

# Towards a Pipeline for Real-Time Visualization of Faces for VR-based Telepresence and Live Broadcasting Utilizing Neural Rendering

Philipp Ladwig\*, Rene Ebertowski\*, Alexander Pech\*, Ralf Dörner†, Christian Geiger\*

\*University of Applied Sciences Düsseldorf, Germany  
Mixed Reality and Visualization Group (MIREVI)  
{philipp.ladwig, rene.ebertowski,  
alexander.pech, geiger}@hs-duesseldorf.de  
www.mirevi.de

†RheinMain University of Applied Sciences  
Faculty of Design – Computer Science – Media, Wiesbaden, Germany  
ralf.doerner@hs-rm.de

## Abstract

While head-mounted displays (HMDs) for Virtual Reality (VR) have become widely available in the consumer market, they pose a considerable obstacle for realistic face-to-face conversation in VR since HMDs hide a significant portion of the participants faces. Even with image streams from cameras directly attached to an HMD, stitching together a convincing image of an entire face remains a challenging task because of extreme capture angles and strong lens distortions due to a wide field of view. Compared to the long line of research in VR, reconstruction of faces hidden beneath an HMD is a very recent topic of research. While the current state-of-the-art solutions demonstrate photo-realistic 3D reconstruction results, many of them require high-cost laboratory equipment and

large computational costs. We present an approach that focuses on low-cost hardware and can be used on a commodity gaming computer with a single GPU. We leverage the benefits of an end-to-end pipeline by means of Generative Adversarial Networks (GAN). Our GAN produces a frontal-facing 2.5D point cloud based on a training dataset captured with an RGBD camera. In our approach, the training process is offline, while the reconstruction runs in real-time. Our results show adequate reconstruction quality within the “learned” expressions. Expressions not learned by the network produce artifacts and can trigger the Uncanny Valley effect.

**Keywords:** Neural Rendering, Telepresence, Face Reconstruction, Virtual Reality, Live Broadcasting, Image-to-Image Translation, Pix2Pix, Generative Adversarial Networks

### Digital Peer Publishing Licence

Any party may pass on this work by electronic means and make it available for download under the terms and conditions of the current version of the Digital Peer Publishing Licence (DPPL). The text of the licence may be accessed and retrieved via internet at <http://www.dipp.nrw.de/>.

*First presented at the Workshop of GI Special Interest group VR/AR 2020, extended and revised for JVRB*

## 1 Introduction

Natural face-to-face communication is three-dimensional. A conversation includes not only the verbal communication channel but also the non-verbal channel. In particular, eye contact, facial expressions as well as gestures performed with arms and hands (kinesics), and even the physical distance between each other (proxemics) are essential informa-




Figure 1: *Our conceptual pipeline: First, we capture several RGBD images with a helmet camera mount. These images are processed and serve as the input data for our GAN. After training, the GAN produces textured point clouds in real time. In this work we improve the data set processing, training and inference stage compared to our previous systems [LPG20, LPDG20]. Building a face-tracking HMD is not part of the present work.*

tion carriers during a conversation [LG18]. Currently, common mainstream technologies for computer-mediated communication are video conferencing applications such as Skype or FaceTime. While these allow reading the facial expressions of the counterpart, no “real” eye contact is possible, wide gestures may be cut off in the camera image, deictic gestures are difficult to interpret spatially, and perceptual physical body distance between the participants does not exist.

Current head-mounted displays for VR are capable of delivering believable and immersive 3D experiences including telepresence. However, this does not fully apply to real-time social interactions in VR. When the face of a person is covered by an HMD, it is impossible to read its non-verbal facial communication cues, which is, in fact, a crucial communication channel between individuals. This is not only relevant for face-to-face meetings in VR (for example in VRChat [Inc22], Altspace [Mic22], or Meta’s Horizon Worlds [Pla22]) but also in VR application scenarios in which only one VR user wears an HMD and tries to engage with their audience. For example, such a VR user could be an architect who presents ideas for a new building in VR to their clients, a Virtual YouTuber (VTuber), a Twitch streamer in front of a green screen who broadcasts themselves from inside a VR environment, or friends playing a VR game together in a living room.

The classic way for creating and rendering photo-realistic humans in real time is costly and requires a lot of manual effort such as scanning, modeling, and manual texturing from a skilled 3D artist. Furthermore, today’s HMDs usually lack adequate sensors for face tracking. Only a few research groups have so far addressed this problem and presented methods that can generate authentic face avatars for VR without extensive manual modeling [TZS<sup>+</sup>18, LSSS18, WSS<sup>+</sup>19, RZS<sup>+</sup>21, GPL<sup>+</sup>22]. These approaches are not available to the public, and many of them require expensive

hardware [LSSS18, WSS<sup>+</sup>19].

Human avatars and their perception have been studied in multiple domains. A systematic review of social presence concludes: “...multiple studies show that the vivid perceptions of another person often lead to greater enjoyment and social influence...” [OBW18]. Although, the Media Richness Theory [DL84] is almost 40 years old, recent studies still confirm it [ILC19]. It suggests that the most comprehensive exchange of information between people happens during face-to-face conversations compared to all other digital or analog communication possibilities. A higher quantity and quality of data shared typically leads to more effective communication. A key issue in this context of face-to-face telepresence in VR is the occurrence of the Uncanny Valley effect [MMK12]. Humans are markedly sensitive to minimal and unnatural discrepancies in faces. As soon as a virtual human does not perfectly resemble a real human, it is often subconsciously classified as unlikeable, unpleasant, or even creepy.

One technique that has successfully bridged the Uncanny Valley in recent years is Generative Adversarial Networks (GANs). Today, GANs serve as the core technology behind *Deepfakes*. They enable such authentic results that their methods and algorithms are the subject of current research to distinguish fake images from real ones, as the human eye is no longer able to reliably do so [RCV<sup>+</sup>19]. Therefore, we use algorithms in the context of this work that are also employed to create Deepfakes. We extend this approach with an additional dimension (textured 2.5D point cloud instead of only an RGB image) to generate realistic representations of 2.5D face avatars. We do not create a full 3D head because we only capture the face of a person from a static frontal position with an RGBD sensor. This implies that we do not generate realistic textures from side views. However, we main-

tain a stereoscopic perception of the reconstructed face during face-to-face conversations in a virtual environment.

We present an end-to-end learning system that has low hardware cost compared to others, requires moderate computational resources, and generates results with frame rates suitable for VR applications. Our research contributes to GANs playing a key role in authentic 3D telepresence applications in the near future. In addition to sharing our insights in this paper, we make the code of our prototype publicly available under: <https://github.com/Mirevi/face-synthesizer-JVRB>.

This work is an extension and improvement of our previous neural rendering pipeline [LPG20] and complements our recent work on how to build an HMD with face tracking capabilities [LPDG20]. The contributions of this work are the creation of a face capture pipeline as well as the introduction of novel Generative Adversarial Networks (GAN) [GPAM<sup>+</sup>14] that are tailored for the authentic reconstruction of faces in three-dimensional telepresence and live broadcasting applications. The motivation is to use a commodity graphics card to capture and reconstruct the individual characteristics of a person's face with a high level of (personal) details and VR-enabled frame rates in order to create an authentic avatar that goes beyond the capabilities of today's avatar creation tools such as VRChat [Inc22], Altspace [Mic22] or Meta's Horizon Worlds [Pla22].

## 2 Related Work

Face reconstruction for telepresence and (live) broadcasting with an HMD occluding a person's face is a young research field. Olszewski et al. [OLSL16] presented a system that uses an RGB camera to transfer facial expressions from the lower face area to an avatar. Li et al. [LTO<sup>+</sup>15] extended this approach with pressure sensors placed in the foam of an HMD capturing a person's facial expressions. The idea of using sensors within the HMD is similar to our concept, but we use personalized avatars that are trained in advance and synthesized in a final step.

Casa et al. [CFA<sup>+</sup>16], Früh et al. [FSK17], and Thies et al. [TZS<sup>+</sup>18] used stationary RGBD cameras to create personalized avatars of users. In the first step, the user was captured by the camera without an HMD in order to create a virtual avatar. When the user wore the HMD, the stationary RGBD cam-

era recognized facial expressions. Due to the fixed position of the camera, the range of head motion was limited. Eye movements were registered by eye-tracking cameras and transferred to the user's face avatar. The approaches of Casa et al. [CFA<sup>+</sup>16] and Früh et al. [FSK17] evoked the Uncanny Valley effect to varying degrees. To mitigate this, Früh et al. [FSK17] did not completely remove the HMD, but rendered it as a semi-transparent object. These systems are similar to our approach in the way that they create a personalized avatar using an RGBD camera and produce almost photo-realistic avatars. The approach of Thies et al. [TZS<sup>+</sup>18] demonstrated better results by using a 3D morphable model (3DMM) as underlying head mesh template and as an inductive bias to their system that provides fundamental data about the composition of a human face. They optimize the visual quality by an *analysis-by-synthesis* approach [BV99] and achieve photo realistic results with only a few image artifacts. While the visual quality is convincing, this approach only provides stereoscopic renderings without the ability to freely choose the perspective around the reconstructed face because the final results are based on a given 2D video. Furthermore, it inherently does not allow for manipulation of the head's rotation, scale, and position in the final result.

The systems of Lombardi et al. [LSSS18], Wei et al. [WSS<sup>+</sup>19], and Raj et al. [RZS<sup>+</sup>21] create photo-realistic avatars with authentic facial expressions. While previous works completed the generation of personalized avatars within a few minutes, the system of Lombardi et al. requires computational time of more than a day. The three-dimensional avatar is generated with the aid of a large number of high-resolution images from different angles and facial expressions with an expensive hardware setup that generates a large amount of data for further processing. The created face avatar can be controlled by three RGB cameras attached to an HMD. A key component of this system is the use of *Variational Autoencoders* (VAEs). Both VAEs and GANs have been proven several times to be suitable for authentic face reconstruction. However, since literature shows that VAEs and only a L1 loss tend to produce blurry results more often, we use GANs [HW19]. The latter concept was first presented by Goodfellow et al. [GPAM<sup>+</sup>14], and Radford et al. [RMC15] improved it in a sustainable way. Furthermore, Karras et al. [KALL17] achieved photo-realistic portrait images that are indistinguishable from

real photographs by using the principle of *Progressive Growing GAN*. However, according to Karras et al., the GAN has little to no external control over the appearance of the generated object or face because the input to the network is a latent vector without any direct relation to a face property such as hair color, facial expression, or gender. In further works Karras et al. [KLA18] enhanced the architecture of the GAN and were able to automatically separate higher-level attributes (e.g. pose, identity) from stochastic variations (e.g. freckles, hair). Nevertheless, this approach does not allow to explicitly control the facial expression.

Conditional GANs (cGANs) have been shown to be able to learn and reproduce specific relationships between inputs and outputs that are understandable for humans. For example, Mirza and Osindero [MO14] have extended the input to the generator and discriminator with a label  $y$ , which makes it possible to generate images from a particular category  $y$ . This method for conditioning GANs was developed further by Radford et al. [RMC15] with the DCGAN and by Isola et al. [IZZE17] with the Pix2Pix GAN. They replaced the noise input vector  $z$  with a user-defined input vector. Without a noise vector, there is no latent space  $Z$  (since  $z \in Z$ ). If the stochastic aspect contained in the noise vector is not compensated, the GAN will only memorize the training examples. Any inputs that deviate from the training data would lead to inadequate results, as described by Isola et al. [IZZE17]. By using a U-net architecture [RFB15] with dropouts in the Pix2Pix GAN, the stochastic aspect as well as the missing latent space can be otherwise integrated into the generator. The discriminator of the Pix2Pix GAN receives the same input image  $x$  as the generator as well as its output image  $y_{fake} = G(x)$  or the image  $y_{real}$  matching  $x$  from the dataset. This is basically equivalent to the idea of cGANs [MO14] where not only the output of the generator is evaluated but also its difference from the input. Unlike the cGAN, the output of the discriminator of the Pix2Pix GAN is not a scalar that decides between “real” or “false” but a matrix. By using convolutional layers (cf. Radford et al. [RMC15]), each entry in the output matrix represents an  $n * m$ -sized region (so-called patch) of the input image. This allows abstract representations to be admitted as matrices (e.g. images) for conditioning the network to have a controlled influence on the output of the generator. This approach was further developed by [WLZ<sup>+</sup>18] with the Pix2PixHD GAN to generate im-

ages with a higher resolution and more details. In this paper, we adapt the idea of cGANs, especially of the Pix2Pix and Pix2PixHD frameworks, and tailor them to our application domain.

### 3 System

In the following, we explain the process and structure of the proposed system, as shown in Fig. 1, and then discuss the steps in more detail in the subsequent sections.

Our process starts with the acquisition of a personal RGBD dataset. The acquired data is preprocessed by an automated procedure. A Facial Landmark Map (FLM) per RGB image is extracted and saved beside the corresponding RGB image. It decodes the facial expression of the respective RGB image in a binary image as so-called landmarks as shown in the second image from left in Fig.1. Our proposed GAN is trained with the captured RGBD images as well as with the corresponding FLM. For each person our GAN must be trained from scratch. We do not use any inductive biases like a 3DMM [BV99] and the system does not learn correspondences between persons. After the training, the system can be used for real-time telepresence or live broadcasting. In the VR application scenario, the user would wear a face-tracking head-mounted display that could create an FLM in real time, which we then feed into the trained generator module of our GAN. This paper does not focus on building and implementing a face-tracking HMD. Further hardware-related implementation details are described in our previous work [LPDG20]. The GAN could create an RGB and a D image of the “learned” person based on the FLM, and finally, we fuse the generated RGB and D images into a textured point cloud.

#### 3.1 Training Data

An RGBD dataset of the respective person forms the basis for the training process. For the acquisition process, an RGBD camera is mounted in a fixed position to a self-made helmet mount, as shown in Figure 1 left and Figure 2. This mount ensures that the entropy in the dataset is as minimal as possible. Varying distances, positions and head rotations do not contribute to “learning” the user’s face. By using the helmet mount, we ensure that the training time of the networks is short and the reconstruction quality of the final results is high, as we learned from previous experiments




Figure 2: *Helmet camera mount for RGBD data acquisition. This mount ensures that head rotations are not included in the training dataset and, therefore, reduces the entropy in the dataset. Moreover, this method significantly reduces the training time and increases the visual quality of the output images. The material price of the helmet mount without RGBD camera is about 60 USD.*

conducted without the helmet mount.

The RGBD sensor in the mount captures facial expressions and stores the color information as a 3-channel image file (8-bit for each channel) with  $2048 \times 1536$  pixels, while it stores the depth information as a 1-channel image file (16-bit grayscale) with  $640 \times 576$  pixels. During the capture process, the user is encouraged to perform a variety of different facial expressions. A dataset of a person should contain about 1500 to 2000 RGBD images, and the acquisition process takes about 10 minutes. The data is then pre-processed for further steps. First, the foreground and background pixels are clipped in front of and behind the face and the depth resolution is reduced from 16 to 8 bit, which reduces the depth range from 65,535 to 255 mm. This speeds up the training process of the GAN and significantly reduces depth noise. With the help of the helmet mount, we ensure that a depth range of 255 mm is sufficient for spatially reproducing the frontal part of the head since the distance between the sensor and the face in the helmet mount is always fixed. Furthermore, the data is normalized and the image data is converted into values between -1 and 1 for the training process. Contrasts are sharpened by normalizing the histogram.

### 3.2 Determination of Facial Landmarks

To control the output of the generator and thus the facial expressions of a person's face, the dataset must be labeled before the training process. We use 70 facial landmarks (68 of the Multi-PIE scheme [GMC<sup>+</sup>08] and two for the location of the irises) as binary images for each tuple of RGB and corresponding depth im-

age in the dataset. As mentioned before, we call these binary images Facial Landmark Maps (FLM). The position of the landmarks identifies the expressions of the person in each image of the dataset. To determine the landmarks, we use the *Face Alignment Network* (FAN) of Bulat and Tzimiropoulos [BT17]. The FAN is not able to determine the position of the person's pupils within an image. Therefore, two additional landmarks were implemented based on an eye-tracking procedure by Timm and Barth [TB11, jon18]. We experimented with an Tobii Eye Tracker 4C, but users reported that additional weight on the helmet mount felt uncomfortable during the capture procedure. Both the tracking results of the Tobii Eye Tracker and Timm and Barth were similar and sufficient for our application.

When the landmarks in each image had been located, a rectangle of the maximal and minimal locations of the landmarks was created and the RGB and D images were cropped to this area and resized to  $512 \times 512$  pixels.

### 3.3 Neural Network Architecture

Previous work with neural networks, such as Wu et al. [WZX<sup>+</sup>16], has shown that a voxel-based approach is associated with high training and execution times of the model. Therefore, an RGBD-based solution was targeted. The advantage of this approach lies in the compact representation of the data as a point cloud and the possibility to adapt previous RGB-based methods. Our underlying network architecture is derived from the Pix2Pix GAN by Isola et al. [IZZE17]. In earlier experiments, we discovered that this architecture was able to produce acceptable results [LPG20, LPDG20], but the images have a low resolution of  $256 \times 256$  pixels, often lack details in high frequency areas such as facial hair and tend to produce time-inconsistent reconstructions with noise. Therefore, we experimented with the Pix2PixHD framework [WLZ<sup>+</sup>18], which is an extension of the Pix2Pix GAN [IZZE17]. Although it produces images with a higher resolution and better quality, the inference time is not capable of retaining real-time frame rates on commodity hardware. Therefore, we kept the Pix2Pix framework as a base and gradually added elements from the Pix2PixHD framework and further improvements from other works until we obtained an acceptable image quality with a reasonable processing speed for interactive frame rates. In summary, we propose the following changes to the Pix2Pix



framework:

1. We added a multi-scale discriminators that receives three different resolutions of the input image and an additional *Feature Matching Loss* as described in Pix2PixHD [WLZ<sup>+</sup>18].
2. We changed the *Sigmoid Cross Entropy Loss* of Pix2Pix's discriminator to the *Least-Squares Loss* of the LSGAN [MLX<sup>+</sup>17] as suggested by Wang et al. [WLZ<sup>+</sup>18].
3. We exchanged the *Perceptual-VGG Loss* [JAFF16] originally suggested by Wang et al. [WLZ<sup>+</sup>18] with the better performing *Learned Perceptual Image Patch Similarity* (LPIPS) by Zhang et al. [ZIE<sup>+</sup>18].

For the discriminator we obtain the following objective:

$$\min_{D_1, D_2, D_3} V_{GAN}(D) = \sum_{k=1,2,3} \mathcal{L}_{cLSGAN\_D}(D_k, G) \quad (1)$$

where  $D_1$ ,  $D_2$ , and  $D_3$  describe the three different resolution of the input image. For the generator we end up with the following objective function:

$$\begin{aligned} \min_G V_{GAN}(G) = & \\ & \sum_{k=1,2,3} \left[ \mathcal{L}_{cLSGAN\_G}(D_k, G) + \lambda_{FM} \mathcal{L}_{FM}(D_k, G) \right] \\ & + \lambda_{L1} \mathcal{L}_{L1}(G) + \lambda_{LPIPS} \mathcal{L}_{LPIPS}(y, G(x)) \end{aligned} \quad (2)$$

where we choose the following hyper-parameters:  $\lambda_{FM} = 10$ ,  $\lambda_{L1} = 100$ ,  $\lambda_{LPIPS} = 10$ . The functions  $\mathcal{L}_{cLSGAN\_D}$  and  $\mathcal{L}_{cLSGAN\_G}$  can be found in Mao et al. [MLX<sup>+</sup>17], whereas the function  $\mathcal{L}_{FM}$  is based on the feature matching loss of Wang et al. [WLZ<sup>+</sup>18]. In order to maintain faster inference time than the Pix2PixHD we did not implement the coarse to fine approach for the generator. We sacrifice very high resolution of the output images for computational speed. We provide further details about the architecture on GitHub.

Using these improvements helps to prevent high-frequency details such as facial hair and significantly increases the reconstruction quality, as explained further in the results section (sec. 4). Because we mainly enhanced the loss function and the discriminator side, we did not need to change the generator. Therefore,




Figure 3: Example convolution for the discriminator input. Each RGBD channel and the FLM are weighted individually.

we are able to maintain high frame rates during inference since only the generator module is used in the telepresence and live broadcasting scenario. As a side effect, the training process requires more memory, but the overall training time decreases by more than a half compared to our Pix2Pix-only approach (from about 19 hours to 8 hours) because the new loss term is more purposeful for our application and, therefore, helps to obtain better results in less time.

The generator of our GAN receives a  $512 \times 512$  pixel FLM of the facial landmarks as input. Compared to the Pix2Pix GAN, the output has been extended by a fourth feature map to be able to generate depth images. In addition, the discriminator receives five feature maps as input instead of only four. While the first four correspond to the four channels of the RGBD image, the remaining feature map contains the corresponding FLM, as visualized in Fig. 3. One of our early hypotheses in [LPG20] was that a higher number of FLMs (e.g. four times) would result in better reconstruction results to start the training process with a balanced ratio between RGBD images and FLMs (cf. Fig. 3). This hypothesis has been disproved. Changing the number of FLMs does not change the quality of the results but only increases the training time.

### 3.4 Training Process

For each person, the discriminator and the generator are each completely trained from scratch. Before the training process, all weights of the two networks are initialized with a random value. The random values follow a Gaussian distribution with an expected value of 0 and a standard deviation of 0.02. The batch size of the input data into the GAN is 1, and the epoch count

is 100. Both generator and discriminator are trained with an initial learning rate of 0.0002, which decreases linearly towards zero over the last 70 epochs. The discriminator  $Loss_D = (Loss_{D_{real}} + Loss_{D_{fake}}) * 0.5$  reduces the learning rate, making it slower to learn relative to the generator. This is necessary because at the beginning of the training phase the discriminator can effortlessly accomplish its task. If the discriminator learns too quickly, the generator has no chance to learn how to create the desired face.

## 4 Results

GANs are difficult to train because of their adversarial training procedure to find the balance between the learning rate and losses of the generator and discriminator networks. As we use many different losses, a long line of hyper parameter tuning was necessary to find the optimal settings. Fig. 5 shows the reconstruction results for four different persons with the best training parameters described in section 3.4.

We did not use the face-tracking HMD for the evaluation, similar to [LPDG20], because it can cause slight tracking errors and could decrease the quality of the results for comparison. Our intention is to directly compare the improvements of our pipeline to our previous system described in [LPG20]. For a comparison in motion, we refer the reader to the corresponding video that can be found via the GitHub link.

As a quantitative metric for the assessment of the reconstruction quality, we use *Structural Similarity* (SSIM) [WBSS04] and *Learned Perceptual Image Patch Similarity* (LPIPS) [ZIE<sup>+</sup>18]. Our previous system [LPG20] performed on average with an SSIM of 0.851 and a value of 0.114 on LPIPS. Our proposed system reaches on average 0.910 for SSIM (higher is better) and 0.082 for LPIPS (lower is better). The numerical results of the SSIM metric are comparable to a JPEG compression of about half the original file size of the images in column 3. In contrast, the previous system [LPG20] only achieved reconstruction quality of less than a quarter of the original file size by means of a comparison with the JPEG compression.

All results and metrics shown in Fig. 5 and 6 are generated from LFMs from the test set. This means, no backpropagation was conducted on the test set. The train/test split was 85% to 15%. The dataset for the first participant in Fig. 5 and 6 was split 1238/207, for the second 1500/250, for the third 1620/271 and for last 2413/403. Please note that the dataset for the last




Figure 4: *Our new pipeline, network architecture and losses significantly improved the quality. Image a) shows a sample from the previous system of Ladwig et al. [LPG20, LPDG20]. Image b) illustrates the enhanced resolution (from  $256 \times 256$  to  $512 \times 512$  pixels) and the improved preservation of high-frequency details.*

participant was around 1000 items larger compared to the previous participants and lead to slightly better quantitative results for the best and worst result according to SSIM and LPIPS (the last two images per participant). This allows the conclusion that a larger data set tends to lead to better image quality in our scenario.

The main issue of our previous approach was the sharpness of the generated images [LPG20, LPDG20], as shown in Fig. 4. Due to the new architecture and loss functions, the system produces images with more details. Even skin pores are reconstructed well, e.g. on the user's forehead, as can be seen in Fig. 5, rows E to H. Furthermore, we noticed a better reconstruction quality in areas with high-frequency details such as facial hair, as shown in Fig. 4. Also, the temporal consistency is improved. Please see the linked video on our GitHub page for details.

The error between the generated and ground truth depth values is mostly below 4 mm, as depicted in column 6. Exceptions are the reconstruction results with the worst SSIM and LPIPS metrics per dataset of a person, such as shown in Fig. 5, row D and H, as well as in Fig. 6, row E and J. Furthermore, outliers can be

seen in column 8. Note that we use the raw depth image of the Kinect and the raw output of our GAN. We do not filter or smooth the images, therefore, we assume that the network has also learned the depth noise of the sensor, which causes additional depth errors.

To compare the faces between the images of columns 2 and 3 without measuring background changes, we determined the facial area based on depth values and rejected all other pixels. At the border area between the faces and the background, large differences in the SSIM and the depth difference visualization can be seen. These differences are caused by the fact that the cropped areas of the real and generated images do not always align perfectly due to minimal differences in the generated faces. In addition, we also apply erosion and clipping to the face to reject parts of the background, which can cause the minimal differences.

Although our quantitative metrics indicate better results, our system still shows limitations in the reconstruction quality of the eyes, lips and oral cavity. Especially teeth and the tongue are often reconstructed with noisy artifacts, as can be seen in Fig. 5, row D, column 2. The error increases with a graceful degradation when the expression moves towards exaggerated expressions that are far from the neutral face expression. The eyes are reconstructed with less artifacts than the oral cavity, but we observed that even a little amount of image artifacts can evoke the Uncanny Valley Effect, as can be seen especially in Fig. 5, row B, column 2. Please zoom in for details.

In our experiments we used PyTorch 1.8 and Python 3.7 on Windows 10. We trained the same data set with approximately 1600 elements (an element is a set comprised of an RGBD image and an FLM) on two different machines. This took 8 hours on a system equipped with an Nvidia RTX2080Ti, an AMD Ryzen Threadripper 2990WX and 128GB RAM. With a newer system, comprised of an Nvidia RTX3090, an Intel i7-9900K and 32GB RAM, training took only 4 hours and 13 minutes. Our previous system took 17-20 hours on the machine with the RTX2080Ti for only 600 elements.

The time required for a forward pass of the generator module with an image size of  $512 \times 512$  pixels is between 3 and 4 ms (333–250 fps) on an Nvidia RTX3090 and between 6 and 7 ms (167–143 fps) on an Nvidia RTX 2080. Our previous system was faster (between 1 and 3 ms) but generates only images with a size of  $256 \times 256$  pixels. The timings of the present

system are still suitable for VR-based application, where 75 to 120 fps are common frame rates. However, note that many face tracking systems only work with 60 or even 30 fps, which can limit the frame rate of the pipeline.

## 5 Limitations and Future Work

As mentioned before, one of the major issues is the low reconstruction quality of exaggerated expressions of the eyes, lips and oral cavity. The artifacts can induce the Uncanny Valley Effect and must be avoided for telepresence or broadcasting applications. As a further step, we plan to use a 3DMM [BV99] such as the FLAME model [LBB<sup>+</sup>17] as a better inductive bias to regularize depth and color information more efficiently. Furthermore, we observed that landmark tracking is not sufficient for faithful lip movement during speech. Therefore, an additional input signal besides the landmarks is necessary. A conditioning of speech as audio signals could provide a solution. Another issue to improve is the uncomfortable helmet mount. A solution with a stationary RGBD camera placed on a tripod or table is a favorable approach for future research. Although, we achieve real-time frame rates on a gaming GPU, the generator network could still be too slow for using on a current stand alone headset like a Meta Quest 2 or a Vive XR Elite. Further experiments, hardware and performance improvements could solve this problem in the future.

## 6 Conclusion

We presented an improved end-to-end pipeline compared to previous approaches [LPG20, LPDG20] and a new GAN architecture that can learn facial identity and individual expressions of a user and reproduce them as a textured point cloud with frame rates that are suitable for Virtual Reality, telepresence and broadcasting environments. We have incorporated and extended the architecture, losses, and processing pipelines of several approaches from the field of neural rendering. Compared to previous works, our proposed system generates higher quality image results with slightly longer run time at inference. We achieved this goal by mainly changing the architecture of the discriminator while keeping the architecture of the generator lean. The reconstruction results partially lie in the Uncanny Valley, but they still con-



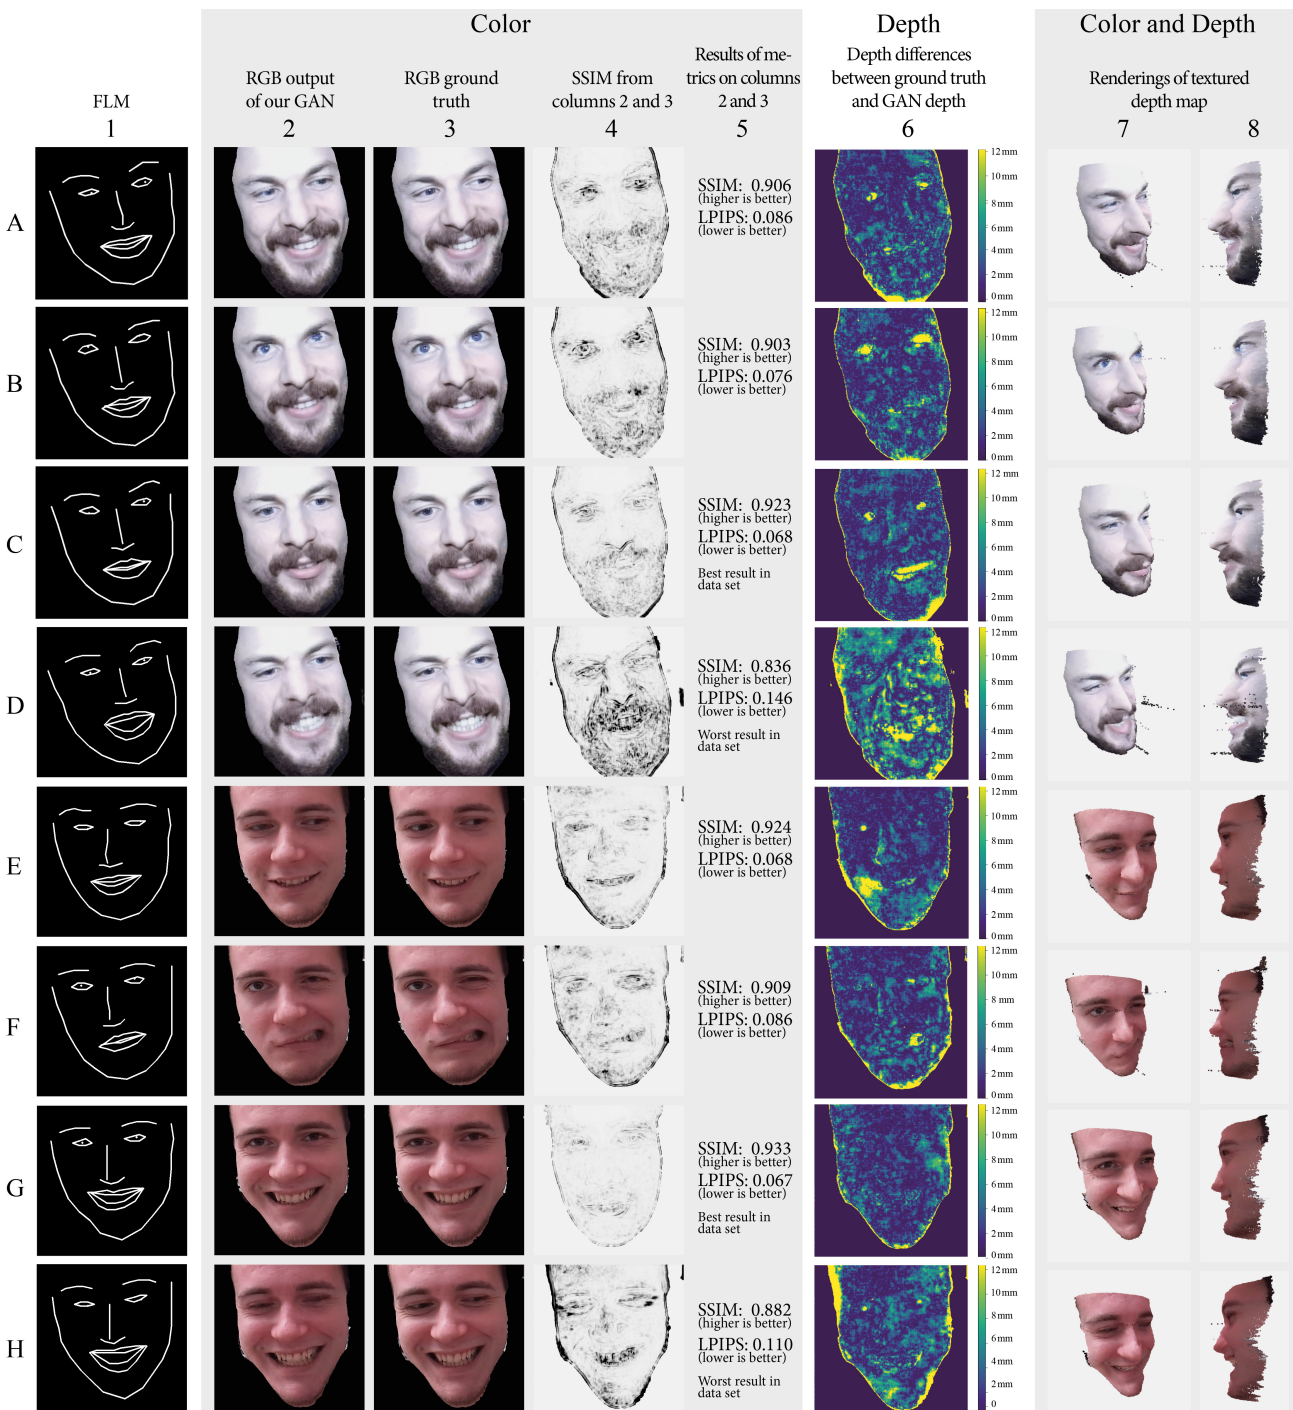


Figure 5: Results 1/2: This overview shows FLMs in column 1 from our evaluation datasets, so the results are based on unseen data for the neural network. The FLMs were created from the images in column 3 – a real image from the evaluation data set. Column 2 shows the results generated by our GAN. The GAN received the FLM from column 1 and generated the images in column 2. Column 4 depicts the SSIM difference. Darker values indicate larger differences between the images in columns 2 and 3. Column 6 visualizes the error between the generated depth and the ground truth depth. The combination of the generated depth and color data can be seen in columns 7 and 8 from an angle of 30 and 90 degrees.

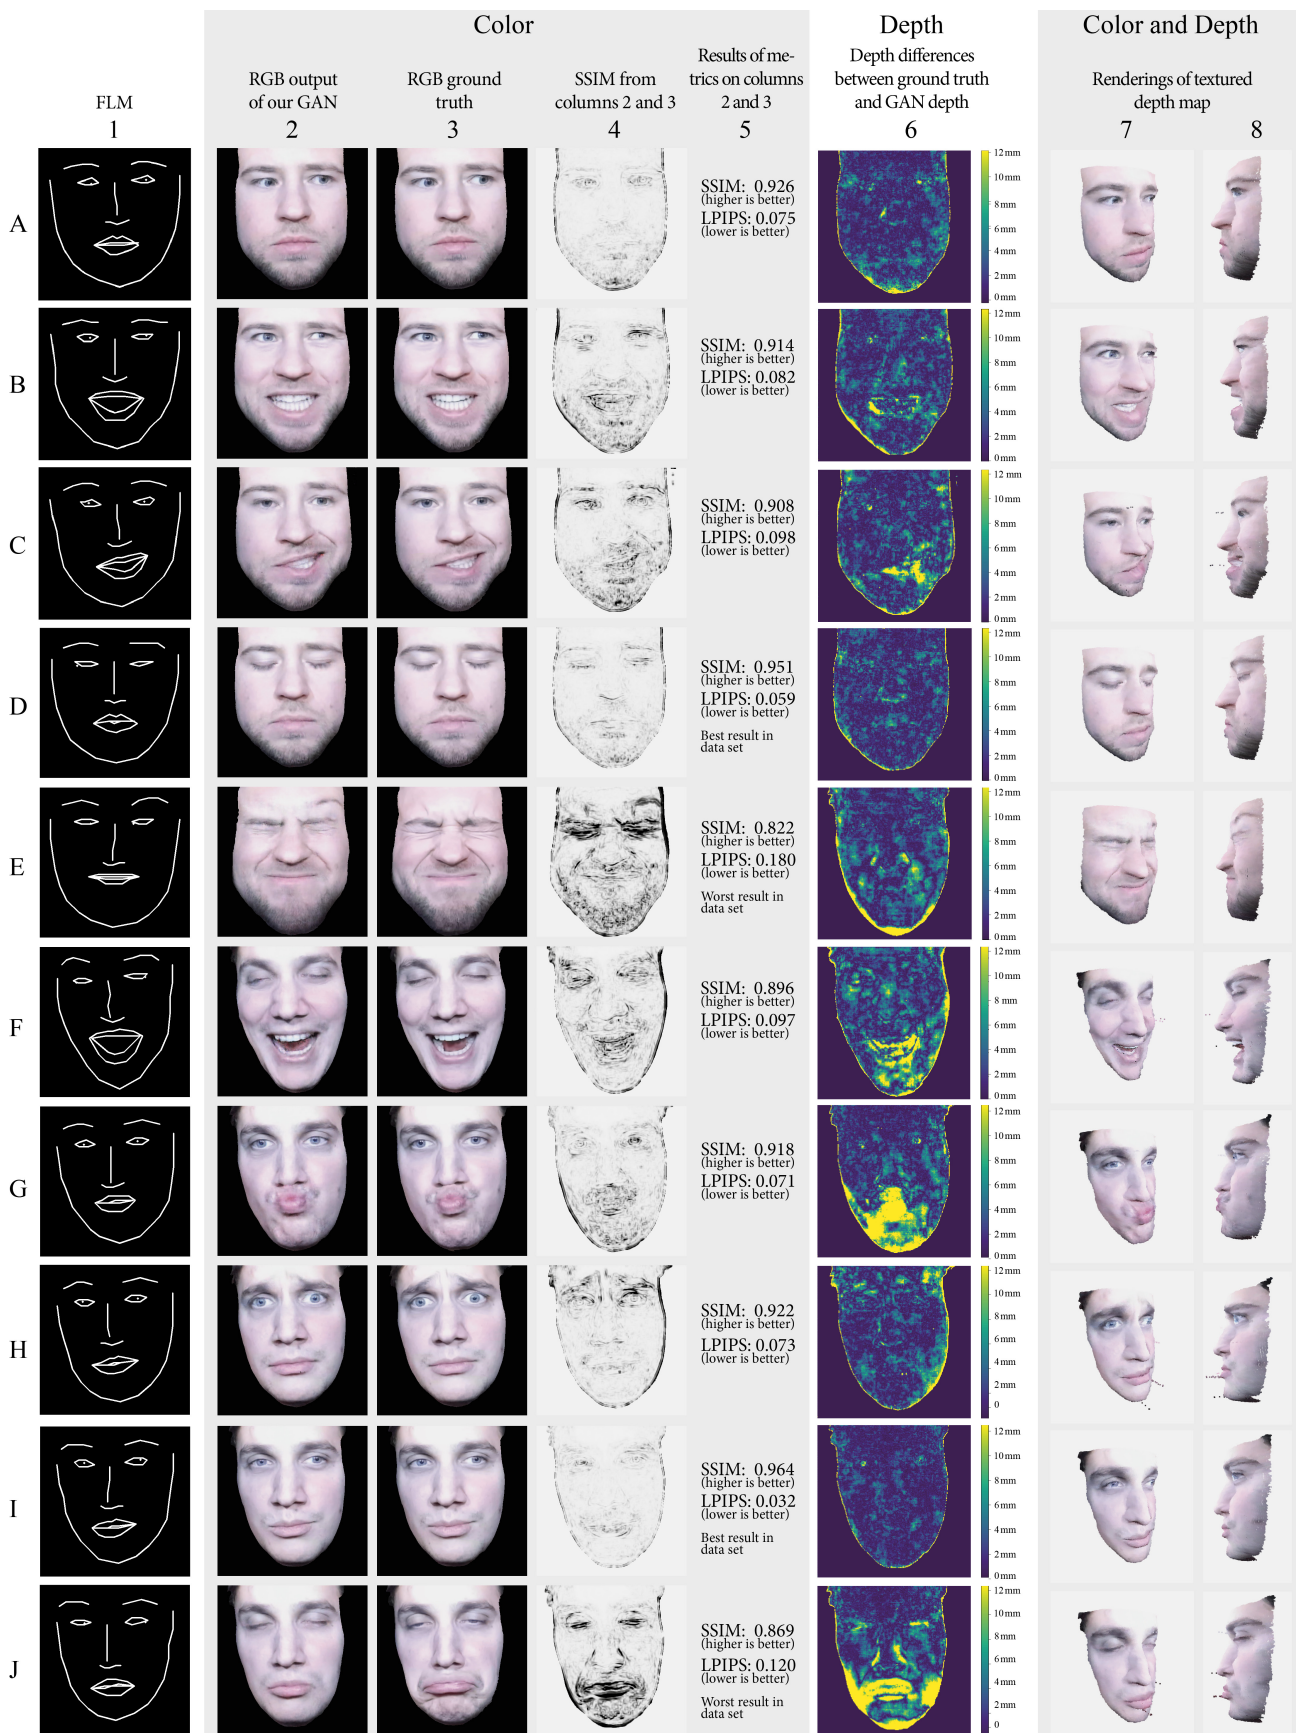


Figure 6: Results 2/2: Two further subjects are shown. The two last samples of each person summarizes the best and the worst images (measured by SSIM and LPIPS) and reflects the range of the reconstruction quality.

vince with an authentic visualization of the respective person's identity and individual facial expressions. We believe that neural rendering will be a crucial part of photo-realistic rendering of humans in real-time applications in the future. Our work is a further step into this direction and hopefully helps to understand, improve and apply this technology.

## 7 Acknowledgments

We thank the MIREVI group at the University of Applied Sciences Düsseldorf and the *Promotionszentrum Angewandte Informatik (PZAI)* in Hessen, Germany. This work is sponsored by the German Federal Ministry of Education and Research (BMBF) under the project numbers 16SV8182 (HIVE-Lab), 13FH022IX6 (iKPT 4.0) and 16SV8756 (AniBot).

## References

- [BT17] Adrian Bulat and Georgios Tzimiropoulos, *How Far are We from Solving the 2D and 3D Face Alignment Problem?*, 2017 IEEE International Conference on Computer Vision (ICCV) (Piscataway, NJ), IEEE, Oct 2017, DOI 10.1109/iccv.2017.116, ISBN 9781538610329.
- [BV99] Volker Blanz and Thomas Vetter, *A Morphable Model for the Synthesis of 3D Faces*, SIGGRAPH99: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (New York, NY), SIGGRAPH, ACM Press/Addison-Wesley Publishing Co., 1999, DOI 10.1145/311535.311556, pp. 187–194, ISBN 978-0-201-48560-8.
- [CFA<sup>+</sup>16] Dan Casas, Andrew Feng, Oleg Alexander, Graham Fyffe, Paul Debevec, Ryosuke Ichikari, Hao Li, Kyle Olaszewski, Evan Suma, and Ari Shapiro, *Rapid Photorealistic Blendshape Modeling from RGB-D Sensors*, CASA'16: Computer Animation and Social Agents (New York, NY), Association for Computing Machinery, 2016, DOI 10.1145/2915926.2915936, pp. 121–129, ISBN 978-1-4503-4745-7.
- [DL84] R. L. Daft and R. H. Lengel, *Information Richness: A New Approach to Managerial Behavior and Organization Design*, *Research in Organizational Behavior* **6** (1984), 191–233, ISSN 0191-3085.
- [FSK17] Christian Frueh, Avneesh Sud, and Vivek Kwatra, *Headset Removal for Virtual and Mixed Reality*, SIGGRAPH '17: ACM SIGGRAPH 2017 Talks (New York, NY), no. Article no. 80, Association for Computing Machinery, 2017, DOI 10.1145/3084363.3085083, ISBN 978-1-4503-5008-2.
- [GMC<sup>+</sup>08] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker, *Multi-PIE*, 2008 8th IEEE International Conference on Automatic Face Gesture Recognition (Piscataway, NJ), IEEE, 2008, DOI 10.1109/AFGR.2008.4813399, ISBN 978-1-4244-2153-4.
- [GPAM<sup>+</sup>14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, *Generative Adversarial Nets*, *Advances in Neural Information Processing Systems (NIPS 2014)* (Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, eds.), vol. 27, Curran Associates, Inc., 2014, Available via proceedings.neurips.cc.
- [GPL<sup>+</sup>22] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies, *Neural head avatars from monocular rgb videos*, June 2022, DOI 10.48550/arXiv.2112.01554, pp. 18653–18664.
- [HW19] John Hany and Greg Walters, *Hands-On Generative Adversarial Networks with PyTorch 1.x*, Packt Publishing, Birmingham; Mumbai, 2019, ISBN 9781789534283.

- [ILC19] Kumi Ishii, Mary Madison Lyons, and Sabrina A. Carr, *Revisiting media richness theory for today and future*, Human Behavior and Emerging Technologies **1** (2019), no. 2, 124–131, ISSN 2578-1863, DOI 10.1002/hbe2.138.
- [Inc22] VRChat Inc., *VRChat*, <https://hello.vrchat.com/>, 2022, Accessed: 2022-03-04.
- [IZZE17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros, *Image-to-Image Translation with Conditional Adversarial Networks*, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (Piscataway, NJ), IEEE, 2017, DOI 10.1109/CVPR.2017.632, pp. 5967–5976.
- [JAFF16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei, *Perceptual Losses for Real-Time Style Transfer and Super-Resolution*, Computer Vision – ECCV 2016 (Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, eds.), Lecture Notes in Computer Science, Cham : Springer International Publishing, 2016, DOI 10.1007/978-3-319-46475-6\_43, pp. 694–711, ISBN 978-3-319-46474-9.
- [jon18] jonnedtc, *Locating eye centers using means of gradients*, <https://github.com/jonnedtc/PupilDetector>, 2018, Accessed 2021-12-28.
- [KALL17] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen, *Progressive Growing of GANs for Improved Quality, Stability, and Variation*, 2017, DOI 10.48550/arXiv.1710.10196.
- [KLA18] Tero Karras, Samuli Laine, and Timo Aila, *A Style-Based Generator Architecture for Generative Adversarial Networks*, 2018, Last Accessed July 31st, 2023 from <https://arxiv.org/abs/1812.04948>, DOI 10.48550/arXiv.1812.04948.
- [LBB<sup>+</sup>17] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero, *Learning a model of facial shape and expression from 4D scans*, ACM Transactions on Graphics **36** (2017), no. 6, ISSN 0730-0301, Article No. 194, DOI <https://doi.org/10.1145/3130800.3130813>.
- [LG18] Philipp Ladwig and Christian Geiger, *A Literature Review on Collaboration in Mixed Reality*, Smart industry & smart education: proceedings of the 15th International Conference on Remote Engineering and Virtual Instrumentation, Lecture Notes in Networks and Systems, vol. 47, Springer International Publishing, 2018, DOI 10.1007/978-3-319-95678-7\_65, pp. 591–600, ISBN 978-3-319-95677-0.
- [LPDG20] Philipp Ladwig, Alexander Pech, Ralf Dörner, and Christian Geiger, *Unmasking Communication Partners: A Low-Cost AI Solution for Digitally Removing Head-Mounted Displays in VR-Based Telepresence*, 2020 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR) (Piscataway, NJ), IEEE, 2020, DOI 10.1109/AIVR50618.2020.00025, ISBN 978-1-72817-463-1.
- [LPG20] Philipp Ladwig, Alexander Pech, and Christian Geiger, *Auf dem Weg zu Face-to-Face-Telepräsenz Anwendungen in Virtual Reality mit generativen neuronalen Netzen*, GI VR/AR Workshop (Benjamin Weyers, Christoph Lürig, and Daniel Zielasko, eds.), Gesellschaft für Informatik e.V., 2020, DOI 10.18420/vrar2020\_15.
- [LSSS18] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh, *Deep Appearance Models for Face Rendering*, ACM Transactions on Graphics **37** (2018), no. 4, ISSN 0730-0301, Article No. 68, DOI 10.1145/3197517.3201401.
- [LTO<sup>+</sup>15] Hao Li, Laura Trutoiu, Kyle Olaszewski, Lingyu Wei, Tristan Trutna,



- Pei-Lun Hsieh, Aaron Nicholls, and Chongyang Ma, *Facial Performance Sensing Head-Mounted Display*, ACM Transactions on Graphics **34** (2015), no. 4, ISSN 0730-0301, Article No. 47, DOI 10.1145/2766939.
- [Mic22] Microsoft, *AltspaceVR*, <https://altvr.com/>, 2022, Accessed: 2022-03-04.
- [MLX<sup>+</sup>17] Xudong Mao, Qing Li, Haoran Xie, Raymond K. Lau, Zhen Wang, and Stephen Paul Smolley, *Least Squares Generative Adversarial Networks*, 2017 IEEE International Conference on Computer Vision (ICCV) (Piscataway, NJ), IEEE, 2017, DOI 10.1109/ICCV.2017.304, pp. 2813–2821, ISBN 9781538610336.
- [MMK12] Masahiro Mori, Karl F. MacDorman, and Norri Kageki, *The Uncanny Valley [From the Field]*, IEEE Robotics & Automation Magazine **19** (2012), no. 2, 98–100, ISSN 1070-9932, DOI 10.1109/MRA.2012.2192811.
- [MO14] Mehdi Mirza and Simon Osindero, *Conditional Generative Adversarial Nets*, 2014, Accessed: 2020-07-07, DOI 10.48550/arXiv.1411.1784.
- [OBW18] Catherine Oh, Jeremy Bailenson, and Gregory Welch, *A Systematic Review of Social Presence: Definition, Antecedents, and Implications*, Frontiers in Robotics and AI **5** (2018), ISSN 2296-9144, DOI 10.3389/frobt.2018.00114.
- [OLSL16] Kyle Olszewski, Joseph J. Lim, Shunsuke Saito, and Hao Li, *High-Fidelity Facial and Speech Animation for VR HMDs*, ACM Transactions on Graphics **35** (2016), no. 6, ISSN 0730-0301, Article No. 221, DOI 10.1145/2980179.2980252.
- [Pla22] Meta Platforms, *Meta Horizon Worlds*, [https://www.oculus.com/horizon-worlds/?locale=de\\_DE](https://www.oculus.com/horizon-worlds/?locale=de_DE), 2022, Accessed: 2022-03-04.
- [RCV<sup>+</sup>19] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Niessner, *FaceForensics++: Learning to Detect Manipulated Facial Images*, 2019 IEEE/CVF International Conference on Computer Vision (ICCV) (Piscataway, NJ), IEEE, 2019, DOI 10.1109/ICCV43118.2019, ISBN 978-1-7281-4803-8.
- [RFB15] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, *U-Net: Convolutional Networks for Biomedical Image Segmentation*, Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015 (Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro Frangi, eds.), LNIP, vol. 9351, Springer International Publishing, 2015, DOI 10.1007/978-3-319-24574-4\_28, ISBN 978-3-319-24573-7.
- [RMC15] Alec Radford, Luke Metz, and Soumith Chintala, *Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks*, 2015, Accessed: 2022-07-07, DOI 10.48550/arXiv.1511.06434.
- [RZS<sup>+</sup>21] Amit Raj, Michael Zollhöfer, Tomas Simon, Jason Saragih, Shunsuke Saito, James Hays, and Stephen Lombardi, *Pixel-Aligned Volumetric Avatars*, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (Piscataway, NJ), IEEE, 2021, DOI 10.1109/CVPR46437.2021.01156, ISBN 978-1-6654-4509-2.
- [TB11] Fabian Timm and Erhardt Barth, *Accurate Eye Centre Localisation by Means of Gradients*, Proceedings of the International Conference on Computer Vision Theory and Applications (VISIGRAPP 2011) - VISAPP (Algarve, Portugal) (José Braz and Leonid Mestetskii, eds.), SciTePress, 2011, DOI 10.5220/0003326101250130, pp. 125–130, ISBN 978-989-8425-47-8.



- [TZS<sup>+</sup>18] Justus Thies, Michael Zollhöfer, Marc Stamminger, Christian Theobalt, and Matthias Nießner, *FaceVR: Real-Time Gaze-Aware Facial Reenactment in Virtual Reality*, ACM Transactions on Graphics **37** (2018), no. 2, Article No. 25, DOI 10.1145/3182644.
- [WBSS04] Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, *Image quality assessment: from error visibility to structural similarity*, IEEE Transactions on Image Processing **13** (2004), no. 4, 600–612, ISSN 1057-7149, DOI 10.1109/TIP.2003.819861.
- [WLZ<sup>+</sup>18] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro, *High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs*, 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (Piscataway, NJ), IEEE, 2018, DOI 10.1109/CVPR.2018.00917, pp. 8798–8807, ISBN 9781538661017.
- [WSS<sup>+</sup>19] 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 Transactions on Graphics **38** (2019), no. 4, ISSN 0730-0301, Article No. 67, DOI 10.1145/3306346.3323030.
- [WZX<sup>+</sup>16] Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T. Freeman, and Joshua B. Tenenbaum, *Learning a Probabilistic Latent Space of Object Shapes via 3D Generative-Adversarial Modeling*, Advances in neural information processing systems 29 (Red Hook, NY) (David D. Lee, Ulrike von Luxburg, Roman Garnett, Masashi Sugiyama, and Isabelle Guyon, eds.), Curran Associates Inc., 2016, pp. 82–90, ISBN 978-1-5108-3881-9.
- [ZIE<sup>+</sup>18] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang, *The Unreasonable Effectiveness of Deep Features as a Perceptual Metric*, 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (Piscataway, NJ), IEEE, 2018, DOI 10.1109/CVPR.2018.00068, ISBN 978-1-5386-6420-9.

Citation
Philipp Ladwig, Rene Ebertowski, Alexander Pech, Ralf Dörner, and Christian Geiger, <i>Towards a Pipeline for Real-Time Visualization of Faces for VR-based Telepresence and Live Broadcasting Utilizing Neural Rendering</i> , Journal of Virtual Reality and Broadcasting, 18(2024), no. 1, December 2024, DOI 10.48663/1860-2037/18.2024.1, ISSN 1860-2037.