

DCFace: Synthetic Face Generation with Dual Condition Diffusion Model

Minchul Kim
kimminc2@msu.edu

Feng Liu
liufeng6@msu.edu

Anil Jain
jain@msu.edu

Xiaoming Liu
liuxm@msu.edu

Michigan State University
East Lansing, MI 48824

Abstract

Generating synthetic datasets for training face recognition models is challenging because dataset generation entails more than creating high fidelity images. It involves generating multiple images of same subjects under different factors (e.g., variations in pose, illumination, expression, aging and occlusion) which follows the real image conditional distribution. Previous works have studied the generation of synthetic datasets using GAN or 3D models. In this work, we approach the problem from the aspect of combining subject appearance (ID) and external factor (style) conditions. These two conditions provide a direct way to control the inter-class and intra-class variations. To this end, we propose a Dual Condition Face Generator (DCFace) based on a diffusion model. Our novel Patch-wise style extractor and Time-step dependent ID loss enables DCFace to consistently produce face images of the same subject under different styles with precise control. Face recognition models trained on synthetic images from the proposed DCFace provide higher verification accuracies compared to previous works by 6.11% on average in 4 out of 5 test datasets, LFW, CFP-FP, CPLFW, AgeDB and CALFW. [Code Link](#)

1. Introduction

What does it take to create a good training dataset for visual recognition? An ideal training dataset for recognition tasks would have 1) large inter-class variation, 2) large intra-class variation and 3) small label noise. In the context of face recognition (FR), it means, the dataset has a large number of unique subjects, large intra-subject variations, and reliable subject labels. For instance, large-scale face datasets such as WebFace4M [73] contain over 1M subjects and large number of images/subject. Both the number of subjects and the number of images per subject are important for training FR models [11, 31]. Also, datasets amassed by crawling the web are not free from label noise [7, 73].

In various domains, synthetic datasets are traditionally




Figure 1. Illustration of three factors that characterize a labeled face dataset. It contains large subject variation, style variation and label consistency. Synthetic face datasets should be created with all three factors in mind. Face images in this figure are samples generated by our proposed method which combines arbitrary ID condition with style condition while preserving subject identity.

used to help generalize deep models when only limited real datasets could be collected [13, 21, 63, 74] or when bias exists in the real dataset [34, 64]. Lately, more attention has been drawn to training with only synthetic datasets in the face domain, as synthetic data can avoid leaking the privacy of real individuals. This is important as real face datasets have been under scrutiny for their lack of informed consent, as web-crawling is the primary means of large-scale data collection [17, 22, 73]. Also, synthetic training datasets can remedy some long-standing issues in real datasets, e.g. the long tail distribution, demographic bias, etc.

When it comes to generating synthetic training datasets, the following questions should be raised. (i) How many novel subjects can be synthesized (ii) How well can we mimic the distribution of real images in the target domain and (iii) How well can we consistently generate multiple images of the same subjects? We start with the hypothesis that face dataset generation can be formulated as a problem that maximizes these criteria together.

Previous efforts in generating synthetic face datasets touch on one of the three aspects but do not consider all of them together [3, 49]. SynFace [49] generates high-fidelity face images based on DiscoFaceGAN [12], coming close to real images in terms of FID metric [18]. However, we were

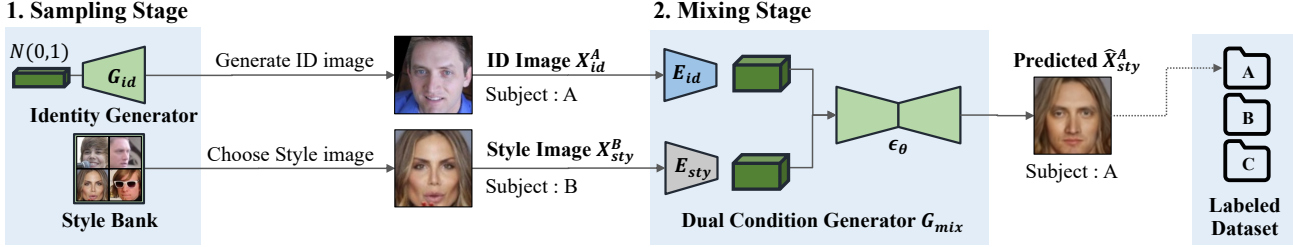


Figure 2. Two stage dataset generation paradigm. In the sampling stage, 1) G_{id} generates a high-quality face image \mathbf{X}_{id} that defines how a person looks and 2) the style bank selects a style image \mathbf{X}_{sty} that defines the overall style of the final image. The mixing stage generates image with identity from \mathbf{X}_{id} and style from \mathbf{X}_{sty} . Repeating this process multiple times, one can generate a labeled synthetic face dataset.

surprised to find that the actual number of unique subjects that can be generated by DiscoFaceGAN is less than 500, a finding that will be discussed in Sec. 3.1. The recent state of the art (SoTA), DigiFace [3], can generate 1M large-scale synthetic face images with many unique subjects based on 3D parametric model rendering. However, it falls short in matching the quality and style of real face images.

We propose a new data generation scheme that addresses all three criteria, *i.e.* the large number of novel subjects (*uniqueness*), real dataset style matching (*diversity*) and label consistency (*consistency*). In Fig. 1, we illustrate the high-level idea by showcasing some of our generated face samples. The key motivation of our paper is that the synthetic dataset generator needs to control the number of unique subjects, match the training dataset’s style distribution and be consistent in the subject label.

In light of this, we formulate the face image generation as a dual condition inverse problem, retrieving the unknown image \mathbf{Y} from the observable Identity condition \mathbf{X}_{id} and Style condition \mathbf{X}_{sty} . Specifically, \mathbf{X}_{id} specifies how a person looks and \mathbf{X}_{sty} specifies how \mathbf{X}_{id} should be portrayed in an image. \mathbf{X}_{sty} contains identity-independent information such as pose, expression, and image quality.

Our choice of dual conditions (identity and style) is important in how we generate a synthetic dataset as ID and style conditions are controllable factors that govern the dataset’s characteristics. To achieve this, we propose a two-stage generation paradigm. First, we generate a high-quality face image \mathbf{X}_{id} using a face image generator and sample a style image \mathbf{X}_{sty} from a style bank. Secondly, we mix these two conditions using a dual condition generator which predicts an image that has the ID of \mathbf{X}_{id} and a style of \mathbf{X}_{sty} . An illustration is given in Fig. 2.

Training the mixing generator in stage 2 is not trivial as it would require a triplet of $(\mathbf{X}_{id}^A, \mathbf{X}_{sty}^B, \mathbf{X}_{sty}^A)$ where \mathbf{X}_{sty}^A is a hypothetical combination of the ID of subject A and the style of subject B . To solve this problem, we propose a new dual condition generator that can learn from $(\mathbf{X}_{id}^A, \mathbf{X}_{sty}^A)$, a tuple of same subject images that can always be obtained in a labeled dataset. The novelty lies in our style condi-

tion extractor and ID loss which prevents the training from falling into a degenerate solution. We modify the diffusion model [19, 55] to take in dual conditions and apply an auxiliary time-dependent ID loss that can control the balance between sample diversity and label consistency.

We show that our Dual Condition Face Dataset Generator (DCFace) is capable of surpassing the previous methods in terms of FR performance, establishing a new benchmark in face recognition with synthetic face datasets. We also show the roles dataset subject uniqueness, diversity and consistency play in face recognition performance.

The followings are the contributions of the paper.

- We propose a two-stage face dataset generator that controls subject uniqueness, diversity and consistency.
- For this, we propose a dual condition generator that mixes the two independent conditions \mathbf{X}_{id} and \mathbf{X}_{sty} .
- We propose uniqueness, consistency and diversity metrics that quantify the respective properties of a given dataset, useful measures that allow one to compare datasets apart from the recognition performance.
- We achieve SoTA in FR with 0.5M image synthetic training dataset by surpassing the previous methods by 6.11% on average in 5 popular test datasets.

2. Related Works

Face Recognition. Face Recognition (FR) is the task of matching query imagery to an enrolled identity database. SoTA FR models are trained on large-scale web-crawled datasets [11, 17, 73] with margin-based softmax losses [11, 24, 31, 39, 65]. The FR performance is measured on various benchmark datasets such as LFW [23], CFP-FP [53], CPLFW [71], AgeDB [43] and CALFW [72]. These datasets are designed to measure factors such as pose changes and age variations. Performance on these datasets for models trained on large-scale datasets such as Web-Face260M is well above 97% [31] in verification accuracy.

Synthetic Face Generation. Recent advances in generative models allow high fidelity synthetic face image generations [6, 8, 19, 27–29, 56]. GANs have been widely used to

manipulate, animate or enhance face images [8, 12, 20, 37, 48, 61, 62, 69]. They typically learn disentangled representations in GAN latent space that control desired face properties. On the contrary, some works leverage the 3D face prior from 3D datasets (*e.g.*, 3DMM [4]) for controllable synthesis [9, 14, 15, 30, 42, 44, 46, 54]. These methods have advantages in the fine-grained control over face generation and 3D consistency yet lack in style or domain variation.

Recent advances in the latent variable models such as diffusion or score-based models have shown great success in high-quality image generation with a more stable and simple objective of MSE loss [19, 45, 55–59]. Diffusion models have advanced the conditional image generation in tasks such as text-conditional image generation, inpainting, *etc* [5, 50, 51, 66]. We adopt the diffusion model as a backbone and explore how the two image characteristics, namely ID and style images, can control complementary information, the subject appearance and the style of an image.

Face Recognition with Synthetic Dataset. Synthetic training datasets offer an advantage over real datasets with regards to ethical issues and class imbalance problems as large-scale face datasets have been criticized for lacking informed consent and reflecting racial biases [3, 11, 70, 73]. Despite the benefit, use of synthetic datasets as the sole training data is not widely adopted due to the resulting low recognition performance. In various domains such as face recognition [3, 38, 49], fingerprint recognition [13, 68], and anti-spoofing [40, 60], synthetic datasets have been shown to improve recognition when combined with real images.

In the face domain, SynFace [49] studied the efficacy of using DiscoFaceGAN [12] for synthetic face generation. Recently, DigiFace-1M [3] studied the efficacy of 3D model based face rendering in combination with image augmentations to create a synthetic dataset. We propose a face dataset generation method that can generate both a large number of subjects and diverse styles that are close to the real dataset.

3. Proposed Approach

We propose Dual Condition Face Dataset Generator (DCFace), a two-stage dataset generator (see Fig. 2). Stage 1 is the Condition Sampling Stage, generating a high-quality ID image (\mathbf{X}_{id}) of a novel subject and selects one arbitrary style image (\mathbf{X}_{sty}) from the bank of real training data. Stage 2 is the Mixing Stage which combines the two images using the Dual Condition Generator.

For trainable models in each stage, Stage 1 requires training an ID image generator G_{id} . For the style bank, we can conveniently use any real face dataset that we wish generated samples to follow. Stage 2 requires training a dual condition mixer G_{mix} . Both G_{id} and G_{mix} are based on diffusion models [19]. We describe each component and the associated training procedure in the following subsections.

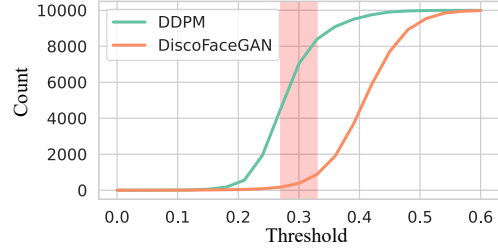


Figure 3. Comparison of the number of unique subjects generated by DiscoFaceGAN [12] and unconditional DDPM [19]. Uniqueness is the number of unique subjects measured by a face recognition model. By varying the threshold which determines a match between two subjects, we plot the number of unique subjects as defined in Eq. 11. Unconditional DDPM and DiscoFaceGAN are trained on FFHQ [28] and each generates 10,000 samples. The ability to generate novel subjects is larger for DDPM. Refer to Supp.E for additional details on the threshold.

3.1. Preliminary

Diffusion models [19, 55] are a class of denoising generative models that are trained to predict an image from random noise through a gradual denoising process. One notable difference from the class of GAN-based generators [16] is in the objective function and the sampling procedure. The forward process as expressed in Eq. 1 corrupts the input \mathbf{X} using variance controlled Gaussian noise over t time-steps,

$$q(\mathbf{X}_t | \mathbf{X}_{t-1}) = \mathcal{N}(\mathbf{X}_t; \sqrt{1 - \beta_t} \mathbf{X}_{t-1}, \beta_t \mathbf{I}), \quad (1)$$

and the denoising is done by training a model $\epsilon_\theta(\mathbf{X}_t, t)$ to predict the initial noise ϵ with an L_2 objective,

$$\mathcal{L} = \mathbb{E}_{t, \mathbf{X}_0, \epsilon} \left[\left\| \underbrace{\epsilon_\theta(\sqrt{\alpha_t} \mathbf{X}_0 + \sqrt{1 - \alpha_t} \epsilon, t)}_{\mathbf{X}_t} - \epsilon \right\|_2^2 \right]. \quad (2)$$

β_t and α_t are pre-set variance scheduling scalars. The denoising diffusion model (DDPM) has shown success in producing diverse samples in text-conditioned image generation [50]. We find that in unconditional face generation, DDPM is also capable of generating many unique subjects. For instance, Fig. 3 compares DiscoFaceGAN [12] with DDPM [19] in their capacity to generate different subjects for every sample. It shows that DDPM [19] is a good model choice for G_{id} and G_{mix} as it can generate many unique subjects. For G_{id} , we adopt the unconditional DDPM trained on FFHQ [28], having observed that it is capable of generating a large number of unique subject images.

3.2. Dual Condition Generator G_{mix}

The two-stage data generation requires Dual Condition Generator G_{mix} which is a conditional DDPM. Two conditions \mathbf{X}_{id} and \mathbf{X}_{sty} are injected into the denoiser $\epsilon_\theta(\mathbf{X}_t, t, E_{id}(\mathbf{X}_{id}), E_{sty}(\mathbf{X}_{sty}))$ using trainable feature

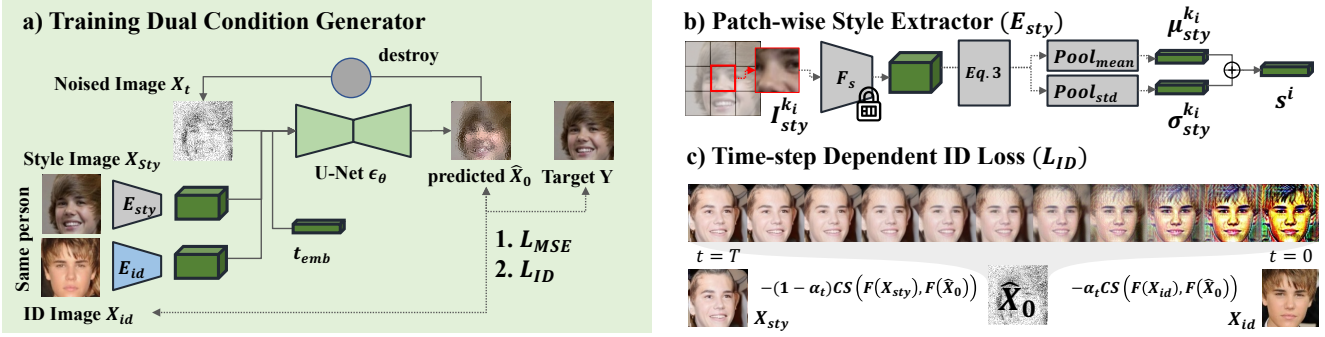


Figure 4. a) A diagram of G_{mix} during training. At each step, we draw two labeled images from the labeled training dataset and use them as \mathbf{X}_{id} and \mathbf{X}_{sty} . We ensure \mathbf{X}_{id} to be the good-quality frontal view image. t_{emb} is the time-step embedding in DDPM [19]. \mathbf{X}_{sty} also serves as a target image and we apply Gaussian noise ϵ to \mathbf{X}_{sty} to create \mathbf{X}_t as DDPM specifies. Then $\epsilon_\theta(\mathbf{X}_t, t, \mathbf{X}_{id}, \mathbf{X}_{sty})$ is trained to predict ϵ using L_{MSE} , conceptually equivalent to the reconstruction loss to recover \mathbf{X}_{sty} . We also apply L_{ID} as in Eq. 10 for the dependence on \mathbf{X}_{id} . b) Patch-wise Style Extractor generates style vectors from small patches of images. Style vectors are architecturally constrained from containing full ID information. c) Time-step dependent ID Loss is a linear interpolation between the \mathbf{X}_{id} and \mathbf{X}_{sty} in the recognition feature space. It forces ϵ_θ to rely on \mathbf{X}_{id} to extract the subject’s appearance and gradually shift the style to \mathbf{X}_{sty} .

extractors E_{id} and E_{sty} and cross-attentions. G_{mix} is responsible for the operation $\mathbf{X}_{id}^A + \mathbf{X}_{sty}^B \rightarrow \mathbf{X}_{sty}^A$, a mixing of an image of a novel subject A and an arbitrary style image of different subject B .

Naive training would require the reference image \mathbf{X}_{sty}^A , an image of subject A in the style of \mathbf{X}_{sty}^B . This reference is absent in the labeled training dataset. As such, we modify the operation to $\mathbf{X}_{id}^A + \mathbf{X}_{sty}^A \rightarrow \mathbf{X}_{sty}^A$, using two different images from the same subject as illustrated in Fig. 4(a). But this formulation is prone to a trivial solution of ignoring \mathbf{X}_{id}^A , making the ID condition unused during test time. To mitigate this issue, we propose the following two elements.

Patch-wise Style Extractor E_{sty} . The motivation of Style Extractor is to map an image \mathbf{X}_{sty} to a feature that contains little ID information, forcing G_{mix} to rely on \mathbf{X}_{id} for ID information. In prior works such as StyleGAN, 1st and 2nd order statistics of a feature are shown to resemble the image style [28, 32, 36]. Yet, resulting statistics are reduced in spatial dimensions and consequently without spatially local informations such as pose.

We propose a module that can extract style information without losing spatial information. Specifically, consider a pretrained and fixed face recognition model F_s and its intermediate feature $F_s(\mathbf{X}_{sty}) = \mathbf{I}_{sty} \in \mathbb{R}^{C \times H \times W}$. We divide the feature into a $k \times k$ grid. For each element in the grid $\mathbf{I}_{sty}^{k_i} \in \mathbb{R}^{C \times \frac{H}{k} \times \frac{W}{k}}$, we perform non-linear mapping on the mean and variance of $\mathbf{I}_{sty}^{k_i}$. Specifically,

$$\hat{\mathbf{I}}^{k_i} = \text{BN}(\text{Conv}(\text{ReLU}(\text{Dropout}(\mathbf{I}_{sty}^{k_i}))))), \quad (3)$$

$$\mu_{sty}^{k_i} = \text{SpatialMean}(\hat{\mathbf{I}}^{k_i}), \quad \sigma_{sty}^{k_i} = \text{SpatialStd}(\hat{\mathbf{I}}^{k_i}), \quad (4)$$

$$\mathbf{s}^{k_i} = \text{LN} \left((\mathbf{W}_1 \odot \mu_{sty}^{k_i} + \mathbf{W}_2 \odot \sigma_{sty}^{k_i}) + \mathbf{P}_{emb} \right), \quad (5)$$

$$E_{sty}(\mathbf{X}_{sty}) := \mathbf{s} = [\mathbf{s}^1, \mathbf{s}^2, \mathbf{s}^{k_i}, \dots, \mathbf{s}^{k \times k}, \mathbf{s}^l], \quad (6)$$

where \mathbf{s}^l corresponds to $\mathbf{I}_{sty}^{k_i}$ being a global feature, where $k = 1$. The final output \mathbf{s} is a concatenation of all style vectors for each patch. Each \mathbf{s}^{k_i} is a mean and variance of local information which is constrained from containing full pixel-level details with the ID information. And \mathbf{P}_{emb} is a learned position embedding to let the model differentiate different patch locations. BN and LN are BatchNorm [25] and LayerNorm [2]. F_s is a shallow CNN taken from the early layers of a pretrained FR model. It is fixed and not updated to prevent it from optimizing \mathbf{I}_{sty} , serving only to create style information. By varying the grid size $k \times k$, we can represent style at different spatial locations. An illustration of E_{sty} can be found in Fig. 4(b).

Time-step Dependent ID Loss. To train Dual Condition Generator G_{mix} , the original DDPM objective of L_2 loss, Eq. 2 is not sufficient to guarantee the consistency in subject identity between the ID condition \mathbf{X}_{id} and the prediction, $\hat{\mathbf{X}}_0$. To ensure the ID consistency, one could devise a loss function to maximize the similarity between \mathbf{X}_{id} and the predicted denoised image $\hat{\mathbf{X}}_0$, in the ID feature space using a pretrained FR model, F . Specifically, following the Eq. 15 of DDPM [19], one-step prediction of the original image is

$$\hat{\mathbf{X}}_0 = (\mathbf{X}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{X}_t, t, \mathbf{X}_{id}, \mathbf{X}_{sty})) / \sqrt{\bar{\alpha}_t}. \quad (7)$$

A simple ID loss to increase cosine similarity (CS) is

$$L_{naive1} = -\text{CS} \left(F(\mathbf{X}_{id}), F(\hat{\mathbf{X}}_0) \right). \quad (8)$$

However, this loss is in conflict with MSE loss and is empirically observed to reduce the predicted image quality. This is because the FR model, F is not invariant to image style; some style of \mathbf{X}_{id} has to match in order to completely reduce L_{naive1} . In contrast, one could also use

$$L_{naive2} = -\text{CS} \left(F(\mathbf{X}_{sty}), F(\hat{\mathbf{X}}_0) \right), \quad (9)$$

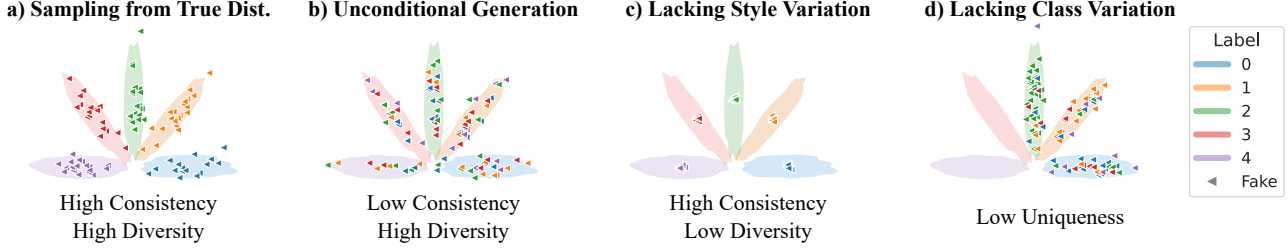


Figure 5. Illustration of conditional distributions in 2D space. Colored regions represent the true data distribution with individual colors representing different labels. Colored triangles represent generated samples with corresponding labels. For each scenario except (a), the generated distribution does not follow the true distribution. Consistency, diversity and uniqueness analysis can quantify the shortcomings.

as during training the label of \mathbf{X}_{sty} and \mathbf{X}_{id} are the same. However, L_{naive2} causes the model to depend on \mathbf{X}_{sty} for ID information. Thus, during evaluation, when \mathbf{X}_{sty} and \mathbf{X}_{id} are different subjects, the label consistency in the generated dataset is compromised. We show this in Tab. 2.

Instead, we propose to interpolate between $F(\mathbf{X}_{id})$ and $F(\mathbf{X}_{sty})$ across diffusion time-steps. Specifically,

$$L_{ID} = -\gamma_t \text{CS} \left(F(\mathbf{X}_{id}), F(\hat{\mathbf{X}}_0) \right) - (1 - \gamma_t) \text{CS} \left(F(\mathbf{X}_{sty}), F(\hat{\mathbf{X}}_0) \right), \quad (10)$$

where $\gamma_t = \frac{t}{T}$ is a time-dependent weight that linearly changes from 0 to 1. When $t = T$, ϵ_θ is predicting \mathbf{X}_{t-1} from random noise, and we let the model fully exploit the ID information of \mathbf{X}_{id} . Gradually as t increases, we let the model’s prediction walk into the direction of \mathbf{X}_{sty} . Note that during training, the actual label of \mathbf{X}_{sty} and \mathbf{X}_{id} are the same. So the interpolation in the loss forces the prediction to be the same in identity but gradually shifting in style toward \mathbf{X}_{sty} . This loss allows $\epsilon_\theta(\mathbf{X}_t, t, \mathbf{X}_{id}, \mathbf{X}_{sty})$ to play different roles depending on t . For $t \approx T$, ϵ_θ will exploit \mathbf{X}_{id} to infer front-view ID rich image. And as $t \rightarrow 0$, it will change the image’s style to match the style of \mathbf{X}_{sty} . The final loss is $L_{MSE} + \lambda L_{ID}$ with λ as a scaling parameter.

E_{id} and Conditioning Mechanism. Following the success text-conditional image generation and inpainting using DDPM [47, 50, 66], we adopt a similar architecture for inserting conditions into the model. We concatenate $E_{id}(\mathbf{X}_{id})$ and $E_{sty}(\mathbf{X}_{sty})$ and put in ϵ_θ using cross-attention and adaptive group normalization layers (AdaGN) [47]. E_{id} is a CNN, with the same architecture as a small FR model (e.g. ResNet50). And E_{id} is trained end-to-end with ϵ_θ to extract useful ID feature for ϵ_θ . More training details can be found in Supp.

3.3. Condition Sampling Strategy

ID Image Sampling. For sampling ID images, we generate 200,000 facial images from G_{ID} , from which we remove faces that are wearing sunglasses or too similar to the subjects in CASIA-WebFace with the Cosine Similarity threshold of 0.3 using F_{eval} . We are left with 105,446

	White	Asian	Others	Black	Indian
CASIA-WebFace	0.634	0.144	0.074	0.074	0.072
DDPM G_{id}	0.660	0.209	0.034	0.046	0.048
Balanced Ethnicity	0.200	0.200	0.200	0.200	0.200

Table 1. Ethnicity Distribution of CASIA-WebFace. Ethnicity prediction is made using [1]. DDPM G_{id} is trained on FFHQ [28].

images. Then we narrow them down to 62,570 images that are unique according to uniqueness, Eq. 11 using F_{eval} and $r = 0.3$. Then we explore two different options, 1) random sampling and 2) gender/ethnicity balanced sampling as G_{id} has a skewed distribution towards White subjects as shown in Tab. 1. We use [1] to classify the ethnicity and use [26, 67] to detect sunglasses. We denote the sampling option 1 as *random* and 2 as *balance*.

Style Image Sampling. For style sampling, for each \mathbf{X}_{id} , we randomly sample \mathbf{X}_{sty} from the style bank. We denote this option as *random*. We also explore the option of sampling \mathbf{X}_{sty} from the pool of images whose gender/ethnicity matches that of \mathbf{X}_{id} . We denote this option as *match*.

4. Dataset Evaluation

In evaluating the synthesized dataset, one often adopts 1) FID [18] for evaluating the distribution similarity to the real images and 2) subsequent recognition performance. In this section, we propose three class-dependent metrics that aid us in understanding the property of generated labeled datasets. We let F_{eval} be an recognition model used for evaluating synthesized face datasets. Note that this is different from F in ID loss. F is a model for training loss and F_{eval} is for evaluating metrics. The more generalizable F_{eval} is, the more accurate the metrics become in capturing the identity and diversity of the synthesized dataset.

Let y_c be a class label, and $f_i = F_{eval}(\mathbf{X}_i)$. Let $d(f_i, f_j)$ be the distance between two images in F_{eval} feature space.

Uniqueness. Consider the following non-overlapping r -ball in F_{eval} space,

$$U = \{f_i : d(f_i, f_j) > r, j < i, i, j \in \{1, \dots, N\}\}, \quad (11)$$

where $d(f_i, f_j)$ is the cosine distance. Then $|U|$ is the count of unique subjects determined by the threshold r in an un-

labeled dataset. Note that the set U is equivalent to sequentially adding a r -ball into F_{eval} -space until you cannot add more without collision. $|U|$ is subject to both r and F_{eval} . In FR, r is a threshold in the FR model that is set to determine match or non-match.

For a labeled synthetic dataset, one generates multiple feature sets $\{f_i^c\}$ for the same label. To count the number of unique subjects, we calculate the number of unique centers, $f^c = \frac{1}{N_c} \sum_i^{N_c} f_i^c$ for $c \in \{1, \dots, C\}$, where C is the number of subjects and N_c is the number of images per subject. Then we define the number of unique subjects in a labeled dataset with $|U_c|$ where U_c is

$$U_c = \{f_c : d(f^{c_n}, f^{c_m}) > r, m < n, n, m \in \{1, \dots, C\}\}, \quad (12)$$

For the metric, we use $U_{class} = |U_c|/C$, the ratio between the number of unique subjects and the number of labels.

Intra-class Consistency. It measures how consistent the generated samples are in adhering to the label condition, as

$$C_{intra} = \frac{1}{C} \sum_{c=1}^C \frac{1}{N_c} \sum_{i=1}^{N_c} d(f_i^c, f^c) < r, \quad (13)$$

which is the ratio of individual features f_i^c being close to the class center f^c . For a given threshold r , higher values of C_{intra} mean the samples are more likely to be the same subject under the same label.

Intra-class Diversity. It measures how diverse the generated samples are under the same label condition. Note that the diversity is in the style of an image, not in the subject's identity. We define the style space as a vector space defined by Inception Network [52] features pretrained on ImageNet [10] following the convention of [35], denoting the real and generated image inception vectors as $\{s_i^c\}, \{\hat{s}_j^c\}$.

For intra-class diversity, we measure how many real images fall into the style space manifold defined by the generated images under the same label condition. We compute this by extending the Improved Recall Metric [35], from comparing the unconditional distributions of real and fake images to comparing the label-conditional distributions. Specifically, for a set of real and generated feature vectors $\{s_i^c\}, \{\hat{s}_j^c\}$ under the same label condition y_c , we define k -nearest feature distance r_k as $r_k = d(\hat{s}_j^c - \text{NN}_k(\hat{s}_j^c, \{s_i^c\}))$, where NN_k returns the k -nearest feature vector in $\{s_i^c\}$ and

$$\mathbf{I}(s_i^c, \{\hat{s}_j^c\}) = \begin{cases} 1, & \exists \hat{s}_j^c \in \{\hat{s}_j^c\} \text{ s.t. } d(s_i^c - \hat{s}_j^c) \leq r_k \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

$d(\cdot)$ is an Euclidean distance. Then diversity is defined by

$$D_{intra} = \frac{1}{C} \frac{1}{N} \sum_{c=1}^C \sum_{i=1}^{N_c} \mathbf{I}(s_i^c, \{\hat{s}_j^c\}), \quad (15)$$

which is the fraction of real image styles manifold covered by the generated image style manifold as defined by

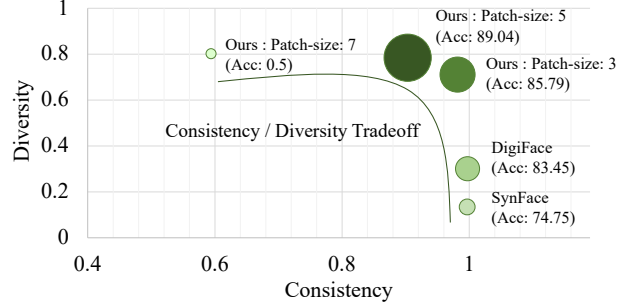


Figure 6. A plot of FR performance on 5 synthetic datasets with respect to Consistency and Diversity metrics. Color intensity and circle size denotes the FR accuracy.

k -nearest neighbor ball. If the style variation is small, then r_k becomes small, reducing the chance of $d(s_i^c - \hat{s}_j^c) \leq r_k$. We compute the recall per class to capture style variation conditional on the subject label.

In Fig. 5, we illustrate different scenarios of conditional generation and how these metrics can capture the shortcomings in each scenario. In Sec. 5 and Fig. 6, we measure the metrics on our generated datasets and compare with previous synthetic datasets [3, 49]. We find that FR performance is at best when consistency and diversity are balanced. Also, we find SynFace and DigiFace have high C_{intra} and low D_{intra} compared to our method in Fig. 5.

5. Experiments

For G_{id} which generates ID images, we adopt the publicly released unconditional DDPM [19] trained on FFHQ [28]. For G_{mix} , we train it on CASIA-WebFace [22] after initializing weights from G_{id} . Although using all of CASIA-WebFace is a valid setting, we split it into a 95-5 split between train and validation sets. The validation set is used as a real dataset in measuring the uniqueness, consistency and diversity metrics. G_{mix} is trained for 10 epochs with a batch-size of 256 using AdamW Optimizer [33, 41] with the learning rate of 0.001. Training takes 8 hours using two A100 GPUs. Once G_{mix} is trained, we use G_{id} , G_{mix} and a style bank to generate a synthetic labeled dataset. The style bank is the CASIA-WebFace training set. For sampling, we use DDIM [56] with 200 intervals. Generating 500K samples takes about 20 hours using one A100 GPU.

To train FR models, for a fair comparison, we adopt the training scheme of [3, 49] using IR-SE-50 [11] as a backbone and AdaFace [31] as a loss function. We evaluate the trained FR models on five datasets, LFW [23], CFP-FP [53], CPLFW [71], AgeDB [43] and CALFW [72]. CFP-FP and CPLFW are designed to measure the FR in the large pose variation and AgeDB and CALFW are for the large age variation. To measure the consistency, diversity and uniqueness during evaluation, we adopt F_{eval} as IR101 [11] model trained on WebFace4M [73] with AdaFace [31] loss.

Grid Size	Loss	Loss Model	U_{class}	C_{intra}	D_{intra}	FR Perf.
SynFace	-	-	0.080	0.9966	0.131	74.75
DigiFace	-	-	0.178	0.9973	0.297	83.45
1×1	L_{ID}	F	0.978	0.9987	0.4418	79.28
3×3			0.956	0.9809	0.7030	85.79
5×5			0.924	0.9035	0.7734	89.04
7×7			0.690	0.5937	0.7950	50.00
5×5			L_{naive1}	F	0.988	0.9996
	L_{naive2}	0.866	0.8046		0.7835	50.00
	L_{ID}	0.924	0.9035		0.7734	89.04
5×5	L_{ID}	F_{bigger}	0.924	0.9035	0.7734	89.04
			0.954	0.9197	0.7715	89.89

Table 2. Model Ablation. For FR performance, we generate a synthetic dataset of 10K subjects with 50 images per subject using (*random*, *random*) ID and style sampling strategy. Blue color indicates the adopted setting for subsequent experiments.

5.1. Model Ablation

To show the efficacy of our proposed modules, we ablate on 1) the grid size in Style extractor E_{sty} , 2) Time-step dependent ID loss and 3) the ID loss backbone F 's. The number of samples we generate for the ablation are 10K subjects with 50 images per subject, similar to CASIA-WebFace image counts. We report the FR performance with the synthetic data by averaging the 5 validation set verification accuracies. To measure U_{class} , C_{intra} and D_{intra} , we use 500 subjects with 20 real images from the held-out validation set of CASIA-WebFace and generate an equivalent number of images from each method.

Grid Size. We choose 4 grid sizes ranging from 1×1 to 7×7. Note that 1×1 corresponds to the style vector of a whole image. We expect to see higher spatial control in X_{sty} as the grid size increases. In Tab. 2, we report the three metrics U_{class} , C_{intra} and D_{intra} . As the grid size increases, E_{sty} features contain more fine-grained information, possibly related to ID, lowering the consistency. However, the diversity increases, making the conditional distribution similar to the real dataset. The subsequent FR performance using the model is the best in the setting 5×5, which is a good compromise between consistency and diversity. In Fig. 7, we show the effect of the grid size with examples.

ID Loss. For ID loss, we compare L_{ID} with L_{naive1} and L_{naive2} in Tab. 2. Using L_{naive1} or L_{naive2} both suffer from lower FR performance, but for different reasons. L_{naive1} has low diversity because it is optimized to be similar to X_{id} of front-view high quality face images. L_{naive2} has low consistency because of the lack of dependence on X_{id} , making the resulting dataset with random labels. FR performance of 0.5 means the model diverged and is returning random predictions. L_{ID} , a linear interpolation of the L_{naive1} and L_{naive2} across time-steps results in the best performance.

ID Loss Backbone F . ID Loss requires a pretrained FR model, F . For all of our experiments, we use F as IR50 trained on CASIA-WebFace. But, we are curious if there is a benefit to have a better representation from F . For this, we

ID	Style	LFW	CFPP	CPLFW	AGEDB	CALFW	AVG
<i>random</i>	<i>random</i>	98.05	84.17	82.20	89.38	91.40	89.04
<i>random</i>	<i>match</i>	98.28	84.61	82.32	89.12	91.28	89.12
<i>balance</i>	<i>random</i>	98.30	83.27	81.60	89.40	91.27	88.77
<i>balance</i>	<i>match</i>	98.38	84.06	82.45	89.30	91.38	89.11
<i>balance</i>	<i>over smpl</i>	98.55	85.33	82.62	89.70	91.60	89.56

Table 3. Sampling Ablation. We generate a synthetic dataset of 10K subjects with 50 images per subject, using the setting indicated by the blue text in Tab. 2. *over smpl* is over-sampling X_{id} during training for showing more front-view faces.

ablate F_{bigger} , a model pretrained on a larger dataset, WebFace4M [73]. Tab. 2 shows that a better FR backbone induce the generator to synthesize better datasets, even without explicitly showing WebFace4M images to generators. But for fairness in comparing to the real CASIA-WebFace dataset, we do not use F_{bigger} for subsequent analysis.

5.2. Sampling Ablation

Using the sampling strategy defined in Sec. 3.3, we ablate on the ID sampling options (*random*, *balance*) and style sampling methods (*random*, *match*) in Tab. 3. We find that either balancing the gender/ethnicity distribution or making the gender/ethnicity of style image equal to that of ID images does not bring significant performance gain.

On the other hand, to compensate for lower label consistency compared to the real dataset, we include the same X_{id} for 5 additional times for each label. This has the effect of oversampling X_{id} during training FR model. When we add the oversampling option to (*balance*, *match*) setting, we observe an average verification accuracy of 89.56%, 0.52% increase over the (*random*, *random*) setting.

5.3. Comparison with Previous Methods

For training FR models with synthetic datasets, we compare with SynFace [49] and DigiFace [3]. We compare 0.5M and 1.2M image count settings. The first setting corresponds to the size of the CASIA-WebFace real dataset. The second setting is to evaluate the effect of increasing the training dataset size. In Tab. 4, we show the verification accuracies of 5 validation sets. In 0.5M regime, our DCFace can surpass DigiFace in 4 out of 5 datasets with an improvement of 6.11% on average. In CFP-FP dataset with extremely large pose variation, DigiFace performs better, showing the merit of 3D consistent face synthesis using 3D models. DCFace has a good balance of consistency and diversity with many unique subjects, leading to a better FR performance in general. Note the larger style variation compared to SynFace and DigiFace in Fig. 7.

The last column of Tab. 4 shows the gap between synthetic and real, calculated as $(REAL - SYN)/SYN$, e.g. $5.65\% = \frac{94.62 - 89.56}{89.56}$. It indicates how much improvement is needed to be on par with the real dataset. In 0.5M setting, DCFace reduces the gap to real performance by 57% over the SoTA. When we use more synthetic data as in 1.2M

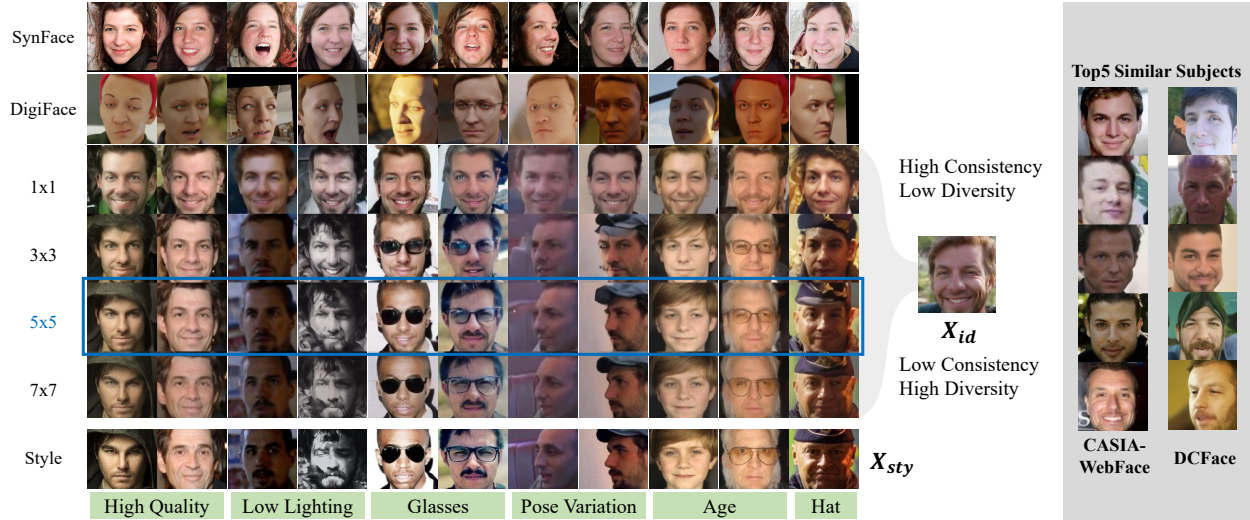


Figure 7. An example of SynFace and DigiFace in rows 1-2 and DCFace with different grid size settings in rows 3-7. SynFace (DiscoFaceGAN) generates mostly frontal-view high-quality images and DigiFace contains synthetic face images with unrealistic texture compared to real images. Our grid size ablation changes the contribution of X_{sty} and X_{id} . A good FR performance is a compromise in-between, 5×5 . Note that our method can have diverse styles such as low lighting, pose, glasses, hat, etc. Using X_{id} to query subjects in CASIA-WebFace and DCFace datasets returns top 5 most similar subjects. We see X_{id} sufficiently different from other (real or fake) subjects.

Methods	Venue	# images (# IDs \times # imgs/ID)	LFW	CFP-FP	CPLFW	AgeDB	CALFW	Avg	Gap to Real
SynFace	ICCV21	0.5M (10K \times 50)	91.93	75.03	70.43	61.63	74.73	74.75	26.58
DigiFace	WACV23	0.5M (10K \times 50)	95.4	87.4	78.87	76.97	78.62	83.45	13.39
DCFace (Ours)	-	0.5M (10K \times 50)	98.55	85.33	82.62	89.70	91.60	89.56	5.65
DigiFace	WACV23	1.2M (10K \times 72 + 100K \times 5)	96.17	89.81	82.23	81.10	82.55	86.37	9.55
DCFace (Ours)	-	1.0M (20K \times 50)	98.83	88.4	84.22	90.45	92.38	90.86	4.14
DCFace (Ours)	-	1.2M (20K \times 50 + 40K \times 5)	98.58	88.61	85.07	90.97	92.82	91.21	3.74
CASIA-WebFace (Real)		0.49M (approx. 10.5K \times 47)	99.42	96.56	89.73	94.08	93.32	94.62	0.0

Table 4. Verification accuracies of FR models trained with SoTA synthetic training datasets. SynFace [49] is a GAN-based dataset with a latent space mixup technique. DigiFace [3] is a 3D model-based dataset with heavy image augmentation. DCFace uses the model setting from the ablation study, Tab. 2, 3 indicated by blue colors. FR backbone is IR-SE50 [11] + AdaFace [31] to match the setting of DigiFace.

regime, the synthetic dataset performance comes closer to that of the real dataset (3.74% in gap), a 60.9% improvement from the previous method (9.55% in gap).

6. Conclusion

This paper presents a method for creating a synthetic training dataset for face recognition. Dataset generation is studied from the perspective of generating many unique subjects with large style diversity and label consistency. We propose the Dual Condition Face Generator to this end and show its large FR performance gain over previous methods on synthetic dataset generation. We believe our approach takes one step towards matching the performance of real training datasets with synthetic training datasets.

Limitations. This work addresses the problem of generating label consistent and diverse datasets for face recognition model training. In our model ablation, we find that sacrificing label consistency for diversity to some degree is beneficial for the FR model training. However, this is not

ideal; for instance, our synthetic face generator lacks 3D consistency across pose, which is an advantage of generative models with 3D priors. Secondly, the goal of our research is to release a synthetic face dataset that alleviates the dependence on large-scale web-crawled images. As shown in our experiments, there is still some performance gap between real and synthetic training datasets. In this work, we take one step towards the goal and hope that the continued research will introduce a standalone synthetic face dataset.

Acknowledgments. This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2022-21102100004. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Gov. is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] Vítor Albiero. Face analysis pytorch. https://github.com/vitoralbiero/face_analysis_pytorch, 2022. 5
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [3] Gwangbin Bae, Martin de La Gorce, Tadas Baltrusaitis, Charlie Hewitt, Dong Chen, Julien Valentin, Roberto Cipolla, and Jingjing Shen. Digiface-1m: 1 million digital face images for face recognition. In *WACV*, 2023. 1, 2, 3, 6, 7, 8
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, 1999. 3
- [5] Andreas Blattmann, Robin Rombach, Kaan Oktay, and Björn Ommer. Retrieval-augmented diffusion models. *arXiv preprint arXiv:2204.11824*, 2022. 3
- [6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 2
- [7] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. VGGFace2: A dataset for recognising faces across pose and age. In *FG*, 2018. 1
- [8] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018. 2, 3
- [9] Jiankang Deng, Shiyang Cheng, Niannan Xue, Yuxiang Zhou, and Stefanos Zafeiriou. UV-GAN: Adversarial facial uv map completion for pose-invariant face recognition. In *CVPR*, 2018. 3
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*. Ieee, 2009. 6
- [11] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 1, 2, 3, 6, 8
- [12] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3D imitative-contrastive learning. In *CVPR*, 2020. 1, 3
- [13] Joshua J Engelsma, Steven A Grosz, and Anil K Jain. Prints-gan: synthetic fingerprint generator. *TPAMI*, 2022. 1, 3
- [14] Baris Gecer, Binod Bhattarai, Josef Kittler, and Tae-Kyun Kim. Semi-supervised adversarial learning to generate photorealistic face images of new identities from 3D morphable model. In *ECCV*, 2018. 3
- [15] Zhenglin Geng, Chen Cao, and Sergey Tulyakov. 3D guided fine-grained face manipulation. In *CVPR*, 2019. 3
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11), 2020. 3
- [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 *ECCV*, 2016. 1, 2
- [18] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30, 2017. 1, 5
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33, 2020. 2, 3, 4, 6
- [20] Qiyang Hu, Attila Szabó, Tiziano Portenier, Paolo Favaro, and Matthias Zwicker. Disentangling factors of variation by mixing them. In *CVPR*, 2018. 3
- [21] Yuan-Ting Hu, Jiahong Wang, Raymond A Yeh, and Alexander G Schwing. Sail-vos 3d: A synthetic dataset and baselines for object detection and 3d mesh reconstruction from video data. In *CVPR*, 2021. 1
- [22] Gary Huang, Marwan Mattar, Honglak Lee, and Erik Learned-Miller. Learning to align from scratch. *NeurIPS*, 25, 2012. 1, 6
- [23] 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
- [24] 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 *CVPR*, 2020. 2
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 4
- [26] Xiaoyi Jiang, Michael Binkert, Bernard Achermann, and Horst Bunke. Towards detection of glasses in facial images. *Pattern Analysis & Applications*, 3(1), 2000. 5
- [27] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018. 2
- [28] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2, 3, 4, 5, 6
- [29] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 2
- [30] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *TOG*, 2018. 3
- [31] Minchul Kim, Anil K Jain, and Xiaoming Liu. AdaFace: Quality adaptive margin for face recognition. In *CVPR*, 2022. 1, 2, 6, 8
- [32] Minchul Kim, Feng Liu, Anil Jain, and Xiaoming Liu. Cluster and aggregate: Face recognition with large probe set. *NeurIPS*, 2022. 4
- [33] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [34] David Kupas and Balazs Harangi. Solving the problem of imbalanced dataset with synthetic image generation for cell classification using deep learning. In *EMBC*, 2021. 1
- [35] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *NeurIPS*, 32, 2019. 6

- [36] HyunJae Lee, Hyo-Eun Kim, and Hyeonseob Nam. Srm: A style-based recalibration module for convolutional neural networks. In *ICCV*, 2019. 4
- [37] Jianxin Lin, Yingce Xia, Tao Qin, Zhibo Chen, and Tie-Yan Liu. Conditional image-to-image translation. In *CVPR*, 2018. 3
- [38] Feng Liu, Minchul Kim, Anil Jain, and Xiaoming Liu. Controllable and guided face synthesis for unconstrained face recognition. In *ECCV*, 2022. 3
- [39] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. SphereFace: Deep hypersphere embedding for face recognition. In *CVPR*, 2017. 2
- [40] Yaojie Liu and Xiaoming Liu. Spoof trace disentanglement for generic face anti-spoofing. *TPAMI*, 45(3), 2023. 3
- [41] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [42] Safa C. Medin, Bernhard Egger, Anoop Cherian, Ye Wang, Joshua B. Tenenbaum, Xiaoming Liu, and Tim K. Marks. MOST-GAN: 3d morphable stylegan for disentangled face image manipulation. In *AAAI*, 2022. 3
- [43] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. AGEDB: the first manually collected, in-the-wild age database. In *CVPRW*, 2017. 2, 6
- [44] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. HoloGAN: Unsupervised learning of 3d representations from natural images. In *ICCV*, 2019. 3
- [45] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, pages 8162–8171. PMLR, 2021. 3
- [46] Jingtian Piao, Chen Qian, and Hongsheng Li. Semi-supervised monocular 3D face reconstruction with end-to-end shape-preserved domain transfer. In *ICCV*, 2019. 3
- [47] Konpat Preechakul, Nattanat Chatthee, Suttisak Widadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *CVPR*, 2022. 5
- [48] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: Anatomically-aware facial animation from a single image. In *ECCV*, 2018. 3
- [49] Haibo Qiu, Baosheng Yu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. SynFace: Face recognition with synthetic data. In *ICCV*, 2021. 1, 3, 6, 7, 8
- [50] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 3, 5
- [51] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 3
- [52] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *NeurIPS*, 29, 2016. 6
- [53] Soumyadip Sengupta, Jun-Cheng Chen, Carlos Castillo, Vishal M Patel, Rama Chellappa, and David W Jacobs. Frontal to profile face verification in the wild. In *WACV*, 2016. 2, 6
- [54] Yujun Shen, Bolei Zhou, Ping Luo, and Xiaoou Tang. Facefeat-GAN: a two-stage approach for identity-preserving face synthesis. *arXiv preprint arXiv:1812.01288*, 2018. 3
- [55] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 2, 3
- [56] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 2, 3, 6
- [57] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *NeurIPS*, 34, 2021. 3
- [58] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *NeurIPS*, 32, 2019. 3
- [59] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *NeurIPS*, 33:12438–12448, 2020. 3
- [60] Joel Stehouwer, Amin Jourabloo, Yaojie Liu, and Xiaoming Liu. Noise modeling, synthesis and classification for generic object anti-spoofing. In *CVPR*, 2020. 3
- [61] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul E Debevec, and Ravi Ramamoorthi. Single image portrait relighting. *TOG*, 2019. 3
- [62] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *CVPR*, 2017. 3
- [63] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *CVPRW*, 2018. 1
- [64] Boris van Breugel, Trent Kyono, Jeroen Berrevoets, and Mihaela van der Schaar. Decaf: Generating fair synthetic data using causally-aware generative networks. *NeurIPS*, 34:22221–22233, 2021. 1
- [65] 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 *CVPR*, 2018. 2
- [66] Tengfei Wang, Ting Zhang, Bo Zhang, Hao Ouyang, Dong Chen, Qifeng Chen, and Fang Wen. Pretraining is all you need for image-to-image translation. *arXiv preprint arXiv:2205.12952*, 2022. 3, 5
- [67] Tianxing Wu. Realtime glasses detection. <https://github.com/TianxingWu/realtime-glasses-detection>, 2022. 5
- [68] Andre Brasil Vieira Wzykowski and Anil K Jain. Synthetic latent fingerprint generator. In *WACV*, 2023. 3
- [69] Taihong Xiao, Jiapeng Hong, and Jinwen Ma. Elegant: Exchanging latent encodings with GAN for transferring multiple face attributes. In *ECCV*, 2018. 3
- [70] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014. 3

- [71] Tianyue Zheng and Weihong Deng. Cross-Pose LFW: A database for studying cross-pose face recognition in unconstrained environments. *Beijing University of Posts and Telecommunications, Tech. Rep*, 5, 2018. [2](#), [6](#)
- [72] Tianyue Zheng, Weihong Deng, and Jiani Hu. Cross-Age LFW: A database for studying cross-age face recognition in unconstrained environments. *CoRR*, abs/1708.08197, 2017. [2](#), [6](#)
- [73] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Jiwen Lu, Dalong Du, et al. WebFace260M: A benchmark unveiling the power of million-scale deep face recognition. In *CVPR*, 2021. [1](#), [2](#), [3](#), [6](#), [7](#)
- [74] Hasib Zunair and A Ben Hamza. Synthesis of covid-19 chest x-rays using unpaired image-to-image translation. *Social network analysis and mining*, 11(1), 2021. [1](#)