

## AvatarArtist: Open-Domain 4D Avatarization

Hongyu Liu<sup>1,2,\*</sup> Xuan Wang<sup>2,§</sup> Ziyu Wan<sup>3</sup> Yue Ma<sup>1</sup> Jingye Chen<sup>1</sup> Yanbo Fan<sup>2</sup>  
 Yujun Shen<sup>2</sup> Yibing Song Qifeng Chen<sup>1,§</sup>

<sup>1</sup>HKUST <sup>2</sup>Ant Group <sup>3</sup>City University of Hong Kong

<https://kumapowerliu.github.io/AvatarArtist>



Figure 1. **Gallery of the proposed AvatarArtist.** Each row features several triplets, where the first column of each triplet is the source image. The subsequent two images in each triplet are results that follow the pose and expression of the driving image, as demonstrated in the bottom right corner of the first three columns. Specifically, our method is applicable to an open domain, encompassing a diverse range of categories including 3D cartoons, video game characters, sculptures, skulls, etc.

### Abstract

*This work focuses on open-domain 4D avatarization, with the purpose of creating a 4D avatar from a portrait image in an arbitrary style. We select parametric triplanes as the intermediate 4D representation, and propose a practical training paradigm that takes advantage of both generative adversarial networks (GANs) and diffusion models. Our design stems from the observation that 4D GANs excel at bridging images and triplanes without*

*supervision yet usually face challenges in handling diverse data distributions. A robust 2D diffusion prior emerges as the solution, assisting the GAN in transferring its expertise across various domains. The synergy between these experts permits the construction of a multi-domain image-triplane dataset, which drives the development of a general 4D avatar creator. Extensive experiments suggest that our model, termed AvatarArtist, is capable of producing high-quality 4D avatars with strong robustness to various source image domains. The code, the data, and the models*

will be made publicly available to facilitate future studies.

## 1. Introduction

Avatarization (dynamic) from one single portrait image has become a fundamental ability of AI content generation. It enables the transfer of motion and expression from a source video to a digital avatar while preserving both motion accuracy and subject identity. This technology has broad applications in virtual reality, social media, gaming, and online education, facilitating efficient character production and enhancing interactive experiences in computer vision and computer graphics.

Studies on avatarization can mainly be categorized as 2D and 4D aspects. The 2D-based methods [20, 56–58, 76, 77, 82] typically employ a self-supervised learning scheme, with monocular video stream data accompanied by facial landmarks or implicit motion representations [57]. More recently, the emergence of powerful generative models, such as diffusion models, which can handle various types of images, has further advanced the field. Some 2D methods [42, 67, 71] incorporate prior knowledge from diffusion models (e.g., Stable Diffusion [54]), enabling them to effectively handle multi-style avatarization (e.g., cartoon, realistic).<sup>\*</sup> Despite achieving impressive results, these 2D methods fail to accurately represent 3D structures. Geometric distortion and content inconsistency often arise when the head pose undergoes significant rotation. Moreover, the iterative computation of diffusion models incurs substantial costs for generating each frame of animated videos, significantly increasing the overall computational burden.

On the other hand, 4D-based methods [9, 14, 33, 43, 84] leverage neural rendering pipelines [27, 45] and 3DMM [32] for efficient avatarization where 3D geometric consistency is maintained across multiple viewpoints. During model inference, these models animate the image feature first then a camera pose to perform neural rendering of target view generations. Despite the demonstrated success, these methods suffer from a lack of 4D data from diverse domains. The human portrait animation is restricted to a limited domain and is difficult to generalize as that of 2D-based methods.

*“Having examined both 2D and 4D-based avatarization methods, we intuitively assume that if sufficient and well-suited 4D datasets covering diverse domains were available, it would be possible to develop a 4D avatarization approach for open-domain inputs using diffusion models.”* Recently, Rodin [65, 78], a diffusion-based single-image-guided static avatar generation method, has demonstrated

impressive performance in the synthetic digital domain. To train this model, a dataset of image-3D representation pairs was constructed using multi-view digital human data. Inspired by this, we believe that an appropriate 4D dataset for our method should consist of image-4D representation pairs spanning multiple domains.

In this work, we propose AvatarArtist, a diffusion-based 4D avatarization model. It is challenging to obtain multi-view, multi-expression 4D captures to create image-4D representation pairs with a fitting process similar to Rodin. Therefore, we resort to synthetic data generation. Fortunately, we found that 4D GANs, such as Next3D [60, 84], can greatly assist in this process. Specifically, Next3D proposed a parametric triplane 4D representation, which divides the traditional triplane [7] into dynamic and static components. The dynamic part is aligned with the 3DMM mesh [32, 64] in UV space, allowing expression changes through mesh rasterization and rendering. With Next3D, we can generate an unlimited amount of image-parametric triplane data simply by sampling, but only for single realistic domain due to the mode collapse issue of GAN. Hence we propose to finetune Next3D to efficiently obtain multiple GANs of diverse domains. While training Next3D only requires 2D images and their corresponding 3DMM meshes, achieving effective multi-domain fine-tuning demands diverse and comprehensive data coverage across various domains. To overcome this limitation, we leverage 2D diffusion models [54] to enrich the diversity of the training data. Specifically, we utilize the SDEdit [44] pipeline and landmark-guided ControlNet to transfer portrait images (e.g., FFHQ) from the realistic domain to other domains. This process ensures coherent pose and expression between the output and input 2D portraits, allowing us to reuse the 3DMM mesh of the 2D portrait from the realistic domain in non-realistic domains. Consequently, we can train 4D GANs for different domains and generate image-parametric triplane datasets across multiple domains. The entire data generation process combines the advantages of both diffusion models and GANs: diffusion models provide multi-domain data for the GAN, while the GAN transforms 2D images into 4D representations in an unsupervised manner.

Using this dataset, we could adopt the latent Diffusion Transformer (DiT) [49] to model its distribution. The process begins with training a VAE to compress triplanes into latent representations, followed by employing a DiT to generate latent guided by a single portrait image. Although the diffusion model is able to generate triplanes effectively, there are still two issues for rendering high-quality frames from these planes. First, not like Rodin which uses a simple MLP decoder for the digital domain, rendering triplanes from multiple diverse domains into high-quality images is much more challenging. Second, Parametric triplanes

<sup>\*</sup> This work is done partially when Hongyu is an intern at Ant Group.

§ Joint corresponding authors.

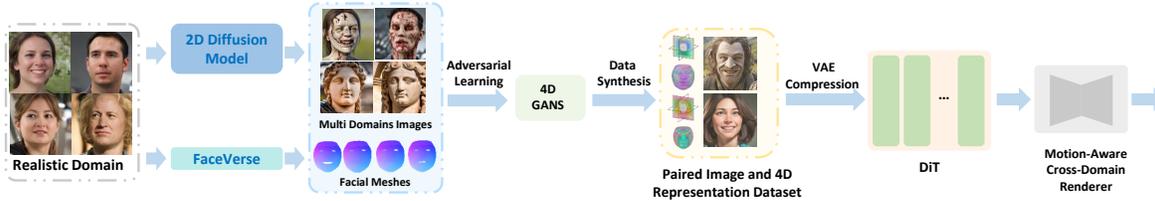


Figure 2. **The overall training pipeline of our method.** We first generate 2D images from different domains using a 2D diffusion model. These images are then used to train 4D GANs for each domain. Subsequently, the trained 4D GANs generate image-4D representation pairs across domains, which are used to train DIT and the rendering model.

primarily focus on motion modeling but are less effective in preserving identity. Next3D employs a CNN to enhance identity preservation, but we found the performance of CNN degrades significantly in the open domain. To address these, we introduce a motion-aware cross-domain renderer based on ViT [69] that incorporates features from the source image, improving cross-domain rendering ability and preserving the identity information. Additionally, we use an implicit motion representation, similar to Portrait4D [14], to avoid artifacts caused by mesh inaccuracies. Compared to baseline methods, our approach delivers superior quantitative and qualitative performance, offering high visual fidelity, accurate identity representation, and precise motion rendering.

## 2. Related Work

We address one-shot, open-domain image-driven talking face generation, which synthesizes a talking head video from a single reference portrait and a sequence of driving expression images. This section provides a concise overview of previous talking head generation methods, broadly categorized into 2D talking face generation and 3D-aware talking portrait synthesis, along with a brief discussion on stylized 3D avatar generation.

### 2.1. 2D Talking Face Generation

Great progress has been made in image-driven 2D talking head generation [4, 16, 19, 36, 41, 56–58, 63, 72, 74, 76, 77]. Numerous approaches harness the capabilities of Generative Adversarial Networks to synthesize high-fidelity talking head videos, most of which fall into the warping-then-rendering scheme. The identity features are first encoded from the reference image and then warped according to the driving signals, finally being rendered into a sequence of talking portraits. More specifically, various types of motion representation, such as landmarks [57, 76], depth [24], and latent code [4], are exploited to deduce the warping field, ensuring that the synthesized portraits exhibit expressions and motions that faithfully correspond to the driving signals. With the advent of diffusion model-based image generation, several methods employ large pre-trained diffusion models to assist in the

task of one-shot talking face generation. By leveraging the powerful prior of pre-trained diffusion models, recent methods [42, 67, 71] have demonstrated that they possess strong generalization capabilities when handling various styles of reference portraits. However, due to a lack of understanding of three-dimensional structures, these 2D-based methods often exhibit obvious geometric distortions when handling larger head movements. Additionally, they lack the ability to control the viewpoint of the generated images with precision.

### 2.2. 3D-aware Talking Portrait Synthesis

To achieve high geometric fidelity in synthesizing portraits with varying head poses, researchers have introduced intermediate 3D representations that capture facial geometry and pose, ensuring structural accuracy across viewpoints. A major breakthrough in novel view synthesis is Neural Radiance Fields (NeRF)[7, 10, 25, 33, 34, 43, 45, 61, 73, 75, 86], which enables precise 3D reconstructions with explicit camera control. NeRF has been widely adopted in 3D-aware one-shot talking head generation, enhancing 3D coherence and pose control for more natural outputs. More recently, GAGAvatar[9] leveraged 3D Gaussian Splatting (3DGS) to accelerate generation while maintaining high quality.

However, most methods [66] rely on in-the-wild video data, making 3D learning from monocular videos highly ill-posed due to depth ambiguity, lighting variations, and facial occlusions. Some approaches incorporate 3D supervision from monocular 3D face reconstruction [11, 13, 17], multi-view lab-captured videos[25, 86], or synthetic multi-view data[14, 15]. While these improve results, they are constrained by limited high-quality 3D data and training challenges. As a result, there remains no open-domain, one-shot 4D portrait generation method capable of generalizing across diverse conditions.

### 2.3. Stylized Avatar Generation

To generate avatars across different domains, some methods [1, 3, 28, 29, 31, 37, 50, 59, 62] use CLIP as a constraint or leverage diffusion models to generate reference images, which are then utilized to create stylized avatars based

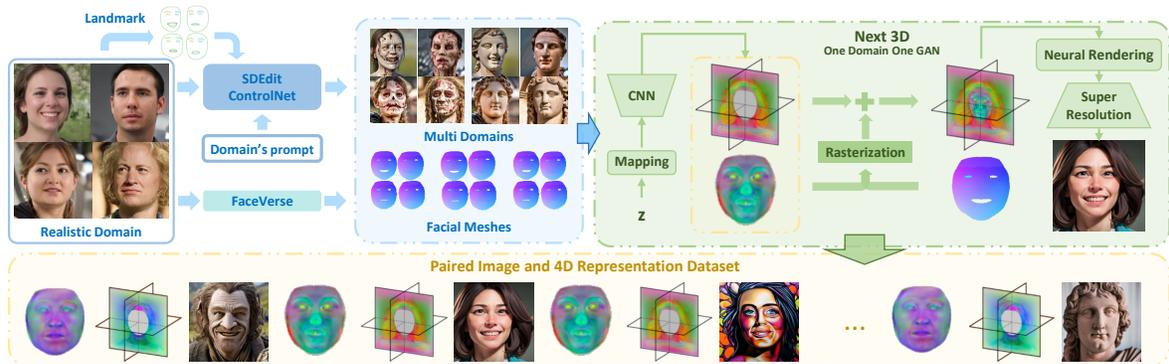


Figure 3. **The pipeline of dataset generation.** We use text prompts to transform images from the realistic domain to the target domain while ensuring pose and expression consistency with SDEdit [44] and landmark-guided ControlNet [80]. This enables direct reuse of the original mesh, avoiding errors in non-realistic domain extraction. After domain transfer, we train 4D GANs to generate image-parametric triplane pairs, which serve as data for the next stage. The parametric triplane comprises dynamic and static components, with the dynamic region aligned to the mesh.

on text prompts. Additionally, StyleAvatar3D [79] and Rodin [65, 78] collect domain-specific datasets to train generative models for stylized avatar synthesis. While these methods significantly advance stylized avatar generation, they do not focus on single-image-guided, animatable avatar creation. Meanwhile, other approaches [2, 18, 48, 85] employ CLIP as a constraint and text as guidance to fine-tune GAN models, enabling the generation of stylized portrait images that align with textual descriptions. Although these methods demonstrate strong manipulation capabilities for stylized portraits, they cannot directly generate avatars.

### 3. Method

We aim to develop a system that generates a 4D avatar from an open-domain image  $I_s$ , driven by the motion of a target individual  $I_t$ . The key to training a deep generative model for open-domain avatarization is a large-scale, high-quality dataset. In Sec.3.1, we introduce how GANs and image generation techniques help construct diverse and consistent training data. With this dataset, we use a latent Diffusion Transformer (DiT) to model the 4D distribution (Sec.3.2). To ensure accurate motion transfer while preserving the source identity, we further employ a motion-aware cross-domain renderer (Sec.3.3). The overall training pipeline is shown in Figure 2. Next, we detail each component.

#### 3.1. Data Curation from 4D GANs

Benefiting from adversarial training, the recent GAN methods have demonstrated great potential in generating high-quality 4D avatars in an unsupervised manner using non-multiview images and 3DMM meshes only. Therefore, we would like to fully leverage this capability of GANs to curate 4D data. Nonetheless, the instability of GAN

training easily caused mode-collapse, failing to cover the distribution of different modes. In this section, we will discuss how to properly use GAN to generate open-domain image-4D representation pairs data.

**Base GAN Model.** We select Next3D [60, 84] as our base GAN model for generating the 4D dataset, given its training efficiency and the proposed robust 4D representation (parametric triplane). Specifically, as shown in Figure 3, given a randomly sampled latent code  $z$ , the mapping network translates  $z$  into an intermediate latent vector, which will modulate conv layers of StyleGAN to generate a parametric triplane  $\in \mathbb{R}^{256 \times 256 \times 4 \times 32}$ . The parametric triplane consists of two parts: a static component representing non-facial regions and a dynamic component aligned with the 3DMM mesh in UV space. During inference, given a specific mesh, the dynamic part is deformed through rasterization and combined with the static part to form a triplane with expressions. Then neural rendering and a super-resolution module are applied to generate the final image. We follow Next3D and use FaceVerse [64] to extract the corresponding parametric mesh. **Multi-Domain Tuning.** To train Next3D models across various domains, the first priority is to obtain diversified images from different domains and extract the corresponding 3DMM meshes. However, it is very challenging to accurately obtain 3DMM meshes for non-realistic portraits. To address this issue, we use the pre-trained 2D diffusion model to generate the target domain from realistic images given specific prompts, so that the 3D meshes from realistic domains can be re-used. Specifically, given a portrait from the realistic domain, we add noise over this image with specific strength [44] and then denoise it by StableDiffusion to generate high-quality and diversified images of the target domain under the guidance of prompts. However, corrupting the images with noise also raises challenges in maintaining the original

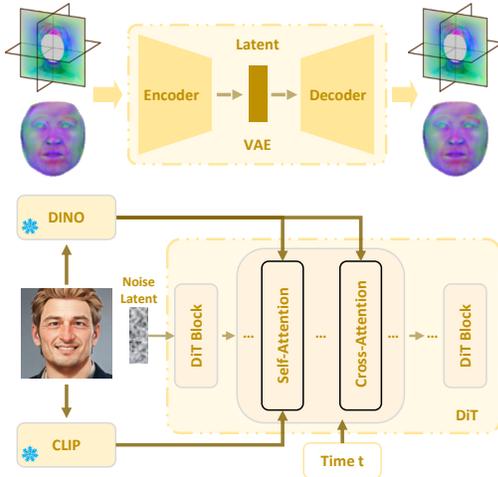


Figure 4. **The pipeline of DiT.** We first train a VAE to compress the parametric triplane into a latent space, and then train a DiT to denoise the noisy latent. We incorporate features from DINO [5] and CLIP [53] into the DiT to guide the generation process.

expression. Hence, we also incorporate the facial landmark through ControlNet [80], which provides the conditional signals of expression. As a result, the newly sampled image from the target domain could closely align with the pose and expression of the source realistic image, enabling the direct transfer of 3DMM mesh labels from the realistic domain to the target domain.

Through the pipeline mentioned above, we totally collected image data from 28 different domains (anime, lego, etc.) transferred from FFHQ [26]. To ensure the efficiency of the data pipeline and avoid model collapse, for each domain, we generate 6,000 images only and use this data to finetune independent GAN from the Next3D model trained with FFHQ [26]. We follow the DATID-3D [28] to use the ADA loss and density regularization loss to guarantee the diverse content generation ability of GAN during tuning.

**Data Synthesis.** We utilize the trained multiple 4D GANs to build two datasets. 1) The image-parametric triplane paired dataset. We randomly sampled poses and meshes of portrait images from the FFHQ dataset, which are then fed into the 4D GAN with a random  $z$  to generate images in different identities along with corresponding triplanes. For each domain, we generated 20K samples, resulting in a total of  $20K \times 28 = 560K$  image-triplane pairs. 2) The multi-view, multi-expression image-parametric triplane dataset. These data assist in learning our motion-aware cross-style renderer. We generate both static and dynamic components following the portrait4D [14]. The dynamic data, which are responsible for head reenactment, consists of synthetic identities with multiple expressions per subject and varying camera poses for each expression. The static data, on the

other hand, are employed to enhance the generalizability of 3D reconstruction and contain only a single expression per identity, also with varying camera poses. Expressions (meshes) in the dynamic dataset are sampled from the VFHQ dataset [70], while those in the static dataset are sampled from the FFHQ dataset. All camera poses are sampled from FFHQ.

### 3.2. 4D Generation

Recently, the latent diffusion model has shown great potential in modeling complex data distributions like images [51], videos [52], and triplanes [6, 55, 65]. Following this trend, this section will depict how we leverage latent diffusion for 4D generation. As shown in Figure 4, we will first introduce a triplane VAE to compress the triplane representations into a latent space, followed by training an image-conditioned DiT [49]. All the data used for training was sourced from our curated datasets.

**Triplane VAE.** For training efficiency, the DiT [49] is trained in a compact latent space by default. To achieve this, we propose a triplane variational autoencoder (VAE) [30] to obtain the latent code of the triplane representations. Specifically, our VAE compresses the triplane  $\in \mathbb{R}^{256 \times 256 \times 4 \times 32}$  to latent  $z_t \in \mathbb{R}^{64 \times 64 \times 4 \times 8}$ . To optimize the VAE model, we compute the  $\mathcal{L}_1$  loss between the reconstructed planes and input triplanes. Meanwhile, we also get depth and rendered images to calculate  $\mathcal{L}_1$  and LPIPS losses, respectively. We did not apply adversarial loss since we found it introduced training instability. For more details, please refer to our supplementary material.

**Image Guided Diffusion Transformer.** We follow the Direct3D [68] and PixArt- $\alpha$  [8] to build our image-guided DiT. For the noised latent  $z_t$  we flatten it to a sequence and send it to the DiT as input. We separately extract the semantic and detailed information from conditional images and inject them into each DiT block. For semantic information, we use CLIP [53] to extract the image’s semantic embeddings, which are then integrated with the model via cross-attention. To capture fine-grained details, we employ the DINO [47] to extract image tokens. In each DiT block, we concatenate the image tokens with the flattened  $z_t$  and feed them into a self-attention layer to model the intrinsic relationship between the image tokens and  $z_t$ . During training, we leverage the objective of IDDPM [46] and predict the noise and variance at each time step  $t$ . We also randomly drop the conditional image with a probability of 10% to enable classifier-free guidance [23] during inference.

### 3.3. Motion-Aware Cross-Domain Renderer

In the rendering process of Next3D, a CNN refines the rasterized parametric triplane to protect the identity information and eliminate identity leakage caused by rasteriza-

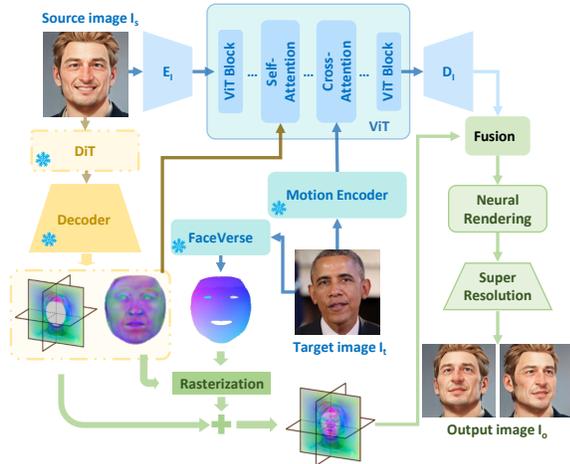


Figure 5. **The pipeline of motion-aware cross-domain renderer.** We use an encoder to extract the feature from the source image. This feature is sent to a ViT to predict results under the guidance of generated parametric triplane and motion embedding. Finally, a decoder processes the output of the ViT and fuses it with the results of rasterization to produce the final output.

tion. However, we found that this rendering approach fails to achieve acceptable quality in our setting (see Figure 8). Since our parametric triplane is generated from different domains, a simple rendering network cannot effectively resolve identity leakage across various domains. Additionally, inaccurate mesh extraction sometimes leads to mismatched expressions in the generated results.

To address these issues, we propose a motion-aware cross-domain renderer. As shown in Figure 5, we first employ an encoder  $E_I$  to extract features from the source images, which are subsequently fed into a Vision Transformer (ViT) model [69]. In the ViT model, we inject the parametric triplane generated by DiT into the self-attention mechanism, which aims to neutralize facial expressions and canonicalize poses inspired by [14], thereby eliminating expression-specific information from the source image. Then, we change the expression with motion embedding [63] by injecting it with cross attention. This embedding is an implicit representation without spatial information, thus preventing identity leakage. The output of the ViT is decoded to match the resolution of the rasterized parametric triplane, after which it is fused with the triplane. Finally, we apply volumetric rendering followed by super-resolution techniques to generate the final output  $I_o$ . The  $I_o$ , rendered from a novel camera viewpoint, preserves the identity features from the source image  $I_s$  meanwhile following the facial expression of the target image  $I_t$ . For training this model, we adopt the loss terms following the [14, 61, 83]. For more details, please refer to our supplementary material.

## 4. Experiments

In this section, we first illustrate our implementation details. Then, we compare our method with existing methods qualitatively and quantitatively. We compare our approach with both 2D and 4D reenactment methods. Specifically, we include comparisons with 2D techniques such as LivePortrait [20] and XPortrait [71], as well as 4D methods like InvertAvatar [84] and Portrait4Dv2 [15]. Finally, an ablation study validates the effectiveness of our contributions. More results are provided in the supplementary files.

### 4.1. Implementation Details

During the training of the Next3D, we extract facial poses and corresponding 3DMM meshes from the FFHQ dataset using FaceVerse [64]. All domains are fine-tuned based on a GAN pre-trained on the FFHQ dataset, with each domain iterating over a total of 300K images. For VAE training, we adopt the same training framework as the VAE used in Stable Diffusion. We utilize the AdamW optimizer [38] with a learning rate of  $1e-4$ . The VAE model is trained on an NVIDIA A100 (80G) GPU for 100K steps with a batch size of 32. Our diffusion model follows the network configuration of DiT-XL/2 [8, 49, 68], consisting of 28 layers of DiT blocks. The diffusion model is trained with 1000 denoising steps using a linear variance scheduler. We employ the AdamW optimizer with a learning rate of  $1e^{-4}$  and train the model for 800K steps. During inference, we use 19 steps of the DPMSolver [39], with a guidance scale set to 4.5. For the motion-aware cross-domain renderer, we train on a total of 12 million images across all domains. For more details, please refer to the supplementary materials.

### 4.2. Qualitative Results

As shown in Figure 6, we present a visual comparison of the results of self-reenactment and cross-reenactment tasks. The first column contains the input images, with the bottom-right corner showing the target image and the larger images representing the source images. The first row displays the self-reenactment results. We observe that InvertAvatar exhibits noticeable artifacts, while XPortrait shows misalignment in pose compared to the target. Although Portrait4D and LivePortrait achieve relatively good results, there are inconsistencies in the expression, particularly in the eye and mouth regions, when compared to the target. In contrast, our method produces more consistent results, achieving better alignment with the target in both pose and expression.

For the cross-reenactment, our source images are from non-realistic domains, while the target expressions are extracted from real-human domains. We observe that both InvertAvatar and Portrait4D struggle to handle portraits that significantly differ from real-human domains effectively. InvertAvatar tends to exhibit severe geometric distortions

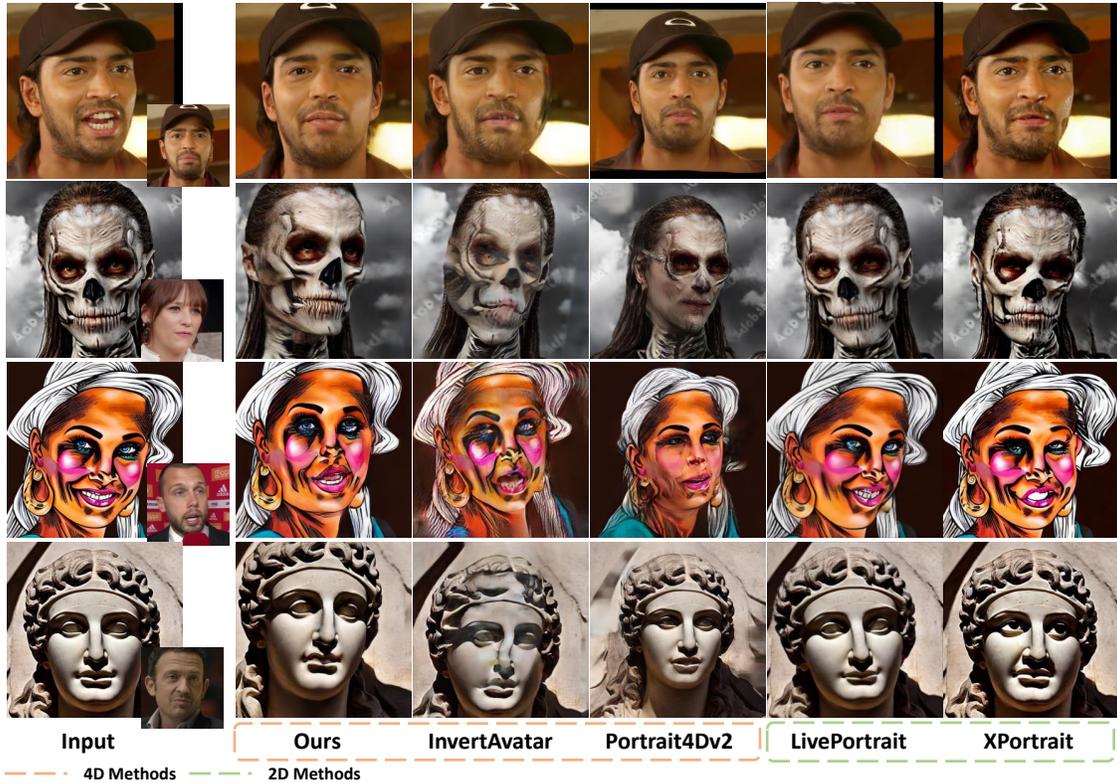


Figure 6. Qualitative comparison with SOTA methods. The leftmost column in the figure shows the input images, with the bottom-right corner representing the target image. The first row displays the results of self-reenactment, while the following three rows show the results of cross-reenactment. It can be observed that our method achieves superior performance in terms of expression and pose consistency, as well as identity preservation.

Table 1. Quantitative evaluation of state-of-the-art methods and our approach on the VFHQ dataset [70]. For self-reenactment, both the source and target images are from the VFHQ dataset. For cross-ID reenactment, the source images are generated from different domains, while the target motions are from VFHQ.  $\downarrow$  indicates lower is better while  $\uparrow$  indicates higher is better. **Red** highlights the best result, and **Blue** highlights the second-best result.

Method	Self reenactment					Cross reenactment			
	LPIPS $\downarrow$	FID $\downarrow$	ID $\uparrow$	AED $\downarrow$	APD $\downarrow$	FID $\downarrow$	CLIP $\uparrow$	AKD $\downarrow$	APD $\downarrow$
LivePortrait [20]	0.27	<b>46.49</b>	<b>0.65</b>	<b>0.025</b>	<b>4.28</b>	100.3	<b>0.91</b>	4.92	139.35
XPortrait [71]	0.31	60.29	0.63	0.036	18.07	<b>78.6</b>	<b>0.89</b>	10.67	237.4
InvertAvatar [84]	0.42	84.71	0.32	0.049	15.58	194.7	0.64	20.78	134.9
Portrait4Dv2 [15]	0.29	66.60	0.58	0.034	<b>5.08</b>	140.5	0.75	7.13	63.3
Ours	<b>0.26</b>	<b>52.62</b>	<b>0.69</b>	<b>0.032</b>	11.72	<b>89.3</b>	0.84	<b>2.58</b>	<b>52.3</b>

and fails to adequately animate the source image. Portrait4D, on the other hand, suffers from identity leakage and generates content that lacks precision. While 2D-based methods preserve the identity of the input image, they fail to ensure that the pose aligns with the target image. In contrast, our method demonstrates exceptional performance when handling non-realistic domains, achieving good accuracy in both expression and pose consistency, as well as identity preservation.

### 4.3. Quantitative Results

The quantitative results are summarized in Table 1. We evaluate our method on 100 VFHQ video clips [70] through self-reenactment and cross-reenactment tests. For self-reenactment, the source image is either the first frame or a random intermediate frame from the video, while for cross-reenactment, we use 50 source images from different domains with target images from VFHQ. To

Table 2. Ablation study on the FFHQ [26] dataset. The source images are generated from different domains, while the target images are from FFHQ. The Next3D rendering means replacing our render model with simple CNN.

Method	Cross reenactment			
	FID ↓	CLIP ↑	AKD ↓	APD ↓
Next3D render	130.72	0.73	5.89	42.93
Ours	<b>68.69</b>	<b>0.86</b>	<b>2.56</b>	<b>40.89</b>



Figure 7. Visualization of ablation study on data generation methods. It is only when combining SDEdit and ControlNet that we can ensure the generated images retain both the same expression and pose as the original images. The corresponding prompts are shown above images.

to assess image quality, we compute LPIPS[81] and FID[22]. Identity consistency is measured using the ID metric[12] for self-reenactment and CLIPScore[53] for cross-reenactment, as the ID metric is unreliable for non-human domains. Expression accuracy is evaluated with Average Expression Distance (AED)[35] for self-reenactment and Average Keypoint Distance (AKD)[40] for cross-reenactment, as 3DMM struggles with non-realistic humans. Additionally, Average Pose Distance (APD) is used to assess pose consistency, with pose information extracted using[21]. As shown in Table 1, our method performs slightly worse than 2D approaches in self-reenactment, but remains comparable while surpassing 4D methods in overall effectiveness. In cross-reenactment, although 2D methods better preserve identity, our approach achieves higher accuracy in capturing pose and expression, demonstrating the advantages of 4D-based techniques.

#### 4.4. Ablation Study

We analyze the impact of different data generation pipelines and the performance of each module in our model.

**Effectiveness of Different Data Generation Methods.** As shown in Figure 7, the leftmost column presents the input images, all from the realistic domain. Without SDEdit, ControlNet provides some control over expressions, but the generated results still deviate significantly from the originals (w/o SDEdit). When using only SDEdit without ControlNet, the results preserve the pose, but the expressions still show noticeable discrepancies (w/o ControlNet).



Figure 8. Visualization of ablation study on motion-aware cross-domain rendering. The Next3D rendering approach involves using a CNN as a replacement for our render model.

By combining ControlNet and SDEdit, we achieve images that maintain both the expression and pose of the original, while shifting entirely to a different domain (Full). This enables the reuse of 3DMM data from the realistic domain to train 4DGANs in various domains.

**Effectiveness of Models.** We design the motion-aware cross-domain renderer to better capture the appearance information from the original image, thereby enhancing fidelity. Additionally, since 3DMM is not perfectly accurate, we incorporate motion embedding to assist in the animation process. As shown in Figure 8, we replaced our renderer with a CNN similar to the one used in Next3D. The results exhibited significant identity leakage (i.e., facial artifacts resembling the target subject’s mesh), and the generated expressions did not accurately match the target (e.g., eye regions in the second row). In contrast, our method better preserves the source identity, and the implicit motion embedding effectively corrects motion inaccuracies from the mesh. The corresponding quantitative metrics in Table 2 show that our approach outperforms all compared methods across all evaluated metrics.

## 5. Conclusion

We propose AvatarArtist, a 4D avatar generation model for open-domain inputs. We use a parameterized triplane as a 4D representation and employ 4D GANs to build an open-domain image-triplane dataset. Specifically, a 2D diffusion model generates images from various domains, which train domain-specific 4DGANs to produce data for our model. Our model consists of DiT and a motion-aware cross-domain renderer. DiT converts the input image into a parameterized triplane, while the renderer refinement module synthesizes and optimizes results. Experiments show that AvatarArtist effectively handles open-domain inputs, successfully transferring target motion while preserving appearance and geometric consistency.

## 6. Acknowledgment

This project was supported by the National Key R&D Program of China under grant number 2022ZD0161501.

## References

- [1] Rameen Abdal, Hsin-Ying Lee, Peihao Zhu, Menglei Chai, Aliaksandr Siarohin, Peter Wonka, and Sergey Tulyakov. 3davatar: Bridging domains for personalized editable avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4552–4562, 2023. 3
- [2] Aibek Alanov, Vadim Titov, and Dmitry P Vetrov. Hyperdomainnet: Universal domain adaptation for generative adversarial networks. *Advances in Neural Information Processing Systems*, 35:29414–29426, 2022. 4
- [3] Qingyan Bai, Zifan Shi, Yinghao Xu, Hao Ouyang, Qiuyu Wang, Ceyuan Yang, Xuan Wang, Gordon Wetzstein, Yujun Shen, and Qifeng Chen. Real-time 3d-aware portrait editing from a single image. In *European Conference on Computer Vision*, pages 344–362. Springer, 2024. 3
- [4] Egor Burkov, Igor Pasechnik, Artur Grigorev, and Victor Lempitsky. Neural head reenactment with latent pose descriptors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13786–13795, 2020. 3
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 5
- [6] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *arXiv*, 2021. 5
- [7] 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, 3
- [8] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis, 2023. 5, 6
- [9] Xuangeng Chu and Tatsuya Harada. Generalizable and animatable gaussian head avatar. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 2, 3
- [10] 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. 3
- [11] Radek Daněček, Michael J Black, and Timo Bolkart. Emoca: Emotion driven monocular face capture and animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20311–20322, 2022. 3
- [12] 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. 8
- [13] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. 3
- [14] Yu Deng, Duomin Wang, Xiaohang Ren, Xingyu Chen, and Baoyuan Wang. Portrait4d: Learning one-shot 4d head avatar synthesis using synthetic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7119–7130, 2024. 2, 3, 5, 6
- [15] Yu Deng, Duomin Wang, and Baoyuan Wang. Portrait4d-v2: Pseudo multi-view data creates better 4d head synthesizer. *arXiv preprint arXiv:2403.13570*, 2024. 3, 6, 7
- [16] Nikita Drobyshev, Jenya Chelishev, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. Megaportraits: One-shot megapixel neural head avatars. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2663–2671, 2022. 3
- [17] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (ToG)*, 40(4): 1–13, 2021. 3
- [18] Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 4
- [19] Yuan Gong, Yong Zhang, Xiaodong Cun, Fei Yin, Yanbo Fan, Xuan Wang, Baoyuan Wu, and Yujiu Yang. Toontalker: Cross-domain face reenactment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7690–7700, 2023. 3
- [20] Jianzhu Guo, Dingyun Zhang, Xiaoqiang Liu, Zhizhou Zhong, Yuan Zhang, Pengfei Wan, and Di Zhang. Liveportrait: Efficient portrait animation with stitching and retargeting control. *arXiv preprint arXiv:2407.03168*, 2024. 2, 6, 7
- [21] Thorsten Hempel, Ahmed A. Abdelrahman, and Ayoub Al-Hamadi. Toward robust and unconstrained full range of rotation head pose estimation. *IEEE Transactions on Image Processing*, 33:2377–2387, 2024. 8
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 8
- [23] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 5
- [24] Fa-Ting Hong, Longhao Zhang, Li Shen, and Dan Xu. Depth-aware generative adversarial network for talking head video generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3397–3406, 2022. 3

- [25] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20374–20384, 2022. 3
- [26] 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. 5, 8
- [27] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 2
- [28] Gwanghyun Kim and Se Young Chun. Datid-3d: Diversity-preserved domain adaptation using text-to-image diffusion for 3d generative model, 2022. 3, 5
- [29] Gwanghyun Kim, Ji Ha Jang, and Se Young Chun. Podia-3d: Domain adaptation of 3d generative model across large domain gap using pose-preserved text-to-image diffusion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 22603–22612, 2023. 3
- [30] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 5
- [31] Biwen Lei, Kai Yu, Mengyang Feng, Miaomiao Cui, and Xuansong Xie. Diffusiongan3d: Boosting text-guided 3d generation and domain adaptation by combining 3d gans and diffusion priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10487–10497, 2024. 3
- [32] 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
- [33] Weichuang Li, Longhao Zhang, Dong Wang, Bin Zhao, Zhigang Wang, Mulin Chen, Bang Zhang, Zhongjian Wang, Liefeng Bo, and Xuelong Li. One-shot high-fidelity talking-head synthesis with deformable neural radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17969–17978, 2023. 2, 3
- [34] Xueting Li, Shalini De Mello, Sifei Liu, Koki Nagano, Umar Iqbal, and Jan Kautz. Generalizable one-shot 3d neural head avatar. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [35] 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. 8
- [36] Hongyu Liu, Xintong Han, Chengbin Jin, Lihui Qian, Huawei Wei, Zhe Lin, Faqiang Wang, Haoye Dong, Yibing Song, Jia Xu, et al. Human motionformer: Transferring human motions with vision transformers. *arXiv preprint arXiv:2302.11306*, 2023. 3
- [37] Hongyu Liu, Xuan Wang, Ziyu Wan, Yujun Shen, Yibing Song, Jing Liao, and Qifeng Chen. Headartist: Text-conditioned 3d head generation with self score distillation. In *ACM SIGGRAPH 2024 Conference Papers*, New York, NY, USA, 2024. Association for Computing Machinery. 3
- [38] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2019. 6
- [39] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022. 6
- [40] 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. 8
- [41] Yue Ma, Yingqing He, Xiaodong Cun, Xintao Wang, Siran Chen, Xiu Li, and Qifeng Chen. Follow your pose: Pose-guided text-to-video generation using pose-free videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4117–4125, 2024. 3
- [42] Yue Ma, Hongyu Liu, Hongfa Wang, Heng Pan, Yingqing He, Junkun Yuan, Ailing Zeng, Chengfei Cai, Heung-Yeung Shum, Wei Liu, et al. Follow-your-emoji: Fine-controllable and expressive freestyle portrait animation. *arXiv preprint arXiv:2406.01900*, 2024. 2, 3
- [43] Zhiyuan Ma, Xiangyu Zhu, Guo-Jun Qi, Zhen Lei, and Lei Zhang. Otavatar: One-shot talking face avatar with controllable tri-plane rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16901–16910, 2023. 2, 3
- [44] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 2, 4
- [45] 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. *Communications of the ACM*, 65(1):99–106, 2021. 2, 3
- [46] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021. 5
- [47] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 5
- [48] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2085–2094, 2021. 4
- [49] William Peebles and Saining Xie. Scalable diffusion models with transformers. *arXiv preprint arXiv:2212.09748*, 2022. 2, 5, 6
- [50] Juan C Pérez, Thu Nguyen-Phuoc, Chen Cao, Artiom Sanakoyeu, Tomas Simon, Pablo Arbeláez, Bernard Ghanem, Ali Thabet, and Albert Pumarola. Styleavatar: Stylizing animatable head avatars. In *Proceedings of the*

- IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 8678–8687, 2024. 3
- [51] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 5
- [52] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024. 5
- [53] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 5, 8
- [54] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [55] J. Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Ankner, Jiajun Wu, and Gordon Wetzstein. 3d neural field generation using triplane diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20875–20886, 2023. 5
- [56] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2377–2386, 2019. 2, 3
- [57] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in neural information processing systems*, 32, 2019. 2, 3
- [58] Aliaksandr Siarohin, Oliver J Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. Motion representations for articulated animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13653–13662, 2021. 2, 3
- [59] Luchuan Song, Lele Chen, Celong Liu, Pinxin Liu, and Chenliang Xu. Texttoon: Real-time text toonify head avatar from single video. *arXiv preprint arXiv:2410.07160*, 2024. 3
- [60] 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 *CVPR*, 2023. 2, 4
- [61] Alex Trevithick, Matthew Chan, Michael Stengel, Eric Chan, Chao Liu, Zhiding Yu, Sameh Khamis, Ravi Ramamoorthi, and Koki Nagano. Real-time radiance fields for single-image portrait view synthesis. 2023. 3, 6
- [62] Ziyu Wan, Despoina Paschalidou, Ian Huang, Hongyu Liu, Bokui Shen, Xiaoyu Xiang, Jing Liao, and Leonidas Guibas. Cad: photorealistic 3d generation via adversarial distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10194–10207, 2024. 3
- [63] Duomin Wang, Yu Deng, Zixin Yin, Heung-Yeung Shum, and Baoyuan Wang. Progressive disentangled representation learning for fine-grained controllable talking head synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17979–17989, 2023. 3, 6
- [64] 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, 4, 6
- [65] 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. 2, 4, 5
- [66] 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. 3
- [67] Huawei Wei, Zejun Yang, and Zhisheng Wang. Aniportrait: Audio-driven synthesis of photorealistic portrait animations. *arXiv:2403.17694*, 2024. 2, 3
- [68] Shuang Wu, Youtian Lin, Feihu Zhang, Yifei Zeng, Jingxi Xu, Philip Torr, Xun Cao, and Yao Yao. Direct3d: Scalable image-to-3d generation via 3d latent diffusion transformer. *arXiv preprint arXiv:2405.14832*, 2024. 5, 6
- [69] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Neural Information Processing Systems (NeurIPS)*, 2021. 3, 6
- [70] Liangbin Xie, Xintao Wang, Honglun Zhang, Chao Dong, and Ying Shan. Vfhq: A high-quality dataset and benchmark for video face super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2022. 5, 7
- [71] You Xie, Hongyi Xu, Guoxian Song, Chao Wang, Yichun Shi, and Linjie Luo. X-portrait: Expressive portrait animation with hierarchical motion attention. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 2, 3, 6, 7
- [72] Sicheng Xu, Guojun Chen, Yu-Xiao Guo, Jiaolong Yang, Chong Li, Zhenyu Zang, Yizhong Zhang, Xin Tong, and Baining Guo. VASA-1: Lifelike audio-driven talking faces generated in real time. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 3
- [73] Zhenhui Ye, Tianyun Zhong, Yi Ren, Jiaqi Yang, Weichuang Li, Jiawei Huang, Ziyue Jiang, Jinzheng He, Rongjie Huang, Jinglin Liu, et al. Real3d-portrait: One-shot realistic 3d talking portrait synthesis. *arXiv preprint arXiv:2401.08503*, 2024. 3
- [74] Fei Yin, Yong Zhang, Xiaodong Cun, Mingdeng Cao, Yanbo Fan, Xuan Wang, Qingyan Bai, Baoyuan Wu, Jue Wang, and Yujiu Yang. Styleheat: One-shot high-resolution editable talking face generation via pre-trained stylegan. In *European*

- conference on computer vision, pages 85–101. Springer, 2022. 3
- [75] Wangbo Yu, Yanbo Fan, Yong Zhang, Xuan Wang, Fei Yin, Yunpeng Bai, Yan-Pei Cao, Ying Shan, Yang Wu, Zhongqian Sun, et al. Nofa: Nerf-based one-shot facial avatar reconstruction. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–12, 2023. 3
- [76] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9459–9468, 2019. 2, 3
- [77] Bowen Zhang, Chenyang Qi, Pan Zhang, Bo Zhang, Hsiang-Tao Wu, Dong Chen, Qifeng Chen, Yong Wang, and Fang Wen. Metaportrait: Identity-preserving talking head generation with fast personalized adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22096–22105, 2023. 2, 3
- [78] Bowen Zhang, Yiji Cheng, Chunyu Wang, Ting Zhang, Jiaolong Yang, Yansong Tang, Feng Zhao, Dong Chen, and Baining Guo. Rodinhd: High-fidelity 3d avatar generation with diffusion models. *arXiv preprint arXiv:2407.06938*, 2024. 2, 4
- [79] 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, 2023. 4
- [80] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 4, 5
- [81] 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. 8
- [82] Jian Zhao and Hui Zhang. Thin-plate spline motion model for image animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3657–3666, 2022. 2
- [83] Xiaochen Zhao, Jingxiang Sun, Lizhen Wang, Jinli Suo, and Yebin Liu. Invertavatar: Incremental gan inversion for generalized head avatars. In *ACM SIGGRAPH 2024 Conference Papers*, New York, NY, USA, 2024. Association for Computing Machinery. 6
- [84] Xiaochen Zhao, Jingxiang Sun, Lizhen Wang, Jinli Suo, and Yebin Liu. Invertavatar: Incremental gan inversion for generalized head avatars. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–10, 2024. 2, 4, 6, 7
- [85] Yiming Zhu, Hongyu Liu, Yibing Song, Ziyang Yuan, Xintong Han, Chun Yuan, Qifeng Chen, and Jue Wang. One model to edit them all: Free-form text-driven image manipulation with semantic modulations. *Advances in Neural Information Processing Systems*, 35:25146–25159, 2022. 4
- [86] Yiyu Zhuang, Hao Zhu, Xusen Sun, and Xun Cao. Mofanerf: Morphable facial neural radiance field. In *European conference on computer vision*, pages 268–285. Springer, 2022. 3