

AVFace: Towards Detailed Audio-Visual 4D Face Reconstruction

Aggelina Chatziagapi, Dimitris Samaras

Stony Brook University

{aggelina,samaras}@cs.stonybrook.edu




Figure 1. Given a monocular talking face video and the corresponding speech segment, AVFace reconstructs the *detailed 4D face geometry* for any input speaker, and it is robust in cases when either modality is insufficient.

Abstract

In this work, we present a multimodal solution to the problem of 4D face reconstruction from monocular videos. 3D face reconstruction from 2D images is an under-constrained problem due to the ambiguity of depth. State-of-the-art methods try to solve this problem by leveraging visual information from a single image or video, whereas 3D mesh animation approaches rely more on audio. However, in most cases (e.g. AR/VR applications), videos include both visual and speech information. We propose AVFace that incorporates both modalities and accurately reconstructs the 4D facial and lip motion of any speaker, without requiring any 3D ground truth for training. A coarse stage estimates the per-frame parameters of a 3D morphable model, followed by a lip refinement, and then a fine stage recovers facial geometric details. Due to the temporal audio and video information captured by transformer-based modules, our method is robust in cases when either modality is insufficient (e.g. face occlusions). Extensive qualitative and quantitative evaluation demonstrates the superiority of our method over the current state-of-the-art.

1. Introduction

Reconstructing the 4D geometry of the human face has been a long standing research problem in computer vision and graphics. Accurate spatio-temporal (4D) face re-

construction has extensive applications in AR/VR, video games, virtual communication, the movie industry etc. However, recovering the per-frame 3D head pose and facial geometry from 2D images is an ill-posed problem due to the ambiguity of depth. Current approaches are largely based on 3D morphable models (3DMMs). Usually, they take a single image [18, 22] or video [60] as input and predict the 3DMM parameters. Some have also tried to predict additional geometric facial details, using the 3DMM fitting as prior [13, 18, 22]. However, most of these video-only methods are either speaker-specific [13, 23, 29], requiring to overfit to a specific speaker to recover their facial details, or fail to accurately capture fine details, like wrinkles and lip movements. In addition, they cannot handle face occlusions, since they solely rely on the visual input. On the other hand, audio-driven 3D mesh animation approaches [17, 21, 52] learn better lip motion, but they require 4D ground truth scans for training, which are rare and expensive to capture. Furthermore, such audio-only methods cannot capture any speaker-specific characteristics or facial expressions, as they do not use any visual information. There is limited work in audio-visual 4D face reconstruction [1, 15, 44], but these methods also require 4D ground truth scans and do not recover any facial geometric details.

In this work, we propose AVFace that learns to reconstruct detailed 4D face geometry from monocular talking face videos, leveraging both audio and video modalities. Without requiring any 3D ground truth scans, it can recover

accurate 4D facial and lip motion for any speaker. A coarse stage estimates a coarse geometry per frame, based on a 3DMM and using both image and speech features. Then, a SIREN MLP [58] further improves the lip position, by learning an implicit representation of the lip shape conditioned on speech. Finally, a fine stage recovers geometric facial details, guided by pseudo-ground truth face normals and producing a high-fidelity reconstruction of the input speaker’s face per frame. Due to the temporal audio and video information captured by transformer-based modules, our method is robust in cases when either modality is insufficient (e.g. face occlusions). To better handle such hard cases that are frequent in talking face videos, we further fine-tune our coarse stage with synthetic face occlusions.

In brief, the contributions of our work are as follows:

- We propose AVFace, a novel audio-visual method for detailed 4D face reconstruction, that follows a coarse-to-fine optimization approach, trained only on monocular talking face videos without any 3D ground truth.
- We introduce an audio-driven lip refinement network, and a fine stage guided by pseudo-ground truth face normals to accurately recover fine geometric details.
- Our temporal modeling, along with fine-tuning on synthetic face occlusions, makes our network robust to cases when either modality is insufficient.

2. Related Work

3D Face Reconstruction. Reconstructing the 3D face geometry from a single image or video has received a lot of attention in the last decades. One of the first methods was the work of Blanz and Vetter [6], which learns a 3DMM from 3D face scans. Following this work, several optimization-based [5,28,49] and learning-based [12,18,19,22,31,38,40,63,71,75] methods regress the parameters of a 3DMM from a single image. However, the PCA-based representation for shape and expression of 3DMMs is limited. Recently, deep learning approaches use the 3DMM fitting as prior and learn additional corrective displacements or normal maps [14,18,22,33,35,36,53,54,61], in order to capture missing facial details. DECA [22] and EMOCA [18] first predict the parameters of FLAME [42], and then learn a UV-map of vertex displacements. Tran et al. [64] borrow details from reference bump maps in order to complete occluded facial regions. Several subject-specific methods leverage multi-view images [35,41,43,45,55,56,63,69] or learn implicit representations using multi-layer perceptrons (MLPs) [13,23,29,46,73]. However, these methods cannot generalize to multiple identities. Monocular video-based optimization techniques [24–27,57,60] leverage the multi-frame consistency to learn facial details. FML [60]

proposes a self-supervised method to learn identity and appearance models from videos. However, all these video-only methods ignore the audio modality that is usually included in most cases. In addition, they do not explicitly handle occlusions that frequently appear in talking face videos.

Audio-Driven 3D Face Animation. Another line of work addresses the problem of audio-driven 3D facial animation. Busso et al. [8,9] use hidden Markov models (HMMs) to model the temporal relation between prosodic audio features and head motion sequences. Several works learn to animate subject-specific face models [10,37,48,51] or artist-designed character rigs [20,59,74] based on input speech. More recently, VOCA [17] maps audio features to 3D vertex displacements from a neutral face mesh. MeshTalk [52] aims to synthesize accurate lip motion and at the same time, plausible animation of parts of the face that are uncorrelated to the audio signal. FaceFormer [21] proposes a transformer-based model that encodes long-term audio context and autoregressively predicts a sequence of 3D face meshes. While these works learn well-synchronized lip movements with the input speech, they cannot capture speaker-specific characteristics or facial expressions, since they do not use any visual input. Furthermore, they require high-quality 4D scans for training. In contrast, our method is trained on monocular talking face videos, without using any 3D ground truth.

Audio-Visual 4D Face Reconstruction. There is limited work for 4D face reconstruction that considers both audio and video modalities. In [44], the authors propose an audio-visual speaker-independent system for real-time facial tracking and animation that is robust to occlusions. Their approach is highly data-driven, using a pre-captured database of 3D mouth shapes, and requires depth information. Chen et al. [15] present an optimization-based approach that learns a phoneme to 3D blendshape mapping, leveraging artist-designed blendshape models. An audio-visual deep learning approach is proposed by Abdelaziz et al. [1], where modality dropout encourages the network to pay attention to the audio input. In contrast to all these methods, AVFace does not require any 3D ground truth. It learns to recover detailed 4D geometry of any speaker and is robust to occlusions, trained only on monocular videos.

3. Method

We present AVFace, an audio-visual method that can accurately reconstruct the 4D face geometry of any input speaker. Given a monocular RGB video of a talking face and the corresponding speech segment, AVFace follows a coarse-to-fine optimization approach, where audio and video modalities complement each other. Leveraging temporal information, it successfully handles occlusions (e.g. hand in front of the speaker’s face) that are frequent in talking face videos. An overview of our method is shown in

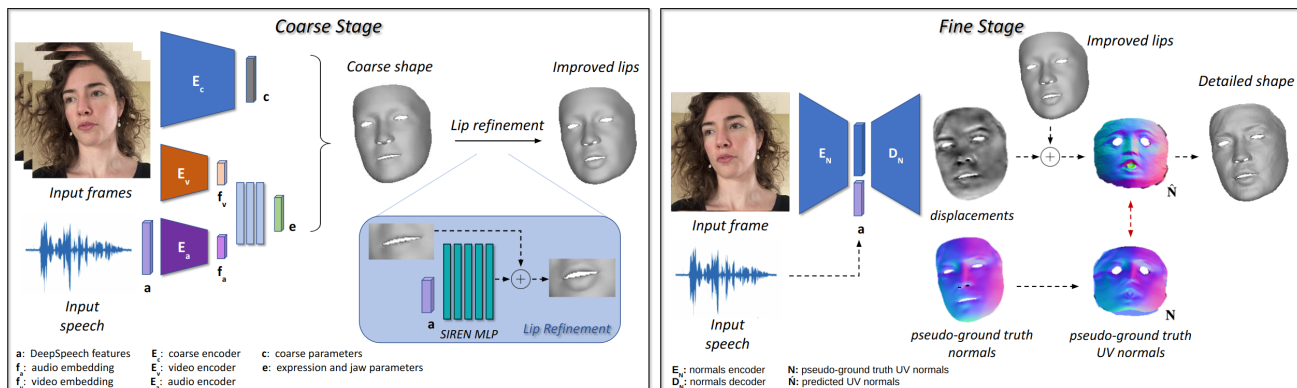


Figure 2. **AVFace overview.** Given an input video of a talking face and the corresponding speech segment, AVFace follows a coarse-to-fine optimization approach, in order to reconstruct the detailed 4D face geometry. A coarse stage estimates the per-frame parameters of a 3DMM. A SIREN MLP further improves the lip position. Finally, a fine stage recovers facial geometric details.

Fig. 2. In the following paragraphs, we describe each individual component in detail.

3.1. Coarse Stage

Coarse Encoder. In the coarse stage, we first estimate a coarse geometry based on the FLAME morphable model [42]. Given x consecutive frames from an input video, a coarse encoder E_c predicts a latent code $\mathbf{c} \in \mathbb{R}^{183}$ for each frame (see Fig. 2), which includes the FLAME parameters for head pose, camera, and shape, as well as the albedo and spherical harmonics coefficients for texture rendering [22]. Although we do not focus on texture prediction in this work, we use the rendered image for photometric loss during training. The coarse encoder consists of a ResNet-50 [34] and a transformer encoder [65]. The ResNet is initialized with pre-trained weights from DECA [22] and fine-tuned during training to predict intermediate embeddings of same size as \mathbf{c} . The transformer encoder is a stack of 3 encoder layers that include multi-head self-attention and feed-forward layers, to output the final \mathbf{c} parameters per frame. Leveraging temporal information, it ensures smooth and accurate prediction of the head pose and camera for the sequence of x input frames. In this way, we successfully handle cases where the face is partially or fully occluded, and single-image based methods like DECA [22] might fail.

Expression and Jaw Prediction. AVFace includes a separate branch for the prediction of the jaw pose and expression parameters $\mathbf{e} \in \mathbb{R}^{53}$, based on both audio and video information. Unlike the expression encoder proposed in EMOCA [18], our expression encoder has the following advantages: (a) we include both speech and visual information, handling cases where either modality is insufficient, (b) our transformer-based module captures temporal information across the x input frames, compared to the single-image input of EMOCA, and (c) we include the jaw pose, which controls the mouth opening and closing and is crucial

for accurate lip movement capture.

Given the raw audio signal of an input video, we extract DeepSpeech [3] features $\mathbf{a} \in \mathbb{R}^{16 \times 29}$ per frame, which correspond to 29-dimensional embeddings for 16 overlapping windows of 20ms length each. DeepSpeech is a large RNN, trained on thousands of hours of data for automatic speech recognition. In this way, it captures useful information for any articulated phoneme and has been used for audio feature extraction by related works [17, 32, 62]. The features \mathbf{a} are given as input to an audio encoder E_a , which is an 1D convolutional network and learns an audio embedding \mathbf{f}_a for each frame. A video encoder E_v outputs the corresponding video embedding \mathbf{f}_v , following a ResNet-50 architecture. The audio and video embeddings are concatenated and passed to a transformer encoder of 3 stacked layers. Leveraging temporal information from the x input frames, the transformer enhances the expression and jaw prediction. In cases of occlusion, the audio input leads to an accurate mouth position based on the recognized phoneme. In order to encourage the network to pay attention to the speech input, and not only to the visual input that is strongly correlated with the output shape, we use modality dropout during training, where \mathbf{f}_a and \mathbf{f}_v are zeroed out with a given probability [1].

Training. In the coarse stage, we minimize a similar objective with DECA, dubbed $\mathcal{L}_{\text{coarse}}$ loss [22]. In addition to the eye closure loss, we add a mouth closure and lip corner loss that measure the relative distances between upper/lower lips and left/right lip corners correspondingly [18].

3.2. Lip Refinement

Following the coarse shape estimation, we further improve the lip shape during a lip refinement stage (see Fig. 2). For this purpose, we use a SIREN MLP [58], which can learn powerful implicit representations of complex natural signals [11, 58] and capture high frequencies due to the pe-

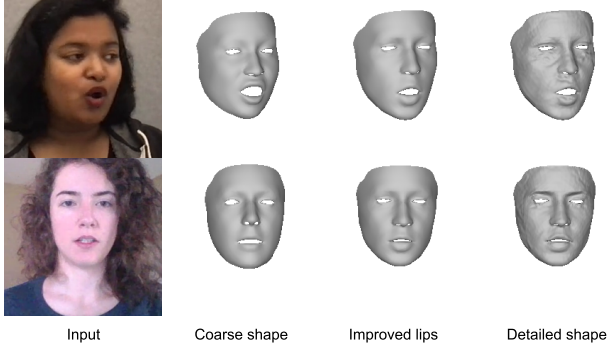


Figure 3. From the coarse stage of AVFace, to the lip refinement, and finally to the detailed 3D reconstruction for an input frame.

riodic activations. Conditioned on the audio features \mathbf{a} , the SIREN MLP predicts 3D vertex displacements for the coarse lip vertices, through a series of FiLM-conditioned linear layers [47]. We follow the architecture of the geometry network of NHA [29]. However, in our case the FiLM layers are conditioned on the audio input and we predict lip offsets for any speaker, in contrast to the speaker-specific approach of NHA that predicts vertex offsets for the entire head conditioned on pose. The predicted lip offsets are added to the coarse lip vertices, in order to better capture the shape and position of the mouth. For example, in the first row of Fig. 3, the lips of the coarse shape resemble an /a/ phoneme, and are modified to the correct /o/ phoneme during lip refinement (denoted as “Improved lips”). The audio conditioning plays a significant role in inferring the lip offsets. For frames without audio, the lip offsets are negligible and the lip shape is deduced based on video only.

We train the lip refinement stage by minimizing the difference between the predicted normals in the lip region \hat{N}_{lip} and the corresponding pseudo-ground truth normals N_{lip} that are extracted from the pre-trained model for face normals by [2]. To better capture the high frequencies, we apply the Laplace operator ∇^2 as follows:

$$\mathcal{L}_{\text{lip}} = \mathcal{L}_{\text{coarse}} + \lambda_n \|\nabla^2 \hat{N}_{\text{lip}} - \nabla^2 N_{\text{lip}}\|_{1,1} \quad (1)$$

3.3. Fine Stage

The fine stage recovers facial geometric details, like wrinkles and skin folds. In this stage, we train a convolutional neural network that predicts offsets of face normals from the coarse output as a one-channel image. The network is based on the encoder-decoder UNet-ResNet architecture proposed by [2], and it is initialized by their pre-trained weights. We experimentally found that adding audio input improves the normals prediction, especially in the mouth region. Thus, we concatenate the extracted audio features \mathbf{a} as the last channel of the output of the encoder’s

third layer. The predicted normal displacements D are combined with the coarse normals N_c in the UV space:

$$V_d = V_c + D \odot N_c \quad (2)$$

where V_c and V_d are the coarse and detailed vertices in the UV space correspondingly. The final detailed face normals \hat{N} , extracted from the detailed shape V_d , are learned by minimizing the normals loss, in order to capture high-frequency details, similarly to the lip refinement stage:

$$\mathcal{L}_{\text{normals}} = \|\nabla^2 \hat{N} - \nabla^2 N\|_{1,1} \quad (3)$$

where N are the pseudo-ground truth normals of the entire face, extracted from the pre-trained model by [2]. We also minimize a photometric loss \mathcal{L}_i on the final detailed rendered image, an ID-MRF loss \mathcal{L}_{mrf} [22, 68], a soft symmetry loss \mathcal{L}_{sym} [22], and a regularization term $\mathcal{L}_{\text{reg}} = \|D\|_{1,1}$. Overall, the final loss of our fine stage is:

$$\mathcal{L}_{\text{detail}} = \mathcal{L}_i + \lambda_n \mathcal{L}_{\text{normals}} + \lambda_m \mathcal{L}_{\text{mrf}} + \lambda_s \mathcal{L}_{\text{sym}} + \lambda_d \mathcal{L}_{\text{reg}} \quad (4)$$

3.4. Fine-tuning with Synthetic Occlusions

Due to the temporal nature of our coarse stage and the audio input, AVFace learns to handle occlusions, predicting the correct head pose, facial expression, and lip position in intermediate frames where part of the face is hidden, e.g. by a hand motion. Our training data include a small percentage of frames with face occlusions, while the speaker is talking. In order to enhance AVFace’s robustness in those cases, we fine-tune the coarse stage with a small number of synthetic data. We created frame sequences with synthetic hand occlusions, using the face occlusion generation method proposed by [66]. For a random sequence of consecutive frames from our real data, we synthesize a hand occlusion for the first frame, transferring the color from the face via Sliced Optimal Transport [7] (see examples in Fig. 7). Then, we randomly rotate and translate the hand from frame to frame to approximate a realistic hand motion.

3.5. Implementation Details

We use the first 100 shape, 50 expression and 50 albedo parameters of FLAME, following DECA [22]. The transformer encoders for both the coarse and expression-jaw prediction include 3 encoder layers, 4 heads for the self-attention mechanism, 256 hidden units for the feed-forward layers, no dropout, and layer normalization [65], followed by 3 feed-forward layers with ReLU activations for projection. During training, we use windows of $x = 16$ frames. During test time, we use overlapping windows of the same size with a step of 1 frame. We experimentally chose a size of 64 for the audio and video embeddings, \mathbf{f}_a and \mathbf{f}_v , correspondingly. The modality dropout probabilities are 0.4 for audio and 0.5 for video. The loss weights are set to

Method	NME ↓	AUC ↑	NME (lips) ↓	NME (occ) ↓
e from Video	2.60	0.82	1.78	2.91
+ Temporal	1.98	0.82	1.33	2.68
e from Audio	2.80	0.81	2.28	2.86
+ Temporal	2.10	0.81	1.56	2.55
e from AV	2.91	0.81	2.05	3.02
+ Dropout	2.14	0.82	1.65	2.79
+ Temporal	1.73	0.83	0.99	2.42
+ Lip Refine	1.63	0.84	0.91	2.39

Table 1. **Ablation study on the coarse stage.** Quantitative results for different variations of our coarse stage on the FaceSet (col. 1-3) and on the synthetic test set with occlusions (col. 4).

Audio Features	NME ↓	AUC ↑	NME (lips) ↓
Mel-spectrogram	2.20	0.82	1.70
Wav2vec 2.0	1.66	0.83	1.11
DeepSpeech (Ours)	1.63	0.84	0.91

Table 2. **Ablation study on the audio features.** Quantitative results on the FaceSet for different audio features, namely mel-spectrogram [50], wav2vec 2.0 [4], and DeepSpeech [3].

$\lambda_l = 5 \times 10^{-1}$, $\lambda_c = 10^{-4}$, $\lambda_n = 10$, $\lambda_m = 5 \times 10^{-2}$, $\lambda_s = 5 \times 10^{-3}$, $\lambda_d = 5 \times 10^{-3}$. We use Adam optimizer [39] with a learning rate of 10^{-4} and weight decay of 5×10^{-4} . See the suppl. material for more implementation details.

4. Experiments

4.1. Datasets

We use a publicly available dataset [16] that includes videos from 6112 identities for training. We use the proposed train-test split, and perform quantitative evaluation on the test set, which we call FaceSet. The videos are sampled at 25 fps and the audio at 16 kHz. For the synthetic data generation (see Sec. 3.4), we randomly sampled 3k and 1k sequences of 16 frames from the train and test set respectively.

For testing, we use Multiface [70], in order to evaluate our 3D reconstruction. Multiface includes videos of 13 identities in a multi-view capture stage, as well as the corresponding 3D ground truth meshes per frame. The videos are captured at 30 fps. For our case, only a single frontal camera view is passed as input to our model.

We additionally evaluate our method on 8 self-captured videos, captured in different environments and from different camera setups. We downsampled them to 25 fps, 640x480 video resolution and 16kHz audio sampling rate.

4.2. Ablation Study

We conduct an ablation study, in order to investigate the contribution of each part of our model. Fig. 3 shows the im-

Method	NME ↓	AUC ↑	NME (lips) ↓
Without synthetic occlusions	2.62	0.74	1.98
With synthetic occlusions	2.39	0.76	1.73

Table 3. **Importance of synthetic occlusions.** Quantitative results on the synthetic test set with occlusions, without and with fine-tuning on synthetic face occlusions.

provement from the coarse to the fine stage. The coarse stage produces a smooth 3D face geometry, capturing a rough estimate of the mouth position. During lip refinement, the lips are corrected. Finally, the fine stage recovers 3D geometric details of the entire face.

Coarse Stage. Table 1 shows the quantitative results for different variations of network. Since FaceSet does not include 3D ground truth, we compute the following metrics: normalized mean error (NME) (%), area under the curve (AUC), NME computed only for the lips, and NME on the synthetic test set with occlusions. In the first and third rows, the jaw pose and expression parameters e are predicted using video only (only E_v) and audio only (only E_a) correspondingly. Temporal modeling boosts their performance. The fifth row considers both modalities and the sixth row adds the modality dropout. We suspect that the simple concatenation of audio and video embeddings confuses the network. Adding the modality dropout encourages the network to pay attention to the individual modalities and extract useful information, thereby improving the performance. Then, the addition of temporal information via the transformer-based modules and lastly the addition of lip refinement further improve the results.

Audio Features. We compared our audio feature extraction with (a) mel-spectrograms and (b) wav2vec 2.0 [4]. For mel-spectrograms, we use the proposed 2D convolutional audio encoder of [50] to get the audio embedding f_a . Wav2vec 2.0 [4] is a state-of-the-art transformer-based ASR model, which learns powerful speech representations in a self-supervised manner. We use the pre-trained wav2vec 2.0 weights¹ and the final output is mapped to a 64-dimensional f_a . The quantitative comparison is shown on Table 2. The DeepSpeech features led to the best performance.

Synthetic Occlusions. Table 3 demonstrates the performance boost in cases of occlusion with the additional fine-tuning on our synthetic data (see Sec. 3.4).

4.3. Qualitative Evaluation

Fig. 4 compares our coarse reconstruction with state-of-the-art methods. Since there is limited work for joint audio-video 4D face reconstruction, we include comparisons with recent video-only, audio-only and audio-video approaches. For a fair comparison with the audio-only methods, which

¹<https://huggingface.co/facebook/wav2vec2-base-960h>

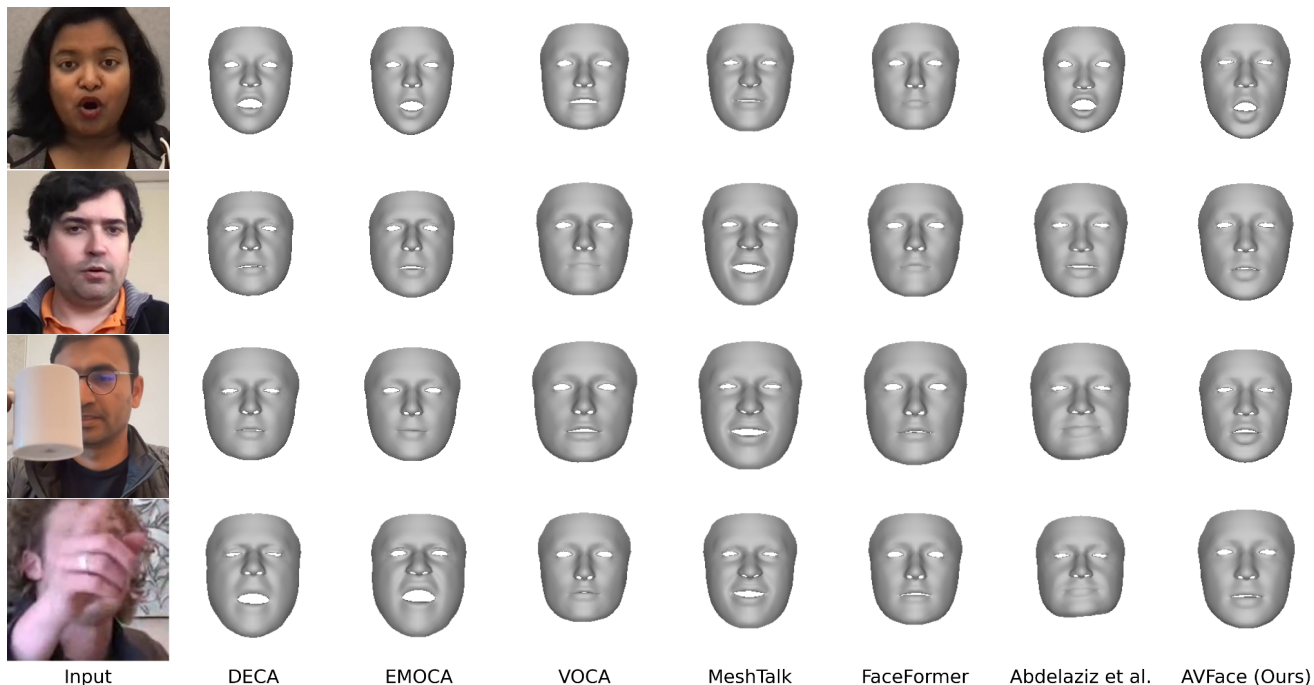


Figure 4. Comparison of our **coarse** reconstruction in frontal head pose with *video-only* (DECA [22], EMOCA [18]), *audio-only* (VOCA [17], MeshTalk [52], FaceFormer [21]) and *audio-video* (Abdelaziz et al. [1]) methods.

animate an input neutral mesh, we turn all results in frontal head pose. As input to the audio-only methods, we use a FLAME mesh fitted to the first frame of the video, in order to capture the speaker identity. As shown, the video-only methods DECA [22] and EMOCA [18] fail in case of occlusion, since they only consider a single image as input. The audio-only methods, namely VOCA [17], MeshTalk [52], and FaceFormer [21], are not affected by occlusions, since they only consider audio input. Abdelaziz et al. [1] have not published their code, so we implemented it ourselves. Overall, our method captures the lip motion more accurately than the other methods (see lips in first two rows) and is robust to occlusions (last two rows). For the last row where the face is fully occluded, the speaker pronounces the word /it/ at that specific moment, and only the lips of FaceFormer and our method resemble this sound.

Fig. 5 compares our final detailed reconstruction with video-only and audio-video methods. Similarly, the video-only methods DECA [22], EMOCA [18] and 3DDFA-V2 [30, 31] fail in case of occlusion, both in terms of the predicted head pose and the lip movement. The audio input in Abdelaziz et al. [1] helps the mouth prediction, but this method cannot handle occlusions either, since it lacks temporal modeling. Our method includes temporal audio and video information, and is fine-tuned to a small number of synthetic occlusions, making it robust in those cases. Furthermore, AVFace recovers accurate facial geometric de-

Method	NME ↓	AUC ↑	NME (lips) ↓	NME (occ) ↓
DECA	1.85	0.82	1.25	3.89
EMOCA	1.98	0.80	1.53	3.89
3DDFA-V2	1.64	0.84	1.15	3.47
VOCA	3.15	0.66	2.66	4.95
MeshTalk	4.48	0.47	4.59	4.89
FaceFormer	3.14	0.65	2.61	4.82
Abdelaziz et al.	2.26	0.82	1.72	2.76
AVFace (Ours)	1.63	0.84	0.91	2.39

Table 4. Quantitative comparison of our method with state-of-the-art methods on the FaceSet (col. 1-3) and on the synthetic test set with occlusions (col. 4). The first 3 rows correspond to *video-only* (DECA [22], EMOCA [18], 3DDFA-V2 [30, 31]), the next 3 to *audio-only* (VOCA [17], MeshTalk [52], FaceFormer [21]) and the last rows to *audio-video* (Abdelaziz et al. [1]) methods.

tails, like wrinkles and skin folds, that are largely missed by the other methods.

4.4. Quantitative Evaluation

Table 4 shows the quantitative evaluation. For the audio-only methods, we use the head pose as predicted by DECA. We notice that the audio-only methods perform poorly in terms of these metrics, as they do not consider visual information, and as a result they miss facial expressions and

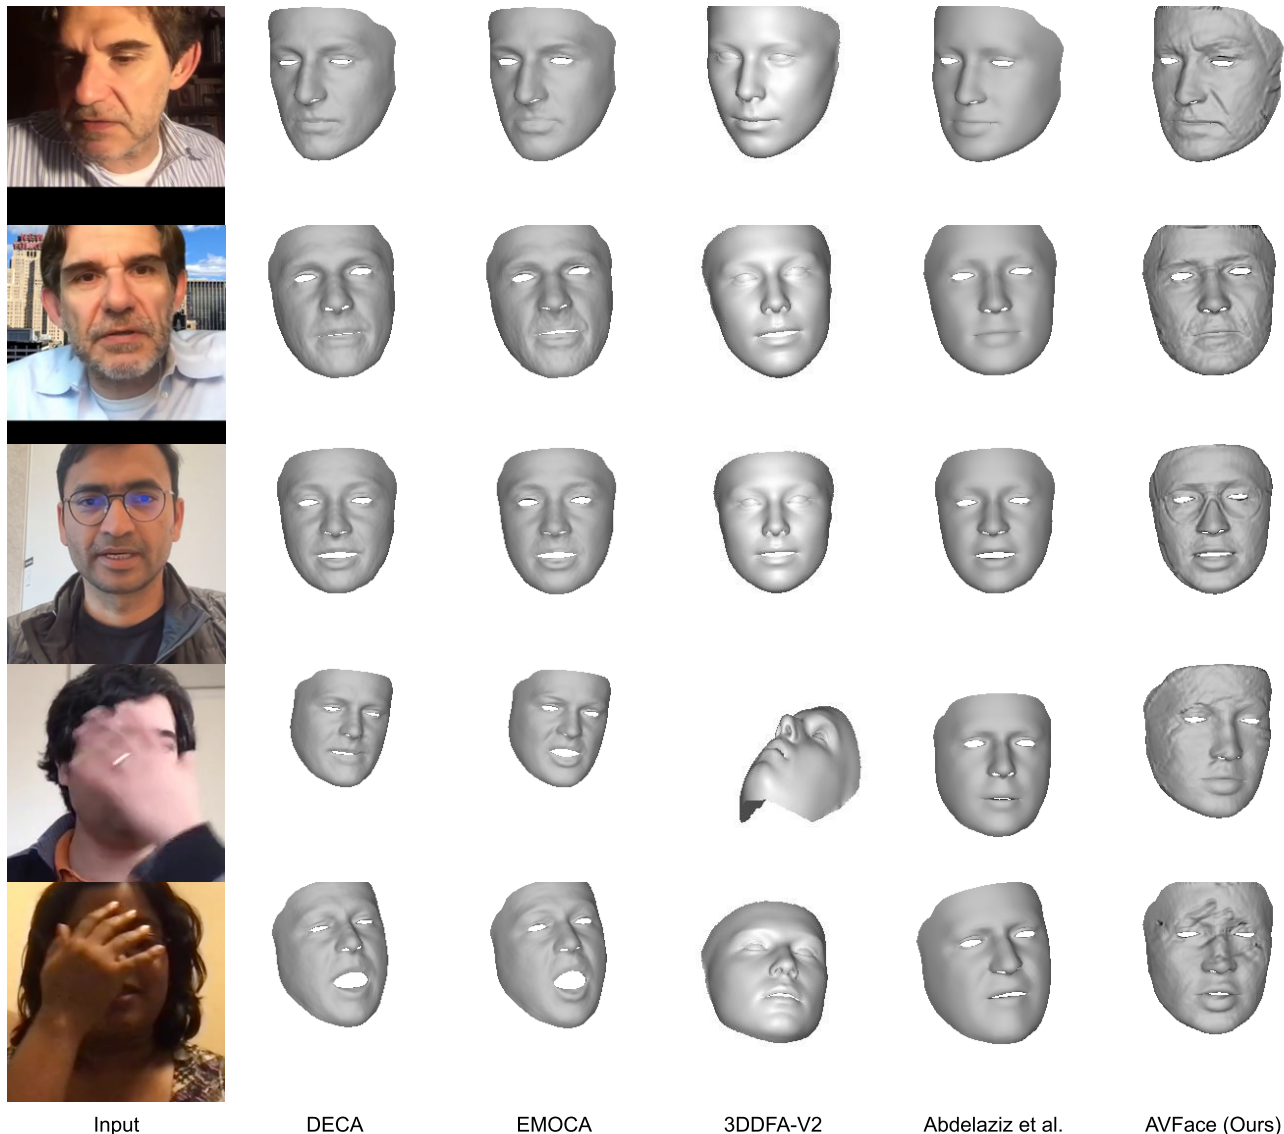


Figure 5. Comparison of our final **detailed** reconstruction with *video-only* (DECA [22], EMOCA [18], 3DDFA-V2 [30,31]) and *audio-video* (Abdelaziz et al. [1]) methods.

large mouth openings. The video-only methods give a high error in case of occlusions. Our method outperforms all these approaches across all metrics, and with a large margin for the test occlusions.

To evaluate the entire 3D mesh, we additionally compute the 3D reconstruction error on Multiface data (see Table 5). We follow the scan-to-mesh error computation of NoW [55] and compare with the video-only and audio-video methods. Our method has the least error, and the error decreases by around 5% from the coarse to the fine stage. Note that these data are not used for training, and they are very different from the training data in terms of illumination and camera pose. In addition, the ground truth meshes are relatively smooth, missing some fine details and as a result, the margin

of our method from DECA is relatively small, in contrast to the qualitative comparison. Fig. 6 illustrates the 3D reconstruction error per vertex on the predicted mesh for 2 examples. We compare with DECA, since it gave the second lowest scores. Our method better reconstructs the speaker’s face shape, facial details, as well as lip position.

We additionally conducted a user study with 15 random participants to evaluate our 4D face reconstruction for full video sequences. We included 2 sections. The first section shows video results of FaceFormer [21], DECA [22], and AVFace in frontal head pose (similar to Fig. 4) and ask the users to choose the method with most accurate lip synchronization (i.e. the lips better follow the speech). In the second section, we compare our final reconstruction with

Method	Median (mm) ↓	Mean (mm) ↓	Std (mm) ↓
DECA	1.68	2.01	1.60
EMOCA	1.75	2.09	1.66
3DDFA-V2	1.83	2.28	2.18
Abdelaziz et al.	1.71	2.06	1.67
AVFace (Ours)	1.61	1.98	1.59

Table 5. 3D reconstruction error on Multiface data comparing with *video-only* (DECA [22], EMOCA [18], 3DDFA-V2 [30, 31]) and *audio-video* (Abdelaziz et al. [1]) methods.

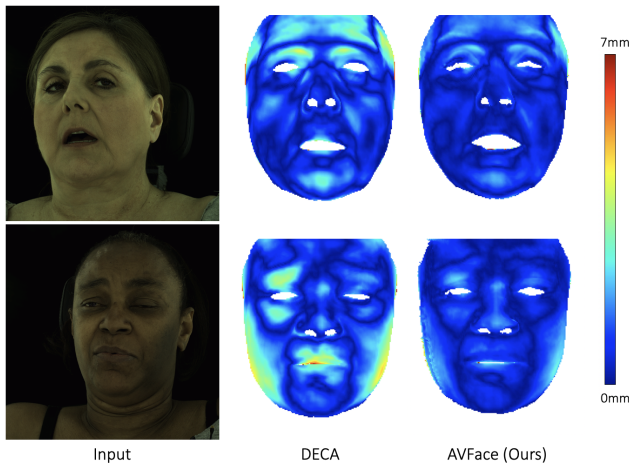


Figure 6. 3D reconstruction error on Multiface data comparing with DECA [22].



Figure 7. **Synthetic occlusions.** Example of our synthetic hand occlusions for a ground truth video frame and the corresponding 3D reconstruction.

DECA, similar to Fig. 5. We ask the users to choose the method that (a) produces better facial details, (b) better handles face occlusions, and (c) has the best quality overall. For each section, we include 5 videos of around 10-30 seconds from FaceSet, making sure that they include face occlusions, are diverse w.r.t. speech variations (e.g. intonation, energy, speaking rate), and are balanced w.r.t. speaker age and gender. In this way, we expect to get a representative sample of talking face videos. The results of our user study are shown in Fig. 8. FaceFormer produces better audio-lip synchronization compared to DECA and a few users prefer DECA over AVFace in terms of details and overall quality.

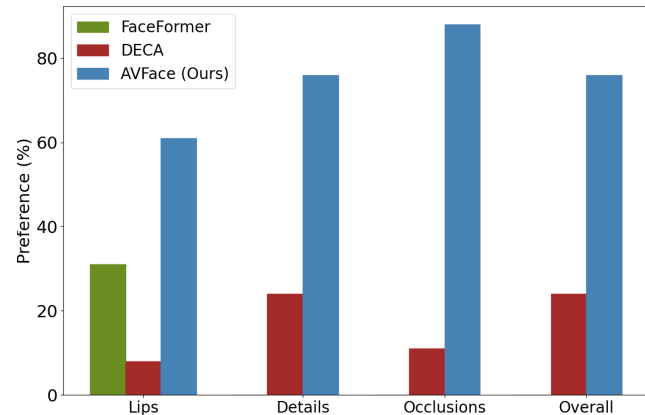


Figure 8. **User study.** Preference (%) in terms of lip synchronization, facial details, robustness in occlusions, and overall quality, comparing FaceFormer [21], DECA [22], and our method.

In all questions, our method leads to the highest preference scores.

5. Discussion

Limitations. Since we predict face normal offsets in the fine stage, face occlusions like hair, glasses and hands can be rendered as additional details, engraved on the speaker’s face (see Fig. 7). This can be handled with additional temporal modeling or depth input which are out of the scope of this work. We also noticed that during occlusions of the whole face, our method might produce a more frontal head pose (last row in Fig. 5). In the future, we plan to further improve our method’s robustness to such extreme cases.

Ethical Considerations. We would like to note the potential misuse of face reconstruction methods. With the advances in detailed 3D reconstruction and neural rendering, it becomes easier to generate photo-realistic fake videos of any speaker, which can be used for malicious purposes (e.g. to spread misinformation). Thus, it is important to ensure fair and safe use of videos, and develop accurate methods for fake content detection and forensics [67, 72].

6. Conclusion

We propose AVFace, a novel audio-visual approach that addresses the problem of 4D face reconstruction from monocular talking face videos. Without requiring any 3D ground truth, AVFace recovers detailed 4D facial and lip motion of any input speaker, following a coarse-to-fine optimization strategy. In addition, leveraging temporal information, it is robust to cases when either modality is insufficient (e.g. face occlusions). Based on a thorough qualitative and quantitative evaluation, we show that AVFace outperforms the current state-of-the-art. In the future, we plan to further enhance its robustness to extreme cases, where the face is fully occluded for multiple consecutive frames.

References

- [1] Ahmed Hussen Abdelaziz, Barry-John Theobald, Paul Dixon, Reinhard Knothe, Nicholas Apostoloff, and Sachin Kajareker. Modality dropout for improved performance-driven talking faces, 2020. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [2] Victoria Fernandez Abrevaya, Adnane Boukhayma, Philip H.S. Torr, and Edmond Boyer. Cross-modal deep face normals with deactivable skip connections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [4](#)
- [3] Dario Amodei, Rishita Anubhai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Jingdong Chen, Mike Chrzanowski, Adam Coates, Greg Diamos, Erich Elsen, Jesse H. Engel, Linxi Fan, Christopher Fougner, Tony Han, Awni Y. Hannun, Billy Jun, Patrick LeGresley, Libby Lin, Sharan Narang, Andrew Y. Ng, Sherjil Ozair, Ryan Prenger, Jonathan Raiman, Sanjeev Satheesh, David Seetapun, Shubho Sengupta, Yi Wang, Zhiqian Wang, Chong Wang, Bo Xiao, Dani Yogatama, Jun Zhan, and Zhenyao Zhu. Deep speech 2: End-to-end speech recognition in english and mandarin. *arXiv preprint arXiv:1512.02595*, 2015. [3](#), [5](#)
- [4] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33:12449–12460, 2020. [5](#)
- [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 Chu-Song Chen, Jiwen Lu, and Kai-Kuang Ma, editors, *Computer Vision – ACCV 2016 Workshops*, pages 377–391, Cham, 2017. Springer International Publishing. [2](#)
- [6] Volker Blanz, Thomas Vetter, et al. A morphable model for the synthesis of 3D faces. 1999. [2](#)
- [7] Nicolas Bonneel and David Coeurjolly. Spot: Sliced partial optimal transport. *ACM Trans. Graph.*, 38(4), jul 2019. [4](#)
- [8] Carlos Busso, Zhigang Deng, Michael Grimm, Ulrich Neumann, and Shrikanth Narayanan. Rigid head motion in expressive speech animation: Analysis and synthesis. *IEEE transactions on audio, speech, and language processing*, 15(3):1075–1086, 2007. [2](#)
- [9] Carlos Busso, Zhigang Deng, Ulrich Neumann, and Shrikanth Narayanan. Natural head motion synthesis driven by acoustic prosodic features. *Computer Animation and Virtual Worlds*, 16(3-4):283–290, 2005. [2](#)
- [10] Yong Cao, Wen C Tien, Petros Faloutsos, and Frédéric Pighin. Expressive speech-driven facial animation. *ACM Transactions on Graphics (TOG)*, 24(4):1283–1302, 2005. [2](#)
- [11] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5799–5809, 2021. [3](#)
- [12] Feng-Ju Chang, Anh Tuan Tran, Tal Hassner, Iacopo Masi, Ram Nevatia, and Gerard Medioni. Expnet: Landmark-free, deep, 3d facial expressions. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 122–129. IEEE, 2018. [2](#)
- [13] Aggelina Chatziagapi, ShahRukh Athar, Francesc Moreno-Noguer, and Dimitris Samaras. SIDER: Single-image neural optimization for facial geometric detail recovery. In *2021 International Conference on 3D Vision (3DV)*, pages 815–824, 2021. [1](#), [2](#)
- [14] Bindita Chaudhuri, Noranart Vesdapunt, Linda Shapiro, and Baoyuan Wang. Personalized face modeling for improved face reconstruction and motion retargeting. In *European Conference on Computer Vision*, pages 142–160. Springer, 2020. [2](#)
- [15] Xin Chen, Chen Cao, Zehao Xue, and Wei Chu. Joint audio-video driven facial animation. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3046–3050, 2018. [1](#), [2](#)
- [16] J. S. Chung, A. Nagrani, and A. Zisserman. Voxceleb2: Deep speaker recognition. In *INTERSPEECH*, 2018. [5](#)
- [17] Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael Black. Capture, learning, and synthesis of 3D speaking styles. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10101–10111, 2019. [1](#), [2](#), [3](#), [6](#)
- [18] Radek Danecek, Michael J. Black, and Timo Bolkart. EMOCA: Emotion driven monocular face capture and animation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [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 *IEEE Conference on Computer Vision and Pattern Recognition Workshop on Analysis and Modeling of Faces and Gestures (CVPRW)*, 2019. [2](#)
- [20] Pif Edwards, Chris Landreth, Eugene Fiume, and Karan Singh. Jali: an animator-centric viseme model for expressive lip synchronization. *ACM Transactions on graphics (TOG)*, 35(4):1–11, 2016. [2](#)
- [21] Yingruo Fan, Zhaojiang Lin, Jun Saito, Wenping Wang, and Taku Komura. FaceFormer: Speech-driven 3D facial animation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [2](#), [6](#), [7](#), [8](#)
- [22] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. volume 40, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [23] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4D facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8649–8658, June 2021. [1](#), [2](#)
- [24] Ravi Garg, Anastasios Roussos, and Lourdes Agapito. Dense variational reconstruction of non-rigid surfaces from monocular video. In *Proceedings of the IEEE Conference on computer vision and pattern recognition*, pages 1272–1279, 2013. [2](#)

- [25] Pablo Garrido, Levi Valgaerts, Chenglei Wu, and Christian Theobalt. Reconstructing detailed dynamic face geometry from monocular video. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 32(6), November 2013. 2
- [26] Pablo Garrido, Michael Zollhöfer, Dan Casas, Levi Valgaerts, Kiran Varanasi, Patrick Pérez, and Christian Theobalt. Reconstruction of personalized 3D face rigs from monocular video. *ACM Transactions on Graphics*, 35(3), May 2016. 2
- [27] Pablo Garrido, Michael Zollhöfer, Dan Casas, Levi Valgaerts, Kiran Varanasi, Patrick Pérez, and Christian Theobalt. Reconstruction of personalized 3D face rigs from monocular video. *ACM Transactions on Graphics (TOG)*, 35(3):1–15, 2016. 2
- [28] Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Luthi, Sandro Schönborn, and Thomas Vetter. Morphable face models—an open framework. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 75–82. IEEE, 2018. 2
- [29] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular RGB videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18653–18664, 2022. 1, 2, 4
- [30] Jianzhu Guo, Xiangyu Zhu, and Zhen Lei. 3ddf. <https://github.com/cleardusk/3DDFA>, 2018. 6, 7, 8
- [31] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3D dense face alignment. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2, 6, 7, 8
- [32] Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. AD-NeRF: Audio driven neural radiance fields for talking head synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5784–5794, 2021. 3
- [33] Yudong Guo, Juyong Zhang, Jianfei Cai, Boyi Jiang, and Jianmin Zheng. Cnn-based real-time dense face reconstruction with inverse-rendered photo-realistic face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018. 2
- [34] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 3
- [35] Loc Huynh, Weikai Chen, Shunsuke Saito, Jun Xing, Koki Nagano, Andrew Jones, Paul Debevec, and Hao Li. Mesoscopic Facial Geometry Inference Using Deep Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [36] L. Jiang, J. Zhang, B. Deng, H. Li, and L. Liu. 3D face reconstruction with geometry details from a single image. *IEEE Transactions on Image Processing*, 27(10):4756–4770, Oct 2018. 2
- [37] Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen. Audio-driven facial animation by joint end-to-end learning of pose and emotion. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017. 2
- [38] Hyeonwoo Kim, Michael Zollhöfer, Ayush Tewari, Justus Thies, Christian Richardt, and Christian Theobalt. Inversefacenet: Deep monocular inverse face rendering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4625–4634, 2018. 2
- [39] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [40] Tatsuro Koizumi and William AP Smith. “look ma, no landmarks!”—unsupervised, model-based dense face alignment. In *European Conference on Computer Vision*, pages 690–706. Springer, 2020. 2
- [41] Samuli Laine, Tero Karras, Timo Aila, Antti Herva, Shunsuke Saito, Ronald Yu, Hao Li, and Jaakko Lehtinen. Production-level facial performance capture using deep convolutional neural networks. In *Eurographics Symposium on Computer Animation*, 2017. 2
- [42] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics*, 36(6), Nov. 2017. 2, 3
- [43] Feng Liu, Ronghang Zhu, Dan Zeng, Qijun Zhao, and Xiaoming Liu. Disentangling features in 3D face shapes for joint face reconstruction and recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [44] Yilong Liu, Feng Xu, Jinxiang Chai, Xin Tong, Lijuan Wang, and Qiang Huo. Video-audio driven real-time facial animation. *ACM Trans. Graph.*, 34(6), nov 2015. 1, 2
- [45] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *ACM Transactions on Graphics*, 37(4), July 2018. 2
- [46] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 2
- [47] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 4
- [48] Hai X Pham, Samuel Cheung, and Vladimir Pavlovic. Speech-driven 3D facial animation with implicit emotional awareness: a deep learning approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 80–88, 2017. 2
- [49] Stylianos Ploumpis, Evangelos Ververas, Eimear O’Sullivan, Stylianos Moschoglou, Haoyang Wang, Nick Pears, William Smith, Baris Gecer, and Stefanos P Zafeiriou. Towards a complete 3D morphable model of the human head. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 2
- [50] K R Prajwal, Rudrabha Mukhopadhyay, Vinay P. Namboodiri, and C.V. Jawahar. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM International Conference on Multimedia*, MM ’20, page 484–492, New York, NY, USA, 2020. Association for Computing Machinery. 5

- [51] Alexander Richard, Colin Lea, Shugao Ma, Jurgen Gall, Fernando De la Torre, and Yaser Sheikh. Audio-and gaze-driven facial animation of codec avatars. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 41–50, 2021. [2](#)
- [52] Alexander Richard, Michael Zollhöfer, Yandong Wen, Fernando de la Torre, and Yaser Sheikh. MeshTalk: 3D face animation from speech using cross-modality disentanglement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1173–1182, October 2021. [1](#), [2](#), [6](#)
- [53] Elad Richardson, Matan Sela, Roy Or-El, and Ron Kimmel. Learning detailed face reconstruction from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [2](#)
- [54] Joseph Roth, Yiyang Tong, and Xiaoming Liu. Adaptive 3D face reconstruction from unconstrained photo collections. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [2](#)
- [55] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael Black. Learning to regress 3D face shape and expression from an image without 3d supervision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#), [7](#)
- [56] 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*, pages 53–70. Springer, 2020. [2](#)
- [57] Fuhao Shi, Hsiang-Tao Wu, Xin Tong, and Jinxiang Chai. Automatic acquisition of high-fidelity facial performances using monocular videos. *ACM Transactions on Graphics (TOG)*, 33(6):1–13, 2014. [2](#)
- [58] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7462–7473. Curran Associates, Inc., 2020. [2](#), [3](#)
- [59] Sarah Taylor, Taehwan Kim, Yisong Yue, Moshe Mahler, James Krahe, Anastasio Garcia Rodriguez, Jessica Hodgins, and Iain Matthews. A deep learning approach for generalized speech animation. *ACM Transactions on Graphics (TOG)*, 36(4):1–11, 2017. [2](#)
- [60] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#), [2](#)
- [61] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [62] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. In *European conference on computer vision*, pages 716–731. Springer, 2020. [3](#)
- [63] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gerard Medioni. Regressing robust and discriminative 3D morphable models with a very deep neural network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [2](#)
- [64] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard Medioni. Extreme 3D face reconstruction: Seeing through occlusions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. [3](#), [4](#)
- [66] Kenny T. R. Voo, Liming Jiang, and Chen Change Loy. Delving into high-quality synthetic face occlusion segmentation datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2022. [4](#)
- [67] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot...for now. In *CVPR*, 2020. [8](#)
- [68] Yi Wang, Xin Tao, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. Image inpainting via generative multi-column convolutional neural networks. *Advances in neural information processing systems*, 31, 2018. [4](#)
- [69] Chenglei Wu, Takaaki Shiratori, and Yaser Sheikh. Deep incremental learning for efficient high-fidelity face tracking. *ACM Transactions on Graphics*, 37(6), Dec. 2018. [2](#)
- [70] Cheng-hsin Wu, Ningyuan Zheng, Scott Ardisson, Rohan Bali, Danielle Belko, Eric Brockmeyer, Lucas Evans, Timothy Godisart, Hyowon Ha, Alexander Hypes, Taylor Koska, Steven Krenn, Stephen Lombardi, Xiaomin Luo, Kevyn McPhail, Laura Millerschoen, Michal Perdoch, Mark Pitts, Alexander Richard, Jason Saragih, Junko Saragih, Takaaki Shiratori, Tomas Simon, Matt Stewart, Autumn Trimble, Xinshuo Weng, David Whitewolf, Chenglei Wu, Shou-I Yu, and Yaser Sheikh. Multiface: A dataset for neural face rendering. In *arXiv*, 2022. [5](#)
- [71] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [2](#)
- [72] Ning Yu, Vladislav Skripniuk, Sahar Abdelnabi, and Mario Fritz. Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In *ICCV*, 2021. [8](#)
- [73] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C. Bühler, Xu Chen, Michael J. Black, and Otmar Hilliges. I M Avatar: Implicit morphable head avatars from videos. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#)
- [74] Yang Zhou, Zhan Xu, Chris Landreth, Evangelos Kalogerakis, Subhansu Maji, and Karan Singh. Visemenet: Audio-

driven animator-centric speech animation. *ACM Transactions on Graphics (TOG)*, 37(4):1–10, 2018. [2](#)

- [75] 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*, 37:523–550, 2018. [2](#)