Middlesex University Research Repository

An open access repository of
Middlesex University research

<http://eprints.mdx.ac.uk>

Gecer, Baris, Ploumpis, Stylianos, Kotsia, Irene ORCID logoORCID:
<https://orcid.org/0000-0002-3716-010X> and Zafeiriou, Stefanos (2019) GANFIT: Generative adversarial network fitting for high fidelity 3D face reconstruction. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). In: 2019 IEEE/CVF International Conference on Computer Vision and Pattern Recognition, 16-20 Jun 2019, Long Beach, California. e-ISBN 9781728132938, pbk-ISBN 9781728132945. ISSN 1063-6919 [Conference or Workshop Item] (doi:10.1109/CVPR.2019.00125)

Final accepted version (with author's formatting)

This version is available at: <https://eprints.mdx.ac.uk/26523/>

Copyright:

Middlesex University Research Repository makes the University's research available electronically.

Copyright and moral rights to this work are retained by the author and/or other copyright owners unless otherwise stated. The work is supplied on the understanding that any use for commercial gain is strictly forbidden. A copy may be downloaded for personal, non-commercial, research or study without prior permission and without charge.

Works, including theses and research projects, may not be reproduced in any format or medium, or extensive quotations taken from them, or their content changed in any way, without first obtaining permission in writing from the copyright holder(s). They may not be sold or exploited commercially in any format or medium without the prior written permission of the copyright holder(s).

Full bibliographic details must be given when referring to, or quoting from full items including the author's name, the title of the work, publication details where relevant (place, publisher, date), pagination, and for theses or dissertations the awarding institution, the degree type awarded, and the date of the award.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Middlesex University via the following email address:

eprints@mdx.ac.uk

The item will be removed from the repository while any claim is being investigated.

See also repository copyright: re-use policy: <http://eprints.mdx.ac.uk/policies.html#copy>

GANFIT: Generative Adversarial Network Fitting for High Fidelity 3D Face Reconstruction

Baris Gecer^{1,2}, Stylianos Ploumpis^{1,2}, Irene Kotsia³, and Stefanos Zafeiriou^{1,2}

¹Imperial College London

²FaceSoft.io

³University of Middlesex

{b.gecer, s.ploumpis, s.zafeiriou}@imperial.ac.uk, drkotsia@gmail.com

Figure 1: The proposed deep fitting approach can reconstruct high quality texture and geometry from a single image with precise identity recovery. The reconstructions in the figure and the rest of the paper are represented by a vector of size 700 floating points and rendered without any special effects. We would like to highlight that the depicted texture is reconstructed by our model and none of the features taken directly from the image.

Abstract

In the past few years, a lot of work has been done towards reconstructing the 3D facial structure from single images by capitalizing on the power of Deep Convolutional Neural Networks (DCNNs). In the most recent works, differentiable renderers were employed in order to learn the relationship between the facial identity features and the parameters of a 3D morphable model for shape and texture. The texture features either correspond to components of a linear texture space or are learned by auto-encoders directly from in-the-wild images. In all cases, the quality of the facial texture reconstruction of the state-of-the-art methods is still not capable of modeling textures in high fidelity. In this paper, we take a radically different approach and harness the power of Generative Adversarial Networks (GANs) and DCNNs in order to reconstruct the facial texture and shape from single images. That is, we utilize GANs to train a very powerful generator of facial texture in UV space. Then, we revisit the original 3D Morphable Models (3DMMs) fitting approaches making use of non-linear optimization to find

the optimal latent parameters that best reconstruct the test image but under a new perspective. We optimize the parameters with the supervision of pretrained deep identity features through our end-to-end differentiable framework. We demonstrate excellent results in photorealistic and identity preserving 3D face reconstructions and achieve for the first time, to the best of our knowledge, facial texture reconstruction with high-frequency details.¹

1. Introduction

Estimation of the 3D facial surface and other intrinsic components of the face from single images (e.g., albedo, etc.) is a very important problem at the intersection of computer vision and machine learning with countless applications (e.g., face recognition, face editing, virtual reality). It is now twenty years from the seminal work of Blanz and Vetter [4] which showed that it is possible to reconstruct shape and albedo by solving a non-linear optimization

¹Project page: <https://github.com/barisgecer/ganfit>

tion problem that is constrained by linear statistical models of facial texture and shape. This statistical model of texture and shape is called a 3D Morphable Model (3DMM). Arguably the most popular publicly available 3DMM is the Basel model built from 200 people [21]. Recently, large scale statistical models of face and head shape have been made publicly available [7, 10].

For many years 3DMMs and its variants were the methods of choice for 3D face reconstruction [33, 46, 22]. Furthermore, with appropriate statistical texture models on image features such as Scale Invariant Feature Transform (SIFT) and Histogram Of Gradients (HOG), 3DMM-based methodologies can still achieve state-of-the-art performance in 3D shape estimation on images captured under unconstrained conditions [6]. Nevertheless, those methods [6] can reconstruct only the shape and not the facial texture. Another line of research in [45, 34] decouples texture and shape reconstruction. A standard linear 3DMM fitting strategy [41] is used for face reconstruction followed by a number of steps for texture completion and refinement. In these papers [34, 45], the texture looks excellent when rendered under professional renderers (e.g., Arnold), nevertheless when the texture is overlaid on the images the quality significantly drops².

In the past two years, a lot of work has been conducted on how to harness Deep Convolutional Neural Networks (DCNNs) for 3D shape and texture reconstruction. The first such methods either trained regression DCNNs from image to the parameters of a 3DMM [42] or used a 3DMM to synthesize images [30, 18] and formulate an image-to-image translation problem using DCNNs to estimate the depth³ [36]. The more recent unsupervised DCNN-based methods are trained to regress 3DMM parameters from identity features by making use of differentiable image formation architectures [9] and differentiable renderers [16, 40, 31].

The most recent methods such as [39, 43, 14] use both the 3DMM model, as well as additional network structures (called correctives) in order to extend the shape and texture representation. Even though the paper [39] shows that the reconstructed facial texture has indeed more details than a texture estimated from a 3DMM [42, 40], it is still unable to capture high-frequency details in texture and subsequently many identity characteristics (please see the Fig. 4). Furthermore, because the method permits the reconstructions to be outside the 3DMM space, it is susceptible to outliers (e.g., glasses etc.) which are baked in shape and texture. Although rendering networks (*i.e.* trained by VAE [26]) generates outstanding quality textures, each network is capable of storing up to few individuals whom should be placed in a

controlled environment to collect ~ 20 millions of images.

In this paper, we still propose to build upon the success of DCNNs but take a radically different approach for 3D shape and texture reconstruction from a single in-the-wild image. That is, instead of formulating regression methodologies or auto-encoder structures that make use of self-supervision [39, 16, 43], we revisit the optimization-based 3DMM fitting approach by the supervision of deep identity features and by using Generative Adversarial Networks (GANs) as our statistical parametric representation of the facial texture.

In particular, the novelties that this paper brings are:

- We show for the first time, to the best of our knowledge, that a large-scale high-resolution statistical reconstruction of the complete facial surface on an unwrapped UV space can be successfully used for reconstruction of arbitrary facial textures even captured in unconstrained recording conditions⁴.
- We formulate a novel 3DMM fitting strategy which is based on GANs and a differentiable renderer.
- We devise a novel cost function which combines various content losses on deep identity features from a face recognition network.
- We demonstrate excellent facial shape and texture reconstructions in arbitrary recording conditions that are shown to be both photorealistic and identity preserving in qualitative and quantitative experiments.

2. History of 3DMM Fitting

Our methodology naturally extends and generalizes the ideas of texture and shape 3DMM using modern methods for representing texture using GANs, as well as defines loss functions using differentiable renderers and very powerful publicly available face recognition networks [12]. Before we define our cost function, we will briefly outline the history of 3DMM representation and fitting.

2.1. 3DMM representation

The first step is to establish dense correspondences between the training 3D facial meshes and a chosen template with fixed topology in terms of vertices and triangulation.

2.1.1 Texture

Traditionally 3DMMs use a UV map for representing texture. UV maps help us to assign 3D texture data into 2D

²Please see the supplementary materials for a comparison with [34, 45].

³The depth was afterwards refined by fitting a 3DMM and then changing the normals by using image features.

⁴In the very recent works, it was shown that it is feasible to reconstruct the non-visible parts a UV space for facial texture completion[11] and that GANs can be used to generate novel high-resolution faces[38]. Nevertheless, our work is the first one that demonstrates that a GAN can be used as powerful statistical texture prior and reconstruct the complete texture of arbitrary facial images.

2.2. Fitting

3D face and texture reconstruction by fitting a 3DMM is performed by solving a non-linear energy based cost optimization problem that recovers a set of parameters $\mathbf{p} = [\mathbf{p}_{s,e}, \mathbf{p}_t, \mathbf{p}_c, \mathbf{p}_l]$ where $\mathbf{p}_c$ are the parameters related to a camera model and $\mathbf{p}_l$ are the parameters related to an illumination model. The optimization can be formulated as:

$$\min_{\mathbf{p}} \mathcal{E}(\mathbf{p}) = \|\mathbf{I}^0(\mathbf{p}) - \mathbf{W}(\mathbf{p})\|_2^2 + \text{Reg}(\{\mathbf{p}_{s,e}, \mathbf{p}_t\}) \quad (3)$$

where $\mathbf{I}^0$ is the test image to be fitted and $\mathbf{W}$ is a vector produced by a physical image formation process (i.e., rendering) controlled by $\mathbf{p}$. Finally, Reg is the regularization term that is mainly related to texture and shape parameters.

Various methods have been proposed for numerical optimization of the above cost functions [19, 2]. A notable recent approach is [6] which uses handcrafted features (i.e., $\mathbf{H}$) for texture representation simplified the cost function as:

$$\min_{\mathbf{p}^r} \mathcal{E}(\mathbf{p}^r) = \|\mathbf{H}(\mathbf{I}^0(\mathbf{p}^r)) - \mathbf{H}(\mathbf{W}(\mathbf{p}^r))\|_{\mathbf{A}}^2 + \text{Reg}(\mathbf{p}_{s,e}) \quad (4)$$

where $\|\mathbf{a}\|_{\mathbf{A}}^2 = \mathbf{a}^T \mathbf{A} \mathbf{a}$, $\mathbf{A}$ is the orthogonal space to the statistical model of the texture and $\mathbf{p}^r$ is the set of reduced parameters $\mathbf{p}^r = \{\mathbf{p}_{s,e}, \mathbf{p}_c\}$. The optimization problem in Eq. 4 is solved by Gauss-Newton method. The main drawback of this method is that the facial texture is not reconstructed.

In this paper, we generalize the 3DMM fittings and introduce the following novelties:

- We use a GAN on high-resolution UV maps as our statistical representation of the facial texture. That way we can reconstruct textures with high-frequency details.
- Instead of other cost functions used in the literature such as low-level ℓ_1 or ℓ_2 loss (e.g., RGB values [29], edges [33]) or hand-crafted features (e.g., SIFT [6]), we propose a novel cost function that is based on feature loss from the various layers of publicly available face recognition embedding network [12]. Unlike others, deep identity features are very powerful at preserving identity characteristics of the input image.
- We replace physical image formation stage with a differentiable renderer to make use of first order derivatives (i.e., gradient descent). Unlike its alternatives, gradient descent provides computationally cheaper and more reliable derivatives through such deep architectures (i.e., above-mentioned texture GAN and identity DCNN).

3. Approach

We propose an optimization-based 3D face reconstruction approach from a single image that employs a high fidelity texture generation network as statistical prior as illustrated in Fig. 2. To this end, the reconstruction mesh is formed by 3D morphable shape model; textured by the generator network's output UV map; and projected into 2D image by a differentiable renderer. The distance between the rendered image and the input image is minimized in terms of a number of cost functions by updating the latent parameters of 3DMM and the texture network with gradient descent. We mainly formulate these functions based on rich features of face recognition network [12, 35, 28] for smoother convergence and landmark detection network [13] for alignment and rough shape estimation.

The following sections introduce firstly our novel texture model that employs a generator network trained by progressive growing GAN framework. After describing the procedure for image formation with differentiable renderer, we formulate our cost functions and the procedure for fitting our shape and texture models onto a test image.

3.1. GAN Texture Model

Although conventional PCA is powerful enough to build a decent shape and texture model, it is often unable to capture high frequency details and ends up having blurry textures due to its Gaussian nature. This becomes more apparent in texture modelling which is a key component in 3D reconstruction to preserve identity as well as photo-realism.

GANs are shown to be very effective at capturing such details. However, they suffer from preserving 3D coherency [17] of the target distribution when the training images are semi-aligned. We found that a GAN trained with UV representation of real textures with per pixel alignment avoids this problem and is able to generate realistic and coherent UVs from 99.9% of its latent space while at the same time generalizing well to unseen data.

In order to take advantage of this perfect harmony, we train a progressive growing GAN [24] to model distribution of UV representations of 10,000 high resolution textures and use the trained generator network

$$\mathcal{G}(\mathbf{p}_t) : \mathbb{R}^{512} \rightarrow \mathbb{R}^{H \times W \times C} \quad (5)$$

as texture model that replaces 3DMM texture model in Eq. 1.

While fitting with linear models, i.e. 3DMM, is as simple as linear transformation, fitting with a generator network can be formulated as an optimization that minimizes per-pixel Manhattan distance between target texture in UV space $\mathbf{I}_{uv}$ and the network output $\mathcal{G}(\mathbf{p}_t)$ with respect to the latent parameter $\mathbf{p}_t$, i.e. $\min_{\mathbf{p}_t} |\mathcal{G}(\mathbf{p}_t) - \mathbf{I}_{uv}|$.

3.2. Differentiable Renderer

Following [16], we employ a differentiable renderer to project 3D reconstruction into a 2D image plane based on deferred shading model with given camera and illumination parameters. Since color and normal attributes at each vertex are interpolated at the corresponding pixels with barycentric coordinates, gradients can be easily backpropagated through the renderer to the latent parameters.

A 3D textured mesh at the center of Cartesian origin $[0, 0, 0]$ is projected onto 2D image plane by a pinhole camera model with the camera standing at $[x_c, y_c, z_c]$, directed towards $[x'_c, y'_c, z'_c]$ and with the focal length f_c . The illumination is modelled by phong shading given 1) direct light source at 3D coordinates $[x_l, y_l, z_l]$ with color values $[r_l, g_l, b_l]$, and 2) color of ambient lighting $[r_a, g_a, b_a]$.

Finally, we denote the rendered image given geometry $(\mathbf{p}_{s,e})$, texture $(\mathbf{p}_t)$, camera $(\mathbf{p}_c = [x_c, y_c, z_c, x'_c, y'_c, z'_c, f_c])$ and lighting parameters $(\mathbf{p}_l = [x_l, y_l, z_l, r_l, g_l, b_l, r_a, g_a, b_a])$ by the following:

$$\mathbf{I}^{\mathcal{R}} = \mathcal{R}(\mathbf{S}(\mathbf{p}_s, \mathbf{p}_e), \mathcal{P}(\mathcal{G}(\mathbf{p}_t)), \mathbf{p}_c, \mathbf{p}_l) \quad (6)$$

where we construct shape mesh by 3DMM as given in Eq. 2 and texture by GAN generator network as in Eq. 5. Since our differentiable renderer supports only color vectors, we sample from our generated UV map to get vectorized color representation as explained in Sec. 2.1.1.

Additionally, we render a secondary image with random expression, pose and illumination in order to generalize identity related parameters well with those variations. We sample expression parameters from a normal distribution as $\hat{\mathbf{p}}_e \sim \mathcal{N}(\mu = 0, \sigma = 0.5)$ and sample camera and illumination parameters from the Gaussian distribution of 300W-3D dataset as $\hat{\mathbf{p}}_c \sim \mathcal{N}(\hat{\mu}_c, \hat{\sigma}_c)$ and $\hat{\mathbf{p}}_l \sim \mathcal{N}(\hat{\mu}_l, \hat{\sigma}_l)$. This rendered image of the same identity as $\mathbf{I}^{\mathcal{R}}$ (i.e., with same $\mathbf{p}_s$ and $\mathbf{p}_t$ parameters) is expressed by the following:

$$\hat{\mathbf{I}}^{\mathcal{R}} = \mathcal{R}(\mathbf{S}(\mathbf{p}_s, \hat{\mathbf{p}}_e), \mathcal{P}(\mathcal{G}(\mathbf{p}_t)), \hat{\mathbf{p}}_c, \hat{\mathbf{p}}_l) \quad (7)$$

3.3. Cost Functions

Given an input image $\mathbf{I}^0$, we optimize all of the aforementioned parameters simultaneously with gradient descent updates. In each iteration, we simply calculate the forthcoming cost terms for the current state of the 3D reconstruction, and take the derivative of the weighted error with respect to the parameters using backpropagation.

3.3.1 Identity Loss

With the availability of large scale datasets, CNNs have shown incredible performance on many face recognition benchmarks. Their strong identity features are robust to many variations including pose, expression, illumination, age etc. These features are shown to be quite effective at

many other tasks including novel identity synthesizing [15], face normalization [9] and 3D face reconstruction [16]. In our approach, we take advantage of an off-the-shelf state-of-the-art face recognition network [12]⁵ in order to capture identity related features of an input face image and optimize the latent parameters accordingly. More specifically, given a pretrained face recognition network $\mathcal{F}^n(\mathbf{I}) : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{512}$ consisting of n convolutional filters, we calculate the cosine distance between the identity features (i.e., embeddings) of the real target image and our rendered images as following:

$$\mathcal{L}_{id} = 1 - \frac{\mathcal{F}^n(\mathbf{I}^0) \cdot \mathcal{F}^n(\mathbf{I}^{\mathcal{R}})}{\|\mathcal{F}^n(\mathbf{I}^0)\|_2 \|\mathcal{F}^n(\mathbf{I}^{\mathcal{R}})\|_2} \quad (8)$$

We formulate an additional identity loss on the rendered image $\hat{\mathbf{I}}^{\mathcal{R}}$ that is rendered with random pose, expression and lighting. This loss ensures that our reconstruction resembles the target identity under different conditions. We formulate it by replacing $\mathbf{I}^{\mathcal{R}}$ by $\hat{\mathbf{I}}^{\mathcal{R}}$ in Eq. 8 and it is denoted as $\hat{\mathcal{L}}_{id}$.

3.3.2 Content Loss

Face recognition networks are trained to remove all kinds of attributes (e.g. expression, illumination, age, pose) other than abstract identity information throughout the convolutional layers. Despite their strength, the activations in the very last layer discard some of the mid-level features that are useful for 3D reconstruction, e.g. variations that depend on age. Therefore we found it effective to accompany identity loss by leveraging intermediate representations in the face recognition network that are still robust to pixel-level deformations and not too abstract to miss some details. To this end, normalized euclidean distance of intermediate activations, namely content loss, is minimized between input and rendered image with the following loss term:

$$\mathcal{L}_{con} = \sum_j^n \frac{\|\mathcal{F}^j(\mathbf{I}^0) - \mathcal{F}^j(\mathbf{I}^{\mathcal{R}})\|_2}{H_{\mathcal{F}^j} \times W_{\mathcal{F}^j} \times C_{\mathcal{F}^j}} \quad (9)$$

3.3.3 Pixel Loss

While identity and content loss terms optimize albedo of the visible texture, lighting conditions are optimized based on pixel value difference directly. While this cost function is relatively primitive, it is sufficient to optimize lighting parameters such as ambient colors, direction, distance and color of a light source. We found that optimizing illumination parameters jointly with others helped to improve albedo of the recovered texture. Furthermore, pixel loss support identity and content loss with fine-grained texture as it supports highest available resolution while images needs to be

⁵We empirically deduced that other face recognition networks work almost equally well and this choice is orthogonal to the proposed approach.

Figure 3: Example fits of our approach for the images from various datasets. Please note that our fitting approach is robust to occlusion (e.g., glasses), low resolution and black-white in the photos and generalizes well with ethnicity, gender and age. The reconstructed textures are very well at capturing high frequency details of the identities; likewise, the reconstructed geometries from 3DMM are surprisingly good at identity preservation thanks to the identity features used, e.g. crooked nose at bottom-left, dull eyes at bottom-right and chin dimple at top-left

downscaled to 112×112 before identity and content loss. The pixel loss is defined by pixel level ℓ_1 loss function as:

$$\mathcal{L}_{pix} = \|\mathbf{I}^0 - \mathbf{I}^R\|_1 \quad (10)$$

3.3.4 Landmark Loss

The face recognition network $\mathcal{F}$ is pre-trained by the images that are aligned by similarity transformation to a fixed landmark template. To be compatible with the network, we align the input and rendered images under the same settings. However, this process disregards the aspect ratio and scale of the reconstruction. Therefore, we employ a deep face alignment network [13] $\mathcal{M}(\mathbf{I}) : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{68 \times 2}$ to detect landmark locations of the input image and align the rendered geometry onto it by updating the shape, expression and camera parameters. That is, camera parameters are optimized to align with the pose of image $\mathbf{I}$ and geometry parameters are optimized for the rough shape estimation. As a natural consequence, this alignment drastically improves the effectiveness of the pixel and content loss, which are sensitive to misalignment between the two images.

The alignment error is achieved by point-to-point euclidean distances between detected landmark locations of the input image and 2D projection of the 3D reconstruction landmark locations that is available as meta-data of the shape model. Since landmark locations of the reconstruction heavily depend on camera parameters, this loss is great a source of information the alignment of the reconstruction

onto input image and is formulated as following:

$$\mathcal{L}_{lan} = \|\mathcal{M}(\mathbf{I}^0) - \mathcal{M}(\mathbf{I}^R)\|_2 \quad (11)$$

3.4. Model Fitting

We first roughly align our reconstruction to the input image by optimizing shape, expression and camera parameters by: $\min_{\mathbf{p}^r} \mathcal{E}(\mathbf{p}^r) = \lambda_{lan} \mathcal{L}_{lan}$. We then simultaneously optimize all of our parameters with gradient descent and backpropagation so as to minimize weighted combination of above loss terms in the following:

$$\begin{aligned} \min_{\mathbf{p}} \mathcal{E}(\mathbf{p}) = & \lambda_{id} \mathcal{L}_{id} + \hat{\lambda}_{id} \hat{\mathcal{L}}_{id} + \lambda_{con} \mathcal{L}_{con} + \lambda_{pix} \mathcal{L}_{pix} \\ & + \lambda_{lan} \mathcal{L}_{lan} + \lambda_{reg} \text{Reg}(\{\mathbf{p}_{s,e}, \mathbf{p}_l\}) \end{aligned} \quad (12)$$

where we weight each of our loss terms with λ parameters. In order to prevent our shape and expression models and lighting parameters from exaggeration to arbitrarily bias our loss terms, we regularize those parameters by $\text{Reg}(\{\mathbf{p}_{s,e}, \mathbf{p}_l\})$.

Fitting with Multiple Images (i.e. Video): While the proposed approach can fit a 3D reconstruction from a single image, one can take advantage of more images effectively when available, e.g. from a video recording. This often helps to improve reconstruction quality under challenging conditions, e.g. outdoor, low resolution. While state-of-the-art methods follow naive approaches by averaging either the reconstruction [42] or features-to-be-regressed [16]

Figure 4: Comparison of our qualitative results with other state-of-the-art methods in MoFA-Test dataset. Rows 2-5 show comparison with textured geometry and rows 6-8 compare only shapes. The Figure is best viewed in colored and under zoom.

before making a reconstruction, we utilize the power of iterative optimization by averaging identity reconstruction parameters ($\mathbf{p}_s, \mathbf{p}_t$) after every iteration. For an image set $\mathbf{I} = \{\mathbf{I}^0, \mathbf{I}^1, \dots, \mathbf{I}^i, \dots, \mathbf{I}^{n_i}\}$, we reformulate our parameters as $\mathbf{p} = [\mathbf{p}_s, \mathbf{p}_e^i, \mathbf{p}_t, \mathbf{p}_c^i, \mathbf{p}_l^i]$ in which we average shape and texture parameters by the following:

$$\mathbf{p}_s = \sum_i^n \mathbf{p}_s^i, \mathbf{p}_t = \sum_i^n \mathbf{p}_t^i \quad (13)$$

4. Experiments

This section demonstrates the excellent performance of the proposed approach for 3D face reconstruction and shape recovery. We verify this by qualitative results in Figures 1, 3, qualitative comparisons with the state-of-the-art in Sec. 4.2 and quantitative shape reconstruction experiment on a database with ground truth in Sec. 4.3.

4.1. Implementation Details

For all of our experiments, a given face image is aligned to our fixed template using 68 landmark locations detected by an hourglass 2D landmark detection [13]. For the identity features, we employ ArcFace [12] network’s pretrained models. For the generator network $\mathcal{G}$, we train a progressive growing GAN [24] with around 10,000 UV maps from [7] at the resolution of 512×512 . We use the Large Scale Face Model [7] for 3DMM shape model with $n_s = 158$ and the expression model learned from 4DFAB database [8] with $n_e = 29$. During fitting process, we optimize parameters using Adam Solver [25] with 0.01 learning rate. And we set our balancing factors as the following: $\lambda_{id} : 2.0, \hat{\lambda}_{id} : 2.0, \lambda_{con} : 50.0, \lambda_{pix} : 1.0, \lambda_{lan} : 0.001, \lambda_{reg} : \{0.05, 0.01\}$. The Fitting converges in around 30 seconds on an Nvidia GTX 1080 TI GPU for a single image.

Method	Cooperative		Indoor		Outdoor	
	Mean	Std.	Mean	Std.	Mean	Std.
Tran <i>et al.</i> [42]	1.93	0.27	2.02	0.25	1.86	0.23
Booth <i>et al.</i> [6]	1.82	0.29	1.85	0.22	1.63	0.16
Genova <i>et al.</i> [16]	1.50	0.13	1.50	0.11	1.48	0.11
Ours	0.95	0.107	0.94	0.106	0.94	0.106

Table 1: Accuracy results for the meshes on the MICC Dataset using point-to-plane distance. The table reports the mean error (Mean), the standard deviation (Std.).

4.2. Qualitative Comparison to the State-of-the-art

Fig. 4 compares our results with the most recent face reconstruction studies [40, 39, 16, 42, 43] on a subset of MoFA test-set. The first four rows after input images show a comparison of our shape and texture reconstructions to [16, 42, 39] and the last three rows show our reconstructed geometries without texture compared to [39, 43]. All in all, our method outshines all others with its high fidelity photorealistic texture reconstructions. Both of our texture and shape reconstructions manifest strong identity characteristics of the corresponding input images from the thickness and shape of the eyebrows to wrinkles around the mouth and forehead.

4.3. 3D shape recovery on MICC dataset

We evaluate the shape reconstruction performance of our method on MICC Florence 3D Faces dataset (MICC) [1] in Table 1. The dataset provides 3D scans of 53 subjects as well as their short video footages under three difficulty settings: ‘cooperative’, ‘indoor’ and ‘outdoor’. Unlike [16, 42] which processes all the frames in a video, we uniformly sample only 5 frames from each video regardless of their zoom level. And, we run our method with multi-image support for these 5 frames for each video separately as shown in Eq. 13. Each test mesh is cropped at a radius of 95mm around the tip of the nose according to [42] in order to evaluate the shape recovery of the inner facial mesh. We perform dense alignment between each predicted mesh and its corresponding ground truth mesh, by implementing an iterative closest point (ICP) method [3]. As evaluation metric, we follow [16] to measure the error by average symmetric point-to-plane distance.

Table 1 reports the normalized point-to-plane errors in millimeters. It is evident that we have improved the absolute error compared to the other two state-of-the-art methods by 36%. Our results are shown to be consistent across all different settings with minimal standard deviation from the mean error.

4.4. Ablation Study

Fig. 5 shows an ablation study on our method where the full model reconstructs the input face better than its variants,

Figure 5: Contributions of the components or loss terms of the proposed approach with an leave-one-out ablation study.

something that suggests that each of our components significantly contributes towards a good reconstruction. Fig. 5(c) indicates albedo is well disentangled from illumination and our model capture the light direction accurately.

While Fig. 5(d-f) shows each of the identity terms contributes to preserve identity, Fig. 5(h) demonstrates the significance identity features altogether. Still, overall reconstruction utilizes pixel intensities to capture better albedo and illumination as shown in Fig. 5(g). Finally, Fig. 5(i) shows the superiority of our textures over PCA-based ones.

5. Conclusion

In this paper, we revisit optimization-based 3D face reconstruction under a new perspective, that is, we utilize the power of recent machine learning techniques such as GANs and face recognition network as statistical texture model and as energy function respectively.

To the best of our knowledge, this is the first time that GANs are used for model fitting and they have shown excellent results for high quality texture reconstruction. The proposed approach shows identity preserving high fidelity 3D reconstructions in qualitative and quantitative experiments.

Acknowledgements: Baris Gecer is funded by the Turkish Ministry of National Education. Stefanos Zafeiriou acknowledges support by EPSRC Fellowship DEFORM (EP/S010203/1) and a Google Faculty Award.

References

- [1] Andrew D Bagdanov, Alberto Del Bimbo, and Iacopo Masi. The florence 2d/3d hybrid face dataset. In *Proceedings of the 2011 joint ACM workshop on Human gesture and behavior understanding*, pages 79–80. ACM, 2011. 8
- [2] Anil Bas, William AP Smith, Timo Bolkart, and Stefanie Wuhrer. Fitting a 3d morphable model to edges: A comparison between hard and soft correspondences. In *ACCV*, 2016. 4
- [3] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, volume 1611, pages 586–607, 1992. 8
- [4] 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. ACM Press/Addison-Wesley Publishing Co., 1999. 1
- [5] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3d morphable model. *TPAMI*, 25(9):1063–1074, 2003. 3
- [6] James Booth, Epameinondas Antonakos, Stylianos Ploumpis, George Trigeorgis, Yannis Panagakis, Stefanos Zafeiriou, et al. 3d face morphable models in-the-wild. In *CVPR*, 2017. 2, 3, 4, 8
- [7] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3d morphable model learnt from 10,000 faces. In *CVPR*, 2016. 2, 7
- [8] Shiyang Cheng, Irene Kotsia, Maja Pantic, and Stefanos Zafeiriou. 4dfab: a large scale 4d facial expression database for biometric applications. *arXiv preprint arXiv:1712.01443*, 2017. 7
- [9] Forrester Cole, David Belanger, Dilip Krishnan, Aaron Sarna, Inbar Mosseri, and William T Freeman. Synthesizing normalized faces from facial identity features. In *CVPR*, 2017. 2, 5
- [10] Hang Dai, Nick Pears, William Smith, and Christian Duncan. A 3d morphable model of craniofacial shape and texture variation. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [11] Jiankang Deng, Shiyang Cheng, Niannan Xue, Yuxiang Zhou, and Stefanos Zafeiriou. Uv-gan: Adversarial facial uv map completion for pose-invariant face recognition. *CVPR*, 2018. 2
- [12] Jiankang Deng, Jia Guo, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *arXiv preprint arXiv:1801.07698*, 2018. 2, 4, 5, 7
- [13] Jiankang Deng, Yuxiang Zhou, Shiyang Cheng, and Stefanos Zafeiriou. Cascade multi-view hourglass model for robust 3d face alignment. In *Automatic Face & Gesture Recognition (FG)*, pages 399–403. IEEE, 2018. 4, 6, 7
- [14] Pablo Garrido, Michael Zollhöfer, Dan Casas, Levi Valgaerts, Kiran Varanasi, Patrick Pérez, and Christian Theobalt. Reconstruction of personalized 3d face rigs from monocular video. *ACM Transactions on Graphics (TOG)*, 35(3):28, 2016. 2, 3
- [15] Baris Gecer, Binod Bhattarai, Josef Kittler, and Tae-Kyun Kim. Semi-supervised adversarial learning to generate photorealistic face images of new identities from 3d morphable model. *ECCV*, 2018. 5
- [16] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T Freeman. Unsupervised training for 3d morphable model regression. In *CVPR*, 2018. 2, 5, 6, 7, 8, 11
- [17] Ian Goodfellow. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016. 4
- [18] Yudong Guo, Juyong Zhang, Jianfei Cai, Boyi Jiang, and Jianmin Zheng. Cnn-based real-time dense face reconstruction with inverse-rendered photo-realistic face images. *IEEE transactions on pattern analysis and machine intelligence*, 2018. 2
- [19] Guosheng Hu, Fei Yan, Josef Kittler, William Christmas, Chi Ho Chan, Zhenhua Feng, and Patrik Huber. Efficient 3d morphable face model fitting. *Pattern Recognition*, 67:366–379, 2017. 4
- [20] Gary B. Huang, Marwan Mattar, Honglak Lee, and Erik Learned-Miller. Learning to align from scratch. In *NIPS*, 2012. 11
- [21] IEEE. *A 3D Face Model for Pose and Illumination Invariant Face Recognition*, 2009. 2
- [22] Luo Jiang, Juyong Zhang, Bailin Deng, Hao Li, and Ligang Liu. 3d face reconstruction with geometry details from a single image. *IEEE Transactions on Image Processing*, 27(10):4756–4770, 2018. 2
- [23] Ian Jolliffe. Principal component analysis. In *International encyclopedia of statistical science*, pages 1094–1096. Springer, 2011. 3
- [24] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *ICLR*, 2018. 4, 7
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [26] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *ACM Transactions on Graphics (TOG)*, 37(4):68, 2018. 2
- [27] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015. 11
- [28] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, et al. Deep face recognition. In *BMVC*, 2015. 4
- [29] Marcel Pietraschke and Volker Blanz. Automated 3d face reconstruction from multiple images using quality measures. In *CVPR*, 2016. 4
- [30] Elad Richardson, Matan Sela, and Ron Kimmel. 3d face reconstruction by learning from synthetic data. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 460–469. IEEE, 2016. 2
- [31] Elad Richardson, Matan Sela, Roy Or-El, and Ron Kimmel. Learning detailed face reconstruction from a single image. In *CVPR*, 2017. 2
- [32] Sami Romdhani, Volker Blanz, and Thomas Vetter. Face identification by fitting a 3d morphable model using linear shape and texture error functions. In *ECCV*, 2002. 3

- [33] Sami Romdhani and Thomas Vetter. Estimating 3d shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *CVPR*, 2005. 2, 4
- [34] Shunsuke Saito, Lingyu Wei, Liwen Hu, Koki Nagano, and Hao Li. Photorealistic facial texture inference using deep neural networks. In *CVPR*, 2017. 2, 12
- [35] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015. 4
- [36] Matan Sela, Elad Richardson, and Ron Kimmel. Unrestricted facial geometry reconstruction using image-to-image translation. In *ICCV*, 2017. 2
- [37] Zhixin Shu, Ersin Yumer, Sunil Hadap, Kalyan Sunkavalli, Eli Shechtman, and Dimitris Samaras. Neural face editing with intrinsic image disentangling. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5444–5453. IEEE, 2017. 11
- [38] Ron Slossberg, Gil Shamaï, and Ron Kimmel. High quality facial surface and texture synthesis via generative adversarial networks. *ECCVW*, 2018. 2
- [39] Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeonwoo Kim, Patrick Pérez, and Christian Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 hz. 2018. 2, 3, 7, 8
- [40] Ayush Tewari, Michael Zollhöfer, Hyeonwoo Kim, Pablo Garrido, Florian Bernard, Patrick Pérez, and Christian Theobalt. Mofa: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In *ICCV*, 2017. 2, 8
- [41] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *CVPR*, pages 2387–2395, 2016. 2
- [42] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gérard Medioni. Regressing robust and discriminative 3d morphable models with a very deep neural network. In *CVPR*, 2017. 2, 6, 7, 8, 11
- [43] Luan Tran and Xiaoming Liu. Nonlinear 3d face morphable model. In *CVPR*, 2018. 2, 3, 7, 8
- [44] Michael J Wilber, Chen Fang, Hailin Jin, Aaron Hertzmann, John Collomosse, and Serge J Belongie. Bam! the behance artistic media dataset for recognition beyond photography. In *ICCV*, pages 1211–1220, 2017. 11
- [45] Shuco 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):162, 2018. 2, 11
- [46] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z Li. Face alignment across large poses: A 3d solution. In *CVPR*, 2016. 2

Appendix A. Experiments on LFW

In order to evaluate identity preservation capacity of the proposed method, we run two face recognition experiments on Labelled Faces in the Wild (LFW) dataset [20]. Following [16], we feed real LFW images and rendered images of their 3D reconstruction by our method to a pretrained face recognition network, namely VGG-Face[27]. We then compute the activations at the embedding layer and measure cosine similarity between 1) real and rendered images and 2) renderings of same/different pairs.

In Fig. 6 and 7, we have quantitatively showed that our method is better at identity preservation and photorealism (i.e., as the pretrained network is trained by real images) than other state-of-the-art deep 3D face reconstruction approaches [16, 42].

Figure 6: Cosine similarity distributions of rendered and real images LFW based on activations at the embedding layer of VGG-Face network[27]. Our method achieves more than 0.5 similarity on average which [16] has 0.35 average similarity and [42] 0.16 average similarity. Camera and lighting parameters are fixed for all renderings.

Figure 7: Our method successfully preserve identity so that distribution of cosine similarity of same/different pairs is separable by thresholding. Camera and lighting parameters are fixed for all renderings.

Figure 8: Our results on BAM dataset[44] compared to [16]. Our method is robust to many image deformations and even capable of recovering identities from paintings thanks to strong identity features.

Figure 9: Qualitative comparison with [45, 37] by overlaying the reconstructions on the input images. Our method can generate high fidelity texture with accurate shape, camera and illumination fitting.

Appendix B. More Qualitative Results

Figures 8, 9, 10, and 11 illustrate the reconstructions of our method under different settings in comparison to the other state-of-the-art methods. Please see figure captions for detailed explanation.

Figure 10: Qualitative comparison with [34] by means of texture maps, whole and partial face renderings. Please note that while our method does not require any particular renderer for special effects, e.g., lighting, [34] produce these renderings with a commercial renderer called Arnold.

Figure 11: Results under more challenging conditions, *i.e.* strong illuminations, self-occlusions and facial hair. (a) Input image. (b) Estimated fitting overlayed including illumination estimation. (c) Overlayed fitting without illumination. (d) Pixel-wise intensity difference of (b) to (c). (e) Estimated shape mesh