

FitDiff: Robust monocular 3D facial shape and reflectance estimation using Diffusion Models

Stathis Galanakis^{1,2}

Alexandros Lattas¹

Stylianos Moschoglou¹

Stefanos Zafeiriou¹

¹Imperial College London

²HUAWEI Noah’s Ark Lab

Abstract

The remarkable progress in 3D face reconstruction has resulted in high-detail and photorealistic facial representations. Recently, Diffusion Models have revolutionized the capabilities of generative methods by surpassing the performance of GANs. In this work, we present FitDiff, a diffusion-based 3D facial avatar generative model. Leveraging diffusion principles, our model accurately generates relightable facial avatars, utilizing an identity embedding extracted from an “in-the-wild” 2D facial image. The introduced multi-modal diffusion model is the first to concurrently output facial reflectance maps (diffuse and specular albedo and normals) and shapes, showcasing great generalization capabilities. It is solely trained on an annotated subset of a public facial dataset, paired with 3D reconstructions. We revisit the typical 3D facial fitting approach by guiding a reverse diffusion process using perceptual and face recognition losses. Being the first 3D LDM conditioned on face recognition embeddings, FitDiff reconstructs relightable human avatars, that can be used as-is in common rendering engines, starting only from an unconstrained facial image, and achieving state-of-the-art performance.

1. Introduction

A fundamental objective of Computer Vision encompasses photorealistic 3D face reconstruction from a single image, which has gained significant attention from the research community over the past few decades. Its numerous applications in computer graphics, virtual reality, and entertainment, include avatar creation, face animation and manipulation [1, 4, 26, 54, 83]. Despite the notable recent progress, accurate replication of personalized facial reconstructions continues to present a challenge. This is primarily due to the inherent ambiguity present in monocular images, associated with handling occlusions and capturing substantial variations in lighting conditions and facial expressions.

On top of that, captured 3D facial datasets are still relatively small [101], including biases and lacking generalization.

Since the introduction of the first 3D Morphable Model (3DMM) [6], there has been tremendous progress in retrieving 3D facial shape information from a monocular image thanks to large-scale statistical models utilizing hundreds of subjects [7, 49, 65]. More recently, Generative Adversarial Networks (GANs) [30], and particularly the style-based generators [37–39], have demonstrated effectiveness in capturing intricate facial frequencies resulting into numerous subsequent studies [24, 26, 27]. A fundamental challenge inherent in optimization fitting methods is their susceptibility to outliers, requiring to heuristically initialize the GAN’s \mathbf{z} or \mathbf{w} embedding for the back-propagation during inference to avoid instabilities [46]. Additionally, they suffer from problems like unstable training and mode collapse [41, 74].

A solution to the aforementioned problems is the integration of the recent emerging Diffusion Models (DMs) [32]. Taking inspiration from the thermodynamics [78], DMs define a T -length Markov Chain by gradually adding normally distributed noise to the data and learn to predict the input noise for each step $t \in \{1, 2, \dots, T\}$. This methodology has already been applied to synthesize and manipulate facial images [61, 91, 97]. Nevertheless, most of these methods employ a conditional mechanism reliant on textual descriptions or other auxiliary information [67], thereby directing their attention not exclusively towards the faithful reconstruction of the input facial identity and the exploration of the potential of robust identity embeddings. Moreover, Relightify [61] requires partially completed facial UV maps and third-party extracted facial shapes as input, thus being prone to these third-party failures.

This study presents FitDiff, the first multi-modal diffusion model that synthesizes high-fidelity facial avatars given only an input facial image, starting from a randomly initialized Gaussian noise. By harnessing the impressive generation capabilities of LDMs [72], we delve into their potential in 3D facial reconstruction. Our approach enables the synthesis of facial avatars by combining facial UV reflectance

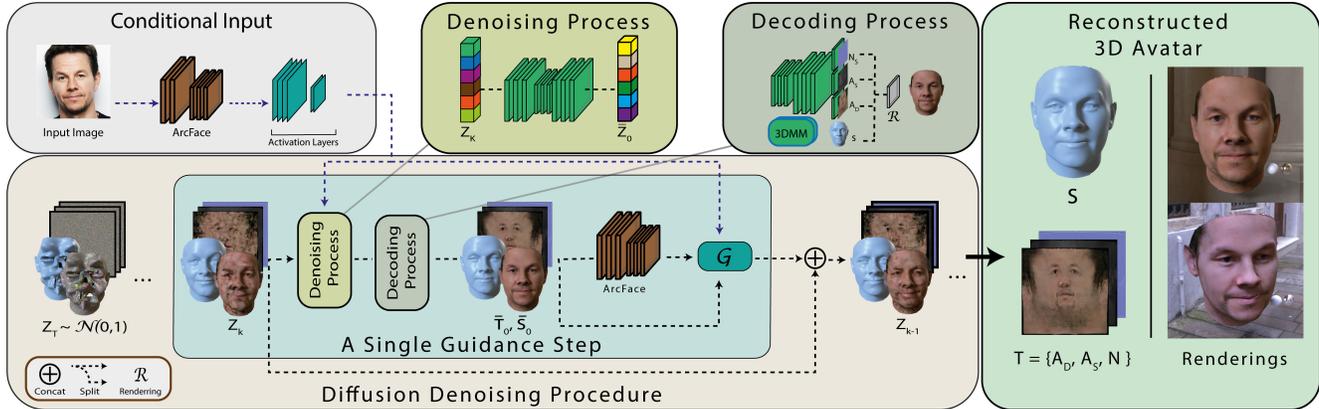


Figure 1. Overview of FitDiff, a diffusion-based 3D facial generative network. Starting from Gaussian noise, our method generates facial avatars with relightable reflectance and shape, conditioned on an identity embedding. During sampling, a novel guidance algorithm (\mathcal{G}) is applied for further control of the resulting identity. \mathbf{Z}_T , \mathbf{Z}_k and \mathbf{Z}_{k-1} are visualized in the actual picture space for illustration purposes.

maps and facial geometry while utilizing an identity embedding layer as a conditioning mechanism. The diffusion process is applied to the concatenation of the latent vectors of the facial texture maps and shape, while a VQGAN AutoEncoder [20] and a 3DMM model (LSFM [7]) are used to decode them. Moreover, we present a novel facial guidance algorithm incorporated into the reverse diffusion process for accurately reconstructing a target facial identity. FitDiff generates high-fidelity facial avatars while achieving a state-of-the-art identity preservation score. The training of our model involves fitting a facial reconstruction network [43] to a manually selected and curated set of images acquired from the CelebA-HQ dataset [36]. The fitting process results in acquiring facial shape and facial reflectance maps. Overall, in this paper:

- We introduce FitDiff, a multi-modal diffusion-based generative model that jointly produces facial geometry and appearance. The facial appearance consists of diffuse albedo, specular albedo, and normal maps, enabling photorealistic rendering.
- We show the first diffusion model conditioned on identity embeddings, acquired from an off-the-shelf face recognition network, whilst introducing a SPADE [63]-conditioned UNet architecture.
- We present unconditional samples of relightable avatars, but most importantly, we achieve facial reconstruction from a single “in-the-wild” image through identity embedding conditioning and guidance.

2. Related Work

2.1. Face Modeling

Various works followed the introduction of the emblematic first 3D Morphable Model (3DMM) [6] trained on

200 distinct subjects. Numerous works have since proposed improvements, especially on realistic expressions [3, 9, 10, 12, 47–49, 86, 94], large-scale dataset and model releases [8, 16, 65, 77]. However, those models could not capture high-frequency details due to their linear nature. To deal with this, many studies integrated 3DMMs with Deep Neural Networks [84, 88], Mesh Convolutions [57], GANs [25–27] and VAEs [5, 51, 70, 92]. Additionally, the photorealistic rendering of implicit representation-based methods [55, 62] led to numerous approaches [4, 22, 23, 34, 54, 59, 95]. However, their applications and editability remain challenging, hence this work focuses on explicit representations. Emphasizing the acquisition of highly detailed facial texture and faithful reflectance maps, Avatarme [42] and AvatarMe++ [44] treat the facial texture as a combination of diffuse albedo, specular albedo and normals UV maps whereas ReflectanceMM [31] models spatially varying BRDF, while trained in low-cost acquired data. In a different vein, the authors of MICA [101] leverage a robust face recognition network [17] to generate facial shapes. In [93], dense facial landmarks were introduced for better shape reconstruction, whilst a transformer-based facial reconstruction network was introduced in [99]. The techniques above and more recent methodologies [14, 21, 45, 68], which utilize displacement maps for finer facial shape details, consistently achieve state-of-the-art performance in facial shape competitions such as NoW [75] and REALY [13]. However, many of these approaches produce facial texture with baked illumination or fail to generate it. Recent works such as [19, 43, 52] aim to concurrently reconstruct facial shape and texture but still rely on separate dedicated models for each component (Fig. 2). Closer to our work is AlbedoGAN [68], which introduces a single network for acquiring both facial shape and texture. However, contrary to our approach, the proposed methodology

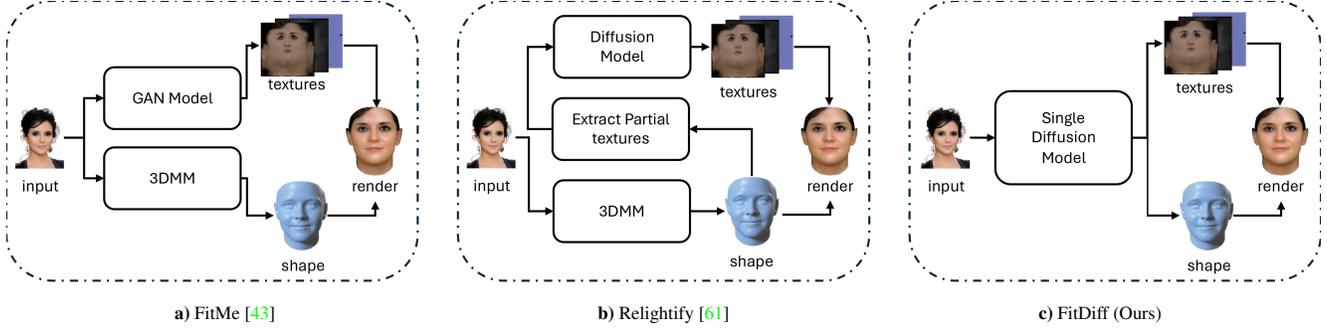


Figure 2. Differences with existing state-of-the-art methods [43, 61]: Prior works rely on multiple separate models, which can fail on challenging inputs (Fig. 10). In contrast, our method uses only a single Latent Diffusion Model for both shape and reflectance texture prediction, achieving simplified architecture, training, and robustness.

generates maps containing baked illumination and cannot handle wearables.

2.2. Diffusion Models

Inspired by [78, 80–82], the authors of [18] showed that DMs can perform better than the widely used GAN-based methods in image synthesis tasks. The high-quality samples and the more stable training have attracted the attention of the research community leading to a great variety of applications in image generation [33, 69, 72, 72, 76, 89], text-to-image [40, 58], text-to-3D [66], Pointcloud [53, 96, 100] and Mesh [50] generation. Closest to our work, other diffusion-based facial models [61, 91, 97] were presented. Unlike our single-stage methodology, the former two employ a coarse-to-fine approach to attain identity information. Furthermore, Rodin [91] uses an implicit representation [15] for facial shape but doesn’t prioritize faithful reconstruction of the input facial identity. On the contrary, FitDiff integrates a fully controllable 3DMM model, facilitating the generation of high-quality facial avatars that accurately represent the input identity. On the other hand, Relightify [61] uses an unconditional denoising network for filling partially completed facial UV maps extracted from third-party models. Although it may seem similar to our proposed methodology, FitDiff has several advantages over Relightify, which we discuss in detail in Section 6.

3. Method

In this work, we propose FitDiff, a latent-diffusion-based approach to reconstruct facial avatars. We harness the power of diffusion models for the generative and fitting process as, by nature, they are very robust in both processes since they directly operate on the image space [72]. On the other hand, GANs suffer from various issues such as mode collapse during training [85] or unrealistic outputs in fitting methods, resulting in unnecessary heuristics [43] to stabilize the outputs at the expense of fidelity.

A facial avatar can be formulated as a combination of a

mesh \mathbf{S} and texture \mathbf{T} which is defined as a combination of facial reflectance UV maps, namely diffuse albedo (\mathbf{A}_D), specular albedo (\mathbf{A}_S) and normals (\mathbf{N}). Also, let us denote an “in-the-wild” image containing a face as \mathbf{I} , and \mathbf{V}_{tgt} as the corresponding target identity embedding of the appearing face. Given \mathbf{V}_{tgt} as input, FitDiff generates a 3D facial avatar of the same identity as the one in \mathbf{I} , alongside with the current scene’s illumination parameters (ambient, diffuse, and specular lighting and lighting direction). An overview of our method is illustrated in Fig. 1, whereas the rest of the section includes a detailed presentation of our method’s architecture (Sec. 3.1), the proposed conditional input (Sec. 3.2), the training scheme (Sec. 3.3), and finally the identity-guidance sampling procedure (Sec. 3.4).

3.1. Model Architecture

FitDiff is a latent-diffusion based approach [72]. This is motivated by the challenges posed by the large number of parameters and the computational expenses associated with simultaneously generating multiple meshes and texture images. Thus, we represent facial avatars as a latent vector containing latent information about the shape, facial texture, and scene illumination: $\mathbf{z} = \{\mathbf{z}_{tex} | \mathbf{z}_{shp} | \mathbf{z}_{ill}\} \in \mathbb{R}^{4288}$, where $\mathbf{z}_{tex} \in \mathbb{R}^{4096}$ signifies the facial reflectance latent vector, $\mathbf{z}_{shp} \in \mathbb{R}^{183}$ the latent shape vector and $\mathbf{z}_{ill} \in \mathbb{R}^9$ the scene illumination parameters. The latent shape vector \mathbf{z}_{shp} can be further separated into the identity parameters $\mathbf{z}_{shp_i} \in \mathbb{R}^{158}$ and expression parameters $\mathbf{z}_{shp_e} \in \mathbb{R}^{25}$.

FitDiff is composed of a 3D statistical model \mathcal{F}_{shp} (LSFM [7]) that generates facial geometry, a branched multi-modal AutoEncoder [20, 72] that generates facial reflectance maps and a denoising UNet AutoEncoder [73]. For a set of identity \mathbf{z}_{shp_i} and expression \mathbf{z}_{shp_e} parameters, the PCA face model \mathcal{F}_{shp} generates a facial mesh \mathbf{S} following the formula:

$$\mathbf{S} = \mathcal{F}_{shp}(\mathbf{z}_{shp}) = \mathbf{U}_i \cdot \mathbf{z}_{shp_i} + \mathbf{U}_e \cdot \mathbf{z}_{shp_e} + \mathbf{m}_i$$

where \mathbf{U}_i and \mathbf{U}_e are the identity and expression bases, respectively, and \mathbf{m}_s is the mean face. Additionally, we

incorporate a robust branched VQGAN [20, 72], to function as a multi-modal texture UV generator. More specifically, the VQGAN encoder \mathcal{E} concurrently encodes facial diffuse albedo \mathbf{A}_D , specular albedo \mathbf{A}_S and normals \mathbf{N} into the same latent vector \mathbf{z}_{tex} , whereas the VQGAN decoder \mathcal{D} reconstructs them given the input latent vector \mathbf{z}_{tex} . Finally, following common diffusion-based methods [18, 32, 72], we utilize a UNet AutoEncoder [73] with self-attention [90] layers $e_\theta(x_t, t)$ conditioned to the input time step $t \in \{1, \dots, T\}$. We train the UNet to predict the injected noise ϵ , sampled from a standard normal distribution, i.e., $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$. The training details of our method are presented in Sec. 3.3.

3.2. Conditional input

In the underlying UNet model, we integrate a powerful conditioning mechanism to effectively learn all the necessary identity information. An ideal identity embedding must contain low and high-frequency information to accurately reconstruct the desired facial avatar. To acquire such an identity embedding, we employ a powerful identity recognition network [17], to which we feed the input facial image. Then, the resulting conditioning vector combines the last feature vector with the intermediate activation layers of the identity recognition network.

Let \mathcal{C}^n be the n -th intermediate layer of the identity recognition network. We extract the intermediate activation maps $\mathcal{C}^2 \in \mathbb{R}^{128 \times 28 \times 28}$, $\mathcal{C}^3 \in \mathbb{R}^{256 \times 14 \times 14}$, $\mathcal{C}^4 \in \mathbb{R}^{512 \times 7 \times 7}$ and concatenate them channel-wise with the identity embedding $\mathbf{V} \in \mathbb{R}^{512}$, which is expanded spatially. Because of the 2D nature of our identity embedding and following [22], we use SPADE layers [63] as a conditioning mechanism to inject the conditioning vector into the intermediate layers of the UNet.

3.3. Model Training

Our training scheme consists of two phases: Initially, we conduct the training for the branched texture AE, followed by the subsequent training for the denoising UNet model. The first part of our training protocol entails the training of the branched multi-modal AE [20], whereby triplets are employed as input data:

$$x = \{\mathbf{A}_D, \mathbf{A}_S, \mathbf{N}\}, \quad \mathbf{A}_D, \mathbf{A}_S, \mathbf{N} \in \mathbb{R}^{512 \times 512 \times 3}$$

where \mathbf{A}_D represents the diffuse albedo, \mathbf{A}_S the specular albedo and \mathbf{N} the normals. Our approach includes a branched multi-modal discriminator [43] in combination with a perceptual loss [98] as the training loss. The discriminator is a path-based discriminator [20] comprising two branches, accommodating their different statistics [43]. The first branch gets as input the concatenation of diffuse and specular albedos $\mathbf{A}_D \oplus \mathbf{A}_S$, whereas the second branch receives the normals \mathbf{N} . Furthermore, we adhere to the

default training parameters outlined in [72]. For a given triplet x_k , the encoder \mathcal{E} projects x_k into a latent representation $\mathbf{z}_{tex} = \mathcal{E}(x_k)$, where $\mathbf{z}_{tex} \in \mathbb{R}^{h \times w \times c}$. Then, the latent vector \mathbf{z}_{tex} is fed into the decoder \mathcal{D} , which produces a reconstructed output triplet \bar{x}_k . We downsample the input texture UVs by a factor of $f = H/h = 512/64 = 8$, following the downsampling investigations in [72] and due to computational limitations. We use the LSFm [7] model for the shape decoder, pre-trained with $\sim 10k$ identities.

After training the texture AutoEncoder, we freeze its weight parameters and embark on training the identity-conditioned UNet. During this phase, we employ multiple components including the identity embedding \mathbf{V}_{tgt} , the shape \mathbf{z}_{shp} , the distinct facial texture maps $\mathbf{A}_D, \mathbf{A}_S$, and \mathbf{N} , and the scene lighting \mathbf{z}_{ill} . At each training step, the facial reflectance maps undergo an initial encoding by the pre-trained encoder \mathcal{E} , thereby yielding the latent texture vector \mathbf{z}_{tex} . Then, the input vectors are concatenated, i.e., $\mathbf{z}_0 = \{\mathbf{z}_{tex} | \mathbf{z}_{shp} | \mathbf{z}_{ill}\}$. Let us denote \mathbf{z}_t the noisy counterparts of \mathbf{z}_0 , resulting from t steps of noise injection. The UNet network gets \mathbf{z}_t as input and learns to predict the injected noise following:

$$L_{noise} := \mathbb{E}_{\mathcal{E}(x), \epsilon \sim \mathcal{N}(0,1), t} [\| \epsilon - \epsilon_\theta(\mathbf{z}_t, t, \mathbf{V}) \|] \quad (1)$$

where ϵ is the ground truth injected noise, ϵ_θ the predicted injected noise from \mathbf{z}_t , t the diffusion time step, and \mathbf{V} signifies the 2D identity embedding vector (Sec. 3.2).

In addition to the primary loss function \mathcal{L}_{noise} , our training scheme integrates additional losses intended to enhance robustness, which are the identity losses [26, 43] \mathcal{L}_{id} , \mathcal{L}_{per} and the shape loss \mathcal{L}_{verts} . A detailed definition of those losses is provided in the supplemental material along with more training details. It is important to note that these auxiliary losses do not apply to the latent variables. Thus, in every training step, the initial latent vector $\bar{\mathbf{z}}_0$ is estimated by: $\bar{\mathbf{z}}_0 = \frac{\mathbf{z}_t - \sqrt{1 - \alpha_t} \epsilon}{\sqrt{\alpha_t}}$. The vector $\bar{\mathbf{z}}_0$ is decoded into the estimated initial avatar. Firstly, $\bar{\mathbf{z}}_0$ is split into the estimated initial latent texture vector $\bar{\mathbf{z}}_{0_{tex}}$, latent shape vector $\bar{\mathbf{z}}_{0_{shp}}$ and scene parameters $\bar{\mathbf{z}}_{0_{ill}}$. The first two vectors are fed into the decoder \mathcal{D} and the PCA model \mathcal{F}_{shp} respectively. In this way, the estimated initial facial reflectance maps $\bar{\mathcal{T}}_0$ and shape $\bar{\mathcal{S}}_0$ are retrieved. Additionally, under the estimated initial scene illumination $\bar{\mathbf{z}}_{0_{ill}}$, we acquire the initial identity rendering \bar{I}_0 using a differentiable renderer [71] while using a differentiable multi-texture map shader as introduced in [43, 44] under a single directional light. Overall, the conditional denoising model is trained using the following formula:

$$\mathcal{L} = \mathcal{L}_{noise} + \mathcal{L}_{id} + \mathcal{L}_{per} + \mathcal{L}_{verts}$$

where \mathcal{L}_{noise} is defined in Eq. 1, \mathcal{L}_{id} is the identity distance, \mathcal{L}_{per} the identity perceptual loss and \mathcal{L}_{verts} the shape loss.

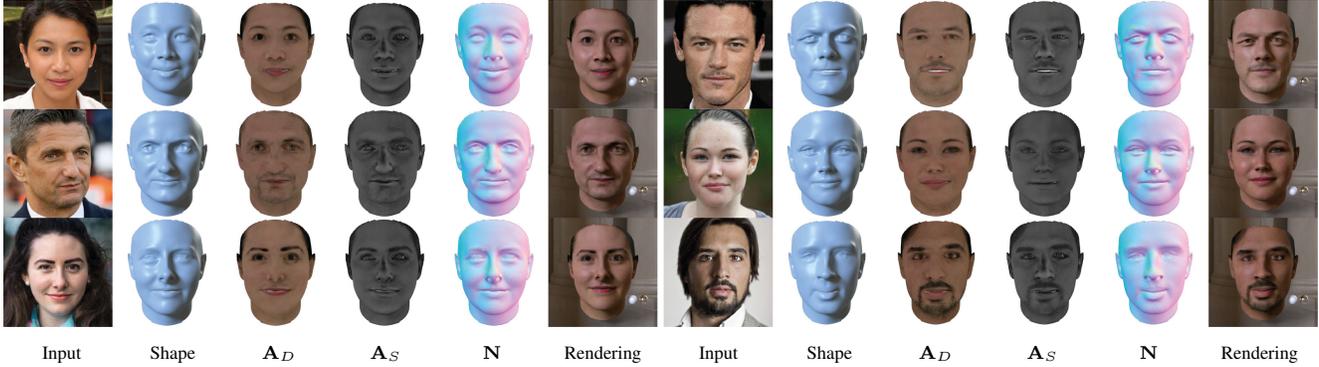


Figure 3. Qualitative results of FitDiff on “in-the-wild” facial images, showing shape, reflectance, and environment map renderings.

Even though both the forward and the reverse diffusion processes can be described via stochastic differential equations in a continuous time [82], they can also be applied in discrete time by choosing a very small step each time. The selection of the appropriate number of diffusion steps is based on the premise that, in the final step, the input data should be completely converted into random noise. We follow the training parameters proposed by the authors of [72], and we choose $T = 1000$ while using a linear noise schedule. Furthermore, FitDiff is trained following the Classifier-Free approach [33], aiming to generate novel identities without any prior input. During training and with a probability of $\mathcal{P}_{uncond} = 0.1$, we randomly set the input identity embedding equal to zero by feeding the face recognition network [17] with an empty image.

3.4. Sampling Procedure

DMs generate new samples by reversing the diffusion process, commencing from an initial random Gaussian noise. In parallel, the authors of [79] introduced Denoising Diffusion Implicit Models (DDIMs), which are implicit probabilistic models [56]. They conduct a modified reverse diffusion process with fewer diffusion steps than those required during the vanilla DDPM sampling.

In our method, we adopt the DDIM sampling technique and integrate it into our trained architecture to generate facial avatars. We select the number of sampling steps to be $T = 50$ for the DDIM sampling process. Aiming to generate accurate photorealistic avatars, we employ a novel guidance algorithm alongside the conditional input. Our approach is inspired by [18], in which the authors introduce a score corrector network conditioned on the diffusion step. In our implementation, we incorporate a face recognition network \mathcal{C} [17] as a score corrector alongside an off-the-shelf facial landmark detector \mathcal{M} [11] and a perceptual loss [98]. The guidance method can be implemented for the vanilla DDPM [18] and the DDIM [79] sampling techniques. For the generation of intermediate images, we employ a differentiable renderer [71] with the modifications

introduced in [43, 44] under a single directional light. The pseudo-code and a detailed presentation of the proposed guidance method are presented in the supplemental while the guidance formula is the following:

$$\mathcal{G} = \mathcal{G}_{id}^{cos} + \lambda_1 \mathcal{G}_{id}^{per} + \lambda_2 \mathcal{G}_{mse} + \lambda_3 \mathcal{G}_{lan} + \lambda_4 \mathcal{G}_{vgg} \quad (2)$$

where \mathcal{G}_{id}^{cos} denotes the cosine similarity between the identity vectors, \mathcal{G}_{id}^{per} the identity perceptual similarity, \mathcal{G}_{mse} the photometric loss, \mathcal{G}_{lan} is the distance between the 3D facial landmarks extracted by using \mathcal{M} [11] and \mathcal{G}_{vgg} is the perceptual similarity using [98]. Fig. 3 showcases examples of our method being applied to “in-the-wild” images.

4. Experiments

4.1. Dataset

Training such a diffusion model in a supervised manner requires a large dataset of labeled sets of facial images \mathbf{I} , facial textures \mathbf{T} , shape parameters \mathbf{z}_{shp} and facial recognition embeddings \mathbf{V} . Although a captured dataset could be used, there are no large enough public datasets [101]. As a workaround, we curate 9000 2D facial images from the CelebA-HQ Dataset [36] $\mathbf{I}_i, i = 0 \dots 9 \cdot 10^3$, on which we performed the following steps to acquire the labeled dataset: A) We use a state-of-the-art face recognition model [17], to extract the identity latent embeddings \mathbf{V} , which captures the facial structure with minimal interference from shading, age, and accessories. B) We train a state-of-the-art StyleGAN-based [39] facial reconstruction network [43] ϕ , on public datasets of facial textures [42, 60], and use the LSFM 3DMM for the facial shape [7]. Following an iterative optimization [43], we fit our model to the CelebA-HQ dataset [36] and acquire pseudo-ground truth facial textures, 3DMM shape weights, and scene illumination parameters $\phi(\mathbf{I}_i) \rightarrow \mathbf{A}_{D_i}, \mathbf{A}_{S_i}, \mathbf{N}_i, \mathbf{z}_{shp_i}, \mathbf{z}_{ill_i}$. After manually filtering out all failed cases, we finalize a dataset consisting of paired images, facial reflectance textures, facial shape, scene illumination and latent vectors: $\{\mathbf{I}_i, \mathbf{A}_{D_i}, \mathbf{A}_{S_i}, \mathbf{N}_i, \mathbf{z}_{shp_i}, \mathbf{z}_{ill_i}, \mathbf{V}_i\}$.

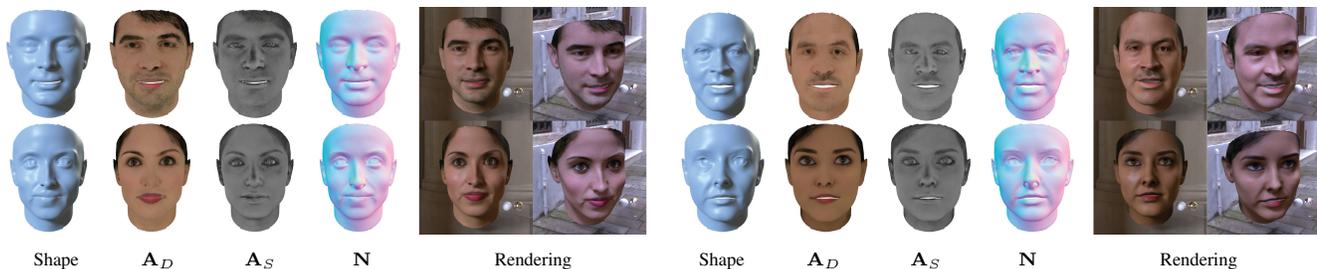


Figure 4. Samples generated by FitDiff with unconditional sampling. Our method can generate diverse facial shapes and reflectance maps.

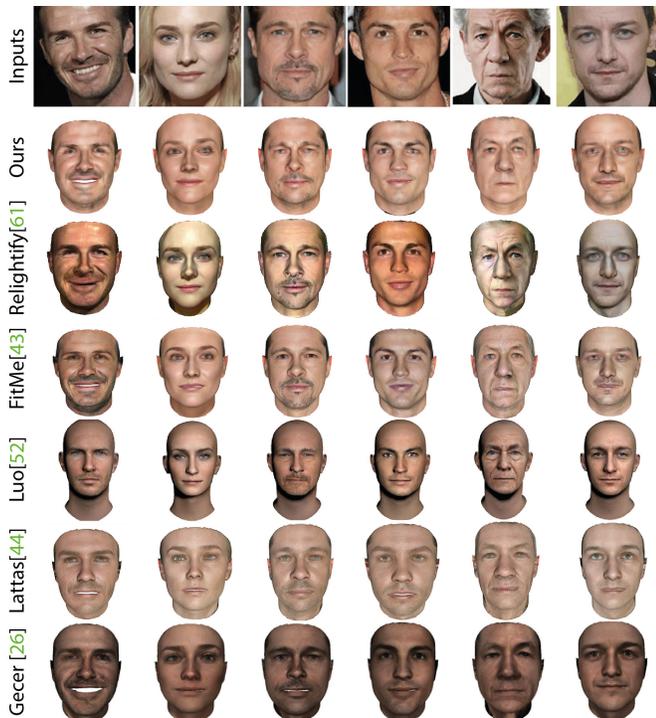


Figure 5. Qualitative comparison between FitDiff and other monocular face reconstruction approaches [26, 43, 44, 52, 61].

4.2. Unconditional Sampling

Following the classifier-free guidance (CFG) [33] training scheme, FitDiff can generate completely random facial identities without any prior input or supervision. We present the unconditional generated diffuse albedos A_D , specular albedos A_S , normals N , facial shapes S and renderings in Fig. 4. This figure illustrates our method’s ability to create distinct shapes and textures.

4.3. Qualitative comparisons

We compare our method’s generated samples with other monocular-image face reconstruction methods [26, 43, 44, 52, 61] and present the generated samples in Fig. 5. Most of these techniques rely on GAN-based methods and em-

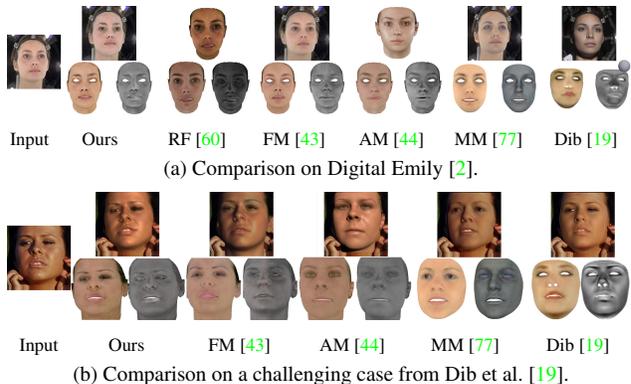


Figure 6. Qualitative comparison on single-image reflectance acquisition against Relightify (RF) [61], FitMe (FM) [43], AvatarMe++ (AM) [44], AlbedoMM (MM) [77] and Dib et al. [19]. Up: overlaid rendering, Left: diffuse, Right: specular.

ploy fitting optimization procedures encompassing lighting, camera pose, and expression parameters, whereas Relightify [61] is the only diffusion-based approach. Our method can capture finer details than most GAN-based methods ([26, 43, 44, 52]), whereas it generates equally detailed avatars like Relightify.

4.4. Quantitative comparisons

4.4.1 Facial Reflectance Acquisition Comparison

We evaluate the quality of our method’s generated facial reflectance maps by reconstructing 6 test subjects captured with a Light Stage [29]. We compare the generated diffuse albedos A_D , specular albedos A_S and normals N with the respective ground truth, and MSE, PSNR, and SSIM distances are measured. We compare our method’s performance with AlbedoMM [77], AvatarMe++ [44] FitMe [43] and Relightify [61] and the results are presented in Tab. 1 and Fig. 6. FitDiff generates state-of-the-art shape normals, whereas it performs on par with Relightify [61] in the diffuse and specular albedo scenarios.

4.4.2 Identity preservation experiment

One of the key elements of our model is the ability to generate the facial identity depicted in the provided “in-

	Diffuse Albedo			Specular Albedo			Normals		
	↓MSE	↑PSNR	↑SSIM	↓MSE	↑PSNR	↑SSIM	↓MSE	↑PSNR	↑SSIM
MM [77]	0.028	15.82	0.595	0.007	21.24	0.608	-	-	-
AM [44]	0.014	18.30	0.635	0.005	19.77	0.640	0.002	27.26	0.723
FM [43]	0.009	21.12	0.645	0.004	23.95	0.642	0.002	26.77	0.719
RF [61]	0.009	22.47	0.647	0.003	27.17	0.710	0.002	26.69	0.719
Ours	0.009	21.16	0.647	0.004	24.30	0.645	0.001	28.74	0.734

Table 1. Quantitative comparison on 6 Light-Stage-captured data [29], between our method, AlbedoMM [77] (MM), AvatarMe++ [44] (AM), FitMe [43] (FM) and Relightify [61] (RF), measuring MSE, PSNR, and SSIM. Our method surpasses prior work in most cases or works on par with the current state-of-the-art in the rest.

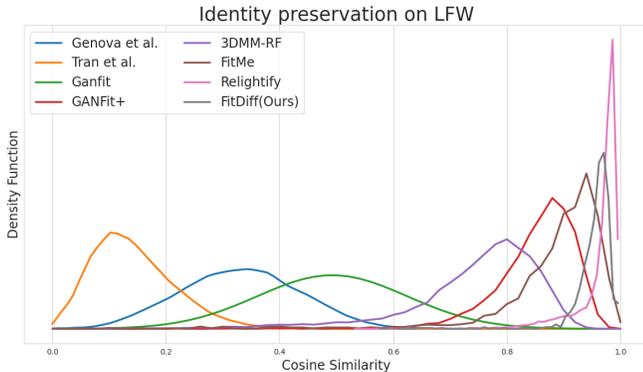


Figure 7. We compare our approach with [22, 26–28, 43, 61, 87]. FitDiff performs on par with the current state-of-the-art method (Relightify [61]), while beating the rest.

the-wild” image. We quantitatively measure this by conducting an identity preservation experiment [22, 26–28]. We reconstruct the facial identities depicted in each image within the Labeled Faces in the Wild (LFW) dataset [35]. The reconstructed identities are fed into a face recognition network [64] and the identity cosine distance is measured by comparing their activation layers. As depicted in Fig. 7, FitDiff outperforms the previous state-of-the-art face-reconstruction methods [26, 43] and performs slightly less than current state-of-the-art method [61].

5. Ablation Study

5.1. Conditioning embedding

The first ablation study focuses on the introduced conditioning embedding vector. Given the face recognition network [17], We examine 3 different types of input identity embeddings a) using only the last layer b) using only the first 3 layers and c) using all four of them, as proposed in FitDiff. We randomly pick about 100 “in-the-wild” images found across the internet, and we fit them using FitDiff. During sampling, we don’t use the guidance algorithm, as we want to let our network be entirely dependent on the identity embedding used. The first two rows of Table 2 show that the proposed condition vector surpasses the other.

Method	Using $n \in \{1, 2, 3\}$ layers	Using n layer	FitDiff	
ID Sim.	0.31	0.39	0.45	
Method	Label Only	CFG (w=2)	CFG (w=9)	Guidance
ID Sim.	0.43	0.49	0.45	0.88

Table 2. Ablation study on the proposed conditioning input vector and the identity similarity performance of our method, using with and without identity guidance.

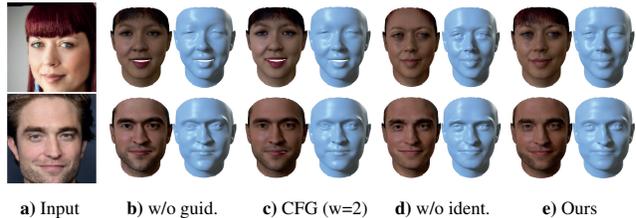


Figure 8. Ablation study: From input images (a), we show results without the guidance algorithm (b), using CFG with $w=2$ (c), without identity embedding (d), and using our full method (e).

5.2. Sampling guidance

Additionally, we examine the importance of the proposed facial guidance during the reverse diffusion process. Using the previously picked images, we consider three scenarios: a) sampling without any guidance, b) sampling using the classifier-free guidance [33] while using guidance scales $w = \{2, 9\}$, and c) our proposed method. The identity similarity scores are presented in Tab. 2, whereas visualizations of the generated examples are included in Fig. 8. The guidance algorithm demonstrates superior reconstruction performance compared to alternative methodologies.

5.3. Use of the conditioning mechanism

Another ablation study includes the necessity of our conditioning mechanism. We compare the texture information between the generated samples with and without our conditioning mechanism, and examples of those are presented in Fig. 8d and 8e, respectively. These examples clearly show that finer details can be generated only when the corresponding identity embedding is used as input.

6. Discussion

FitDiff is a latent diffusion model which concurrently carries out facial shape and texture generation as a combination of diffuse albedo, specular albedo, and normals. A close work to ours is Relightify [61], which only generates facial texture UV maps modeled like ours. It treats the texture reconstruction problem as an in-painting approach by copying the visible part of the facial textures acquired from third-party off-the-shelf approaches. Then, a multi-modal diffusion model is used to complete the non-visible parts of the facial texture maps and reflectance. Relightify retains the visible input during inference, resulting in a great per-

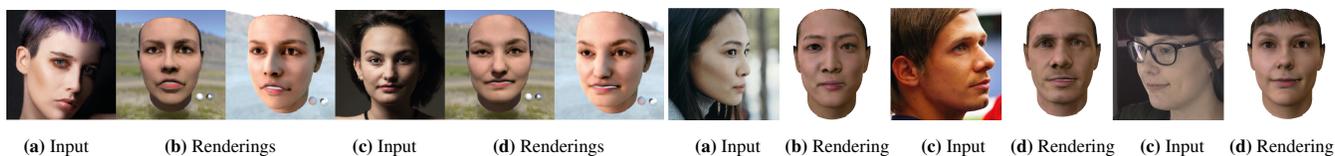


Figure 9. Examples of FitDiff under extreme illumination (left) and extreme angles (right)

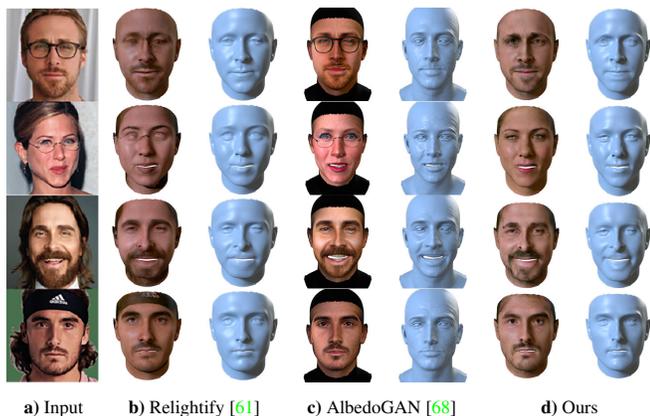


Figure 10. Qualitative comparison between FitDiff, Relightify [61], and AlbedoGAN [68] highlights the strengths of our approach. While Relightify is sensitive to occlusions and AlbedoGAN generates a single texture by replicating the input image, FitDiff produces precise and high-quality facial reflectance maps.

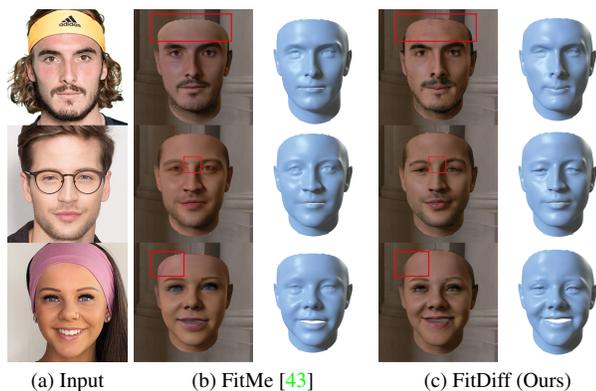


Figure 11. We applied our method to challenging images from FitMe [43], provided by the authors. Despite occlusions in the input facial images, our reconstructions remain unaffected.

formance in identity preservation (Sec. 4.4.2), which, however, introduces certain limitations: a) low-resolution images result in low-resolution textures, and b) wearable items (glasses, headbands) partially leak into the reflectance textures, as shown in Fig. 10. On the contrary, FitDiff concurrently generates facial texture maps and geometry from scratch and initializes the process from Gaussian noise. This results in a) robust shape and texture reconstruction even under occlusions and accessories, and b) always generating high-quality texture maps. Despite these advantages, our approach performs slightly worse than Relightify in fa-

cial texture benchmarks. This is due to Relightify’s albedo-copying approach and the fact that FitDiff was trained in solely synthetic data, contrary to the light-stage training data used in Relightify. We sincerely believe that the robustness benefits above make FitDiff an important alternative approach with multiple important use cases. In fact, our model can be used in tandem with Relightify to provide more accurate shape priors to further boost its performance.

On the other hand, the very recent AlbedoGAN [68] is the first model that concurrently generates facial texture and shape. It provides FLAME [49] parameters alongside displacement maps for accurate facial shape reconstruction. While finetuning on the image input, it also produces a facial texture UV map, often incorporating baked illumination and accessories. This consequently leads to reduced relightability in the generated avatars. It also replicates the limitations observed in Relightify, where wearable items are incorporated into the final facial texture (Fig. 10c).

Compared with FitMe [43], our experimentation results and qualitative assessments show that our generated texture maps achieve superior performance. This is due to the necessity in GAN-based fitting methods to meticulously initialize the GAN’s z or w embedding for back-propagation during inference to mitigate optimization instabilities. The proposed method excels without the need for heuristic priors or regularization terms.

7. Conclusion

In this paper, we introduced FitDiff, a diffusion-based 3D facial generative model conditioned on identity embeddings from a pre-trained facial recognition system. This approach captures diverse attributes like ethnicity, age, and gender from a single 2D image, with no restriction on quality, pose, or illumination. Our method jointly generates facial shapes, facial reflectance maps, and scene illumination parameters. Through a series of experiments, FitDiff showcases state-of-the-art performance in preserving identity and reconstructing facial reflectance, matching or surpassing established methods. Finally, it can generate unconditional samples and handles occlusions, further highlighting its versatility and effectiveness.

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).

References

- [1] V. Abrevaya, A. Boukhayma, S. Wuhrer, and E. Boyer. A decoupled 3d facial shape model by adversarial training. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9418–9427, Los Alamitos, CA, USA, nov 2019. IEEE Computer Society. [1](#)
- [2] Oleg Alexander, Mike Rogers, William Lambeth, Jen-Yuan Chiang, Wan-Chun Ma, Chuan-Chang Wang, and Paul Debevec. The digital emily project: Achieving a photorealistic digital actor. *IEEE Computer Graphics and Applications*, 30(4):20–31, 2010. [6](#)
- [3] Brian Amberg, Reinhard Knothe, and Thomas Vetter. Expression invariant 3D face recognition with a morphable model. In *2008 8th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2008*, pages 1–6. IEEE, 2008. [2](#)
- [4] ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. Rignerf: Fully controllable neural 3d portraits. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [2](#)
- [5] Timur Bagautdinov, Chenglei Wu, Jason Saragih, Pascal Fua, and Yaser Sheikh. Modeling facial geometry using compositional vaes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3877–3886, 2018. [2](#)
- [6] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH '99*, 1999. [1](#), [2](#)
- [7] James Booth, Anastasios Roussos, Allan Ponniah, David Dunaway, and Stefanos Zafeiriou. Large scale 3d morphable models. *International Journal of Computer Vision*, 126(2):233–254, 2018. [1](#), [2](#), [3](#), [4](#), [5](#)
- [8] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3D morphable model learnt from 10,000 faces. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2016-December, pages 5543–5552, 2016. [2](#)
- [9] Sofien Bouaziz, Yangang Wang, and Mark Pauly. Online modeling for realtime facial animation. *ACM Transactions on Graphics*, 32(4):40, 2013. [2](#)
- [10] Martin Breidt, Heinrich H. Bülthoff, and Cristobal Curio. Robust semantic analysis by synthesis of 3D facial motion. In *2011 IEEE International Conference on Automatic Face and Gesture Recognition and Workshops, FG 2011*, pages 713–719. IEEE, 2011. [2](#)
- [11] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *International Conference on Computer Vision*, 2017. [5](#)
- [12] Chen Cao, Yanlin Weng, Shun Zhou, Yiyong Tong, and Kun Zhou. FaceWarehouse: A 3D facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):413–425, 2014. [2](#)
- [13] Zenghao Chai, Haoxian Zhang, Jing Ren, Di Kang, Zhengzhuo Xu, Xuefei Zhe, Chun Yuan, and Linchao Bao. Realy: Rethinking the evaluation of 3d face reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [2](#)
- [14] Zenghao Chai, Tianke Zhang, Tianyu He, Xu Tan, Tadas Baltrusaitis, HsiangTao Wu, Runnan Li, Sheng Zhao, Chun Yuan, and Jiang Bian. Hiface: High-fidelity 3d face reconstruction by learning static and dynamic details. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9087–9098, October 2023. [2](#)
- [15] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J. Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16123–16133, June 2022. [3](#)
- [16] Hang Dai, Nick Pears, William Smith, and Christian Duncan. A 3D morphable model of craniofacial shape and texture variation. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2017-October, pages 3104–3112, 2017. [2](#)
- [17] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. [2](#), [4](#), [5](#), [7](#)
- [18] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794. Curran Associates, Inc., 2021. [3](#), [4](#), [5](#)
- [19] Abdallah Dib, Cedric Thebault, Junghyun Ahn, Philippe-Henri Gosselin, Christian Theobalt, and Louis Chevallier. Towards high fidelity monocular face reconstruction with rich reflectance using self-supervised learning and ray tracing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12819–12829, 2021. [2](#), [6](#)
- [20] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12873–12883, June 2021. [2](#), [3](#), [4](#)
- [21] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Trans. Graph.*, 40(4), jul 2021. [2](#)
- [22] Stathis Galanakis, Baris Gecer, Alexandros Lattas, and Stefanos Zafeiriou. 3dmm-rf: Convolutional radiance fields for 3d face modeling. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3536–3547, January 2023. [2](#), [4](#), [7](#)
- [23] Chen Gao, Yichang Shih, Wei-Sheng Lai, Chia-Kai Liang, and Jia-Bin Huang. Portrait neural radiance fields from a single image. *arXiv preprint arXiv:2012.05903*, 2020. [2](#)

- [24] Baris Gecer, Jiankang Deng, and Stefanos Zafeiriou. Osted: One-shot texture completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7628–7638, June 2021. [1](#)
- [25] Baris Gecer, Alexander Lattas, Stylianos Ploumpis, Jiankang Deng, Athanasios Papaioannou, Stylianos Moschoglou, and Stefanos Zafeiriou. Synthesizing Coupled 3D Face Modalities by Trunk-Branch Generative Adversarial Networks. In *European Conference on Computer Vision (ECCV)*, 2020. [2](#)
- [26] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [1](#), [2](#), [4](#), [6](#), [7](#)
- [27] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos P Zafeiriou. Fast-ganfit: Generative adversarial network for high fidelity 3d face reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [1](#), [2](#), [7](#)
- [28] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T. Freeman. Unsupervised training for 3d morphable model regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [7](#)
- [29] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. *ACM Transactions on Graphics (TOG)*, 30(6):1–10, 2011. [6](#), [7](#)
- [30] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. [1](#)
- [31] Yuxuan Han, Zhibo Wang, and Feng Xu. Learning a 3d morphable face reflectance model from low-cost data. In *CVPR*, 2023. [2](#)
- [32] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arxiv:2006.11239*, 2020. [1](#), [4](#)
- [33] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. [3](#), [5](#), [6](#), [7](#)
- [34] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. 2022. [2](#)
- [35] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007. [7](#)
- [36] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. [2](#), [5](#)
- [37] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Proc. NeurIPS*, 2020. [1](#)
- [38] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *CoRR*, abs/1812.04948, 2018. [1](#)
- [39] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. [1](#), [5](#)
- [40] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2426–2435, June 2022. [3](#)
- [41] Naveen Kodali, Jacob Abernethy, James Hays, and Zsolt Kira. On convergence and stability of gans. *arXiv preprint arXiv:1705.07215*, 2017. [1](#)
- [42] Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Triantafyllou, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme: Realistically renderable 3d facial reconstruction “in-the-wild”. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [2](#), [5](#)
- [43] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Jiankang Deng, and Stefanos Zafeiriou. FitMe: Deep photorealistic 3D morphable model avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [44] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos P Zafeiriou. Avatarme++: Facial shape and brdf inference with photorealistic rendering-aware gans. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (01):1–1, 2021. [2](#), [4](#), [5](#), [6](#), [7](#)
- [45] Biwen Lei, Jianqiang Ren, Mengyang Feng, Miaomiao Cui, and Xuansong Xie. A hierarchical representation network for accurate and detailed face reconstruction from in-the-wild images, 2023. [2](#)
- [46] Chunlu Li, Andreas Morel-Forster, Thomas Vetter, Bernhard Egger, and Adam Kortylewski. Robust model-based face reconstruction through weakly-supervised outlier segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 372–381, 2023. [1](#)
- [47] Hao Li, Thibaut Weise, and Mark Pauly. Example-based facial rigging. *ACM SIGGRAPH 2010 Papers, SIGGRAPH 2010*, 29(4):32, 2010. [2](#)
- [48] R. Li, K. Bladin, Y. Zhao, C. Chinara, O. Ingraham, P. Xiang, X. Ren, P. Prasad, B. Kishore, J. Xing, and H. Li. Learning formation of physically-based face attributes. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3407–3416, Los Alamitos, CA, USA, jun 2020. IEEE Computer Society. [2](#)
- [49] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. 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. 1, 2, 8
- [50] Zhen Liu, Yao Feng, Michael J. Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling. In *International Conference on Learning Representations*, 2023. 3
- [51] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *ACM Transactions on Graphics*, 37(4):68, 2018. 2
- [52] Huiwen Luo, Koki Nagano, Han-Wei Kung, Qingguo Xu, Zejian Wang, Lingyu Wei, Liwen Hu, and Hao Li. Normalized avatar synthesis using stylegan and perceptual refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11662–11672, June 2021. 2, 6
- [53] Zhaoyang Lyu, Zhifeng Kong, Xudong Xu, Liang Pan, and Dahua Lin. A conditional point diffusion-refinement paradigm for 3d point cloud completion. *arXiv*, abs/2112.03530, 2021. 3
- [54] Zhiyuan Ma, Xiangyu Zhu, Guojun Qi, Zhen Lei, and Lei Zhang. Otavatar: One-shot talking face avatar with controllable tri-plane rendering. *arXiv preprint arXiv:2303.14662*, 2023. 1, 2
- [55] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020. 2
- [56] Shakir Mohamed and Balaji Lakshminarayanan. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016. 5
- [57] Stylianos Moschoglou, Stylianos Ploumpis, Mihalis Nicolaou, Athanasios Papaioannou, and Stefanos Zafeiriou. 3DFaceGAN: Adversarial nets for 3D face representation, generation, and translation. *arXiv preprint arXiv:1905.00307*, 2019. 2
- [58] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 3
- [59] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13503–13513, June 2022. 2
- [60] Athanasios Papaioannou, Baris Gecer, Shiyang Cheng, Grigorios Chrysos, Jiankang Deng, Eftychia Fotiadou, Christos Kampouris, Dimitrios Kollias, Stylianos Moschoglou, Kritaphat Songsri-In, et al. Mimicme: A large scale diverse 4d database for facial expression analysis. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII*, pages 467–484. Springer, 2022. 5, 6
- [61] Foivos Paraperas Papanioniou, Alexandros Lattas, Stylianos Moschoglou, and Stefanos Zafeiriou. Relightify: Relightable 3d faces from a single image via diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 1, 3, 6, 7, 8
- [62] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [63] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2, 4
- [64] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In Mark W. Jones Xianghua Xie and Gary K. L. Tam, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 41.1–41.12. BMVA Press, September 2015. 7
- [65] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3D Face Model for Pose and Illumination Invariant Face Recognition. In *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance*, pages 296–301, Sept. 2009. 1, 2
- [66] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 3
- [67] Alec Radford, Jong Wook Kim, Chris Hallacy, A. Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1
- [68] Aashish Rai, Hiresh Gupta, Ayush Pandey, Francisco Vicente Carrasco, Shingo Jason Takagi, Amaury Aubel, Daeil Kim, Aayush Prakash, and Fernando De la Torre. Towards realistic generative 3d face models. *arXiv preprint arXiv:2304.12483*, 2023. 2, 8
- [69] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *CoRR*, abs/2102.12092, 2021. 3
- [70] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J. Black. Generating 3D faces using convolutional mesh autoencoders. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11207 LNCS:725–741, 2018. 2
- [71] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 4, 5
- [72] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022. 1, 3, 4, 5

- [73] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. 3, 4
- [74] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. 1
- [75] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael Black. Learning to regress 3D face shape and expression from an image without 3D supervision. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 7763–7772, June 2019. 2
- [76] Abhishek Sinha, Jiaming Song, Chenlin Meng, and Stefano Ermon. D2c: Diffusion-decoding models for few-shot conditional generation. *Advances in Neural Information Processing Systems*, 34:12533–12548, 2021. 3
- [77] William A. P. Smith, Alassane Seck, Hannah Dee, Bernard Tiddeman, Joshua Tenenbaum, and Bernhard Egger. A Morphable Face Albedo Model. *arXiv:2004.02711 [cs]*, Apr. 2020. 2, 6, 7
- [78] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *32nd International Conference on Machine Learning, ICML 2015*, 3:2246–2255, 2015. 1, 3
- [79] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv:2010.02502*, October 2020. 5
- [80] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. 3
- [81] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 3
- [82] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. 3, 5
- [83] Fariborz Taherkhani, Aashish Rai, Quankai Gao, Shaunak Srivastava, Xuanbai Chen, Fernando de la Torre, Steven Song, Aayush Prakash, and Daeil Kim. Controllable 3d generative adversarial face model via disentangling shape and appearance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 826–836, January 2023. 1
- [84] Ayush Tewari, Florian Bernard, Pablo Garrido, Gaurav Bharaj, Mohamed Elgharib, Hans-Peter Seidel, Patrick Perez, Michael Zollhofer, and Christian Theobalt. FML: Face Model Learning From Videos. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10804–10814, Long Beach, CA, USA, June 2019. IEEE. 2
- [85] Hoang Thanh-Tung and Truyen Tran. Catastrophic forgetting and mode collapse in gans. In *2020 international joint conference on neural networks (ijcnn)*, pages 1–10. IEEE, 2020. 3
- [86] Justus Thies, Michael Zollhöfer, Matthias Nießner, Levi Valgaerts, Marc Stamminger, and Christian Theobalt. Real-time expression transfer for facial reenactment. *ACM Transactions on Graphics*, 34(6):181–183, 2015. 2
- [87] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gérard Medioni. Regressing robust and discriminative 3D morphable models with a very deep neural network. In *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, volume 2017-January, pages 1493–1502, 2017. 7
- [88] Luan Tran and Xiaoming Liu. On learning 3D face morphable model from in-the-wild images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019. 2
- [89] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In *Neural Information Processing Systems (NeurIPS)*, 2021. 3
- [90] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 4
- [91] Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrušaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, et al. Rodin: A generative model for sculpting 3d digital avatars using diffusion. *arXiv preprint arXiv:2212.06135*, 2022. 1, 3
- [92] Shih En Wei, Jason Saragih, Tomas Simon, Adam W. Harley, Stephen Lombardi, Michal Perdoch, Alexander Hypes, Dawei Wang, Hernan Badino, and Yaser Sheikh. VR facial animation via multiview image translation. *ACM Transactions on Graphics*, 38(4):67, 2019. 2
- [93] Erroll Wood, Tadas Baltrušaitis, Charlie Hewitt, Matthew Johnson, Jingjing Shen, Nikola Milosavljević, Daniel Wilde, Stephan Garbin, Toby Sharp, Ivan Stojiljković, et al. 3d face reconstruction with dense landmarks. In *European Conference on Computer Vision*, pages 160–177. Springer, 2022. 2
- [94] Fei Yang, Dimitri Metaxas, Jue Wang, Eli Shechtman, and Lubomir Bourdev. Expression flow for 3D-Aware face component transfer. *ACM Transactions on Graphics*, 30(4):1–10, 2011. 2
- [95] Tarun Yenamandra, Ayush Tewari, Florian Bernard, Hans-Peter Seidel, Mohamed Elgharib, Daniel Cremers, and Christian Theobalt. i3dmm: Deep implicit 3d morphable model of human heads. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12803–12813, 2021. 2
- [96] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 3

- [97] Longwen Zhang, Qiwei Qiu, Hongyang Lin, Qixuan Zhang, Cheng Shi, Wei Yang, Ye Shi, Sibe Yang, Lan Xu, and Jingyi Yu. Dreamface: Progressive generation of animatable 3d faces under text guidance. *arXiv preprint arXiv:2304.03117*, 2023. [1](#), [3](#)
- [98] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [4](#), [5](#)
- [99] Tianke Zhang, Xuangeng Chu, Yunfei Liu, Lijian Lin, Zhendong Yang, Zhengzhuo Xu, Chengkun Cao, Fei Yu, Changyin Zhou, Chun Yuan, and Yu Li. Accurate 3d face reconstruction with facial component tokens. pages 8999–9008, 10 2023. [2](#)
- [100] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5826–5835, October 2021. [3](#)
- [101] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Towards metrical reconstruction of human faces. In *European Conference on Computer Vision*, pages 250–269. Springer, 2022. [1](#), [2](#), [5](#)