

Neural Head Avatars from Monocular RGB Videos

Philip-William Grassal^{1*}

Malte Prinzler^{1*}

Titus Leistner¹

Carsten Rother¹


Matthias Nießner²

Justus Thies³

¹Heidelberg University

²Technical University of Munich

³Max Planck Institute for Intelligent Systems



Monocular RGB Video

Neural Head Avatar with Articulated Geometry & Photorealistic Texture

Figure 1. Given a monocular portrait video of a person, we reconstruct a *Neural Head Avatar*. Such a 4D avatar will be the foundation of applications like teleconferencing in VR/AR, since it enables novel-view synthesis and control over pose and expression.

Abstract

We present *Neural Head Avatars*, a novel neural representation that explicitly models the surface geometry and appearance of an animatable human avatar that can be used for teleconferencing in AR/VR or other applications in the movie or games industry that rely on a digital human.¹ Our representation can be learned from a monocular RGB portrait video that features a range of different expressions and views. Specifically, we propose a hybrid representation consisting of a morphable model for the coarse shape and expressions of the face, and two feed-forward networks, predicting vertex offsets of the underlying mesh as well as a view- and expression-dependent texture. We demonstrate that this representation is able to accurately extrapolate to unseen poses and view points, and generates natural expressions while providing sharp texture details. Compared to previous works on head avatars, our method provides a disentangled shape and appearance model of the complete human head (including hair) that is compatible with the standard graphics pipeline. Moreover, it quantitatively and qualitatively outperforms current state of the art in terms of reconstruction quality and novel-view synthesis.

* Both authors contributed equally to the paper

¹ philgras.github.io/neural_head_avatars/neural_head_avatars.html

1. Introduction

Reconstructing and reenacting human heads has been a long studied research problem and will be a key driver for future applications in VR/AR, teleconferencing, games and the movie industry. For those applications, it is of particular interest to get an accurate 3D shape and appearance model that provides 3D consistency and strong identity preservation under novel view points, poses and expressions. Reconstructing such a model, especially from monocular input data (e.g., from a webcam), is difficult due to the complex geometry of facial dynamics and the missing 3D information [96]. Indeed, several state-of-the-art methods for talking head synthesis avoid explicit geometry reconstruction and rely on image or feature-based warping for motion control and generative networks for image synthesis [73,84,92]. These methods are generalized and deliver impressive reenactment results even with only a single input image of the subject. However, the quality of the approaches drops significantly for larger changes in pose or view point as no 3D-consistent geometry representation is used. Shape proxies such as a 3D morphable model [9,28] can be utilized to improve the 3D consistency of synthetic faces [15,44,79,81] since the facial information is embedded on the proxy surface. Besides image- or surface-based representations, volumetric representations are used [31,52]. While they show promising results without an explicit surface prior, these methods still lack a consistent full head shape reconstruct-

tion from single view inputs and are not compatible with the standard rasterization pipelines.

In this work, we present *Neural Head Avatars*, an explicit shape and appearance representation of the complete human head (including hair) which can be used in existing graphics pipelines that use triangular meshes. Specifically, we employ coordinate-based multi-layer perceptrons (MLPs) to predict the 3D meshes and dynamic textures, depending on the facial expression and pose of real humans. These networks are embedded on the surface of the FLAME morphable head model [49] which also serves as a coarse shape and expression proxy. We show that one can optimize such an explicit head representation based on a short monocular RGB video sequence. Using color-dependent and color-independent energy terms during optimization, we disentangle the reconstruction of surface geometry and color detail. The resulting controllable 4D avatar (3D model + motion) is subject-specific and generates novel poses and expressions while preserving high photo-realism. Moreover, it demonstrates great visual quality under large view point changes and, therefore, addresses one of the main drawbacks of related approaches.

In summary, our contributions are:

- Neural Head Avatars, a novel, subject-specific representation for articulated human heads that explicitly reconstructs the full head geometry and produces photo-realistic results even under large view point changes,
- A fully differentiable optimization pipeline to optimize Neural Head Avatars from a short, monocular RGB video with color-dependent and color-independent energy terms that allow for the disentanglement of the surface shape and color detail.

2. Related Work

Reconstructing controllable 4D head or facial avatars is an actively studied field at the intersection of computer vision and computer graphics. For an extensive review of methods, we refer to the state of the art report of Zollhöfer et al. [96].

Image-based models. Image-based models synthesize the face of a subject without relying on any (explicit or implicit) representation in 3D space. These methods either utilize (learned) warping fields [6, 73] to deform an input image to match new poses or expressions, or deploy encoder-decoder architectures, where the encoder extracts an identity code from a given source image and a decoder synthesizes the output image [57, 83, 92, 94]. The decoder may be conditioned by facial landmarks [57, 92], facial contours [83], or parsing maps [94]. Even though these methods produce high-quality results and even allow for real-time synthesis [92], they suffer from artifacts for strong

pose and expression changes, and lack geometric and temporal consistency. This is mainly due to the fact, that the appearance of deformations in three-dimensional space (e.g., yaw opening, head rotation) must be learned in 2D by these models.

Implicit models. Implicit models represent the geometry using implicit surface functions (e.g., signed distance functions) or by volumetric representations. A common approach is to represent the appearance of a target person in a discrete latent feature voxel grid that can be deformed to synthesize dynamic deformations [52, 84]. Motivated by their recent success in 3D scene reconstruction [77], neural radiance fields (NeRF) in combination with volumetric rendering [55] have been used to replace the discrete feature voxel grids [5, 31, 43, 50, 58–60, 63, 64, 68, 75, 90]. Articulated head avatars can be synthesized by conditioning the NeRF on low-dimensional parameters of a face model [31, 86] or audio signals [38]. Pixel-aligned Volumetric avatars [68] are generalized across subjects and can generate novel views, based on single or multiple input images. Even though solving geometric and temporal inconsistencies, the proposed methods either fail to disentangle pose and expression [84], are limited to static reconstructions [43, 68, 69, 85] or fail to generalize to unseen poses and expressions [31].

Explicit models. The majority of head reconstruction methods relies on explicit scene representations, i.e., triangular meshes [7, 8, 13, 19, 20, 22, 32–34, 39, 40, 44, 48, 78–82, 88, 89, 96]. For these methods, morphable models are used as a prior to reconstruct the face from incomplete (e.g., partially occluded) or noisy data (e.g., from depth maps). Morphable models are computed from a population of 3D head scans [28], and provide statistical information on physiologically plausible head shapes and facial movements [12, 25, 49, 62]. In addition to the geometry, these models can provide a statistical linear model for the texture [3, 9, 10, 29, 62] which can be used to reconstruct faces from RGB data only [9, 81, 96]. Recent work utilizes generative adversarial networks (GAN) [37, 56] to generate and optimize albedo and normal maps for specific subjects [36, 47]. Other approaches utilize 2D neural rendering [76, 77], to learn how to render photo-realistic imagery of a specific subject from a short training dataset [44, 78, 79]. These approaches are based on deep neural networks which can be conditioned on coarse RGB renderings based on a linear texture model [44], uv-maps [15], latent feature maps [79] or point clouds [67]. While these methods produce geometrically consistent avatars that can be easily controlled, they either are limited to craniofacial structures and do not include the synthesis of hair [36, 47] or suffer from temporal and spatial inconsistencies due to their loose bound to

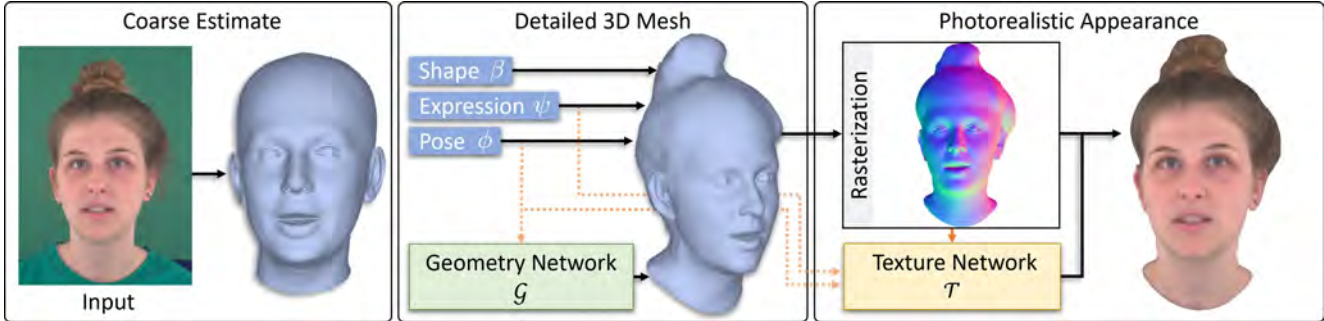


Figure 2. Overview of our method. Given an RGB input frame, we use an adaption of the real-time face tracker [81] to estimate low-dimensional shape, expression, and pose parameters of the linear head model FLAME [49]. We adapt the geometry generated by FLAME with a geometry refinement network \mathcal{G} . The resulting mesh is rasterized with a standard computer graphics pipeline. The texture network \mathcal{T} synthesizes the mesh texture of the rasterized surface. Conditioning \mathcal{T} on the canonical surface position, a local normal patch, and flame parameters enables the synthesis of view- and expression-dependent effects. The networks and FLAME parameters are optimized in an analysis-by-synthesis fashion with color-dependent and color-independent energy terms that allow for the disentanglement of the surface shape and color detail. The reconstructed *Neural Head Avatar* can be animated using the expression and pose parameters ψ, θ , and can be rendered under novel viewpoints.

the geometric backbone [15, 67, 79]. Other methods model mesh displacements to reconstruct fine detail wrinkles and hair [4, 18, 30], therefore, maintaining a close bound to the underlying geometry. However, none of them produces photo-realistic outputs under novel views.

In contrast to previous work, we jointly optimize a photo-realistic dynamic texture together with subject specific geometry that includes facial detail and hair structure, stored in a coordinate-based, fully connected neural network. By deploying a texture-based approach, we tightly link appearance and underlying geometry which ensures spatial consistency and generalization to unseen poses and expressions. Articulated jaw, neck, and eyes allow for intuitive avatar control. We demonstrate that an avatar can be optimized from a short monocular RGB sequence without the need of special (multi-view) camera setups as in [51, 54].

3. Method

Given an RGB video sequence of a talking person consisting of N consecutive frames I_1, I_2, \dots, I_N , we reconstruct a 4D neural avatar based on an explicit representation that allows for pose- and expression-dependent novel viewpoint synthesis. Specifically, our model outputs a classical triangle mesh, i.e., vertices $V = (v_1, v_2, \dots, v_n), v_i \in \mathbb{R}^3$, connecting faces F and a texture function \mathcal{T} that assigns an RGB color value to each point on the surface defined by V and F (see Section 3.1). Thus, the standard graphics pipeline can be utilized to obtain a rendered image \hat{I} assuming a full perspective camera projection. Based on this image formation model, we optimize our avatar representation in an analysis-by-synthesis-based fashion (see Section 3.2). An overview of our method is depicted in Figure 2.

3.1. Explicit Neural Head Representation

Our explicit surface representation is embedded on the FLAME [49] template surface and shares its topology F . Specifically, we employ a multi-layer perceptron (MLP) \mathcal{G} which models the pose-dependent offsets w.r.t. the template surface. To generate the view-, pose-, and expression-dependent texture of the face, we use an MLP \mathcal{T} which predicts a color value at any surface point of the mesh.

Template Model. We deploy the parametric FLAME head model [49] as a geometric backbone of our method:

$$V_{\text{flame}} : \mathbb{R}^{300}, \mathbb{R}^{100}, \mathbb{R}^{3k} \rightarrow \mathbb{R}^{16227 \times 3}$$

$$\beta, \psi, \phi \mapsto V_{\text{flame}}(\beta, \psi, \phi)$$

where β, ψ and ϕ describe shape, expression and $k = 4$ joint pose parameters, respectively. We perform minor adjustments to the FLAME topology, namely, we uniformly subdivide the faces (four-way subdivision), remove the faces belonging to the lower neck region and add additional faces to close the mouth cavity. This increases the original number of vertices from 5023 to 16227.

Geometry Refinement Network \mathcal{G} . To model facial detail and hair which is not represented by the FLAME head model, we introduce a pose-dependent offset function for geometry corrections:

$$\mathcal{G} : \mathbb{R}^{3k} \rightarrow \mathbb{R}^{16227 \times 3}$$

$$\phi \mapsto \mathcal{G}(\phi).$$

Using this offset function, the mesh geometry is given by:

$$V(\beta, \psi, \phi) = V_{\text{flame}}(\beta, \psi, \phi) + \mathcal{G}(\phi).$$




Figure 3. Our optimization is based on the color data of a short video sequence, the corresponding detected facial landmarks (a), predicted semantic labels (b) and predicted normal maps (c).

Texture Network \mathcal{T} . While [29] provides a linear texture space for the FLAME head model, due to its Gaussian nature it lacks fine detail and photo-realism. We introduce a novel appearance model \mathcal{T} which generates a photo-realistic texture, including the synthesis of expression and view dependent effects. In order to predict the color of a point on the mesh, \mathcal{T} receives the 3D coordinates of that point on the canonical FLAME template mesh, the expression and pose of the current frame, and a local patch of the rendered normals as input, and returns the estimated color value. This conditioning of \mathcal{T} enables the synthesis of expression- and view-dependent effects (compare Figure 8 (b)). Formally, \mathcal{T} performs the mapping:

$$\mathcal{T} : \mathbb{R}^3, \mathbb{R}^{100}, \mathbb{R}^{3k}, \mathbb{R}^{n \times n \times 3} \rightarrow \mathbb{R}^3$$

$$p_i, \psi, \phi, \hat{N}_i \mapsto c_i$$

with c_i denoting the predicted color at pixel i , \hat{N}_i a local patch of the rendered normal map around i and p_i being the 3D location on the canonical FLAME template mesh depicted in i .

We approximate both functions \mathcal{G} and \mathcal{T} using two subject-specific, coordinate-based multi-layer perceptrons [21]. Please refer to Appendix A.2 for additional implementation details.

3.2. Optimization based on Monocular RGB Data

The joint optimization of head geometry and texture is a highly underconstrained optimization problem for short monocular video sequences. Besides the data terms that are based on the RGB inputs, we employ regularization strategies that ensure smooth reconstructed surfaces, and view-consistent texture synthesis. The objective of the joint optimization E_{joint} is defined as:

$$E_{\text{joint}} = E_{\text{geom}} + E_{\text{app}}, \quad (1)$$

where E_{geom} measures the data and regularization terms w.r.t. the geometry and E_{app} contains the terms w.r.t. the appearance, i. e., texture and color reproduction.

Geometry Objective E_{geom} . To disentangle appearance and geometry, we define a geometry energy term which is independent of the actual appearance:

$$E_{\text{geom}} = w_{\text{lmk}} \cdot E_{\text{lmk}} + w_{\text{normal}} \cdot E_{\text{normal}} + w_{\text{semantic}} \cdot E_{\text{semantic}} + w_{\text{reg,geom}} \cdot E_{\text{reg,geom}}. \quad (2)$$

The landmark energy E_{lmk} measures the ℓ_1 distance of detected 2D facial landmarks [16, 53] and the projected counterparts on the mesh surface. Besides the absolute positions of the landmarks, it also measures the relative distances of the eye landmarks at the upper and lower lid to improve the reconstruction of eye lid closure [30] (see ablation study in Figure 8 (c)).

The energy term E_{normal} is based on pseudo-normal maps $N \in \mathbb{R}^{H \times W \times 3}$ using the pretrained model of Abrevaya et al. [1]. Based on these predictions, we formulate a reconstruction energy term for fine geometric detail. As our focus lies on high-frequency geometry detail, instead of minimizing the absolute difference between pseudo-normals N and our predicted normals \hat{N} , we instead match their image-space Laplacians $\lambda(\cdot)$:

$$E_{\text{normal}} = |\lambda(\hat{N}) - \lambda(N)|_1.$$

We employ E_{semantic} , to match the semantic regions S_k of the input and \hat{S}_k of the reconstructed mesh for *facial skin and neck, eyes, ears, and hair*:

$$E_{\text{semantic}} = \sum_{k=1}^4 S_k \oplus \hat{S}_k,$$

where \oplus denotes an xor on the region. The respective semantic maps S_k are computed using [91, 95] (see Figure 3 (b)).

Besides the data terms, we employ a regularization term $E_{\text{reg,geom}}$ which regularizes the FLAME parameters, as well as the geometry MLP \mathcal{G} :

$$E_{\text{reg,geom}} = w_{\text{reg,flame}} \cdot E_{\text{reg,flame}} + w_{\text{reg,offset}} \cdot E_{\text{reg,offset}}.$$

Following [9, 81], $E_{\text{reg,flame}}$ uses the statistical properties of the linear shape model, and regularizes the prediction towards the canonical template head using an ℓ_2 -norm on β, ψ, ϕ . The offsets are regularized using $E_{\text{reg,offset}}$ which consists of a Laplacian regularizer and a regularizer that controls the pose-consistency and distribution of the predicted offsets. We refer to Appendix A.3 for additional details.

Appearance Objective E_{app} . The appearance term E_{app} measures the reproduction of the color image I . It depends on both the geometry as well as the texture parameters of our neural head model. We use dense per-pixel energy terms

E_{phot} , as well as an energy term that measures the perceptual distance E_{perc} [42]:

$$E_{\text{app}} = w_{\text{phot}} \cdot E_{\text{phot}} + w_{\text{perc}} \cdot E_{\text{perc}}. \quad (3)$$

E_{perc} compares image features of predicted and ground truth images extracted by the face detector from [26].

Initialization and Optimization Strategy To initialize our reconstruction method, we adopt the tracking algorithm proposed by [81] which optimizes camera, shape, expression, and pose parameters based on an analysis-by-synthesis approach, using the FLAME model [49] with a linear texture space [29]. The resulting reconstruction is in coarse alignment with the training sequence, but the FLAME model is limited to bald heads and lacks fine, subject-specific geometric detail (see Figure 2). We initialize the geometry refinement network \mathcal{G} , by optimizing only for the E_{geom} term defined in Equation (2). Once, we obtained an estimate of the full head geometry, we optimize for the texture MLP parameters w.r.t. E_{app} (Equation (3)). Based on this initialization scheme, we optimize jointly the geometry and texture parameters to minimize E_{joint} (Equation (1)). For implementation details, we refer to Appendix A.²

4. Results

We quantitatively and qualitatively evaluate the performance of our model on the tasks of geometry reconstruction as well as novel pose-, expression-, and view synthesis, and compare it to state-of-the-art methods.

4.1. Datasets

Since the current literature does not provide suitable datasets for the evaluation of dynamic full head approaches, we created two datasets.

Synthetic Data. Our synthetic dataset has been generated with the open source *MakeHuman* project [14] which allows to model fully animatable and texturized human models with high variance in appearance and facial geometry. We generated two female and two male subjects with different ethnicity and head geometry, and render animated sequences (200 training frames, 210 validation frames). The resulting sequences provide ground truth RGB-, normal- and semantic maps as well as landmarks and 3D meshes which is used to quantitatively evaluate the geometry reconstruction of our method. Consequently, for evaluation purposes, we will not rely on predicted pseudo-ground truth (normal and semantic maps, landmarks) for experiments with this dataset. Note that the synthetic meshes have different topologies compared to FLAME and our model.

²Additionally, we will release the code for research purposes.

Metric	Female 1	Female 2	Male 1	Male 2
Normal: FLAME	15.8°	13.3°	15.0°	14.8°
Normal: Ours	14.4°	12.2°	13.7°	13.7°
Hausdorff: FLAME (Face)	1.4	0.9	1.5	1.1
Hausdorff: Ours (Face)	1.2	0.9	1.4	1.2
Hausdorff: FLAME	5.5	4.8	6.1	5.7
Hausdorff: Ours	2.6	2.5	3.0	3.1

Table 1. For four synthetic characters shown in Figure 4, we evaluate our shape prediction using the validation sequence (210 frames) of our dataset. We list the averaged normal error (angular error) and the average mesh alignment error (Hausdorff distance in mm). Normal vectors are compared per pixel in the rendered image, masked by the head region. The single-sided Hausdorff distance is computed from prediction to ground truth meshes, on either the full head or the facial region only. We compare against the reconstruction by FLAME [49] as a baseline.

Real Data. Our real dataset contains four sequences of humans, two with male actors, two with a female actor. The sequences capture various hair styles such as short hair, long hair and a hair bun. All sequences show the subjects during a natural conversation in front of a green screen in an environment with uniform lighting. We capture 750 training frames and 750 validation frames for each sequence and ensure that both sides of the head are visible at least once in the training partition. We complement the obtained RGB ground truth with detected facial landmarks as well as normal- and semantic maps by deploying pretrained models [1, 16, 95]. The resulting dataset is used to qualitatively and quantitatively evaluate our model on the task of novel pose-, expression-, and view synthesis, and to compare to state-of-the-art methods.

4.2. Geometry Reconstruction Quality

To quantify the head shape reconstruction quality, we utilize the rendered ground truth normals, as well as the meshes provided by the synthetic recordings. Those are compared with the predicted meshes and normal maps by computing their single-sided Hausdorff distance in millimeters (mesh alignment error, as in [66]) and their normal angular error on the validation sequence. The Hausdorff distance is computed once for the full head and once for the facial region only. We compare with the reconstructions by FLAME [49] obtained from our tracker as a baseline. The quantitative results are reported in Table 1. Our reconstruction and the misalignment error w.r.t. the ground truth are visualized in Figure 4. We can see that our approach reconstructs the talking head faithfully, and even regions, where only a few silhouette views are available (side of the bun or front/back of the neck) can be estimated, however, with a slightly higher reconstruction error. The reconstruction of real subjects in neutral pose is compared with multi-view stereo (MVS) recordings in Figure 18 in the Appendix.




Figure 4. From left to right: Ground truth image of the synthetic subject, its ground truth geometry, and our predicted geometry. On the right, we show the average Hausdorff distance in mm (from prediction to the ground-truth mesh) over all validation frames.

4.3. Novel Pose and Expression Synthesis

The animatable geometric backbone of our model allows to synthesize new expressions and poses for an optimized avatar. We quantitatively and qualitatively evaluate our model on this task by optimizing it on the training partition of the real dataset and using the resulting avatar to reconstruct the validation frames. For reconstructing the validation frames, we optimize the expression and pose parameters of our method in an analysis-by-synthesis manner, i.e., minimizing Equation (1) only for ψ and θ , while keeping all other components fixed. We compare the results to recent works on talking-head synthesis: NerFACE [31], Deep Video Portraits (DVP) [44], First-Order Motion Model (FOMM) [73], Bi-Layer [92], and VariTex [15]. All one-shot approaches (FOMM, Bi-Layer and VariTex) are given the first frontal frame of the training sequence, while the subject-specific methods (NerFACE and DVP) utilize the whole training sequence. For evaluation, we deploy the pixel-wise ℓ_1 metric; the reference-free cumulative probability blur detection metric (CPBD) [11]; the multi-scale structural similarity metric (MS-SSIM) [87]; the learned perceptual image patch similarity (LPIPS) [93]; and the cosine distance of a pretrained face recognition network (CSIM) [26]. Table 2 shows the resulting scores of perceptual and photometric metrics averaged over all sequences

Method	L1 ↓	MS-SSIM ↑	LPIPS ↓	CSIM ↑	CPBD ↑
VariTex [15]	0.207	0.640	0.316	0.347	0.565
Bi-Layer [92]	0.108	0.805	0.283	0.757	0.403
FOMM [73]	0.055	0.911	0.201	0.800	0.370
DVP [44]	0.042	0.904	0.098	0.815	0.632
NerFACE [31]	0.057	0.897	0.156	0.864	0.485
Ours	0.033	0.923	0.079	0.884	0.674

Table 2. Quantitative evaluation of the appearance reconstruction on the real dataset. Reported are the average scores over all validation frames for all subjects in our real dataset (see Section 4.1).

of our real dataset, and demonstrates that we outperform related approaches consistently. The qualitative comparison in Figure 5 confirms that our model synthesizes images with higher detail and better expression conservation than related methods.

4.4. Novel View Synthesis

The explicit geometry representation for each subject (shown in Figure 6) enables 3D-consistent novel viewpoint synthesis. We evaluate the novel viewpoint synthesis, by rendering frames from the validation set of the real sequences under novel yaw angles. We compare our method against NerFACE [31] and DeepVideoPortraits (DVP) [44] as these methods rely on a representation in 3D space that can be rotated accordingly. The underlying 3D representations are centered on the image plane before rendering to account for varying coordinate origin definitions. Figure 12 demonstrates that while related methods suffer from significant artifacts, our method maintains its high visual quality.

4.5. Ablation Study

We evaluate the influence of different architecture choices and optimization terms on the real dataset. Our geometry is based on the FLAME model, and extends it to capture person-specific detail. For deformations of the face, we rely on the blendshape-based expression model and the linear blend-skinning of the yaw bone. Nevertheless, in our experiments, we found that a static geometry refinement of the underlying FLAME model does not reconstruct the neck region realistically as the joint poses and global rotation of the initial FLAME estimate are highly ambiguous. To this end, we condition the geometry network \mathcal{G} with the joint poses of the FLAME model which compensates errors in the neck region, see Figure 8 (a).

Similar to the dynamic geometry network, we use a pose- and expression-dependent network \mathcal{T} to predict the surface radiance. Figure 8 (b) demonstrates that the use of a static texture, i.e. when the dynamic conditionings of \mathcal{T} are fixed to zero, results in less authentic synthesis results, especially for the highly dynamic mouth region.

Our model also supports eye blinks, which are modelled by the geometry of the FLAME model. To enforce the re-




Figure 5. Comparison of novel pose & expression synthesis results. VariTex, Bi-Layer and FOMM are one-shot approaches and estimate the avatar from the first frontal training frame. All other methods are optimized subject-specifically on the respective training set.




Figure 6. Our 3D mesh reconstruction for the subjects in the real dataset. The meshes align accurately with the real head shapes, even for longer hair, and preserve fine facial details.

constructions of eye blinks, we introduced a specific energy term on the landmarks in the eye region (see E_{lmk} in Section 3.2). Using this energy term in the shape optimization, the geometry faithfully reconstructs blinking eyes, see Figure 8 (c). Without this term, eye blinks are not recovered. Further energy term ablations are shown in the Appendix.

4.6. Discussion

We have demonstrated that our method produces high quality results even for large head rotations due to our explicit geometry reconstruction of the face and hair region. As such, our method addresses one of the main issues of

other learning-based approaches which suffer from significant artifacts when synthesizing novel views. Similar to all baselines, we do not address the synthesis of physical effects like floating or deforming hair. Incorporating physics into a 4D avatar is an interesting field for future work.

Moreover, our method exhibits limitations in regions where the explicit geometry is greatly unconstrained, most prominently the mouth cavity. As a consequence, the visual quality of the synthesized inner mouth region, especially teeth, may decrease if expressions and poses lie far outside of the training corpus (see Figure 9). The most natural approach to overcome this issue would be the integration of a well-aligned geometry of the inner mouth cavity. However, so far none of the publicly available parametric head models provides such geometry due to the difficulties in the acquisition of ground truth data. Similar to NerFACE [31] and DVP [44], our method is person-specific and, thus, requires optimization of the neural network for every new actor which takes 7 hours using two Nvidia A100 GPUs, when optimizing on images with a resolution of 512×512 px. The optimization time and necessary computational resources can be reduced greatly when optimizing against smaller images or image patches. Generalizing our approach is future work which can benefit from the findings of VariTex [15] and pi-GAN [21].




Figure 7. Qualitative novel viewpoint synthesis comparison. While related methods suffer significant artifacts during novel viewpoint synthesis, our method exhibits high robustness even under large rotations and preserves its high texture detail.

Ethical Considerations With the advances in the synthesis of photo-realistic human avatars, the potential misuse (e.g., misinformation) becomes an increasingly important ethical concern. While active watermarking of generated content can be employed, there is no guarantee that this watermark can not be removed. To this end, there exists the field of multi-media forensics which analyzes methods for active and passive forgery detection. Passive forgery detection [2, 23, 24, 71, 72] is able to detect manipulated or synthetic imagery without any explicit watermarking. While these methods can be trained to detect specific manipulation methods [71, 72], generalized methods [2, 23, 24] have problems in reliably detecting fakes. Also, forgery detectors can be used in an adversarial training to improve the quality of the synthetic renderings. Thus, besides forgery detection algorithms, cryptographical approaches for signing the authenticity of video material have to be used in the future (which requires a trust network).

5. Conclusion

In this work, we presented Neural Head Avatars, a method that accurately reconstructs geometry and appearance of the human head from a monocular RGB sequence. Our approach combines a parametric head model with multi-layer perceptrons that refine geometry and synthesize a photorealistic texture. The resulting 4D avatar is robust with respect to large pose-, view- and expression changes, and we show that it outperforms state-of-the-art head avatar methods qualitatively and quantitatively.

While recent work on head avatar synthesis moved to-

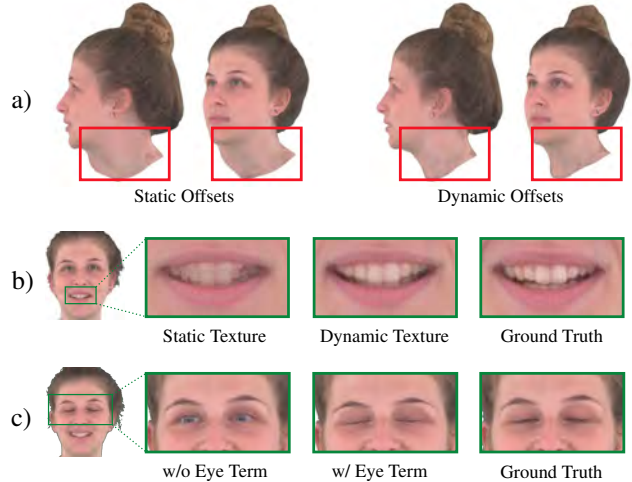


Figure 8. Ablations: **a)** Static vertex offsets lead to bulky neck reconstructions compared to dynamic (pose-conditioned) geometry. **b)** Static texture representations fail to reconstruct highly dynamic face regions such as the interior of the mouth cavity. **c)** Explicit optimization for lid closing reconstructs eye closure faithfully.




Figure 9. Failure cases. Misaligned eye lids in the reconstructed mesh can cause rare eye artifacts (a). Extreme mouth poses may degrade the synthesized mouth interior (b). The reconstruction of ear shapes that deviate strongly from the statistical average yields local geometric artifacts (c).

wards implicit representations of spatial geometry, our work demonstrates the benefits of an explicit geometry reconstruction in combination with a deep appearance network for dynamic surface textures in terms of photorealism and generalizability. We hope that our work inspires further research at the intersection of explicit geometry reconstruction and deep appearance representations.

Acknowledgements

This project has received funding from the DFG in the joint German-Japan-France grant agreement (RO 4804/3-1), the ERC Starting Grant Scan2CAD (804724), and a Microsoft Research Grant. We thank the recorded actors and video narrator for their help, as well as, Mohit Mendiratta and Guy Gafni for providing the synthesis results of [44] and [31] for our comparisons. We gratefully acknowledge the GWK support for funding this project by providing computing time through the Center for Information Services and HPC (ZIH) at TU Dresden.

References

- [1] Victoria Fernandez Abrevaya, Adnane Boukhayma, Philip H.S. Torr, and Edmond Boyer. Cross-modal deep face normals with deactivable skip connections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 4, 5
- [2] Shruti Agarwal, Hany Farid, Tarek El-Gaaly, and Ser-Nam Lim. Detecting deep-fake videos from appearance and behavior. pages 1–6, 12 2020. 8
- [3] Oswald Aldrian and WA Smith. A linear approach of 3d face shape and texture recovery using a 3d morphable model. In *Proceedings of the British Machine Vision Conference*, pages 75–1. Citeseer, 2010. 2
- [4] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *2018 International Conference on 3D Vision (3DV)*, pages 98–109. IEEE, 2018. 3
- [5] Thiemo Alldieck, Hongyi Xu, and Cristian Sminchisescu. imghum: Implicit generative models of 3d human shape and articulated pose. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [6] Hadar Averbuch-Elor, Daniel Cohen-Or, Johannes Kopf, and Michael F. Cohen. Bringing portraits to life. *ACM Transactions on Graphics (Proceeding of SIGGRAPH Asia 2017)*, 36(6):196, 2017. 2
- [7] Volker Blanz, Curzio Basso, Tomaso Poggio, and Thomas Vetter. Reanimating faces in images and video. In *Proc. EUROGRAPHICS*, volume 22, pages 641–650, 2003. 2
- [8] Volker Blanz, Kristina Scherbaum, Thomas Vetter, and Hans-Peter Seidel. Exchanging faces in images. *CGF*, 23(3):669–676, 2004. 2
- [9] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. 1, 2, 4, 15
- [10] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3d morphable model. *IEEE Transactions on pattern analysis and machine intelligence*, 25(9):1063–1074, 2003. 2
- [11] P Bohr, Rupali Gargote, Rupali Vhorkate, RU Yawle, and VK Bairagi. A no reference image blur detection using cumulative probability blur detection (cpbd) metric. *International Journal of Science and Modern Engineering*, 1(5), 2013. 6
- [12] James Booth, Anastasios Roussos, Allan Ponniah, David Dunaway, and Stefanos Zafeiriou. Large scale 3d morphable models. *International Journal of Computer Vision*, 126(2):233–254, 2018. 2
- [13] Sofien Bouaziz, Yangang Wang, and Mark Pauly. Online modeling for realtime facial animation. *ACM TOG*, 32:40:1–40:10, 2013. 2
- [14] Leyde Briceno and Gunther Paul. Makehuman: a review of the modelling framework. In *Congress of the International Ergonomics Association*, pages 224–232. Springer, 2018. 5
- [15] Marcel C Bühler, Abhimitra Meka, Gengyan Li, Thabo Beeler, and Otmar Hilliges. Varitex: Variational neural face textures. *arXiv preprint arXiv:2104.05988*, 2021. 1, 2, 3, 6, 7, 17
- [16] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *International Conference on Computer Vision*, 2017. 4, 5, 15, 17
- [17] Andrei Burov, Matthias Nießner, and Justus Thies. Dynamic surface function networks for clothed human bodies. *arXiv preprint arXiv:2104.03978*, 2021. 13, 15
- [18] Chen Cao, Derek Bradley, Kun Zhou, and Thabo Beeler. Real-time high-fidelity facial performance capture. *ACM Trans. Graph.*, 34(4), jul 2015. 3
- [19] Chen Cao, Qiming Hou, and Kun Zhou. Displaced dynamic expression regression for real-time facial tracking and animation. *ACM Trans. Graph.*, 33(4), jul 2014. 2
- [20] Chen Cao, Yanlin Weng, Stephen Lin, and Kun Zhou. 3d shape regression for real-time facial animation. *ACM Trans. Graph.*, 32(4), jul 2013. 2
- [21] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proc. CVPR*, 2021. 4, 7, 13, 14
- [22] Yen-Lin Chen, Hsiang-Tao Wu, Fuhao Shi, Xin Tong, and Jinxiang Chai. Accurate and robust 3d facial capture using a single rgbd camera. *Proc. ICCV*, pages 3615–3622, 2013. 2
- [23] Davide Cozzolino, Andreas Rössler, Justus Thies, Matthias Nießner, and Luisa Verdoliva. Id-reveal: Identity-aware deepfake video detection. 2021. 8
- [24] Davide Cozzolino, Justus Thies, Andreas Rössler, Christian Riess, Matthias Nießner, and Luisa Verdoliva. Forensictransfer: Weakly-supervised domain adaptation for forgery detection. *arXiv*, 2018. 8
- [25] Hang Dai, Nick Pears, William Smith, and Christian Duncan. Statistical modeling of craniofacial shape and texture. *International Journal of Computer Vision*, 128(2):547–571, Nov 2019. 2
- [26] Jiankang Deng, Jia Guo, Tongliang Liu, Mingming Gong, and Stefanos Zafeiriou. Sub-center arcface: Boosting face recognition by large-scale noisy web faces. In *Proceedings of the IEEE Conference on European Conference on Computer Vision*, 2020. 5, 6, 17
- [27] Vincent Dumoulin, Ethan Perez, Nathan Schucher, Florian Strub, Harm de Vries, Aaron Courville, and Yoshua Bengio. Feature-wise transformations. *Distill*, 3(7):e11, 2018. 13, 14
- [28] Bernhard Egger, William A. P. Smith, Ayush Tewari, Stefanie Wuhler, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, Christian Theobalt, Volker Blanz, and Thomas Vetter. 3d morphable face models - past, present and future. *ACM Transactions on Graphics*, 39(5), Aug. 2020. 1, 2
- [29] Haven Feng and Timo Bolkart. Photometric flame fitting. https://github.com/HavenFeng/photometric_optimization, 2020. Accessed: 2021-08-10. 2, 4, 5
- [30] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. volume 40, 2021. 3, 4, 15

- [31] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. 2021. [1](#), [2](#), [6](#), [7](#), [8](#), [16](#), [17](#)
- [32] Pablo Garrido, Levi Valgaerts, Ole Rehmsen, Thorsten Thormählen, Patrick Perez, and Christian Theobalt. Automatic face reenactment. In *Proc. CVPR*, 2014. [2](#)
- [33] Pablo Garrido, Levi Valgaerts, Hamid Sarmadi, Ingmar Steiner, Kiran Varanasi, Patrick Perez, and Christian Theobalt. VDub - modifying face video of actors for plausible visual alignment to a dubbed audio track. In *Computer Graphics Forum (Proceedings of EUROGRAPHICS)*, 2015. [2](#)
- [34] Pablo Garrido, Levi Valgaerts, Chenglei Wu, and Christian Theobalt. Reconstructing detailed dynamic face geometry from monocular video. In *ACM TOG*, volume 32, pages 158:1–158:10, 2013. [2](#)
- [35] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. A neural algorithm of artistic style. *CoRR*, abs/1508.06576, 2015. [16](#)
- [36] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1155–1164, 2019. [2](#)
- [37] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [2](#)
- [38] Yudong Guo, Keyu Chen, Sen Liang, Yongjin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [39] Pei-Lun Hsieh, Chongyang Ma, Jihun Yu, and Hao Li. Unconstrained realtime facial performance capture. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. [2](#)
- [40] Alexandru Eugen Ichim, Sofien Bouaziz, and Mark Pauly. Dynamic 3d avatar creation from hand-held video input. *ACM TOG*, 34(4):45:1–45:14, 2015. [2](#)
- [41] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. [16](#)
- [42] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [5](#)
- [43] Petr Kellnhofer, Lars C Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4287–4297, 2021. [2](#)
- [44] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018. [1](#), [2](#), [6](#), [7](#), [8](#), [16](#), [17](#)
- [45] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015. [17](#)
- [46] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. In *International conference on machine learning*, pages 1558–1566. PMLR, 2016. [16](#)
- [47] Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Triantafyllou, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme: Realistically renderable 3d facial reconstruction in-the-wild”. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 760–769, 2020. [2](#)
- [48] Hao Li, Jihun Yu, Yuting Ye, and Chris Bregler. Realtime facial animation with on-the-fly correctives. In *ACM TOG*, volume 32, 2013. [2](#)
- [49] 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, (Proc. SIGGRAPH Asia)*, 36(6):194:1–194:17, 2017. [2](#), [3](#), [5](#), [13](#)
- [50] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Trans. Graph.(ACM SIGGRAPH Asia)*, 2021. [2](#)
- [51] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018. [3](#)
- [52] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. [1](#), [2](#)
- [53] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019. [4](#), [15](#)
- [54] Shugao Ma, Tomas Simon, Jason Saragih, Dawei Wang, Yuecheng Li, Fernando De La Torre, and Yaser Sheikh. Pixel codec avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 64–73, 2021. [3](#)
- [55] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020. [2](#)
- [56] 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. pages 1–12, 12 2018. [2](#), [17](#), [18](#)
- [57] Yuval Nirkin, Yosi Keller, and Tal Hassner. FSGAN: Subject agnostic face swapping and reenactment. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7184–7193, 2019. [2](#)
- [58] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *International Conference on Computer Vision (ICCV)*, 2021. [2](#)

- [59] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. [2](#)
- [60] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021. [2](#)
- [61] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. [17](#)
- [62] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *2009 sixth IEEE international conference on advanced video and signal based surveillance*, pages 296–301. Ieee, 2009. [2](#)
- [63] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Animatable neural radiance fields for human body modeling. *arXiv preprint arXiv:2105.02872*, 2021. [2](#)
- [64] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9054–9063, June 2021. [2](#)
- [65] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. [13](#), [14](#)
- [66] Rohith Krishnan Pillai, Laszlo Attila Jeni, Huiyuan Yang, Zheng Zhang, Lijun Yin, and Jeffrey F Cohn. The 2nd 3d face alignment in the wild challenge (3dfaw-video): Dense reconstruction from video. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3082–3089. IEEE Computer Society, 2019. [5](#)
- [67] Sergey Prokudin, Michael J Black, and Javier Romero. Smpix: Neural avatars from 3d human models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1810–1819, 2021. [2](#), [3](#)
- [68] Amit Raj, Michael Zollhofer, Tomas Simon, Jason Saragih, Shunsuke Saito, James Hays, and Stephen Lombardi. Pixel-aligned volumetric avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11733–11742, 2021. [2](#)
- [69] Eduard Ramon, Gil Triginer, Janna Escur, Albert Pumarola, Jaime Garcia, Xavier Giro-i Nieto, and Francesc Moreno-Noguer. H3d-net: Few-shot high-fidelity 3d head reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5620–5629, 2021. [2](#)
- [70] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020. [17](#)
- [71] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv*, 2018. [8](#)
- [72] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics++: Learning to detect manipulated facial images. In *ICCV 2019*, 2019. [8](#)
- [73] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *Conference on Neural Information Processing Systems (NeurIPS)*, December 2019. [1](#), [2](#), [6](#), [7](#)
- [74] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020. [13](#), [14](#)
- [75] Shih-Yang Su, Frank Yu, Michael Zollhoefer, and Helge Rhodin. A-nerf: Surface-free human 3d pose refinement via neural rendering. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. [2](#)
- [76] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, K. Sunkavalli, R. Martin-Brualla, T. Simon, J. Saragih, M. Nießner, R. Pandey, S. Fanello, G. Wetzstein, J.-Y. Zhu, C. Theobalt, M. Agrawala, E. Shechtman, D. B Goldman, and M. Zollhöfer. State of the art on neural rendering. *EG*, 2020. [2](#)
- [77] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Yifan Wang, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, Tomas Simon, Christian Theobalt, Matthias Niessner, Jonathan T. Barron, Gordon Wetzstein, Michael Zollhoefer, and Vladislav Golyanik. Advances in neural rendering, 2021. [2](#)
- [78] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. *ECCV 2020*, 2020. [2](#)
- [79] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. [1](#), [2](#), [3](#)
- [80] Justus Thies, Michael Zollhöfer, Matthias Nießner, Levi Valgaerts, Marc Stamminger, and Christian Theobalt. Real-time expression transfer for facial reenactment. *ACM TOG*, 34(6), 2015. [2](#)
- [81] J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2016. [1](#), [2](#), [3](#), [4](#), [5](#), [15](#)
- [82] Justus Thies, Michael Zollhöfer, Christian Theobalt, Marc Stamminger, and Matthias Niessner. Headon. *ACM Transactions on Graphics*, 37(4):1–13, Aug 2018. [2](#)
- [83] 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. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. [2](#)

- [84] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10039–10049, 2021. 1, 2
- [85] Xueying Wang, Yudong Guo, Zhongqi Yang, and Juyong Zhang. Prior-guided multi-view 3d head reconstruction. *IEEE Transactions on Multimedia*, 2021. 2
- [86] Ziyang Wang, Timur Bagautdinov, Stephen Lombardi, Tomas Simon, Jason Saragih, Jessica Hodgins, and Michael Zollhöfer. Learning Compositional Radiance Fields of Dynamic Human Heads. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 2
- [87] Z. Wang, E.P. Simoncelli, and A.C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems Computers, 2003*, volume 2, pages 1398–1402 Vol.2, 2003. 6
- [88] Thibaut Weise, Sofien Bouaziz, Hao Li, and Mark Pauly. Real-time performance-based facial animation. In *ACM TOG*, volume 30, 2011. 2
- [89] Thibaut Weise, Hao Li, Luc J. Van Gool, and Mark Pauly. Face/Off: live facial puppetry. In *Proc. SCA*, pages 7–16, 2009. 2
- [90] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [91] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 325–341, 2018. 4
- [92] Egor Zakharov, Aleksei Ivakhnenko, Aliaksandra Shysheya, and Victor Lempitsky. Fast bi-layer neural synthesis of one-shot realistic head avatars. In *European Conference of Computer vision (ECCV)*, August 2020. 1, 2, 6, 7
- [93] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
- [94] Yunxuan Zhang, Siwei Zhang, Yue He, Cheng Li, Chen Change Loy, and Ziwei Liu. One-shot face reenactment. *arXiv preprint arXiv:1908.03251*, 2019. 2
- [95] zllrunning. Face parsing pytorch. <https://github.com/zllrunning/face-parsing.PyTorch>, 2021. 4, 5, 15
- [96] M. Zollhöfer, J. Thies, P. Garrido, D. Bradley, T. Beeler, P. Pérez, M. Stamminger, M. Nießner, and C. Theobalt. State of the Art on Monocular 3D Face Reconstruction, Tracking, and Applications. *Computer Graphics Forum (Eurographics State of the Art Reports)*, 37(2), 2018. 1, 2




Figure 10. As our template, we uniformly subdivide the FLAME [49] mesh and simplify the mouth cavity.

A. Implementation Details

A.1. Template Model

We utilize the FLAME head model [49] in its updated version from 2020 as the geometric backbone of our method. As mentioned in the main paper, we perform minor adjustments to the FLAME topology, namely, we uniformly subdivide the faces (four-way subdivision), remove the faces belonging to the lower neck region, and add faces to close the mouth cavity, see Figure 10. This increases the original number of vertices from 5023 to 16227. Inspired by [17], we use \tanh activation functions to limit the joint rotations of FLAME to physiologically plausible ranges.

A.2. Network Architectures

Our method relies on two multi-layer perceptrons (MLPs). The Geometry MLP \mathcal{G} refines the mesh resulting from the linear FLAME head model and adds facial detail and hair structure. The Texture MLP \mathcal{T} synthesizes a dynamic, photo-realistic texture that is able to reproduce view- and expression dependent effects, e.g., wrinkles and reflections. This section details the model architecture of both networks. An overview is presented in Figure 11.

A.2.1 Geometry MLP \mathcal{G}

The Geometry MLP \mathcal{G} adds facial detail and hair geometry to the mesh resulting from the FLAME head model. The inputs to the network are the 3D coordinates of the vertices on the template mesh (see Figure 10), normalized to the range $[-1, 1]$, as well as vertex-specific embedding vectors that are optimized during training. Dynamic geometry effects are enabled by also passing the FLAME pose parameters to \mathcal{G} . In practice, to avoid overfitting only the 3 pose parameters of the neck joint are used.

The core of \mathcal{G} is a SIREN-based fully connected MLP [74] which takes the vertex positions and embeddings as input, processes them by a sequence of FiLM-conditioned linear layers [27, 65], and predicts the three-dimensional ver-

tex offsets. The linear, FiLM-conditioned layers first perform an affine transformation on the input signal, followed by a sinusoidal activation function. The phase shifts and frequencies of the sinusoidal activation functions are generated by a mapping network which is a concatenation of linear layers with leaky ReLU activations. This architecture was greatly inspired by [21]. The input to the mapping network is the three-dimensional pose of the FLAME neck joint which enables the synthesis of dynamically changing geometry refinements. We ensure spatial geometry consistency by only allowing dynamic geometry refinements for the neck region. To this end, we compute the offsets twice: once conditioned on the original neck pose values and once with all conditions set to zero. We smoothly blend both offset predictions according to a fading body region mask.

Configuration details. The vertex embedding vectors have a feature dimension of 32 (for each vertex of template mesh). The mapping network comprises 3 hidden layers with 256 neurons each, followed by leaky ReLU activation functions with slope 0.2. The SIREN MLP consists of 6 consecutive linear FiLM layers with 128 neurons each. The final layer is a linear layer with 3 output neurons, followed by a \tanh activation function.

A.2.2 Texture MLP \mathcal{T}

The architecture of our texture representation \mathcal{T} is similar to \mathcal{G} . For obtaining the color value of a point on the mesh surface, its 3D coordinates on the FLAME template mesh as well as a surface embedding vector are fed into a SIREN MLP. The surface embedding vector is sampled from a discrete 2D grid in uv space. A mapping network takes features extracted from the FLAME pose and expression parameters and predicts the phase shifts and frequencies of the sinusoidal activation functions in the linear FiLM layers of the SIREN MLP to synthesize a dynamic texture. The dynamic texture is needed in the regions where the reconstructed geometry does not align well with the target surface. This is especially the case for the mouth cavity. Consequently, we process the FLAME pose and expression parameters into features that are highly correlated with the mouth articulation. More specifically, we determine the effective rotation that applies to the mouth region, i.e., combine global rotation and neck rotation, and use its axis-angle representation together with 10 pair-wise distances between vertices on the inner side of the lips as inputs to the mapping network. These conditions only cover mouth-related information and, therefore, are only used as input to the mapping network for the respective surface regions. For all other regions, we fill the conditioning vector with zeros. Effectively, this approach greatly limits the dynamic capacities of the \mathcal{T} to ensure a close bound to the underly-

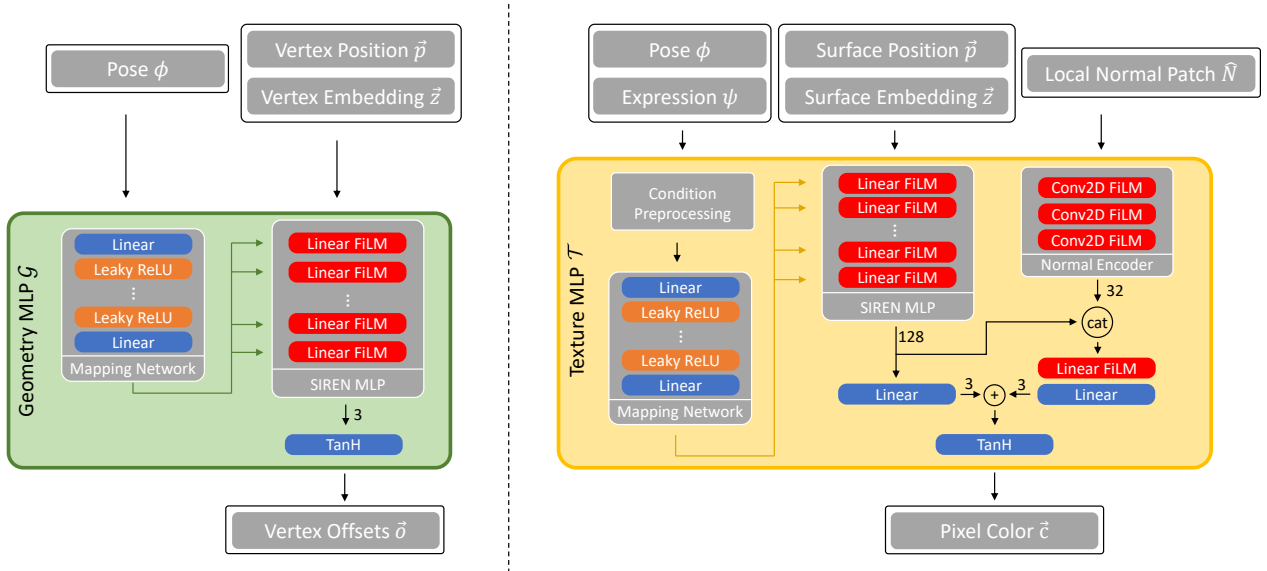


Figure 11. Overview of our model architectures. Our *Neural Head Avatar* relies on SIREN-based MLPs [74] with fully connected linear layers, periodic activation functions and FiLM conditionings [27, 65]. Inspired by [21], surface coordinates and spatial embeddings (either vertex-wise for \mathcal{G} , or as an interpolatable grid in uv-space for \mathcal{T}) are used as an input to the SIREN MLP. The dynamic frequencies and phase shifts of the Linear FiLM layers are predicted by fully connected mapping networks which are conditioned on the FLAME parameters. For the Texture MLP \mathcal{T} , a fully convolutional normal encoder generates additional conditions from a local patch of the predicted normal map.

ing geometry which in turn enables extrapolation to unseen poses and expressions. The outputs of the mapping network are used to dynamically adjust the phase shifts and frequencies of the linear FiLM layers in the SIREN MLP. Given these conditions and the surface position and surface embedding, the SIREN MLP produces a latent space vector. This vector serves as input to two network heads. In the first head, the latent vector is fed through a linear layer with 3 output neurons. In the second head, the latent vector is concatenated with the outputs of a normal encoder network. This network takes a local patch of the rendered normal map as input, passes it through a sequence of 2D-convolutional FiLM layers and outputs a feature vector in latent space. 2D-convolutional FiLM layers have the same structure as linear FiLM layers, but instead of a linear layer, a 2D-convolutional layer is used. The latent vector produced by the normal encoder contains information on the local geometry configuration and enables the synthesis of expression- and view-dependent effects (e.g., ambient occlusions and specular highlights). The output vectors of the SIREN MLP and the normal encoder are concatenated and fed through a sequence of linear layers to produce 3 output activations. The activations of both heads are summed up and a tanh activation is applied to achieve the final RGB predictions in a range of (-1, 1).

Configuration details. The surface embedding is sampled in uv space from a discrete feature grid with 256×256

feature vectors with 64 channels each via bilinear interpolation. We use a separate uv map for the inner mouth region with 64×64 vectors. The mapping network has the same architecture as for the Geometry MLP \mathcal{G} . The SIREN MLP consists of 8 consecutive linear FiLM layers with 256 neurons each, except for the last which has 128 neurons. The normal encoder is designed as a fully convolutional network with 3 consecutive 2D-convolutional layers with stride 1 and kernel size 3 and with periodic activations. All layers have 128 feature channels except for the last one which has 32. The first model head contains one fully connected layer with 3 output neurons. The second model head contains one linear FiLM layer with 128 neurons and one linear layer with 3 neurons.

A.3. Optimization from Monocular RGB Data

Detailed Geometry Objective E_{geom} (Eq. 2) As defined in Eq. 2 in the main paper, the geometry energy term is:

$$E_{\text{geom}} = w_{\text{lmk}} \cdot E_{\text{lmk}} + w_{\text{normal}} \cdot E_{\text{normal}} + w_{\text{semantic}} \cdot E_{\text{semantic}} + w_{\text{reg, geom}} \cdot E_{\text{reg, geom}}. \quad (4)$$

E_{normal} and E_{semantic} are already defined in the main paper, for the other terms we provide additional explanation.

The landmark energy E_{lmk} measures the distance of detected 2D facial landmarks $\mathbf{l}_i \in \mathbb{R}^2$ and the projected coun-

	$w_{\text{reg,flame}}$		$w_{\text{reg,surface}}$		$w_{\text{reg,edge}}$		$w_{\text{reg,lapl}}$	
	Geometry Optim.	Joint Optim.	Geometry Optim.	Joint Optim.	Geometry Optim.	Joint Optim.	Geometry Optim.	Joint Optim.
Eyeballs	1.0E-03	1.0E-03	1.0E-04	1.0E-04	0	0	0	0
Eye Surrounding	1.0E-03	1.0E-03	1.0E-04	1.0E-04	0	0	5	10
Forehead	1.0E-03	1.0E-03	1.0E-04	1.0E-04	0	0	0.05	0.1
Face	1.0E-03	1.0E-03	1.0E-04	1.0E-04	0	0	0.05	0.1
Ears	1.0E-03	1.0E-03	1.0E-04	1.0E-04	0	0	25	50
Scalp	1.0E-03	1.0E-03	1.0E-04	1.0E-04	10	10	0.05	0.1
Neck	1.0E-03	1.0E-03	1.0E-04	1.0E-04	0	0	0.1	0.2
Lower Neck	1.0E-03	1.0E-03	1.0E-04	1.0E-04	0	0	0.25	0.5
Nose	1.0E-03	1.0E-03	1.0E-04	1.0E-04	0	0	2.50E-02	5.00E-02

Table 3. Geometry regularization weights. The regularization weights differ for the individual head regions. However, they do not change among target subjects, i.e., they can be used to reconstruct a wide variety of geometries such as different hair styles.

terparts on the mesh surface $\hat{\mathbf{l}}_i \in \mathbb{R}^2$ and is given by:

$$E_{\text{lmk}} = \sum_{i=1}^{70} \|\mathbf{l}_i - \hat{\mathbf{l}}_i\|_1 + w_{\text{lid}} \cdot \sum_{i \in \{\text{left, right}\}} \|\mathbf{d}_i - \hat{\mathbf{d}}_i\|_1.$$

Besides the absolute positions of the landmarks $\hat{\mathbf{l}}_i$, we measure relative distances $\hat{\mathbf{d}}_i$ of the eye landmarks at the upper and lower lid to improve the reconstruction of eye lid closure [30]. However, we found that target lid distances \mathbf{d}_i are less noisy when being computed on the facial segmentation (computed by [95]) rather than the detected landmarks. The 2D facial landmarks \mathbf{l}_i are detected using [16, 53] and also contain two iris landmarks.

To avoid convergence to local minima, we employ several geometry regularization strategies summing up to the term $E_{\text{reg, geom}}$ which is given by:

$$E_{\text{reg, geom}} = w_{\text{reg, flame}} \cdot E_{\text{reg, flame}} + w_{\text{reg, lapl}} \cdot E_{\text{reg, lapl}} + w_{\text{reg, surface}} \cdot E_{\text{reg, surface}} + w_{\text{reg, edge}} \cdot E_{\text{reg, edge}}. \quad (5)$$

Following [9, 81], $E_{\text{reg, flame}}$ uses the statistical properties of the linear shape model, and regularizes the prediction towards the average face:

$$E_{\text{reg, flame}} = w_\beta \cdot |\beta|_2^2 + w_\theta \cdot |\theta|_2^2 + w_\psi \cdot |\psi|_2^2.$$

The Laplacian regularizer $E_{\text{reg, lapl}}$ computes the relative loss of the Laplace values $\lambda^*(V)$ of the predicted mesh vertices V w.r.t. the surface of the FLAME model V_{Flame} , and controls the smoothness of the predicted offsets:

$$E_{\text{reg, lapl}} = |\mathcal{W}_{\text{reg, lapl}} \circ (\lambda^*(V) - \lambda^*(V_{\text{Flame}}))|_1.$$

Note that $\lambda^*(\cdot)$ denotes the discretized Laplace-Beltrami operator on the 1-ring neighborhood. $\mathcal{W}_{\text{reg, lapl}} \in \mathbb{R}_+^V$ defines vertex specific weights which allow to control the smoothness in specific regions (e.g., the neck region has a higher regularization). \circ denotes the component-wise Hadamard product. As in [17], we regularize pose-dependent offset variations by adding a pair-wise surface consistency term. For two randomly picked frames i, j , the ℓ_1 -norm is computed over the difference between $\mathcal{G}(\phi_i)$ and $\mathcal{G}(\phi_j)$:

$$E_{\text{reg, surface}} = |\mathcal{G}(\phi_i) - \mathcal{G}(\phi_j)|_1.$$

For large geometry corrections (e.g., long hair), we noticed an irregular vertex distribution over the surface which resulted in areas with large triangles and coarse shape modelling. To mitigate this issue, we regularize the length of edges e_i in the scalp region if they deviate too much from the average edge length \bar{e} .

$$E_{\text{reg, edge}} = \sum_{e_i} \begin{cases} e_i - \bar{e} & \text{if } e_i > 1.5 \cdot \bar{e} \\ 0 & \text{otherwise} \end{cases}$$

As the individual facial regions are subject to different requirements, the applied regularization weights differ. Table 3 presents the weights for the individual regions. Please note that while the weights differ for the individual face regions, they are the same for all avatars. No subject-specific fine-tuning of regularization parameters is required while still being able to reconstruct a wide set of structures (e.g. long and short hair).




Figure 12. Additional qualitative novel view synthesis comparisons. We observe that the advantages of our method, namely spatial consistency and high texture detail even under extreme head rotations, apply to a variety of identities. Also see Figure 1 for additional subjects.

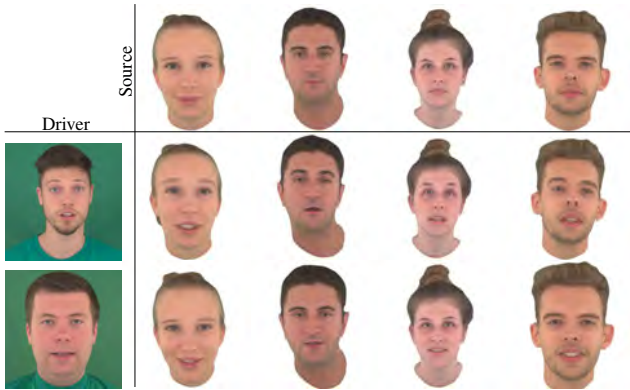


Figure 13. Expressions ψ and poses ϕ from the driving frames on the left are added to the neutral poses of optimized source avatars at the top. The reenacted avatars are displayed below.

Detailed Appearance Objective E_{app} (Eq. 3) The appearance energy term is defined in Eq. 3 in the main paper as:

$$E_{\text{app}} = w_{\text{phot}} \cdot E_{\text{phot}} + w_{\text{perc}} \cdot E_{\text{perc}}. \quad (6)$$

The photo-metric energy term E_{phot} is defined on the intersection $\mathcal{V} = \mathcal{S} \cap \hat{\mathcal{S}}$ of the foreground segmentation of the input I and the region that is generated by our head model \hat{I} :

$$E_{\text{phot}} = |\mathcal{V} \cdot (\hat{I} - I)|_1.$$

To generate sharp textures [41, 46], we employ the style-based perceptual energy term E_{perc} proposed by [35].

Initialization and Optimization. As discussed in Section 3.2 of the main paper, texture and geometry have to be initialized before starting the joint optimization against RGB




Figure 14. Qualitative novel view synthesis comparison for pitch rotations. We report the frontal view as well as synthesis results under a pitch angle of $\pm 15^\circ$.

sequences to avoid converging to bad local minima. To this end, the FLAME head model is aligned with the training sequence to obtain a coarse geometry initialization. Based on this coarse geometry, the weights of \mathcal{G} and the frame-specific parameters of the FLAME model, i.e., pose and expressions are optimized w.r.t. E_{geom} (Eq. (2) in main paper). As a result of this optimization, we obtain a geometry estimate that aligns well with the target silhouette. Using this refined geometry, \mathcal{T} is trained to minimize E_{app} (Eq. (3) in main paper). In a final optimization step, all components are optimized jointly to minimize E_{joint} (Eq. (1) in main paper).

Given the 750 input frames of our real sequences, we initialize \mathcal{G} for 150 epochs, \mathcal{T} is initialized for another 100 epochs, and we jointly optimize both components as long as




Figure 15. Quantitative novel view synthesis comparison. We report the cosine similarity score between latent features predicted by a pretrained face recognition network [26]. The feature vectors are compared between the front facing ground truth and the predicted image under different rotation angles. We report the average scores over all validation frames from two sequences in our real dataset together with the respective 1σ regions. Our method consistently outperforms related approaches for pitch angles between -15° and $+15^\circ$. For extreme yaw rotations, we observe significantly reduced CSIM scores even though qualitative comparisons demonstrate high identity preservation under these conditions (see Figure 7 in main paper).




Figure 16. Qualitative comparison with paGAN [56] on the validation sequences. Note that only the facial region is synthesized by paGAN.

the perceptual loss on the holdout validation set decreases. This takes approximately 50 epochs. For longer sequences, we reduce the epoch count such that the number of iterations per optimization stage remains the same.

We deploy a standard Adam optimizer [45] for all frame-agnostic parameters and a standard SGD optimizer for frame-specific components (e.g. expression and pose parameters). Weight decay is applied to all texture-related

components. The entire pipeline is implemented with Pytorch [61] and Pytorch3D [70].

B. Additional Results

B.1. Avatar Reenactment

Neural Head Avatars are controllable by the disentangled pose and expression spaces of the FLAME head model. This naturally enables the reenactment of an optimized avatar via a driving sequence. We demonstrate this capability in Figure 13. In this experiment, we utilize expressions ψ and poses ϕ from a driver’s video and add them to the neutral expression of optimized avatars. ψ and ϕ are extracted using our tracking algorithm discussed in the main paper. We observe that our method is able to faithfully transfer pose and expression between various subjects.

B.2. Novel View Synthesis

In Section 4.4 of the main paper, we qualitatively compare the synthesis from novel viewpoints against related methods. Figure 14 and Figure 12 provide further qualitative comparisons on different subjects and demonstrate that the advantages of our method apply to a variety of identities.

We also provide a quantitative evaluation of the novel view synthesis results in Figure 15. For a quantitative analysis, given that no ground truth is available for novel views, we evaluate the cosine similarity (CSIM) of the latent feature vectors predicted by a pretrained face recognition network [16] between front-facing ground truth and novel view prediction. We report the CSIM scores averaged over all validation frames of two sequences in our real dataset. We




Figure 17. Synthesis results for a non-caucasian subject. The high quality of the synthesis results of our method is consistent also for people of color.

find that our method outperforms related approaches consistently for pitch angles in a range of $\pm 15^\circ$. However, we observe that for large yaw angles ($\geq \pm 30^\circ$), the scores for all considered models decrease rapidly. Still, qualitative comparisons demonstrate that our method exhibits high identity preservation even under viewpoint changes in that range.

B.3. Comparison with paGAN [56]

In Figure 16, we compare our generated images of the real identities with the outputs of paGAN [56] provided by the authors. As their results are overlays on top of the original video, only the synthesized parts (facial region) should be considered when comparing to it.

B.4. Synthesis Results for Non-Caucasian Subjects

To validate that our method also performs well for non-caucasian subjects, we include synthesis results for a person of color in Figure 17. Also in this case, our method reconstructs the head geometry faithfully and renders visually plausible images.

B.5. Geometry Evaluation

We compare the geometry predicted by our approach on the real dataset with multi-view stereo (MVS) recordings of the respective identities, see Figure 18. The MVS data was captured separately with a handheld DSLR camera. As hair styles and face dynamics can not be reproduced reliably in separate recordings, only the neutral pose of the face region is compared.

B.6. Energy Ablation

Figure 19 demonstrates the effect of further energy terms on the synthesis results. We observe that E_{phot} prevents color shifts, E_{semantic} improves alignment of overlapping semantic regions within the avatar’s silhouette and E_{perc} results in textures with more detail.




Figure 18. Geometry evaluation on the real identities. The right most column visualizes the Hausdorff distance from our predicted face mesh to the recorded GT. From top to bottom, the total alignment errors of the identities are 1.5, 1.6, 1.6, and 1.6 mm.




Figure 19. Energy Term Ablation. Reference normals are the inputs we use for optimization.