

OrthoPlanes: A Novel Representation for Better 3D-Awareness of GANs

Honglin He^{1,2†*} Zhuoqian Yang^{1,3‡*} Shikai Li¹ Bo Dai¹ Wayne Wu¹✉
¹ Shanghai AI Laboratory ² Tsinghua University
³ School of Computer and Communication Sciences, EPFL

hehl21@mails.tsinghua.edu.cn zhuoqian.yang@epfl.ch lishikai@pjlab.org.cn

daibo@pjlab.org.cn wuwenyan0503@gmail.com



Figure 1: **We build a new hybrid explicit-implicit 3D representation - OrthoPlanes.** Our model can synthesize diverse objects with reasonable geometry without target-specific 3D priors (Left). Furthermore, as a general representation, our approach can be translated to conditional tasks expediently (Right).

Abstract

We present a new method for generating realistic and view-consistent images with fine geometry from 2D image collections. Our method proposes a hybrid explicit-implicit representation called **OrthoPlanes**, which encodes fine-grained 3D information in feature maps that can be efficiently generated by modifying 2D StyleGANs. Compared to previous representations, our method has better scalability and expressiveness with clear and explicit information. As a result, our method can handle more challenging view-angles and synthesize articulated objects with high spatial degree of freedom. Experiments demonstrate that our method achieves state-of-the-art results on FFHQ and SHHQ datasets, both quantitatively and qualitatively. Project page: <https://orthoplanes.github.io/>.

* Equal contribution.

† Interns at Shanghai AI Laboratory.

‡ Work done as research engineer at Shanghai AI Laboratory.

1. Introduction

Recovering 3D world from 2D images, known as inverse rendering, is a typical problem in computer vision and computer graphics. It has many practical uses in VR/AR and other domains like movie production and virtual try-on. Physically-based methods [9, 12] can yield superior outcomes, but they are associated with notable costs and intensive labor requirements. On the other hand, data-centric approaches stand out for their adaptability, ease of use, and photorealism [56]. In this study, our objective is to develop a novel data-driven representation that enhances the 2D GAN’s understanding of 3D geometry, enabling it to produce more detailed and realistic images of diverse objects.

3D-aware GANs have advanced rapidly to synthesize images that are consistent across multiple views from 2D image collections [6, 5, 18, 43, 41, 49, 62, 22, 65, 51]. These methods combine unsupervised learning, neural rendering and super resolution to produce realistic results. However, they continue to face challenges in accurately reconstructing 3D shapes from 2D images, a crucial step in rendering 2D GAN genuinely aware of 3D contexts. Certain methods attempt to address this issue utilizing a tri-plane representation

[5] or through the use of multiple parallel images [64] for the reconstruction of 3D entities. However, these strategies fall short when dealing with non-rigid entities like human figures that exhibit asymmetrical poses and varied appearances, which cannot be adequately represented by merely three orthogonal or parallel projections.

Inspired by successful scene-specific representations [38], we present a method for efficiently upscaling the generalizable tri-plane representation [5]. The main idea of our work is that a representation indexed via sectional projection can better represent 3D world than one that is indexed via orthogonal or parallel projection, because it has more explicit 3D information such as projection direction and distance. Based on this intuition, we introduce a novel hybrid explicit-implicit 3D representation, **OrthoPlanes**, as shown in Fig. 2 (d). We introduce a pre-defined location embedding to each feature map [28] to make them location-aware. The renderings exhibit better view consistency owing to this explicit data. Moreover, the geometry can be more detailed since the codebook size and density surpasses that of the representations based on orthogonal projections.

With acceptable increase of computing costs over EG3D [5], our approach can improve image quality, deal with more difficult viewing perspectives and generate articulated objects with high spatial variability, *e.g.*, human bodies, as shown in Fig. 1. Our experiment show state-of-the-art results for 3D-aware image synthesis from 2D collections on diverse datasets including FFHQ [27] and SHHQ [15].

We summarize our contributions as the follows: 1) We present a 3D representation termed *orthoplanes* aimed at enhancing the 3D awareness of 2D GANs, which significantly improves view-consistency and geometry. 2) we add a new branch to the 2D generator ensuring both training efficiency and heightened scalability.

2. Related Works

Neural scene representation. Neural scenes optimized through differentiable neural representations from collections of 2D-images and corresponding camera pose [37, 1, 19, 55, 47, 14, 31, 32, 29, 50, 24, 63, 60] and rendered via volume rendering [35] has been gaining attention. The essence of these scene representations revolves around associating each spacial location with a feature vector, which can range from color, density, to signed distance [44], among others. Generally, this mapping technique, or the 3D-scene representation, can be categorized into two primary types: point-wise processing [37, 63, 44] and methods based on lookup procedures [14, 5, 38].

Point-wise methods include fully-implicit and local-implicit representations. The fully-implicit representations utilize a coordinate neural network and depict the scene as a continuous function. Although they tend to be efficient in memory usage, rendering a scene using these implicit

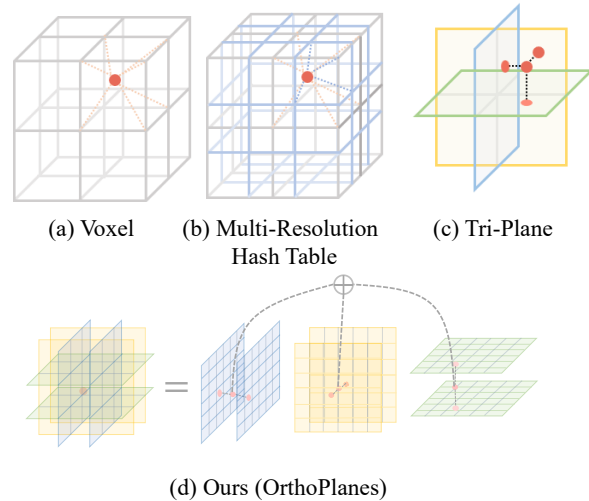


Figure 2: **Look-up based 3D representations.** Dense voxel grid [53] (a) uses explicitly located anchor-point table to store the scene. Multi-resolution hash table [38] (b) uses multiple hash tables to store the at with different resolutions. Tri-plane [5] (c) projects 3D spatial points onto three orthogonal planes, each of which stores compressed features of scene. Our orthoplanes (d) uses three orthogonal plane-groups to store the scene, each plane stores compressed features of the scene but preserves information associated with projection distance.

approaches can be time-consuming in practice. This is because every queried point must be processed through deep neural networks. On the other hand, local-implicit methods [47, 23, 7, 4] rely on manually designed partitions. A point within a specific partition can query its corresponding shallow neural network to ascertain its features, which offers numerous advantages.

Look-up based methods include explicit and hybrid representations. Explicit representations such as voxel grids [53], are fast to query but memory unfriendly. Hybrid explicit-implicit representations [38, 5, 34] which bring together the advantages of both forms, are often favored. Conceptually, these approaches serve as codebooks storing the spatial or temporal features [58, 13] of a scene. The representational capacity of these methods is intrinsically tied to the size of their codebooks and spatial distribution of the codes. As illustrated in Fig. 2, our proposed representation employs three orthogonal groups of parallel planes to encode the scene. On the one hand, the increased number of planes result in increased density of 3D information. On the other hand, the dispersion of this 3D data facilitates the spread of gradients throughout the 3D scene during training.

3D-aware image generation. Generative adversarial networks [17] have achieved photo-realistic 2D image quality [27, 28, 26]. A recent trend in image synthesis is to generate high-quality multi-view-consistent images with high-fidelity




Figure 3: **Our unconditional GAN framework.** Our generator comprises of: a codebook generator with a mapping network, a neural renderer implemented by **orthoplanes** representation with a lightweight MLP-based feature decoder, a volumetric renderer and a super resolution module. The discriminator employs dual-discriminator design to ensure view-inconsistency.

geometry from 2D image collections. Previous works used voxel-based GANs, in which 2D-CNN are replaced by 3D-CNN [16, 20, 39, 40]. Due to the use of voxel grids and 3D convolutions, the memory requirements of these approaches are so high that it makes high-resolution training infeasible. These limitations seen a lift thanks to the development of neural implicit representations [37, 49, 6]. Nevertheless, since these methods process each spatial points by a deep MLP, it is still challenging to generate images at a resolution achieved by state-of-the-art 2D methods.

For the creation of an effective 3D GAN capable of producing high-resolution images with precise geometry, EG3D [5] was introduced. Leveraging StyleGAN2 as the foundation to generate 2D feature maps that serve as codebooks for point features, and combined with upsampling layers and a dual-discrimination approach, it proves both efficient and expressive. While it successfully delivers high-quality images with detailed geometry for rigid objects, the mere three projections fall short in adequately representing articulated objects possessing a greater spatial complexity. We discover that the quality of generated image and geometry is associated with the size of the codebooks, not only the resolution but also the spatial distribution features. We are thus motivated to improve the expressiveness of this hybrid explicit-implicit representation by increasing the number of feature planes and placing them evenly across the 3D scene.

Approaches for 3D-aware image synthesis like GMPI [64] and EpiGRAF [54] demonstrate impressive image quality. The core difference between these works and ours is that we design a novel 3D representation used for volumetric rendering of generator motivated by enlarging the size of neural codebook while GMPI is based on multiplane image rendering which is 2.5D and EpiGRAF paid attention to developing efficient discriminators for training to synthesize high-resolution images directly based on existing representations.

3. Methodology

Our goal is to use orthoplanes in GAN to synthesize high-quality multi-view-consistent image with high-fidelity geometry on diverse datasets. In this section, we first introduce our motivation and the design of orthoplanes (Sec. 3.1). While improving the expressiveness based on well-defined planes, we maintain the efficiency based on our new design of generator (Sec. 3.2). We discuss inference of our model on both unconditional (Sec. 3.3) and conditional (Sec. 3.4) generation tasks.

3.1. OrthoPlanes 3D Representation

As discussed in EG3D [5], to train a high-resolution 3D-aware GAN, the representation should be both efficient and expressive. The tri-plane representation projects each 3D coordinate to three axis-aligned orthogonal feature planes, each with size $N \times N \times C$ where N is the resolution and C is the number of channels. The feature associated with the location $x \in R^3$ is given by $F_{xy} + F_{yz} + F_{zx}$, each summed term extracted from its corresponding plane via bilinear interpolation. The color and density used for volume rendering [37] is obtained by processing this feature with a lightweight MLP.

From our perspective, a tri-plane is essentially a codebook for storing 3D information efficiently like other methods [53, 38]. For a 3D space with a resolution of N , it scales at $O(N^2)$ complexity, which is more efficient than the $O(N^3)$ complexity of voxel [53]. However, if one wish to enlarge the size of the codebook from N^2 to KN^2 , K must be the power of 4, and if $K = 4^i$, the backbone of StyleGAN2 must be with i more layers, which is inefficient based on the design of StyleGAN2.

In our orthoplanes formulation, we use three axis-aligned orthogonal groups of feature planes, each with K parallel planes with resolution of $N \times N \times C$, as shown in Fig. 2d. Each plane is anchored along its projection direction, and

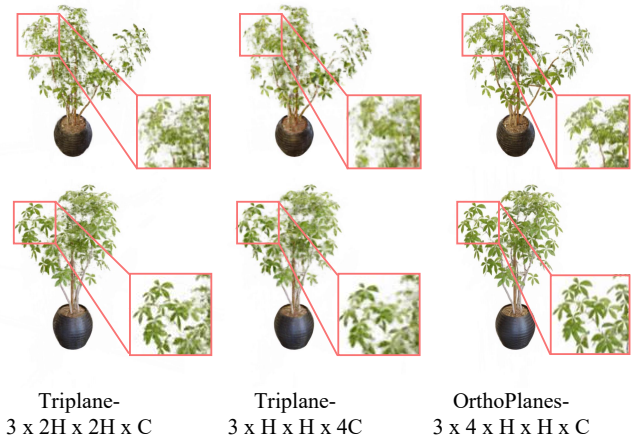


Figure 4: **Single-scene overfitting on the ficus scene.** Please zoom in for details.

	<i>lego</i>	<i>ficus</i>	<i>family</i>	<i>caterpillar</i>
Tri-plane ($H^2, 4C$)	26.56	23.13	30.59	22.03
Tri-plane ($(2H)^2, C$)	28.66	23.50	30.48	22.32
Ours ($H^2, C; K = 4$)	31.12	26.83	32.06	23.93

Table 1: **Quantitative evaluation of PSNR on SSO cases.**

distributed evenly along the axis. We associate each spatial location with a combination of three sectional projections onto the nearest planes on three orthogonal directions. Take $F_{X_f Y_f Z_c}$ as an example, it denotes the feature extracted from plane group $X_f Y_f$, projected along Z_c , using trilinear interpolation.

We demonstrate the expressiveness of orthoplanes with a single-scene overfitting (SSO) experiment. We show quantitative results in Tab. 1 and one example scene in Fig. 4. As shown in Tab. 1, our model archives better PSNR in all cases compared to tri-plane, while maintaining consistent parameter values among methods. As shown in Fig. 4, our approach recovers the color tone of *ficus*, especially asymmetric parts like sparse leaves and protruding branches.

Note that our method scales at $O(KN^2)$ complexity, which means we can control the codebook size by manipulating K . Meanwhile, just like a voxel with inhomogeneous resolution, it’s still much more efficient than voxel grids when K is much smaller than N . Furthermore, based on the StyleGAN2 backbone, we can generate these planes efficiently using the method detailed in Sec. 3.2.

3.2. Codebook Generator

The feature codebook of our orthoplanes representation is generated with a StyleGAN2 [28] generator. The random latent code z sampled from normal distribution is processed by the mapping network to get the style vector w , used to modulate the convolution kernels in synthesis layers.

To make the plane location-aware, we introduce a location embedding l_i to mark the position of each plane. Each plane-

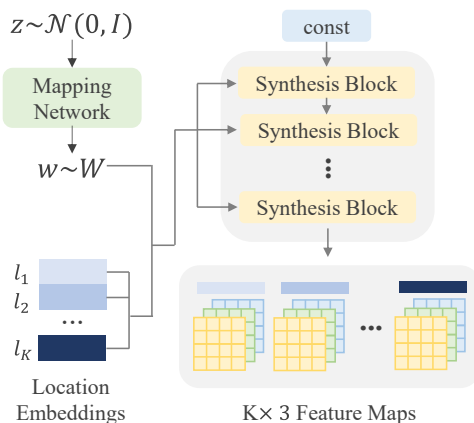


Figure 5: **Architecture of orthoplanes generator.** The structure of synthesis blocks follow StyleGAN2 [28]. We replace the ToRGB layer by our ToFeature layer to generate location-embedded feature maps in parallel

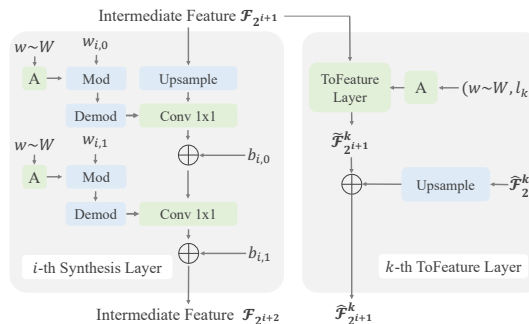


Figure 6: **Implementation of our generator.** There are K parallel ToFeature layers with shared weights in a single synthesis layer. We reuse the intermediate feature map synthesized by StyleGAN2 synthesis layer to ensure efficiency.

group is a set of tuples (I_i, l_i) , $i \in \{1, 2, \dots, K\}$, where $I_i \in R^{H \times W \times C}$ denotes the feature map of the i -th plane, while $l_i \in [-1, 1]$ denotes the location embeddings of the i -th plane along corresponding projection direction. In practice, we use a StyleGAN2 generator to generate tensors of shape $K \times N \times N \times 3C$, and reshape it to $3 \times K \times N \times N \times C$ by split and permutation to obtain the three orthogonal plane groups, where N is the resolution and $3C$ the number of channels.

The same process can be viewed as generating K tri-planes in parallel. Inspired by GMPI [64], we choose to reuse the same intermediate feature map synthesized in each synthesis block across all tri-planes, repeat them K times, and modulate the results in *ToRGBLayer* using the style vector w with location embeddings l_i . Specifically, we use $\mathcal{R} = \{4, 8, 16, \dots, N\}$ to denote the resolution of the results synthesized by each synthesis block. The intermediate feature maps at each resolution $h \in \mathcal{R}$ are computed with a synthesis

	FID ↓	#FLOPs ↓	T.Speed ↑	T.Mem. ↓
EG3D (256 ² , 32)	14.1	126.25G	1×	1×
EG3D (512 ² , 32)	14.2 (+0.7%)	150.14G(+18.9%)	0.51×	1.55×
EG3D (256 ² , 128)	11.6 (-17.7%)	149.38G(+18.3%)	0.41×	1.75×
Ours (256 ² , 32; K = 4)	11.2 (-20.6%)	129.57G(+2.6%)	0.82×	1.16×

Table 2: **Quantitative evaluation on 100K-SHHQ.** *CB. size* denotes relative codebook size, i.e. $K \times H^2 \times C$. *T.Speed* and *T.Mem* denote the speed and GPU memory usage during training, respectively. *I.Speed* denotes inference speed. *CB.Size*, *T.Speed*, *T.Mem* and *I.Speed* are all reported relative to the EG3D (256², 32) model.

block $f_{Syn,h}$, i.e.,

$$\mathcal{F}_h = f_{Syn,h}(\mathbf{w}), \forall h \in \mathcal{R} \quad (1)$$

The i -th feature codebook $\hat{\mathcal{F}}_N^i \in R^{N \times N \times 3C}$ used in volumetric rendering, with resolution of N , is obtained by accumulating intermediate results at all resolutions, i.e.,

$$\hat{\mathcal{F}}_h^i = \begin{cases} \tilde{\mathcal{F}}_h^i + Upsample(\hat{\mathcal{F}}_{\frac{h}{2}}^i), & h \in \mathcal{R} \setminus \{4\} \\ \tilde{\mathcal{F}}_h^i, & h = 4 \end{cases} \quad (2)$$

where $\tilde{\mathcal{F}}_h^i$ is the residual given by *ToFeature* Layer $f_{ToFeat,h}$ at resolution h , and $\tilde{\mathcal{F}}_h^i$ the feature codebook. $f_{ToFeat,h}$ takes \mathcal{F}_h and concatenated vector $(\mathbf{w}, \mathbf{l}_i)$ as input, to get the location-specific feature codebook. Based on this design, $f_{ToFeat,h}$ operates on each tri-plane individually, and the operation can be done in parallel. Furthermore, it uses the same intermediate feature map \mathcal{F}_h as input. These characteristics preserve the efficiency while expanding the expressive power of 3D representation. In Table. 2, compared to enlarging the resolution or the number of channels of the generated feature map, our method can generate larger codebooks with less memory and time.

3.3. Unconditional GAN Framework

Fig. 3 gives an overview of our generator architecture for unconditional 3D-aware image synthesis. Orthoplanes representation is used as the codebook for spatial query during volumetric rendering. Following [41, 5], the output of volumetric renderer is a feature image \mathbf{I}_F and a RGB image \mathbf{I}_{RGB}^- , both in low resolution. A super resolution module upsamples these images to yield the final high-resolution RGB image \mathbf{I}_{RGB} .

Orthoplanes Generator. The codebook for orthoplanes is produced using a StyleGAN2 generator, as outlined in [28], incorporating our supplementary branch detailed in Sec. 3.2. Initially, the random latent code \mathbf{z} is sampled from a standard normal distribution $\mathcal{N}(0, I)$. This is subsequently transformed by a mapping network to obtain the intermediate latent code \mathbf{w} from the distribution \mathcal{W} . This code then modulates the weights of the convolution kernels within the codebook synthesis network.

The output of the synthesis network is of shape of $K \times 256 \times 256 \times 96$. We obtain orthoplanes with shape $3 \times K \times 256 \times 256 \times 32$ through channel-wise split and permutation.

Feature Aggregation. To render an image of the 3D scene, we sample spatial locations along camera rays calculated with camera parameters \mathbf{P} , and query their features from the orthoplanes codebook.

The feature of a spatial location from each plane-group is obtained through trilinear interpolation. The features from three orthogonal plane-groups are summed to yield the aggregated feature. This aggregate feature is then processed using a streamlined MLP decoder. Similar to the approach in EG3D [5], the MLP input excludes any point-wise data such as positional encoding.

Volumetric Rendering and Super Resolution. The volumetric renderer is developed with the importance sampling strategy, as highlighted in [37]. In line with other studies [41, 5], the renderer yields a low-resolution RGB image, denoted as \mathbf{I}_{RGB}^- , accompanied by a feature image \mathbf{I}_F . These images are subsequently refined by a super-resolution module to produce the final RGB image, \mathbf{I}_{RGB} . This super-resolution module comprises of two StyleGAN2 synthesis blocks, which are modulated using the style vector \mathbf{w} . In all our tests, the input to the super-resolution module has a resolution of 128² and a channel size of 32.

Dual Discrimination. As introduced in EG3D [5], we use dual discrimination to enforce view inconsistency, and use the camera parameters \mathbf{P} as condition of the discriminator.

Our models in all unconditional experiments use non-saturating GAN Loss [17] with R1 regularization [36].

$$\mathcal{L}_{adv,D} = \mathbf{E}_{\mathbf{z} \sim P_z, \mathbf{c} \sim P_{data}} [f(D(G(\mathbf{z}; \mathbf{c})))] \quad (3)$$

$$+ \mathbf{E}_{\mathbf{I} \sim P_{data}} [f(D(\mathbf{I})) + \lambda \|\nabla_{\mathbf{I}} D(\mathbf{I})\|_2^2] \quad (4)$$

$$\mathcal{L}_{adv,G} = \mathbf{E}_{\mathbf{z} \sim P_z, \mathbf{c} \sim P_{data}} [f(-D(G(\mathbf{z}; \mathbf{c})))] \quad (5)$$

where $f(x) = -\log(1 + \exp(-x))$, P_{data} is the real data distribution and \mathbf{c} the camera parameter.

3.4. Conditional GAN Framework

We also examine the effectiveness of our representation for generating articulated objects. Specifically, we build a pose-conditional full-body human image generation model with orthoplanes.

The model architecture of our pose-conditional generator is given in the supplementary materials. The pose parameters θ defined by SMPL [33] are used as the condition for both of generator and discriminator. We follow the practice of transforming queried coordinates to canonical human-body space [45, 57, 2, 22, 62].

Pose-Canonical Human Generation. We leverage the orthoplanes representation for generation of clothed humans in the canonical space \mathcal{C} . Specifically, the random latent code \mathbf{z}




Figure 7: **Curated examples at 512^2** . (we clip images of human bodies to the resolution of 512×256), synthesized by models trained with SHHQ [15], FFHQ [27] and AFHQv2-Cat [8].



Figure 8: **Curated examples on conditional SHHQ [15]**. (we clip images to the resolution of 512×256), synthesized by models trained with SHHQ [15]. Meshes on the left side are visualization of SMPL [33].

is used to generate the canonical feature of clothed humans via codebook generator as described in Sec. 3.2. Every point sampled in observation space \mathcal{O} will be transformed to the canonical space \mathcal{C} before querying the codebook.

Delta SDF Learning Strategy. Like concurrent works [22, 62], we use employ the Delta SDF [61] learning strategy instead of predicting density directly. Specifically, the spatial point \mathbf{x} will first compute its signed distance $d_0(\mathbf{x})$ from the SMPL Template mesh under given pose θ , instead of output the density $\sigma(\mathbf{x})$, we just output the offset $\Delta d(\mathbf{x})$ from $d_0(\mathbf{x})$ [61], and translate SDF value $d_0(\mathbf{x}) + \Delta d(\mathbf{x})$ to the density value $\sigma(\mathbf{x})$ using method proposed by StyleSDF [43].

Training. Please see the supplement for implementation details of our conditional GAN framework.

4. Experiments and Results

Datasets. To illustrate the effectiveness of our method, we compare with state-of-the-art methods in both unconditional and conditional 3D-aware image generation. For the task of unconditional generation, we compare performance on three datasets FFHQ [27], AFHQv2-Cat [8] and SHHQ [15] separately. FFHQ is a real-world human-face dataset. AFHQv2-Cat is a small cat-face dataset. SHHQ is a real-world human-body dataset with about 220K images. For FFHQ and AFHQv2-Cat, we use the version with horizontal flips provided by EG3D [5]. For AFHQv2-Cat, we apply transfer learning [25] from FFHQ checkpoint provided by EG3D [5] with adaptive data augmentation [25].

For the task of conditional generation, we compare on

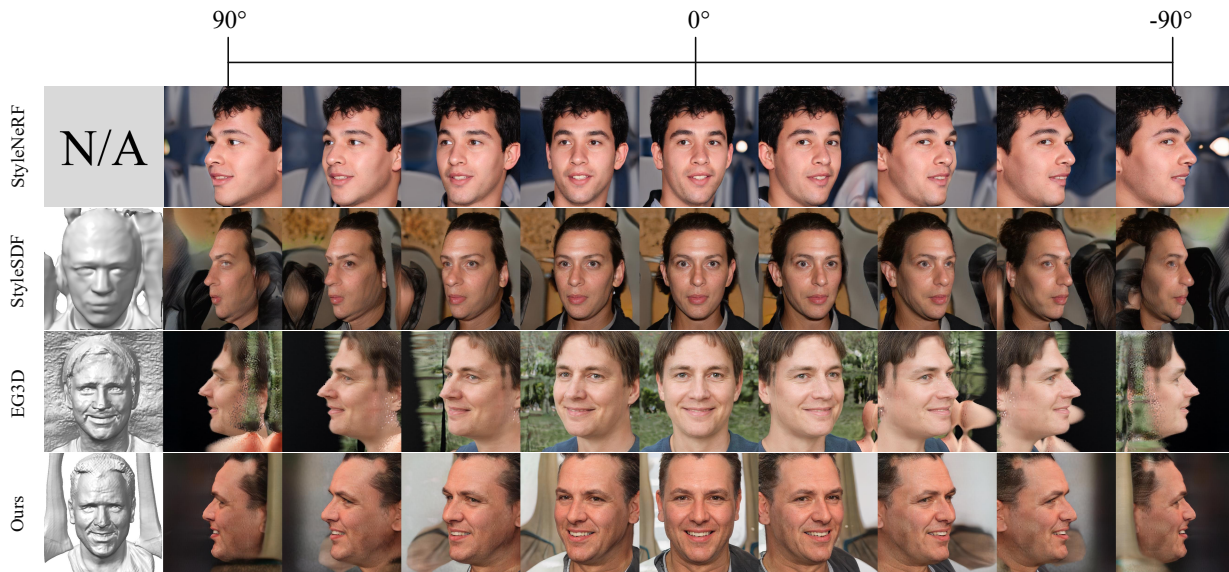


Figure 9: **Images rendered on various camera poses.** There are obvious distortions on extreme yaw angles in previous works. In contrast, our method is better on both of 2D image quality and 3D geometry.

the full-body human image dataset SHHQ, SMPL pose-condition for each image is estimated using an off-the-shelf model PARE [30].

Baselines. We compare our method against nine state-of-the-art methods for 3D-aware image synthesis. For unconditional generation, we compare against five methods: π -GAN [6], StyleSDF [43], StyleNeRF [18], GMPI [64] and EG3D [5]. For conditional GAN, we compare our method against four methods: ENARF-GAN [42], 3DHumanGAN [59], EVA3D [22] and AvatarGen [62].

Evaluation Metrics. We follow prior work to assess the results using common metrics. To measure 2D qualities, we report Fréchet Inception Distance [21] (FID) and Kernel Inception Distance [3] (KID). For unconditional FFHQ, we report the identity consistency (ID) [10, 11, 5] to measure the view-consistency of generated results. For unconditional SHHQ, we report the depth accuracy (Depth) by calculating Mean Squared Error (MSE) between pseudo depth-maps estimated from generated RGB images by MiDaS [46] and the generated ones. For conditional SHHQ, we additionally report Percentage of Correct Keypoints (PCK) [42] to evaluate the effectiveness of the pose controllability.

4.1. Comparisons

Qualitative Results. Fig. 7 presents selected examples synthesized by our model with different datasets. And Fig. 8 presents selected examples synthesized by our conditional human body generative model. Fig. 10 and Fig. 9 provide qualitative comparisons against baselines. StyleSDF, EG3D and ours are all based the two-stage approach, which uses low-resolution neural rendering followed by super resolution modules. StyleSDF is a fully-implicit method while EG3D

	FFHQ			AFHQv2-Cat	
	FID ↓	KID ↓	ID ↑	FID ↓	KID ↓
pi-GAN 128 ²	29.9	35.73	0.67	16.0	14.92
StyleSDF 256 ²	11.5	2.65	-	12.8*	4.47*
StyleNeRF 512 ²	7.80	2.20	-	13.2*	3.60*
GMPI 512 ²	8.29	4.54	0.74	7.79	4.74
EG3D 512 ²	4.70 [‡]	1.32	0.77	2.77[†]	0.41[†]
Ours-S 512 ²	<u>4.11[‡]</u>	1.05	0.71	<u>2.82[†]</u>	<u>0.46[†]</u>
Ours-R 512 ²	4.01[‡]	<u>1.23</u>	0.73	-	-

Table 3: **Quantitative evaluation for FFHQ and AFHQv2-Cats.** For KID, we report all results $\times 10^3$.

* Metrics are reported on the whole AFHQv2 dataset.

† Trained with adaptive discriminator augmentation [25].

‡ Trained with camera condition [5].

is hybrid explicit-implicit. Geometry generated by SDF is overly smooth and lack of high-frequency details. The geometry of EG3D on SHHQ lack of details because the dataset is extremely unbalanced on viewpose, the majority of data is front view that it will converge to the state that feature only depends on one projection plane. As shown in Fig. 10, the geometry of EG3D is too flat while ours is with the feeling of depth thanks to our design of sectional projection.

When it comes to view-consistency, as shown in Fig. 9. Our work performs significantly better than previous works. Our work is capable of rendering at a consistent high quality at large yaw angles while previous methods fail, showing different levels of degradation and artifacts. StyleNeRF [18] and StyleSDF [43] can only capture the information within the distribution of dataset. EG3D [5] can synthesize reasonable results within larger range. However, limited by the

	FID ↓	KID ↓	Depth ↓
StyleSDF 512^2	33.29	25.2	0.036
StyleNeRF 512^2	7.60	3.96	-
EG3D 512^2	5.79	2.26	0.082
Ours 512^2	4.18	2.05	0.082

Table 4: **Quantitative evaluation for SHHQ.** All models are trained from scratch until convergence.

	FID ↓	KID ↓	Depth ↓	PCK ↑
ENARF-GAN 128^2	20.09	16.94	0.086	79.40
3DHumanGAN \dagger 512^2	9.31	<u>5.16</u>	-	-
EVA3D \dagger 512^2	11.99	9.00	0.017	88.95
AvatarGen $\dagger \ddagger$ 512^2	4.29	-	0.365	99.49
Ours 512^2	<u>9.00</u>	4.37	<u>0.032</u>	<u>99.12</u>

Table 5: **Quantitative evaluation for conditional SHHQ.**

\dagger We quota results from their papers.

\ddagger AvatarGen uses 50K fake images and less than 50K real images to compute FID, which may cause difference.

expressiveness of orthogonal projection and the sparse distribution of features and gradients in space, it fails on extreme yaw angles. With the predefined sectional projection with projection distance, our method captures more spatial information. As an additional benefit, the geometry of human face is separated from the background without additional process, as shown in Fig. 7 and Fig. 9.

Quantitative Results. Table. 3 and Table. 4 provide quantitative metrics comparing the proposed approach and baselines on unconditional generation. For FFHQ and AFHQv2-Cat, Ours-S and Ours-R are models trained from scratch and fine-tuned on pretrained checkpoint provided by EG3D [5], respectively. In the config-R, the weights of our new branch are randomly initialized, after training about 14% of EG3D’s full pipeline training time, it demonstrates notable improvements in FID scores. For SHHQ, our method is the best on FID and KID, for the Depth metrics, StyleSDF is better thanks to higher resolution of neural renderer. To eliminate the possible impact of data volume, we train a model with a small subset of SHHQ with 40K images, results are given in the supplementary materials. For AFHQv2-Cats, since we use the checkpoint of FFHQ from EG3D [5], there are some key weights of ToFeature layer that cannot be transferred from a larger dataset, limiting its performance. Additionally, Table. 5 provides quantitative metrics on conditional generation, more results and analysis are given in the supplementary materials.

4.2. Ablation Study

Location embeddings. As described in Sec. 3.2, we attach a location embedding to each style vector in the orthoplane codebook. Using a linear embedding $\mathbf{l}_k = -1 + \frac{k-1}{K} \cdot 2$ is able to distinguish these planes. If we want to learn finer geometry details, high-frequency information between these



Figure 10: **Qualitative comparison on FFHQ and SHHQ.** Shapes are extracted by marching cubes.




Figure 11: **Ablation study.** With frequency encoding, there will be less artifacts. By increasing the number of planes, there will be more details and none artifacts on geometry.

planes is needed, for which we use $FE(\mathbf{l}_k)$ as location embeddings in majority of our experiments. Intuitively, frequency encoding is equivalent to assigning higher weights for the location embeddings. It’s a simple but useful solution to represent the difference among images with different projection distance. As shown in Table. 6 and Fig. 11, frequency encoding provides improvement in FID and geometry.

Number of planes. Intuitively, network with more planes can capture more information. What’s more, based on the flexible design of ToFeature branch, the number of planes can be scaled up easily even after training. As shown in Table. 7, using checkpoint provided by EG3D [5] (1-plane) with some minor additional training, we can obtain better results. As shown in Table. 6 and Fig. 11, the increase of planes brings improvement on FID and geometry.

#.Plane	Loc.Embed	FID ↓	KID ↓	Depth ↓
4	Linear	12.85	8.81	0.098
4	Frequency	11.52	7.04	0.110
8	Frequency	10.87	7.29	0.082

Table 6: **Ablating the number of planes and location embeddings.** Measured in terms of FID, KID and Depth scores on SHHQ.

	FID ↓	T.C	Extra T.C
EG3D	4.70	68	-
Ours (From Scratch)	4.11	68	-
EG3D + 8 plane	4.39	-	10
EG3D + 12 plane	4.01	-	10

Table 7: **Ablating the number of planes and location embeddings.** Measured in terms of FID and training cost (T.C) of V100 GPU days on FFHQ. EG3D + K plane denotes that we train from the checkpoint provided by EG3D.

4.3. Applications

Single image 3D reconstruction. Through pivotal tuning inversion (PTI) [48], our model can be used to reconstruct 3D shapes from single views of objects. In this manuscript, we provide results of human body reconstruction with known pose conditions. We use off-the-shelf pose estimator PARE [30] to extract body-pose parameters and camera parameters. Following PTI [48], we optimize the latent codes for 500 iterations, followed by fine-tuning the generator weights for an additional 500 iterations.

As shown in Fig. 12, the reconstructed results are with fine geometry. Novel views of the same identity can also be synthesized. Furthermore, thanks to our decoupled design, we can use other poses to generate images of the same identity with novel poses, which enables the reconstructed person to be used as an animatable avatar.

Style mixing. Our model inherits the property of StyleGAN [27, 28, 26] latent space, although our models are all trained without style-mixing regularization. We use style-mixing in both of unconditional and pose-conditional tasks, while using the same pose as condition, the mixed samples are all with the correct pose, showing that our conditional model decouples the appearance and pose. More results and analysis are given in the supplementary.

Interpolation on latent space. Our model generates a smooth transition between latent codes, showing that latent space learned by the model is semantically meaningful. Results are given in the supplementary materials.

5. Discussion

Limitations and future work. Although our work improves the performance of view-consistency and geometry effec-

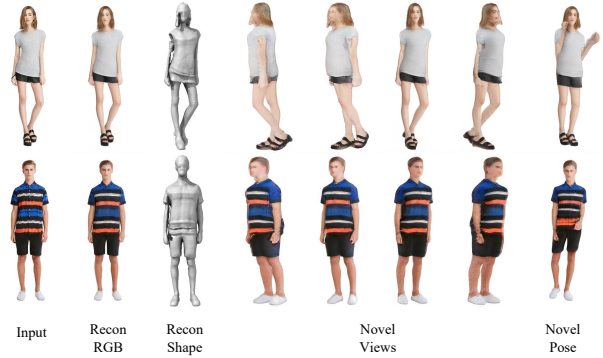


Figure 12: **Inversion on SHHQ [15] (Conditional).** Based on PTI [48], our model can be used to reconstruct 3D shapes with single image. Furthermore, based on our conditional model, using optimized style vector and model parameters, we can synthesize images with novel view and novel pose of the given identity.

tively compared to previous works. There are still some artifacts at background. To tackle this problem, some modeling assumptions [63, 52] can be used.

Furthermore, based on the two-stage strategy, *i.e.*, rendering low-resolution feature map and upsampling, there are some inconsistencies of the generated results under view variation, occasionally. One possible solution is to rendering RGB image directly. To tackle the problem of the unbearable computing overhead of rendering high-resolution images, some works [54] introduce other strategies, where our new representation can be incorporated into.

Conclusion. We design a novel representation orthoplanes to make 2D GAN more 3D-aware. Our model generates multi-view-consistent results and fine geometry even under extreme view angles. By introducing more explicit spatial information, our method can synthesize diverse objects and is not confined to rigid objects. This could foster innovative methods for 3D reconstruction and facilitate the creation of 3D models for various applications.

References

- [1] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 2
- [2] Alexander W Bergman, Petr Kellnhofer, Yifan Wang, Eric R Chan, David B Lindell, and Gordon Wetzstein. Generative neural articulated radiance fields. *arXiv preprint arXiv:2206.14314*, 2022. 5
- [3] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018. 7
- [4] Rohan Chabra, Jan E Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed

- 3d reconstruction. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pages 608–625. Springer, 2020. [2](#)
- [5] 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. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [6] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. 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*, pages 5799–5809, 2021. [1](#), [3](#), [7](#)
- [7] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8628–8638, 2021. [2](#)
- [8] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020. [6](#)
- [9] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (ToG)*, 34(4):1–13, 2015. [1](#)
- [10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. [7](#)
- [11] Yu Deng, Jialong 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. [7](#)
- [12] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d: Real-time performance capture of challenging scenes. *ACM Transactions on Graphics (ToG)*, 35(4):1–13, 2016. [1](#)
- [13] Sara Fridovich-Keil, Giacomo Meanti, Frederik Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. *arXiv preprint arXiv:2301.10241*, 2023. [2](#)
- [14] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. [2](#)
- [15] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen Change Loy, Wayne Wu, and Ziwei Liu. Stylegan-human: A data-centric odyssey of human generation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVI*, pages 1–19. Springer, 2022. [2](#), [6](#), [9](#)
- [16] Matheus Gadelha, Subhransu Maji, and Rui Wang. 3d shape induction from 2d views of multiple objects. In *2017 International Conference on 3D Vision (3DV)*, pages 402–411. IEEE, 2017. [3](#)
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [2](#), [5](#)
- [18] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985*, 2021. [1](#), [7](#)
- [19] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021. [2](#)
- [20] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Escaping plato’s cave: 3d shape from adversarial rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9984–9993, 2019. [3](#)
- [21] 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. [7](#)
- [22] Fangzhou Hong, Zhaoxi Chen, Yushi Lan, Liang Pan, and Ziwei Liu. Eva3d: Compositional 3d human generation from 2d image collections. *arXiv preprint arXiv:2210.04888*, 2022. [1](#), [5](#), [6](#), [7](#)
- [23] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020. [2](#)
- [24] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1251–1261, 2020. [2](#)
- [25] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33:12104–12114, 2020. [6](#), [7](#)
- [26] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. [2](#), [9](#)
- [27] 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. [2](#), [6](#), [9](#)
- [28] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the*

- IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. [2](#), [4](#), [5](#), [9](#)
- [29] Petr Kellnhofer, Lars C Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4287–4297, 2021. [2](#)
- [30] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11127–11137, 2021. [7](#), [9](#)
- [31] David B Lindell, Julien NP Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14556–14565, 2021. [2](#)
- [32] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. [2](#)
- [33] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. [5](#), [6](#)
- [34] Julien NP Martel, David B Lindell, Connor Z Lin, Eric R Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. *arXiv preprint arXiv:2105.02788*, 2021. [2](#)
- [35] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. [2](#)
- [36] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International conference on machine learning*, pages 3481–3490. PMLR, 2018. [5](#)
- [37] B Mildenhall, PP Srinivasan, M Tancik, JT Barron, R Ramamoorthi, and R Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, 2020. [2](#), [3](#), [5](#)
- [38] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. [2](#), [3](#)
- [39] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7588–7597, 2019. [3](#)
- [40] Thu H Nguyen-Phuoc, Christian Richardt, Long Mai, Yongliang Yang, and Niloy Mitra. Blockgan: Learning 3d object-aware scene representations from unlabelled images. *Advances in neural information processing systems*, 33:6767–6778, 2020. [3](#)
- [41] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021. [1](#), [5](#)
- [42] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Unsupervised learning of efficient geometry-aware neural articulated representations. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*, pages 597–614. Springer, 2022. [7](#)
- [43] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13503–13513, 2022. [1](#), [6](#), [7](#)
- [44] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. [2](#)
- [45] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021. [5](#)
- [46] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3):1623–1637, 2020. [7](#)
- [47] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14335–14345, 2021. [2](#)
- [48] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on Graphics (TOG)*, 42(1):1–13, 2022. [9](#)
- [49] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems*, 33:20154–20166, 2020. [1](#), [3](#)
- [50] Katja Schwarz, Axel Sauer, Michael Niemeyer, Yiyi Liao, and Andreas Geiger. Voxgraf: Fast 3d-aware image synthesis with sparse voxel grids. *arXiv preprint arXiv:2206.07695*, 2022. [2](#)
- [51] Zifan Shi, Yujun Shen, Yinghao Xu, Sida Peng, Yiyi Liao, Sheng Guo, Qifeng Chen, and Dit-Yan Yeung. Learning 3d-aware image synthesis with unknown pose distribution. *arXiv preprint arXiv:2301.07702*, 2023. [1](#)
- [52] Minjung Shin, Yunji Seo, Jeongmin Bae, Young Sun Choi, Hyunsu Kim, Hyeran Byun, and Youngjung Uh. Ballgan: 3d-aware image synthesis with a spherical background. *arXiv preprint arXiv:2301.09091*, 2023. [9](#)
- [53] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvox-

- els: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019. [2](#), [3](#)
- [54] Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. Epigraf: Rethinking training of 3d gans. *arXiv preprint arXiv:2206.10535*, 2022. [3](#), [9](#)
- [55] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. [2](#)
- [56] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering. In *Computer Graphics Forum*, volume 39, pages 701–727. Wiley Online Library, 2020. [1](#)
- [57] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16210–16220, 2022. [5](#)
- [58] Eric Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Wenqing Zhang, Song Bai, Jiashi Feng, and Mike Zheng Shou. Pv3d: A 3d generative model for portrait video generation. *arXiv preprint arXiv:2212.06384*, 2022. [2](#)
- [59] Zhuoqian Yang, Shikai Li, Wayne Wu, and Bo Dai. 3dhuman: Towards photo-realistic 3d-aware human image generation. *arXiv preprint arXiv:2212.07378*, 2022. [7](#)
- [60] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. [2](#)
- [61] Wang Yifan, Lukas Rahmann, and Olga Sorkine-Hornung. Geometry-consistent neural shape representation with implicit displacement fields. *arXiv preprint arXiv:2106.05187*, 2021. [6](#)
- [62] Jianfeng Zhang, Zihang Jiang, Dingdong Yang, Hongyi Xu, Yichun Shi, Guoxian Song, Zhongcong Xu, Xinchao Wang, and Jiashi Feng. Avatargen: a 3d generative model for animatable human avatars. *arXiv preprint arXiv:2211.14589*, 2022. [1](#), [5](#), [6](#), [7](#)
- [63] Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. [2](#), [9](#)
- [64] Xiaoming Zhao, Fangchang Ma, David Güera, Zhile Ren, Alexander G Schwing, and Alex Colburn. Generative multiplane images: Making a 2d gan 3d-aware. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 18–35. Springer, 2022. [2](#), [3](#), [4](#), [7](#)
- [65] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. Cips-3d: A 3d-aware generator of gans based on conditionally-independent pixel synthesis. *arXiv preprint arXiv:2110.09788*, 2021. [1](#)