# Hybrid Approach for 3D Head Reconstruction: Using Neural Networks and Visual Geometry

Oussema Bouafif, Bogdan Khomutenko, Mohamed Daoudi

# Hybrid Approach for 3D Head Reconstruction: Using Neural Networks and Visual Geometry

Oussema Bouafif*†, Bogdan Khomutenko* and Mohamed Daoudi†

*MCQ-Scan, Lille, France

†IMT Lille Douai, Univ. Lille, CNRS UMR 9189 CRIStAL, Lille, France

Email: o.bouafif@mcq-scan.com, b.khomutenko@mcq-scan.com, mohamed.daoudi@imt-lille-douai.fr

*Abstract*—**Recovering the 3D geometric structure of a face from a single input image is a challenging active research area in computer vision. In this paper, we present a novel method for reconstructing 3D heads from a single or multiple image(s) using a hybrid approach based on deep learning and geometric techniques. We propose an encoder-decoder network based on the U-net architecture and trained on synthetic data only. It predicts both pixel-wise normal vectors and landmarks maps from a single input photo. Landmarks are used for the pose computation and the initialization of the optimization problem, which, in turn, reconstructs the 3D head geometry by using a parametric morphable model and normal vector fields. State-of-the-art results are achieved through qualitative and quantitative evaluation tests on both single and multi-view settings. Despite the fact that the model was trained only on synthetic data, it successfully recovers 3D geometry and precise poses for real-world images.**

## I. INTRODUCTION

In the last few years, several computer vision applications have used 3D face models. Three-dimensional face model reconstruction encounters different challenges, such as difficult orientations and illumination variations [1]. Several applications such as facial recognition [2], facial animation [3], gender classification [4], and facial expression recognition [5] have achieved better results using 3D face models. There is a number of different techniques to tackle the 3D face reconstruction problem. Some approaches are based on structure from motion, optical flow, or shape from shading [6]. However, these approaches all have limitations since most of them are sensitive to lighting conditions, reflections, shadows, and image quality.

Recently, machine learning has been used to solve these problems [7], [8]. Neural networks are good at implicit modeling of complex variations of lighting, shading, and so on, which leads to a robust geometry estimation. On the other hand, one of the most well-known difficulties in applying neural networks is the lack of 3D face data sets. In some cases, this problem makes end-to-end learning systems less efficient than geometric methods. The use of synthetic data or fitted 3D models from [6] has been proposed in many approaches to meet this need, even if the latter produces only an approximation of the ground truth, which may lead to a low reconstruction accuracy. Another drawback of end-to-end machine-learning-based approaches is the lack of control over the reconstruction process.

Another challenge in 3D face reconstruction is geometry representation. One possible way is a 3D Morphable Model [9] (3DMM). Such models guarantee that the result will be a plausible head reconstruction, and at the same time they are flexible and expressive enough to represent a wide variety of face morphologies. The challenge is to estimate the model parameters. Several approaches to fit morphable models have been proposed [3], [10], [11]. They are usually based either on a set of discrete landmarks [12] or on contours [13]. We argue that pixel-wise fitting (used in some other 3D reconstruction applications [14]) is the best way to exploit as much information present in the image as possible. But the application of this technique to 3D head reconstruction is problematic. In the classical photometric approach, we would require texture and lighting model, which are generally not available.

In this paper, we propose a hybrid approach composed of deep learning and geometric optimization methods (See Fig. 1), capable of reconstructing a 3D head model from one or multiple facial images. The key concept is that surface normals have a high degree of invariance (scale, translation), they don't depend on texture, shadows, and lighting; and they convey rich information about geometry. First, we use an encoder-decoder network that translates a facial input image into a landmarks map ($\mathcal{Z}$) and a facial normal map ($\mathcal{N}$). Then, using these maps in a parametric regression algorithm, we reconstruct the 3D facial model.

The main contributions of this paper are:

- The proposed approach predicts both pixel-wise normal vectors and landmarks maps from a single input photo.
- We present a mesh fitting strategy based on surface normal vectors which address the reconstruction of the head directly from a single or multiple images and aims to recover a complete craniofacial human head as well as its pose.
- We show that our approach achieves state-of-the-art performance on the BU-3DFE [15] data set. We show that the proposed model generalizes well to real-world images, even though it has been trained only on synthetic data.
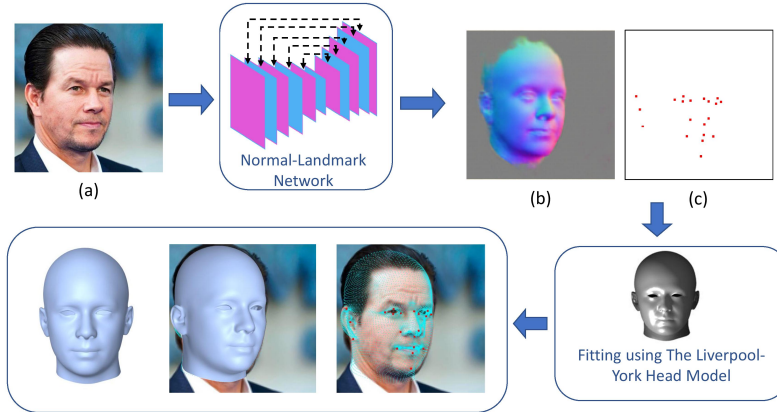
Fig. 1. Overview of our proposed 3D face reconstruction method. Given an input facial image (a), we estimate two different maps (normal surface map $\mathcal{N}$ (b), Landmarks map $\mathcal{Z}$ (c)) used to reconstruct the 3D facial shape via a fitting process with the LYHM [16] morphable model.

## II. RELATED WORK

### A. Morphable Models

The use of 3D morphable models (3DMM) is one of the fundamental approaches proposed to address 3D face reconstruction from a single input image. The most well-known model is perhaps the BFM model [9]. This model is constructed using Principal Component Analysis (PCA) given a large set of 3D faces. By changing the model parameters, one can obtain a variety of 3D facial geometries.

### B. Deep Learning Face Reconstruction Methods

In the last few years, Convolutional Neural Networks (CNNs) have demonstrated better performances with fitting the 3DMM to an input image. For example, fitting via inverse rendering as it was shown in [3], [11], or regressing the 3DMM shape parameters directly from an input image [17].

Other techniques directly map the input photo onto diverse kinds of images which contain information on the geometry, such as [10], which produces a depth map and a 3D correspondence map. After that, a template mesh is deformed in order to fit these produced maps. In [7] Richardson *et al.* proposed CoarseNet and FineNet to provide both coarse and fine details of facial shape. Another method was proposed by [18], which learns 3D face curves from horizontal and vertical epipolar plane images (EPI) of light field images utilizing a densely connected network (FaceLFnet). This model-free approach produced curves that are combined together to obtain a more accurate combined 3D point cloud. An end-to-end 3D face reconstruction method has been proposed in [8], where the problem was divided into two sub-tasks: predicting the neutral 3D face and the expression parameters. In [19] Wu *et al.* proposed a multi-view facial approach based on an end-to-end trainable CNN to regress 3DMM parameters via a differentiable dense optical flow estimator for the alignment and the photometric losses. In [20], a transfer function between an input image and the UV map was learned using a CNN. This UV map is a 2D representation of the 3D shape of the

3DMM. A direct approach was proposed in [21] to reconstruct the 3D face with the use of a volumetric CNN regression.

To be able to use an unlimited amount of data for training, unsupervised learning methods were proposed. A differentiable renderer was used in [22] for the optimization task after learning 3DMM parameters from facial identity encoding. In [23], RingNet was trained using a shape consistency loss similar to the triplet loss to regress all images of the same person to a unique latent form vector. A novel deep 3D face reconstruction approach in [24] uses a hybrid loss function for weakly-supervised learning.

More recently, Wang *et al.* [25] uses an off-the-shelf face recognition neural network trained on high-quality synthetic data to fully reconstruct facial geometries based on a single selfie. Inspired from [26] and based on a GAN model, Face-Normal-Net [27] learned from a synthetic data set to regress a normal map, used in a weighted least square solver to reconstruct a surface approximating the 3D shape. The DF2Net network [28] uses three combined modules, which are each trained on a separate type of data set with different training strategies. The model progressively refines the subtle facial details like small crow's feet and wrinkles. MMFace [29] is a multi-metric regression network composed of two sub-networks: a volumetric sub-network to estimate an intermediate face geometry and a parametric sub-network to infer the corresponding 3DMM parameters.

### C. Normal Map Prediction

Multiple works in computer vision have been proposed to solve tasks such as depth map or surface curvature prediction, semantic segmentation, and edge detection. Among these works, prediction of the normal map has been studied by different learning-based approaches since producing high-quality normal maps for complex objects represents a challenging task. Several methods have been dedicated to the generation of normal maps from (outdoor/indoor) scenes or objects. In [30], a skip-network model has been trained to predict surface normal which is the input to another two networks that estimate the pose and the style of the object

image. In [31], $DeepLiDAR$ has been proposed to predict depth map with a normal surface of outdoor scenes from sparse LiDAR Data and a gray-scale image.

Another category of normal-map-prediction methods has been proposed, where the network takes sketches as input images. In [32], Hudon *et al.*propose a multi-scale representation of the input images to ensure accuracy of produced maps. An interactive method was proposed by Su *et al.* [33] where they proposed a conditional GAN framework based on U-Net model [34]. As we present here, another type of approach has proposed to predict normal maps from facial images to use it in a 3D reconstruction task. In [26], a fully convolutional network trained on various synthetic and real-world data sets. The obtained normal map is then used in the Frankot-Chellappa method [35] to reconstruct 3D facial shape. In [36], a novel architecture SfSNet learns from a mixture of labeled synthetic and unlabeled real images to solve the problem of inverse face rendering. It utilizes residual blocks to disentangle normal and albedo maps into separate subspaces. More recently, Bouafif *et al.* [27] trained a GAN model to generate normal maps from synthetic facial images which are used in a variational 3D reconstruction method. In [37], a new cross-modal learning architecture was proposed to resolve the limited available ground truth data. Core to this approach is a novel module called deactivable skip connections, which allows integrating both the auto-encoded and image-to-normal branches within the same architecture that can be trained end-to-end. This allows learning of a rich latent space that can accurately capture the normal information.

### D. Training With Synthetic Data

Several facial analysis approaches like facial recognition [38], face detection [39] and facial expression analysis [40] have utilized synthetic data to overcome the lack of large-scale training data during the use of neural networks. 3D face reconstruction is among these applications. In [17], [41], a semi-synthetic data set was produced by implementing an optimization-based algorithm with proven precision to a database containing real-world faces. Other approaches [11] have also reinforced this track by integrating a technique of inverse rendering. Another option was brought up by [7], [8], [10], [28], [42], which consists in producing a full synthetic data set. For example, they use the BFM morphable model [9] for identities with Face Warehouse [43] or 3DDFA [44] model for expressions and producing with it, several types of images according to the need of the method. However, the performance of these methods might be limited by lighting models used for rendering, as well as the representativity of the subjects participated in the acquisition to create the model.

To deal with some of these problems, a short time ago, novel approaches have been proposed to boost realism in 3D synthetic face data sets. To achieve that, more productive models and more accurate rendering techniques are suggested. [18] proposed a photo-realistic light field image synthesis method to generate a large-scale Epipolar Plane Images (EPI) data set selecting some examples from the BU-3DFE [15].

In [25], Wang *et al.*proposed an augmented 3D head data set with UV texture and a high-quality engine for rendering. More recently, the LYHM [16] morphable model was used in [27] with hair model [45] and texture to construct a synthetic data head generator that can be used in various approaches related to facial analysis.

## III. PROPOSED METHOD

In this section, we describe the details of our proposed framework as shown in Fig.1. Our 3D reconstruction method takes a facial image as the only input and produces two output maps from an encoder-decoder model. These two maps are a normal surface map $\mathcal{N}$ and a map which contains facial landmarks $\mathcal{Z}$. We use these outputs in a 3D fitting algorithm to recover the 3D Model of the head. In Section III-A, we describe our synthetic 3D head generator which allowed us to produce our training data set. In Section III-B, we show our network architecture which regresses two different maps that are aligned with the input image. The first is a normal surface map ($\mathcal{N}$). The second output is a map that contains facial landmarks ($\mathcal{Z}$). Finally, we explain the details of 3DMM fitting in Section III-C.



Fig. 2. Training data samples. From top to bottom: Synthetic facial images, Normal surface maps $\mathcal{N}$ and Landmarks maps $\mathcal{Z}$.

### A. Synthetic Data Generation

We propose a fully synthetic data set for deep neural network training. To achieve this, we have set up a synthetic data generator based on the one described in [27]. The main source of information for geometric fitting is the normal map $\mathcal{N}$ produced by the face generator (second row in Fig. 2).

To ensure alignment during the adjustment process, we have selected 24 vertices on the head model which are distributed over the eyes, nose, mouth, chin, and ears. After that, the points that are not hidden by the hair or difficult poses are projected, and the landmark map ($\mathcal{Z}$) is generated (third row in Fig. 2). This map is composed of 24 channels, each one containing a positive value at the corresponding landmark projection and zeros everywhere else.

### B. Network Architecture

We propose for the network architecture an encoder-decoder based on the Su *et al.* [33] model which trained a GAN model
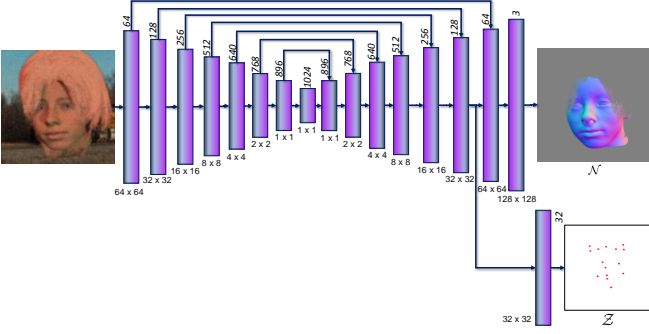
Fig. 3. Details of Normal-Landmark Network. Our encoder-decoder architecture produces two different maps ($\mathcal{N}$ and $\mathcal{Z}$) (shown on the right) from a facial input image (shown on the left). The spatial size and number of layers are shown below and above each block, respectively.

to map a normal surface from a sketch and binary point mask inputs. This model uses the symmetric connection between the encoder and the decoder to reduce the information loss among successive layers, thus implementing the U-Net model [34].

In order to adapt the network to our problem, a certain number of changes have been made. The discriminator has been discarded. For the training data, we give for the input the facial image (RGB), while we have only two generated maps $\mathcal{N}$ (3 channels) and $\mathcal{Z}$ (24 channels). For the loss function, we use a composed pixel-wise loss to measure the difference between the generated maps and the real input ones as described here:

$$\text{Loss} = L_{\mathcal{N}} + L_{\mathcal{Z}}$$
$$L_{\mathcal{N}} = ||\mathcal{N}_{\text{GT}} - \mathcal{N}||_2^2 \quad , \quad L_{\mathcal{Z}} = ||\mathcal{Z}_{\text{GT}} - \mathcal{Z}||_2^2 \tag{1}$$

Where it enforces the rendered images $\mathcal{N}$ and $\mathcal{Z}$ to be similar to the inputs $\mathcal{N}_{\text{GT}}$ and $\mathcal{Z}_{\text{GT}}$ respectively. The architecture of Normal-Landmark Network is illustrated in Fig. 3.

*1) Training Details:* To train our encoder-decoder model, we generate 60,000 facial images scaled to 128 x 128 (equally distributed between males and females) and their corresponding $\mathcal{N}$ and $\mathcal{Z}$ maps. We train the model for about 3000 epochs with a learning rate of $1e-5$, 32 as batch size, and we adopt RMSprop as the optimizer. To make the training images more realistic we add random blur effect and Gaussian noise as data augmentation. When minimizing enery function in (7), we use $\lambda_{\mathcal{N}} = 1$ , $\lambda_{\mathcal{Z}} = 0.8$ and, $\lambda_P = 0.4$.

### C. 3D Morphable Model fitting

The core idea of the present work is the geometry-based fitting of a morphable model to a predicted field of normals, using a camera projection model. This subsection describes in details the problem formulation.

*1) 3D Morphable Model:* In our reconstruction method, we use the LYHM [16] morphable model. By adding parameterized deformations to the mean face model $X_0$, we can generate variability in head identity, as follows:

$$X = X_0 + W\boldsymbol{y} \tag{2}$$

Where $X \in \mathbb{R}^{3 \times N_X}$ is the generated 3D head with $N_X = 11,510$ vertices; $X_0 \in \mathbb{R}^{3 \times N_X}$ is the mean face, computed over the aligned facial 3D scans in the Headspace [46] collection; $\boldsymbol{y} \in \mathbb{R}^{N_y}$ is the identity parameter vector with $N_y = 100$; $W \in \mathbb{R}^{3 \times N_X \times N_y}$ is a tensor representing principal components of the shape model.

*2) Morphable Model's Normals:* In order to fit the normal map $\mathcal{N}$, the normal vector $\boldsymbol{n}_i$ is calculated for each vertex location $\boldsymbol{p}_i \in \mathbb{R}^3$, given the set of adjacent vertices $\{\boldsymbol{q}_{i,1}, \boldsymbol{q}_{i,2}, ..., \boldsymbol{q}_{i,M}\} \subset \mathbb{R}^3$. First an unnormalized version $\tilde{\boldsymbol{n}}$ is computed:

$$\tilde{\boldsymbol{n}}_i = \sum_{j=1}^{N_X} \frac{(\boldsymbol{q}_{i,j} - \boldsymbol{p}_i) \times (\boldsymbol{q}_{i,j+1} - \boldsymbol{p}_i)}{||(\boldsymbol{q}_{i,j} - \boldsymbol{p}_i) \times (\boldsymbol{q}_{i,j+1} - \boldsymbol{p}_i)||} \tag{3}$$

Then the actual normal $\boldsymbol{n}_i$ is obtained by a simple vector normalization:

$$\boldsymbol{n}_i = \frac{\tilde{\boldsymbol{n}}_i}{||\tilde{\boldsymbol{n}}_i||} \tag{4}$$

This is not the most exact method of computing normals, but it is fast, which is important since it is computed at every optimization iteration.

*3) Projection Model:* In order to get the reference normal for each vertex $\boldsymbol{p}$, the latter is projected onto the image plane of $\mathcal{N}$ using the Pinhole camera model. The bi-cubic interpolation is used to get the corresponding reference normal from $\mathcal{N}$ (which is a discrete grid). The projection process can be expressed as follows:

$$\boldsymbol{a} = \frac{1}{z} KR(\boldsymbol{p} + \boldsymbol{t}) \tag{5}$$

Where $\boldsymbol{a} \in \mathbb{R}^2$ is the projected point; $R \in \text{SO}(3)$ is the rotation matrix, parametrized by the three rotation angles of roll, pitch, and yaw, and denoted by $\boldsymbol{r}$; and $\boldsymbol{t} \in \mathbb{R}^3$ is a translation vector. In this context $z$ in the denominator represents the normalization of a 3D point to bring it to the normal plane ($z$ is taken after multiplying the vector by $R$ but before multiplying it by $K$). The projection matrix K contains three parameters $f, u_0, v_0$ and defines the projection model:

$$K = \begin{pmatrix} f & 0 & u_0 \\ 0 & f & v_0 \end{pmatrix} \tag{6}$$

Since the camera calibration is approximative, we think that it is enough to have only three intrinsic camera parameters.

*4) Parametric Regression:* Our implementation follows standard practices (for example, [7], [10], [11], [17]) which use the morphable model in a fitting process. It aims to form an image as close as possible to a target image by finding the most suitable combination of parameters. The target images ($\mathcal{N}$ and $\mathcal{Z}$) are produced from our encoder-decoder network described in Section III-B. The problem can be expressed as the minimization of an energy function that represents the error between the produced maps and those generated by the morphable model. Our energy function contains three main components $E_{\mathcal{N}}$, $E_{\mathcal{Z}}$ and $E_P$:

$$E = \lambda_{\mathcal{N}} E_{\mathcal{N}} + \lambda_{\mathcal{Z}} E_{\mathcal{Z}} + \lambda_P E_P \tag{7}$$

Let us describe each term. $E_{\mathcal{N}}$ represents the difference between interpolated normals $\mathcal{N}(\boldsymbol{a})$ from the 2D projection of the morphable model onto $\mathcal{N}$ and vertex normals $\boldsymbol{n}$ which are computed from the Morphable model using (3). $E_{\mathcal{N}}$ is defined by:

$$E_{\mathcal{N}} = \frac{1}{2}\sum_{i=1}^{N_X}||\mathcal{N}(\boldsymbol{a}_i) - \boldsymbol{n}_i||^2 \qquad (8)$$

The landmarks loss $E_{\mathcal{Z}}$ is defined as the distance between the detected landmarks $\boldsymbol{z}_j$ and the projections $\boldsymbol{b}_j$ of the corresponding subset of 3D vertices using the same projection model described in Eq.5.

$$E_{\mathcal{Z}} = \frac{1}{2}\sum_{j=1}^{N_z}||\boldsymbol{z}_j - \boldsymbol{b}_j||^2 \qquad (9)$$

Finally, the shape prior loss $E_P$ ensures plausibility of reconstructed heads by assuming priors given by the statistical head model represented by singular values $\sigma_k$. In most 3DMM reconstruction methods [9]–[11], [19], [22], [25], this term is used to prevent the degeneration of the reconstructed geometry:

$$E_P = \frac{1}{2}\boldsymbol{y}^T \mathcal{C}_y^{-1} \boldsymbol{y} \qquad (10)$$

$\mathcal{C}_y$ is the covariance matrix, which in this case is a diagonal matrix containing $\sigma_1^2, \sigma_2^2, ..\sigma_{N_y}^2$.

The goal of the parameter regressor is to predict morphable model parameters $y$ and those related to the projection of the model from 3D space to 2D one. This can be formulated as a minimization problem as the following:

$$\boldsymbol{y}^*, \boldsymbol{t}^*, \boldsymbol{r}^*, K^* = \underset{\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{t}, K}{\arg\min} E \qquad (11)$$

where $\boldsymbol{y}$ is the head shape parameter vector, $\boldsymbol{t}$ and $\boldsymbol{r}$ are the translation and rotation parameters defining the head pose with respect to the camera, $K$ is the intrinsic camera matrix. This is a non-linear least-square problem which can be solved efficiently using the Levenberg-Marquardt algorithm. The Ceres solver [47] is used as the optimization back-end.

*5) Head Pre-Alignment:* If we use only normals to fit the head model, we may encounter a certain number of convergence issues. First, it might take some time for the optimizer to converge to the desired position and orientation. Second, The optimization might get stuck in a local minimum. It is a good thing for all local-search optimizers to have a good initial guess. We can achieve it by computing a plausible initial pose. It is done by optimizing $E_{\mathcal{Z}}$ over $\boldsymbol{t}_0, \boldsymbol{r}_0$, while fixing $\boldsymbol{y}$ to 0 and $K$ to a certain prior value, in other words, by using the mean head model and landmarks given by the network. Since it is a purely geometric optimization problem, it almost always converges to the global solution and does it quickly, considering the low number of errors (up to 48 in our case) and unknowns (6 DoF). This approach allows us to make the system robust with respect to the exact head pose and orientation in the image.
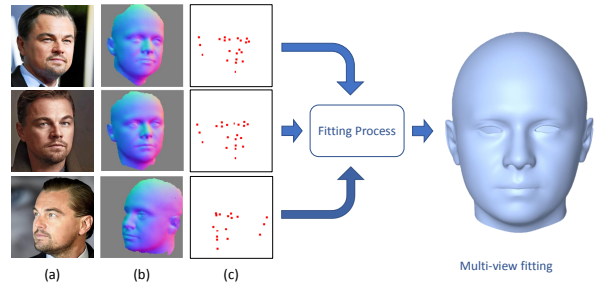
## D. Multi-view 3D Face Reconstruction



Fig. 4. An illustration of our method for multi-view fitting. We estimate the two maps $\mathcal{N}$ (b) and $\mathcal{Z}$ (c) from each input image (a), then we use them in the same fitting process to get a unique 3D head reconstruction.

Our approach can be employed to regress 3DMM parameters from multiple facial images of the same person in different views. Note that our approach can be generalized to any number of input images. We produce $\mathcal{N}_i$ and $\mathcal{Z}_i$ maps for all input images, then we use output maps in the same fitting process.

Poses and camera parameters are computed independently for each image while $\boldsymbol{y}$ is shared for all images. In Fig. 4, we show an example of 3DMM fitting from multiple images of the same person.

## IV. EXPERIMENTS AND RESULTS

In this section, we evaluate the quality of our 3D head reconstruction results obtained using the proposed framework. Firstly, we show the qualitative results in Section IV-A. Section IV-B is dedicated to the ability of the system to reconstruct 3D heads from real-world images by quantitative experiments on a benchmark database with available ground truth, where we compare our results to state-of-the-art methods.

## A. Qualitative Evaluation

In Fig. 5, we show our monocular 3D reconstruction process on the photos of certain celebrities captured from different poses. This evaluation demonstrates how the network learns to represent the head structures of real people. From row 2, we notice our network produces high-quality normal maps as well as accurate landmarks, although it was trained with fully synthetic data. It successfully separates the head from the hair and the background. One can see our method produces high-quality results, which fit the overall structure well. As we use the LYHM [16] morphable model that includes the cranial part, our method allows us to recover the 3D model of the head. The cranial region approximates the areas hidden by the hair whereas we do not estimate the invisible parts of the face skin.

A comparison between the monocular and a multi-view processes using an example of the BU-3DFE [15] data set is illustrated in Fig. 6. We can notice a significant visual resemblance between the 3D reconstruction and the ground truth using the multi-image process—the 3D reconstruction better captures the shape of the face. We notice that for some of the other examples in BU-3DFE, the results are less accurate.
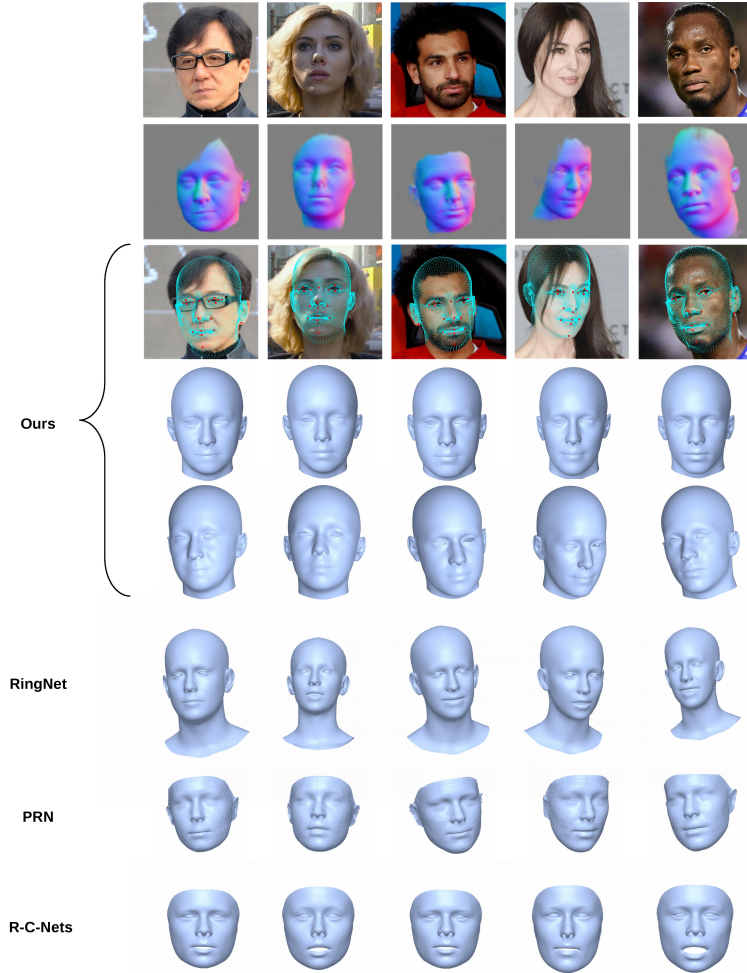
Fig. 5. Examples of visual comparisons from some celebrities facial images with other methods. Rows contain in order; input image, predicted $\mathcal{N}$ map, input image with predicted landmarks (red squares) and dense alignment results (projected vertices of the 3DMM produced by our fitting process), Ours (frontal view), Ours (aligned) RingNet [23], PRN [20] and R-C-Nets [24]

## B. Quantitative Evaluation

For quantitative evaluation, we demonstrate the effectiveness of our approach using the BU-3DFE [15] data set, which contains 100 subjects (56 female, 44 male), with different ages and a variety of ethnic/racial ancestries, and where each 3D model has corresponding 2D images captured under controlled settings. In our evaluation, we use both frontal and profile images of all subjects to evaluate our algorithm's performance with both mono and multi-view reconstruction. Since our 3DMM does not contain expressions, we use only neutral expressions faces. We performed a rough pre-alignment process between the reconstructed model and the ground truth using 6 pre-selected vertices. Next, the alignment and registration processes with the ground truth model are performed using the iterative closest point (ICP) solver, after that, we compute the point-to-plane distances and the absolute depth errors between the reconstructed 3D models and the ground-truth 3D meshes. We eliminate examples when the alignment process fails (which is about 3-4%).

We use the point-to-plane Root Mean Square Error (RMSE) as the performance measure. We argue that it is a good way of evaluating performance, which is not sensitive to either the number of vertices in the model or the way of how they are distributed. It is computed as follows. Once the models are aligned, for each vertex $\boldsymbol{p}_i$ of the reference model, we compute the nearest neighbor $\boldsymbol{q}_i$ from the reconstructed model to be evaluated, as well as normal vectors $\boldsymbol{n}_i$ of the reference model. Second, we compute the error projected to the normal, and compute the RMSE:

$$\varepsilon = \sqrt{\frac{\sum_{i=1}^{N} \left((\boldsymbol{p}_i - \boldsymbol{q}_i) \cdot \boldsymbol{n}_i\right)^2}{N}} \qquad (12)$$

where $N$ is the number of vertices of the reference model. Once $\varepsilon$ is computed for each subject, we can find the mean and standard deviation across the benchmark data set.

*1) Results and Discussion:* In Table I, we report the numeric results in the form of $\mathrm{mean} \pm \mathrm{std}$ for the point-to-plane (RMSE) and for the absolute depth errors, we report mean ($\mu$),

standard deviation ($\sigma$), median ($\widetilde{m}$), and the average ninety percent most significant error ($\delta_{90\%}$). We have compared our system in both monocular (Mono) and multi-view (Multi) settings to the state-of-the-art systems with code available online [20], [23], [24]. We used the exact same procedure for all the tested systems to make sure that all of them are in the same conditions. We can see that the mean error of the
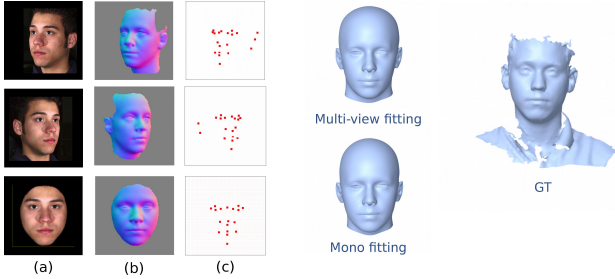


Fig. 6. An illustration of stereo and mono fitting from BU-3DFE [15] data set example. Input images (a), normals surface map $\mathcal{N}$ (b), landmarks map $\mathcal{Z}$ (c). (Multi-view fitting): 3D head reconstruction using all images in the same fitting process. (Mono fitting) : 3D head fitted exploiting only the frontal image in the fitting process (third row). (GT) : the ground-truth of the 3D head mesh.

multi-view approach is lower than that which we have with the monocular configuration. In most cases, the morphology of the entire head is better captured using several images. In some other cases, this multi-view approach gives less good result. And since the evaluation is only done on the facial part, it does not contribute to the precision.

One can see that the proposed method is performing at the level of the state of the art, being clearly outperformed only by the R-C-Nets. The precision obtained by us for the latter is somehow worse than the one reported in the corresponding publication [24], which is $1.40 \pm 0.31$ for the point-to-plane (RMSE) and which can be due to some differences in the evaluation technique.

TABLE I
QUANTITATIVE COMPARISON ON THE BU-3DFE [15] DATA SET. LOWERS ARE BETTER

| Method | Ours (Mono) | Ours (Multi) | RingNet [23] | PRN [20] | R-C-Nets [24] |
|---|---|---|---|---|---|
| RMSE | $1.74 \pm 0.44$ | $1.67 \pm 0.43$ | $1.90 \pm 0.49$ | $1.86 \pm 0.47$ | $\mathbf{1.60 \pm 0.41}$ |
| $\mu$ | 2.21 | 2.17 | 3.42 | 1.83 | **1.64** |
| $\sigma$ | 1.08 | **1.04** | 1.58 | 1.70 | 1.69 |
| $\widetilde{m}$ | 2.09 | 2.05 | 3.23 | 1.43 | **1.27** |
| $\delta_{90\%}$ | 3.61 | 3.53 | 5.60 | 3.46 | **3.00** |

## V. CONCLUSIONS

In this paper, we presented a novel approach composed of both deep learning and visual-geometry-based methods for estimating the complete 3D human head shape from a single or multiple images. Our method employs an encode-decoder network that maps the input image to a rasterized normal map $\mathcal{N}$ and a landmark map $\mathcal{Z}$. These maps are then used in a fitting process to regress the 3DMM parameters of the face

identity from the LYHM model. To our knowledge, pixel-wise fitting is generally based on photometric information and not on output from a neural network, we strongly believe this process is intrinsically better than integration-based surface-reconstruction methods.

Our network has been trained only on a synthetic facial data set. The network has shown good results in terms of both accuracy and generalization on real-world images. In addition to the 3D reconstruction of the head, the landmark map $\mathcal{Z}$ can be directly used for face tracking and pose estimation, an essential part of our reconstruction pipeline. The landmarks are used to find a good initial guess for the pose before the fine-grained fitting. It improves the convergence rate of the regression process and minimizes chances of reaching a local minimum.

We have performed quantitative and qualitative experiments to evaluate our pipeline performance. We demonstrate that our proposed framework achieves state-of-the-art performance in 3D face reconstruction for both single and multi-view settings. Overall, the multi-view setting gives promising results but further work is required to take full advantage of it.

Despite the robust performance in many cases, our method has some limitations. The used 3DMM does not include facial expressions and has a limited age range. This is why it is difficult to reconstruct the finest details in 3DMM because the accuracy of the recovered geometry is limited to the flexibility of this model. This limitation is not fundamental to our proposed method. It can be overcome by adopting a more expressive morphable model for the synthetic data generator and for the fitting process. Another limitation is that synthetic data can have unrealistic features, which in turn can introduce certain biases in the learning process. The use of GAN architectures in combination with classic 3D rendering may improve the realism of generated photos.

## REFERENCES

[1] A. F. Abate, M. Nappi, D. Riccio, and G. Sabatino, "2D and 3D face recognition: A survey," *Pattern Recognition Letters*, vol. 28, no. 14, pp. 1885–1906, 2007. [Online]. Available: https://doi.org/10.1016/j.patrec.2006.12.018

[2] H. Drira, B. Ben Amor, A. Srivastava, M. Daoudi, and R. Slama, "3D face recognition under expressions, occlusions, and pose variations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 9, pp. 2270–2283, 2013. [Online]. Available: https://doi.org/10.1109/TPAMI.2013.48

[3] J. Thies, M. Zollhofer, M. Stamminger, C. Theobalt, and M. Nießner, "Face2face: Real-time face capture and reenactment of rgb videos," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2387–2395.

[4] B. Xia, B. Ben Amor, H. Drira, M. Daoudi, and L. Ballihi, "Combining face averageness and symmetry for 3D-based gender classification," *Pattern Recognit.*, vol. 48, no. 3, pp. 746–758, 2015. [Online]. Available: https://doi.org/10.1016/j.patcog.2014.09.021

[5] B. Ben Amor, H. Drira, S. Berretti, M. Daoudi, and A. Srivastava, "4-D facial expression recognition by learning geometric deformations," *IEEE Trans. Cybernetics*, vol. 44, no. 12, pp. 2443–2457, 2014. [Online]. Available: https://doi.org/10.1109/TCYB.2014.2308091

[6] I. Kemelmacher-Shlizerman and R. Basri, "3D face reconstruction from a single image using a single reference face shape," *IEEE transactions on pattern analysis and machine intelligence*, vol. 33, no. 2, pp. 394–405, 2010.

[7] E. Richardson, M. Sela, R. Or-El, and R. Kimmel, "Learning detailed face reconstruction from a single image," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1259–1268.

[8] P. Dou, S. K. Shah, and I. A. Kakadiaris, "End-to-end 3D face reconstruction with deep neural networks," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 21–26.

[9] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3D faces," in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 1999, pp. 187–194.

[10] M. Sela, E. Richardson, and R. Kimmel, "Unrestricted facial geometry reconstruction using image-to-image translation," *arxiv*, 2017.

[11] Y. Guo, J. Cai, B. Jiang, J. Zheng *et al.*, "Cnn-based real-time dense face reconstruction with inverse-rendered photo-realistic face images," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 6, pp. 1294–1307, 2018.

[12] Y. Yanga, X.-J. Wu, and J. Kittler, "Landmark weighting for 3dmm shape fitting," *arXiv preprint arXiv:1808.05399*, 2018.

[13] A. Bas, W. A. Smith, T. Bolkart, and S. Wuhrer, "Fitting a 3d morphable model to edges: A comparison between hard and soft correspondences," in *Asian Conference on Computer Vision*. Springer, 2016, pp. 377–391.

[14] J. Engel, T. Schöps, and D. Cremers, "Lsd-slam: Large-scale direct monocular slam," in *European conference on computer vision*. Springer, 2014, pp. 834–849.

[15] L. Yin, X. Wei, Y. Sun, J. Wang, and M. J. Rosato, "A 3D facial expression database for facial behavior research," in *7th international conference on automatic face and gesture recognition (FGR06)*. IEEE, 2006, pp. 211–216.

[16] H. Dai, N. Pears, W. A. P. Smith, and C. Duncan, "A 3D morphable model of craniofacial shape and texture variation," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.

[17] A. Tuan Tran, T. Hassner, I. Masi, and G. 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*, 2017, pp. 5163–5172.

[18] M. Feng, S. Zulqarnain Gilani, Y. Wang, and A. Mian, "3D face reconstruction from light field images: A model-free approach," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 501–518.

[19] F. Wu, L. Bao, Y. Chen, Y. Ling, Y. Song, S. Li, K. N. Ngan, and W. Liu, "Mvf-net: Multi-view 3D face morphable model regression," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 959–968.

[20] Y. Feng, F. Wu, X. Shao, Y. Wang, and X. Zhou, "Joint 3d face reconstruction and dense alignment with position map regression network," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 534–551.

[21] A. S. Jackson, A. Bulat, V. Argyriou, and G. Tzimiropoulos, "Large pose 3d face reconstruction from a single image via direct volumetric cnn regression," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1031–1039.

[22] K. Genova, F. Cole, A. Maschinot, A. Sarna, D. Vlasic, and W. T. Freeman, "Unsupervised training for 3D morphable model regression," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8377–8386.

[23] S. Sanyal, T. Bolkart, H. Feng, and M. J. Black, "Learning to regress 3D face shape and expression from an image without 3D supervision," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7763–7772.

[24] Y. Deng, J. Yang, S. Xu, D. Chen, Y. Jia, and X. Tong, "Accurate 3D face reconstruction with weakly-supervised learning: From single image to image set," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.

[25] R. Wang, C.-F. Chen, H. Peng, X. Liu, O. Liu, and X. Li, "Digital twin: Acquiring high-fidelity 3D avatar from a single image," *arXiv preprint arXiv:1912.03455*, 2019.

[26] G. Trigeorgis, P. Snape, S. Zafeiriou, and I. Kokkinos, "Normal Estimation For "in-the-wild" Faces Using Fully Convolutional Networks," in *Computer Vision and Pattern Recognition (CVPR)*, 2017.

[27] O. Bouafif, B. Khomutenko, and M. Daoudi, "Monocular 3D head reconstruction via prediction and integration of normal vector field," in *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications-Volume 5: VISAPP, Valletta, Malta*, 2020, pp. 359–369.

[28] X. Zeng, X. Peng, and Y. Qiao, "Df2net: A dense-fine-finer network for detailed 3D face reconstruction," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 2315–2324.

[29] H. Yi, C. Li, Q. Cao, X. Shen, S. Li, G. Wang, and Y.-W. Tai, "Mmface: A multi-metric regression network for unconstrained face reconstruction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7663–7672.

[30] A. Bansal, B. Russell, and A. Gupta, "Marr revisited: 2D-3D alignment via surface normal prediction," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5965–5974.

[31] J. Qiu, Z. Cui, Y. Zhang, X. Zhang, S. Liu, B. Zeng, and M. Pollefeys, "Deeplidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3313–3322.

[32] M. Hudon, M. Grogan, R. Pagés, and A. Smolić, "Deep normal estimation for automatic shading of hand-drawn characters," in *European Conference on Computer Vision*. Springer, 2018, pp. 246–262.

[33] W. Su, D. Du, X. Yang, S. Zhou, and H. Fu, "Interactive sketch-based normal map generation with deep neural networks," *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 1, no. 1, p. 22, 2018.

[34] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.

[35] R. T. Frankot and R. Chellappa, "A method for enforcing integrability in shape from shading algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 10, no. 4, pp. 439–451, 1988. [Online]. Available: https://doi.org/10.1109/34.3909

[36] S. Sengupta, A. Kanazawa, C. D. Castillo, and D. W. Jacobs, "Sfsnet: Learning shape, reflectance and illuminance of facesin the wild'," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6296–6305.

[37] V. F. Abrevaya, A. Boukhayma, P. H. Torr, and E. Boyer, "Cross-modal deep face normals with deactivable skip connections," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4979–4989.

[38] A. Kortylewski, A. Schneider, T. Gerig, B. Egger, A. Morel-Forster, and T. Vetter, "Training deep face recognition systems with synthetic data," *arXiv preprint arXiv:1802.05891*, 2018.

[39] J. Han, S. Karaoglu, H.-A. Le, and T. Gevers, "Improving face detection performance with 3D-rendered synthetic data," *arXiv preprint arXiv:1812.07363*, 2018.

[40] I. Abbasnejad, S. Sridharan, D. Nguyen, S. Denman, C. Fookes, and S. Lucey, "Using synthetic data to improve facial expression analysis with 3D convolutional networks," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2017, pp. 1609–1618.

[41] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li, "Face alignment across large poses: A 3D solution," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 146–155.

[42] E. Richardson, M. Sela, and R. Kimmel, "3D face reconstruction by learning from synthetic data," in *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, 2016, pp. 460–469.

[43] C. Cao, Y. Weng, S. Zhou, Y. Tong, and K. Zhou, "Facewarehouse: A 3D facial expression database for visual computing," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 3, pp. 413–425, 2013.

[44] X. Zhu, Z. Lei, J. Yan, D. Yi, and S. Z. Li, "High-fidelity pose and expression normalization for face recognition in the wild," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 787–796.

[45] L. Hu, C. Ma, L. Luo, and H. Li, "Single-view hair modeling using a hairstyle database," *ACM Transactions on Graphics (Proceedings SIGGRAPH 2015)*, vol. 34, no. 4, July 2015.

[46] C. Duncan, H. Dai, N. Pears, and W. Smith, "A novel shape based outcomes analysis tool for craniofacial surgeons," in *17th Biennial Congress of the International Society of Craniofacial Surgery (ISCFS)*, vol. 1, no. 2, 2017.

[47] S. Agarwal, K. Mierle, and Others, "Ceres solver," http://ceres-solver.org.