

# GRAM-HD: 3D-Consistent Image Generation at High Resolution with Generative Radiance Manifolds

Jianfeng Xiang<sup>\*1,2</sup> Jiaolong Yang<sup>2</sup> Yu Deng<sup>\*1,2</sup> Xin Tong<sup>2</sup>

<sup>1</sup>Tsinghua University <sup>2</sup>Microsoft Research Asia

{t-jxiang, jiaoyan, xtong}@microsoft.com dengyu2008@hotmail.com



Figure 1: Curated generation samples of  $1024^2$  human faces  $512^2$  cats from GRAM-HD. Our method can generate high-resolution, high-quality and strict 3D-consistent images with explicitly controllable poses. Note the realistic geometry details such as the whisker of cats. (Best viewed with zoom-in; see also the [project page](#) for more results with videos.)

## Abstract

Recent works have shown that 3D-aware GANs trained on unstructured single image collections can generate multiview images of novel instances. The key underpinnings to achieve this are a 3D radiance field generator and a volume rendering process. However, existing methods either cannot generate high-resolution images (e.g., up to  $256 \times 256$ ) due to the high computation cost of neural volume rendering, or rely on 2D CNNs for image-space upsampling which jeopardizes the 3D consistency across different views. This paper proposes a novel 3D-aware GAN that can generate high resolution images (up to  $1024 \times 1024$ ) while keeping strict 3D consistency as in volume rendering. Our motivation is to achieve super-resolution directly in the 3D space to preserve 3D consistency. We avoid the otherwise prohibitively-expensive computation cost by applying 2D convolutions on a set of 2D radiance manifolds

defined in the recent generative radiance manifold (GRAM) approach, and apply dedicated loss functions for effective GAN training at high resolution. Experiments on FFHQ and AFHQv2 datasets show that our method can produce high-quality 3D-consistent results that significantly outperform existing methods. It makes a significant step towards closing the gap between traditional 2D image generation and 3D-consistent free-view generation.<sup>1</sup>

## 1. Introduction

While generative modeling of 2D images have achieved tremendous success [5, 22–24] with generative adversarial networks (GANs) [17], 3D-aware GANs that aims to generate photorealistic multiview images begun to emerge in recent years [8, 12, 41, 47, 54, 67]. Despite both being trained on unstructured 2D image collections, the latter is capable of synthesizing the images of an object at different 3D viewpoints. The key to achieve this is to generate an un-

<sup>\*</sup>Work done when JX and YD were interns at MSRA.

<sup>1</sup>Project page: <https://jeffrexyang.github.io/GRAM-HD/>

derlying 3D representation, for which neural radiance field (NeRF) [37] has been the cornerstone for recent methods. With volumetric rendering, NeRF can produce realistic images while enforcing strong 3D consistency across views.

However, the high computational cost of neural volumetric rendering greatly limits the affordable image resolution for GAN training. It also introduces hurdles in fine detail generation due to insufficient point sampling. A workaround is only judging whether a patch of the generated image is real to not during training [47, 54] rather than a whole image. But using a patch discriminator may lack global perception of the images and lead to inferior image generation quality. The recent generative radiance manifold (GRAM) method [12] significantly improved the generation quality by sampling points on a set of learned surface manifolds. Still, it can only be trained on images up to  $256 \times 256$  resolution on modern GPUs, which is in sheer contrast to state-of-the-art 2D GANs that can easily model  $1024 \times 1024$  images with moderate computing cost.

Along a different axis, many methods resort to 2D convolutions to tackle the dilemma [7, 18, 41, 44, 64]. A straightforward idea shared by these methods is to render low-resolution images or feature maps and apply 2D CNNs to increase the resolution. With this strategy, they have demonstrated higher-resolution generation (e.g.,  $512 \times 512$  for [7, 64] and  $1024 \times 1024$  for [18, 44]). Unfortunately, image-space upsampling with 2D CNNs inevitably incurs 3D inconsistency among the generated multiview images. As such, these methods can be used in user-interactive image generation and manipulation but are not suitable for video synthesis and animation.

We propose *GRAM-HD*, a GAN method that can synthesize strongly 3D-consistent images at high resolution. Our motivation is to do upsampling or super-resolution in the 3D space and keep the volume rendering paradigm to retain strict 3D consistency. But how to achieve this efficiently is not straightforward (e.g., upsampling a discretized low-resolution volume using 3D convolutions quickly becomes untractable for high-resolution output). In this paper, we leverage the GRAM [12] method, which defines a set of surface manifolds, to handle high resolution generation. Our key insight is that the surface manifolds can be upsampled using 2D CNNs for efficient super-resolution. We flatten and sample each learned surface to regular 2D image grids and apply a shared 2D CNN for upsampling and feature-to-radiance translation. This way, our method not only ensures multiview consistency by generating a high-resolution 3D representation for rendering, but also enjoys the computational efficiency of 2D CNNs. In essence, we tackle a 3D super-resolution task with object-space 2D CNN.

We evaluate our method on the FFHQ [23] and AFHQv2-CATS [11] datasets. We show that GRAM-HD can generate photorealistic images that are both of

high resolution (up to  $1024 \times 1024$ ) and strongly multiview-consistent, which cannot be achieved by any previous method. It also outperforms GRAM in terms of both generation quality and speed at the same image resolution by upsampling low-resolution manifolds.

**The contributions of this work** are summarized below:

- We present a novel 3D-aware image generation approach that can generate high-resolution images (up to  $1024 \times 1024$ ) with strong multiview-consistency and highly-realistic geometry details.
- We introduce a method for 3D space super-resolution using efficient 2D CNN under the radiance manifold representation.
- We significantly reduced the computation cost of the radiance manifold based 3D-aware generation method while obtaining higher quality images (e.g., 76%↓ memory cost, 58%↓ inference time, 21%↓ FID on FFHQ-256<sup>2</sup>; 95%↓ inference time for 1024<sup>2</sup>)

## 2. Related Work

**Neural scene representations** Neural scene representations [14, 15, 36, 37, 42, 43, 46, 50–53, 56, 57] have seen tremendous progress in the past several years and shown promising results for modeling complex 3D scenes. Among these methods, NeRF [37] and its variants [1, 2, 34] excel at learning detailed scene structures from a collection of posed images and synthesizing 3D-consistent novel views. Nevertheless, the high computational complexity of NeRF restricts its power when applied to reconstruction of numerous instances [21, 32] or generative modeling of object categories [8, 47]. Several methods have been proposed to reduce the computation cost of NeRF, including utilizing sparse data structures [19, 31, 62], introducing multi-resolution feature encoding [33, 38], or reducing sampling points during rendering [30, 51]. However, it remains unclear how these approaches can be applied to a generative modeling paradigm. Very recently, GRAM [12] proposes to regulate radiance field learning on 2D manifolds and prominently improves the image generation quality at a resolution of  $256 \times 256$ . This paper inherits the manifolds representation of GRAM, and extends it to 3D-consistent image generation of high resolution.

**Generative 3D-aware image synthesis** 3D-aware generative models [8, 28, 39, 41, 47, 49] aim to learn multiview image synthesis of an object category given uncontrolled 2D images collections. They achieve this by incorporating 3D representation learning into GAN training. Earlier works [39, 40, 55] utilize mesh or voxel to represent the underlying 3D scenes. Recently, plenty of works [8, 45, 47, 60] leverage NeRF’s volumetric representation to achieve image generation with higher 3D-consistency. However,

NeRF’s high computational cost prevents them from generating high-resolution images with fine details. Most of the follow-up works [7, 18, 41, 44, 61, 64, 68] handle this problem by first rendering low-resolution images or feature maps and then applying 2D CNNs for super-resolution. They suffer from a common 3D inconsistency issue when varying camera viewpoints, due to the black-box rendering of CNNs. Some methods [7, 18] alleviate this problem by enforcing constraints between the generated high-resolution images and their low-resolution counterparts, but still cannot guarantee the consistency of high-frequency details. Different from the above methods, [12] learns radiance manifolds and applies manifold rendering [12, 69] to achieve high-quality and multiview-consistent image synthesis, but it has difficulties to generate images with a resolution beyond  $256 \times 256$ . In this paper, we propose a novel approach to learn high-resolution radiance manifolds, and achieve strongly multiview-consistent image generation at high resolution.

**Image super-resolution** Image super-resolution is a longstanding task in computer vision and has seen enormous developments [9, 10, 13, 16, 26, 29, 58, 59, 63, 65]. We leverage the techniques from this field, but for super-resolution of radiance manifolds in 3D space instead of 2D images.

### 3. Approach

Given a collection of 2D images, our method aims to learn, through adversarial learning, a 3D-aware image generator  $G$  which takes a latent code  $\mathbf{z} \in \mathbb{R}^{d_z} \sim p_z$  and an explicit camera pose  $\boldsymbol{\theta} \in \mathbb{R}^3 \sim p_\theta$  as inputs, and outputs a synthesized image  $I$  of an virtual instance determined by  $\mathbf{z}$  under pose  $\boldsymbol{\theta}$ :

$$G : (\mathbf{z}, \boldsymbol{\theta}) \in \mathbb{R}^{d_z+3} \rightarrow I \in \mathbb{R}^{H \times W \times 3}. \quad (1)$$

Figure 2 depicts the overall framework of our method. We first use a radiance manifold generator to generate the radiance and feature manifolds representing a low-resolution scene. Then we flatten and discretize the manifolds and upsample them using a shared CNN to get an high-resolution representation. After training, GRAM-HD can render high-definition and 3D-consistent images.

#### 3.1. Radiance Manifolds

The concept of radiance manifolds is introduced in GRAM [12] and we briefly review it here. The original neural radiance field [37] models the radiance of a continuous 3D scene with a neural network, and volumetric rendering is done by sampling points along each viewing ray and integrating the radiance. However, the high memory and computation cost greatly restricts the number of point samples for GAN training. The GRAM method regulates point sampling and radiance field learning on a set of learned surface

manifolds which can significantly improve image generation quality.

The manifolds are embodied as a set of iso-surfaces in a 3D scalar field. This scalar field is shared for all generated instances and represented by a light-weight MLP called the *manifold predictor*  $\mathcal{M}$ :

$$\mathcal{M} : \mathbf{x} \in \mathbb{R}^3 \rightarrow s \in \mathbb{R}. \quad (2)$$

With  $\mathcal{M}$ , the iso-surfaces  $\{S_i\}$  can be extracted with  $N$  pre-defined levels  $\{l_i\}$ :

$$S_i = \{\mathbf{x} | \mathcal{M}(\mathbf{x}) = l_i\}. \quad (3)$$

For rendering, point samples are calculated as the intersections between each viewing ray  $\mathbf{r}$  and the extracted surfaces:

$$\{\mathbf{x}_i\} = \{\mathbf{x} | \mathbf{x} \in \mathbf{r} \cap \{S_j\}\}. \quad (4)$$

Radiance of these point samples are then generated by the *radiance generator*  $\Phi$ , which is an MLP modulated by latent code  $\mathbf{z}$ :

$$\Phi : (\mathbf{z}, \mathbf{x}, \mathbf{d}) \in \mathbb{R}^{d_z+6} \rightarrow (\mathbf{c}, \alpha) \in \mathbb{R}^4, \quad (5)$$

where  $\mathbf{d}$  is a view direction,  $\mathbf{c}$  is the color, and  $\alpha$  is the occupancy. The final color value for each ray can be computed with the rendering equation [43, 69]:

$$C(\mathbf{r}) = \sum_{i=1}^N \prod_{j<i} (1 - \alpha(\mathbf{x}_j)) \alpha(\mathbf{x}_i) \mathbf{c}(\mathbf{x}_i). \quad (6)$$

Note the latent code  $\mathbf{z}$  and view direction  $\mathbf{d}$  are omitted here for brevity.

Our radiance manifolds are adapted from GRAM with a few twists. First, we only use these MLP-generated radiance manifolds to represent a coarse, low-resolution scene trained to render  $64 \times 64$  images. Second, since the intermediate feature also bears useful information, we concatenate them with the final output color and occupancy as the input to the subsequent manifold super-resolution.

#### 3.2. Manifold Super-Resolution

Our key insight for manifold super-resolution is that efficient 2D CNNs can be applied. To this end, we first map the manifolds to regular image grids.

**Manifold gridding** The manifold gridding operation  $\Pi$  flattens the learned surfaces  $\{S_i\}$  and samples them to get a set of low-resolution radiance and feature maps:

$$\Pi : \{S_i\} \rightarrow \{R_i\} \doteq \mathbf{R}_{\text{lr}} \in \mathbb{R}^{N, H_{\text{lr}}, W_{\text{lr}}, 4}, \quad (7)$$

where  $H_{\text{lr}}$  and  $W_{\text{lr}}$  are the spatial resolution for which we use  $64 \times 64$ . Such a gridding operation can be implemented in various ways. GRAM has shown that the learned surfaces are nearly planar. So in this work we simply use orthogonal projection to flatten the surfaces.

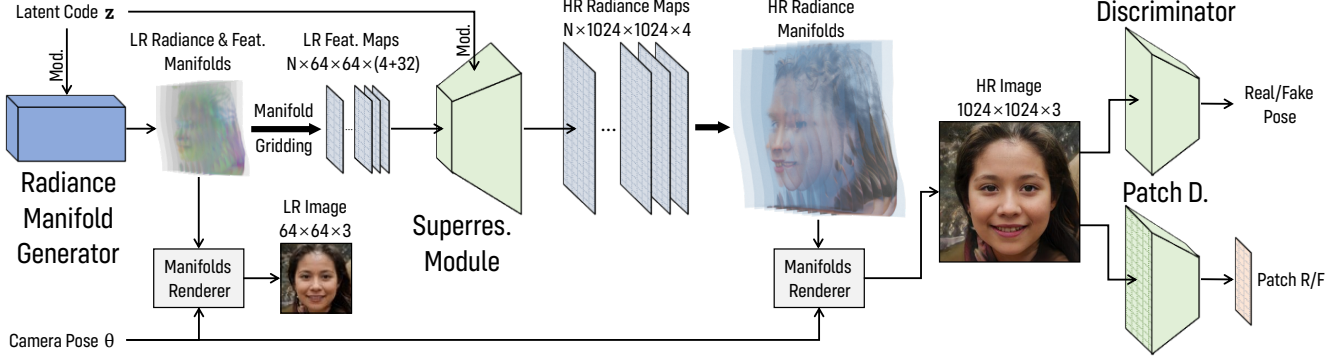


Figure 2: The overall framework of our GRAM-HD method. The generator consists of two components: the radiance manifold generator and the manifold super-resolution module. The former generates radiance and feature manifolds that represent an LR 3D scene. Through manifold gridding, the manifolds are sampled to discrete 2D feature maps. The super-resolution module then processes these feature maps and output HR radiance maps. Finally, an HR image can be rendered by computing ray-manifold intersections and integrating their radiance sampled from the HR radiance maps. Note that *not like* previous works utilizing *2D image super-resolution* for HR image generation, we directly do *3D representation super-resolution* and keep the volume rendering paradigm, thus keep the strong 3D consistency of the output images.

Specifically, we sample the surfaces by shooting  $H_{lr} \times W_{lr}$  rays that are parallel to Z-axis. For the first  $N-1$  surfaces, the sampling area encompasses the foreground objects in 3D. For the last surface which is the background plane in GRAM, we use a larger area as it spans a much wider region for covering the background under extreme views. We also apply an additional nonlinear mapping to sample denser points around the center region (see *suppl. material* for details). After concatenating the sampled color, occupancy, and intermediate features from  $\Phi$ , the manifold gridding produces a 4D tensor  $\mathbf{R}_{lr}^+ \in \mathbb{R}^{N, H_{lr}, W_{lr}, 4+d_f}$  where  $d_f$  is the channel number of features.

**Super-resolution CNN** To upsample the LR radiance maps, we condition the super-resolution CNN  $\mathcal{U}$  with latent code  $z$  such that the added details are also controlled by  $z$  to the extent possible:

$$\mathcal{U}: \begin{aligned} & (z \in \mathbb{R}^{d_z}, \mathbf{R}_{lr}^+ \in \mathbb{R}^{N, H_{lr}, W_{lr}, 4+d_f}) \\ & \rightarrow \mathbf{R}_{hr} \in \mathbb{R}^{N, H_{hr}, W_{hr}, 4}, \end{aligned} \quad (8)$$

where  $H_{hr}$  and  $W_{hr}$  are the HR resolution. Considering that the distribution of background radiance significantly differ from the foreground, we employ *two CNNs*, one for the first  $N-1$  foreground maps and another much smaller one (with half channels of the former) for the last background map. Adding this dedicated background network leads to higher image quality as we will show in the experiment. See also the *suppl. material* for more details.

Our CNN architecture is adapted from previous image super-resolution networks [48, 59, 66]. Specifically, we first apply Residual-in-Residual Dense Block (RRDB) [59] to process the LR radiance and feature maps, and then use sub-pixel convolutions [48, 66] for upsampling. RRDB is a CNN equipped with massive residual and dense connec-

tions which can effectively process the features, while sub-pixel convolution is a learnable upsampling method proven to work well for super-resolution. We empirically found that their combination leads to better performance.

We inject style information encoded by  $z$  using the weight modulation method of [24]. An MLP is employed to map  $z$  to  $s_i$ , and apply weight modulation to all convolutional layers after RRDB. More details of our CNN architecture can be found in the *suppl. material*.

**Image rendering** To render the final image after obtaining the HR radiance maps, we first calculate the intersections between the view rays and surface manifolds as in [12]. Then we use the same mapping function as in manifold gridding to get their projected 2D positions on the radiance maps and obtain radiance via bilinear interpolation. Pixel colors can then be calculated per Eq. 6.

### 3.3. Network Training

To train our model, we randomly sample latent code  $z$ , camera pose  $\theta$  and real image  $I$  from prior distributions  $p_z$ ,  $p_\theta$ , and  $p_{real}$ . We first use the following two loss functions as in GRAM [12].

**Adversarial loss** For adversarial learning [17], a discriminator  $D$  is used to distinguish the generated images from the real ones and compete with the generator  $G$ . We use the non-saturating GAN loss with R1 regularization [35] for training:

$$\mathcal{L}_{adv} = \mathbb{E}_{z \sim p_z, \theta \sim p_\theta} [f(D(G(z, \theta)))] + \mathbb{E}_{I \sim p_{real}} [f(-D(I)) + \lambda \|\nabla D(I)\|^2], \quad (9)$$

where  $f(x) = \log(1 + \exp x)$  is the softplus function. Two discriminators are applied for the low and high reso-




Figure 3: **Left:** Rendered images and radiance maps on the surface manifolds. Three sampled manifolds are shown here. The corresponding LR results before super-resolution are presented at bottom right. **Right:** Extracted proxy 3D shapes at HR (top row) and LR (bottom row). The rendered images are shown at bottom right for reference.

lution, respectively. The network architectures are similar to GRAM discriminator [12].

**Pose loss** Following GRAM [12], we add a loss to regularize the generate poses and ensure correct 3D geometry. Specifically, the discriminators also estimate a camera pose for each image, and  $L_2$  loss is calculated between given camera pose and the estimated one:

$$\mathcal{L}_{\text{pose}} = \mathbb{E}_{\mathbf{z} \sim p_z, \theta \sim p_\theta} \|D^p(G(\mathbf{z}, \theta)) - \theta\|^2 + \mathbb{E}_{I \sim p_{\text{real}}} \|D^p(I) - \hat{\theta}\|^2. \quad (10)$$

We also introduce two new losses below, which are designed for our high resolution generation task.

**Patch adversarial loss** We apply an extra patch discriminator [70] to the generated final images, which can help remove the checkerboard artifacts arisen due to upsampling. The patch adversarial loss can be written as:

$$\mathcal{L}_{\text{patch}} = \mathbb{E}_{\mathbf{z} \sim p_z, \theta \sim p_\theta} [f(D_{\text{patch}}(G(\mathbf{z}, \theta)))] + \mathbb{E}_{I \sim p_{\text{real}}} [f(-D_{\text{patch}}(I))] \quad (11)$$

**Cross-resolution consistency loss** Finally, we enforce the consistency between the generated LR and HR contents. We apply an  $L_2$  difference to penalize the difference in both the rendered images and radiance maps:

$$\mathcal{L}_{\text{cons}} = \mathbb{E}_{\mathbf{z} \sim p_z, \theta \sim p_\theta} [\|\Gamma(G_{\text{hr}}(\mathbf{z}, \theta)) - G_{\text{lr}}(\mathbf{z}, \theta)\|^2 + \|\Gamma(\mathbf{R}_{\text{hr}}) - \mathbf{R}_{\text{lr}}\|^2], \quad (12)$$

where  $\Gamma$  denotes the bicubic downsampling operator. This loss enforces the super-resolution module to focus on fine detail generation without introducing dramatic change.

We employ a two-stage training strategy in our implementation. At the first stage, we train  $G_{\text{lr}}$  and  $D_{\text{lr}}$  (i.e., the GRAM model) for low-resolution image generation using  $\mathcal{L}_{\text{adv}}$  and  $\mathcal{L}_{\text{pose}}$ . Then we froze  $G_{\text{lr}}$  and train high resolution generation using all the four loss functions.

## 4. Experiments

**Implementation details** We train our method on two datasets: FFHQ [23] and AFHQv2-CATS [11], which contain 70K human face images of  $1024^2$  resolution and 5.5K cat face images of  $512^2$  resolution, respectively. Camera poses are estimated using off-the-shelf landmark detectors [6, 27]. 24 surface manifolds are used in GRAM-HD for both human face and cats as in GRAM [12]. All our models are trained on 8 NVIDIA Tesla V100 GPUs with 16GB memory. More training details such as learning rate and batchsize are presented in the *suppl. material*.

### 4.1. Visual Results

**Generated images** Figure 1 shows some images samples generated by our method. GRAM-HD can generate high-quality images at high resolution with rich details. Moreover, it allows explicit manipulation of camera pose while maintains strong 3D consistency across different views. For some thin structures like human hair, glasses, and whiskers of cats, our method produces realistic fine details and correct parallax effect viewed from different angles.

**Radiance maps and 3D geometry** Figure 3 (left) shows the rendered images and radiance maps on manifold surfaces before and after super-resolution. Visually inspected, the HR images and radiance maps are consistent with the LR counterparts in general but contain much more details and thin structures, especially for those high-frequency area like hair and fur. Figure 3 (right) presents the 3D proxy shapes extracted for some generated instances. Here we use a multi-view depth fusion method to extract the shapes since at high resolution the grid sampling and mesh extraction strategy used in [12] is not applicable; details can be found in the *suppl. material*. As we can see, the shapes extracted at high resolution contain much more shape details.

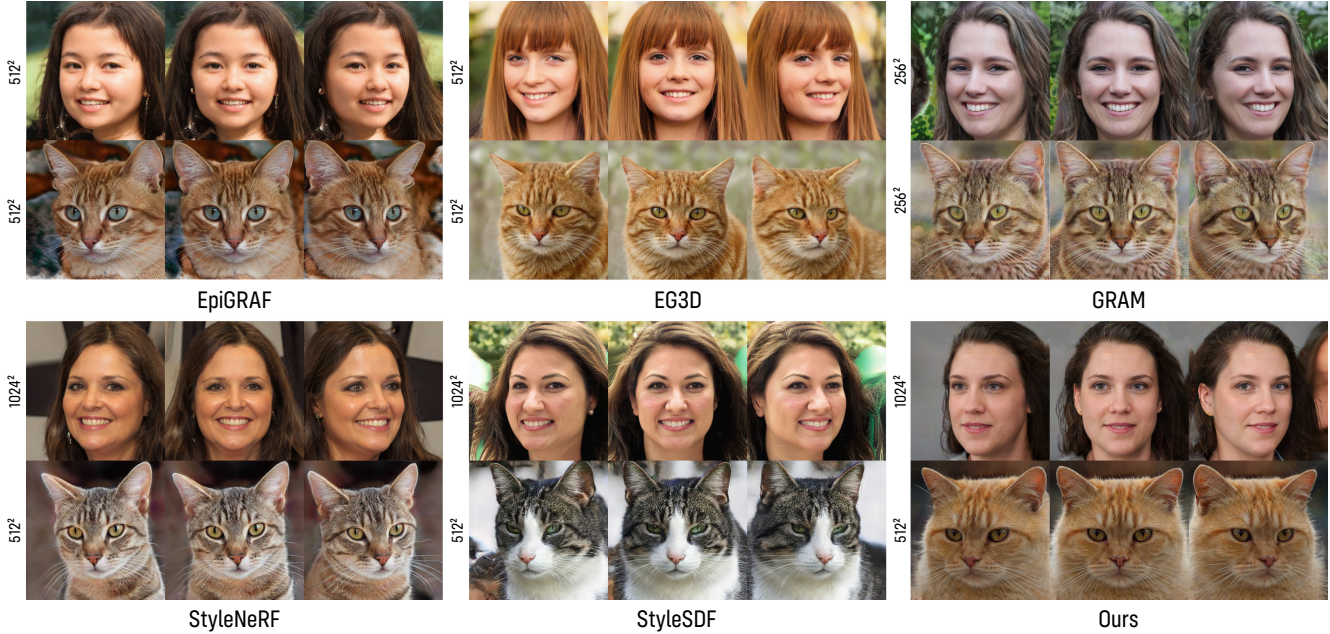


Figure 4: Qualitative comparison with recent 3D-aware GANs. The cat images of StyleNeRF are taken from their paper which is produced by a model trained on all images in AFHQv2; our training of StyleNeRF on cat images failed. (**Best viewed with zoom-in**)




Figure 5: Comparison of 3D consistency using spatiotemporal line textures akin to the Epipolar Line Images (EPI) [4]. We rotate the camera horizontally and stack the texture of a fixed horizontal line segment. Our method leads to a natural and smooth texture pattern, whereas others yield distorted and/or noisy patterns, indicating different degrees of 3D inconsistency. Detailed explanations can be found in the text. (**See the accompanying video for results under continuous view change.**)

## 4.2. Comparison with Previous Methods

We compare our GRAM-HD with several recent 3D-aware GANs that can generate high resolution images, including StyleNeRF [18], StyleSDF [44], EG3D [7] and EpiGRAF [54]. For StyleNeRF and StyleSDF, we use their publicly-released models trained on FFHQ for comparison, and train them with the released code on AFHQv2-CATS. However, our training of StyleNeRF on AFHQv2-CATS

failed even though multiple trials were made, so no comparison is made for this setting. For EG3D, we use their released models. For EpiGRAF, we train with official implementation using identical datasets for comparison.

**Qualitative comparison** Figure 4 shows some generated images of different methods: StyleNeRF [18], StyleSDF [44], EG3D [7], EpiGRAF [54], GRAM [12], and our GRAM-HD. Visually inspected, the images from

Table 1: Quantitative comparison of image quality using FID and KID ( $\times 1000$ ) between 20K generated and real images.

Method	FFHQ1024		FFHQ512		CATS512		FFHQ256	
	FID	KID	FID	KID	FID	KID	FID	KID
StyleNeRF [18]	9.45	<b>2.65</b>	-	-	-	-	<b>9.24</b>	<b>3.19</b>
StyleSDF [44]	<b>9.44</b>	2.83	-	-	7.91	3.90	-	-
EG3D [7]	-	-	<b>8.72</b>	<b>3.61</b>	<b>6.28</b>	<b>1.67</b>	-	-
EpiGRAF [54]	-	-	14.7	6.93	9.40	4.11	16.1	7.73
<i>Ours</i>	12.0	5.23	12.2	5.41	7.67	3.41	11.8	4.72

Table 2: Quantitative comparison of 3D consistency measured by the multiview reconstruction quality of the NeuS [57] method.

Method	FFHQ1024		FFHQ512		CATS512		FFHQ256	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
StyleNeRF [18]	30.0	0.80	-	-	-	-	31.9	0.92
StyleSDF [44]	31.1	0.84	-	-	26.6	0.75	-	-
EG3D [7]	-	-	33.7	0.88	28.4	0.78	-	-
<i>Ours</i>	<b>33.8</b>	<b>0.87</b>	<b>34.0</b>	<b>0.90</b>	<b>28.8</b>	<b>0.81</b>	<b>36.5</b>	<b>0.96</b>

Table 3: Quantitative comparison of image quality between GRAM and our GRAM-HD. Due to its high memory cost, GRAM can only train on up to  $256^2$  resolution.

Method	FFHQ256		CATS256	
	FID	KID	FID	KID
GRAM [12]	15.0	6.55	12.9	7.37
<i>Ours</i>	<b>11.8</b>	<b>4.72</b>	<b>7.05</b>	<b>2.53</b>

GRAM-HD have similar quality to other method that adopts 2D CNN for image-space upsampling and superior to EpiGRAF which is though a pure 3D representation, but lack details due to volume rendering.

However, the results of other methods except for EpiGRAF exhibit different levels of 3D inconsistency under view change, as can be observed in our accompanying video. To better visualize 3D consistency here, we present the spatiotemporal textures of different methods in Figure 5. Specifically, we smoothly rotate the camera horizontally and stack the texture of a fixed horizontal line segment, forming some spatiotemporal texture images similar to the Epipolar Line Images (EPI) [4]. For 3D-consistent generation, the resultant spatiotemporal texture should appear smooth and natural. Figure 5 shows that the resultant texture image of StyleNeRF contain both distorted and noisy regions, indicating both low-frequency inconsistency and high-frequency texture flicking. The line texture from StyleSDF barely changes across views, indicating the texture sticking artifact [22]. EG3D’s spatiotemporal texture also contains some non-smooth, high-frequency patterns which are caused by texture flicking. In contrast, the texture

Table 4: Comparison of running time (*left*) and memory cost (*right*) on a NVIDIA V100 GPU. Whole computational graph retained. “Strict 3D” means methods that do not use 2D upsampling thus have strong multiview consistency.

	Methods	$256^2$			$512^2$			$1024^2$		
		256 <sup>2</sup>	512 <sup>2</sup>	1024 <sup>2</sup>	256 <sup>2</sup>	512 <sup>2</sup>	1024 <sup>2</sup>	256 <sup>2</sup>	512 <sup>2</sup>	1024 <sup>2</sup>
Strict 3D	piGAN	0.29s	1.15s	4.59s	21.1G	OOM	OOM			
	GRAM	0.43s	1.54s	6.69s	22.2G	OOM	OOM			
	<i>Ours</i>	0.18s	0.22s	0.36s	5.4G	6.3G	9.0G			
Not strict 3D	StyleNeRF	0.06s	0.08s	0.16s	0.48G	0.85G	3.0G			
	StyleSDF	0.11s	0.11s	0.13s	0.40G	0.53G	0.80G			
	EG3D	-	0.10s	-	-	2.96G	-			

Table 5: Ablation study on FFHQ256.

Method	FID
StyleGAN2 architecture	18.2
Base architecture	14.4
+ Sub-pixel convolutions	14.1
+ Style modulation	14.0
+ LR generator feature	13.3
+ Background-net	12.0
+ $\mathcal{L}_{patch}$ ( <i>Ours</i> )	<b>11.8</b>
- $\mathcal{L}_{cons}$	16.9

from GRAM-HD appears natural with no noticeable noise or distortion, demonstrating its strong 3D consistency.

**Quantitative comparison** Table 1 shows the FID [20] and KID [3] metrics of different methods. Our scores are slightly higher than those 3D-inconsistent methods, which indicate again the comparable image generation quality. We further compare GRAM-HD to 3D-consistent methods EpiGRAF and GRAM in Table 1 and Table 3, respectively. GRAM-HD consistently outperforms on all experiment settings and can be applied to higher resolution. To evaluate computation efficiency, we record the running time and memory cost for each method in Table 4, where GRAM-HD is at most 3 times slower than 3D-inconsistent methods but significantly faster than 3D-consistent methods, especially at high resolution. This demonstrates that GRAM-HD not only has superior computation efficiency but also generates higher-quality images.

To quantitatively evaluate 3D consistency, for each method we generate 30 images under different views, and train the multiview reconstruction method NeuS [57] on them. We report the PSNR and SSIM scores of the images reconstructed by NeuS. In theory, the more consistent the input multiview images are, the higher the reconstruction quality will be. Table 2 presents the scores averaged on 50 randomly generated instances. As expected, our GRAM-HD leads to consistently better reconstruction than StyleNeRF, StyleSDF and EG3D, demonstrating its superior multiview consistency compared to the competing methods.




Figure 6: **Left:** Sample results w/ and w/o the patch adversarial loss  $\mathcal{L}_{\text{patch}}$  on 1024<sup>2</sup> resolution.  $\mathcal{L}_{\text{patch}}$  can effectively eliminate checkerboard artifacts. **Right:** Sample results w/ and w/o cross-resolution consistency loss  $\mathcal{L}_{\text{cons}}$ . Some unwanted floaters appear in front of the faces without  $\mathcal{L}_{\text{cons}}$ .

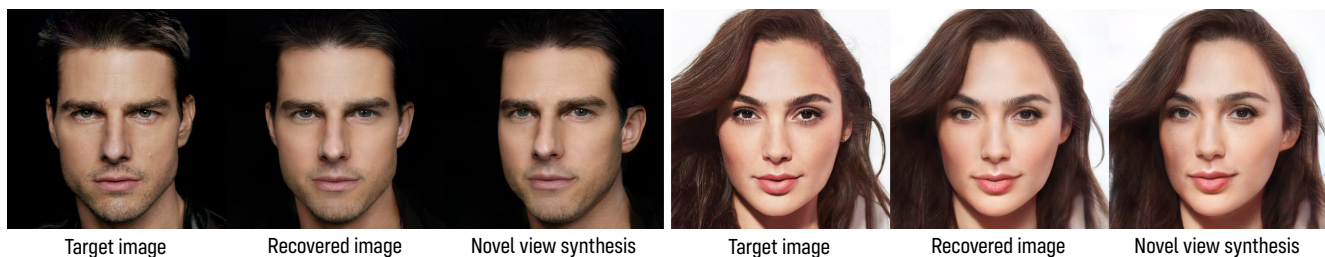


Figure 7: High-resolution (1024<sup>2</sup>) image embedding and editing results (**Best viewed with zoom-in**)

### 4.3. Ablation Study

We further conduct ablation studies to validate the efficacy of our super-resolution CNN architecture design and the loss functions for high-resolution training. Unless otherwise specified, the experiments are conducted on FFHQ with 256<sup>2</sup> resolution for efficiency.

**Network architecture** The first 5 rows of Table 5 shows the FID scores with different architectures we tested for the super-resolution CNN. Here the base architecture is the original network of ESRGAN [59], which we found to perform better than a decoder structure of StyleGAN2 [24]. We then add different components, including subpixel convolutions for upsampling, style modulation, intermediate feature from LR radiance generator, and a separate CNN for background radiance, which all lead to lower FID scores.

**Loss functions** We verify the efficacy of the two new loss functions  $\mathcal{L}_{\text{patch}}$  and  $\mathcal{L}_{\text{cons}}$  with both numerical metrics and visual quality. The results are presented in Table 5 (last two rows) and Figure 6, respectively. We find that adding the patch discriminator and applying  $\mathcal{L}_{\text{patch}}$  lead to lower FID score at low resolution (256<sup>2</sup> in Table 5), and can effectively eliminate the checkerboard artifact at high resolution (1024<sup>2</sup> in Figure 6). Without the cross-resolution consistency loss  $\mathcal{L}_{\text{cons}}$ , the FID score increases significantly and the super-resolution CNN generates some floaters in the front as shown in Figure 6, which might be caused by less stable training.

### 4.4. Applications

**High resolution image embedding and editing** Like other GAN methods, one can embed an image into the latent space of our trained GRAM-HD and achieve pose editing by rendering images at novel views. As shown in Figure 7, our method can faithfully reconstruct a high-resolution image through GAN inversion and generate high-fidelity novel view rendering results.

**Real-time free-view synthesis** Similar to GRAM [12], we can cache the manifold surfaces and HR radiance maps as textured 3D meshes and then run fast free-view synthesis with mesh rendering. With an efficient mesh rasterizer from [25], we achieve free-view synthesis of 1024<sup>2</sup> images at 90 FPS on a Nvidia Tesla V100 GPU.

## 5. Conclusion

We have presented a novel GAN approach for high-resolution, 3D-consistent multiview image generation, which is trained on unstructured single image collections. Our key idea is to tackle a 3D super-resolution task using efficient 2D CNNs by leveraging the recent radiance manifold representation. The experiments show that our generation results at high resolution (*e.g.*, 1024<sup>2</sup>) are not only of high quality but also strongly 3D consistent, significantly outperforming recent 3D-ware GANs. We believe our method bring 3D-consistent image generation closer to the traditional 2D GANs, and it paves the way for high-quality 3D content creation applications such as 3D video generation



and animation.

**Limitations and future work** Our method still have several limitations. Based on the radiance manifold representation, it is difficult to handle objects with complex 3D geometry. Its view extrapolation capability is also limited compared to dense 3D radiance fields. Besides, the generation quality of our method still lags behind traditional 2D image generative models. Better representations or training strategies could be further explored to close the gap.

**Ethics consideration** The goal of this paper is to generate images of virtual objects for applications such as photorealistic virtual avatar creation. However, it could be misused to create misleading or harmful contents. We condemn any behavior to use our method for fraud purposes. On the positive side, our method could be used for forgery detection system testing and for generating privacy-free contents. The face dataset we use may contain certain biases in race, gender, etc., which can be inherited by our trained model and results. Bias-free generative modeling is important direction worth further exploration by the community.

## References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *IEEE/CVF International Conference on Computer Vision*, 2021. [2](#)
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. [2](#)
- [3] Mikolaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. In *International Conference on Learning Representations*, 2018. [7](#)
- [4] Robert C Bolles, H Harlyn Baker, and David H Marimont. Epipolar-plane image analysis: An approach to determining structure from motion. *International Journal of Computer Vision*, 1(1):7–55, 1987. [6](#), [7](#)
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019. [1](#)
- [6] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks). In *IEEE International Conference on Computer Vision*, pages 1021–1030, 2017. [5](#)
- [7] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *IEEE/CVF International Conference on Computer Vision*, 2022. [2](#), [3](#), [6](#), [7](#)
- [8] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5799–5809, 2021. [1](#), [2](#)
- [9] Hong Chang, Dit-Yan Yeung, and Yimin Xiong. Super-resolution through neighbor embedding. In *2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. IEEE, 2004. [3](#)
- [10] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8628–8638, 2021. [3](#)
- [11] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020. [2](#), [5](#)
- [12] Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. Gram: Generative radiance manifolds for 3d-aware image generation. In *IEEE/CVF International Conference on Computer Vision*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [13] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2015. [3](#)
- [14] Alexey Dosovitskiy, Jost Tobias Springenberg, Maxim Tatarchenko, and Thomas Brox. Learning to generate chairs, tables and cars with convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):692–705, 2016. [2](#)
- [15] SM Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S Morcos, Marta Garnelo, Avraham Ruderman, Andrei A Rusu, Ivo Danihelka, Karol Gregor, et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018. [2](#)
- [16] Daniel Glasner, Shai Bagon, and Michal Irani. Super-resolution from a single image. In *2009 IEEE 12th International Conference on Computer Vision*, pages 349–356. IEEE, 2009. [3](#)
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2014. [1](#), [4](#)
- [18] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. In *International Conference on Learning Representations*, 2021. [2](#), [3](#), [6](#), [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 *IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021. [2](#)
- [20] 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. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017. [7](#)
- [21] Wobong Jang and Lourdes Agapito. Codenerf: Disentangled neural radiance fields for object categories. In *IEEE/CVF International Conference on Computer Vision*, pages 12949–12958, 2021. [2](#)

- [22] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Advances in Neural Information Processing Systems*, 2021. 1, 7
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 2, 5
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 1, 4, 8
- [25] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics*, 39(6):1–14, 2020. 8
- [26] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4681–4690, 2017. 3
- [27] Taehee Brad Lee. Cat hipsterizer. [https://github.com/kairess/cat\\_hipsterizer](https://github.com/kairess/cat_hipsterizer), 2018. 5
- [28] Yiyi Liao, Katja Schwarz, Lars Mescheder, and Andreas Geiger. Towards unsupervised learning of generative models for 3d controllable image synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5871–5880, 2020. 2
- [29] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 136–144, 2017. 3
- [30] David B Lindell, Julien NP Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14556–14565, 2021. 2
- [31] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *Advances in Neural Information Processing Systems*, 2020. 2
- [32] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *IEEE/CVF International Conference on Computer Vision*, pages 5773–5783, 2021. 2
- [33] Julien N. P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. *ACM Trans. Graph. (SIGGRAPH)*, 40(4), 2021. 2
- [34] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 2
- [35] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International Conference on Machine Learning*, pages 3481–3490, 2018. 4
- [36] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 2
- [37] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020. 2, 3
- [38] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. 2
- [39] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3D representations from natural images. In *IEEE/CVF International Conference on Computer Vision*, pages 7588–7597, 2019. 2
- [40] Thu Nguyen-Phuoc, Christian Richardt, Long Mai, Yong-Liang Yang, and Niloy Mitra. BlockGAN: Learning 3d object-aware scene representations from unlabelled images. In *Advances in Neural Information Processing Systems*, 2020. 2
- [41] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021. 1, 2, 3
- [42] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. 2
- [43] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *IEEE/CVF International Conference on Computer Vision*, 2021. 2, 3
- [44] 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 *IEEE/CVF International Conference on Computer Vision*, 2022. 2, 3, 6, 7
- [45] Xingang Pan, Xudong Xu, Chen Change Loy, Christian Theobalt, and Bo Dai. A shading-guided generative implicit model for shape-accurate 3d-aware image synthesis. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [46] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 2
- [47] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In *Advances in Neural Information Processing Systems*, 2020. 1, 2
- [48] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan

- Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016. 4
- [49] Yichun Shi, Divyansh Aggarwal, and Anil K Jain. Lifting 2d stylegan for 3d-aware face generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6258–6266, 2021. 2
- [50] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33, 2020. 2
- [51] Vincent Sitzmann, Semon Rezchikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [52] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019.
- [53] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, pages 1121–1132, 2019. 2
- [54] Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. EpiGRAF: Rethinking training of 3d GANs. In *Advances in Neural Information Processing Systems*, 2022. 1, 2, 6, 7
- [55] Attila Szabó, Givi Meishvili, and Paolo Favaro. Unsupervised generative 3d shape learning from natural images. *arXiv preprint arXiv:1910.00287*, 2019. 2
- [56] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Multi-view 3d models from single images with a convolutional network. In *European Conference on Computer Vision*, pages 322–337, 2016. 2
- [57] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 2021. 2, 7
- [58] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *IEEE/CVF International Conference on Computer Vision*, pages 1905–1914, 2021. 3
- [59] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pages 0–0, 2018. 3, 4, 8
- [60] Xudong Xu, Xingang Pan, Dahua Lin, and Bo Dai. Generative occupancy fields for 3d surface-aware image synthesis. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [61] Yinghao Xu, Sida Peng, Ceyuan Yang, Yujun Shen, and Bolei Zhou. 3d-aware image synthesis via learning structural and textural representations. *arXiv preprint arXiv:2112.10759*, 2021. 3
- [62] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [63] Kai Zhang, Luc Van Gool, and Radu Timofte. Deep unfolding network for image super-resolution. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3217–3226, 2020. 3
- [64] Xuanmeng Zhang, Zhedong Zheng, Daiheng Gao, Bang Zhang, Pan Pan, and Yi Yang. Multi-view consistent generative adversarial networks for 3d-aware image synthesis. In *IEEE/CVF International Conference on Computer Vision*, 2022. 2, 3
- [65] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *European Conference on Computer Vision*, pages 286–301, 2018. 3
- [66] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2472–2481, 2018. 4
- [67] Xiaoming Zhao, Fangchang Ma, David Güera, Zhile Ren, Alexander G. Schwing, and Alex Colburn. Generative multiplane images: Making a 2d gan 3d-aware. In *Proc. ECCV*, 2022. 1
- [68] 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. 3
- [69] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: learning view synthesis using multiplane images. *ACM Transactions on Graphics*, 37(4):1–12, 2018. 3
- [70] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision*, pages 2223–2232, 2017. 5