

EMOCA: Emotion Driven Monocular Face Capture and Animation

Radek Daněček

rdanecek@tue.mpg.de

Michael Black

black@tue.mpg.de

Timo Bolkart

tbolkart@tue.mpg.de

Max Planck Institute for Intelligent Systems, Tübingen, Germany




Figure 1. **EMOCA** regresses 3D faces from images with facial geometry that captures the original emotional content. Top row: images of people with challenging expressions. Middle row: coarse shape reconstruction. Bottom row: reconstruction with detailed displacements.

Abstract

As 3D facial avatars become more widely used for communication, it is critical that they faithfully convey emotion. Unfortunately, the best recent methods that regress parametric 3D face models from monocular images are unable to capture the full spectrum of facial expression, such as subtle or extreme emotions. We find the standard reconstruction metrics used for training (landmark reprojection error, photometric error, and face recognition loss) are insufficient to capture high-fidelity expressions. The result is facial geometries that do not match the emotional content of the input image. We address this with **EMOCA** (**EMOTION** Capture and Animation), by introducing a novel deep perceptual emotion consistency loss during training, which helps ensure that the reconstructed 3D expression matches the expression depicted in the input image. While **EMOCA** achieves 3D reconstruction errors that are on par with the current best methods, it significantly outperforms them in terms of the quality of the reconstructed expression and the perceived emotional content. We also directly regress levels of valence and arousal and classify basic expressions from the estimated 3D face parameters. On the task of in-the-wild emotion recognition, our purely geometric approach is on par with the best image-based methods, highlighting the value of 3D geometry in analyzing human behavior. The model and code are publicly available at <https://emoca.is.tue.mpg.de>.

1. Introduction

Teaching computers to see humans and understand their behavior is a long-standing goal of computer vision. To accomplish this, computers need to understand how humans look, how they move, and what they feel. Faces and their emotional expressions provide an important source of information about a person’s internal emotional state. To support automated analysis of emotional state, we capture a person’s face, including its 3D shape, pose, and facial expression, given a single RGB image. To do so, we go beyond prior work to extract 3D geometry that carries rich emotional content. We focus on parametric methods (i.e., animatable and model-based) due to their wide applicability for 3D avatar creation [36], image synthesis [32, 81], video editing [41, 84] and face recognition [9, 68].

The field of 3D face reconstruction has rapidly advanced over the last two decades; see Egger et al. [22] for a review. Existing methods that estimate 3D face models struggle to capture facial expressions in detail and often produce 3D shapes that do not carry the emotional content of the input image. This has several causes. First, some 3D face models lack sufficient expressiveness to capture subtle or extreme expressions. Second, reconstruction metrics like landmark reprojection loss [8], photometric loss [10], face recognition loss [30], or multi-image consistency losses [73, 80], are either not affected by facial expressions, or require perfect image alignment to capture subtle cues. Subtle changes

in geometry, however, can result in large differences in the perceived emotion. We argue that, to recover 3D expression accurately, we need a new reconstruction metric that measures differences in expressions between the 3D reconstruction and the input image.

To that end, we describe EMOCA (EMotion Capture and Animation), a neural network that learns an animatable face model from in-the-wild images without 3D supervision. The design of our method is inspired by advances in the field of facial emotion recognition, which has made tremendous progress to date on estimating affect (or emotion) from in-the-wild-images [50]. Specifically, we train a state-of-the-art emotion recognition model, and leverage this during training of EMOCA as supervision. EMOCA introduces a novel perceptual *emotion consistency loss* that encourages the similarity of emotional content between the input and rendered reconstruction.

While the new emotion consistency loss results in better reconstructed emotions, this alone is insufficient. Large image datasets used by previous 3D reconstruction methods, while containing a large number of subjects of diverse ethnicities, lack emotional expressivity [14, 17, 44, 92]. Large datasets with facial expressions, valence, and arousal in-the-wild, on the other hand, while rich in emotions, do not provide multiple images per subject in diverse conditions [7, 13, 46–48, 57, 91] and smaller datasets in controlled settings are not suitable for deep learning [54–56, 60, 78]. Multiple images of the same person, however, are required to train current state-of-the-art 3D face reconstruction methods [19, 27, 73]. To overcome this, EMOCA builds on top of DECA [27], a publicly available 3D face reconstruction framework that achieves state-of-the-art identity shape reconstruction accuracy [29, 73]. Specifically, we augment DECA’s architecture with an additional trainable prediction branch for facial expression, while keeping other parts fixed. This enables us to only train the expression part of EMOCA on emotion-rich image data [57], which results in improved emotion reconstruction performance, while retaining DECA’s identity face shape quality.

Once trained, EMOCA reconstructs a 3D face from a single image (Fig. 1), it significantly outperforms previous state-of-the-art methods in terms of the reconstructed expression quality, it preserves the state-of-the-art identity shape reconstruction accuracy, and the reconstructed face can be readily animated. Further, the expression parameters regressed by EMOCA convey sufficient information for in-the-wild emotion recognition, with on-par performance with the best image-based methods [86].

In summary, our main contributions are: 1) The first approach to reconstruct an animatable 3D face model from an in-the-wild image, that is capable of recovering facial expressions that convey the correct emotional state. 2) A novel perceptual emotion-consistency loss that rewards the

accuracy of the reconstructed emotion. 3) The first 3D geometry-based framework for in-the-wild emotion recognition, with comparable performance to current state-of-the-art image-based methods. 4) The code and model are publicly available for research purposes at <https://emoca.is.tue.mpg.de>.

2. Related work

Monocular face reconstruction: Reconstructing 3D face shape from images has been studied extensively for more than two decades [22, 99]. Model-free approaches directly regress 3D meshes [18, 20, 28, 33, 40, 70, 74, 79, 93, 95, 97] or voxels [38] from an image, or optimize a Signed Distance Function (SDF) [61] to fit a face image. Most of these methods require explicit 3D supervision during training. While the output is model-free, acquiring the training data typically relies on a 3D face model (3D Morphable Model, or 3DMM). Thus their ability to reconstruct expressive faces may be limited by the 3DMM-based reconstruction used to generate the paired training data [18, 28, 33, 38, 40, 70, 93], the domain gap between 3DMM-based synthetic training data and real images [20, 74, 97], or the regularization towards a fixed 3DMM fitting result [16]. Instead, EMOCA is trained in a self-supervised fashion without any explicit 3D supervision, which enables it to capture less constrained expressions. Other self-supervised methods do not leverage face-domain-specific knowledge, which makes them applicable to general objects, but also limits the reconstruction quality [79, 95]. Unlike EMOCA, none of these model-free methods separate facial identity from facial expression, making them inappropriate for applications like expression re-targeting or animation.

Several works reconstruct the parameters of fixed statistical models like the Basel Face Model (BFM) [62], FaceWarehouse [12], or FLAME [51], or jointly learn a model and reconstruct faces from images [80, 82, 89]. Existing methods can be categorized into optimization-based [4, 5, 9, 10, 31, 45, 63, 69, 84, 90] and learning-based. The latter are trained fully supervised [15, 34, 42, 67, 87, 88, 98] or self-supervised with predicted 2D keypoints [19, 27, 52, 73, 76, 80, 82, 83, 96], 2D face contours [52], photometric constraints [19, 27, 30, 76, 80, 82, 83, 96], face recognition features [19, 27, 30, 76], multi-view constraints [76], or multi-image constraints [19, 27, 30, 73, 80]. Each supervision signal impacts the reconstructed 3D face in a unique way. Explicit 3D mesh or model parameter supervision induces a bias towards the method used to generate the pseudo-ground truth. Using face recognition features or leveraging multiple images of the same identity during training mainly impacts identity shape and appearance. Keypoint losses impact the facial geometry and image alignment (global transformation, identity and expression shape parameters), but predicted keypoints are sparse (commonly 51–68 points),

often inaccurate - especially for extreme expressions and head poses - and obtaining the optimal embedding of the corresponding keypoints on the model’s surface is challenging. Photometric losses impact all model parameters (global transformation, identity and expression shape, appearance, and lighting), but, as with the keypoint losses, are strongly affected by misalignments between the predicted 3D face and the image. While using multi-view data during training has the potential to reconstruct more accurate 3D faces, there are no large datasets with a large number of identities and large diversity in expression, ethnicity, age, lighting conditions, etc. Consequently, while the field of monocular in-the-wild face capture has made tremendous progress, there are still limitations, particularly in the accuracy of the reconstructed expressions, which limit the emotions that can be perceived from the reconstructed 3D shapes. EMOCA instead learns to reconstruct expressive faces by combining emotion features that mainly propagate to the reconstructed expression, with a unique self-supervised framework that enables us to leverage a large dataset of diverse expressions.

Emotion analysis from images: Emotion analysis is a long-standing problem in computer vision and related fields (see [3, 11] for comprehensive surveys). Emotional states are commonly represented as discrete basic [23, 24] (e.g., Happiness, Surprise, ...) or compound expression categories [21] (e.g., happily surprised), continuous valence (positive-negative) and arousal (relaxed-intensive) values [71], or Facial Action Units (FACS) activations [25], where each action unit (AU) corresponds to a particular emotion-related facial muscle movement.

Early work on expression recognition extracts geometric features defining shape and location of face components [59, 85], appearance features [26, 75], or combinations of these [39, Chapter 19]. Over the last decade, the availability of large datasets for single-image expression analysis [7, 57] and audio-visual videos [46–48] shifted the focus from manually designed features to end-to-end trained models [50]. While early work like Wen and Huang [94] uses 3D non-rigid surface tracking to extract features for expression reconstruction, the majority of 3D-based methods focus on recognizing expressions from 3D scans [58, 72]. Among these, the most relevant to EMOCA, is [65] as they use 3DMM features to classify three expressions (obtained by fitting the 3DMM to the scans); most other methods use diverse 2D and 3D features extracted from the textured 3D scans.

Few 3DMM-based methods exist to recognize expressions from images. Bejaoui et al. [6] fit a 3DMM to images, while Chang et al. [15] and Koujan et al. [49] train a 3DMM parameter regressor, fully-supervised by parameters obtained by fitting a 3DMM to images and videos. From the 3DMM expression parameters, they then learn to

classify different expressions. Most related to EMOCA, Shi et al. [77] use an expression recognition loss during training, but with the goal of obtaining a more discriminative latent representation. These methods focus on recognizing expressions, not improving 3D reconstruction. In contrast, EMOCA leverages recent advances in emotion recognition to reconstruct more expressive 3D faces.

3. Preliminaries

Face model: FLAME [51] is a statistical 3D head model with parameters for identity shape $\beta \in \mathbb{R}^{|\beta|}$, facial expression $\psi \in \mathbb{R}^{|\psi|}$, and pose parameters $\theta \in \mathbb{R}^{3k+3}$ for rotations around $k = 4$ joints (neck, jaw, and eyeballs) and the global rotation. Given all parameters, FLAME outputs a mesh with $n_v = 5023$ vertices. Formally, FLAME is:

$$M(\beta, \theta, \psi) \rightarrow (\mathbf{V}, \mathbf{F}), \quad (1)$$

with vertices $\mathbf{V} \in \mathbb{R}^{n_v \times 3}$ and $n_f = 9976$ faces $\mathbf{F} \in \mathbb{R}^{n_f \times 3}$. FLAME comes with an appearance model, converted from Basel Face Model’s albedo space [62] to FLAME’s UV layout [1]. Given parameters $\alpha \in \mathbb{R}^{|\alpha|}$, this model outputs a FLAME texture map $A(\alpha) \in \mathbb{R}^{d \times d \times 3}$.

Face reconstruction: DECA [27] is a publicly available framework to reconstruct a detailed, animatable 3D face model from a single image. We follow DECA’s notation for simplicity. Given an image I , the coarse encoder

$$E_c(I) \rightarrow (\beta, \theta, \psi, \alpha, \mathbf{l}, \mathbf{c}) \quad (2)$$

outputs FLAME geometry parameters β, θ, ψ , albedo α , Spherical Harmonics (SH) [64] lighting $\mathbf{l} \in \mathbb{R}^{27}$, and camera $\mathbf{c} \in \mathbb{R}^3$, which is the concatenation of isotropic scale $s \in \mathbb{R}$ and translation $\mathbf{t} \in \mathbb{R}^2$. The detail encoder

$$E_d(I) \rightarrow \delta \quad (3)$$

encodes I to a subject-specific detail vector $\delta \in \mathbb{R}^{128}$. To reconstruct dynamic expression wrinkles, the detail decoder

$$F_d(\delta, \psi, \theta_{jaw}) \rightarrow D \quad (4)$$

uses δ to parametrize static person-specific details, and FLAME’s expression ψ and jaw-pose parameters θ_{jaw} to generate an expression-dependent detail UV displacement map $D \in \mathbb{R}^{d \times d \times 3}$.

Denoting the rendering function with R [66], the coarse shape can be rendered to a 2D image as $R(M(\beta, \theta, \psi), \alpha, \mathbf{l}, \mathbf{c}) \rightarrow I_{Rc}$. To render the FLAME mesh, with expression-dependent details, to an image I_{Rd} , the D are converted to a detailed normal map N_d , and provided as additional parameters to R ; formally $R(M(\beta, \theta, \psi), \alpha, \mathbf{l}, \mathbf{c}, N_d) \rightarrow I_{Rd}$.

Relative keypoint loss: Given 2D face keypoints $\mathbf{k}_i \in \mathbb{R}^2$ and the corresponding keypoints on FLAME’s mesh surface $M_i \in \mathbb{R}^3$, the relative keypoint loss [27] computes

offset vectors between pairs of 2D keypoints and between the corresponding pairs of projected model keypoints, and penalizes the difference. Formally, the loss computes as

$$L_{rk}^E = \sum_{(i,j) \in E} \|\mathbf{k}_i - \mathbf{k}_j - s\Pi(M_i - M_j)\|_1, \quad (5)$$

where E is a set of landmark index pairs, and $\Pi \in \mathbb{R}^{2 \times 3}$ is the orthographic 3D-2D projection matrix.

Emotion recognition: For the emotion network, we use ResNet-50 as the backbone, with a fully connected prediction head that outputs expression classification, valence, and arousal. See Sup. Mat. for experiments with other backbones. The network is trained on AffectNet [57], a large-scale annotated emotion dataset. We adapt the training setting from Toisoul et al. [86] with minor modifications as described in the Sup. Mat. The loss function consists of several terms such as categorical cross entropy for expression classification, mean squared error and correlation coefficient losses for valence and arousal; see Sup. Mat. for details of the losses. After the network is trained, prediction heads are discarded, and the features of the final layer of the backbone network serve as our emotion feature $\epsilon \in \mathbb{R}^{|\epsilon|}$. We denote the emotion network as $A(I) \rightarrow \epsilon$.

4. Method: EMOCA

The main goal of EMOCA is to address a significant limitation of the prior art - to recover 3D face shapes from single images that convey the full spectrum of emotion. Our technical contribution is twofold, first, we introduce a novel *emotion consistency loss* that is designed to encourage *emotion similarity* between the input image and the output rendering as training supervision. Second, we leverage parts of DECA’s [27] trained model in order to only train the expression part of EMOCA on emotion-rich image data, while preserving DECA’s identity shape reconstruction performance.

Architecture: EMOCA’s architecture is based on DECA [27]. As with many state-of-the-art methods, DECA takes an input image and uses several neural networks to factor it into shape, albedo, lighting, etc. Given these factors, one can differentially render an output image that should look like the input. Here we exploit this output image in a novel way by encouraging it to have the same *expression* as the input image.

Training models like DECA [27] on emotion-rich image data [57] is infeasible, due to DECA’s requirement of multiple training images per subject to regularize the training of the identity shape reconstruction of E_c (Eq. 2). Instead, we augment DECA’s architecture with an additional expression encoder

$$E_e(I) \rightarrow \psi_e, \quad (6)$$

and keep the weights of E_c fixed during training, thereby retaining the predictions of $\beta, \theta, \alpha, \mathbf{l}$ and \mathbf{c} from DECA, but discarding DECA’s ψ . Further, let $R(M(\beta, \theta, \psi_e), \alpha, \mathbf{l}, \mathbf{c}) \rightarrow I_{Re}$ denote the rendering of the output of E_c with the expression of the input image, $E_e(I)$.

For an overview of the model architecture, see Figure 2. Training E_e only has several advantages: 1) Training datasets are not required to contain multiple images per subject. 2) Not training the identity prediction enables us to remove the face recognition loss. 3) Having fixed pose, shape, and camera parameters allows us to remove the landmark reprojection loss. 4) This results in reduced training resources, faster training time, and reduced memory consumption due to the lower number of training parameters.

Loss function: In total, we optimize:

$$L = \lambda_{emo}L_{emo} + \lambda_{pho}L_{pho} + \lambda_{eye}L_{eye} + \lambda_{mc}L_{mc} + \lambda_{lc}L_{lc} + \lambda_{\psi}L_{\psi} \quad (7)$$

with emotion consistency loss L_{emo} , photometric loss L_{pho} , eye closure loss L_{eye} , mouth closure loss L_{mc} , lip corner loss L_{lc} , and expression regularizer L_{ψ} , each weighted by a factor λ_x .

Emotion consistency loss: The emotion consistency loss computes the difference between the emotion features of the input image $\epsilon_I = A(I)$ and those of the rendered image, $\epsilon_{Re} = A(I_{Re})$ as:

$$L_{emo} = d(\epsilon_I, \epsilon_{Re}), \quad (8)$$

with $d(\epsilon_1, \epsilon_2) = \|\epsilon_1 - \epsilon_2\|_2$. Instead of measuring a geometric error, L_{emo} computes a perceptual difference between the input image and the rendered image. Optimizing this loss during training ensures that the reconstructed 3D face conveys the emotional content of the input image.

Photometric loss: The photometric loss computes the pixel error between the input image I and the output rendering I_{Re} . $L_{pho} = \|V_I \odot (I - I_{Re})\|_{1,1}$. V_I denotes a rendered mask of the output face shape, with each pixel located in the face skin region is equal to 1, and 0 elsewhere. The operator \odot denotes the Hadamard product.

Eye closure loss: The eye closure loss computes as $L_{eye} = L_{rk}^{E_{eye}}$, where E_{eye} is a set of upper/lower eyelid keypoint pairs. Due to slight misalignment between image landmarks and projected 3D landmarks, enforcing standard landmark reprojection losses produces incorrect predictions. Instead, using (translation-invariant) relative keypoint losses (for eye closure, mouth closure, and mouth width) is less susceptible to misalignments.

Mouth closure loss: The loss computes as $L_{mc} = L_{rk}^{E_{mc}}$, where E_{mc} is a set of upper/lower lip keypoint pairs.

Lip corner loss: The lip corner loss computes as $L_{lc} = L_{rk}^{E_{lc}}$, where E_{lc} is the pair of left and right lip corners.

Expression regularization: The expression is regularized as $L_{\psi} = \|\psi\|_2^2$.




Figure 2. EMOCA overview. For the coarse training stage (green box), the input image is fed to the coarse shape encoder (initialized from DECA [27] and fixed) and EMOCA’s trainable expression shape encoder. A textured 3D mesh is then reconstructed from the regressed identity shape, expression shape, pose, and albedo parameters with FLAME’s geometry and albedo models as fixed decoders. This textured mesh is rendered by a differentiable renderer with the regressed camera and spherical harmonics lighting. Our novel emotion consistency loss (Eq. 8) penalizes the difference between the emotion features of the input image and those of the rendered coarse shape, after passing both images through a fixed emotion recognition network. For the detail training stage (yellow box), EMOCA’s expression encoder is fixed, and the regressed expression (and jaw-pose) parameters are used to condition the detail decoder.

Detailed stage: The detail training stage adds wrinkle details that are animatable. Here we follow DECA’s design, and use the same architecture and losses.

5. Experiments

5.1. Training setting

The first stage (coarse part) of EMOCA is trained with AffectNet [57] for a maximum of 20 epochs, with early stopping, using the Adam optimizer [43] and a learning rate of $5e - 5$. We use the same training/validation/testing split as proposed by [86]. We set $\lambda_{emo} = 1$, $\lambda_{pho} = 2$, $\lambda_{eye} = \lambda_{lc} = \lambda_{mc} = 0.5$ and $\lambda_{\psi} = 1e - 4$. EMOCA’s second stage (detail part) training is comparable to DECA’s second stage training. We use the same training data [14, 17] and train with the same settings. Please refer to Sup. Mat. for more training details.

5.2. Quantitative evaluation

While, for the task of 3D face reconstruction, standard benchmarks exist to quantitatively evaluate the identity face shape [29, 73], no such benchmark exists to assess the accuracy of the reconstructed expression. Unlike the identity shape benchmarks, quantitatively measuring the difference between a reconstructed 3D facial expression and a ground truth scan is less meaningful. The errors would be dom-

inated by errors of the reconstructed identity face shape, and a low geometric error would not necessarily correspond to a small difference in human perception of the emotion. Instead, we evaluate EMOCA 1) qualitatively, 2) quantitatively for the task of in-the-wild emotion recognition, and 3) perceptually in an Amazon Mechanical Turk (AMT) study.

Emotion recognition: Our goal is to quantify how much of the input emotion is conveyed in the reconstructed 3D face. For this, we apply 3D face reconstruction methods to in-the-wild emotional face images and evaluate emotion recognition accuracy based on the 3D reconstruction. Here we focus on methods that reconstruct a parametric model of the face, i.e. a 3DMM. To that end, for each 3D face reconstruction method, we train a 4-layer MLP with Batch Normalization [37] and LeakyReLUs to regress valence and arousal levels, and classify expression labels directly from the predicted 3DMM parameters. The training details are described in Sup. Mat.

We evaluate emotion recognition on the AffectNet test set [57] and the AFEW-VA test set [47]. For each method, we report Concordance correlation coefficients (CCC \uparrow), Pearson correlation coefficients (PCC \uparrow), root mean squared error (RMSE \downarrow), and sign agreement (SAGR \uparrow) for valence (V) and arousal (A) regression and accuracy for expression (E) classification on the test set defined by [86]. EMOCA outperforms all 3D face reconstruction methods, and is on

Model	V-PCC \uparrow	V-CCC \uparrow	V-RMSE \downarrow	V-SAGR \uparrow	A-PCC \uparrow	A-CCC \uparrow	A-RMSE \downarrow	A-SAGR \uparrow	E-ACC \uparrow
EmoNet [86]	0.75	0.73	0.32	0.80	0.68	0.65	0.29	0.78	0.68
Deep3DFace [19]	0.75	0.73	0.33	0.80	0.66	0.65	0.31	0.78	0.65
ExpNet [15]	0.45	0.42	0.43	0.73	0.39	0.36	0.38	0.64	0.46
MGCNet [76]	0.71	0.69	0.35	0.80	0.59	0.58	0.34	0.77	0.60
3DDFA_V2 [34]	0.63	0.62	0.39	0.75	0.53	0.50	0.34	0.73	0.52
DECA [27]	0.70	0.69	0.36	0.76	0.59	0.58	0.33	0.74	0.59
DECA w/ details [27]	0.70	0.69	0.37	0.77	0.59	0.57	0.33	0.77	0.58
EMOCA (Ours)	0.78	0.77	0.31	0.81	<i>0.69</i>	<i>0.68</i>	<i>0.30</i>	<i>0.81</i>	<i>0.68</i>
EMOCA w/ details (Ours)	0.77	0.76	0.31	0.81	0.70	0.69	0.29	0.83	0.69

Table 1. **Emotion recognition performance on the AffectNet test set [57]**. The EmoNet performance is measured using the model that is publicly released by the authors. For EMOCA and the other 3D baselines, we train the recognition module as described in Sec. 5.2. DECA w/ detail means that DECA’s detail code prediction was included in the input to the regressor, along with the 3DMM parameters. Please note that EMOCA’s performance is on par with EmoNet and it outperforms all other 3D reconstruction-based methods.




Figure 3. Comparison of **coarse reconstruction** methods, from left to right: Input, 3DDFA_V2 [34], MGCNet [76], Deng et al. [19], DECA [27] (coarse), and EMOCA (coarse). EMOCA conveys the emotions of the input images better than other methods.

par with the image-based state-of-the-art [86]. For details, see Tab. 1 for results on the AffectNet dataset and Tab. 1 of Sup. Mat. for the AFEW-VA dataset.

Note that EMOCA performs on par with EmoNet [86], which is a recent method for estimating emotion from images. This confirms that the emotional content is present in our 3D reconstruction and that 3D shape is sufficient to understand emotion. This has implications for future research on emotion recognition.

Perceptual study: The 3D geometry reconstructed from an image must convey the emotion of the input image. Directly comparing rendered geometry with an image is difficult due to the domain gap. Instead, we perform a perceptual study using AMT to assess the perceived expression of emotion from rendered 3D reconstructions. Specifically,

given an image, we ask participants to categorize the perceived expression of emotion into one of the 7 basic emotions (Anger, Disgust, Fear, Happiness, Sadness, Surprise, and Contempt) or as a neutral expression (no emotion). A single evaluation task contains 75 images in random order; 35 real images, the 35 corresponding rendered reconstructions (from one method), and 5 qualification samples. The 5 qualification samples are duplicates sampled from the 35 real images, and they are chosen to be of easily recognizable emotion. Each task is performed by 10 participants. Participants that either misclassify the emotion, or inconsistently label the duplicated images, for at least 2 (of the 5) qualification samples are discarded from further analysis to filter out inattentive and/or uncooperative participants. For each method, we measure the classification consistency between

each participant’s labels for the rendered images and their labels for the corresponding real images. If the rendered 3D meshes contain the emotional content of the images, then the scores given to both should be consistent.

We select 35 images with balanced emotional content (i.e., 5 images per basic emotion) from the AffectNet test set [57]. For each image, we reconstruct 3D faces using EMOCA, DECA [27], Deep3DFace [19], MGCNet [76] and 3DDFA_V2 [34]. The classification consistency averaged across participants for each method are: EMOCA (coarse) 0.68, EMOCA (detail) 0.65, Deep3DFace 0.37, DECA (coarse) 0.33, DECA (detail) 0.31, MGCNet 0.32, 3DDFA_V2 0.31. In summary, EMOCA preserves the emotional content of images better than the other methods. Note that, perhaps surprisingly, there is little difference between the scores for the coarse meshes of EMOCA/DECA and the detailed ones. Despite more wrinkle detail, our perceptual experiments suggest that the detailed meshes do not convey more emotional content. One possible explanation is that, in addition to adding valid wrinkle details, the detail generator sometimes adds artifacts in the lip region (e.g. Fig. 1, col. 1 & 3), and hallucinates details in the forehead (e.g. Fig. 1, col. 8). These could negatively impact participants’ perception. For the full confusion tables, see the Sup. Mat.

Emotion recognition vs. perceptual study: There is a considerable discrepancy between the results of the automatic emotion recognition results and the perceptual study results, in particular for Deep3DFace [19]. Deep3DFace performs much better on the emotion recognition task (slightly below SOTA), than on the perceptual study. Unlike EMOCA, it is not capable of producing highly emotional reconstructions (see Fig. 3). We hypothesize that the automatic predictors are capable of detecting more subtle cues than humans. We investigate this by measuring the agreement (i.e., percentage of matching predictions) between the method’s classifier (from the reconstructed face parameters) and the participant’s annotation of the *input images* from the perceptual study. The results are: EMOCA 62% and Deep3DFace 62%. This indicates that the predicted parameters for both methods contain a similar amount of information about the emotions compared to the annotations of the input images. However, the agreement between the method’s classifier and the participant’s annotation of the *rendered reconstructions* is for EMOCA 48%, and for Deep3DFace 26%. In other words, EMOCA is significantly more in agreement with human perception.

5.3. Qualitative evaluation

We provide a visual comparison of the coarse shape reconstruction methods in Fig. 3. Observe that EMOCA outperforms all the previous methods in terms of capturing the emotional content of the original image in the reconstructed expression. In Fig. 4 we compare our detail reconstruction




Figure 4. Comparison of 3D reconstructions with **detail displacements**. Top: Input, Middle: DECA [27], Bottom: EMOCA. EMOCA results contain more expression-dependent details that better convey the emotion of the input images than DECA.

tions to DECA’s detail reconstruction. Compared to DECA, our detailed displacement better captures the fine details of highly emotional input images.

5.4. Ablation experiment

Table 2 shows the effect of ablating the training data and the emotion consistency loss. The table summarizes the effect of EMOCA trained w/ and w/o the emotion consistency loss, and using the DECA data only [14, 17] instead of the AffectNet training data [57].

6. Discussion and limitations

Baseline: EMOCA builds on top of DECA due to its state-of-the-art identity shape reconstruction performance. We found in our experiments that the recently released Deep3DFaceRecon [2] gives better 3D face reconstructions than reported in the paper [19], and in some cases, it outperforms DECA in terms of the reconstructed expression. Combining our emotion consistency loss with the Deep3DFaceRecon framework to further improve their reconstructed expressions is worth further investigation.

Image alignment: DECA sometimes predicts 3D faces that are slightly misaligned with the input images. EMOCA inherits this limitation due to the fixed coarse shape encoder. Further, while EMOCA reconstructs more expressive faces that better convey the emotion of the input image, expressions are also sometimes misaligned. Mitigating these artifacts, and better balancing the trade-off between geometric alignment and emotion similarity, requires further work.

Emotion embedding analysis: We assume that the emotion embedding extracted by the emotion recognition network has desirable properties to guide the optimization of FLAME’s expression parameters. We found that the emotion recognition loss is more difficult to optimize and it requires more careful weighting of the loss compared to the identity recognition losses used in previous work [19, 27, 30]. Directly using the pre-trained EmoNet [86]

Model	V-PCC \uparrow	V-CCC \uparrow	V-RMSE \downarrow	V-SAGR \uparrow	A-PCC \uparrow	A-CCC \uparrow	A-RMSE \downarrow	A-SAGR \uparrow	E-ACC \uparrow
DECA [27]	0.70	0.69	0.36	0.76	0.59	0.58	0.33	0.74	0.59
EMOCA DS w/o Emo	0.70	0.69	0.37	0.78	0.61	0.58	0.32	0.79	0.60
EMOCA w/o Emo	0.68	0.66	0.36	0.74	0.59	0.58	0.32	0.77	0.59
EMOCA DS	0.77	0.76	0.31	0.82	0.69	0.67	0.29	0.79	0.68
EMOCA	0.78	0.77	0.31	0.81	0.69	0.68	0.30	0.81	0.68

Table 2. **Ablation experiment.** Effect of ablating the data and the emotion consistency loss on EMOCA evaluated on the emotion recognition task. From top to bottom, we see the performance of DECA, EMOCA trained on the DECA dataset w/o emotion loss, EMOCA w/o emotion loss, EMOCA trained on the DECA dataset, and EMOCA. We refer to DECA’s training data as DECA dataset (DS), which is a combination of VGGFace2 [14], and VoxCeleb2 [17]. The key finding is that novel emotion consistency loss is critical for the performance of this task, as with emotion loss, EMOCA’s performance improves. Finetuning on AffectNet, which has a much richer variety of facial expressions, only marginally increases in performance over training on DECA’s original training data (DS).

for instance did not provide sufficient supervision. However, our work is the first to demonstrate how to use emotion recognition features to guide the task of 3D geometry reconstruction. In addition using our emotion consistency loss to train EMOCA, we have experimented with the applicability of emotion features for the tasks of emotion retrieval and emotion retargetting via FLAME expression parameter optimization (see Sup. Mat.).

Emotion network architecture: Using a pre-trained state-of-the-art emotion recognition network [86] does not provide satisfactory supervision during optimization or training. Instead, it produces strong artifacts in the reconstructed geometry. To overcome this, we investigate different ResNet [35] and Swin Transformer [53] based emotion network architectures, and show the effect of different networks in the Sup. Mat. Based on this analysis, we use a ResNet-50 backbone for our emotion network.

Jaw rotations: While FLAME’s jaw rotation parameters θ_{jaw} contribute to facial expressions, we found the optimization of θ_{jaw} to be unstable while training EMOCA. We hypothesize, that this is due to the lack of a good prior for the jaw rotation. However, using different simplified priors for the jaw pose like a simple L2 regularizer did not give satisfactory results. We offer a more detailed discussion in Sup. Mat. Investigating the effect of more advanced data-driven jaw priors when optimizing the emotion loss is subject to future work.

Implementation details: For details on all hyper parameters and discussion on design choices see the Sup. Mat.

7. Conclusions

We have presented EMOCA, a method that takes a single in-the-wild image and reconstructs a 3D face with sufficient facial expression detail to convey the emotional state of the input image. EMOCA is trained in a self-supervised fashion from a large dataset of emotion-rich images. A novel *emotion similarity* loss provides supervision on the reconstructed expressions during training. The emotion similarity relies on deep features extracted from a neural network

trained for single-image affect (emotion) recognition in-the-wild. EMOCA reconstructs 3D face shape on par with current state-of-the-art methods but outperforms them in terms of the quality of the reconstructed expression. Further, using the reconstructed expression parameters for the task of in-the-wild emotion recognition, EMOCA outperforms existing 3DMM-based face reconstruction methods and gives on par results with the best purely image-based method.

In summary, this is the first in-the-wild monocular face reconstruction work that puts explicit emphasis on the *perceptual quality* of the expression and the emotion it communicates instead of standard geometric and photometric losses. This presents a new direction for the monocular face reconstruction community. This work has potential to further combine the fields of monocular 3D face reconstruction and emotion analysis. Further, downstream application of this work can be employed in the industry, including but not limited to gaming, movies, AR/VR and communication.

Of course, any improvement to 3D face acquisition and animation may also enable more realistic ‘deep fakes.’ Subtle emotional cues are individualistic and reproducing these could make it harder to detect such fakes. While cognizant of the risks, we are also sensitive to the importance of facial emotion in human communication. The trend towards emotional avatars in games and communication is clear. If communicative avatars do not properly communicate emotion, that, in itself presents a risk of misunderstandings.

Acknowledgement: We thank Y. Feng and H. Feng for DECA support and helpful discussions, T. Alexiadis and T. McConnell for help with the perceptual study, S. Zuffi, S. Sanyal, O. Bendov, N. Andreou, P. Patel and P. Forte for proofreading, and A. Toisoul for EmoNet discussions. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No.860768 (CLIFE project).

Disclosure: MJB has received research gift funds from Adobe, Intel, Nvidia, Facebook, and Amazon. MJB has financial interests in Amazon, Datagen Technologies, and Meshcapade GmbH. While TB is part-time employee of Amazon, this research was performed solely at, and funded solely by, MPI.

References

- [1] BFM_to_FLAME. https://github.com/TimoBolkart/BFM_to_FLAME, 2021. 3
- [2] Deep3DFaceRecon_PyTorch. <https://github.com/microsoft/Deep3DFaceReconstruction>, 2021. 7
- [3] Xavier Alameda-Pineda, Elisa Ricci, and Nicu Sebe. Multimodal behavior analysis in the wild: An introduction. In Xavier Alameda-Pineda, Elisa Ricci, and Nicu Sebe, editors, *Multimodal Behavior Analysis in the Wild*, Computer Vision and Pattern Recognition, pages 1–8. Academic Press, 2019. 3
- [4] Oswald Aldrian and William AP Smith. Inverse rendering of faces with a 3D morphable model. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 35(5):1080–1093, 2013. 2
- [5] Anil Bas, William A. P. Smith, Timo Bolkart, and Stefanie Wuhrer. Fitting a 3D morphable model to edges: A comparison between hard and soft correspondences. In *Asian Conference on Computer Vision Workshops*, pages 377–391, 2017. 2
- [6] Hela Bejaoui, Haythem Ghazouani, and Walid Barhoumi. Fully automated facial expression recognition using 3D morphable model and mesh-local binary pattern. In *Advanced Concepts for Intelligent Vision Systems*, volume 10617, pages 39–50, 2017. 3
- [7] Carlos Fabian Benítez-Quiroz, Ramprakash Srinivasan, and Aleix M. Martínez. EmotioNet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5562–5570, 2016. 2, 3
- [8] Volker Blanz, Curzio Basso, Tomaso A. Poggio, and Thomas Vetter. Reanimating faces in images and video. *Computer Graphics Forum (Proc. EUROGRAPHICS)*, 22(3):641–650, 2003. 1
- [9] Volker Blanz, Sami Romdhani, and Thomas Vetter. Face identification across different poses and illuminations with a 3D morphable model. In *International Conference on Automatic Face & Gesture Recognition (FG)*, pages 202–207, 2002. 1, 2
- [10] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, pages 187–194, 1999. 1, 2
- [11] Rafael A. Calvo, Sidney D’Mello, Jonathan Gratch, and Arvid Kappas. *The Oxford Handbook of Affective Computing*. Oxford University Press, Inc., USA, 1st edition, 2014. 3
- [12] Chen Cao, Yanlin Weng, Shun Zhou, Yiyang Tong, and Keliang Zhou. FaceWarehouse: A 3D facial expression database for visual computing. *Transactions on Visualization and Computer Graphics*, 20:413–425, 2014. 2
- [13] Houwei Cao, David Cooper, Michael Keutmann, Ruben Gur, Ani Nenkova, and Ragini Verma. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on affective computing*, 5:377–390, 10 2014. 2
- [14] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. VGGFace2: A dataset for recognising faces across pose and age. In *International Conference on Automatic Face & Gesture Recognition (FG)*, pages 67–74, 2018. 2, 5, 7, 8
- [15] Feng-Ju Chang, Anh Tuan Tran, Tal Hassner, Iacopo Masi, Ram Nevatia, and Gerard Medioni. ExpNet: Landmark-free, deep, 3D facial expressions. In *International Conference on Automatic Face & Gesture Recognition (FG)*, pages 122–129, 2018. 2, 3, 6
- [16] Aggelina Chatziagapi, ShahRukh Athar, Francesc Moreno-Noguer, and Dimitris Samaras. SIDER: single-image neural optimization for facial geometric detail recovery. In *International Conference on 3D Vision (3DV)*, pages 815–824, 2021. 2
- [17] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. VoxCeleb2: Deep speaker recognition. In *Conference of the International Speech Communication Association (INTER-SPEECH)*, pages 1086–1090, 2018. 2, 5, 7, 8
- [18] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. RetinaFace: Single-shot multi-level face localisation in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5202–5211, 2020. 2
- [19] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. 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)*, pages 285–295, 2019. 2, 6, 7
- [20] Pengfei Dou, Shishir K Shah, and Ioannis A Kakadiaris. End-to-end 3D face reconstruction with deep neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5908–5917, 2017. 2
- [21] Shichuan Du, Yong Tao, and Aleix M Martínez. Compound facial expressions of emotion. *National Academy of Sciences*, 111(15):E1454–E1462, 2014. 3
- [22] Bernhard Egger, William A. P. Smith, Ayush Tewari, Stefanie Wuhrer, Michael Zollhöfer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, Christian Theobalt, Volker Blanz, and Thomas Vetter. 3D morphable face models - past, present, and future. *Transactions on Graphics (TOG)*, 39(5):157:1–157:38, 2020. 1, 2
- [23] Paul Ekman. An argument for basic emotions. *Cognition and Emotion*, 6(3-4):169–200, 1992. 3
- [24] Paul Ekman and Wallace V Friesen. Constants across cultures in the face and emotion. *Journal of personality and social psychology*, 17(2):124, 1971. 3
- [25] Paul Ekman and Wallace V. Friesen. Facial action coding system: A technique for the measurement of facial movement. *Palo Alto: Consulting Psychologists Press*, 1978. 3
- [26] Xiaoyi Feng, M Pietikainen, and Abdenour Hadid. Facial expression recognition with local binary patterns and linear programming. *Pattern Recognition And Image Analysis*, 15(2):546, 2005. 3
- [27] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. *Transactions on Graphics, (Proc. SIGGRAPH)*, 40(4):88:1–88:13, 2021. 2, 3, 4, 5, 6, 7, 8

- [28] Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. Joint 3D face reconstruction and dense alignment with position map regression network. In *European Conference on Computer Vision (ECCV)*, pages 534–551, 2018. 2
- [29] Zhen-Hua Feng, Patrik Huber, Josef Kittler, Peter Hancock, Xiao-Jun Wu, Qijun Zhao, Paul Koppen, and Matthias Rätzsch. Evaluation of dense 3D reconstruction from 2D face images in the wild. In *International Conference on Automatic Face & Gesture Recognition (FG)*, pages 780–786, 2018. 2, 5
- [30] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T. Freeman. Unsupervised training for 3D morphable model regression. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8377–8386, 2018. 1, 2, 7
- [31] Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Luthi, Sandro Schönborn, and Thomas Vetter. Morphable face models-an open framework. In *International Conference on Automatic Face & Gesture Recognition (FG)*, pages 75–82, 2018. 2
- [32] Partha Ghosh, Pravir Singh Gupta, Roy Uziel, Anurag Ranjan, Michael J. Black, and Timo Bolkart. GIF: Generative interpretable faces. In *International Conference on 3D Vision (3DV)*, pages 868–878, 2020. 1
- [33] Riza Alp Güler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, and Iasonas Kokkinos. DenseReg: Fully convolutional dense shape regression in-the-wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6799–6808, 2017. 2
- [34] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3D dense face alignment. In *European Conference on Computer Vision (ECCV)*, pages 152–168, 2020. 2, 6, 7
- [35] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 8
- [36] Liwen Hu, Shunsuke Saito, Lingyu Wei, Koki Nagano, Jaewoo Seo, Jens Fursund, Iman Sadeghi, Carrie Sun, Yen-Chun Chen, and Hao Li. Avatar digitization from a single image for real-time rendering. *Transactions on Graphics (TOG)*, 36(6):195:1–195:14, 2017. 1
- [37] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, volume 37, pages 448–456, 2015. 5
- [38] Aaron S Jackson, Adrian Bulat, Vasileios Argyriou, and Georgios Tzimiropoulos. Large pose 3D face reconstruction from a single image via direct volumetric CNN regression. In *International Conference on Computer Vision (ICCV)*, pages 1031–1039, 2017. 2
- [39] Anil K Jain and Stan Z Li. *Handbook of face recognition*, volume 1. Springer, 2011. 3
- [40] Harim Jung, Myeong-Seok Oh, and Seong-Whan Lee. Learning free-form deformation for 3D face reconstruction from in-the-wild images. In *International Conference on Systems, Man, and Cybernetics (SMC)*, pages 2737–2742, 2021. 2
- [41] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Nießner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *Transactions on Graphics (TOG)*, 37(4):163:1–163:14, 2018. 1
- [42] Hyeonwoo Kim, Michael Zollhöfer, Ayush Tewari, Justus Thies, Christian Richardt, and Christian Theobalt. Inverse-FaceNet: deep monocular inverse face rendering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4625–4634, 2018. 2
- [43] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 5
- [44] Brendan F. Klare, Ben Klein, Emma Taborsky, Austin Blanton, Jordan Cheney, Kristen Allen, Patrick Grother, Alan Mah, Mark James Burge, and Anil K. Jain. Pushing the frontiers of unconstrained face detection and recognition: IARPA janus benchmark A. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1931–1939, 2015. 2
- [45] Tatsuro Koizumi and William A. P. Smith. "look ma, no landmarks!" - unsupervised, model-based dense face alignment. In *European Conference on Computer Vision (ECCV)*, volume 12347, pages 690–706, 2020. 2
- [46] Dimitrios Kollias and Stefanos Zafeiriou. Aff-Wild2: Extending the Aff-Wild database for affect recognition. *CoRR*, abs/1811.07770, 2018. 2, 3
- [47] Jean Kossaifi, Georgios Tzimiropoulos, Sinisa Todorovic, and Maja Pantic. A few-va database for valence and arousal estimation in-the-wild. *Image and Vision Computing*, 65:23–36, 2017. 2, 3, 5
- [48] Jean Kossaifi, Robert Walecki, Yannis Panagakis, Jie Shen, Maximilian Schmitt, Fabien Ringeval, Jing Han, Vedhas Pandit, Antoine Toisoul, Björn W. Schuller, Kam Star, Elmar Hajiyev, and Maja Pantic. SEWA DB: A rich database for audio-visual emotion and sentiment research in the wild. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 43(3):1022–1040, 2021. 2, 3
- [49] Mohammad Rami Koujan, Luma Alharbawee, Giorgos Giannakakis, Nicolas Pugeault, and Anastasios Roussos. Real-time facial expression recognition "in the wild" by disentangling 3d expression from identity. In *International Conference on Automatic Face & Gesture Recognition (FG)*, pages 24–31, 2020. 3
- [50] Shan Li and Weihong Deng. Deep facial expression recognition: A survey. *Transactions on Affective Computing*, 2020. 2, 3
- [51] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6):194:1–194:17, 2017. 2, 3
- [52] Yaojie Liu, Amin Jourabloo, William Ren, and Xiaoming Liu. Dense face alignment. In *International Conference on Computer Vision Workshops (ICCV-W)*, pages 1619–1628, 2017. 2
- [53] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical vision transformer using shifted windows. In

- International Conference on Computer Vision (ICCV)*, pages 9992–10002, 2021. [8](#)
- [54] Patrick Lucey, Jeffrey F. Cohn, Takeo Kanade, Jason Saragih, Zara Ambadar, and Iain Matthews. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *Conference on Computer Vision and Pattern Recognition Workshops (CVPR-W)*, pages 94–101, 2010. [2](#)
- [55] Mohammad Mavadati, Peyton Sanger, and Mohammad H. Mahoor. Extended DISFA dataset: Investigating posed and spontaneous facial expressions. In *Conference on Computer Vision and Pattern Recognition Workshops (CVPR-W)*, pages 1452–1459, 2016. [2](#)
- [56] S. Mohammad Mavadati, Mohammad H. Mahoor, Kevin Bartlett, Philip Trinh, and Jeffrey F. Cohn. Disfa: A spontaneous facial action intensity database. *IEEE Transactions on Affective Computing*, 4(2):151–160, 2013. [2](#)
- [57] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1):18–31, 2017. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [58] Francesca Nonis, Nicole Dagnes, Federica Marcolin, and Enrico Vezzetti. 3D approaches and challenges in facial expression recognition algorithms — a literature review. *Applied Sciences*, 9(18):3904, 2019. [3](#)
- [59] Maja Pantic and Léon J. M. Rothkrantz. Expert system for automatic analysis of facial expressions. *Image and Vision Computing*, 18(11):881–905, 2000. [3](#)
- [60] M. Pantic, Michel Valstar, R. Rademaker, and L. Maat. Web-based database for facial expression analysis. In *International Conference on Multimedia and Expo (ICME)*, pages 317–321, 2005. [2](#)
- [61] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019. [2](#)
- [62] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3D face model for pose and illumination invariant face recognition. In *International Conference on Advanced Video and Signal based Surveillance (AAAI)*, pages 296–301, 2009. [2](#), [3](#)
- [63] Stylianos Ploumpis, Evangelos Ververas, Eimear O’ Sullivan, Stylianos Moschoglou, Haoyang Wang, Nick E. Pears, William A. P. Smith, Baris Gecer, and Stefanos Zafeiriou. Towards a complete 3D morphable model of the human head. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 43(11):4142–4160, 2021. [2](#)
- [64] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. *Annual Conference on Computer Graphics and Interactive Techniques*, pages 497–500, 2001. [3](#)
- [65] Subramanian Ramanathan, Ashraf A. Kassim, Y. V. Venkatesh, and Wu Sin Wah. Human facial expression recognition using a 3D morphable model. In *International Conference on Image Processing (ICIP)*, pages 661–664, 2006. [3](#)
- [66] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3D deep learning with pytorch3d. *arXiv:2007.08501*, 2020. [3](#)
- [67] E. Richardson, M. Sela, and R. Kimmel. 3D face reconstruction by learning from synthetic data. In *International Conference on 3D Vision (3DV)*, pages 460–469, 2016. [2](#)
- [68] Sami Romdhani, Volker Blanz, and Thomas Vetter. Face identification by fitting a 3D morphable model using linear shape and texture error functions. In *European Conference on Computer Vision (ECCV)*, pages 3–19, 2002. [1](#)
- [69] S. Romdhani and T. Vetter. Estimating 3D shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 986–993, 2005. [2](#)
- [70] Zeyu Ruan, Changqing Zou, Longhai Wu, Gangshan Wu, and Limin Wang. SADRNet: Self-aligned dual face regression networks for robust 3D dense face alignment and reconstruction. *Transactions on Image Processing*, 30:5793–5806, 2021. [2](#)
- [71] James A Russell. A circumplex model of affect. *Journal of personality and social psychology*, 39(6):1161, 1980. [3](#)
- [72] Georgia Sandbach, Stefanos Zafeiriou, Maja Pantic, and Lijun Yin. Static and dynamic 3D facial expression recognition: A comprehensive survey. *Image and Vision Computing*, 30(10):683–697, 2012. [3](#)
- [73] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael Black. Learning to regress 3D face shape and expression from an image without 3D supervision. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7763–7772, 2019. [1](#), [2](#), [5](#)
- [74] Matan Sela, Elad Richardson, and Ron Kimmel. Unrestricted facial geometry reconstruction using image-to-image translation. In *International Conference on Computer Vision (ICCV)*, pages 1576–1585, 2017. [2](#)
- [75] Caifeng Shan, Shaogang Gong, and Peter W McOwan. Facial expression recognition based on local binary patterns: A comprehensive study. *Image and vision Computing*, 27(6):803–816, 2009. [3](#)
- [76] Jiaxiang Shang, Tianwei Shen, Shiwei Li, Lei Zhou, Mingmin Zhen, Tian Fang, and Long Quan. Self-supervised monocular 3D face reconstruction by occlusion-aware multi-view geometry consistency. In *European Conference on Computer Vision (ECCV)*, volume 12360, pages 53–70, 2020. [2](#), [6](#), [7](#)
- [77] Yingyan Shi, Qiaosha Zou, and Yiyun Zhang. Pose-robust facial expression recognition by 3D morphable model learning. In *International Conference on Computer and Communications (ICCC)*, pages 2458–2462, 2020. [3](#)
- [78] Ian Sneddon, Margaret McRorie, Gary Mckeown, and Jennifer Hanratty. The belfast induced natural emotion database. *Transactions on Affective Computing*, 3:32–41, 08 2013. [2](#)
- [79] Attila Szabó, Givi Meishvili, and Paolo Favaro. Unsupervised generative 3D shape learning from natural images. *CoRR*, abs/1910.00287, 2019. [2](#)

- [80] Ayush Tewari, Florian Bernard, Pablo Garrido, Gaurav Bharaj, Mohamed Elgharib, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. FML: face model learning from videos. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10812–10822, 2019. 1, 2
- [81] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. StyleRig: Rigging stylegan for 3D control over portrait images. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6141–6150, 2020. 1
- [82] Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeonwoo Kim, Patrick Pérez, and Christian Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 Hz. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2549–2559, 2018. 2
- [83] Ayush Tewari, Michael Zollhöfer, Hyeonwoo Kim, Pablo Garrido, Florian Bernard, Patrick Perez, and Christian Theobalt. MoFA: model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In *International Conference on Computer Vision (ICCV)*, pages 1274–1283, 2017. 2
- [84] Justus Thies, Michael Zollhöfer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2Face: Real-time face capture and reenactment of RGB videos. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2387–2395, 2016. 1, 2
- [85] Ying-li Tian, Takeo Kanade, and Jeffrey F. Cohn. Recognizing action units for facial expression analysis. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 23(2):97–115, 2001. 3
- [86] Antoine Toisoul, Jean Kossaifi, Adrian Bulat, Georgios Tzimiropoulos, and Maja Pantic. Estimation of continuous valence and arousal levels from faces in naturalistic conditions. *Nature Machine Intelligence*, 3(1):42–50, 2021. 2, 4, 5, 6, 7, 8
- [87] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gerard Medioni. Regressing robust and discriminative 3D morphable models with a very deep neural network. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1599–1608, 2017. 2
- [88] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard Medioni. Extreme 3D face reconstruction: Seeing through occlusions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3935–3944, 2018. 2
- [89] Luan Tran, Feng Liu, and Xiaoming Liu. Towards high-fidelity nonlinear 3D face morphable model. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1126–1135, 2019. 2
- [90] Thomas Vetter and Volker Blanz. Estimating coloured 3D face models from single images: An example based approach. In *European Conference on Computer Vision (ECCV)*, pages 499–513, 1998. 2
- [91] Kaisiyuan Wang, Qianyi Wu, Linsen Song, Zhuoqian Yang, Wayne Wu, Chen Qian, Ran He, Yu Qiao, and Chen Change Loy. MEAD: A large-scale audio-visual dataset for emotional talking-face generation. In *European Conference on Computer Vision (ECCV)*, pages 700–717, 2020. 2
- [92] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *International Conference on Computer Vision (ICCV)*, October 2019. 2
- [93] Huawei Wei, Shuang Liang, and Yichen Wei. 3D dense face alignment via graph convolution networks. *arXiv preprint arXiv:1904.05562*, 2019. 2
- [94] Zhen Wen and Thomas S. Huang. Capturing subtle facial motions in 3D face tracking. In *International Conference on Computer Vision (ICCV)*, pages 1343–1350, 2003. 3
- [95] Shangzhe Wu, Christian Rupperecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3D objects from images in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–10, 2020. 2
- [96] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. FaceScape: a large-scale high quality 3D face dataset and detailed riggable 3D face prediction. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 601–610, 2020. 2
- [97] Xiaoxing Zeng, Xiaojiang Peng, and Yu Qiao. DF2Net: A dense-fine-finer network for detailed 3D face reconstruction. In *International Conference on Computer Vision (ICCV)*, 2019. 2
- [98] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z. Li. Face alignment across large poses: A 3D solution. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 146–155, 2016. 2
- [99] Michael Zollhöfer, Justus Thies, Pablo Garrido, Derek Bradley, Thabo Beeler, Patrick Pérez, Marc Stamminger, Matthias Nießner, and Christian Theobalt. State of the art on monocular 3D face reconstruction, tracking, and applications. *Computer Graphics Forum*, 2018. 2