Robust Egocentric Photo-realistic Facial Expression Transfer for Virtual Reality

Amin Jourabloo[†] Baris Gecer[‡] Fernando De la Torre^{§†} Jason Saragih[†] Shih-En Wei[†] Stephen Lombardi[†] Te-Li Wang[†] Danielle Belko[†] Autumn Trimble[†] Hernan Badino[†] [†]Facebook Reality Labs, Pittsburgh, PA [‡]Imperial College London

[§]Robotics Institute, Carnegie Mellon University

Abstract

Social presence, the feeling of being there with a "real" person, will fuel the next generation of communication systems driven by digital humans in virtual reality (VR). The best 3D video-realistic VR avatars that minimize the uncanny effect rely on person-specific (PS) models. However, these PS models are time-consuming to build and are typically trained with limited data variability, which results in poor generalization and robustness. Major sources of variability that affects the accuracy of facial expression transfer algorithms include using different VR headsets (e.g., camera configuration, slop of the headset), facial appearance changes over time (e.g., beard, make-up), and environmental factors (e.g., lighting, backgrounds). This is a major drawback for the scalability of these models in VR.

This paper makes progress in overcoming these limitations by proposing an end-to-end multi-identity architecture (MIA) trained with specialized augmentation strategies. MIA drives the shape component of the avatar from three cameras in the VR headset (two eyes, one mouth), in untrained subjects, using minimal personalized information (i.e., neutral 3D mesh shape). Similarly, if the PS texture decoder is available, MIA is able to drive the full avatar (shape+texture) robustly outperforming PS models in challenging scenarios. Our key contribution to improve robustness and generalization, is that our method implicitly decouples, in an unsupervised manner, the facial expression from nuisance factors (e.g., headset, environment, facial appearance). We demonstrate the superior performance and robustness of the proposed method versus state-of-the-art PS approaches in a variety of experiments.

1. Introduction

Our experience with communication systems is twodimensional, mostly via video teleconferencing (e.g., Messenger), that includes both audio and video transmissions. Recent studies on videoconferencing have shown that the more closely technology can simulate a face-to-face inter-



Figure 1. 3D Photo-realistic avatar driven by three headsetmounted camera (HMC) images in a VR headset. This paper presents a system to drive photo-realistic avatars robustly with variability in headsets, lighting, environmental background, head pose, and facial appearance.

action, the more participants are able to focus, engage, and retain information [43]. A more advanced level of communication with virtual reality (VR) via telepresence [5, 8, 12, 19, 30, 34, 44, 45] will allow virtual presence at a distant location and a more authentic interaction. If successful, this new form of face-to-face interaction can reduce the time and financial commitments of travel, make sales meetings or family meetings more immersive, with a huge impact for the environment and use of personal time.

Today most real-time systems for avatars in AR/VR are cartoon-like (e.g., Hyprsense, Loom AI); on the other hand, Hollywood has animated nearly uncanny digital humans as virtual avatars using advanced computer graphics technology and person-specific models (e.g., Siren). While some of these avatars can be driven in real-time from cameras, building the PS model is an extremely time-consuming and hand-tuned process that prevents democratization of this technology. This paper makes progress in this direction by generating video-realistic avatars by transferring subtle facial expressions from the headset mounted cameras (HMC) images in a VR headset to a 3D talking head (see Fig. 1).

We build on recent work on codec avatars (CA) [26] that learn a PS model from a Plenoptic study. Recall that driving an avatar from HMC cameras is typically more



Figure 2. Comparing the HMC images of a subject in multiple HMC captures with variations in headset, environment and facial appearance. (a) The training HMC capture, (b-e) the testing HMC captures. The **blue bold** font shows the variations respect to the training capture (a). The red circles in (f) show the locations of the cameras inside the headset.

challenging than driving it from regular cameras (e.g., iPhone) [25, 37, 52], due to the domain difference between IR cameras and the texture/shape of the avatar, variability in HMC images due to headset variability (e.g., camera location, IR LED illumination), high-distortion introduced by the near-camera views, and partial visibility of the face (Fig. 2). Wei *et al.* [49] proposed an end-to-end deep learning network for learning the mapping between the HMC images and the parameterized avatar. First, this model solves the unknown correspondence between HMC images and the avatar parameters in an unsupervised manner using an eleven-view HMC headset. Second, to animate the CA in real time from three HMC images(i.e., inference), Wei *et al.* [49] learns an encoder network to regress from 3view HMC images to CA's parameters (Fig. 3).

While previous work has reported compelling photorealistic facial expression transfer results, the existing method has limitations due to the PS nature of the approach. It is time-consuming, expensive and error-prone to capture sufficient statistical variability when collecting PS samples to learn a robust model. It will typically require recording several sessions with variability across lighting, headsets and iconic changes (e.g., makeup, beard), which limits its scalability. To build generic models (conditioned to the neutral shape), the most import contribution of this paper is to propose multi-identity architecture (MIA), an architecture that factorizes nuisance parameters such as camera parameters, facial aesthetic changes (e.g., beard, makeup) and environmental factors (e.g., lighting) from the facial motion (i.e., facial expression). This is critical because the encoder is able to extract from the HMC images *only* the information that is relevant to the final task, which is transferring subtle facial expressions, and it is able to marginalize information that is not relevant (headset, facial appearance, environment). Implicitly, this results in an algorithm that aligns facial expressions (3D shape + texture) across users in an unsupervised manner. Recall that is a very difficult problem to align subtle facial behavior (using both 3D shape + texture) across users in a supervised or unsupervised manner. That is, how can we find the correspondence of expression across subjects? Even if done manually, it is an extremely challenging problem and MIA (to the best of our knowledge) is the first algorithm that solves this problem in an unsupervised and discriminative manner (see subsection 4.3). MIA results in an algorithm for facial expression transfer for VR, that improves upon PS models in realistic scenarios.

2. Prior Work

2.1. Animating Stylized and Codec Avatars

Animating stylized avatars from video has a long history, for instance [7] fits a generic 3DMM to the face and use it to retarget the facial motion to a 3D characters. To improve the accuracy, Chaudhuri *et al.* [6] proposed to learn personspecific expression blendshapes and dynamic albedo maps from the input video of subjects. In [42], facial action unit intensity is estimated in a self-supervise manner by utilizing a differentiable rendering layer for fitting the expression and



Figure 3. Training and testing pipeline for animating the face codec avatar. In the data collection stage, we perform a face capture to generate the codec avatar of the subject [26] and a HMC capture. We utilize [49] to find the correspondence between the avatar and the HMC capture. Finally, we can train a model to animate the codec avatar (CA) from the HMC images in real-time.



Figure 4. The proposed multi-identity architecture (MIA). It consists of three main parts: the backbone network **B**, the 3D shape network **G**, and the texture networks \mathbf{F}_i . The identity selector module pass the features to the corresponding texture network.

to retarget the expression to the character. In contrast, expression transfer from a VR headset [12, 18, 27, 33] is more challenging due to partial visibility of face in HMC images, the specific hardware, and limited existing data.

CAs animate avatars by estimating the parameters of a PS shape and texture model from HMC [9, 26, 39, 49], see Fig. 3. In [26], combination of real and synthetic HMC images are utilized for reducing the domain gap between real HMC images in IR spectrum and rendered images for training encoder and reducing the HMC-avatar domain gap. Wei *et al.* [49] utilize a cycle-GAN to achieve accurate cycle consistency between 11-view HMC images and CA. Then, they train a person specific regressor from 3-view HMC images to the CA's parameter. Chu *et al.* [9] propose to use modular CA to have more freedom for animating the eyes and mouth. In a different approach, Richard *et al.* [36] animate the CA based on the gaze direction and audio inputs. The aforementioned methods rely on PS models, are typically not robust to variations in headsets and environments.

2.2. 3D Shape Estimation

Early approaches for model-based shape and texture estimation are based on active shape model [11] (ASM) and Active Appearance Model [10, 28] (AAM). The AAM methods learn a joint holistic model of shape and appearance. 3D Morphable Model (3DMM) provides a dense 3D representation for faces *e.g.* the Basel Face Model [35] and the FaceWarehouse [3]. In [16, 21], 3DMM is incorporated in an end-to-end CNN training to dicriminatively estimate the 3D shape of faces given single input image. Tran *et al.* [47] propose to learn a nonlinear 3DMM via deep neural network from in-the-wild images, and in this way 3DMMs are capable of representing non-linear facial expressions. The proposed method in [13] can extract expression-dependent details of the 3D shape from a single image. [14] proposed to use the GAN generator

for 3DMM fitting and estimating high-fidelity UV texture. Similarly, [20] proposed to utilize the volumetric representation of face instead of using 3DMM. An unsupervised method proposed in [15] for, identity 3DMM fitting, regressing the 3D shape and texture. Also, in [38] the identity constraints utilized among the images of the same subject. Similar to [47] we learn a non-linear discriminative 3DMM, but we extend it to learn the model from HMC images given a neutral 3D shape, and align the expressions across subjects in an unsupervised manner. To the best of our knowledge, this is the first work that solves the correspondence of expression across subjects in unsupervised manner.

3. Multi-Identity Model

This section describes the proposed multi-identity architecture (MIA) and augmentation techniques to robustify and generalize existing encoder models for driving CAs.

3.1. Multi-Identity Architecture (MIA)

Given 3-view HMC images of the eyes and mouth (see Fig. 1), our goal is to estimate the facial expression of a CA (shape+texture), and render it in an arbitrary view in VR. The MIA has three main parts (see Fig. 4): the backbone network, the 3D shape network, and the texture branch. **Backbone network:** The backbone network, \mathbf{B}_{ψ} in Fig. 4, is shared among subjects. Its goal is to factorize the expression from other nuisance factors such as lighting, background, or camera views, and build an internal representation that is invariant to those factors. As we will show in the experiments section, MIA naturally finds that the best way to encode HMC images across subjects, is by marginalizing out person-specific factors in addition to the nuisance factors mentioned. This results in learning an embedding that only preserves expression without the need of solving for correspondence across expression among subjects.

3D shape network: MIA assumes that the neutral shape



Figure 5. Examples of applying 3D augmentation layer to HMC images. First row: Real HMC images, second row: Augmented images by changing 3D pose, focal length and background.

of the test subject, $\mathbf{S}^N \in \mathbb{R}^{7306 \times 3}$, in given¹. This is the only information MIA needs to generalize the shape component of the network to untrained subjects. A network \mathbf{G}_{γ} is trained to estimate 3D shape, $\hat{\mathbf{S}} \in \mathbb{R}^{7306 \times 3}$ from HMC images. The network \mathbf{G}_{γ} takes both of the output of the backbone network \mathbf{B}_{ψ} and \mathbf{S}^N to estimate the person specific 3D shape expression residual. The neutral 3D shape is used to re-inject person-specific information that was factored out in \mathbf{B}_{ψ} . For instance, eye openness, which varies across identities, can be extracted from the neutral 3D shape \mathbf{S}^N of each subject. With this, we reconstruct the 3D shape of subject *i* as:

$$\hat{\mathbf{S}}_i = \mathbf{S}_i^N + \mathbf{G}_{\gamma}(\mathbf{B}_{\psi}(\mathbf{H}_i^0, \mathbf{H}_i^1, \mathbf{H}_i^2), \mathbf{S}_i^N).$$
(1)

The network \mathbf{G}_{γ} is trained by minimizing the Euclidean distance between the target \mathbf{S}_i and estimated $\hat{\mathbf{S}}_i$ 3D shapes,

$$\mathbf{L}_{S}^{i} = \|\mathbf{W}^{S} \odot (\mathbf{S}_{i} - \hat{\mathbf{S}}_{i})\|_{2}^{2}, \tag{2}$$

where \mathbf{W}^{S} is the weight mask for the visible areas.

Texture network: When the pre-trained PS texture decoder is available for each identity, our goal is to be able to animate the CA from HMC images robustly and with minimal adaptation effort. In this paper, we presume pre-trained decoders \mathbf{D}_{ϕ} from [26] are available, but our work can be similarly applied to other PS models as well (e.g. [24, 46]). The network \mathbf{D}_{ϕ} takes, as input, an expression parameter $\mathbf{z} \in \mathbb{R}^{256}$ and a view vector $\mathbf{v} \in \mathbb{R}^3$, and generates personspecific and view-specific texture $\mathbf{T}^{\mathbf{v}} \in \mathbb{R}^{1024 \times 1024 \times 3}$ that, together with shape, can be used to render the avatar,

$$\mathbf{T}^{\mathbf{v}} = \mathbf{D}_{\phi}(\mathbf{z}, \mathbf{v}). \tag{3}$$

However, since each PS model is trained independently of all others, the structure of the latent space, z, is not consistent across identities. We would like to utilize the shared backbone encoder \mathbf{B}_{ψ} across identities to encourage robustness via joint training. Inspired by multi-task learning techniques [4, 32], we additionally learn person-specific adaptation layers, \mathbf{F}_{θ} , that transform the identity-consistent expression embedding produced by \mathbf{B}_{ψ} to each identity's personalized latent space. Finally, to eliminate unnecessary dimensions in \mathbf{z} , non-informative dimensions, we apply PCA dimensionality reduction, denoted $\mathbf{P} \in \mathbb{R}^{256 \times 80}$ to each identity's latent space and fix it during training. Together, these components are used to generate PS expression parameters as follows:

$$\hat{\mathbf{z}}_i = \mathbf{P}_i(\mathbf{F}_{\theta_i}(\mathbf{B}_{\psi}(\mathbf{H}_i^0, \mathbf{H}_i^1, \mathbf{H}_i^2))) + \bar{\mathbf{z}}_i, \qquad (4)$$

where *i* is subject index, and $\overline{\mathbf{z}}_i$ is the average expression parameter for subject *i*. Then, we use Eqn. 3 to generate the estimated texture $\mathbf{T}_i^{\mathbf{v}}$ from view **v**. To guide the network, we minimize the Euclidean loss between the estimated and the target expression parameters and textures:

$$\mathbf{L}_{T}^{i} = \|\mathbf{z}_{i} - \hat{\mathbf{z}}_{i}\|_{2}^{2} + \lambda_{\mathbf{T}} \|\mathbf{W}^{T} \odot (\mathbf{T}_{i}^{\mathbf{v}} - \hat{\mathbf{T}}_{i}^{\mathbf{v}})\|_{2}^{2}, \quad (5)$$

where \mathbf{W}^T is the weight mask for the visible areas from the HMC images and λ_T is the weight for the texture loss.

Total Loss: The entire MIA network is trained end-to-end to optimize the networks' parameters by minimizing:

$$\underset{\psi,\gamma,\{\theta_i\}_{i=0}^{K}}{\text{minimize}} \sum_{i=0}^{K} \mathbf{L}_{T}^{i} + \lambda_{\mathbf{S}} \mathbf{L}_{S}^{i}, \tag{6}$$

where K is the number of subjects and λ_{s} is the weight for the shape loss.

3.2. Augmentation

Data augmentation is a wildly practiced heuristic in many deep learning tasks. The main goal is to make the distribution of variations in training data more similar to those in the test set. Most common data augmentations techniques include scaling [41], color augmentation [23], simple geometric transformations [40], and utilizing synthetic data [22, 29]. However, a major source of variability in our task stems from headset factors, such as variations in camera placement and focus, as well as the slop of the headset relative to the face which varies during usage. These variations are not easily modeled using standard augmentation techniques that do not take the 3D shape of the face into account. In this paper, we simulate headset-based variations by perturbing the 3D rotation and translation of the face shape in the training set, and use it to re-render augmented views of each HMC image on random backgrounds. Some examples are shown in Fig. 5. As demonstrated in the experiments section below, this simple augmentation technique substantially improves the robustness of our method to real world variations.

¹Extracting neutral face from a single or few-shot phone-captured images is a well studied problem [25, 37, 52], and there are a number of commercial solutions available [1, 2].



Figure 6. The testing results for estimating the 3D shapes for six untrained subjects and their ground truth based on 11-view results in [49].

4. Experimental Results

This section reports experimental results and analysis on MIA. The first experiment shows how MIA can estimate accurate 3D shapes directly from HMC images of untrained subjects. In the second experiment, we evaluate the quality of MIA's texture prediction for identities with pre-trained avatars under challenging testing scenarios. In the third experiment, we show how MIA can incorporate new subjects with minimal training. In addition, we also present further analysis about what MIA learns prior and during adaptation.

Data: We used 120 HMC captures of different subjects for training and 21 HMC captures for testing. Training and testing HMC captures do not overlap. Each HMC capture is a 45 minutes long video (30fps) of 11-views HMC images, and contains 73 peak expressions, two sets of continuous range-of-motion, recitation of 50 sentences and 5-10 minutes of conversion. The HMC images are in the IR spectrum with a resolution of 480×640 . During testing only 3-views are available. For each subject, we have a pre-trained decoder to generate PS texture for various expressions from arbitrary views. For more information, of how to build the PS decoder see [26] and Eqn. 3.

Ground Truth: We utilize the result of the method in [49], that solves for the correspondence between 11-views HMC images and the CA parameters as the ground truth. Recall that the training data is captured with 11-views to achieve more precise results in the correspondence between HMC and CA, while the testing data has only 3-views.

Baseline method: We compare MIA with the person specific (PS) encoder in [49]. The PS encoder is trained with one HMC capture (3-view images) and uses a CNN architecture with the same number of parameters as ours.

Evaluation metrics: We report the average Euclidean error for the eyes, mouth and face areas separately for both 3D shape and the texture. The 3D shape errors are measured in millimeters and the texture errors in raw intensity values (i.e. 0-255). We report the localized error metrics to analyze failure modes better. For example, the 3D shape error

Subject	3D Shape Error (mm)					
Subject	Face	Eyes	Mouth			
person 1	1.68	1.08	2.90			
person 2	1.51	1.21	2.32			
person 3	1.07	0.74	1.82			
person 4	1.84	1.20	2.92			
person 5	1.45	0.89	2.23			
person 6	1.47	0.97	2.56			
person 7	1.73	1.20	2.98			
person 8	1.57	0.89	2.52			
person 9	1.21	0.91	1.84			
person 10	1.58	1.00	2.68			
overall	1.51 ± 0.23	1.00 ± 0.16	2.47 ± 0.42			

Table 1. Test results for 3D shape estimation in untrained subjects.

in the eyes capture openness and blinking errors, while in the mouth, they capture deviations in lip shapes important for visual-speech. Similarly, texture error in the eyes is typically due to the errors in gaze direction, and in mouth, it corresponds to incorrect teeth and tongue estimation.

Implementation details: In training, we use the Adam optimizer, setting the batch size to 32 and the initial learning rate to 1e-3. We decrease learning rate by 8e-1 after each 25K iterations. In total, we train the encoder for 250K iterations, and set both of $\lambda_{\rm T}$ and $\lambda_{\rm S}$ to 100. We crop and resize the HMC images to 192×192 to focus on the face areas.

The backbone network, \mathbf{B}_{ψ} , consists of two residual networks [17], one for eye images $[\mathbf{H}^0, \mathbf{H}^1] \in \mathbb{R}^{192 \times 192 \times 2}$ and another for the mouth $\mathbf{H}^2 \in \mathbb{R}^{192 \times 192}$. Each network consists of a Res-Net head module, five BottleNeck blocks and a 64-way fully connected layer. Each BottleNeck block consists of ten convolutional layers with 3×3 and 1×1 filters. We add shortcut connections among the convolutional layers, and each layer is followed by ReLU [31] and Instance normalization [48] layers. To extract the final identity invariant features, we apply a global average pooling and a 64-way fully connected layer to the activations of the last BottleNeck block. The architecture of the 3D shape network, \mathbf{G}_{γ} , consists of four fully connected layers where each one is followed by a leaky ReLU [51] layer with negative slope of 0.2. We normalize the extracted features



Figure 7. Test results to estimate the 3D shape from HMC images of an untrained subject for a wide range of expressions.

from the HMC images and the neutral 3D shape, to account for their different domains, by employing group normalization [50] after concatenating the features. Finally, for the texture network \mathbf{F}_{θ} , we utilize the combination of a ReLU layer and a fully connected layer without bias.

4.1. Quantitative Evaluations

This section quantifies the performance of MIA using three experiments: (1) driving the 3D shape of untrained subjects. (2) robustness of shape and texture estimation for subjects with trained PS models. (3) generalization of learned features on new subjects.

Driving 3D Shape: Inputs to the shape generation network, \mathbf{G}_{γ} , are the HMC images and the corresponding identitiy's neutral 3D shape. We train the network with 120 subjects using the loss function in Eqn. 2 as guidance. Fig. 6 shows the estimated 3D shape for extreme expression examples from six untrained subjects along with their ground truth. Our 3D shape estimator captures subtle details in expressions necessary for inferring social signal. Table 1 shows the 3D shape errors for face, eyes and mouth areas of the whole sequence for ten untrained subjects. The error is less than 2mm in the face/eyes and 3mm in the mouth. Recall that MIA does not use any sample from the test subject other than the neutral shape and has never seen any HMC images for these subjects during training. Fig. 7 shows testing results of one untrained subject for a wide range of expressions. Note that PS [49] is not able to estimate the 3D

shape for untrained subjects.

Comparing Table 1 with PS's results for the 3D shape error in Table 2 (different capture), we find that MIA outperforms PS, despite PS having access to subject-specific HMC images, and their target shapes, during training. We suspect the reason for this is that MIA learns to marginalize the extrinsic variability of the problem (i.e. environment, headset) from the 120 subjects that is trained on, while the PS tends to overfit to the specific HMC capture session used for training. More comparative results can be found in the video in the supplementary material.

Driving Full Avatars: In this experiment, we evaluate the ability of MIA to generate both shape and texture and its robustness against extrinsic factors such as headset, environment and facial appearance variations. Here, data for test subject is available during training, but from a different HMC capture. The selected subject was captured on five different dates; examples of the HMC images are shown in Fig. 2. These samples show large appearance variations due to facial hair, pose changes in the headset slop, and camera assembly differences across headsets; it also contains background variation due to changing environment and overall lighting differences. We use one HMC capture (Fig. 2(a)) of the subject with 119 HMC captures of other subjects for training, and test on the remaining four HMC captures of that subject. Table 2 compares the testing errors of MIA against PS [49]. On test capture 1, which is very similar to the training capture, PS [49] performs better than MIA. But,

Table 2. Testing results for a subject with multiple testing HMC captures with different variations. One HMC capture of the subject is inside the training set. The 3D shape errors are in mm and the texture errors are in intensity.

Test	Sample	Variations	Method	3D Shape Error			Texture Error		
Capture	Image	variations	Wiethou	Face	Eyes	Mouth	Face	Eyes	Mouth
1 Fig. 2.(b)	Headset	PS [49]	0.85	0.65	1.33	1.13	1.93	1.50	
		MIA	1.20	0.85	1.90	1.33	2.34	1.79	
2 Fig. 2.(c)	Headset	PS [49]	2.04	0.77	4.71	1.84	2.47	3.51	
	Facial appearance	MIA	1.28	0.79	2.22	1.49	2.52	2.00	
3 Fig. 2.(d)	Headset	PS [49]	1.90	0.98	3.68	1.65	2.73	2.84	
	Facial appearance	MIA	1.26	0.86	2.23	1.32	2.52	2.03	
4 Fig. 2.(e)	Environment	PS [49]	2.21	0.86	4.92	1.92	2.33	3.39	
	Fig. 2.(e)	Facial appearance	MIA	1.14	0.73	1.94	1.45	2.11	2.05
		overall	PS [49]	1.75 ± 0.61	0.81 ± 0.13	3.66 ± 1.64	1.63 ± 0.35	$\textbf{2.36} \pm \textbf{0.33}$	2.81 ± 0.92
			MIA	1.22 ± 0.06	$\textbf{0.80} \pm \textbf{0.06}$	$\textbf{2.07} \pm \textbf{0.17}$	1.39 ± 0.08	2.37 ± 0.19	1.96 ± 0.12

Table 3. Testing results for training and testing on new subjects with pre-trained fixed backbone network.

Method	3D Shape Error			Texture Error		
	Face	Eyes	Mouth	Face	Eyes	Mouth
PS [49]	1.12 ± 0.26	0.74 ± 0.11	1.98 ± 0.65	$\textbf{2.22} \pm \textbf{0.61}$	2.90 ± 0.95	2.77 ± 0.83
MIA	1.05 ± 0.19	0.74 ± 0.09	1.71 ± 0.37	2.22 ± 0.62	2.88 ± 0.81	2.65 ± 0.78

its performance declines significantly when testing on the other captures, where variations in environment and facial appearance are more extreme. Note that the overall errors for MIA, for all areas of 3D shape and texture, are more stable and are similarly low across all test captures. The first two rows of Fig. 10 shows visual comparison of methods on the test HMC captures, where a significant reduction in expressive detail is noticeable in results for PS [49]. We refer the reader to the supplementary material for more results.

Adaptation to New Identities: We evaluated the generalization of MIA's feature extraction to new subjects on HMC captures of 6 subjects that are not trained in MIA. Each of the 6 subjects has more than one HMC capture exhibiting variations in extrinsic factors. We used the pre-trained MIA network with 120 subjects (excluding the test 6 subjects), and fix the shape generation network, G, and backbone network, \mathbf{B}_{ψ} . For each new subject, we trained a new small texture network \mathbf{F}_{θ} . During the testing on HMC captures with variations, we used the newly trained texture estimation branch for estimating the texture parameters, and decode both the texture and the 3D shape by utilizing Eqn. 3. Table 3 shows the overall errors for 3D shape and texture for different areas of 7 testing HMC captures of the 6 subjects. MIA achieves lower errors for all areas with smaller variability, demonstrating the effectiveness of the features extracted from the fixed backbone network. The last three rows of Fig. 10 show visualizations of this case for.

4.2. Ablation Study

3D augmentation layer: To analyze the advantage of using the 3D augmentation layer, we compare the errors of the PS [49] model, MIA with 3D augmentation trained with 1 subject, MIA without 3D augmentation (3D Aug) trained with 30 subjects, and MIA with 3D Aug trained with 30 subjects. Fig. 8 shows the average errors for the four test captures in Table 2. It shows that even using the 3D Aug layer with 1 subject reduces errors slightly in comparison to PS [49]. However, there is huge drop in errors by using the 3D Aug layer with 30 subjects. This reduction of er-



Figure 8. The advantage of the 3D augmentation layer, the errors drop significantly by using both of 3D augmentation and MIA.



Figure 9. The influence of number of training subjects. The shape errors are decreasing by increasing number of training subjects.

ror is more significant in the mouth area. It shows that the combination of MIA and 3D Aug is effective.

Influence of number of subjects: We evaluate the influence of number of training subjects in the performance of MIA during testing. We train MIA with 30, 60, 100 and 120 subjects, and test them on ten untrained subjects for estimating the 3D shapes. Fig. 9 shows that by increasing number of training subjects the 3D shape errors are decreasing, especially for the mouth area.

4.3. Unsupervised Expression Correspondence

The MIA implicitly learns to solve for correspondence across expressions in order to marginalize nuisance parameters (e.g., lighting). It naturally discovers that the best way to encode HMC multi-identity data is finding a latent space that only contains expression information. Fig. 11 illustrates how MIA learns to solve for correspondence across expressions. The first column shows the input HMC images and the second column is the CA of the subject in the first column. The remaining columns are the CAs of other subjects driven from the HMC images in first column, that is, the same extracted features from HMC images are utilized to estimate (by using the corresponding \mathbf{F}_{θ}) a new expression parameter (with the same facial expression meaning), z, in the latent space of each of the remaining subjects. As we can observe, MIA is able to *align* the expression across all of the subjects in an unsupervised manner, and creating a common expression-only space. Please pay attention to the mouth area in the second row of Fig. 11 that shows the same expression with different mouth interior.

5. Conclusion and Future Directions

This paper proposes MIA to robustify and generalize existing PS methods for driving CAs. MIA learns to extract identity invariant features related to facial expression while marginalizing nuisance factors (headset, environment, facial expression) in an unsupervised manner. We show that MIA is able to drive the shape component in untrained subjects, and if the PS texture decoder is available, with a minimum training, MIA can drive CAs for new subjects. For future directions, first, we will design new loss function based on the closeness of 3D surfaces to model the lips and eyes closure. Second, we will work on texture-conditional decoders to make the texture part of the method generalizable for new subjects without pre-trained decoders.



Figure 10. The comparison of PS [49] and MIA methods for animating the codec avatar from HMC images. The MIA can estimate more expressive and accurate expressions.



Figure 11. The examples of synchronized expressions of training subjects by using the same input HMC images.

6. Acknowledgement

We thank "Baris Gecer" for early implementation of the 3D augmentation approach in Section 3.2, and help brainstorming ideas during his internship at Facebook in the summer of 2019. By error, his name was first omitted from the CVPR 2022 publication of this work.

References

- [1] FacePlusPlus. https://www.faceplusplus.com/ 3dface/.4
- [2] KeenTools. https://keentools.io/. 4
- [3] Chen Cao, Yanlin Weng, Shun Zhou, Yiying Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Trans. Vis. Comput. Graph.*, 20(3):413–425, 2013. 3
- [4] Jiajiong Cao, Yingming Li, and Zhongfei Zhang. Partially shared multi-task convolutional neural network with local constraint for face attribute learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4290–4299, 2018. 4
- [5] Prashanth Chandran, Derek Bradley, Markus Gross, and Thabo Beeler. Semantic deep face models. In *Int. Conf.* 3D Vis., pages 345–354. IEEE, 2020. 1
- [6] Bindita Chaudhuri, Noranart Vesdapunt, Linda Shapiro, and Baoyuan Wang. Personalized face modeling for improved face reconstruction and motion retargeting. In *Eur. Conf. Comput. Vis.*, pages 142–160. Springer, 2020. 2
- [7] Bindita Chaudhuri, Noranart Vesdapunt, and Baoyuan Wang. Joint face detection and facial motion retargeting for multiple faces. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9719–9728, 2019. 2
- [8] Lele Chen, Chen Cao, Fernando De la Torre, Jason Saragih, Chenliang Xu, and Yaser Sheikh. High-fidelity face tracking for ar/vr via deep lighting adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 13059–13069, 2021.
- [9] Hang Chu, Shugao Ma, Fernando De la Torre, Sanja Fidler, and Yaser Sheikh. Expressive telepresence via modular codec avatars. In *Eur. Conf. Comput. Vis.*, pages 330–345. Springer, 2020. 3
- [10] Timothy F Cootes, Gareth J Edwards, and Christopher J Taylor. Active appearance models. In *Eur. Conf. Comput. Vis.*, pages 484–498. Springer, 1998. 3
- [11] Timothy F Cootes, Christopher J Taylor, and Andreas Lanitis. Active shape models: Evaluation of a multi-resolution method for improving image search. In *Brit. Mach. Vis. Conf.*, volume 1, pages 327–336. Citeseer, 1994. 3
- [12] Mohamed Elgharib, Mallikarjun BR, Ayush Tewari, Hyeongwoo Kim, Wentao Liu, Hans-Peter Seidel, and Christian Theobalt. Egoface: Egocentric face performance capture and videorealistic reenactment. arXiv preprint arXiv:1905.10822, 2019. 1, 3
- [13] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-thewild images. ACM Trans. Graph., 40(4):1–13, 2021. 3
- [14] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1155–1164, 2019. 3
- [15] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T Freeman. Unsupervised training for 3d morphable model regression. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8377–8386, 2018. 3
- [16] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3d dense

face alignment. In *Eur. Conf. Comput. Vis.*, pages 152–168. Springer, 2020. 3

- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. 5
- [18] Steven Hickson, Nick Dufour, Avneesh Sud, Vivek Kwatra, and Irfan Essa. Eyemotion: Classifying facial expressions in VR using eye-tracking cameras. In *IEEE Winter Conference on Applications of Computer Vision*, pages 1626–1635. IEEE, 2019. 3
- [19] Tao Hu, Kripasindhu Sarkar, Lingjie Liu, Matthias Zwicker, and Christian Theobalt. Egorenderer: Rendering human avatars from egocentric camera images. In *Int. Conf. Comput. Vis.*, pages 14528–14538, 2021. 1
- [20] Aaron S Jackson, Adrian Bulat, Vasileios Argyriou, and Georgios Tzimiropoulos. Large pose 3d face reconstruction from a single image via direct volumetric cnn regression. In *Int. Conf. Comput. Vis.*, pages 1031–1039, 2017. 3
- [21] Amin Jourabloo, Mao Ye, Xiaoming Liu, and Liu Ren. Poseinvariant face alignment with a single cnn. In *Int. Conf. Comput. Vis.*, pages 3200–3209, 2017. 3
- [22] Martin Klaudiny, Steven McDonagh, Derek Bradley, Thabo Beeler, and Kenny Mitchell. Real-time multi-view facial capture with synthetic training. In *Computer Graphics Forum*, volume 36, pages 325–336. Wiley Online Library, 2017. 4
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Adv. Neural Inform. Process. Syst., pages 1097– 1105, 2012. 4
- [24] Gun-Hee Lee and Seong-Whan Lee. Uncertainty-aware mesh decoder for high fidelity 3D face reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6100–6109, 2020. 4
- [25] Jiangke Lin, Yi Yuan, Tianjia Shao, and Kun Zhou. Towards high-fidelity 3d face reconstruction from in-the-wild images using graph convolutional networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5891–5900, 2020. 2, 4
- [26] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. ACM Trans. Graph., 37(4):1–13, 2018. 1, 2, 3, 4, 5
- [27] Jianwen Lou, Yiming Wang, Charles Nduka, Mahyar Hamedi, Ifigeneia Mavridou, Fei-Yue Wang, and Hui Yu. Realistic facial expression reconstruction for VR hmd users. *IEEE Trans. Multimedia*, 22(3):730–743, 2019. 3
- [28] Iain Matthews and Simon Baker. Active appearance models revisited. Int. J. Comput. Vis., 60(2):135–164, 2004. 3
- [29] Steven McDonagh, Martin Klaudiny, Derek Bradley, Thabo Beeler, Iain Matthews, and Kenny Mitchell. Synthetic prior design for real-time face tracking. In *Int. Conf. 3D Vis.*, pages 639–648. IEEE, 2016. 4
- [30] Koki Nagano, Jaewoo Seo, Jun Xing, Lingyu Wei, Zimo Li, Shunsuke Saito, Aviral Agarwal, Jens Fursund, and Hao Li. pagan: real-time avatars using dynamic textures. ACM Trans. Graph., 37(6):1–12, 2018. 1
- [31] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Int. Conf. Machine Learning*, 2010. 5

- [32] J Naruniec, L Helminger, C Schroers, and RM Weber. Highresolution neural face swapping for visual effects. In *Computer Graphics Forum*, volume 39, pages 173–184. Wiley Online Library, 2020. 4
- [33] Kyle Olszewski, Joseph J Lim, Shunsuke Saito, and Hao Li. High-fidelity facial and speech animation for vr hmds. ACM Trans. Graph., 35(6):1–14, 2016. 3
- [34] Sergio Orts-Escolano, Christoph Rhemann, Sean Fanello, Wayne Chang, Adarsh Kowdle, Yury Degtyarev, David Kim, Philip L Davidson, Sameh Khamis, Mingsong Dou, et al. Holoportation: Virtual 3d teleportation in real-time. In *Proceedings of Symposium on User Interface Software and Technology*, pages 741–754, 2016. 1
- [35] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *International Conference on Advanced Video and Signal Based Surveillance*, pages 296–301. IEEE, 2009. 3
- [36] Alexander Richard, Colin Lea, Shugao Ma, Juergen Gall, Fernando de la Torre, and Yaser Sheikh. Audio-and gazedriven facial animation of codec avatars. In *IEEE Winter Conference on Applications of Computer Vision*, 2020. 3
- [37] Joseph Roth, Yiying Tong, and Xiaoming Liu. Adaptive 3d face reconstruction from unconstrained photo collections. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4197–4206, 2016. 2, 4
- [38] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael J Black. Learning to regress 3d face shape and expression from an image without 3d supervision. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7763–7772, 2019. 3
- [39] Gabriel Schwartz, Shih-En Wei, Te-Li Wang, Stephen Lombardi, Tomas Simon, Jason Saragih, and Yaser Sheikh. The eyes have it: an integrated eye and face model for photorealistic facial animation. ACM Trans. Graph., 39(4):91–1, 2020. 3
- [40] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. CNN features off-the-shelf: an astounding baseline for recognition. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 806–813, 2014. 4
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 4
- [42] Xinhui Song, Tianyang Shi, Zunlei Feng, Mingli Song, Jackie Lin, Chuanjie Lin, Changjie Fan, and Yi Yuan. Unsupervised learning facial parameter regressor for action unit intensity estimation via differentiable renderer. In *Proceedings of ACM International Conference on Multimedia*, pages 2842–2851, 2020. 2
- [43] Oddvar Hagen Stein R Bolle, Frank Larsen and Mads Gilbert. Video conferencing versus telephone calls for team work across hospitals: a qualitative study on simulated emergencies. *Emergency Medicine*, 22, 2009. 1
- [44] Ayush Tewari, Florian Bernard, Pablo Garrido, Gaurav Bharaj, Mohamed Elgharib, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. Fml: Face model learning from videos. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10812–10822, 2019. 1

- [45] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2387–2395, 2016. 1
- [46] Luan Tran, Feng Liu, and Xiaoming Liu. Towards highfidelity nonlinear 3D face morphable model. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1126–1135, 2019. 4
- [47] Luan Tran and Xiaoming Liu. On learning 3d face morphable model from in-the-wild images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019. 3
- [48] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. arXiv preprint arXiv:1607.08022, 2016. 5
- [49] Shih-En Wei, Jason Saragih, Tomas Simon, Adam W Harley, Stephen Lombardi, Michal Perdoch, Alexander Hypes, Dawei Wang, Hernan Badino, and Yaser Sheikh. VR facial animation via multiview image translation. ACM Trans. Graph., 38(4):1–16, 2019. 2, 3, 5, 6, 7, 8
- [50] Yuxin Wu and Kaiming He. Group normalization. In Eur. Conf. Comput. Vis., pages 3–19, 2018. 6
- [51] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015. 5
- [52] Wenbin Zhu, HsiangTao Wu, Zeyu Chen, Noranart Vesdapunt, and Baoyuan Wang. Reda: Reinforced differentiable attribute for 3d face reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4958–4967, 2020. 2, 4