

# NeRF-GAN Distillation for Efficient 3D-Aware Generation with Convolutions

Mohamad Shahbazi<sup>1</sup>    Evangelos Ntavelis<sup>1,3</sup>    Alessio Tonioni<sup>2</sup>    Edo Collins<sup>2</sup>  
 Danda Pani Paudel<sup>1</sup>    Martin Danelljan<sup>1</sup>    Luc Van Gool<sup>1</sup>

<sup>1</sup>Computer Vision Lab, ETH Zürich    <sup>2</sup>Google Zürich    <sup>3</sup>ML & Robotics, CSEM, Switzerland

{mshahbazi, entavelis, paudel, martin.danelljan, vangool}@vision.ee.ethz.ch,

{alessiot, edocollins}@google.com

## Abstract

*Pose-conditioned convolutional generative models struggle with high-quality 3D-consistent image generation from single-view datasets, due to their lack of sufficient 3D priors. Recently, the integration of Neural Radiance Fields (NeRFs) and generative models, such as Generative Adversarial Networks (GANs), has transformed 3D-aware generation from single-view images. NeRF-GANs exploit the strong inductive bias of neural 3D representations and volumetric rendering at the cost of higher computational complexity. This study aims at revisiting pose-conditioned 2D GANs for efficient 3D-aware generation at inference time by distilling 3D knowledge from pretrained NeRF-GANs. We propose a simple and effective method, based on re-using the well-disentangled latent space of a pre-trained NeRF-GAN in a pose-conditioned convolutional network to directly generate 3D-consistent images corresponding to the underlying 3D representations. Experiments on several datasets demonstrate that the proposed method obtains results comparable with volumetric rendering in terms of quality and 3D consistency while benefiting from the computational advantage of convolutional networks. The code is available at: <https://github.com/mshahbazi72/NeRF-GAN-Distillation>*

## 1. Introduction

Generative Adversarial Networks (GANs) [14] have undergone outstanding progress in photo-realistic image generation and manipulation in a variety of applications [3, 26, 39, 21, 22, 8, 43, 42]. Recently, there has been an increasing interest in extending GANs to the task of 3D-aware generation from single-view image datasets, with the goal of providing disentangled control over the content and the view-point of the generated images.

Volumetric Rendering (EG3D)



Our Convolutional Rendering

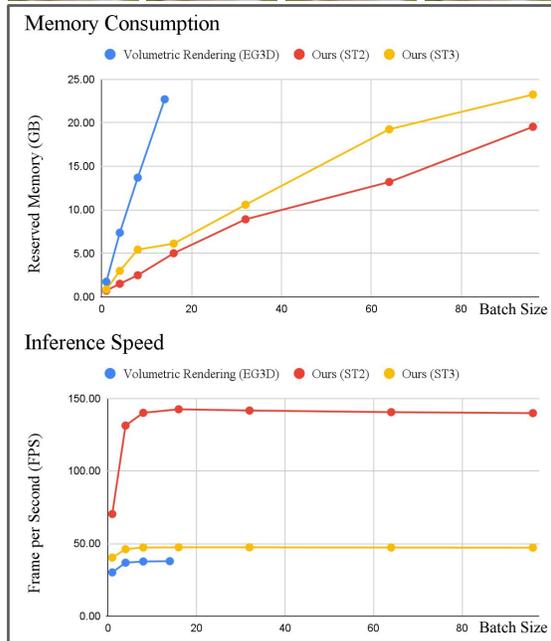


Figure 1. Top: views of the same subject cat generated by a volumetric rendering generator (EG3D) and by our convolutional generator. Bottom: comparison of the inference memory consumption and speed (on a fixed GPU budget) for the two methods.

Image GAN models have been historically based on convolutional architectures, enabling efficient training and generation for 2D tasks. However, pose-conditioned convolutional GANs (pcGANs) struggle with 3D-consistent image generation, due to their lack of sufficient 3D priors [34]. Therefore, some studies have previously attempted to disentangle the pose from the content in pcGANs using explicit 3D supervision [50, 27, 11], which, however, is not readily available for most datasets. As a result, later methods moved away from fully convolutional GANs by incorporating 3D inductive biases in the architecture and training pipeline, such as 3D neural representations and differentiable rendering methods [34, 35, 46, 38].

The advent of Neural Radiance Fields (NeRFs) [30] has recently transformed the neural 3D representation and the task of novel-view synthesis [1, 2, 41, 31, 29, 16, 54]. For this reason, NeRFs have been successfully integrated with GANs to achieve promising results in 3D-aware generation [44, 5, 37, 15, 48, 36, 45]. Nerf-GANs, in their general form, map a latent space to a 3D representation of objects and generate images from queried viewpoints using volumetric rendering. However, volumetric rendering is computationally demanding due to its ray-casting process, making high-resolution generation slow and memory-expensive. Recent works have proposed different approaches to improve the computational efficiency of NeRF-GANs using more efficient 3D representations [36, 5, 45] and training protocols [5, 37]. Nevertheless, volumetric rendering remains an integral part of these models.

In recent NeRF-GANs [36, 5, 45, 37], convolutional networks have been reintroduced in the generator architecture as super-resolution networks or as 3D-representation generators, in order to scale up NeRF-GANs for high-resolution generation. In this study, we take a different approach to integrating NeRF-GANs and convolutional GANs for 3D-aware generation from single-view images. In particular, we investigate the capacity of convolutional generators to achieve 3D-consistent rendering with explicit pose control when learning from a pretrained NeRF-GAN without any additional explicit 3D supervision. A convolutional generator that fairly preserves the 3D consistency, image quality, and the correspondence between the generated images and the underlying 3D representation can be used for efficient multi-view inference in setups where volumetric rendering is not affordable, such as in mobile applications. However, balancing and minimizing the trade-off between efficiency and 3D consistency is a highly challenging task, which we set out to explore in this work.

We propose a simple but effective method for distilling a pretrained NeRF-GAN into a pose-conditioned fully convolutional generator. The main component of our approach is based on exploiting the well-disentangled intermediate latent space of the NeRF-GAN in the convolutional gen-

erator. In particular, our convolutional generator learns to map each latent code from the 3D generator, along with the target viewpoint, to the corresponding obtained images by explicit volumetric rendering. By doing so, we aim to distill the NeRF-GAN’s underlying 3D knowledge into the convolutional generator, as well as to establish a correspondence between the images of the generator and the 3D representation of the NeRF-GAN. As demonstrated in Fig. 1, our experiments on three different datasets indicate that the convolutional generator trained with our method is capable of achieving results comparable to volumetric rendering in terms of image quality and 3D-consistency, while benefiting from the superior efficiency of convolutional networks.

Our contributions are summarized as follows:

- We propose a method to distill NeRF-GANs into convolutional generators for efficient 3D-aware inference.
- We provide a simple and effective method to condition the convolutional generator on the well-disentangled intermediate latent space of the NeRF-GAN.
- Through experiments on three different datasets, we show that the generator trained by our distillation method well preserves the 3D consistency, image quality, and semantics of the pretrained NeRF-GAN.

## 2. Related Works

**3D-aware Generation from Single-View Images.** Prior works have attempted to create 3D awareness in 2D GANs using explicit 3D supervision, such as 3D models [50, 11], pose and landmark annotations [47, 20], and synthetic data [27]. In many applications, obtaining such 3D supervision is not practical. As a result, later works aimed at unsupervised methods by introducing 3D inductive biases in GANs, including 3D neural representations and differentiable rendering [34, 38, 46, 35]. These methods, although promising, lag far behind 2D GANs in terms of image quality or struggle with high-resolution generation due to the additional computational complexity.

**NeRF-GANs.** NeRFs have shown outstanding potential in compactly representing 3D scenes for novel view synthesis. GRAF [44] and Pi-GAN [5] are the first works to integrate NeRFs and GANs. While achieving highly consistent 3D-aware generation, the computational restrictions of NeRF framework make these methods impractical for high-resolution generation or environments with constrained resources. In order to extend NeRF-GANs to higher resolutions, convolutional super-resolution networks were used in later studies [36, 15, 37] at the expense of some multi-view inconsistencies. EpiGRAF [48], in contrast, adopts an efficient multi-scale patch training protocol, but still requires full high-resolution sampling rendering for inference, which makes it comparatively more computationally demanding than competitors.

Other studies aim at bringing the recent advances in the efficiency of NeRFs to NeRF-GANs. Although there exist numerous works on efficient 3D representations [32, 41, 49, 7, 53, 52] and volumetric sampling [55, 13, 17, 33, 19] in NeRFs, only a subset of them [4, 45, 56] have been successfully applied to NeRF-GANs. This is because they are mainly designed for the single-scene setup, making their adaptation to the generative setup not trivial. The use of sparse voxel grids in VoxGRAF [45] and multi-plane image representations in [56] result in efficient and 3D-consistent generation while compromising the image quality and 3D geometry. EG3D [4] proposes using tri-planes to represent the geometry of the generated objects. Exploiting tri-planes, coupled with carefully designed techniques to enforce 3D consistency, allows EG3D to significantly improve both computational efficiency and image quality. Live 3D Portrait [51] is a concurrent work based on EG3D that aims at real-time one-shot reconstruction of faces by estimating the canonical tri-planes of a pre-trained EG3D. However, Live 3D Portrait is computationally limited by the underlying volumetric rendering of EG3D. Most similar to our study, SURF-GAN [28] aims to discover directions for pose control in a pretrained 2D GAN by generating multi-view images using a pretrained NeRF-GAN. However, using NeRF-GANs only as multi-view supervision does not fully exploit their underlying 3D knowledge. Moreover, the 2D generator obtained using this method does not preserve any correspondence between the NeRF-GAN’s 3D representations and the generated images. Different from SURF-GAN, we exploit the intermediate latent space of retrained NeRF-GANs to distill 3D knowledge into a 2D generator and establish correspondence between the convolutional generator and the NeRF-GAN’s 3D representations.

### 3. Method

In this section, we first provide a brief overview of the formulation of NeRF-GANs and then explain the proposed formulation in detail.

#### 3.1. Preliminaries

The general formulation of NeRF-GANs consists of a 3D-representation generator  $G^{3D}(z)$ , which maps a latent variable  $z$  (usually drawn from a normal distribution) to a 3D representation of an object. Then, in order to render an image  $I_{z,c}$  from the target viewpoint (camera parameters)  $c \in R^{25}$ , volumetric rendering is applied to the generated 3D representation. We base our method on EG3D [4], as it provides a strong trade-off in image quality, 3D consistency, and efficiency, among recent NeRF-GANs.

EG3D represents 3D scenes using tri-planes, which are three axis-aligned orthogonal feature planes, each with a size of  $N \times N \times C$ , where  $N$  is spatial resolution and  $C$  is the number of channels. To represent a 3D position  $x \in R^3$ ,  $x$

is projected onto each of the three feature planes, retrieving the corresponding feature vector  $(F_{xy}, F_{xz}, F_{yz})$  via bilinear interpolation, and aggregating the three feature vectors via summation. To obtain the color and density at position  $x$ , a lightweight MLP decodes the feature vector obtained for the queried position to a density and color value.

The tri-plane generator  $G^{3D}(z, c)$  in EG3D consists of a mapping network  $M^{3D}(z, c)$ , which maps the input latent code and the target viewpoint to an intermediate latent variable  $w$ , namely the style code. The style code then is used to modulate a convolutional synthesis network  $S^{3D}(w)$  to generate the tri-planes  $I_{z,c}^{3p}$ . In order to render an image  $I_{z,c}$  from the target viewpoint  $c$ , hierarchical volumetric rendering is applied to the tri-planes. Since volumetric rendering at high resolutions is computationally too expensive, EG3D does so at a lower resolution and uses a convolutional super-resolution network to obtain a final image. More specifically, the low-resolution output of volumetric rendering in EG3D consists of a 32-channel feature map  $I_{z,c}^f$ , the first three of which represent the low-resolution RGB image  $I_{z,c}^{LR}$ , which is given as input to the super-resolution network. EG3D is trained in an adversarial fashion with a viewpoint-conditioned dual discriminator  $D$  that ensures the photorealism of the generated images from the target viewpoints, as well as the consistency between high-resolution and low-resolution images.

#### 3.2. Convolutional Rendering of Pretrained NeRF-GANs

The aim of the proposed method is to distill a pre-trained NeRF-GAN  $G_{z,c}^{3D}$  into a 2D image generator  $G_{z,c}^{2D}$ , such that  $G_{z,c}^{2D}$  directly predicts 3D-consistent multi-view images  $I'_{z,c}$ , corresponding to the volumetric renderings obtained by the underlying 3D representation of  $G_{z,c}^{3D}$ . To this end, we propose to exploit the well-disentangled style space of  $G^{3D}$  to distill the underlying 3D representation into  $G^{2D}$ . Sharing the style space  $w$  of the pre-trained 3D generator with the convolutional renderer is the first step towards establishing a correspondence between the 3D representations and the generated images. Secondly, it allows training the convolutional generator for 3D-consistent generation without the need for generating multiple views of the same objects and enforcing multi-view consistency.

The overall architecture of EG3D and our convolutional generator is visualized in 2. The convolutional generator is based on StyleGAN architecture [23, 24], consisting of a mapping network, a low-resolution convolutional feature prediction, and a convolutional super-resolution network. The mapping network transforms the style code  $w$  of  $G^{3D}$ , and the target viewpoint  $c$  to the style code  $w'$  of  $G^{2D}$ . Then, the low-resolution feature predictor  $S^{2D}$  estimates the EG3D features and low-resolution image obtained by the volumetric rendering. Estimated features and images are

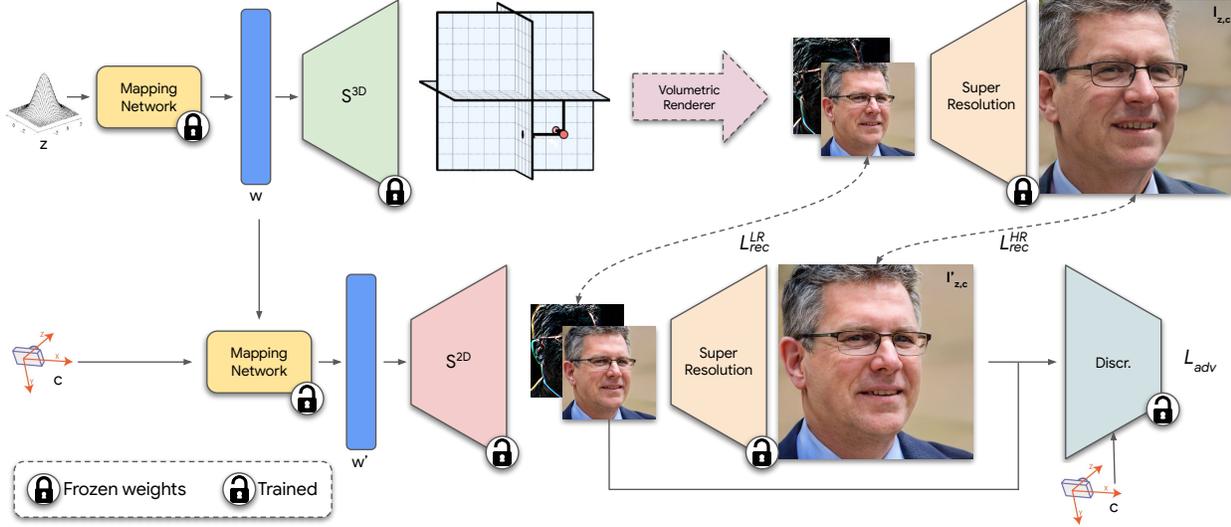


Figure 2. Using a student-teacher framework, we distill 3D consistency from a frozen volumetric rendering based Nerf-GAN (top) to a 2D convolutional renderer (bottom). A loss consisting of high- and low-resolution image reconstruction and an adversarial component allows us to retain good image quality and 3D consistency.

then mapped to the high-resolution outputs using the super-resolution network. In our setup, the super-resolution network is initialized with EG3D’s super-resolution network and is jointly optimized with the feature predictor network.

### 3.3. Training

In order to train the proposed convolutional renderer, we use a teacher-student framework, where the volumetric rendering is used to supervise  $G^{2D}$  on the viewpoint-conditioned mapping of  $(z, c)$  to  $I'_{z,c}$ . A schematic representation of our training regime is reported in Fig. 2. Specifically, for each training sample, we randomly sample  $z_i$  and  $c_i$  and use the pretrained NeRF-GAN to obtain: the corresponding style code  $w_i$ , the low-resolution image  $I_{z_i, c_i}^{LR}$ , feature maps  $I_{z_i, c_i}^f$  rendered by volumetric rendering, as well as, the high-resolution image  $I_{z_i, c_i}^{HR}$  generated by the super-resolution network. These together form a training sample  $i$  for the proposed convolutional renderer.

We provide  $z$  and  $c$  to  $G^{2D}$  and compute a loss function composed of three parts. We first add a reconstruction term  $L_{rec}^{LR}$  between the low-resolution outputs of the volumetric and convolutional renderers. A second reconstruction loss  $L_{rec}^{HR}$  is applied between the super-resolved outputs of the two renderers. Lastly, we apply an adversarial term  $L_{adv}$ .

In the following, we drop the subscripts  $z_i, c_i$  to reduce the clutter in notation. The low-resolution reconstruction loss  $L_{rec}^{LR}$  consists of a pixel-wise smooth L1 loss between the two feature maps, as well as a perceptual loss between the generated and target low-resolution images,

$$L_{rec}^{LR} = \lambda_{L1}^{LR} * \text{SmoothL1}(I^f, I^f) + \lambda_{perc}^{LR} * \text{PerceptualLoss}(I^{LR}, I^{LR}). \quad (1)$$

Here,  $\lambda_{L1}^{LR}$  and  $\lambda_{perc}^{LR}$  are the weights for the low-resolution smooth L1 and perceptual loss, respectively.

Similarly, the high-resolution reconstruction loss  $L_{rec}^{HR}$  is defined as:

$$L_{rec}^{HR} = \lambda_{L1}^{HR} * \text{SmoothL1}(I'^f, I^f) + \lambda_{perc}^{HR} * \text{PerceptualLoss}(I'^{HR}, I^{HR}). \quad (2)$$

where  $\lambda_{L1}^{HR}$  and  $\lambda_{perc}^{HR}$  are the weights for the high-resolution smooth L1 and perceptual loss, respectively.

The adversarial term  $L_{adv}$  is similar to the one used in EG3D. We use the same dual-discriminator architecture  $D$  as in EG3D to ensure the realism of the high-resolution images, their consistency with the low-resolution version, and the compliance of the generated image with the queried viewpoints. The total loss  $L_{total}$  for training  $G^{2D}$  is,

$$L_{total} = L_{rec}^{LR} + L_{rec}^{HR} + \lambda_{adv} * L_{adv}, \quad (3)$$

where  $\lambda_{adv}$  is the weight for the adversarial loss.

**Two-stage training:** in practice, empirical experiments show that training the convolutional renderer using the full objective from the beginning will lead to high-quality but 3D-inconsistent images. Therefore, we instead propose a 2-stage training curriculum. In the first stage,  $G^{2D}$  is only optimized by pure distillation of the volumetric rendering  $G^{3D}$  using  $L_{rec}^{LR}$  and  $L_{rec}^{HR}$  until the renderer achieves reasonable generation quality. Then, the adversarial loss  $L_{adv}$  is added to the training to further improve the performance. By applying this 2-stage curriculum, we are able to counter the 3D inconsistency induced by the adversarial training.

**Pose-correlated dataset bias:** As shown in EG3D [4], adversarial training of the convolutional network is prone to

learning pose-correlated dataset biases, such as more smiling in non-frontal viewpoints in FFHQ dataset [25], which in turn results in 3D attribute inconsistencies. To mitigate such biases in FFHQ dataset, we use both real images and the images rendered from EG3D as the real examples shown to the discriminator. The proportion of the EG3D-rendered images shown to  $D$  is controlled by the hyper-parameter  $\alpha$  ( $0 \leq \alpha \leq 1$ ), which is set to 0.5 in our experiments. As we will discuss in section 4.6,  $\alpha$  can be used to control the trade-off between image quality and 3D consistency.

## 4. Experiments

In this section, we first describe our experimental setup for the evaluation of our method. Then, we compare the proposed method with baselines in terms of visual quality, 3D consistency, and computational efficiency. Moreover, we provide an ablation study and a discussion on the benefits and trade-offs of the proposed method.

### 4.1. Datasets

Following EG3D [5], we evaluate our method on three different datasets:

**Flickr-Faces-HQ (FFHQ)** [25]: a collection of 70k high-quality images of real-world human faces, as well as corresponding approximate camera extrinsics estimated using an off-the-shelf pose estimator.

**AFHQ Cats**: a sub-category of the Animal-Face-HQ (AFHQ) [9], consisting of around 5k high-quality images of cat faces, as well as corresponding camera extrinsics estimated using an off-the-shelf pose estimator.

**ShapeNet Cars**: a category of ShapeNet [6] consisting of synthetic images of cars rendered from different viewpoints, as well as the corresponding camera extrinsics annotations.

### 4.2. Baselines

We consider EG3D [4] and the method proposed in SURF-GAN [28] as our main baselines for this study. For a more complete evaluation, we also compare our method to additional relevant baselines:

**EG3D** [4]: the NeRF-GAN used for distilling 3D knowledge in the convolutional generator. EG3D serves as the upper bound for the 3D consistency of the proposed method.

**Pose-Conditioned StyleGAN (PC-GAN)**: a standard conditional 2D GAN, conditioned on the pose annotations without any knowledge distillation.

**SURF**: Inspired by the proposed method in SURF-GAN [28], we create a baseline called SURF, where multi-view images of EG3D are used to discover pose-control in a pre-trained 2D StyleGAN.

**LiftGAN** [46]: a method predating EG3D and SURF baselines based on differentiable rendering that distills a 2D GAN in order to train a 3D generator.

### 4.3. Implementation and Evaluation Details

We implement and evaluate the proposed generator using both StyleGAN2 (ST2) [23] and StyleGAN3 (ST3) [24] architectures. For the pretrained NeRF-GAN, we use the official models from EG3D [4] (for Shapenet Cars, we re-train the model as the official model does not match the results reported by EG3D). We train our experiments using a batch size of 16. The rendering resolution and the final resolution are (128, 512) for FFHQ and AFHQ and (64, 128) for Shapenet Cars. Both training and inference experiments were conducted using NVidia RTX 3090 GPUs. In all experiments, we set all of the weights of reconstruction loss terms ( $\lambda_{L1}^{LR}$ ,  $\lambda_{perc}^{LR}$ ,  $\lambda_{L1}^{HR}$ ,  $\lambda_{perc}^{HR}$ ) to the value 1 and the weight of the adversarial loss ( $\lambda_{adv}$ ) to the value 0.1.

### 4.4. Metrics

We evaluate our method quantitatively in terms of visual quality, and 3D consistency.

**Fréchet Inception Distance (FID)** [18]: The most common metric to assess the quality and diversity of generation.

**Kernel Inception Distance (KID)** [18]: An unbiased alternative to FID for smaller datasets.

**Pose Accuracy (PA)**: following previous works [4], we measure the ability of the model in generating images of the query poses by calculating the mean squared error (MSE) between the query poses and the pose of the generated images, estimated using an off-the-shelf pose estimator [12].

**Identity Preservation (ID)**: As a metric for 3D consistency, we measure the degree of face identity preservation between different viewpoints with respect to the canonical pose using ArcFace [10] cosine similarity for the FFHQ setup.

**3D Landmark Consistency**: As another 3D consistency metric, we measure the change in facial landmarks between different viewpoints in FFHQ using MSE. The 3D landmarks are estimated using an off-the-shelf estimator [12].

### 4.5. Quantitative Comparison

In the following, we quantitatively compare the proposed method with the baselines described in Sec. 4.2 in terms of inference efficiency, visual quality, and 3D consistency.

#### 4.5.1 Efficiency

The efficiency of fully-convolutional networks compared to the rendering-based methods is well-known. To better assess the practical benefit of the proposed method, we provide a comparison of inference efficiency between EG3D. Fig 1 visualizes an example of the inference memory consumption and speed of the two methods using different batch sizes on a fixed GPU budget (in this case, on RTX 3090 GPU with 24G of memory). As shown, EG3D is restricted to small batch sizes (a maximum of 14) due to its costly memory consumption, where our method can scale

Table 1. Comparison of image quality on three datasets in terms of FID and KID metrics. \*The value is borrowed from [46].

Method	FFHQ		AFHQ		ShapeNET Cars	
	FID ↓	KID ↓	FID ↓	KID ↓	FID ↓	KID ↓
EG3D [4]	5.0	0.0018	2.9	0.0003	3.5	0.0017
PC-GAN	19.3	0.0085	4.5	0.0009	6.1	0.0018
LiftGAN [46]	29.8*	-	-	-	-	-
SURF	31.1	0.0153	-	-	-	-
Ours (ST2)	6.6	0.0019	3.8	0.0011	3.1	0.0013
Ours (ST3)	6.8	0.0023	3.2	0.0007	3.1	0.0012

Table 2. Comparison of 3D consistency metrics on FFHQ.

Method	Pose Acc. ↓	3D Landmark ↓	ID ↑
EG3D [4]	0.002	0.018	0.75
PC-GAN	0.009	0.062	0.56
SURF	0.044	0.014	0.86
Ours (ST2)	0.002	0.023	0.75
Ours (ST3)	0.002	0.022	0.75

up to a maximum of 96 samples per batch ( $\sim 7\times$ ). As for the speed, our generator achieves a better frame-per-second, especially when using StyleGAN2 as its backbone ( $> 3\times$ ).

#### 4.5.2 Image Quality

To assess the trade-off brought about by our convolutional generator, we evaluate the quality of the generated images. Table 1 shows the FID and KID scores for our method and the baselines on different datasets. Compared to the PC-GAN and SURF baselines, our method constantly achieves higher quality. This confirms that exploiting the style space of the pretrained NeRF-GAN contributes to the ability of the convolutional renderer in pose-conditioned generation. Although our method does not fully match the visual quality of EG3D, it is still able to fairly maintain high image quality and significantly reduce the compromise in the quality compared to the other convolutional counterparts.

#### 4.5.3 3D Consistency

While, unlike volumetric rendering, 2D convolutions do not guarantee 3D consistency, we show that our approach achieves a good performance in this regard. We assess the 3D consistency of generated images on FFHQ by measuring the pose accuracy, 3D landmark consistency, and face identity preservation, as discussed in Sec. 4.4, which are provided in Table 2. Based on the results, Our method achieves comparable 3D consistency with EG3D, while PC-GAN and SURF struggle. Note that the high values for identity preservation and 3D landmark consistency in SURF are due to the limited pose variations, and hence generating similar images regardless of the input pose, as reflected by the pose accuracy (and the visual examples in Fig. 3).

## 4.6. Ablation

**Ablation on loss functions:** the proposed training objective in section 3.3 consists of different loss terms to ensure both image quality and consistency with the output of volumetric rendering. In this section, we ablate the importance of these components. Table 3 shows the FID scores for the following experiments on the loss terms on AFHQ dataset:

- **LR:** only the low-resolution reconstruction loss. the super-resolution network is frozen.
- **HR:** only the high-resolution reconstruction loss.
- **LR + HR** both low-resolution and high-resolution reconstruction losses.
- **HR + ADV:** reconstruction and adversarial losses on the high-resolution images.
- **Full (LR + HR + ADV)** full training objective, including the reconstruction and adversarial terms.

As shown by the ablation study, the combination of all the proposed loss terms leads to the best FID scores.

**Single-stage v.s. two-stage training:** As mentioned in 3.3, we find out that single-stage training by jointly optimizing for both reconstruction and adversarial losses results in subtle inconsistencies such as color shifts and geometry warps, which can be mitigated using the proposed 2-stage training in section 3. As the observed inconsistencies are difficult to capture using our quantitative 3D consistency metrics, we provide a visual comparison between the examples of single-stage and two-stage training on AFHQ in Fig. 6.

**Mitigating Pose-Attribute Correlation:** In Table 4, we provide an ablation on the parameter  $\alpha$  introduced in section 3.3 for FFHQ dataset. As shown, including EG3D-generated images ( $\alpha > 0$ ) improves the 3D consistency at the cost of a lower generation quality.

Table 3. Ablation on different loss functions for training the ST3 convolutional renderer on AFHQ Cats dataset.

Method	FID ↓
LR	30.55
HR	10.39
LR + HR	9.1
HR + ADV	6.58
Full (LR + HR + ADV)	3.2

Table 4. The effect of mixing real images and EG3D-rendered images as real examples for adversarial training, controlled by the parameter  $\alpha$ , on FFHQ dataset.

Method	Ours (ST2)		Ours (ST3)	
	FID ↓	3D Landmark ↓	FID ↓	3D Landmark ↓
$\alpha = 0$	5.5	0.027	5.7	0.027
$\alpha = 0.3$	6.3	0.023	6.1	0.024
$\alpha = 0.5$	6.8	0.022	6.6	0.023

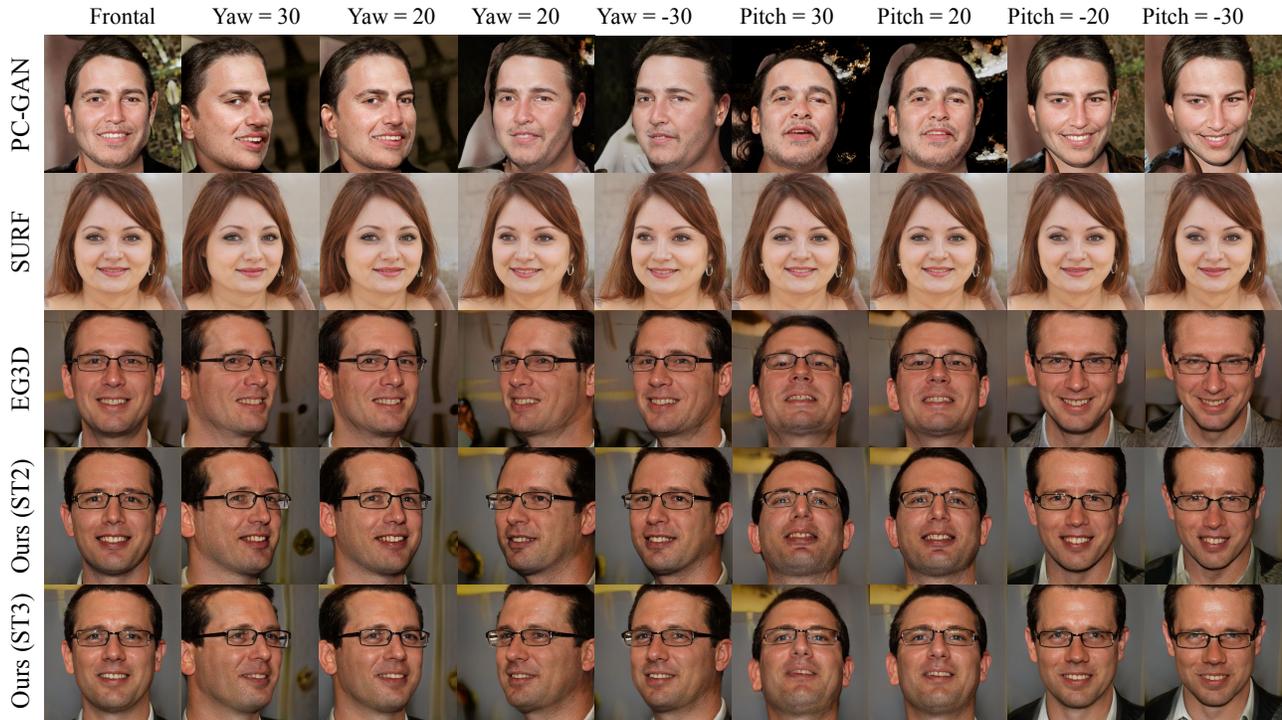


Figure 3. Qualitative examples of variations in yaw and pitch for FFHQ. Compared to the pose-conditioned GAN and SURF baseline, our proposed method nearly matches the 3D consistency and image quality of volumetric rendering (EG3D).

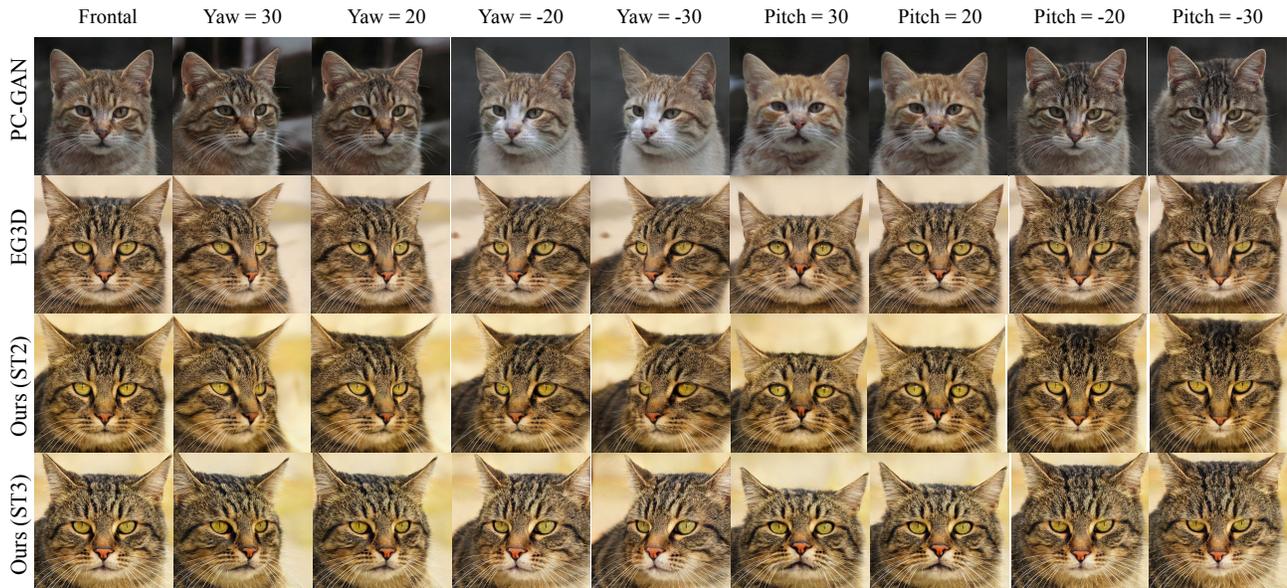


Figure 4. Qualitative examples of yaw and pitch variations for AFHQ cats. In line with our quantitative experiments, the pose-conditioned convolutional baseline (PC-GAN) fails to preserve the identity of the subject under different poses. In contrast, our method exhibits similar preservation of identity to the volume rendering approach (EG3D), despite the difference in computational resources and time.

#### 4.7. Qualitative Comparison

In this section, We provide a visual comparison of our method with the baselines. In Fig. 3 and 4, we provide visual examples of variations in yaw and pitch for FFHQ

and AFHQ Cats. Compared to a PC-GAN and SURF, our proposed method closely matches the 3D consistency and maintains the image quality of volumetric rendering. Fig. 5 additionally provides examples of Shapenet Cars generated using our method and their corresponding images from

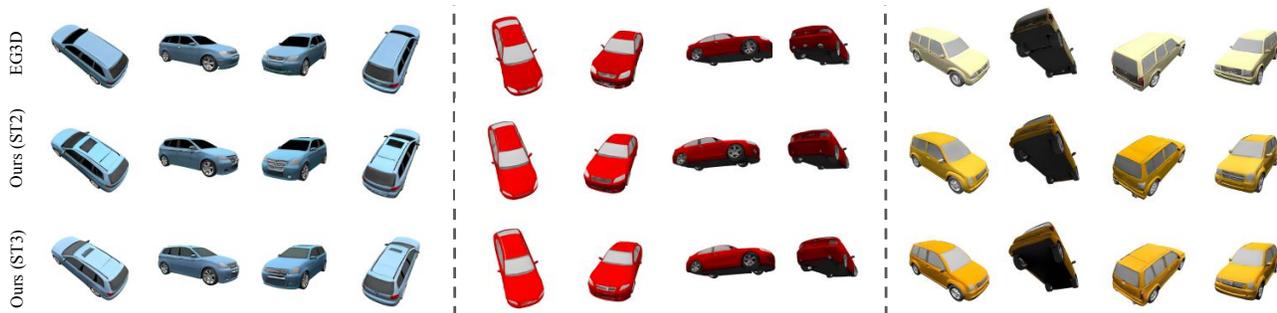


Figure 5. Qualitative examples of different camera poses in Shapenet Cars for three different car models.



Figure 6. one-stage training causes subtle color and geometry inconsistencies (first row). Such inconsistencies can be resolved using our proposed 2-stage training (second row).

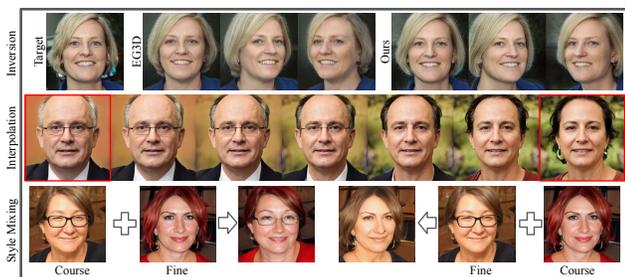


Figure 7. First row: Inversion using PTI [40] for EG3D and our method; second row: interpolation in the latent space of our method; third row: style mixing in the latent space of our method.

EG3D. Similarly, our method exhibits preservation of 3D consistency and image quality, despite the difference in required computational resources.

#### 4.8. Inversion, Interpolation, and Style Mixing

As the proposed generator follows a StyleGAN architecture, it can easily benefit from most of the editing techniques common in the GANs’ literature. Fig. 7 shows examples of inversion using Pivotal Tuning Inversion (PTI) [40], latent space interpolation, and style mixing.

#### 4.9. Discussion: StyleGAN2 V.S. StyleGAN3

StyleGAN2 is more computationally efficient than StyleGAN3. Based on the provided quantitative evaluations, our method reaches comparable image quality and 3D consistency with both architectures. However, StyleGAN2 is known to suffer from more texture stitching and artifacts [24], which we also observe in the generated images.

#### 4.10. Correspondence between Convolutional and Volumetric Rendering

As mentioned before, exploiting the style space of the pretrained NeRF-GAN also provides an opportunity for establishing a direct correspondence between the 3D representation of the 3D generator and the generated images using the convolutional generator. A close comparison of images generated using the convolutional and volumetric rendering in Figures 3, 4, and, 5 indicates that the convolutional render is able to infer and match many attributes of the underlying 3D representation from the shared latent space and generate images similar in content to those of volumetric rendering. However, there still remains a gap in the full correspondence of the two rendering methods, as semantic and identity changes are visible between the corresponding images generated by the two methods. Investigating more explicit approaches for enforcing correspondence could be an interesting direction for improving the convolutional rendering for NeRF-GAN models.

### 5. Conclusion

We presented a method to distill a pretrained NeRF-GAN into a pose-conditioned convolutional generator. The proposed method enables considerably higher efficiency, which is crucial if 3D neural rendering is to become ubiquitous and deployed at scale. To do so, we proposed exploiting the intermediate latent space of the pretrained NeRF-GAN as a conditioning input of the convolutional generator. We additionally provided a training protocol to further improve the visual quality and 3D consistency of the images generated using our generator. Through our experiments, we showed that our method maintains good image quality and 3D consistency, significantly better than previous existing fully-convolutional methods and approaching that of the baseline NeRF-GAN with volumetric rendering. Finally, while our method takes steps toward achieving full correspondence between the two rendering methods, there remains a gap in terms of image semantics. Improving this aspect remains a subject for further research.

## References

- [1] H. Baatz, J. Granskog, M. Pappas, F. Rousselle, and J. Novák. Nerf-tex: Neural reflectance field textures. *Computer Graphics Forum*, 41(6):287–301, 2022.
- [2] 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. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct 2021.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [4] 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.
- [5] 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.
- [6] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [7] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, 2022.
- [8] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [9] 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.
- [10] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- [11] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *IEEE Computer Vision and Pattern Recognition*, 2020.
- [12] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *IEEE Computer Vision and Pattern Recognition Workshops*, 2019.
- [13] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. 2021.
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [15] 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*, 2022.
- [16] Yuan-Chen Guo, Di Kang, Linchao Bao, Yu He, and Song-Hai Zhang. Nerfren: Neural radiance fields with reflections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18409–18418, June 2022.
- [17] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. *ICCV*, 2021.
- [18] 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.
- [19] Tao Hu, Shu Liu, Yilun Chen, Tiancheng Shen, and Jiaya Jia. Efficientnerf efficient neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12902–12911, June 2022.
- [20] Yibo Hu, Xiang Wu, Bing Yu, Ran He, and Zhenan Sun. Pose-guided photorealistic face rotation. *CVPR*, 2018.
- [21] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017.
- [22] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [23] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Advances in Neural Information Processing Systems*, 2020.
- [24] 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.
- [25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. pages 4396–4405, 06 2019.
- [26] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020.
- [27] Marek Kowalski, Stephan J. Garbin, Virginia Estellers, Tadas Baltrušaitis, Matthew Johnson, and Jamie Shotton. Config: Controllable neural face image generation. In *European Conference on Computer Vision (ECCV)*, 2020.
- [28] Jeong-gi Kwak, Yuanming Li, Dongsik Yoon, Donghyeon Kim, David Han, and Hanseok Ko. Injecting 3d perception of controllable nerf-gan into stylegan for editable portrait image synthesis. In *European Conference on Computer Vision*, pages 236–253. Springer, 2022.
- [29] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duck-

- worth. Nerf in the wild: Neural radiance fields for unconstrained photo collections, 2020.
- [30] 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 *ECCV*, 2020.
- [31] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022.
- [32] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022.
- [33] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Joerg H. Mueller, Chakravarty R. Alla Chaitanya, Anton S. Kaplanyan, and Markus Steinberger. DONeRF: Towards Real-Time Rendering of Compact Neural Radiance Fields using Depth Oracle Networks. *Computer Graphics Forum*, 40(4), 2021.
- [34] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *The IEEE International Conference on Computer Vision (ICCV)*, Nov 2019.
- [35] 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 33*, Nov 2020.
- [36] 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.
- [37] 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 (CVPR)*, pages 13503–13513, June 2022.
- [38] Xingang Pan, Bo Dai, Ziwei Liu, Chen Change Loy, and Ping Luo. Do 2d gans know 3d shape? unsupervised 3d shape reconstruction from 2d image gans. In *International Conference on Learning Representations*, 2021.
- [39] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [40] 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.
- [41] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [42] Axel Sauer, Kashyap Chitta, Jens Muller, and Andreas Geiger. Projected gans converge faster. In *NeurIPS*, 2021.
- [43] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. volume abs/2201.00273, 2022.
- [44] 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 (NeurIPS)*, 2020.
- [45] 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.
- [46] Yichun Shi, Divyansh Aggarwal, and Anil Jain. Lifting 2d stylegan for 3d-aware face generation. pages 6254–6262, 06 2021.
- [47] Alon Shoshan, Nadav Bhonker, Igor Kviatkovsky, and Gérard Medioni. Gan-control: Explicitly controllable gans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14083–14093, October 2021.
- [48] Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. Epigraf: Rethinking training of 3d gans. *arXiv preprint arXiv:2206.10535*, 2022.
- [49] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022.
- [50] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Perez, Michael Zollhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images. pages 6141–6150, 06 2020.
- [51] Alex Trevithick, Matthew Chan, Michael Stengel, Eric R. Chan, Chao Liu, Zhiding Yu, Sameh Khamis, Manmohan Chandraker, Ravi Ramamoorthi, and Koki Nagano. Real-time radiance fields for single-image portrait view synthesis. In *ACM Transactions on Graphics (SIGGRAPH)*, 2023.
- [52] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [53] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [54] Liwen Wu, Jae Yong Lee, Anand Bhattad, Yuxiong Wang, and David Forsyth. Diver: Real-time and accurate neural radiance fields with deterministic integration for volume rendering, 2021.
- [55] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021.
- [56] Xiaoming Zhao, Fangchang Ma, David Güera, Zhile Ren, Alexander G. Schwing, and Alex Colburn. Generative mul-

tiplane images: Making a 2d gan 3d-aware. In *Proc. ECCV*, 2022.