

FitMe: Deep Photorealistic 3D Morphable Model Avatars

Alexandros Lattas Stylianos Moschoglou Stylianos Ploumpis
 Baris Gecer Jiankang Deng Stefanos Zafeiriou

Imperial College London, UK

{a.lattas, s.moschoglou, s.ploumpis, b.gecer, j.deng16, s.zafeiriou}@imperial.ac.uk

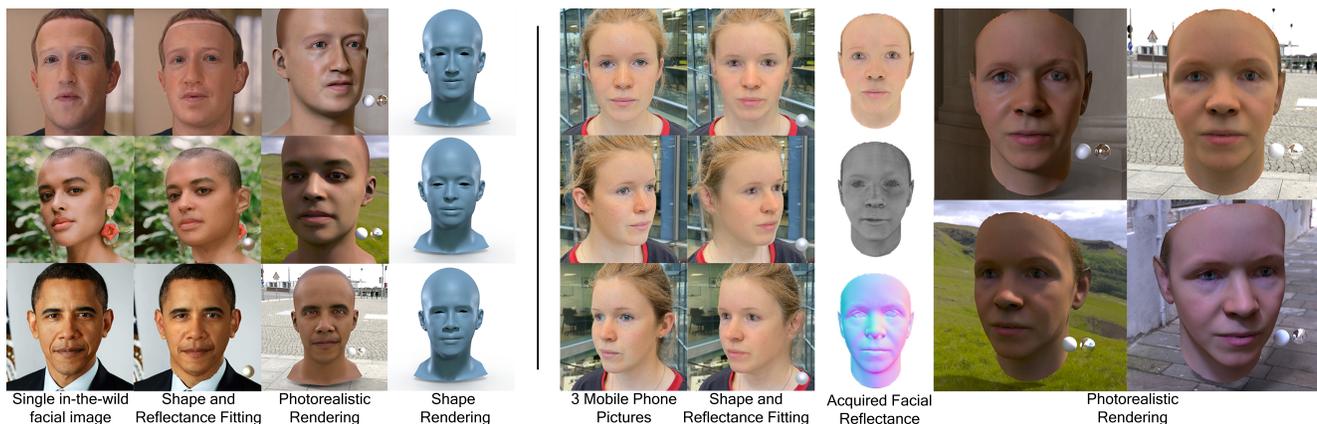


Figure 1. FitMe reconstructs relightable shape and reflectance maps for facial avatars, from a single (left) or multiple (right) unconstrained facial images, using a reflectance model and differentiable rendering. The results can be photorealistically rendered in common engines.

Abstract

In this paper, we introduce FitMe, a facial reflectance model and a differentiable rendering optimization pipeline, that can be used to acquire high-fidelity renderable human avatars from single or multiple images. The model consists of a multi-modal style-based generator, that captures facial appearance in terms of diffuse and specular reflectance, and a PCA-based shape model. We employ a fast differentiable rendering process that can be used in an optimization pipeline, while also achieving photorealistic facial shading. Our optimization process accurately captures both the facial reflectance and shape in high-detail, by exploiting the expressivity of the style-based latent representation and of our shape model. FitMe achieves state-of-the-art reflectance acquisition and identity preservation on single “in-the-wild” facial images, while it produces impressive scan-like results, when given multiple unconstrained facial images pertaining to the same identity. In contrast with recent implicit avatar reconstructions, FitMe requires only one minute and produces relightable mesh and texture-based avatars, that can be used by end-user applications.

1. Introduction

Despite the tremendous steps forward witnessed in the last decade, 3D facial reconstruction from a single unconstrained image remains an important research problem with an active presence in the computer vision community. Its applications are now wide-ranging, including but not limited to human digitization for virtual and augmented reality applications, social media and gaming, synthetic dataset creation, and health applications. However, recent works come short of accurately reconstructing the identity of different subjects and usually fail to produce assets that can be used for photorealistic rendering. This can be attributed to the lack of diverse and big datasets of scanned human geometry and reflectance, the limited and ambiguous information available on a single facial image, and the limitations of the current statistical and machine learning methods.

3D Morphable Models (3DMM) [18] have been a standard method of facial shape and appearance acquisition from a single “in-the-wild” image. The seminal 3DMM work in [7] used Principal Component Analysis (PCA), to model facial shape and appearance with variable identity and expression, learned from about 200 subjects. Since

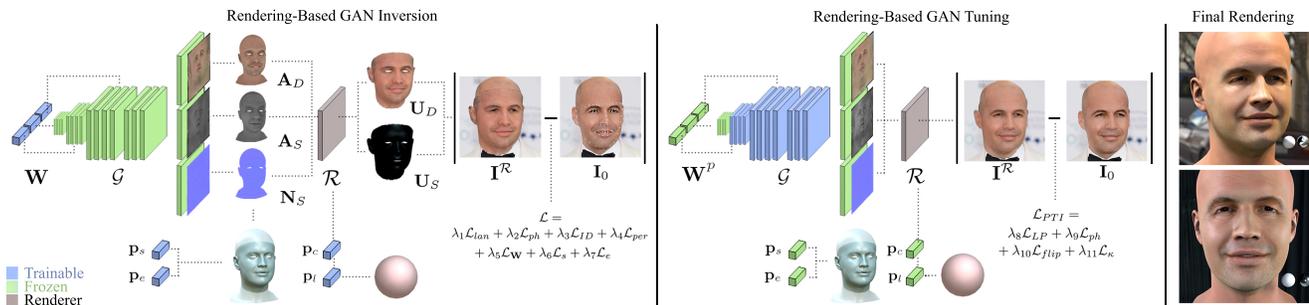


Figure 2. FitMe method overview. For a target image I_0 , we optimize the latent vector \mathbf{W} of the generator \mathcal{G} , and the shape identity \mathbf{p}_s , expression \mathbf{p}_e , camera \mathbf{p}_c and illumination \mathbf{p}_l parameters, by combining 3DMM fitting and GAN inversion methods, through accurate differentiable diffuse \mathbf{U}_D and specular \mathbf{U}_S rendering \mathcal{R} . Then, from the optimized \mathbf{W}^p , we tune the generator \mathcal{G} weights, through the same rendering process. The reconstructed shape \mathbf{S} and facial reflectance (diffuse albedo \mathbf{A}_D , specular albedo \mathbf{A}_S and normals \mathbf{N}_S), achieve great identity similarity and can be directly used in typical renderers, to achieve photorealistic rendering, as shown on the right-hand side.

then, larger models have been introduced, i.e. the LSFM [10], Basel Face Model [52] and Facescape [71], with thousands of subjects. Moreover, recent works have introduced 3DMMs of complete human heads [43, 54, 55] or other facial parts such as ears [54] and tongue [53]. Finally, recent works have introduced extensions ranging from non-linear models [49, 66, 67] to directly regressing 3DMM parameters [61, 68]. However, such models cannot produce textures capable of photorealistic rendering.

During the last decade we have seen considerable improvements in deep generative models. Generative Adversarial Networks (GANs) [30], and specifically progressive GAN architectures [34] have achieved tremendous results in learning distributions of high-resolution 2D images of human faces. Recently, style-based progressive generative networks [35–38] are able to learn meaningful latent spaces, that can be traversed in order to reconstruct and manipulate different attributes of the generated samples. Some methods have also been shown effective in learning a 2D representation of 3D facial attributes, such as UV maps [21, 23, 24, 45].

3D facial meshes generated by 3DMMs can be utilized in rendering functions, in order to create 2D facial images. Differentiating the rendering process is also required in order to perform iterative optimization. Recent advances in differentiable rasterization [44], photorealistic facial shading [41] and rendering libraries [20, 25, 56], enable the photorealistic differentiable rendering of such assets. Unfortunately, 3DMM [10, 23, 45] works rely on the lambertian shading model which comes short of capturing the complexity of facial reflectance. The issue being, photorealistic facial rendering requires various facial reflectance parameters instead of a single RGB texture [41]. Such datasets are scarce, small and difficult to capture [27, 46, 57], despite recent attempts to simplify such setups [39].

Several recent approaches have achieved either high-fidelity facial reconstructions [6, 23, 45] or relightable facial

reflectance reconstructions [16, 17, 19, 40, 41, 65], including infra-red [47], however, high-fidelity and relightable reconstruction still remains elusive. Moreover, powerful models have been shown to capture facial appearance with deep models [22, 42], but they fail to show single or multi image reconstructions. A recent alternative paradigm uses implicit representations to capture avatar appearance and geometry [11, 69], the rendering of which depends on a learned neural rendering. Despite their impressive results, such implicit representations cannot be used by common renderers and are not usually relightable. Finally, the recently introduced Albedo Morphable Model (AlbedoMM) [65] captures facial reflectance and shape with a linear PCA model, but per-vertex color and normal reconstruction is too low-resolution for photorealistic rendering. AvatarMe++ [40, 41] reconstructs high-resolution facial reflectance texture maps from a single “in-the-wild” image, however, its 3-step process (reconstruction, upsampling, reflectance), cannot be optimized directly with the input image.

In this work, we introduce *FitMe*, a fully renderable 3DMM with high-resolution facial reflectance texture maps, which can be fit on unconstrained facial images using accurate differentiable renderings. FitMe achieves identity similarity and high-detailed, fully renderable reconstructions, which are directly usable by off-the-shelf rendering applications. The texture model is designed as a multi-modal style-based progressive generator, which concurrently generates the facial *diffuse-albedo*, *specular-albedo* and *surface-normals*. A meticulously designed branched discriminator enables smooth training with modalities of different statistics. To train the model we create a capture-quality facial reflectance dataset of 5k subjects, by fine-tuning AvatarMe++ on the MimicMe [50] public dataset, which we also augment in order to balance skin-tone representation. For the shape, we use interchangeably a face and head PCA model [54], both trained on large-scale ge-

ometry datasets. We design a single or multi-image fitting method, based on style-based generator projection [35] and 3DMM fitting. To perform efficient iterative fitting (in under 1 minute), the rendering function needs to be differentiable and fast, which makes models such as path tracing unusable. Prior works in the field [10, 12, 23] use simpler shading models (e.g. Lambertian), or much slower optimization [16]. We add a more photorealistic shading than prior work, with plausible diffuse and specular rendering, which can acquire shape and reflectance capable of photorealistic rendering in standard rendering engines (Fig. 1). The flexibility of the generator’s extended latent space and the photorealistic fitting, enables FitMe to reconstruct high-fidelity facial reflectance and achieve impressive identity similarity, while accurately capturing details in diffuse, specular albedo and normals. Overall, in this work we present:

- the first 3DMM capable of generating high-resolution facial reflectance and shape, with an increasing level of detail, that can be photorealistically rendered,
- the first branched multi-modal style-based progressive generator of high-resolution 3D facial assets (diffuse albedo, specular albedo and normals), and a suitable multi-modal branched discriminator,
- a method to acquire and augment a vast facial reflectance dataset of, using assets from a public dataset,
- a multi-modal generator projection, optimized with diffuse and specular differentiable rendering.

2. Related Work

2.1. 3D Morphable Models

Early facial modeling and fitting methods, starting from the seminal 3DMM work of Blanz and Vetter in [7], have always used a linear model for facial shape and appearance [10, 52]. They are thoroughly analyzed in the recent review of Egger et al. [18]. Moreover, 3DMMs have also been extended to the entire head in [54]. Finally, the recent AlbedoMM [65] uses separate PCA models of the diffuse and specular albedo, however the per-vertex albedo is low resolution and inadequate for photorealistic rendering.

Recent works have been replacing parts of the linear shape/texture models with neural networks in order to capture non-linearities. Bagautdinov et al. [5] introduced the first method for non-linear facial shape modeling based on variational autoencoders. 3DFaceGAN [49] introduced a GAN-based approach for facial geometry based on UV maps which captured important non-linearities of the facial shape. Closest to our work, [22, 42] designed models that combines facial albedo, shape, and normals.

2.2. Deep Generative Networks

GANs [30] first achieved high-resolution facial generation with the progressive growing of GANs [34], which introduced a generator-discriminator pair trained on progressively growing resolutions. StyleGAN [37] also introduced a noise-injection technique and a mapping network that learns meaningful latent representations. StyleGAN2 [38] further optimized the architecture and projection method, while StyleGAN2ADA [35] introduced online data augmentation methods, that help when training with limited data, where the discriminator overfits. Finally, [36] improved the signal flow in the generator. In our work, we build upon these powerful models in order to achieve multi-modal reflectance generation.

2.3. Facial Reflectance Acquisition

The first device to accurately acquire 3D facial scans was the LightStage [13], a room-sized dome equipped with programmable illumination and high-end cameras. The diffuse and specular components of the reflectance can be separated by exploiting polarization in multiple captures of gradient illumination [27]. Simplified methods have also been proposed, using unpolarized binary patterns [33] or passive illumination [57]. Although highly accurate, the above methods require large and expensive equipment. A recent practical system made of commodity devices [39] enables significantly faster and cheaper facial reflectance acquisition.

Inverse Rendering approaches have also been successful in acquiring facial reflectance. Recent works employ differentiable ray tracing algorithms [16, 17], to solve an inverse rendering optimization that yields facial diffuse and specular components, as well as specular roughness. However, such methods are computationally expensive and the optimization is susceptible to ambiguity in the subject images.

To overcome such ambiguities, multiple approaches use linear or deep models as priors, or directly regress the facial reflectance. Early approaches proposed deep image-translation models [12, 62, 63], while later methods [32, 60, 70] first acquired realistic facial albedos and displacement textures, using one or more deep neural networks. GANFIT [23] was the first method to use a linear shape model in combination with a GAN-based texture model. Nevertheless, the inferred texture contained baked environment illumination, making photorealistic rendering impossible. AvatarMe [40] introduced a super-resolution and an image-translation network that transformed textures generated from GANFIT into high-resolution facial reflectance. Its extension AvatarMe++ [41] introduced a generalized facial model that works with arbitrary 3DMM fitting algorithms or scanning methods, by using a differentiable shader. However, both methods are separated from the initial fitting method and cannot pick-up all the facial and illumination cues of the input image. An alternative approach [6] produces highly re-

alistic facial shape and albedo, but requires a video stream of RGB and depth data for optimization. Moreover, TRUST [19] propose a method to overcome the ambiguity between skin-tone and illumination optimization, by using cues from the background illumination or multiple faces in the same image. Closest to our work, Luo et al. [45] acquire normalized avatars using a style-based generator together with an iterative optimization. Although similar, our approach has the following advantages: a) our generative model is extended to complete facial reflectance, and our renderer is extended with a more photorealistic shading model, b) our branched generator and discriminator approach generates additional modalities, c) our fitting pipeline introduces a latent space and shading regularization, as well as identity loss, and d) our fine-tuning results in an editable latent space, which can be further optimized through rendering.

3. Method

In this work, we present FitMe, a deep facial reflectance 3DMM, which is based on a branched multi-modal style-based generative network (Sec. 3.1). It is trained on a big capture-quality facial reflectance dataset, which is augmented to balance skin tones (Sec. 3.2). The nature of the data allows accurate diffuse and specular differentiable rendering, using an appropriate shader (Sec. 3.3). Finally, we combine a style-based generator latent space projection with 3DMM fitting methods (Sec. 3.4) and achieve high fidelity facial reconstruction, from a single or multiple images. The acquired shape and reflectance avatars achieve high identity similarity (given the model’s flexibility, diverse dataset and rendering), have expression blendshapes and can be directly used by common rendering applications.

3.1. The FitMe Deep 3D Morphable Model

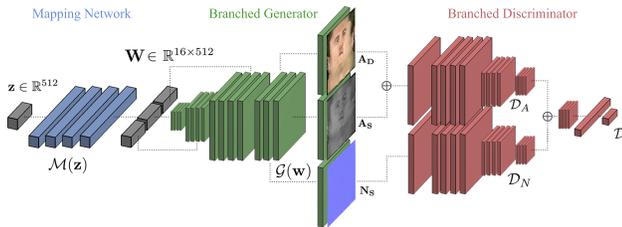


Figure 3. Overview of BRDF-GAN, a style-based generator [38], consisting of a mapping network \mathcal{M} which translates the latent vector \mathbf{z} , to style \mathbf{W} . The branched, multi-modal synthesis network \mathcal{G} generates a diffuse albedo \mathbf{A}_D , specular albedo \mathbf{A}_S and normals \mathbf{N}_S at 1024×1024 resolution. The model is trained in tandem with a branched multi-modal discriminator \mathcal{D} . Each branch enables \mathcal{D} to model the distribution of the albedo and normals separately. The output of each branch is concatenated before entering the last convolutional block and the fully connected layers of \mathcal{D} , in order to maintain feature consistency between modalities.

We introduce BRDF-GAN, a multi-modal style-based generative network, which concurrently generates facial diffuse albedo \mathbf{A}_D , specular albedo \mathbf{A}_S and surface normals \mathbf{N}_S in a UV parameterisation, at 1024×1024 resolution. The model is trained by a novel branched discriminator, that ensures the consistency between the modalities. The shape is modeled with a 3DMM and we model both the facial and the head mesh topology similar to [54].

Generating the shape as another branch of BRDF-GAN, requires many redundant parameters, is prone to quantization, and its fitting is problematic. A recent approach generates 3DMM offsets as a UV texture map, but still requires 3DMM projection to generate the final shape [45]. We opt instead to use a 3DMM with a mesh representation. We use the Universal Head Model [54], with a facial UV topology, which is completed after fitting. For a set of identity \mathbf{p}_s and expression \mathbf{p}_e parameters, with identity and expression bases \mathbf{U}_s and \mathbf{U}_e , respectively, and mean \mathbf{m}_s , the facial geometry $\mathbf{S}(\mathbf{p}_s, \mathbf{p}_e) \in \mathbb{R}^{106317 \times 3}$ can be reconstructed as:

$$\mathbf{S}(\mathbf{p}_s, \mathbf{p}_e) = \mathbf{m}_s + \mathbf{U}_s \mathbf{p}_s + \mathbf{U}_e \mathbf{p}_e \quad (1)$$

Based on StyleGAN2 [38], BRDF-GAN consists of a mapping network $\mathcal{M} : \mathbf{z} \rightarrow \mathbf{W}$, which translates a latent vector $\mathbf{z} \in \mathbb{R}^{512}$ into the latent space $\mathbf{W} \in \mathbb{R}^{16 \times 512}$, a multi-modal branched synthesis network $\mathcal{G} : \mathbf{W} \rightarrow \mathbf{T}$, where $\mathbf{T}_R = \{\mathbf{A}_D, \mathbf{A}_S, \mathbf{N}_S\} \in \mathbb{R}^{7 \times 1024 \times 1024}$. Both the diffuse albedo and normals have 3 channels, however, we use a monochrome specular albedo, which is typical practice in shading [33, 41]. The synthesis network \mathcal{G} follows the skip-connection convolutional blocks architecture of StyleGAN2ADA [35], however, the last block of each up-sampled resolution is branched per reflectance mode, as shown in Fig. 3. For resolutions of 512×512 and up, we find this necessary, in order to achieve satisfactory FID scores. Using a single generator, while branching the final block for each modality, enables the details to be accurately captured between modalities and ensures the final output corresponds to the same identity.

BRDF-GAN is also trained using a branched discriminator \mathcal{D} . We observed that the distributions of the albedos $\mathbf{A}_D, \mathbf{A}_S$ with mean 0.24 and std. 0.12, and the normals \mathbf{N}_S with mean 0.61 and std 0.24, were significantly different, something which inhibited the GAN training using a vanilla \mathcal{D} (i.e., a discriminator with no branches). Therefore, we designed \mathcal{D} as a long-branched network, based on the residual-based discriminator of StyleGAN2-ADA [35]. One branch \mathcal{D}_A receives as input the concatenation of the diffuse and specular albedos $\mathbf{A}_D \oplus \mathbf{A}_S$, which have similar statistics, while the other branch \mathcal{D}_N receives the surface normals \mathbf{N}_S . After running all-but-the-last convolutional blocks, the output of both branches is concatenated and given to the last convolutional block and the fully connected layers of \mathcal{D} . This ensures that \mathcal{D} can accurately cap-

ture the much different distribution of the albedo and normals, while ensuring consistent facial features per generated subject. A detailed visualization is shown in Fig. 2. Luo et al. [45] recently proposed a similar model, where the generator outputs an albedo and a shape, and is trained by three complete and separate discriminators. Our approach, not only requires a smaller network, but is also able to generate facial reflectance modalities of highly different distributions, and, as we demonstrate in the experiments, outperforms their generalization capabilities.

3.2. Dataset Acquisition and Augmentation

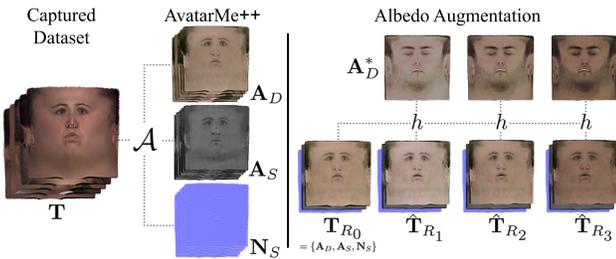


Figure 4. Dataset and augmentation. Left: creation of diffuse \mathbf{A}_D and specular \mathbf{A}_S albedo and normals \mathbf{N}_S , using finetuned a AvatarMe++ [41] network \mathcal{A} on the MimicMe dataset [50] of textures \mathcal{T} . Right: albedo skin-tone augmentation using our masked histogram matching h on different target albedos \mathbf{A}_D^* (Sec. 3.2).

State-of-the-art 3DMMs and GANs require thousands of data samples [10, 35], however no facial reflectance datasets exist at this scale. Instead, we utilize a recently published large dataset of facial captures [50], and a state-of-the-art facial appearance generation model [41], to create such a dataset at capture-like quality and high-resolution. Finally, we alleviate the skin-tone imbalance in the dataset, by augmenting the resulting data with histogram matching [29].

Firstly, we acquire the MimicMe dataset [50], consisting of $\mathcal{T} = \{\mathbf{T}_0, \dots, \mathbf{T}_{n_T}\}$, $n_T = 4,700$ neutral facial textures, which are projected to a common UV map parameterization. These textures include diverse subjects and expressions [50], however, they have baked illumination from the capturing system and do not constitute an albedo. We train an image-to-image translation network $\mathcal{A} : \mathcal{T} \rightarrow \{\mathbf{A}_D, \mathbf{A}_S, \mathbf{N}_S\}$ that transforms textures from \mathcal{T} to reflectance textures $\{\mathbf{A}_D, \mathbf{A}_S, \mathbf{N}_S\}$, following AvatarMe++ [41], while making the following changes: (a) we approximate the capturing environment of [50] using AvatarMe++ [41], and train the AvatarMe++ network \mathcal{A} for this environment only, (b) we generate only the surface (i.e., specular) normals \mathbf{N}_S , in tangent space, which aids the FitMe fitting, as discussed in 3.3 and (c) we do not use a super-resolution network, since the textures in \mathcal{T} are high-resolution captured textures. In this manner, we can use \mathcal{A} to transform \mathcal{T} into a dataset of reflectance textures \mathcal{T}_R .

Regarding the shape, we choose not to represent it as a UV map, and instead we use an existing PCA model. Although we can extract the shape from the above datasets, our experiments show no improvement over using PCA, while also making the network training more challenging. Therefore, with our method being shape-model agnostic, we use the public UHM head shape model [54].

Biased skin tone prediction is a significant problem [19], and therefore balanced skin tone representation in the dataset is of paramount importance. Skin tone augmentation can be done by manipulating the melanin and hemoglobin values of the albedo, which requires expensive data collection [28], pre-training [3] or potentially noisy inverse rendering [4]. Instead, we propose an alternative method based on histogram matching [29] (Overview in Fig. 4), which is trivial to use and noise-free. We acquire 10 scanned albedos [27], with each subject matched to the Monk Skin-Tone Scale (MST) [48] and transform them to our facial UV topology. Then, for each albedo \mathbf{A}_D in the reflectance dataset, we perform histogram matching h [29] with a number of randomly selected target albedos \mathbf{A}_D^* from the MST scale albedos. To avoid non-skin features that could affect the transformation, we calculate a mask \mathbf{M} using the distance of each pixel from the average skin tone $\bar{\mathbf{A}}_D$ in the forehead region $\mathbf{M} = \|\mathbf{A}_D - \bar{\mathbf{A}}_D\|$. Then, we acquire an augmented albedo $\hat{\mathbf{A}}_D$ as:

$$\hat{\mathbf{A}}_D = \mathbf{M} \odot h(\mathbf{A}_D, \mathbf{A}_D^*) + (1 - \mathbf{M}) \odot \mathbf{A}_D \quad (2)$$

3.3. Fast Differentiable Photorealistic Rendering

Typically, 3DMM models rely on a single RGB per-vertex texture color, a pinhole camera projection model and a diffuse-only lambertian shading model, based on spherical harmonics illumination [18]. Although such modeling provides enough cues to the fitting function for the shape and texture parameters to be optimized, the texture model cannot capture the facial reflectance parameters required for photorealistic rendering and the lambertian shading model cannot handle both the diffuse and specular components.

On the contrary, differentiable local-reflection models such as Blinn-Phong [8] can achieve more accurate facial shading [41]. We define a viewing direction \mathbf{v} , an ambient intensity \mathbf{c}_a , and a set of n_l light sources with direction \mathbf{l}_j and intensity \mathbf{c}_j . For a rasterized pixel i with diffuse albedo \mathbf{A}_D , diffuse normals \mathbf{N}_D , specular albedo \mathbf{S}_A and surface normals \mathbf{N}_S , the diffuse \mathbf{U}_{iD} and specular \mathbf{U}_{iS} shading components are defined as:

$$\mathbf{U}_{D_i} = \mathbf{c}_a \mathbf{A}_{D_i} \sum_{j=1}^{n_l} (\mathbf{N}_{D_i} \cdot \mathbf{l}_j) \mathbf{c}_j \quad (3)$$

$$\mathbf{U}_{S_i} = \mathbf{A}_{S_i} \sum_{j=1}^{n_l} (\chi^+(\mathbf{N}_{S_i} \cdot \mathbf{h}_j))^s \mathbf{c}_j, \quad \mathbf{h}_j = \frac{\mathbf{l}_j + \mathbf{v}_j}{\|\mathbf{l}_j + \mathbf{v}_j\|} \quad (4)$$

where $\chi^+(x)$ is a piece-wise function that returns $\max\{0, x\}$, and s is the shininess coefficient. The final shading is acquired by addition $\mathbf{U}_i = \mathbf{U}_{iD} + \mathbf{U}_{iS}$. We also perform subsurface scattering using [41].

For photorealistic shading, separate diffuse and surface normals are required [27, 41, 46], so that the diffuse component approximates the smoother subsurface scattering. In order to correctly optimize the shape mesh during fitting, we use the smoother geometric normals as diffuse normals \mathbf{N}_D . The detailed specular normals are generated by BRDF-GAN in tangent space \mathbf{N}_S , and during rendering are added to the shape normals \mathbf{N} , to acquire the specular normals in object space.

Our differentiable rendering implementation, not only renders skin photorealistically compared to recent fitting methods [23, 45, 65], but it is also implemented directly in image-space, compared to [41], so that direct supervision can be achieved. Moreover, the 3DMM shape \mathbf{S} and GAN-generated reflectance texture maps $\mathbf{T}_R = \{\mathbf{A}_D, \mathbf{A}_S, \mathbf{N}_S\}$ can be efficiently optimized by photometric, identity and perceptual losses. The rendering function \mathcal{R} with optimizable camera \mathbf{p}_c and illumination \mathbf{p}_l can be formulated as:

$$\mathbf{I}_{\mathcal{R}} = \mathcal{R}(\mathbf{S}, \mathbf{T}_R, \mathbf{p}_c, \mathbf{p}_l) \quad (5)$$

3.4. Fitting by Inversion through Rendering

Following conventional 3DMM fitting approaches [7, 10, 18, 23, 59], we build an optimization pipeline with two major improvements: (a) we replace the statistical texture model and its PCA-based optimization, with BRDF-GAN textures and GAN inversion optimization [38, 58], b) we implement an accurate differentiable renderer in the image-space, to enable BRDF-GAN inversion, by rendering its results in high-fidelity. As described below, our optimization pipeline combines the merits of both 3DMM fitting and GAN inversion and achieves photorealistic reflectance reconstruction.

3.4.1 BRDF-GAN Inversion

Generator inversion refers to the acquisition of latent code that recreates the target image and its application on style-based generators is a widely studied problem [38, 58]. However, all the above assume a 2D target image and can directly optimize the generator’s output on it. Moreover, FitMe requires to pass the generated reflectance through rasterization and rendering, while preserving the meaningful properties of the texture map, for which a standard inversion does not suffice. Inspired by both previous 3DMM fitting methods [10, 16, 23, 41], and inversion [35, 58], we design a rendering-based inversion method for our multi-modal generator.

Style-based generators [35, 38] feed a latent code \mathbf{z} into a mapping network, to acquire the native latent space \mathbf{W} ,

which is fed to the synthesis network. Inversion methods optimize \mathbf{W} directly, whilst also optimizing the noise vector \mathbf{n} [35, 58], by using a perceptual LPIPS loss [72]. However, the multi-modal generator faces three challenges: a) being multi-modal makes the concurrent optimization of the different modes more difficult, b) requires rasterization and as such not all pixels are visible to the loss function and c) requires shading, shape, camera and light source optimization, which all remain sensitive to the optimization with LPIPS. We find that an additional set of losses make the inversion more accurate and stable:

The terms of our objective function includes both primitive supervision such as landmark, photometric, and regularization loss functions, as well as identity and perceptual losses. For a target image \mathbf{I}_0 and a rendered reconstruction $\mathbf{I}_{\mathcal{R}}$, we used the following losses to optimize the shape 3DMM, the camera and illumination parameters, and also to guide the rasterized BRDF-GAN results:

Landmark Loss: We estimate 3D landmarks using a deep alignment network [14] $\mathcal{M}(\mathbf{I}) : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{68 \times 2}$ and penalize the distance by $\mathcal{L}_{lan} = \|\mathcal{M}(\mathbf{I}_0) - \mathcal{M}(\mathbf{I}_{\mathcal{R}})\|_2$.

Photometric Loss: It is defined as the distance between per-pixel color intensity values, to capture the skin color and illumination from the target image as $\mathcal{L}_{ph} = \|\mathbf{I}_0 - \mathbf{I}_{\mathcal{R}}\|_1$.

Identity Loss: Following [24, 26], we supervise fitting by identity features extracted by a face recognition network [15] with n layers, $\mathcal{F}^n(\mathbf{I}) : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{512}$ to provide strong identity similarity to the target image:

$$\mathcal{L}_{ID} = 1 - \frac{\mathcal{F}^n(\mathbf{I}_0) \cdot \mathcal{F}^n(\mathbf{I}_{\mathcal{R}})}{\|\mathcal{F}^n(\mathbf{I}_0)\|_2 \cdot \|\mathcal{F}^n(\mathbf{I}_{\mathcal{R}})\|_2} \quad (6)$$

Along with the abstract level supervision provided by identity features, we also optimize for intermediate activations of the face recognition network to ensure mid-level perceptual information can be reconstructed as well:

$$\mathcal{L}_{per} = \sum_j^n \frac{\|\mathcal{F}^j(\mathbf{I}_0) - \mathcal{F}^j(\mathbf{I}_{\mathcal{R}})\|_2}{H_{\mathcal{F}^j} \cdot W_{\mathcal{F}^j} \cdot C_{\mathcal{F}^j}} \quad (7)$$

BRDF-GAN Regularization: Rasterization results in various pixels and facial areas not being used during optimization, while shading may be imperfect in complex illumination environments, both of which introduce noise and environmental features to our albedo and normals during \mathbf{W} ’s optimization. To avoid such artifacts, we follow the \mathbf{W} initialization protocol introduced in [37], and on top we apply an L_2 constraint to make sure \mathbf{W} ’s values do not greatly deviate from their feasible space.

Shape Regularization: We follow the literature [9] in constraining the shape parameters weighted by their inverse eigenvalues: $\mathcal{L}_s = \|\mathbf{p}_s\|_{\Sigma_s^{-1}}^2$ and $\mathcal{L}_e = \|\mathbf{p}_e\|_{\Sigma_e^{-1}}^2$, where Σ denotes a diagonal matrix with the eigenvalues.

Overall loss: It is defined as follows, where all λ ’s are hyper-parameters suitably chosen to balance the different

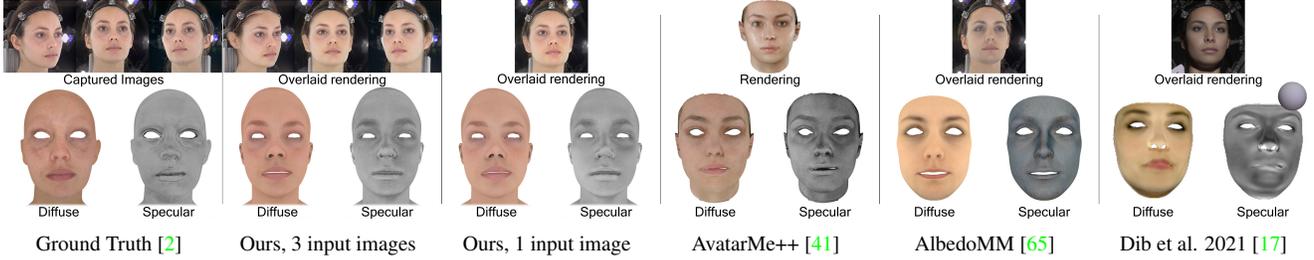


Figure 5. Comparison of diffuse and specular albedo reconstruction and rendering, of Digital Emily [2] with prior works. Both our single and three images reconstruction, achieve similar results to the captured data, which need specialized hardware and hundreds of images.

loss terms during optimization:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{lan} + \lambda_2 \mathcal{L}_{ph} + \lambda_3 \mathcal{L}_{ID} + \lambda_4 \mathcal{L}_{per} + \lambda_5 \mathcal{L}_{\mathbf{W}} + \lambda_6 \mathcal{L}_s + \lambda_7 \mathcal{L}_e, \quad (8)$$

3.4.2 BRDF-GAN Tuning

Recent works show that \mathbf{W} -based inversion cannot fully recover the target image [58]. One approach is to optimize the extended latent space $\mathbf{W} + [1, 45]$, however this inhibits latent code manipulations [58], while, in our experiments, ends up in an outright copy of the target image into the albedo, inhibiting also the re-rendering of the avatar. On the other hand, Pivotal Tuning Inversion (PTI) [58], fine-tunes the generator on the target image, after finding and freezing \mathbf{W} . We find this approach much more accurate in our case, as only the visible texture parts produce gradients and are optimized. Again, given imperfect illumination optimization and shape fitting, PTI transfers high-frequency noise, illumination and non-skin features on the albedo, when tuning the whole generator.

Our experiments show that during fine-tuning, the first layers of BRDF-GAN change the color of the albedo but absorb illumination, the middle layers change the mesostructure of the skin, while the last layers change fine details, but absorb noise from the target image. Therefore, we found that it is optimal to only tune the 4 middle layers, while keeping the rest frozen. We use the proposed LPIPS \mathcal{L}_{LP} [72], using a pre-trained VGG network [64], and photometric \mathcal{L}_{ph} losses. Moreover, we add two regularization losses often used in inverse rendering [17]: i) a horizontal flip loss \mathcal{L}_{flip} , and ii) a chromaticity loss \mathcal{L}_{κ} , in order restrict the fine-tuning to add shadows and highlights, that are not reproduced by the rendering, to the albedo:

$$\mathcal{L}_{PTI} = \lambda_8 \mathcal{L}_{LP} + \lambda_9 \mathcal{L}_{ph} + \lambda_{10} \mathcal{L}_{flip} + \lambda_{11} \mathcal{L}_{\kappa} \quad (9)$$

where λ 's are hyper-params chosen before the fine-tuning.

4. Experiments

4.1. Implementation Details

Our generator code builds on the public repository of StyleGAN2-ADA [35]. The branches follow the same ar-

chitecture, and are connected through concatenation. For the differentiable rendering, we use a blinn-phong shader on Pytorch3D [56] following [41]. Our fitting code requires on average 50 seconds on a machine with one NVIDIA 2080 GPU. We run the fitting for 200 iterations for inversion (Sec. 3.4) and 20 iterations for tuning (Sec. 3.4). Detailed loss parameters are included in the supplemental.

4.2. Single-Image Reconstruction

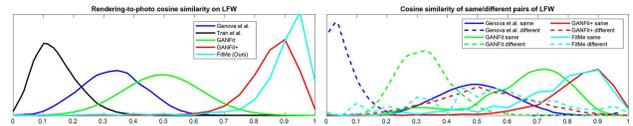


Figure 6. Left: Cosine similarity distributions between real image and rendered reconstruction of LFW [31], using VGG-Face [51]. Right: Cosine similarity distribution between same and different pairs of reconstructions. Compared with [23–25, 68].

We evaluate the identity preservation of our single-image reconstruction, by comparing our results in the LFW dataset [31] with previous literature [23–25, 68]. Using a pre-trained face recognition network [51], we measure the cosine similarity between real image and rendering, and plot the results in Fig. 6, which clearly illustrates the effectiveness of our method. We additionally plot the cosine similarity of between reconstruction pairs of the same identity versus different identity. Moreover, in Fig. 8, we show a qualitative comparison with prior work on single-image reconstruction and in Fig. 9, with single-image reflectance acquisition methods [16, 40, 41, 65]. Finally, in Fig. 7 we show our results on challenging examples. All the above show our method’s ability to preserve the identity of the subject and produce high-quality shape and reflectance.

Additionally, to show the capabilities of the learned latent space, we perform interpolations between different fit-



Figure 7. Results of our method on challenging examples.

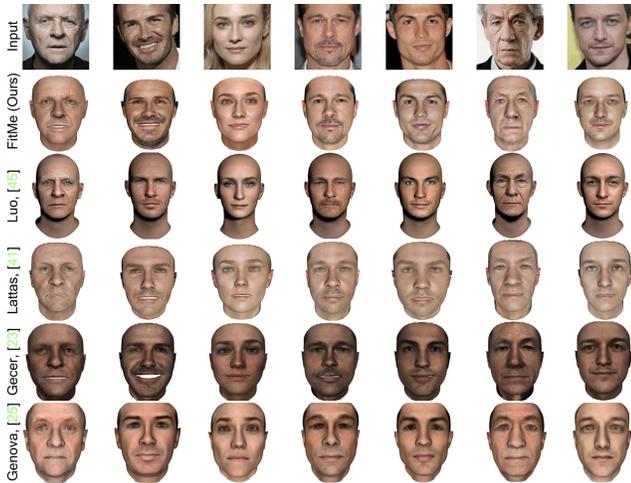


Figure 8. Comparison with single-image reconstruction methods [23, 26, 41, 45] Rows 3,5,6 acquired from [45] benchmark. FitMe achieves higher likeness and acquires reliable reconstructions.

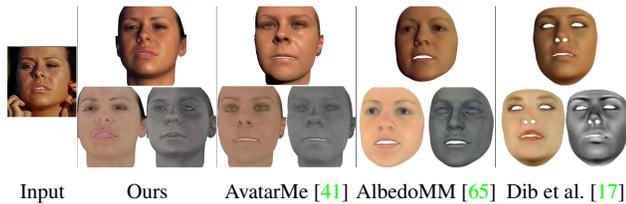


Figure 9. Comparison with single-image reflectance acquisition [17, 41, 65]. Top: rendering, bottom: diffuse and specular albedo.

tings, which we visualize in Fig. 10, by linearly blending their parameters. Our model is able to transition smoothly between fittings, even after generator tuning. In supplemental, we also show PCA-based editing of the latent space.

4.3. Multi-Image Reconstruction

So far, we have been assuming a single target image, however, our method can be used as is for multi-view reconstruction. To do so, we optimize along a batch of N images $\mathbf{I}_{0_i}, i = 0, 1, \dots, N$. For each batch, we separately optimize the camera \mathbf{p}_{e_i} , and illumination \mathbf{p}_{l_i} parameters, while optimizing a single latent vector \mathbf{W} , and shape vector \mathbf{p}_s . The expression vector \mathbf{p}_e can be varied if needed. Moreover, we only use one instance of the generator \mathcal{G} . During optimization, we average the loss across the batch.

We find that a frontal and two side images produce high quality reconstructions that resemble facial capture. Fig. 5 shows a comparison of our 3-image reconstruction and our 1-image reconstruction, with a Light Stage captured Digital Emily [2], and prior work [17, 41, 65]. As can be seen, our method can successfully be used for fast shape and reflectance acquisition from multi-view sets. Finally, Fig. 1 and multiple supplemental figures, show examples of our

reconstruction, using unconstrained mobile phone images, enabling quick avatar generation.



Figure 10. Interpolation of the latent parameters between fittings. For each row, we interpolate the latent \mathbf{W} , and shape parameters $\mathbf{p}_s, \mathbf{p}_e$, between the left-most (A) and the right-most fitting (B).

4.4. Dataset Augmentation

Our dataset augmentation achieves comparable performance to physics-based skin models [3, 28], without requiring the calculation of such complex interactions. Moreover, compared to the more practical method [4], we avoid inverse-rendering based noise and achieve similar results in 2.5 seconds per 1024×1024 albedo map, compared to the 46.1 seconds of [4]. We provide a comparison in the supplemental materials.

5. Limitations

Our method is limited in certain aspects pertaining the data and the fitting process. The data limitations include a) the imbalance of features in the training data, which we attempt to overcome by the proposed augmentation and generator tuning. Moreover, the reflectance data are captured with the assumption of skin-like materials and therefore the eyes exhibit noisy reflections. Finally, there exists ambiguity between skin-tone and illumination intensity during fitting, which could be alleviated by combining our method with the findings of TRUST [19] in future work.

6. Conclusion

In this paper we introduced FitMe, a method that is able to produce highly accurate, renderable human avatars based on a single or multiple “in-the-wild” images. To achieve this, we introduced a deep facial reflectance model that consists of a multi-branched style-based GAN and a PCA-based shape model, as well as an easy skin-tone augmentation method. Moreover, we presented a novel iterative optimization procedure that is based on a differential rendering process. As we showed in a series of experiments, our method helps to bridge the uncanny valley and creates pleasing results, directly renderable in common renderers.

Acknowledgments

AL and SZ were supported by EPSRC Project DEFORM (EP/S010203/1) and SM by an Imperial College DTA.

References

- [1] Rameen Abdal, Peihao Zhu, Niloy J Mitra, and Peter Wonka. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (ToG)*, 40(3):1–21, 2021. **7**
- [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. **7, 8**
- [3] Carlos Aliaga, Christophe Hery, and Mengqi Xia. Estimation of spectral biophysical skin properties from captured rgb albedo. *arXiv preprint arXiv:2201.10695*, 2022. **5, 8**
- [4] Sarah Alotaibi and William AP Smith. A biophysical 3d morphable model of face appearance. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 824–832, 2017. **5, 8**
- [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. **3**
- [6] Linchao Bao, Xiangkai Lin, Yajing Chen, Haoxian Zhang, Sheng Wang, Xuefei Zhe, Di Kang, Haozhi Huang, Xinwei Jiang, Jue Wang, et al. High-fidelity 3d digital human head creation from rgb-d selfies. *ACM Transactions on Graphics (TOG)*, 41(1):1–21, 2021. **2, 3**
- [7] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. **1, 3, 6**
- [8] James F Blinn. Models of light reflection for computer synthesized pictures. In *Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, pages 192–198, 1977. **5**
- [9] James Booth, Epameinondas Antonakos, Stylianos Ploumpis, George Trigeorgis, Yannis Panagakis, and Stefanos Zafeiriou. “3d face morphable models” in-the-wild”. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 48–57, 2017. **6**
- [10] 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 conference on computer vision and pattern recognition*, pages 5543–5552, 2016. **2, 3, 5, 6**
- [11] Chen Cao, Tomas Simon, Jin Kyu Kim, Gabe Schwartz, Michael Zollhoefer, Shun-Suke Saito, Stephen Lombardi, Shih-En Wei, Danielle Belko, Shoou-I Yu, et al. Authentic volumetric avatars from a phone scan. *ACM Transactions on Graphics (TOG)*, 41(4):1–19, 2022. **2**
- [12] Anpei Chen, Zhang Chen, Guli Zhang, Kenny Mitchell, and Jingyi Yu. Photo-realistic facial details synthesis from single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9429–9439, 2019. **3**
- [13] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156, 2000. **3**
- [14] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5203–5212, 2020. **6**
- [15] 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. **6**
- [16] Abdallah Dib, Gaurav Bharaj, Junghyun Ahn, Cédric Thébault, Philippe Gosselin, Marco Romeo, and Louis Chevallier. Practical face reconstruction via differentiable ray tracing. In *Computer Graphics Forum*, volume 40, pages 153–164. Wiley Online Library, 2021. **2, 3, 6, 7**
- [17] Abdallah Dib, Cédric Thébault, 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, 3, 7, 8**
- [18] Bernhard Egger, William AP Smith, Ayush Tewari, Stefanie Wuhler, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, et al. 3d morphable face models—past, present, and future. *ACM Transactions on Graphics (TOG)*, 39(5):1–38, 2020. **1, 3, 5, 6**
- [19] Haiwen Feng, Timo Bolkart, Joachim Tesch, Michael J Black, and Victoria Abrevaya. Towards racially unbiased skin tone estimation via scene disambiguation. In *European Conference on Computer Vision*, pages 72–90. Springer, 2022. **2, 4, 5, 8**
- [20] Clement Fuji Tsang, Maria Shugrina, Jean Francois Lafache, Towaki Takikawa, Jiehan Wang, Charles Loop, Wenzheng Chen, Krishna Murthy Jatavallabhula, Edward Smith, Artem Rozantsev, Or Perel, Tianchang Shen, Jun Gao, Sanja Fidler, Gavriel State, Jason Gorski, Tommy Xi-ang, Jianing Li, Michael Li, and Rev Lebaredian. Kaolin: A pytorch library for accelerating 3d deep learning research. <https://github.com/NVIDIAGameWorks/kaolin>, 2022. **2**
- [21] Baris Gecer, Jiankang Deng, and Stefanos Zafeiriou. Osetec: One-shot texture completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7628–7638, 2021. **2**
- [22] Baris Gecer, Alexandros 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*, pages 415–433. Springer, 2020. **2, 3**
- [23] 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*, pages 1155–1164, 2019. 2, 3, 6, 7, 8
- [24] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. Fast-ganfit: Generative adversarial network for high fidelity 3d face reconstruction. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):4879–4893, 2021. 2, 6, 7
- [25] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T. Freeman. Unsupervised training for 3d morphable model regression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 7
- [26] 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*, pages 8377–8386, 2018. 6, 8
- [27] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, pages 1–10, 2011. 2, 3, 5, 6
- [28] Yuliya Gitlina, Giuseppe Claudio Guarnera, Daljit Singh Dhillon, Jan Hansen, Alexander Lattas, Dinesh Pai, and Abhijeet Ghosh. Practical measurement and reconstruction of spectral skin reflectance. In *Computer graphics forum*, volume 39, pages 75–89. Wiley Online Library, 2020. 5, 8
- [29] Rafael C Gonzalez. *Digital image processing*. Pearson education india, 2009. 5
- [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. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014. 2, 3
- [31] 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
- [32] Loc Huynh, Weikai Chen, Shunsuke Saito, Jun Xing, Koki Nagano, Andrew Jones, Paul Debevec, and Hao Li. Mesoscopic facial geometry inference using deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8407–8416, 2018. 3
- [33] Christos Kampouris, Stefanos Zafeiriou, and Abhijeet Ghosh. Diffuse-specular separation using binary spherical gradient illumination. In *EGSR (EI&I)*, pages 1–10, 2018. 3, 4
- [34] 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, 3
- [35] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33:12104–12114, 2020. 2, 3, 4, 5, 6, 7
- [36] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. 2, 3
- [37] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 2, 3, 6
- [38] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 2, 3, 4, 6
- [39] Alexandros Lattas, Yiming Lin, Jayanth Kannan, Ekin Ozturk, Luca Filipi, Giuseppe Claudio Guarnera, Gaurav Chawla, and Abhijeet Ghosh. Practical and scalable desktop-based high-quality facial capture. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VI*, pages 522–537. Springer, 2022. 2, 3
- [40] 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*, pages 760–769, 2020. 2, 3, 7
- [41] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme++: Facial shape and brdf inference with photorealistic rendering-aware gans. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9269–9284, 2021. 2, 3, 4, 5, 6, 7, 8
- [42] Ruilong Li, Karl Bladin, Yajie Zhao, Chinmay Chinara, Owen Ingraham, Pengda Xiang, Xinglei Ren, Pratusha Prasad, Bipin Kishore, Jun Xing, et al. Learning formation of physically-based face attributes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3410–3419, 2020. 2, 3
- [43] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017. 2
- [44] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7708–7717, 2019. 2
- [45] 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*, pages 11662–11672, 2021. 2, 4, 5, 6, 7, 8
- [46] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, Paul E Debevec, et al. Rapid ac-

- quisition of specular and diffuse normal maps from polarized spherical gradient illumination. *Rendering Techniques*, 2007(9):10, 2007. 2, 6
- [47] Yunqi Miao, Alexandros Lattas, Jiankang Deng, Jungong Han, and Stefanos Zafeiriou. Physically-based face rendering for NIR-VIS face recognition. In *Advances in Neural Information Processing Systems*, 2022. 2
- [48] Ellis Monk. Monk skin-tone (MST) scale. <https://skintone.google/>. Accessed: 2012-10-11. 5
- [49] Stylianos Moschoglou, Stylianos Ploumpis, Mihalis A Nicolaou, Athanasios Papaioannou, and Stefanos Zafeiriou. 3dfacegan: adversarial nets for 3d face representation, generation, and translation. *International Journal of Computer Vision*, 128(10):2534–2551, 2020. 2, 3
- [50] 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. 2, 5
- [51] O Parkhi, A Vedaldi, and A Zisserman. Deep face recognition. pages 1–12. British Machine Vision Association, 2015. 7
- [52] 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. Ieee, 2009. 2, 3
- [53] Stylianos Ploumpis, Stylianos Moschoglou, Vasileios Triantafyllou, and Stefanos Zafeiriou. 3d human tongue reconstruction from single “in-the-wild” images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2771–2780, 2022. 2
- [54] Stylianos Ploumpis, Evangelos Ververas, Eimear O’Sullivan, Stylianos Moschoglou, Haoyang Wang, Nick Pears, William AP Smith, Baris Gecer, and Stefanos Zafeiriou. Towards a complete 3d morphable model of the human head. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):4142–4160, 2020. 2, 3, 4, 5
- [55] Stylianos Ploumpis, Haoyang Wang, Nick Pears, William AP Smith, and Stefanos Zafeiriou. Combining 3d morphable models: A large scale face-and-head model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10934–10943, 2019. 2
- [56] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgios Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 2, 7
- [57] Jérémy Riviere, Paulo Gotardo, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. Single-shot high-quality facial geometry and skin appearance capture. *ACM Transactions on Graphics (TOG)*, 39(4):81–1, 2020. 2, 3
- [58] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on Graphics (TOG)*, 42(1):1–13, 2022. 6, 7
- [59] Sami Romdhani and Thomas Vetter. Estimating 3d shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 2, pages 986–993. IEEE, 2005. 6
- [60] Shunsuke Saito, Lingyu Wei, Liwen Hu, Koki Nagano, and Hao Li. Photorealistic facial texture inference using deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5144–5153, 2017. 3
- [61] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael J Black. Learning to regress 3d face shape and expression from an image without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7763–7772, 2019. 2
- [62] Soumyadip Sengupta, Angjoo Kanazawa, Carlos D Castillo, and David W Jacobs. Sfsnet: Learning shape, reflectance and illuminance of faces in the wild’. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6296–6305, 2018. 3
- [63] Zhixin Shu, Ersin Yumer, Sunil Hadap, Kalyan Sunkavalli, Eli Shechtman, and Dimitris Samaras. Neural face editing with intrinsic image disentangling. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5541–5550, 2017. 3
- [64] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 7
- [65] William AP Smith, Alassane Seck, Hannah Dee, Bernard Tiddeman, Joshua B Tenenbaum, and Bernhard Egger. A morphable face albedo model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5011–5020, 2020. 2, 3, 6, 7, 8
- [66] Luan Tran, Feng Liu, and Xiaoming Liu. Towards high-fidelity nonlinear 3d face morphable model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1126–1135, 2019. 2
- [67] Luan Tran and Xiaoming Liu. Nonlinear 3d face morphable model. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7346–7355, 2018. 2
- [68] 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 of the IEEE conference on computer vision and pattern recognition*, pages 5163–5172, 2017. 2, 7
- [69] Daoye Wang, Prashanth Chandran, Gaspard Zoss, Derek Bradley, and Paulo Gotardo. Morf: Morphable radiance fields for multiview neural head modeling. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–9, 2022. 2
- [70] Shugo Yamaguchi, Shunsuke Saito, Koki Nagano, Yajie Zhao, Weikai Chen, Kyle Olszewski, Shigeo Morishima, and Hao Li. High-fidelity facial reflectance and geometry inference from an unconstrained image. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018. 3
- [71] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: a large-scale

high quality 3d face dataset and detailed riggable 3d face prediction. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition*, pages 601–610, 2020.

2

- [72] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6, 7