

# Towards Metrical Reconstruction of Human Faces

Wojciech Zielonka, Timo Bolkart, and Justus Thies

Max Planck Institute for Intelligent Systems, Tübingen

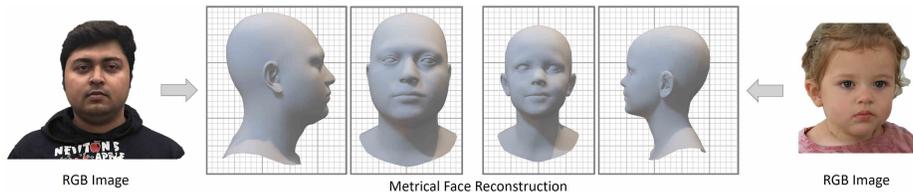


Fig. 1: An RGB image of a subject serves as input to MICA, which predicts a metrical reconstruction of the human face. Images from NoW [59], StyleGan2 [38].

**Abstract.** Face reconstruction and tracking is a building block of numerous applications in AR/VR, human-machine interaction, as well as medical applications. Most of these applications rely on a metrically correct prediction of the shape, especially, when the reconstructed subject is put into a metrical context (i.e., when there is a reference object of known size). A metrical reconstruction is also needed for any application that measures distances and dimensions of the subject (e.g., to virtually fit a glasses frame). State-of-the-art methods for face reconstruction from a single image are trained on large 2D image datasets in a self-supervised fashion. However, due to the nature of a perspective projection they are not able to reconstruct the actual face dimensions, and even predicting the average human face outperforms some of these methods in a metrical sense. To learn the actual shape of a face, we argue for a supervised training scheme. Since there exists no large-scale 3D dataset for this task, we annotated and unified small- and medium-scale databases. The resulting unified dataset is still a medium-scale dataset with more than 2k identities and training purely on it would lead to overfitting. To this end, we take advantage of a face recognition network pretrained on a large-scale 2D image dataset, which provides distinct features for different faces and is robust to expression, illumination, and camera changes. Using these features, we train our face shape estimator in a supervised fashion, inheriting the robustness and generalization of the face recognition network. Our method, which we call MICA (Metric fAce), outperforms the state-of-the-art reconstruction methods by a large margin, both on current non-metric benchmarks as well as on our metric benchmarks (15% and 24% lower average error on NoW, respectively).

**Project website:** <https://zielon.github.io/mica/>

## 1 Introduction

Learning to reconstruct 3D content from 2D imagery is an ill-posed inverse problem [4]. State-of-the-art RGB-based monocular facial reconstruction and tracking methods [18, 23] are based on self-supervised training, exploiting an underlying metrical face model which is constructed using a large-scale dataset of registered 3D scans (e.g., 33000 scans for the FLAME [47] model). However, when assuming a perspective camera, the scale of the face is ambiguous since a large face can be modeled by a small face that is close to the camera or a gigantic face that is far away. Formally, a point  $\mathbf{x} \in \mathbb{R}^3$  of the face is projected to a point  $\mathbf{p} \in \mathbb{R}^2$  on the image plane with the projective function  $\pi(\cdot)$  and a rigid transformation composed of a rotation  $\mathbf{R} \in \mathbb{R}^{3 \times 3}$  and a translation  $\mathbf{t} \in \mathbb{R}^3$ :

$$\mathbf{p} = \pi(\mathbf{R} \cdot \mathbf{x} + \mathbf{t}) = \pi(s \cdot (\mathbf{R} \cdot \mathbf{x} + \mathbf{t})) = \pi(\mathbf{R} \cdot (s \cdot \mathbf{x}) + (s \cdot \mathbf{t})).$$

The perspective projection is invariant to the scaling factor  $s \in \mathbb{R}$ , and thus, if  $\mathbf{x}$  is scaled by  $s$ , the rigid transformation can be adapted such that the point still projects onto the same pixel position  $\mathbf{p}$  by scaling the translation  $\mathbf{t}$  by  $s$ . In consequence, face reconstruction methods might result in a good 2D alignment but can fail to reconstruct the metrical 3D surface and the meaningful metrical location in space. However, a metric 3D reconstruction is needed in any scenario where the face is put into a metric context. E.g., when the reconstructed human is inserted into a virtual reality (VR) application or when the reconstructed geometry is used for augmented reality (AR) applications (teleconferencing in AR/VR, virtual try-on, etc.). In these scenarios, the methods mentioned above fail since they do not reproduce the correct scale and shape of the human face. In the current literature [25, 59, 83], we also observe that methods use evaluation measurements not done in a metrical space. Specifically, to compare a reconstructed face to a reference scan, the estimation is aligned to the scan via Procrustes analysis, including an optimal scaling factor. This scaling factor favors the estimation methods that are not metrical, and the reported numbers in the publications are misleading for real-world applications (relative vs. absolute/metrical error). In contrast, we aim for a metrically correct reconstruction and evaluation that directly compares the predicted geometry to the reference data without any scaling applied in a post-processing step which is fundamentally different. As discussed above, the self-supervised methods in the literature do not aim and cannot reconstruct a metrically correct geometry. However, training these methods in a supervised fashion is not possible because of the lack of data (no large-scale 3D dataset is available). Training on a small- or medium-scale 3D dataset will lead to overfitting of the networks (see study in the supplemental document). To this end, we propose a hybrid method that can be trained on a medium-scale 3D dataset, reusing powerful descriptors from a pretrained face recognition network (trained on a large-scale 2D dataset). Specifically, we propose the usage of existing 3D datasets like LYHM [16], FaceWarehouse [10], Stirling [26], etc., that contain RGB imagery and corresponding 3D reconstructions to learn a metrical reconstruction of the human head. To use these 3D datasets, significant work

has been invested to unify the 3D data (i.e., to annotate and non-rigidly fit the FLAME model to the different datasets). This unification provides us with meshes that all share the FLAME topology. Our method predicts the head geometry in a neutral expression, only given a single RGB image of a human subject in any pose or expression. To generalize to unseen in the wild images, we use a state-of-the-art face recognition network [17] that provides a feature descriptor for our geometry-estimating network. This recognition network is robust to head poses, different facial expressions, occlusions, illumination changes, and different focal lengths, thus, being ideal for our task (see Figure 3). Based on this feature, we predict the geometry of the face with neutral expression within the face space spanned by FLAME [47], effectively disentangling shape and expression. As an application, we demonstrate that our metrical face reconstruction estimator can be integrated in a new analysis-by-synthesis face tracking framework which removes the requirement of an identity initialization phase [70]. Given the metrical face shape estimation, the face tracker is able to predict the face motion in a metrical space.

In summary, we have the following contributions:

- a dataset of 3D face reference data for about 2300 subjects, built by unifying existing small- and medium-scale datasets under common FLAME topology.
- a metrical face shape predictor – MICA – which is invariant to expression, pose and illumination, by exploiting generalized identity features from a face recognition network and supervised learning.
- a hybrid face tracker that is based on our (learned) metrical reconstruction of the face shape and an optimization-based facial expression tracking.
- a metrical evaluation protocol and benchmark, including a discussion on the current evaluation practise.

## 2 Related Work

Reconstructing human faces and heads from monocular RGB, RGB-D, or multi-view data is a well-explored field at the intersection of computer vision and computer graphics. Zollhöfer et al. [85] provide an extensive review of reconstruction methods, focusing on optimization-based techniques that follow the principle of analysis-by-synthesis. Primarily, the approaches that are based on monocular inputs are based on a prior of face shape and appearance [6, 7, 27, 28, 40, 66–71, 77, 78]. The seminal work of Blanz et al. [8] introduced such a 3D morphable model (3DMM), which represents the shape and appearance of a human in a compressed, low-dimensional, PCA-based space (which can be interpreted as a decoder with a single linear layer). There is a large corpus of different morphable models [21], but the majority of reconstruction methods use either the Basel Face Model [8, 52] or the Flame head model [47]. Besides using these models for an analysis-by-synthesis approach, there is a series of learned regression-based methods. An overview of these methods is given by Morales et al. [50]. In the following, we will discuss the most relevant related work for monocular RGB-based reconstruction methods.

*Optimization-based Reconstruction of Human Faces.* Along with the introduction of a 3D morphable model for faces, Blanz et al. [8] proposed an optimization-based reconstruction method that is based on the principle of analysis-by-synthesis. While they used a sparse sampling scheme to optimize the color reproduction, Thies et al. [69, 70] introduced a dense color term considering the entire face region that is represented by a morphable model using differentiable rendering. This method has been adapted for avatar digitization from a single image [36] including hair, is used to reconstruct high-fidelity facial reflectance and geometry from a single images [79], for reconstruction and animation of entire upper bodies [71], or avatars with dynamic textures [51]. Recently, these optimization-based methods are combined with learnable components such as surface offsets or view-dependent surface radiance fields [32]. In addition to a photometric reconstruction objective, additional terms based on dense correspondence [35] or normal [1, 32] estimations of neural network can be employed. Optimization-based methods are also used as a building block for neural rendering methods such as deep video portraits [40], deferred neural rendering [68], or neural voice puppetry [67]. Note that differentiable rendering is not only used in neural rendering frameworks but is also a key component for self-supervised learning of regression-based reconstruction methods covered in the following.

*Regression-based Reconstruction of Human Faces.* Learning-based face reconstruction methods can be categorized into supervised and self-supervised approaches. A series of methods are based on synthetic renderings of human faces to perform a supervised training of a regressor that predicts the parameters of a 3D morphable model [20, 41, 56, 57]. Genova et al. [31] propose a 3DMM parameter regression technique that is based on synthetic renderings (where ground truth parameters are available) and real images (where multi-view identity losses are applied). It uses FaceNet [60] to extract features for the 3DMM regression task. Tran et al. [72] and Chang et al. [11] (ExpNet) directly regress 3DMM parameters using a CNN trained on fitted 3DMM data. Tu et al. [75] propose a dual training pass for images with and without 3DMM fittings. Jackson et al. [37] propose a model-free approach that reconstructs a voxel-based representation of the human face and is trained on paired 2D image and 3D scan data. PRN [24] is trained on 'in-the-wild' images with fitted 3DMM reconstructions [84]. It is not restricted to a 3DMM model space and predicts a position map in the UV-space of a template mesh. Instead of working in UV-space, Wei et al. [76] propose to use graph convolutions to regress the coordinates of the vertices. MoFA [65] is a network trained to regress the 3DMM parameters in a self-supervised fashion. As a supervision signal, it uses the dense photometric losses of Face2Face [70]. Within this framework, Tewari et al. proposed to refine the identity shape and appearance [64] as well as the expression basis [63] of a linear 3DMM. In a similar setup, one can also train a non-linear 3DMM [74] or personalized models [12]. RingNet [59] regresses 3DMM parameters and is trained on 2D images using losses on the reproduction of 2D landmarks and shape consistency (different images of the same subject) and shape inconsistency (images of different subjects) losses. DECA [23] extends RingNet with

expression dependent offset predictions in UV space. It uses dense photometric losses to train the 3DMM parameter regression and the offset prediction network. This separation of a coarse 3DMM model and a detailed bump map has been introduced by Tran et al. [73]. Chen et al. [13] use a hybrid training composed of self-supervised and supervised training based on renderings to predict texture and displacement maps. Deng et al. [18] train a 3DMM parameter regressor based on multi-image consistency losses and 'hybrid-level' losses (photometric reconstruction loss with skin attention masks, and a perception-level loss based on FaceNet [60]). On the NoW challenge [59], DECA [23] and the method of Deng et al. [18] show on-par state-of-the-art results. Similar to DECA's offset prediction, there are GAN-based methods that predict detailed color maps [29, 30] or skin properties [44, 45, 58, 79] (e.g., albedo, reflectance, normals) in UV-space of a 3DMM-based face reconstruction. In contrast to these methods, we are interested in reconstructing a metrical 3D representation of a human face and not fine-scale details. Self-supervised methods suffer from the depth-scale ambiguity (the face scale, translation away from the camera, and the perspective projection are ambiguous) and, thus, predict a wrongly scaled face, even though 3DMM models are by construction in a metrical space. We rely on a strong supervision signal to learn the metrical reconstruction of a face using high-quality 3D scan datasets which we unified. In combination with an identity encoder [17] trained on in-the-wild 2D data, including occlusions, different illumination, poses, and expressions, we achieve robust geometry estimations that significantly outperform state-of-the-art methods.

### 3 Metrical Face Shape Prediction

Based on a single input RGB image  $I$ , MICA aims to predict a metrical shape of a human face in a neutral expression. To this end, we leverage both 'in-the-wild' 2D data as well as metric 3D data to train a deep neural network, as shown in Figure 2. We employ a state-of-the-art face recognition network [17] which is trained on 'in-the-wild' data to achieve a robust prediction of an identity code, which is interpreted by a geometry decoder.

*Identity Encoder.* As an identity encoder, we leverage the ArcFace [17] architecture which is pretrained on Glint360K [2]. This ResNet100-based network is trained on 2D image data using an additive angular margin loss to obtain highly discriminative features for face recognition. It is invariant to illumination, expression, rotation, occlusion, and camera parameters which is ideal for a robust shape prediction. We extend the ArcFace architecture by a small mapping network  $\mathcal{M}$  that maps the ArcFace features to our latent space, which can then be interpreted by our geometry decoder:

$$\mathbf{z} = \mathcal{M}(\text{ArcFace}(I)),$$

where  $\mathbf{z} \in \mathbb{R}^{300}$ . Our mapping network  $\mathcal{M}$  consists of three fully-connected linear hidden layers with ReLU activation and the final linear output layer.

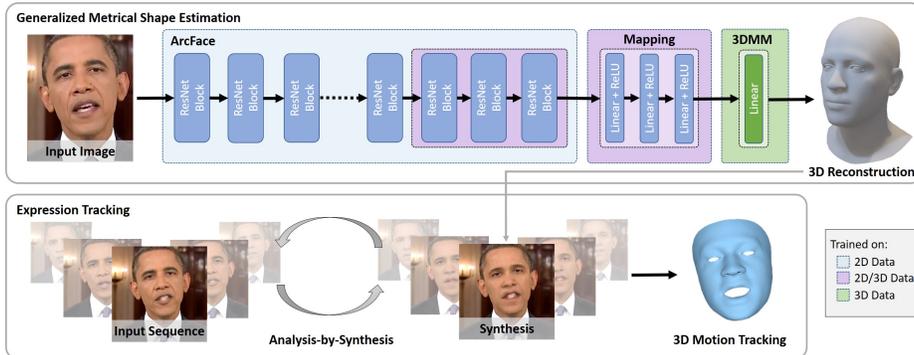


Fig. 2: We propose a method for metrical human face shape estimation from a single image which exploits a supervised training scheme based on a mixture of different 2D, 2D/3D and 3D datasets. This estimation can be used for facial expression tracking using analysis-by-synthesis which optimizes for the camera intrinsics, as well as the per-frame illumination, facial expression and pose.

*Geometry Decoder.* There are essentially two types of geometry decoders used in the literature, model-free and model-based. Throughout the project of this paper, we conducted experiments on both types and found that both perform similarly on the evaluation benchmarks. Since a 3DMM model efficiently represents the face space, we focus on a model-based decoder. Specifically, we use FLAME [47] as a geometry decoder, which consists of a single linear layer:

$$\mathcal{G}_{3DMM}(z) = \mathbf{B} \cdot z + \mathbf{A},$$

where  $\mathbf{A} \in \mathbb{R}^{3N}$  is the geometry of the average human face and  $\mathbf{B} \in \mathbb{R}^{3N \times 300}$  contains the principal components of the 3DMM and  $N = 5023$ .

*Supervised Learning.* The networks described above are trained using paired 2D/3D data from existing, unified datasets  $\mathcal{D}$  (see Section 5). We fix large portions of the pre-trained ArcFace network during the training and refine the last 3 ResNet blocks. Note that ArcFace is trained on a much larger amount of identities, therefore, refining more hidden layers results in worse predictions due to overfitting. We found that using the last 3 ResNet blocks gives the best generalization (see supplemental document). The training loss is:

$$\mathcal{L} = \sum_{(I, \mathcal{G}) \in \mathcal{D}} |\kappa_{mask}(\mathcal{G}_{3DMM}(\mathcal{M}(\text{ArcFace}(I))) - \mathcal{G})|, \quad (1)$$

where  $\mathcal{G}$  is the ground truth mesh and  $\kappa_{mask}$  is a region dependent weight (the face region has weight 150.0, the back of the head 1.0, and eyes with ears 0.01). We use AdamW [49] for optimization with fixed learning rate  $\eta = 1e-5$  and weight decay  $\lambda = 2e-4$ . We select the best performing model based on the validation set loss using the Florence dataset [3]. The model was trained for 160k steps on Nvidia Tesla V100.

## 4 Face Tracking

Based on our shape estimate, we demonstrate optimization-based face tracking on monocular RGB input sequences. To model the non-rigid deformations of the face, we use the linear expression basis vectors and the linear blendskinning of the FLAME [47] model, and use a linear albedo model [22] to reproduce the appearance of a subject in conjunction with a Lambertian material assumption and a light model based on spherical harmonics. We adapt the analysis-by-synthesis scheme of Thies et al. [70]. Instead of using a multi-frame model-based bundling technique to estimate the identity of a subject, we use our one-shot shape identity predictor. We initialize the albedo and spherical harmonics based on the same first frame using the energy:

$$E(\phi) = w_{dense}E_{dense}(\phi) + w_{lmk}E_{lmk}(\phi) + w_{reg}E_{reg}(\phi), \quad (2)$$

where  $\phi$  is the vector of unknown parameters we are optimizing for. The energy terms  $E_{dense}(\phi)$  and  $E_{reg}(\phi)$  measure the dense color reproduction of the face ( $\ell_1$ -norm) and the deviation from the neutral pose respectively. The sparse landmark term  $E_{lmk}(\phi)$  measures the reproduction of 2D landmark positions (based on Google’s mediapipe [33, 39] and Face Alignment [9]). The weights  $w_{dense}$ ,  $w_{lmk}$  and  $w_{reg}$  balance the influence of each sub-objectives on the final loss. In the first frame vector  $\phi$  contains the 3DMM parameters for albedo, expression, and rigid pose, as well as the spherical harmonic coefficients (3 bands) that are used to represent the environmental illumination [54]. After initialization, the albedo parameters are fixed and unchanged throughout the sequence tracking.

*Optimization.* We optimize the objective function Equation (2) using Adam [42] in PyTorch. While recent soft-rasterizers [48, 55] are popular, we rely on a sampling based scheme as introduced by Thies et al. [70] to implement the differentiable rendering for the photo-metric reproduction error  $E_{dense}(\phi)$ . Specifically, we use a classical rasterizer to render the surface of the current estimation. The rasterized surface points that survive the depth test are considered as the set of visible surface points  $\mathcal{V}$  for which we compute the energy term  $E_{dense}(\phi) = \sum_{i \in \mathcal{V}} |I(\pi(\mathbf{R} \cdot p_i(\phi) + \mathbf{t})) - c_i(\phi)|$  where  $p_i$  and  $c_i$  being the  $i$ -th vertex and color of the reconstructed model, and  $I$  the RGB input image.

## 5 Dataset Unification

In the past, methods and their training scheme were limited by the availability of 3D scan datasets of human faces. While several small and medium-scale datasets are available, they are in different formats and do not share the same topology. To this end, we unified the available datasets such that they can be used as a supervision signal for face reconstruction from 2D images. Specifically, we register the FLAME [47] head model to the provided scan data. In an initial step, we fit the model to landmarks and optimize for the FLAME parameters based on an iterative closest point (ICP) scheme [5]. We further jointly optimize FLAME’s

Table 1: Overview of our unified datasets. The used datasets vary in the capture modality and the capture protocol. Here, we list the number of subject, the minimum number of images per subjects, and whether the dataset includes facial expressions. In total our dataset contains 2315 subjects with FLAME topology.

Dataset		#Subj.	#Min. Img.	Expr.
Stirling [26]		133	8	✓
D3DFACS [15]		10	videos	✓
Florence 2D/3D [3]		53	videos	✓
BU-3DFE [81]		100	83	✓
LYHM [16]		1211	2	✗
FaceWarehouse [10]		150	119	✓
FRGC [53]		531	7	✓
BP4D+ [82]		127	videos	✓

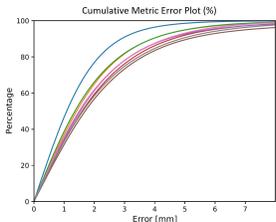
model parameters, and refine the fitting with a non-rigid deformation regularized by FLAME, similar to Li and Bolkart et al. [47]. In Table 1, we list the datasets that we unified for this project. We note that the datasets vary in the capturing modality and capturing script (with and without facial expressions, with and without hair caps, indoor and outdoor imagery, still images, and videos), which is suitable for generalization. The datasets are recorded in different regions of the world and are often biased towards ethnicity. Thus, combining other datasets results in a more diverse data pool. In the supplemental document, we show an ablation on the different datasets. *Upon agreement of the different dataset owners, we will share our unified dataset, i.e., for each subject one registered mesh with neutral expression in FLAME topology.* Note that in addition to the datasets listed in Table 1, we analyzed the FaceScape dataset [80]. While it provides a large set of 3D reconstructions ( $\sim 17k$ ), which would be ideal for our training, the reconstructions are not done in a metrical space. Specifically, the data has been captured in an uncalibrated setup and faces are normalized by the eye distance, which has not been detailed in their paper (instead, they mention sub-millimeter reconstruction accuracy which is not valid). This is a fundamental flaw of this dataset, and also questions their reconstruction benchmark [83].

## 6 Results

Our experiments mainly focus on the metrical reconstruction of a human face from in the wild images. In the supplemental document, we show results for the sequential tracking of facial motions using our metrical reconstruction as initialization. The following experiments are conducted with the original models of the respective publications including their reconstructions submitted to the given benchmarks. Note that these models are trained on their large-scale datasets, training them on our medium-scale 3D dataset would lead to overfitting.

Table 2: Quantitative evaluation of the face shape estimation on the *NoW Challenge* [59]. Note that we list two different evaluations: the non-metrical evaluation from the original NoW challenge and our new metrical evaluation (including a cumulative error plot on the left). The original NoW challenge cannot be considered metrical since Procrustes analysis is used to align the reconstructions to the corresponding reference meshes, including scaling. We list all methods from the original benchmark and additionally show the performance of the average human face of FLAME [47] as a reference (first row).

NoW-Metric Challenge		Non-Metrical [59]			Metrical (mm)		
Method		Median	Mean	Std	Median	Mean	Std
Average Face (FLAME [47])		1.21	1.53	1.31	1.49	1.92	1.68
3DMM-CNN [72]		1.84	2.33	2.05	3.91	4.84	4.02
PRNet [24]		1.50	1.98	1.88	-	-	-
Deng et al [18] (TensorFlow)		1.23	1.54	1.29	2.26	2.90	2.51
● Deng et al [18] (PyTorch)		1.11	1.41	1.21	1.62	2.21	2.08
● RingNet [59]		1.21	1.53	1.31	1.50	1.98	1.77
● 3DDFA-V2 [34]		1.23	1.57	1.39	1.53	2.06	1.95
● MGCNet [62]		1.31	1.87	2.63	1.70	2.47	3.02
● UMDFA [43]		1.52	1.89	1.57	2.31	2.97	2.57
● Dib et al. [19]		1.26	1.57	1.31	1.59	2.12	1.93
● DECA [23]		1.09	1.38	1.18	1.35	1.80	1.64
● FOCUS [46]		1.04	1.30	1.10	1.41	1.85	1.70
● Ours		<b>0.90</b>	<b>1.11</b>	<b>0.92</b>	<b>1.08</b>	<b>1.37</b>	<b>1.17</b>



## 6.1 Face Shape Estimation

In recent publications, face shape estimation is evaluated on datasets where reference scans of the subjects are available. The NoW Challenge [59] and the benchmark of Feng et al. [25] which is based on Stirling meshes [26] are used in the state-of-the-art methods [18, 23, 59]. We conduct several studies on these benchmarks and propose different evaluation protocols.

**Non-Metrical Benchmark.** The established evaluation methods on these datasets are based on an optimal scaling step, i.e., to align the estimation to the reference scan, they optimize for a rigid alignment and an additional scaling factor which results in a non-metric/relative error. This scaling compensates for shape mispredictions, e.g., the mean error evaluated on the NoW Challenge for the average FLAME mesh (Table 2) drops from 1.92mm to 1.53mm because of the applied scale optimization. This is an improvement of around 20% which has nothing to do with the reconstruction quality and, thus, creates a misleading benchmark score where methods appear better than they are. Nevertheless, we evaluate our method on these benchmarks and significantly outperform all state-of-the-art methods as can be seen in Tables 2 and 4 (‘Non-Metrical’ column).

**Metrical Benchmark.** Since for a variety of applications, actual metrical reconstructions are required, we argue for a new evaluation scheme that uses a purely rigid alignment, i.e., without scale optimization (see Figure 5). The error is calculated using an Euclidean distance between each scan vertex and the closest point on the mesh surface. This new evaluation scheme enables a comparison of methods based on metrical quantities (see Tables 2 and 4) and, thus,

Table 3: Quantitative evaluation of the face shape estimation on the *Stirling Reconstruction Benchmark* [25] using the NoW protocol [59]. We list two different evaluations: the non-metric evaluation from the original benchmark and the metric evaluation. *Note that for this experiment, we exclude the Stirling dataset from our training set.*

Stirling (NoW Protocol)	Non-Metrical						Metrical (mm)					
	Median		Mean		Std		Median		Mean		Std	
	LQ	HQ	LQ	HQ	LQ	HQ	LQ	HQ	LQ	HQ	LQ	HQ
Average Face (FLAME [47])	1.23	1.22	1.56	1.55	1.38	1.35	1.44	1.40	1.84	1.79	1.64	1.57
RingNet [59]	1.17	1.15	1.49	1.46	1.31	1.27	1.37	1.33	1.77	1.72	1.60	1.54
3DDFA-V2 [34]	1.26	1.20	1.63	1.55	1.52	1.45	1.49	1.38	1.93	1.80	1.78	1.68
Deng et al. [18] (TensorFlow)	1.22	1.13	1.57	1.43	1.40	1.25	1.85	1.81	2.41	2.29	2.16	1.97
Deng et al. [18] (PyTorch)	1.12	0.99	1.44	1.27	1.31	1.15	1.47	1.31	1.93	1.71	1.77	1.57
DECA [23]	1.09	1.03	1.39	1.32	1.26	1.18	1.32	1.22	1.71	1.58	1.54	1.42
<b>Ours w/o. Stirling</b>	<b>0.96</b>	<b>0.92</b>	<b>1.22</b>	<b>1.16</b>	<b>1.11</b>	<b>1.04</b>	<b>1.15</b>	<b>1.06</b>	<b>1.46</b>	<b>1.35</b>	<b>1.30</b>	<b>1.20</b>

Table 4: Quantitative evaluation of the face shape estimation on the *Stirling Reconstruction Benchmark* [25]. We list two different evaluations: the non-metric evaluation from the original benchmark and the metric evaluation. This benchmark is based on an alignment protocol that only relies on reference landmarks and, thus, is very noisy and dependent on the landmark reference selection (in our evaluation, we use the landmark correspondences provided by the FLAME [47] model). We use the image file list from [59] to compute the scores (i.e., excluding images where a face is not detectable). *Note that for this experiment, we exclude the Stirling dataset from our training set.*

Stirling/ESRC Benchmark	Non-Metrical [25]						Metrical (mm)					
	Median		Mean		Std		Median		Mean		Std	
	LQ	HQ	LQ	HQ	LQ	HQ	LQ	HQ	LQ	HQ	LQ	HQ
Average Face (FLAME [47])	1.58	1.62	2.06	2.08	1.82	1.83	1.70	1.62	2.19	2.09	1.96	1.85
RingNet [59]	1.56	1.60	2.01	2.05	1.75	1.76	1.67	1.64	2.16	2.09	1.90	1.81
3DDFA-V2 [34]	1.58	1.49	2.03	1.90	1.74	1.63	1.70	1.56	2.16	1.98	1.88	1.70
Deng et al. [18] (TensorFlow)	1.56	1.41	2.02	1.84	1.77	1.63	2.13	2.14	2.71	2.65	2.33	2.12
Deng et al. [18] (PyTorch)	1.51	1.29	1.95	1.64	1.71	1.39	1.78	1.54	2.28	1.97	1.97	1.68
DECA [23]	1.40	1.32	1.81	1.72	1.59	1.50	1.56	1.45	2.03	1.87	1.81	1.64
<b>Ours w/o. Stirling</b>	<b>1.26</b>	<b>1.22</b>	<b>1.62</b>	<b>1.55</b>	<b>1.41</b>	<b>1.34</b>	<b>1.36</b>	<b>1.26</b>	<b>1.73</b>	<b>1.60</b>	<b>1.48</b>	<b>1.37</b>

is *fundamentally* different from the previous evaluation schemes. In addition, the benchmark of Feng et al. [25] is based on the alignment using sparse facial (hand-selected) landmarks. Our experiments showed that this scheme is highly dependent on the selection of these markers and results in inconsistent evaluation results. In our listed results, we use the marker correspondences that come with the FLAME model [47]. To get a more reliable evaluation scheme, we evaluate the benchmark of Feng et al. using the dense iterative closest point (ICP) technique from the NoW challenge, see Table 3. On all metrics, our proposed method significantly improves the reconstruction accuracy. Note that some methods are even performing worse than the mean face [47].

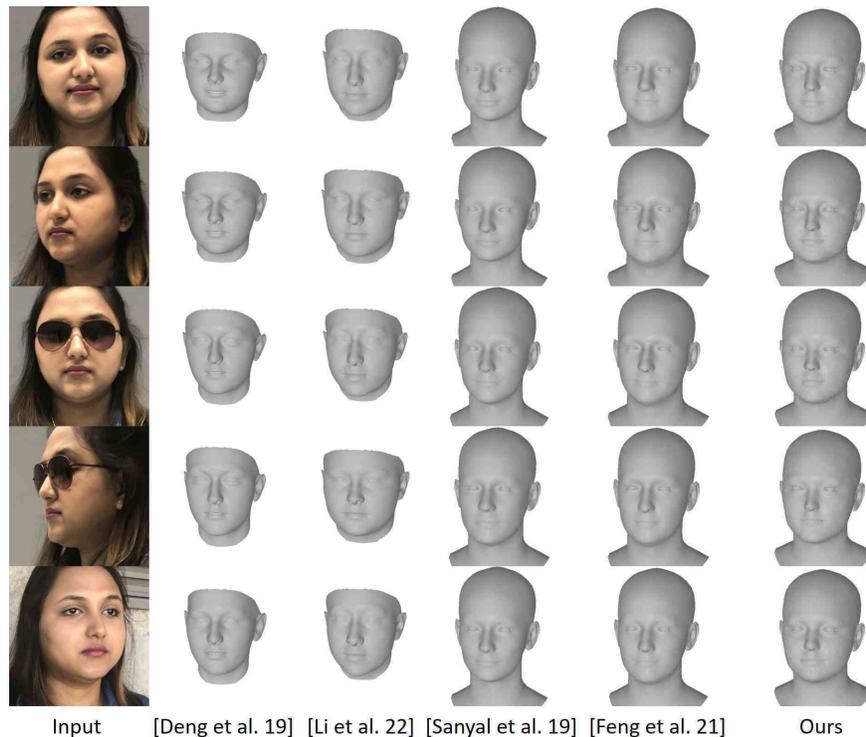


Fig. 3: Qualitative results on NoW Challenge [59] to show the invariance of our method to changes in illumination, expression, occlusion, rotation, and perspective distortion in comparison to other methods.

**Qualitative Results.** In Figure 3, we show qualitative results to analyze the stability of the face shape prediction of a subject across different expressions, head rotation, occlusions, or perspective distortion. As can be seen, our method is more persistent compared to others, especially, in comparison to Deng et al. [18] where shape predictions vary the most. Figure 4 depicts the challenging scenario of reconstructing toddlers from single images. Instead of predicting a small face for a child, the state of the art methods are predicting faces of adults. In contrast, MICA predicts the shape of a child with a correct scale.

In Figure 6 reconstructions for randomly sampled identities from the VoxCeleb2 [14] dataset are shown. Some of the baselines, especially, RingNet [59], exhibits strong bias towards the mean human face. In contrast, our method is able to not only predict better overall shape but also to reconstruct challenging regions like nose or chin, even though the training dataset contains a much smaller identity and ethnicity pool. Note that while the reconstructions of the baseline methods look good under the projection, they are not metric as shown in Tables 2 and 4.

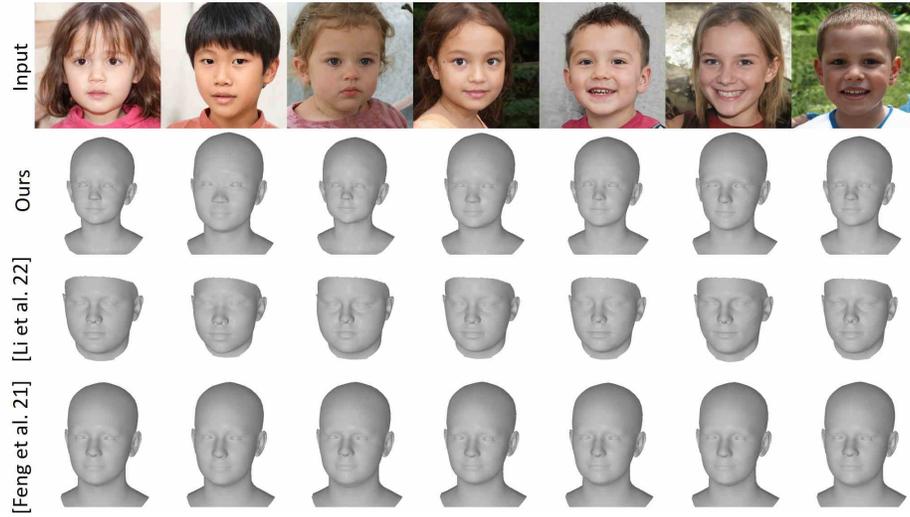


Fig. 4: Current methods are not predicting metrical faces, which becomes visible when displaying them in a metrical space and not in their image spaces. To illustrate we render the prediction of the faces of toddlers in a common metrical space using the same projection. State-of-the-art approaches trained in a self-supervised fashion like DECA [23] or weakly-supervised like FOCUS [46] scale the face of an adult to fit the observation in the image space, thus, the prediction in 3D is non-metrical. In contrast, our reconstruction method is able to recover the physiognomy of the toddlers. Input images are generated by StyleGan2 [38].

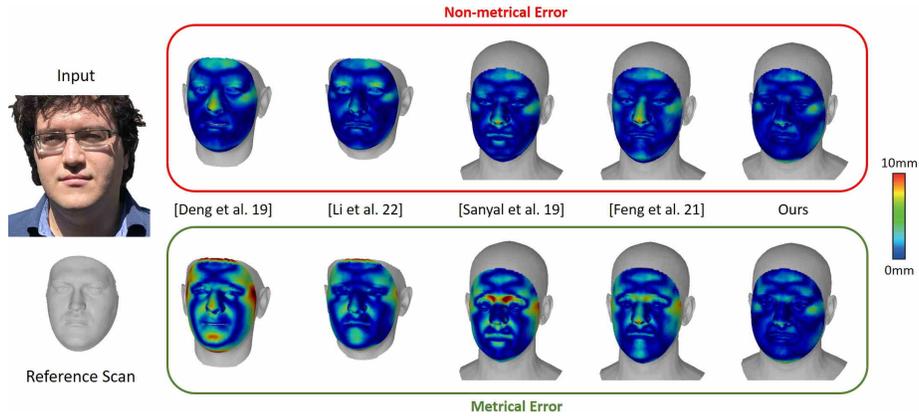


Fig. 5: Established evaluation benchmarks like [25, 59] are based on a non-metrical error metric (top-row). We propose a new evaluation protocol which measures reconstruction errors in a metrical space (bottom row) (c.f. Table 2). Image from the NoW [59] validation set.

## 6.2 Limitations

Our method is not designed to predict shape and expressions in one forward pass, instead, we reconstruct the expression separately using an optimization-based tracking method. However, this optimization-based tracking leads to temporally coherent results, as can be seen in the suppl. video. In contrast to DECA [23] or Deng et al. [18], the focus of our method is the reconstruction of a metrical 3D model, reconstructing high-frequent detail on top of our prediction is an interesting future direction. Our method fails, when the used face detector [61] does not recognize a face in the input.

## 7 Discussion & Conclusion

A metrical reconstruction is key for any application that requires the measurement of distances and dimensions. It is essential for the composition of reconstructed humans and scenes where objects of known size are in, thus, it is especially important for virtual reality and augmented reality applications. However, we show that recent methods and evaluation schemes are not designed for this task. While the established benchmarks report numbers in millimeters, they are computed with an optimal scale to align the prediction and the reference. We strongly argue against this practice, since it is misleading and the errors are not absolute metrical measurements. To this end, we propose a simple, yet fundamental adjustment of the benchmarks to enable metrical evaluations. Specifically, we remove the optimal scaling, and only allow rigid alignment of the prediction with the reference shape. As a stepping stone towards metrical reconstructions, we unified existing small- and medium-scale datasets of paired 2D/3D data. This allows us to establish 3D supervised losses in our novel shape prediction framework. While our data collection is still comparably small (around 2k identities), we designed MICA that uses features from a face recognition network pretrained on a large-scale 2D image dataset to generalize to in-the-wild image data. We validated our approach in several experiments and show state-of-the-art results on our newly introduced metrical benchmarks as well as on the established scale-invariant benchmarks. We hope that this work inspires researchers to concentrate on metrical face reconstruction.

*Acknowledgement.* We thank Haiwen Feng for support with NoW and Stirling evaluations, and Chunlu Li for providing FOCUS results. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Wojciech Zielonka.

*Disclosure.* While TB is part-time employee of Amazon, his research was performed solely at, and funded solely by MPI. JT is supported by Microsoft Research gift funds.

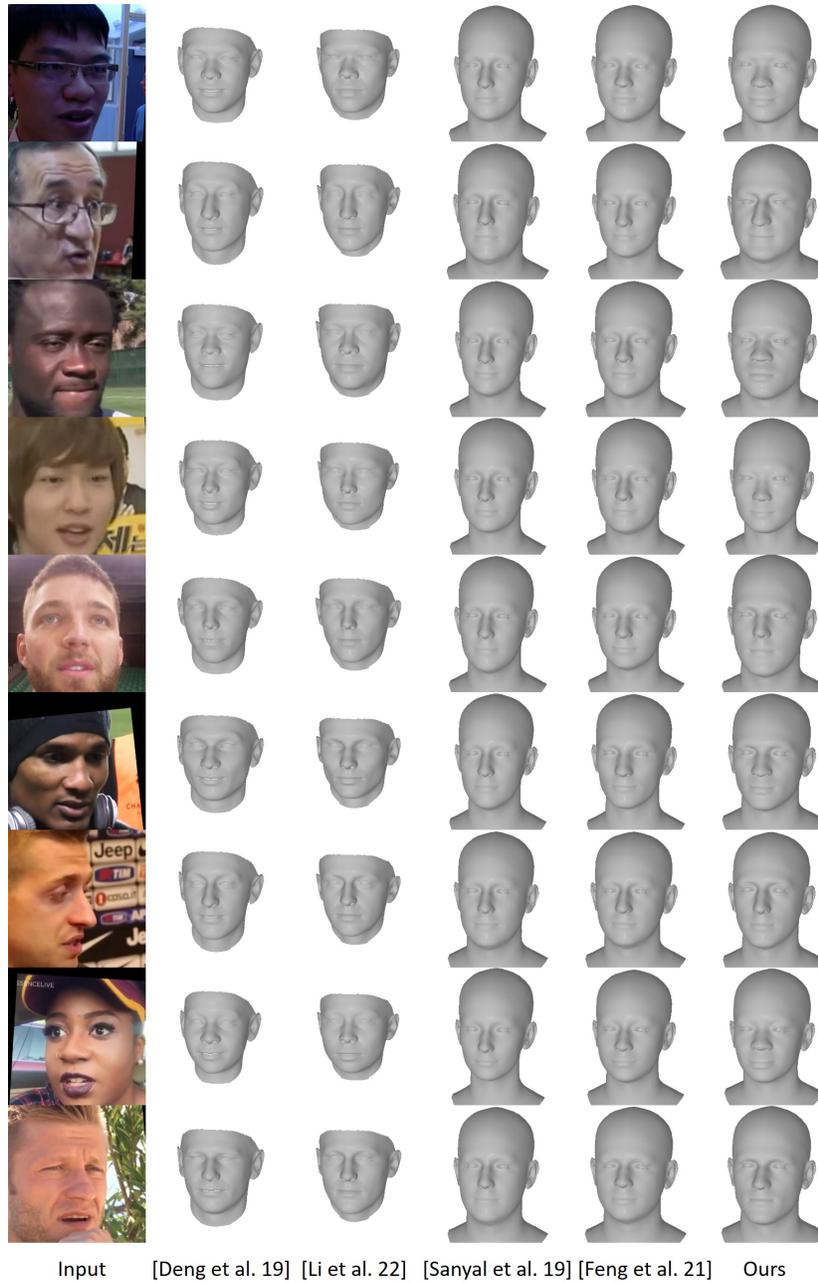


Fig. 6: Qualitative comparison on randomly sampled images from the VoxCeleb2 [14] dataset. Our method is able to capture face shape with intricate details like nose and chin, while being metrical plausible (c.f., Tables 2 and 4).

## Bibliography

- [1] Abrevaya, V.F., Boukhayma, A., Torr, P.H., Boyer, E.: Cross-modal deep face normals with deactivable skip connections. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4978–4988 (2020) [4](#)
- [2] An, X., Zhu, X., Xiao, Y., Wu, L., Zhang, M., Gao, Y., Qin, B., Zhang, D., Ying, F.: Partial fc: Training 10 million identities on a single machine. In: Arxiv 2010.05222 (2020) [5](#)
- [3] Bagdanov, A.D., Del Bimbo, A., Masi, I.: The florence 2D/3D hybrid face dataset. In: Proceedings of the 2011 Joint ACM Workshop on Human Gesture and Behavior Understanding. p. 79–80. J-HGBU '11, Association for Computing Machinery, New York, NY, USA (2011). <https://doi.org/10.1145/2072572.2072597>, <https://doi.org/10.1145/2072572.2072597> [6](#), [8](#)
- [4] Bas, A., Smith, W.A.P.: What does 2D geometric information really tell us about 3D face shape? International Journal of Computer Vision (IJCV) **127**(10), 1455–1473 (2019) [2](#)
- [5] Besl, P.J., McKay, N.D.: Method for registration of 3-d shapes. In: Sensor fusion IV: control paradigms and data structures. vol. 1611, pp. 586–606. International Society for Optics and Photonics (1992) [7](#)
- [6] Blanz, V., Basso, C., Poggio, T., Vetter, T.: Reanimating faces in images and video. In: EUROGRAPHICS (EG). vol. 22, pp. 641–650 (2003) [3](#)
- [7] Blanz, V., Scherbaum, K., Vetter, T., Seidel, H.P.: Exchanging faces in images. Computer Graphics Forum **23**(3), 669–676 (2004) [3](#)
- [8] Blanz, V., Vetter, T.: A morphable model for the synthesis of 3D faces. In: SIGGRAPH. pp. 187–194 (1999) [3](#), [4](#)
- [9] Bulat, A., Tzimiropoulos, G.: How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In: International Conference on Computer Vision (2017) [7](#)
- [10] Cao, C., Weng, Y., Zhou, S., Tong, Y., Zhou, K.: FaceWarehouse: A 3D facial expression database for visual computing. Transactions on Visualization and Computer Graphics **20**, 413–425 (01 2013) [2](#), [8](#)
- [11] Chang, F.J., Tran, A.T., Hassner, T., Masi, I., Nevatia, R., Medioni, G.: Expnet: Landmark-free, deep, 3d facial expressions. In: International Conference on Automatic Face & Gesture Recognition (FG). pp. 122–129 (2018) [4](#)
- [12] Chaudhuri, B., Vedapant, N., Shapiro, L., Wang, B.: Personalized face modeling for improved face reconstruction and motion retargeting (2020) [4](#)
- [13] Chen, A., Chen, Z., Zhang, G., Mitchell, K., Yu, J.: Photo-realistic facial details synthesis from single image. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9429–9439 (2019) [5](#)
- [14] Chung, J.S., Nagrani, A., Zisserman, A.: VoxCeleb2: Deep speaker recognition. In: "INTER\_SPEECH" (2018) [11](#), [14](#)

- [15] Cosker, D., Krumhuber, E., Hilton, A.: A faces valid 3d dynamic action unit database with applications to 3d dynamic morphable facial modeling. In: 2011 International Conference on Computer Vision. pp. 2296–2303 (2011). <https://doi.org/10.1109/ICCV.2011.6126510> **8**
- [16] Dai, H., Pears, N., Smith, W., Duncan, C.: Statistical modeling of cranio-facial shape and texture. *International Journal of Computer Vision (IJCV)* **128**(2), 547–571 (2019). <https://doi.org/10.1007/s11263-019-01260-7> **2, 8**
- [17] Deng, J., Guo, J., Liu, T., Gong, M., Zafeiriou, S.: Sub-center arcface: Boosting face recognition by large-scale noisy web faces. In: European Conference on Computer Vision (ECCV). pp. 741–757 (2020) **3, 5**
- [18] Deng, Y., Yang, J., Xu, S., Chen, D., Jia, Y., Tong, X.: Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In: Conference on Computer Vision and Pattern Recognition Workshops (CVPR-W) (2019) **2, 5, 9, 10, 11, 13**
- [19] Dib, A., Thebault, C., Ahn, J., Gosselin, P., Theobalt, C., Chevallier, L.: Towards high fidelity monocular face reconstruction with rich reflectance using self-supervised learning and ray tracing. In: International Conference on Computer Vision (ICCV). pp. 12819–12829 (2021) **9**
- [20] Dou, P., Shah, S.K., Kakadiaris, I.A.: End-to-end 3d face reconstruction with deep neural networks (2017) **4**
- [21] Egger, B., Smith, W.A.P., Tewari, A., Wuhler, S., Zollhoefer, M., Beeler, T., Bernard, F., Bolkart, T., Kortylewski, A., Romdhani, S., Theobalt, C., Blanz, V., Vetter, T.: 3D morphable face models - past, present and future. *Transactions on Graphics (TOG)* **39**(5) (2020). <https://doi.org/10.1145/3395208> **3**
- [22] Feng, H., Bolkart, T.: Photometric FLAME fitting (2020), [https://github.com/HavenFeng/photometric\\_optimization](https://github.com/HavenFeng/photometric_optimization) **7**
- [23] Feng, Y., Feng, H., Black, M.J., Bolkart, T.: Learning an animatable detailed 3D face model from in-the-wild images. *Transactions on Graphics, (Proc. SIGGRAPH)* **40**(8) (2021) **2, 4, 5, 9, 10, 12, 13**
- [24] Feng, Y., Wu, F., Shao, X., Wang, Y., Zhou, X.: Joint 3D face reconstruction and dense alignment with position map regression network. In: European Conference on Computer Vision (ECCV). pp. 534–551 (2018) **4, 9**
- [25] Feng, Z., Huber, P., Kittler, J., Hancock, P.J.B., Wu, X., Zhao, Q., Koppen, P., Räscht, M.: Evaluation of dense 3D reconstruction from 2D face images in the wild. In: International Conference on Automatic Face & Gesture Recognition (FG). pp. 780–786 (2018). <https://doi.org/10.1109/FG.2018.00123> **2, 9, 10, 12**
- [26] Feng, Z., Huber, P., Kittler, J., Hancock, P.J.B., Wu, X., Zhao, Q., Koppen, P., Räscht, M.: Evaluation of dense 3d reconstruction from 2d face images in the wild. *CoRR* **abs/1803.05536** (2018), <http://arxiv.org/abs/1803.05536> **2, 8, 9**
- [27] Garrido, P., Valgaerts, L., Rehmsen, O., Thormaehlen, T., Perez, P., Theobalt, C.: Automatic face reenactment. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4217–4224 (2014) **3**

- [28] Garrido, P., Valgaerts, L., Sarmadi, H., Steiner, I., Varanasi, K., Perez, P., Theobalt, C.: VDub - modifying face video of actors for plausible visual alignment to a dubbed audio track. In: EUROGRAPHICS (EG). pp. 193–204 (2015) [3](#)
- [29] Gecer, B., Ploumpis, S., Kotsia, I., Zafeiriou, S.: Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019) [5](#)
- [30] Gecer, B., Ploumpis, S., Kotsia, I., Zafeiriou, S.P.: Fast-ganfit: Generative adversarial network for high fidelity 3d face reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021) [5](#)
- [31] Genova, K., Cole, F., Maschinot, A., Sarna, A., Vlastic, D., Freeman, W.T.: Unsupervised training for 3d morphable model regression (2018) [4](#)
- [32] Grassal, P.W., Prinzler, M., Leistner, T., Rother, C., Nießner, M., Thies, J.: Neural head avatars from monocular rgb videos (2021) [4](#)
- [33] Grishchenko, I., Ablavatski, A., Kartynnik, Y., Raveendran, K., Grundmann, M.: Attention mesh: High-fidelity face mesh prediction in real-time (2020) [7](#)
- [34] Guo, J., Zhu, X., Yang, Y., Yang, F., Lei, Z., Li, S.Z.: Towards fast, accurate and stable 3d dense face alignment. In: European Conference on Computer Vision (ECCV) (2020) [9](#), [10](#)
- [35] Güler, R.A., Trigeorgis, G., Antonakos, E., Snape, P., Zafeiriou, S., Kokkinos, I.: Densereg: Fully convolutional dense shape regression in-the-wild (2017) [4](#)
- [36] Hu, L., Saito, S., Wei, L., Nagano, K., Seo, J., Fursund, J., Sadeghi, I., Sun, C., Chen, Y.C., Li, H.: Avatar digitization from a single image for real-time rendering. *ACM Trans. Graph.* **36**(6) (nov 2017). <https://doi.org/10.1145/3130800.31310887>, <https://doi.org/10.1145/3130800.31310887> [4](#)
- [37] Jackson, A.S., Bulat, A., Argyriou, V., Tzimiropoulos, G.: Large pose 3d face reconstruction from a single image via direct volumetric cnn regression (2017) [4](#)
- [38] Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of StyleGAN. In: Proc. CVPR (2020) [1](#), [12](#)
- [39] Kartynnik, Y., Ablavatski, A., Grishchenko, I., Grundmann, M.: Real-time facial surface geometry from monocular video on mobile gpus (2019) [7](#)
- [40] Kim, H., Garrido, P., Tewari, A., Xu, W., Thies, J., Niessner, M., Pérez, P., Richardt, C., Zollhöfer, M., Theobalt, C.: Deep video portraits. *Transactions on Graphics (TOG)* **37**(4), 1–14 (2018) [3](#), [4](#)
- [41] Kim, H., Zollhöfer, M., Tewari, A., Thies, J., Richardt, C., Theobalt, C.: InverseFaceNet: Deep monocular inverse face rendering. In: Conference on Computer Vision and Pattern Recognition (CVPR) (Jun 2018) [4](#)
- [42] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *CoRR abs/1412.6980* (2015) [7](#)

- [43] Koizumi, T., Smith, W.A.P.: "look ma, no landmarks!" - unsupervised, model-based dense face alignment. In: European Conference on Computer Vision (ECCV). vol. 12347, pp. 690–706 (2020) [9](#)
- [44] Lattas, A., Moschoglou, S., Gecer, B., Ploumpis, S., Triantafyllou, V., Ghosh, A., Zafeiriou, S.: AvatarMe: Realistically renderable 3D facial reconstruction" in-the-wild". In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 760–769 (2020) [5](#)
- [45] Lattas, A., Moschoglou, S., Ploumpis, S., Gecer, B., Ghosh, A., Zafeiriou, S.P.: AvatarMe++: Facial shape and BRDF inference with photorealistic rendering-aware GANs. Transactions on Pattern Analysis and Machine Intelligence (PAMI) (2021) [5](#)
- [46] Li, C., Morel-Forster, A., Vetter, T., Egger, B., Kortylewski, A.: To fit or not to fit: Model-based face reconstruction and occlusion segmentation from weak supervision. CoRR [abs/2106.09614](#) (2021), <https://arxiv.org/abs/2106.09614> [9](#), [12](#)
- [47] Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a model of facial shape and expression from 4D scans. Transactions on Graphics, (Proc. SIGGRAPH Asia) **36**(6), 194:1–194:17 (2017), <https://doi.org/10.1145/3130800.3130813> [2](#), [3](#), [6](#), [7](#), [8](#), [9](#), [10](#)
- [48] Liu, S., Li, T., Chen, W., Li, H.: Soft rasterizer: A differentiable renderer for image-based 3d reasoning. International Conference on Computer Vision (ICCV) (Oct 2019) [7](#)
- [49] Loshchilov, I., Hutter, F.: Fixing weight decay regularization in adam. CoRR [abs/1711.05101](#) (2017), <http://arxiv.org/abs/1711.05101> [6](#)
- [50] Morales, A., Piella, G., Sukno, F.M.: Survey on 3d face reconstruction from uncalibrated images (2021) [3](#)
- [51] Nagano, K., Seo, J., Xing, J., Wei, L., Li, Z., Saito, S., Agarwal, A., Fursund, J., Li, H.: Pagan: Real-time avatars using dynamic textures. ACM Trans. Graph. **37**(6) (dec 2018). <https://doi.org/10.1145/3272127.3275075>, <https://doi.org/10.1145/3272127.3275075> [4](#)
- [52] Paysan, P., Knothe, R., Amberg, B., Romdhani, S., Vetter, T.: A 3D face model for pose and illumination invariant face recognition. In: International conference on advanced video and signal based surveillance. pp. 296–301 (2009) [3](#)
- [53] Phillips, P., Flynn, P., Scruggs, T., Bowyer, K., Chang, J., Hoffman, K., Marques, J., Min, J., Worek, W.: Overview of the face recognition grand challenge. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). vol. 1, pp. 947–954 vol. 1 (2005). <https://doi.org/10.1109/CVPR.2005.268> [8](#)
- [54] Ramamoorthi, R., Hanrahan, P.: An efficient representation for irradiance environment maps. In: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques. p. 497–500. SIGGRAPH '01, Association for Computing Machinery, New York, NY, USA (2001). <https://doi.org/10.1145/383259.383317>, <https://doi.org/10.1145/383259.383317> [7](#)

- [55] Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. arXiv:2007.08501 (2020) 7
- [56] Richardson, E., Sela, M., Kimmel, R.: 3d face reconstruction by learning from synthetic data (2016) 4
- [57] Richardson, E., Sela, M., Or-El, R., Kimmel, R.: Learning detailed face reconstruction from a single image (2017) 4
- [58] Saito, S., Wei, L., Hu, L., Nagano, K., Li, H.: Photorealistic facial texture inference using deep neural networks (2016) 5
- [59] Sanyal, S., Bolkart, T., Feng, H., Black, M.: Learning to regress 3d face shape and expression from an image without 3d supervision. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 1, 2, 4, 5, 9, 10, 11, 12
- [60] Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (Jun 2015). <https://doi.org/10.1109/cvpr.2015.7298682>, <http://dx.doi.org/10.1109/CVPR.2015.7298682> 4, 5
- [61] Serengil, S.I., Ozpinar, A.: Hyperextended lightface: A facial attribute analysis framework. In: 2021 International Conference on Engineering and Emerging Technologies (ICEET). pp. 1–4. IEEE (2021). <https://doi.org/10.1109/ICEET53442.2021.9659697>, <https://doi.org/10.1109/ICEET53442.2021.9659697> 13
- [62] Shang, J., Shen, T., Li, S., Zhou, L., Zhen, M., Fang, T., Quan, L.: Self-supervised monocular 3D face reconstruction by occlusion-aware multi-view geometry consistency. In: European Conference on Computer Vision (ECCV). vol. 12360, pp. 53–70 (2020) 9
- [63] Tewari, A., Bernard, F., Garrido, P., Bharaj, G., Elgharib, M., Seidel, H.P., Pérez, P., Zollhofer, M., Theobalt, C.: Fml: Face model learning from videos. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10812–10822 (2019) 4
- [64] Tewari, A., Zollhöfer, M., Garrido, P., Bernard, F., Kim, H., Pérez, P., Theobalt, C.: Self-supervised multi-level face model learning for monocular reconstruction at over 250 hz. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) 4
- [65] Tewari, A., Zollöfer, M., Kim, H., Garrido, P., Bernard, F., Perez, P., Christian, T.: MoFA: Model-based Deep Convolutional Face Autoencoder for Unsupervised Monocular Reconstruction. In: The IEEE International Conference on Computer Vision (ICCV) (2017) 4
- [66] Thies, J., Zollhöfer, M., Stamminger, M., Theobalt, C., Nießner, M.: Facevr: Real-time gaze-aware facial reenactment in virtual reality. ACM Transactions on Graphics 2018 (TOG) (2018) 3
- [67] Thies, J., Elgharib, M., Tewari, A., Theobalt, C., Nießner, M.: Neural voice puppetry: Audio-driven facial reenactment. European Conference on Computer Vision (ECCV) (2020) 4

- [68] Thies, J., Zollhöfer, M., Nießner, M.: Deferred neural rendering: Image synthesis using neural textures. *Transactions on Graphics (TOG)* **38**(4), 1–12 (2019) [4](#)
- [69] Thies, J., Zollhöfer, M., Nießner, M., Valgaerts, L., Stamminger, M., Theobalt, C.: Real-time expression transfer for facial reenactment. *Transactions on Graphics (TOG)* **34**(6) (2015) [4](#)
- [70] Thies, J., Zollhöfer, M., Stamminger, M., Theobalt, C., Nießner, M.: Face2Face: Real-time face capture and reenactment of RGB videos. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2387–2395 (2016) [3](#), [4](#), [7](#)
- [71] Thies, J., Zollhöfer, M., Theobalt, C., Stamminger, M., Niessner, M.: Headon: Real-time reenactment of human portrait videos. *ACM Transactions on Graphics* **37**(4), 1–13 (Aug 2018). <https://doi.org/10.1145/3197517.3201350>, <http://dx.doi.org/10.1145/3197517.3201350> [3](#), [4](#)
- [72] Tran, A.T., Hassner, T., Masi, I., Medioni, G.: 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 (CVPR)*. pp. 1599–1608 (2017) [4](#), [9](#)
- [73] Tran, A.T., Hassner, T., Masi, I., Paz, E., Nirkin, Y., Medioni, G.: Extreme 3D face reconstruction: Seeing through occlusions. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2018) [5](#)
- [74] Tran, L., Liu, F., Liu, X.: Towards high-fidelity nonlinear 3d face morphable model. In: *In Proceeding of IEEE Computer Vision and Pattern Recognition*. Long Beach, CA (June 2019) [4](#)
- [75] Tu, X., Zhao, J., Jiang, Z., Luo, Y., Xie, M., Zhao, Y., He, L., Ma, Z., Feng, J.: Joint 3D face reconstruction and dense face alignment from a single image with 2D-assisted self-supervised learning. *arXiv preprint arXiv:1903.09359* (2019) [4](#)
- [76] Wei, H., Liang, S., Wei, Y.: 3d dense face alignment via graph convolution networks (2019) [4](#)
- [77] Weise, T., Bouaziz, S., Li, H., Pauly, M.: Realtime performance-based facial animation. In: *Transactions on Graphics (TOG)*. vol. 30 (2011) [3](#)
- [78] Weise, T., Li, H., Gool, L.J.V., Pauly, M.: Face/Off: live facial puppetry. In: *SIGGRAPH/Eurographics Symposium on Computer Animation (SCA)*. pp. 7–16 (2009) [3](#)
- [79] Yamaguchi, S., Saito, S., Nagano, K., Zhao, Y., Chen, W., Olszewski, K., Morishima, S., Li, H.: High-fidelity facial reflectance and geometry inference from an unconstrained image. *ACM Trans. Graph.* **37**(4) (jul 2018). <https://doi.org/10.1145/3197517.3201364>, <https://doi.org/10.1145/3197517.3201364> [4](#), [5](#)
- [80] Yang, H., Zhu, H., Wang, Y., Huang, M., Shen, Q., Yang, R., Cao, X.: Facescape: A large-scale high quality 3d face dataset and detailed riggable 3d face prediction. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2020) [8](#)

- [81] Yin, L., Wei, X., Sun, Y., Wang, J., Rosato, M.: A 3d facial expression database for facial behavior research. In: 7th International Conference on Automatic Face and Gesture Recognition (FGR06). pp. 211–216 (2006). <https://doi.org/10.1109/FGR.2006.6> 8
- [82] Zhang, Z., Girard, J.M., Wu, Y., Zhang, X., Liu, P., Ciftci, U., Canavan, S., Reale, M., Horowitz, A., Yang, H., Cohn, J.F., Ji, Q., Yin, L.: Multi-modal spontaneous emotion corpus for human behavior analysis. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3438–3446 (2016). <https://doi.org/10.1109/CVPR.2016.374> 8
- [83] Zhu, H., Yang, H., Guo, L., Zhang, Y., Wang, Y., Huang, M., Shen, Q., Yang, R., Cao, X.: Facescape: 3d facial dataset and benchmark for single-view 3d face reconstruction. arXiv preprint arXiv:2111.01082 (2021) 2, 8
- [84] Zhu, X., Lei, Z., Liu, X., Shi, H., Li, S.Z.: Face alignment across large poses: A 3D solution. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 146–155. IEEE Computer Society, Los Alamitos, CA, USA (jun 2016). <https://doi.org/10.1109/CVPR.2016.23>, <https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.23> 4
- [85] Zollhöfer, M., Thies, J., Garrido, P., Bradley, D., Beeler, T., Pérez, P., Stamminger, M., Nießner, M., Theobalt, C.: State of the art on monocular 3D face reconstruction, tracking, and applications. *Computer Graphics Forum (Eurographics State of the Art Reports)* **37**(2) (2018) 3