieved 98.00%.

2. Related Work

The vast majority of DGMs that are proposed in the literature are often designed for generating synthetic images of random identities [22, 28, 30, 10] or editing certain facial attributes of existing reference images [55, 51, 49, 54], and thus, such models are not by design capable of generating synthetic identities with multiple different images per identity. While other DGMs proposed to generate images of

synthetic identities, most of them either rely on pre-existing semantic attribute annotations [55, 50, 51, 49, 15, 34, 21, 52, 54], meticulously constructed training batches [17, 15, 52], or the supervision by sophisticated 3D Morphable Face Models (3DMM) [50, 51, 15, 34, 21, 54]. In some cases, even all of these requirements are necessary in order to explicitly model specific parametric factors, such as age or illumination, and thus, gain control over the generative process [15]. These approaches presented impressive results in generating high-quality and realistic images [15, 52] with an unprecedented level of control. However, the visual appearances in their generated images are limited to a pre-defined set of attributes [15, 52], and thus, their synthetic images might not contain natural real-world variations or large diversities in terms of utility [20]. Utilizing such synthetic data for FR training might lead to suboptimal verification accuracies. For that reason, SynFace [43] proposed a synthetic-based FR approach based on DiscoFaceGAN [15] with synthetic identity mix-up to enhance the intra-class diversity. USynthFace [7] proposed the use of unlabelled synthetic data for unsupervised FR training, achieving competitive accuracies. During FR training, USynthFace [7] proposed to utilize intensive data augmentation, which significantly improved the overall verification accuracy. On the other hand, SFace [6] and IDnet [33] proposed to train a StyleGAN-ADA [29] under a class-conditional setting. SFace [6] images contain a higher intra-class variation, which, comes at the cost of low identity discrimination, when compared to the other approaches [7]. ExFaceGAN [8] presented a framework to disentangle identity information in learned latent spaces of GANs to generate multiple samples of any synthetic identity. DigiFace-1M [2] utilized a digital rendering pipeline to generate synthetic images based on a learned model of facial geometry and a collection of textures, hairstyles, and 3D accessories. DigiFace-1M [2] also applied aggressive data augmentations for their FR training. However, it comes at a considerable computational cost during the rendering process. All previously mentioned approaches presented promising accuracies. Nonetheless, these accuracies are still significantly lower than the ones achieved by FR trained on authentic data, where for example, the average accuracy (on five benchmarks, see Table 3) of SOTA synthetic-based FR was 83.45%, far away from 94.82% achieved by authentic-based FR.

Motivated by the remarkable text-to-image results of recent approaches that leverage DMs [45, 44, 47] and in an effort to bridge the gap between synthetic- and authentic-based FR performance, this paper is the first to propose an identity-conditioned approach based on DM to generate synthetic identity-specific images with a more realistic intra-class diversity for FR training, outperforming all recent SOTA synthetic-based FRs with an obvious margin.

3. Methodology

This section describes our proposed IDiff-Face approach to generate identity-specific yet realistic synthetic face images. IDiff-Face is based on a DM that is conditioned on identity contexts. During the training stage, our proposed IDiff-Face is conditionally trained on authentic embeddings obtained from the authentic training dataset. After training, our IDiff-Face can be used either to generate variations of existing authentic images by using authentic embeddings or to generate novel synthetic identities by using synthetic embeddings. In Figure 1, an overview of the proposed method is provided, with a clear distinction between the training and the sampling stages. In order to generate samples of synthetic identities, a synthetic identity representation has to be created, e.g. synthetic uniform or synthetic two-stage contexts, as explored in this work.

3.1. Identity-Conditioned Latent Diffusion

Our IDiff-Face is based on a DDPM that is trained in the latent space of a pre-trained AE [45] and conditioned on identity-contexts i.e. feature representations extracted using a FR model. A DDPM [25] is a DM that discretizes the diffusion processes into a finite number T of steps and learns to reverse this process by training a conditional DNN to estimate the noise that has been added to a sample \mathbf{x}_t at time step t . The architecture of our IDiff-Face is a modified U-Net [46] based on [25] that includes residual and attention blocks. We incorporate the identity context of a sample \mathbf{x} in the DM to encourage the DM to learn to generate identity-specific face images. This has been achieved by mapping $\mathbf{x} \in \mathbb{R}^{W \times H \times C}$ into a feature representation $f(\mathbf{x}) = \mathbf{c} \in \mathbb{R}^d$ using a pre-trained face recognition model f optimized to learn discriminant identity information, which is then injected to the DM using CA mechanism as proposed by [45]. An ablation study on two different conditional mechanisms, including conditional CA [45] and Adaptive Group Normalization (AdaGN) [16] for identity-specific images generation is provided in the supplementary material.

Let $0 < \beta_1, \beta_2, \dots, \beta_T < 1$ be a fixed (linear) variance schedule and \mathbf{x}_0 be the encoded image \mathbf{x} in the pre-trained latent space. The learnable *reverse diffusion process* is defined as a Markov chain with Gaussian transitions $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ starting at a prior distribution $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. As proposed by the DDPM authors, this transition kernel is parameterized by the DNN with parameters θ that learn to predict the noise at time step t based on the current estimate \mathbf{x}_t and, in our case, the additional condition \mathbf{c} . With the notational abbreviations $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t := \sum_{i=1}^t \alpha_i$ [25], our conditional variant of the original DDPM training objective is given as:

$$\begin{aligned} \mathcal{L}(\theta) &:= \mathbb{E}_{t, \mathbf{x}_t, \epsilon} \left[\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t, \mathbf{c})\|_2^2 \right] \\ &= \mathbb{E}_{t, \mathbf{x}_0, \epsilon} \left[\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t, \mathbf{c})\|_2^2 \right]. \end{aligned}$$

This expectation can be minimized through SGD by randomly drawing samples $(\mathbf{x}_0, t, \epsilon, \mathbf{c})$ with $\mathbf{x}_0 \in \mathcal{D}, t \sim \mathcal{U}(1, T), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and $\mathbf{c} = f(\mathbf{x})$ and then minimizing the MSE between the true and the predicted noise.

Analogously, the original DDPM [25] sampling is slightly modified to incorporate identity contexts \mathbf{c} into the iterative sampling process, which mirrors a score-based sampling chain with *Langevin dynamics* using additionally added noise vectors $\zeta_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ at each time step t .

$$\begin{aligned} \mathbf{x}_{t-1} &= \mu_\theta(\mathbf{x}_t, t, \mathbf{c}) + \sigma_t \zeta_t \\ &= \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(\mathbf{x}_t, t, \mathbf{c}) \right) + \sigma_t \zeta_t. \end{aligned}$$

At the end of the iterative process, the final estimate of \mathbf{x}_0 from the last iteration has to be mapped back from the latent space to the pixel space by using the pre-trained decoder.

3.2. Synthetic Image Generation

As the proposed IDiff-Face is conditionally trained on identity contexts, i.e. feature representations of input face images, such feature representations are also required to generate synthetic samples. Feature representations \mathbf{c} can be obtained from any authentic image \mathbf{x} and, in this case, the IDiff-Face will generate variations of \mathbf{x} using randomly sampled $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and different seeds for the non-deterministic sampling process. To generate images of synthetic identities, we explored two different approaches. First, one can randomly sample the identity context from a uniform distribution over the hypersphere, because face representations are typically L2-normalized [4, 56]. In order to generate samples from this distribution, a naive way of doing it would be to simply generate a random direction from a uniform distribution $\mathcal{U}(-1, 1)$ for each of the components before normalizing the resulting vector to get a point on the surface of the sphere. Unfortunately, this does not result in a uniform distribution over the surface. Instead, the components of the initial direction have to be sampled from a spherical Gaussian distribution, e.g. $\mathcal{N}(\mathbf{0}, \mathbf{I})$, to obtain a uniform distribution after the normalization process [38]. After sampling, one can utilize it as a fixed input to our IDiff-Face to generate different samples from the same identity by varying the noise seeds. The second approach is to utilize an additional unconditional DM to generate a synthetic image \mathbf{x}' . After that, a feature representation \mathbf{c}' can be extracted using f as $\mathbf{c}' = f(\mathbf{x}')$. Analogously to the first approach, different samples of the same identity can be obtained by fixing \mathbf{c}' throughout different samples.

The latter approach will be referred to as Two-Stage IDiff-Face, where a random synthetic image from a random identity is generated first using an unconditional DM model and then it is used as input to our conditional IDiff-Face to generate different samples of the same identity.

3.3. Enhancing Intra-Class Variation via CPD

To prevent the model from overfitting to identity contexts, which limits the intra-class variation in the generated samples, we introduced Dropout [53] to the IDiff-Face training. If the network is overfitted to the identity-context condition, the generated images using the same identity context will be almost identical i.e. small intra-class variations, regardless of the initial starting noise or random seed. This has been achieved by dropping out components of a context embedding, each with a certain probability. This method is referred to as Contextual Partial Dropout (CPD) as the context is only partially dropped out during training.

4. Experimental Setup

4.1. IDiff-Face Training Dataset

The proposed DDPM was trained using the FFHQ dataset [30], which consists of 70,000 high-quality images of human faces, showcasing a diverse range of attributes such as age, lightning, and facial expressions. The face images used for training have a resolution of 128×128 pixels.

4.2. IDiff-Face Training Setups

The experiments in this paper were conducted on a cluster of 8 NVIDIA A100-SXM4-40GB GPUs. The proposed ϵ -prediction model (U-Net) [45] takes a three-channel input and 96 channels for the initial image projection layer, which are then used as input for the first residual block [24]. The network has four resolution levels, and the multipliers for the number of channels used on those levels are 1, 2, 2, and 2 respectively. Attention mechanisms use a fixed number of 32 channels per head, with the number of heads calculated based on the number of incoming channels. Attention blocks are applied in all residual blocks except the first resolution level [45]. Each resolution level has 2 consecutive residual blocks. For the experiments, three different levels of CPD probability have been explored: 0%, 25%, or 50%. We trained and evaluated three instances of our IDiff-Face with these three dropout probabilities, which will be denoted as CPD0, CPD25, and CPD50.

The training process follows a batch-wise training loop over the training dataset \mathcal{D} with half-precision (float16) and a learning rate schedule based on the number of global steps. An annealing cosine learning rate schedule with warm restarts [36] is used, which reduces the learning rate in repeating phases where the first phase is 10,000 steps long and each subsequent phase is twice as long as the previous one. Therefore, the number of global steps per training is set to 150,000 steps, which corresponds to 4 full phases. Also, the Adam optimizer [32] is used with an initial learning rate of $\gamma = 1e - 4$. Following [25, 45], we apply an EMA to the weights of the model with a negative exponential factor of 0.75, leading to an effective decay factor of 0.999 at 10,000 steps. Therefore, a second EMA

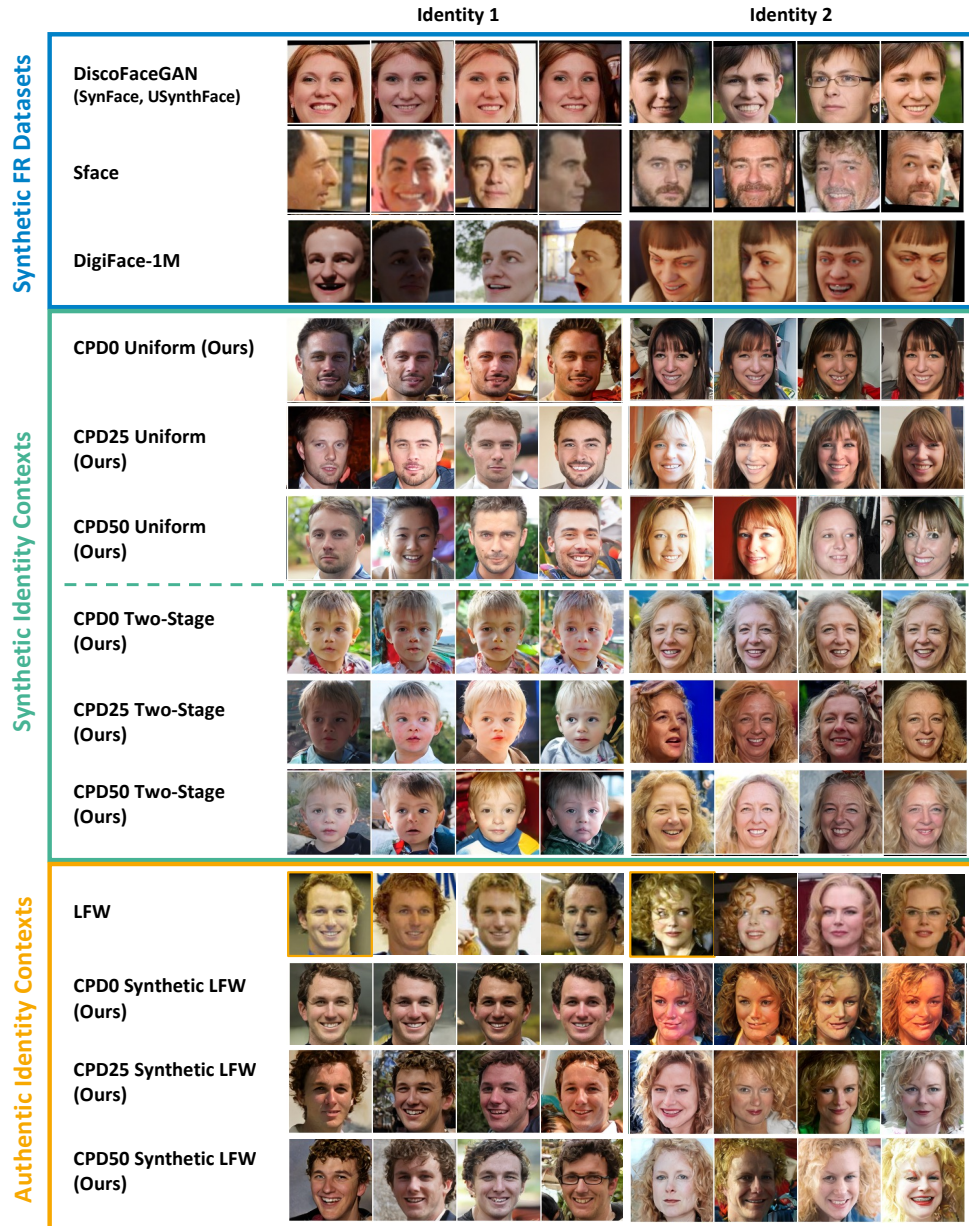


Figure 2: Samples from synthetic FR training datasets and our IDiff-Face . The upper box (in blue) shows samples from the three DGMs that are used in the SOTA synthetic-based FRs. The next group (in green) presents samples from our IDiff-Face models with different CPD probabilities and different types of synthetic embeddings. The last group (in yellow) demonstrates the generation of variations for existing LFW identities, whose reference images are framed in yellow colour. There are four consecutive images per identity, and two identities in total are presented for each method. Zoom in for the best view.

copy of the currently trained model is maintained for inference purposes. The batch size is fixed at 512 and is equally split across 8 GPUs. The training dataset is augmented with horizontal flipping with a 50% probability. Regarding the diffusion process, $T = 1,000$ time steps and a linear diffusion variance schedule are used throughout all experiments.

Autoencoder: The latent space of the pre-trained AE is used to lower the computational demands and facilitate

the learning process by learning a perceptually similar but less complex space for the DGM before the LDM is trained. For the AE, a pre-trained VQGAN model (*vq-f4* first stage model) from the official LDM [45] repository is used. A VQGAN can be understood as a regularized AE, which enforces a discrete latent space by learning a codebook of latent vector and replacing the encoder’s outputs with its nearest-neighbor codes during inference. During training,

VQGANs are guided by an additional patch-based adversarial objective \mathcal{L}_{adv} to better match the realism of the training data with the reconstructions. The pre-trained VQGAN provides an effective trade-off between sample compression and generative abilities as demonstrated in recent work [45].

Identity Context: The identity context is obtained by projecting the training image into its feature representation using a SOTA pre-trained FR, namely ElasticFace [4]. We used the official pre-trained model released by [4]. The network architecture of ElasticFace [4] is ResNet-100 [24] trained on the MS1MV2 [23, 14] dataset using ElasticFace-Arc loss and the output feature dimensionality is 512.

5. Results

Assessment of Intra-Class Diversity and Identity Separability: Figure 2 presents samples of recent works that used synthetic datasets for FR model training. Synthetic face recognition models, SynFace [43] and USynthFace [7], utilized synthetic images generated by DiscoFaceGAN [15]. DiscoFaceGAN is based on disentangled representation learning to generate images from synthetic identities with predefined attributes e.g. pose, illumination, or expression. As generated images are explicitly controlled by a predefined set of attributes, such images might lack the intra-class diversity that exists in real-world face images, which is needed to successfully train FR models. SFace [6], on the other hand, is a class-conditional GAN model that does not explicitly model these attributes. It is conditionally trained to generate synthetic images with a specific label. SFace provided images with a larger degree of intra-class variations, however, at the cost of less identity separability. In contrast to that, the DigiFace-1M [2] images are produced by a 3DMM rendering process. The identities in DigiFace-1M are artificially defined as a combination of facial geometry, texture, and most notably hairstyle. Unfortunately, this approach is less suitable for research purposes, as it is extremely computationally expensive for generating a large dataset with an advanced computational rendering pipeline.

These observations are quantitatively supported by the genuine and imposter comparison score distribution plots in Figure 3 and corresponding verification performance results in Table 3. The verification performance metrics include FMR100, and FMR1000, which are the lowest false non-match rate (FNMR) for a false match rate (FMR) $\leq 1.0\%$ and $\leq 0.1\%$, respectively, along with the Equal Error Rate (EER) [27]. We additionally report the mean and standard deviation of the genuine and imposter scores. Further, we report the Fisher Discriminant Ratio (FDR) [41] to provide an in-depth analysis of the separability of genuine and imposters scores. We used a pre-trained ElasticFace-Arc [4] by the corresponding author (model publicly available) to extract the feature embeddings of CASIA-WebFace [59], LFW [26], DiscoFaceGAN [15] (used in SynFace [43] and

USynthFace [7]), DigiFace-1M [2], and our IDiff-Face for the investigations described in this subsection. CASIA-WebFace [59] and LFW [26] are authentic datasets, and they are commonly used to train or evaluate FR models [14, 4], respectively. We made the following observations:

1): In comparison to the strong identity discrimination in the authentic LFW and CASIA-WebFace datasets, our IDiff-Face (CPD of 0%) with two-stage and uniform identity-context sampling approach generates synthetic samples that clearly maintain identity discrimination, where, for example, the achieved EER on LFW was 0.002 and by our models with CPD of 0% were 0.007 (Uniform) and 0.003 (Two-Stage). It should be noted that CASIA-WebFace is an FR training dataset, and contains some noisy labels, as reported in [57], which contributes to the higher EER.

2): Training our proposed models with CPD demonstrates that with an increase of the CPD probability, the intra-class diversity becomes larger as well, leading to more variation, in the head pose, the illumination, the expression, the accessories, and sometimes even in the age of the depicted individual, as shown in Figure 2. This can be seen in the shift in the genuine distributions to left (i.e. genuine distribution shifts to the left as more challenging and realistic variations are generated) when the CPD is increased, as shown in Figure 3 and the mean genuine score in Table 3. The higher the CPD probability, the larger intra-class and realistic real-world variations can be achieved. It also comes at the cost of slightly losing identity discrimination. This observation can be seen in the EER value increase and the decrease in FDR values induced by incorporating CPD into the model training (Table 1). Hence, the probability of the CPD can be interpreted as a trade-off between identity discrimination and intra-class diversity. Achieving a good balance for this trade-off is required to generate useful data for successful FR model training and to achieve high accuracies, as we will empirically test in the next section.

3): Among the recent SOTA synthetic FR training datasets, the synthetic dataset that is generated by DiscoFaceGAN [15] (used in SynFace and USynthFace) maintains identity discrimination with EER of 0.01069. However, it comes at the high cost of low intra-class diversity as the generated images are controlled by a predefined set of attributes, making it less optimal to train FR models. On the contrary, SFace possesses a large intra-class diversity, however, with a lower degree of identity discrimination in comparison to other synthetic datasets. DigiFace-1M maintains to some degree identity discrimination. However, it is obtained using a computationally expensive digital rendering process, leading to very high costs in the generation process. Our proposed IDiff-Face with CPD achieved the best (most representative of reality and thus suitable for training) trade-off between identity discrimination and intra-class di-

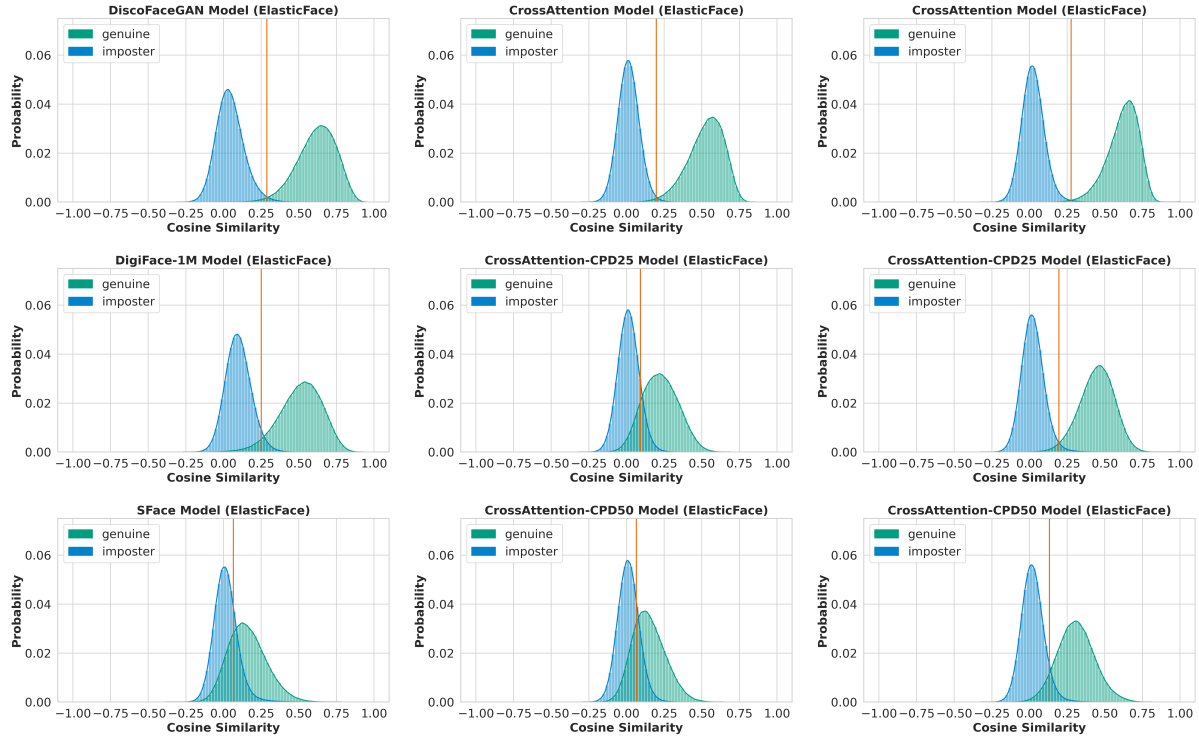


Figure 3: Syn-vs-syn genuine and imposter comparison score distributions for different DGMs. The first column presents the distributions obtained from samples of DGMs used in SOTA synthetic-based FRs. The second column presents the results from our three models based on synthetic uniform embeddings, while the last column shows our distributions based on synthetic two-stage embeddings. For each, all genuine comparison scores have been computed ($\approx 1,200,000$ pairs) based on 5 K identities with 16 images each, and the same number of imposter scores has been randomly sampled.

versity and leads to SOTA synthetic-based FR, as we will present in the next section.

Synthetic-based FR: We evaluate first the different variants of the proposed solution regarding their applicability to generate synthetic data for FR. For that, we start by generating three sets of synthetic data using a uniform identity-context sampling, each with different CPD probabilities of 0, 25% and 50%. We also generated other three sets of synthetic data using a two-stage identity-context sampling, with the same three CPD probabilities. All training datasets contain 80,000 samples (5,000 identities with 16 samples each). We utilized ResNet-18 [24] with the CosFace [56] loss to train six FR models on each of our synthetic datasets. For this small-scale FR model training, a fixed 40 iteration over the entire training data, a step-based learning rate schedule with an initial rate of 0.1 and reduces it by a factor of 0.1 after the 22nd, the 30th, and the 35th epoch, and an SGD optimizer with 0.9 momentum and $5e - 4$ weight decay were used, following [7]. Evaluations are reported on the five benchmarks, LFW [26], AgeDb-30 [39], CA-LFW [61], CFP-FP [48], and CP-LFW [60], following their official evaluation protocols (see Table 2).

It can be observed from Table 2 that FR models trained

on our IDiff-Face datasets achieved high verification accuracies, even only using a small synthetic dataset (80K samples) for training. For the models trained with datasets generated by IDiff-Face with uniform identity-context sampling, the best verification accuracies were achieved using the FR model trained on the dataset generated by IDiff-Face with CPD25 (79.54% average accuracies). For the models trained with datasets generated by IDiff-Face with two-stage identity-context sampling, both FR models trained on the datasets generated by IDiff-Face with CPD25 and CPD50 achieved very competitive results. The best overall verification accuracy was achieved by IDiff-Face using two-stage identity-context sampling with CPD50 (average accuracies of 79.73%) which is very close to two-stage identity-context sampling with CPD25 (average accuracies of 79.72%) and uniform identity-context sampling and CPD25 (Average accuracies of 79.54%). Although the CPD approach slightly affects, to some degree, the identity discrimination in the synthetically generated samples (as we discussed in the previous section), it leads to significant improvements in the FR verification accuracies as it increases the intra-class variations in the generated samples. As we mentioned in the previous section, achieving a realis-

Method	Operation Metrics			Score Distributions				
	EER ↓	FMR100 ↓	FMR1000 ↓	genuine		imposter		
				mean	std	mean	std	FDR ↑
CASIA-WebFace [59]	0.076	0.092	0.107120	0.536	0.215	0.003	0.070	5.5409
LFW [26]	0.002	0.002	0.002	0.708	0.099	0.003	0.070	33.301
DiscoFaceGAN [15]								
(SynFace & USynthFace)	0.011	0.011	0.051	0.619	0.128	0.044	0.092	13.378
DigiFace-1M [2, 58]	0.042	0.087	0.199	0.512	0.140	0.099	0.084	6.372
SFace [6]	0.236	0.768	0.380	0.159	0.125	0.016	0.079	0.941
IDiff-Face								
Synthetic Uniform (Ours)								
CPD0	<u>0.007</u>	<u>0.005</u>	<u>0.019</u>	0.528	0.117	0.014	0.069	14.243
CPD25	0.130	0.385	0.607	0.226	0.117	0.014	0.070	2.427
CPD50	0.225	0.660	0.857	0.149	0.108	0.013	0.070	1.120
w/ Synthetic Two-Stage (Ours)								
CPD0	0.003	0.001	0.009	0.621	0.102	0.024	0.075	22.172
CPD25	0.018	0.030	0.190	0.448	0.114	0.023	0.075	9.733
CPD50	0.069	0.238	0.659	0.309	0.122	0.021	0.075	4.064

Table 1: Evaluation of identity-separability in synthetic FR datasets proposed in the literature. The first two rows present the results on authentic LFW and CASIA-WebFace datasets. These results are provided as a reference. In the next rows, we provide the evaluation results on SOTA synthetic FR datasets and our IDiff-Face. Except for the first row, which shows the results obtained by computing the comparison scores on pre-defined pairs of authentic LFW, all the synthetic evaluations are based on a synthetically generated dataset with 5,000 identities and 16 sampled images per identity. The lowest errors and the highest genuine-imposter separability scores (FDR) on synthetic datasets are marked in **bold**. The second best per column is underlined.

Training Dataset		Verification Benchmarks ↑						
		Cross-Age			Cross-Pose		Avg.	
		LFW	AgeDB-30	CA-LFW	CFP-FP	CP-LFW		
Identity-context model	Uniform	CPD0	86.50	60.30	72.13	62.26	62.40	68.72
		CPD25	95.07	74.87	84.60	70.09	73.08	79.54
		CPD50	92.92	71.43	80.48	66.79	69.03	76.13
Two-Stage		CPD0	86.08	61.50	71.68	61.64	62.20	68.62
		CPD25	<u>95.10</u>	<u>76.18</u>	<u>84.17</u>	<u>70.67</u>	<u>72.47</u>	<u>79.72</u>
		CPD50	95.42	76.55	83.22	70.90	<u>72.57</u>	79.73

Table 2: Verification Accuracies (in %) on five FR benchmarks for FRMs (ResNet-18) trained on 80,000 samples (5,000 identities with 16 images per identity) generated with different levels of CPD probability. The best accuracies are marked in **bold** and the second best is underlined.

tic trade-off between identity discrimination and intra-class diversity is required for the FR training dataset to achieve high verification accuracies on evaluation benchmarks.

Comparison to the SOTA Synthetic-based FR: We compare the verification performance of FR models trained with our synthetic data with the recent works that proposed the use of synthetic data for FR training. To provide a fair comparison, we followed the SOTA synthetic FR approaches [43, 6, 7] to generate 500,000 images from 10,000 synthetic identities with 50 samples each, and utilized these images to train ResNet-50 [24]. We also followed the exact training setups described in [7] to train 15 FR models using different dataset sizes, three levels of CPD, and two

identity-context sampling approaches to investigate the effect of training dataset width and depth and identity-context sampling mechanisms on FR verification accuracy. The achieved results by SOTA synthetic-based FR and our proposed IDiff-Face are presented in Table 3. One can observe the following from the reported results in Table 3:

(1): With a clear margin, FR models trained with our IDiff-Face outperformed all previous synthetic-based FRs. The best achieved average accuracy by our models was 88.20% and the average accuracy of SOTA synthetic-based FR was 83.45% (achieved by DigiFace-1M [2]).

(2): Increasing our training dataset size did increase the FR verification accuracies in all experimental settings. This indicates that increasing our synthetic training dataset size could further improve the FR verification accuracy.

(3): Increasing our dataset width (the number of identities) led to higher accuracies in comparison to the case when the dataset depth (the number of images per identity) is increased. For example, IDiff-Face with uniform sampling (CPD25) led to an average accuracy of 82.86% when using 160K samples (5K identities with 32 images per identity). This accuracy improved to 83.87% when we trained on 160K (10K images with 16 images per identity). This observation can be concluded from all of our experiments.

(4): Introducing data augmentation [7] to model training improved the verification accuracies, where we additionally trained the best model in each setting with RandAugment [12] using settings provided in [7]. This experiment is highlighted with + *RandAugment(4, 16)*. It should be noted that previous SOTA synthetic-based FR models (DigiFace-1M [2], ExFaceGAN [8], IDnet [33] and USynthFace [7]) also applied extensive data augmentation in their model training.

(5): With only 160K training samples (10K identities with 16 images each), our approach outperformed all previous synthetic-based FR that used GAN-generated training face images (SynFace, SFace, ExFaceGAN, IDnet and USynthFace) and the one that used digital face rendering (DigiFace-1M), where the average accuracy by our best model trained on 160K images was 83.87% and the average accuracies by SynFace (500K training samples), SFace (634K training samples), USynthFace (400K training samples) and DigiFace-1M (500K training samples) were 74.75%, 77.71%, 79.30% and 83.45%, respectively.

(6): Although it is out of our paper’s comparison scope, we compared our synthetic-based FRs with the SOTA FRs trained on authentic datasets (CASIA-WebFace and MS1MV2). Using only our synthetic data, our proposed models achieved very competitive results to models trained on the large authentic training dataset. This clearly indicates that synthetically generated data is growing to become a valid alternative to authentic data to train FR. We achieved for example 98.00% accuracy on LFW [26], which is competitive to SOTA FR accuracies that are trained on authentic

Method	Dataset				Verification Benchmarks \uparrow					
					Cross-Age			Cross-Pose		
name	Aug.	Id	N per Id	images	LFW	AgeDB-30	CA-LFW	CFP-FP	CP-LFW	Avg.
ElasticFace [4] (CASIA-WebFace [59])	✗	10.5 K	-	494 K	99.52	94.77	93.93	95.52	90.38	94.82
ElasticFace [4] (MS1MV2 [23])	✗	85 K	-	5,800 K	99.82	98.27	96.03	98.61	93.17	97.18
SynFace [43]	✗	10 K	50	500 K	91.93	61.63	74.73	75.03	70.43	74.75
SFace [6]	✗	10.5 K	10	105 K	87.13	63.30	73.47	68.84	66.82	71.91
	✗	10.5 K	20	211 K	90.50	69.17	76.35	73.33	71.17	76.10
	✗	10.5 K	40	423 K	91.43	69.87	76.92	73.10	73.42	76.95
	✗	10.5 K	60	634 K	91.87	71.68	77.93	73.86	73.20	77.71
USynthFace [7] + RandAugment [7]	✓	100 K	1	100 K	92.12	71.08	76.15	78.19	71.95	77.90
USynthFace [7] + RandAugment [7]	✓	200 K	1	200 K	91.93	71.23	76.73	78.03	72.27	78.04
USynthFace [7] + RandAugment [7]	✓	400 K	1	400 K	92.23	71.62	77.05	78.56	72.03	79.30
DigiFace-1M [2]	✗	10 K	50	500 K	88.07	60.92	69.23	70.99	66.73	71.19
	+ Augmentation[2]	✓	10 K	50	500 K	95.40	76.97	78.62	87.40	78.87
ExFaceGAN(SG3) [8] + RandAugment [7]	✓	10 K	50	500 K	90.47	72.85	78.60	72.70	69.27	76.78
ExFaceGAN(Con) [8] + RandAugment [7]	✓	10 K	50	500 K	93.50	78.92	82.98	73.84	71.60	80.17
IDnet [33]	✗	10.5 K	50	528 K	84.83	63.58	71.50	70.43	67.35	71.54
	+ RandAugment [7]	✓	10.5 K	50	528 K	92.58	73.53	79.90	75.40	74.25
IDiff-Face CPD25 (Uniform) - Ours	✗	5 K	16	80 K	94.37	76.70	85.02	70.31	72.18	79.72
	✗	5 K	32	160 K	96.23	80.50	88.30	73.87	75.42	82.86
	✗	10 K	16	160 K	96.72	81.52	88.83	75.43	76.85	83.87
	✗	10 K	50	500 K	97.68	84.63	90.58	82.39	79.70	87.00
	+ RandAugment [7]	✓	10 K	50	500 K	98.00	86.43	90.65	<u>85.47</u>	80.45
IDiff-Face CPD25 (Two-Stage) - Ours	✗	5 K	16	80 K	95.27	75.72	84.63	70.07	72.65	79.67
	✗	5 K	32	160 K	95.92	77.85	85.40	72.97	74.18	81.26
	✗	10 K	16	160 K	96.60	79.87	87.03	74.47	75.47	82.69
	✗	10 K	50	500 K	97.52	81.65	87.77	80.03	77.60	84.91
	+ RandAugment [7]	✓	10 K	50	500 K	97.32	83.45	87.98	81.80	77.92
IDiff-Face CPD50 (Two-Stage) - Ours	✗	5 K	16	80 K	95.58	74.00	83.57	70.43	72.53	79.22
	✗	5 K	32	160 K	96.40	78.23	85.87	73.00	75.38	81.78
	✗	10 K	16	160 K	97.00	80.85	86.38	74.24	76.73	83.04
	✗	10 K	50	500 K	97.87	83.53	89.05	80.00	78.95	85.88
	+ RandAugment [7]	✓	10 K	50	500 K	<u>97.97</u>	84.40	88.52	83.87	<u>79.88</u>

Table 3: Verification Accuracies (in %) on five FR benchmarks for SOTA synthetic-based FRs. the first two rows present the results of FRs trained on authentic data. These results are provided as references. All results of previous work are copied from their corresponding works. The synthetic-based FR utilized ResNet-50 as network architecture. The best verification accuracies of synthetic-based FR are marked in **bold** and the second best are underlined.

datasets, 99.52% using CASIA-WebFace and 99.82% using MS1MV2. This achieved accuracy (98.00%) on LFW is even higher than human-level performance in face verification on LFW (97.5% [35]).

6. Conclusion

An ideal training dataset for FR would strongly ensure identity discrimination and exhibit a realistic and large intra-class variation. Authentic face datasets proposed in the literature hold, to a large degree, these properties, leading to breakthroughs in verification accuracy. However, future developments using these authentic datasets might be infeasible due to increased legal and ethical concerns about the use and distribution of sensitive authentic data for FR development. In this work, we proposed IDiff-Face, an identity-conditioned generative model based on a DM to

generate synthetic identity-specific face images. We also proposed CPD as a simple, yet effective mechanism to prevent the model from overfitting to the identity context and to control the trade-off between identity-separability and intra-class variation. Utilizing our synthetically generated data for FR training led to new SOTA verification accuracies on five mainstream FR benchmarks, outperforming all recent SOTA synthetic-based FR approaches.

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] Ziga Babnik, Peter Peer, and Vitomir Struc. Diffiqa: Face image quality assessment using denoising diffusion probabilistic models. *CoRR*, abs/2305.05768, 2023. [1](#)
- [2] Gwangbin Bae, M. D. L. Gorce, Tadas Baltrušaitis, Charlie Hewitt, Dong Chen, Julien P. C. Valentin, Roberto Cipolla, and JingJing Shen. Digiface-1m: 1 million digital face images for face recognition. *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3515–3524, 2022. [1](#), [3](#), [6](#), [8](#), [9](#)
- [3] Ankan Bansal, Anirudh Nanduri, Carlos Domingo Castillo, Rajeev Ranjan, and Rama Chellappa. Umdfaces: An annotated face dataset for training deep networks. *2017 IEEE International Joint Conference on Biometrics (IJCB)*, pages 464–473, 2016. [1](#)
- [4] Fadi Boutros, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Elasticface: Elastic margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1578–1587, 6 2022. [1](#), [4](#), [6](#), [9](#)
- [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. *2022 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–11, 2022. [1](#), [2](#), [3](#), [6](#), [8](#), [9](#)
- [7] Fadi Boutros, Marcel Klemm, Meiling Fang, Arjan Kuijper, and Naser Damer. Unsupervised face recognition using unlabeled synthetic data. *2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–8, 2022. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#), [9](#)
- [8] Fadi Boutros, Marcel Klemm, Meiling Fang, Arjan Kuijper, and Naser Damer. Exfacegan: Exploring identity directions in gan’s learned latent space for synthetic identity generation. In *IEEE International Joint Conference on Biometrics, IJCB 2023*, September 2023. [3](#), [8](#), [9](#)
- [9] Fadi Boutros, Patrick Siebke, Marcel Klemm, 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] Fadi Boutros, Vitomir Struc, Julian Fierrez, and Naser Damer. Synthetic data for face recognition: Current state and future prospects. *Image Vis. Comput.*, 135:104688, 2023. [1](#), [2](#)
- [11] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 67–74, 2017. [1](#)
- [12] Ekin Dogus Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. Randaugment: Practical automated data augmentation with a reduced search space. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 3008–3017, 2019. [8](#)
- [13] 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](#)
- [14] Jiankang Deng, J. Guo, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4685–4694, 2018. [1](#), [2](#), [6](#)
- [15] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5153–5162, 2020. [1](#), [3](#), [6](#), [8](#)
- [16] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 8780–8794, 2021. [2](#), [3](#)
- [17] Chris Donahue, Zachary C. Lipton, Akshay Balsubramani, and Julian J. McAuley. Semantically decomposing the latent spaces of generative adversarial networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [3](#)
- [18] European Commission. 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) (Text with EEA relevance), 2016. [1](#)
- [19] Meiling Fang, Marco Huber, and Naser Damer. Synthaspoof: Developing face presentation attack detection based on privacy-friendly synthetic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1061–1070, June 2023. [1](#)
- [20] Biying Fu, Marcel Klemm, Fadi Boutros, and Naser Damer. On the quality and diversity of synthetic face data and its relation to the generator training data. In *11th International Workshop on Biometrics and Forensics, IWBF 2023, Barcelona, Spain, April 19-20, 2023*, pages 1–6. IEEE, 2023. [3](#)
- [21] Partha Ghosh, Pravir Singh Gupta, Roy Uziel, Anurag Ranjan, Michael J. Black, and Timo Bolkart. GIF: Generative interpretable faces. In *International Conference on 3D Vision (3DV)*, pages 868–878, 2020. [1](#), [3](#)
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahra-

- mani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. [1](#), [2](#)
- [23] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *European Conference on Computer Vision*, 2016. [1](#), [2](#), [6](#), [9](#)
- [24] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2015. [1](#), [4](#), [6](#), [7](#), [8](#)
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. 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. [2](#), [3](#), [4](#)
- [26] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, 11 2007. [2](#), [6](#), [7](#), [8](#)
- [27] ISO/IEC JTC1 SC37 Biometrics. ISO/IEC 19795-1:2021 Information technology — Biometric performance testing and reporting — Part 1: Principles and framework. International Organization for Standardization, 2021. [6](#)
- [28] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. [2](#)
- [29] 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. [1](#), [3](#)
- [30] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4396–4405, 2018. [2](#), [4](#)
- [31] Minchul Kim, Anil K. Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18729–18738, 2022. [1](#)
- [32] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [4](#)
- [33] Jan Niklas Kolf, Tim Rieber, Jurek Elliesen, Fadi Boutros, Arjan Kuijper, and Naser Damer. Identity-driven three-player generative adversarial network for synthetic-based face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 806–816, June 2023. [1](#), [3](#), [8](#), [9](#)
- [34] Marek Kowalski, Stephan J. Garbin, Virginia Estellers, Tadas Baltrusaitis, Matthew Johnson, and Jamie Shotton. CONFIG: controllable neural face image generation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XI*, volume 12356 of *Lecture Notes in Computer Science*, pages 299–315. Springer, 2020. [3](#)
- [35] Neeraj Kumar, Alexander C. Berg, Peter N. Belhumeur, and Shree K. Nayar. Attribute and simile classifiers for face verification. In *IEEE 12th International Conference on Computer Vision, ICCV 2009, Kyoto, Japan, September 27 - October 4, 2009*, pages 365–372. IEEE Computer Society, 2009. [2](#), [9](#)
- [36] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. [4](#)
- [37] Richard T. Marriott, Safa Madiouni, Sami Romdhani, Stéphane Gentic, and Liming Chen. An assessment of gans for identity-related applications. *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10, 2020. [1](#)
- [38] George Marsaglia. Choosing a point from the surface of a sphere. *Annals of Mathematical Statistics*, 43:645–646, 1972. [4](#)
- [39] 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 (CVPRW)*, pages 1997–2005, 2017. [7](#)
- [40] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8162–8171. PMLR, 2021. [2](#)
- [41] Norman Poh and Samy Bengio. A study of the effects of score normalisation prior to fusion in biometric authentication tasks. Technical report, IDIAP, 2004. [6](#)
- [42] Deepti S Prakash, Lucia E Ballard, Jerrold V Hauck, Feng Tang, Etai Littwin, Pavan Kumar Ansalu Vasu, Gideon Littwin, Thorsten Gernoth, Lucie Kucerova, Petr Kostka, et al. Biometric authentication techniques, Feb. 23 2021. US Patent 10,929,515. [1](#)
- [43] Haibo Qiu, Baosheng Yu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. Synface: Face recognition with synthetic data. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10860–10870, 2021. [1](#), [2](#), [3](#), [6](#), [8](#), [9](#)
- [44] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv*, abs/2204.06125, 2022. [2](#), [3](#)
- [45] Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis

- with latent diffusion models. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, 2021. [2](#), [3](#), [4](#), [5](#), [6](#)
- [46] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015. (available on arXiv:1505.04597 [cs.CV]). [3](#)
- [47] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022. [2](#), [3](#)
- [48] 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. *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9, 2016. [7](#)
- [49] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9240–9249, 2020. [1](#), [2](#), [3](#)
- [50] 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 *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018*, 2018. [3](#)
- [51] Yujun Shen, Bolei Zhou, Ping Luo, and Xiaoou Tang. Facefeat-gan: a two-stage approach for identity-preserving face synthesis, 2018. [2](#), [3](#)
- [52] Alon Shoshan, Nadav Bhonker, Igor Kviatkovsky, and Gerard Medioni. Gan-control: Explicitly controllable gans. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14063–14073, 2021. [1](#), [3](#)
- [53] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. [4](#)
- [54] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zöllhofer, and Christian Theobalt. Stylerig: Rigging style-gan for 3d control over portrait images, cvpr 2020. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 6 2020. [1](#), [2](#), [3](#)
- [55] Luan Tran, Xi Yin, and Xiaoming Liu. Representation learning by rotating your faces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(12):3007–3021, 12 2019. [2](#), [3](#)
- [56] H. Wang, Yitong Wang, Zheng Zhou, Xing Ji, Zhifeng Li, Dihong Gong, Jin Zhou, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018. [1](#), [4](#), [7](#)
- [57] 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. [6](#)
- [58] Erroll Wood, Tadas Baltruvsaitis, Charlie Hewitt, Sebastian Dziadzio, Matthew Johnson, Virginia Estellers, Thomas J. Cashman, and Jamie Shotton. Fake it till you make it: face analysis in the wild using synthetic data alone. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3661–3671, 2021. [8](#)
- [59] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z. Li. Learning face representation from scratch. *CoRR*, abs/1411.7923, 2014. [2](#), [6](#), [8](#), [9](#)
- [60] Tianyue Zheng and Weihong Deng. Cross-pose lfw : A database for studying cross-pose face recognition in unconstrained environments. 2018. [7](#)
- [61] Tianyue Zheng, Weihong Deng, and Jiani Hu. Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments. 08 2017. [7](#)