


DiffusionAvatars: Deferred Diffusion for High-fidelity 3D Head Avatars

Tobias Kirschstein

Simon Giebenhain

Matthias Nießner

Technical University of Munich



(a) Avatar creation

(b) Avatar animation

Figure 1. Given a set of multi-view videos and corresponding fitted meshes, we build a *DiffusionAvatar* of a person. Our method translates expressions of a morphable model into realistic facial appearances of a person while also providing control over the viewpoint. Project website: <https://tobias-kirschstein.github.io/diffusion-avatars/>

Abstract

DiffusionAvatars synthesizes a high-fidelity 3D head avatar of a person, offering intuitive control over both pose and expression. We propose a diffusion-based neural renderer that leverages generic 2D priors to produce compelling images of faces. For coarse guidance of the expression and head pose, we render a neural parametric head model (NPHM) from the target viewpoint, which acts as a proxy geometry of the person. Additionally, to enhance the modeling of intricate facial expressions, we condition DiffusionAvatars directly on the expression codes obtained from NPHM via cross-attention. Finally, to synthesize consistent surface details across different viewpoints and expressions, we rig learnable spatial features to the head’s surface via TriPlane lookup in NPHM’s canonical space. We train DiffusionAvatars on RGB videos and corresponding fitted NPHM meshes of a person and test the obtained avatars in both self-reenactment and animation scenarios. Our experiments demonstrate that DiffusionAvatars generates temporally consistent and visually appealing videos for novel poses and expressions of a person, outperforming existing approaches.

1. Introduction

A significant problem in computer vision and graphics is the creation of photorealistic animatable avatars. Ideally, these avatars permit consistent video renderings while providing free control over desired expression, pose, and viewpoint. A digital 3D head avatar may then be used in various scenarios such as VR/AR applications, immersive teleconferencing, gaming, movie animation, or as a virtual assistant.

Unfortunately, creating a 3D avatar from which one can render photorealistic images from arbitrary viewpoints is challenging. The task becomes even more difficult when the avatar has to simultaneously offer precise control over poses and expressions, for instance given by a driving sequence from a different video or controlled manually. Technically, this leads to a 4D photometric reconstruction problem which is typically underconstrained. This presents animation systems with a challenging task that traditionally involves extensive efforts by 3D artists.

Many existing methods have tackled the challenge of facial image synthesis via 2D neural networks [26, 31, 52, 58, 70, 73]. In particular, Diffusion Models have recently been demonstrated to produce visually appealing imagery [23, 59, 61, 86, 89]. 2D models are great at generating photo-realistic images but generally do not offer

the degree of control and temporal consistency needed for a high-quality 3D head avatar. On the other side, methods that reconstruct an animatable 3D representation of the head [21, 92, 93, 95] provide much better consistency but in many cases the resulting renderings do not have the same photo-realism as 2D models.

To this end, we introduce DiffusionAvatars, a novel approach for 3D head avatar creation, by combining the strong image synthesis capabilities of 2D diffusion models with the view consistency of a detailed 3D head representation. As a controllable 3D representation, we leverage the recent neural parametric head model (NPHM) [18] due to its detailed prediction of human head geometry, which can give us better geometric cues than established mesh-based 3DMMs [38, 55]. Inspired by neural textures [72], we further base face synthesis on person-specific learnable features rigged to the model’s surface via NPHM’s canonical space. Such surface-mapped features can effectively compensate for imperfect geometry and improve consistency across views by giving the neural renderer cues about corresponding surface points. To avoid learning a neural renderer from scratch, we base our architecture on a pre-trained latent diffusion model (LDM) [59] and convert it into an image-to-image translation model by following the ControlNet [89] paradigm to condition on the rasterized images of NPHM meshes. That way, we not only inherit a powerful image synthesis backbone but also benefit from the learned facial prior obtained from large 2D image datasets, which facilitates generalization to unseen expressions. Finally, we insert cross-attention blocks into the pre-trained LDM to condition the diffusion-based neural renderer on the expression codes of NPHM directly. This allows our model to produce more expressive renderings. Intuitively, the direct conditioning helps the diffusion model to distinguish subtle expression details while the rasterized NPHM renderings encode the viewpoint and the overall shape of the head. In summary, our contributions are as follows:

- We present DiffusionAvatars, a diffusion-based neural renderer that leverages ControlNet to create animatable 3D head avatars.
- We design a method for rigging learnable spatial features to the surface of the underlying NPHM via TriPlanes.
- We propose direct expression conditioning via cross-attention to transfer detailed expressions from NPHM to the synthesized 3D head avatar.

2. Related Work

2.1. 3D Face Animation

It is natural to approach the task of 3D Face Animation with an actual 3D representation. In the seminal work by [4], a 3D morphable model (3DMM) was introduced, which paved the way towards reconstructing animatable 3D head

geometry from images or depth observations. More recent 3DMMs [7, 38, 55, 74, 82] still build upon this paradigm and have since been the drivers of face animation oriented applications with most methods leveraging a 3DMM to gain control over the expression and pose of the avatar. For example, Neural Head Avatars [21] and ROME [33] finetune the mesh topology of FLAME [38] to obtain more faithful and realistic mesh-based avatars. NeRFace [16], INSTA [95] and RigNeRF [1] build a radiance field that is controlled by a 3DMM. Other 3D representations such as points [93] or signed distance functions [40, 92] have also been explored for avatar creation.

In our work, we do not aim to learn a complete 3D representation. Instead, we utilize the 3D prior of the recent neural parametric head model (NPHM) [18] with a 2D rendering network for high-quality image synthesis.

2.2. Face Synthesis in 2D

Ever since the impactful Pix2Pix [94] work, it has been a popular approach to obtain control over synthesized images via 2D networks. In the face synthesis domain, methods based on Generative Adversarial Networks (GANs) [20] have received a lot of attention [26, 31, 52, 58, 70, 73]. To obtain better control over head pose and camera viewpoint, a common technique is to combine powerful 2D image synthesis networks with a controllable 3D head proxy [2, 8, 11, 39, 49, 67, 69, 71, 72, 79]. Another line of work aims at directly generating consistent videos by video-to-video translation [34, 76] or even video stylization [15, 28, 83]. Most related to our method is Deferred Neural Rendering [72] which allows facial animation by learning a neural renderer in order to decode neural textures rigged to Basel Face Model [55]. While 2D-based approaches produce visually appealing frames, they often struggle with synthesizing consistent images across both views and time or provide only limited viewpoint control. Our method aims to improve consistency by utilizing a more powerful implicit 3D representation as geometric proxy and provides better expression generalization by utilizing a diffusion model as a synthesis backbone.

2.3. Controllable 2D Diffusion

Recently, Diffusion models emerged as powerful 2D image generation models. While they demonstrate superior capabilities in generating high-fidelity and diverse 2D content from text prompts [23, 59, 61], it has been a key challenge to leverage their power for other tasks. Several works extend 2D diffusion models to text-guided video generation by employing a multi-stage approach [5, 24, 64, 77], generating frames autoregressively [25] or stylizing an input video [17, 78, 81, 84]. In the human domain, diffusion models have already been applied to face editing [27], animation [66, 86] and bodies [68]. Other works, such as

ControlNet [89], IPAdapter [85] or T2I-Adapter [50], allow fine-tuning a pre-trained diffusion model on additional controls such as landmarks or depth cues, opening a wide space of possible applications for 2D face editing. Most similar to our method, DiffusionRig [12] proposes a diffusion model conditioned on rasterized FLAME meshes to provide 3D animation control over an avatar created from a personal photo album. While the synthesized images or videos of these approaches already show great controllability and visual quality, they typically lack consistency due to the absence of an underlying accurate 3D representation.

2.4. Diffusion for 3D Faces

Following a different approach, several works explore diffusion for generic 3D generation [43, 51, 56]. This has also been extended to 3D head generation [22, 75, 88], text-guided 3D head editing [53] or 3D body generation [87]. Some works also generate animatable 3D asset animatable for heads [3] and bodies [6, 29, 57] by rigging the 3D asset to a parameterized FLAME [38] or SMPL [45] mesh. Such generated 3D avatars are view consistent by design but lack the visual quality of 2D-based diffusion models. Furthermore, animation control is limited by the underlying mesh template. In our work, we employ a more accurate, implicit 3D morphable head model for better expression control and follow ControlNet [89] and IPAdapter [85] to build a diffusion-based neural renderer for high-quality image synthesis.

3. Method

Our goal is to create a temporally consistent 3D head avatar of a person with explicit control over viewpoint and expression. We approach this task by designing a diffusion-based neural renderer that decodes learnable features rigged to the surface of an implicitly defined proxy geometry. In analogy to Deferred Neural Rendering [72], we dub this approach *Deferred Diffusion*. We begin by introducing the foundations of diffusion models (Sec. 3.1) and NPHM (Sec. 3.2) that we use as proxy geometry before going into detail about our method: In the rasterization (Sec. 3.3) and surface feature mapping (Sec. 3.4) stage, we generate the input images for our diffusion-based neural renderer. These renderings encode the desired viewpoint, head shape, and global expression. For more fine-grained expression control, we add NPHM’s expression codes as an additional input to the diffusion model (Sec. 3.5). Finally, our pipeline leverages a pre-trained diffusion model’s image synthesis and generalization capabilities for better quality.

3.1. Diffusion

We build our work upon latent diffusion models (LDM) [59], which operate in the latent space of an au-

toencoder. Formally, LDMs employ an encoder \mathcal{E} to map RGB images into a lower-dimensional latent space, facilitating generative tasks [13]. In the following, any mention of an image x_0 always refers to an RGB image I that was mapped into the autoencoder’s latent space via $x_0 = \mathcal{E}(I)$ to obtain a latent image.

In diffusion, the fixed forward process iteratively adds noise to the latent image:

$$q(x_\tau|x_{\tau-1}) = \mathcal{N}(x_\tau; \sqrt{\alpha_\tau}x_{\tau-1}, (1 - \alpha_\tau)\mathbf{I}) \quad (1)$$

where $\tau = 1 \dots N$ indicates the denoising step and the forward process variances α_τ define the noise scheduling [23]. In practice, noisy samples x_τ are obtained with the standard Gaussian reparameterization:

$$x_\tau = \sqrt{\alpha_\tau}x_0 + \sqrt{1 - \alpha_\tau}\epsilon \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (2)$$

Typically, diffusion models are trained to predict the original noise ϵ given the noisy sample x_τ . We empirically find a different method, known as the v -parameterization [61], to work better in our scenario, where v is defined as:

$$v = \sqrt{\alpha_\tau}\epsilon - \sqrt{1 - \alpha_\tau}x_0 \quad (3)$$

This has two advantages. First, v -prediction ensures that the loss can always guide the model to learn something meaningful, even when the input already contains a lot of noise, leading to faster convergence [24]. Second, it allows us to train the model also on pure noise inputs, which during inference is the most challenging denoising step. To this end, we rescale the noise schedule α to ensure zero signal-to-noise ratio inputs during training as proposed by [42].

3.2. NPHM

NPHM [18] is a morphable head model that generates a signed distance field (SDF) of a human head given an identity code z_{id} and an expression code z_{exp} . We employ COLMAP [62, 63] to obtain pointclouds $\{\mathcal{P}_t\}$ for each timestep t of a multi-view video sequence of the person. Subsequently, we fit NPHM to each pointcloud \mathcal{P}_t to obtain z_{id} and z_{exp}^t . We use a variant of NPHM, namely MonoNPHM [19], that uses a backward deformation field instead of the originally proposed forward deformation field as it simplifies the fitting process. Formally, we obtain the NPHM fitting as follows:

$$z_{id}, z_{exp}^t = \arg \min_{z_{id}, z_{exp}} \sum_{x \in \mathcal{P}_t} |\mathcal{F}_{id}(\mathcal{F}_{exp}(x))| \quad (4)$$

where \mathcal{F}_{id} is NPHM’s identity network implemented as a signed distance field, and \mathcal{F}_{exp} is the backward deformation network. Note that we drop the dependency of \mathcal{F}_{id} and \mathcal{F}_{exp} on their respective latent codes for simplicity. The fitted SDF representation can then be translated into a mesh

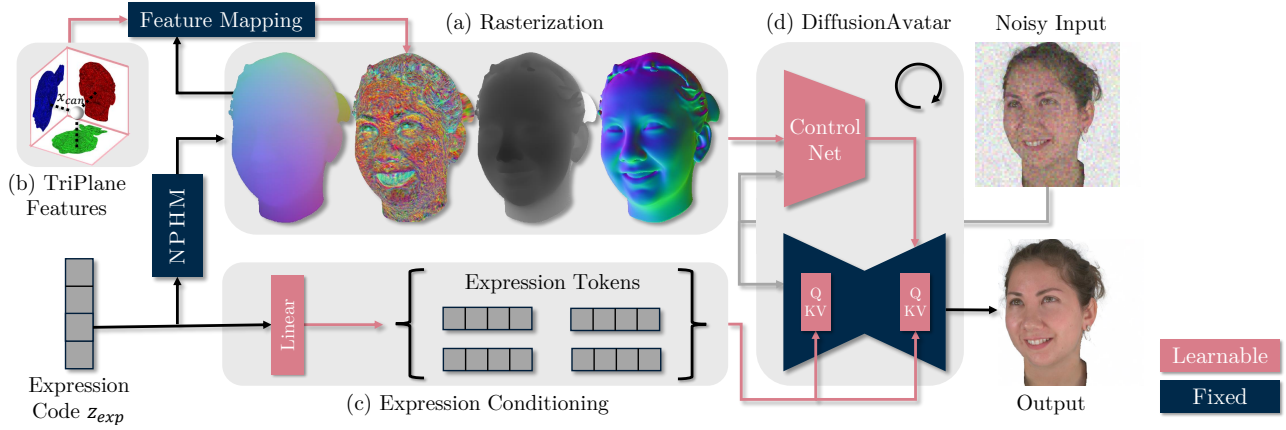


Figure 2. **Method overview:** We decode an NPHM expression code z_{exp} in two ways to obtain a realistic image: We first extract an NPHM mesh and rasterize it from the desired viewpoint in (a), giving us canonical coordinates, depths, and normal renderings for the head mesh. In (b), the canonical coordinates x_{can} are used to look up spatial features in a TriPlanes structure, rigging the features onto the mesh surface. Together with the rasterizer output, these mapped features form the input for the ControlNet part of DiffusionAvatar. The second route for the expression code goes through a linear layer depicted in (c). It yields expression tokens that are subsequently used in a newly added cross-attention layer inside the pre-trained latent diffusion model. Intuitively, the rasterized inputs should encode pose, shape and rough expression while the direct expression conditioning hints at more detailed facial expressions. The final image is synthesized in (d) by iteratively denoising Gaussian noise using the original DDPM denoising schedule [23].

M_t via marching cubes [46]. Furthermore, for each vertex $x \in M_t$ of the extracted mesh, we can obtain its canonical coordinates $x_{can} \in \mathbb{R}^{3+2}$ via the backward deformation field:

$$x_{can} = \mathcal{F}_{exp}(x) \quad (5)$$

where the first 3 coordinates of x_{can} represent the usual spatial dimensions while the remaining 2 coordinates are ambient dimensions that can help resolve topological issues in mouth and eye regions [54]. We utilize the meshes M_t in the subsequent rasterization stage, whereas the canonical coordinates x_{can} are necessary for the spatial feature lookup.

3.3. Rasterization

We employ nvdiffrast [37] to generate the actual input images for the diffusion model. Let M_t be a facial mesh of a person at timestep t and π_t a camera pose. We then obtain a set of renderings $R^t \in \mathbb{R}^{H \times W \times C}$ for this timestep using the rasterizer \mathcal{R} :

$$R^t = \mathcal{R}(M_t, \pi_t) \quad (6)$$

In our case, the channels of R^t are normals, depths, and a canonical coordinate rendering $R_{can}^t \in \mathbb{R}^{H \times W \times (3+2)}$ of the NPHM mesh.

3.4. TriPlane Feature Mapping

It has been shown that tying learnable features to a surface helps neural renderers to synthesize more detailed images, a technique known as neural textures [72]. Since NPHM

lacks a consistent UV space due to its implicit nature, we propose a simple extension of neural textures: We tie learnable features to the surface of the extracted mesh M_t by querying a spatial structure with the canonical coordinates x_{can} . We use TriPlanes [8] for the 3 spatial dimensions and a regular 2D feature map for the ambient dimensions:

$$R_{feat}^t = \text{TRIPLANE}(R_{can,0-3}^t) \quad (7)$$

$$R_{feat,amb}^t = \text{AMBIENTMAP}(R_{can,3-5}^t) \quad (8)$$

These learnable feature maps form the 2D input for our diffusion-based neural renderer, together with the other buffers R^t from the rasterizer:

$$R^t \leftarrow R^t, R_{feat}^t, R_{feat,amb}^t \quad (9)$$

3.5. Direct Expression Conditioning

Conditioning on renderings guides the head pose as well as coarse expressions. However, subtle details may not be easy for the network to decode from renderings. To facilitate the synthesis of detailed expressions, we add new cross-attention layers to the U-Net, following IPAdapter [85]. Let Z be the intermediate feature map computed by an existing cross-attention operation in the pre-trained LDM: $Z = \text{ATTENTION}(Q, K, V)$. We then perform direct expression conditioning by adding another cross-attention layer:

$$f_{exp}^t = \text{EXP}(z_{exp}^t) \quad (10)$$

$$Z \leftarrow Z + \text{ATTENTION}(Q, W^k f_{exp}^t, W^v f_{exp}^t) \quad (11)$$

where EXP is our expression conditioning module that linearly maps z_{exp}^t into a sequence $f_{exp}^t \in \mathbb{R}^{4 \times d}$ of 4 expression tokens forming the keys and values for cross-attention. In total, we add 15 cross-attention layers to the LDM.

3.6. Deferred Diffusion

The final rendering is obtained by iteratively denoising full noise $x_T \sim \mathcal{N}(0, I)$ with our diffusion-based neural renderer \mathcal{D} conditioned on the rasterized NPHM meshes R^t and expression tokens f_{exp} :

$$x_{\tau-1} \sim \mathcal{D}(x_\tau) = \mathcal{S}(x_\tau, f_{exp}^t, \mathcal{C}(x_\tau, R^t)) \quad (12)$$

where \mathcal{C} is the ControlNet architecture and \mathcal{S} a pre-trained LDM with our expression conditioning modules inserted. The denoised prediction $x_{\tau-1}$ is obtained via the sampling procedure of [23].

During training, we minimize the following loss:

$$\mathcal{L} = \mathbb{E}_{\epsilon, \tau, x_0^t, R^t, f_{exp}^t} \left[\left\| \mathcal{D}(x_\tau^t, R^t, f_{exp}^t) - v \right\|_2 \right] \quad (13)$$

where ϵ is the sampled noise, τ is the denoising step and $x_0^t = \mathcal{E}(I^t)$ is the ground truth latent image. We use \mathcal{L} to optimize (i) the ControlNet \mathcal{C} , (ii) the parameters of expression conditioning W^k, W^v , and EXP, and (iii) the spatial feature maps TRIPLANE and AMBIENTMAP.

4. Experimental Results

4.1. Training

For our experiments, we employ the recently released multi-view video dataset from NeRSemble [36]. The dataset contains 26 different multi-view sequences per person captured with 16 cameras, featuring different facial expressions, sentences and emotions. We use BackgroundMattingV2 [41] to remove the background and a segmentation network to remove the torso [91]. We fit NPHM on the video sequences of 8 different persons to create pairings of multi-view videos and fitted meshes. In total, we use about 3300 timesteps per person. We train one DiffusionAvatar for each person on these pairings with Adam [35] using a learning rate of $1e-4$ for ControlNet and the expression conditioning layers, and a learning rate of $1e-2$ for the TriPlane feature lookups. We use a batch size of $B = 8$ and spatial resolutions $H = W = 512$. We employ random cropping and resizing of the target images to avoid overfitting. The size of our spatial feature maps is $512 \times 512 \times 16$ for each TriPlane and the ambient map. In total, this yields $16 \times (3 + 1) + 5 + 1 + 3 = 73$ channels for the input renderings R_t fed to our diffusion-based renderer. The pre-trained LDM is a Stable Diffusion v2.1 network [59]. Training takes roughly two days on a single RTX A6000 for 100k gradient descent steps.

Method	PSNR↑	LPIPS↓	JOD↑	AKD↓	AED↓	APD↓	CSIM↑
NeRFace [16]	23.0	0.279	6.76	5.37	1.06	0.053	0.787
DiffusionRig [12]	19.6	0.220	6.41	2.74	0.55	0.029	0.887
DNR [72]	24.5	0.226	7.32	2.06	0.63	0.027	0.903
DNR+GAN [72]	23.0	0.114	7.08	2.14	0.69	0.028	0.868
MVP [44]	23.6	0.221	7.02	3.42	0.78	0.034	0.882
Ours	24.9	0.081	7.55	1.79	0.50	0.023	0.917

Table 1. **Quantitative Comparison:** We report metrics for a self-reenactment scenario on unseen expressions and unseen views averaged over 8 persons from the NeRSemble dataset [36].

4.2. Experiment Setup

Tasks. To measure the quality of a 3D avatar, we conduct experiments on two tasks: Self-reenactment and avatar animation. For self-reenactment, we hold out 2 complete sequences and the frontal camera of every sequence during training to measure both image quality and view consistency. During testing, we provide the holdout fitted meshes of the same person and render from the novel view. For avatar animation, we first generate NPHM meshes using the same identity code z_{id} as the avatar and a sequence of expression codes $\{z_{exp}^t\}$ from a different person. We then render a video of the avatar from a novel view. To evaluate the realism of this expression transfer, we conduct a user study where we show the animated avatars to participants.

Metrics. We employ two paired-image metrics to measure the quality of individual generated frames: Peak Signal-to-Noise Ratio (PSNR) and Learned Perceptual Image Patch Similarity (LPIPS) [90]. Furthermore, we measure visual similarity of a rendered video to its ground truth counterpart with the perceptual video metric JOD [47]. Finally, we make use of several face-specific metrics: Average Keypoint Distance (AKD) measured in pixels with keypoints estimated from PIPNet [30], and cosine similarity (CSIM) of identity embeddings based on ArcFace [10]. We further follow [9] and utilize DECA [14] to measure Average Expression Distance (AED) and Average Pose Distance (APD) between the rendered and the ground truth images.

Baselines. We compare our method against the following state-of-the-art 3D head avatar creation methods:

NeRFace [16]. A NeRF-based method that reconstructs a radiance field conditioned on expression codes. The original implementation is for monocular videos and uses FLAME. We extend it to our multi-view training scenario and provide NPHM expression codes for better comparison.

Mixture of Volumetric Primitives (MVP) [44]. An encoder-decoder architecture that generates a set of small feature grids rigged to a template mesh. We provide a FLAME tracking to learn the driving.



Figure 3. **Qualitative results for self-reenactment.** We compare against 3D methods (NeRFace [16], Mixture of Volumetric Primitives [44]) and methods employing 2D renderers (Deferred Neural Rendering [72], DiffusionRig [12]). Note that the slightly gray background for DiffusionRig is caused by their training scheme using ϵ -prediction (see Sec. 3.1). Our method consistently produces more expressive facial performances while simultaneously providing more detailed renderings.

Deferred Neural Rendering (DNR) [72]. A screen-space decoder built on top of Pix2Pix [94] that uses learnable neural textures rigged to a tracked Basel Face Model [55]. We provide our fitted NPHM meshes as input for better comparison. Additionally, we compare against a version of DNR that also uses an adversarial loss for sharper images that we call DNR+GAN in the following.

DiffusionRig [12]. A recent diffusion-based method that uses albedo, normal, and shaded color renderings of FLAME meshes tracked by DECA [14] as input to control pose and expression. While their method was originally proposed to be fine-tuned on a small personal photo album of 20 images, we fine-tune their pre-trained model on all training frames of our dataset.




Figure 4. **Qualitative results for avatar animation.** Our method faithfully transfers the source actor’s expression and consistently produces compelling renderings, even for complex performances.

4.3. Self-Reenactment

In Tab. 1, we compare DiffusionAvatars to other state-of-the-art approaches on the self-reenactment task. Results are averaged over avatars generated from eight different persons. Our method shows the overall best performance with the biggest improvement in perceptual and face-specific metrics. The results indicate that DiffusionAvatars generates sharper images (LPIPS) that are more temporally consistent (JOD). Furthermore, our method better maintains the person’s identity when rendered from a novel view (CSIM) than the baselines. Finally, facial features such as the position of mouth corners and eyebrows synthesized by our method match the ground truth expressions more closely (AKD, AED, APD).

We further provide a qualitative comparison in Fig. 3. The baselines struggle especially in facial regions that are poorly explained by the NPHM mesh, such as the mouth interior. In contrast, our method can plausibly fill these regions, providing a consistently realistic appearance.

4.4. Avatar Animation

Due to the task’s inherent lack of any ground truth, we perform a user study to validate avatar animation. We collect 735 responses from 49 participants on two questions:

	NeRFace	DiffusionRig	DNR	DNR+GAN	Ours
VQ↑	2.19	2.47	2.87	3.06	4.02
DF↑	2.35	3.52	3.94	3.97	4.14

Table 2. **User Study:** For our user study on avatar animation, we collect 735 responses from 49 participants to measure Visual Quality (VQ) and Driving Fidelity (DF) on a scale of 1-5.

(i) Visual Quality (VQ): What is the overall quality of the avatar animation?

(ii) Driving Fidelity (DF): How closely does the animated avatar match the expressions of the driving sequence?

We present two videos side-by-side: One sequence of a source actor performing various expressions, and a video of another person’s avatar with the expressions transferred to it. The results of our user study can be seen in Tab. 2. Our method achieves the highest scores for both visual quality and similarity to the driving sequence. The difference is particularly pronounced for the visual quality score. We attribute this to the fact that users heavily penalize any visual artifacts immediately obvious in video renderings. These perceptual differences in the task of avatar animation are also apparent in the qualitative comparison shown in Fig. 4.

4.5. Ablations

In Tab. 3, we examine the design choices of our method. All ablation experiments are conducted on 3 avatars on the self-reenactment task. We refer to our supplementary material for a qualitative comparison.

	PSNR \uparrow	LPIPS \downarrow	JOD \uparrow	AKD \downarrow	CSIM \uparrow
w/o diffusion	25.1	0.133	7.69	1.96	0.892
FLAME	24.2	0.083	7.36	2.46	0.900
w/o exp. cond.	24.5	0.081	7.65	1.98	0.911
w/o LDM prior	24.5	0.078	7.67	1.89	0.913
w/o spatial features	24.9	0.078	7.77	1.87	0.918
spherical UV	24.9	0.075	7.77	1.95	0.920
Ours	25.3	0.074	7.85	1.91	0.918

Table 3. Ablation of architectural choices.

Effect of Diffusion. To study whether diffusion is necessary, we train our architecture without the pre-trained LDM to directly predict the output in a single forward pass, similar to Deferred Neural Rendering. Such a method performs notably worse, especially in the LPIPS metric, indicating that diffusion is important for sharp renderings.

Effect of NPHM vs FLAME. We compare DiffusionAvatars with a version trained on renderings of a fitted FLAME instead of NPHM. Since NPHM models facial geometry as an SDF, it possesses a much higher representational power than FLAME. As a result, the renderings are much closer to the actual desired appearance of the head. This is also reflected in our ablation study, where using FLAME gives consistently worse performance. This is especially true for the AKD metric, indicating that our diffusion-based neural renderer cannot reenact certain complex expressions based on FLAME alone.

Effect of Expression Conditioning. Turning off expression conditioning worsens overall performance, presumably because the expression codes help to distinguish subtle expression details that can hardly be inferred from the NPHM mesh alone. It also helps DiffusionAvatars to synthesize expressions not captured by NPHM, such as tongue movement, by enabling the network to correlate such appearances with certain expression code combinations.

Effect of 2D LDM Prior. We train our pipeline from scratch without resorting to a pre-trained LDM. In this case, we remove the ControlNet module and condition the U-Net directly on the NPHM renderings. This architecture, trained from scratch, performs competitively but slightly worse. We see the biggest advantage of pre-training in faster convergence and better coherence of the generated images.

Effect of Spatial Features. We compare three spatial feature mapping approaches: (i) Our Triplanes, (ii) a simpler approach based on spherical feature mapping, and (iii) no spatial features at all. For spherical feature mapping, we create an approximate UV map for each NPHM mesh by projecting the 3D canonical coordinates for each vertex onto a sphere centered inside the head. Tab. 3 shows that using TriPlanes boosts most metrics, whereas the simpler spherical mapping brings almost no benefit.

5. Limitations

DiffusionAvatars can create photo-realistic 3D head avatars with pose and expression control. More challenges must be solved, however, before an avatar can be used in production. For example, to put an avatar into a realistic environment, it is necessary to have control over the lighting properties. Currently, DiffusionAvatars bakes the lighting into the generated images. Since we employ a detailed 3D head geometry underneath, one could imagine directly modeling shadow effects as a composite of the synthesized images. Further, our current architecture is not (yet) amenable to real-time applications due to the comparatively compute-intensive denoising loop. Here, recent advances in the distillation of diffusion models may provide a remedy [48, 65].

6. Conclusion

This paper presents a novel method DiffusionAvatars, which creates a photo-realistic 3D head avatar from multi-view videos using a tracked parametric head model of a person. DiffusionAvatars can be animated by a source actor or directly via NPHM. We use rendered NPHM meshes with rigged spatial features as input to a diffusion-based neural renderer. We leverage the image synthesis capabilities of a pre-trained latent diffusion model with ControlNet, facilitating generalization to unseen expressions. While this pipeline is already capable of synthesizing appealing images, we show that additionally conditioning the neural renderer on NPHM’s expression codes further improves the model’s ability to generate complex facial expressions. Our experiments suggest that the proposed architecture can effectively synthesize high-quality renderings and animate the 3D avatar with high fidelity. We believe our approach demonstrates an exciting application for high-fidelity 3D avatar creation.

Acknowledgements

This work was supported by the ERC Starting Grant Scan2CAD (804724) and the German Research Foundation (DFG) Research Unit “Learning and Simulation in Visual Computing”. We would also like to thank Karla Weighart for proof-reading and Angela Dai for the video voice-over.

References

- [1] ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. Rignerf: Fully controllable neural 3d portraits. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 20364–20373, 2022. 2
- [2] Alexander Bergman, Petr Kellnhofer, Wang Yifan, Eric Chan, David Lindell, and Gordon Wetzstein. Generative neural articulated radiance fields. *Advances in Neural Information Processing Systems*, 35:19900–19916, 2022. 2
- [3] Alexander W Bergman, Wang Yifan, and Gordon Wetzstein. Articulated 3d head avatar generation using text-to-image diffusion models. *arXiv preprint arXiv:2307.04859*, 2023. 3
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 157–164, 2023. 2
- [5] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22563–22575, 2023. 2
- [6] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K Wong. Dreamavatar: Text-and-shape guided 3d human avatar generation via diffusion models. *arXiv preprint arXiv:2304.00916*, 2023. 3
- [7] Zenghao Chai, Tianke Zhang, Tianyu He, Xu Tan, Tadas Baltrusaitis, HsiangTao Wu, Runnan Li, Sheng Zhao, Chun Yuan, and Jiang Bian. Hiface: High-fidelity 3d face reconstruction by learning static and dynamic details. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9087–9098, 2023. 2
- [8] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16123–16133, 2022. 2, 4
- [9] Xuangeng Chu, Yu Li, Ailing Zeng, Tianyu Yang, Lijian Lin, Yunfei Liu, and Tatsuya Harada. GPAvatar: Generalizable and precise head avatar from image(s). In *The Twelfth International Conference on Learning Representations*, 2024. 5
- [10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. 5
- [11] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *IEEE Computer Vision and Pattern Recognition*, 2020. 2
- [12] Zheng Ding, Xuaner Zhang, Zhihao Xia, Lars Jebe, Zhuowen Tu, and Xiuming Zhang. Diffusionrig: Learning personalized priors for facial appearance editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12736–12746, 2023. 3, 5, 6, 7, 13, 15
- [13] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 3
- [14] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. 2021. 5, 6
- [15] Jakub Fišer, Ondřej Jamriška, David Simons, Eli Shechtman, Jingwan Lu, Paul Asente, Michal Lukáč, and Daniel Šykora. Example-based synthesis of stylized facial animations. *ACM Transactions on Graphics (TOG)*, 36(4):1–11, 2017. 2
- [16] Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8649–8658, 2021. 2, 5, 6, 7, 13
- [17] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Tokenflow: Consistent diffusion features for consistent video editing. *arXiv preprint arXiv:2307.10373*, 2023. 2
- [18] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Learning neural parametric head models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21003–21012, 2023. 2, 3
- [19] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Monophm: Dynamic head reconstruction from monocular videos. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [20] 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
- [21] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular rgb videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18653–18664, 2022. 2
- [22] Jiatao Gu, Qingzhe Gao, Shuangfei Zhai, Baoquan Chen, Lingjie Liu, and Josh Susskind. Learning controllable 3d diffusion models from single-view images. *arXiv preprint arXiv:2304.06700*, 2023. 3
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1, 2, 3, 4, 5
- [24] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 2, 3
- [25] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv:2204.03458*, 2022. 2
- [26] Fa-Ting Hong, Longhao Zhang, Li Shen, and Dan Xu. Depth-aware generative adversarial network for talking head

- p>video generation. In
- Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*
- , pages 3397–3406, 2022. 1, 2
- [27] Ziqi Huang, Kelvin C.K. Chan, Yuming Jiang, and Ziwei Liu. Collaborative diffusion for multi-modal face generation and editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 2
 - [28] Ondřej Jamriška, Šárka Sochorová, Ondřej Texler, Michal Lukáč, Jakub Fišer, Jingwan Lu, Eli Shechtman, and Daniel Šykora. Stylizing video by example. *ACM Transactions on Graphics (TOG)*, 38(4):1–11, 2019. 2
 - [29] Ruixiang Jiang, Can Wang, Jingbo Zhang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Avatarcraft: Transforming text into neural human avatars with parameterized shape and pose control. *arXiv preprint arXiv:2303.17606*, 2023. 3
 - [30] Haibo Jin, Shengcai Liao, and Ling Shao. Pixel-in-pixel net: Towards efficient facial landmark detection in the wild. *International Journal of Computer Vision*, 2021. 5
 - [31] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 1, 2
 - [32] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in neural information processing systems*, 34:852–863, 2021. 14
 - [33] Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. Realistic one-shot mesh-based head avatars. In *European Conference on Computer Vision*, pages 345–362. Springer, 2022. 2
 - [34] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Nießner, Patrick Pérez, Christian Richardt, Michael Zollhofer, and Christian Theobalt. Deep video portraits. *ACM Transactions on Graphics (TOG)*, 37(4):163, 2018. 2
 - [35] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
 - [36] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Trans. Graph.*, 42(4), 2023. 5, 13, 14
 - [37] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics*, 39(6), 2020. 4
 - [38] 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, 13
 - [39] C.Z. Lin, D.B. Lindell, E.R. Chan, and G. Wetzstein. 3d gan inversion for controllable portrait image animation. In *ECCV Workshop on Learning to Generate 3D Shapes and Scenes*, 2022. 2
 - [40] Connor Z. Lin, Koki Nagano, Jan Kautz, Eric R. Chan, Umar Iqbal, Leonidas Guibas, Gordon Wetzstein, and Sameh Khamis. Single-shot implicit morphable faces with consistent texture parameterization. In *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. 2
 - [41] Shanchuan Lin, Andrey Ryabtsev, Soumyadip Sengupta, Brian L Curless, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Real-time high-resolution background matting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8762–8771, 2021. 5
 - [42] Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and sample steps are flawed. *arXiv preprint arXiv:2305.08891*, 2023. 3
 - [43] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object, 2023. 3
 - [44] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhofer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Transactions on Graphics (ToG)*, 40(4):1–13, 2021. 5, 6, 14, 15
 - [45] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866. 2023. 3
 - [46] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. 4
 - [47] Rafał K. Mantiuk, György Denes, Alexandre Chapiro, Anton Kaplanyan, Gizem Rufo, Romain Bachy, Trisha Lian, and Anjul Patney. Fovvideovdp: A visible difference predictor for wide field-of-view video. *ACM Trans. Graph.*, 40(4), 2021. 5
 - [48] Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14297–14306, 2023. 8
 - [49] Dann Mensah, Nam Hee Kim, Miika Aittala, Samuli Laine, and Jaakko Lehtinen. A hybrid generator architecture for controllable face synthesis. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–10, 2023. 2
 - [50] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, Ying Shan, and Xiao Hu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023. 3
 - [51] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulo, Peter Kotschieder, and Matthias Nießner. Diffirf: Rendering-guided 3d radiance field diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4328–4338, 2023. 3
 - [52] Yotam Nitzan, Kfir Aberman, Qiurui He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. Mystyle: A personalized generative prior. *arXiv preprint arXiv:2203.17272*, 2022. 1, 2
 - [53] Mohit Mendiratta Pan, Mohamed Elgharib, Kartik Teotia, Ayush Tewari, Vladislav Golyanik, Adam Kortylewski,

- Christian Theobalt, et al. Avatarstudio: Text-driven editing of 3d dynamic human head avatars. *arXiv preprint arXiv:2306.00547*, 2023. 3
- [54] 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. 4
- [55] 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, 6
- [56] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 3
- [57] Bosheng Qin, Wentao Ye, Qifan Yu, Siliang Tang, and Yuet-ing Zhuang. Dancing avatar: Pose and text-guided human motion videos synthesis with image diffusion model. *arXiv preprint arXiv:2308.07749*, 2023. 3
- [58] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2
- [59] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2, 3, 5, 16
- [60] 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 preprint arXiv:1803.09179*, 2018. 14
- [61] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022. 1, 2, 3
- [62] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [63] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 3
- [64] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 2
- [65] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. 2023. 8
- [66] Michał Stypułkowski, Konstantinos Vougioukas, Sen He, Maciej Zieba, Stavros Petridis, and Maja Pantic. Diffused heads: Diffusion models beat gans on talking-face generation. *arXiv preprint arXiv:2301.03396*, 2023. 2
- [67] Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. Next3d: Generative neural texture rasterization for 3d-aware head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20991–21002, 2023. 2
- [68] David Svitov, Dmitrii Gudkov, Renat Bashirov, and Victor Lempitsky. Dinar: Diffusion inpainting of neural textures for one-shot human avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7062–7072, 2023. 2
- [69] Junshu Tang, Bo Zhang, Binxin Yang, Ting Zhang, Dong Chen, Lizhuang Ma, and Fang Wen. 3dfaceshop: Explicitly controllable 3d-aware portrait generation. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 2
- [70] Ayush Tewari, Mohamed Elgharib, Mallikarjun B R, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zöllhofer, and Christian Theobalt. Pie: Portrait image embedding for semantic control. 2020. 1, 2
- [71] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zöllhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images, cvpr 2020. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2020. 2
- [72] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics 2019 (TOG)*, 2019. 2, 3, 4, 5, 6, 7, 13, 14, 15
- [73] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4): 1–14, 2021. 1, 2
- [74] Lizhen Wang, Zhiyua Chen, Tao Yu, Chenguang Ma, Liang Li, and Yebin Liu. Faceverse: a fine-grained and detail-controllable 3d face morphable model from a hybrid dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR2022)*, 2022. 2
- [75] Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, et al. Rodin: A generative model for sculpting 3d digital avatars using diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4563–4573, 2023. 3
- [76] 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 Conference on Computer Vision and Pattern Recognition*, 2021. 2
- [77] Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, et al. Lavie: High-quality video generation with cascaded latent diffusion models. *arXiv preprint arXiv:2309.15103*, 2023. 2
- [78] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023. 2
- [79] Yue Wu, Yu Deng, Jiaolong Yang, Fangyun Wei, Chen Qifeng, and Xin Tong. Anifacegan: Animatable 3d-aware

- face image generation for video avatars. In *Advances in Neural Information Processing Systems*, 2022. 2
- [80] Cheng-hsin Wu, Ningyuan Zheng, Scott Ardisson, Rohan Bali, Danielle Belko, Eric Brockmeyer, Lucas Evans, Timothy Godisart, Hyowon Ha, Xuhua Huang, Alexander Hypes, Taylor Koska, Steven Krenn, Stephen Lombardi, Xiaomin Luo, Kevyn McPhail, Laura Millerschoen, Michal Perdoch, Mark Pitts, Alexander Richard, Jason Saragih, Junko Saragih, Takaaki Shiratori, Tomas Simon, Matt Stewart, Autumn Trimble, Xinshuo Weng, David Whitewolf, Chenglei Wu, Shoou-I Yu, and Yaser Sheikh. Multiface: A dataset for neural face rendering. In *arXiv*, 2022. 13
- [81] Hanshu Yan, Jun Hao Liew, Long Mai, Shanchuan Lin, and Jiashi Feng. Magicprop: Diffusion-based video editing via motion-aware appearance propagation. *arXiv preprint arXiv:2309.00908*, 2023. 2
- [82] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: a large-scale high quality 3d face dataset and detailed riggable 3d face prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [83] Shuai Yang, Liming Jiang, Ziwei Liu, and Chen Change Loy. Vtoonify: Controllable high-resolution portrait video style transfer. *ACM Transactions on Graphics (TOG)*, 41(6):1–15, 2022. 2
- [84] Shuai Yang, Yifan Zhou, Ziwei Liu, , and Chen Change Loy. Rerender a video: Zero-shot text-guided video-to-video translation. In *ACM SIGGRAPH Asia Conference Proceedings*, 2023. 2
- [85] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 3, 4
- [86] Bohan Zeng, Xuhui Liu, Sicheng Gao, Boyu Liu, Hong Li, Jianzhuang Liu, and Baochang Zhang. Face animation with an attribute-guided diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 628–637, 2023. 1, 2
- [87] Yifei Zeng, Yuanxun Lu, Xinya Ji, Yao Yao, Hao Zhu, and Xun Cao. Avatarbooth: High-quality and customizable 3d human avatar generation. *arXiv preprint arXiv:2306.09864*, 2023. 3
- [88] Chi Zhang, Yiwen Chen, Yijun Fu, Zhenglin Zhou, Gang Yu, Billzb Wang, Bin Fu, Tao Chen, Guosheng Lin, and Chunhua Shen. Styleavatar3d: Leveraging image-text diffusion models for high-fidelity 3d avatar generation. *arXiv preprint arXiv:2305.19012*, 2023. 3
- [89] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. 1, 2, 3, 16
- [90] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 5
- [91] Yinglin Zheng, Hao Yang, Ting Zhang, Jianmin Bao, Dongdong Chen, Yangyu Huang, Lu Yuan, Dong Chen, Ming Zeng, and Fang Wen. General facial representation learning in a visual-linguistic manner. *arXiv preprint arXiv:2112.03109*, 2021. 5
- [92] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C Bühler, Xu Chen, Michael J Black, and Otmar Hilliges. Im avatar: Implicit morphable head avatars from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13545–13555, 2022. 2
- [93] Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J Black, and Otmar Hilliges. Pointavatar: Deformable point-based head avatars from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21057–21067, 2023. 2
- [94] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2, 6
- [95] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Instant volumetric head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4574–4584, 2023. 2

Appendix

A. Additional Results

We provide additional qualitative results for the self-reenactment task on three more people from the NeRSemable [36] dataset in Fig. 11.

A.1. Novel View Synthesis Results

Tab. 4 contains results for novel view synthesis. We measure the performance from the unseen frontal view and average over the 24 train sequences for each of our 5 avatars. DiffusionAvatars scores the best metrics except for PSNR, which favors the overly blurry renderings of NeRFace. These results show that our method also provides good 3D control.

Method	PSNR \uparrow	LPIPS \downarrow	JOD \uparrow	AKD \downarrow	CSIM \uparrow
NeRFace [16]	28.12	0.284	7.93	2.72	0.841
DiffusionRig [12]	19.73	0.231	6.02	9.71	0.917
DNR [72]	26.90	0.213	7.91	1.58	0.933
DNR+GAN [72]	23.56	0.104	7.30	1.70	0.906
Ours	26.93	0.067	8.28	1.16	0.939

Table 4. **Quantitative Comparison:** We report metrics for novel-view-synthesis averaged over five persons from the NeRSemable dataset.

A.2. Ablation of Number of Training Views

We analyze the data efficiency of DiffusionAvatars on one person of the Multiface dataset [80] in Fig. 5. Our method only suffers negligibly from the removal of train views, as opposed to MVP. We attribute this to NPHM’s strong prior over 3D head shapes which allows fitting to sparse point-clouds, while direct 3D methods like MVP struggle more to recover accurate surfaces from sparse views.




Figure 5. Ablation of number of cameras on the Multiface dataset.

A.3. Qualitative Comparison of Ablations

As an addition to Tab. 3 of the main paper, we provide qualitative comparisons of three ablations in Fig. 6, Fig. 7 and Fig. 8.




Figure 6. **Effect of FLAME vs. NPHM.** Using FLAME instead of NPHM as the underlying morphable head model makes the generated faces less expressive. This is due to FLAME’s limited geometry and expression space, which does not provide enough geometric cues for our diffusion-based neural renderer.




Figure 7. **Effect of expression conditioning.** Directly providing z_{exp} to the diffusion model helps synthesize areas not modeled by NPHM’s geometry, such as eye movement.




Figure 8. **Effect of LDM prior.** Using a pre-trained LDM instead of training the diffusion model from scratch steers the generation towards more realistic faces.

A.4. Experiments on the Multiface dataset

We conduct additional experiments on 16 cameras from the Multiface dataset [80] in Tab. 5. Although MVP has the advantage of a high-end textured mesh tracking pre-computed from 100+ cameras there, DiffusionAvatars performs competitively showing that our method (i) can compete with recent 3D methods and (ii) generalizes to different datasets. In Fig. 9, we show that our method produces sharper renderings than MVP on the Multiface dataset.

	PSNR↑	LPIPS↓	JOD↑	AKD↓	AED↓	APD↓	CSIM↑
MVP	26.4	0.126	7.92	2.25	0.655	0.026	0.904
Ours	25.5	0.077	7.77	1.86	0.503	0.025	0.914

Table 5. Quantitative comparison for self-reenactment averaged over 2 persons from the Multiface dataset.




Figure 9. Qualitative comparison for self-reenactment on the Multiface dataset.

B. User Study Setup

In Fig. 10, we show the interface of our user study. We ask users to rate visual quality and driving fidelity by providing them with the following definitions:

Visual Quality. The overall quality of the generated video: Is there any flickering, visual artifacts, blurriness, etc., that feel unnatural?

Driving Fidelity. How closely does the generated video follow the desired facial expressions: Are the displayed expressions the same? Is the emotion displayed by the avatar recognizable?

During the study, each user is presented with 15 randomly chosen pairs of driving sequences (left) and avatar animations (right). The pairs were drawn from 135 different avatar animations generated by the baselines and our method. In total, we received 735 responses for each question from 49 participants.

C. Analysis of Temporal Consistency

Since DiffusionAvatars is based on 2D diffusion models, the generated images may be susceptible to view-inconsistencies and screen-space stitching artifacts. Fol-

DiffusionAvatar

Thank you for your participation! In the following you will see examples of expression transfer from a person performing the desired expressions to an avatar.

Figure 10. **Setup of our User Study.** We ask participants to evaluate the avatar animation task in which an avatar is controlled by a sequence of a different person.

lowing StyleGAN3 [32], we therefore perform an analysis of the temporal consistency of generated videos. In figure Fig. 12, we compare our method to MVP [44] and Deferred Neural Rendering [72]. Our finding is that Diffusion-Avatars generates images with more high-frequency detail than 3D methods while being significantly less prone to texture-sticking than other 2D methods.

D. Model Architecture

Fig. 13 depicts DiffusionAvatars’s model architecture.

E. Societal Impact

DiffusionAvatars provides the means to generate realistic images of faces with control over viewpoint, pose, and expression. As such, an avatar could be misused to the original person’s disadvantage. However, we focus on 3D applications such as immersive teleconferencing or character animation, not indistinguishable 2D video synthesis. Our method does not model lighting, the background, or the torso. Therefore, we believe a detection system can be built to identify generations of our method by exploiting such characteristics [60].

All subjects shown in the paper and the supplemental video are part of the NeRSemble dataset [36] and have consented to their recordings being used for research purposes.




Figure 11. **Qualitative results for self-reenactment.**




Figure 12. **Analysis of Temporal Consistency.** We extract a short vertical segment of pixels from each frame of a generated video and concatenate them left-to-right (bottom). The desired result is that texture details move together with the underlying head geometry. It can be clearly seen that combining Deferred Neural Rendering with an adversarial loss (DNR+GAN) causes severe texture sticking as the person rotates their head to the left. On the other hand, MVP and our method do not suffer from this issue. Additionally, as opposed to MVP, our method generates much of the high-frequency detail that is also present in the ground truth video.

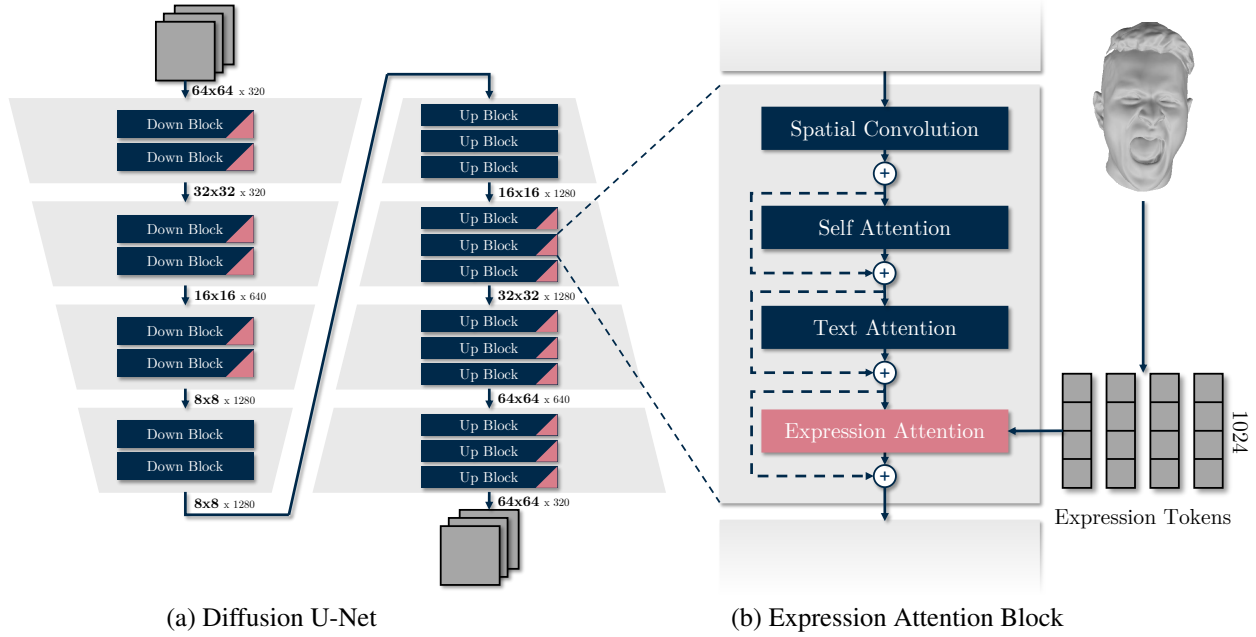


Figure 13. **Architecture of DiffusionAvatars.** We build on the U-Net architecture of a pre-trained Stable Diffusion [59] model. While we keep all original layers fixed, we introduce an additional cross attention layer in each block that allows spatial features to attend to expression tokens extracted from NPHM’s expression codes. That way, subtle expression details can influence the image formation process. Additionally, following the ControlNet [89] paradigm, the encoder part of the diffusion U-Net is copied and conditioned on the NPHM rasterizations. The feature maps produced by this ControlNet are added to the existing skip connections of the U-Net, allowing the diffusion model to generate articulated facial images while keeping the knowledge from its pre-training.