3D Face From X: Learning Face Shape from Diverse Sources

Yudong Guo,   Lin Cai,   Juyong Zhang†

Abstract—We present a novel method to jointly learn a 3D face parametric model and 3D face reconstruction from diverse sources. Previous methods usually learn 3D face modeling from one kind of source, such as scanned data or in-the-wild images. Although 3D scanned data contain accurate geometric information of face shapes, the capture system is expensive and such datasets usually contain a small number of subjects. On the other hand, in-the-wild face images are easily obtained and there are a large number of facial images. However, facial images do not contain explicit geometric information. In this paper, we propose a method to learn a unified face model from diverse sources. Besides scanned face data and face images, we also utilize a large number of RGB-D images captured with an iPhone X to bridge the gap between the two sources. Experimental results demonstrate that with training data from more sources, we can learn a more powerful face model.

Index Terms—Face modeling, shape from X, 3D face reconstruction.

I. INTRODUCTION

This paper considers the problem of jointly learning a nonlinear 3D Face Morphable Model (3DMM) and 3D face reconstruction with training data from three different sources: scanned 3D face, RGB-D images, and RGB images. 3D face modeling from scanned 3D face dataset or face modeling and capture from RGB images have been well studied in recent years. Since Blanz et al. [6] build a PCA model with 200 scanned 3D faces, there have been several works that build a parametric model from scanned 3D faces [12], [37]. However, 3D scanned faces are usually expensive to obtain, and the training data usually contain a small number of subjects. On the other hand, in-the-wild RGB images are very easy to obtain and there are many works to learn face modeling and face reconstruction on such dataset [51], [61], [60], [20]. However, compared to scanned data, RGB images do not contain geometric information, and it is an ill-posed problem to infer geometry from RGB images. Thus these works have to use facial priors for regularization.

Although 3D face modeling from RGB input have been well studied in recent years, real-time RGB based methods still have difficulty to achieve high-quality and robust facial performance capture due to limited information from the input RGB images. For example, RGB-based methods only have appearance information, and they cannot perform well or even fail in a dark environment. RGB-D sensors, on the other hand, provide not only appearance information but also depth information that is

robust to illumination changes and occlusions. With the help of additional depth information, RGB-D based methods could achieve better and more robust performance compared to RGB-based methods. Hence RGB-D based capturing has recently received increasing attention, especially with consumer-grade depth cameras becoming popular. For example, some newly released smartphones like Apple's iPhone X are equipped with an RGB-D camera, which stimulates the call for more efficient and pragmatic RGB-D face capturing solutions on mobile devices.

To bridge the gap between the reconstruction results of different sources and take advantage of all types of data, we propose to use datasets from all kinds of sources to train a unified face model. Our work is inspired by the great success of deep learning-based technology for face modeling and 3D face reconstruction from one kind of source. For example, many deep learning-based methods have been proposed for 3D face reconstruction based on RGB inputs. [51] and [20] propose a coarse-to-fine network architecture where a coarse-scale CNN is trained to regress the parameters of a parametric face model and a fine-scale CNN is trained to recover the face detail. [61] proposes an unsupervised CNN structure including a CNN encoder and a fixed rendering module as the decoder and trains the structure end-to-end for 3D face reconstruction from RGB input. [60] further improves the reconstruction accuracy by learning the corrective basis from sparsely labeled in-the-wild images. Very recently, [39] presents a deep learning-based method to learn face modeling from diverse raw scan data, in which a PointNet [48] is used as feature encoder. Similar to 3DMM [6], nonlinear 3DMM [65] and FLAME [37], we learn a low-dimensional representation for face modeling. Low-dimensional models can reconstruct global facial structure while can not recover fine-scale details due to the limited representation ability of the low-dimensional parametric model. The advantage of our method over existing low-dimensional models is to jointly learn a powerful face model and 3D face reconstruction from diverse sources.

It is however non-trivial to directly adopt existing methods to learn 3D face from diverse sources. There are two challenges that make it difficult to learn a unified face model from diverse sources. The first challenge is the gap between different sources: the scanned data live in 3D space while the facial images live in 2D space, which makes it difficult to handle different sources with the same network architecture. The second challenge is the lack of dense correspondence. As pointed out in [39], dense correspondence is crucial and challenging for face modeling from different databases. In our case, the problem is more challenging since our data are from different domains.

Yudong Guo and Juyong Zhang are with School of Mathematical Sciences, University of Science and Technology of China.

Lin Cai is with Beijing Dilusense Technology Corporation.

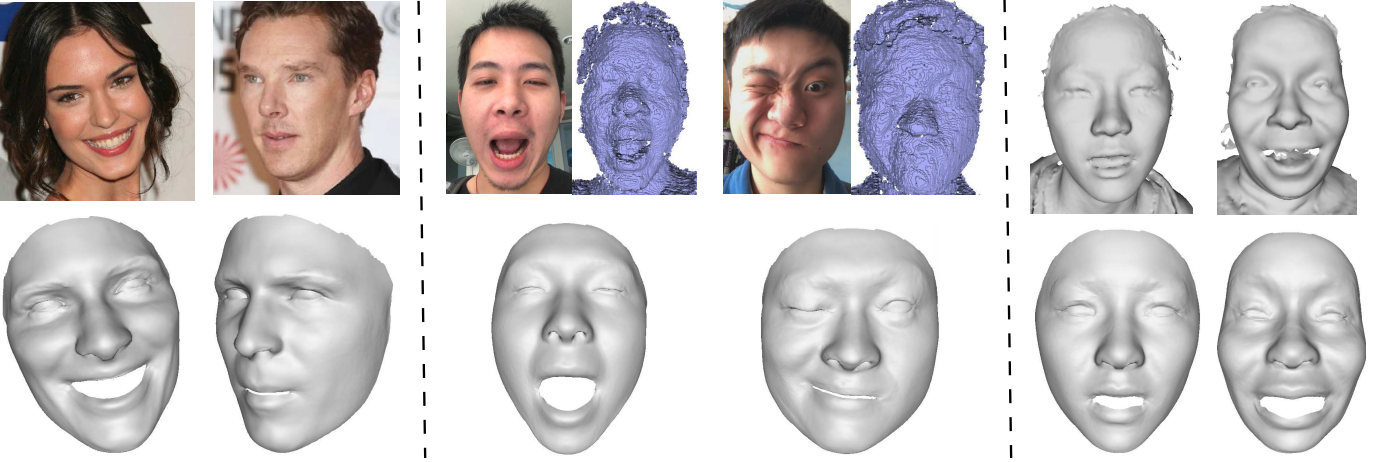†Corresponding author. Email: `juyong@ustc.edu.cn`.

Figure 1: Several 3D face reconstruction results from diverse sources. On the top are different inputs from RGB, RGB-D and raw scan data respectively, and on the bottom are corresponding reconstruction results. Our method is the first deep learning based 3D facial performance capture work with inputs from diverse sources.

We adopt several strategies to address these challenges. To bridge the gap between the two sources, we utilize a large number of RGB-D data, which contain information from both types while live in the same domain. For domain consistency, we render scanned 3D face to depth maps and use a three-branches CNN to encode data from different sources to geometric parameters. To make different sources benefit each other, we use cross-domain supervision losses to make the dataset with weak signals benefit from data with stronger supervision signals. Specifically, to make the RGB branch benefit from additional supervision signal from the RGB-D dataset, we feed the RGB images from RGB-D data to the RGB branch and use the depth information as supervision signal. To make the RGB-D branch benefit from more accurate and complete supervision signal from the scan dataset, we render the scan data to RGB-D images with artificial speckle noise by simulating the capturing process of depth cameras. And then feed the synthesized RGB-D images to the RGB-D branch while using original point clouds from scan data as supervision signals.

To tackle the second challenge of building dense correspondence, we first use a fixed face model by fixing the geometry decoder and only updating the encoder network to fit different sources. For scan and RGB-D data, we build dense correspondence between face model and input data by finding the closest point in the input point cloud. For RGB data, we build correspondence by projection and rasterization. We then update the face model with fixed correspondence obtained in the last step. The two steps are taken iteratively. To sum up, the paper has the following contributions:

- We propose a three-branches CNN learning framework to jointly learn face modeling and 3D face reconstruction with training data from three different sources. To the best of our knowledge, this is the first method to learn a unified 3D face model from diverse sources. Experiments demonstrate that with training data from more sources, we can learn a more powerful face model.

- We bridge the gap of supervision signals from different sources by using novel cross-domain supervision losses to make dataset with weak signals benefit from data with stronger supervision signals. Experimental results demonstrate that with additional supervision from RGB-D data during training, our model can infer more accurate geometry for RGB inputs, and that better supervision signals from scan data lead to more accurate reconstruction for RGB-D inputs.

- To bridge the gap between the scanned face dataset and RGB images dataset, we collect a large scale RGB-D dataset from a variety of identities with different facial expressions, poses, and lighting conditions. [1]

We show several reconstruction results of our method on different types of input data in Fig. I. It shows that our method can handle diverse sources well within a unified face model. The qualitative and quantitative comparison with the state-of-the-art RGB and RGB-D based methods in Section IV demonstrate that our method performs well on both types of data. We show more results and the learned face model in the supplementary video.

II. RELATED WORK

3D Face Shape Models. Human faces have similar global characteristics such as the location of eyes, nose and mouth. Based on this prior, some researchers propose to build low-dimensional face models to reduce the reconstruction problem into searching within the parameter space. Since Blanz and Vetter [6] propose to build a 3D Morphable Model (3DMM) to represent a textured 3D face with principal components analysis (PCA), low-dimensional face models have been widely used for 3D face representation. Among them, Paysan et al. [46] build the publicly available Basel Face Model (BFM2009) with 200 face scans in a neutral expression. Booth et al. [7] utilize

---

[1] Our constructed RGB-D facial dataset and the reconstructed mesh of each frame will be publicly available.

neutral scans of 9663 subjects to improve the representation ability of 3DMM. To also cover facial expressions, Vlasic et al. [68] adopt a multilinear face model representation that jointly parameterizes the identity and expression variations. Cao et al. [12] use a similar representation to build a multilinear 3D face model from 3000 scans of 150 subjects each with 20 different expressions. Li et al. [37] build the FLAME model that additionally represents head rotation and yaw motion with linear blend skinning.

Besides above traditional statistical method for 3D face modeling from scan data, recently there are also some works that apply CNN to learn powerful face shape models from RGB images. Tran et al. [66] propose to use a multilayer perceptron (MLP) to build a nonlinear form of 3DMM and learn the model with in-the-wild images. [65] improves the work by extending to consider facial geometric details. Sengupta et al. [54] propose to jointly learn the decomposition of shape, albedo and lighting conditions of a face from a collection of RGB images.

Optimization-based Facial Performance Capture from RGB Inputs. Monocular RGB images or videos can be easily obtained and many methods have been proposed to reconstruct faces from monocular RGB video. A popular way is to fit the input images using a low-dimensional parametric face model such as 3DMM with analysis-by-synthesis optimization strategy. The reconstructed model can be further refined by recovering 3D shape from shading variation [28], [27]. For instance, in [16], a template is deformed to a 3D model created by a binocular stereo approach. In [67], a blendshape model is built and fitted to a monocular video off-line, and then the surface detail is added by shading-based shape refinement under general lighting conditions. [55] uses a similar approach and refines the results by iteratively optimizing the large-scale facial geometry and the fine-scale facial detail. [23] proposes to create fully rigged, personalized 3D facial avatars from a video captured from a hand-held camera. The method performs the tracking and reconstruction computation on PC though the input images are captured by mobile devices. [17] fits a 3D face using a multi-layer approach and extracts a high-fidelity parameterized 3D face rig that contains a generative wrinkle formation model capturing the person-specific idiosyncrasies. As the low-rank parametric face model limits its expressiveness and ability to capture fine details, [57] proposes to derive a person-specific face model from all available images, and each frame of the video is reconstructed by a novel 3D optical flow approach coupled with shading cues. In general, these methods work off-line and are not suitable for real-time 3D facial performance capture. [63] presents a method to jointly fit a parametric model for identity, expression and skin reflectance to the input color, which achieves real-time 3D face tracking and facial reenactment with GPU acceleration.

Learning-based 3D Face Reconstruction from RGB Inputs. To achieve real-time capturing of 3D faces, many methods utilize machine learning to shift complex optimization to the off-line training stage. [11] presents a learning-based regression approach to fit a generic identity and expression model to an RGB face video on the fly. [70] extends it to run on mobile devices at real-time frame rates by directly regressing the head

poses and expression coefficients in one-step. The approach is also extended to include the learning of fine-scale facial wrinkles [10].

With the powerfulness of convolution neural networks, various deep learning-based methods have been developed for 3D face reconstruction, face alignment, face recognition and dense facial correspondences from RGB images. Examples are [26], [4], [74], [29], [36], [64], [25], [18], [73], [15]. In these methods, 3D Morphable Model (3DMM) [6], [52], [7], [41] is used to represent 3D faces and CNN is applied to learn the 3DMM and pose parameters. Most of these methods train the CNN in a fully supervised manner. [50] presents a method for 3D face reconstruction from an RGB image by training a CNN with synthetic data to learn the 3DMM parameters and [51] further trains a fine-scale CNN to regress per-pixel depth displacement for detailed geometry. [33] proposes to directly learn the vertex positions from RGB images by building an initial PCA basis. [20] proposes a fully supervised deep learning framework for reconstructing a detailed 3D face from monocular RGB video in real time. [24] proposes to regress volumes with CNN for a single face image directly. One problem of these supervised methods is the gap between the synthesized images and real-captured images. To tackle this problem, unsupervised [61] and weakly-supervised [60], [14] methods have been proposed to train CNNs with real-captured images by using an analysis-by-synthesis energy function. While these CNN-based methods for 3D face reconstruction from RGB images can achieve impressive real-time performance, they have inherent limitations due to the RGB inputs: easily affected by illumination, unable to handle images captured in a dark environment.

Optimization-based Facial Performance Capture from RGB-D Inputs. Compared to RGB inputs, RGB-D inputs provide additional depth information that benefits more robust facial performance capture with better quality. In fact, Weise et al. [69] developed a facial performance capture system that combines 3D and 2D non-rigid registration in optimization and achieves real-time robust 3D face tracking. However, this system needs to build an accurate 3D expression blendshape by scanning and processing a predefined set of facial expressions in advance. [34] describes a system that only requires pre-building a neutral face model, generates initial blendshapes using deformation transfer [56], and then trains PCA-based correctives for the blendshapes during tracking with examples obtained from per-vertex Laplacian deformations. [8] introduces a calibration-free system that requires no user-specific prepro-cessing by jointly solving for a detailed 3D expression model of the user and the corresponding dynamic tracking parameters via an optimization procedure. [13] proposes an approach that incrementally deforms a 3D template mesh model to best match the input RGB-D data and facial landmarks detected from the single RGB-D image. [22] presents a real-time facial tracking system in unconstrained settings using a consumer-level RGB-D sensor. The system personalizes the tracking model on-the-fly by progressively refining the user's identity, expressions, and texture with reliable samples and temporal filtering.

To generate high-quality 3D face models from RGB-D data, [35] proposes an approach consisting of two stages: offline

construction of personalized wrinkled blendshape and online 3D facial performance capturing. [62] presents a method that reconstructs high-quality facial performance of each actor by jointly optimizing the unknown head pose, identity parameters, facial expression parameters, face albedo values, and the illumination. The method focuses on photo-realistic capture and re-rendering of facial templates to achieve expression transfer between real actors. Different from fitting the RGB-D data to the parametric model, [38] proposes to divide the input depth frame into several semantic regions and search for the best matching shape for each region in the database. The reconstructed 3D mesh is constructed by combining the input depth frame with the matched shapes in the database. [42] proposes a direct approach that reconstructs and tracks the 3D face without templates or prior models.

III. PROPOSED METHOD

This section first introduces face representation and rendering procedure used in our method, and then presents the architecture of our three-branches CNN.

A. Face Representation

The face is basically represented as a vector $\mathbf{p} = [\mathbf{p}_1^T, \mathbf{p}_2^T, \cdots, \mathbf{p}_n^T]^T \in \mathcal{R}^{3n}$ via a mesh of n vertices v_i $(i = 1, 2, \ldots, n)$ with fixed connectivity, where $\mathbf{p}_i$ denotes the position of vertex v_i. We use a nonlinear 3DMM [66] to represent the facial geometry. Specifically, the facial geometry is represented as:

$$\mathbf{p} = D_s(\boldsymbol{\alpha}_s), \tag{1}$$

where D_s is the shape decoder network and $\boldsymbol{\alpha}_s$ is the latent code that control the face shape. Since our method concentrates on geometry, for albedo representation we just follow [6] that represents albedo with PCA:

$$\mathbf{b} = \bar{\mathbf{b}} + \mathbf{A}_{\mathrm{alb}}\boldsymbol{\alpha}_{\mathrm{alb}}, \tag{2}$$

where $\bar{\mathbf{b}}$ denotes average face albedo, $\mathbf{A}_{\mathrm{alb}}$ denotes the principal axes extracted from a set of textured 3D meshes, and $\boldsymbol{\alpha}_{\mathrm{alb}}$ is the albedo coefficient vector. We use the BFM2009 [46] for our albedo representation.

B. Rendering Process

The rendering process of a face image depends on several factors: face geometry, albedo, pose, lighting and camera parameters. To generate a realistic synthetic image of a face, we render the facial imagery using a standard perspective pinhole camera, which can be parameterized as follows:

$$\mathbf{q}_i = \boldsymbol{\Pi}(\mathbf{R}\mathbf{p_i} + \mathbf{t}) \tag{3}$$

where $\mathbf{p}_i$ and $\mathbf{q}_i$ are the locations of vertex v_i in the world coordinate system and in the image plane, respectively, $\mathbf{R}$ is the rotation matrix constructed from Euler angles $pitch, yaw, roll$, $\mathbf{t}$ is the translation vector and $\boldsymbol{\Pi} : \mathbb{R}^3 \to \mathbb{R}^2$ is a perspective projection.

To model the lighting condition, we approximate the global illumination using the SH basis functions. Furthermore, we assume that a face is a Lambertian surface. Under this assumption, the irradiance of a vertex v_i is determined by vertex normal $\mathbf{n}_i$ and scalar albedo b_i:

$$\mathbf{I}(\mathbf{n}_i, b_i \mid \boldsymbol{\gamma}) = b_i \sum_{k=1}^{B^2} \gamma_k \phi_k(\mathbf{n}_i) = b_i \phi(\mathbf{n}_i) \cdot \boldsymbol{\gamma} \tag{4}$$

where $\phi(\mathbf{n}_i) = [\phi_1(\mathbf{n}_i), \ldots, \phi_{B^2}(\mathbf{n}_i)]^T$ is the SH basis computed with normal $\mathbf{n}_i$, and $\boldsymbol{\gamma} = [\gamma_1, \ldots, \gamma_{B^2}]^T$ is the SH coefficients. In this paper, we use the first three bands $(B = 3)$ of SHs for the illumination computation.

C. Proposed CNN learning framework for 3D Face Learning From Various Sources

The overall pipeline of the proposed CNN learning framework is shown in Fig. III-C. The framework consists of three branches with different sources as input: RGB image, RGB-D image and scanned 3D data. The three branches regress parameters which are then fed into the shape decoder to generate facial shape for corresponding input.

Training Data Preparation. To train the three-branches CNN, we collect training data from three different kinds of sources. For RGB-D data, we capture 800 RGB-D videos with iPhone X (resolution 480×640 at 30 fps). During data capturing, we ask the person to rotate head while varying facial expressions freely under different lighting conditions (with light turning on or off). Fig. 3 shows samples of the captured data. For RGB data, we use about 300K images from CelebA [40] and VGG-Face [44] dataset. For raw scan data, we use a collection of 7450 scans from BU-3DFE [72] and FRGC [47] face database.

For depth information input, we convert the depth value at every pixel to 3-channel point coordinates in the camera space with the camera intrinsic parameters, before feeding it to CNN. In this way, the network can directly infer Euler angles and the translation vector from the 3-channel point coordinates. For raw scan input, we use a virtual camera to render the 3D face to a depth map and convert it to 3-channel as described above.

Three-Branches CNN. The three-branches CNN is targeted at 3D face reconstruction from different sources. That is, given data from different sources as input, our network regresses the parameter set $\chi = \{\boldsymbol{\alpha}_s, \boldsymbol{\alpha}_{\mathrm{alb}}, pitch, yaw, roll, \mathbf{t}, \mathbf{r}\}$, where $\mathbf{r} = (\boldsymbol{\gamma}_r^T, \boldsymbol{\gamma}_g^T, \boldsymbol{\gamma}_b^T)^T$ denotes the SH illumination coefficients of RGB channels. We denote $\chi_{scan}, \chi_{rgb}, \chi_{rgbd}$ as the regressed parameters of data from scan, RGB image and RGB-D image respectively.

As shown in Fig. III-C, the network consists of three branches. Each of the top and bottom branches contains a ResNet-50 [21] to extract RGB and XYZ features from RGB data and scan data respectively. The middle branch is used to handle RGB-D data. It uses the first 22 layers from the other two branches to extract features from RGB and D respectively, and then concatenates the features and feeds them to the last 28 layers of a ResNet-50. The extracted features from all sources are fed into corresponding fully-connected layers to output parameter sets.

We train the three-branches CNN end-to-end. Since we do not have ground truth labels for all three types of data, we use a weakly-supervised approach that utilizes signals from
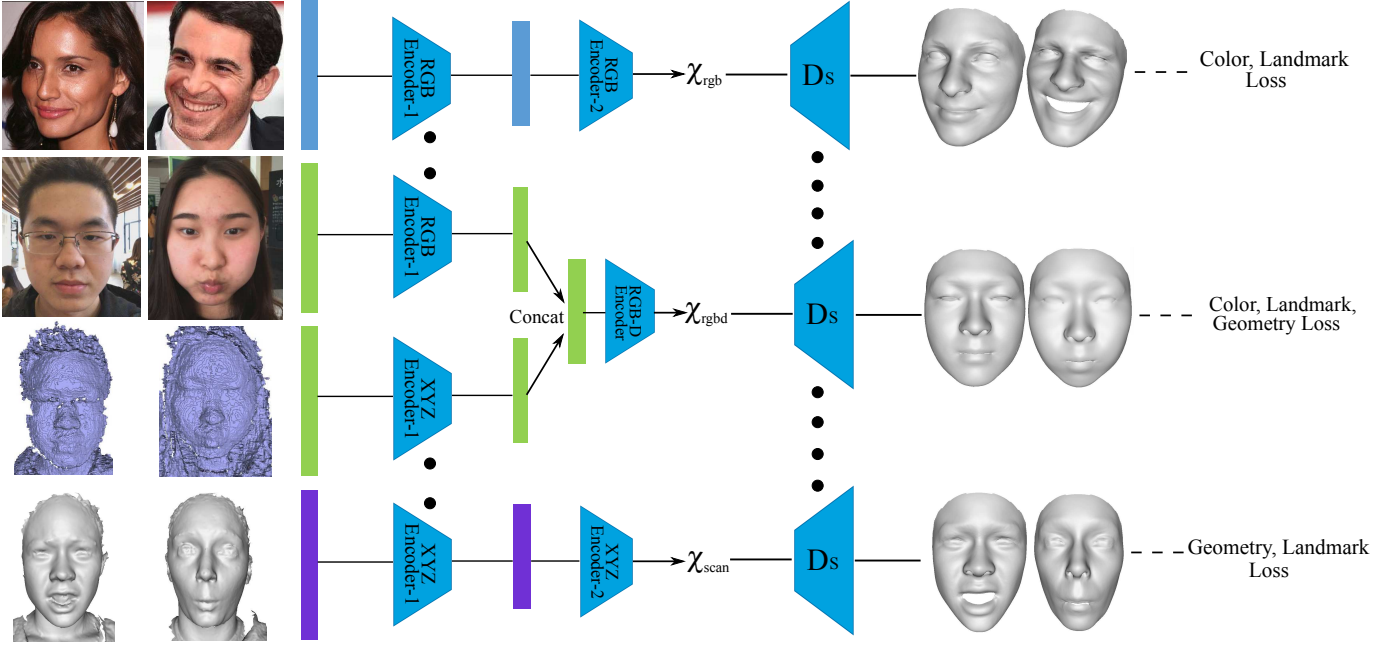
Figure 2: The pipeline of our proposed three-branches CNN learning framework for 3D face from diverse sources. The rectangles with different colors indicate features of different branches extracted from data belonging to different sources. The black dots indicate weight sharing.
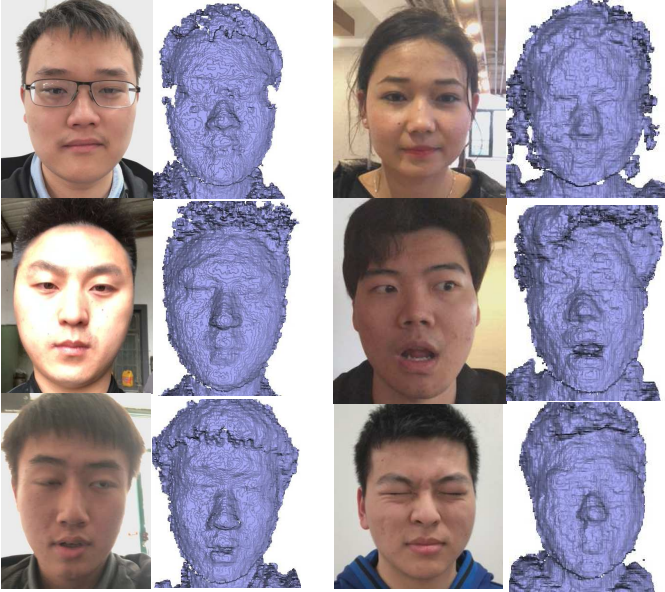

Figure 3: Samples of the captured RGB-D training data. We capture subjects with different poses and expressions under different lighting conditions.

training data themselves and some automatic detected sparse landmarks. Since different data have different signals, the loss functions are different. According to the signals, there are three kinds of loss functions.

Geometry Term. For scanned data and RGB-D data that have geometric information, we use a point-to-point distance term to measure how well the reconstructed mesh matches the point cloud:

$$E_{\mathrm{geo}}(\chi) = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p} \in \mathcal{P}} \|\mathbf{R}\mathbf{p} + \mathbf{t} - \mathbf{p}'\|_2, \qquad (5)$$

where $\|\cdot\|_2$ is the $\ell_{2,1}$ norm, $\mathcal{P}$ is the set of all vertices, $\mathbf{p}$ is a vertex in the reconstructed mesh, and $\mathbf{p}'$ is the closest point in the point cloud.

Color Term. For RGB data and RGB-D data that have color information, we use a color term to evaluates how well the rendered face image based on the regressed parameters matches the input RGB image:

$$E_{\mathrm{col}}(\chi) = \frac{1}{|\mathcal{F}|} \sum_{m \in \mathcal{F}} \|\mathbf{I}_{\mathrm{syn}}(m) - \mathbf{I}_{\mathrm{real}}(m)\|_2, \qquad (6)$$

where $\mathcal{F}$ is the set of all pixels in the projected facial region, $\mathbf{I}_{\mathrm{syn}}(m)$ and $\mathbf{I}_{\mathrm{real}}(m)$ are the synthetic color and the real color at pixel m, respectively. To enable back propagation, we do a standard rasterization process and then fix the pixel-triangle correspondence and barycentric coordinates during back propagation. We also use the $\ell_{2,1}$ norm for robustness since we do not account for specular reflection and person-specific albedo, which are considered as outliers in our model.

Landmark Term. For all three kinds of data, we use a landmark term to measure how close the sparse vertices are to the corresponding landmarks in the input image or point cloud. For RGB and RGB-D images, we compute the landmark loss in 2D:

$$E_{\mathrm{lan}}(\chi) = \frac{1}{|\mathcal{L}|} \sum_{i \in \mathcal{L}} \|\mathbf{q}_i - \mathbf{\Pi}(\mathbf{R}\mathbf{p}_i + \mathbf{t})\|^2, \qquad (7)$$

where $\mathcal{L}$ is the set of landmarks, $\mathbf{q}_i$ is a detected landmark position in the input image, $\mathbf{p}_i$ is the corresponding vertex location in the 3D mesh. We adopt the method proposed in [9] to detect landmarks for facial images. For scan data, we compute the landmark loss in 3D:

$$E_{\text{lan3d}}(\chi) = \frac{1}{|\mathcal{L}|} \sum_{i \in \mathcal{L}} \|\mathbf{p}_i' - (\mathbf{R}\mathbf{p}_i + \mathbf{t})\|^2, \qquad (8)$$

where $\mathbf{p}_i'$ is the 3D landmark in the point cloud. Similar to [39], the 3D landmarks are obtained based on rendered images.

Cross-domain Supervision Term. Different sources in our training data have different types of supervision signals. Among the three sources, the scan dataset have the most complete and accurate 3D signals. The RGB-D dataset have partial 3D signals since they live in 2.5D space and contain some noises produced in capturing process of depth cameras. The RGB dataset have the weakest signals with 2D information only. To bridge the gap of supervision signals from different sources, we use cross-domain supervision losses to make dataset with weak signals benefit from data with stronger supervision signals. Specifically, to make the RGB branch benefit from additional supervision signal from RGB-D dataset, we feed the RGB images from RGB-D data to the RGB branch during training and use a cross-domain supervision term:

$$E_{\text{cross}}(\chi) = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p} \in \mathcal{P}} \|\mathbf{p} - \mathbf{p}_{rgbd}\|_2, \qquad (9)$$

where p_{rgbd} is the corresponding point in the mesh reconstructed from the corresponding RGB-D image. The process is shown in Fig. 4 (a).

To make the RGB-D branch benefit from more accurate and complete supervision signal from the scan dataset, we render the scan data to RGB-D images by simulating the capturing process of depth cameras. Specifically, we use a virtual camera to generate a simulated speckle image of the model from scan dataset and then recover the depth map from the speckle image. We feed the synthesized RGB-D images to the RGB-D branch and use original point clouds from scan data as supervision signals to compute geometric loss:

$$E_{\text{cross\_geo}}(\chi) = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p} \in \mathcal{P}} \|\mathbf{R}\mathbf{p} + \mathbf{t} - \mathbf{p}'\|_2, \qquad (10)$$

where $\mathbf{p}$ is a vertex in the reconstructed mesh, and $\mathbf{p}'$ is the closest point in the point cloud of scan data. The process is shown in Fig. 4 (b). Note that the scan data also benefit the RGB branch indirectly by improving the RGB-D branch and providing more accurate cross-domain supervision signals to the RGB branch.

Implementation. We implement the network architecture using PyTorch [45]. When feeding the input image into the network, we crop the face region to 256×256 with bilinear interpolation. The network parameters are updated using Adam solver [30] with the mini-batch size of 20. We initialize the geometry decoder with the parametric face model used in [20], which uses BFM for identity and builds a 79-dim expression model from FaceWarehouse [12] dataset. We update our network encoders and shape decoder in an iterative manner.
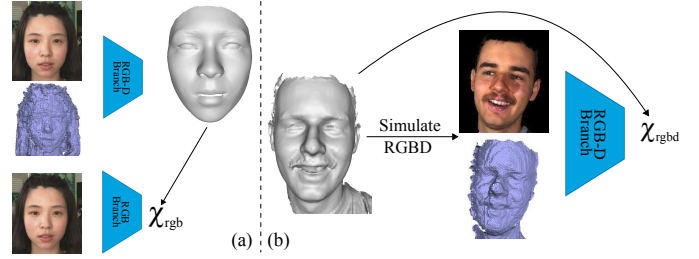

Figure 4: The processes of two cross-domain supervision losses. The arrows pointing from face models to regressed parameters indicate cross-domain supervision signals.

In each iteration, there are two steps that update the encoders of three branches and the shape decoder respectively. In the first step, we fix the shape decoder and update the encoders of three branches with the loss terms described above until convergence (5 epochs in the experiment). The second step is to learn the shape decoder, and we first build dense correspondence between our face model and different sources. Specifically, for scan data and RGB-D data, we find the corresponding point $\mathbf{p}'$ used in Eq. 5 by searching the closest point in the input point cloud for each vertex $\mathbf{p}$ in our model based on the reconstruction of last step. For RGB data, we find the pixel-triangle correspondence and compute barycentric coordinates by a standard rasterization process also based on the reconstruction of last step. Then we fix the dense correspondence obtained above and network encoders and update the shape decoder with the loss terms described above until convergence (2 epochs in the experiment). The whole network is trained with 3 iterations.

IV. RESULTS

In this section, we first evaluate the performance of our learning-based 3D face from various sources in different aspects, and then compare it with the state-of-the-art RGB and RGB-D based methods.

A. Evaluation of Our Method

Consistence of different sources. The three branches of our network process different sources respectively, and they should be consistent with each other. In other words, the results from three sources of the same subject should be similar to each other. We demonstrate this by feeding the RGB image, RGB-D pair and depth map from an RGB-D image into the three branches respectively. Fig. 5 shows the results with different sources. We can see that the results of different sources are very close to each other, and that the results with RGB input only are comparable to that with RGB-D input. This is due to two reasons. First, the RGB-D branch shares weights with the other two branches in the first several layers of the network. Second, since the RGB images from RGB-D dataset are also fed into the RGB branch, the RGB-D dataset provide additional supervision for the RGB branch.

Dense correspondence accuracy across diverse domains. During network training, we adopt a simple yet effective iterative manner to build dense correspondence between our
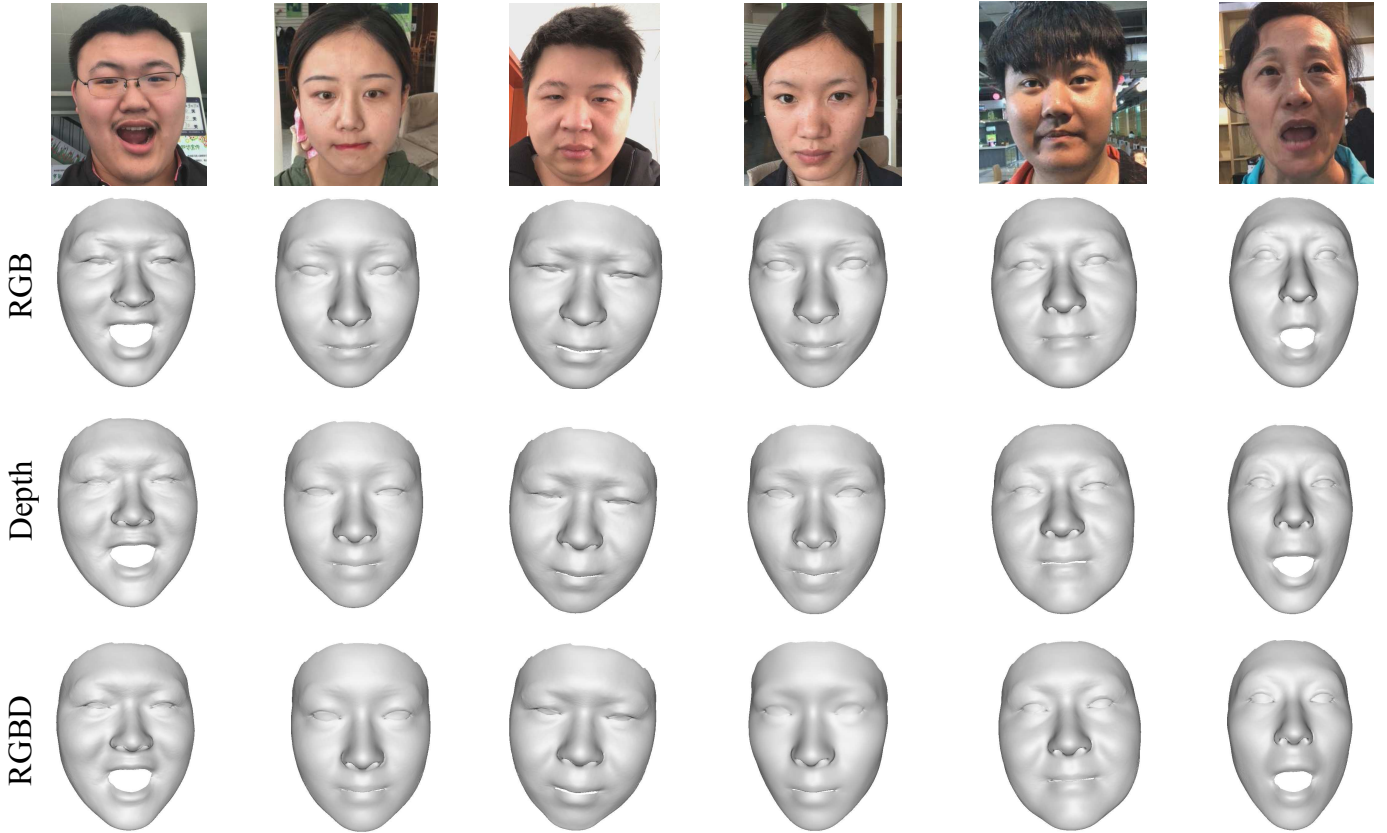
Figure 5: Reconstruction results of our three branches CNN with inputs from different sources of the same sample. From top to bottom: input subjects, results with RGB input, results with depth input, results with RGB-D input. Please refer to the accompanying video for complete sequences.

face model and different sources. We evaluate this strategy by computing the correspondence accuracy of different iterations on BU3DFE database. BU3DFE contains 100 subjects and for each subject, it contains one neutral and six expression scans with four levels of strength and corresponding RGB images. Thus the database could be fed into all our three branches. All the scans are manually labeled with 83 landmarks and we use the landmarks of all neutral scans and expression scans in the highest level for evaluation. Specifically, for scan branch and RGB-D branch, we compute the distance of 3D landmarks between the reconstructed shape and manually labeled ones. For RGB branch, we compute the distance of 2D landmarks in image domain between the projected reconstructed faces and the labeled landmarks in 2D. Note that for this experiment we do not include BU3DFE in our training data. Also note that the performance gains are not achieved by more training epochs since the loss terms converge in each iteration. The landmark errors of different sources in different iterations are reported in Tab. I. It can be easily observed that more accurate correspondences are built with more iterations for all three branches, demonstrating the effectiveness of the proposed strategy.

Representation Power of Model Trained with Different Sources. In our three-branches CNN, we learn a unified face model with training data from three different sources. We

	Scan (mm)	RGB-D (mm)	RGB (pixel)
Iter #1	7.52	7.63	4.37
Iter #2	6.71	6.95	3.74
Iter #3	**6.50**	**6.73**	**3.53**

Table I: Mean semantic landmark error of different branches in different iterations on BU3DFE database.

demonstrate the advantage of this choice by comparing the representation power of the learned model with other models that are not trained with all of the sources. Specifically, we train 6 models with the same network architecture as ours but trained only with one kind of source (scan data, RGB data, RGB-D data) and two kinds of sources (scan and RGBD, scan and RGB, RGBD and RGB) respectively. To further evaluate the performance of different relative ratios of data sources, we train other 3 models each with half of data from one kind of source removed. We compare the representation ability of the trained shape decoders of different models on 3000 meshes (150 subjects with 20 expressions each) from FaceWarehouse dataset and 1440 meshes (10 meshes each sequence) from COMA [49] dataset. Specifically, we compare the 10 models by optimizing shape parameters and rigid transform parameters to fit the shape model to the meshes and then compute point-to-point distances. We also compare with the face model used in [20], which is used as the initialization of our shape decoder.
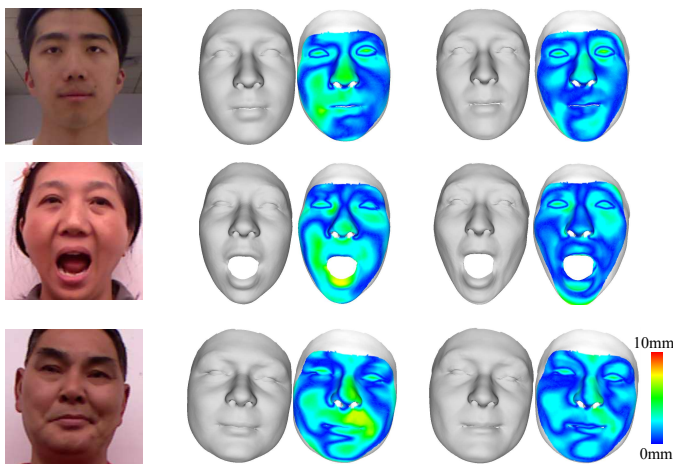
Figure 6: Comparisons on RGB inputs with the model trained without the cross-domain supervision loss. From left to right: input images, results of the model trained without the cross-domain supervision loss, results of our model.
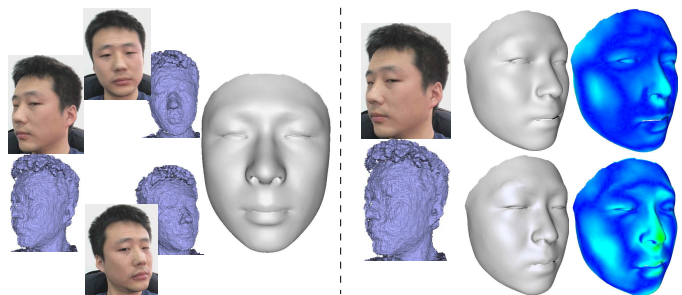

Figure 7: Comparisons on RGB-D input with the model trained without the cross-domain supervision loss. On the left are multi-view RGB-D images and corresponding optimized face shape. On the right are single-view RGB-D input and results shown in another view from our model (top) and the model trained without the cross-domain supervision loss (bottom).

The average distances of all models are listed in Tab. II. From the results of single-branch, two-branches and three-branches models we can see that better results are achieved if trained with more kinds of data sources. From the results of last 7 columns we can see that more training data lead to more powerful models. By analyzing the behavior of different sources we can observe that the RGB-D data and scan data have more significant impact on the performance gains than the RGB data while the best result is achieved with all data from all sources. Please refer to the accompanying video that shows the behavior of the learned shape decoder.

Advantage of Cross-domain Supervision Loss. In our training architecture, we make different sources benefit each other by using a cross-domain supervision loss. We demonstrate the advantage of this loss by comparing our model with the models that are trained without cross-domain supervision loss.

First, we evaluate the two models with RGB images as input on 180 samples of 9 subjects from FaceWarehouse dataset. Specifically, we feed the 180 RGB facial images to the two models and compute point-to-point distances between the reconstructed meshes with the corresponding ground-truth meshes after doing ICP alignment with uniform scaling. The average distance of our model and the model trained without cross-domain supervision loss are 1.56 mm and 1.73 mm respectively. Some samples are shown in Fig. 6. With additional supervision from RGB-D data during training, our model can infer more accurate geometry for RGB inputs.

We also evaluate the two models with RGB-D images as input to show the benefits of synthesized RGB-D data from scan data to the RGB-D branch in Fig. 7. We first optimize a face shape from multi-view RGB-D images as ground-truth. Then we feed an RGB-D image captured in a side view to the two models and compute the point-to-point to the ground-truth shape. As shown in Fig. 7, the result in unseen regions of our model is more accurate due to more complete and accurate supervision signals from scan data.

B. Comparisons with State-of-the-Art Approaches on RGB-D inputs

In this subsection, we compare our trained three-branches CNN model with state-of-the-art approaches [62], [34], [1] that can reconstruct 3D faces from RGB-D images.

Comparison on test RGB-D data captured with Iphone X. We compare our method with the state-of-the-art optimization-based method [62], [34]. The method of Thies et al. [62] achieves state-of-the-art performance on real-time 3D face reconstruction and tracking from a single RGB-D camera. It basically recovers 3D face geometry, albedo and lighting from an RGB-D sequence simultaneously by solving an optimization problem. It achieves real-time tracking on a desktop PC with the help of GPU computing. Li et al. [34] proposes a system that first builds initial blendshapes by doing deformation transfer [56] on a scanned neutral face, and then trains PCA-based correctives from input depth maps during tracking to represent more accurate facial geometry that is not spanned by initial blendshapes. We compare with the two methods on 5 test RGB-D videos captured with an Iphone X camera. We compute the point-to-point distance of the results of all methods to the point clouds generated by the input depth maps. The geometric errors are listed in Tab. III and we show three samples in Fig. 8. Our method outperforms [62] and [34] on all the test videos. Please refer to the accompanying video for complete sequences.

Comparison on Range 7 data. We also compare our method with the methods of ARKit [1] and Thies et al. [62] on 4 samples that have ground-truth models captured by a Range 7 3D Scanner [32]. ARKit [1] released by Apple Inc is based on the method of [8] and has the feature of capturing a user's 3D face with iPhone X smartphone. All the three methods take the RGB-D data scanned by iPhone X as input. We quantitatively evaluate the geometry accuracy of all methods, which is defined as the point-to-point distance between a reconstruction mesh and the corresponding ground-truth geometry captured by the Range 7 3D Scanner. Note that we only compare the facial part since the laser scanner fails to capture hair. To align the reconstructed 3D face model with the ground-truth 3D

	[20]	R	D	S	R&D	R&S	D&S	R_h&D&S	R&D_h&S	R&D&S_h	R&D&S
FaceWarehouse	0.583	0.556	0.508	0.513	0.501	0.509	0.483	0.475	0.490	0.486	**0.468**
COMA	0.478	0.445	0.412	0.401	0.398	0.390	0.379	0.372	0.377	0.383	**0.366**

Table II: Average geometric errors (mm) on FaceWarehouse dataset and COMA dataset of different models trained with different sources. R: RGB data. D: RGB-D data. S: Scan data. X_h: half of the data from X. Our model that are trained from all data from all sources has the best representation ability.
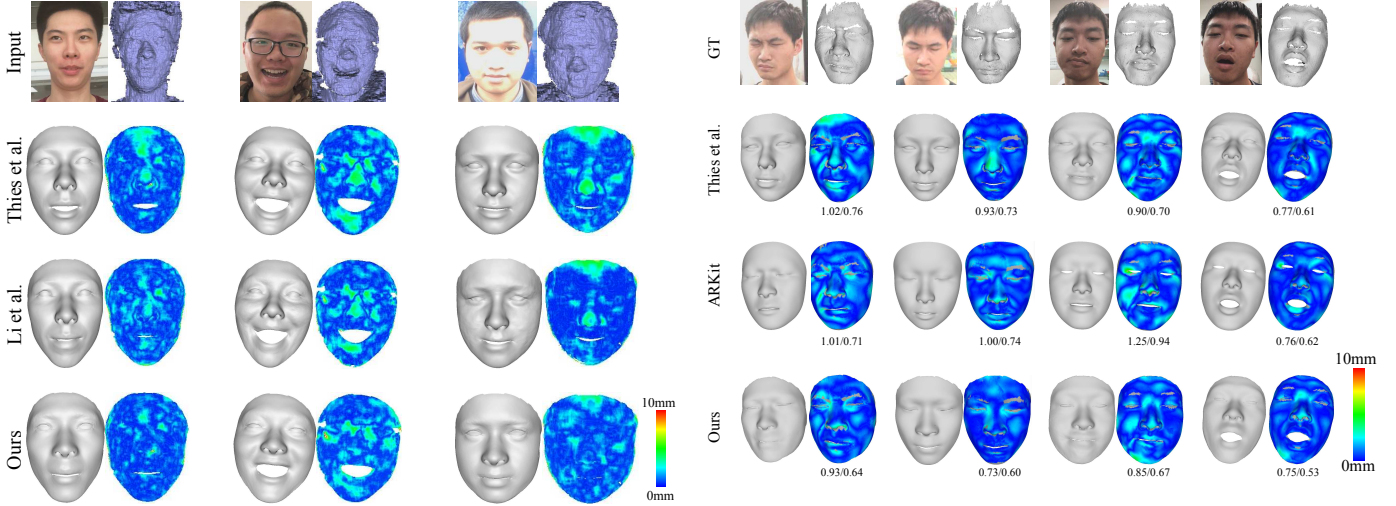

Figure 8: Comparisons with Thies et al. [62] and Li et al. [34] on RGB-D inputs. From top to bottom: input samples, results of Thies et al. [62], results of Li et al. [34], results of our method. The error map images show the fitting error between the reconstructed models and the point clouds generated by depth maps.

Video	1	2	3	4	5	Ave
Thies et al. [62]	1.36	1.29	0.95	1.40	1.26	1.26
Li et al. [34]	1.26	1.24	0.89	1.35	1.19	1.19
Ours	**1.17**	**1.10**	**0.84**	**1.25**	**1.16**	**1.11**

Table III: Geometric errors of [62], [34] and our method on 5 test RGB-D videos.

point cloud, we first manually label several landmarks to do rigid registration and then apply iterative closest point (ICP) algorithm [3] for dense registration. From Fig. 9, we can see that the geometry produced by our method is visually comparable or better than ARKit and [62]. The fitting errors in mean and standard deviation of our method and the other two methods against the ground-truth geometry are also given for the examples in Fig. 9, where our method obtains the smallest fitting errors.

C. Comparisons with State-of-the-Art Approaches on RGB inputs

For RGB inputs, we compare our method with the state-of-the-art RGB-based methods [14], [29], [61], [60], [58], [59], [17], [24], [20], [15], [53].

Quantitative comparison on FaceWarehouse dataset with [14], [29], [61], [60], [58], [59], [17]. With the same setting as [60], we compare our results on 180 meshes of 9 subjects from FaceWarehouse with five learning-based

Figure 9: Comparisons with Thies et al. [62] and ARKit [1] on Range 7 dataset. 1st row: ground-truth geometry captured by Range 7 3D Scanner and input images. 2nd row: results of [62] and geometry accuracy (w.r.t. ground-truth geometry). 3rd row: results of ARKit. 4th row: results of our method. The error map images show the fitting error between the reconstructed models and the groundtruth models. The mean / standard deviations of errors (mm) are listed at the bottom.

approaches of Tewari et al. [61], [60], [58], Deng et al. [14], Kim et al. [29], an optimization-based approach of Garrido et al. [17] and a hybrid approach of Tewari et al. [59]. Following the evaluation protocol of [14], we compute the point-to-point distances between the reconstructed meshes and the ground-truth meshes after alignment by running ICP with uniform scaling. The point-to-point errors are listed in Tab. IV. Our method achieves similar accuracy as the optimization-based method [17], while is orders of magnitude faster. And our method greatly outperforms other learning-based methods and the hybrid approach [59]. We also show the reconstruction results visually in Fig. 10. It can be seen that our method outperforms the methods of Tewari et al. [60] and Deng et al. [14] in terms of shape and expression reconstruction.

Quantitative comparison on IPhone X dataset with [24], [20]. We also compare our method with two RGB-based methods [24], [20] by using the source codes from the authors. We compare our results on 5 videos captured with an Iphone X camera that have depth maps as ground-truth. We compute the fitting errors by aligning the reconstructed meshes with the point cloud generated from depth map by doing ICP with uniform scaling and then compute a point-to-point distance. The point-to-point errors are listed in Tab. V and we show sample results in Fig. 11. Our method outperforms the other

Table IV: Mean reconstruction error on 180 meshes of 9 subjects from FaceWarehouse. Our geometric error is the lowest among all methods, even lower than the optimization-based method [17]. The times of other methods are quoted from [14], [58].

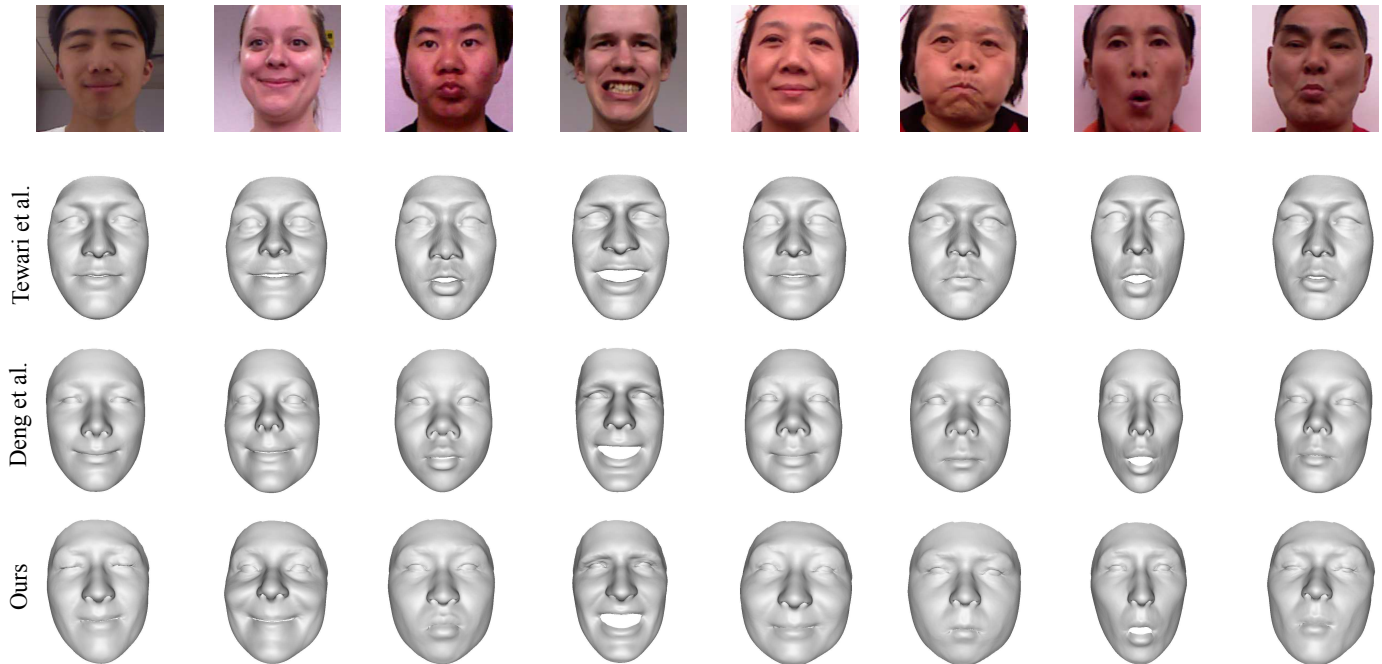|	Learning							Optimization	Hybrid
	Ours	[14]	[58]	[60]-F	[60]-C	[61]	[29]	[17]	[59]
Mean	**1.56**	1.81	2.01	1.84	2.03	2.19	2.11	1.59	1.87
SD	**0.25**	0.50	0.41	0.38	0.52	0.54	0.46	0.30	0.42
Time	20 ms	20 ms	5.2ms	4 ms	4 ms	4 ms	4 ms	120 s	110 ms

Figure 10: Reconstruction results with RGB inputs (top row) on FaceWarehouse of the method of Tewari et al. [60], Deng et al. [14] and ours. Our method represents the shapes and expressions better, especially in the eye and mouth regions. Please refer to the accompanying video for all comparison results on the 180 samples.

two methods greatly both quantitatively and qualitatively thanks to more accurate supervision signals with cross-domain supervision loss during training.

Video	1	2	3	4	5	Ave
Jackson et al. [24]	1.43	2.31	2.09	2.50	1.97	2.08
Guo et al. [20]	1.80	2.05	1.49	2.37	1.62	1.88
Ours	**1.39**	**1.56**	**1.07**	**1.63**	**1.54**	**1.44**

Table V: Geometric errors of [24], [20] and our method on 5 test RGB-D videos with RGB inputs.

Quantitative comparison on MICC dataset with [15], [53]. We also compare our method on the MICC dataset [2] with the methods of Feng et al. [15] and Soubhik et al. [53]. The MICC dataset contains 53 subjects with its ground-truth textured 3D mesh acquired from a structured-light scanning system. We render each subject to an RGB image and use it as input to all the three methods. For equal comparison, we crop all the reconstructed meshes to 95mm around the nose tip and align the reconstructed shapes to the corresponding ground-truth meshes by doing ICP with uniform scaling and we delete the inner eye and inner mouth part in the reconstructed meshes of [53]. Then we compute point-to-plane distances from the reconstructed meshes of all methods to ground-truth
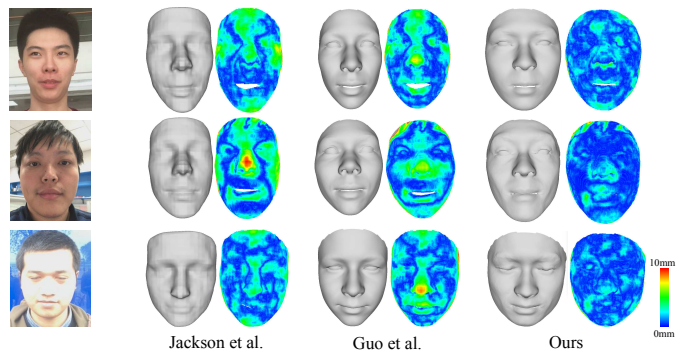

Figure 11: Comparisons with the RGB-based methods Jackson et al. [24] and Guo et al. [20]. From left to right: input facial images, results of [24], results of [20], results of ours. Our results greatly outperforms the other methods in terms of geometric errors.

meshes. The average geometric errors are listed in Tab. VI and we show three samples in Fig. 12. Our method outperforms the other two methods in terms of geometric errors.

Quantitative comparison on AFLW2000-3D dataset [74]

Method	[15]	[53]	Ours
Mean (mm)	1.51	1.59	**1.38**

Table VI: Geometric errors of Feng et al. [15], Soubhik et al. [53] and ours on MICC dataset.
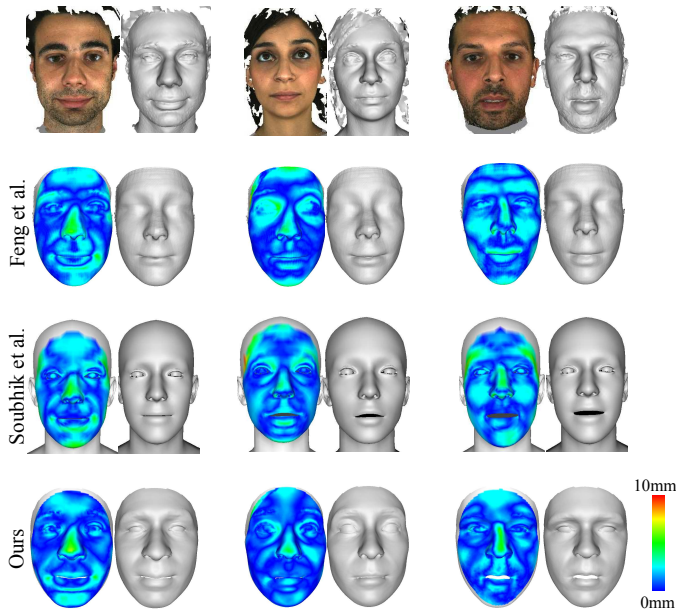

Figure 12: Reconstruction results with RGB inputs on MICC dataset of the method of Feng et al. [15], Soubhik et al. [53] and ours. From top to bottom: input images and corresponding ground-truth meshes, results of [15], results of [53], results of our method.

with [74], [5], [9], [15], [19], [71]. We compare the 3D face alignment accuracy of our method with other 3D facial landmark detection and reconstruction methods [74], [5], [9], [15], [19], [71] on AFLW2000-3D dataset. AFLW2000-3D is constructed by Zhu et al. [74] for evaluating 3D face alignment performance, which contains the ground truth 3D faces and the corresponding 68 landmarks of the first 2,000 samples from AFLW dataset [31]. Since AFLW2000-3D has many facial images with large poses, we fine-tune our RGB branch with 300W-LP dataset [74] for this comparison. Following the protocol of Guo et al. [19], we compute the Normalized Mean Error (NME) by bounding box size for each sample. The errors are listed in Tab. VII. It can be seen that our method achieves comparable results with MMFace-ICP-192 [71] while outperforms MMFace-ICP-128 (the authors of [71] used MMFace-ICP-128 by default in their paper) and all other methods. Thanks to the more powerful face model learned from diverse sources and more accurate semantic correspondence achieved by the iteratively dense correspondence updating strategy, our method also behaves well in this 3D facial landmark detection task.

Qualitative comparison with [50], [61], [66], [14]. We conduct qualitative comparison with Richardson et al. [50], Tewari et al. [61], Tran and Liu [66] and Deng et al. [14] in Fig. 13. Our method produces the best rendering results and the most accurate reconstruction of closed eyes. The reason

Method	AFLW2000-3D (68 pts)			
	[0, 30]	[30, 60]	[60, 90]	Mean
3DDFA [74]	3.78	4.54	7.93	5.42
3DDFA+SDM [74]	3.43	4.24	7.17	4.94
3DSTN [5]	3.15	4.33	5.98	4.49
3DFAN [9]	2.77	3.48	4.61	3.62
PRNet [15]	2.75	3.51	4.61	3.62
3DDFA_V2 [19]	2.63	3.42	4.48	3.51
MMFace-ICP-128 [71]	2.61	3.65	4.43	3.56
MMFace-ICP-192 [71]	**2.50**	3.63	**4.25**	**3.46**
Ours	2.60	**3.41**	4.43	3.48

Table VII: The NME (%) of different methods on AFLW2000-3D. The numbers of other methods are quoted from [19], [71].
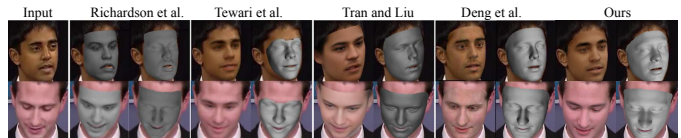

Figure 13: Qualitative comparison with Richardson et al. [50], Tewari et al. [61], Tran and Liu [66], and Deng et al. [14]. Images are from [14].

why our method outperforms existing methods is owing to the cross-domain supervision loss, which helps improving the reconstruction accuracy for RGB images and thus enabling better rendering results.

V. LIMITATIONS.

We have demonstrated high-quality 3D face reconstruction from diverse sources. Still, our approach has a few limitations which can be addressed in future work. First, we can not guarantee robust reconstruction of occluded faces. As shown in Fig. 14, our method produces inaccurate shapes or poses when the input faces are occluded by hands. One potential solution to this problem is to detect the occluded regions in training images with the face segmentation method [43] and exclude these regions in the loss terms during training. Second, we enforce low-dimensionality of the face model for robust learning. Thus our approach can not recover fine-scale facial details. We consider this as a solved problem since the coarse-to-fine strategies used in [20], [51], [65] could be easily extended to our framework.

VI. CONCLUSION

We have presented a deep learning method for learning a face model and 3D face reconstruction from diverse sources. The underlying technique is a self-supervised CNN framework that trains a three-branches encoder and a shape decoder from a combination of RGB-D, RGB and raw scan data. The framework is flexible to support facial performance capture with RGB-D, depth or RGB input during testing. We demonstrate the advantage of our framework by showing the representation power of the learned shape decoder and comparing the face reconstruction results with state-of-art methods. Experiments show that our method is suitable for consumer-level RGB-D
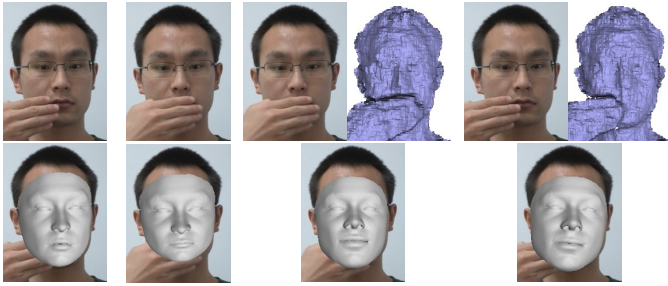
Figure 14: Reconstruction results of occluded faces. On the top are input RGB or RGB-D images, and on the bottom are the corresponding reconstructed faces.

or RGB sensors and capturing 3D faces in the wild, which is of great practical value.

ACKNOWLEDGMENTS

REFERENCES

[1] Apple Inc. Arkit, 2017.
[2] A. D. Bagdanov, A. Del Bimbo, and I. Masi. The florence 2d/3d hybrid face dataset. In *Proceedings of the 2011 Joint ACM Workshop on Human Gesture and Behavior Understanding*. ACM, 2011.
[3] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:239–256, 1992.
[4] C. Bhagavatula, C. Zhu, K. Luu, and M. Savvides. Faster than real-time facial alignment: A 3d spatial transformer network approach in unconstrained poses. In *IEEE International Conference on Computer Vision*, pages 4000–4009, 2017.
[5] C. Bhagavatula, C. Zhu, K. Luu, and M. Savvides. Faster than real-time facial alignment: A 3d spatial transformer network approach in unconstrained poses. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3980–3989, 2017.
[6] 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*, SIGGRAPH '99, pages 187–194, 1999.
[7] J. Booth, A. Roussos, A. Ponniah, D. Dunaway, and S. Zafeiriou. Large scale 3d morphable models. *International Journal of Computer Vision*, 126(2-4):233–254, 2018.
[8] S. Bouaziz, Y. Wang, and M. Pauly. Online modeling for realtime facial animation. *ACM Transactions on Graphics (TOG)*, 32(4):40, 2013.
[9] A. Bulat and G. Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *IEEE International Conference on Computer Vision*, 2017.
[10] C. Cao, D. Bradley, K. Zhou, and T. Beeler. Real-time high-fidelity facial performance capture. *ACM Transactions on Graphics (TOG)*, 34(4):46, 2015.
[11] C. Cao, Q. Hou, and K. Zhou. Displaced dynamic expression regression for real-time facial tracking and animation. *ACM Transactions on graphics (TOG)*, 33(4):43, 2014.
[12] 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*, 20(3):413–425, 2014.
[13] Y.-L. Chen, H.-T. Wu, F. Shi, X. Tong, and J. Chai. Accurate and robust 3d facial capture using a single rgbd camera. In *IEEE International Conference on Computer Vision*, pages 3615–3622, 2013.
[14] 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 *IEEE Computer Vision and Pattern Recognition Workshops*, 2019.
[15] 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)*, pages 534–551, 2018.
[16] P. Garrido, L. Valgaerts, C. Wu, and C. Theobalt. Reconstructing detailed dynamic face geometry from monocular video. 32(6):158:1–158:10, 2013.
[17] P. Garrido, M. Zollhöfer, D. Casas, L. Valgaerts, K. Varanasi, P. Pérez, and C. Theobalt. Reconstruction of personalized 3d face rigs from monocular video. *ACM Transactions on Graphics (TOG)*, 35(3):28, 2016.
[18] R. A. Güler, G. Trigeorgis, E. Antonakos, P. Snape, S. Zafeiriou, and I. Kokkinos. Densereg: Fully convolutional dense shape regression in-the-wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2614–2623, 2017.
[19] J. Guo, X. Zhu, Y. Yang, F. Yang, Z. Lei, and S. Z. Li. Towards fast, accurate and stable 3d dense face alignment. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
[20] Y. Guo, J. Zhang, J. Cai, B. Jiang, and J. Zheng. Cnn-based real-time dense face reconstruction with inverse-rendered photo-realistic face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(6):1294–1307, 2019.
[21] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
[22] P. Hsieh, C. Ma, J. Yu, and H. Li. Unconstrained realtime facial performance capture. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1675–1683, 2015.
[23] A. E. Ichim, S. Bouaziz, and M. Pauly. Dynamic 3d avatar creation from hand-held video input. *ACM Transactions on Graphics (ToG)*, 34(4):45, 2015.
[24] 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 *IEEE International Conference on Computer Vision*, pages 1031–1039, 2017.
[25] X. Jin and X. Tan. Face alignment in-the-wild: A survey. *Computer Vision and Image Understanding*, 162:1–22, 2017.
[26] A. Jourabloo and X. Liu. Pose-invariant face alignment via cnn-based dense 3d model fitting. *International Journal of Computer Vision*, 124(2):187–203, 2017.
[27] I. Kemelmacher-Shlizerman and R. Basri. 3d face reconstruction from a single image using a single reference face shape. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(2):394–405, 2011.
[28] I. Kemelmacher-Shlizerman and S. M. Seitz. Face reconstruction in the wild. In *IEEE International Conference on Computer Vision*, pages 1746–1753, 2011.
[29] H. Kim, M. Zollhöfer, A. Tewari, J. Thies, C. Richardt, and C. Theobalt. InverseFaceNet: Deep monocular inverse face rendering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
[30] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
[31] M. Koestinger, P. Wohlhart, P. M. Roth, and H. Bischof. Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)*, pages 2144–2151. IEEE, 2011.
[32] Konica Minolta. Konica minolta range 7 3d scanner, 2018.
[33] S. Laine, T. Karras, T. Aila, A. Herva, S. Saito, R. Yu, H. Li, and J. Lehtinen. Production-level facial performance capture using deep convolutional neural networks. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 10. ACM, 2017.
[34] H. Li, J. Yu, Y. Ye, and C. Bregler. Realtime facial animation with on-the-fly correctives. *ACM Transactions on Graphics (TOG)*, 32(4), 2013.
[35] J. Li, W. Xu, Z. Cheng, K. Xu, and R. Klein. Lightweight wrinkle synthesis for 3d facial modeling and animation. *Computer-Aided Design*, 58:117–122, 2015.
[36] M. Li, W. Zuo, and D. Zhang. Convolutional network for attribute-driven and identity-preserving human face generation. *CoRR*, abs/1608.06434, 2016.
[37] T. Li, T. Bolkart, M. J. Black, H. Li, and J. Romero. Learning a model of facial shape and expression from 4d scans. *ACM Transactions on Graphics (ToG)*, 36(6):194, 2017.
[38] S. Liang, I. Kemelmacher-Shlizerman, and L. G. Shapiro. 3d face hallucination from a single depth frame. In *2nd International Conference on 3D Vision, 3DV 2014, Tokyo, Japan, December 8-11, 2014, Volume 1*, pages 31–38, 2014.
[39] F. Liu, L. Tran, and X. Liu. 3d face modeling from diverse raw scan data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9408–9418, 2019.
[40] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

[41] M. Lüthi, T. Gerig, C. Jud, and T. Vetter. Gaussian process morphable models. *IEEE transactions on pattern analysis and machine intelligence*, 40(8):1860–1873, 2017.

[42] R. A. Newcombe, D. Fox, and S. M. Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 343–352, 2015.

[43] Y. Nirkin, I. Masi, A. T. Tran, T. Hassner, and G. G. Medioni. On face segmentation, face swapping, and face perception. In *IEEE Conference on Automatic Face and Gesture Recognition*, 2018.

[44] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015.

[45] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.

[46] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter. A 3d face model for pose and illumination invariant face recognition. In *IEEE International Conference on Advanced video and signal based surveillance*, pages 296–301, 2009.

[47] P. J. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. K. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2005.

[48] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.

[49] A. Ranjan, T. Bolkart, S. Sanyal, and M. J. Black. Generating 3D faces using convolutional mesh autoencoders. In *European Conference on Computer Vision (ECCV)*, 2018.

[50] E. Richardson, M. Sela, and R. Kimmel. 3d face reconstruction by learning from synthetic data. In *International Conference on 3D Vision (3DV)*, pages 460–469. IEEE, 2016.

[51] E. Richardson, M. Sela, R. Or-El, and R. Kimmel. Learning detailed face reconstruction from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition*, July 2017.

[52] S. Romdhani and T. Vetter. Estimating 3d shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 986–993, 2005.

[53] S. Sanyal, T. Bolkart, H. Feng, and M. Black. Learning to regress 3D face shape and expression from an image without 3D supervision. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[54] S. Sengupta, A. Kanazawa, C. D. Castillo, and D. W. Jacobs. Sfsnet: Learning shape, refectance and illuminance of faces in the wild. In *Computer Vision and Pattern Regognition (CVPR)*, 2018.

[55] F. Shi, H.-T. Wu, X. Tong, and J. Chai. Automatic acquisition of high-fidelity facial performances using monocular videos. *ACM Transactions on Graphics (TOG)*, 33(6):222, 2014.

[56] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. *ACM Transactions on Graphics (TOG)*, 23(3):399–405, 2004.

[57] S. Suwajanakorn, I. Kemelmacher-Shlizerman, and S. M. Seitz. Total moving face reconstruction. In *ECCV*, pages 796–812, 2014.

[58] A. Tewari, F. Bernard, P. Garrido, G. Bharaj, M. Elgharib, H.-P. Seidel, P. Pérez, M. Zöllhofer, and C. Theobalt. Fml: Face model learning from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[59] A. Tewari, M. Zollhöfer, F. Bernard, P. Garrido, H. Kim, P. Pérez, and C. Theobalt. High-fidelity monocular face reconstruction based on an unsupervised model-based face autoencoder. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):357–370, 2020.

[60] A. Tewari, M. Zollhöfer, P. Garrido, F. Bernard, H. Kim, P. Pérez, and C. Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 hz. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[61] A. Tewari, M. Zollöfer, H. Kim, P. Garrido, F. Bernard, P. Perez, and T. Christian. MoFA: Model-based Deep Convolutional Face Autoencoder for Unsupervised Monocular Reconstruction. In *IEEE International Conference on Computer Vision*, 2017.

[62] J. Thies, M. Zollhöfer, M. Nießner, L. Valgaerts, M. Stamminger, and C. Theobalt. Real-time expression transfer for facial reenactment. *ACM Transactions on Graphics (TOG)*, 34(6):183:1–183:14, 2015.

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

[64] A. T. Tran, T. Hassner, I. Masi, and G. G. Medioni. Regressing robust and discriminative 3d morphable models with a very deep neural network. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1493–1502, 2017.

[65] L. Tran, F. Liu, and X. Liu. Towards high-fidelity nonlinear 3d face morphable model. In *In Proceeding of IEEE Computer Vision and Pattern Recognition*, 2019.

[66] L. Tran and X. Liu. Nonlinear 3d face morphable model. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, UT, June 2018.

[67] L. Valgaerts, C. Wu, A. Bruhn, H. Seidel, and C. Theobalt. Lightweight binocular facial performance capture under uncontrolled lighting. *ACM Transactions on graphics (TOG)*, 31(6):187:1–187:11, 2012.

[68] D. Vlasic, M. Brand, H. Pfister, and J. Popovic. Face transfer with multilinear models. *ACM Trans. Graph.*, 24(3):426–433, 2005.

[69] T. Weise, S. Bouaziz, H. Li, and M. Pauly. Realtime performance-based facial animation. In *ACM transactions on graphics (TOG)*, volume 30, page 77. ACM, 2011.

[70] Y. Weng, C. Cao, Q. Hou, and K. Zhou. Real-time facial animation on mobile devices. *Graphical Models*, 76(3):172–179, 2014.

[71] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 7663–7672, 2019.

[72] 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)*, pages 211–216. IEEE, 2006.

[73] R. Yu, S. Saito, H. Li, D. Ceylan, and H. Li. Learning dense facial correspondences in unconstrained images. In *IEEE International Conference on Computer Vision*, pages 4733–4742, 2017.

[74] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li. Face alignment across large poses: A 3d solution. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 146–155, 2016.