

## HeadGAN: One-shot Neural Head Synthesis and Editing

Michail Christos Doukas<sup>1,2</sup>, Stefanos Zafeiriou<sup>1,2</sup>, Viktoriia Sharmanska<sup>1,3</sup>

<sup>1</sup>Imperial College London, UK   <sup>2</sup>Huawei Technologies, London, UK   <sup>3</sup>University of Sussex, UK

{michail-christos.doukas16, s.zafeiriou, sharmanska.v}@imperial.ac.uk



Figure 1: Our proposed *HeadGAN* method performs reenactment (b), by fully transferring the facial expressions and head pose from a driving frame to a reference image. When the driving and reference identities coincide (c), it can be used for facial video compression and reconstruction. In addition, *HeadGAN* can be applied to facial expression editing (d), novel view synthesis (e) and face frontalisation (f). Project’s page: <https://michaildoukas.github.io/HeadGAN/>

### Abstract

Recent attempts to solve the problem of head reenactment using a single reference image have shown promising results. However, most of them either perform poorly in terms of photo-realism, or fail to meet the identity preservation problem, or do not fully transfer the driving pose and expression. We propose *HeadGAN*, a novel system that conditions synthesis on 3D face representations, which can be extracted from any driving video and adapted to the facial geometry of any reference image, disentangling identity from expression. We further improve mouth movements, by utilising audio features as a complementary input. The 3D face representation enables *HeadGAN* to be further used as an efficient method for compression and reconstruction and a tool for expression and pose editing.

### 1. Introduction

Visual data synthesis [44, 43], including talking head animation [45, 49, 48, 17, 33, 34, 23] are particularly exciting and thriving research areas, with countless applications in editing, games, social media, VR, teleconference and virtual assistance. Over the past years, solutions were mainly given by the graphics community. For instance, *Face2Face* [40] method performs face reenactment, by recovering facial expressions from a driving video and overwriting them to the source frames. Some recent learning-based approaches [23, 25, 13] have sought to solve the problem of full head reenactment, which aims to transfer not only the expression, but also the pose, from a driving person to the source identity. The shortcoming of such methods is their dependence on long video footage of the source, as they train person-specific models. At the same time, various

methods have been proposed for reenacting human heads under a few-shot setting [45, 43, 48, 17, 34, 14], where only a limited number of reference images are available, even a single one. Most state-of-the-art approaches use facial key-points to guide synthesis [49, 48, 43, 34], which usually leads to *identity preservation problems* during reenactment, as key-points encode appearance information. The problem becomes more prominent when the head geometry of the source differs from that of the person in the driving video.

In this paper we propose *HeadGAN*, a novel one-shot GAN-based method for head animation and editing. We take a different approach from most existing few-shot methods and use a *3D face representation* similar to PNCC [52] to condition synthesis. We capitalise on prior knowledge of expression and identity disentanglement, enclosed within *3D Morphable Models (3DMMs)* [3, 5, 4, 6]. Our decision to model faces with 3DMMs enables *HeadGAN* to operate as: 1) a real-time *reenactment system* operating at  $\sim 20$  fps, 2) an efficient method for facial video *compression and reconstruction*, 3) a facial *expression editing method*, 4) a *novel view synthesis system*, including face *frontalisation*. Fig. 1 illustrates the tasks supported by our method. Apart from 3D faces, we optionally condition the generative process on speech features coming from the audio signal, enabling our method to perform accurate mouth synthesis, as suggested by our automated lipreading experiment.

We perform extensive comparisons with state-of-the-art methods [45, 43, 48, 34, 31, 51] and report superior image quality and performance, in terms of standard GAN metrics [19, 41], on the tasks of reconstruction, reenactment and frontalisation, even when compared to models [48] trained on the larger VoxCeleb2 [9] dataset. Lastly, we conduct an ablation study in order to demonstrate the contribution of each component of our system.

## 2. Related Work

**Model-free methods for face synthesis.** *X2Face* [45] is among the earliest learning-based methods for animating human heads that does not rely on any prior knowledge of faces. In some cases their warping operation causes unnatural head deformations, leading to poor photo-realism. *MonkeyNet* [33] is a more recent deep learning framework that proposes to infer motion via key-point detection from the driving video. Then, the appearance extracted from the reference image along with motion information are used to generate the output. In the follow-up work, *First Order Motion Model (FOMM)* [34] significantly improves the results of single image animation. *FOMM* uses relative key-point locations in order to preserve the identity of the source, which requires the object in the first frame of the driving video to be in the same pose with the one in the source image. As such assumption is not always met, the head pose of generated samples is not guaranteed to follow the driver.

**Landmark-based face modeling and generation.** *Bringing Portraits to Life* [1] is one of the initial attempts to animate still images. 2D warps are applied on the source image in order to imitate the facial transformations in the driving video. It shows promising results when the source head pose is close to the one appearing in the target image and only a small deformation is required. *Warp-Guided GANs* [14] is a more recent work that uses 2D facial landmarks and 2D warps to animate an image. It requires a photo captured in a frontal pose with neutral expression. Much research on head animation assumes a few-shot setting, where a small number of reference images are available. Zakharov *et al.* [49] extracts identity related embeddings from the reference images and injects them into the generator through adaptive instance normalisation layers (AdaIN) [20]. Their image-based method performs best after fine-tuning on the new identity. *Bi-layer Neural Avatars* [48] is another one-shot model that capitalises on SPADE [30], while operating in real-time speeds during inference. Using SPADE layers [30] for the adaptation of the generative process to the appearance of the source has been proposed earlier, as part of the *few-shot vid2vid* model [43], which is a video-based method extending *vid2vid* [44]. Most aforementioned methods are not able to address the identity preservation problem in reenactment, since facial landmarks allow identity related information from the driver to be transferred into the generated samples. *MarioNETte* [17] tries to solve this problem, by proposing a method for landmark transformation that adapts the driving landmarks to the reference head shape.

**Head animation assisted by 3D faces.** 3DMMs [3, 5, 4, 6] have been proven to be very effective for modeling human faces and have been widely used to drive face synthesis [40, 23, 35, 39]. Fitting 3DMMs, enables recovering accurate pose and expressions from the target frames, as well as identity related parameters from the reference image(s). Then, the rendered 3D faces are used to condition neural networks, which complete the texture and fill in the areas of missing information (hair, body, background, etc.). *Deep video portraits (DVP)* [23] and *Head2Head* [25] are examples of such reenactment systems driven by 3D information. Both methods train *person specific* models, using a long video footage of the source subject. On the contrary, our proposed approach is *person generic*, *i.e.* it can perform video synthesis for any unseen person, using a single reference image. Other methods such as *StyleRig* [38] and *GIF* [15] use 3DMMs to control *StyleGAN* and *StyleGAN2* [21, 22], but fail to preserve parts of the scene that are not explained by the face models, such as hair, background.

**Audio-driven head synthesis.** Apart from the video-driven techniques discussed above, there exists an extensive body of literature focusing on audio-driven talking face synthesis [37, 8, 35, 7, 42, 39]. Much different from these meth-

ods, our system can optionally use audio signals to enhance speech quality and realism within the mouth area, while the pose and expressions are guided by the target video footage.

### 3. Methodology

#### 3.1. 3D Face Representation

In order to accurately transfer the expressions of the driving person while preserving the facial geometry of the source identity, we take advantage of prior knowledge of human faces, contained within 3DMMs [3, 5, 4, 6]. Given a driving video of  $T$  frames,  $\mathbf{y}_{1:T} = \{\mathbf{y}_t \mid t = 1, \dots, T\}$ , the 3DMM fitting stage produces a sequence of *camera parameters*  $\mathbf{c}_{1:T}$  and *shape parameters*  $\mathbf{p}_{1:T}$ , with  $\mathbf{p}_t = [\mathbf{p}_t^{id\top}; \mathbf{p}_t^{exp\top}]^\top$ . That is, for each frame  $t$ , we obtain two types of shape parameters: a) identity related parameters  $\mathbf{p}_t^{id} \in \mathbb{R}^{n_{id}}$ , encoding facial geometry and b) expression parameters  $\mathbf{p}_t^{exp} \in \mathbb{R}^{n_{exp}}$ , representing facial deformations. This enables disentangling facial shape attributes that depend on identity from shape deformations caused by motion. We recover very accurate facial expressions, as our 3DMM fitting stage relies on a dense set of 3D points (around 1K). These points are regressed from frames with *RetinaFace* [12], which is pre-trained on WIDER FACE dataset [47]. Additionally, given a reference image of the source identity  $\mathbf{y}_{ref}$ , we perform 3DMM fitting to obtain the source’s shape parameters  $\mathbf{p}_{ref}^{id}$ ,  $\mathbf{p}_{ref}^{exp}$  and camera parameters  $\mathbf{c}_{ref}$ . For details on the 3DMM fitting algorithm please refer to the supplementary material.

Next, for each frame  $t$ , we compute the 3D facial shape (3D mesh)  $\mathbf{s}_t = [x_1, y_1, z_1, \dots, x_N, y_N, z_N]^\top \in \mathbb{R}^{3N}$ , as

$$\mathbf{s}_t = \bar{\mathbf{x}} + \mathbf{U}^{id} \mathbf{p}_{ref}^{id} + \mathbf{U}^{exp} \mathbf{p}_t^{exp}. \quad (1)$$

Here  $\bar{\mathbf{x}} \in \mathbb{R}^{3N}$  is the mean shape,  $\mathbf{U}_{id}$  is the identity orthonormal basis and  $\mathbf{U}_{exp}$  is the expression orthonormal basis of LSFM morphable model [6]. By construction, this 3D shape  $\mathbf{s}_t$  reflects the facial structure  $\mathbf{p}_{ref}^{id}$  of the source with the facial expressions  $\mathbf{p}_t^{exp}$  of the driving identity. In this way, we address the source identity preservation problem. Finally, we render a 3D face representation  $\mathbf{x}_t = \mathcal{R}(\mathbf{s}_t, \mathbf{c}_t)$ , using the 3D shape  $\mathbf{s}_t$  and camera parameters  $\mathbf{c}_t$ . This is an RGB image similar to PNCC [52], as shown in Fig. 2. We also render  $\mathbf{x}_{ref}$ , which is the 3D face recovered from the reference image  $\mathbf{y}_{ref}$ , using camera parameters  $\mathbf{c}_{ref}$  and shape  $\mathbf{s}_{ref}$  that is obtained from  $\mathbf{p}_{ref}^{id}$  and  $\mathbf{p}_{ref}^{exp}$  using Eq. 1. In Sec. 4.1 we discuss 3D Face Rendering in more detail.

Summarising, given a driving video  $\mathbf{y}_{1:T}$  and a source image  $\mathbf{y}_{ref}$ , the data pre-processing pipeline recovers a sequence of images  $\mathbf{x}_{1:T}$ , which depict the 3D face extracted from the driver and adapted to the facial geometry of the source, as well as the 3D face of the reference image  $\mathbf{x}_{ref}$ . These facial representations are used to condition image synthesis with *HeadGAN*’s Generator network.

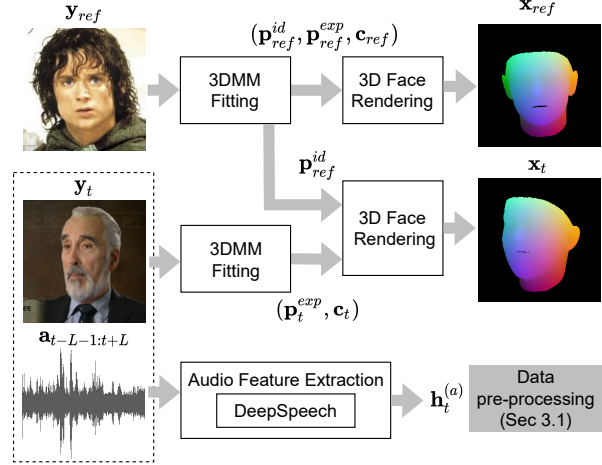


Figure 2: Data pre-processing stage. We recover and render the 3D face of the reference image  $\mathbf{y}_{ref}$ , as well as the driving frame  $\mathbf{y}_t$ , after adapting the identity parameters.

#### 3.2. Audio features

As opposed to previous one-shot head reenactment systems, our method takes advantage of the driving audio stream and its correlation with facial and mouth movements. We split the audio signal into  $T$  parts  $\mathbf{a}_{1:T}$ , where each part  $\mathbf{a}_t$  is aligned and corresponds to frame  $\mathbf{y}_t$  of the driving video with length  $T$ . Then, we apply audio feature extraction to a window of  $2L$  audio parts  $\mathbf{a}_{t-L-1:t+L} = \{\mathbf{a}_{t-L-1}, \dots, \mathbf{a}_t, \dots, \mathbf{a}_{t+L}\}$ , centred around frame  $t$ , to obtain a feature vector  $\mathbf{h}_t^{(a)}$ , which contains information from the past and future time steps. We employ [16] for the extraction of low level features, such as MFCCs, signal energy and entropy, which yields a feature vector  $\mathbf{h}_t^{(aL)} \in \mathbb{R}^{84}$ . Then, we use DeepSpeech [18] for the extraction of character level logits from each part  $\mathbf{a}_{t'} \in \mathbf{a}_{t-L-1:t+L}$ . This results in  $2L$  logits, which after concatenation gives a feature vector  $\mathbf{h}_t^{(aH)} \in \mathbb{R}^{2L \cdot 27}$ . Our final audio feature vector is given as  $\mathbf{h}_t^{(a)} = [\mathbf{h}_t^{(aL)\top}; \mathbf{h}_t^{(aH)\top}]^\top \in \mathbb{R}^{300}$ , for  $L = 4$ .

#### 3.3. HeadGAN Framework

Our GAN-based head reenactment system is equipped with a Generator driven by two modalities: 1) the *3D face representation* extracted from the driving video and the reference image, 2) optionally the *audio features* coming from the driver. Given  $\mathbf{x}_{t-k:t}$ , the driving 3D face representation from frame  $t$ , concatenated channel-wise with the 3D faces coming from the past  $k = 2$  frames, the reference image  $\mathbf{y}_{ref}$  with the corresponding 3D face representation  $\mathbf{x}_{ref}$  and the audio feature vector  $\mathbf{h}_t^{(a)}$ , our Generator hallucinates a photo-realistic image, given as

$$\tilde{\mathbf{y}}_t = G(\mathbf{x}_{t-k:t}, \mathbf{y}_{ref}, \mathbf{x}_{ref}, \mathbf{h}_t^{(a)}; \theta_G). \quad (2)$$

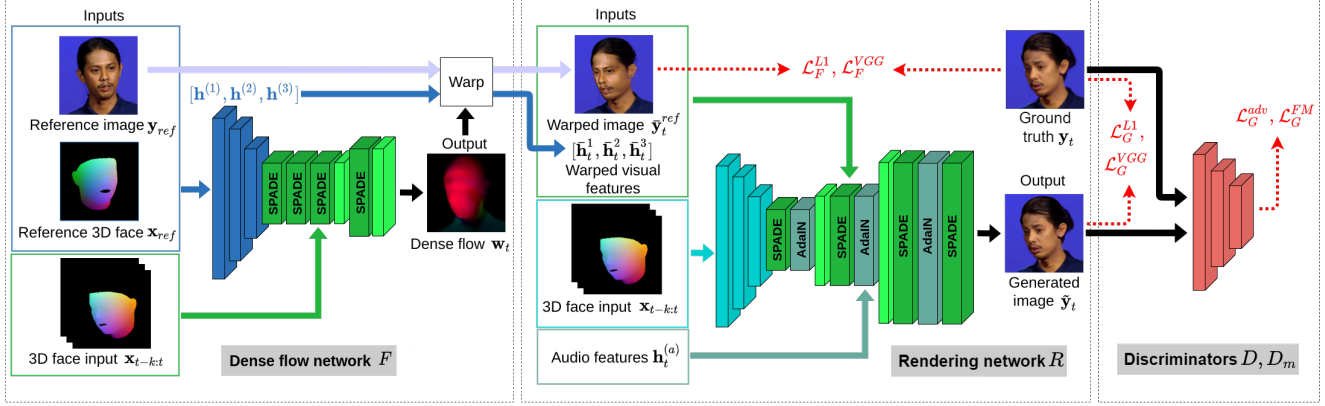


Figure 3: Overview of *HeadGAN*. The dense flow network  $F$  computes a flow field for warping the reference image and features, according to the 3D face input. Then, the rendering network  $R$  uses this visual information along with the audio features in order to translate the 3D face input into a photo-realistic image of the source.

Conditioning synthesis on the spatio-temporal volume  $\mathbf{x}_{t-k:t}$  helps to achieve temporal coherence across frames. The reference image  $\mathbf{y}_{ref}$  provides information on the texture and appearance of the source person, while audio features enhance the generative ability of  $G$  across the face, and mainly the mouth area. In more detail, the Generator consists of two sub-networks: a *dense flow network*  $F$  and a *rendering network*  $R$ . For an overview of the Generator  $G$ , please refer to Fig. 3.

**Dense flow network  $F$ .** Our rendering network  $R$  relies on high quality visual features that reflect the appearance of the source identity. Nonetheless, we observed that simply using an encoder to extract such features from the reference image  $\mathbf{y}_{ref}$ , does not capitalise well on the potential of the rendering network’s architecture. It has been proved more meaningful to align the visual feature maps with the desired head pose, which is reflected in the 3D face representation  $\mathbf{x}_t$ , coming from the driving video. With this in mind, we propose a dense flow network, which learns a flow  $\mathbf{w}_t$  that can be used to warp visual features. For that, we pass the concatenation of the reference image and its corresponding 3D face ( $\mathbf{y}_{ref}, \mathbf{x}_{ref}$ ) through an encoder, for the extraction of visual feature maps in three spatial scales  $\mathbf{h}^{(1)}, \mathbf{h}^{(2)}, \mathbf{h}^{(3)}$ , which represent the appearance of the source identity. Then, a decoder predicts the flow  $\mathbf{w}_t$ , guided by the driving 3D face representation  $\mathbf{x}_{t-k:t}$ , which is injected into  $F$  through SPADE blocks [30]. Ideally, when applied on the reference image  $\mathbf{y}_{ref}$ , this dense flow should yield a warped image of the source person, with the same head pose and expression, as shown in the driving 3D face representation  $\mathbf{x}_t$ . By applying the flow field on each visual feature map, we obtain the warped visual features  $\bar{\mathbf{h}}_t^{(1)}, \bar{\mathbf{h}}_t^{(2)}, \bar{\mathbf{h}}_t^{(3)}$  and the warped reference image  $\tilde{\mathbf{y}}_t^{ref}$ , all of which depend on the driving head pose at frame  $t$ .

**Rendering network  $R$ .** At the core of our Generator, the rendering network aims to translate the 3D face representation  $\mathbf{x}_{t-k:t}$  into a photo-realistic image  $\tilde{\mathbf{y}}_t$  of the source. This is achieved with the assistance of high quality audio features  $\mathbf{h}_t^{(a)}$  and visual feature maps  $\bar{\mathbf{h}}_t^{(1)}, \bar{\mathbf{h}}_t^{(2)}, \bar{\mathbf{h}}_t^{(3)}$ . First, an encoder receives  $\mathbf{x}_{t-k:t}$  as input and applies a sequence of convolutional layers with down-sampling. Then, a decoder consisting of alternating SPADE [30] and AdaIN [20] layers generates the desired frame  $\tilde{\mathbf{y}}_t$ . These adaptive normalisation layers enable injecting 2D visual feature maps into the rendering network through SPADE blocks, as well as 1D audio features through AdaIN blocks. As opposed to the original work on SPADE [30], where the conditional input of all SPADE layers is the same segmentation map down-sampled to match the spatial size of each layer, we capitalise on visual feature maps of multiple spatial scales  $\bar{\mathbf{h}}_t^{(1)}, \bar{\mathbf{h}}_t^{(2)}, \bar{\mathbf{h}}_t^{(3)}, \tilde{\mathbf{y}}_t^{ref}$  as modulation inputs to SPADE blocks. On the contrary, we pass the same audio feature vector  $\mathbf{h}_t^{(a)}$  to AdaIN blocks of all spatial scales. The decoder is further equipped with PixelShuffle layers [32] for up-sampling, which contribute to the quality of generated samples.

**Discriminators  $D$  and  $D_m$ .** The image Discriminator receives a synthetic pair  $(\mathbf{x}_t, \tilde{\mathbf{y}}_t)$ , or a real one  $(\mathbf{x}_t, \mathbf{y}_t)$  and learns to distinguish between them. We use a second Discriminator  $D_m$ , which focuses on the mouth region. Apart from the real  $\mathbf{y}_t^m$  or generated  $\tilde{\mathbf{y}}_t^m$  cropped mouth area, this network is conditioned on the audio feature vector  $\mathbf{h}_t^{(a)}$ , which is spatially replicated and then concatenated to the cropped images channel-wise.

**Training Objective.** Networks  $F$  and  $R$  which constitute the Generator are optimised jointly. We train *HeadGAN* on reconstruction, by applying perceptual and pixel losses  $\mathcal{L}_F^{VGG}, \mathcal{L}_G^{VGG}$  and  $\mathcal{L}_F^{L1}, \mathcal{L}_G^{L1}$ , both on the warped and generated images, as seen in Fig. 3 (red arrows). A GAN Hinge

loss  $\mathcal{L}_G^{adv}$  [27] along with a feature matching loss  $\mathcal{L}_G^{FM}$  [46] further increase the photo-realism of results. An extended discussion on the objective functions and networks architecture can be found in the supplementary material.

### 3.4. Advantages of 3D face modeling

The *semantic information* enclosed within the 3D face representation allows our dense flow network  $F$  to learn a precise flow of the facial region, as it provides a dense correspondence of facial points, between the reference and driving images. Compared to scene flow [28, 26] that can be obtained from 3D meshes, our flow field hallucinated by  $F$  exists in areas where 3D representation is missing, such as hair and upper body, where warping is equally important. Furthermore, as 3DMMs allow to disentangle identity from expression, our choice to condition the rendering network  $R$  on a 3D face extracted from the driving person and adapted to the identity characteristics of the source, enables *HeadGAN* to tackle the identity preservation problem on the task of reenactment.

Lastly, modeling faces with 3DMMs makes *HeadGAN* very efficient for video conferencing applications. It enables a "sender" to efficiently compress a driving frame  $\mathbf{y}_t$ , in the form of expression parameters  $\mathbf{p}_t^{exp} \in \mathbb{R}^{28}$  and camera parameters  $\mathbf{c}_t \in \mathbb{R}^7$ , giving a total of 35 floating point values. Then, these parameters can be used by a "receiver" to render the 3D face and use the Generator to reconstruct the driving frame. A single reference image needs to be sent once in the beginning of the session.

## 4. Experiments

### 4.1. Implementation details

**3D Face Rendering.** Given a set of camera parameters  $\mathbf{c}$  and a 3D facial shape  $\mathbf{s} \in \mathbb{R}^{3N}$  (see Eq. 1), we rasterise the 3D mesh and produce a visibility mask  $I \in \mathbb{R}^{H \times W}$  in the image plane. Each spatial location of  $I$  stores the index of the corresponding visible triangle on the 3D face seen from this pixel. Then, we use the mean shape  $\bar{\mathbf{x}}$  of the 3DMM, in order to find the normalised x-y-z coordinates of the center of each visible triangle. In this way we obtain a 3D face representation  $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ , where each pixel contains

three coordinates. These values can be interpreted as colors and therefore texture of the 3D face [52].

**Dataset and Training.** We train and evaluate *HeadGAN* on VoxCeleb [29] dataset, which contains over 100,000 videos of 1,251 identities, at  $256 \times 256$  resolution. We maintain the original train and test split. As a pre-processing step, we compute a 3D face image for each video frame in the dataset and extract per-frame audio feature vectors. During training, we perform self-reenactment, as we randomly sample the reference image from the target video. This provides access to ground truth data, enabling us to design reconstruction loss terms to train the Generator. For the optimisation of *HeadGAN*, we use the ADAM [24] with  $\beta_1 = 0.5$ ,  $\beta_2 = 0.999$  and learning rate  $\eta = 0.0002$ , both for the Generator and Discriminator.

### 4.2. Comparison with baselines

**Reconstruction (self-reenactment).** First, we compare our approach with four state-of-the-art methods on the problem of self-reenactment, where the reference identity coincides with the driving one. Here, the task for *HeadGAN* is to reconstruct the driving video from the 3D face representation sequence, using a single reference image to access appearance information. We perform both qualitative and quantitative comparisons with *X2Face* [45], *few-shot vid2vid* [43], *Bi-layer Neural Avatars* [48] and *First Order Motion Model* [34]. For [45] and [34], we use the pre-trained models provided by the authors, trained on VoxCeleb. We trained [43] from scratch, since no checkpoint was available. Lastly, we used the model provided by the authors of [48], trained on the larger VoxCeleb2\* [9].

For the numerical evaluation of reconstruction we use L1 distance between the generated and ground truth frames, as well as Peak signal-to-noise ratio (PSNR) and Learned Perceptual Image Patch Similarity (LPIPS) [50]. We access the realism of frames with Fréchet Inception Distance (FID) [19] and Fréchet Video Distance (FVD) [41] metrics. We use Cosine Similarity (CSIM) [11] to measure identity preservation. The results are presented in Table 1 and reveal that *HeadGAN* outperforms all four baselines by a noteworthy margin, in every single metric.

Method	Reconstruction						Reenactment			
	L1 ↓	PSNR ↑	LPIPS ↓	FID ↓	FVD ↓	CSIM ↑	FID ↓	CSIM ↑	ARD ↓	AU-H ↓
<i>X2Face</i> [45]	13.49	20.69	0.260	130.2	697	0.600	122.1	0.520	4.39	0.346
<i>fs-vid2vid</i> [43]	17.15	18.52	0.197	62.8	471	0.542	-	-	-	-
<i>Bi-layer*</i> [48]	12.18	20.19	0.152	92.2	394	0.590	172.8	0.563	<b>1.01</b>	<b>0.296</b>
<i>FOMM</i> [34]	12.34	20.93	0.153	64.9	338	0.754	63.7	<b>0.765</b>	12.53	0.400
<i>HeadGAN</i>	<b>11.32</b>	<b>21.46</b>	<b>0.112</b>	<b>36.1</b>	<b>254</b>	<b>0.807</b>	<b>58.0</b>	0.688	1.35	0.326

Table 1: Quantitative results on the tasks of reconstruction and reenactment for VoxCeleb [29] test set.



Figure 4: Qualitative comparison with baselines, on the task of reconstruction (self-reenactment).

The samples illustrated in Fig. 4 show that our method generates far more realistic images than the baselines with better preserved appearance traits. We urge the reader visually inspect the results in the supplementary video.

**Reenactment.** The objective of reenactment is to fully transfer the head pose and facial expressions of the target sequence to the person shown in the source image, while preserving the latter identity, as the driving and reference subjects are now different. To that end, we selected 15 random (video, image) pairs from VoxCeleb test set and performed reenactment, generating around 6K frames in total with each method. Apart from image quality (FID) and identity preservation (CSIM), we further evaluate the pose and expression transferability of systems, with Average Rotation Distance (ARD) in degrees and Action Units Hamming distance (AU-H) [2], respectively. Details on metrics can be found in the supplementary material. As can be seen in Table 1, *HeadGAN* creates superior samples in terms of visual quality. *Bi-layer Neural Avatars* [48] performs slightly better on the task of facial expression transfer, which could be attributed to the fact that it was trained on the larger VoxCeleb2 [9]. However, it performs poorly on identity preservation, as it conditions synthesis on facial landmarks, which unavoidably pass on identity related information from the driving subject. On the other hand, *FOMM* [34] uses relative key-point locations to account for the identity preser-

vation problem that seems to increase CSIM. Nonetheless, this comes at the expense of pose transfer, as the model requires the face in the first frame of the driving video to have the same pose with the reference face, which is very rarely the case, also confirmed by the large ARD. When *FOMM* uses absolute key-point locations instead, in order to accurately transfer pose, the identity preservation problem becomes apparent and CSIM drops to 0.587. This behaviour can be observed visually in Fig. 5, where the head geometry of the driver is reflected in the generated samples of *FOMM-abs*. Different than the baselines, *HeadGAN* performs well on all three requirements of successful reenactment (pose, expression transfer and identity preservation). Here we note that we have omitted comparisons with *MarioNETte* [17] and *Warp-guided GANs* [14] since the source codes are not publicly available.

**Frontalisation.** Our choice to use a 3D face representation to condition synthesis allows us to set the desired head pose manually, without the need of driving frames. By re-setting the camera parameters to frontal, we are able to generate a frontal view of the reference image. We compare *HeadGAN* on the task of frontalisation with *pixel2style2pixel (pSp)* [31] and *Rotate-and-Render (RaR)* [51]. For that, we randomly select one frame from each video of VoxCeleb test split and frontalise it. Since there is no audio input for this task, we trained a varia-



Figure 5: Qualitative comparison with baselines, on the task of reenactment. *FOMM-Rel* refers to relative key-point coordinates strategy and *FOMM-Abs* to absolute key-point coordinates, used in *First Order Motion Model* [34].

tion of *HeadGAN* without AdaIN layers for audio features. Then, we measure the photo-realism of generated samples with FID, the identity preservation with CSIM and the Average Rotation Error (ARE) as the deviation from the frontal pose, in degrees. We define frontal pose as zero Euler angles in camera rotation parameters coming from 3DMM fitting. Therefore, ARE is only used for reference here, as a sanity check. The results are displayed in Table 2. *HeadGAN* performs equally well on CSIM with *RaR* and surpasses baselines on image quality and frontalisation accuracy. Please see Fig. 6 for a visual comparison.

Method	FID ↓	CSIM ↑	ARE ↓
<i>psp</i> [31]	147.8	0.130	2.66
<i>RaR</i> [51]	88.4	0.753	2.65
<i>HeadGAN</i>	<b>30.1</b>	<b>0.766</b>	<b>0.76</b>

Table 2: Quantitative results on frontalisation.




Figure 6: Qualitative comparison on frontalisation.

### 4.3. Image Expression and Pose Editing

Our model can be further used as an image editing tool. Given a source image  $\mathbf{y}_{ref}$  and its shape and camera parameters  $\mathbf{p}_{ref}^{id}$ ,  $\mathbf{p}_{ref}^{exp}$  and  $\mathbf{c}_{ref}$ , first we render the corresponding 3D face representation  $\mathbf{x}_{ref}$ . Then, we re-adjust the expression or camera parameters manually and render a pseudo-driving 3D face  $\mathbf{x}_t$ . We pass  $\mathbf{x}_t$ ,  $\mathbf{y}_{ref}$ ,  $\mathbf{x}_{ref}$  through the Generator to obtain a synthetic image with a novel expression or pose, reflecting the adjusted parameters. In Fig. 7 we show the results after manipulating the first three principal components (p. c.) of expression and the camera angles.

### 4.4. Ablation Study

We conducted an ablation study in order to assess 1) the significance of the dense flow network  $F$ , 2) the advantage of the 3D face representation compared to sketches of 2D

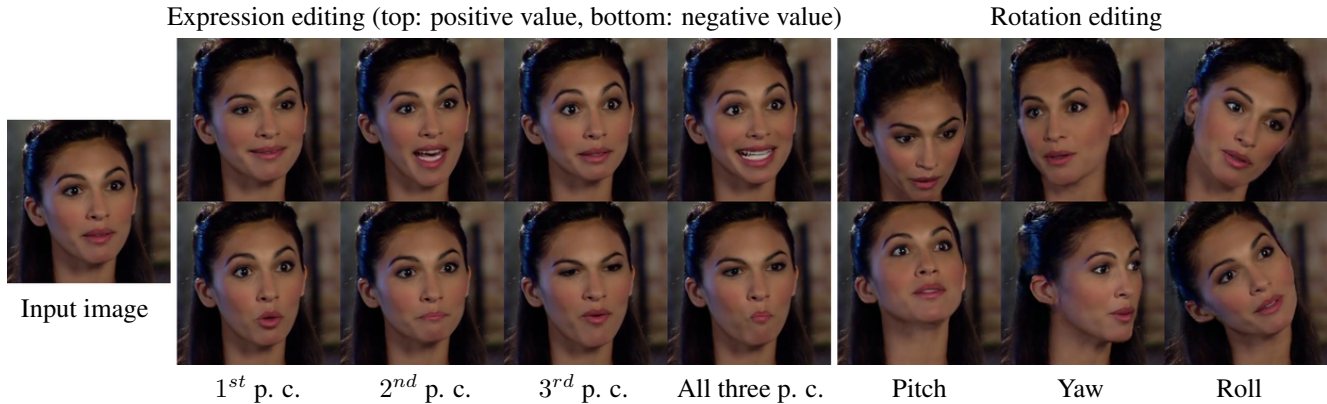


Figure 7: Expression and Pose Editing. Please note that we use our model’s variation without audio (AdaIN) for this task.

landmarks, that is the input used by [43] and [48], 3) the contribution of audio modality. In order to evaluate the importance of  $F$ , we removed its decoding layers. We kept its encoder, which extracts appearance feature maps from the reference image. Then, instead of warping these features maps and reference image, we passed them directly to the SPADE layers of the rendering network  $R$ . In addition, we implemented a *HeadGAN* variation with landmarks, by conditioning synthesis on sketch images, drawn by connecting 2D landmarks with edges [43]. As can be seen in Table 3, the full model outperforms all variations. In terms of scores, we observe that flow network  $F$  is an essential component of our system. The use of a 3D face representation (instead of landmarks) alleviates the identity preservation problem and can be better understood visually, in Fig. 8b.

Method	FID ↓	FVD ↓	CSIM ↑
<i>HeadGAN</i> w/o network $F$	63.3	473	0.307
<i>HeadGAN</i> w/ landmarks	55.7	371	0.699
<i>HeadGAN</i> w/o audio input	55.1	356	0.687
<i>HeadGAN</i>	<b>50.9</b>	<b>334</b>	<b>0.716</b>

Table 3: Ablation study numerical results. Please note that we trained models for half epochs in this experiment.

**Lipreading experiment.** We further evaluate the contribution of audio input to our system quantitatively, by employing an external lipreading network to classify synthetic videos. To that end, we chose 25 word classes of BBC dataset [10] and trained a lipreading classifier [36] on the default training split. After that, we reconstructed the test split of BBC dataset, using a random frame from the video as reference. We report a lipreading accuracy of 97% on real test samples, 82% on samples generated using the the full model and 73% on synthetic data produced by the variation without considering audio input. These results suggest that



(a) Significance of dense flow network  $F$  in image quality.



(b) The identity preservation problem becomes prominent when conditioning on facial landmarks, instead of the 3D face representation.

Figure 8: The importance of *HeadGAN* components.

the audio modality contributes largely on the generation of more plausible lip movements.

## 5. Conclusion

We presented *HeadGAN*, a novel one-shot method for animating heads, driven by 3D facial data and audio features. Compared to SOTA methods, our framework exhibits superior reenactment performance and higher photorealism. Our method can be further used for reconstruction, pose and facial expression editing, as well as frontalisation.



## References

- [1] Hadar Averbuch-Elor, Daniel Cohen-Or, Johannes Kopf, and Michael F. Cohen. Bringing portraits to life. *ACM Transactions on Graphics (Proceeding of SIGGRAPH Asia 2017)*, 36(6):196, 2017. [2](#)
- [2] T. Baltrušaitis, M. Mahmoud, and P. Robinson. Cross-dataset learning and person-specific normalisation for automatic action unit detection. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 06, pages 1–6, 2015. [6](#)
- [3] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '99*, page 187–194, USA, 1999. ACM Press/Addison-Wesley Publishing Co. [2](#), [3](#)
- [4] James Booth, Anastasios Roussos, Allan Ponniah, David Dunaway, and Stefanos Zafeiriou. Large scale 3d morphable models. *Int. J. Comput. Vision*, 126(2–4):233–254, Apr. 2018. [2](#), [3](#)
- [5] J. Booth, A. Roussos, E. Ververas, E. Antonakos, S. Ploumpis, Y. Panagakis, and S. Zafeiriou. 3d reconstruction of “in-the-wild” faces in images and videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(11):2638–2652, 2018. [2](#), [3](#)
- [6] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3d morphable model learnt from 10,000 faces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [2](#), [3](#)
- [7] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7832–7841, 2019. [2](#)
- [8] Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. You said that? In *British Machine Vision Conference*, 2017. [2](#)
- [9] J. S. Chung, A. Nagrani, and A. Zisserman. Voxceleb2: Deep speaker recognition. In *INTERSPEECH*, 2018. [2](#), [5](#), [6](#)
- [10] J. S. Chung and A. Zisserman. Lip reading in the wild. In *Asian Conference on Computer Vision*, 2016. [8](#)
- [11] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. [5](#)
- [12] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotzia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [3](#)
- [13] Michail Christos Doukas, Mohammad Rami Koujan, Viktoriia Sharmanska, Anastasios Roussos, and Stefanos Zafeiriou. Head2head++: Deep facial attributes re-targeting. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(1):31–43, 2021. [1](#)
- [14] Jiahao Geng, T. Shao, Youyi Zheng, Y. Weng, and K. Zhou. Warp-guided gans for single-photo facial animation. *ACM Transactions on Graphics (TOG)*, 37:1 – 12, 2018. [2](#), [6](#)
- [15] Partha Ghosh, Pravir Singh Gupta, Roy Uziel, Anurag Ranjan, Michael J. Black, and Timo Bolkart. GIF: Generative interpretable faces. In *International Conference on 3D Vision (3DV)*, pages 868–878, 2020. [2](#)
- [16] Theodoros Giannakopoulos. pyaudioanalysis: An open-source python library for audio signal analysis. *PLoS one*, 10(12), 2015. [3](#)
- [17] Sungjoo Ha, Martin Kersner, Beomsu Kim, Seokjun Seo, and Dongyoung Kim. Marionette: Few-shot face reenactment preserving identity of unseen targets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. [1](#), [2](#), [6](#)
- [18] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Ng. Deep-speech: Scaling up end-to-end speech recognition. 12 2014. [3](#)
- [19] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 6626–6637. Curran Associates, Inc., 2017. [2](#), [5](#)
- [20] X. Huang and S. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1510–1519, 2017. [2](#), [4](#)
- [21] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [22] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. [2](#)
- [23] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Nießner, Patrick Pérez, Christian Richardt, Michael Zollöfer, and Christian Theobalt. Deep video portraits. *ACM Transactions on Graphics (TOG)*, 37(4):163, 2018. [1](#), [2](#)
- [24] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. [5](#)
- [25] M. Koujan, M. Doukas, A. Roussos, and S. Zafeiriou. Head2head: Video-based neural head synthesis. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020) (FG)*, pages 319–326, Los Alamitos, CA, USA, may 2020. IEEE Computer Society. [1](#), [2](#)
- [26] M. Koujan, A. Roussos, and S. Zafeiriou. Deepfaceflow: In-the-wild dense 3d facial motion estimation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6617–6626, Los Alamitos, CA, USA, jun 2020. IEEE Computer Society. [5](#)
- [27] Jae Hyun Lim and Jong Chul Ye. Geometric gan, 2017. [5](#)
- [28] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. Flownet3d: Learning scene flow in 3d point clouds. *CVPR*, 2019. [5](#)

- [29] A. Nagrani, J. S. Chung, and A. Zisserman. Voxceleb: a large-scale speaker identification dataset. In *INTER-SPEECH*, 2017. 5
- [30] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2, 4
- [31] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. *arXiv preprint arXiv:2008.00951*, 2020. 2, 6, 7
- [32] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1874–1883, 2016. 4
- [33] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2
- [34] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *Conference on Neural Information Processing Systems (NeurIPS)*, December 2019. 1, 2, 5, 6, 7
- [35] Yang Song, Jingwen Zhu, Dawei Li, Andy Wang, and Hairong Qi. Talking face generation by conditional recurrent adversarial network. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 919–925. International Joint Conferences on Artificial Intelligence Organization, 7 2019. 2
- [36] Themis Stafylakis and Georgios Tzimiropoulos. Combining residual networks with lstms for lipreading. *CoRR*, abs/1703.04105, 2017. 8
- [37] Supasorn Suwajanakorn, Steven Seitz, and Ira Kemelmacher. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics*, 36:1–13, 07 2017. 2
- [38] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zöllhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images, cvpr 2020. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, june 2020. 2
- [39] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. *ECCV 2020*, 2020. 2
- [40] J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2016. 1, 2
- [41] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 2, 5
- [42] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with gans. *International Journal of Computer Vision*, 10 2019. 2
- [43] Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. Few-shot video-to-video synthesis. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019. 1, 2, 5, 6, 8
- [44] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 1, 2
- [45] Olivia Wiles, A. Sophia Koepke, and Andrew Zisserman. X2face: A network for controlling face generation using images, audio, and pose codes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 1, 2, 5, 6, 7
- [46] Xiangyu Xu, Deqing Sun, Jinshan Pan, Yujin Zhang, Hanspeter Pfister, and Ming-Hsuan Yang. Learning to super-resolve blurry face and text images. In *Proceedings of the IEEE international conference on computer vision*, pages 251–260, 2017. 5
- [47] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Wider face: A face detection benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [48] Egor Zakharov, Aleksei Ivakhnenko, Aliaksandra Shysheya, and Victor Lempitsky. Fast bi-layer neural synthesis of one-shot realistic head avatars. In *European Conference of Computer vision (ECCV)*, August 2020. 1, 2, 5, 6, 7, 8
- [49] E. Zakharov, Aliaksandra Shysheya, Egor Burkov, and V. Lempitsky. Few-shot adversarial learning of realistic neural talking head models. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9458–9467, 2019. 1, 2
- [50] Richard Yi Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings - 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2018*, Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pages 586–595. IEEE Computer Society, 2018. 5
- [51] Hang Zhou, Jihao Liu, Ziwei Liu, Yu Liu, and Xiaogang Wang. Rotate-and-render: Unsupervised photorealistic face rotation from single-view images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 6, 7
- [52] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li. Face alignment across large poses: A 3d solution. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 146–155, Los Alamitos, CA, USA, jun 2016. IEEE Computer Society. 2, 3, 5