

RodinHD: High-Fidelity 3D Avatar Generation with Diffusion Models

Bowen Zhang^{1*}, Yiji Cheng^{2*}, Chunyu Wang^{3†}, Ting Zhang³, Jiaolong Yang³, Yansong Tang², Feng Zhao¹, Dong Chen³, and Baining Guo³

¹ University of Science and Technology of China

² Tsinghua University

³ Microsoft Research Asia

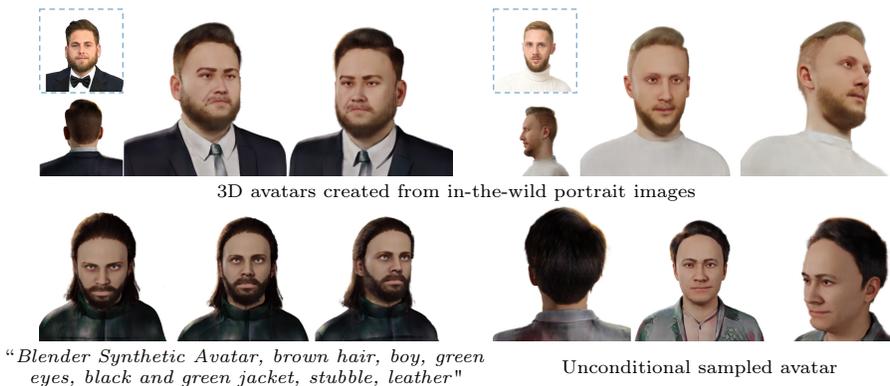


Fig. 1: RodinHD generates detailed 3D avatars from single portrait images (dashed boxes) without compromising cross-view consistency (first row). It also supports text-conditioned (second row left) or unconditional (second row right) generation.

Abstract. We present RodinHD, which can generate high-fidelity 3D avatars from a portrait image. Existing methods fail to capture intricate details such as hairstyles which we tackle in this paper. We first identify an overlooked problem of catastrophic forgetting that arises when fitting triplanes sequentially on many avatars, caused by the MLP decoder sharing scheme. To overcome this issue, we raise a novel data scheduling strategy and a weight consolidation regularization term, which improves the decoder’s capability of rendering sharper details. Additionally, we optimize the guiding effect of the portrait image by computing a finer-grained hierarchical representation that captures rich 2D texture cues, and injecting them to the 3D diffusion model at multiple layers via cross-attention. When trained on 46K avatars with a noise schedule optimized for triplanes, the resulting model can generate 3D avatars with notably better details than previous methods and can generalize to in-the-wild portrait input. See Fig. 1 for some examples. Project page: <https://rodinhd.github.io/>.

Keywords: 3D avatar generation · Diffusion · Catastrophic forgetting

* Interns at Microsoft Research Asia. Equal contribution. †Corresponding author.

1 Introduction

High-fidelity 3D avatar generation has many applications in fields such as gaming and metaverse. Recent development of generative diffusion models [19, 21, 24, 46, 49] and implicit neural radiance fields [39, 65] has opened up new opportunities for automatic generation of 3D avatar [20, 40, 52, 65] at scale. However, current methods struggle to generate fine details, which is a core challenge that has to be addressed. Otherwise, a “toy-like” avatar is less effective in delivering practical values.

The work of 3D generative diffusion models [2, 20, 28, 40, 44, 52, 65, 72, 73] usually follow a two-stage framework. First, they compute a proxy 3D representation of fixed length such as triplanes or volumes from the original unstructured meshes or point clouds so that they can be handled by diffusion models. A paired decoder is jointly learned to render 360° images from the representation. Different from the dominant NeRF fitting methods, the decoder here is shared among all avatars to decode novel generated triplanes. Then, they train diffusion models on the proxy representation to generate diverse avatars. However, they struggle to generate fine details such as sharp cloth textures and hair strands. To alleviate the problem, Rodin [65] applies a convolution refiner [66] to complement the missing details for each rendered image. Although the 2D refiner improves the visual quality in one view, it significantly compromises 3D consistency, which is not tolerable in many applications.

In this work, we introduce RodinHD, which aims to improving the fidelity of avatars without any refiners. We begin by empirically showing that fitting triplanes sequentially on a large number of avatars suffers from catastrophic forgetting, which can result in under-fitted decoders incapable of generating intricate details on novel triplanes. This occurs because the triplane of an avatar is typically trained for many iterations before switching to the next one, in order to reduce the data transfer costs between CPUs and GPUs. However, this process gradually causes the shared decoder to forget knowledge learned from previous avatars, leading to a lack of generalizability. Fig. 2 illustrates this issue with typical renderings of the resulting decoder. This problem has been largely overlooked in the literature, hindering the development of high-fidelity generation based on neural radiance fields.

To address this issue, we propose a novel data scheduling strategy called task replay, and a weight consolidation regularization term that effectively preserves the decoder’s capability of rendering sharp details. The idea of task-replay is to switch avatars more frequently so that each can be seen periodically for multiple times, preventing the decoder from over-fitting to a single avatar. The weight consolidation regularization term prevents the critical weights from deviating far from its consolidated values. As a result, knowledge learned from previous data can be retained during training. The method effectively alleviates the forgetting problem and improve the model’s capability to encode intricate details, paving the way for the subsequent generation step.

We train a cascaded diffusion model on the triplanes for conditional generation. It consists of a *base model*, which generates a low-resolution triplane



Fig. 2: Catastrophic forgetting. As training proceeds, decoder gradually forgets the knowledge learned on the previous avatars of 1&4 and is overly adapted to avatar 9.

conditioned on a portrait image, and an *upsample model*, which subsequently generates a high-resolution triplane. Differing from the previous work [65] which uses CLIP [48] to compute a global semantic token for the portrait image as conditions, we maximize its guiding effect by computing a finer-grained hierarchical representation to provide more detailed cues for the 3D diffusion model using a VAE-based image encoder. The multi-scale features are injected into different layers of U-Net via cross-attention, which significantly improves the coherence between the generated avatars and the portrait images. Fig. 3 shows the process. Besides, inspired by [8], we also optimize the noise schedule for the triplane considering its high redundancy in both spatial and channel dimensions.

We train the model on 46K digital avatars [67] of diverse identities, expressions, hairstyles, and clothing. We render high-quality images at the resolution of 1024×1024 . The resulting model is capable of generating highly detailed avatars with clear clothing textures and hairstyles using a simple diffusion model without extra refinement models. While only validated on avatars, the proposed techniques are general and can be applied to other 3D generation tasks.

2 Related Work

Early works in 3D generation have primarily focused on generating coarse 3D shapes, which are typically represented as meshes [34, 57], point clouds [1, 32], voxel grids [4, 68], and implicit neural representations [45, 54], using either GANs [16, 29, 74] or VAEs [30]. However, it remains unclear whether these methods can be effectively applied to generate complex 3D avatars with rich details.

3D-aware GANs [5, 6, 10, 12, 15, 43, 51, 60, 70, 71] are able to generate high-resolution images with the aid of 2D upsampler and patch-based image discriminator. Nevertheless, they suffer from cross-view consistency [69]. Gram [70] performs upsampling on the surface manifold to promote multiview consistency and efficiency but it cannot handle large viewpoint changes and complex geometry. Moreover, they are prone to mode collapse due to training instabilities of GANs. Although score-distillation-based methods [9, 35, 47, 56, 58, 59] are proposed to distill the 2D diffusion prior to a 3D representation with score function, they suffer from the Janus problem which prevents them from generating accurate geometry because the problem is under-determined.

Recently, diffusion models [13, 23, 55] have achieved notable success in text-to-image synthesis, with a few on high-resolution image generation [14, 18, 38, 46, 49, 50, 62]. Inspired by this, many recent works apply diffusion models for 3D generation [2, 7, 11, 20, 28, 37, 40, 44, 52, 61, 65, 72]. They first compute a proxy 3D representation from the raw data, and then train a diffusion model on the proxy data with either text or image as conditions. However, most of them struggle to generate fine details such as cloth textures and hair strands, due to the lack of high-quality proxy data and a well-designed upsample network. Rodin [65] exploits a convolution refiner [66] to complement the missing details in each individually rendered image, but this compromises cross-view consistency.

Some recent work [14, 18, 38, 46, 49, 50, 62] have investigated high-resolution image generation, and obtained some interesting findings about optimal noise schedules, model architectures and capacity scaling principles. Despite this, we find the triplanes are essentially different from images, and directly transferring their findings barely works in our scenario. For instance, Stable Diffusion [49] uses a VAE to compress an image into a lower-resolution latent for diffusion. However, we observe that compressing triplanes by training a VAE will lose many high-frequency details in the renderings. Moreover, the considerations for designing the noise schedules are also different from those in the images.

3 Method

Our framework comprises two primary steps: fitting and modeling. In the first step, we fit a high-resolution triplane $\mathbf{x}_0 \in \mathbb{R}^{3 \times H_{\mathbf{x}} \times W_{\mathbf{x}} \times C}$ for each avatar, and learn a decoder \mathcal{F} , which is shared by all avatars, to render high-fidelity images from the triplanes. The parameters of $H_{\mathbf{x}}$ and $W_{\mathbf{x}}$ denote the height and width of triplanes, taking value on 512, and C is the number of channels, taking value on 32. In the second step, we train a 3D diffusion model \mathcal{G} that can generate triplanes from random noises $\epsilon \in \mathbb{R}^{3 \times H_{\mathbf{x}} \times W_{\mathbf{x}} \times C}$ conditioned on a portrait image $\mathbf{I}_{\text{front}} \in \mathbb{R}^{H_{\mathbf{I}} \times W_{\mathbf{I}} \times 3}$. Fig. 3 shows the overall framework. During inference, we can render 360° high-resolution images $\{\hat{\mathbf{I}} \in \mathbb{R}^{H_{\mathbf{I}} \times W_{\mathbf{I}} \times 3}\}$ from the following process:

$$\mathcal{G} : (\epsilon, \mathbf{I}_{\text{front}}) \xrightarrow{\text{Denosing}} \mathbf{x}_0, \quad (1)$$

$$\mathcal{F} : \mathbf{x}_0 \xrightarrow{\text{Rendering}} \{\hat{\mathbf{I}}\}. \quad (2)$$

The primary challenge arises from the high-resolution nature of the triplane, presenting difficulties in both fitting and modeling. We will elaborate on several key design considerations in the following.

3.1 Triplane Fitting

The increased resolution of the triplanes introduces a higher computational load, leading to potential bottlenecks in terms of both time and memory. Therefore we split this task into two stages to reduce computation cost. In the first

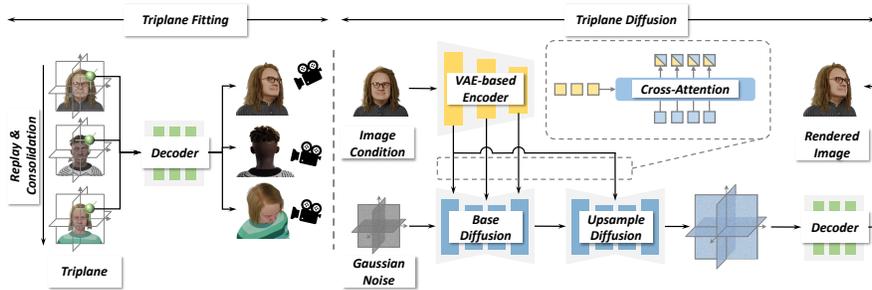


Fig. 3: Overview of our method.

stage, we jointly train the MLP decoder and the triplanes on a smaller subset of avatars. In the second stage, we fix the decoder’s weights and fine-tune the triplane of each avatar independently, which can be performed in parallel. The first stage of getting an accurate yet generalizable MLP decoder is critical in order to generate details for all avatars including those generated ones.

The decoder \mathcal{F} parameterized by ω learns identity-agnostic priors from a set of S avatars. Each avatar is represented as N multi-view renderings $\mathcal{D}^S = \{(\mathbf{I}^{s,n}, c^{s,n})\}_{n=1}^N$, where $\mathbf{I}^{s,n}$ is the RGBA image and $c^{s,n}$ is the corresponding camera configuration. There are two critical issues overlooked in previous work [65]. 1) During training, batches typically comprise rays sampled from a single avatar due to limited GPU memory. Consequently, each avatar is treated as an independent task. To minimize data transfer between CPU and GPU, each avatar undergoes multiple training iterations until convergence before being replaced by the next one. As a result, the MLP may gradually forget the previously learned knowledge and become overly adapted to the current avatar. This is known as catastrophic forgetting in continual learning. Fig. 2 visualizes the phenomenon. 2) Additionally, the approach encounters training instabilities when switching between avatars due to the substantial gaps between them. This usually results in an under-fitted MLP that is incapable of decoding high-frequency details, even with the second triplane finetuning stage. Fig. 10 compares the differences. In the following, we formally describe the task-replay strategy and weight consolidation regularizer to address the above challenges.

Task replay. The core idea is to switch avatars more frequently, allowing each avatar to be seen periodically for multiple times and preventing the decoder from over-fitting to a single avatar. As a result, the decoder is exposed to triplanes of a variety of avatars, ensuring its generalization ability. To implement this strategy, each avatar is trained multiple times, with each time fitting for a shorter time without requiring convergence by tuning “outer_loop_iteration” and “inner_loop_iteration” in Algorithm 1, while the naive method trains each avatar only once, and fits it to convergence before switching to the next avatar. The task replay strategy proves effective in mitigating the risk of over-fitting to a single avatar, which may otherwise result in poor generalization performance.

Algorithm 1 The First Stage of Triplane Fitting

Require: dataset $\{\mathcal{D}^s\}_{s=1}^S$, triplane parameters $\{\mathbf{x}^s\}_{s=1}^S$, shared MLP \mathcal{F} parameters ω , IWC states $\{\Omega^s = \mathbf{0}, \omega^{s,*}\}_{s=1}^S$, learning rate α, β .

- 1: **repeat**
- 2: Sample from \mathcal{D}^s and load triplane \mathbf{x}^s from disk
- 3: **repeat**
- 4: Sample rays $r \sim \mathcal{D}^s$
- 5: $\mathcal{L}(\mathbf{x}^s, \omega; \Omega^s, \omega^{s,*}) \leftarrow \mathcal{L}_{\mathcal{D}^s} + \mathcal{L}_{\text{IWC}}$
- 6: $\mathbf{x}^s \leftarrow \mathbf{x}^s - \alpha \nabla_{\mathbf{x}^s} \mathcal{L}(\mathbf{x}^s, \omega; \Omega^s, \omega^{s,*})$
- 7: $\omega \leftarrow \omega - \beta \nabla_{\omega} \mathcal{L}(\mathbf{x}^s, \omega; \Omega^s, \omega^{s,*})$
- 8: **until** inner_loop_iteration
- 9: $\{\Omega^s, \omega^{s,*}\} \leftarrow \{(\nabla_{\omega} \mathcal{L}_{\mathcal{D}^s})^2, \omega\}$ ▷ update state
- 10: **until** outer_loop_iteration

Weight consolidation. Learning the triplanes and the decoder also suffers from instabilities caused by the occasionally occurring large gradients when we switch between the avatars. To mitigate this issue, we introduce an Identity-aware Weight Consolidation (IWC) regularizer, a technique that stabilizes learning by consolidating knowledge and reducing drastic shifts in the learning landscape. Specifically, it uses the elastic weight consolidation regularizer [31] to prevent the most important weights of this avatar from deviating far from its consolidated values during training:

$$\begin{aligned} \mathcal{L}_{\text{Fitting}} &= \mathcal{L}_{\mathcal{D}^s} + \mathcal{D}_{\text{IWC}} \\ &= \mathcal{L}_{\mathcal{D}^s} + \frac{\lambda}{2} \sum_i \Omega_i^s (\omega_i - \omega_i^{s,*})^2, \end{aligned} \quad (3)$$

where $\mathcal{L}_{\mathcal{D}^s}$ is the conventional rendering loss between rendered images and ground-truth images, $\omega_i^{s,*}$ is the i^{th} MLP weight obtained when training the avatar in the previous epoch, and Ω_i^s is the importance of i^{th} weight, calculated by the i^{th} diagonal element of the Fisher Information Matrix Ω^s . Since the Fisher Information Matrix is equivalent to the second derivatives of the loss near a minimum and it can be computed from first-order derivatives, computing the IWC loss is efficient.

We find that our approach not only improves the details but also reduces the high-frequency components in the triplanes, making them easier to learn by diffusion models [33, 38]. See Fig. 4 for comparison. We also use TV loss and L_2 regularization on the triplanes to promote smoothness as in the prior work [65].

3.2 Triplane Diffusion

We train a cascaded diffusion model to generate high-resolution triplanes. It consists of a *base model*, which generates a low-resolution triplane conditioned on a portrait image $\mathbf{I}_{\text{front}}$, and an *upsample model*, which subsequently generates a high-resolution triplane. During training, we first obtain the destructed triplane

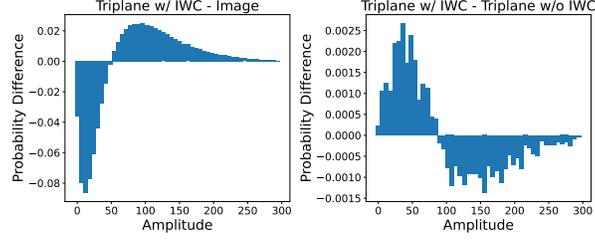


Fig. 4: Frequency difference between two sources. **Left:** Triplanes have more high-frequency components than images. **Right:** Triplanes learned with our proposed IWC have fewer high-frequency components.

\mathbf{x}_t according to $\mathbf{x}_t := \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}$, where α_t, σ_t define the noise schedule which determines the destruction strength of the signal, and t is a continuous number ranged from 0 to 1. We train the base and upsample diffusion models in two separate steps. For base diffusion training, we parameterize the diffusion model $\hat{\boldsymbol{\epsilon}}_\theta$ to predict the noise [23] added to the low-resolution triplane \mathbf{x}_t^{LR} :

$$\mathcal{L}_{\text{simple}}^{\text{LR}} = \mathbb{E}_{t, \mathbf{x}_0^{\text{LR}}, \boldsymbol{\epsilon}} \left[\left\| \hat{\boldsymbol{\epsilon}}_\theta \left(\alpha_t \mathbf{x}_0^{\text{LR}} + \sigma_t \boldsymbol{\epsilon}, t, \mathbf{I}_{\text{front}} \right) - \boldsymbol{\epsilon} \right\|_2^2 \right]. \quad (4)$$

Despite the $\mathcal{L}_{\text{simple}}^{\text{LR}}$, we also use the variational lower bound to optimize the negative log-likelihood of estimated distribution by incorporating $\mathcal{L}_{\text{vib}}^{\text{LR}}$ following [42] for higher generation quality. Our upsample diffusion model is learned to enhance the high-fidelity details of the low-resolution triplane, which is conditioned on both \mathbf{x}_0^{LR} and $\mathbf{I}_{\text{front}}$. We directly parameterize the model to predict the noiseless input \mathbf{x}_0^{HR} :

$$\mathcal{L}_{\text{simple}}^{\text{HR}} = \mathbb{E}_{t, \mathbf{x}_0^{\text{HR}}, \boldsymbol{\epsilon}} \left[\left\| \hat{\mathbf{x}}_\theta \left(\alpha_t \mathbf{x}_0^{\text{HR}} + \sigma_t \boldsymbol{\epsilon}, t, \mathbf{I}_{\text{front}}, \mathbf{x}_0^{\text{LR}} \right) - \mathbf{x}_0^{\text{HR}} \right\|_2^2 \right]. \quad (5)$$

To ensure the rendered images of our generated triplanes have compelling visual quality, we also adopt image-level supervision when training the upsample model inspired by previous works [65]. Specifically, we penalize the discrepancy between the rendered image patch $\hat{\mathbf{I}}_{\text{patch}}$ from predicted triplane $\hat{\mathbf{x}}_\theta$ and the ground truth image patch $\mathbf{I}_{\text{patch}}$:

$$\begin{aligned} \mathcal{L}_{\text{Image}}^{\text{HR}} &= \mathcal{L}_{\text{Pixel}} + \mathcal{L}_{\text{Perc}} \\ &= \mathbb{E}_{t, \hat{\mathbf{I}}_{\text{patch}}} \left(\sum_l \left\| \boldsymbol{\Psi}^l(\hat{\mathbf{I}}_{\text{patch}}) - \boldsymbol{\Psi}^l(\mathbf{I}_{\text{patch}}) \right\|_2^2 \right) \\ &\quad + \mathbb{E}_{t, \hat{\mathbf{I}}_{\text{patch}}} \left(\left\| \hat{\mathbf{I}}_{\text{patch}} - \mathbf{I}_{\text{patch}} \right\|_2^2 \right), \end{aligned} \quad (6)$$

where $\boldsymbol{\Psi}^l$ denotes the multi-scale feature extracted using a pre-trained VGG [53]. Ensuring the diffusion model scales well with the increased resolution is crucial. As high-resolution triplanes involve a large number of parameters, processing

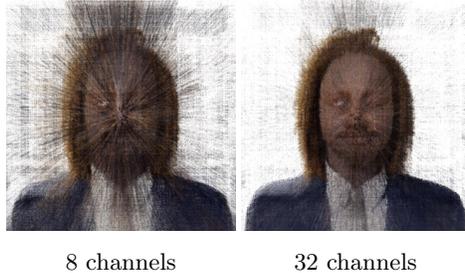


Fig. 5: Rendered images from triplanes with 8 and 32 channels, respectively. The triplanes are destructed with the same noise level ($\log\text{SNR}(t) = 0.57$). The 32-channel triplane has larger redundancy so it is less destructed.

such vast amounts of information efficiently can be a substantial hurdle. Next, we introduce improvements to our diffusion model.

Multi-scale image feature conditioning. It is difficult to hallucinate detailed 3D avatars from scratch. So, we propose to supply ample details from the portrait image to alleviate the difficulty. Previous works such as [65] compute a global semantic token using the pre-trained CLIP image encoder [48], which results in a substantial loss of detailed information.

To fully utilize the information of the portrait, we compute a multi-scale feature representation using a pre-trained Variational Autoencoder [46]. Since the VAE is trained to accurately reconstruct the input images, the low-level visual details are well preserved in the latent features. Formally, we denote the VAE encoder as E , and the frontal portrait image $\mathbf{I}_{\text{front}} \in \mathbb{R}^{H \times W \times 3}$. We compute the conditional signals for our diffusion model by:

$$\{\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3\} = E(\mathbf{I}_{\text{front}}), \quad (7)$$

where $\mathbf{y}_1 \in \mathbb{R}^{\frac{H}{2} \times \frac{W}{2} \times C}$, $\mathbf{y}_2 \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 2C}$, and $\mathbf{y}_3 \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times 4C}$ denote the multi-scale spatial features, and C is the based channel dimension.

Since the 2D portrait is not aligned with the 3D triplane, directly harmonizing the conditional signals with diffusion U-Net features via concatenation [24, 41] or addition [75] is problematic. Therefore, we elect to perform cross-attention [63]. In this way, the network is learned to automatically discover the spatial correspondence between triplanes and 2D images.

Formally, let \mathbf{x}_i be the feature maps of the i -th attention resolution in U-Net, we inject \mathbf{y}_i by

$$\begin{aligned} \{\mathbf{y}_i^k\}_{k=1}^K &= \text{PatchPartition}(\mathbf{y}_i), \\ \mathbf{x}'_i &= \text{CrossAttn}(\mathbf{x}_i; \{\mathbf{y}_i^k\}_{k=1}^K), \end{aligned} \quad (8)$$

where K and \mathbf{x}'_i are the number of patches and the output features, respectively. For the base model, the conditions are injected into both encoder and decoder layers that have the same resolutions. For the upsample model, the conditions are only injected into the middle latent features to reduce computation costs.

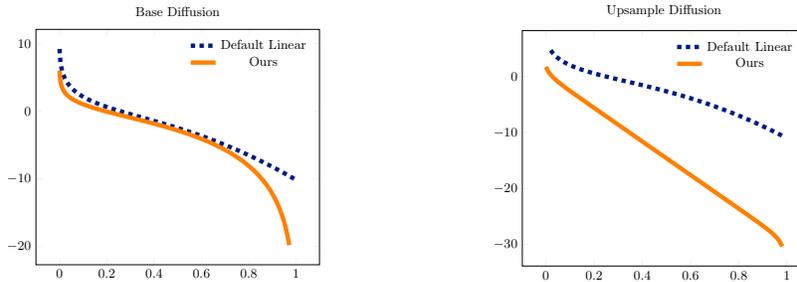


Fig. 6: Optimized noise schedule. LogSNR comparison between the default and our optimized noise schedules for the base (left) and upsample (right) networks.

Optimized noise schedule. In high-resolution image generation, some work [8, 17, 25, 36] find that using a stronger noise is critical to retain the learning difficulty for images with higher resolutions. We similarly study the optimal noise schedule for triplanes in 3D generation. Compared to images, triplanes have large spatial resolutions and channel numbers, and even long-range dependencies introduced by 3D correspondence. All these factors increase the redundancy in the representation. Fig. 5 shows an example. We fit two triplanes with the same spatial resolution but different channel dimensions (8 vs. 32), and add the same level of noises ($\text{SNR}(t) = \alpha_t^2 / \sigma_t^2$) to them, respectively. However, we can see that the rendered images are destructed differently. The triplane with fewer channels suffers from more disruption compared with the one with larger channels.

Considering the larger redundancy in triplanes, we propose to apply a stronger noise to fully destruct triplanes to prevent the model from under-training. To be more specific, we utilize the adjusted cosine noise schedule [8, 42] in our base diffusion and a much stronger sigmoid noise schedule [8, 27] in the upsample diffusion stage. The adjusted noise schedules are shown in Fig. 6. They are much stronger than the default linear noise scheduling [23] designed for low-resolution images *e.g.* $32 \times 32, 64 \times 64$.

4 Experiments

4.1 Dataset and Metrics

We conduct experiments on 46K avatars created from Blender [67]. For each avatar, we uniformly render 300 multi-view images at the resolution of 1024×1024 . We evaluate our model’s conditional and unconditional generation capability, respectively. For image-conditioned generation, we compute numerical results using the common metrics of FID [22], LPIPS [76], and Structural Similarity Index Measure (SSIM) between 5K rendered multi-view images from generated avatars with ground-truth images. For unconditional generation, we report FID of 5K rendered images from randomly sampled avatars. We also measure cross-view consistency by fitting a NeuS model [64] from the multi-view renderings of the generated avatars.

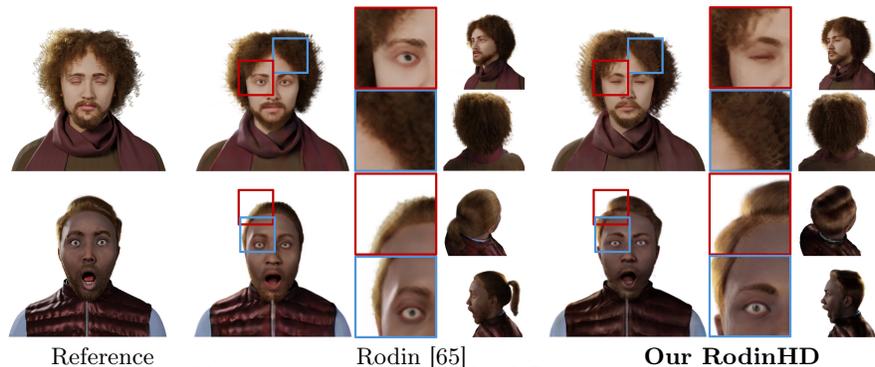


Fig. 7: Avatars generated by different methods.

Table 1: Quantitative results of conditional avatar generation.

Models	FID↓	LPIPS↓	SSIM↑
Rodin [65]	33.20	0.323	0.758
Ours	26.49	0.299	0.765

4.2 Implementation Details

The triplane resolution is 512×512 and its channel number is 32. We randomly select 64 avatars for jointly training the triplanes and the NeRF decoder for two days on 8 Tesla V100 GPUs, with task replay and IWC regularization applied. Then, we fit the triplanes for all avatars independently, with each taking about 15 minutes. We adopt the U-Net architecture [13] as diffusion backbone. For the configuration of conditional feature injection, please refer to supplementary material. We utilize conditional augmentation [24, 65] when training the upsample diffusion to reduce the domain gap between training and inference. Both our base and upsample models are trained on 32 Tesla V100 (32G) GPUs, with batch sizes 96 and 32 respectively. We randomly drop the conditional features with 20% probability, which enables us to perform classifier-free guidance and unconditional generation.

4.3 Main Results

Conditional generation. We compare our approach with Rodin [65]. We report the results of Rodin without 2D refinement for fair comparison. As shown in Fig. 7, our model captures the detailed appearance and vivid expression of the given portrait, benefiting from high-quality triplane fitting and strong guidance of the proposed multi-scale image feature conditioning. On the contrary, Rodin fails to synthesize details of avatars, *e.g.*, hairstyles and closed eyes. Moreover, our model enables us to directly render compelling images in high resolution without any 2D refinement, demonstrating the strong capacity of the proposed

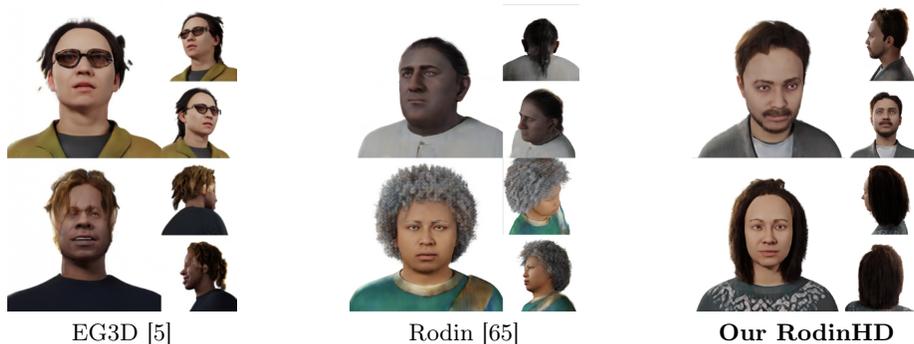


Fig. 8: Qualitative results of unconditional avatar generation.

Table 2: Quantitative results of unconditional avatar generation. The subscript * indicates that 2D refinement is applied to the rendered images.

	Pi-GAN	GIRAFFE	EG3D*	Rodin*	Rodin	Ours
FID ↓	78.3	64.6	40.5	30.29	45.70	32.62

model. Furthermore, we achieve the best results across all metrics as shown in Tab. 1, demonstrating the superiority of our model.

Unconditional generation. We compare our method with the state-of-the-art methods, including both 3D-aware GANs [5, 6, 43] and 3D-diffusion-based Rodin [65]. We present the results of Rodin before and after the 2D refiner (denoted as Rodin*), respectively. Tab. 2 shows the results. Our approach achieves significantly better results than the methods except Rodin* which uses a 2D refiner. However, as we will discuss in the following, our method achieves notably better 3D consistency. We also provide visual comparison in Fig. 8. While prior arts tend to generate blur rendering, our model is able to provide complex and diverse avatars with rich details, *e.g.*, hair and clothing.

3D consistency. We evaluate the 3D consistency of the generated results by visualizing the spatiotemporal textures following [70]. We include Rodin*, ours and ground-truth triplane renderings for comparison in Fig. 9. As the camera moving smoothly, the texture of the fixed horizontal line is expected to be smooth and nature as GT triplane renderings. However, the spatiotemporal textures of Rodin* suffer from obvious flickering, which suggests inconsistency across views (see video in supplementary for more intuitive comparison). In contrast, our model produces smooth and natural results as GT, demonstrating the strong consistency of the proposed model. Moreover, the renderings of Rodin* fail to maintain the skin tone of the input while our model provides more faithful results. We also quantitatively measure the 3D consistency by training a multi-view reconstruction method NeuS [64] on 300 rendering images of Rodin* and ours respectively. The results of our model achieve significantly better metrics

Table 3: 3D consistency measured by the fitting quality of NeuS [64].

Models	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
Rodin* [65]	31.73	0.051	0.973
Ours	35.46	0.041	0.975

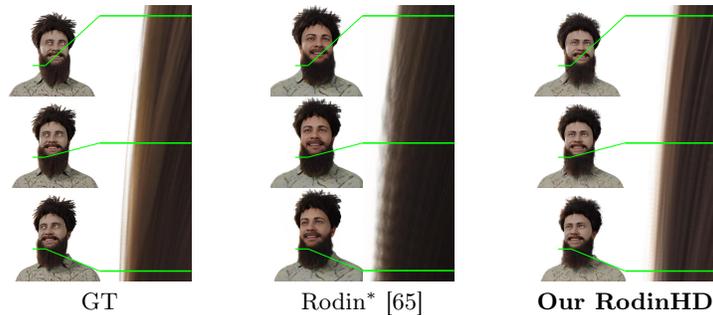


Fig. 9: Visual comparison of 3D consistency akin to the Epipolar Line Images [3]. We visualize the stacked texture of a fixed line as the camera smoothly rotates horizontally. Our model yields smooth and natural texture as GT, whereas a clear flickering pattern is shown in Rodin’s results, indicating the 3D inconsistency using 2D refinement.

thanks to the 3D consistency, whereas Rodin* obtains worse results since the 2D refiner breaks the 3D consistency.

4.4 Ablation Study

We first study the factors that affect triplane fitting. As shown in Tab. 4(a), we present results for five baselines. Rodin (256) is the method proposed in [65]. In Rodin (512), we naively increase the triplane resolution to 512. We can see that directly increasing the triplane resolution provides very marginal gains due to the forgetting problem. With our proposed bunch of techniques, scaling the resolution can notably increase PSNR to 31.54.

Fig. 10 visualizes the differences of the baselines. Rodin (256) cannot obtain a sharp beard due to the low-resolution triplanes. Rodin (512) obtains sharp details but there are many noises caused by the under-fitted decoder. In contrast, our method with task replay and consolidation weight regularization is able to get clean and sharp details. Also, see the areas around the silver tie for differences.

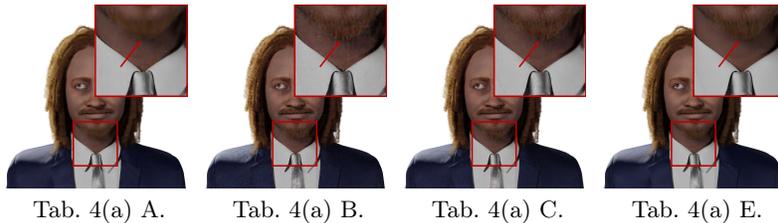
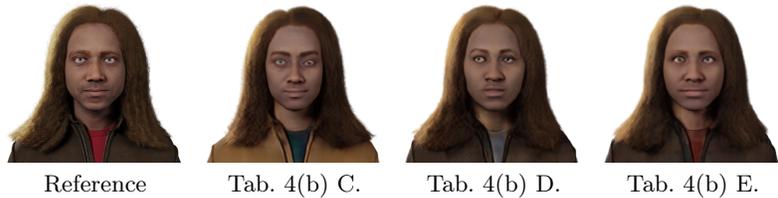
We further evaluate the factors that affect the generation results of our diffusion model in Tab. 4(b). Starting from Rodin [65] which trains diffusion on $256 \times 256 \times 32$ triplanes, we only scale the triplane to $512 \times 512 \times 32$ with other factors unchanged. We can see that the results even become worse. We think this is because the original noises are not suitable for larger triplanes. Using a stronger noise schedule obtains more reasonable results. Replacing the original CLIP image encoder with the VAE encoder (single-scale) also improves

Table 4: Ablation study of the proposed components.

Training strategy	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	Model configuration	FID \downarrow	LPIPS \downarrow	SSIM \uparrow
A. Rodin (256)	30.31	0.131	0.862	A. Rodin (256)	33.20	0.323	0.758
B. Rodin (512)	30.56	0.129	0.863	B. + Naively scale to 512	35.84	0.324	0.740
C. + Task replay	30.93	0.106	0.890	C. + Strong noise schedule	27.74	0.318	0.745
D. + Weight decay	31.00	0.094	0.909	D. + Single-scale cond.	27.13	0.301	0.763
E. + Consolidation	31.45	0.086	0.911	E. + Multi-scale cond.	26.49	0.299	0.765

(a)

(b)

**Fig. 10:** Qualitative ablation for the representation fitting.**Fig. 11:** Qualitative ablation for 3D diffusion model. Multi-scale feature conditioning results in significantly faithful and detailed generation results.

the results. Our method with multi-scale features achieves the best results. We also provide visual comparison as shown in Fig. 11, the introduced multi-scale features of the input portrait provide significantly more detailed texture cues, enabling high-fidelity avatar creation.

4.5 Applications

Avatar creation from in-the-wild portrait. Despite training on the synthetic dataset [67], our model is robust to create 3D avatars conditioned on single in-the-wild portraits. The results in Fig. 12 validate our method’s capability for generalization to real-world images, which outperforms Rodin [65] in retaining both identity and details.

Text-conditioned avatar creation. To enable high-quality avatar creation from text, we first convert the text prompt to reference portrait leveraging strong 2D diffusion network [49], thereafter generate a high-fidelity avatar with the



Fig. 12: Avatars generated conditioned on single in-the-wild portraits.

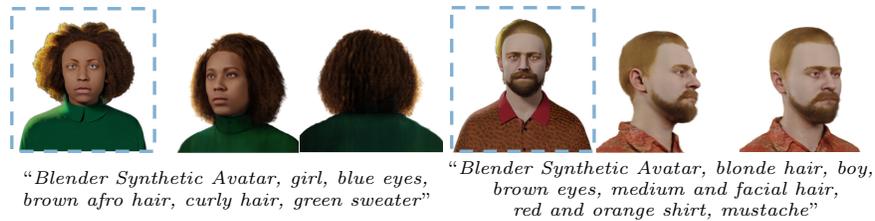


Fig. 13: RodinHD can create high-fidelity avatars based on the reference portraits (dashed boxes) generated by finetuned 2D text-to-image diffusion models.

reference portrait. To be more specific, we use 40 selected frontal images and corresponding text prompts in our synthetic dataset to perform LoRA-based [26] finetuning for Stable Diffusion. We employ the term “Blender Synthetic Avata” as our trigger words within prompts to generate frontally aligned image inputs, which is a widely adopted technique for steering the style of images produced by fine-tuned SD. We provide the samples of text-to-avatar creation in Fig. 13.

5 Conclusion

We present RodinHD for high-fidelity 3D avatar generation by improving both the data fitting and the diffusion modeling. To maintain the compelling visual details of avatars, we propose a task-relay strategy and identity-aware weight consolidation regularizer for high-quality and robust data fitting in large-scale. Furthermore, to model the distribution of highly detailed avatars, we introduce a multi-scale visual feature conditioning mechanism in our cascaded diffusion model, which provides fine-grind guidance to diffusion generation. In addition, we study previous 2D optimized noise scheduling in both high-resolution and high-dimensional 3D diffusion training. Our optimized noise schedules effectively enhance the details of the generated avatars. Extensive experiments show the proposed framework can generate high-fidelity 3D avatars with rich details, which is also promising to apply our model to general 3D scene modeling.

Acknowledgments

This work was supported in part by the Anhui Provincial Natural Science Foundation under Grant 2108085UD12. We acknowledge the support of GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC.

References

1. Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L.: Learning representations and generative models for 3d point clouds. In: International Conference on Machine Learning. pp. 40–49. PMLR (2018)
2. Bautista, M.A., Guo, P., Abnar, S., Talbott, W., Toshev, A., Chen, Z., Dinh, L., Zhai, S., Goh, H., Ulbricht, D., et al.: Gaudi: A neural architect for immersive 3d scene generation. arXiv preprint arXiv:2207.13751 (2022)
3. Bolles, R.C., Baker, H.H., Marimont, D.H.: Epipolar-plane image analysis: An approach to determining structure from motion. *International Journal of Computer Vision* **1**(1), 7–55 (1987)
4. Brock, A., Lim, T., Ritchie, J.M., Weston, N.: Generative and discriminative voxel modeling with convolutional neural networks. arXiv preprint arXiv:1608.04236 (2016)
5. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16123–16133 (2022)
6. Chan, E.R., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5799–5809 (2021)
7. Chen, H., Gu, J., Chen, A., Tian, W., Tu, Z., Liu, L., Su, H.: Single-stage diffusion nerf: A unified approach to 3d generation and reconstruction. arXiv preprint arXiv:2304.06714 (2023)
8. Chen, T.: On the importance of noise scheduling for diffusion models. arXiv preprint arXiv:2301.10972 (2023)
9. Chen, Y., Wang, T., Wu, T., Pan, X., Jia, K., Liu, Z.: Comboverse: Compositional 3d assets creation using spatially-aware diffusion guidance. arXiv preprint arXiv:2403.12409 (2024)
10. Cheng, Y., Yin, F., Huang, X., Yu, X., Liu, J., Feng, S., Yang, Y., Tang, Y.: Efficient text-guided 3d-aware portrait generation with score distillation sampling on distribution. arXiv preprint arXiv:2306.02083 (2023)
11. Dai, W., Chen, L.H., Wang, J., Liu, J., Dai, B., Tang, Y.: Motionlcm: Real-time controllable motion generation via latent consistency model. arXiv preprint arXiv:2404.19759 (2024)
12. Deng, Y., Yang, J., Xiang, J., Tong, X.: Gram: Generative radiance manifolds for 3d-aware image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10673–10683 (2022)
13. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* **34**, 8780–8794 (2021)

14. Ding, Z., Zhang, M., Wu, J., Tu, Z.: Patched denoising diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2308.01316 (2023)
15. Gao, J., Shen, T., Wang, Z., Chen, W., Yin, K., Li, D., Litany, O., Gojcic, Z., Fidler, S.: Get3d: A generative model of high quality 3d textured shapes learned from images. arXiv preprint arXiv:2209.11163 (2022)
16. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020)
17. Gu, J., Zhai, S., Zhang, Y., Bautista, M.A., Susskind, J.: f-dm: A multi-stage diffusion model via progressive signal transformation. arXiv preprint arXiv:2210.04955 (2022)
18. Gu, J., Zhai, S., Zhang, Y., Susskind, J., Jaitly, N.: Matryoshka diffusion models. arXiv preprint arXiv:2310.15111 (2023)
19. Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector quantized diffusion model for text-to-image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10696–10706 (2022)
20. Gupta, A., Xiong, W., Nie, Y., Jones, I., Oğuz, B.: 3dgen: Triplane latent diffusion for textured mesh generation. arXiv preprint arXiv:2303.05371 (2023)
21. Hang, T., Gu, S., Li, C., Bao, J., Chen, D., Hu, H., Geng, X., Guo, B.: Efficient diffusion training via min-snr weighting strategy. arXiv preprint arXiv:2303.09556 (2023)
22. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems* **30** (2017)
23. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
24. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. *The Journal of Machine Learning Research* **23**(1), 2249–2281 (2022)
25. Hoogeboom, E., Heek, J., Salimans, T.: simple diffusion: End-to-end diffusion for high resolution images. arXiv preprint arXiv:2301.11093 (2023)
26. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685 (2021)
27. Jabri, A., Fleet, D., Chen, T.: Scalable adaptive computation for iterative generation. arXiv preprint arXiv:2212.11972 (2022)
28. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. arXiv preprint arXiv:2305.02463 (2023)
29. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4401–4410 (2019)
30. Kingma, D.P., Welling, M., et al.: An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning* **12**(4), 307–392 (2019)
31. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* **114**(13), 3521–3526 (2017)
32. Li, R., Li, X., Hui, K.H., Fu, C.W.: Sp-gan: Sphere-guided 3d shape generation and manipulation. *ACM Transactions on Graphics (TOG)* **40**(4), 1–12 (2021)

33. Li, Z., Tucker, R., Snavely, N., Holynski, A.: Generative image dynamics. arXiv preprint arXiv:2309.07906 (2023)
34. Liao, Y., Schwarz, K., Mescheder, L., Geiger, A.: Towards unsupervised learning of generative models for 3d controllable image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5871–5880 (2020)
35. Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3d: High-resolution text-to-3d content creation. arXiv preprint arXiv:2211.10440 (2022)
36. Lin, S., Liu, B., Li, J., Yang, X.: Common diffusion noise schedules and sample steps are flawed. arXiv preprint arXiv:2305.08891 (2023)
37. Liu, J., Dai, W., Wang, C., Cheng, Y., Tang, Y., Tong, X.: Plan, posture and go: Towards open-world text-to-motion generation. arXiv preprint arXiv:2312.14828 (2023)
38. Ma, H., Zhang, L., Zhu, X., Zhang, J., Feng, J.: Accelerating score-based generative models for high-resolution image synthesis. arXiv preprint arXiv:2206.04029 (2022)
39. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
40. Müller, N., Siddiqui, Y., Porzi, L., Bulo, S.R., Kotschieder, P., Nießner, M.: Diffrf: Rendering-guided 3d radiance field diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4328–4338 (2023)
41. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
42. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International Conference on Machine Learning. pp. 8162–8171. PMLR (2021)
43. Niemeyer, M., Geiger, A.: Giraffe: Representing scenes as compositional generative neural feature fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11453–11464 (2021)
44. Ntavelis, E., Siarohin, A., Olszewski, K., Wang, C., Van Gool, L., Tulyakov, S.: Autodecoding latent 3d diffusion models. arXiv preprint arXiv:2307.05445 (2023)
45. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 165–174 (2019)
46. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023)
47. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022)
48. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
49. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10684–10695 (2022)
50. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Ayan, B.K., Mahdavi, S.S., Lopes, R.G., et al.: Photorealistic text-

- to-image diffusion models with deep language understanding. arXiv preprint arXiv:2205.11487 (2022)
51. Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems* **33**, 20154–20166 (2020)
 52. Shue, J.R., Chan, E.R., Po, R., Ankner, Z., Wu, J., Wetzstein, G.: 3d neural field generation using triplane diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20875–20886 (2023)
 53. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
 54. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems* **33**, 7462–7473 (2020)
 55. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456 (2020)
 56. Sun, J., Zhang, B., Shao, R., Wang, L., Liu, W., Xie, Z., Liu, Y.: Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. arXiv preprint arXiv:2310.16818 (2023)
 57. Szabó, A., Meishvili, G., Favaro, P.: Unsupervised generative 3d shape learning from natural images. arXiv preprint arXiv:1910.00287 (2019)
 58. Tang, J., Wang, T., Zhang, B., Zhang, T., Yi, R., Ma, L., Chen, D.: Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. arXiv preprint arXiv:2303.14184 (2023)
 59. Tang, J., Zeng, Y., Fan, K., Wang, X., Dai, B., Chen, K., Ma, L.: Make-it-vivid: Dressing your animatable biped cartoon characters from text. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6243–6253 (2024)
 60. Tang, J., Zhang, B., Yang, B., Zhang, T., Chen, D., Ma, L., Wen, F.: Explicitly controllable 3d-aware portrait generation. arXiv preprint arXiv:2209.05434 (2022)
 61. Tang, Y., Liu, J., Liu, A., Yang, B., Dai, W., Rao, Y., Lu, J., Zhou, J., Li, X.: Flag3d: A 3d fitness activity dataset with language instruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 22106–22117 (2023)
 62. Teng, J., Zheng, W., Ding, M., Hong, W., Wangni, J., Yang, Z., Tang, J.: Relay diffusion: Unifying diffusion process across resolutions for image synthesis. arXiv preprint arXiv:2309.03350 (2023)
 63. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in Neural Information Processing Systems* **30** (2017)
 64. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689 (2021)
 65. Wang, T., Zhang, B., Zhang, T., Gu, S., Bao, J., Baltrusaitis, T., Shen, J., Chen, D., Wen, F., Chen, Q., 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*. pp. 4563–4573 (2023)
 66. Wang, X., Li, Y., Zhang, H., Shan, Y.: Towards real-world blind face restoration with generative facial prior. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9168–9178 (2021)

67. Wood, E., Baltrušaitis, T., Hewitt, C., Dziadzio, S., Cashman, T.J., Shotton, J.: Fake it till you make it: face analysis in the wild using synthetic data alone. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3681–3691 (2021)
68. Wu, J., Zhang, C., Xue, T., Freeman, B., Tenenbaum, J.: Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *Advances in Neural Information Processing Systems* **29** (2016)
69. Xia, W., Xue, J.H.: A survey on deep generative 3d-aware image synthesis. *ACM Computing Surveys* **56**(4), 1–34 (2023)
70. Xiang, J., Yang, J., Deng, Y., Tong, X.: Gram-hd: 3d-consistent image generation at high resolution with generative radiance manifolds. *arXiv preprint arXiv:2206.07255* (2022)
71. Yin, F., Zhang, Y., Wang, X., Wang, T., Li, X., Gong, Y., Fan, Y., Cun, X., Shan, Y., Oztireli, C., et al.: 3d gan inversion with facial symmetry prior. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 342–351 (2023)
72. Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: Lion: Latent point diffusion models for 3d shape generation. *arXiv preprint arXiv:2210.06978* (2022)
73. Zhang, B., Cheng, Y., Yang, J., Wang, C., Zhao, F., Tang, Y., Chen, D., Guo, B.: Gaussiancube: Structuring gaussian splatting using optimal transport for 3d generative modeling. *arXiv preprint arXiv:2403.19655* (2024)
74. Zhang, B., Gu, S., Zhang, B., Bao, J., Chen, D., Wen, F., Wang, Y., Guo, B.: Styleswin: Transformer-based gan for high-resolution image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11304–11314 (2022)
75. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)
76. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 586–595 (2018)