

Learning an Animatable Detailed 3D Face Model from In-The-Wild Images

YAO FENG*, Max Planck Institute for Intelligent Systems and Max Planck ETH Center for Learning System, Germany
HAIWEN FENG*, Max Planck Institute for Intelligent Systems, Germany
MICHAEL J. BLACK, Max Planck Institute for Intelligent Systems, Germany
TIMO BOLKART, Max Planck Institute for Intelligent Systems, Germany

While current monocular 3D face reconstruction methods can recover fine geometric details, they suffer several limitations. Some methods produce faces that cannot be realistically animated because they do not model how wrinkles vary with expression. Other methods are trained on high-quality face scans and do not generalize well to in-the-wild images. We present the first approach that regresses 3D face shape and animatable details that are specific to an individual but change with expression. Our model, DECA (Detailed Expression Capture and Animation), is trained to robustly produce a UV displacement map from a low-dimensional latent representation that consists of person-specific detail parameters and generic expression parameters, while a regressor is trained to predict detail, shape, albedo, expression, pose and illumination parameters from a single image. To enable this, we introduce a novel detail-consistency loss that disentangles person-specific details from expression-dependent wrinkles. This disentanglement allows us to synthesize realistic person-specific wrinkles by controlling expression parameters while keeping person-specific details unchanged. DECA is learned from in-the-wild images with no paired 3D supervision and achieves state-of-the-art shape reconstruction accuracy on two benchmarks. Qualitative results on in-the-wild data demonstrate DECA's robustness and its ability to disentangle identity- and expression-dependent details enabling animation of reconstructed faces. The model and code are publicly available at <https://deca.is.tue.mpg.de>.

CCS Concepts: • **Computing methodologies** → *Mesh models*.

Additional Key Words and Phrases: Detailed face model, 3D face reconstruction, facial animation, detail disentanglement

ACM Reference Format:

Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. 2021. Learning an Animatable Detailed 3D Face Model from In-The-Wild Images. *ACM Trans. Graph.* 40, 4, Article 88 (August 2021), 13 pages. <https://doi.org/10.1145/3450626.3459936>

1 INTRODUCTION

Two decades have passed since the seminal work of Vetter and Blanz [1998] that first showed how to reconstruct 3D facial geometry

*Both authors contributed equally to the paper

Authors' addresses: Yao Feng, Max Planck Institute for Intelligent Systems, Tübingen, Max Planck ETH Center for Learning System, Tübingen, Germany, yfeng@tuebingen.mpg.de; Haiwen Feng, Max Planck Institute for Intelligent Systems, Tübingen, Germany, hifeng@tuebingen.mpg.de; Michael J. Black, Max Planck Institute for Intelligent Systems, Tübingen, Germany, black@tuebingen.mpg.de; Timo Bolkart, Max Planck Institute for Intelligent Systems, Tübingen, Germany, tbolkart@tuebingen.mpg.de.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).
0730-0301/2021/8-ART88

<https://doi.org/10.1145/3450626.3459936>




Fig. 1. DECA. Example images (row 1), the regressed coarse shape (row 2), detail shape (row 3) and reposed coarse shape (row 4), and reposed with person-specific details (row 5) where the source expression is extracted by DECA from the faces in the corresponding colored boxes (row 6). DECA is robust to in-the-wild variations and captures person-specific details as well as expression-dependent wrinkles that appear in regions like the forehead and mouth. Our novelty is that this detailed shape can be reposed (*animated*) such that the wrinkles are specific to the source shape and target expression. Images are taken from Pexels [2021] (row 1; col. 5), Flickr [2021] (bottom left) @ Gage Skidmore, Chicago [Ma et al. 2015] (bottom right), and from NoW [Sanyal et al. 2019] (remaining images).

from a single image. Since then, 3D face reconstruction methods have rapidly advanced (for a comprehensive overview see [Morales et al. 2021; Zollhöfer et al. 2018]) enabling applications such as 3D avatar creation for VR/AR [Hu et al. 2017], video editing [Kim et al. 2018a; Thies et al. 2016], image synthesis [Ghosh et al. 2020; Tewari et al. 2020] face recognition [Blanz et al. 2002; Romdhani et al. 2002], virtual make-up [Scherbaum et al. 2011], or speech-driven facial animation [Cudeiro et al. 2019; Karras et al. 2017; Richard et al. 2021]. To make the problem tractable, most existing methods incorporate prior knowledge about geometry or appearance by

leveraging pre-computed 3D face models [Brunton et al. 2014; Egger et al. 2020]. These models reconstruct the coarse face shape but are unable to capture geometric details such as expression-dependent wrinkles, which are essential for realism and support analysis of human emotion.

Several methods recover detailed facial geometry [Abrevaya et al. 2020; Cao et al. 2015; Chen et al. 2019; Guo et al. 2018; Richardson et al. 2017; Tran et al. 2018, 2019], however, they require high-quality training scans [Cao et al. 2015; Chen et al. 2019] or lack robustness to occlusions [Abrevaya et al. 2020; Guo et al. 2018; Richardson et al. 2017]. None of these explore how the recovered wrinkles change with varying expressions. Previous methods that learn expression-dependent detail models [Bickel et al. 2008; Chaudhuri et al. 2020; Yang et al. 2020] either use detailed 3D scans as training data and, hence, do not generalize to unconstrained images [Yang et al. 2020], or model expression-dependent details as part of the appearance map rather than the geometry [Chaudhuri et al. 2020], preventing realistic mesh relighting.

We introduce DECA (Detailed Expression Capture and Animation), which learns an *animatable* displacement model from in-the-wild images without 2D-to-3D supervision. In contrast to prior work, these *animatable expression-dependent wrinkles are specific to an individual* and are regressed from a single image. Specifically, DECA jointly learns 1) a geometric detail model that generates a UV displacement map from a low-dimensional representation that consists of subject-specific detail parameters and expression parameters, and 2) a regressor that predicts subject-specific detail, albedo, shape, expression, pose, and lighting parameters from an image. The detail model builds upon FLAME’s [Li et al. 2017] coarse geometry, and we formulate the displacements as a function of subject-specific detail parameters and FLAME’s jaw pose and expression parameters.

This enables important applications such as easy avatar creation from a single image. While previous methods can capture detailed geometry in the image, most applications require a face that can be animated. For this, it is not sufficient to recover accurate geometry in the input image. Rather, we must be able to animate that detailed geometry and, more specifically, the details should be person specific.

To gain control over expression-dependent wrinkles of the reconstructed face, while preserving person-specific details (i.e. moles, pores, eyebrows, and expression-independent wrinkles), the person-specific details and expression-dependent wrinkles must be disentangled. Our key contribution is a novel *detail consistency loss* that enforces this disentanglement. During *training*, if we are given two images of the same person with different expressions, we observe that their 3D face shape and their person-specific details are the same in both images, but the expression and the intensity of the wrinkles differ with expression. We exploit this observation during training by swapping the detail codes between different images of the same identity and enforcing the newly rendered results to look similar to the original input images. Once trained, DECA reconstructs a detailed 3D face from a *single image* (Fig. 1, third row) in real time (about 120fps on a Nvidia Quadro RTX 5000), and is able to animate the reconstruction with realistic adaptive expression wrinkles (Fig. 1, fifth row).

In summary, our main contributions are: 1) The first approach to learn an *animatable displacement model* from in-the-wild images that can synthesize plausible geometric details by varying expression parameters. 2) A novel detail consistency loss that disentangles identity-dependent and expression-dependent facial details. 3) Reconstruction of geometric details that is, unlike most competing methods, robust to common occlusions, wide pose variation, and illumination variation. This is enabled by our low-dimensional detail representation, the detail disentanglement, and training from a large dataset of in-the-wild images. 4) State-of-the-art shape reconstruction accuracy on two different benchmarks. 5) The code and model are available for research purposes at <https://deca.is.tue.mpg.de>.

2 RELATED WORK

The reconstruction of 3D faces from visual input has received significant attention over the last decades after the pioneering work of Parke [1974], the first method to reconstruct 3D faces from multi-view images. While a large body of related work aims to reconstruct 3D faces from various input modalities such as multi-view images [Beeler et al. 2010; Cao et al. 2018a; Pighin et al. 1998], video data [Garrido et al. 2016; Ichim et al. 2015; Jeni et al. 2015; Shi et al. 2014; Suwajanakorn et al. 2014], RGB-D data [Li et al. 2013; Thies et al. 2015; Weise et al. 2011] or subject-specific image collections [Kemelmacher-Shlizerman and Seitz 2011; Roth et al. 2016], our main focus is on methods that use only a single RGB image. For a more comprehensive overview, see Zollhöfer et al. [2018].

Coarse reconstruction: Many monocular 3D face reconstruction methods follow Vetter and Blanz [1998] by estimating coefficients of pre-computed statistical models in an analysis-by-synthesis fashion. Such methods can be categorized into optimization-based [Aldrian and Smith 2013; Bas et al. 2017; Blanz et al. 2002; Blanz and Vetter 1999; Gerig et al. 2018; Romdhani and Vetter 2005; Thies et al. 2016], or learning-based methods [Chang et al. 2018; Deng et al. 2019; Genova et al. 2018; Kim et al. 2018b; Ploumpis et al. 2020; Richardson et al. 2016; Sanyal et al. 2019; Tewari et al. 2017; Tran et al. 2017; Tu et al. 2019]. These methods estimate parameters of a statistical face model with a fixed linear shape space, which captures only low-frequency shape information. This results in overly-smooth reconstructions.

Several works are model-free and directly regress 3D faces (i.e. voxels [Jackson et al. 2017] or meshes [Dou et al. 2017; Feng et al. 2018b; Güler et al. 2017; Wei et al. 2019]) and hence can capture more variation than the model-based methods. However, all these methods require explicit 3D supervision, which is provided either by an optimization-based model fitting [Feng et al. 2018b; Güler et al. 2017; Jackson et al. 2017; Wei et al. 2019] or by synthetic data generated by sampling a statistical face model [Dou et al. 2017] and therefore also only capture coarse shape variations.

Instead of capturing high-frequency geometric details, some methods reconstruct coarse facial geometry along with high-fidelity textures [Gecer et al. 2019; Saito et al. 2017; Slossberg et al. 2018; Yamaguchi et al. 2018]. As this “bakes” shading details into the texture, lighting changes do not affect these details, limiting realism and the range of applications. To enable animation and relighting, DECA captures these details as part of the geometry.

Detail reconstruction: Another body of work aims to reconstruct faces with “mid-frequency” details. Common optimization-based methods fit a statistical face model to images to obtain a coarse shape estimate, followed by a shape from shading (SfS) method to reconstruct facial details from monocular images [Jiang et al. 2018; Li et al. 2018; Riviere et al. 2020], or videos [Garrido et al. 2016; Suwajanakorn et al. 2014]. Unlike DECA, these approaches are slow, the results lack robustness to occlusions, and the coarse model fitting step requires facial landmarks, making them error-prone for large viewing angles and occlusions.

Most regression-based approaches [Cao et al. 2015; Chen et al. 2019; Guo et al. 2018; Lattas et al. 2020; Richardson et al. 2017; Tran et al. 2018] follow a similar approach by first reconstructing the parameters of a statistical face model to obtain a coarse shape, followed by a refinement step to capture localized details. Chen et al. [2019] and Cao et al. [2015] compute local wrinkle statistics from high-resolution scans and leverage these to constrain the fine-scale detail reconstruction from images [Chen et al. 2019] or videos [Cao et al. 2015]. Guo et al. [2018] and Richardson et al. [2017] directly regress per-pixel displacement maps. All these methods only reconstruct fine-scale details in non-occluded regions, causing visible artifacts in the presence of occlusions. Tran et al. [2018] gain robustness to occlusions by applying a face segmentation method [Nirkin et al. 2018] to determine occluded regions, and employ an example-based hole filling approach to deal with the occluded regions. Further, model-free methods exist that directly reconstruct detailed meshes [Sela et al. 2017; Zeng et al. 2019] or surface normals that add detail to coarse reconstructions [Abrevaya et al. 2020; Sengupta et al. 2018]. Tran et al. [2019] and Tewari et al. [2019; 2018] jointly learn a statistical face model and reconstruct 3D faces from images. While offering more flexibility than fixed statistical models, these methods capture limited geometric details compared to other detail reconstruction methods. Lattas et al. [2020] use image translation networks to infer the diffuse normals and specular normals, resulting in realistic rendering. Unlike DECA, none of these detail reconstruction methods offer animatable details after reconstruction.

Animatable detail reconstruction: Most relevant to DECA are methods that reconstruct detailed faces while allowing animation of the result. Existing methods [Bickel et al. 2008; Golovinskiy et al. 2006; Ma et al. 2008; Shin et al. 2014; Yang et al. 2020] learn correlations between wrinkles or attributes like age and gender [Golovinskiy et al. 2006], pose [Bickel et al. 2008] or expression [Shin et al. 2014; Yang et al. 2020] from high-quality 3D face meshes [Bickel et al. 2008]. Fyffe et al. [2014] use optical flow correspondence computed from dynamic video frames to animate static high-resolution scans. In contrast, DECA learns an animatable detail model solely from in-the-wild images without paired 3D training data. While FaceScape [Yang et al. 2020] predicts an animatable 3D face from a single image, the method is not robust to occlusions. This is due to a two step reconstruction process: first optimize the coarse shape, then predict a displacement map from the texture map extracted with the coarse reconstruction.

Chaudhuri et al. [2020] learn identity and expression corrective blendshapes with dynamic (expression-dependent) albedo maps [Nagano et al. 2018]. They model geometric details as part of the albedo map, and therefore, the shading of these details does not

adapt with varying lighting. This results in unrealistic renderings. In contrast, DECA models details as geometric displacements, which look natural when re-lit.

In summary, DECA occupies a unique space. It takes a single image as input and produces person-specific details that can be realistically animated. While some methods produce higher-frequency pixel-aligned details, these are not animatable. Still other methods require high-resolution scans for training. We show that these are not necessary and that animatable details can be learned from 2D images without paired 3D ground truth. This is not just convenient, but means that DECA learns to be robust to a wide variety of real-world variation. We want to emphasize that, while elements of DECA are built on well-understood principles (dating back to Vetter and Blanz), our core contribution is new and essential. The key to making DECA work is the detail consistency loss, which has not appeared previously in the literature.

3 PRELIMINARIES

Geometry prior: FLAME [Li et al. 2017] is a statistical 3D head model that combines separate linear identity shape and expression spaces with linear blend skinning (LBS) and pose-dependent corrective blendshapes to articulate the neck, jaw, and eyeballs. Given parameters of facial identity $\beta \in \mathbb{R}^{|\beta|}$, pose $\theta \in \mathbb{R}^{3k+3}$ (with $k = 4$ joints for neck, jaw, and eyeballs), and expression $\psi \in \mathbb{R}^{|\psi|}$, FLAME outputs a mesh with $n = 5023$ vertices. The model is defined as

$$M(\beta, \theta, \psi) = W(T_P(\beta, \theta, \psi), \mathbf{J}(\beta), \theta, \mathcal{W}), \quad (1)$$

with the blend skinning function $W(\mathbf{T}, \mathbf{J}, \theta, \mathcal{W})$ that rotates the vertices in $\mathbf{T} \in \mathbb{R}^{3n}$ around joints $\mathbf{J} \in \mathbb{R}^{3k}$, linearly smoothed by blendweights $\mathcal{W} \in \mathbb{R}^{k \times n}$. The joint locations \mathbf{J} are defined as a function of the identity β . Further,

$$T_P(\beta, \theta, \psi) = \mathbf{T} + B_S(\beta; \mathcal{S}) + B_P(\theta; \mathcal{P}) + B_E(\psi; \mathcal{E}) \quad (2)$$

denotes the mean template \mathbf{T} in “zero pose” with added shape blendshapes $B_S(\beta; \mathcal{S}) : \mathbb{R}^{|\beta|} \rightarrow \mathbb{R}^{3n}$, pose correctives $B_P(\theta; \mathcal{P}) : \mathbb{R}^{3k+3} \rightarrow \mathbb{R}^{3n}$, and expression blendshapes $B_E(\psi; \mathcal{E}) : \mathbb{R}^{|\psi|} \rightarrow \mathbb{R}^{3n}$, with the learned identity, pose, and expression bases (i.e. linear subspaces) \mathcal{S}, \mathcal{P} and \mathcal{E} . See [Li et al. 2017] for details.

Appearance model: FLAME does not have an appearance model, hence we convert the Basel Face Model’s linear albedo subspace [Paysan et al. 2009] into the FLAME UV layout to make it compatible with FLAME. The appearance model outputs a UV albedo map $A(\alpha) \in \mathbb{R}^{d \times d \times 3}$ for albedo parameters $\alpha \in \mathbb{R}^{|\alpha|}$.

Camera model: Photographs in existing in-the-wild face datasets are often taken from a distance. We, therefore, use an orthographic camera model \mathbf{c} to project the 3D mesh into image space. Face vertices are projected into the image as $\mathbf{v} = s\Pi(M_i) + \mathbf{t}$, where $M_i \in \mathbb{R}^3$ is a vertex in M , $\Pi \in \mathbb{R}^{2 \times 3}$ is the orthographic 3D-2D projection matrix, and $s \in \mathbb{R}$ and $\mathbf{t} \in \mathbb{R}^2$ denote isotropic scale and 2D translation, respectively. The parameters s , and \mathbf{t} are summarized as \mathbf{c} .

Illumination model: For face reconstruction, the most frequently-employed illumination model is based on Spherical Harmonics (SH) [Ramamoorthi and Hanrahan 2001]. By assuming that the light source is distant and the face’s surface reflectance is Lambertian,

the shaded face image is computed as:

$$B(\boldsymbol{\alpha}, \mathbf{I}, N_{uv})_{i,j} = A(\boldsymbol{\alpha})_{i,j} \odot \sum_{k=1}^9 \mathbf{I}_k H_k(N_{i,j}), \quad (3)$$

where the albedo, A , surface normals, N , and shaded texture, B , are represented in UV coordinates and where $B_{i,j} \in \mathbb{R}^3$, $A_{i,j} \in \mathbb{R}^3$, and $N_{i,j} \in \mathbb{R}^3$ denote pixel (i, j) in the UV coordinate system. The SH basis and coefficients are defined as $H_k : \mathbb{R}^3 \rightarrow \mathbb{R}$ and $\mathbf{I} = [\mathbf{I}_1^T, \dots, \mathbf{I}_9^T]^T$, with $\mathbf{I}_k \in \mathbb{R}^3$, and \odot denotes the Hadamard product. **Texture rendering:** Given the geometry parameters $(\boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\psi})$, albedo $(\boldsymbol{\alpha})$, lighting (\mathbf{I}) and camera information \mathbf{c} , we can generate the 2D image I_r by rendering as $I_r = \mathcal{R}(M, B, \mathbf{c})$, where \mathcal{R} denotes the rendering function.

FLAME is able to generate the face geometry with various poses, shapes and expressions from a low-dimensional latent space. However, the representational power of the model is limited by the low mesh resolution and therefore mid-frequency details are mostly missing from FLAME's surface. The next section introduces our expression-dependent displacement model that augments FLAME with mid-frequency details, and it demonstrates how to reconstruct this geometry from a single image and animate it.

4 METHOD

DECA learns to regress a parameterized face model with geometric detail solely from in-the-wild training images (Fig. 2 left). Once trained, DECA reconstructs the 3D head with detailed face geometry from a single face image, I . The learned parametrization of the reconstructed details enables us to then animate the detail reconstruction by controlling FLAME's expression and jaw pose parameters (Fig. 2, right). This synthesizes new wrinkles while keeping person-specific details unchanged.

Key idea: The key idea of DECA is grounded in the observation that an individual's face shows different details (i.e. wrinkles), depending on their facial expressions but that other properties of their shape remain unchanged. Consequently, facial details should be separated into static person-specific details and dynamic expression-dependent details such as wrinkles [Li et al. 2009]. However, disentangling static and dynamic facial details is a non-trivial task. Static facial details are different across people, whereas dynamic expression dependent facial details even vary for the same person. Thus, DECA learns an expression-conditioned detail model to infer facial details from both the person-specific detail latent space and the expression space.

The main difficulty in learning a detail displacement model is the lack of training data. Prior work uses specialized camera systems to scan people in a controlled environment to obtain detailed facial geometry. However, this approach is expensive and impractical for capturing large numbers of identities with varying expressions and diversity in ethnicity and age. Therefore we propose an approach to learn detail geometry from in-the-wild images.

4.1 Coarse reconstruction

We first learn a coarse reconstruction (i.e. in FLAME's model space) in an analysis-by-synthesis way: given a 2D image I as input, we encode the image into a latent code, decode this to synthesize a

2D image I_r , and minimize the difference between the synthesized image and the input. As shown in Fig. 2, we train an encoder E_c , which consists of a ResNet50 [He et al. 2016] network followed by a fully connected layer, to regress a low-dimensional latent code. This latent code consists of FLAME parameters $\boldsymbol{\beta}, \boldsymbol{\psi}, \boldsymbol{\theta}$ (i.e. representing the coarse geometry), albedo coefficients $\boldsymbol{\alpha}$, camera \mathbf{c} , and lighting parameters \mathbf{I} . More specifically, the coarse geometry uses the first 100 FLAME shape parameters $(\boldsymbol{\beta})$, 50 expression parameters $(\boldsymbol{\psi})$, and 50 albedo parameters $(\boldsymbol{\alpha})$. In total, E_c predicts a 236 dimensional latent code.

Given a dataset of 2D face images I_i with multiple images per subject, corresponding identity labels c_i , and 68 2D keypoints \mathbf{k}_i per image, the coarse reconstruction branch is trained by minimizing

$$L_{coarse} = L_{lmk} + L_{eye} + L_{pho} + L_{id} + L_{sc} + L_{reg}, \quad (4)$$

with landmark loss L_{lmk} , eye closure loss L_{eye} , photometric loss L_{pho} , identity loss L_{id} , shape consistency loss L_{sc} and regularization L_{reg} .

Landmark re-projection loss: The landmark loss measures the difference between ground-truth 2D face landmarks \mathbf{k}_i and the corresponding landmarks on the FLAME model's surface $M_i \in \mathbb{R}^3$, projected into the image by the estimated camera model. The landmark loss is defined as

$$L_{lmk} = \sum_{i=1}^{68} \|\mathbf{k}_i - s\Pi(M_i) + \mathbf{t}\|_1. \quad (5)$$

Eye closure loss: The eye closure loss computes the relative offset of landmarks \mathbf{k}_i and \mathbf{k}_j on the upper and lower eyelid, and measures the difference to the offset of the corresponding landmarks on FLAME's surface M_i and M_j projected into the image. Formally, the loss is given as

$$L_{eye} = \sum_{(i,j) \in E} \|\mathbf{k}_i - \mathbf{k}_j - s\Pi(M_i - M_j)\|_1, \quad (6)$$

where E is the set of upper/lower eyelid landmark pairs. While the landmark loss, L_{lmk} (Eq. 5), penalizes the absolute landmark location differences, L_{eye} penalizes the relative difference between eyelid landmarks. Because the eye closure loss L_{eye} is translation invariant, it is less susceptible to a misalignment between the projected 3D face and the image, compared to L_{lmk} . In contrast, simply increasing the landmark loss for the eye landmarks affects the overall face shape and can lead to unsatisfactory reconstructions. See Fig. 10 for the effect of the eye-closure loss.

Photometric loss: The photometric loss computes the error between the input image I and the rendering I_r as

$$L_{pho} = \|V_I \odot (I - I_r)\|_{1,1}.$$

Here, V_I is a face mask with value 1 in the face skin region, and value 0 elsewhere obtained by an existing face segmentation method [Nirkin et al. 2018], and \odot denotes the Hadamard product. Computing the error in only the face region provides robustness to common occlusions by e.g. hair, clothes, sunglasses, etc. Without this, the predicted albedo will also consider the color of the occluder, which may be far from skin color, resulting in unnatural rendering (see Fig. 10).

Identity loss: Recent 3D face reconstruction methods demonstrate the effectiveness of utilizing an identity loss to produce more realistic face shapes [Deng et al. 2019; Gecer et al. 2019]. Motivated by




Fig. 2. DECA training and animation. During training (left box), DECA estimates parameters to reconstruct face shape for each image with the aid of the shape consistency information (following the blue arrows) and, then, learns an expression-conditioned displacement model by leveraging detail consistency information (following the red arrows) from multiple images of the same individual (see Sec. 4.3 for details). While the analysis-by-synthesis pipeline is, by now, standard, the yellow box region contains our key novelty. This displacement consistency loss is further illustrated in Fig. 3. Once trained, DECA animates a face (right box) by combining the reconstructed source identity’s shape, head pose, and detail code, with the reconstructed source expression’s jaw pose and expression parameters to obtain an animated coarse shape and an animated displacement map. Finally, DECA outputs an animated detail shape. Images are taken from NoW [Sanyal et al. 2019]. Note that NoW images are not used for training DECA, but are just selected for illustration purposes.




Fig. 3. Detail consistency loss. DECA uses multiple images of the same person during training to disentangle static person-specific details from expression-dependent details. When properly factored, we should be able to take the detail code from one image of a person and use it to reconstruct another image of that person with a different expression. See Sec. 4.3 for details. Images are taken from NoW [Sanyal et al. 2019]. Note that NoW images are not used for training, but are just selected for illustration purposes.

this, we also use a pretrained face recognition network [Cao et al. 2018b], to employ an identity loss during training.

The face recognition network f outputs feature embeddings of the rendered images and the input image, and the identity loss then measures the cosine similarity between the two embeddings. Formally, the loss is defined as

$$L_{id} = 1 - \frac{f(I)f(I_r)}{\|f(I)\|_2 \cdot \|f(I_r)\|_2}. \quad (7)$$

By computing the error between embeddings, the loss encourages the rendered image to capture fundamental properties of a person’s

identity, ensuring that the rendered image looks like the same person as the input subject. Figure 10 shows that the coarse shape results with L_{id} look more like the input subject than those without.

Shape consistency loss: Given two images I_i and I_j of the same subject (i.e. $c_i = c_j$), the coarse encoder E_c should output the same shape parameters (i.e. $\beta_i = \beta_j$). Previous work encourages shape consistency by enforcing the distance between β_i and β_j to be smaller by a margin than the distance to the shape coefficients corresponding to a different subject [Sanyal et al. 2019]. However, choosing this fixed margin is challenging in practice. Instead, we propose a different strategy by replacing β_i with β_j while keeping all other parameters unchanged. Given that β_i and β_j represent the same subject, this new set of parameters must reconstruct I_i well. Formally, we minimize

$$L_{sc} = L_{coarse}(I_i, \mathcal{R}(M(\beta_j, \theta_i, \psi_i), B(\alpha_i, \mathbf{l}_i, N_{uv,i}), c_i)). \quad (8)$$

The goal is to make the rendered images look like the real person. If the method has correctly estimated the shape of the face in two images of the same person, then swapping the shape parameters between these images should produce rendered images that are indistinguishable. Thus, we employ the photometric and identity loss on the rendered images from swapped shape parameters.

Regularization: L_{reg} regularizes shape $E_\beta = \|\beta\|_2^2$, expression $E_\psi = \|\psi\|_2^2$, and albedo $E_\alpha = \|\alpha\|_2^2$.

4.2 Detail reconstruction

The detail reconstruction augments the coarse FLAME geometry with a detailed UV displacement map $D \in [-0.01, 0.01]^{d \times d}$ (see Fig. 2). Similar to the coarse reconstruction, we train an encoder E_d

(with the same architecture as E_c) to encode I to a 128-dimensional latent code δ , representing subject-specific details. The latent code δ is then concatenated with FLAME’s expression ψ and jaw pose parameters θ_{jaw} , and decoded by F_d to D .

Detail decoder: The detail decoder is defined as

$$D = F_d(\delta, \psi, \theta_{jaw}), \quad (9)$$

where the detail code $\delta \in \mathbb{R}^{128}$ controls the static person-specific details. We leverage the expression $\psi \in \mathbb{R}^{50}$ and jaw pose parameters $\theta_{jaw} \in \mathbb{R}^3$ from the coarse reconstruction branch to capture the dynamic expression wrinkle details. For rendering, D is converted to a normal map.

Detail rendering: The detail displacement model allows us to generate images with mid-frequency surface details. To reconstruct the detailed geometry M' , we convert M and its surface normals N to UV space, denoted as $M_{uv} \in \mathbb{R}^{d \times d \times 3}$ and $N_{uv} \in \mathbb{R}^{d \times d \times 3}$, and combine them with D as

$$M'_{uv} = M_{uv} + D \odot N_{uv}. \quad (10)$$

By calculating normals N' from M' , we obtain the detail rendering I'_r by rendering M with the applied normal map as

$$I'_r = \mathcal{R}(M, B(\alpha, \mathbf{1}, N'), \mathbf{c}). \quad (11)$$

The detail reconstruction is trained by minimizing

$$L_{detail} = L_{phoD} + L_{mrf} + L_{sym} + L_{dc} + L_{regD}, \quad (12)$$

with photometric detail loss L_{phoD} , ID-MRF loss L_{mrf} , soft symmetry loss L_{sym} , and detail regularization L_{regD} . Since our estimated albedo is generated by a linear model with 50 basis vectors, the rendered coarse face image only recovers low frequency information such as skin tone and basic facial attributes. High frequency details in the rendered image result mainly from the displacement map, and hence, since L_{detail} compares the rendered detailed image with the real image, F_d is forced to model detailed geometric information.

Detail photometric losses: With the applied detail displacement map, the rendered images I'_r contain some geometric details. Equivalent to the coarse rendering, we use a photometric loss $L_{phoD} = \|\mathbb{V}_I \odot (I - I'_r)\|_{1,1}$, where, recall, \mathbb{V}_I is a mask representing the visible skin pixels.

ID-MRF loss: We adopt an Implicit Diversified Markov Random Field (ID-MRF) loss [Wang et al. 2018] to reconstruct geometric details. Given the input image and the detail rendering, the ID-MRF loss extracts feature patches from different layers of a pre-trained network, and then minimizes the difference between corresponding nearest neighbor feature patches from both images. Larsen et al. [2016] and Isola et al. [2017] point out that L1 losses are not able to recover the high frequency information in the data. Consequently, these two methods use a discriminator to obtain high-frequency detail. Unfortunately, this may result in an unstable adversarial training process. Instead, the ID-MRF loss regularizes the generated content to the original input at the local patch level; this encourages DECA to capture high-frequency details.

Following Wang et al. [2018], the loss is computed on layers $conv3_2$ and $conv4_2$ of VGG19 [Simonyan and Zisserman 2014] as

$$L_{mrf} = 2L_M(conv4_2) + L_M(conv3_2), \quad (13)$$

where $L_M(layer_{th})$ denotes the ID-MRF loss that is employed on the feature patches extracted from I'_r and I with layer $layer_{th}$ of VGG19. As with the photometric losses, we compute L_{mrf} only for the face skin region in UV space.

Soft symmetry loss: To add robustness to self-occlusions, we add a soft symmetry loss to regularize non-visible face parts. Specifically, we minimize

$$L_{sym} = \|V_{uv} \odot (D - flip(D))\|_{1,1}, \quad (14)$$

where V_{uv} denotes the face skin mask in UV space, and $flip$ is the horizontal flip operation. Without L_{sym} , for extreme poses, boundary artifacts become visible in occluded regions (Fig. 9).

Detail regularization: The detail displacements are regularized by $L_{regD} = \|D\|_{1,1}$ to reduce noise.

4.3 Detail disentanglement

Optimizing L_{detail} enables us to reconstruct faces with mid-frequency details. Making these detail reconstructions animatable, however, requires us to disentangle person specific details (i.e. moles, pores, eyebrows, and expression-independent wrinkles) controlled by δ from expression-dependent wrinkles (i.e. wrinkles that change for varying facial expression) controlled by FLAME’s expression and jaw pose parameters, ψ and θ_{jaw} . Our key observation is that the same person in two images should have both similar coarse geometry and personalized details.

Specifically, for the rendered detail image, *exchanging the detail codes between two images of the same subject should have no effect on the rendered image*. This concept is illustrated in Fig. 3. Here we take the the jaw and expression parameters from image i , extract the detail code from image j , and combine these to estimate the wrinkle detail. When we swap detail codes between different images of the same person, the produced results must remain realistic.

Detail consistency loss: Given two images I_i and I_j of the same subject (i.e. $c_i = c_j$), the loss is defined as

$$L_{dc} = L_{detail}(I_i, \mathcal{R}(M(\beta_i, \theta_i, \psi_i), A(\alpha_i), F_d(\delta_j, \psi_i, \theta_{jaw,i}), \mathbf{1}_i, \mathbf{c}_i)), \quad (15)$$

where β_i , θ_i , ψ_i , $\theta_{jaw,i}$, α_i , $\mathbf{1}_i$, and \mathbf{c}_i are the parameters of I_i , while δ_j is the detail code of I_j (see Fig. 3). The detail consistency loss is essential for the disentanglement of identity-dependent and expression-dependent details. Without the detail consistency loss, the person-specific detail code, δ , captures identity and expression dependent details, and therefore, reconstructed details cannot be re-posed by varying the FLAME jaw pose and expression. We show the necessity and effectiveness of L_{dc} in Sec. 6.3.

5 IMPLEMENTATION DETAILS

Data: We train DECA on three publicly available datasets: VGGFace2 [Cao et al. 2018b], BUPT-Balancedface [Wang et al. 2019] and VoxCeleb2 [Chung et al. 2018a]. VGGFace2 [Cao et al. 2018b] contains images of over 8k subjects, with an average of more than 350 images per subject. BUPT-Balancedface [Wang et al. 2019] offers 7k subjects per ethnicity (i.e. Caucasian, Indian, Asian and African), and VoxCeleb2 [Chung et al. 2018a] contains 145k videos of 6k subjects. In total, DECA is trained on 2 Million images.

All datasets provide an identity label for each image. We use FAN [Bulat and Tzimiropoulos 2017] to predict 68 2D landmarks \mathbf{k}_i on each face. To improve the robustness of the predicted landmarks, we run FAN for each image twice with different face crops, and discard all images with non-matching landmarks. See Sup. Mat. for details on data selection and data cleaning.

Implementation details: DECA is implemented in PyTorch [Paszke et al. 2019], using the differentiable rasterizer from Pytorch3D [Ravi et al. 2020] for rendering. We use Adam [Kingma and Ba 2015] as optimizer with a learning rate of $1e-4$. The input image size is 224^2 and the UV space size is $d = 256$. See Sup. Mat. for details.

6 EVALUATION

6.1 Qualitative evaluation

Reconstruction: Given a single face image, DECA reconstructs the 3D face shape with mid-frequency geometric details. The second row of Fig. 1 shows that the coarse shape (i.e. in FLAME space) well represents the overall face shape, and the learned DECA detail model reconstructs subject-specific details and wrinkles of the input identity (Fig. 1, row three), while being robust to partial occlusions.

Figure 5 qualitatively compares DECA results with state-of-the-art coarse face reconstruction methods, namely PRNet [Feng et al. 2018b], RingNet [Sanyal et al. 2019], Deng et al. [2019], FML [Tewari et al. 2019] and 3DDFA-V2 [Guo et al. 2020]. Compared to these methods, DECA better reconstructs the overall face shape with details like the nasolabial fold (rows 1, 2, 3, 4, and 6) and forehead wrinkles (row 3). DECA better reconstructs the mouth shape and the eye region than all other methods. DECA further reconstructs a full head while PRNet [Feng et al. 2018b], Deng et al. [2019], FML [Tewari et al. 2019] and 3DDFA-V2 [Guo et al. 2020] reconstruct tightly cropped faces. While RingNet [Sanyal et al. 2019], like DECA, is based on FLAME [Li et al. 2017], DECA better reconstructs the face shape and the facial expression.

Figure 6 compares DECA visually to existing detailed face reconstruction methods, namely Extreme3D [Tran et al. 2018], Cross-modal [Abrevaya et al. 2020], and FaceScape [Yang et al. 2020]. Extreme3D [Tran et al. 2018] and Cross-modal [Abrevaya et al. 2020] reconstruct more details than DECA but at the cost of being less robust to occlusions (rows 1, 2, 3). Unlike DECA, Extreme3D and Cross-modal only reconstruct static details. However, using static details instead of DECA’s animatable details leads to visible artifacts when animating the face (see Fig. 4). While FaceScape [Yang et al. 2020] provides animatable details, unlike DECA, the method is trained on high-resolution scans while DECA is solely trained on in-the-wild images. Also, with occlusion, FaceScape produces artifacts (rows 1, 2) or effectively fails (row 3).

In summary, DECA produces high-quality reconstructions, outperforming previous work in terms of robustness, while enabling animation of the detailed reconstruction. To demonstrate the quality of DECA and the robustness to variations in head pose, expression, occlusions, image resolution, lighting conditions, etc., we show results for 200 randomly selected ALFW2000 [Zhu et al. 2015] images in the Sup. Mat. along with more qualitative coarse and detail reconstruction comparisons to the state-of-the-art.

Detail animation: DECA models detail displacements as a function of subject-specific detail parameters δ and FLAME’s jaw pose θ_{jaw} and expression parameters ψ as illustrated in Fig. 2 (right). This formulation allows us to animate detailed facial geometry such that wrinkles are specific to the source shape and expression as shown in Fig. 1. Using static details instead of DECA’s animatable details (i.e. by using the reconstructed details as a static displacement map) and animating only the coarse shape by changing the FLAME parameters results in visible artifacts as shown in Fig. 4 (top), while animatable details (middle) look similar to the reference shape (bottom) of the same identity. Figure 7 shows more examples where using static details results in artifacts at the mouth corner or the forehead region, while DECA’s animated results look plausible.

6.2 Quantitative evaluation

We compare DECA with publicly available methods, namely 3DDFA-V2 [Guo et al. 2020], Deng et al. [2019], RingNet [Sanyal et al. 2019], PRNet [Feng et al. 2018b], 3DMM-CNN [Tran et al. 2017] and Extreme3D [Tran et al. 2018]. Note that there is no benchmark face dataset with ground truth shape detail. Consequently, our quantitative analysis focuses on the accuracy of the coarse shape. Note that DECA achieves SOTA performance on 3D reconstruction without any paired 3D data in training.

NoW benchmark: The NoW challenge [Sanyal et al. 2019] consists of 2054 face images of 100 subjects, split into a validation set (20 subjects) and a test set (80 subjects), with a reference 3D face scan per subject. The images consist of indoor and outdoor images, neutral expression and expressive face images, partially occluded faces, and varying viewing angles ranging from frontal view to profile view, and selfie images. The challenge provides a standard evaluation protocol that measures the distance from all reference scan vertices to the closest point in the reconstructed mesh surface, after rigidly aligning scans and reconstructions. For details, see [NoW challenge 2019].

We found that the tightly cropped face meshes predicted by Deng et al. [2019] are smaller than the NoW reference scans, which would result in a high reconstruction error in the missing region. For a fair comparison to the method of Deng et al. [2019], we use the Basel Face Model (BFM) [Paysan et al. 2009] parameters they output, reconstruct the complete BFM mesh, and get the NoW evaluation for these complete meshes. As shown in Tab. 1 and the cumulative error plot in Figure 8 (left), DECA gives state-of-the-art results on NoW, providing the reconstruction error with the lowest mean, median, and standard deviation.

To quantify the influence of the geometric details, we separately evaluate the coarse and the detail shape (i.e. w/o and w/ details) on the NoW *validation set*. The reconstruction errors are, median: 1.18/1.19 (coarse / detailed), mean: 1.46/1.47 (coarse / detailed), std: 1.25/1.25 (coarse / detailed). This indicates that while the detail shape improves visual quality when compared to the coarse shape, the quantitative performance is slightly worse.

To test for gender bias in the results, we report errors separately for female (f) and male (m) NoW test subjects. We find that recovered female shapes are slightly more accurate. Reconstruction errors are, median: 1.03/1.16 (f/m), mean: 1.32/1.45 (f/m), and std:

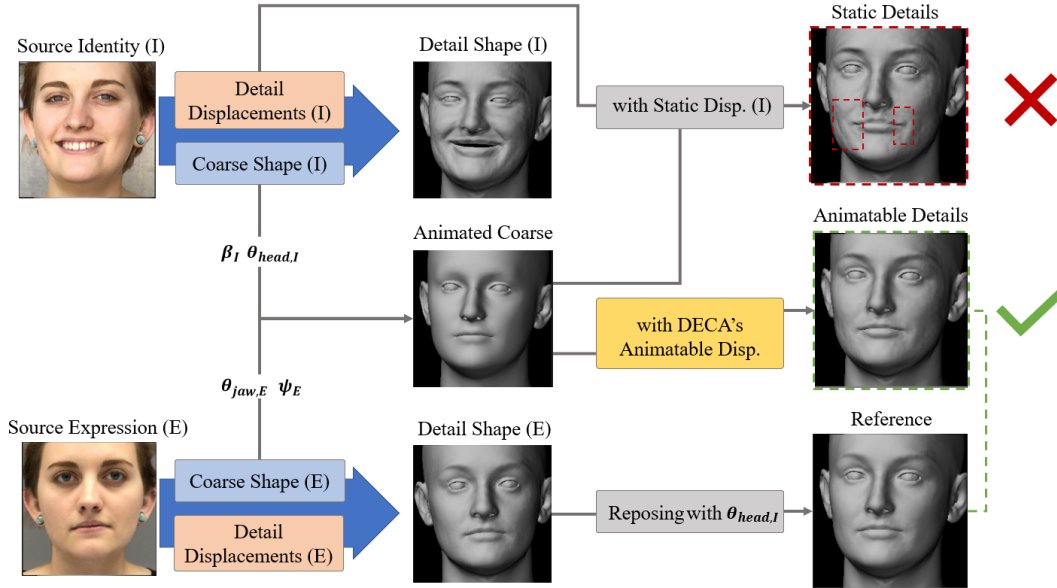


Fig. 4. Effect of DECA’s animatable details. Given images of source identity I and source expression E (left), DECA reconstructs the detail shapes (middle) and animates the detail shape of I with the expression of E (right, middle). This synthesized DECA expression appears nearly identical to the reconstructed same subject’s reference detail shape (right, bottom). Using the reconstructed details of I instead (i.e. static details) and animating the coarse shape only, results in visible artifacts (right, top). See Sec. 6.1 for details. Input images are taken from NoW [Sanyal et al. 2019].




Fig. 5. Comparison to other **coarse reconstruction** methods, from left to right: PRNet [Feng et al. 2018b], RingNet [Sanyal et al. 2019] Deng et al. [2019], FML [Tewari et al. 2019], 3DDFA-V2 [Guo et al. 2020], DECA (ours). Input images are taken from VoxCeleb2 [Chung et al. 2018b].

1.16/1.20 (f/m). The cumulative error plots in Fig. 1 of the Sup. Mat. demonstrate that DECA gives state-of-the-art performance for both genders.

Feng et al. benchmark: The Feng et al. [2018a] challenge contains 2000 face images of 135 subjects, and a reference 3D face scan for each subject. The benchmark consists of 1344 low-quality (LQ)

Method	Median (mm)	Mean (mm)	Std (mm)
3DMM-CNN [Tran et al. 2017]	1.84	2.33	2.05
PRNet [Feng et al. 2018b]	1.50	1.98	1.88
Deng et al.19 [2019]	1.23	1.54	1.29
RingNet [Sanyal et al. 2019]	1.21	1.54	1.31
3DDFA-V2 [Guo et al. 2020]	1.23	1.57	1.39
MGCNet [Shang et al. 2020]	1.31	1.87	2.63
DECA (ours)	1.09	1.38	1.18

Table 1. Reconstruction error on the NoW [Sanyal et al. 2019] benchmark.

Method	Median (mm)		Mean (mm)		Std (mm)	
	LQ	HQ	LQ	HQ	LQ	HQ
3DMM-CNN [Tran et al. 2017]	1.88	1.85	2.32	2.29	1.89	1.88
Extreme3D [Tran et al. 2018]	2.40	2.37	3.49	3.58	6.15	6.75
PRNet [Feng et al. 2018b]	1.79	1.59	2.38	2.06	2.19	1.79
RingNet [Sanyal et al. 2019]	1.63	1.59	2.08	2.02	1.79	1.69
3DDFA-V2 [Guo et al. 2020]	1.62	1.49	2.10	1.91	1.87	1.64
DECA (ours)	1.48	1.45	1.91	1.89	1.66	1.68

Table 2. Feng et al. [2018a] benchmark performance.

images extracted from videos, and 656 high-quality (HQ) images taken in controlled scenarios. A protocol similar to NoW is used for evaluation, which measures the distance between all reference scan vertices to the closest points on the reconstructed mesh surface, after rigidly aligning scan and reconstruction. As shown in Tab. 2 and the cumulative error plot in Fig. 8 (middle & right), DECA provides state-of-the-art performance.

6.3 Ablation experiment

Detail consistency loss: To evaluate the importance of our novel detail consistency loss L_{dc} (Eq. 15), we train DECA with and without L_{dc} . Figure 9 (left) shows the DECA details for detail code δ_I from



Fig. 6. Comparison to other detailed face reconstruction methods, from left to right: Extreme3D [Tran et al. 2018], FaceScape [Yang et al. 2020], Cross-modal [Abrevaya et al. 2020], DECA (ours). See Sup. Mat. for many more examples. Input images are taken from AFLW2000 [Zhu et al. 2015] (rows 1-3) and VGGFace2 [Cao et al. 2018b] (rows 4-6).



Fig. 7. Effect of DECA's animatable details. Given a single image (left), DECA reconstructs a coarse mesh (second column) and a detailed mesh (third column). Using static details and animating (i.e. reposing) the coarse FLAME shape only (fourth column) results in visible artifacts as highlighted by the red boxes. Instead, reposing with DECA's animatable details (right) results in a more realistic mesh with geometric details. The reposing uses the source expression shown in Fig. 1 (bottom). Input images are taken from NoW [Sanyal et al. 2019] (top), and Pexels [2021] (bottom).

the source identity, and expression ψ_E and jaw pose parameters $\theta_{jaw,E}$ from the source expression. For DECA trained with L_{dc} (top), wrinkles appear in the forehead as a result of the raised eyebrows of the source expression, while for DECA trained without L_{dc} (bottom),

no such wrinkles appear. This indicates that without L_{dc} , person-specific details and expression-dependent wrinkles are not well disentangled. See Sup. Mat. for more disentanglement results.

ID-MRF loss: Figure 9 (right) shows the effect of L_{mrf} on the detail reconstruction. Without L_{mrf} (middle), wrinkle details (e.g. in the forehead) are not reconstructed, resulting in an overly smooth result. With L_{mrf} (right), DECA captures the details.

Other losses: We also evaluate the effect of the eye-closure loss L_{eye} , segmentation on the photometric loss, and the identity loss L_{id} . Fig. 10 provides a qualitative comparison of the DECA coarse model with/without using these losses. Quantitatively, we also evaluate DECA with and without L_{id} on the NoW validation set; the former gives a mean error of 1.46mm, while the latter is worse with an error of 1.59mm.

7 LIMITATIONS AND FUTURE WORK

While DECA achieves SOTA results for reconstructed face shape and provides novel animatable details, there are several limitations. First, the rendering quality for DECA detailed meshes is mainly limited by the albedo model, which is derived from BFM. DECA requires an albedo space without baked in shading, specularities, and shadows in order to disentangle facial albedo from geometric details. Future work should focus on learning a high-quality albedo model with a sufficiently large variety of skin colors, texture details, and no illumination effects. Second, existing methods, like DECA, do not explicitly model facial hair. This pushes skin tone into the lighting model and causes facial hair to be explained by shape deformations. A different approach is needed to properly model this. Third, while robust, our method can still fail due to extreme head pose and lighting. While we are tolerant to common oclusions in existing face datasets (Fig. 6 and examples in Sup. Mat.), we do not address extreme occlusion, e.g. where the hand covers large portions of the face. This suggests the need for more diverse training data.

Further, the training set contains many low-res images, which help with robustness but can introduce noisy details. Existing high-res datasets (e.g. [Karras et al. 2018, 2019]) are less varied, thus training DECA from these datasets results in a model that is less robust to general in-the-wild images, but captures more detail. Additionally, the limited size of high-resolution datasets makes it difficult to disentangle expression- and identity-dependent details. To further research on this topic, we also release a model trained using high-resolution images only (i.e. DECA-HR). Using DECA-HR increases the visual quality and reduces noise in the reconstructed details at the cost of being less robust (i.e. to low image resolutions, extreme head poses, extreme expressions, etc.).

DECA uses a weak perspective camera model. To use DECA to recover head geometry from “selfies”, we would need to extend the method to include the focal length. For some applications, the focal length may be directly available from the camera. However, inferring 3D geometry and focal length from a single image under perspective projection for in-the-wild images is unsolved and likely requires explicit supervision during training (cf. [Zhao et al. 2019]).




Fig. 8. Quantitative comparison to state-of-the-art on two 3D face reconstruction benchmarks, namely the NoW [Sanyal et al. 2019] challenge (left) and the Feng et al. [2018a] benchmark for low-quality (middle) and high-quality (right) images.

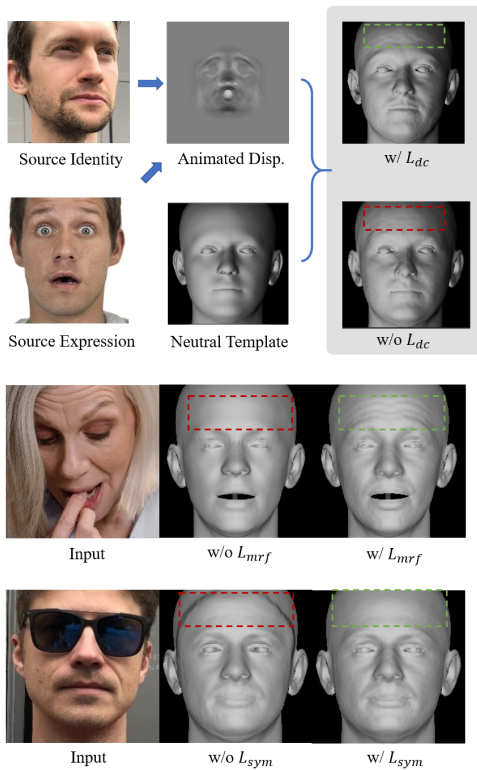


Fig. 9. Ablation experiments. Top: Effects of L_{dc} on the animation of the source identity with the source expression visualized on a neutral expression template mesh. Without L_{dc} , no wrinkles appear in the forehead despite the “surprise” source expression. Middle: Effect of L_{mrf} on the detail reconstruction. Without L_{mrf} , fewer details are reconstructed. Bottom: Effect of L_{sym} on the reconstructed details. Without L_{sym} , boundary artifacts become visible. Input images are taken from NoW [Sanyal et al. 2019] (rows 1 & 4), Chicago [Ma et al. 2015] (row 2), and Pexels [2021] (row 3).

Finally, in future work, we want to extend the model over time, both for tracking and to learn more personalized models of individuals from video where we could enforce continuity of intrinsic wrinkles over time.




Fig. 10. More ablation experiments. Left: estimated landmarks and reconstructed coarse shape from DECA (first column) and DECA without L_{eye} (second column), and without L_{id} (third column). When trained without L_{eye} , DECA is not able to capture closed-eye expressions. Using L_{id} helps reconstruct coarse shape. Right: rendered image from DECA and DECA without segmentation. Without using the skin mask in the photometric loss, the estimated result bakes in the color of the occluder (e.g. sunglasses, hats) into the albedo. Input images are taken from NoW [Sanyal et al. 2019].

8 CONCLUSION

We have presented DECA, which enables detailed expression capture and animation from single images by learning an animatable detail model from a dataset of in-the-wild images. In total, DECA is trained from about 2M in-the-wild face images without 2D-to-3D supervision. DECA reaches state-of-the-art shape reconstruction performance enabled by a shape consistency loss. A novel detail consistency loss helps DECA to disentangle expression-dependent wrinkles from person-specific details. The low-dimensional detail latent space makes the fine-scale reconstruction robust to noise and occlusions, and the novel loss leads to disentanglement of identity and expression-dependent wrinkle details. This enables applications like animation, shape change, wrinkle transfer, etc. DECA is publicly available for research purposes. Due to the reconstruction accuracy, the reliability, and the speed, DECA is useful for applications like face reenactment or virtual avatar creation.

ACKNOWLEDGMENTS

We thank S. Sanyal for providing us the RingNet PyTorch implementation, support with paper writing, and fruitful discussions, M.

Kocabas, N. Athanasiou, V. Fernández Abrevaya, and R. Danecek for the helpful suggestions, and T. McConnell and S. Sorce for the video voice over. This work was partially supported by the Max Planck ETH Center for Learning Systems.

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

REFERENCES

- Victoria Fernández Abrevaya, Adnane Boukhayma, Philip HS Torr, and Edmond Boyer. 2020. Cross-modal Deep Face Normals with Deactivable Skip Connections. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4979–4989.
- Oswald Aldrian and William AP Smith. 2013. Inverse Rendering of Faces with a 3D Morphable Model. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 35, 5 (2013), 1080–1093.
- Anil Bas, William A. P. Smith, Timo Bolkart, and Stefanie Wuhrer. 2017. Fitting a 3D Morphable Model to Edges: A Comparison Between Hard and Soft Correspondences. In *Asian Conference on Computer Vision Workshops*. 377–391.
- Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. 2010. High-quality single-shot capture of facial geometry. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 40.
- Bernd Bickel, Manuel Lang, Mario Botsch, Miguel A. Otaduy, and Markus H. Gross. 2008. Pose-Space Animation and Transfer of Facial Details. In *Eurographics/SIGGRAPH Symposium on Computer Animation (SCA)*, Markus H. Gross and Doug L. James (Eds.). 57–66.
- Volker Blanz, Sami Romdhani, and Thomas Vetter. 2002. Face identification across different poses and illuminations with a 3D morphable model. In *International Conference on Automatic Face & Gesture Recognition (FG)*. 202–207.
- Volker Blanz and Thomas Vetter. 1999. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*. 187–194.
- Alan Brunton, Augusto Salazar, Timo Bolkart, and Stefanie Wuhrer. 2014. Review of statistical shape spaces for 3D data with comparative analysis for human faces. *Computer Vision and Image Understanding (CVIU)* 128, 0 (2014), 1–17.
- Adrian Bulat and Georgios Tzimiropoulos. 2017. How far are we from solving the 2D & 3D face alignment problem? (and a dataset of 230,000 3D facial landmarks). In *IEEE International Conference on Computer Vision (ICCV)*. 1021–1030.
- Chen Cao, Derek Bradley, Kun Zhou, and Thabo Beeler. 2015. Real-time high-fidelity facial performance capture. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–9.
- Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. 2018b. VG-Face2: A dataset for recognising faces across pose and age. In *International Conference on Automatic Face & Gesture Recognition (FG)*. 67–74.
- Xuan Cao, Zhang Chen, Anpei Chen, Xin Chen, Shiyang Li, and Jingyi Yu. 2018a. Sparse Photometric 3D Face Reconstruction Guided by Morphable Models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4635–4644.
- Feng-Ju Chang, Anh Tuan Tran, Tal Hassner, Iacopo Masi, Ram Nevatia, and Gerard Medioni. 2018. ExpNet: Landmark-free, deep, 3D facial expressions. In *International Conference on Automatic Face & Gesture Recognition (FG)*. 122–129.
- Bindita Chaudhuri, Noranart Vesdapunt, Linda G. Shapiro, and Baoyuan Wang. 2020. Personalized Face Modeling for Improved Face Reconstruction and Motion Retargeting. In *European Conference on Computer Vision (ECCV)*. 142–160.
- Anpei Chen, Zhang Chen, Guli Zhang, Kenny Mitchell, and Jingyi Yu. 2019. Photo-Realistic Facial Details Synthesis from Single Image. In *IEEE International Conference on Computer Vision (ICCV)*. 9429–9439.
- J. S. Chung, A. Nagrani, and A. Zisserman. 2018a. VoxCeleb2: Deep Speaker Recognition. In *INTERSPEECH*.
- Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. 2018b. VoxCeleb2: Deep Speaker Recognition. In *Annual Conference of the International Speech Communication Association (Interspeech)*, B. Yegnanarayana (Ed.). ISCA, 1086–1090.
- Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael Black. 2019. Capture, Learning, and Synthesis of 3D Speaking Styles. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 10101–10111.
- Yu Deng, Jialong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. 2019. Accurate 3D Face Reconstruction with Weakly-Supervised Learning: From Single Image to Image Set. In *Computer Vision and Pattern Recognition Workshops*. 285–295.
- Pengfei Dou, Shishir K Shah, and Ioannis A Kakadiaris. 2017. End-to-end 3D face reconstruction with deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5908–5917.
- 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. 2020. 3D Morphable Face Models - Past, Present, and Future. *ACM Transactions on Graphics (TOG)* 39, 5 (2020), 157:1–157:38.
- Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. 2018b. Joint 3D Face Reconstruction and Dense Alignment with Position Map Regression Network. In *European Conference on Computer Vision (ECCV)*. 534–551.
- Zhen-Hua Feng, Patrik Huber, Josef Kittler, Peter Hancock, Xiao-Jun Wu, Qijun Zhao, Paul Koppen, and Matthias Rätzsch. 2018a. Evaluation of dense 3D reconstruction from 2D face images in the wild. In *International Conference on Automatic Face & Gesture Recognition (FG)*.
- Flickr image. 2021. <https://www.flickr.com/photos/gageskidmore/14602415448/>.
- Graham Fyfe, Andrew Jones, Oleg Alexander, Ryosuke Ichikari, and Paul E. Debevec. 2014. Driving High-Resolution Facial Scans with Video Performance Capture. *ACM Transactions on Graphics (TOG)* 34, 1 (2014), 8:1–8:14.
- Pablo Garrido, Michael Zollhöfer, Dan Casas, Levi Valgaerts, Kiran Varanasi, Patrick Pérez, and Christian Theobalt. 2016. Reconstruction of personalized 3D face rigs from monocular video. *ACM Transactions on Graphics (TOG)* 35, 3 (2016), 28.
- Baris Geceer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. 2019. GANFIT: Generative Adversarial Network Fitting for High Fidelity 3D Face Reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1155–1164.
- Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T. Freeman. 2018. Unsupervised Training for 3D Morphable Model Regression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 8377–8386.
- Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Luthi, Sandro Schönborn, and Thomas Vetter. 2018. Morphable face models—an open framework. In *International Conference on Automatic Face & Gesture Recognition (FG)*. 75–82.
- Partha Ghosh, Pravir Singh Gupta, Roy Uziel, Anurag Ranjan, Michael J. Black, and Timo Bolkart. 2020. GIF: Generative Interpretable Faces. In *International Conference on 3D Vision (3DV)*. 868–878.
- Aleksey Golovinskiy, Wojciech Matusik, Hanspeter Pfister, Szymon Rusinkiewicz, and Thomas A. Funkhouser. 2006. A statistical model for synthesis of detailed facial geometry. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 1025–1034.
- Riza Alp Güler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, and Iasonas Kokkinos. 2017. DenseReg: Fully convolutional dense shape regression in-the-wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 6799–6808.
- Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. 2020. Towards Fast, Accurate and Stable 3D Dense Face Alignment. In *European Conference on Computer Vision (ECCV)*. 152–168.
- Yudong Guo, Jianfei Cai, Boyi Jiang, Jianmin Zheng, et al. 2018. CNN-based real-time dense face reconstruction with inverse-rendered photo-realistic face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 41, 6 (2018), 1294–1307.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778.
- Liwen Hu, Shunsuke Saito, Lingyu Wei, Koki Nagano, Jaewoo Seo, Jens Fursund, Iman Sadeghi, Carrie Sun, Yen-Chun Chen, and Hao Li. 2017. Avatar Digitization from a Single Image for Real-time Rendering. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 195:1–195:14.
- Alexandru Eugen Ichim, Sofien Bouaziz, and Mark Pauly. 2015. Dynamic 3D avatar creation from hand-held video input. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 45.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5967–5976.
- Aaron S Jackson, Adrian Bulat, Vasileios Argyriou, and Georgios Tzimiropoulos. 2017. Large pose 3D face reconstruction from a single image via direct volumetric CNN regression. In *IEEE International Conference on Computer Vision (ICCV)*. 1031–1039.
- László A Jeni, Jeffrey F Cohn, and Takeo Kanade. 2015. Dense 3D face alignment from 2D videos in real-time. In *International Conference on Automatic Face & Gesture Recognition (FG)*, Vol. 1. 1–8.
- Luo Jiang, Juyong Zhang, Bailin Deng, Hao Li, and Ligang Liu. 2018. 3D face reconstruction with geometry details from a single image. *Transactions on Image Processing* 27, 10 (2018), 4756–4770.
- Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen. 2017. Audio-driven facial animation by joint end-to-end learning of pose and emotion. *ACM Transactions on Graphics, (Proc. SIGGRAPH)* 36, 4 (2017), 94:1–94:12.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations (ICLR)*.
- Tero Karras, Samuli Laine, and Timo Aila. 2019. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4401–4410.

- Ira Kemelmacher-Shlizerman and Steven M Seitz. 2011. Face reconstruction in the wild. In *IEEE International Conference on Computer Vision (ICCV)*. 1746–1753.
- Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Nießner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. 2018a. Deep video portraits. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 163:1–163:14.
- Hyeonwoo Kim, Michael Zollhöfer, Ayush Tewari, Justus Thies, Christian Richardt, and Christian Theobalt. 2018b. InverseFaceNet: Deep Monocular Inverse Face Rendering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4625–4634.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*.
- Martin Köstinger, Paul Wohlhart, Peter M. Roth, and Horst Bischof. 2011. Annotated Facial Landmarks in the Wild: A large-scale, real-world database for facial landmark localization. In *IEEE International Conference on Computer Vision Workshops (ICCV-W)*. 2144–2151.
- Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. 2016. Autoencoding beyond pixels using a learned similarity metric. In *International Conference on Machine Learning (ICML)*, Vol. 48. 1558–1566.
- Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Triantafyllou, Abhijeet Ghosh, and Stefanos Zafeiriou. 2020. AvatarMe: Realistically Renderable 3D Facial Reconstruction "In-the-Wild". In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 757–766.
- Hao Li, Bart Adams, Leonidas J. Guibas, and Mark Pauly. 2009. Robust single-view geometry and motion reconstruction. *ACM Transactions on Graphics (TOG)* 28 (2009), 175.
- Hao Li, Jihun Yu, Yuting Ye, and Chris Bregler. 2013. Realtime facial animation with on-the-fly correctives. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 42–1.
- Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* 36, 6 (2017), 194:1–194:17.
- Yue Li, Liqian Ma, Haoqiang Fan, and Kenny Mitchell. 2018. Feature-preserving detailed 3D face reconstruction from a single image. In *European Conference on Visual Media Production*. 1–9.
- Debbie S. Ma, Joshua Correll, and Bernd Wittenbrink. 2015. The Chicago face database: A free stimulus set of faces and norming datan. *Behavior Research Methods volume* 47 (2015), 1122–1135.
- Wan-Chun Ma, Andrew Jones, Jen-Yuan Chiang, Tim Hawkins, Sune Frederiksen, Pieter Peers, Marko Vukovic, Ming Ouhyoung, and Paul E. Debevec. 2008. Facial performance synthesis using deformation-driven polynomial displacement maps. *ACM Transactions on Graphics (TOG)* 27, 5 (2008), 121.
- Araceli Morales, Gemma Piella, and Federico M Sukno. 2021. Survey on 3D face reconstruction from uncalibrated images. *Computer Science Review* 40 (2021), 100400.
- Koki Nagano, Jaewoo Seo, Jun Xing, Lingyu Wei, Zimo Li, Shunsuke Saito, Aviral Agarwal, Jens Fursund, and Hao Li. 2018. paGAN: real-time avatars using dynamic textures. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 258:1–258:12.
- Yuval Nirkin, Iacopo Masi, Anh Tran Tuan, Tal Hassner, and Gerard Medioni. 2018. On face segmentation, face swapping, and face perception. In *International Conference on Automatic Face & Gesture Recognition (FG)*. 98–105.
- NoW challenge. 2019. <https://ringnet.is.tue.mpg.de/challenge>.
- Frederick Ira Parke. 1974. *A parametric model for human faces*. Technical Report. University of Utah.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. 2009. A 3D face model for pose and illumination invariant face recognition. In *International Conference on Advanced Video and Signal Based Surveillance*. 296–301.
- Pexels. 2021. <https://www.pexels.com>.
- Frédéric Pighin, Jamie Hecker, Dani Lischinski, Richard Szeliski, and David H. Salesin. 1998. Synthesizing Realistic Facial Expressions from Photographs. In *SIGGRAPH*. 75–84.
- Stylianos Ploumpis, Evangelos Ververas, Eimear O’Sullivan, Stylianos Moschoglou, Haoyang Wang, Nick Pears, William Smith, Baris Gecer, and Stefanos P Zafeiriou. 2020. Towards a complete 3D morphable model of the human head. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* (2020).
- R. Ramamoorthi and P. Hanrahan. 2001. An efficient representation for irradiance environment maps. *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001).
- Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. 2020. PyTorch3D. <https://github.com/facebookresearch/pytorch3d>.
- Alexander Richard, Michael Zollhöfer, Yandong Wen, Fernando De la Torre, and Yaser Sheikh. 2021. MeshTalk: 3D Face Animation from Speech using Cross-Modality Disentanglement. *CoRR* abs/2104.08223 (2021).
- E. Richardson, M. Sela, and R. Kimmel. 2016. 3D Face Reconstruction by Learning from Synthetic Data. In *International Conference on 3D Vision (3DV)*. 460–469.
- Elad Richardson, Matan Sela, Roy Or-El, and Ron Kimmel. 2017. Learning Detailed Face Reconstruction From a Single Image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1259–1268.
- Jérémy Riviere, Paulo F. U. Gotardo, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. 2020. Single-shot high-quality facial geometry and skin appearance capture. *ACM Transactions on Graphics, (Proc. SIGGRAPH)* 39, 4 (2020), 81.
- Sami Romdhani, Volker Blanz, and Thomas Vetter. 2002. Face identification by fitting a 3D morphable model using linear shape and texture error functions. In *European Conference on Computer Vision (ECCV)*. 3–19.
- Sami Romdhani and Thomas Vetter. 2005. Estimating 3D shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 2. 986–993.
- Joseph Roth, Yiyang Tong, and Xiaoming Liu. 2016. Adaptive 3D face reconstruction from unconstrained photo collections. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4197–4206.
- Shunsuke Saito, Lingyu Wei, Liwen Hu, Koki Nagano, and Hao Li. 2017. Photorealistic Facial Texture Inference Using Deep Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5144–5153.
- Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael Black. 2019. Learning to Regress 3D Face Shape and Expression from an Image without 3D Supervision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 7763–7772.
- Kristina Scherbaum, Tobias Ritschel, Matthias Hullin, Thorsten Thormählen, Volker Blanz, and Hans-Peter Seidel. 2011. Computer-suggested facial makeup. *Computer Graphics Forum* 30, 2 (2011), 485–492.
- Matan Sela, Elad Richardson, and Ron Kimmel. 2017. Unrestricted facial geometry reconstruction using image-to-image translation. In *IEEE International Conference on Computer Vision (ICCV)*. 1576–1585.
- Soumyadip Sengupta, Angjoo Kanazawa, Carlos D. Castillo, and David W. Jacobs. 2018. S3Net: Learning Shape, Reflectance and Illuminance of Faces in the Wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 6296–6305.
- Jiaxiang Shang, Tianwei Shen, Shiwei Li, Lei Zhou, Mingmin Zhen, Tian Fang, and Long Quan. 2020. Self-Supervised Monocular 3D Face Reconstruction by Occlusion-Aware Multi-view Geometry Consistency. In *European Conference on Computer Vision (ECCV)*, Vol. 12360. 53–70.
- Fuhao Shi, Hsiang-Tao Wu, Xin Tong, and Jinxiang Chai. 2014. Automatic acquisition of high-fidelity facial performances using monocular videos. *ACM Transactions on Graphics (TOG)* 33, 6 (2014), 222.
- Il-Kyu Shin, A Cengiz Öztireli, Hyeon-Joong Kim, Thabo Beeler, Markus Gross, and Soo-Mi Choi. 2014. Extraction and transfer of facial expression wrinkles for facial performance enhancement. In *Pacific Conference on Computer Graphics and Applications*. 113–118.
- Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556 (2014).
- Ron Slossberg, Gil Shamai, and Ron Kimmel. 2018. High quality facial surface and texture synthesis via generative adversarial networks. In *European Conference on Computer Vision Workshops (ECCV-W)*.
- Supasorn Suwajanakorn, Ira Kemelmacher-Shlizerman, and Steven M Seitz. 2014. Total moving face reconstruction. In *European Conference on Computer Vision (ECCV)*. 796–812.
- Ayush Tewari, Florian Bernard, Pablo Garrido, Gaurav Bharaj, Mohamed Elgharib, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. 2019. FML: Face Model Learning from Videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 10812–10822.
- Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. 2020. StyleRig: Rigging StyleGAN for 3D Control Over Portrait Images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 6141–6150.
- Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeonwoo Kim, Patrick Pérez, and Christian Theobalt. 2018. Self-supervised multi-level face model learning for monocular reconstruction at over 250 Hz. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2549–2559.
- Ayush Tewari, Michael Zollhöfer, Hyeonwoo Kim, Pablo Garrido, Florian Bernard, Patrick Pérez, and Christian Theobalt. 2017. MoFA: Model-Based Deep Convolutional Face Autoencoder for Unsupervised Monocular Reconstruction. In *IEEE International Conference on Computer Vision (ICCV)*. 1274–1283.
- Justus Thies, Michael Zollhöfer, Matthias Nießner, Levi Valgaerts, Marc Stamminger, and Christian Theobalt. 2015. Real-time expression transfer for facial reenactment. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 183–1.
- Justus Thies, Michael Zollhöfer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. 2016. Face2Face: Real-time face capture and reenactment of RGB videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2387–2395.

- Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gerard Medioni. 2017. Regressing Robust and Discriminative 3D Morphable Models With a Very Deep Neural Network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1599–1608.
- Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard Medioni. 2018. Extreme 3D face reconstruction: Seeing through occlusions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3935–3944.
- Luan Tran, Feng Liu, and Xiaoming Liu. 2019. Towards High-Fidelity Nonlinear 3D Face Morphable Model. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1126–1135.
- Xiaoguang Tu, Jian Zhao, Zihang Jiang, Yao Luo, Mei Xie, Yang Zhao, Linxiao He, Zheng Ma, and Jiashi Feng. 2019. Joint 3D Face Reconstruction and Dense Face Alignment from A Single Image with 2D-Assisted Self-Supervised Learning. *IEEE International Conference on Computer Vision (ICCV)* (2019).
- Thomas Vetter and Volker Blanz. 1998. Estimating coloured 3D face models from single images: An example based approach. In *European Conference on Computer Vision (ECCV)*. 499–513.
- Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. 2019. Racial Faces in the Wild: Reducing Racial Bias by Information Maximization Adaptation Network. In *IEEE International Conference on Computer Vision (ICCV)*.
- Yi Wang, Xin Tao, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. 2018. Image inpainting via generative multi-column convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*. 331–340.
- Huawei Wei, Shuang Liang, and Yichen Wei. 2019. 3D Dense Face Alignment via Graph Convolution Networks. *arXiv preprint arXiv:1904.05562* (2019).
- Thibaut Weise, Sofien Bouaziz, Hao Li, and Mark Pauly. 2011. Realtime performance-based facial animation. *ACM Transactions on Graphics, (Proc. SIGGRAPH)* 30, 4 (2011), 77.
- Shugo Yamaguchi, Shunsuke Saito, Koki Nagano, Yajie Zhao, Weikai Chen, Kyle Olszewski, Shigeo Morishima, and Hao Li. 2018. High-fidelity Facial Reflectance and Geometry Inference from an Unconstrained Image. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 162:1–162:14.
- Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. 2020. FaceScape: a Large-scale High Quality 3D Face Dataset and Detailed Riggable 3D Face Prediction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 601–610.
- Xiaoxing Zeng, Xiaojiang Peng, and Yu Qiao. 2019. DF2Net: A Dense-Fine-Finer Network for Detailed 3D Face Reconstruction. In *IEEE International Conference on Computer Vision (ICCV)*.
- Yajie Zhao, Zeng Huang, Tianye Li, Weikai Chen, Chloe LeGendre, Xinglei Ren, Ari Shapiro, and Hao Li. 2019. Learning perspective undistortion of portraits. In *IEEE International Conference on Computer Vision (ICCV)*. 7849–7859.
- Xiangyu Zhu, Zhen Lei, Junjie Yan, Dong Yi, and Stan Z Li. 2015. High-fidelity pose and expression normalization for face recognition in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 787–796.
- Michael Zollhöfer, Justus Thies, Pablo Garrido, Derek Bradley, Thabo Beeler, Patrick Pérez, Marc Stamminger, Matthias Nießner, and Christian Theobalt. 2018. State of the Art on Monocular 3D Face Reconstruction, Tracking, and Applications. *Computer Graphics Forum (Eurographics State of the Art Reports 2018)* 37, 2 (2018).