

# VR Facial Animation for Immersive Telepresence Avatars

Andre Rochow, Max Schwarz, Michael Schreiber, and Sven Behnke

**Abstract**—VR Facial Animation is necessary in applications requiring clear view of the face, even though a VR headset is worn. In our case, we aim to animate the face of an operator who is controlling our robotic avatar system. We propose a real-time capable pipeline with very fast adaptation for specific operators. In a quick enrollment step, we capture a sequence of source images from the operator without the VR headset which contain all the important operator-specific appearance information. During inference, we then use the operator keypoint information extracted from a mouth camera and two eye cameras to estimate the target expression and head pose, to which we map the appearance of a source still image. In order to enhance the mouth expression accuracy, we dynamically select an auxiliary expression frame from the captured sequence. This selection is done by learning to transform the current mouth keypoints into the source camera space, where the alignment can be determined accurately. We, furthermore, demonstrate an eye tracking pipeline that can be trained in less than a minute, a time efficient way to train the whole pipeline given a dataset that includes only complete faces, show exemplary results generated by our method, and discuss performance at the ANA Avatar XPRIZE semifinals.

## I. INTRODUCTION

Virtual Reality (VR) Facial Animation is a challenging sub-problem of facial animation, which aims to generate realistic face images. It has many applications, such as in computer games, training data generation, virtual reality, and video compression. VR allows virtual immersion in another world. When interacting virtually with a VR user, it is often desirable to perceive all their characteristics and facial expressions. However, VR headsets occlude a large area of the face, which makes direct face capture impossible.

Especially in recent years, deep learning techniques have made great progress in generative modeling [1] and rendering [2]. In this work, we demonstrate a deep learning-based approach to the VR Facial Animation problem. Our main motivation is the ANA Avatar XPRIZE Competition<sup>1</sup>, where judges interact through avatar systems developed by the participant teams. Here, a specific imposition for a facial animation system is that adaption to the operator has to be finished in one hour. This time has to be shared with the operator training time, so that available time for adaptations is even shorter. Our avatar robotic system allows an operator to directly perceive the world through the eyes of the avatar. We animate the operator on a display which is attached via a 6DoF arm to the avatar. For more information about our avatar system we refer to Schwarz *et al.* [3]. Inspired by the method of Siarohin *et al.* [4], our pipeline is trained



Fig. 1. Operator interacting through the NimbRo Avatar system with a human recipient at the ANA Avatar XPRIZE semifinals. Top: Operator at remote site. Bottom: Two different facial expressions of the operator animated with our method. See also the supplementary video for an animated version.

with the VoxCeleb [5] dataset to animate a face from a source frame, driven by keypoints from a target frame. We demonstrate how the pipeline can be adapted to generalize to the VR facial animation sub-problem. To enhance the modeling capabilities, we select a third frame, i.e. the expression frame. Using this frame, we directly embed operator-specific information, which is captured by a camera attached below the HMD (mouth camera), into the animation process. We propose an inference pipeline that dynamically and automatically chooses the optimal expression frame via keypoint-driven image retrieval. Our pipeline runs in real time and can be adapted to an unknown operator with only 15 min preprocessing.

In addition to the full real-time VR facial animation pipeline, our contributions include:

- 1) A method allowing fast adaptation to new operators,
- 2) a training regime that allows offline training on the well-known VoxCeleb [5] dataset,
- 3) a fast approximate approach to solve the alignment problem between facial image sequences captured with and without a HMD, allowing retrieval of matching frames,
- 4) an efficient eye tracking method for challenging camera perspectives, trainable in under a minute, and
- 5) methods encouraging temporal continuity at inference.

All authors are with the Autonomous Intelligent Systems group of University of Bonn, Germany; rochow@ais.uni-bonn.de

<sup>1</sup><https://www.xprize.org/prizes/avatar>

## II. RELATED WORK

*a) Facial Image Animation:* Facial Image Animation is a long-standing problem in computer graphics and aims to generate new facial images with controllable expressions. Wiles *et al.* [6] propose a method that learns the sampling coordinates from a source image to an embedded image and then from the embedded image to a driving frame using a separate network. Whereas other warping-based methods exist [4], [7]–[9], there are also indirect approaches [10]–[12] that rely on generative modeling to perform facial expression synthesis. Zakharov *et al.* [12] learn an embedding vector from few source images which then modulate a keypoint-driven generator network via Adaptive Instance Normalization (AdaIN) [13], which has demonstrated to be especially well-suited to perform style transfer. For VR facial animation, those approaches, however, must be adapted to work with partially occluded driving frames, e.g. captured from a mouth camera.

*b) Keypoint-Driven Motion Transfer:* Siarohin *et al.* [7] use deep neural networks to decouple appearance and motion information. They combine the appearance extracted from a source image and the motion derived from the driving video. Their pipeline is separated into a keypoint detector, a dense motion network, and a generator network. Based on this architecture, Siarohin *et al.* [4] encode motions based on keypoints and introduce local affine transformations that allow to model a larger family of transformations. More recently, Zhao *et al.* [8] proposed an encoder/decoder dense motion network that employs AdaIn [13] to transfer source face keypoint geometry to the encoder and driving face keypoint geometry to the decoder. They separate the dense motion network into a global branch and multiple local branches that have a limited visibility—to focus on generating a more accurate motion for the eyes and the mouth area. Furthermore, they investigate how to improve the temporal alignment of the keypoint detector as proposed by Bulat and Tzimiropoulos [14], which is also the default choice in our approach. Our method is based on Siarohin *et al.* [4], but faces a much more difficult problem where the driving frame has occlusion caused by the VR headset and is captured from a perspective different from the source frame. To address these issues, we forgo using local affine transformations [4] and rely on a keypoint-based image retrieval to simulate the lower face region more precisely.

*c) Virtual Reality Facial Animation:* VR Facial Animation is a special case of facial animation where large parts of the face are occluded by an HMD. Thies *et al.* [15] use image retrieval to obtain the most similar source views. They use blending to combine the retrieved images to a photo-realistic output. Several methods [16]–[18] render a virtual avatar based on operator-specific geometry. Lombardi *et al.* [17] propose to learn a Variational Autoencoder (VAE) with an encoder that predicts a viewpoint-independent latent variable and a decoder that can be conditioned with extrinsic and intrinsic variables controlling the camera pose, speech, identity, and gaze. Wei *et al.* [16] extend this idea and generate




Fig. 2. The modified Valve Index VR headset. We attached three additional cameras to capture the eyes and the mouth expression of the operator. The inside of the VR headset is lit using IR LEDs. We show the corresponding camera views at the bottom.

ground truth data with expression-preserving style transfer networks, which map between the VR camera domain and the avatar domain. More recently, Richard *et al.* [18] predict facial coefficients that parametrize a codec avatar with audio and gaze information only. This is especially useful if the lower face region is occluded as well, e.g. with a medical mask. In contrast to these methods, our pipeline does not assume pre-trained personalized parametric face models and can therefore be adapted to a specific operator with much less preparation time.

*d) Neural Rendering:* Very recently, Gafni *et al.* [19] proposed to learn facial image animation using dynamic neural radiance fields [2]. Whereas they generate impressive results, their pipeline must be trained for a specific operator, which is impracticable in our application due to the long training time.

## III. METHOD

### A. Avatar System and Modified VR Headset

Our robotic avatar system is described in [3]. Briefly put, it allows a human operator to immerse themselves into a remote robot and to interact and cooperate with humans at the remote site.

For visualization on the operator side, we use the Valve Index VR headset. In order to capture the mouth expression, we attach a Logitech Brio webcam below the HMD (see Fig. 2). Furthermore, we allow for eye tracking by mounting two additional cameras and IR LEDs inside the HMD.

### B. Basic Expression Mapping Pipeline

We propose a pipeline that allows to map the appearance  $\Lambda_S$  from a source image  $I_S$  to the expression  $\xi_D$  and the head pose  $\theta_D$  (motion) present in the driving frame, which may be present as keypoints only (see Figs. 3 and 4).

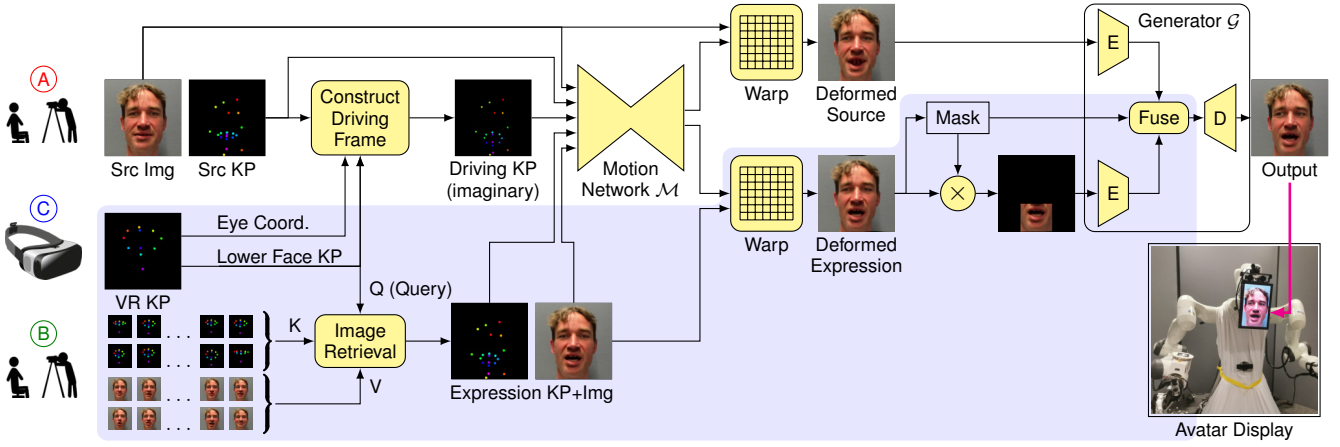


Fig. 3. Inference pipeline for VR Facial Animation. We select a still image from a portrait video of the operator shot before the run as source image (A). The remaining frames are used as a key-value storage of expression keypoints and corresponding images (B). The live keypoints measured inside and outside the VR headset (C) are then approximately projected to the source image frame, where they are used to retrieve the closest expression keypoints and image from the storage. The source keypoints, a constructed set of driving keypoints, and the retrieved expression keypoints then enter the motion network  $\mathcal{M}$ , which estimates warping vectors that deform the source image and the expression image to match the driving keypoints. The two deformed images then enter an encoder-decoder architecture that fuses them and generates the output image.

Similar to Siarohin *et al.* [4], we separate our pipeline into a keypoint detector  $\mathcal{K}$ , a motion network  $\mathcal{M}$ , and an image generator  $\mathcal{G}$ . The motion network  $\mathcal{M}$  uses the keypoints extracted by  $\mathcal{K}$  to generate a deformation grid  $\mathcal{D}_{S \leftarrow D}$  which warps the source image  $I_S$  to the head pose and expression of the driving frame  $I_D$ .

Solely warping is insufficient to generate a realistic output, though. To address this issue, we add a generator network  $\mathcal{G}$  which then creates the final output image  $I_O$ , given the initial motion estimate  ${}^D I_S$ . For precise architectural information, we refer to Siarohin *et al.* [4].

Note that there are approaches that forego using a source image by embedding appearance features directly in the network weights [16], [19], however, making only implicit use of the appearance  $\Lambda_S$  as encoded in a given source image allows us to generalize to unseen operators.

a) *Keypoint Detector Network*: Our keypoint detector  $\mathcal{K}$  extracts keypoints from the source image  $I_S$  and from the driving image  $I_D$ . We obtain two sequences of  $k$  keypoints, respectively:

$$\mathcal{K}(I_S) = [kp_S^{(1)}, kp_S^{(2)}, \dots, kp_S^{(k)}] \text{ and} \quad (1)$$

$$\mathcal{K}(I_D) = [kp_D^{(1)}, kp_D^{(2)}, \dots, kp_D^{(k)}]. \quad (2)$$

Unlike Siarohin *et al.* [4], we separate our keypoint detector into two models:

- (i) One model  $\mathcal{K}_{VR}$  for the expression of the lower face part, and
- (ii) a global keypoint detector  $\mathcal{K}_F$  that has access to images of whole faces, which detects all other keypoints (eyes, head pose, etc.).

$\mathcal{K}_{VR}$  and  $\mathcal{K}_F$  are both based on an Hourglass network [20]. This separation is necessary because during inference the operator’s facial expression, i.e. the driving frame, cannot be reconstructed from a single image due to occlusions by the VR headset.

The global keypoint detector  $\mathcal{K}_F$  is trained to detect keypoints, which are primarily encoding the head pose  $\theta$  and information about the operator’s eyes. We obtain these keypoints from a pretrained first-order-model keypoint detector [4], which was trained in a self-supervised manner.

In contrast,  $\mathcal{K}_{VR}$  is trained in a supervised manner using annotated images from the Vox-Celeb dataset [5]. We annotate Vox-Celeb images by cropping the face and extracting keypoints using the method proposed by Bulat and Tzimiropoulos [14]. However, we are only interested in the keypoints of the lower face which are visible in our mouth camera (see Fig. 2). In order to simulate the lower-face image  $I_M$  during training, we crop a random quadratic region with the only constraint that all lower-face keypoints must fit into this region. The cropped region is then resized to  $128 \times 128$  pixels. We therefore implicitly train  $\mathcal{K}_{VR}$  to extract keypoints in partially visible faces—as captured at inference time by the mouth camera.

For simplification, we define

$$\mathcal{K}(I) := \mathcal{K}_{VR}(I) \oplus \mathcal{K}_F(I) \quad (3)$$

to be the concatenation  $\oplus$  of both keypoint sequences.

b) *Motion Network*: The motion network  $\mathcal{M}$  is also based on an Hourglass architecture and produces a deformation of the source frame appearance  $\Lambda_S$ , represented with the driving frame’s facial expression and head pose  $(\xi_D, \theta_D)$ . Similar to Siarohin *et al.* [4], we first create for each keypoint  $kp_S^{(i)} \in \mathcal{K}(I_S)$  a shifted source image that aligns  $kp_S^{(i)}$  with the corresponding driving keypoint  $kp_D^{(i)} \in \mathcal{K}(I_D)$ . These  $k$  shifted versions are then fed into  $\mathcal{M}$  together with the heatmap representation of the driving keypoints. The motion network then predicts a deformation grid  $\mathcal{D}_{S \leftarrow D}$  which can be used to sample the source image deformed to the driving keypoints  ${}^D I_S$ . For a broader explanation we refer to Siarohin *et al.* [4].

To enhance the modeling capabilities, Siarohin *et al.* [4]



propose to use local affine transformations for each keypoint instead of just shifting. However, this assumes the existence of a complete driving frame. Since our driving frame, i.e. the mouth image, is partially occluded and captured from a different perspective, we rely on a motion network which processes shifted source images and keypoint heatmaps only. This allows us to create imaginary driving frame keypoints with perspective corrections and arbitrary head poses (see Section III-C).

c) *Generator Network*: The basic generator network has an encoder-decoder architecture and predicts the output image  $I_O$ , given the deformed source image  ${}^D I_S$ .

### C. Inference with VR Headset

In order to animate an operator controlling our avatar, we first need to capture a source image of them without any occlusions. The major challenge during inference is that we do not have access to a driving frame corresponding to a complete face image. Instead, we have to work with a mouth camera and two cameras capturing the eyes. The basic inference pipeline is illustrated in the non-blue area of Figure 3.

To obtain valid driving keypoints to which we can map the source frame’s appearance  $\Lambda_S$ , we construct the keypoints  $\hat{k}p_D^{(i)} \in \mathcal{K}(\hat{I}_D)$  of an imaginary driving frame  $\hat{I}_D$ . The constructed keypoints  $\mathcal{K}(\hat{I}_D)$  should encode (i) the same (frontal facing) head pose as in the source image, (ii) the gaze direction and eye openness of the operator, and (iii) the lower-face expression which is captured by the mouth camera. Note that (ii) and (iii) include nearly all the expressions that we can capture using the mouth and eye cameras.

Taking (i)-(iii) into account, we therefore build the driving keypoints (as illustrated by “Construct Driving Frame” in Figure 3) from

$$\mathcal{K}(\hat{I}_D) = \Pi_S(\mathcal{K}_{\mathcal{VR}}(I_M)) \oplus \rho(\mathcal{K}_{\mathcal{F}}(I_S), \hat{k}p_{eye}), \quad (4)$$

where  $I_M$  is the mouth image,  $\Pi_S(\cdot)$  (see Eq. (5)) maps each lower-face keypoint  $kp_M^{(i)} \in \mathcal{K}_{\mathcal{VR}}(I_M)$  into the source image  $I_S$ , and  $\rho(\cdot)$  replaces the eye key points detected in  $I_S$  with the modified values  $\hat{k}p_{eye}$  in order to include the operator’s current gaze direction and eye openness. Furthermore, keypoints which encode the face pose are simply copied from the source image. This is sufficient in our use case, since we move the avatar’s head display following the operator’s head motions (see Figure 1).

### D. Learned Keypoint Mapping $\Pi_S$

Constructing the keypoints of our imaginary driving frame  $\mathcal{K}(\hat{I}_D)$  requires mapping keypoints from the mouth camera space to the source frame. A naïve approach would be to simply perform a translation and scale adjustment [3], however, this does not consider perspective differences. One could also attempt to estimate the transformation matrix that maps from the mouth camera to the source camera. This would force us to estimate exact depth in the mouth image, which is not very robust. Furthermore, there are facial

distortions caused by the weight of the HMD, which would not be modeled by such an approach.

Instead, we propose to learn a separate homogeneous transformation matrix  ${}^S T_M^{(i)}$  that maps from each lower-face keypoint  $kp_M^{(i)}$  to the corresponding keypoint  $\hat{k}p_D^{(i)}$  in the source frame head pose. Estimating this transformation is not trivial, since the lower face keypoints exhibit a high variance and learning such a mapping requires corresponding mouth and source image pairs. We therefore capture not just a single source image of the operator, but a whole video sequence including mouth movements and a similar second video when the operator wears the VR headset. Note that finding corresponding pairs is challenging since we cannot capture both videos simultaneously.

Because establishing correspondences manually is expensive, we solve this alignment problem approximately by (i) capturing two videos with roughly the same mouth expressions and (ii) iteratively refining the learned keypoint transformations based on the currently associated pairs. This iterative process has the following steps:

- 1) Extract the keypoints of all source and mouth video frames.
- 2) Initialize each homogeneous transformation  ${}^S T_M^{(i)}$  from the mean scale difference between the two sets of keypoints.
- 3) Map the keypoints of each mouth image into each source frame.
- 4) For each mapped keypoint sequence, search for the best corresponding source image, yielding  $N$  pairs of images.
- 5) Optimize each  ${}^S T_M^{(i)}$  to minimize the Euclidean distance of the current keypoint pairs and goto step 3) unless the maximum number of 1000 iterations is met.

Note that the optimization runs in approx. 10s for a reasonable number of 250 calibration images. In order to create robustness against head movements while capturing the source video and changes of the VR headset relative to the operator, we define the mapping in a coordinate system relative to the centroid of  $I_M$  and  $I_S$ :

$$\Pi_S(kp_M^{(i)}) = {}^S T_M^{(i)} (kp_M^{(i)} - \overline{\mathcal{K}_{\mathcal{VR}}}(I_M)) + \overline{\mathcal{K}_{\mathcal{VR}}}(I_S), \quad (5)$$

where  $\overline{\mathcal{K}_{\mathcal{VR}}}(I_S), \overline{\mathcal{K}_{\mathcal{VR}}}(I_M)$  are the mean values of the lower-face keypoints in the source image and the VR camera mouth image, respectively. As demonstrated in Figure 8, this normalization also generates robustness to switching the operator, who is controlling our system, at inference time.

### E. Extended Inference Pipeline with Image Retrieval

Whereas the head and eye positions can be encoded well using only keypoints, it is highly challenging to generate proper mouth expressions using only one source image and few keypoints. To address this issue, we capture not just a single source image but multiple source frames in a video. This allows us to dynamically change the source image to the one which is closest to the projected lower-face keypoints of the current mouth image. In such an approach, the pipeline

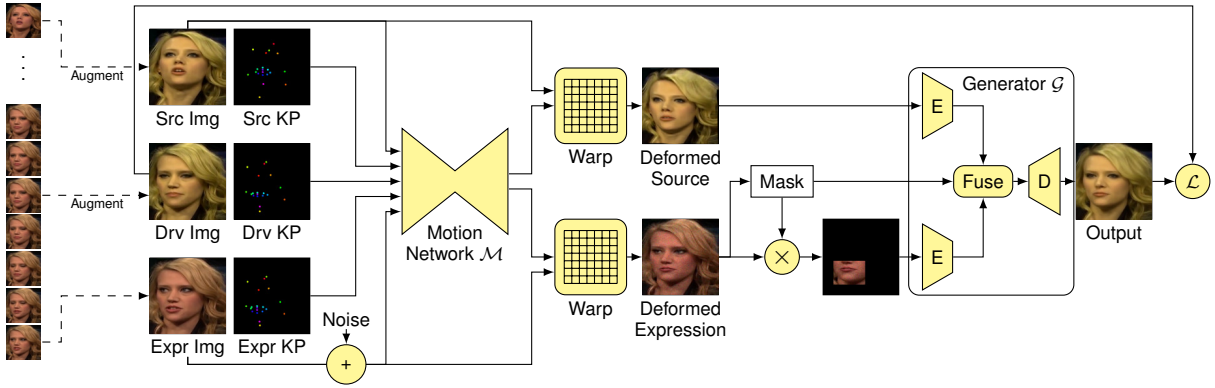


Fig. 4. Training the facial animation network from videos. The training loss  $\mathcal{L}$  is minimized when the network reconstructs the driving image from source image and keypoints, as well as the driving keypoints. The expression image is chosen from a close time interval around the driving image and is available as auxiliary input which is already close to the target expression.

would then need to apply only small corrections to the facial expression and could, thus, generate more accurate mouth expressions. However, in this naïve setup, non-negligible flickering effects would appear whenever we change the source view.

Instead, we propose to train a modified generator network that allows to decode the information of not just one source frame but also the lower-face information of a second source frame, which should be closer to the target. Therefore, the primary source frame  $I_S$  always remains constant while the second source frame, which we call the expression frame  $I_E$ , can change arbitrarily (see Figure 3). This has a positive effect on temporal continuity and counteracts flickering during source frame changes significantly.

Following standard information retrieval, we compare the current query

$$Q = \mathcal{K}_{\mathcal{VR}}(I_M)$$

with all keys

$$K = [\mathcal{K}_{\mathcal{VR}}(I_{S_1}), \dots, \mathcal{K}_{\mathcal{VR}}(I_{S_n})]$$

via

$$\sum_l \|\Pi_{S_i}(Q^{(l)}) - k_i^{(l)}\|_2, k_i \in K$$

to retrieve the optimal index of the (image, keypoints) tuples

$$V = [(I_{S_1}, \mathcal{K}(I_{S_1})), \dots, (I_{S_n}, \mathcal{K}(I_{S_n}))].$$

We modify the pipeline and generator network accordingly:

- (i) Use the motion network to generate a deformed image of both the source frame and the expression frame ( $^D I_S, ^D I_E$ ).
- (ii) Split the generator into two separate encoder networks, where the first encoder  $\mathcal{G}_S^{Enc}$  extracts the source image features

$$^D f_S = \mathcal{G}_S^{Enc}(^D I_S) \quad (6)$$

and the second encoder  $\mathcal{G}_E^{Enc}$  extracts the expression image features

$$^D f_E = \mathcal{G}_E^{Enc}(m \odot ^D I_E), \quad (7)$$

where  $m \in \{0, 1\}^{N \times N}$  is a mask that hides all information except the lower face region of  $^D I_E$ .

- (iii) Fuse the activation by

$$^D f_{S,E} = \frac{m_{\downarrow}}{2} \odot (^D f_S + ^D f_E) + (1 - m_{\downarrow}) \odot ^D f_S, \quad (8)$$

where  $\odot$  denotes element-wise multiplication and  $m_{\downarrow}$  is a down-scaled version of the binary mask  $m$ .

- (iv) Decode the fused representation with the decoder  $\mathcal{G}^{Dec}$  of the generator network

$$I_O = \mathcal{G}^{Dec}(^D f_{S,E}). \quad (9)$$

This additional branch to our inference pipeline is visualized in the blue-marked area of Figure 3.

When capturing the two videos (with and without HMD), it is important to have a high variety of mouth expression in the set of source images. We therefore propose to capture two videos of the operator, while reading a sentence that covers many phonemes. The sentence "That quick beige fox jumped in the air over each thin dog, look out he shouts for he's foiled you again, creating chaos", is known to match these requirements and gave us good results during testing.

## F. Training

The training pipeline is illustrated in Figure 4.

During inference, we use information retrieval to select the current expression frame, based on the optimized keypoint transformations  $^S T_M$ . During training, however, we use the Vox-Celeb dataset [5], which mainly consists of celebrities being interviewed. In this setup, it is not trivial to select an expression frame: If we would just set the expression image to be the driving image itself, the network would learn to simply copy the information.

To avoid this, we select the expression frame from a small interval around the driving image. This makes the assumption that temporally close frames also exhibit similar expressions.

The generator network (especially the feature encoder sub-modules) can thus learn to mostly ignore the lower-face region of the motion-transmitted source image, since the expression image is generally closer to the driving frame. This, however, leads to temporal instabilities whenever the

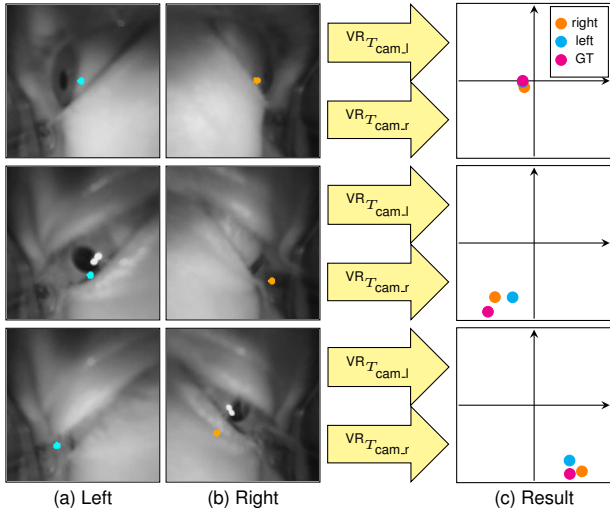


Fig. 5. Eye tracking. We show the eye keypoints in the left and right image frame, which are learned without direct supervision (a, b). The resulting transformed gaze direction prediction is shown in (c) in cyan and orange, with the ground truth in magenta. The transformations  $VR_{T_{cam}}$  are learned.

expression frame is switched. We counteract this by choosing a reasonable interval around the driving frame and augment the chosen expression frame using color jittering and injection of several types of random noise as proposed by Carlson *et al.* [21]. We argue that in this setup, the generator network is now explicitly guided to keep the mapped source frame information  $^DI_S$  to generate a proper facial animation.

### G. Eye Tracking and Animation

We introduce an image-driven eye tracking pipeline that needs to be calibrated once and can be trained in less than a minute for an operator. To obtain training data, we request the operator to follow with their gaze a red dot that moves in the VR display. This gives us calibration/training triplets of left eye images, right eye images and 3D gaze directions.

*a) Network Architecture:* We build a very lightweight hourglass network with only two downsampling and two upsampling layers that takes an input eye image and outputs a heatmap which is used to generate a single keypoint (see Fig. 5). We map this keypoint coordinate into the VR space using a learned homogeneous transformation  $VR_{T_{cam}}$ , which is jointly optimized with the hourglass network. While the homogeneous transformation is trained with supervision, we train the hourglass end-to-end in a self-supervised manner.

*b) Inference:* During inference, we take the mean prediction  $p$  of both the left eye prediction  $p_L \in [-1, 1]^2$  and the right eye prediction  $p_R \in [-1, 1]^2$ . We, furthermore, estimate a normalized confidence measure

$$c = 1 - \frac{1}{2\sqrt{2}} \|p_L - p_R\|_2 \in [0, 1] \quad (10)$$

which is large when both predicted eye coordinates are close to each other.

At this stage, only the recognition of the eye openness remains. We found that the eye openness strongly correlates with the gaze direction, i.e. the more an operator looks down




Fig. 6. Illustration of the eyes reacting to keypoint manipulations. The center image defines the origin and the other images define the boundaries of the normalized eye coordinate system.

the less open the eyes are. This property was also learned by our networks. To simulate both eyes, the networks only use one keypoint at the upper eyelid directly above the pupil center of the right eye (see Fig. 6). During inference, we can thus control the gaze and eye openness by modifying a single keypoint of the source image.

The assumptions above apply as long as the eyes are not completely closed, i.e. not when blinking. It is beneficial to detect this case without requiring additional annotations. Our experiments showed that whenever the eyes are closed, there is a very low confidence value. This is due to the fact that the predictions of both eyes are very different, since we did not equip the networks with the capabilities of handling such cases. Hence, we detect closed eyes whenever the confidence parameter is below a threshold  $\lambda_C$ . This also has the benefit of hiding implausible eye configurations from the viewer by showing the operator with closed eyes.

*c) Eye Coordinate System:* One remaining problem is still to define the region in the source image in which the eye keypoint can move, i.e. a coordinate system mapping. One possibility is to automatically find these boundaries. However, this requires to capture a second video of the operator (without VR headsets) in which the eyes are moving.

To decrease required capture time, we built an interactive annotation tool, which renders the source view with eye keypoints at the current cursor position. This allows us to manually define the boundaries of the possible eye keypoints in the source image, which are illustrated in Figure 6. The annotated boundaries then define a normalized coordinate system which is centered in the frontal-facing gaze direction.

### H. Temporal Consistency Filters

We apply several filters to the image retrieval process and facial keypoints to enhance the temporal consistency.

*a) Expression Frame Filtering:* We apply a straightforward filter that regulates the expression frame retrieval process: We only allow to change the expression frame  $I_E$ , when there is another frame  $1+\lambda_{swap}$  times closer to the current projected lower-face keypoints. This hysteresis avoids fast switching.




Fig. 7. Image retrieval process experiments. Unlike the first two columns, the last two columns are examples for image retrievals where the operator keys and values are different from the operator in the query.

We also apply a recursive low-pass filter parameterized by  $\lambda_E, \lambda_{\tilde{E}}, \lambda_O$ . Instead of simply using the raw expression frame  $I_E$ , we propose to build the current expression frame  $I_{\tilde{E}}^{(t)}$  at time step  $t$  according to

$$I_{\tilde{E}}^{(t)} = \lambda_E I_E^{(t)} + \lambda_{\tilde{E}}^E I_{\tilde{E}}^{(t-1)} + \lambda_O^E I_O^{(t-1)}, \quad (11)$$

where  $I_{\tilde{E}}^{(t)}$  is the current (raw) expression frame,  $^E I_{\tilde{E}}^{(t-1)}$  is the last expression frame  $I_{\tilde{E}}^{(t-1)}$  (recursively) mapped to the current expression frame using the motion network, and  $^E I_O^{(t-1)}$  is the last prediction  $I_O^{(t-1)}$  mapped to the current expression frame, respectively. In practice, we choose  $\lambda_E = 0.7$ ,  $\lambda_{\tilde{E}} = 0.1$ , and  $\lambda_O = 0.2$ .

*b) Eye Keypoint Filtering:* As explained in Section III-G, we can obtain a confidence value of the current eye position by comparing the left and right eye. We then recursively low-pass filter the eye position and parametrize the filter with the confidence  $c \in [0, 1]$ . We propose to derive the filtered eye position  $\tilde{p}$  from

$$\tilde{p}^{(t)} = \lambda_G p^{(t)} + (1 - \lambda_G) \tilde{p}^{(t-1)}, \quad (12)$$

where we determine the contribution  $\lambda_G \in [0, 1]$  of the current eye position estimate  $p^{(t)}$  according to:

$$\lambda_G = \begin{cases} 5c - 4, & \text{if } c \geq 0.8 \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

This low-pass filter is especially useful when the prediction  $p$  of the eye tracking network is noisy.

#### IV. EXPERIMENTS

We report several qualitative results and discuss our performance at the ANA Avatar XPRIZE semifinals.

##### A. Qualitative Results

*a) VR Facial Animation:* Figure 8 illustrates exemplary results of our full pipeline when (i) mapping from one operator to the appearance of the same operator and (ii) mapping to the appearance of another operator. Note that our complete forward pass runs with 33 fps on a single NVIDIA RTX 3090 GPU and an image resolution of  $256 \times 256$ .

*b) Image Retrieval:* Figure 7 visualizes the accuracy of our proposed keypoint-based image retrieval. Experiments have shown that image retrieval from a set of source images allows us to generate much more accurate mouth expressions compared to just a single source. As shown in Figure 7, the image retrieval also works well between different operators.

##### B. The ANA Avatar XPRIZE Semifinals

At the semifinals of the ANA Avatar XPRIZE Competition, three scenarios had to be accomplished with our avatar system by a previously unknown operator. The scenarios were repeated in a second run, with the better score persisting. Before each run, a preparation time of one hour was allotted in which we could introduce the operator to the system. We also used this time to prepare the facial animation pipeline, including recording two source videos (with and without the HMD), optimizing the keypoint transformations, calibrating and training the eye tracking pipeline, and annotating eye coordinates using our interactive tool. On average, operator adaptation took us about 15 minutes, which could mostly be done in parallel to the general operator introduction. The scenarios all required the operator to interact though the avatar system with another person from the jury, i.e. the recipient. For the recipient, facial animation is particularly important, as it helps to convey the focus of attention and the emotions of the operator and confirms the recipients that they are interacting with a real person.

XPRIZE defined rigorous scoring criteria. Specific criteria regarding facial animation were:

- 1) The Recipient was able to identify the remote Operator and felt the Operator was present in space.
- 2) The Recipient was able to understand the Operator's emotions through the Avatar.
- 3) The Recipient felt a sense of shared experience with the remote Operator.
- 4) The Recipient was able to understand the Operator's gestures through the Avatar.
- 5) The Operator was able to express their emotions.

Our team NimbRo performed very well at these and all other criteria and was ranked first in the semifinals with an overall score of 99 out of 100 points. Figure 1 shows our avatar system and facial animations during the challenge.

#### V. CONCLUSIONS

We have demonstrated an efficient and real-time capable VR Facial Animation pipeline that generalizes well to operators unknown *a priori* with a fast adaptation process. We show how to accurately solve the alignment problem between images captured with and without an HMD using learned keypoint transformations and iterative refinement. Furthermore, we demonstrate how to ensure temporal continuity. Using the proposed method, our team reached an almost perfect score of 99/100 points and was ranked first in the ANA Avatar XPRIZE semifinals event in Miami with a total of 28 teams. A useful extension of our method could be to forgo selecting an expression frame, but directly use the mouth image as captured by the mouth camera.



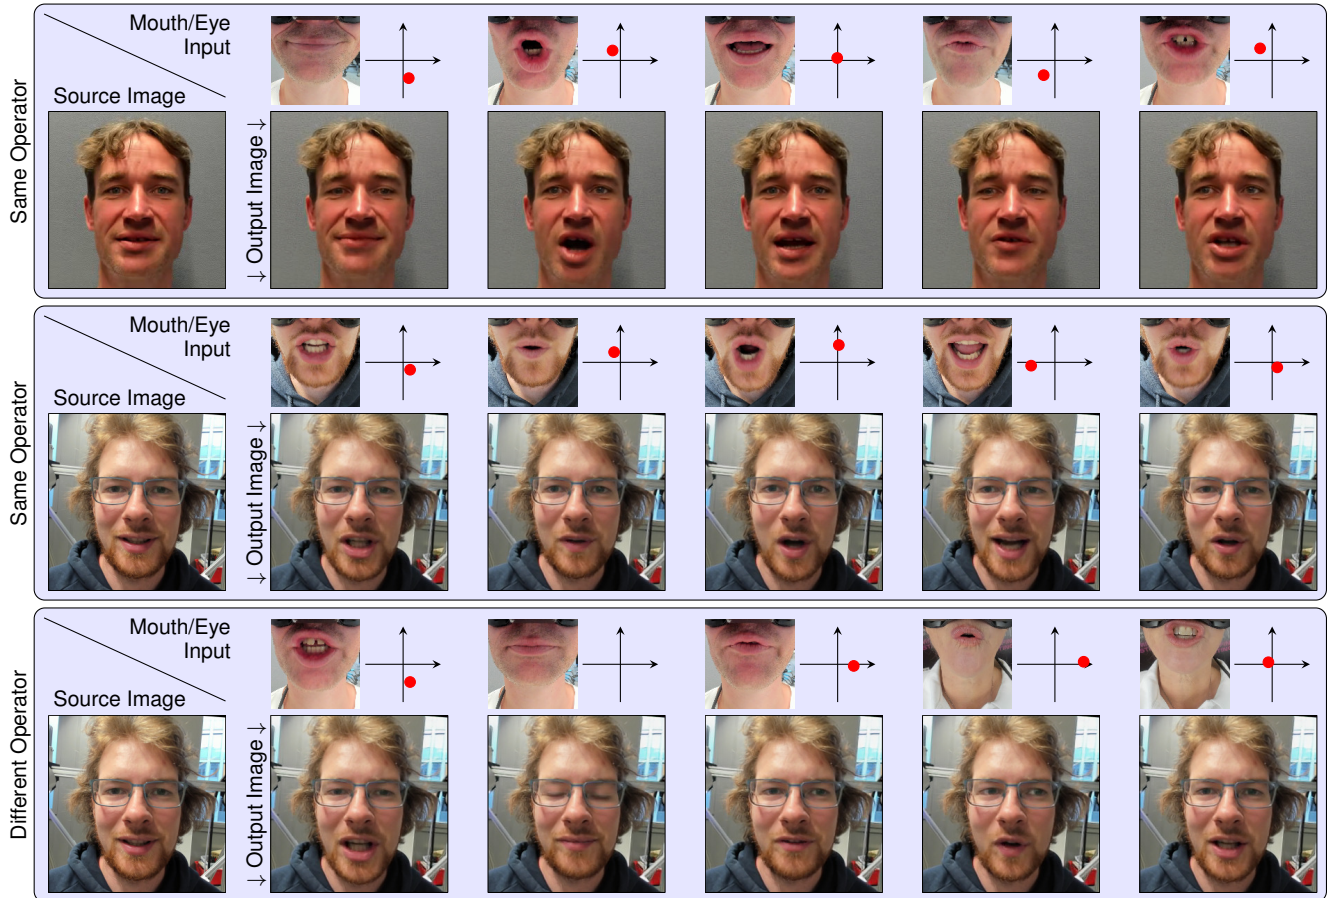


Fig. 8. Generated faces. In each box we use the same source image (left) to generate a facial reconstruction given the mouth camera image and eye coordinates. Note how the system matches mouth and eye configurations closely (all rows) and even performs inter-person animation (bottom row).

## REFERENCES

- [1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Neural Information Processing Systems (NeurIPS)*, vol. 27, 2014.
- [2] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," in *European Conference on Computer Vision (ECCV)*, 2020.
- [3] M. Schwarz, C. Lenz, A. Rochow, M. Schreiber, and S. Behnke, "NimbRo Avatar: Interactive immersive telepresence with force-feedback telemanipulation," in *International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [4] A. Siarohin, S. Lathuilière, S. Tulyakov, E. Ricci, and N. Sebe, "First order motion model for image animation," in *Neural Information Processing Systems (NeurIPS)*, 2019.
- [5] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "Voxceleb: Large-scale speaker verification in the wild," *Computer Speech & Language*, vol. 60, 2020.
- [6] O. Wiles, A. Koepke, and A. Zisserman, "X2Face: A network for controlling face generation using images, audio, and pose codes," in *European Conference on Computer Vision (ECCV)*, 2018.
- [7] A. Siarohin, S. Lathuilière, S. Tulyakov, E. Ricci, and N. Sebe, "Animating arbitrary objects via deep motion transfer," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [8] R. Zhao, T. Wu, and G. Guo, "Sparse to dense motion transfer for face image animation," in *International Conference on Computer Vision (ICCV)*, 2021.
- [9] T.-C. Wang, M.-Y. Liu, A. Tao, G. Liu, J. Kautz, and B. Catanzaro, "Few-shot video-to-video synthesis," in *Neural Information Processing Systems (NeurIPS)*, 2019.
- [10] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [11] A. Pumarola, A. Agudo, A. M. Martinez, A. Sanfeliu, and F. Moreno-Noguer, "GANimation: Anatomically-aware facial animation from a single image," in *European Conf. on Computer Vision (ECCV)*, 2018.
- [12] E. Zakharov, A. Shysheya, E. Burkov, and V. Lempitsky, "Few-shot adversarial learning of realistic neural talking head models," in *International Conference on Computer Vision (ICCV)*, 2019.
- [13] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *International Conference on Computer Vision (ICCV)*, 2017.
- [14] A. Bulat and G. Tzimiropoulos, "How far are we from solving the 2D & 3D face alignment problem? (and a dataset of 230,000 3D facial landmarks)," in *Int. Conference on Computer Vision (ICCV)*, 2017.
- [15] J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner, "Facevr: Real-time gaze-aware facial reenactment in virtual reality," *ACM Transactions on Graphics (ToG)*, vol. 37, 2018.
- [16] S.-E. Wei, J. Saragih, T. Simon, A. W. Harley, S. Lombardi, M. Perdoch, A. Hypes, D. Wang, H. Badino, and Y. Sheikh, "VR facial animation via multiview image translation," *Transactions on Graphics (ToG)*, vol. 38, no. 4, 2019.
- [17] S. Lombardi, J. Saragih, T. Simon, and Y. Sheikh, "Deep appearance models for face rendering," *Transactions on Graphics (ToG)*, vol. 37, no. 4, 2018.
- [18] A. Richard, C. Lea, S. Ma, J. Gall, F. De la Torre, and Y. Sheikh, "Audio-and gaze-driven facial animation of codec avatars," in *Winter Conference on Applications of Computer Vision (WACV)*, 2021.
- [19] G. Gafni, J. Thies, M. Zollhofer, and M. Nießner, "Dynamic neural radiance fields for monocular 4D facial avatar reconstruction," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [20] A. Newell, K. Yang, and J. Deng, "Stacked hourglass networks for human pose estimation," in *European Conference on Computer Vision (ECCV)*, 2016.
- [21] A. Carlson, K. A. Skinner, R. Vasudevan, and M. Johnson-Roberson, "Modeling camera effects to improve visual learning from synthetic data," in *European Conference on Computer Vision (ECCV)*, 2018.