A Perceptual Shape Loss for Monocular 3D Face Reconstruction

C. Otto^{1,2}, P. Chandran¹, G. Zoss¹, M. Gross^{1,2}, P. Gotardo^{*1}, D. Bradley¹

¹DisneyResearchlStudios, Switzerland ²ETH Zürich, Switzerland

Abstract

Monocular 3D face reconstruction is a wide-spread topic, and existing approaches tackle the problem either through fast neural network inference or offline iterative reconstruction of face geometry. In either case carefully-designed energy functions are minimized, commonly including loss terms like a photometric loss, a landmark reprojection loss, and others. In this work we propose a new loss function for monocular face capture, inspired by how humans would perceive the quality of a 3D face reconstruction given a particular image. It is widely known that shading provides a strong indicator for 3D shape in the human visual system. As such, our new 'perceptual' shape loss aims to judge the quality of a 3D face estimate using only shading cues. Our loss is implemented as a discriminator-style neural network that takes an input face image and a shaded render of the geometry estimate, and then predicts a score that perceptually evaluates how well the shaded render matches the given image. This 'critic' network operates on the RGB image and geometry render alone, without requiring an estimate of the albedo or illumination in the scene. Furthermore, our loss operates entirely in image space and is thus agnostic to mesh topology. We show how our new perceptual shape loss can be combined with traditional energy terms for monocular 3D face optimization and deep neural network regression, improving upon current state-of-the-art results.

CCS Concepts

• Computing methodologies \rightarrow Mesh geometry models;

1. Introduction

Monocular 3D face reconstruction is a primary research problem in computer vision and provides the ability to recover 3D geometry and motion of a person's face from images or videos, and opens up many applications in entertainment, telepresence and interaction in virtual environments. This problem has been studied for a long time, and many methods and benchmarks have been proposed [SBFB19, CZR*22]. Methods for monocular 3D face reconstruction range from inference-based approaches that focus on training a neural network to predict the face shape from an image to others that directly optimize for the face shape given an image.

In either case, the one common theme is that most methods formulate an energy minimization problem (either during training, or during optimization) with loss terms that aim to make the resulting geometry match the shape of the person's face in the image. Common loss terms include a pixel-wise photo-consistency loss, a landmark alignment loss, and sometimes more specialized loss terms like identity or emotion recognition losses. In this work we propose a new loss term based on human perception of shape that can be added to a wide variety of techniques used in monocular face capture.

Our work is motivated by the human perception of shape and our ability to use shading cues to estimate 3D orientations of surfaces in images. For example, Ramachandran [Ram88] showed that humans can perceive shape from shading using intensity alone, assuming the object is lit from a single light source. Based on this inspiration, we design a new neural loss function that takes an RGB image of a face and a shaded render of the geometry estimate as input, and predicts a scalar value (e.g. a score) to indicate how well the geometry matches the shape of the face, including identity, pose and expression. Following the findings of Ramachandran [Ram88], we render the geometry estimate with diffuse gray shading, lit from a single light source (e.g. using PyTorch3D [RRN*20]). We implement our perceptual shape loss as a discriminator (or critic) network, based on the Deep Convolutional GAN (DCGAN) architecture [RMC16], and train it with both real and fake image-render pairs, obtained from a dataset of 3D facial scans in a studio setting augmented with in-the-wild synthetic data.

Our critic network operates entirely in image space and it is agnostic of the mesh topology, which is beneficial since different algorithms rely on different 3D face models, like the Basel face model (BFM) [GMFB*18, PKA*09], FLAME [LBB*17], FaceWarehouse [CWZ*14] or other custom topology meshes [YZW*20, BLC*21, BRZ*16, DPSD20, CZR*22, WCY*22, WBH*21, CBGB20]. As such, our trained critic network

^{*}Now at Google

^{© 2023} Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd.

can be easily used as a loss term in energy functions for 3D face reconstruction.

We demonstrate the use of our perceptual shape loss in the scenario of offline optimization-based monocular 3D face reconstruction as well as in the scenario of deep neural network 3D face regression. Specifically, in an optimization framework we show that the addition of our loss improves overall reconstruction accuracy as compared to optimizing a standard set of loss functions used in literature, when evaluated using the NoW benchmark [SBFB19]. Similarly, in an inference framework we show that using our loss to train a new 3D face regression method (which we call PSL) by fine-tuning DECA [FFBB21], improves the results of identity and expression as judged by both the NoW benchmark [SBFB19] and the REALY benchmark [CZR*22]. To summarize:

- We present a new perceptual shape loss function for monocular face reconstruction, inspired by the human ability to distinguish shape from shading.
- Our loss function is topology-agnostic, and can be used with any 3D face model.
- Our loss is generic, and can be readily incorporated into energy functions used for 3D face reconstruction from images in both optimization and regression frameworks.
- We train a novel Perceptual Shape Loss (PSL) method for 3D face regression by fine-tuning DECA with our loss function and show that its reconstructions improve identity and expression accuracy on the NoW and the REALY benchmarks.

2. Related Work

High-quality 3D face reconstructions have traditionally required expensive multi-camera studio setups [BBB*10]. Great research effort has gone into exploring alternative approaches that can alleviate many of the traditional constraints (e.g. multiple cameras, controlled lighting, expensive hardware) by investigating singlecamera, in-the-wild 3D face reconstruction. Most monocular methods use a 3D Morphable Model (3DMM) [BV99] as a statistical face prior for reconstruction, to account for ambiguities, such as the missing depth information [WBGB16]. 3DMMs are created from a corpus of 3D face scans, and rely on statistics to represent plausible human face shapes and expressions in a data-driven way. Detailed surveys of 3DMMs and their use in monocular face capture are given in Egger et al. [EST*20], Zollhöfer et al. [ZTG*18], and Morales et al. [MPS21]. Most state-of-the-art 3D face reconstruction methods are based on either the FLAME face model [LBB*17] or the BFM [PKA*09, GMFB*18]. The best 3DMM shape, expression and pose parameters for a given RGB image are searched either via analysis-by-synthesis optimization [BV99, GZC*16, DBA^{*}21, GPKZ19, TZS^{*}18, TZN^{*}15] or via deep neural network regression [LRF*23, RGP*23, TZK*17, aTaHMM16, GCM*18, RSK16,RSOEK17,KZT*18,DSK17,TZX*21,TZG*18].

One way to evaluate the quality of state-of-the-art methods is to compare their 3D face reconstructions with the corresponding ground truth face scans. The *not quite in-the-wild* (NoW) benchmark [SBFB19] and the benchmark proposed by Feng et al. [FHK*18] evaluate the quality of the neutral face shape (i.e. the identity). The REALY benchmark [CZR*22, DPSD20, YZW*20, BLC*21] evaluates both identity and expression, because their ground truth face scans also include non-neutral faces. Both NoW and REALY follow ICP-based protocols to align the face reconstructions with ground truth face scans, establish correspondences, and compute vertex errors. Two of the best ranking methods on the NoW benchmark to date are MICA [ZBT22] and DECA [FFBB21], both of which are based on the FLAME face model. Further state-of-the-art methods as judged by the NoW benchmark are either BFM-based [LMFV*21, DTA*21, WXN21, GZY*20, DYX*19, SSL*20, THMM17], FLAME-based [SBFB19, ZTZ*22], based on other face models [WBH*22, KS20] or are model-free [FWS*18]. Most of these methods use traditional loss functions when training inference networks, like pixel-wise photoconsistency and landmark alignment. Our proposed perceptual loss function could readily complement the majority of these approaches.

EMOCA [DBB22] and SPECTRE [FRPP*22] are two recent methods that introduce new perceptual loss terms to improve upon the quality of the reconstructions predicted by DECA. EMOCA uses an emotion CNN trained on an annotated large scale emotion dataset [MHM19] to recognize emotions in RGB images and rendered reconstructions. They compare the predicted features in the last network layer, effectively creating a perceptual loss by enforcing consistency between the two. SPECTRE improves upon the lip reconstruction by running a pre-trained lip-reading network [MPP22] for detecting features from lip RGB image crops and rendered lip reconstructions. Similar to these two works, we also utilize a CNN to arrive at our perceptual shape judgements. In contrast to both methods, our perceptual shape loss maximizes a global score based on a shaded render input and the target image, drawing inspiration from the human ability of leveraging shading cues to perceive shape [Ram88]. In terms of quality, the reconstructions by EMOCA display more faithful emotions than previous methods, and SPECTRE improves upon the visual quality of the lip reconstruction. However, the visual improvements of both methods are reflected only in certain categories of the RE-ALY benchmark [CZR^{*}22] and not reflected in the NoW benchmark [SBFB19]. One reason for this is that both methods do not improve upon the identity component initialized by DECA, focusing only on the expression or jaw pose components. On the other hand, the visual improvements gained with our perceptual shape loss are also reflected by quantitative improvements in the identity and the expression reconstruction, which we measure on both the NoW and the REALY benchmarks (Section 5).

3. Perceptual Shape Loss

Our key idea is to create a loss function for monocular 3D face reconstruction, implemented as a neural network that takes a face image and a gray-shaded render of face geometry as input, and outputs a scalar value to indicate how well the render matches the image in terms of shape. The network should intuitively critique the inputs, and provide continuous feedback about the 'goodness' of match between the image and the render, for any image-render pair. Once such a critic network is trained, its output can be interpreted as a perceptual shape loss that can be used for 3D face reconstruction tasks. In the following, we describe our network design, training data, and other implementation details.

3.1. Critic Network

Our network is implemented as a CNN, inspired by the Deep Convolutional GAN (DCGAN) [RMC16] discriminator. The DC-GAN discriminator is a simple and well-explored architecture that is commonly used for training critic neural networks [GAA*17, ACB17]. It consists only of several strided convolution layers, normalization layers and leaky ReLU activation functions and is therefore fast to train. There is no activation function after the last layer, allowing for a real valued scalar output. To accommodate for our input resolution of 256×256 we extend the original DCGAN (which only supported 64×64 images) by two additional convolution layers. The input to the network is an RGB face image I and a single channel gray-shaded render of a geometry estimate I_R , rendered from the same camera viewpoint and shaded by a single point light in front of the face. The image I and the render I_R are stacked to make a 4-channel input which is fed to the critic \mathcal{D} . The output of our critic network is a 1-dimensional scalar S, which will represent the perceptual shape loss value. An overview of the network is shown in Figure 1.



Figure 1: Our convolutional network D takes an aligned RGB face image and a gray-shaded geometry render as input, and predicts a score S that indicates how well the render matches the image in terms of shape.

We train the network as a discriminator to regress higher values when the image-render pair match well, and lower values when they match poorly. Our training data includes a set of *real* imagerender pairs where the render matches the image perfectly, and a set of *fake* image-render pairs where the render does not match the image (see Section 3.2 for more details on the training data). We formulate our training objective so that the score for real examples is maximized and for fake examples it is minimized, using the WGAN-GP loss [GAA*17, ACB17]:

$$\mathcal{L}_{\text{critic}} = \mathcal{D}\left(\mathbf{\tilde{x}}\right) - \mathcal{D}\left(\mathbf{x}\right) + \lambda \left(\|\nabla_{\mathbf{\hat{x}}}\mathcal{D}\left(\mathbf{\hat{x}}\right)\|_{2} - 1\right)^{2}, \quad (1)$$

where \mathcal{D} is our network, \mathbf{x} is a real sample and $\mathbf{\tilde{x}}$ is a fake sample. A gradient penalty term encourages the gradient norm $\|\nabla_{\mathbf{\hat{x}}} \mathcal{D}(\mathbf{\hat{x}})\|_2$



for Computer Graphics and John Wiley & Sons Ltd.

to stay close to 1 (1-Lipschitz constraint), where $\hat{\mathbf{x}} = \varepsilon \mathbf{x} + (1-\varepsilon)\tilde{\mathbf{x}}$ and ε is a random number drawn from U[0, 1] [GAA*17]. The loss $\mathcal{L}_{\text{critic}}$ is calculated during training after passing both a batch of real and fake examples through the critic network. Although our critic outputs a continuous real valued scalar during training, we re-normalize its output at inference time such that the median of real examples is 1.0 and median of fake examples is -1.0, to have an intuitive loss.

3.2. Training Data

For our training dataset, we use high-quality 3D face scans acquired from a multi-view studio capture setup [CBGB20]. The dataset contains 358 different identities, displaying 24 different expressions under 12 different camera viewpoints. The geometry is reconstructed using the passive stereo method of Beeler et al. [BBB*10], and fit to a common template topology [CBGB20].

To create real samples \mathbf{x} we can select from the identities, expressions and cameras and use the corresponding reconstructed geometry to create the render, as it perfectly matches the identity and expression in the face skin area (excluding eyes, inner mouth and hair). Some examples of these real samples used for training are shown in Figure 2 (left side). To create fake examples $\mathbf{\tilde{x}}$ we purposely select mismatching images and geometry renders (e.g. an image of one identity but the geometry render of a different identity). Examples of such fake samples for training are illustrated in Figure 3, where the fake render may have an incorrect expression, identity or both.

One problem with the studio dataset is that all the images are acquired from a fixed set of cameras, under uniform white illumination. To encourage closure of the domain gap between studio and in-the-wild data, we create additional data using the approach of Chandran et al. [CWZ*21], who render the textured geometry in various scenes and inpaint the eyes, hair and background using a 2D human face image generator (StyleGAN2 [KLA*20]). Their approach allows control over viewpoint and lighting to imitate inthe-wild conditions. The resulting dataset contains a large variety of quasi in-the-wild images for which we have the corresponding high-quality geometry serving as ground truth. Given this synthetic data, we can create real and fake samples using the same approach as for studio data. Some real examples from this in-the-wild synthetic dataset are shown in Figure 2 (right side). Note that the approach of Chandran et al. [CWZ*21] is limited to mostly frontal viewpoints and does not allow for the creation of extreme side pose examples.

In total, we train our network on 92736 real studio examples and 6276 real synthetic examples. We generate three times the number of fake examples in the following way: One third of the fake examples have renders displaying an incorrect identity, another third display an incorrect expression and the last third display both, incorrect identity and expression. To encourage our loss to critique the head pose, we create another 99012 fake examples where we apply spurious head rotations and translations to the geometry before rendering. This leads to a total of 495060 unique training examples (approximately 100K real and 400K fake). During training we oversample the real examples by a factor of 4 to equalize the

C. Otto et al. / A Perceptual Shape Loss for Monocular 3D Face Reconstruction

Studio Data (Real Images) In-the-Wild Data (Synth. Images)



Figure 2: Ground truth 'real' samples we use for training. Our dataset includes both studio images of real people (left) and in-the-wild synthetic images (right).



Figure 3: We create 'fake' samples for training by intentionally mixing face images with geometry renders that have incorrect expression, identity or both. Additional fake images are created with incorrect head pose (not shown in the figure).

number of real and fake examples seen by the network. We create an additional 20736 examples (10368 real and 10368 fake) to serve as a validation set.

3.3. Implementation Details

Prior to training we normalize the 4D input (I, I_R) images by cropping the faces and aligning the eyes, mouth and nose to approximately the same position across all inputs [NHSW20]. As we aim to operate on any face topology including both full head geometry as well as frontal face meshes only, we ensure a consistent render area by discarding vertices at a distance greater than 0.7 × (outer eye dist + nose dist) from the center of the face, as suggested by related methods [SBFB19, FHK*18], before rendering the geometry. We further crop away any vertices from the interior of the



Figure 4: *Histogram of score separation on our held-out validation set. Note: These show the scores prior to* -1 *to* 1 *normalization.*

mouth and eye region that exist in the geometry from Chandran et al. [CBGB20], because many methods do not reconstruct these. For our training dataset, the final cropped mesh topology consists of 38799 vertices.

We follow Gulrajani et al. [GAA^{*}17] and do not use batch normalization in our critic network and replace it with instance normalization [UVL16]. We implement our method in Py-Torch [PGM^{*}19]. We train our critic network for 4 epochs using the Adam optimizer [KB15] with a learning rate of 0.001 and $\lambda = 10$ on a TITAN X GPU. Training takes around eight hours (2 hours per epoch) with input images of size 256² and batch size 64 to reach the best separation between fake and real data on a validation set (refer to Section 4 and Figure 4). We create our renders with PyTorch3D [RRN^{*}20]. We apply several color transformations to the RGB face images using Kornia data augmentations [ERB20, SRM^{*}20] in order to increase the variance of our dataset. Similar to McDonagh et al. [MKB^{*}16], we use uniform noise as the background of our render so that the model learns the signal only from the foreground.

4. Validation

We now present experiments to validate our critic network as a perceptual shape loss. We first validate that our network learned to identify real versus fake input examples by plotting the distribution of output scores for our held-out validation set. As illustrated in Figure 4, our network has learned a reasonable separation of the two distributions. Note that we plot the original output scores Sbefore normalizing the medians to -1 to 1.

Equally important, we need to validate that our critic network is suitable for use as a loss function in gradient-based optimizations. A loss function should be differentiable, continuous, smooth, and ideally have a minimum at the optimal parameter inputs. As our perceptual loss is implemented as a CNN in PyTorch, it is trivially differentiable. In order to demonstrate continuity and smoothness



Figure 5: We demonstrate the smooth deterioration of the perceptual shape score predicted by our critic network as the input shape gradually moves away from the ground truth. We show examples of the critic's response to deviation in expression, identity, rotation, and translation on four validation examples that were not seen by the critic during training. The smooth and intuitive degradation of the critic's score highlights the appropriateness of our perceptual shape loss for face reconstruction.

of the function, we perform a series of experiments where we plot the normalized output score while perturbing the geometry render in various ways. Specifically, starting with the ground truth (real) render, we slowly edit the expression of the geometry so that it becomes increasingly farther from the ground truth. Figure 5 (top row) shows several steps of this experiment, along with a plot of the per-step network output values. We repeat this experiment three more times on different example images, this time varying the identity, head rotation and head translation, respectively (rows 2 to 4 in Figure 5). As the plots indicate, our function is smooth and continuous, and generally decreases as we deviate from the optimal parameters. While not performing an exhaustive evaluation of the local neighborhood around the optimal parameters to show a strict maximum, our experiments show evidence of a local maximum for the optimal geometry renders.

Given these successful validation experiments, we hypothesize that our critic network is suitable to use as a loss function.

5. Face Reconstruction

We now demonstrate the application of our perceptual shape loss in the context of monocular 3D face reconstruction. At a high level, the scenario is the following: given a face image and a user-defined

```
© 2023 Eurographics - The European Association
for Computer Graphics and John Wiley & Sons Ltd.
```

 Input
 DECA [19]
 L
 L
 L
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 I
 <

Figure 6: Here we show qualitative optimization results for face reconstruction on our held-out validation set, where the inferred geometry from DECA [FFBB21] fails to faithfully reproduce the expressions seen in the input image. Even when this result is further optimized upon using traditional losses (\mathcal{L}_{base}), the visual quality of the result does not improve substantially. In contrast, the addition of the proposed perceptual loss (\mathcal{L}_{full}) enables the optimization to faithfully reproduce the expressions seen in the input image.

deformable face model, we want to optimize or regress the model parameters in order to best match the pose and shape of the face in the image. Whether optimizing the single image directly (Section 5.1) or training a neural network for regression (Section 5.2), our perceptual loss can be used by simply rendering the evaluated face geometry at each iteration and passing it through our critic network. For these experiments we will use the FLAME [LBB*17] face model, and we will compare different energy terms, with and without our new loss function.

The FLAME model consists of $V \in \mathbb{R}^{5023 \times 3}$ vertices and 9976 faces *F*, which are evaluated by the model function \mathcal{M} , as

$$\mathcal{M}(\beta, \theta, \psi) \to (V, F),$$
 (2)

where β , ψ and θ are the identity (shape), expression and pose parameters, respectively.

As a framework for both the single-image optimization and the network training, we consider a set of standard energy terms defined by a recent regression-based face reconstruction method, DECA [FFBB21]. Specifically, we formulate the following baseline energy function to minimize:

$$\mathcal{L}_{\text{base}} = \lambda_{l} \mathcal{L}_{l} + \lambda_{p} \mathcal{L}_{p} + \lambda_{id} \mathcal{L}_{id} + \lambda_{r} \mathcal{L}_{r}, \qquad (3)$$

where \mathcal{L}_{l} , \mathcal{L}_{p} , \mathcal{L}_{id} and \mathcal{L}_{r} are the landmark reprojection loss, photometric loss, identity loss, and regularization term defined in DECA. As a comparison, we define a second energy function as

$$\mathcal{L}_{\text{full}} = \mathcal{L}_{\text{base}} + \lambda_{\text{s}} \mathcal{L}_{\text{s}},\tag{4}$$

where \mathcal{L}_s is our new perceptual shape loss that uses our pre-trained critic network \mathcal{D} , and is defined as

$$\mathcal{L}_{\rm s} = -D(I, I_R) \tag{5}$$

given the RGB face image I and the estimated geometry render I_R . We further demonstrate the power of our perceptual loss term by defining a third, more compact loss function as

$$\mathcal{L}_{\text{compact}} = \lambda_1 \mathcal{L}_1 + \lambda_s \mathcal{L}_s + \lambda_r \mathcal{L}_r, \tag{6}$$

which contains only our perceptual loss, the landmark reprojection loss and the regularization losses, dropping the photometric and identity losses. We keep the landmark loss because, like most face reconstruction methods, we rely on landmarks to robustly ground the reconstruction and achieve stable results over a wide range of images. In the following sections we evaluate the quality of different face reconstruction methods using the three loss functions \mathcal{L}_{base} , \mathcal{L}_{full} and $\mathcal{L}_{compact}$.

5.1. Single-Image Optimization

In the context of offline single-image face reconstruction, we compare reconstruction accuracy while using each of the three loss functions during parameter optimization. For evaluation, we use the common non-metrical NoW benchmark [SBFB19] and reconstruct both the NoW validation set and the NoW test set. For every optimization, we initialize the model parameters using DECA's pre-trained coarse inference network and perform 60 optimization iterations with the following hyper-parameters: $\lambda_l=0.8, \lambda_p=2.0, \lambda_{id}=0.2, \lambda_r=0.0001,$ and $\lambda_s=0.5.$ Following DECA, we optimize for the first 100 identity parameters $\beta \in \mathbb{R}^{100}$ and the first 50 expression parameters $\psi \in \mathbb{R}^{50}$. We also optimize for a subset of $\theta \in \mathbb{R}^9$ pose parameters describing global rotation, jaw joint and translation. The face albedo and illumination parameters required for the photometric and identity losses are fixed after initialization.

Optimization results on the validation set are shown in the upper section of Table 1, where we can see that using our new loss function (\mathcal{L}_{full}) increases the accuracy of reconstruction over the



Figure 7: *Here we show additional qualitative optimization results for face reconstructions on images in the wild, comparing the initial DECA* [*FFBB21*] *fit with optimizing parameters without* (\mathcal{L}_{base}) *and with* (\mathcal{L}_{full}) *our perceptual loss.*

baseline. The greatest improvement over the baseline can be seen on the NoW selfie challenge subset, which contains mostly frontal faces. Note that all optimizations also naturally improve over the initialization (DECA), which is an inference-based method. Results on the NoW test set show similar improvements (upper section of

C. Otto et al. / A Perceptual Shape Loss for Monocular 3D Face Reconstruction

NoW val Optimization	Full (Med/Mean/Std)↓	Neutral (Med/Mean/Std)↓	Expression (Med/Mean/Std) \downarrow	$Occlusion~(Med/Mean/Std) {\downarrow}$	Selfie (Med/Mean/Std) \downarrow
Initialization (DECA)	1.1757 / 1.4611 / 1.2511	1.1546 / 1.4433 / 1.2408	1.1351 / 1.4251 / 1.2384	1.2181 / 1.4890 / 1.2528	1.2562 / 1.5445 / 1.2996
\mathcal{L}_{base}	1.1241 / 1.4033 / 1.2062	1.1029 / 1.3749 / 1.1843	1.0700 / 1.3456 / 1.1788	1.1634 / 1.4403 / 1.2125	1.2566 / 1.5505 / 1.2990
$\mathcal{L}_{compact}$ (Ours)	1.1228 / 1.3987 / 1.1936	1.1046 / 1.376 / 1.1741	1.0699 / 1.3441 / 1.1678	1.1870 / 1.4673 / 1.2318	1.1901 / 1.4649 / 1.2235
\mathcal{L}_{full} (Ours)	1.1226 / 1.3983 / 1.1932	1.1067 / 1.3778 / 1.1747	1.0694 / 1.3432 / 1.1665	1.1831 / 1.4650 / 1.2327	1.1907 / 1.4637 / 1.2209
NoW val Inference-based	Full (Med/Mean/Std)↓	Neutral (Med/Mean/Std) \downarrow	Expression (Med/Mean/Std) \downarrow	$Occlusion~(Med/Mean/Std) {\downarrow}$	Selfie (Med/Mean/Std) \downarrow
EMOCA [DBB22]	1.1757 / 1.4611 / 1.2511	1.1546 / 1.4433 / 1.2408	1.1351 / 1.4251 / 1.2384	1.2181 / 1.4890 / 1.2528	1.2562 / 1.5445 / 1.2996
EMOCA_V2 [DBB22]	1.1757 / 1.4611 / 1.2511	1.1546 / 1.4433 / 1.2408	1.1351 / 1.4251 / 1.2384	1.2181 / 1.4890 / 1.2528	1.2562 / 1.5445 / 1.2996
SPECTRE [FRPP*22]	1.1757 / 1.4611 / 1.2511	1.1546 / 1.4433 / 1.2408	1.1351 / 1.4251 / 1.2384	1.2181 / 1.4890 / 1.2528	1.2562 / 1.5445 / 1.2996
DECA [FFBB21]	1.1757 / 1.4611 / 1.2511	1.1546 / 1.4433 / 1.2408	1.1351 / 1.4251 / 1.2384	1.2181 / 1.4890 / 1.2528	1.2562 / 1.5445 / 1.2996
\mathcal{L}_{base}	1.0805 / 1.3489 / 1.1616	1.0544 / 1.3226 / 1.1463	1.0442 / 1.3162 / 1.1553	1.1300 / 1.3861 / 1.1647	1.1497 / 1.4267 / 1.2027
PSL - L _{compact} (Ours)	1.0780 / 1.3538 / 1.1588	1.0454 / 1.3234 / 1.1418	1.0455 / 1.3289 /1.1558	1.1717 / 1.4266 / 1.1782	1.0707 / 1.3546 / 1.1645
PSL - \mathcal{L}_{full} (Ours)	1.0812 / 1.3573 / 1.1712	1.0440 / 1.3204 / 1.1488	1.0335 / 1.3154 / 1.1601	1.1850 / 1.4436 / 1.1970	1.1119 / 1.3920 / 1.1943

Table 1: The NoW validation set results displaying the full score and the scores when evaluated on each of the four sub-challenges in isolation (multiview neutral, multiview expression, multiview occlusion and selfie). The NoW benchmark evaluates only the identity component of the reconstruction (no expression). For the inference-based comparison, we also compare our method with other state-of-the-art methods that fine-tune the DECA regression network (i.e. EMOCA [DBB22] and SPECTRE [FRPP*22]). As EMOCA [DBB22] and SPECTRE [FRPP*22] use the DECA [FFBB21] identity parameters, they achieve the same score as DECA. Errors are reported in mm.

Table 2). Qualitative comparisons are also illustrated in Figure 6 on our own held-out validation set of studio and in-the-wild data, where we can see that using our perceptual loss greatly helps to match the expression in the face image. Additional optimization results are shown on an image dataset in-the-wild in Figure 7.

We perform the validation and test set optimizations again using the compact loss function $\mathcal{L}_{compact}$. As we can see in Table 1, using only this compact loss can achieve favorable results, better than the baseline and nearly as good as the full set of loss terms. Note, that we choose to set $\lambda_r = 0.0$ for $\mathcal{L}_{compact}$ during optimization, while still obtaining good results. This shows that our perceptual loss term is quite powerful on its own, and could replace several traditional losses.

5.2. Regression Network Training

We now use our perceptual shape loss to train a new inferencebased parameter regression method, based on the network architecture of DECA [FFBB21], which we call PSL (for *perceptual shape loss*). We train two versions of PSL - one using \mathcal{L}_{full} and one using $\mathcal{L}_{compact}$. Additionally, we train a baseline inference network using \mathcal{L}_{base} . Similar to related work [DBB22, FRPP*22] our training proceeds by fine-tuning a pre-trained DECA network.

The optimization experiments in Section 5.1 show that our perceptual shape loss improves the most on frontal poses (Table 1 - Selfie) and expressions (Figure 6). Therefore, we fine-tune DECA using 2000 frontal images extracted from the CelebAMask-HQ dataset [LLWL20]. We use a learning rate of 0.00001 with the Adam optimizer [KB15] and train for up to 1700 iterations. We set the following hyper-parameters: $\lambda_l = 1.0$, $\lambda_p = 2.0$, $\lambda_{id} = 0.2$, $\lambda_r = 0.0001$, and $\lambda_s = 0.01$.

Once training is completed we evaluate the results quantitatively on the NoW benchmark to investigate identity improvement. We additionally evaluate PSL on the REALY benchmark [CZR*22], which evaluates both identity and expressions. We compare models trained with \mathcal{L}_{base} , \mathcal{L}_{full} and $\mathcal{L}_{compact}$. We show qualitative results

NoW test - Optimization	$\textbf{Median} \downarrow$	$Mean \downarrow$	Std \downarrow
Initialization (DECA)	1.0936	1.3815	1.1829
$\mathcal{L}_{\mathrm{base}}$	1.0865	1.3760	1.1780
$\mathcal{L}_{compact}$ (Ours)	1.0825	1.3700	1.1724
\mathcal{L}_{full} (Ours)	1.0818	1.3693	1.1724
NoW test - Inference-based	$\textbf{Median} \downarrow$	$Mean \downarrow$	Std \downarrow
EMOCA [DBB22]	1.0936	1.3815	1.1829
SPECTRE [FRPP*22]	1.0936	1.3815	1.1829
DECA [FFBB21]	1.0936	1.3815	1.1829
$\mathcal{L}_{\mathrm{base}}$	1.0662	1.3432	1.1405
PSL - $\mathcal{L}_{compact}$ (Ours)	1.0890	1.3680	1.1477
PSL - $\mathcal{L}_{\text{full}}$ (Ours)	1.0902	1.3661	1.1453

Table 2: Results on the non-metrical NoW benchmark test set [SBFB19] of the initial DECA reconstruction versus our different loss terms for optimization (in mm). The optimization leads to best results when including our perceptual shape loss (\mathcal{L}_{full}), and even using our shape loss and landmark loss alone achieves good results ($\mathcal{L}_{compact}$). Similarly, we compare our the fine-tuned DECA regression network methods [SBFB19] to the initial DECA reconstruction and other fine-tuned DECA regression networks (i.e. EMOCA [DBB22] and SPECTRE [FRPP*22]). While \mathcal{L}_{base} performs best, PSL - \mathcal{L}_{full} and PSL - $\mathcal{L}_{compact}$ also improve over the initialization.

produced by the different loss settings in Figure 8. We also compare with the state-of-the-art inference-based methods EMOCA [DBB22] and SPECTRE [FRPP*22]. Note that while other approaches (e.g. MICA [ZBT22] or HRN [LRF*23]) may show better overall performance on the benchmarks, we focus our comparison on EMOCA and SPECTRE because, similar to us, both methods fine-tune DECA and make use of their own perceptual losses during fine-tuning.

The results in the lower part of Table 1 (Inference-based) show

C. Otto et al. / A Perceptual Shape Loss for Monocular 3D Face Reconstruction

REALY frontal-view	All (avg)↓	Nose (avg/med/std)↓	Mouth (avg/med/std) \downarrow	Forehead (avg/med/std)↓	Cheek (avg/med/std)↓
EMOCA [DBB22]	2.103	1.868 / 1.821 / 0.387	2.679/2.419/1.112	2.426 / 2.383 / 0.641	1.438 / 1.294 / 0.501
EMOCA_V2 [DBB22]	2.230	1.704 / 1.680 / 0.411	3.290 / 3.190 / 0.741	2.228 / 2.160 / 0.504	1.698 / 1.577 / 0.612
SPECTRE [FRPP*22]	2.218	2.027 / 2.015 / 0.401	2.695 / 2.623 / 0.969	2.335 / 2.263 / 0.581	1.817 / 1.703 / 0.585
DECA [FFBB21]	2.010	1.697 / 1.654 / 0.355	2.516 / 2.465 / 0.839	2.394 / 2.256 / 0.576	1.479 / 1.400 / 0.535
$\mathcal{L}_{ ext{base}}$	2.223	1.804 / 1.768 / 0.434	2.899 / 2.801 / 0.937	2.391 / 2.321 / 0.535	1.796 / 1.715 / 0.564
PSL - $\mathcal{L}_{compact}$ (Ours)	1.840	1.658 / 1.591 / 0.353	1.832 / 1.775 / 0.549	2.394 / 2.395 / 0.552	1.475 / 1.386 / 0.515
PSL - \mathcal{L}_{full} (Ours)	1.882	1.708 / 1.688 / 0.349	1.876 / 1.777 / 0.563	2.350 / 2.343 / 0.551	1.593 / 1.482 / 0.540
REALY side-view	All (avg)↓	Nose (avg/med/std) \downarrow	Mouth (avg/med/std) \downarrow	Forehead (avg/med/std) \downarrow	Cheek (avg/med/std)↓
EMOCA [DBB22]	2.125	1.867 / 1.548 / 0.554	2.636 / 1.738 / 1.284	2.448 / 2.369 / 0.708	1.548 / 1.424 / 0.590
EMOCA_V2 [DBB22]	2.301	1.849 / 1.713 / 0.726	3.283 / 3.238 / 0.830	2.316 / 2.256 / 0.653	1.758 / 1.621 / 0.669
SPECTRE [FRPP*22]	2.578	2.334 / 1.941 / 1.322	3.210 / 2.839 / 1.467	2.567 / 2.433 / 0.725	2.203 / 1.885 / 1.580
DECA [FFBB21]	2.107	1.903 / 1.700 / 1.050	2.472 / 2.348 / 1.079	2.423 / 2.308 / 0.720	1.630 / 1.456 / 1.135
\mathcal{L}_{base}	2.170	1.847 / 1.731 / 0.683	2.579 / 2.466 / 1.083	2.465 / 2.426 / 0.638	1.789 / 1.675 / 0.744
PSL - L _{compact} (Ours)	1.931	1.715 / 1.621 / 0.591	1.989 / 1.836 / 0.786	2.535 / 2.492 / 0.634	1.486 / 1.358 / 0.592
PSL - \mathcal{L}_{full} (Ours)	1.857	1.685 / 1.629 / 0.475	1.820 / 1.757 / 0.557	2.454 / 2.426 / 0.608	1.469 / 1.378 / 0.495

Table 3: Comparing different DECA-based neural network parameter regression methods on the REALY single image frontal-view and side-view benchmarks [CZR*22]. For both categories, fine-tuning DECA [FFBB21] with our perceptual shape loss (PSL) shows the lowest average error ('All' column) among the compared methods. We can observe large improvements especially in the mouth region, indicating high-quality expressions. Methods that fine-tune DECA with other perceptual loss terms (i.e. EMOCA [DBB22] and SPECTRE [FRPP*22]) do not perform as well. The numbers report the mean NMSE in mm.



Figure 8: *Here we show qualitative inference-based results for face reconstructions, comparing the initial DECA* [*FFBB21*] *fit with* \mathcal{L}_{base} *and the results achieved when using our perceptual shape loss during neural network training:* PSL - $\mathcal{L}_{compact}$ and PSL - \mathcal{L}_{full} .



Figure 9: Here we show qualitative inference-based results for face reconstructions, comparing the initial DECA [FFBB21] fit with our new method PSL - \mathcal{L}_{full} and related state-of-the-art methods EMOCA [DBB22], EMOCA_V2 [DBB22] and SPECTRE [FRPP*22].

that using our perceptual shape loss improves upon the DECA initialization in terms of identity. We observe, similar to the optimization case, that the \mathcal{L}_{full} and $\mathcal{L}_{compact}$ settings show the greatest improvement over \mathcal{L}_{base} on the NoW Challenge selfie category. On the NoW test set (Table 2), \mathcal{L}_{base} achieves the best score. However, both \mathcal{L}_{full} and $\mathcal{L}_{compact}$ still improve upon the DECA initialization. In contrast, EMOCA and SPECTRE keep the DECA identity parameters fixed and do not improve upon the DECA results on the NoW benchmark.

Our PSL methods show large improvements on the REALY benchmark, when compared to other DECA-based parameter regression methods. Table 3 displays the results for both the frontal-view and the side-view evaluation categories. In both categories, PSL-based networks outperform \mathcal{L}_{base} , EMOCA and SPECTRE. Among the compared methods, PSL - $\mathcal{L}_{compact}$ reaches the lowest error for the frontal-view category and PSL - \mathcal{L}_{full} for the side-view category ('All' column). Furthermore, the results show large quantitative improvements in the mouth area for PSL-based methods, indicating high-quality expressions. These improvements can also be observed when looking at the qualitative results in Figure 9.

Summary: In this section we have shown that our perceptual shape loss can be used in the context of monocular face reconstruction, and that employing our loss term in optimization or neural network training scenarios improves the accuracy and visual quality of the results.



Figure 10: We employ our loss function on four different mesh topologies (5072, 19577, 35709 and 38799 vertices). Specifically, on our three PCA face models and the Basel Face Model (BFM) [PKA*09] used in Deep3D [DYX*19]. The results of monocular face capture with our PCA face models show negligible differences across topologies. The BFM result differs slightly due to the different shape manifold of this model.

5.3. Topology Independence

As our critic network operates only on images and shaded geometry renders it is agnostic to the actual face mesh topology. This means that the loss function is not tied to any particular deformable model, vertex count or UV layout. To highlight this, we introduce an experiment where we use our perceptual loss on four different mesh topologies and show that our perceptual shape loss can be applied independent of the topology.

The experiment is again in the domain of optimization-based

monocular face capture, similar to Section 5.1. Here we isolate a single face image with neutral expression for illustration, and instead of using the FLAME face model we first create three separate PCA face models at three different resolutions: 5072 vertices, 19577 vertices, and 38799 vertices (refer to Figure 10). The PCA bases are created from the 3D reconstructions of neutral faces from the dataset of Chandran et al. [CBGB20], after resampling the data to the aforementioned vertex counts.

For the sake of this experiment, we optimize for the parameters of each of the three PCA models using only our perceptual loss function \mathcal{L}_s , starting from the mean PCA face shape. Figure 10 shows that each optimization of our PCA face models produces the same result, with nearly identical error maps when compared to the ground truth shape scan.

Additionally, we use our perceptual shape loss function \mathcal{L}_s to optimize parameters of the Basel Face Model (BFM) [PKA*09] used in Deep3D [DYX*19], which includes expression parameters from FaceWarehouse [CWZ*14, GZC*19]. Its topology consists of 35709 vertices and it operates in a different parameter space than the FLAME model or each of our PCA face models. We initialize the optimization from the mean BFM face with shape and expression parameters set to zero. Naturally, optimizing in the BFM parameter space leads to slightly different results when compared to the PCA face models, due to the different shape manifolds of each model. However, the optimization succeeds to capture the shape of the face in the image (Figure 10).

These experiments indicate that our loss function is topologyagnostic, and our critic network does not need to be retrained when the topology (or model) changes. Further evidence can be found by the fact that the experiments in Section 5.1 and Section 5.2 were performed on the FLAME face model, which has yet a different topology (5023 vertices) than the four topologies used in this experiment.

6. Conclusion

In this work we propose a novel perceptual shape loss function for monocular 3D face reconstruction. Inspired by a study on the human perception of shape [Ram88], we design a discriminator (critic) network that learns to distinguish 'good' (real) face reconstructions from 'bad' (fake) ones, by only looking at pairs of input images and a gray shaded geometry renders. This is accomplished by creating a dataset of real and fake face reconstructions, and by training the critic network to output higher scores for the real examples than for the fake examples. Once trained, our critic's score, which we refer to as the perceptual shape loss, can be plugged into any optimization or neural network regression framework for evaluating the quality of a given face reconstruction. We demonstrate the ability of our shape loss to successfully improve upon face reconstructions during parameter optimization. We also show that our shape loss is useful when fine-tuning a parameter regression neural network [FFBB21]. The result is a novel method (PSL) that yields improved results on standard single-image face reconstruction benchmarks [SBFB19, CZR*22], when compared to the initialization and other state-of-the-art methods that do similar finetuning with different perceptual loss terms [DBB22, FRPP*22].

Further, as our critic operates only on shaded renders of a 3D face shape, our perceptual shape loss is agnostic to topology of the 3D face mesh, thereby making it suitable for use across a large variety of face models.

6.1. Limitations and Future Work

Our method in its current state provides limited benefits for optimizing or training on side poses, primarily because of the lack of synthetic in-the-wild examples used to train the critic. Currently when using our loss on side poses, we found that the additional landmark term is required to keep the mesh in the correct pose. We however note that this is a limitation of the dataset generation technique [CWZ*21] and not of our method itself. Furthermore, on some occasions we notice that our method does not fully close the mouth of individuals with closed mouth expressions. We think this might come from the mouth crop that we apply to the ground truth geometries and might be alleviated by cropping less of the inner mouth region away in future work. Even though we followed Ramachandran [Ram88] and based our metric on the perception of shading and used gray shaded renders as input to our critic, one might imagine that a more thorough exploration of the light, and material properties used to create the geometry render could lead to improved results. We leave such an exploration to future work. Finally another line of research could explore the design space of critic networks to perhaps obtain local score estimates [SSK20] which could provide finer signals for supervision.

References

- [ACB17] ARJOVSKY M., CHINTALA S., BOTTOU L.: Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning* (06–11 Aug 2017), Precup D., Teh Y. W., (Eds.), vol. 70 of *Proceedings of Machine Learning Research*, PMLR, pp. 214–223. 3
- [aTaHMM16] AN TRÂN A. T., HASSNER T., MASI I., MEDIONI G.: Regressing robust and discriminative 3D morphable models with a very deep neural network. arXiv preprint arXiv:1612.04904 (2016). 2
- [BBB*10] BEELER T., BICKEL B., BEARDSLEY P., SUMNER B., GROSS M.: High-quality single-shot capture of facial geometry. ACM Trans. on Graphics (Proc. SIGGRAPH) 29, 3 (2010), 40:1–40:9. 2, 3
- [BLC*21] BAO L., LIN X., CHEN Y., ZHANG H., WANG S., ZHE X., KANG D., HUANG H., JIANG X., WANG J., YU D., ZHANG Z.: Highfidelity 3d digital human head creation from rgb-d selfies. ACM Transactions on Graphics (2021). 1, 2
- [BRZ*16] BOOTH J., ROUSSOS A., ZAFEIRIOU S., PONNIAH A., DUNAWAY D.: A 3d morphable model learnt from 10,000 faces. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016), pp. 5543–5552. 1
- [BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3d faces. In Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (USA, 1999), SIGGRAPH '99, ACM Press/Addison-Wesley Publishing Co., p. 187–194. 2
- [CBGB20] CHANDRAN P., BRADLEY D., GROSS M., BEELER T.: Semantic deep face models. In 2020 International Conference on 3D Vision (3DV) (2020), pp. 345–354. 1, 3, 4, 10
- [CWZ*14] CAO C., WENG Y., ZHOU S., TONG Y., ZHOU K.: Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics* 20, 3 (2014), 413–425. 1, 10

- [CWZ*21] CHANDRAN P., WINBERG S., ZOSS G., RIVIERE J., GROSS M., GOTARDO P., BRADLEY D.: Rendering with style: Combining traditional and neural approaches for high-quality face rendering. ACM Trans. Graph. 40, 6 (2021). 3, 10
- [CZR*22] CHAI Z., ZHANG H., REN J., KANG D., XU Z., ZHE X., YUAN C., BAO L.: Realy: Rethinking the evaluation of 3d face reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2022). 1, 2, 7, 8, 10
- [DBA*21] DIB A., BHARAJ G., AHN J., THÉBAULT C., GOSSELIN P., ROMEO M., CHEVALLIER L.: Practical face reconstruction via differentiable ray tracing. *Computer Graphics Forum 40*, 2 (2021), 153–164.
- [DBB22] DANECEK R., BLACK M. J., BOLKART T.: EMOCA: Emotion driven monocular face capture and animation. In *Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), pp. 20311–20322. 2, 7, 8, 9, 10
- [DPSD20] DAI H., PEARS N., SMITH W., DUNCAN C.: Statistical modeling of craniofacial shape and texture. Int. J. Comput. Vision 128, 2 (2020), 547–571. 1, 2
- [DSK17] DOU P., SHAH S. K., KAKADIARIS I. A.: End-to-end 3d face reconstruction with deep neural networks. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (Los Alamitos, CA, USA, jul 2017), IEEE Computer Society, pp. 1503–1512. 2
- [DTA*21] 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 Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2021). 2
- [DYX*19] 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 *IEEE Computer Vision and Pattern Recognition Workshops* (2019). 2, 9, 10
- [ERB20] E. RIBA D. MISHKIN D. P. E. R., BRADSKI G.: Kornia: an open source differentiable computer vision library for pytorch. In *Winter Conference on Applications of Computer Vision* (2020). 4
- [EST*20] EGGER B., SMITH W. A. P., TEWARI A., WUHRER S., ZOLLHOEFER M., BEELER T., BERNARD F., BOLKART T., KO-RTYLEWSKI A., ROMDHANI S., THEOBALT C., BLANZ V., VETTER T.: 3d morphable face models—past, present, and future. ACM Trans. Graph. 39, 5 (2020). 2
- [FFBB21] FENG Y., FENG H., BLACK M. J., BOLKART T.: Learning an animatable detailed 3D face model from in-the-wild images. vol. 40. 2, 5, 6, 7, 8, 9, 10
- [FHK*18] FENG Z.-H., HUBER P., KITTLER J., HANCOCK P. J. B., JUN WU X., ZHAO Q., KOPPEN P., RÄTSCH M.: Evaluation of dense 3d reconstruction from 2d face images in the wild. In *13th IEEE International Conference on Automatic Face and Gesture Recognition, FG* 2018, Xi'an, China, May 15-19, 2018 (2018), IEEE Computer Society, pp. 780–786. 2, 4
- [FRPP*22] FILNTISIS P. P., RETSINAS G., PARAPERAS-PAPANTONIOU F., KATSAMANIS A., ROUSSOS A., MARAGOS P.: Visual speech-aware perceptual 3d facial expression reconstruction from videos. 2, 7, 8, 9, 10
- [FWS*18] FENG Y., WU F., SHAO X., WANG Y., ZHOU X.: Joint 3d face reconstruction and dense alignment with position map regression network. In ECCV (2018). 2
- [GAA*17] GULRAJANI I., AHMED F., ARJOVSKY M., DUMOULIN V., COURVILLE A. C.: Improved training of wasserstein gans. In Advances in Neural Information Processing Systems (2017), Guyon I., Luxburg U. V., Bengio S., Wallach H., Fergus R., Vishwanathan S., Garnett R., (Eds.), vol. 30, Curran Associates, Inc. 3, 4
- [GCM*18] GENOVA K., COLE F., MASCHINOT A., SARNA A., VLA-SIC D., FREEMAN W. T.: Unsupervised training for 3d morphable model regression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018). 2

- [GMFB*18] GERIG T., MOREL-FORSTER A., BLUMER C., EGGER B., LUTHI M., SCHOENBORN S., VETTER T.: Morphable face models an open framework. In 2018 13th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2018) (2018), pp. 75–82. 1, 2
- [GPKZ19] 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). 2
- [GZC*16] GARRIDO P., ZOLLHÖFER M., CASAS D., VALGAERTS L., VARANASI K., PÉREZ P., THEOBALT C.: Reconstruction of personalized 3d face rigs from monocular video. 2
- [GZC*19] GUO Y., ZHANG J., CAI J., JIANG B., ZHENG J.: Cnn-based real-time dense face reconstruction with inverse-rendered photo-realistic face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence 41*, 6 (2019), 1294–1307. 10
- [GZY*20] GUO J., ZHU X., YANG Y., YANG F., LEI Z., LI S. Z.: Towards fast, accurate and stable 3d dense face alignment. In *Proceedings* of the European Conference on Computer Vision (ECCV) (2020). 2
- [KB15] KINGMA D. P., BA J.: Adam: A method for stochastic optimization. In 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings (2015). 4, 7
- [KLA*20] KARRAS T., LAINE S., AITTALA M., HELLSTEN J., LEHTI-NEN J., AILA T.: Analyzing and improving the image quality of Style-GAN. In Proc. CVPR (2020). 3
- [KS20] KOIZUMI T., SMITH W. A. P.: "look ma, no landmarks!" unsupervised, model-based dense face alignment. In *European Conference* on Computer Vision (ECCV) (2020), vol. 12347, pp. 690–706. 2
- [KZT*18] KIM H., ZOLLÖFER M., TEWARI A., THIES J., RICHARDT C., THEOBALT C.: Inverse face net: Deep single-shot inverse face rendering from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018). 2
- [LBB*17] LI T., BOLKART T., BLACK M. J., LI H., ROMERO J.: Learning a model of facial shape and expression from 4d scans. ACM Trans. Graph. 36, 6 (nov 2017). 1, 2, 5
- [LLWL20] LEE C.-H., LIU Z., WU L., LUO P.: Maskgan: Towards diverse and interactive facial image manipulation. In *IEEE Conference* on Computer Vision and Pattern Recognition (CVPR) (2020). 7
- [LMFV*21] LI C., MOREL-FORSTER A., VETTER T., EGGER B., KO-RTYLEWSKI A.: To fit or not to fit: Model-based face reconstruction and occlusion segmentation from weak supervision. arXiv preprint arXiv:2106.09614 (2021). 2
- [LRF*23] LEI B., REN J., FENG M., CUI M., XIE X.: A hierarchical representation network for accurate and detailed face reconstruction from in-the-wild images, 2023. 2, 7
- [MHM19] MOLLAHOSSEINI A., HASANI B., MAHOOR M. H.: Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing 10*, 01 (2019), 18–31. 2
- [MKB*16] MCDONAGH S., KLAUDINY M., BRADLEY D., BEELER T., MATTHEWS I., MITCHELL K.: Synthetic prior design for real-time face tracking. In 2016 Fourth International Conference on 3D Vision (3DV) (2016), pp. 639–648. 4
- [MPP22] MA P., PETRIDIS S., PANTIC M.: Visual Speech Recognition for Multiple Languages in the Wild. *Nature Machine Intelligence 4* (2022), 930–939. 2
- [MPS21] MORALES A., PIELLA G., SUKNO F. M.: Survey on 3d face reconstruction from uncalibrated images. *Computer Science Review 40* (2021), 100400. 2
- [NHSW20] NARUNIEC J., HELMINGER L., SCHROERS C., WEBER R.: High-resolution neural face swapping for visual effects. *Computer Graphics Forum* 39, 4 (2020), 173–184. 4

^{© 2023} Eurographics - The European Association

for Computer Graphics and John Wiley & Sons Ltd.

- [PGM*19] PASZKE A., GROSS S., MASSA F., LERER A., BRADBURY J., CHANAN G., KILLEEN T., LIN Z., GIMELSHEIN N., ANTIGA L., DESMAISON A., KÖPF A., YANG E., DEVITO Z., RAISON M., TE-JANI A., CHILAMKURTHY S., STEINER B., FANG L., BAI J., CHIN-TALA S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library. Curran Associates Inc., Red Hook, NY, USA, 2019. 4
- [PKA*09] PAYSAN P., KNOTHE R., AMBERG B., ROMDHANI S., VET-TER T.: A 3d face model for pose and illumination invariant face recognition. IEEE. 1, 2, 9, 10
- [Ram88] RAMACHANDRAN V. S.: Perception of shape from shading. *Nature 331*, 6152 (1988), 163–166. 1, 2, 10
- [RGP*23] RAI A., GUPTA H., PANDEY A., CARRASCO F. V., TAKAGI S. J., AUBEL A., KIM D., PRAKASH A., DE LA TORRE F.: Towards realistic generative 3d face models. arXiv preprint arXiv:2304.12483 (2023). 2
- [RMC16] RADFORD A., METZ L., CHINTALA S.: Unsupervised representation learning with deep convolutional generative adversarial networks. In 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings (2016). 1, 3
- [RRN*20] 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). 1, 4
- [RSK16] RICHARDSON E., SELA M., KIMMEL R.: 3d face reconstruction by learning from synthetic data. In 2016 Fourth International Conference on 3D Vision (3DV) (2016), pp. 460–469. 2
- [RSOEK17] RICHARDSON E., SELA M., OR-EL R., KIMMEL R.: Learning detailed face reconstruction from a single image. 2
- [SBFB19] SANYAL S., BOLKART T., FENG H., BLACK M.: 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)* (June 2019), pp. 7763–7772. 1, 2, 4, 6, 7, 10
- [SRM*20] SHI J., RIBA E., MISHKIN D., MORENO F., NICOLAOU A.: Differentiable data augmentation with kornia, 2020. 4
- [SSK20] SCHONFELD E., SCHIELE B., KHOREVA A.: A u-net based discriminator for generative adversarial networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2020), pp. 8207–8216. 10
- [SSL*20] 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)* (2020), vol. 12360, pp. 53–70. 2
- [THMM17] 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)* (2017), pp. 1599–1608. 2
- [TZG*18] 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). 2
- [TZK*17] 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). 2
- [TZN*15] THIES J., ZOLLHÖFER M., NIESSNER M., VALGAERTS L., STAMMINGER M., THEOBALT C.: Real-time expression transfer for facial reenactment. 2
- [TZS*18] THIES J., ZOLLHÖFER M., STAMMINGER M., THEOBALT C., NIESSNER M.: Face2face: Real-time face capture and reenactment of rgb videos. *Commun. ACM* 62, 1 (dec 2018), 96–104. 2

- [TZX*21] TU X., ZHAO J., XIE M., JIANG Z., BALAMURUGAN A., LUO Y., ZHAO Y., HE L., MA Z., FENG J.: 3d face reconstruction from a single image assisted by 2d face images in the wild. *IEEE Transactions* on Multimedia 23 (2021), 1160–1172. 2
- [UVL16] ULYANOV D., VEDALDI A., LEMPITSKY V.: Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022* (2016). 4
- [WBGB16] WU C., BRADLEY D., GROSS M., BEELER T.: An anatomically-constrained local deformation model for monocular face capture. ACM Trans. Graph. 35, 4 (jul 2016). 2
- [WBH*21] WOOD E., BALTRUŠAITIS T., HEWITT C., DZIADZIO S., CASHMAN T. J., SHOTTON J.: Fake it till you make it: Face analysis in the wild using synthetic data alone. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2021), pp. 3681–3691. 1
- [WBH*22] WOOD E., BALTRUSAITIS T., HEWITT C., JOHNSON M., SHEN J., MILOSAVLJEVIC N., WILDE D., GARBIN S., RAMAN C., SHOTTON J., SHARP T., STOJILJKOVIC I., CASHMAN T., VALENTIN J.: 3d face reconstruction with dense landmarks. *European Conference* on Computer Vision (2022). 2
- [WCY*22] WANG L., CHEN Z., YU T., MA C., LI L., LIU Y.: Faceverse: a fine-grained and detail-controllable 3d face morphable model from a hybrid dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR2022)* (June 2022). 1
- [WXN21] WU C.-Y., XU Q., NEUMANN U.: Synergy between 3DMM and 3D landmarks for accurate 3D facial geometry. In *International Conference on 3D Vision (3DV)* (2021), pp. 453–463. 2
- [YZW*20] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020). 1, 2
- [ZBT22] ZIELONKA W., BOLKART T., THIES J.: Towards metrical reconstruction of human faces. *European Conference on Computer Vision* (2022). 2, 7
- [ZTG*18] 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 37*, 2 (2018), 523–550. 2
- [ZTZ*22] ZHANG H., TIAN Y., ZHANG Y., LI M., AN L., SUN Z., LIU Y.: Pymaf-x: Towards well-aligned full-body model regression from monocular images. arXiv preprint arXiv:2207.06400 (2022). 2