







Arc2Face: A Foundation Model for ID-Consistent Human Faces

Foivos Paraperas Papantoniou¹, Alexandros Lattas¹, Stylianos Moschoglou¹, Jiankang Deng¹, Bernhard Kainz^{1,2}, and Stefanos Zafeiriou¹

¹ Imperial College London, UK

² FAU Erlangen–Nürnberg, Germany

<https://arc2face.github.io/>

Abstract. This paper presents Arc2Face, an identity-conditioned face foundation model, which, given the ArcFace embedding of a person, can generate diverse photo-realistic images with an unparalleled degree of face similarity than existing models. Despite previous attempts to decode face recognition features into detailed images, we find that common high-resolution datasets (*e.g.* FFHQ) lack sufficient identities to reconstruct any subject. To that end, we meticulously upsample a significant portion of the WebFace42M database, the largest public dataset for face recognition (FR). Arc2Face builds upon a pretrained Stable Diffusion model, yet adapts it to the task of ID-to-face generation, conditioned solely on ID vectors. Deviating from recent works that combine ID with text embeddings for zero-shot personalization of text-to-image models, we emphasize on the compactness of FR features, which can fully capture the essence of the human face, as opposed to hand-crafted prompts. Crucially, text-augmented models struggle to decouple identity and text, usually necessitating some description of the given face to achieve satisfactory similarity. Arc2Face, however, only needs the discriminative features of ArcFace to guide the generation, offering a robust prior for a plethora of tasks where ID consistency is of paramount importance. As an example, we train a FR model on synthetic images from our model and achieve superior performance to existing synthetic datasets.

Keywords: Face synthesis · ID-embeddings · Subject-driven generation

1 Introduction

Learning statistical priors from data to generate new facial images is a highly popular problem at the intersection of Computer Vision and Machine Learning. Arguably, the most well-known early method was the so-called Eigenfaces [78], which applies the Karhunen-Loeve expansion, also known as Principal Component Analysis (PCA), to a set of facial images. The success of Eigenfaces with well-aligned frontal facial images captured under controlled conditions spurred

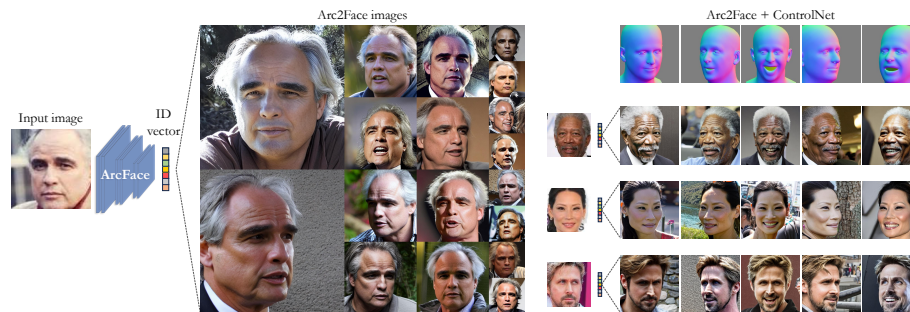


Fig. 1: Given the ID-embedding from [14], Arc2Face can generate high-quality images of any subject with compelling similarity. Using popular extensions, such as ControlNet [94], we can explicitly control facial attributes such as the pose or expression.

a significant line of research into linear and nonlinear component analysis over 25 years ago. Probably the pinnacle of PCA applications was in 3D Morphable Models (3DMMs) [6], where it was used to learn a prior for facial textures represented in a UV map. However, when directly applied to 2D images, PCA had many shortcomings, including the requirement for images to be perfectly pixel-aligned. Furthermore, it could not describe hair (facial or not) and also struggled with non-linearities introduced by facial expressions and lighting conditions.

The recent evolution of deep neural networks (DNNs) has made it feasible to overcome many of the limitations inherent in simple linear statistical models, such as PCA, enabling the learning of statistical priors from large-scale facial data captured under varied recording conditions. Arguably, the most successful method, based on Generative Adversarial Networks (GANs) [25], is StyleGAN [38] and its variants [37, 39]. StyleGAN learns a decoder governed by low-dimensional latent codes, which can be sampled to generate images. Moreover, StyleGAN can be used to fit images to its distributions (*i.e.*, find the best latent code to generate a specific image) [84], proving highly useful as a prior for many tasks, including 3D face reconstruction for facial texture generation [24]. Additionally, several studies have discussed how StyleGAN’s latent spaces can be utilized for controllable manipulations (*i.e.*, changing the expressions or the pose of facial images) [27, 73, 100]. However, to date, identity drift remains a challenge in latent space manipulations using StyleGAN.

One of the first problems to find a robust solution with the advent of deep learning is face recognition (FR), generally approached by identifying features that represent facial identities (also known as ID-embeddings). These features are so powerful that they are used in smartphones for face verification, leading to a vast research field. Arguably, the most widely used ID-embeddings are those produced by the ArcFace [14] loss. ArcFace features have been utilized as perceptual features for 3D face reconstruction and generation. However, robustly controlling the identity within the StyleGAN framework remains challenging,

mainly due to the necessity of large high-resolution datasets with significant intra-class variability.

In the past couple of years, another revolution concerning generative methods has occurred in deep learning. The emergence of the so-called Diffusion Models [74] has demonstrated not only the possibility of modeling the distribution of images and sampling from it [17, 29, 76], but also the feasibility of guiding the generation process with features that correlate images with textual descriptions, such as CLIP [59, 62]. Recently, it has been shown that other features, *e.g.*, ID-embeddings, can be incorporated to steer the generation process for faces, to not only adhere to textual descriptions but also to the facial characteristics of a person’s image [11, 56, 79, 82, 89].

A challenge arises when combining CLIP features with ID-embeddings because CLIP features contain ID-related information (otherwise, generating the image of a famous celebrity from a textual prompt would not be possible), as well as possess contradictory features to ID-embeddings (*e.g.*, when requesting a photo of a “Viking” with the ID-embeddings of a Chinese person). Hence, although these models are useful for entertainment purposes and demonstrate the power of diffusion models, they are neither suitable for the controlled generation of facial images nor for sampling conditioned to the ID features of a specific subject.

In this paper, we meticulously study the problem of high-resolution facial image synthesis conditioned on ID-embeddings and propose a large-scale foundation model. Developing such a model poses a significant challenge due to the limited availability of high-quality facial image databases. Specifically:

- We show that smaller, single-image-per-person databases [38] are insufficient to train a robust foundation model, thus, we introduce a large dataset of high-resolution facial images with consistent identity and intra-class variability derived from WebFace42M.
- We adapt a large facial identity encoder trained *solely* on ID-embeddings, in contrast to text-based models where ID interferes with language.
- We introduce the first ID-conditioned face foundation model, which we make available to the public.
- We propose a robust benchmark for evaluating ID-conditioned models. Our experiments demonstrate superior performance compared to existing approaches.

2 Related Work

2.1 Generative Models for the Human Face

The task of facial image generation has witnessed great success in recent years, with the advent of style-based GANs [37–39, 51], even for view-consistent synthesis [3, 8, 9, 16, 26, 53, 70]. Despite early attempts to condition GANs on modalities such as text [33, 34, 67], the recent Diffusion Models [17, 29, 76, 87] have shown tremendous progress. Trained on large-scale datasets, several prominent models,

including DALL-E-2 [60], Imagen [66], or Stable Diffusion [62], have emerged, marking a breakthrough in creative image generation. Later models have enabled an even higher level of detail [57]. Nevertheless, dataset memorization remains an issue in such models [75]. Moreover, recent methods have achieved even 3D human generation using general [42,55] and text-guided diffusion models [32,93].

In terms of controlling the generation output, adding images alongside text has been proposed. For example, ILVR [12] conditions sampling on a target image through iterative refinement, while SDEdit [50] adds noise to the input before editing it according to the target description. DiffusionRig [18] enables explicit face editing, using a crude 3DMM [20] rendering to condition generation on pose, illumination, and expression. Perhaps the most notable approach, ControlNet [94], first achieved accurate spatial control of text-to-image models, with the addition of trainable copies of network parts. In a different direction, universal guidance was introduced [5] to avoid re-training the model during conditioning. While the combined use of text and spatial control can enable manipulation of an input photo to some extent, it cannot accurately generate one’s identity under any conditions, which is addressed by subject-driven generative models.

2.2 Subject-Specific Facial Generation

When considering subject-conditioned generation, face recognition (FR) models provide a powerful platform for identity features extraction from facial images, since they typically extract comprehensive facial embeddings in order to measure identity similarity (*e.g.* [7, 14, 15, 40]). Therefore, the inversion of such models in a black-box setting has been shown capable of producing facial images from an identity embedding [49, 61, 80, 88], even when using GAN and diffusion architectures for zero-shot generation [19, 35, 77]. However, current inversion methods are either trained on low-resolution datasets [19, 77] designed for face recognition or use high-quality but limited images [35], constraining their generalizability.

The most impressive results have recently been shown as extensions of Stable Diffusion [62]. The main paradigm emerged from Textual Inversion [22] and especially DreamBooth [64], which fine-tunes a diffusion model on several images of a subject, to learn a subject-specific class identifier reproducing that specific subject. Follow-up works reduced the optimization time, such as HyperDreamBooth [65], which utilizes LoRA [30] and a hypernetwork to perform tuning on a single image. Similarly, E4T [23] and ProFusion [99] propose encoder-based approaches, while CustomDiffusion [43] optimizes only a subset of network’s parameters. Moreover, Celeb-Basis [92] and StableIdentity [81] learn an embedding basis from a set of celebrities that can be used to condition the text-based model. In a generalized manner, Kosmos-G [54] introduced a multi-modal perception model that accepts various inputs, including facial images.

Closest to ours lies a series of recent works that condition Diffusion Models directly on facial features for tuning-free personalization. Various approaches, including FastComposer [85], PhotoVerse [10] and PhotoMaker [45], employ features from the CLIP [59] image encoder to represent the input subject. However,




Fig. 2: Overview of Arc2Face. We use a straightforward design to condition Stable Diffusion on ID features. The ArcFace embedding is processed by the text encoder using a frozen pseudo-prompt for compatibility, allowing projection into the CLIP latent space for cross-attention control. Both the encoder and UNet are optimized on a million-scale FR dataset [101] (after upsampling), followed by additional fine-tuning on high-quality datasets [36, 38], without any text annotations. **The resulting model exclusively adheres to ID-embeddings, disregarding its initial language guidance.**

this representation is constrained by the CLIP’s facial encoding abilities. Methods such as Face0 [79], DreamIdentity [11], and PortraitBooth [56] condition the model additionally on FR embeddings for improved fidelity. Similarly, IP-Adapter [89] uses a decoupled cross-attention mechanism to separate text from subject conditioning and has released, post-publication, [90] an impressive model based on ID features. InstantID [82] extends [90] with an additional network for stronger ID guidance and facial landmarks conditioning. Finally, FaceStudio [86] learns combined CLIP and ID-embeddings, achieving impressive stylizations.

A common concept in prior works is to use subject embeddings as an additional conditioning mechanism on top of text. Therefore, they exhibit a deficiency in identity retention due to the ambiguous relation of text description and fine identity characteristics. In contrast, in this work, we introduce a model solely conditioned on robust identity features, achieving state-of-the-art control and identity retention, as shown in Fig. 1.

3 Method

Our objective is to develop a foundation model that accurately generates images of any subject independent of pose, expression, or contextual scene information. To that end, we employ the ArcFace [14] network as our feature extractor, owing to its intrinsic ability to filter out such information. Our model extends a powerful Stable Diffusion [62] backbone, enabling efficient sampling of high-quality images. Specific details of our methodology are provided in the subsequent sections.

3.1 Preliminaries

Face Recognition Features: The utilization of pre-trained FR networks to constrain face-related optimization or generation tasks has seen significant success in recent years. Typically, ID similarity is employed as an auxiliary loss,

often by extracting features from multiple layers of these networks. In this work, we aim to utilize these networks as frozen feature extractors in order to condition a generative model of facial images. In particular, let ϕ denote the forward function of a pre-trained ArcFace [14] network. Given a cropped and aligned face image $x \in \mathbb{R}^{H \times W \times C}$, ϕ extracts a high-level embedding $\mathbf{w} = \phi(x) \in \mathbb{R}^{512}$, designed to separate the face from other subjects. We rely only on this vector to recover novel images of the subject, without access to any intermediate layers of [14]. As shown by Kansy *et al.* [35], it is possible to learn this mapping via a conditional diffusion model without any ID-specific loss functions. Similarly, we employ \mathbf{w} to condition our model by projecting it onto the cross-attention layers of the generator, following the successful paradigm of Stable Diffusion.

Latent Diffusion Models: Diffusion models [17, 29, 76] employ a denoising mechanism to approximate the distribution of real images x . During training, images undergo distortion through the addition of Gaussian noise via a predetermined diffusion schedule at different timesteps t , given by $x_t = \sqrt{\alpha_t}x_0 + (1 - \alpha_t)\epsilon$. At the same time, a denoising autoencoder $\epsilon_\theta(x_t, t)$ is trained to recover the normally distributed noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ by minimizing the prediction error :

$$\mathcal{L} = \mathbb{E}_{x_t, t, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\epsilon - \epsilon_\theta(x_t, t)\|_2^2] \quad (1)$$

The Latent Diffusion (LD) model [62] is a widely adopted architecture within the diffusion framework. Notably, a Variational Autoencoder (VAE), \mathcal{E} , is employed to compress images into a lower-dimensional latent space $z = \mathcal{E}(x)$ for efficient training. Moreover, LD introduces a universal conditioning mechanism, using a UNet [63] as the backbone for ϵ_θ , with cross-attention layers to map auxiliary features C to its intermediate layers for conditional noise prediction $\epsilon_\theta(z_t, t, C)$. Particularly, Stable Diffusion (SD) represents a text-conditioned LD model, which uses text embeddings C , produced by a CLIP encoder [59], to enable stochastic text-to-image (T2I) generation. Trained on the LAION-5B database [68], SD stands as a foundation open-source T2I model which is commonly utilized by the research community as a powerful image prior for downstream tasks.

3.2 Base Models

To demonstrate the effectiveness of ID-embeddings in face reconstruction and highlight the necessity of extensive datasets, we conducted initial experiments training an ID-conditioned model from scratch on low-resolution images with varying data sizes. This initial model follows the conditional LD format [62], *i.e.* a perceptual autoencoder and a UNet denoiser equipped with cross-attention layers, and is trained on image-embedding pairs using the standard loss of Eq. 1.

Fig. 3 compares the performance of the model when trained on three different datasets: FFHQ [38] (70K single-person images, downsampled to 256×256), WebFace42M-10% (a subset with ~ 4 M images from ~ 200 K identities), and the complete WebFace42M [101] dataset. For each model, we generate images based on 400 real and 500 synthetic input faces (see Sec. 4.1 for the images used in

our experiments) and measure the ID similarity between the input and generations. Results reveal that a model trained on FFHQ [38] exhibits limited ID retention due to the relatively modest number of images and IDs. In contrast, WebFace42M [101], despite comprising low-resolution face images (112×112 pixels) cropped around the facial region for FR training, proves particularly suitable for our task, even when reduced to 10% size. Notably, it constitutes the largest public FR dataset, consisting of approximately 42M images, and, in contrast to other facial datasets (*e.g.* [38, 69, 98]), it contains significant and consistent intra-class variation, which is crucial for diverse generations of the same ID.

Expanding upon this finding, we then develop Arc2Face, which enables the synthesis of photo-realistic images at higher resolutions. This model is derived by fine-tuning SD on a restored version of WebFace42M [101], along with FFHQ [38] and CelebA-HQ [36] for enhanced quality, as detailed below.




Fig. 3: ArcFace [14] similarity distributions between input and generated faces from LD models trained on the ID-to-image task. We use two different datasets of input IDs for evaluation (500 and 400 IDs respectively) and generate 5 images per ID. We compare models trained on three datasets: FFHQ, WebFace42M-10%, and WebFace42M.

3.3 Arc2Face

Motivated by the ID retention offered by large-scale databases, we structure our model in two ways: 1) we employ a pre-trained SD [62] as our prior, 2) we automatically generate a high-resolution dataset from [101], enabling efficient training of our backbone without compromising its superior quality priors.

ID-Conditioning: Our model is implemented with `stable-diffusion-v1-5`, which uses a CLIP text encoder [59] to guide image synthesis. Our goal is to condition it on ArcFace embeddings, while directly harnessing the generative power of its UNet. Thus, it is necessary to project ArcFace embeddings to the space of CLIP embeddings, used by the original model. We achieve this by feeding them into the same encoder and fine-tuning it to swiftly adapt itself to the ArcFace input. This approach offers a more seamless projection than replacing

CLIP with an MLP as in recent works [45, 81, 86] (see Sec. 4.4 for an ablation comparison). To ensure compatibility with CLIP, we employ a simple prompt, “photo of a <id> person”. Following tokenization, we substitute the placeholder <id> token embedding with the ArcFace vector \mathbf{w} , yielding a sequence of token embeddings $s = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \hat{\mathbf{w}}, \mathbf{e}_5\}$. Here, $\hat{\mathbf{w}} \in \mathbb{R}^{768}$ corresponds to $\mathbf{w} \in \mathbb{R}^{512}$ after zero-padding to match the dimension of $\mathbf{e}_i \in \mathbb{R}^{768}$. The resulting sequence is fed to the encoder τ , which maps it to the CLIP output space $C = \tau(s) \in \mathbb{R}^{N \times 768}$ (with N denoting the tokenizer’s maximum sentence length). This process is illustrated in Fig. 2. Note that as a byproduct of this operation, the ID information is shared across multiple embeddings in the output of τ , offering more detailed guidance to the UNet. This concept is also used in recent works [11, 79]. During training, we consistently employ this default pseudo-prompt for all images. This intentional choice directs the encoder’s attention solely to the ID vector, disregarding any irrelevant contextual information. Consequently, through extensive fine-tuning, we effectively transform the text encoder into a face encoder specifically tailored for projecting ArcFace embeddings into the CLIP latent space.

Dataset: We retain WebFace42M [101] for its vast size and intra-class variability, which, however, suffers from low-resolution data and a tight facial crop. Even more, the pre-trained SD backbone is designed for a resolution of 512×512 and a similar resolution is required to fine-tune Arc2Face. To alleviate this, we meticulously upsample them using GFPGAN (v1.4), a state-of-the-art blind face restoration network [83], denoted as γ . We perform degradation removal and $4 \times$ upscaling to 448×448 , so that $\hat{x} = \gamma(x)$. GFPGAN is a GAN-based upsampling method, with a strong face prior using both adversarial and ID-preserving losses, achieving crisp and faithful restoration. We follow this process on a large portion of the original database, given our computational limits, acquiring approximately 21M images of 1M identities at 448×448 pixels. Using the restored images \hat{x} , we train Arc2Face. Albeit of higher quality, this dataset is still limited to a tightly cropped facial area for FR training. While it allows to learn a robust ID prior, a complete face image is usually preferred. Thus, we further fine-tune it on FFHQ [38] and CelebA-HQ [36], which consist of less constrained face images. Our final model generates FFHQ-aligned images at 512×512 pixels.

4 Experiments

4.1 ID-Consistent Generation

We perform extensive quantitative comparisons to evaluate the performance of recent ID-conditioned models in generating both diverse and faithful images of a subject. Details about the methods and comparison process are provided below.

Methods We compare against recent zero-shot methods that condition synthesis on identity information. Typically, these methods employ either CLIP image features or FR features to achieve tuning-free customization of SD on a given




Fig. 4: Distribution of ArcFace similarity between input IDs, synthetic (a) or real (b), and generated images of them by different models. As all non-CLIP-based methods use [14] for conditioning, we evaluate them with [14]. For an evaluation with a different network, please refer to the Supp. Material, where similar observations can be made.




Fig. 5: Percentage of user votes received by our method and InstantID re. ID fidelity.

face. In particular, we compare against the following open-source methods:

- 1) *FastComposer* [85], which combines text with features extracted from the CLIP image encoder. It further uses a localized attention mechanism and proposes delayed subject conditioning during sampling for improved text editability. It is trained on FFHQ-wild [38], automatically annotated with text prompts.
- 2) *PhotoMaker* [45], which extracts CLIP features from one or a few images of a subject and combines them with text to condition the model. It is trained on a custom dataset with 112K images from 13K celebrities collected from the web.
- 3) *IP-Adapter-FaceID (IPA-FaceID)* [89, 90], which uses a decoupled attention mechanism for ID features in addition to text. Subsequent versions (*IPA-FaceID-Plus* and *Plusv2*) also use a combination of ID with CLIP image embeddings.
- 4) *InstantID* [82], which extends *IPA-FaceID* with an IdentityNet, akin to ControlNet, conditioned on FR embeddings and sparse landmarks. It is trained on LAION-Face [98], plus 10M automatically annotated images from the web.

Prompt Selection Since the aforementioned methods are tailored for text-driven synthesis, their ability to generate consistent IDs depends heavily on the input prompt. Although this allows creative stylizations, it demands prompt engineering and user inspection. To eliminate subjective bias, we suggest an automatic evaluation using a simple prompt, “photo of a person”, for all samples. This enables us to assess ID-retention performance based on face features without elaborate text descriptions, ensuring a fair comparison across all methods.

Datasets and Metrics We construct two datasets for our comparison. The first, hereafter denoted as Synth-500, includes 500 images of never-before-seen identities, generated by [1]. The second is a selection of 400 real images from the public AgeDB [52] database, chosen based on higher resolution. For each image, we generate 5 samples using all methods and calculate the following metrics: 1) ID similarity: we compute the cosine similarity of ID features between the input and generated faces. The corresponding distributions are plotted in Fig. 4. 2)

LPIPS distance: similarly to [45], we assess the diversity of generation by calculating the pairwise perceptual distance (LPIPS) [95] between images of the same ID. The average distance for all pairs is reported in Tab 1. To focus on the actual facial diversity, we first detect and align all faces and remove the background using an off-the-shelf parser [2]. 3) Exp./Pose diversity: we predict the expression and pose (jaw/neck articulation) parameters of the FLAME model [44] for each sample using EMOCA v2 [13, 21] and compute the average pairwise ℓ_2 expression and pose distances between images of the same ID. 4) FID: we use the FID [28, 71] metric to assess the quality of samples (after face detection and alignment) with respect to the input images. Results are presented in Tab 1.

	LPIPS \uparrow		Exp. (ℓ_2) \uparrow		Pose (ℓ_2) \uparrow		FID \downarrow	
	Synth-500	AgeDB	Synth-500	AgeDB	Synth-500	AgeDB	Synth-500	AgeDB
FastComposer	0.389	0.487	3.597	4.678	0.163	0.225	13.517	31.736
Photomaker	0.410	0.424	3.920	4.283	0.167	0.165	13.295	8.410
InstantID	0.386	0.437	3.733	4.569	0.059	0.082	22.859	18.598
IPA-FaceID (SDXL)	0.402	0.462	4.648	5.812	0.181	0.197	7.104	24.105
IPA-FaceID-Plus	0.320	0.384	2.706	3.518	0.150	0.194	14.880	11.817
IPA-FaceID-Plusv2	0.356	0.429	3.147	4.092	0.185	0.236	9.752	10.798
Arc2Face (Ours)	0.506	0.508	6.375	5.966	0.317	0.273	5.673	6.628

Table 1: Quantitative comparison between Arc2Face and [45, 82, 85, 89] on 500 synthetic and 400 real IDs. We produce 5 samples per ID for all methods and assess the diversity of generated faces using perceptual and 3DMM-based distances, as well as their quality based on FID. Bold values denote the best results in each metric.

As depicted in Fig. 4, existing methods, particularly [45, 85] that rely on CLIP features, struggle to preserve the ID without detailed text descriptions of the subject. InstantID [82] attains the highest similarity, behind Arc2Face. However, this results in limited pose/exp. diversity as it further uses the input facial landmarks to constrain generation. Arc2Face does not require any text or spatial conditions and achieves the highest facial similarity as well as diversity and realism across both datasets. This particularly highlights the efficiency of ID-embeddings in the context of face generation against CLIP image or text features. A visual comparison with the above methods is provided in Fig. 6, whereas additional qualitative results and visualizations are provided in the Supp. Material. We further conducted a user study to compare with the second best-performing method in terms of ID similarity, *i.e.* InstantID [82]. In particular, we asked 50 users to choose the method whose result best resembles the input face, regardless of quality or realism, for a randomly selected set of 30 IDs from our datasets. The two methods were randomly presented side-by-side. Fig. 5 confirms a strong preference for Arc2Face in terms of ID resemblance.




Fig. 6: Visual comparison of Arc2Face with state-of-the-art methods [45, 82, 85, 89] using the abstract prompt “photo of a person” to focus on their ID-conditioning ability.

4.2 Face Recognition with Synthetic Data

To further demonstrate the potential of our model, we use it to create synthetic images for FR training. In particular, we sample novel ID vectors from the distribution of ArcFace embeddings learned by PCA on WebFace42M [101], while ensuring sufficient intra-class variance for the vectors of the same ID. To guarantee the uniqueness of generated subjects, we keep only those with ID similarity between them below 0.3. We compare against recent methods that train FR models on synthetic data, namely SynFace [58], DigiFace [4], and DC-Face [41], following their training scheme with IR-SE-50 [14] as a backbone and the AdaFace [40] loss function. The trained models are evaluated on LFW [31], CFP-FP [72], CPLFW [96], AgeDB [52] and CALFW [97] datasets, ensuring large pose [72, 96] and age [52, 97] variation. In Tab. 2, we present results for 0.5M and 1.2M settings, corresponding to the size of CASIA-WebFace [91] and an increased size, respectively. Arc2Face surpasses DCFace in all five test datasets with an average improvement of 2.17% in the 0.5M regime and 1.93% in the 1.2M regime. Especially on the CFP-FP dataset, Arc2Face significantly outperforms DigiFace and DCFace, showing high ID-consistency under large pose variation.

4.3 Pose/Expression Control

Our model can be trivially combined with ControlNet [94] for spatial control of the output. In particular, we use EMOCA v2 [13, 21] to perform 3D reconstruc-

Methods	Venue	# images (# IDs× # imgs/ID)	LFW	CFP-FP	CPLFW	AgeDB	CALFW	Avg
SynFace	ICCV21	0.5M (10K × 50)	91.93	75.03	70.43	61.63	74.73	74.75
DigiFace	WACV23	0.5M (10K × 50)	95.4	87.4	78.87	76.97	78.62	83.45
DCFace	CVPR23	0.5M (10K × 50)	98.55	85.33	82.62	89.70	91.60	89.56
Arc2Face	-	0.5M (10K × 50)	98.81	91.87	85.16	90.18	92.63	91.73
DigiFace	WACV23	1.2M (10K × 72 + 100K × 5)	96.17	89.81	82.23	81.10	82.55	86.37
DCFace	CVPR23	1.2M (20K × 50 + 40K × 5)	98.58	88.61	85.07	90.97	92.82	91.21
Arc2Face	-	1.2M (20K × 50 + 40K × 5)	98.92	94.58	86.45	92.45	93.33	93.14
CASIA-WebFace (Real)		0.49M (approx. 10.5K × 47)	99.42	96.56	89.73	94.08	93.32	94.62

Table 2: Verification accuracies of FR models trained with synthetic datasets. SynFace [58] is a GAN-based dataset with a latent space mixup technique. DigiFace [4] is a 3D model-based dataset with image augmentation. DCFace [41] is a diffusion-based dataset with separate ID and style as dual conditions.

tion on FFHQ [38], and train a ControlNet module, conditioned on the rendered face normals. During inference, we can render the 3D face normals of a source person under the expression and pose extracted from a reference image. This rendering is used to guide the synthesis of the source identity as shown in Fig. 7.




Fig. 7: Samples from Arc2Face, conditioned on a 3DMM [13, 44] using ControlNet [94].

4.4 Ablation Studies

ID-Conditioning via MLP To achieve ID-conditioning in SD, we map ArcFace [14] features to the CLIP embedding space by adapting the pre-trained text encoder [59] through extensive fine-tuning. Recent approaches combining text and ID-conditioning typically employ a simple MLP for subject embedding projection [11, 45, 56, 79, 81, 85, 86]. To validate our choice, we conducted an ablation study, replacing the CLIP encoder with a 4-layer MLP while maintaining the same training setting. Fig. 8 shows that our MLP-based model yields lower face similarity in generated images. Since SD is trained with CLIP embeddings, using the same encoder is a more natural choice than training an MLP to learn the CLIP latent space. Additionally, the former represents a well-established ar-

chitecture, while the latter would necessitate a more extensive hyperparameter search, rather than a simple selection of a shallow MLP, as seen in most works. **Principal Component Analysis** To explore the intrinsic dimensionality of our facial ID representation, we conducted PCA on the ArcFace embeddings derived from WebFace42M images. Fig. 9a depicts the cumulative percentage of variance explained by the principal components. Notably, maintaining facial fidelity requires at least 300-400 components, as fewer components result in noticeable distortion, as shown in Fig. 9b. This reveals the challenge of compressing ArcFace embeddings significantly, emphasizing the inherent complexity of our facial ID representation.




Fig. 8: ID similarity [14] distributions between input and generated images from our model. We compare the proposed approach with that of using an MLP for the projection of ID vectors to the CLIP latent space.




Fig. 9: PCA on ID-embeddings: We show the cumulative percentage of variance across components (a) and samples from Arc2Face when projecting ID vectors to a varying number of components (b).

Generalization Ability By training on an extensive dataset with more than 1M IDs and significant intra-class variance, our model is capable of reproducing photos of any person with high fidelity. To further assess its capacity, we examine whether or not Arc2Face has learned to replicate its training IDs by exploring the output images for the 500 previously unseen identities from Synth-500. In particular, for each of the 2.5K samples (5 per ID), we identified the closest training image - defined as the one exhibiting the highest ID similarity within the combined set of restored WebFace42M, FFHQ, and CelebA-HQ. Fig. 10b shows some examples of generated images and their closest matches from the training data. We further plot the distribution of similarities in Fig. 10a. Notably, the generated images exhibit an average cosine similarity of 0.37 to their closest training sample and 0.74 to their (unseen) input features. This observation indicates that our model indeed does not memorize images from the training set, confirming its ability to faithfully reconstruct faces of novel identities.

4.5 Implementation Details

We start from `stable-diffusion-v1-5` and fine-tune both the UNet and CLIP encoder with AdamW [46] and a learning rate of $1e-6$, using 8 NVIDIA A100 GPUs and a batch size of 4 per GPU. First, we train on 21M restored images




Fig. 10: We show that Arc2Face does not replicate its training data. We generate images for 500 unseen IDs from Synth-500 and retrieve the train image with the highest ID similarity for each one. The average similarity between each output and its closest match is 0.37, whereas the similarity between the output and input features is 0.74.

from WebFace42M for 5 epochs with a resolution of 448×448 , and then fine-tune on 512×512 - sized FFHQ and CelebA-HQ images for another 15 epochs. All our results shown in this paper are generated using DPM-Solver [47, 48] with 25 inference steps and a classifier-free guidance scale of 3, which we empirically found to produce highly realistic images. For ID-embedding extraction, we use a frozen IR-100 ArcFace [14] trained on WebFace42M and normalize embeddings to unit magnitude. Regarding the methods we compare against, we use the official implementations and, in all cases, select the default hyper-parameters.

5 Limitations

In this work, we introduce a model capable of generating high-quality facial images from facial embeddings. Our method is limited by the capacity of the face embeddings encoder [14], as well as the datasets used [101], both of which are fortunately state-of-the-art. Moreover, only one person per image can be generated. Despite these limitations, we provide a foundation model, that can be further fine-tuned to other datasets and modalities, aiding follow-up research.

6 Conclusion

In this work, we explore the power of ID features derived from FR networks as a comprehensive representation for face generation in the context of large-scale diffusion models. We show that million-scale face recognition datasets are required to effectively train an ID-conditioned model. To that end, we fine-tune the pre-trained SD on carefully restored images from WebFace42M. Our ID-conditioning mechanism transforms the model into an ArcFace-to-Image model, deliberately disregarding text information in the process. Our experiments demonstrate its ability to faithfully reproduce the facial ID of any individual, generating highly realistic images with a greater degree of similarity compared to any existing method, all while preserving diversity in the output. We hope our model encourages further research in ID-preserving generative AI for human faces.

Acknowledgments. S. Zafeiriou and part of the research was funded by the EPSRC Fellowship DEFORM (EP/S010203/1) and EPSRC Project GNOMON (EP/X011364/1). B. Kainz received support from the ERC, project MIA-NORMAL 101083647, DFG 512819079, and by the State of Bavaria (HTA). HPC resources were provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) under the NHR project b180dc. NHR@FAU hardware is partially funded by the German Research Foundation (DFG) – 440719683.

References

1. <https://thispersondoesnotexist.com/>
2. <https://github.com/zllrunning/face-parsing.PyTorch>
3. An, S., Xu, H., Shi, Y., Song, G., Ogras, U.Y., Luo, L.: Panohead: Geometry-aware 3d full-head synthesis in 360deg. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20950–20959 (2023)
4. Bae, G., de La Gorce, M., Baltrusaitis, T., Hewitt, C., Chen, D., Valentin, J., Cipolla, R., Shen, J.: Digiface-1m: 1 million digital face images for face recognition. In: WACV (2023)
5. Bansal, A., Chu, H.M., Schwarzschild, A., Sengupta, S., Goldblum, M., Geiping, J., Goldstein, T.: Universal guidance for diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 843–852 (June 2023)
6. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques. p. 187–194. SIGGRAPH ’99, ACM Press/Addison-Wesley Publishing Co., USA (1999). <https://doi.org/10.1145/311535.311556>, <https://doi.org/10.1145/311535.311556>
7. Boutros, F., Damer, N., Kirchbuchner, F., Kuijper, A.: Elasticface: Elastic margin loss for deep face recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1578–1587 (2022)
8. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16123–16133 (2022)
9. Chan, E.R., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5799–5809 (2021)
10. Chen, L., Zhao, M., Liu, Y., Ding, M., Song, Y., Wang, S., Wang, X., Yang, H., Liu, J., Du, K., et al.: Photoverse: Tuning-free image customization with text-to-image diffusion models (2023)
11. Chen, Z., Fang, S., Liu, W., He, Q., Huang, M., Zhang, Y., Mao, Z.: Dreamidentity: Improved editability for efficient face-identity preserved image generation. arXiv preprint arXiv:2307.00300 (2023)
12. Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: Ilvr: Conditioning method for denoising diffusion probabilistic models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14367–14376 (2021)

13. Daněček, R., Black, M.J., Bolkart, T.: Emoca: Emotion driven monocular face capture and animation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20311–20322 (2022)
14. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4690–4699 (2019)
15. Deng, J., Guo, J., Yang, J., Lattas, A., Zafeiriou, S.: Variational prototype learning for deep face recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11906–11915 (2021)
16. Deng, Y., Yang, J., Xiang, J., Tong, X.: Gram: Generative radiance manifolds for 3d-aware image generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10673–10683 (2022)
17. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
18. Ding, Z., Zhang, X., Xia, Z., Jebe, L., Tu, Z., Zhang, X.: Diffusionrig: Learning personalized priors for facial appearance editing. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12736–12746 (2023)
19. Duong, C.N., Truong, T.D., Luu, K., Quach, K.G., Bui, H., Roy, K.: Vec2face: Unveil human faces from their blackbox features in face recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6132–6141 (2020)
20. Feng, Y., Feng, H., Black, M.J., Bolkart, T.: Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (ToG)* **40**(4), 1–13 (2021)
21. Filntisis, P.P., Retsinas, G., Paraperas-Papantoniou, F., Katsamanis, A., Roussos, A., Maragos, P.: Visual speech-aware perceptual 3d facial expression reconstruction from videos. *arXiv preprint arXiv:2207.11094* (2022)
22. Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-Or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618* (2022)
23. Gal, R., Arar, M., Atzmon, Y., Bermano, A.H., Chechik, G., Cohen-Or, D.: Encoder-based domain tuning for fast personalization of text-to-image models. *ACM Transactions on Graphics (TOG)* **42**(4), 1–13 (2023)
24. Gecer, B., Ploumpis, S., Kotsia, I., Zafeiriou, S.: Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
25. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* **27** (2014)
26. Gu, J., Liu, L., Wang, P., Theobalt, C.: Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In: *International Conference on Learning Representations* (2022)
27. Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. *Advances in neural information processing systems* **33**, 9841–9850 (2020)
28. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)

29. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 6840–6851 (2020)
30. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021)
31. Huang, G.B., Mattar, M., Berg, T., Learned-Miller, E.: Labeled Faces in the Wild: A database for studying face recognition in unconstrained environments. In: *Tech. Report* (2008)
32. Huang, X., Shao, R., Zhang, Q., Zhang, H., Feng, Y., Liu, Y., Wang, Q.: Human-norm: Learning normal diffusion model for high-quality and realistic 3d human generation. *arXiv preprint arXiv:2310.01406* (2023)
33. Kang, M., Shin, J., Park, J.: Studiogan: a taxonomy and benchmark of gans for image synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
34. Kang, M., Zhu, J.Y., Zhang, R., Park, J., Shechtman, E., Paris, S., Park, T.: Scaling up gans for text-to-image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10124–10134 (2023)
35. Kansy, M., Raël, A., Mignone, G., Naruniec, J., Schroers, C., Gross, M., Weber, R.M.: Controllable inversion of black-box face recognition models via diffusion. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*. pp. 3167–3177 (2023)
36. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196* (2017)
37. Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems* **34**, 852–863 (2021)
38. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4401–4410 (2019)
39. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8110–8119 (2020)
40. Kim, M., Jain, A.K., Liu, X.: Adaface: Quality adaptive margin for face recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 18750–18759 (2022)
41. Kim, M., Liu, F., Jain, A., Liu, X.: Dcfac: Synthetic face generation with dual condition diffusion model. In: *CVPR* (2023)
42. Kirschstein, T., Giebenhain, S., Nießner, M.: Diffusionavatars: Deferred diffusion for high-fidelity 3d head avatars. *arXiv preprint arXiv:2311.18635* (2023)
43. Kumari, N., Zhang, B., Zhang, R., Shechtman, E., Zhu, J.Y.: Multi-concept customization of text-to-image diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1931–1941 (2023)
44. Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* **36**(6), 194:1–194:17 (2017), <https://doi.org/10.1145/3130800.3130813>
45. Li, Z., Cao, M., Wang, X., Qi, Z., Cheng, M.M., Shan, Y.: Photomaker: Customizing realistic human photos via stacked id embedding. *arXiv preprint arXiv:2312.04461* (2023)
46. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017)

47. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems* **35**, 5775–5787 (2022)
48. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095* (2022)
49. Mai, G., Cao, K., Yuen, P.C., Jain, A.K.: On the reconstruction of face images from deep face templates. *IEEE transactions on pattern analysis and machine intelligence* **41**(5), 1188–1202 (2018)
50. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: SDEdit: Guided image synthesis and editing with stochastic differential equations. In: *International Conference on Learning Representations* (2022)
51. Mensah, D., Kim, N.H., Aittala, M., Laine, S., Lehtinen, J.: A hybrid generator architecture for controllable face synthesis. In: *ACM SIGGRAPH 2023 Conference Proceedings*. pp. 1–10 (2023)
52. Moschoglou, S., Papaioannou, A., Sagonas, C., Deng, J., Kotsia, I., Zafeiriou, S.: Agedb: the first manually collected, in-the-wild age database. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop* (2017)
53. Niemeyer, M., Geiger, A.: Giraffe: Representing scenes as compositional generative neural feature fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11453–11464 (2021)
54. Pan, X., Dong, L., Huang, S., Peng, Z., Chen, W., Wei, F.: Kosmos-g: Generating images in context with multimodal large language models. *arXiv preprint arXiv:2310.02992* (2023)
55. Papantoniou, F.P., Lattas, A., Moschoglou, S., Zafeiriou, S.: Relightify: Relightable 3d faces from a single image via diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 8806–8817 (October 2023)
56. Peng, X., Zhu, J., Jiang, B., Tai, Y., Luo, D., Zhang, J., Lin, W., Jin, T., Wang, C., Ji, R.: Portraitbooth: A versatile portrait model for fast identity-preserved personalization. *arXiv preprint arXiv:2312.06354* (2023)
57. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952* (2023)
58. Qiu, H., Yu, B., Gong, D., Li, Z., Liu, W., Tao, D.: SynFace: Face recognition with synthetic data. In: *ICCV* (2021)
59. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PMLR (2021)
60. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* **1**(2), 3 (2022)
61. Razzhigaev, A., Kireev, K., Kaziakhmedov, E., Tursynbek, N., Petiushko, A.: Black-box face recovery from identity features. In: *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part V* **16**. pp. 462–475. Springer (2020)
62. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)

63. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
64. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22500–22510 (2023)
65. Ruiz, N., Li, Y., Jampani, V., Wei, W., Hou, T., Pritch, Y., Wadhwa, N., Rubinstein, M., Aberman, K.: Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. arXiv preprint arXiv:2307.06949 (2023)
66. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022)
67. Sauer, A., Karras, T., Laine, S., Geiger, A., Aila, T.: Stylegan-t: Unlocking the power of gans for fast large-scale text-to-image synthesis. arXiv preprint arXiv:2301.09515 (2023)
68. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems* **35**, 25278–25294 (2022)
69. Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., Komatsuzaki, A.: Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. arXiv preprint arXiv:2111.02114 (2021)
70. Schwarz, K., Sauer, A., Niemeyer, M., Liao, Y., Geiger, A.: Voxgraf: Fast 3d-aware image synthesis with sparse voxel grids. *Advances in Neural Information Processing Systems* **35**, 33999–34011 (2022)
71. Seitzer, M.: pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid> (August 2020), version 0.3.0
72. Sengupta, S., Chen, J.C., Castillo, C., Patel, V.M., Chellappa, R., Jacobs, D.W.: Frontal to profile face verification in the wild. In: WACV (2016)
73. Shen, Y., Yang, C., Tang, X., Zhou, B.: Interfacegan: Interpreting the disentangled face representation learned by gans. *IEEE transactions on pattern analysis and machine intelligence* **44**(4), 2004–2018 (2020)
74. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: Proceedings of the 32nd International Conference on Machine Learning. vol. 37, pp. 2256–2265. Lille, France (07–09 Jul 2015)
75. Somepalli, G., Singla, V., Goldblum, M., Geiping, J., Goldstein, T.: Diffusion art or digital forgery? investigating data replication in diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6048–6058 (2023)
76. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=PxTIG12RRHS>
77. Truong, T.D., Duong, C.N., Le, N., Savvides, M., Luu, K.: Vec2face-v2: Unveil human faces from their blackbox features via attention-based network in face recognition. arXiv preprint arXiv:2209.04920 (2022)

78. Turk, M., Pentland, A.: Eigenfaces for recognition. *Journal of cognitive neuroscience* **3**(1), 71–86 (1991)
79. Valevski, D., Lumen, D., Matias, Y., Leviathan, Y.: Face0: Instantaneously conditioning a text-to-image model on a face. In: *SIGGRAPH Asia 2023 Conference Papers*. pp. 1–10 (2023)
80. Vendrow, E., Vendrow, J.: Realistic face reconstruction from deep embeddings. In: *NeurIPS 2021 Workshop Privacy in Machine Learning* (2021), <https://openreview.net/forum?id=-WsBmzWwPee>
81. Wang, Q., Jia, X., Li, X., Li, T., Ma, L., Zhuge, Y., Lu, H.: Stableidentity: Inserting anybody into anywhere at first sight. *arXiv preprint arXiv:2401.15975* (2024)
82. Wang, Q., Bai, X., Wang, H., Qin, Z., Chen, A.: Instantid: Zero-shot identity-preserving generation in seconds. *arXiv preprint arXiv:2401.07519* (2024)
83. Wang, X., Li, Y., Zhang, H., Shan, Y.: Towards real-world blind face restoration with generative facial prior. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021)
84. Xia, W., Zhang, Y., Yang, Y., Xue, J.H., Zhou, B., Yang, M.H.: Gan inversion: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(3), 3121–3138 (2022)
85. Xiao, G., Yin, T., Freeman, W.T., Durand, F., Han, S.: Fastcomposer: Tuning-free multi-subject image generation with localized attention. *arXiv preprint arXiv:2305.10431* (2023)
86. Yan, Y., Zhang, C., Wang, R., Zhou, Y., Zhang, G., Cheng, P., Yu, G., Fu, B.: Facestudio: Put your face everywhere in seconds. *arXiv preprint arXiv:2312.02663* (2023)
87. Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Zhang, W., Cui, B., Yang, M.H.: Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys* **56**(4), 1–39 (2023)
88. Yang, Z., Zhang, J., Chang, E.C., Liang, Z.: Neural network inversion in adversarial setting via background knowledge alignment. In: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. pp. 225–240 (2019)
89. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721* (2023)
90. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *GitHub repository* <https://github.com/tencent-ailab/IP-Adapter> (2024)
91. Yi, D., Lei, Z., Liao, S., Li, S.Z.: Learning face representation from scratch. *arXiv preprint arXiv:1411.7923* (2014)
92. Yuan, G., Cun, X., Zhang, Y., Li, M., Qi, C., Wang, X., Shan, Y., Zheng, H.: Inserting anybody in diffusion models via celeb basis. *arXiv preprint arXiv:2306.00926* (2023)
93. Zhang, L., Qiu, Q., Lin, H., Zhang, Q., Shi, C., Yang, W., Shi, Y., Yang, S., Xu, L., Yu, J.: Dreamface: Progressive generation of animatable 3d faces under text guidance. *arXiv preprint arXiv:2304.03117* (2023)
94. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3836–3847 (2023)

95. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
96. Zheng, T., Deng, W.: Cross-Pose LFW: A database for studying cross-pose face recognition in unconstrained environments. Tech. Report (2018)
97. Zheng, T., Deng, W., Hu, J.: Cross-Age LFW: A database for studying cross-age face recognition in unconstrained environments. Tech. Report (2017)
98. Zheng, Y., Yang, H., Zhang, T., Bao, J., Chen, D., Huang, Y., Yuan, L., Chen, D., Zeng, M., Wen, F.: General facial representation learning in a visual-linguistic manner. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18697–18709 (2022)
99. Zhou, Y., Zhang, R., Sun, T., Xu, J.: Enhancing detail preservation for customized text-to-image generation: A regularization-free approach. arXiv preprint arXiv:2305.13579 (2023)
100. Zhu, J., Feng, R., Shen, Y., Zhao, D., Zha, Z.J., Zhou, J., Chen, Q.: Low-rank subspaces in gans. *Advances in Neural Information Processing Systems* **34**, 16648–16658 (2021)
101. Zhu, Z., Huang, G., Deng, J., Ye, Y., Huang, J., Chen, X., Zhu, J., Yang, T., Lu, J., Du, D., et al.: Webface260m: A benchmark unveiling the power of million-scale deep face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10492–10502 (2021)